JSART: JavaScript Assertion-Based Regression Testing

Shabnam Mirshokraie and Ali Mesbah

University of British Columbia
Vancouver, BC, Canada
{shabnam, amesbah}@ece.ubc.ca

Abstract. Web 2.0 applications rely heavily on JAVASCRIPT and client-side runtime manipulation of the DOM tree. One way to provide assurance about the correctness of such highly evolving and dynamic applications is through regression testing. However, JAVASCRIPT is loosely typed, dynamic, and notoriously challenging to analyze and test. We propose an automated technique for JAVASCRIPT regression testing, which is based on on-the-fly JAVASCRIPT source code instrumentation and dynamic analysis to infer invariant assertions. These obtained assertions are injected back into the JAVASCRIPT code to uncover regression faults in subsequent revisions of the web application under test. Our approach is implemented in a tool called JSART. We present our case study conducted on nine open source web applications to evaluate the proposed approach. The results show that our approach is able to effectively generate stable assertions and detect JAVASCRIPT regression faults with a high degree of accuracy and minimal performance overhead.

Keywords: JavaScript, web, regression testing, assertions, dynamic analysis.

1 Introduction

JAVASCRIPT is increasingly being used to create modern interactive web applications that offload a considerable amount of their execution to the client-side. JAVASCRIPT is a notoriously challenging language for web developers to use, maintain, analyze and test. It is dynamic, loosely typed, and asynchronous. In addition, it is extensively used to interact with the DOM tree at runtime for user interface state updates.

Web applications usually evolve fast by going through rapid development cycles and are, therefore, susceptible to regressions, i.e., new faults in existing functionality after changes have been made to the system. One way of ensuring that such modifications (e.g., bug fixes, patches) have not introduced new faults in the modified system is through systematic regression testing. While regression testing of classical web applications has been difficult [25], dynamism and non-determinism pose an even greater challenge [22] for Web 2.0 applications.

In this paper, we propose an automated technique for JAVASCRIPT regression testing, which is based on dynamic analysis to infer invariant assertions. These obtained assertions are injected back into the JAVASCRIPT code to uncover regression faults in subsequent revisions of the web application under test. Our technique automatically (1) intercepts and instruments JAVASCRIPT on-the-fly to add tracing code (2) navigates the web application to produce execution traces, (3) generates dynamic invariants from the

```
1  function setDim(height, width) {  
2      var h = 4*height, w = 2*width;  
3      ...  
4      return {h:h, w:w};  
5  }  
  
7  function play() {  
8      $('#end').css("height", setDim($('#body').width(), $('#body').height()).↵  
9          h + 'px');  
10     ...  
11 }
```

Fig. 1. Motivating JAVASCRIPT example

trace data, (4) transforms the invariants into stable assertions and injects them back into the web application for regression testing.

Our approach is orthogonal to server-side technology, and it requires no manual modification of the source code. It is implemented in an open source tool called JSART (JAVASCRIPT Assertion-based Regression Testing). We have empirically evaluated the technique on nine open-source web applications. The results of our evaluation show that the approach generates stable invariant assertions, which are capable of spotting injected faults with a high rate of accuracy.

2 Motivation and Challenges

Figure 1 shows a simple JAVASCRIPT code snippet. Our motivating example consists of two functions, called `setDim` and `play`. The `setDim` function has two parameters, namely `height` and `width`, with a simple mathematical operation (Line 2). The function returns local variables, `h` and `w` (Line 4). `setDim` is called in the `play` function (Line 8) to set the `height` value of the CSS property of the DOM element with ID `end`. Any modification to the values of `height` or `width` would influence the returned values of `setDim` as well as the property of the DOM element. Typical programmatic errors include swapping the order of `height` and `width` when they are respectively assigned to local variables `h` and `w` or calling `setDim` with wrong arguments, i.e., changing the order of function arguments.

Detecting such regression errors is a daunting task for web developers, especially in programming languages such as JAVASCRIPT, which are known to be challenging to test. One way to check for these regressions is to define invariant expressions of expected behaviour [17] over program variables and assert their correctness at run-time. This way any modification to `height`, `width`, `h`, or `w` that violates the invariant expression will be detected. However, manually expressing such assertions for web applications with thousands of lines of JAVASCRIPT code and several DOM states, is a challenging and time-consuming task. Our aim in this work is to provide a technique that automatically captures regression faults through generated JAVASCRIPT assertions.

