

Precise and Scalable Static Analysis of jQuery using a Regular Expression Domain

Changhee Park

KAIST, Republic of Korea
changhee.park@kaist.ac.kr

Hyeonseung Im

Kangwon National University,
Republic of Korea
hsim@kangwon.ac.kr

Sukyoung Ryu

KAIST, Republic of Korea
sryu.cs@kaist.ac.kr

Abstract

jQuery is the most popular JavaScript library but the state-of-the-art static analyzers for JavaScript applications fail to analyze simple programs that use jQuery. In this paper, we present a novel abstract string domain whose elements are simple regular expressions that can represent prefix, infix, and postfix substrings of a string and even their sets. We formalize the new domain in the abstract interpretation framework with abstract models of strings and objects commonly used in the existing JavaScript analyzers. For practical use of the domain, we present polynomial-time inclusion decision rules between the regular expressions and prove that the rules exactly capture the actual inclusion relation. We have implemented the domain as an extension of the open-source JavaScript analyzer, SAFE, and we show that the extension significantly improves the scalability and precision of the baseline analyzer in analyzing programs that use jQuery.

Categories and Subject Descriptors F.3.2 [*Semantics of Programming Languages*]: Program Analysis

Keywords JavaScript, static analysis, regular expressions

1. Introduction

Static analysis of JavaScript web applications often requires static analysis of jQuery [11] because it is the most popular JavaScript library. However, the dynamic natures of JavaScript that are heavily used in jQuery make static analysis impractical by causing scalability problems due to imprecise analysis results. For example, an imprecise approximate value of a property name x in a dynamic property access $\text{obj}[x]$ may yield accessing all properties of obj at the worst case leading to the analysis state explosion.

To alleviate such problems in dynamic property accesses, static analysis should provide an elaborate string analysis model for property names, but existing analyzers use simple string domains that are not enough to represent first-class property names widely used in jQuery. Thus, the state-of-the-art JavaScript static analyzers such as JSAI [13], SAFE [19, 20], and TAJIS [3] often fail to analyze simple applications that use jQuery. In such string domains, for instance, concatenation of a statically indeterminate value like a random value and a constant string is approximated as any string value, which may lead to the state explosion when a property access uses it as a property name.

In this paper, we present a novel string domain that consists of simple regular expressions. The domain is expressive enough to represent prefix, infix, and postfix substrings of a string and even their sets. We formalize the domain in abstract interpretation [6] for the core semantics of objects, and describe how it improves analysis precision in various abstract operations such as string concatenation and dynamic property read and write, compared to a constant string domain commonly used in existing JavaScript static analyzers.

For a sound abstraction, we define an order relation between regular expressions in the domain using the inclusion relation [8, 17] between them, which states that one regular expression R_1 includes another R_2 if a set of all strings that R_2 matches is a subset of the set of all strings that R_1 matches. While deciding the inclusion relation for general regular expressions has been known as an intractable problem [24], we present polynomial-time decision rules for the inclusion relation between the regular expressions in the domain and prove that the rules exactly capture the actual inclusion relation. Finally, we demonstrate that the new domain significantly improves the analysis scalability and precision of an existing JavaScript static analyzer.

The contributions of this paper are as follows:

- We formally describe how to extend an analysis model of string and object domains, commonly used in the existing static analyzers for JavaScript, with the new domain and prove its soundness theorem.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive version was published in the following publication:

DLS'16, November 1, 2016, Amsterdam, Netherlands
© 2016 ACM. 978-1-4503-4445-6/16/11...
<http://dx.doi.org/10.1145/2989225.2989228>

```

1 <script type="text/javascript"
2   src="jquery-1.7.0.js">
3 <ul id="ulUsers"></ul>
4 <script type="text/javascript">
5   $(function() {
6     $.get("http://tests.jquery-tutorial.net/
7         json.php?callback=?",
8         function(data, textStatus) {
9           $.each(data, function(index, user){
10            $("#ulUsers").append(
11              $("<li></li>").text(
12                user.name + " is " +
13                user.age + " years old"));
14            });
15          },
16          "json" );
17        });
18 </script>

```

Figure 1: An example from a jQuery tutorial

- We present polynomial-time inclusion decision rules for the regular expressions and prove their soundness and completeness with respect to the actual inclusion relation. To the best of our knowledge, no work has presented tractable inclusion decision rules for them.
- We have implemented the new domain as an extension of SAFE. We make the implementation of the extended SAFE open to the public [12].
- We evaluate the extension by analyzing 61 programs that use jQuery and show that the technique significantly improves the scalability and precision of SAFE.

The rest of the paper is organized as follows. Section 2 gives an example showing challenges in existing static analyzers, and informally describes the new domain. In Section 3, we formalize the concrete semantics and the existing analysis model of JavaScript strings and objects. Sections 4 and 5 formally present the regular expression domain and its inclusion rules, and Section 6 evaluates the domain with respect to analysis scalability and precision. Finally, we discuss related work in Section 7 and conclude in Section 8.

2. Motivation

This section shows an example code that significantly degrades analysis precision and scalability of a state-of-the-art JavaScript static analyzer and describes our solution.

2.1 Example

Figure 1 shows a simple program from a jQuery tutorial¹, which contains 18 lines of code that updates the contents in the `` tag (line 3) with data received from a server using an AJAX request. The program first loads a jQuery library (lines 1~2) and registers an anonymous function as a callback (line 5) to be called once a browser has completed loading the document; when this callback is called, using

¹<http://www.jquery-tutorial.net/ajax/requesting-a-file-from-a-different-domain-using-jsonp>

```

1   expando: "jQuery" +
2     (jQuery.fn.jquery + Math.random()) ...
3   jsc = jQuery.now(), ...
4   jsonpCallback: function() {
5     return jQuery.expando + "_" + (jsc++);
6   } ...
7   jQuery.ajaxPrefilter("json jsonp",
8     function (...) { ...
9     jsonpCallback = ... s.jsonpCallback() ...
10    previous = window[jsonpCallback], ...
11    window[jsonpCallback] = function() { ... }; ...
12    jqXHR.always(function() { ...
13      window[jsonpCallback] = previous; ... }) ...
14  });

```

Figure 2: An excerpt from jQuery 1.7.0

`$.get` (line 6) it receives JSON data from a server url at lines 6 and 7 and passes the data to the handler at lines 8~15; then the handler creates `` tags with the user information inside the `` tag (lines 11~13). Hence, if a server provides the following JSON data,

```

[{"name": "John Doe", "age": 42},
 {"name": "Jane Doe", "age": 39}]

```

then, the program creates the following tags with the data as the result of the program:

```

<li> John Doe is 42 years old </li>
<li> Jane Doe is 39 years old </li>

```

2.2 Challenges in Static Analysis

Statically analyzing the `$` function in jQuery requires high analysis precision because it behaves differently for different argument values. For example, the `$` calls at lines 5, 10, and 11 in Figure 1 perform different tasks; the first registers its argument function as a load event handler, the second finds a DOM element with the `ulUsers` id, and the third creates a new DOM element with the `` tag. Since SAFE supports flow-sensitive and context-sensitive analyses, it can analyze different behaviors of the `$` function with high precision.

