

MENG INDIVIDUAL PROJECT

DEPARTMENT OF COMPUTING

**Towards Automatic Verification of JavaScript
Programs**

Author:
Beatrix de Wilde

Supervisor:
Prof. Philippa Gardner
Second Marker:
Prof. Alessio R. Lomuscio

June 19, 2017

Abstract

JavaScript, due to its dynamic nature and complex semantics, has fewer verification tools than languages such as C and Java. In order to tackle this challenge, we introduce for the first time a bi-abductive symbolic analysis for JSIL, an intermediate language for JavaScript verification. Our analysis is scalable and can fully automatically infer specifications of JSIL procedures that are non-recursive and do not contain loops.

Bi-Abduction provides the foundation of our analysis; we give the rules of our bi-abductive symbolic execution and prove their soundness. We further present an algorithm for generating procedure specifications using these rules.

We give an implementation of our bi-abductive analysis for JSIL and use it to generate specifications for JSIL programs. Using the tool, we generate and verify specifications for some of the JavaScript internal functions. Finally, we evaluate the quality of the automatically generated specifications by comparing them with their handwritten counterparts.

Acknowledgements

First, I would like to thank my supervisor, Prof. Philippa Gardner, for her wholehearted support throughout this project.

I would also like to thank José Fragoso Santos for all his input and feedback, it has been a privilege to work with him. I am grateful to Petar Maksimović and Azalea Raad, for sharing their insight during the project.

For her inspiration and advice across my time at Imperial, I would like to thank my tutor, Prof. Susan Eisenbach.

Finally, I would like to thank my family and friends for their endless support across these past four years.

Contents

1	Introduction	5
2	Background	7
2.1	Program Verification	7
2.1.1	Program Correctness	7
2.1.2	Static Verification Techniques	8
2.1.3	Related Tools	9
2.2	Separation Logic and Tools	10
2.2.1	Hoare Logic	10
2.2.2	Separation Logic	10
2.2.3	Bi-Abduction	11
2.2.4	Related Tools	11
2.3	The JavaScript Language	12
2.4	The JavaScript Intermediate Language	13
2.4.1	Motivation for JSIL	14
2.4.2	The JavaScript Intermediate Language	14
2.4.3	The JavaScript Verification Toolchain	17
3	Symbolic Analysis	21

3.1	Axioms for Basic Commands	21
3.2	Symbolic Execution Rules for Control Flow Commands	25
4	Bi-Abductive Symbolic Analysis	31
4.1	Bi-Abductive Analysis: High Level Description	31
4.2	Bi-Abductive Axioms for Basic Commands	32
4.3	Bi-Abductive Proof Rules for Control Flow Commands	35
4.4	Bi-Abductive Algorithm for JSIL Programs	39
5	Implementation for Bi-Abductive Symbolic Analysis	43
5.1	Architecture	43
5.2	Implementation Challenges	45
6	Evaluation	49
6.1	Evaluation Infrastructure	49
6.2	Generating Specifications for the Internal Functions	50
6.2.1	Specification Comparison	51
7	Conclusion	55
7.1	Objectives Achieved	55
7.2	Future Work	55
	Bibliography	57
A	JSIL Semantics and Substitution Lemma	61
A.0.1	JSIL Substitution Lemma	61
A.0.2	JSIL Semantics	61

Chapter 1

Introduction

JavaScript is the developer's language of choice for Web applications and is supported by all major browsers. It is an untyped, dynamic language, with complicated semantics and notoriously confusing corner cases. Among other features, it supports extensible objects, prototype inheritance, dynamic property access, dynamic functions calls, and type coercion for delaying error reporting. These features interact in an intricate and sometimes unexpected fashion and can lead to developers unknowingly introducing bugs and vulnerabilities. For example, even simple addition can perplex the seasoned developer: adding an object to an empty string results in the number 0 rather than an error, whereas reversing the arguments leads to a completely different result and type. In particular, the JavaScript program `'' + {}` results in the string `'[object_Object]'` whereas the program `{ } + ''` results in the number 0.

In addition to the substantial economic cost resulting from vulnerabilities, program errors in medical software have led to fatalities in the past. Developers aim to reduce bugs dominantly through testing and code reviews. This requires significant developer effort and is subject to human error. Therefore program verification tools have seen a rise in the recent years. For practical applications, it is vital that these tools can scale to industrial-level codebases. Compared to other languages, such as Java, C, and C++, JavaScript has far fewer tools available for developers.

As part of an endeavour to verify JavaScript programs, the Program Specification and Verification Group at Imperial College London has developed JaVerT [35], which is a semi-automatic verification toolchain for JavaScript based on separation logic [40]. JaVerT can verify functional correctness properties of JavaScript programs annotated with preconditions and postconditions for functions, loop invariants, and folding and unfolding instructions for user-defined predicates, written in an assertion language, JS Logic. JaVerT translates JavaScript programs to the simple intermediate goto language JSIL, translates JS Logic assertions to JSIL Logic assertions, and performs verification at the level of JSIL. Correctness results have been established to prove that verification at the level of JSIL lifts to verification at the level of JavaScript.