3 Our Approach

Our regression testing approach is based on *dynamic analysis* of JAVASCRIPT code to infer invariants from a given web application. We use the thus obtained invariants as

Fig. 2. Overview of the JAVASCRIPT tracing and invariant generation steps (web application version n)

runtime assertions in the JAVASCRIPT code to automatically uncover regression errors that can be introduced after changes have been made to the web application in a subsequent reversion. Our approach is largely based on two assumptions (1) the current version of the web application, from which invariants are being generated, is bug-free (2) the inferred invariants capture program specifications that are unlikely to change frequently in the following revisions (we revisit these two assumptions in Section 6). Our regression testing technique is composed of the following four main steps: (1) JavaScript tracing, (2) Invariant generation, (3) Filtering unstable assertions, and (4) Regression testing through assertions. In the following subsections, we will describe each step in details.

3.1 JAVASCRIPT Tracing

In order to infer useful program invariants, we need to collect execution traces of the JAVASCRIPT code. The idea is to log as much program variable value changes at runtime as possible. Figure 2 depicts a block diagram of the tracing step. Our approach automatically generates trace data in three subsequent steps: (i) JAVASCRIPT interception and instrumentation, (ii) navigation, and (iii) trace collection. In the following, we explain each step in details.

JAVASCRIPT Interception and Instrumentation. The approach we have chosen for logging variables is on-the-fly JAVASCRIPT source code transformation to add instrumentation code. We intercept all the JAVASCRIPT code of a given web application, both in JAVASCRIPT files and HTML pages, by setting up a proxy [3] between the server and the browser. We first parse the intercepted source code into an Abstract Syntax Tree (AST). We then traverse the AST in search of *program variables* as well as *DOM modifications* as described below.

Tracing Program Variables. Our first interest is the range of values of JAVASCRIPT program variables. We probe function entry and function exit points, by identifying function definitions in the AST and injecting statements at the start, end, and before every return statement. We instrument the code to monitor value changes of *global variables*, *function arguments*, and *local variables*. Per program point, we yield information on *script name*, *function name*, and *line number*, used for debugging purposes. Going back to our running example (Figure 1), our technique adds instrumentation code

to trace width, height, h, and w. For each variable, we collect information on *name*, *runtime type*, and *actual values*. The runtime type is stored because JAVASCRIPT is a loosely typed language, i.e., the types of variables cannot be determined syntactically, thus we log the variable types at runtime.

Tracing DOM Modifications. In modern web applications, JAVASCRIPT code frequently interacts with the DOM to update the client-side user interface state. Our recent study [18] of four bug-tracking systems indicated that DOM-related errors form 80% of all reported JAVASCRIPT errors. Therefore, we include in our execution trace how DOM elements and their attributes are modified by JAVASCRIPT at runtime. For instance, by tracing how the CSS property of the ‘end’ DOM element in Figure 1 is changed during various execution runs, we can infer the range of values for the height attribute.

Based on our observations, JAVASCRIPT DOM modifications usually follow a certain pattern. Once the pattern is reverse engineered, we can add proper instrumentation code around the pattern to trace the changes. In the patterns that we observed, first a JAVASCRIPT API is used to find the desired DOM element. Next, a function is called on the returned object responsible for the actual modification of the DOM-tree. After recognizing a pattern in the parsed AST, we add instrumentation code that records the value of the DOM attribute before and after the actual modification. Hence, we are able to trace DOM modifications that happen programmatically through JAVASCRIPT.

Navigation. Once the AST is instrumented, we serialize it back to the corresponding JAVASCRIPT source code file and pass it to the browser. Next, we navigate the application in the browser to produce an execution trace. The application can be navigated in different ways including (1) manual clicking (2) test case execution (3) or using a web crawler. To automate the approach, our technique is based on automated dynamic crawling [15]. The execution needs to run as much of the JAVASCRIPT code as possible and execute it in various ways. This can be achieved through visiting as many DOM state changes as possible as well as providing different values for function arguments.

Trace Collection. As the web application is navigated, the instrumented JAVASCRIPT code produces trace data, which needs to be collected for further analysis. Keeping the trace data in the browser’s memory during the program execution can make the browser slow when a large amount of trace data is being produced. On the other hand, sending data items to the proxy as soon as the item is generated, can put a heavy load on the proxy, due to the frequency of HTTP requests. In order to tackle the aforementioned challenges, we buffer a certain amount of trace data in the memory in an array, post the data as an HTTP request to a proxy server when the buffer’s size reaches a predefined threshold, and immediately clear the buffer in the browser’s memory afterwards. Since the data arrives at the server in a synchronized manner, we concatenate the tracing data into a single trace file on the server side, which is then seeded into the next step (See Figure 2).