However, SAFE fails to analyze the simple program in Figure 1 within the timeout of 10 hours. Figure 2 shows an excerpt from jQuery 1.7.0 where the anonymous function at line 8 is called to install a callback function with an appointed name before the actual AJAX request. In the case of the example in Figure 1, since the program does not appoint any name for a callback function by `"callback=?"` (line 7 of Figure 1), jQuery gives a random name by calling `jsonpCallback` (line 4), which returns a string that results from concatenating the values of `jQuery.expando`, `"_"`, and `jsc`; lines 1,2 and 3 show the definitions of `jQuery.expando` and `jsc` that have string values with a random number postfixed and a number representation of the current time, respectively. Then, it makes a backup for the old value with the name `jsonpCallback` (line 10), installs a new callback function (line 11), and restores the old value after the request is done (line 13).

Note that existing static analyzers for JavaScript applications approximate concatenation of a string and a statically indeterminate number string like the results of `Math.random()` and `jQuery.now()` as the string value that may become any string; from now on, we call such a value \top_s . Then, the analysis result of `jsonpCallback` at line 9 is \top_s , and that of `previous` at line 10 is a combined value of all global variable values including global functions, because the `window` object represents a global scope object in JavaScript and thus `window[jsonpCallback]` accesses all global variables. What is worse, this imprecise value is assigned to each global variable at line 13 polluting all variables at the global scope. We found that a subsequent global function call is resolved into 69 possible calls incurring a huge computation overhead during analysis to reach a fix-point. The problem lies in that the JavaScript static analyzers use simple variants of the constant string lattice domain. Using such simple domains, analyzers often approximate the value of `jsonpCallback` as \top_s losing the information from the concatenation that the value should include "-".

A fundamental solution may be to use a more expressive domain but we found that all 12 string domains suggested by Madsen and Andreasen [16] are not sufficient to solve the problem. A character inclusion domain [16] can express the inclusion of "-" in `jsonpCallback`, but it cannot express `jQuery.expando` at line 1, which `jQuery` frequently uses as a key to property accesses. In this case, the prefix domain would be helpful since we can use the prefix "jQuery". However, unlike the character inclusion domain, it cannot express the inclusion of "-" in `jsonpCallback`. All 12 string domains in the previous work [16] fail to express both information for `jsonpCallback` and `expando`.

2.3 Regular Expression Domain

As a solution, we present a novel abstract string domain that has regular expressions as its elements. The domain is expressive enough to represent prefix, infix, and postfix substrings and even their sets. It contains a special element \star , which is regarded as \top_s but may be concatenated with other constant strings improving the expressiveness. For example, the abstract values for `jsonpCallback` and `expando` in the previous example have " \star_* " and "jQuery \star " in our domain, respectively, where \star matches any strings. Then, when an analyzer uses the values as keys to property accesses, it accesses only such properties that match the given regular expressions. In this way, the analyzer, remaining sound, can increase not only the analysis precision but also its performance getting benefits from the improved precision. As for the example in Figure 1, when we apply the new domain to SAFE, it completes the analysis within 9 minutes with precise analysis results that approximate each global function call as one correct function call. In the next sections, we formalize the domain in abstract interpretation and show that it significantly improves the scalability and precision of an existing static analyzer in analyzing programs that use `jQuery`.

$$\begin{aligned} obj, o &\in \mathbf{Object} = \mathbf{String} \rightarrow \mathbf{Value} \\ \text{PropRead} &\in \mathbf{Object} \times \mathbf{String} \rightarrow \mathbf{Value} \\ \text{PropWrite} &\in \mathbf{Object} \times \mathbf{String} \times \mathbf{Value} \rightarrow \mathbf{Object} \end{aligned}$$

$$\begin{aligned} \text{PropRead}(obj, x) &= \begin{cases} obj(x) & \text{if } x \in \text{dom}(obj) \\ \text{undefined} & \text{otherwise} \end{cases} \\ \text{PropWrite}(obj, x, v) &= \begin{cases} (obj \setminus [x \mapsto obj(x)]) * [x \mapsto v] & \text{if } x \in \text{dom}(obj) \\ obj * [x \mapsto v] & \text{otherwise} \end{cases} \end{aligned}$$

Figure 3: JavaScript objects and object property read/write

$$\begin{aligned} \text{PropReadSet} &\in \wp(\mathbf{Object}) \times \wp(\mathbf{String}) \rightarrow \wp(\mathbf{Value}) \\ \text{PropWriteSet} &\in \wp(\mathbf{Object}) \times \wp(\mathbf{String}) \times \wp(\mathbf{Value}) \rightarrow \wp(\mathbf{Object}) \\ \text{StrConcatSet} &\in \wp(\mathbf{String}) \times \wp(\mathbf{String}) \rightarrow \wp(\mathbf{String}) \\ \text{PropReadSet}(oset, sset) &= \{\text{PropRead}(o, x) \mid o \in oset \wedge x \in sset\} \\ \text{PropWriteSet}(oset, sset, vset) &= \{\text{PropWrite}(o, x, v) \mid o \in oset \wedge x \in sset \wedge v \in vset\} \\ \text{StrConcatSet}(sset_1, sset_2) &= \{s_1 \oplus s_2 \mid s_1 \in sset_1 \wedge s_2 \in sset_2\} \end{aligned}$$

Figure 4: Property read/write operations on objects and string concatenation in collecting semantics

3. Formal Description

We formally present a simplified semantics of property read and write operations of objects in JavaScript and describe an existing static analysis model for them, which are commonly used in JSAI, SAFE, and TAJIS. For presentation brevity, we omit internal properties in objects and prototype-related parts in the formalization. Note that our implementation addresses them as well using the existing support from the baseline SAFE analyzer.

3.1 Concrete Semantics of JavaScript Objects

Figure 3 presents definitions of JavaScript objects and their property read and write operations using the following:

- $\text{dom}(obj)$: returns a set of all domain elements of obj
- $obj * [x \mapsto v]$: adds a map element $[x \mapsto v]$ to obj
- $obj \setminus [x \mapsto v]$: deletes a map element $[x \mapsto v]$ from obj .