As a first step towards fully automated verification of JavaScript programs, we aim to generate automatic specifications for JSIL procedures. Bi-Abduction [12] is a technique for automatically generating specifications through the use of separation logic inference rules. The properties that can be proven using bi-abduction are weaker than those provable semi-automatically, but bi-abduction allows verification tools to scale to industry-grade code bases.

Motivation

There are two key motivations for having a fully automatic tool: to automatically verify general properties and to support the developer in verifying functional properties of JavaScript programs.

Specifications provide the footprint of a program. Therefore, using these specifications it is possible to automatically verify properties of programs. In particular, verification of confinement properties is of interest for JavaScript programs. Confinement involves preventing untrusted code from modifying or accessing security-sensitive resources. As web applications are usually composed of many programs coming from different origins executing in the same global context, it would be beneficial to automatically verify whether sensitive resources get tainted or leaked.

In addition to automatic verification of general properties, generating specifications can support developers in verifying functional properties. As it can be taxing for developers to manually annotate code and write specifications for a large codebase, automatically generating specifications allows the developer to annotate an interesting subset of procedures or adapt generated specifications to verify more expressive functional properties.

Contribution

We formalise the symbolic analysis used for the semi-automatic JaVerT tool in a way that closely follows the implementation of the tool (Section 3). In particular, and in contrast to a previous formalisation of JSIL logic [35], our analysis is specified in an algorithmic fashion, meaning that it can be straightforwardly implemented and serve as a basis for the bi-abductive analysis we are aiming at.

Building upon our symbolic analysis for JSIL, we introduce a bi-abductive symbolic analysis that generates verifiable specifications for JSIL procedures (Section 4). We establish the soundness of both analyses. Furthermore, we present an algorithm, that given a JSIL program, infers the specifications of its procedures. Our analysis is a first step to a fully-fledged bi-abductive analysis for JSIL. In particular, we do not support the inference of loop invariants and specifications of recursive procedures.

Finally, we introduce the tool, JSIL Abduce, that implements the bi-abductive symbolic analysis for JSIL. The tool is able to produce procedure specifications for a JSIL program (Section 5). We evaluate the tool by using JSIL Abduce to automatically generate specifications for the JSIL implementations of the JavaScript internal functions (Section 6). We check that the generated specifications are correct by checking that they are satisfied by their corresponding implementations, using the existing verification tool for JSIL, JSIL Verify [35]. We analyse the quality of the bi-abductive analysis by comparing the inferred specifications with the prior handwritten specifications.

In the conclusion (Section 7) we discuss how our contributions can be used to achieve fully automatic verification of JavaScript as well as possible applications of such an analysis.

Chapter 2

Background

First, we provide a background in program verification and static analysis techniques (Section 2.1). We give a summary of related work in this area. We then focus on separation logic, bi-abduction and a range of program verification tools based on separation logic (Section 2.2). Additionally, we give high level details of the JavaScript language (Section 2.3).

Finally, we discuss work by the Program Specification and Verification Group at Imperial College London in JavaScript verification. In particular, we describe the JavaScript intermediate language (Section 2.4.2) and the JavaScript verification tool chain (Section 2.4.3).

2.1 Program Verification

2.1.1 Program Correctness

Program correctness is vital, not only to safety-critical code, but to any industrial level system. The impact of avoidable bugs is evident in today's society, with recent vulnerabilities such as a bug in Microsoft Windows [34] which was exploited by the ransomware WannaCry [41]. Affecting over 150 countries, the ransomware encrypted the files of the infected computers requesting a ransom in Bitcoins. The impact of software bugs can be catastrophic and even lead to loss of lives, as shown by a bug in Therac-25, a radiation therapy machine, which led to the death of six patients from radiation overdose [28].

In practice, developers ensure the correctness of code through testing and code reviews. This is at the expense of developer time and is subject to human error. Dijkstra noted "Testing shows the presence, not the absence of bugs" [9]. Therefore, the effectiveness of testing is based on the coverage of the test suite, which, in turn, depends on the manual effort of the developer. Increasingly, developers are turning to automatic verification tools to analyse the correctness of programs and support them in the development process. Program correctness can be split into verifying two types of properties: functional and general.

Verifying functional properties requires checking that a program meets a specification [27]. Program specifications typically consist of a precondition and postcondition. The precondition outlines the state required before the code is run and the postcondition specifies the state which holds after execution. We give an example specification in Figure 2.1. The program increments the i -th element of the array. The precondition states that the i -th element must equal some value u and the postcondition states the i -th element must equal to $u+1$ after executing the program.

```
1 \\ Precondition: arr[i] = u
2 arr[i] += 1;
3 \\ Postcondition: arr[i] = u + 1
```

Figure 2.1: Specification example

On the other hand, verifying general properties involves proving whether a program will produce a certain type of error, such as an integer overflow or a null pointer exception. In the example in Figure 2.1, a general property which may be checked is whether the index i is in the bounds of the array. General properties are often easier to automatically verify as they require less developer input.

Generally, verification is split into static and dynamic analysis [18]. Static analysis involves evaluating the program without running the code, whereas dynamic analysis requires an environment in which the program is run. Typically, dynamic analysis [2, 36] is scalable as it only explores realistic paths. However, dynamic analysis requires the entire execution environment and often the entire codebase, therefore it is hard to analyse isolated sections of the system. Developer tests, such as unit tests, can be considered a form of dynamic analysis. In comparison, static analysis is able to explore more execution paths, gaining much higher coverage. Additionally, if the analysis is performed in a modular style, then it is able to analyse incomplete sections of code.

