

Dependent Types for JavaScript

Ravi Chugh

University of California, San Diego
rchugh@cs.ucsd.edu

David Herman

Mozilla Research
dherman@mozilla.com

Ranjit Jhala

University of California, San Diego
jhala@cs.ucsd.edu

Abstract

We present Dependent JavaScript (DJS), a statically typed dialect of the imperative, object-oriented, dynamic language. DJS supports the particularly challenging features such as run-time type-tests, higher-order functions, extensible objects, prototype inheritance, and arrays through a combination of *nested refinement types*, *strong updates* to the heap, and *heap unrolling* to precisely track prototype hierarchies. With our implementation of DJS, we demonstrate that the type system is expressive enough to reason about a variety of tricky idioms found in small examples drawn from several sources, including the popular book *JavaScript: The Good Parts* and the SunSpider benchmark suite.

Categories and Subject Descriptors D.3.3 [*Programming Languages*]: Language Constructs and Features – Inheritance, Polymorphism; F.3.1 [*Logics and Meanings of Programs*]: Specifying and Verifying and Reasoning about Programs – Logics of Programs

Keywords Refinement Types, JavaScript, Strong Updates, Prototype Inheritance, Arrays

1. Introduction

Dynamic languages like JavaScript, Python, and Ruby are widely popular for building both client and server applications, in large part because they provide powerful sets of features — run-time type tests, mutable variables, extensible objects, and higher-order functions. But as applications grow, the lack of static typing makes it difficult to achieve reliability, security, maintainability, and performance. In response, several authors have proposed type systems which provide static checking for various subsets of dynamic languages [5, 16, 22, 23, 30, 36].

Recently, we developed System D [9], a core calculus for dynamic languages that supports the above dynamic id-

ioms but in a purely functional setting. The main insight in System D is to *dependently* type all values with formulas drawn from an SMT-decidable refinement logic. We use an SMT solver to reason about the properties it tracks well, namely, control-flow invariants and dictionaries with dynamic keys that bind scalar values. But to describe dynamic keys that bind rich values like functions, System D encodes function types as logical terms and *nests* the typing relation as an uninterpreted predicate within the logic. By dividing work between syntactic subtyping and SMT-based validity checking, the calculus supports fully automatic checking of dynamic features like run-time type tests, value-indexed dictionaries, higher-order functions, and polymorphism.

In this paper, we scale up the System D calculus to Dependent JavaScript (abbreviated to DJS), an explicitly typed dialect of a real-world, imperative, object-oriented, dynamic language. We bridge the vast gap between System D and JavaScript in three steps.

Step 1: Imperative Updates. The types of variables in JavaScript are routinely “changed” either by assignment or by incrementally adding or removing fields to objects bound to variables. The presence of mutation makes it challenging to assign precise types to variables, and the standard method of assigning a single “invariant” reference type that overapproximates all values held by the variable is useless in the JavaScript setting. We overcome this challenge by extending our calculus with *flow-sensitive heap types* (in the style of [2, 12, 15, 31, 32]) which allow the system to precisely track the heap location each variable refers to as well as aliasing relationships, thereby enabling *strong updates* through mutable variables. Our formulation of flow-sensitive heaps combined with higher-order functions and refinement types is novel, and allows DJS to express precise pre- and post-conditions of heaps, as in separation logic [17].

Step 2: Prototype Inheritance. Each JavaScript object maintains an implicit link to the “prototype” object from which it derives. To resolve a key lookup from an object at run-time, JavaScript *transitively* follows its prototype links until either the key is found or the root is reached without success. Thus, unlike in class-based languages, inheritance relationships are *computed* at run-time rather than provided as declarative specifications. The semantics of prototypes is

challenging for static typing, because to track the type of a key binding, the system must statically reason about a potentially unbounded number of prototype links! In DJS, we solve this problem with a novel decomposition of the heap into a “shallow” part, for which we precisely track a finite number of prototype links, and a “deep” part, for which we do not have precise information, represented abstractly via a logical heap variable. We *unroll* prototype hierarchies in shallow heaps to precisely model the semantics of object operations, and we use *uninterpreted heap predicates* to reason abstractly about deep parts. In this way, we reduce the reasoning about unbounded, imperative, prototype hierarchies to the underlying decidable, first-order, refinement logic.

Step 3: Arrays. JavaScript arrays are simply objects whose keys are string representations of integers. Arrays are commonly used both as *heterogeneous* tuples (that have a fixed number of elements of different types) as well as *homogeneous* collections (that have an unbounded number of elements of the same type). The overloaded use of arrays, together with the fact that arrays are otherwise syntactically indistinguishable and have the same prototype-based semantics as non-array objects, makes it hard to statically reason about the very different ways in which they are used. In DJS, we use nested refinements to address the problem neatly by uniformly encoding tuples and collections with refinement predicates, and by using *intersection* types that simultaneously encode the semantics of tuples, collections, and objects.

Expressiveness. We have implemented DJS (available at ravichugh.com/djs) and demonstrated its expressiveness by checking a variety of properties found in small but subtle examples drawn from a variety of sources, including the popular book *JavaScript: The Good Parts* [10] and the SunSpider benchmark suite [34]. Our experiments show that several examples *simultaneously* require the gamut of features in DJS, but that many examples conform to recurring patterns that rely on particular aspects of the type system. We identify several ways in which future work can handle these patterns more specifically in order to reduce the annotation burden and performance for common cases, while falling back to the full expressiveness of DJS in general. Thus, we believe that DJS provides a significant step towards truly retrofitting JavaScript with a practical type system.

2. Overview

Let us begin with an informal overview of the semantics of JavaScript. We will emphasize the aspects that are the most distinctive and challenging from the perspective of type system design, and describe the key insights in our work that overcome these challenges.

JavaScript Semantics by Desugaring. Many corner cases of JavaScript are clarified by λ_{JS} [21], a syntax-directed translation, or *desugaring*, of JavaScript programs to a mostly-standard lambda-calculus with explicit references.




Figure 1. Architecture of DJS

As λ_{JS} is a core language with well-understood semantics and proof techniques, the translation paves a path to a typed dialect of JavaScript: define a type system for the core language and then type check desugared JavaScript programs.

We take this path by developing System !D (pronounced “D-ref”), a new calculus based on λ_{JS} . Although the operational semantics of System !D is straightforward, the *dynamic* features of the language ensure that building a type system expressive enough to support desugared JavaScript idioms is not. We solve this problem by scaling the purely functional technique of nested refinement types up to the imperative, object-oriented, setting of real-world JavaScript.

Figure 1 depicts the architecture of our approach: we desugar a Dependent JavaScript (DJS) file `program.js` to the System !D file `program.dref`, which is analyzed by the type checker along with a standard prelude comprising three files (`basics.dref`, `objects.dref`, and `prelude.dref`) that model JavaScript semantics.

Terminology. JavaScript has a long history and an evolving specification. In this paper, we say “JavaScript” to roughly mean ECMAScript Edition 3, the standard version of the language for more than a decade [24]. We say “ES5” to refer to Edition 5 of the language, recently released by the JavaScript standards committee [13]; Edition 4 was never standardized. We say “ES6” to refer to features proposed for the next version of the language, scheduled to be finalized within the next one or two years. DJS includes a large set of core features common to all editions.

2.1 Base Types, Operators, and Control Flow

Consider the following function adapted from [9] and annotated in DJS. A function type annotation is written just above the definition inside a JavaScript comment demarcated by an additional `:` character. We typeset annotations in math mode for clarity, but the ASCII versions parsed by our type checker are quite similar.

```

/*:  $x: \text{Top} \rightarrow \{\nu \mid \text{ite } \text{Num}(x) \text{ Num}(\nu) \text{ Bool}(\nu)\}$  */
function negate(x) {
  if (typeof x == "number") { return 0 - x; }
  else { return !x; } }
  
```

The `typeof` operator is a facility used pervasively to direct control flow based on the run-time “tag” of a value. If the input to `negate` is a number, so is the return value. If not, the function uses an interesting feature of JavaScript, namely,

```

val typeof :: (* x:Top → {ν = tag(x)} *)
val      ! :: (* x:Top → {ν iff falsy(x)} *)
val      (||) :: (* x:Top → y:Top → {ite falsy(x) (ν = y) (ν = x)} *)
val      (&&) :: (* x:Top → y:Top → {ite truthy(x) (ν = y) (ν = x)} *)
val      (===) :: (* x:Top → y:{tag(ν) = tag(x)} → {ν iff (x = y ∧ x ≠ NaN)} *)
val      (==) :: (* x:Top → y:Top → {Bool(ν) ∧ (tag(x) = tag(y) ⇒ ν iff (x = y ∧ x ≠ NaN))} *)
val      (+) :: (* x:Str → y:Str → Str *)
val      (+) :: (* x:Num → y:Num → {Num(ν) ∧ ((Int(x) ∧ Int(y)) ⇒ (Int(ν) ∧ ν = x + y))} *)
val      fix :: (* ∀A. (A → A) → A *)

```

Figure 2. Excerpt from basics.dref

that *all* values have a boolean interpretation. The values false, null, undefined, the empty string “”, 0, and the “not-a-number” value NaN are considered *falsy*, and evaluate to false when used in a boolean context; all other values are *truthy*. The operator ! inverts “truthiness,” so the else branch returns a boolean no matter what the type of *x* is. The ability to treat arbitrary values as booleans is commonly used, for example, to guard against null values.

The `negate` function demonstrates that even simple JavaScript programs depend heavily on sophisticated control-flow based reasoning. Syntactic type systems are capable of tracking control flow to a limited degree [22, 36], but none can handle complex invariants like the relationship between the input and output of `negate`. To have any chance of capturing such invariants, types must be able to *depend* on other program values. Powerful dependent type systems like Coq can express extremely rich invariants, but are too heavyweight for our goals since they require the programmer to discharge type checking obligations interactively.

Refinement Types. We adopt a more lightweight mechanism called *refinement types* that has been previously applied to purely functional dynamic languages [5, 9]. We demonstrate that refinement types afford us the expressiveness needed to precisely track control-flow invariants in the JavaScript setting and, unlike more powerful dependent systems, without sacrificing decidable type checking. In particular, once the programmer has written type annotations for function definitions, type checking is carried out automatically via a combination of syntactic subtyping and SMT-based [11] logical validity checking.

In System !D, every value is described by a refinement type of the form $\{\nu | p\}$, read “ ν such that p ”, where p is a formula that can mention ν . For example, 3 can be given the type $\{\nu | \text{tag}(\nu) = \text{“number”}\}$ and `true` the type $\{\nu | \text{tag}(\nu) = \text{“boolean”}\}$, where *tag* is an uninterpreted function symbol in the refinement logic, *not* a function in the programming language. We define abbreviations in Figure 3 to make the refinement binder implicit and the types concise.

Primitives. We use refinements to assign precise, and sometimes exact, types to System !D primitive functions,

$$\begin{aligned}
\{p\} &\stackrel{\circ}{=} \{\nu | p\} & Num(x) &\stackrel{\circ}{=} \text{tag}(x) = \text{“number”} \\
Top(x) &\stackrel{\circ}{=} true & Bool(x) &\stackrel{\circ}{=} \text{tag}(x) = \text{“boolean”} \\
T &\stackrel{\circ}{=} \{T(\nu)\} & Str(x) &\stackrel{\circ}{=} \text{tag}(x) = \text{“string”} \\
\text{if } p \text{ then } q_1 \text{ else } q_2 &\stackrel{\circ}{=} \text{ite } p \ q_1 \ q_2 &&\stackrel{\circ}{=} (p \Rightarrow q_1) \wedge (\neg p \Rightarrow q_2) \\
x \text{ iff } p &\stackrel{\circ}{=} \text{ite } p \ (x = true) \ (x = false) \\
falsy(x) &\stackrel{\circ}{=} x \in \{\text{false} \vee \text{null} \vee \text{undefined} \vee \text{“”} \vee 0 \vee \text{NaN}\} \\
truthy(x) &\stackrel{\circ}{=} \neg falsy(x)
\end{aligned}$$

Figure 3. Abbreviations for common types

defined in the file `basics.dref` (Figure 2). Notice that `typeof` returns the tag of its input. Some examples beyond ones we have already seen include $\text{tag}(\text{null}) = \text{“object”}$ and $\text{tag}(\text{undefined}) = \text{“undefined”}$. The type of the negation operator ! inverts “truthiness.” The types of the operators `&&` and `||` are interesting, because as in JavaScript, they do not necessarily return booleans. The “guard” operator `&&` returns its second operand if the first is truthy, which enables the idiom `if (x && x.f) { ... }` that checks whether the object *x* and its “*f*” field are non-null. Dually, the “default” operator `||` returns its second operand if the first is falsy, which enables the idiom `x = x || default` to specify a default value. The `+` operator is specified as an *intersection* of function types and captures the fact that it performs both string concatenation and numerical addition, but does *not* type check expressions like `3 + “hi”` that rely on the *implicit coercion* in JavaScript. We choose types for System !D primitives that prohibit implicit coercions since they often lead to subtle programming errors.

Equality. JavaScript provides two equality operators: `==` implicitly coerces the second operand if its tag differs from the first, and strict equality `===` does not perform any coercion. To avoid reasoning about implicit coercions, we give a relatively weaker type to `==`, where the boolean result relates its operands *only if* they have the same tag.