An object maps strings to values, \mathbf{Value} , and the property read and write operations are functions from a pair of an object and a string to a value (PropRead) and from a tuple of an object, a string, and a value to an updated object (PropWrite), respectively. Note that, in JavaScript, an absent property read gives the undefined value and an absent property write adds a new property to the target object. Figure 4 shows lifted property read and write operations for objects and string concatenation in collecting semantics; we use $\wp(S)$ for the power set of the set S and \oplus for string con-

$$\begin{aligned}
& \text{(abstraction)} \quad \alpha_s \in \wp(\mathbf{String}) \mapsto \widehat{\mathbf{String}} \\
& \text{(concretization)} \quad \gamma_s \in \widehat{\mathbf{String}} \mapsto \wp(\mathbf{String}) \\
& \alpha_s(sset) = \begin{cases} \perp_s & \text{if } sset = \{\} \\ x & \text{if } sset = \{x\} \\ \top_s & \text{otherwise} \end{cases} \\
& \gamma_s(\hat{s}) = \begin{cases} \{\} & \text{if } \hat{s} = \perp_s \\ \{x\} & \text{if } \hat{s} = x \\ \mathbf{String} & \text{if } \hat{s} = \top_s \end{cases} \\
& \text{StrConcat}(\hat{s}_1, \hat{s}_2) = \begin{cases} \perp_s & \text{if } \hat{s}_1 = \perp_s \vee \hat{s}_2 = \perp_s \\ x_1 \oplus x_2 & \text{if } \gamma_s(\hat{s}_1) = \{x_1\} \wedge \gamma_s(\hat{s}_2) = \{x_2\} \\ \top_s & \text{otherwise} \end{cases}
\end{aligned}$$

Figure 5: Galois connection of strings and abstract string concatenation

catenation. The collecting semantics approximates all possible result values of operations on a given program point in all reachable execution traces. For example, on a certain program point with a PropRead operation, if possible input values for PropRead in all execution are two objects, $[a \mapsto 1]$ and $[b \mapsto 2]$, and two property names, “a” and “b”, then PropReadSet produces $\{1, 2, \text{undefined}\}$ that contains result values from all combinations of the input values. Throughout this paper, we present only the string concatenation \oplus as a representative string operation.

3.2 Static Analysis Model

We present an analysis model of strings and objects commonly used in existing analyzers, and we formalize them in the abstract interpretation framework [6].

Abstract strings. An abstract string domain denoted by $\widehat{\mathbf{String}}$ is the constant string lattice domain whose element is either \top_s , a constant string, or \perp_s ; intuitively, \top_s means any string, a constant means a single concrete string, and \perp_s no string. The elements have the following order relation \sqsubseteq_s :

$$\hat{s}_1 \sqsubseteq_s \hat{s}_2 \Leftrightarrow (\hat{s}_1 = \perp_s \vee \hat{s}_2 = \top_s \vee \hat{s}_1 = \hat{s}_2).$$

The domain is the one that JSAI, SAFE, and TAJs commonly support while they provide variants of it. All of them have additional elements that represent any number strings and any non-number strings to distinguish property names for arrays, JSAI has another element to represent built-in property names of JavaScript objects. We omit the details of the variants because they still have the same problem described in Section 2.2 as the base constant domain does. Figure 5 shows the definitions of two translation functions α_s and γ_s . It is trivial to show that these two functions satisfy the following Galois connection condition:

THEOREM 1. (*Galois connection of strings*)
 $\forall s \in \wp(\mathbf{String}), \hat{s} \in \widehat{\mathbf{String}} : \alpha_s(s) \sqsubseteq_s \hat{s} \iff s \subseteq \gamma_s(\hat{s}).$

$$\begin{aligned}
& \perp_o, \hat{o} = (\text{propmap}, \text{s\hat{u}m}, \text{defset}) \in \mathbf{Object} = \\
& \quad \{\perp_o\} \cup ((\mathbf{String} \mapsto \mathbf{Value}) \times \mathbf{Value} \times \wp(\mathbf{String})) \\
& \alpha_v \in \wp(\mathbf{Value}) \rightarrow \mathbf{Value} \\
& \gamma_v \in \mathbf{Value} \rightarrow \wp(\mathbf{Value}) \\
& \alpha_o \in \wp(\mathbf{Object}) \rightarrow \mathbf{Object} \\
& \gamma_o \in \mathbf{Object} \rightarrow \wp(\mathbf{Object}) \\
& \text{Propmap}(\langle \text{propmap}, \text{s\hat{u}m}, \text{defset} \rangle) = \text{propmap} \\
& \text{Summary}(\langle \text{propmap}, \text{s\hat{u}m}, \text{defset} \rangle) = \text{s\hat{u}m} \\
& \text{Defset}(\langle \text{propmap}, \text{s\hat{u}m}, \text{defset} \rangle) = \text{defset} \\
& \hat{o}_1 \sqsubseteq_o \hat{o}_2 \Leftrightarrow (\hat{o}_1 = \perp_o \vee (\hat{o}_1 \neq \perp_o \wedge \hat{o}_2 \neq \perp_o \wedge \\
& \quad \sqcup_v \{ \text{Propmap}(\hat{o}_1)(s) \mid \\
& \quad \quad s \in \text{dom}(\text{Propmap}(\hat{o}_1)) \setminus \text{dom}(\text{Propmap}(\hat{o}_2)) \} \\
& \quad \sqsubseteq_v \text{Summary}(\hat{o}_2) \\
& \quad \wedge \forall s \in (\text{dom}(\text{Propmap}(\hat{o}_1)) \cap \text{dom}(\text{Propmap}(\hat{o}_2))) : \\
& \quad \quad \text{Propmap}(\hat{o}_1)(s) \sqsubseteq_v \text{Propmap}(\hat{o}_2)(s) \\
& \quad \wedge \text{Summary}(\hat{o}_1) \sqsubseteq_v \text{Summary}(\hat{o}_2) \\
& \quad \wedge \text{Defset}(\hat{o}_2) \subseteq \text{Defset}(\hat{o}_1)) \\
& \hat{o}_1 \sqcup_o \hat{o}_2 = \begin{cases} \text{if } \hat{o}_1 = \perp_o, & \hat{o}_2 \\ \text{if } \hat{o}_2 = \perp_o, & \hat{o}_1 \\ \text{otherwise,} & \\ \{ \{ [s \mapsto (\text{Propmap}(\hat{o}_1)(s) \sqcup_v \text{Propmap}(\hat{o}_2)(s))] \mid \\ & s \in \text{dom}(\text{Propmap}(\hat{o}_1)) \vee s \in \text{dom}(\text{Propmap}(\hat{o}_2)) \} \}, \\ & \text{Summary}(\hat{o}_1) \sqcup_v \text{Summary}(\hat{o}_2), \\ & \text{Defset}(\hat{o}_1) \cap \text{Defset}(\hat{o}_2) \} \\
& \alpha_o(\text{oset}) = \begin{cases} \text{if } \text{oset} = \{\}, & \perp_o \\ \text{otherwise,} & \\ \{ \{ [s \mapsto \alpha_v(\{o_2(s) \mid o_2 \in \text{oset} \wedge s \in \text{dom}(o_2)\})] \\ & \mid o \in \text{oset} \wedge s \in \text{dom}(o) \} \}, \\ & \perp_v, \cap \{ \text{dom}(o) \mid o \in \text{oset} \} \} \\
& \gamma_o(\hat{o}) = \cup \{ \text{oset} \mid \alpha_o(\text{oset}) \sqsubseteq \hat{o} \}
\end{aligned}$$