2.1.2 Static Verification Techniques

We will now explore the main techniques used in static analysis tools.

SAT and SMT solvers are required for many backends of static analysis tools [15, 3, 8]. Satisfiability (SAT) solvers take a boolean formula and try to solve for a variable assignment where the formula evaluates to true. Satisfiability modulo theories (SMT) solvers are similar to SAT solvers; however, they take a formula in the first-order logic with equality. Competitions are run in order to encourage advancement in SAT and SMT solvers [32, 30].

Model checking is a lightweight technique to check whether the program satisfies a property instead of verifying the entire system [14, 22]. Given a model, M , which is a representation of a system and a temporal logic formula, ϕ , relating to the property to be verified, then we check that the model satisfies the formula. Formally, $M \models \phi$.

Bounded model checking (BMC) takes a program transition system and unrolls the loops to a predefined depth [6, 13]. A SAT solver can then check whether there is a property violation up to that depth. BMC is an under-approximation and can scale for large programs. The developer need only provide a correctness property.

Symbolic execution evaluates program instructions on symbolic values instead of concrete values [25]. Symbolic values represent an entire range of values that a variable can take, and get constrained by conditions resulting from the program path covered so far. Therefore, symbolic execution analysis does not need to perform the infeasible task of exhaustively testing all possible concrete input values. Additionally, symbolic execution has the ability to produce concrete test cases, and is therefore usually the basis of many testing tools.

In order to illustrate symbolic execution of a program, we give a snippet of code for a binary search algorithm in Figure 2.2. The code calculates the middle of the upper and lower bounds, then it checks if the search key is at that index. Otherwise, it checks if the search key is smaller or greater than the middle index and returns.

```

1 index = (lower_index + upper_index) / 2;
2
3 if (arr[index] = key) {
4     return index;
5 } else if (arr[index] < key) {
6     lower_index = index + 1;
7 } else {
8     upper_index = index - 1;
9 }

```

Figure 2.2: Code Snippet for Binary Search

We give the symbolic execution graph for this code snippet in Figure 2.3. Branching of the symbolic execution occurs when an `if` statements is reached, the path conditions are recorded for each branch. For the branch, where the search key is greater than the middle index, the path conditions `not (arr[index] = key)` and `not (arr[index] < key)` are combined to the single path condition `arr[index] > key`. It was discovered that almost all binary searches contain a bug [24], if the code snippet given in Figure 2.2, is written in Java then the first line contains a bug. An integer overflow occurs if the sum of `lower_index` and `upper_index` exceeds the maximum positive value for type `int`.

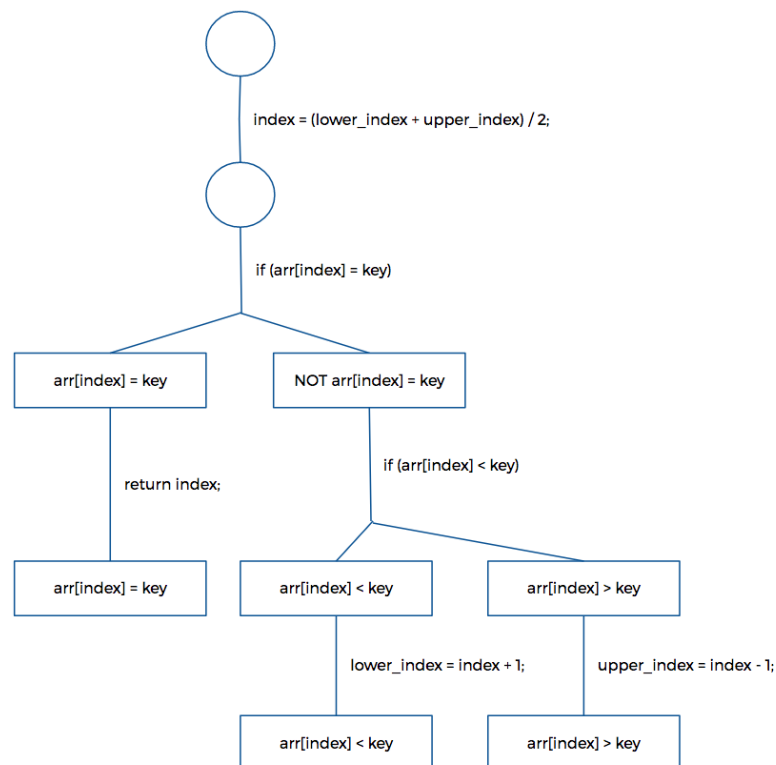


Figure 2.3: Symbolic Execution Graph of Code Snippet for Binary Search

2.1.3 Related Tools

KLEE [10], a symbolic execution tool, automatically generates tests for the intermediate language LLVM. The tool uses environment models written in C in order to stub system calls. The core KLEE engine uses the STP constraint solver [31] to solve branch conditions and search heuristics, random path selection and coverage-optimised search to decide which paths to explore. The focus of KLEE is on general properties and it succeeds in finding many bugs, mostly memory-related errors, in well established-code.