Integers. JavaScript provides a single number type that has no minimum or maximum value. However, programmers and optimizing JIT compilers [38] often distinguish integers from arbitrary numbers. In System !D, we describe integers with the abbreviation $\text{Int}(x) \stackrel{\circ}{=} \text{Num}(x) \wedge \text{integer}(x)$. We introduce the uninterpreted predicate *integer*(*x*) in the types

<pre>function also_negate(x) { if (typeof x == "number") x = 0 - x; else x = !x; return x; }</pre>	<pre>1 let also_negate = fun x -> (* Γ₁ = ∅; Σ₁ = ∅ *) 2 let _x = ref x in (* Γ₂ = x:Top; Σ₂ = (ℓ_x ↦ x) *) 3 if typeof (deref _x) == "number" then (* Γ₃ = Γ₂, _x:Ref ℓ_x; Σ₃ = Σ₂ *) 4 _x := 0 - (deref _x) (* Γ₄ = Γ₃, Num(x); Σ₄ = ∃x₄:Num. (ℓ_x ↦ x₄) *) 5 else 6 _x := !(deref _x) (* Γ₆ = Γ₃, ¬Num(x); Σ₆ = ∃x₆:{ν iff falsy(x)}. (ℓ_x ↦ x₆) *) 7 ; (* Γ₇ = Γ₃; Σ₇ = ∃x':{ite Num(x) Num(ν) Bool(ν)}. (ℓ_x ↦ x') *) 8 deref _x in (* Γ₈ = Γ₃; Σ₈ = Σ₇ *) 9 let _also_negate = ref {"__code__": also_negate}</pre>
--	--

Figure 4. DJS function `also_negate`; Desugared to System !D; Verifying $x:Top \rightarrow \{\text{ite } Num(x) \text{ Num}(\nu) \text{ Bool}(\nu)\}$

of integer literals, and functions like `+` propagate “integer-ness” where possible. Furthermore, numeric functions use the (decidable) theory of linear arithmetic to precisely reason about integers, which is important for dealing with arrays.

Tracking Control Flow. System !D tracks control flow precisely by recording that the guard of an if-expression is *truthy* (resp. *falsy*) along the then-branch (resp. else-branch), enabling System !D to verify the annotation for `negate` as follows. Because of the call to `typeof`, System !D tracks that $Num(x)$ holds along the then-branch, so `x` can be safely passed to the subtraction operator which produces a number as required. For the else-branch, System !D records that $\neg Num(x)$. The negation operator, which can be applied to any value, produces a value of type $\{\nu \text{ iff } \neg falsy(x)\}$ which is a subtype of $Bool$. Thus, both branches satisfy the specification provided by the programmer.

2.2 Imperative Updates

JavaScript is an imperative language where variables can be reassigned arbitrary values. Consider the DJS function `also_negate` in Figure 4 that is like `negate` but first assigns the eventual result in the variable `x`, and its translation to System !D on the right (ignore the comments for now).

Several aspects of the translation warrant attention. First, since the formal parameter `x`, like all JavaScript variables, is mutable, the translation of the function body begins with an *explicit reference* cell `_x` initialized with `x`, and each read of `x` is desugared to a dereference of `_x`. Presentations of imperative languages often model assignable variables directly rather than with explicit references. Both approaches are equivalent; we choose the latter to make the presentation more similar to λ_{JS} [21] and System D [9]. Second, notice that scalar constants like `0` and `true` and operators like `typeof` and `==` are translated directly to corresponding ones in System !D. Third, notice that each assignment to the variable `x` translates to a set reference (*i.e.*, assignment) operation to update the contents of the heap cell. Finally, since every JavaScript function is actually an object, the translation stores the function value in a distinguished “`__code__`” field, which we assume is inaccessible to source programs. Instead of this assumption, we could treat each function as a pair of a function value and an associated object, but we

follow the λ_{JS} encoding for simplicity. For System !D to verify that `also_negate` satisfies the specification, it must precisely reason about heap updates in addition to control-flow as before.

Reference Types. The traditional way to handle references in the λ -calculus [28] is to (a) assign a reference cell some type $Ref\ T$, (b) require that only values of type T be stored in it, and then (c) conclude that dereferences produce values of type T . This approach supports so-called *weak updates*, because even if a stored value satisfies a stronger type S than T (*i.e.*, if S is a subtype of T), subsequent dereferences produce values of the original, weaker type T . Put another way, this approach requires that the type assigned to a reference cell be a supertype of all the values written to the cell. Unfortunately, weak updates would preclude System !D from verifying `also_negate`. The initialization of `_x` on line 2 stores the parameter `x` which has type Top , so `_x` would be assigned type $Ref\ Top$. The assignments on lines 4 and 6 type check because the updated values satisfy the trivial type Top , but the dereference on line 8 produces a value with type Top , which does *not* satisfy the specified return type. Thus, we need a way to reason more precisely about heap updates.

Strong Updates. Allowing assignment to *change* the type of a reference is called *strong update*, which is sound only when a reference is guaranteed to point to a *single* heap cell and when there are no accesses through other *aliases* that refer to the same cell. The *Alias Types* approach [32] provides a means of managing these concerns. Rather than $Ref\ T$, a reference type is written $Ref\ \ell$, where ℓ is the (compile-time) name of a location in the heap, and a separate (compile-time) heap maps locations to types, for example, $(\ell \mapsto T)$. Strong updates are realized by allowing heaps to change *flow-sensitively*, and the aliasing problem is mitigated by maintaining the invariant that distinct location names ℓ and ℓ' do not alias. System !D employs this approach by using a *type environment* Γ that grows and shrinks as usual during type checking but remains flow-insensitive, and a *heap environment* Σ that can be strongly updated per program point.

Figure 4 shows how System !D checks the desugared version of `also_negate`. The figure shows, at each line i , the type environment Γ_i used to check the expression on the line, and the heap environment Σ_i that exists *after* checking

the expression. After starting with the empty heap $\Sigma_1 = \emptyset$, the allocation on line 2 creates a fresh location ℓ_x in the new heap $\Sigma_2 \stackrel{\circ}{=} \Sigma_1 \oplus (\ell_x \mapsto \mathbf{x})$ and adds $_x: \text{Ref } \ell_x$ to the type environment. We use the symbol \oplus to construct unordered sets of heap bindings. To exploit the precision of dependent types, we map locations to *values* rather than types (i.e., $(\ell \mapsto \mathbf{x})$ rather than $(\ell \mapsto \text{Top})$).

When checking the if-expression guard on line 3, the deference retrieves the initial value \mathbf{x} from the heap Σ_2 , so as a result of the tag-test, System !D adds $\text{Num}(\mathbf{x})$ to the type environment Γ_4 along the true-branch and $\neg \text{Num}(\mathbf{x})$ to Γ_6 along the false-branch. In the true-branch, the subtraction on line 4 is well-typed because $\text{Num}(\mathbf{x})$, and produces a number x_4 that is stored in the heap Σ_4 at location ℓ_x . In the false-branch, \mathbf{x} is negated on line 6, producing a boolean x_6 that is stored in the heap Σ_6 at location ℓ_x . System !D combines the branches by *joining* the heaps Σ_4 and Σ_6 , producing Σ_7 that describes the heap no matter which branch is taken. The dereference on line 8 retrieves x' , a value of type $\{\text{ite } \text{Num}(x) \text{ Num}(\nu) \text{ Bool}(\nu)\}$, as required by the return type annotation.

In this way, System !D syntactically tracks strong updates to the heap, while reducing subtyping obligations to implication queries in an ordinary, pure refinement logic [9] that does *not* model imperative updates.

2.3 Simple Objects

JavaScript’s objects exhibit several interesting semantic properties. Consider the following object initialized with a single key (also known as field or property). We assume that `assert` is a pre-defined function that aborts when its argument is falsy; JavaScript does not provide such a function as built-in, but it is trivial to define.

```
var x = {"f": 1};
assert (x.f == 1 && x.g == undefined);
x.g = 2; delete x.f;
assert (x.g == 2 && x.f == undefined);
x.f.g; // raises exception
var k = "h"; x[k] = 3; assert (x[k] == 3);
```

Notice that when retrieving the non-existent “g” key from `x`, JavaScript returns `undefined` as opposed to raising an exception. Attempting to retrieve a key from `undefined`, or `null`, however, does raise an exception. Keys can be added or removed to objects, and can even be arbitrary *dynamically-computed* values, not just string literals, that are converted to strings if necessary. Dynamic keys are pervasive — objects are commonly used as hash tables with unknown sets of keys — but they are quite challenging to track inside a static type system.

Nested Refinements. To support dynamic keys, we adopt the System D primitives [9] for (functional) dictionary operations, shown in the first four lines of the file `objects.dref` (Figure 5). The primitive function application `get d k re-`

trieves the key k from dictionary d , where $\text{sel}(d, k)$ describes the exact binding as a value in the refinement logic; `set d k y` produces a new dictionary that extends d with a binding for k , shadowing previous bindings, if any, where $\text{upd}(d, k, y)$ describes the new dictionary; `del d k` produces a new dictionary with a binding removed, using the logical symbol bot (distinct from all source-level values) to denote its absence; and `mem d k` indicates the presence or absence of a binding, where we write the abbreviation $\text{has}(d, k) \stackrel{\circ}{=} \text{sel}(d, k) \neq \text{bot}$.

The key innovation of *nested refinements* in System D allows *syntactic type terms* U (like function types) to be written within refinement formulas using an uninterpreted “has-type” predicate $x :: U$, while staying within the decidable McCarthy theory of arrays [27]. The has-type predicate allows System D to describe dictionaries that map *dynamic* keys to *arbitrary* values. For example, we write $\{\text{Dict}(\nu) \wedge \text{sel}(\nu, k) :: \text{Bool} \rightarrow \text{Bool}\}$ to describe a dictionary d with key k that binds a boolean-to-boolean function, and $\{\nu = \text{upd}(d, \text{“g”}, 4)\}$ to describe a dictionary d' that is just like d but with an additional binding. Prior approaches such as [5] were limited to dynamic keys that store first-order (non-function) values. We refer the reader to [9] for the technical development of nested refinements.

Mutability and Dynamic Keys. The combination of nested refinements and strong updates allows us to precisely track objects with dynamic keys despite the presence of imperative updates. Consider the desugaring of our example above; we omit the assertions for clarity.

```
1 let _x = ref (ref {"f": 1}) in
2 _x := set (deref (deref _x)) "g" 2;
3 _x := del (deref (deref _x)) "f";
4 let _k = "h" in
5 _x := set (deref (deref _x))
6           (coerceToStr (deref _k)) 3;
```

The allocation on line 1 adds three bindings to the type environment — $d: \{\nu = \text{upd}(\text{empty}, \text{“f”}, 1)\}$, $\text{ptr}: \text{Ref } \ell$, and $_obj: \text{Ref } \ell'$, where ℓ and ℓ' are fresh locations — and produces the heap $\Sigma_1 \stackrel{\circ}{=} (\ell' \mapsto \text{ptr}) \oplus (\ell \mapsto d)$. Notice that the dictionary is stored via an additional level of indirection to facilitate the encoding of *side-effecting* JavaScript object operations. The object extension on line 2 adds $d': \{\nu = \text{upd}(d, \text{“g”}, 2)\}$ to the type environment and strongly updates the heap to $\Sigma_2 \stackrel{\circ}{=} (\ell' \mapsto \text{ptr}) \oplus (\ell \mapsto d')$. The deletion on line 3 and the extension on line 5 (through a dynamic key) have similar effects on the static heap, thereby statically verifying the assertions.

2.4 Function Types

In order to fully understand the challenges of JavaScript objects, we must pause to take a closer look at function types. The function types we have seen so far — for `negate` and the primitives in `basics.dref` — have not mentioned

```

val set :: (* d:Dict → k:Str → y:Top → {ν = upd(d,k,y)} *)
val del :: (* d:Dict → k:Str → {ν = upd(d,k,bot)} *)
val has :: (* d:Dict → k:Str → {ν iff has(d,k)} *)
val get :: (* d:Dict → k:Str → {ite has(d,k) (ν = sel(d,k)) (ν = undefined)} *)

val setPropObj :: (* (x:Ref, k:Str, y:Top)/(x ↦ ⟨d:Dict, x̂⟩) → {ν = y}/(x ↦ ⟨d':{ν = upd(d,k,y)}, x̂⟩) *)
val delPropObj :: (* (x:Ref, k:Str)/(x ↦ ⟨d:Dict, x̂⟩) → Bool/(x ↦ ⟨d':{ν = upd(d,k,bot)}, x̂⟩) *)
val hasPropObj :: (* (x:Ref, k:Str)/(x ↦ ⟨d:Dict, x̂⟩) → {ν iff ObjHas(d,k,cur,x̂)}/same *)
val getPropObj :: (* (x:Ref, k:Str)/(x ↦ ⟨d:Dict, x̂⟩)
→ {ite ObjHas(d,k,cur,x̂) (ν = ObjSel(d,k,cur,x̂)) (ν = undefined)}/same *)

val getIdxArr :: (* ∀A. (x:Ref, i:Int)/(x ↦ ⟨a:Arr(A), x̂⟩)
→ {ite ¬packed(a) (ν :: A ∨ Undef(ν)) (ite (0 ≤ i < len(a)) (ν :: A) (Undef(ν)))}/same *)
val getLenArr :: (* ∀A. (x:Ref, k:{ν = "length"})/(x ↦ ⟨a:Arr(A), x̂⟩)
→ {ite packed(a) (ν = len(a)) Int(ν)}/same *)
val getPropArr :: (* ∀A. (x:Ref, k:{Str(ν) ∧ ν ≠ "length"})/(x ↦ ⟨a:Arr(A), x̂⟩)
→ {ite HeapHas(H,x̂,k) (ν = HeapSel(H,x̂,k)) (ν = undefined)}/same *)

val getElem :: (and (type getPropObj) (type getIdxArr) (type getLenArr) (type getPropArr))

val setIdxArr :: (* ∀A. (x:Ref, i:Int, y:A)/(x ↦ ⟨a:Arr(A), x̂⟩)
→ {ν = y}/(x ↦ ⟨a':{ν :: Arr(A) ∧ arrSet(ν,a,i)}, x̂⟩) *)

```

Figure 5. Excerpt from `objects.dref`

heaps, because their inputs and outputs are scalar values. However, JavaScript objects are reference values, and are passed to and returned from functions through the heap. Thus, to account for heaps and side-effects, a System !D function type has the following form.

$$\forall[\bar{A}; \bar{L}; \bar{H}] x_1:T_1/\hat{\Sigma}_1 \rightarrow x_2:T_2/\hat{\Sigma}_2$$

This type describes a function that, given an argument x_1 of type T_1 in a calling context that satisfies the input heap type $\hat{\Sigma}_1$, produces an output value x_2 of type T_2 and a modified heap type $\hat{\Sigma}_2$. A function type can be parameterized by sequences of *type variables* A , *location variables* L , and *heap variables* H . A heap type $\hat{\Sigma}$ is like a heap environment Σ but maps locations to binder-type pairs rather than values (e.g., $(\ell \mapsto y:T)$ rather than $(\ell \mapsto v)$); the binders name the values stored in the heap before and after a function call. The binder x_1 and all of the binders in the input heap $\hat{\Sigma}_1$ are in scope in the output *world* $x_2:T_2/\hat{\Sigma}_2$. Intuitively, a world describes a tuple where every component except the first resides in a heap. We often omit binders when they are not referred to. To match the structure of function types, function applications must instantiate type, location, and heap variables. However, our implementation infers instantiations in almost all cases using standard local type inference techniques (§6). When we write DJS examples in the sequel, we omit instantiations at applications wherever our current implementation infers them. We sweeten function type syntax with some sugar:

- When used as an output heap, the token *same* refers to the sequence of locations in the corresponding input

heap, where each binding records that the final value is exactly equal to the initial value.