Figure 6: Abstract objects and Galois connection of objects

StrConcat in Figure 5 defines the abstract string concatenation on the abstract string domain. Note that while this operation gives a precise result only when both of two arguments indicate constant string values, it returns the imprecise \top_s value when either of two arguments is \top_s . Then the imprecise value may flow into a key for property accesses in objects and incur a huge performance overhead due to the state explosion such as the one described in Section 2. We can prove that the function StrConcat is monotone and a sound approximation of StrConcatSet. Formal theorems and their proofs are publicly available [12].

Abstract objects. Figure 6 shows the definition of abstract objects. An abstract object \hat{o} is either a bottom object (\perp_o) or a tuple of a property map (*propmap*), a summary abstract value (*s\hat{u}m*), and a definite set of strings (*defset*):

can be represented using regular expressions in our new domain. Hence, all abstract operations defined in **String** are re-usable in the new regular expression domain and it is possible to refine the operations incrementally as needed. This point can save huge amounts of modification in the design and implementations of the existing analyzers that use constant string domains to use the new domain.

Figure 8 shows re-defined operations in the new domain where a regular expression matches a finite number of strings unless it contains $*$. $\widehat{\text{NewStrConcat}}$ simply concatenates all atomic regular expressions in input regular expressions; in this way, it can maintain substring information even when it concatenates a random string and a constant string unlike $\widehat{\text{StrConcat}}$ that loses the information. $\widehat{\text{NewPropRead}}$ returns no value (\perp_v) when the input values include no object (\perp_o) or no regular expression ($\{\}$) in the first rule and re-uses PropRead when the set of strings that the given regular expression matches is non-empty and finite. The function in the last rule retrieves all property values whose names match the given regular expression and soundly joins all the values with the summary value of the input object and undefined. Similarly, $\widehat{\text{NewPropWrite}}$ returns no object (\perp_o) with abnormal input values and re-uses PropWrite when the set of strings that the given regular expression matches is non-empty and finite. In the other case, it *weakly* updates the values of the properties whose names match the regular expression ($s \in \llbracket R \rrbracket$). Note that for the properties whose names do not match the regular expression ($s \notin \llbracket R \rrbracket$), it keeps the original values ($\text{Propmap}(\hat{o})(s)$).

The operations in Figure 8 are monotone and sound approximation of their corresponding concrete operations:

THEOREM 5. (*Soundness of $\widehat{\text{NewPropRead}}$ and $\widehat{\text{NewPropWrite}}$*)

$$\begin{aligned}
 & (\forall o \in \wp(\mathbf{Object}), s \in \wp(\mathbf{String}) : \\
 & \alpha_v(\text{PropReadSet}(o, s)) \sqsubseteq_v \widehat{\text{NewPropRead}}(\alpha_o(o), \alpha_r(s)) \\
 & \wedge \\
 & (\forall o \in \wp(\mathbf{Object}), s \in \wp(\mathbf{String}), v \in \wp(\mathbf{Value}) : \\
 & \alpha_o(\text{PropWriteSet}(o, s, v)) \\
 & \sqsubseteq_o \widehat{\text{NewPropWrite}}(\alpha_o(o), \alpha_r(s), \alpha_v(v))).
 \end{aligned}$$

The proof is publicly available [12].

5. Inclusion Decision Rules for Regular Expressions

Recall that we define the order relation \sqsubseteq_r between our regular expressions using the inclusion relation as follows:

$$R_1 \sqsubseteq_r R_2 \iff \llbracket R_1 \rrbracket \subseteq \llbracket R_2 \rrbracket.$$

Implementing an inclusion decision algorithm is necessary to practically use regular expression domains in sound static analysis, but as we mentioned in the previous section, the decision problem for general regular expressions is intractable. While the literature [8, 17] discovered some tractable cases with restricted subsets of regular expressions, we present

Auxiliary Functions

$$\begin{aligned}
 \text{pruneNonStar}(R) &= \{\ast aA \mid \ast aA \in R\} \\
 \text{pruneHead}(a, R) &= \text{pruneNonStar}(R) \cup \\
 & \quad \{A \mid aA \in R \vee \ast aA \in R\}
 \end{aligned}$$

Inclusion Decision Rules

$R_1 \sqsubseteq_r R_2$ means R_2 includes R_1 .

$$\begin{aligned}
 \text{INCLSTAR} & \frac{\ast \in R_2}{R_1 \sqsubseteq_r R_2} & \text{INCLEMPTYSET} & \frac{\ast \notin R}{\{\} \sqsubseteq_r R} \\
 \text{INCLLSET} & \frac{\ast \notin R_2 \quad |R_1| \geq 2 \quad \forall A \in R_1 : \{A\} \sqsubseteq_r R_2}{R_1 \sqsubseteq_r R_2} \\
 \text{INCLEMPTY} & \frac{\ast \notin R \quad \text{Empty} \in R}{\{\text{Empty}\} \sqsubseteq_r R} \\
 \text{INCLAREG} & \frac{\ast \notin R \quad \{A\} \sqsubseteq_r \text{pruneHead}(a, R)}{\{aA\} \sqsubseteq_r R} \\
 \text{INCLSTARREG} & \frac{\ast \notin R \quad \{aA\} \sqsubseteq_r \text{pruneNonStar}(R)}{\{\ast aA\} \sqsubseteq_r R}
 \end{aligned}$$

Figure 9: Inclusion decision rules for regular expressions in \mathbb{R}

novel decision rules for another subset of regular expressions, which has the polynomial-time complexity.