3.2 Invariant Generation

The assertion generation phase is involved with analyzing the collected execution traces to extract invariants. Substantial amount of research has been carried out on detecting dynamic program invariants [5,7,9,13]. Our approach is based on Daikon [9] to infer

Fig. 3. Overview of the filtering step to remove unstable invariant assertions, for web application version n

likely invariants. As indicated with the dotted line in Figure 2, we cycle through the navigation and invariant generation phases until the size of generated invariant file remains unchanged, which is an indication that all possible invariants have been detected.

3.3 Filtering Unstable Invariant Assertions

The next step is to make sure that the generated invariants are truly invariants. An invariant assertion is called unstable when it is falsely reported as a violated assertion. Such assertions result in producing a number of false positive errors during the testing phase. To check the stability of the inferred invariants, we use them in the same version of the web application as assertions. Theoretically, no assertion violations should be reported because the web application has not changed. Hence, any assertion violation reported as such is a false positive and should be eliminated. Our filtering process, shown in Figure 3, consists of the following four processes:

- Converting the inferred invariants into checkable assertions;
- Injecting the invariant assertions into the same version of the web application;
- Navigating the web application;
- Collecting assertion violations and removing them;

From each of the inferred invariants, we generate an assertion in JAVASCRIPT format. We use on-the-fly transformation to inject the assertions directly into the source code of the same version of the web application. Since we have all the information about the program points and the location of the assertions, we can inject the assertions at the correct location in the JAVASCRIPT source code through the proxy, while the code is being sent to the client by the server. This way the assertions can run as part of the client-side code and gain access to the values of all program variables needed at runtime. Once the assertions are injected, we execute the web application in the browser and log the output. Next we collect and remove any violated assertions. The output of this step is a set of *stable* invariant assertions, used for automated regression testing in the next step.

Fig. 4. Overview of the JavaScript regression testing step through invariant assertions, for web application version $n+1$

3.4 Regression Testing through Assertions

Once a set of stable invariant assertions are derived from version n of a web application, they can be used for automatic regression testing a subsequent version $(n+1)$ of the web application. The regression testing phase is depicted in Figure 4.

We inject the inferred stable assertions to the JAVASCRIPT source code of the modified web application, in a similar fashion to the filtering step in Section 3.3. Once the assertions are injected, the new version of the web application is ready for regression testing. Any failed assertion during the testing phase generates an entry in the test report, which is presented to the tester at the end of the testing step. The generated test report provides precise information on the failed assertion, the file name, the line number, and the function name of the assertion.

```

1  function setDim(height, width) {
2      assert((width < height), 'example.js:setDim:ENTER:POINT1');
3      var h = 4*height, w = 2*width;
4      ...
5      assert((width < height), 'example.js:setDim:EXIT:POINT1');
6      assert((w < h), 'example.js:setDim:EXIT:POINT1');
7      return{h:h, w:w};
8  }

10 function play(){
11     $('#end').css('height', setDim($('#body').width(), ($('#body').height()).h + ←
12         'px');
13     assert(isIn($('#end').css('height'), {100, 200,
14         300}), 'example.js:play:POINT3');
15 }

```

Fig. 5. Invariant assertion code for JAVASCRIPT function parameters, local variables and DOM modifications. Injected assertions are shown in bold.

Figure 5 shows the automatically injected invariant assertions for our running example of Figure 1. Note that we do not show all the assertions as they clutter the figure. Each `assert` call has the invariant as the first parameter and the corresponding debugging information in the second parameter, which includes information about script name, function name, and line number. In this example, the inferred invariants yield information about the inequality relation between function arguments, `width` and `height`, as well as local variables, `w` and `h`. The assertions in lines 2 and 5-6 check the

corresponding inequalities, at entry and exit points of the `setDim` function at runtime. The example also shows the assertion that checks the `height` attribute of the DOM element, after the JAVASCRIPT DOM modification in the `play` function. The assertion that comes after the DOM manipulation (Line 12) checks the `height` value by calling the auxiliary `isIn` function. `isIn` checks the value of `height` to be in the given range, i.e., either 100, 200, or 300. Any values out of the specified range would violate the assertion.