- In an input world, a reference binding $x:Ref$ without a location introduces a location variable L that is quantified by the type, and x (a value of type $Ref L$) can be used as a location in heaps to refer to this variable L . Further, the dotted variable \dot{x} introduces a location parameter, corresponding to the prototype of x .
- A heap variable H is implicitly added to a function type when it contains none, and H is added to both the input and output heaps; this variable corresponds to the “frame” from separation logic [17]. In this case, the token *cur* refers to H .

For example, compare the type for `hasPropObj` (Figure 5) followed by its expansion.

$$\begin{aligned}
& (x:Ref, k:Str)/(x \mapsto \langle d:Dict, \dot{x} \rangle) \\
& \rightarrow \{\nu \text{ iff } ObjHas(d,k,cur,\dot{x})\}/same \\
& \forall L, L', H. (x:Ref L, k:Str)/H \oplus (L \mapsto \langle d:Dict, L' \rangle) \\
& \rightarrow \{\nu \text{ iff } ObjHas(d,k,H,L')\}/H \oplus (L \mapsto \langle d':\{\nu = d\}, L' \rangle)
\end{aligned}$$

2.5 Prototype-Based Objects

JavaScript sports a special form of inheritance, where each base object is equipped with a link to its *prototype object*. This link is set when the base object is created and cannot be changed or accessed by the program. When trying to retrieve a key k not stored in an object x itself, JavaScript *transitively* searches the *prototype chain* of x until it either finds k or it reaches the root of the object hierarchy without finding k .

The prototype chain does not play a role in the semantics of key update, addition, or deletion.¹

For example, consider the initially empty object `child` created by the function `beget` (described in the sequel) with prototype object `parent`. The prototype of object literals, like `parent`, is the object stored in `Object.prototype` (note that the “prototype” key of `Object` is *not* the same as its prototype object). Thus, all keys in `parent` and `Object.prototype` are transitively accessible via `child`.

```
var parent = {"last": "Doe"};
var child = beget(parent);
child.first = "John";
assert (child.first + child.last == "John Doe");
assert ("last" in child == true);
assert (child.hasOwnProperty("last") == false);
```

The JavaScript operator `k in x` tests for the presence of `k` *anywhere* along the prototype chain of `x`, whereas the native function `Object.prototype.hasOwnProperty` tests only the “own” object itself. Keys routinely resolve through prototypes, so a static type system must precisely track them. Unfortunately, we cannot encode prototypes directly within the framework of refinement types and strong update, as the semantics of transitively traversing mutable and unbounded prototype hierarchies is beyond the reach of decidable, first-order reasoning.

Shallow and Deep Heaps. We solve this problem by syntactically reducing reasoning about prototype-based objects to the refinement logic. Our key insight is to decompose the heap into a “shallow” part, the bounded portion of the heap for which we have explicit locations, and a “deep” part, which is the potentially unbounded portion which we can represent by uninterpreted heap variables H . We explicitly track prototype links in the “shallow” heap by using bindings of the form $(\ell \mapsto \langle d, \ell' \rangle)$, where the prototype of the object at ℓ is stored at ℓ' . We cannot track prototype links explicitly in the “deep” heap, so instead we *summarize* information about deep prototype chains by using the abstract (uninterpreted) heap predicate $HeapHas(H, \ell, k)$ to encode the proposition that the object stored at location ℓ in H *transitively* has the key k , and the abstract (uninterpreted) heap function $HeapSel(H, \ell, k)$ to represent the corresponding value retrieved by lookup.

As an example, recall the `child` object and its prototype `parent`. Suppose that the prototype of `parent` is an unknown object `grandpa`, rather than `Object.prototype` as written. If `child`, `parent`, and `grandpa` are stored at locations ℓ_1 , ℓ_2 , and ℓ_3 with underlying “own” dictionary values d_1 , d_2 , and d_3 , then we write the heap $\{\ell_1 \mapsto \langle d_1, \ell_2 \rangle, \ell_2 \mapsto \langle d_2, \ell_3 \rangle, \ell_3 \mapsto \langle d_3, \ell_4 \rangle, H\}$ — we use set notation to abbreviate the concatenation of heap bindings

with \oplus . Despite not knowing what value is the prototype of `grandpa`, we name its location ℓ_4 that is somewhere in the deep part of the heap H .

Key Membership and Lookup. When describing simple objects, we used the original System D primitives (`mem` and `get`) to desugar key membership and lookup operations. But in fact, to account for the transitive semantics of key membership and lookup facilitated by prototype links, System !D uses new primitives `hasPropObj` and `getPropObj` defined in `objects.dref` (Figure 5). These primitives differ from their purely functional System D counterparts in two ways: each operation goes through a *reference* to a dictionary on the heap, and the abstract predicates $ObjHas$ and $ObjSel$ are used in place of has and sel . These abstract predicates are defined over the *disjoint union* of the shallow and deep heaps as follows and, intuitively, summarize whether an object transitively has a key and, if so, the value it binds.

$$ObjHas(d, k, H, \dot{x}) \stackrel{\circ}{=} has(d, k) \vee HeapHas(H, \dot{x}, k)$$

$$\nu = ObjSel(d, k, H, \dot{x}) \stackrel{\circ}{=} \text{if } has(d, k) \text{ then } \nu = sel(d, k) \\ \text{else } \nu = HeapSel(H, \dot{x}, k)$$

Transitive Semantics via Unrolling. Let us return to the example of the `child`, `parent` and `grandpa` prototype chain to understand how unrolling captures the semantics of transitive lookup. The DJS key membership test on the left desugars to System !D on the right as follows.

<code>k in child</code>	<code>hasPropObj (deref _child, deref _k)</code>
-------------------------	--

The result of the function call has the following type.

$$\{\nu \text{ iff } ObjHas(d_1, k, \{\ell_2 \mapsto \langle d_2, \ell_3 \rangle, \ell_3 \mapsto \langle d_3, \ell_4 \rangle, H\}, \ell_2)\}$$

We expand this type by unrolling $ObjHas$ to the following.

$$\{\nu \text{ iff } has(d_1, k) \vee has(d_2, k) \vee has(d_3, k) \vee HeapHas(H, \ell_4, k)\}$$

The first three disjuncts correspond to looking for `k` in the shallow heap, and the last is the uninterpreted predicate that summarizes whether `k` exists in the deep heap. Similarly, key lookup in DJS on the left is desugared as follows.

<code>child[k]</code>	<code>getPropObj (deref _child, deref _k)</code>
-----------------------	--

We unroll the type of the System !D expression as follows.

$$\{\text{if } has(d_1, k) \text{ then } \nu = sel(d_1, k) \text{ else} \\ \text{if } has(d_2, k) \text{ then } \nu = sel(d_2, k) \text{ else} \\ \text{if } has(d_3, k) \text{ then } \nu = sel(d_3, k) \text{ else} \\ \text{ite } HeapHas(H, \ell_4, k) (\nu = HeapSel(H, \ell_4, k)) \text{ } Undef(\nu)\}$$

Thus, our technique of decomposing the heap into shallow and deep parts, followed by heap unrolling, captures the *exact* semantics of prototype-based object operations modulo the unknown portion of the heap. Thus, System !D precisely tracks objects in the presence of mutation and prototypes.

¹ Many implementations expose the prototype of an object `x` with a non-standard `x...proto...` property, and prototypes do affect key update in ES5. We discuss these issues further in §7.

```

var __hasOwn =
  /*: (this:Ref, k:Str)/(this ↦ {d:Dict, this}) → {ν iff has(d,k)}/same
  ∧ ∀A. (this:Ref, i:Int)/(this ↦ {a:Arr(A), this}) → {ite packed(a) (ν iff 0 ≤ i < len(a)) Bool(ν)}/same
  ∧ ∀A. (this:Ref, k:Str)/(this ↦ {a:Arr(A), this}) → {ν iff k = "length"}/same */ "#extern";

function Object() { ... }; Object.prototype = {"hasOwnProperty": __hasOwn, "constructor": Object, ... };

var __push = /*: ∀A. (this:Ref, x:A)/(this ↦ {a:Arr(A), this})
  → Int/(this ↦ {a':{ν :: Arr(A) ∧ arrSize(ν, a, 1)}, this}) */ "#extern";
var __pop = /*: ∀A. (this:Ref, x:A)/(this ↦ {a:Arr(A), this})
  → {ite packed(a) (ν :: A) (ν :: A ∨ Undef(ν))}
  / (this ↦ {a':{ν :: Arr(A) ∧ arrSize(ν, a, -1)}, this}) */ "#extern";

function Array() { ... }; Array.prototype = {"push": __push, "pop": __pop, "constructor": Array, ... };

```

Figure 6. Excerpt from `prelude.js`, which desugars to the `prelude.dref` file in the standard prelude

Additional Primitives. The new update and deletion primitives `setPropObj` and `delPropObj` (Figure 5) affect only the “own” object, since the prototype chain does not participate in the semantics. We model native JavaScript functions like `Object.prototype.hasOwnProperty` with type annotations in the file `prelude.js` (Figure 6). Notice that the function type for objects (the first in the intersection) checks only the “own” object for the given key.

Constructors. JavaScript provides the expression form `new Foo(args)` as a second way of *constructing* objects, in addition to object literals whose prototypes are set to `Object.prototype`. The semantics are straightforward, but quite different than the traditional `new` syntax suggests. Here, if `Foo` is any function (object), then a fresh, empty object is created with prototype object `Foo.prototype`, and `Foo` is called with the new object bound to `this` (along with the remaining arguments) to finish its initialization. We desugar constructors and `new` with standard objects and functions (following λ_{JS} [21]) *without* adding any special System !D constructs or primitive functions.

Inheritance. Several inheritance relationships, including ones that simulate traditional classes, can be encoded with the construction mechanism, as shown in the popular book *JavaScript: The Good Parts* [10]. Here, we examine the *prototypal pattern*, a minimal abstraction which wraps construction to avoid the unusual syntax and semantics that leads to common errors; we discuss the rest in § 6. The function `beget` (the basis for `Object.create` in ES5) returns a fresh empty object with prototype `o`.

```

1 /*: ∀L. o:Ref/(o ↦ {d:Dict, o})
2   → Ref L/(L ↦ {{ν = empty}, o}) ⊕ (o ↦ same) */
3 function beget(o) {
4   /*: #ctor this:Ref → {ν = this} */
5   function F() { return this; };
6   F.prototype = o;
7   return new /*:L*/ F(); }

```

The `#ctor` on line 4 instructs desugaring to: initialize the function object with a “prototype” key that stores an empty object literal (since it will be called as a constructor); and expand the type annotation as follows to require that `this` initially be empty, as is common for all constructors.

$$this:Ref / (this \mapsto \{\nu = empty\}, this) \rightarrow \{\nu = this\} / same$$

The assignment on line 6 strongly updates `Foo.prototype` (overwriting its initial empty object) with the argument `o`. Thus, the object constructed (at location L) on line 7 has prototype `o`, so `beget` has the ascribed type. In most cases, `new` can be used without a location annotation and a fresh one is chosen. In this case, we annotate line 7 with L (from the type of `beget`), which our implementation does not infer because there is no input corresponding to L .

2.6 Arrays

The other workhorse data structure of JavaScript is the array, which is really just an object with integer “indices” converted to ordinary string keys. However, arrays pose several tricky challenges as they are commonly used both as finite tuples as well as unbounded collections.

```

var arr = [17, "hi", true];
arr[3] = 3; arr.push(4);
assert (arr.length == 5 && arr[5] == undefined);

```

As for any object, retrieving a non-existent key returns `undefined` rather than raising an “out-of-bounds” exception. Like other objects, arrays are extensible simply by writing “past the end.” Array literal objects have prototype `Array.prototype`, which includes a `push` (resp. `pop`) function for adding an element to (resp. removing an element from) the end of an array.

Loops are used to iterate over arrays of unknown size. But since lookups may return `undefined`, it is important to track when an access is “in-bounds.” JavaScript bestows upon arrays an unusual “length” property, rather than a method,

to help. Reading it returns the largest integer key of the array, which is *not* necessarily its “size” because it may contain “holes” or even non-integer keys. Furthermore, assigning a number n to the “length” of an array either *truncates* it if n is less than its current length, or *extends* it (by padding with holes) if it is greater. Despite the unusual semantics, programmers commonly use arrays as if they are traditional “packed” arrays with integer “indices” zero to “size” minus one. The type system must reconcile this discrepancy.