Figure 9 shows the definition of the decision relation \sqsubseteq_r . Let us consider core rules first using the $\{\ast aa\} \sqsubseteq_r \{\ast a, b\}$ example. The **INCLSTARREG** rule prunes the atomic regular expression b in $\{\ast a, b\}$ using pruneNonStar because b can never include $\ast aa$, and then it tries to decide $\{aa\} \sqsubseteq_r \{\ast a\}$ since if $\ast a$ includes aa , it also includes $\ast aa$. Next, the **INCLAREG** rule matches the first character a in aa with $\ast a$ and removes the matched parts in them, and then the decision depends on that of the remaining parts, $\{a\} \sqsubseteq_r \{\ast a, \text{Empty}\}$. Note that pruning the matched part a in $\{\ast a\}$ by pruneHead produces $\{\ast a, \text{Empty}\}$ because $\{\ast a\}$ is semantically equivalent to $\{\top \ast a, a\}$ if \top means a regular expression that matches any single character, and pruning a in $\{\top \ast a, a\}$ produces $\{\ast a, \text{Empty}\}$. Likewise, the **INCLAREG** rule re-applies to $\{a\} \sqsubseteq_r \{\ast a, \text{Empty}\}$ again and prunes the head a in both $\{a\}$ and $\{\ast a, \text{Empty}\}$ leaving the decision $\{\text{Empty}\} \sqsubseteq_r \{\ast a, \text{Empty}\}$. Finally, the **INCLEMPTY** rule accepts the decision since $\text{Empty} \in \{\ast a, \text{Empty}\}$.

The rest are trivial. The **INCLSTAR** rule states that a regular expression that contains $*$ includes any regular expression, and the other 5 rules handle the cases where R_2 in $R_1 \sqsubseteq_r R_2$ does not have $*$. The **INCLEMPTYSET** and **INCLLSET** rules are when the size of R_1 in $R_1 \sqsubseteq_r R_2$ is 0 and bigger than 1, respectively; the former states that any regular expression includes no regular expression ($\{\}$) and the latter states that if R_2 includes all atomic regular expressions in R_1 , it includes R_1 as well. Note that $\{\ast\} \sqsubseteq_r R_2$ does not hold unless R_2 contains $*$. This is because of the restriction, $\forall R \in \mathbb{R} : \text{getAlphabetSet}(R) \not\subseteq \Sigma$, that we defined in

the previous section; thanks to the restriction, it is impossible without \star to make syntactically different but semantically equivalent regular expressions with $\{\star\}$. We found that the restriction makes the decision rules and related proofs simpler with the smaller time complexity than without the restriction. The decision rules exactly capture the inclusion relation between regular expressions in the domain:

THEOREM 6. (*Soundness and completeness of $R_1 \leq_r R_2$*)
 $\forall R_1 \in \mathbb{R}, R_2 \in \mathbb{R} : \llbracket R_1 \rrbracket \subseteq \llbracket R_2 \rrbracket \iff R_1 \leq_r R_2.$

Now, we state the complexity of the decision rules. First, auxiliary functions $\text{pruneHead}(a, R)$ and $\text{pruneNonStar}(R)$ can be computed in $O(|R|)$ time by choosing a suitable set data structure whose membership and union operations take $O(1)$ and $O(n)$ time, respectively, when n is the size of the set. Consider a relation $\{A\} \leq_r R$. If $|A| \leq 1$, then we can decide it in $O(1)$ time (by `INCLEMPTYSET` or `INCLEMPTY`). Otherwise, $|A| \geq 2$ and we use the rule `INCLAREG` or `INCLSTARREG`. The maximum size of the set obtained by repeatedly applying pruneHead and pruneNonStar to R is bounded by $|R| \cdot \max_{A' \in R} |A'|$ where $\max_{A' \in R} |A'|$ chooses the maximum of the lengths of atomic regular expressions in R . Since the size of A decreases by 1 in the rules `INCLAREG` and `INCLSTARREG` and both pruneHead and pruneNonStar functions can be computed in $O(|R| \cdot \max_{A' \in R} |A'|)$ time, we can decide a relation $\{A\} \leq_r R$ in $O(|A| \cdot |R| \cdot \max_{A' \in R} |A'|)$ time. Therefore, we can decide a relation $R_1 \leq_r R_2$ in $O(|R_1| \cdot \max_{A \in R_1} |A| \cdot |R_2| \cdot \max_{A' \in R_2} |A'|)$ time in the worst case. More detailed and formal analysis for the time complexity is available in our companion report [12].

6. Evaluation

In this section, we evaluate our new domain with its implementation in terms of analysis scalability and precision.

6.1 Methodology

Implementation. The prototype implementation of our domain extends SAFE [19, 20], which is a flow-sensitive and context-sensitive static analyzer for JavaScript web applications that give approximate heap states at each program point as analysis results. While SAFE supported only a variant of a constant string domain at first with two special elements, `NumStr` and `OtherStr`, that abstract all number strings and all non-number strings, respectively, it has been extended later to support a set domain whose elements are \top_s (any string), \perp_s (no string), `NumStr`, `OtherStr`, and a set with the maximum k number of concrete strings. Note that the formalization of the regular expression domain in Section 4 does not contain `NumStr` and `OtherStr` but the implementation incorporates them, which can be also incorporated in the formalization in a similar way that \star is added in the domain.

While the regular expression domain can represent all elements in the set domain of SAFE so that it does not require any modification in the existing abstract semantics of string operations, we refined 5 operations to utilize the

enhanced expressivity of the new domain elements: string concatenation, object property read, object property write, string conversion from an object, and string conversion from a function. We refined the first three operations as we described in Section 4. The last two operations return string values resulted from string conversion of an object and a function, respectively. According to the ECMAScript specification [1], the conversion for an object returns “[Object classname]” where `classname` is the internal class name of the object, and the conversion for a function returns a source string for the function that starts with “function” followed by an implementation-dependent string. Since SAFE does not fully support internal class names of objects and implementation-dependent features, it conservatively models them as \top_s in the set domain. However, we refined them as “[Object \star]” and “function \star ” in the new domain. Note that the refinements improve the precision of the existing abstract semantics while preserving its soundness. Other string operations can be refined further to achieve even more improvements but we found that the refinements in the 5 operations already show significant improvements in analysis scalability and precision for SAFE. Throughout this section, we denote by `SAFEreg` the implementation with the new domain and distinguish it from the original SAFE. We make the implementation open to the public [12].

Evaluation Methods. For evaluation target programs, we use a program group `BENCH` that includes 61 JavaScript programs from a jQuery tutorial². The literature [3, 19] also used `BENCH` to evaluate JavaScript static analyzers. However, while all 61 programs in the previous work load a modified version of jQuery 1.10.0 where some statically indeterminate values are replaced with constant values, our target programs all use non-modified jQuery 1.7.0 because it is the highest version that both SAFE and TAJs can successfully finish analysis within the timeout of 10 minutes.