4 Tool Implementation

We have implemented our JAVASCRIPT regression testing approach in a tool called JSART. JSART is written in Java and is available for download.¹

JSART extends and builds on top of our InvarScope [10] tool. For JAVASCRIPT code interception, we use an enhanced version of Web-Scarab's proxy [3]. This enables us to automatically analyze and modify the content of HTTP responses before they reach the browser. To instrument the intercepted code, Mozilla Rhino² is used to parse JAVASCRIPT code to an AST, and back to the source code after instrumentation. The AST generated by Rhino's parser has traversal API's, which we use to search for program points where instrumentation code needs to be added. For the invariant generation step, we have extended Daikon [9] with support for accepting input and generating output in JAVASCRIPT syntax. The input files are created from the trace data and fed through the enhanced version of Daikon to derive dynamic invariants. The navigation step is automated by making JSART operate as a plugin on top of our dynamic AJAX crawler, CRAWLJAX [15].³

5 Empirical Evaluation

To quantitatively assess the accuracy and efficiency of our approach, we have conducted a case study following guidelines from Runeson and Höst [23]. In our evaluation, we address the following research questions:

RQ1 How successful is JSART in generating stable invariant assertions?

RQ2 How effective is our overall regression testing approach in terms of correctly detecting faults?

RQ3 What is the performance overhead of JSART?

The experimental data produced by JSART is available for download.¹

5.1 Experimental Objects

Our study includes nine web-based systems in total. Six are game applications, namely, SameGame, Tunnel, TicTacToe, Symbol, ResizeMe, and GhostBusters. Two of the web

¹ <http://salt.ece.ubc.ca/content/jsart/>

² <http://www.mozilla.org/rhino/>

³ <http://www.crawljax.com>

Table 1. Characteristics of the experimental objects

App ID	Name	JS LOC	# Functions	# Local Vars	# Global Vars	CC	Resource
1	SameGame	206	9	32	5	37	http://crawljax.com/same-game
2	Tunnel	334	32	18	13	39	http://arcade.christianmontoya.com/tunnel
3	TicTacToe	239	11	22	23	83	http://www.dynamicdrive.com/dynamicindex12/tictactoe.htm
4	Symbol	204	20	28	16	32	http://10k.aneventapart.com/2/Uploads/652
5	ResizeMe	45	5	4	7	2	http://10k.aneventapart.com/2/Uploads/594
6	GhostBusters	277	27	75	4	52	http://10k.aneventapart.com/2/Uploads/657
7	Jason	107	8	4	8	6	http://jasonjulien.com
8	Sofa	102	22	2	1	5	http://www.madebysofa.com/archive
9	TuduList	2767	229	199	31	28	http://tudu.ess.ch/tudu

applications are Jason and Sofa, which are a personal and a company homepage, respectively. We further include TuduList, which is a web-based task management application. All these applications are open source and use JAVASCRIPT on the client-side.

Table 1 presents each application’s ID, name, and resource, as well as the characteristics of the custom JAVASCRIPT code, such as JAVASCRIPT lines of code (LOC), number of functions, number of local and global variables, as well as the cyclomatic complexity (CC). We use Eclipse IDE to count the JAVASCRIPT lines of code, number of functions, number of local as well as global variables. JSmeter⁴ is used to compute the cyclomatic complexity. We compute the cyclomatic complexity across all JAVASCRIPT functions in the application.

5.2 Experimental Setup

To run the experiment, we provide the URL of each experimental object to JSART. In order to produce representative execution traces, we navigate each application several times with different crawling settings. Crawling settings differ in the number of visited states, depth of crawling, crawling time, and clickable element types. To obtain representative data traces, each of our experimental objects is navigated three times on average. Although JSART can easily instrument the source code of imported JAVASCRIPT libraries (e.g., jQuery, Prototype, etc), in our experiments we are merely interested in custom code written by developers, since we believe that is where most programming errors occur.

To evaluate our approach in terms of inferring stable invariant assertions (RQ1), we count the number of stable invariant assertions generated by JSART before and after performing the filtering step. As a last check, we execute the initial version of the application using the stable assertions to see whether our filtered invariant assertions are reporting any false positives.