Array Types. We introduce a new syntactic type term $Arr(T)$ and maintain the following four properties for every value a that satisfies the has-type predicate $a :: Arr(T)$. We refer to strings that do *not* coerce to integers as “safe,” and we use an uninterpreted predicate *safe* to describe such strings (e.g., *safe*(“foo”) whereas \neg *safe*(“17”).

- (A1) a contains the special “length” key.
- (A2) All other “own” keys of a are (strings that coerce to) integers.
- (A3) For all integers i , either a maps the key i to a value of type T , or it has no binding for i .
- (A4) All inherited keys of a are safe (i.e., non-integer) strings.

An array can have arbitrary objects in its prototype chain, so to ensure (A4), we require that *all* non-array objects bind only safe strings. This sharp distinction between array objects (that bind integer keys) and non-array objects (that bind safe string keys) allows System !D to avoid reasoning about string coercions, and does not significantly limit expressiveness because, in our experience, programs typically conform to this division anyway. To enforce this restriction, the type for keys manipulated by primitives in `objects.dref` and `prelude.js` is actually *SafeStr*, rather than *Str* as shown in Figure 5 and Figure 6, where $SafeStr \stackrel{\circ}{=} \{Str(\nu) \wedge safe(\nu)\}$. We discuss an alternative approach in §7 that allows non-array objects to bind unsafe strings.

Packed Arrays. Arrays a that additionally satisfy the uninterpreted predicate $packed(a)$ enjoy the following property, where $len(a)$ is an uninterpreted function symbol.

- (A5) For all integers i , if i is between zero and $len(a)$ minus one, then a maps i to a value of type T . Otherwise, a has no binding for i .

Tuple Arrays. Using additional predicates, System !D gives precise types to array literals, which are often used as finite tuples in idiomatic code. For example, we can describe pairs as follows:

$$\begin{aligned} (Int, Int) &\stackrel{\circ}{=} \{\nu :: Arr(Int) \wedge packed(\nu) \wedge len(\nu) = 2\} \\ (Bool, Str) &\stackrel{\circ}{=} \{\nu :: Arr(Top) \wedge packed(\nu) \wedge len(\nu) = 2 \\ &\quad \wedge Str(sel(\nu, 0)) \wedge Bool(sel(\nu, 1))\} \end{aligned}$$

Thus, the technique of nested refinements allows us to smoothly reason about arrays both as packed homogeneous collections as well as *heterogeneous* tuples.

Array Primitives. We define several array-manipulating primitives in `objects.dref` (some of which we show in Figure 5) that maintain and use the array invariants above. For key lookup on arrays, we define three primitives: `getIdxArr` looks for the integer key i on the own object a and ignores the prototype chain of a because (A4) guarantees that a will not inherit i , and returns a value subject to the properties (A3) and (A5) that govern its integer key bindings; `getLenArr` handles the special case when the string key k is “length”, which (A1) guarantees is bound by a , and returns the true length of the array only if it is packed; and `getPropArr` deals with all other (safe) string keys k by reading from the prototype chain of the array (re-using the heap unrolling mechanism) ignoring its own bindings because of (A2).

For array updates, we define `setIdxArr` that uses the following macros to preserve packedness (A5) when possible.

$$\begin{aligned} arrSet(a', a, i) &\stackrel{\circ}{=} \text{if } 0 \leq i < len(a) \text{ then } arrSize(a', a, 0) \text{ else} \\ &\quad \text{if } i = len(a) \text{ then } arrSize(a', a, 1) \text{ else } true \\ arrSize(a', a, n) &\stackrel{\circ}{=} packed(a) \Rightarrow \\ &\quad (packed(a') \wedge len(a') = len(a) + n) \end{aligned}$$

In particular, the updated array a' is packed if: (1) the original array a is packed; and (2) the updated index i is either within the bounds of a (in which case, the length of a' is the same as a) or just past the end (so the length of a' is one greater than a). In similar fashion, we specify the remaining primitives for update and deletion to maintain the array invariants, and the ones for key membership to use them, but we do not show them in Figure 5.

In `prelude.js` (Figure 6), we use precise types to model the native `push` and `pop` methods of `Array.prototype` (which maintain packedness, as above), as well as the behavior of `Object.prototype.hasOwnProperty` on arrays (the last two cases of the intersection type). Thus, the precise dependent types we ascribe to array-manipulating operations maintain invariants (A1) through (A5) and allow DJS to precisely track array operations.

Desugaring. It may seem that we need to use separate primitive functions for array and non-array object operations, even though they are syntactically indistinguishable in JavaScript. Nevertheless, we are able to desugar DJS based purely on expression syntax (and *not* type information) by unifying key lookup within a single primitive `getElem` and giving it a type that is the *intersection* of the (three) array lookup primitives and the (one) non-array lookup primitive `getPropObj`. We define `getElem` in Figure 5, where we specify the intersection type using `and` and `type` as syntactic sugar to refer to the previous type annotations. We define similar unified primitives for `setElem`, `hasElem`, and `delElem` (not shown in Figure 5). Desugaring uniformly translates object operations to these unified general primitives, and type checking of function calls ensures that the appropriate cases of the intersection type apply.

2.7 Collections

As discussed in §2.2, strong updates are sound only for references that point to *exactly one* object, which is far too restrictive as real programs manipulate collections of objects. In this section, we describe *weak references* in DJS to refer to multiple objects, a facility that enables programming with arrays of mutable objects as well as recursive types.

Weak References. In the following example, we iterate over an array of passenger objects and compute the sum of their weights; we use a default value `max_weight` when a passenger does not list his weight (ignore the unfamiliar annotations for now).

```

1  /*: (ℓpass ↦ frzn) → same */
2  for (i=0; i < passengers.length; i++) {
3    var p = passengers[i];
4    /*: #thaw p */
5    if (p.weight) { sum += p.weight; }
6    else { sum += max_weight; }
7    /*: #freeze p */
8  }

```

We could describe the array `passengers` with the type $Ref\ell$ for a location ℓ . However, this type is not very useful as it denotes an array of references to a *single* object.

Weak Locations. To refer to an *arbitrary number* (zero or more) objects of the same type, we adopt the Alias Types [32] solution, which categorizes some locations as *weak* to describe an arbitrary number of locations that satisfy the same type, and syntactically ensures that weak locations are weakly updated.

We introduce a new kind of heap binding ($\tilde{\ell} \mapsto \langle T, \ell' \rangle$), where $\tilde{\ell}$ is a weak location, all objects that might reside there satisfy T , and ℓ' is the strong location of the prototype of all objects that reside at location $\tilde{\ell}$. There is *no* heap binder for weak locations since there is not a single value to describe.

In our example, we can use ($\tilde{\ell}_{pass} \mapsto \langle T_{pass}, \ell_{op} \rangle$) to describe passenger objects, where ℓ_{op} is the location of `Object.prototype` and T_{pass} is the dictionary type $\{Dict(\nu) \wedge has(\nu, \text{"weight"}) \Rightarrow Num(sel(\nu, \text{"weight"}))\}$. If we assign the type $\{\nu :: Arr(Ref\tilde{\ell}_{pass}) \wedge packed(\nu)\}$, to `passengers`, then `p` has type $Ref\tilde{\ell}_{pass}$, and thus each (desugared) use of `p` is a dictionary of type T_{pass} . This type is quite unsatisfying, however, because the conditional establishes that along the then-branch, `p` *does* possess the key and therefore should be assigned the more precise type $\{Num(sel(d, \text{"weight"}))\}$.

Thaw and Freeze. To solve this problem, we adopt a mechanism found in derivatives of Alias Types (e.g., [2, 12, 15, 31]) that allows a weak location to be *temporarily* treated as strong. A weak location $\tilde{\ell}$ is said to be *frozen* if all references $Ref\tilde{\ell}$ use the location only at its weak (invariant) type. The type system can *thaw* a location, producing a strong reference $Ref\ell_k$ (with a fresh name) that can be used

to strongly update the type of the cell. While a location is thawed, the type system prohibits the use of weak references to the location, and does not allow further thaw operations. When the thawed (strong) reference is no longer needed, the type system checks that the original type has been restored, *re-freezes* the location, and discards the thawed location. Soundness of the approach depends on the invariant that each weak location has at most one corresponding thawed location at a time.

In our example, we do not need to temporarily violate the type of `p`, but the thaw/freeze mechanism does help us *relate* the two accesses to `p`. The *thaw state* annotation above the loop declares that before each iteration of the loop (including the first one), the location $\tilde{\ell}_{pass}$ must be frozen. The thaw annotation on line 4 changes the type of `p` to a strong reference to a fresh thawed location ℓ_1 , which stores a *particular* dictionary on the heap (named with a binder) that is retrieved by both subsequent uses of `p`. Thus, we can relate the key membership test to the lookup, and track that `p.weight` produces a number. The freeze annotation on line 7 restores the invariant required before the next iteration. We describe this technique further in §4.

Recursive Types. We reuse the weak location mechanism to describe recursive data structures. Consider the following adapted from the SunSpider [34] benchmark `access-binary-trees.js`, annotated in DJS.

```

1  /*: #define Ttn {"i":Num, "l","r":Refℓtn} */
2  /*: #weak (ℓtn ↦ ⟨Ttn, ℓtnp⟩) */
3  /*: #ctor (this:Ref, left:right:Refℓtn, item:Num)
4     / (ℓtn ↦ frzn) → Refℓtn/same */
5  function TreeNode(left, right, item) {
6    this.l = left; this.r = right; this.i = item;
7    /*: #freeze this */
8    return this;
9  }
10 /*: this:Refℓtn → Num */
11 TreeNode.prototype.itemCheck = function f() {
12   // thaw/freeze annotations inferred
13   if (this.l == null) return this.item;
14   else { return this.i
15         + f.apply(this.l)
16         - f.apply(this.r); }
17 }

```

The source-level macro on line 1 introduces T_{tn} to abbreviate the type of `TreeNodes`, using traditional record type syntax instead of the underlying McCarthy operators. Line 2 defines the weak location for `TreeNodes`, using the predictable location ℓ_{tnp} created by desugaring for the object `TreeNode.prototype`. The constructor annotation itself declares that the return type is a reference to one of these recursive objects, which System !D verifies by checking that on line 6 the appropriate fields are added to the strong, initially-empty object `this` before it is frozen and returned.

Recursive Traversal. There are two differences in the `itemCheck` function above compared to the original version, which cannot be type checked in DJS. First, we name the function being defined (notice the `f` on line 11), a JavaScript facility for recursive definitions. Second, we write `f.apply(this.r)` instead of `this.r.itemCheck()` as in the original, where the native JavaScript function `apply` allows a caller to explicitly supply a receiver argument. The trouble with the original call is that it goes through the heap (in particular, the prototype chain of `this`) to resolve the recursive function being defined. This function will be stored in a strong object, and we have no facility (e.g., mu-types) for strong objects with recursive types; our only mechanism is for weak objects. If we write `f.apply(this.r)`, however, the recursive function `f` is syntactically manifest, and we can translate the definition with a call to the standard `fix` primitive (Figure 2). In §5, we describe how we handle a limited form of `apply` that is sufficient for our idiomatic recursive definitions in DJS. We expect that we can add a more powerful mechanism for recursive types that supports the original code as written, but we leave this to future work.

2.8 Rest of the Paper

We have now completed our tour of Dependent JavaScript. Next, we formally define the syntax of System !D in §3 and the type system in §4. In §5, we present the syntax of DJS and its desugaring to System !D. We discuss our implementation and results in §6, directions for future work in §7, and related work in §8. Additional details may be found in an accompanying technical report [7].

3. Syntax and Semantics of System !D

We now introduce the formal syntax of values, expressions, and types of System !D, defined in Figure 7.

Values. Values v include variables x , constants c , lambdas $\lambda x. e$, (functional) dictionaries $v_1 ++ v_2 \mapsto v_3$, and run-time heap locations r . The set of constants c includes base values (numbers, booleans, strings, the empty dictionary `{}`, `null`, `undefined`, `NaN`, etc.) and the primitive functions from `basics.dref` and `objects.dref` (`typeof`, `get`, `getElem`, etc.). We use tuple syntax (v_0, \dots, v_n) as sugar for the dictionary with fields “0” through “n” bound to the component values. Logical values w are all values and applications of primitive function symbols F , such as addition `+` and dictionary selection `sel`, to logical values.

Expressions. We use an A-normal expression syntax so that we need only define substitution of values (not arbitrary expressions) into types. We use a more general syntax for examples throughout this paper, and our implementation desugars expressions into A-normal form. Expressions e include values, function application, if-expressions, and let-bindings. The ascription form $e \text{ as } T$ allows source-level type annotations. Since function types will be parameterized

by type, location, and heap variables, the syntax of function application requires that these be instantiated. Reference operations include reference allocation, dereference, and update, and the run-time semantics maintains a separate heap that maps locations to values. The expression `newobj $\ell v v'$` stores the value v at a fresh location r — where the name ℓ is a compile-time abstraction of a set of run-time location names that includes r — with its prototype link set to v' , which should be a location. The `thaw ℓv` operation converts a weak reference to a strong one; `freeze $\tilde{\ell} \theta v$` converts a strong reference to a weak one, where `thaw` state θ is used by the type system for bookkeeping.

The operational semantics is standard, based on λ_{JS} with minor differences. For example, we make prototype links manifest in the syntax of heaps (to facilitate heap unrolling in the type system), whereas λ_{JS} stores them inside objects in a distinguished “`_proto_`” field. We refer the reader to [21] for the full details.