For analysis scalability, we compare the analysis results of our target programs with three static analyzers, SAFE, `SAFEreg`, and TAJs. We refer interested readers to Section 7 for detailed explanation about SAFE and TAJs. As in the evaluation of previous work [3, 19], we run three analyzers with the target programs in the timeout of 10 minutes and count the number of programs whose analyses successfully complete within the timeout with normal exit heaps. In addition, we measure analysis time in analyzing 10 sample programs in `BENCH` to check how the new domain affects the analysis performance of SAFE in detail.

For analysis precision, we compare SAFE and `SAFEreg` with the analysis results of 10 sample programs. We exclude TAJs for the comparison of analysis precision because fair measurement of analysis precision requires knowledge of analysis internals. We use the following three criteria that a previous work [20] used to measure analysis precision:

²<http://www.jquery-tutorial.net>

Group (# of programs)	Success (#)		
	SAFE	SAFE _{reg}	TAJS
BENCH (61)	27	40	2

Table 1: Numbers of successful analysis results in 10 min.

- MD : the number of object dereference points with multiple object values
- MC : the number of calls with multiple function values
- PR : the number of object property accesses with a property name approximated as non-constant

Lower MD, MC, and PR mean more precise analysis results. For example, for a method invocation `obj[x]()`, if the analysis result of the object dereference point `obj` contains multiple objects, MD includes it; if the analysis result of the method `obj[x]` contains multiple functions, MC includes it; if the analysis result of the object property name `x` is not a constant, PR includes it.

We performed all experiments on a Linux x64 machine with 3.4GHz Intel Core i7 CPU and 32GB memory, and all figures in the results are averaged over those from 5 runs.

6.2 Experimental Results

Scalability. Table 1 shows how many programs among 61 target programs in the BENCH group SAFE, SAFE_{reg}, and TAJS successfully finish analysis within the timeout of 10 minutes. While SAFE and TAJS complete the analysis of 27 and 2 programs, respectively, SAFE_{reg} performs the best with the complete analysis results of 40 programs. Table 2 shows detailed analysis results of 10 sample programs in the BENCH group. The first 5 programs are those that both SAFE and SAFE_{reg} successfully analyze within the timeout and the last 5 are those that only SAFE_{reg} successfully analyzes. The second and sixth columns of the table show analysis time that SAFE and SAFE_{reg} took to analyze each target program. We denote the timeout result by **X**. On average, the new domain reduced the analysis time of SAFE from 101.86 to 65.05 seconds for the first 5 programs and from minimum 10 minutes to 198.75 seconds for the rest 5 programs. Note that `4.html` is the motivating example that we presented in Section 2 and we described that its analysis did not complete within 10 hours in SAFE. The regular expression domain can cause performance overhead as we discuss in the next section. However, this result demonstrates that the new domain significantly improves the scalability of an existing analyzer in analyzing jQuery programs rather than causing extra performance overhead.

Precision. Table 2 shows analysis precision of SAFE and SAFE_{reg} in analyzing 10 sample programs. For the first 5 programs that SAFE successfully analyzes within the timeout, SAFE_{reg} shows relatively small improvement compared to the results for the rest 5 programs because SAFE already shows reasonable analysis precision. While MC remains the

same in both analyzers, MD and PR decreased from 21 and 13 to 12 and 5, respectively, on average in SAFE_{reg}. On the other hand, for the rest 5 programs that only SAFE_{reg} successfully analyzes within the timeout, SAFE_{reg} shows great improvement. On average, it reduces MD, MC, and PR from 218, 50, and 40 to 23, 3, and 16, respectively. This result demonstrates not only that improving analysis precision also improves analysis scalability but also that the regular expression domain significantly improves analysis precision of an existing analyzer. Note that refining more string operations may further enhance the improvement in analysis scalability and precision shown in this section.

6.3 Discussion

Limitations and threats to validity. Analysis results in the regular expression domain are always more precise than or as precise as those in constant and set string domains since the domain includes all elements in the other two domains with additional regular expression elements that refine the τ_s element of the other two domains. However, generally speaking, the domain can cause extra performance overhead; while the decision complexity for order relations between elements in constant and set domains is constant and linear, respectively, the complexity for the regular expression domain is biquadratic depending on the number and the maximum length of atomic regular expressions. But because more imprecise results can incur more performance overhead, if the performance benefits from improved precision are bigger than the performance overhead by the order relation decision, the new domain leads to faster analysis than the other two. We demonstrated such cases in the previous section.

A possible threat to validity of our results is that we used simple programs that use jQuery on the experiment and thus conclusions drawn from the results may not hold in real-world applications. However, even state-of-the-art static analyzers such as SAFE and TAJS are not able to analyze such simple programs because of the dynamic natures of jQuery. We believe that making existing analysis scalable enough to analyze simple programs is an essential step towards scalable analysis of real-world applications, and our technique contributes to it. Another threat is that analysis results regarded as successful in SAFE and TAJS may not be sound because of possible unsoundness in the implementation of SAFE and TAJS. Although SAFE and TAJS provide comprehensive models of browser environments [10, 20], analysis of programs that use unmodeled features may yield unsound results. However, since it is orthogonal to the technique presented in this paper, we assume that the base analyzers provide sound analysis results for target programs.

Other causes for scalability problems. We investigated why SAFE_{reg} still could not complete the analysis of some programs within 10 minutes and observed two reasons: 1) a sound event system in SAFE and 2) statically indeterminate values in AJAX communication. The SAFE event

Target Program	SAFE				SAFE _{reg}			
	Time (s)	MD (#)	MC (#)	PR (#)	Time (s)	MD (#)	MC (#)	PR (#)
1.html	120.07	24	3	12	82.72	16	3	5
10.html	112.90	24	3	12	75.73	16	3	5
20.html	93.72	20	3	12	57.04	9	3	5
30.html	93.86	18	3	12	57.54	9	3	5
54.html	88.75	17	3	16	52.23	9	3	6
Average	101.86	21	3	13	65.05	12	3	5
4.html	×	321	159	41	483.60	16	3	8
29.html	×	208	55	29	113.83	23	3	10
35.html	×	200	32	42	99.58	26	3	22
57.html	×	180	3	43	147.93	25	3	20
60.html	×	181	3	43	148.84	25	3	20
Average	×	218	50	40	198.75	23	3	16

Table 2: Analysis time and precision of SAFE and SAFE_{reg} for 10 jQuery tutorial programs with the timeout of 10 minutes.

system that approximates all possible execution sequences of registered event handlers and AJAX communication that receives arbitrary values produce over-approximate values, which cause scalability problems. We believe that more elaborate event analysis systems and developers’ hints for values from AJAX communication would alleviate the problems.