Once the stable invariant assertions are obtained for each web application, we perform regression testing on modified versions of each application (RQ2). To that end, in order to mimic regression faults, we produce twenty different versions for each web

⁴ <http://jsmeter.info>

Table 2. Properties of the invariant assertions generated by JSART

App ID	Trace Data (MB)	# Total Assertions	# Entry Assertions	# Exit Assertions	# DOM Assertions	# Total Unstable Assertions	# Unstable Entry Assertions	# Unstable Exit Assertions	# Unstable DOM Assertions	# Total Stable Assertions	# Stable Entry Assertions	# Stable Exit Assertions	# Stable DOM Assertions
1	8.6	303	120	171	12	0	0	0	0	303	120	171	12
2	124	2147	1048	1085	14	14	9	5	0	2133	1039	1080	14
3	1.2	766	387	379	0	16	8	8	0	750	379	371	0
4	31.7	311	138	171	2	14	7	7	0	297	131	164	2
5	0.4	55	20	27	8	0	0	0	0	55	20	27	8
6	2.3	464	160	266	38	3	1	2	0	461	159	264	38
7	1.2	29	4	6	19	0	0	0	0	29	4	6	19
8	0.1	20	2	2	16	0	0	0	0	20	2	2	16
9	2.6	163	58	104	1	0	0	0	0	163	58	104	1

application by injecting twenty faults into the original version, one at a time. We categorize our faults according to the following fault model:

1. **Modifying Conditional Statements:** This category is concerned with swapping consecutive conditional statements, changing the upper/lower bounds of loop statements, as well as modifying the condition itself;
2. **Modifying Global/Local Variables:** In this category, global/local variables are changed by modifying their values at any point of the program, as well as removing or changing their names;
3. **Changing Function Parameters/Arguments:** This category is concerned with changing function parameters or function call arguments by swapping, removing, and renaming parameters/arguments. Changing the sequence of consecutive function calls is also included in this category;
4. **DOM modifications:** Another type of fault, which is introduced in our fault model is modifying DOM properties at both JAVASCRIPT code level and HTML code level.

For each fault injection step, we randomly pick a JAVASCRIPT function in the application code and seed a fault according to our fault model. We seed five faults from each category.

To evaluate the effectiveness of JSART (RQ2), we measure the precision and recall as follows:

Precision is the rate of injected faults found by the tool that are correct: $\frac{TP}{TP+FP}$

Recall is the rate of correct injected faults that the tool finds: $\frac{TP}{TP+FN}$

where TP (true positives), FP (false positives), and FN (false negatives) respectively represent the number of faults that are correctly detected, falsely reported, and missed.

To evaluate the performance of JSART (RQ3), we measure the extra time needed to execute the application while assertion checks are in place.

Table 3. Precision and Recall for JSART fault detection

App ID	# FN	# FP	# TP	Precision (%)	Recall (%)
1	2	0	18	100	90
2	4	0	16	100	80
3	1	0	19	100	95
4	2	0	18	100	90
5	0	0	20	100	100
6	1	0	19	100	95
7	0	0	20	100	100
8	0	0	20	100	100
9	1	0	19	100	95

5.3 Results

In this section, we discuss the results of the case study with regard to our three research questions.

Generated Invariant Assertions. Table 2 presents the data generated by our tool. For each web application, the table shows the total size of collected execution traces (MB), the total number of generated JAVASCRIPT assertions, the number of assertions at entry point of the functions, the number of assertions at exit point of the functions, and the number of DOM assertions. The unstable assertions before the filtering as well as the stable assertions after the filtering step are also presented. As shown in the table, for applications 1, 5, 7, 8, and 9, all the generated invariant assertions are stable and the filtering step does not remove any assertions. For the remaining four applications (2, 3, 4, 6), less than 5% of the total invariant assertions are seen as unstable and removed in the filtering process. Thus, for all the experimental objects, the resulting stable assertions found by the tool is more than 95% of the total assertions. Moreover, we do not observe any unstable DOM assertions. In order to assure the stability of the resulting assertions, we examine the obtained assertions from the filtering step across multiple executions of the original application. The results show that all the resulting invariant assertions are truly stable since we do not observe any false positives.

As far as RQ1 is concerned, our findings indicate that (1) our tool is capable of automatically generating a high rate of JAVASCRIPT invariant assertions, (2) the unstable assertions are less than 5% of the total generated assertions, (3) the filtering technique is able to remove the few unstable assertions, and (4) all the remaining invariant assertions that JSART outputs are stable, i.e., they do not produce any false positives on the same version of the web application.

Effectiveness. Since applications 3, 4, and 9 do not contain many DOM assertions, we were not able to inject 5 faults from the DOM modification category. Therefore, we randomly chose faults from the other fault model categories.