Types and Formulas. Values in System !D are described by *refinement types* of the form $\{x|p\}$ where x may appear free in the formula p and *existential types* $\exists x:T. S$ where x may appear free in S . Rather than introduce additional syntactic categories, we assume, by convention, that existential types do not appear in source programs; they are created only during type checking in a controlled fashion that does not preclude algorithmic type checking [25]. When the choice of refinement binder does not matter, we write $\{p\}$ as shorthand for $\{v|p\}$.

The language of *refinement formulas* includes predicates P , such as equality and the dictionary predicate *has*, and the usual logical connectives. Similar to the syntax for expression tuples, we use (T_1, \dots, T_n) as sugar for the dictionary type with fields “0” through “n” with the corresponding types. As in System D, we use an uninterpreted *has-type predicate* $w :: U$ in formulas to describe values that have complex types, represented by *type terms* U , which includes function types, type variables, null, reference, and array types. A reference type names a strong or weak location in the heap, where a strong location ℓ is either a constant a or a variable L and a weak location $\tilde{\ell}$ is a constant \tilde{a} . We discussed function types in §2.4; now we use the metavariable W to range over *worlds*. A world $x:T/\hat{\Sigma}$ describes the binders and types of a tuple of values, where every component except the first (x of type T) resides in the heap $\hat{\Sigma}$.

Heap Types. A *heap type* $\hat{\Sigma}$ is an unordered set of *heap variables* H and *heap bindings* \hat{h} concatenated with the \oplus operator. To simplify the presentation, we syntactically require that each heap has exactly one heap variable, so we write a heap type as the pair (H, \hat{h}) , where H is the “deep” part for which we have no information and \hat{h} is the “shallow” part for which we have precise location information. The heap binding $(\ell \mapsto x:T)$ represents the fact that the value at location ℓ has type T ; the binder x refers to this value in the types of other heap bindings. The binding $(\ell \mapsto \langle x:T, \ell' \rangle)$

Values	$v ::= x \mid c \mid v_1 ++ v_2 \mapsto v_3 \mid \lambda x. e \mid r$		
Expressions	$e ::= v \mid [\overline{T}; \overline{\ell}; \overline{\Sigma}] v_1 v_2 \mid \text{if } v \text{ then } e_1 \text{ else } e_2 \mid \text{let } x = e_1 \text{ in } e_2 \mid e \text{ as } T$ $\mid \text{ref } \ell v \mid \text{deref } v \mid v_1 := v_2 \mid \text{newobj } \ell v v' \mid \text{freeze } \tilde{\ell} \theta v \mid \text{thaw } \ell v$		
Types	$S, T ::= \{x \mid p\} \mid \exists x:T. S$		
Formulas	$p, q ::= P(\overline{w}) \mid w :: U \mid \text{HeapHas}(H, \ell, w) \mid p \wedge q \mid p \vee q \mid \neg p$		
Logical Values	$w ::= v \mid F(\overline{w}) \mid \text{HeapSel}(H, \ell, w)$		
Syntactic Type Terms	$U ::= \forall[\overline{A}; \overline{L}; \overline{H}] W_1 \rightarrow W_2 \mid A \mid \text{Null} \mid \text{Ref } \ell \mid \text{Ref } \tilde{\ell} \mid \text{Arr}(T)$		
Heap Bindings	$\hat{h} ::= (\ell \mapsto x:T) \mid (\ell \mapsto \langle x:T, \ell' \rangle) \mid (\tilde{\ell} \mapsto \theta) \mid \hat{h}_1 \oplus \hat{h}_2 \mid \emptyset$		
Thaw States	$\theta ::= \text{frzn} \mid \text{thwd } \ell$		
Worlds	$W ::= x:T/\hat{\Sigma} \quad x, y, z \in \text{Identifiers} \quad A, B \in \text{TypeVariables}$		
Heaps	$\hat{\Sigma} ::= (H, \hat{h}) \quad a \in \text{LocationConstants} \quad H \in \text{HeapVariables}$		
Strong Locations	$\ell ::= a \mid L \quad r \in \text{DynamicHeapLocations} \quad L \in \text{LocationVariables}$		
Weak Locations	$\tilde{\ell} ::= \tilde{a} \quad F \in \text{LogicalFunctionSymbols} \quad P \in \text{LogicalPredicateSymbols}$		
$c \in \text{ValueConstants} \supset \{ \text{true}, \text{false}, \text{null}, \text{undefined}, 1, 2, \text{"hanna"}, (=), !, \text{typeof}, \text{get}, \text{getElem}, \text{fix} \}$			

Figure 7. Syntax of System !D

additionally records a prototype link ℓ' . The binding $(\tilde{\ell} \mapsto \theta)$ records the current *thaw state* of weak location $\tilde{\ell}$, to help maintain the invariant that it has at most one thawed location at a time. We abuse the notation $(H, \hat{h}_1) \oplus \hat{h}_2$ to mean $(H, \hat{h}_1 \oplus \hat{h}_2)$.

Uninterpreted Heap Symbols. To describe invariants about the deep part of a heap, System !D introduces two uninterpreted *heap symbols*. The predicate $\text{HeapHas}(H, \ell, k)$ represents the fact that the chain of objects in H starting with ℓ has the key k . Similarly, the function symbol $\text{HeapSel}(H, \ell, k)$ refers to the value retrieved when looking up key k in the heap H starting with ℓ .

4. Type Checking

In this section, we discuss the well-formedness, typing, and subtyping relations of System !D. The type system reuses the System D [9] subtyping algorithm to factor subtyping obligations between a first-order SMT solver and syntactic subtyping rules. The novel technical developments here are: the formulation of flow-sensitive heap types in a dependent setting; the use of uninterpreted heap symbols to regain precision in the presence of imperative, prototype-based objects; the encoding of array primitives to support idiomatic use of JavaScript arrays; and the use of refinement types to assign precise types to JavaScript operators.

Environments. The type checking relations make use of type environments Γ and heap environments Σ .

$$\begin{aligned} \Gamma &::= \emptyset \mid \Gamma, x:T \mid \Gamma, p \mid \Gamma, A \mid \Gamma, L \mid \Gamma, H \mid \Gamma, (\tilde{\ell} \mapsto \langle T, \ell \rangle) \\ \Sigma &::= (H, h) \\ h &::= \emptyset \mid h_1 \oplus h_2 \mid (\ell \mapsto v) \mid (\ell \mapsto \langle v, \ell' \rangle) \mid (\tilde{\ell} \mapsto \theta) \end{aligned}$$

A type environment binding records either: the derived type for a variable; a formula p to track control flow along a conditional branch; a polymorphic variable introduced by a function type; or the description of a weak location (which does not change flow-sensitively), namely, that every object stored at $\tilde{\ell}$ satisfies type T and has prototype link ℓ . A heap environment is just like a heap type, except a strong location ℓ binds the value v it stores (as opposed to the *type* of v).

4.1 Well-Formedness

As usual in a refinement type system, we define well-formedness relations (see [7]) that govern how values may be used inside formulas. The key intuition is that formulas are boolean propositions and mention only variables that are currently in scope. Locations in a heap type $\hat{\Sigma}$ must either be location constants or location variables bound by the type environment, and may not be bound multiple times. All heap binders may refer to each other. Thus, the values in a heap can be regarded as a dependent tuple. For the input world $x_1:T_1/\hat{\Sigma}_1$ of a function type, the binder x_1 and the binders in $\hat{\Sigma}_1$ may appear in the output world W_2 .

4.2 Subtyping

Several relations (see Figure 8 and [7]) comprise subtyping.

Subtyping and Implication. As in System D, subtyping on refinement types reduces to implication of refinement formulas, which is discharged by a combination of uninterpreted, first-order reasoning and syntactic subtyping. Our treatment of existential types follows the algorithmic (decidable) approach in [25]. In particular, when on the left side of an obligation, the S-EXISTS rule adds the existential binding to the environment; there is *no* support for existentials on

Subtyping

$$\begin{array}{c}
\boxed{\Gamma \vdash T_1 \sqsubseteq T_2} \\
\text{[S-REFINE]} \quad \frac{y \text{ fresh} \quad \Gamma, p[y/x] \Rightarrow q[y/x]}{\Gamma \vdash \{x|p\} \sqsubseteq \{x|q\}} \quad \text{[S-EXISTS]} \quad \frac{\Gamma, x:T \vdash S_1 \sqsubseteq S_2}{\Gamma \vdash \exists x:T. S_1 \sqsubseteq S_2} \\
\text{[U-ARRAY]} \quad \frac{}{\Gamma \vdash \text{Arr}(T) <: \text{Arr}(T)} \quad \text{[U-VAR]} \quad \frac{}{\Gamma \vdash A <: A} \quad \text{[U-STRONGREF]} \quad \frac{}{\Gamma \vdash \text{Ref } \ell <: \text{Ref } \ell} \\
\text{[U-NULL]} \quad \frac{}{\Gamma \vdash \text{Null} <: \text{Ref } \bar{\ell}} \quad \text{[U-WEAKREF]} \quad \frac{}{\Gamma \vdash \text{Ref } \bar{\ell} <: \text{Ref } \bar{\ell}} \\
\boxed{\Gamma \vdash U_1 <: U_2}
\end{array}$$

Figure 8. Subtyping for System !D

Value Typing (selected rules)

$$\begin{array}{c}
\boxed{\Gamma; \Sigma \vdash v :: T} \\
\text{[T-CONST]} \quad \frac{}{\Gamma; \Sigma \vdash c :: ty(c)} \quad \text{[T-VAR]} \quad \frac{\Gamma(x) = S}{\Gamma; \Sigma \vdash x :: \{y|y = x\}} \\
\text{[T-FUN]} \quad \frac{U = \forall[\bar{A}; \bar{L}; \bar{H}] x:T_1/\hat{\Sigma}_1 \rightarrow W_2 \quad \Gamma \vdash U \quad \text{HeapEnv}(\hat{\Sigma}_1) = (\bar{z}:\bar{S}, \Sigma_1) \quad \Gamma_1 = \Gamma, \bar{A}, \bar{L}, \bar{H}, x:T_1, \bar{z}:\bar{S} \quad \Gamma_1; \Sigma_1 \vdash e :: T_2/\Sigma_2 \quad \Gamma_1 \vdash T_2/\Sigma_2 \models W_2}{\Gamma; \Sigma \vdash \lambda x. e :: \{y|y :: U\}}
\end{array}$$

Figure 9. Value type checking for System !D

the right. The way in which type checking introduces existentials guarantees that they always appear on the left.

References and Arrays. As in Alias Types [32], we enforce the invariant that distinct strong locations do not alias, so references to them are *never* related by subtyping. In contrast, weak locations describe zero or more locations, and it is safe to treat null as a subtype of *any* weak location (U-NULL). That is, weak references are *nullable* but strong ones are not. Arrays are invariant in their type parameter (U-ARRAY), as usual, but can be related with additional predicates. For example, $\{\nu :: \text{Arr}(\text{Int}) \wedge \text{len}(\nu) = 2\}$ is a subtype of $\{\nu :: \text{Arr}(\text{Int})\}$.

Heaps. The heap subtyping judgement (defined in [7]) relates two heap types (H_1, \hat{h}_1) and (H_2, \hat{h}_2) if: (1) the heaps agree on the “deep” part, that is, if $H_1 = H_2$; (2) the structure of the “shallow” parts \hat{h}_1 and \hat{h}_2 match modulo permutation; and (3) the heap bindings in \hat{h}_1 and \hat{h}_2 , which can be thought of as dependent tuples, are related by subtyping.

4.3 Value Typing

The value typing judgement $\Gamma; \Sigma \vdash v :: T$ (defined in Figure 9 and [7]) verifies that the value v has type T in the given environments. Since values do not produce any effects, this judgement does *not* produce an output heap environment. Each primitive constant c has a type, denoted by $ty(c)$, that is used by T-CONST. In our implementation,

$ty(c)$ is defined in the standard prelude files (`basics.dref`, `objects.dref`, and `prelude.dref`). The standard T-VAR rule assigns *singleton* types to variables. The rule T-FUN uses the procedure `HeapEnv` that takes a “snapshot” of the input heap type $\hat{\Sigma}_1$ by collecting all of its binders $\bar{z}:\bar{S}$ to add to the type environment and producing a heap environment Σ_1 for type checking the body. Dually, the world satisfaction judgement $\Gamma_1 \vdash T_2/\Sigma_2 \models W_2$ (defined in [7]) checks that the resulting type and heap environment T_2/Σ_2 satisfies W_2 , modulo permutation of heap bindings.

4.4 Expression Typing

The expression typing judgement $\Gamma; \Sigma \vdash e :: T/\Sigma'$ (in Figure 10 and [7]) verifies that the evaluation of expression e produces a value of type T and a new heap environment Σ' . We write U as shorthand for the type $\{\nu :: U\}$, and $\Sigma \equiv \Sigma'$ for heap equality modulo permutation of bindings.

Prenex Quantified Types. The T-LET rule uses an existential to describe the type T_1 of the variable x that goes out of scope after the body expression is checked. Alternatively, the more traditional approach (*e.g.*, [9]) requires that the variable be eliminated (*e.g.*, via subsumption), but we use existentials because it simplifies several other typing rules.

So that existentials appear only on the left side of subtyping obligations, we ensure that the typing rules derive *prenex quantified types* of the form $\exists \bar{x}:\bar{T}. S$, where all the types \bar{T} and S are refinement types, not existential types. In particular, to combine the worlds of two branches, the Join operator (defined in [7]), rearranges existentials to ensure that the resulting world is in prenex form. For example, for a conditional with guard b , the join of $(\exists x_1:T_1. \text{Top}/(\ell \mapsto x_1))$ and $(\exists x_2:T_2. \text{Top}/(\ell \mapsto x_2))$ is $(\exists y:T_{12}. \text{Top}/(\ell \mapsto y))$ where $T_{12} \stackrel{\circ}{=} \{\text{if } b \text{ then } T_1(\nu) \text{ else } T_2(\nu)\}$.