Analysis of non-jQuery programs. We further evaluated our regular expression domain with 21 non-jQuery benchmarks from a previous work [13]; the benchmarks consist of 7 benchmarks from SunSpider and V8, 7 Firefox browser addons, and 7 programs from open source JavaScript frameworks and their test suites³. While the regular expression domain can cause performance overhead, we found that it only causes negligible overhead in analyzing the target benchmarks. Both SAFE and SAFE_{reg} could not complete 3 programs within the timeout of 10 minutes, but the average analysis time over the results finished in 10 minutes were 49.02 (min.: 0.64, max.: 575.58, median: 4.04, standard deviation: 136.20) and 49.45 (min.: 0.63, max.: 576.67, median: 4.09, standard deviation: 136.37) seconds in SAFE and SAFE_{reg}, respectively. As for precision, the analysis results from all targets showed the same MD, MC, and PR in SAFE and SAFE_{reg}. We believe that this is because we refined only 5 string operations where only 3 operations introduce regular expression elements and the targets do not use them much.

We found that evaluating our domain with real-world applications such as web pages is premature because analysis of such applications involves many challenges such as AJAX communication and heavy user interaction by events, which are beyond the scope of this paper. Although static analysis for JavaScript does not scale to real-world applications yet, we believe that our domain itself in SAFE_{reg} would not cause much performance overhead compared to constant and set domains in deciding order relations between elements.

As we discussed before, the complexity of the order decision depends on the number and maximum length of atomic regular expressions, which we expect to be small numbers due to only 3 operations that introduce simple regular expressions. However, with more operations that can introduce complex regular expressions, the domain can cause more overhead, where it is possible to restrict such numbers as some maximum values to minimize the complexity in the domain.

7. Related Work

JavaScript static analysis frameworks. SAFE [14, 19], JSAI [13], and TAJs [3, 9] are JavaScript static analyzers that produce approximate heap information at each program point. All of them faithfully model the ECMAScript semantics with some built-in function modeling. In addition, SAFE and TAJs provide a comprehensive modeling of browser environments [10, 20] to analyze web applications. While they use different lattice domains for abstract values, we formalized the conceptually common model of strings, objects, and main operations on them in Section 3.

SAFE and TAJs have made significant progress in analyzing jQuery. Loop-sensitive analysis (LSA) of SAFE [19] significantly improves the analysis precision and the scalability enough to analyze all released versions of jQuery within a practical time limit. We have applied the new domain to SAFE with LSA, and we showed that the domain further improves its precision and scalability. TAJs [3] can analyze 11 out of all 12 released jQuery 1.x versions within one minute. But the experiments used modified versions of jQuery where some statically indeterminate values like `Math.random()` were replaced with constant values. We found that TAJs can analyze 4 out of 12 jQuery 1.x versions within 10 minutes when it analyzes non-modified jQuery.

WALA [22, 23] supports static pointer analysis of Java and JavaScript programs. It can analyze 3 versions of jQuery. While we can apply our regular expression domain to WALA

³<http://www.defensivejs.com>, <http://linqjs.codeplex.com>

as well, we did not consider WALA because SAFE and WALA provide different information as analysis results as heap approximation and pointer information, respectively.

Dynamic object and jQuery typing. Politz *et al.* [21] presented a type system for dynamic objects commonly used in dynamic languages such as JavaScript, Python, and Ruby to guarantee that well-typed programs do not read absent properties of objects during execution. Their object type is a map from property names to property types where the property names can be any patterns such as regular expressions unlike our object model that has concrete string names for properties. Their type system parameterizes the pattern domain with a restriction that the inclusion between patterns should be decidable. Our regular expression domain is an instance of the type system satisfying the restriction.

Lerner *et al.* [15] proposed a type system for jQuery considering it as a language. The type system has two novel concepts, *multiplicity* that represents the possible size of object collections and *local structure* that represents types only for local parts of a web document that the target jQuery program uses. The type system detects misuses of jQuery APIs such as operations on empty object collections. However, the type system requires user annotation for the local structure, which subsequently requires understanding of the interaction between web documents and jQuery programs. While our analysis does not require any user annotation, it has higher time complexity by fixpoint computation than the type system and it cannot handle events precisely.

String and regular expression domains. Madsen and Andreasen [16] suggested 12 string domains, among which 7 are new, for precise analysis of dynamic property accesses in JavaScript. They showed that a hybrid domain of string set, character inclusion, and string hash domains produced the best analysis precision and scalability. However, even the hybrid domain cannot represent string values resulted from concatenation of a constant string with multiple characters and a statically indeterminate value, which are frequently used as keys to dynamic property accesses in jQuery.

Researchers have used regular expressions for Java string analysis. Christensen *et al.* [5] presented a string analysis that produces, for each string operation in a program, a regular expression that approximates string values resulted from the operation. The analysis represents a program as a def-use graph of string data in string operations and constructs a multi-level finite automaton (MLFA) using the context-free and regular grammars translated from the graph. Multiple deterministic finite-state automata (DFA) can be derived for each string operation from MLFA and then the regular expression that each DFA represents is an analysis result for the result string values of the operation. The analysis does not consider inclusion relations between regular expressions during analysis because regular expressions appear only as analysis results. However, the complexity of extracting DFA is doubly exponential in the worst case while our domain re-

quires the polynomial-time complexity. In addition, the analysis requires a def-use graph for a target program, which is still a research problem for JavaScript.

Choi *et al.* [4] formalized a string analysis using regular expressions in abstract interpretation like our approach. Their regular expression domain is more expressive than ours in that their domain allows the Kleene star operation (*) on arbitrary combinations of regular expressions while ours allows it only on disjunctions over all alphabets. They also presented complex widening rules to minimize precision loss by join operations on regular expressions. While we define the order relation between regular expressions using the inclusion relation between them, they do not present how to decide the relation. Note that, without the order relation decision, the analysis may perform unnecessary widening for regular expressions already in the order relation.

Regular expression inclusion. To use regular expressions in a sound way as elements of abstract string domains, one should define the order relation between regular expressions using the inclusion relation between them. If one regular expression R_1 includes another R_2 , it means that a set of all strings that R_1 matches includes a set of all strings that R_2 matches. However, deciding inclusion for general regular expressions has been known as a problem in PSPACE-complete [24]. Fortunately, previous work [2, 18] has found tractable cases for some restricted subset of regular expressions and Martens *et al.* [17] studied the complexity of inclusion for various subsets of regular expressions including previously known tractable cases. According to the study, surprisingly, inclusion decision between elements in a simple domain whose element is a sequence of a and $(a_1|...|a_n)^*$ (a , a_i , and a_n are alphabets with $n \geq 1$ and $|$ and $*$ are usual disjunction and Kleene star operations) already reaches the maximum complexity, PSPACE-complete. To the best of our knowledge, the complexity of inclusion for our regular expressions is previously unknown, and we presented polynomial-time inclusion decision rules for them.