Symbolic PathFinder (SPF) [38, 39] focuses on the verification of Java. The tool symbolically executes the intermediate language Java ByteCode. This symbolic execution is built on top of a model checking tool, Java PathFinder [42]. When branching, the path conditions are checked by constraint solvers. The tool uses three constraint solvers, Choco [1], CVC3 [33] and IASolver [20], for different types of constraints. JUnit test cases can be generated by the tool.

CBMC [26] is a bounded model checker for annotated C and C++ programs. First, it translates the program into an intermediate goto language. It then performs symbolic execution unrolling the loops to a predefined depth. It passes the resulting formulae, corresponding to paths through the program control flow graph, into the SAT solver MiniSat 2.2.0 [17]. CBMC can be run with increasing unrolling depths, in order to reduce false negatives.

2.2 Separation Logic and Tools

2.2.1 Hoare Logic

Hoare logic, a system to reason about the correctness of programs, was developed by Tony Hoare in the 1960s [21]. Hoare Triples describe the connection between a program's precondition P , the code C , and postcondition Q . A Hoare Triple, $\{P\} C \{Q\}$, can be interpreted as "if the precondition P holds, then after executing the program C , the postcondition Q will hold". Proof of a Hoare Triple $\{P\} C \{Q\}$ gives partial correctness. In order to establish total correctness, termination of the Hoare Triple must also be proven. Although Hoare Logic is capable of reasoning about programs that alter the variable state, extending it to account for heap state leads to issues with scalability when it comes to reasoning about programs.

2.2.2 Separation Logic

Separation Logic [40] extends Hoare Logic in order to scalably reason about the correctness of programs with heap state. It extends the formulae of predicate calculus with assertions and connectives to describe the heap. Four new constructs are added to describe the heap: the empty heap emp ; the singleton heap $e_1 \mapsto e_2$; the separating conjunction $a_1 * a_2$; and the separating implication $a_1 \multimap a_2$.

The assertion emp asserts the heap is empty. The assertion $e_1 \mapsto e_2$ asserts the heap has exactly one cell, with the address given by the value of the expression e_1 and the contents given by the value of the expression e_2 . The assertion $a_1 * a_2$ asserts the heap can be split into two disjoint parts, where one part satisfies a_1 and the other satisfies a_2 . The assertion $a_1 \multimap a_2$ asserts if the heap is extended with a part where a_1 holds, then in the resulting heap a_2 holds.

The Frame Rule, given in Figure 2.4, allows reasoning about portions of the heap that are affected by the program code C , leaving the framed section R of the heap unmodified. This allows for local reasoning about the heap.

$$\frac{\text{FRAME} \quad \{P * R\} C \{Q * R\}}{\{P\} C \{Q\}}$$

where no variable occurring free in R is modified by C .

Figure 2.4: Frame Rule

2.2.3 Bi-Abduction

Abductive inference, described by Charles Peirce [37], finds the most likely explanation for an observation. Notably adopted by the fictional detective Sherlock Holmes in the works of Sir Arthur Conan Doyle, Holmes is able to solve cases by abducting important crime information from clues. Calcagno et al. [12] proposed using abduction in order generate program specifications.

Formally, abductive inference is when, given an assumption A and an observation O , we abduce the missing assumption M . In first-order logic, this is presented by the entailment $A \wedge M \vdash O$. This problem has been adapted for separation logic and is presented by the entailment $A * M \vdash O$. This entailment states that given two spatially disjoint assumptions, A and M , we are able to entail the observation O .

Bi-Abduction [12] generalises the abductive inference problem. It solves the following question: Given an assumption A and an observation O , what is the required missing assumption M and additional untouched conclusion F not captured by the observation O ? Formally, the problem can be given by the entailment $A * M \vdash O * F$. We call the missing state M the anti-frame and F the frame axiom.

Upon calling a procedure in a program, we are able to use bi-abduction to solve for both the additional state required before calling the procedure and any additional state not touched by the procedure call. Given the precondition of the procedure being called and the current state we solve the problem $current\ state * M \vdash precondition * F$.

We give an example of a simple addition procedure $add(x,y)$ which adds the two inputs x and y together. The precondition of the procedure is $\{ x \mapsto a * y \mapsto b \}$ and the postcondition is $\{ x \mapsto a * y \mapsto b * z \mapsto a + b \}$. When calling this procedure in a state $\{ x \mapsto 4 * w \mapsto 2 \}$ we pose the bi-abduction question:

$$x \mapsto 4 * w \mapsto 2 * ?M \vdash x \mapsto a * y \mapsto b * ?F$$

A possible solution is $?M = y \mapsto b * a \doteq 4$ and $?F = w \mapsto 2$, where $?M$ is required for the precondition and $?F$ is untouched by the precondition.

2.2.4 Related Tools

Separation logic allows verification tools to reason about programs at a larger scale. We discuss some of these tools.

An initial experiment into whether hand written separation logic proofs could work in an automatic environment was explored with the tool Smallfoot [4]. Smallfoot is a semi-automatic tool, as annotations for preconditions and postconditions are required. A specially designed input language was developed for the tool. The assertion language provides inductive predicates for trees and lists. They note the potential for symbolic execution in automating separation logic proofs.

Expanding on Smallfoot, Space Invader [43] adds a widening operator to the symbolic execution in order to guarantee termination of fixed-point calculations. These calculations are performed in the semantics of while loops. Therefore, the tool is able to find and then prove while loop invariants.