In Table 3, we present the accuracy of JSART in terms of its fault finding capability. The table shows the number of false negatives, false positives, true positives, as well as the percentages for precision and recall. As far as RQ2 is concerned, our results show that JSART is very accurate in detecting faults. The precision is 100%, meaning that all the injected faults, which are reported by the tool, are correct. This also implies that our filtering mechanism successfully eliminates unstable assertions as we do not observe any false positives. The recall oscillates between 80-100%, which is caused by a low

Fig. 6. Performance plot of JSART

rate of missed faults (discussed in Section 6 under Limitations). Therefore, as far as RQ2 is concerned, JSART is able to successfully spot the injected faults with a high accuracy rate.

Performance. Figure 6 depicts the total running time needed for executing each web application with and without the assertion code. Checking a fairly large number of assertions at runtime can be time consuming. Thus, to capture the effect of the added assertions on the execution time, we exploit a 2-scale diagram. As shown in Figure 6, each experimental object is associated with two types of data. The left-hand Y-axis represents the running time (seconds), whereas the right-hand Y-axis shows the number of assertions. This way we can observe how the number of assertions relates to the running time. As expected, the figure shows that by increasing the number of assertions, the running time increases to some degree. While the time overhead of around 20 seconds is more evident for the experimental object 2 (i.e., Tunnel with 2147 number of assertions), it is negligible for the rest of the experimental objects. Considering that Tunnel has 260 statements in total, the number of assertions instrumented in the code is eight times more than the number of statements in the original version. Therefore, it is reasonable to observe a small amount of overhead. Though assertions introduce some amount of overhead, it is worth mentioning that we have not experienced a noticeable change (i.e., freezing or slowed down execution) while running the application in the browser.

Thus, as far as RQ3 is concerned, the amount of overhead introduced by our approach is 6 seconds on average for our experimental objects, which is negligible during testing. Furthermore, based on our observations, the assertions do not negatively affect the observable behaviour of the web applications in the browser.

Table 4. Manual effort imposed by our approach for deriving stable invariant assertions.

App ID	Total Time (min)	Manual Effort (min)
1	13	4
2	11.5	3
3	15.5	5
4	11	3
5	6.5	2.5
6	9	4.5
7	7.5	3.5
8	6.5	2
9	18	13

6 Discussion

Unstable Assertions. As mentioned in Section 3.3, we observe a few number of unstable invariant assertions initially, which are removed by our filtering mechanism. By analyzing our trace data, we observe that such unstable assertions arise mainly because of the multiple runtime types of JAVASCRIPT variables. This is based on the fact that in JAVASCRIPT it is possible to change the type of a variable at runtime. However, Daikon treats variables as single type, selects the first observed type, and ignores the subsequent types in the trace data. This results in producing a few number of unstable invariant assertions for JAVASCRIPT. We remove such unstable assertions in our filtering step. A drawback of removing these assertions, is that our tool might miss a fault during the regression testing phase. However, according to our observations, such unstable assertions form only around 5% of the total generated assertions. Thus, we are still able to achieve high accuracy as presented in the previous section.

Limitations. Our approach is not able to detect syntax errors that are present in the JAVASCRIPT code. Furthermore, tracing DOM manipulations using APIs other than the standard DOM API or jQuery is currently not supported by JSART. Further, a regression fault either directly violates an invariant assertion, or it can violate closely related assertions, which have been affected by the fault. However, if the tool is not able to infer any invariants in the affected scope of the error, it fails to detect the fault. This results in observing a low rate of false negatives as illustrated in Section 5.

Revisiting the Assumptions. As we mentioned in Section 3, we assume that the current version of the web application is bug-free. This is based on the fact that in regression testing a gold standard is always needed as a trusted version for comparing the test results against [4] to detect regression faults. However, if the original version of the application does contain an error, the generated assertions might reflect the error as well, and as such they are not able to detect the fault. Our second assumption states that the program specifications are unlikely to change frequently in revisions. Here we assume that software programs evolve gradually and regression faults are mostly due to small changes. However, if major upgrades occur in subsequent revisions such that the core specification of the application is affected, the inferred invariants from the original version may not be valid any longer and new invariant assertions need to be generated.

Automation Level. While the testing phase of JSART is fully automated, the navigation part requires some manual effort. Although the crawling is performed

automatically, we do need to manually setup the tool with different crawling configurations per application execution. Moreover, for each application run, we manually look at the size of the invariant output to decide whether more execution traces (and thus more crawling sessions) are needed. We present the manual effort involved with detecting stable invariant assertions in Table 4. The table shows the total time, which is the duration time of deriving stable assertions including both automatic and manual parts. The reported manual effort contains the amount of time required for setting up the tool as well as the manual tasks involved with the navigation part. The results show the average manual effort is less than 5 minutes.