Imperative Operations. We use two kinds of reference cells in System !D: *simple references* that store base values and functions without prototype links, and *object references* that store dictionaries paired with prototype links. We require that *all* imperative operations go through strong locations. We do not need weak, simple locations since they cannot appear in DJS programs and they are not needed for desugaring; and we do not need weak, object locations because we can use our thawing mechanism instead.

Three rules manipulate simple references. To check the reference allocation `ref ℓ v`, the rule T-REF ensures that ℓ is not already bound in the heap, and then adds a binding that records exactly the value being stored. The rule T-DEREF checks that the given value is a reference to a simple location, and then retrieves the stored value; this is the imperative analog to the “selfifying” T-VAR rule. The rule T-SETREF strongly updates a simple location.

The rule T-NEWOBJ stores the dictionary v_1 in the heap at location ℓ_1 *along with* a prototype link to the location ℓ_2 that v_2 refers to. Although no typing rules manipulate object locations, several primitives (`getElem`, `setElem`, *etc.*) do.

$$\begin{array}{c}
 \text{[T-VAL]} \quad \frac{\Gamma; \Sigma \vdash v :: T}{\Gamma; \Sigma \vdash v :: T/\Sigma} \quad \text{[T-LET]} \quad \frac{\Gamma; \Sigma \vdash e_1 :: T_1/\Sigma_1 \quad \Gamma, x:T_1; \Sigma_1 \vdash e_2 :: T_2/\Sigma_2}{\Gamma; \Sigma \vdash \text{let } x = e_1 \text{ in } e_2 :: \exists x:T_1. T_2/\Sigma_2} \\
 \text{[T-IF]} \quad \frac{\Gamma; \Sigma \vdash v :: S \quad \Gamma, \text{trathy}(v); \Sigma \vdash e_1 :: T_1/\Sigma_1 \quad \Gamma, \text{falsy}(v); \Sigma \vdash e_2 :: T_2/\Sigma_2 \quad T/\Sigma' = \text{Join}(v, T_1/\Sigma_1, T_2/\Sigma_2)}{\Gamma; \Sigma \vdash \text{if } v \text{ then } e_1 \text{ else } e_2 :: T/\Sigma'} \\
 \text{[T-REF]} \quad \frac{\ell \notin \text{dom}(\Sigma) \quad \Gamma; \Sigma \vdash v :: T \quad \Sigma' = \Sigma \oplus (\ell \mapsto v)}{\Gamma; \Sigma \vdash \text{ref } \ell v :: \text{Ref } \ell/\Sigma'} \quad \text{[T-DEREF]} \quad \frac{\Gamma; \Sigma \vdash v :: \text{Ref } \ell \quad \Sigma \equiv \Sigma_0 \oplus (\ell \mapsto v')}{\Gamma; \Sigma \vdash \text{deref } v :: \{y | y = v'\}/\Sigma} \\
 \text{[T-SETREF]} \quad \frac{\Gamma; \Sigma \vdash (v_1, v_2) :: (\text{Ref } \ell, T) \quad \Sigma \equiv \Sigma_0 \oplus (\ell \mapsto v) \quad \Sigma' = \Sigma_0 \oplus (\ell \mapsto v_2)}{\Gamma; \Sigma \vdash v_1 := v_2 :: \{x | x = v_2\}/\Sigma'} \quad \text{[T-NEWOBJ]} \quad \frac{\ell_1 \notin \text{dom}(\Sigma) \quad \Gamma; \Sigma \vdash (v_1, v_2) :: (\text{Dict}, \text{Ref } \ell_2) \quad \Sigma \equiv \Sigma_0 \oplus (\ell_2 \mapsto \langle v', \ell_3 \rangle) \quad \Sigma' = \Sigma \oplus (\ell_1 \mapsto \langle v_1, \ell_2 \rangle)}{\Gamma; \Sigma \vdash \text{newobj } \ell_1 v_1 v_2 :: \text{Ref } \ell_1/\Sigma'} \\
 \text{[T-APP]} \quad \frac{\Gamma; \Sigma \vdash v_1 :: \forall [\bar{A}; \bar{L}; \bar{H}] W_1 \rightarrow W_2 \quad \Gamma; \Sigma \vdash v_2 :: T_2 \quad \Gamma \vdash [\bar{T}/\bar{A}] \quad \Gamma \vdash [\bar{m}/\bar{M}] \quad \Gamma \vdash [\bar{\Sigma}/\bar{H}] \quad W'_2 = \text{Freshen}(W_2) \quad (W'_1, W'_2) = \text{Unroll}(\text{HInst}(\text{LInst}(\text{TInst}((W_1, W'_2), \bar{A}, \bar{T}), \bar{L}, \bar{\ell}), \bar{H}, \bar{\Sigma}))) \quad \Gamma \vdash T_2/\Sigma \vDash W'_1; \pi \quad W'_1 = x:T_{11}/\hat{\Sigma}_{11} \quad \pi' = \pi[v_2/x] \quad \pi' W'_2 = x':T_{12}/\hat{\Sigma}_{12} \quad \text{HeapEnv}(\hat{\Sigma}_{12}) = (\bar{y}:\bar{S}, \Sigma_{12})}{\Gamma; \Sigma \vdash [\bar{T}; \bar{\ell}; \bar{\Sigma}] v_1 v_2 :: \exists x':T_{12}. \exists \bar{y}:\bar{S}. \{z | z = x'\}/\Sigma_{12}}
 \end{array}$$

Figure 10. Expression type checking for System !D

$$\begin{array}{l}
 \text{Unroll}(\text{HeapHas}((H, \hat{h}), \ell, k)) = \text{UnrollHas}(H, \hat{h}, \ell, k) \\
 \text{UnrollHas}(H, \hat{h}, \ell, k) = \\
 \begin{cases} \text{has}(d, k) \vee \text{UnrollHas}(H, \hat{h}, \ell', k) & \text{if } (\ell \mapsto \langle d:T, \ell' \rangle) \in \hat{h} \\ \text{HeapHas}(H, \ell, k) & \text{else if } \ell \neq \circ \\ \text{false} & \text{else (i.e., } \ell = \circ) \end{cases} \\
 \text{Unroll}(\psi(\text{HeapSel}((H, \hat{h}), \ell, k))) = \text{UnrollSel}(\psi, H, \hat{h}, \ell, k) \\
 \text{UnrollSel}(\psi, H, \hat{h}, \ell, k) = \\
 \begin{cases} \text{ite } \text{has}(d, k) \psi(\text{sel}(d, k)) & (\text{UnrollSel}(\psi, H, \hat{h}, \ell', k)) \\ & \text{if } (\ell \mapsto \langle d:T, \ell' \rangle) \in \hat{h} \\ \psi(\text{HeapSel}(H, \ell, k)) & \text{else if } \ell \neq \circ \\ \psi(\text{undefined}) & \text{else (i.e., } \ell = \circ) \end{cases}
 \end{array}$$

Figure 11. Heap unrolling

Function Application. To type check $[\bar{T}; \bar{m}; \bar{\Sigma}] v_1 v_2$, the T-APP rule must perform some heavy lifting. Three well-formedness checks ensure that the number of type, location, and heap parameters must match the number of type, location, and heap variables of the function type, and that the sequence of locations $\bar{\ell}$ contains no duplicates to ensure the soundness of strong updates [32]. The procedure Freshen generates fresh binders for the output world so that bindings at different call sites do not collide.

The substitution of parameters for polymorphic variables proceeds in three steps. First, the type variables \bar{A} inside has-type predicates are instantiated with the type parameters \bar{T} using the procedure TInst. Second, the location variables \bar{L} are replaced with the parameters $\bar{\ell}$ by ordinary substitution. Third, the heap variables \bar{H} are instantiated with heap parameters using a procedure HInst that substitutes heap bind-

ings for heap variables. As a result, *HeapHas* and *HeapSel* may refer to arbitrary heaps rather than just heap variables, as required. These *pre-types* (and *pre-formulas*, *pre-heaps*, *etc.*) are expanded using the procedure Unroll, defined in Figure 11, that transitively follows prototype links in heap bindings, precisely matching the semantics of object key membership and lookup. We write the location \circ for the root of the prototype hierarchy. We use the notation $\psi(p)$ to refer to a *formula context* ψ , a formula with a hole, filled with p .

At this point, the polymorphic variables have been fully instantiated. Next, the argument type T_2 and current heap Σ are checked to satisfy the input world $x:T_{11}/\hat{\Sigma}_{11}$. If so, the substitution π maps binders from the input heap $\hat{\Sigma}_{11}$ to the corresponding ones in the current heap Σ . The substitution is extended with a binding from x to the argument v_2 and applied to the output world. Then, like in the T-FUN rule, we use HeapEnv to collect the bindings $\bar{y}:\bar{S}$ in $\hat{\Sigma}_{12}$ and convert it to a heap environment Σ_{12} . Finally, the derived type uses existentials to describe the values in the output world.

Thaw and Freeze. To safely allow a weak, object location $\tilde{\ell}$ to be treated *temporarily* as strong, System !D ensures that $\tilde{\ell}$ has at most one corresponding *thawed* location ℓ at a time (if there is none, we say $\tilde{\ell}$ is *frozen*) by recording its *thaw state* — either `thwd` ℓ or `frzn`. One interesting aspect of our formulation is that, to facilitate idiomatic programming, we choose to allow the use of possibly-null references. Instead, to require that all references be provably non-null before use, we can simply update the type signatures for object primitives. We omit the typing rules for thawing and freezing from Figure 10; we refer the reader to [7] for more details.

Location Polymorphism. To simplify the presentation of System !D in this paper, we have limited location polymor-

phism in two ways. First, we allow location variables L to refer only to *strong* locations. In [7], we describe how to add *weak location polymorphism* to function types and update the typing rules appropriately. Second, we offer only a single mechanism — namely, universal quantification — to abstract over both simple locations as well as object locations. As a result, functions must be quantified over all simple locations inserted by desugaring (to model imperative JavaScript variables), which clutters function types and, worse, requires explicit declaration and instantiation of locations that are “internal” to the desugaring translation and *not* accessible in the original DJS program. Instead, in [7], we show how to use existential quantification in the output types of functions to describe simple locations, and use universal quantification in the input types of functions only to describe object locations, which *are* visible in DJS.

4.5 Type Soundness

Many standard *stuck* states are *not* stuck in JavaScript: a function can be applied with any number of arguments; an operator can be used with any values because of implicit coercion; and, all property lookups succeed (possibly producing `undefined`). Nonetheless, several (non-exceptional) stuck states remain: applying a non-function value; and retrieving a property for a non-object value. System !D is designed to ensure that well-typed programs do not get stuck and can only fail with exceptions due to retrieving a property from `undefined` or `null`. We can also provide the stronger guarantees that only bound keys are retrieved and only non-null objects are accessed (thus ruling out the possibility of null dereference exceptions) simply by changing the types of object primitives appropriately.

We expect that System !D satisfies progress and preservation theorems, but we have not yet proven them. The process will likely be tedious but not require new proof techniques. Unlike System D, which introduced the problematic *nesting* of syntactic types inside uninterpreted formulas, System !D does not introduce any new mechanisms in the refinement logic. Furthermore, several variations of Alias Types [23, 32, 37], even in a dependent setting [31], have been proven sound, and we expect to re-use their techniques to prove the soundness of System !D.

5. Desugaring DJS to System !D

In Figure 12, we present a selection of the explicitly typed abstract syntax of DJS along with desugaring rules $\llbracket e \rrbracket = e$ that translate DJS expressions e to System !D expressions e . Most of the desugaring rules follow λ_{JS} [21] closely, so we limit our discussion to the aspects most relevant to DJS; we refer the reader to [7] and their work for more details. We use the metavariable $I \triangleq [\bar{T}; \bar{m}; \bar{\Sigma}]$ to range over instantiation parameters for function application. Instantiation parameters are usually inferred by the type checker (§6).