Similar to Space Invader, SLayer [5], builds upon Smallfoot's assertion language. The tool

verifies C programs and targets memory safety errors. No annotations are required. The tool's approach to loop invariants is similar to Space Invader. The Z3 SMT solver [15] is used to reason about pure formulae. SLayer is able to verify sizeable code bases, including Windows device drivers.

Abductor [12], goes one step further by generating specifications and is, therefore, fully automatic and able to scale to larger programs. Bi-Abduction is used to generate preconditions for procedures. The specifications Abductor generates are small specifications, as they only contain the footprint of the procedures.

A specification table is gradually built up as the functions are split into partitions where every function is in a lower or the same partition with respect to its caller. The tool starts by inferring the specifications of the functions in the first partition. Then, it moves on to inferring the specifications of the functions in the second partition. The tool continues through all the partitions, where in each partition it uses the specifications already generated from the lower partitions.

Abductor was developed further into the industrial level tool Infer [11]. The compositional nature of the tool allowed it to properly scale. This is shown through the ability for Infer to be integrated into the Facebook development cycle. Initially aimed at C, the tool has also been expanded to Java.

Verifast [23] is a semi-automatic verification tool based on separation logic. The tool requires annotations for preconditions and postconditions, but not for loop invariants. It can take programs written in C and Java and supports folding and unfolding of user defined predicates. The toolchain JaVerT [35] is inspired by Verifast, but targets JavaScript verification and is discussed in Section 2.4.3.

2.3 The JavaScript Language

JavaScript is the predominant language for Web developers. It is conventionally coupled with CSS and HTML for the writing of client-side web applications. Untyped, dynamic-natured and supporting concepts from multiple programming paradigms, JavaScript makes the detection of vulnerabilities difficult. In addition to the nature of the language, most applications include external libraries such as jQuery, React, and Async, as well as other third-party code. Having multiple scripts from different sources running in the same global environment can easily lead to leaking sensitive information.

JavaScript is standardised by the ECMAScript Committee, which provides and regularly updates the ECMAScript Language Standard. The standard is currently in its sixth edition (ES6). The standard also defines a strict mode of the language, which improves error reporting and has better behavioural properties. For this project, we will be working with the strict mode of the fifth edition of the standard (ES5 Strict) [16], as JaVerT is targeting ES5 Strict. There have been several attempts at formalising the semantics of JavaScript. Small-step operational semantics for ECMAScript have been defined by Maffeis et al [29] for ES3. Drawing inspiration from that work, JSCert [7] provides a formalisation of ES5 semantics in the Coq theorem prover.

JavaScript Objects. As defined in the ECMAScript specification, objects in JavaScript are collections of properties. There are two types of object properties: *internal* and *named*. Internal properties capture the inner workings of the language, such as prototype inheritance and object extensibility, and are not available to the programmer. Named properties can be either *data properties* or *accessor properties* and, unlike in C++ or Java, they are not associated with values, but rather with *property descriptors*. Property descriptors are lists of *attributes*, which describe the ways in which a property can be accessed and/or modified. Depending on the attributes

they contain, named properties can either be *data properties* or *accessor properties*. Data properties contain the value, writable, enumerable, and configurable attributes (denoted by [V], [W], [E], and [C]), whereas accessor properties contain get and set attributes (denoted by [G] and [S]), as well as [E] and [C]. The attributes have the following semantics: [V] holds the actual value of the property; [W] describes whether or not the property's value can be changed; [E] indicates whether or not the property will be included in a for-in enumeration; [C] allows or disallows any change to the other attributes (except for value, which it does not affect), as well as any change in the type of the property (data to accessor and vice versa); [G] and [S] play a role similar to getters and setters of Java and provide property encapsulation.

JavaScript Initial Heap. Before the execution of any JavaScript program, an initial heap must be established, as described in Chapter 15 of the ECMAScript standard. It must contain a unique global object as well as the constructors and prototypes of all JavaScript built-in libraries, such as Object, Function, Array, and String, which are widely used by JavaScript programmers.

The global object is critical for the correct functioning of JavaScript programs, as it holds: all global variables; value properties such as NaN, infinity and undefined; and function properties such as eval, which executes the supplied string argument containing arbitrary JavaScript code. Some properties of the global object are fixed; it is, for instance, not possible to reassign to undefined, Array, or String.prototype. It is, however, possible to reassign to eval, as shown in figure 2.5, which may introduce serious vulnerabilities. It is also possible to perform *prototype poisoning* by, for instance, adding or overriding the properties in the built-in object prototypes, thereby altering their functionality and possibly gaining access to sensitive information.

```
1 eval('2+2') // returns 4
2 // alters the global objects eval function
3 this.eval = new Function('return_this;')
4 // which now returns the global object
5 eval('2+2') // returns {global: ...}
```

Figure 2.5: Altering the eval property of the global object

In addition to built-in library functions, JavaScript has internal functions not accessible to developers. These functions provide core functionality, including: type conversion methods such as toPrimitive and toBoolean; prototype chain traversal methods such as getProperty; equality comparison methods; and the methods getValue and putValue.