Other applications. Regular expressions are widely used in various applications including pattern matching with strings. Regular expressions for string matching are often called regexes and supported by many languages such as Perl, JavaScript, and Ruby. Regexes are used to detect SQL injection attacks⁴ by checking if a user input contains SQL specific meta-characters and commands. Because most regexes implementations include non-regular features, they are more expressive than ours but with higher complexity.

Another well-studied example is XML type checking using regular expression types to describe XML types such as DTD and XML Schema. Regular expression types [7] consist of primitive types, pair types, recursive types, and union types, and they subsume our regular expressions. While we

⁴<http://www.symantec.com/connect/articles/detection-sql-injection-and-cross-site-scripting-attacks>

can use regular expression types, the complexity of their inclusion relation is much higher as EXPTIME-complete. Instead, we devised a subclass of regular expressions that is simple but powerful enough for our problem and whose inclusion relation can be efficiently decided in PTIME.

8. Conclusion

We proposed a novel abstract string domain whose elements are simple regular expressions for precise and scalable static analysis of jQuery, the most popular JavaScript library. The domain is expressive enough to represent prefix, infix, and postfix substrings of a string and even their sets. In a theoretical aspect, we formalized the domain and related operations in abstract interpretation with the soundness proofs. In a practical aspect, we formally presented polynomial-time inclusion decision rules between the regular expressions with the soundness and completeness theorems for the implementation of the domain. We have implemented the domain as an extension of SAFE and demonstrated that the technique significantly improves the scalability and precision of SAFE in analyzing 61 programs that use jQuery. We make the formal proofs and the implementation open to the public.

Acknowledgments

This work was supported in part by Korea Ministry of Education, Science and Technology (MEST) / National Research Foundation of Korea (NRF) (Grants NRF-2014R1A2A2A01003235 and 2016R1C1B1015095) and Samsung Electronics.

References

- [1] ECMAScript Language Specification. Edition 5.1, 2011.
- [2] P. A. Abdulla, A. Bouajjani, and B. Jonsson. On-the-fly analysis of systems with unbounded, lossy fifo channels. In *Proceedings of the 10th International Conference on Computer Aided Verification*. Springer Berlin Heidelberg, 1998.
- [3] E. Andreassen and A. Møller. Determinacy in static analysis for jQuery. In *Proceedings of the International Conference on Object Oriented Programming Systems Languages and Applications*, 2014.
- [4] T.-H. Choi, O. Lee, H. Kim, and K.-G. Doh. A practical string analyzer by the widening approach. In *Proceedings of the Asian Symposium on Programming Languages and Systems*, 2006.
- [5] A. S. Christensen, A. Møller, and M. I. Schwartzbach. Precise analysis of string expressions. In *Proceedings of the International Symposium on Static Analysis*, 2003.
- [6] P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proceedings of the Symposium on Principles of Programming Languages*, 1977.
- [7] H. Hosoya, J. Vouillon, and B. C. Pierce. Regular expression types for XML. *ACM Transactions on Programming Languages and Systems*, 27(1):46–90, 2005.
- [8] D. Hovland. The inclusion problem for regular expressions. *Journal of Computer and System Sciences*, 78(6), 2012.
- [9] S. H. Jensen, A. Møller, and P. Thiemann. Type analysis for JavaScript. In *Proceedings of the International Symposium on Static Analysis*. Springer-Verlag, 2009.
- [10] S. H. Jensen, M. Madsen, and A. Møller. Modeling the HTML DOM and browser API in static analysis of JavaScript web applications. In *Proceedings of the International Symposium on Foundations of Software Engineering*. ACM, 2011.
- [11] jQuery Foundation. jQuery. <http://jquery.com>.
- [12] KAIST PLRG. <http://plrg.kaist.ac.kr/pch>.
- [13] V. Kashyap, K. Dewey, E. A. Kuefner, J. Wagner, K. Gibbons, J. Sarracino, B. Wiedermann, and B. Hardekopf. JSAI: A static analysis platform for JavaScript. In *Proceedings of the International Symposium on Foundations of Software Engineering*, 2014.
- [14] H. Lee, S. Won, J. Jin, J. Cho, and S. Ryu. SAFE: Formal specification and implementation of a scalable analysis framework for ECMAScript. In *Workshop on Foundations of Object Oriented Languages*, 2012.
- [15] B. S. Lerner, L. Elberty, J. Li, and S. Krishnamurthi. Combining form and function: Static types for jQuery programs. In *Proceedings of the European Conference on Object-Oriented Programming*, 2013.
- [16] M. Madsen and E. Andreassen. String analysis for dynamic field access. In *Proceedings of the International Conference on Compiler Construction*, 2014.
- [17] W. Martens, F. Neven, and T. Schwentick. Complexity of decision problems for simple regular expressions. In *Proceedings of the International Symposium on Mathematical Foundations of Computer Science*, 2004.
- [18] T. Milo and D. Suciu. Index structures for path expressions. In *Proceedings of the International Conference on Database Theory*, 1999.
- [19] C. Park and S. Ryu. Scalable and precise static analysis of JavaScript applications via loop-sensitivity. In *Proceedings of the European Conference on Object-Oriented Programming*, 2015.
- [20] C. Park, S. Won, J. Jin, and S. Ryu. Static analysis of JavaScript web applications in the wild via practical DOM. In *Proceedings of the International Conference on Automated Software Engineering*, 2015.
- [21] J. G. Politz, A. Guha, and S. Krishnamurthi. Semantics and types for objects with first-class member names. In *Workshop on Foundations of Object Oriented Languages*, 2012.
- [22] M. Schäfer, M. Sridharan, J. Dolby, and F. Tip. Dynamic determinacy analysis. In *Proceedings of the Conference on Programming Language Design and Implementation*, 2013.
- [23] M. Sridharan, J. Dolby, S. Chandra, M. Schäfer, and F. Tip. Correlation tracking for points-to analysis of JavaScript. In *Proceedings of the European Conference on Object-Oriented Programming*, 2012.
- [24] L. J. Stockmeyer and A. R. Meyer. Word problems requiring exponential time (preliminary report). In *Proceedings of the Annual ACM Symposium on Theory of Computing*, 1973.