Desugaring (selected rules)

$\llbracket e \rrbracket = e$	$\llbracket e \rrbracket = e$
$\llbracket c \rrbracket = c$	[DS-CONST]
$\llbracket x \rrbracket = \text{deref } x$	[DS-DEREF]
$\llbracket e_1 = e_2 \rrbracket = \llbracket e_1 \rrbracket := \llbracket e_2 \rrbracket$	[DS-SETREF]
$\llbracket \text{var } x = e; e' \rrbracket = \text{let } x = \text{ref } a_x \llbracket e \rrbracket \text{ in } \llbracket e' \rrbracket$	[DS-REF]
$\llbracket /* : I * / e_1[e_2] \rrbracket =$ $/* : I * / \text{getElem} (\llbracket e_1 \rrbracket, \llbracket e_2 \rrbracket)$	[DS-GETELEM]
$\llbracket /* : I * / e_1[e_2] = e_3 \rrbracket =$ $/* : I * / \text{setElem} (\llbracket e_1 \rrbracket, \llbracket e_2 \rrbracket, \llbracket e_3 \rrbracket)$	[DS-SETELEM]
$\llbracket /* : I * / \text{delete } e_1[e_2] \rrbracket =$ $/* : I * / \text{delElem} (\llbracket e_1 \rrbracket, \llbracket e_2 \rrbracket)$	[DS-DELELEM]
$\llbracket /* : I * / e_2 \text{ in } e_1 \rrbracket =$ $/* : I * / \text{hasElem} (\llbracket e_1 \rrbracket, \llbracket e_2 \rrbracket)$	[DS-HASELEM]
$\llbracket /* : \ell * / \{ \bar{e}_x; \bar{e}_v \} \rrbracket =$ $\text{newobj } \ell \{ \llbracket e_{x0} \rrbracket \mapsto \llbracket e_{v0} \rrbracket \} ++ \dots \} (\text{pro}(\text{Object}))$	[DS-OBJLIT]
$\llbracket /* : \ell * / [\bar{e}] \rrbracket =$ $\text{newobj } \ell \{ \text{"0"} \mapsto \llbracket e_0 \rrbracket \} ++ \dots \} (\text{pro}(\text{Array}))$	[DS-ARRLIT]
$\llbracket /* : I * / e(e_1, \dots, e_n) \rrbracket =$ $/* : I * / \llbracket e \rrbracket (\text{window}, (\llbracket e_1 \rrbracket, \dots, \llbracket e_n \rrbracket))$	[DS-FUNCCALL]
$\llbracket /* : I * / e[e'](e_1, \dots, e_n) \rrbracket =$ $\text{let } obj = \llbracket e \rrbracket \text{ in}$ $\text{let } m = \text{getElem} (obj, \llbracket e' \rrbracket) \text{ in}$ $/* : I * / \llbracket m \rrbracket (obj, (\llbracket e_1 \rrbracket, \dots, \llbracket e_n \rrbracket))$	[DS-METHCALL]
$\llbracket \text{new } /* : \ell_{\text{new}} I * / e(e_1, \dots, e_n) \rrbracket =$ $\text{let } foo = \llbracket e \rrbracket \text{ in}$ $\text{let } obj = \text{newobj } \ell_{\text{new}} \{ (foo \text{ "prototype"}) \} \text{ in}$ $/* : I * / \llbracket foo \rrbracket (obj, (\llbracket e_1 \rrbracket, \dots, \llbracket e_n \rrbracket))$	[DS-NEW]
$\llbracket /* : I * / e.\text{apply}(e_1, \dots, e_n) \rrbracket =$ $/* : I * / \llbracket e \rrbracket (\llbracket e_1 \rrbracket, \llbracket e_2 \rrbracket, \dots, \llbracket e_n \rrbracket)$	[DS-APPLY]

Figure 12. Desugaring DJS to System !D

The DS-CONST rule for constants, the three DS-*REF rules for mutable variables, and the four DS-*ELEM rules for object and array operations are straightforward.

Object Literals. In the rules DS-OBJLIT and DS-ARRLIT, we write $\text{pro}(e) \triangleq \text{getProp} (\llbracket e \rrbracket, \text{"prototype"})$ to set the prototypes of fresh object and array literals, translated to `newobj` which creates values with prototype links. Our implementation inserts a fresh location if none is provided.

Function Application. The last four rules in Figure 12 handle different kinds function calls in JavaScript. The rules DS-FUNCCALL and DS-METHCALL desugar “direct calls” and “method calls”, where $\llbracket e' \rrbracket \triangleq \text{getElem} (e, \text{"_code_"})$. Notice that non-receiver arguments are packed into a single “arguments” tuple. In DS-FUNCCALL, we write `window` for the “global object” supplied as the receiver for direct calls. The DS-NEW rule handles object construction by creating a fresh object with `newobj` whose prototype is set to the object in the constructor object’s “prototype” field, and calls the function in the “_code_” field to finish the initialization.

JavaScript provides two native functions `apply` and `call` in `Function.prototype` that allow the caller to explicitly supply the receiver argument. We do not provide general support `apply` and `call` in DJS, because they require mechanisms beyond the scope of our (already-large) type system; for example, the latter accepts an arbitrary number of arguments. The primary benefit of (non-constructor) functions as objects in JavaScript is that they inherit `apply` and `call` from `Function.prototype`, but since we do not support them, we sacrifice little expressiveness if the type system treats every non-constructor function as a *scalar* function value, rather than an object with the function stored in “`__code__`”. Furthermore, we can then support the limited use of `apply` required for our recursive function idioms in § 2.7 using the rule DS-APPLY that *syntactically* looks for “`apply`” and explicitly sets the receiver. Because the type system prohibits (non-constructor) functions from being used as objects, there is no danger that the `apply` be “hijacked” by overwriting the “`apply`” property.

Control Operators. Throughout the paper, we wrote only JavaScript functions that have either a single `return` statement or a `return` statement along every control flow path. In general, however, `return` statements — as well as looping constructs and other control operators — can appear in arbitrary positions. In System !D, we handle the general case using *break* and *label* expressions, following λ_{JS} [21]. We omitted the formulation from our presentation for ease of exposition; see [7] for details.

6. Evaluation

In this section, we describe our implementation, the benchmarks we have annotated and type checked so far that demonstrate the expressiveness of DJS, and identify several ways for future work to improve the tool.

6.1 Implementation

We have implemented a type checker for DJS, available at ravichugh.com/djs, that is currently approximately 6,600 (non-whitespace, non-comment) lines of OCaml code. We borrow the λ_{JS} [21] JavaScript parser, use their desugaring as a starting point for our own, and use the Z3 SMT solver [11] to discharge logical validity queries. We specify the System !D primitive functions in the files `basics.dref` and `objects.dref`, and JavaScript built-in functions like `Object.prototype.hasOwnProperty` in `prelude.js` (desugared to `prelude.dref`). These three files comprise a standard prelude included with every desugared DJS program for type checking.

Local Inference. Function definitions require explicit type annotations, and we employ *bidirectional type checking* [29] techniques to infer types for local expressions. At a function application, we infer missing type and location parameters by “greedily” matching the types of arguments against any

Adapted Benchmark	Un	Ann	Queries	Time
<i>JS: The Good Parts</i>				
<code>prototypal</code>	18	36	731	2
<code>pseudoclassical</code>	15	23	706	2
<code>functional</code>	19	43	862	8
<code>parts</code>	11	20	605	3
<i>SunSpider</i>				
<code>string-fasta</code>	10	18	263	1
<code>access-binary-trees</code>	34	50	2389	23
<code>access-nbody</code>	129	201	4225	39
<i>V8</i>				
<code>splay</code>	17	36	571	1
<i>Google Closure Library</i>				
<code>typeOf</code>	15	31	1975	52
<i>Other</i>				
<code>negate</code>	9	9	296	1
<code>passengers</code>	9	19	310	3
<code>counter</code>	16	24	272	1
<code>dispatch</code>	4	8	219	1
Totals	306	518	13424	137

Figure 13. Benchmarks (Un: LOC without annotations; Ann: LOC with annotations; Queries: Number of Z3 queries; Time: Running time in seconds)

$Arr(T)$ and $Ref L$ type terms in the the declared input type and input heap. Because these type terms are invariant in their parameters (recall the U-ARRAY and U-STRONGREF rules from Figure 8), the greedy choice is always the right one. For a function type with exactly one heap variable H (like all the ones we have encountered) and input heap type (H, \hat{h}) , we infer the corresponding heap argument by simply collecting all locations in the current heap environment that do *not* match the explicit location bindings in \hat{h} . In our benchmarks, we are able to omit most type and location arguments and all heap arguments.

6.2 Benchmarks

To demonstrate the expressiveness of DJS, we have annotated and checked several small examples — inspired by *JavaScript: The Good Parts* [10], the Google Closure Library [19], and the SunSpider [34] and V8 [20] benchmarks — that exercise a variety of invariants, besides those demonstrated by previous examples (*e.g.*, `negate`, `passengers`, *etc.*). We also ported the `counter` and `dispatch` examples from System D [9] to DJS to demonstrate the nesting of function types inside objects with dynamic keys. Figure 13 summarizes our results, where for each example: “Un” is the number of (non-whitespace, non-comment) lines of code in the *unannotated* benchmark; “Ann” is the lines of code in the annotated DJS version (including comments because they contain annotations); “Time” is the running time rounded to the nearest second, tested on a 2.66GHz machine with 4GB of RAM running Ubuntu; and “Queries” is the number of validity queries issued to Z3 during type checking.

Expressiveness. We highlight some of the features of DJS that our benchmarks leverage. Besides the `prototypal` pattern discussed in § 2.5, Crockford [10] presents three additional

inheritance patterns using JavaScript’s construction mechanism. Each of these examples relies on the support for imperative, prototype-based objects in DJS.

The behavior of the `typeof` function is like the `typeof` operator except that it returns the more informative result “null” for null and “array” for arrays; the operator returns “object” in both cases. The type specification for `typeof` depends on the ability to express intersections of function types in DJS, and verifying it requires control-flow tracking in the presence of mutation as well as a precise specification for the native (ES5) function `Array.isArray`, which we model in `prelude.js`.

The `makeCumulative` function in `string-fasta.js` iterates over an object with an unknown number of keys that all store integers, and sums them in place within the object. While iterating over the keys of the object, the function uses a variable to store the key from the *previous* iteration, a subtle invariant that DJS is able to express by describing the heap before and after each iteration. Compared to the original version, we allow the bindings to store arbitrary values and use a tag-test to sum only the integer bindings. To specify the original version requires universally quantified formulas, which DJS avoids to retain decidable type checking.

The `splay` benchmark defines the following interesting tree node constructor. Rather than initializing each “own” object with null left and right subtrees, the constructor’s prototype object stores the defaults.

```
function Node(k,v) { this.k = k; this.v = v; }
Node.prototype.left = null;
Node.prototype.right = null;
```

After construction, however, Nodes are often extended with explicit subtrees. Using the flexibility of refinements, we assign each Node a type with the predicate $has(\nu, \text{“left”}) \Rightarrow sel(\nu, \text{“left”}) :: Ref \ell$, where ℓ is the weak location that describes Nodes, to ensure that retrieving the “left” key produces another Node regardless of whether it is stored on the object or not (and similarly for “right”).

Our largest example is `access-nbody`, which defines a constructor function `NBodySystem` that creates a container object to store an array of `Body` objects. The prototypes of both constructors are augmented with methods, and the thaw/freeze mechanism is heavily used while iterating over the array of `Body` objects to read and write their fields.

6.3 Annotation Burden

As Figure 13 shows, our annotated benchmarks are approximately 1.7 times as large (70% overhead) as their unannotated versions on average. In our experience, a significant majority of the annotation burden is boilerplate — unrelated to the interesting typing invariants — that fall into the following five patterns. Our implementation includes preliminary support for automatically inserting several common patterns of annotations by tracking a limited amount of type

information during desugaring (that require *no* changes to type checking). This effort has already significantly reduced the annotation overhead, but there is plenty of room for further improvements in future work.

Closures. If a function refers to a variable from an outer scope, its heap type must explicitly list its location and type. In the following example, the desugarer uses the predictable locations a_{pi} and a_e when desugaring `pi` and `e`, and the function type must contain the binding for a_{pi} .

```
var pi = 3.14, e = 2.718;
/*: Top / (a_pi ↦ n: Num) → Num / same */
function getPi() { return pi; }
```

To ease this burden, we collect the free variables in each function definition and automatically add the corresponding heap bindings that are missing. In situations where we cannot insert a suitably precise type for a location, we allow the programmer to annotate a variable declaration `var i = /* : T */ e` and we propagate T to functions that refer to i .

Loops. Because loops desugar to functions, they require a heap type annotation (like for arbitrary closures) to describe invariants that hold before and after every iteration. We infer heap types for basic patterns like the following.

```
/*: (a_i ↦ {Int(ν) ∧ i ≥ 0}) ⊕ (a_sum ↦ Num)
   ⊕ (a_ns ↦ Ref ℓ)
   ⊕ (ℓ ↦ {ν :: Arr(Num) ∧ packed(ν)}, ℓ') */
for (i=0; i < ns.length; i++) { sum += ns[i]; }
```

Thaw and Freeze. Every weak reference must first be thawed before access, which quickly becomes burdensome. As a simple aid, we surround an access to a weak reference with `thaw` and `freeze` operations, which is sufficient for simple cases involving reads and weak updates. For more complex invariants, like the relationship *between* accesses to a weak reference (as in the `passengers` example from §2.7), a single `thaw` and `freeze` pair must surround *both* accesses. In the future, we plan to insert these operations at basic block and function boundaries in the style of [31] so that objects are tracked with strong references as long as possible.

Untampered Natives. Functions that use JavaScript primitive functions like `Object.prototype.hasOwnProperty` and `Array.prototype.push` and expect them *not* to be overwritten, must explicitly constrain their input heaps as such. In most cases, programmers expect natives to remain “untampered,” so desugaring could augment all function types with these constraints.

Constructor Prototypes. The purpose of a constructor `C` is to allow its instances to inherit properties of `C.prototype` (stored at location a_{Cpro}), but functions like `useC` that use such an instance must then explicitly list the type of `C.prototype`.

```

/*: #define TC {Dict( $\nu$ )  $\wedge$  ... } */
/*: #ctor this:Ref / (this  $\mapsto$   $\langle$ Emp,  $a_{Cpro}$  $\rangle$ )
    $\rightarrow$  { $\nu = this$ } / (this  $\mapsto$   $\langle$ TC,  $a_{Cpro}$  $\rangle$ ) */
function C() { ...; return this; }
C.prototype.f = /*: Tf */ ...;
C.prototype.g = /*: Tg */ ...;

/*: x:Ref / (x  $\mapsto$   $\langle$ TC,  $a_{Cpro}$  $\rangle$ )  $\oplus$ 
   ( $a_{Cpro} \mapsto$   $\langle$ Tf(sel( $\nu$ , "f"))  $\wedge$  Tg(sel( $\nu$ , "g")),  $\ell_{op}$ ))
    $\rightarrow$  Top/same */
function useC(x) { ... x.f ... }

```

This is a predictable pattern that should be easy to incorporate into desugaring, though we have not yet done so.