2.4 The JavaScript Intermediate Language

A separation logic (JS Logic) was developed for a subset of JavaScript, however it is only able to model a simplified JavaScript heap and is difficult for automation purposes. Therefore, an intermediate goto language, JSIL, was developed. JSIL, is able to overcome verification difficulties with JavaScript and JS Logic. The toolchain, JaVerT, then compiles JavaScript programs and JS logic annotations to JSIL and verifies the JSIL programs.

In section 2.4.1 we discuss the complexities involved with JavaScript and the separation logic developed for JavaScript. In section 2.4.2 we discuss the JSIL language and finally in section 2.4.3 we discuss the toolchain, JaVerT.

2.4.1 Motivation for JSIL

A separation logic (JS Logic) has been developed for a subset of JavaScript by Gardner et al. [19], built upon a big-step operational semantics of ES3. JS Logic simplifies the memory model and targets only a fragment of JavaScript. Nonetheless, it is very complex, as illustrated by the function call rule in figure 2.6, which captures the JavaScript function call.

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

Figure 2.6: Function Call Rule in JS Logic [19]

JS Logic assertions feature standard boolean assertions and operators on expressions, the separating conjunction $*$ and implication \multimap , and the heap cell assertion $(E_1, E_2) \mapsto E_3$, which asserts that the object at location E_1 has field E_2 with value E_3 (or has no field E_2 , if E_3 equals \emptyset). It also features a non-standard separation logic connective, $P \wp Q$ (read: sepish), which describes a heap that can be split into two parts that do not need to be disjoint. This connective, as well as the \multimap separating implication are known to be very difficult for automation. Expanding JS Logic, as formulated in [19] to the full memory model of JavaScript and automating it is, therefore, not feasible.

2.4.2 The JavaScript Intermediate Language

The JSIL language is simpler compared to JavaScript. It does not contain dynamic function calls, also conditional branches and loops are performed by goto commands and commands are explicit with no corner cases. Additionally, the separation logic for JSIL (JSIL Logic) is comparatively simpler than JS Logic: it does not include the \wp connective or the separating implication \multimap and therefore can be easily automated. Additionally, reasoning in JSIL Logic does not require the simplification of the JavaScript memory model.

JSIL Syntax. The syntax of JSIL is given in Figure 2.7. A JSIL program $p \in \mathcal{P}$ is a set of top-level procedures. Procedures consist of three elements: a procedure name, $m \in \text{Str}$; a list of parameters, \bar{x} and a list of commands, \bar{c} . The list of commands is numbered and we use the notation $p_m(i)$ to refer to the i -th command of procedure m in program p . JSIL does not provide an explicit return command; a procedure terminates when it reaches one of two dedicated indexes, i_{nm} and i_{er} . When the i_{nm} -th command is reached, the procedure returns normally and returns the value of the dedicated variable x_{ret} ; when the i_{er} -th command is reached, the procedure returns an error and returns the value of the dedicated variable x_{err} . JSIL commands are divided into basic commands, which do not affect the control flow of programs, and control flow commands, which do affect it.

JSIL basic commands include support for object creation, variable assignment and field manip-

Strings: $m \in Str$	Numbers: $n \in Num$	Booleans: $b \in Bool$	Locations: $l \in \mathcal{L}$
Variables: $x \in \mathcal{X}_{JSIL}$	Literals: $\lambda \in Lit ::= n \mid b \mid m \mid \text{undefined} \mid \text{null}$		
Types: $t \in Types ::= Num \mid Bool \mid Str \mid Undef \mid Null \mid Empty \mid Obj \mid List \mid Type$			
Values: $v \in \mathcal{V}_{JSIL} ::= \lambda \mid l \mid \text{empty} \mid \text{error} \mid t \mid \bar{v}$			
Expressions: $e \in \mathcal{E}_{JSIL} ::= v \mid x \mid \ominus e \mid e \oplus e \mid \text{typeOf}(e) \mid \bar{e} \mid \text{nth}(e, e)$			
Basic Commands:			
$bc \in BCmd ::= \text{skip} \mid x := e \mid x := \text{new}() \mid x := [e, e] \mid [e, e] := e \mid \text{delete}(e, e) \mid$ $x := \text{hasField}(e, e) \mid x := \text{getFields}(e)$			
Commands: $c \in Ccmd ::= bc \mid \text{goto } i \mid \text{goto } [e] i, j \mid x := e(\bar{e}) \text{ with } j \mid x := \phi(\bar{x})$			
Procedures: $proc \in Proc ::= \text{proc } m(\bar{x})\{\bar{c}\}$			
Notation: $\bar{x}, \bar{v}, \bar{e}$, respectively, denote lists of variables, values, and expressions.			

Figure 2.7: JSIL Syntax

ulation including deletion, lookup and assignment. Additionally, there is support to check if an object has a field and to get all the fields of an object. The notation $[e_1, e_2]$ denotes the field e_2 of object e_1 .

JSIL control flow commands use numbered command labels to transfer control to other commands in the current procedure. The unconditional `goto` command transfers control to the command labelled i . The conditional `goto` command transfers control to the command labelled i if the expression e evaluates to true; otherwise, control shifts to the command labelled j . The procedure call command obtains the procedure name and the arguments by evaluating, respectively, the expression e and the list of expressions \bar{e} . The result of the procedure call is assigned to the variable x . If the program raises an error, control is transferred to the j -th command otherwise, control is transferred to the following command. The `phi`-assignment command assigns an element of the list \bar{x} to the variable x . Each variable in the list \bar{x} relates to a path taken to the current command. If there are n variables in the list \bar{x} then there are n paths to the current command. If the i -th path was taken then the i -th element of \bar{x} is assigned to x .