7 Related Work

Automated testing of modern web applications is becoming an active area of research [2,14,16,19]. Most of the existing work on JAVASCRIPT analysis is, however, focused on spotting errors and security vulnerabilities through static analysis [11,12,27]. We classify related work into two broad categories: web application regression testing and program invariants.

Web Application Regression Testing. Regression testing of web applications has received relatively limited attention from the research community [26,25]. Alshahwan and Harman [1] discuss an algorithm for regression testing of web applications that is based on session data [24,8] repair. Roest et al. [22] propose a technique to cope with the dynamism in Ajax web interfaces while conducting automated regression testing. None of these works, however, target regression testing of JAVASCRIPT in particular.

Program Invariants. The concept of using invariants to assert program behaviour at runtime is as old as programming itself [6]. A more recent development is the automatic detection of program invariants through dynamic analysis. Ernst et al. have developed Daikon [9], a tool capable of inferring likely invariants from program execution traces. Other related tools for detecting invariants include Agitator [5], DIDUCE [13], and DySy [7]. Recently, Ratcliff et al. [20] have proposed a technique to reuse the trace generation of Daikon and integrate it with genetic programming to produce useful invariants. Conceptually related to our work, Rodríguez-Carbonell and Kapur [21] use inferred invariant assertions for program verification.

Mesbah et al. [16] proposed a framework called ATUSA for manually specifying generic and application-specific invariants on the DOM-tree and JAVASCRIPT code. These invariants were subsequently used as test oracles to detect erroneous behaviours in modern web applications. Pattabiraman and Zorn proposed DoDOM [19], a tool for inferring invariants from the DOM tree of web applications for reliability testing.

To the best of our knowledge, our work in this paper is the first to propose an automated regression testing approach for JAVASCRIPT, which is based on JAVASCRIPT invariant assertion generation and runtime checking.

8 Conclusions and Future Work

JAVASCRIPT is playing a prominent role in modern Web 2.0 applications. Due to its loosely typed and dynamic nature, the language is known to be error-prone and difficult

to test. In this paper, we present an automated technique for JAVASCRIPT regression testing based on generated invariant assertions. The contributions of this work can be summarized as follows:

- A method for detecting JAVASCRIPT invariants across multiple application executions through on-the-fly JAVASCRIPT instrumentation and tracing of program variables and DOM manipulations;
- A technique for automatically converting the inferred invariants into stable assertions, and injecting them back into the web application for regression testing;
- The implementation of our proposed technique in an open source tool called JSART;
- An empirical study on nine open source JAVASCRIPT applications. The results of our study show that our tool is able to effectively infer stable assertions and detect regression faults with minimal performance overhead;

Our future work encompasses conducting more case studies to generalize the findings as well as extending the current JAVASCRIPT DOM modifications detector so that it is capable of coping with more patterns in other JAVASCRIPT libraries. In addition, we will explore ways of fully automating the navigation part by generating crawling specifications.

References

1. Alshahwan, N., Harman, M.: Automated session data repair for web application regression testing. In: Proceedings of the Int. Conf. on Software Testing, Verification, and Validation (ICST 2008), pp. 298–307. IEEE Computer Society (2008)
2. Artzi, S., Dolby, J., Jensen, S., Møller, A., Tip, F.: A framework for automated testing of JavaScript web applications. In: Proceedings of the Intl. Conference on Software Engineering (ICSE), pp. 571–580. ACM (2011)
3. Bezemer, C.-P., Mesbah, A., van Deursen, A.: Automated security testing of web widget interactions. In: Proceedings of the 7th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering (ESEC-FSE 2009), pp. 81–91. ACM (2009)
4. Binder, R.: Testing Object-Oriented Systems: Models, Patterns, and Tools. Addison-Wesley (2000)
5. Boshernitsan, M., Doong, R., Savoia, A.: From Daikon to Agitator: lessons and challenges in building a commercial tool for developer testing. In: Proc. Int. Sym. on Software Testing and Analysis (ISSTA 2006), pp. 169–180. ACM (2006)
6. Clarke, L.A., Rosenblum, D.S.: A historical perspective on runtime assertion checking in software development. ACM SIGSOFT Software Engineering Notes 31(3), 25–37 (2006)
7. Csallner, C., Tillmann, N., Smaragdakis, Y.: DySy: Dynamic symbolic execution for invariant inference. In: Proceedings of the 30th International Conference on Software Engineering (ICSE 2008), pp. 281–290. ACM (2008)
8. Elbaum, S., Rothermel, G., Karre, S., Fisher, M.: Leveraging user-session data to support web application testing. IEEE Trans. Softw. Eng. 31, 187–202 (2005)
9. Ernst, M., Perkins, J., Guo, P., McCamant, S., Pacheco, C., Tschantz, M., Xiao, C.: The Daikon system for dynamic detection of likely invariants. Science of Computer Programming 69(1-3), 35–45 (2007)
10. Groeneveld, F., Mesbah, A., van Deursen, A.: Automatic invariant detection in dynamic web applications. Technical Report TUD-SERG-2010-037, TUDelft (2010)