6.4 Performance

The running time of our type checker is acceptable for small examples, but less so as the number of queries to the SMT solver increases. We have not yet spent much effort to improve performance, but we have implemented a few optimizations that have already reduced the number of SMT queries. First, even though desugaring works perfectly well *without* any type information, we use DJS type annotations to translate object and array operations to specific primitives, where possible, rather than the more general ones (e.g., `getPropObj` and `getIdxArr` rather than `getElement`) so that type checking has fewer cases to try, and we insert type and location parameters so that they need not be inferred. Second, we modify the T-VAR rule, which normally assigns the “selfified” type $\{\nu = x\}$ to variable x that is already bound in Γ . Although this precision is crucial, the variable x often has a simple syntactic type (e.g., $Ref\ \ell$) that is “hidden” behind the equality. Instead, if $\Gamma(x)$ is of the form $\{\nu :: U \wedge p\}$, we assign $\{\nu :: U \wedge p \wedge \nu = x\}$ so that subsequent typing rules can *syntactically* look for U rather than going through additional SMT queries as in the general case [9].

We expect that syntactically handling more common cases will further improve performance. For example, even though the dynamic keys are crucial in certain situations, many examples use objects with finite and fixed key names, which we should be able to handle with far fewer queries to the SMT solver than in the current implementation.

7. Conclusion and Future Work

In this paper, we have shown how to scale up prior work on System D — a type system for dynamic languages in a functional setting — to the real-world JavaScript setting — with imperative updates, prototype-based objects, and arrays — through a combination of strong updates and prototype chain unrolling. We have demonstrated that our new system, System !D, is expressive enough to support the invariants from a series of small but varied examples drawn from existing JavaScript benchmarks. We have found that the full range of features in DJS are indeed required, but that many

examples fall into patterns that do not simultaneously exercise all features. Therefore, we believe that future work on desugaring and on type checking can treat common cases specially in order to reduce the annotation burden and running time, and fall back to the full expressiveness of the system when necessary. We believe that Dependent JavaScript is the most promising approach, to date, for supporting real-world dynamic languages like JavaScript.

Features for Future Work. DJS already supports a large subset of JavaScript that can be used for projects where all the code is controlled (e.g., server-side applications), and future work on integrating with run-time environments could allow DJS code to run alongside full untyped JavaScript. Next, we describe several features that we currently do not support, in addition to general use of `apply` and `call` as discussed in §5.

To allow mutation of prototype links via the non-standard “`__proto__`” property, we could add a `setproto` expression to the language and detect cycles during heap unrolling.

The `eval` statement allows a string to be parsed and executed, which is useful but dangerous if misused. Since DJS is flow-sensitive, we can constraint `eval` with heap invariants before and after the statement, and then perform *staged type checking* in the style of [8] at run-time.

ES5 introduces optional per-object and per-property attributes (for example, to prevent modifications or deletions) that can likely be incorporated into our encoding of dictionaries. One benefit of such an extension is that the type system could reason more precisely about *which* objects are in a prototype chain. For example, we could then allow non-array objects to bind unsafe strings as long as we prevent them from appearing in the prototype chain of arrays, thus weakening the distinction we impose between array and non-array objects (§2.6). A second benefit is that native objects could be marked as unmodifiable, statically enforcing the pattern they are usually “untampered” as discussed in §6.

ES5 getters and setters interpose on object reads and writes. Since this is a deep change to the semantics of object operations (invoking arbitrary functions), adding general support for these will likely be heavyweight. Interestingly, one can think of our treatment of the special array “length” property (§2.6) as a built-in getter/setter.

Each function has an implicit `arguments` array that binds all parameters supplied by the caller, regardless of how many formals the function defines. Current ES6 proposals include a modified version, where an *explicit* parameter can bind a variable number of arguments *beyond* those named by formals, similar in style to Python.

The `x instanceof Foo` operator checks whether or not `Foo.prototype` is *somewhere* along the prototype chain of `x`. We could add a primitive to match these semantics.

Scalar values can be explicitly coerced by wrapper functions, such as `Boolean`, in addition to the implicit coercion we have discussed.

Undesirable Features. The last three features we discuss regularly compete for the title of worst among several “warts” in the language (*e.g.*, [10]) that lead to confusing code and hard-to-detect bugs. Incidentally, the λ_{JS} translations of all three are straightforward and can be supported in DJS, but we see no reason to given their demerits.

The `with` statement adds the fields of an object to the current scope of a block, allowing them to be accessed without qualification. There is hardly a good reason to use this feature, and it is banned in ES5 “strict” mode.

All `var` declarations are implicitly lifted to the top of the enclosing function, resulting in “function scope” rather than lexical scope. Although simple to detect when `var`-lifting kicks in, we opt for the latter. ES6 will likely add an explicit `let` binding form that is not subject to lifting. In DJS, `var` is essentially the new `let` form, but we stick with the traditional syntax for familiarity.

For a “method call” `x.f(y)`, the receiver `x` is supplied for the `this` argument to the function, but for a “direct call” `x(y)`, JavaScript implicitly supplies the global object for `this`, masking common errors. We choose to statically reject direct calls to functions that require a `this` parameter.

8. Related Work

In this section, we discuss topics related to types for imperative dynamic languages, and hence strong updates and inheritance. The reader may refer to [9] for background on the challenging idioms of even functional dynamic languages and the solution that nested refinements provide.

Location Sensitive Types. The way we handle reference types draws from the approach of *Alias Types* [32], in which strong updates are enabled by describing reference types with abstract location names and by factoring reasoning into a flow-insensitive tying environment and a flow-sensitive heap. *Low-level liquid types* [31] employs their approach in the setting of a first-order language with dependent types. In contrast, our setting includes higher-order functions, and our formulation of heap types gives variable names to unknown heaps to reason about prototypes and gives names to *all* heap values, which enables the specification of precise relationships between values of *different* heaps; the heap binders of [31] allow only relationships between values in a *single* heap to be described.

The original *Alias Types* work also includes support for *weak references* that point to zero or more values, for which strong updates are not sound. Several subsequent proposals [2, 3, 14, 15, 31, 33] allow strong updates to weak references under certain circumstances to support temporary invariant violations. We adapt the *thaw/freeze* mechanism from [2] and [31] with mostly cosmetic changes.

Prototype Inheritance. Unlike early class-based languages, such as Smalltalk and C++, the (untyped) language Self allows objects to be extended after creation and feature prototype, or delegation, inheritance. Static typing disciplines

for class-based languages (*e.g.*, [1]) explicitly preclude object extension to retain soundness in the presence of *width subtyping*, the ability to forget fields of an object. To mitigate the tension between object extension and subtyping, several proposals [6, 18] feature quite a different flavor: the fields of an object are split into a “reservation” part, which may be added to an object but cannot be forgotten, and a “sealed part” that can be manipulated with ordinary subtyping. Our approach provides additional precision in two important respects. First, we precisely track prototype hierarchies, whereas the above approaches flatten them into a single collection of fields. Second, we avoid the separation of reservation and sealed fields but still allow subtyping, since “width subtyping” in System !D is simply logical implication over refinement formulas; forgetting a field — discarding a $has(d, k)$ predicate — does not imply that $\neg has(d, k)$, which guards the traversal of the prototype chain.

Typed Subsets of JavaScript. Several (syntactic) type systems for various JavaScript subsets have been proposed. Among the earliest is [35], which identifies silent errors that result from implicit type coercion and the fact that JavaScript returns `undefined` when trying to look up a non-existent key from an object. The approach in [4] distinguishes between *potential* and *definite* keys, similar to the reservation and sealed discussed above; this general approach has been extended with flow-sensitivity and polymorphism [39]. The notion of *recency types*, similar to *Alias Types*, was applied to JavaScript in [23], in which typing environments, in addition to heap types, are flow-sensitive. Prototype support in [23] is limited to the finite number of prototype links tracked by the type system, whereas the *heap symbols* in System !D enable reasoning about *entire* prototype hierarchies. Unlike System !D, all of the above systems provide global type inference; our system does not have principal types, so we can only provide local type inference [29]. AD-safety [30] is a type system for ADsafe, a JavaScript sandbox, that restricts access to some fields. Although expressive enough to check ADsafe, which heavily uses large object literals, they do not support strong update and so cannot reason about object extension. Unlike System !D, none of the above systems include dependent types, which are required to express truly dynamic object keys and precise control-flow based invariants.

Recent work on JavaScript verification uses separation logic [17] to track precise flow-sensitive invariants. They support only first-order programs, and the expressiveness of their logic takes them beyond automatic verification, thus requiring properties to be manually proved.

JavaScript Semantics. We chose the JavaScript “semantics-by-translation” of λ_{JS} [21] since it targets a conventional core language that has been convenient for our study. An alternate semantics [26] inherits unconventional aspects of the language specification [24] (*e.g.*, “scope objects”), which complicates the formulation of static reasoning.

Acknowledgments

This work was supported by NSF Grants CCF-0644361, CNS-0964702, and gifts from Microsoft Research. Part of this work was done while the first author was at Mozilla.

References

- [1] M. Abadi and L. Cardelli. *A Theory of Objects*. Springer-Verlag, 1996.
- [2] A. Ahmed, M. Fluet, and G. Morrisett. L^3 : A Linear Language with Locations. *Fundamenta Informaticae*, 77(4), June 2007.
- [3] A. Aiken, J. Kodumal, J. S. Foster, and T. Terauchi. Checking and Inferring Local Non-Aliasing. In *PLDI*, 2003.
- [4] C. Anderson, S. Drossopoulou, and P. Giannini. Towards Type Inference for JavaScript. In *ECOOP*, 2005.
- [5] G. M. Bierman, A. D. Gordon, C. Hritcu, and D. E. Langworthy. Semantic Subtyping with an SMT Solver. In *ICFP*, 2010.
- [6] V. Bono and K. Fisher. An Imperative, First-Order Calculus with Object Extension. In *ECOOP*, 1998.
- [7] R. Chugh, D. Herman, and R. Jhala. Dependent Types for JavaScript — Appendix. arxiv.org/abs/1112.4106.
- [8] R. Chugh, J. A. Meister, R. Jhala, and S. Lerner. Staged Information Flow for JavaScript. In *PLDI*, 2009.
- [9] R. Chugh, P. M. Rondon, and R. Jhala. Nested Refinements: A Logic for Duck Typing. In *POPL*, 2012.
- [10] D. Crockford. *JavaScript: The Good Parts*. Yahoo! Press, 2008.
- [11] L. de Moura and N. Bjørner. Z3: An Efficient SMT solver. In *TACAS*, 2008.
- [12] R. DeLine and M. Fähndrich. Enforcing High-level Protocols in Low-level Software. In *PLDI*, 2001.
- [13] ECMA. TC-39 Committee. www.ecmascript.org/community.php.
- [14] M. Fähndrich and R. DeLine. Adoption and Focus: Practical Linear Types for Imperative Programming. In *PLDI*, 2002.
- [15] J. Foster, T. Terauchi, and A. Aiken. Flow-sensitive Type Qualifiers. In *PLDI*, 2002.
- [16] M. Furr, J. hoon (David) An, J. S. Foster, and M. W. Hicks. Static type inference for ruby. In *SAC*, 2009.
- [17] P. Gardner, S. Maffei, and G. D. Smith. Towards a Program Logic for JavaScript. In *POPL*, 2012.
- [18] P. D. Gianantonio, F. Honsell, and L. Liquori. A Lambda Calculus of Objects with Self-Inflicted Extension. In *OOPSLA*, 1998.
- [19] Google. Closure library. <https://developers.google.com/closure/library>.
- [20] Google. V8 benchmark. <http://v8.googlecode.com/svn/data/benchmarks/>.
- [21] A. Guha, C. Saftoiu, and S. Krishnamurthi. The Essence of JavaScript. In *ECOOP*, 2010.
- [22] A. Guha, C. Saftoiu, and S. Krishnamurthi. Typing Local Control and State Using Flow Analysis. In *ESOP*, 2011.
- [23] P. Heidegger and P. Thiemann. Recency Types for Analyzing Scripting Languages. In *ECOOP*, 2010.
- [24] E. International. *ECMAScript Language Specification, ECMA-262, 3rd ed.* 1999.
- [25] K. W. Knowles and C. Flanagan. Compositional Reasoning and Decidable Checking for Dependent Contract Types. In *PLPV*, 2009.
- [26] S. Maffei, J. Mitchell, and A. Taly. An operational semantics for JavaScript. In *APLAS*, 2008.
- [27] J. McCarthy. Towards a Mathematical Science of Computation. In *IFIP*, 1962.
- [28] B. C. Pierce. *Types and Programming Languages*. 2002.
- [29] B. C. Pierce and D. N. Turner. Local Type Inference. In *POPL*, 1998.
- [30] J. G. Politz, S. A. Eliopoulos, A. Guha, and S. Krishnamurthi. Adsafety: Type-based Verification of JavaScript Sandboxing. In *USENIX Security*, 2011.
- [31] P. Rondon, M. Kawaguchi, and R. Jhala. Low-Level Liquid Types. In *POPL*, 2010.
- [32] F. Smith, D. Walker, and G. Morrisett. Alias Types. In *ESOP*, 2000.
- [33] J. Sunshine, K. Naden, S. Stork, J. Aldrich, and E. Tanter. First-class State Change in Plaid. In *OOPSLA*, 2011.
- [34] SunSpider. Javascript benchmark. <http://www.webkit.org/perf/sunspider/sunspider.html>.
- [35] P. Thiemann. Towards a Type System for Analyzing JavaScript Programs. In *ESOP*, 2005.
- [36] S. Tobin-Hochstadt and M. Felleisen. Logical Types for Untyped Languages. In *ICFP*, 2010.
- [37] D. Walker and G. Morrisett. Alias Types for Recursive Data Structures. In *TIC*. 2000.
- [38] S. yu Guo and B. Hackett. Fast and Precise Hybrid Type Inference for JavaScript. In *PLDI*, 2012.
- [39] T. Zhao. Polymorphic Type Inference for Scripting Languages with Object Extensions. In *DLS*, 2011.