JSIL expressions include JSIL values, JSIL variables, lists of expressions and operators on expressions. There are various unary and binary operators that can be applied to JSIL expressions as well as the `typeOf`(e) operator which returns the type of an expression and the `nth`(e_1, e_2) operator which returns the e_2 -th element of the list e_1 .

In order to illustrate some aspects of JSIL, we give the example below, which demonstrates a JSIL program with one procedure that attempts to retrieve an interval value from a timeout object. In the program, "else", "then" and "rlab" are command labels. The basic command `hasField` checks if the timeout object has the field `interval` and assigns the result to variable x . Then, the control flow command `goto` transfers control to the command labelled "then" if the variable x evaluated to true, i.e. if the field `interval` exists; otherwise, it transfers control to the command labelled "else".

The basic command at the label "then" assigns the timeout interval to the variable y , control is then transferred to the unconditional `goto` which jumps to the return label `rlab`. The program then terminates with the return variable y .

The basic command at the label "else" assigns the default timeout interval to the variable y , control is then transferred to the `skip` command which again terminates with the return variable y .

```

proc getTimeoutInterval (timeout, defVal) {
  x := hasField(timeout, "interval");
  goto [x] then else;
then: y := [timeout, "interval"];
  goto rlab;
else: y := defVal;
  rlab: skip
} with {ret: y, rlab};

```

JSIL Semantics. The JSIL memory model includes:

- A JSIL store, $\rho \in Sto : \mathcal{X}_{JSIL} \rightarrow \mathcal{V}_{JSIL}$, which maps variables to values.
- A JSIL heap, $h \in \mathcal{H}_{JSIL} : \mathcal{L} \times \mathcal{X}_{JSIL} \rightarrow \mathcal{V}_{JSIL}$, which maps locations and variables to values.

Each JSIL procedure is executed in an individual store. The semantics of JSIL expressions are denoted by the judgement $\llbracket e \rrbracket_{\rho} = v$, where evaluating the expression e with respect to the store ρ results in the value v .

The semantics of JSIL basic commands are denoted by the judgement $\llbracket bc \rrbracket_{h,\rho} = (h', \rho', v)$ where evaluating the basic command bc with respect to the store ρ and heap h results in the value v , store ρ' and heap h' . The semantics of JSIL control flow commands are denoted by the judgement $p \vdash \langle h, \rho, i, j \rangle \Downarrow_m \langle h', \rho', v \rangle$ where starting from the i -th command in procedure m in program p with heap h and store ρ , where we have come from the j -th command, results in the heap h' , store ρ' and value v . The full semantics for JSIL are given in Appendix A.0.2.

JSIL Logic Assertions. JSIL logic assertions are described in Figure 2.8. JSIL assertions include existential quantification, heap assertions, and pure boolean assertions, excluding disjunction. The JSIL heap assertions consist of the separating conjunction, the empty heap, the heap cell $(E_1, E_2) \mapsto E_3$ where the field E_2 in object E_1 maps to the value E_3 and the empty fields assertion $\text{emptyFields}(E \mid \bar{E})$ where the object E can only possibly have the fields \bar{E} .

LOGICAL VALUES : $V \in \mathcal{V}_{JSIL}^L$	$\triangleq v \mid \emptyset \mid \bar{V}$
LOGICAL EXPRESSIONS : $E \in \mathcal{E}_{JSIL}^L$	$\triangleq V \mid x \mid X \mid \ominus E \mid E \oplus E \mid \bar{E}$
JSIL ASSERTIONS : $P \in \mathcal{AS}_{JSIL}$	\triangleq true false $\neg P$ $P \wedge P$ $\exists X.P$ PURE BOOLEAN $ E = E \mid E \leq E$ EQUALITIES $ \text{emp} \mid (E, E) \mapsto E$ HEAP $ P * P \mid \text{emptyFields}(E \mid \bar{E})$

Figure 2.8: JSIL Logic Assertions

JSIL logical expressions used in JSIL assertions are an extension of JSIL expressions with logical variables and the value \emptyset . The special value \emptyset is used with the heap cell assertion $(E_1, E_2) \mapsto \emptyset$ to denote that the object E_1 does not have the field E_2 . The semantics of JSIL logical expressions are similar to the semantics of JSIL expressions. The judgement $\llbracket E \rrbracket_{\rho}^{\epsilon}$ describes the evaluation of the expression E with respect to the store ρ and logical environment ϵ . Logical values consist of JSIL values and the special value \emptyset .