11. Guarnieri, S., Livshits, B.: Gatekeeper: mostly static enforcement of security and reliability policies for JavaScript code. In: Conference on USENIX Security Symposium, SSYM 2009, pp. 151–168 (2009)
12. Guha, A., Krishnamurthi, S., Jim, T.: Using static analysis for Ajax intrusion detection. In: Intl. Conference on World Wide Web (WWW), pp. 561–570 (2009)
13. Hangal, S., Lam, M.S.: Tracking down software bugs using automatic anomaly detection. In: Proceedings of the 24th International Conference on Software Engineering (ICSE 2002), pp. 291–301. ACM Press (2002)
14. Marchetto, A., Tonella, P., Ricca, F.: State-based testing of Ajax web applications. In: Proc. 1st Int. Conference on Sw. Testing Verification and Validation (ICST 2008), pp. 121–130. IEEE Computer Society (2008)
15. Mesbah, A., van Deursen, A., Lenselink, S.: Crawling Ajax-based web applications through dynamic analysis of user interface state changes. *ACM Transactions on the Web (TWEB)* 6(1), 3:1–3:30 (2012)
16. Mesbah, A., van Deursen, A., Roest, D.: Invariant-based automatic testing of modern web applications. *IEEE Transactions on Software Engineering (TSE)* 38(1), 35–53 (2012)
17. Meyer, B.: Applying design by contract. *Computer* 25(10), 40–51 (1992)
18. Ocariza, F.J., Pattabiraman, K., Mesbah, A.: AutoFLox: An automatic fault localizer for client-side JavaScript. In: Proceedings of the 5th IEEE International Conference on Software Testing, Verification and Validation (ICST 2012), pp. 31–40. IEEE Computer Society (2012)
19. Pattabiraman, K., Zorn, B.: DoDOM: Leveraging DOM invariants for Web 2.0 application robustness testing. In: Proc. Int. Conf. Sw. Reliability Engineering (ISSRE 2010), pp. 191–200. IEEE Computer Society (2010)
20. Ratcliff, S., White, D., Clark, J.: Searching for invariants using genetic programming and mutation testing. In: Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation (GECCO). ACM (2011)
21. Rodríguez-Carbonell, E., Kapur, D.: Program Verification Using Automatic Generation of Invariants. In: Liu, Z., Araki, K. (eds.) *ICTAC 2004*. LNCS, vol. 3407, pp. 325–340. Springer, Heidelberg (2005)
22. Roest, D., Mesbah, A., van Deursen, A.: Regression testing Ajax applications: Coping with dynamism. In: Proc. 3rd Int. Conf. on Sw. Testing, Verification and Validation (ICST 2010), pp. 128–136. IEEE Computer Society (2010)
23. Runeson, P., Höst, M.: Guidelines for conducting and reporting case study research in software engineering. *Empirical Software Engineering* 14(2), 131–164 (2009)
24. Sprenkle, S., Gibson, E., Sampath, S., Pollock, L.: Automated replay and failure detection for web applications. In: ASE 2005: Proc. 20th IEEE/ACM Int. Conf. on Automated Sw. Eng., pp. 253–262. ACM (2005)
25. Tarhini, A., Ismail, Z., Mansour, N.: Regression testing web applications. In: Int. Conf. on Advanced Comp. Theory and Eng., pp. 902–906. IEEE Computer Society (2008)
26. Xu, L., Xu, B., Chen, Z., Jiang, J., Chen, H.: Regression testing for web applications based on slicing. In: Proc. of Int. Conf. on Computer Software and Applications (COMPSAC), pp. 652–656. IEEE Computer Society (2003)
27. Zheng, Y., Bao, T., Zhang, X.: Statically locating web application bugs caused by asynchronous calls. In: Proceedings of the Intl. Conference on the World-Wide Web (WWW), pp. 805–814. ACM (2011)