In order to illustrate some aspects of JSIL assertions, we give the example below, which is a JSIL assertion describing one object whose location in the heap is l . This object has three fields, interval, function, and @proto, described by the first three heap cell assertions. There, we see that the value of interval is the value of variable time, that the value of function is undefined,

and that the value of @proto is null. The emptyFields assertion states that the only fields that the object has are interval, function and @proto. The equality assertion states that the variable time has value 400.

$$\left\{ \begin{array}{l} (l, \text{interval}) \mapsto \text{time} * (l, \text{function}) \mapsto \text{undefined} * (l, \text{@proto}) \mapsto \text{null} * \\ \text{time} = 400 * \text{emptyFields}(l \mid \{\{\text{interval}, \text{function}, \text{@proto}\}\}) \end{array} \right\}$$

JSIL abstract heaps $H \in \mathcal{H}_{\text{JSIL}}^{\emptyset} : \mathcal{L} \times \text{Str} \rightarrow \mathcal{V}_{\text{JSIL}}^L$ are an extension of JSIL heaps with the special value \emptyset . An abstract heap is well-formed, if all objects have a non-none @proto field. Formally, $\forall H. \text{wf}(H) \Leftrightarrow (\forall l. (l, -) \in \text{dom}(H) \Rightarrow \exists V. V \neq \emptyset \wedge (l, \text{@proto}) \mapsto V \in H)$. A JSIL logical environment $\epsilon \in \mathcal{Env}_{\text{JSIL}} : \mathcal{X}_{\text{JSIL}}^L \rightarrow \mathcal{V}_{\text{JSIL}}^L$ is a mapping from logical variables to logical values.

A JSIL assertion $P \in \mathcal{AS}_{\text{JSIL}}$ is satisfiable if $H, \rho, \epsilon \models P$ where H is an abstract heap, ρ is a JSIL store and ϵ is a logical environment. The satisfiability relation of JSIL assertions is given in Figure 2.9.

$H, \rho, \epsilon \models \text{true}$	$\Leftrightarrow \text{always}$
$H, \rho, \epsilon \models \text{false}$	$\Leftrightarrow \text{never}$
$H, \rho, \epsilon \models \neg P$	$\Leftrightarrow H, \rho, \epsilon \not\models P$
$H, \rho, \epsilon \models P \wedge Q$	$\Leftrightarrow H, \rho, \epsilon \models P \wedge Q \text{ and } H, \rho, \epsilon \models Q$
$H, \rho, \epsilon \models \exists X. P$	$\Leftrightarrow \exists V \in \mathcal{V}_{\text{JSIL}}^L. H, \rho, \epsilon[X \mapsto V] \models P$
$H, \rho, \epsilon \models E_1 = E_2$	$\Leftrightarrow \llbracket E_1 \rrbracket_{\rho}^{\epsilon} = \llbracket E_2 \rrbracket_{\rho}^{\epsilon}$
$H, \rho, \epsilon \models E_1 \leq E_2$	$\Leftrightarrow \llbracket E_1 \rrbracket_{\rho}^{\epsilon} \leq \llbracket E_2 \rrbracket_{\rho}^{\epsilon}$
$H, \rho, \epsilon \models \text{emp}$	$\Leftrightarrow H = \text{emp}$
$H, \rho, \epsilon \models (E_1, E_2) \mapsto E_3$	$\Leftrightarrow H = (\llbracket E_1 \rrbracket_{\rho}^{\epsilon}, \llbracket E_2 \rrbracket_{\rho}^{\epsilon}) \mapsto \llbracket E_3 \rrbracket_{\rho}^{\epsilon}$
$H, \rho, \epsilon \models P * Q$	$\Leftrightarrow \exists H_1, H_2. H = H_1 \uplus H_2 \wedge (H_1, \rho, \epsilon \models P) \wedge (H_2, \rho, \epsilon \models Q)$
$H, \rho, \epsilon \models \text{emptyFields}(E \mid E_1, \dots, E_n)$	$\Leftrightarrow H = \uplus_{m \notin \{\llbracket E_1 \rrbracket_{\rho}^{\epsilon}, \dots, \llbracket E_n \rrbracket_{\rho}^{\epsilon}\}} (\llbracket E \rrbracket_{\rho}^{\epsilon}, m) \mapsto \emptyset$

Figure 2.9: JSIL Logical Assertions (Satisfiability Relation)

2.4.3 The JavaScript Verification Toolchain

The JavaScript Verification Toolchain (JaVerT), shown in Figure 2.10, verifies JavaScript programs annotated with JS Logic specifications. It is based on an infrastructure which has three components: (1) JS-2-JSIL, the compiler from JavaScript to JSIL; (2) JSIL Verify, a semi-automatic tool for verifying JSIL code annotated with JSIL logic specifications; and (3) JSIL implementations and verified specifications of the JavaScript internal functions. JaVerT works in the following way: first, JS-2-JSIL compiles JavaScript to JSIL and translates JS logic annotations to JSIL logical annotations. Then the semi-automatic tool JSIL Verify verifies the compiled JSIL programs with respect to the corresponding JSIL logic annotations, using the JSIL logic specifications of the JavaScript internal functions.

Annotations for JSIL programs are JSIL logic assertions. Annotations are given for procedure specifications which include pre- and postconditions. Additionally, annotations for loop invariants, and folding and unfolding user-defined predicates can be written using JSIL logic.

JS-2-JSIL compiles JavaScript code to JSIL code, line-by-line following the ECMAScript English standard. In order to show this correspondence, as well as how the use of the internal functions in the standard is reflected in compiled JSIL code, we give a snippet of JSIL code which is compiled from the JavaScript program `0 == ''`.

As described in the standard, the abstract equality operator first calls `getValue` on both arguments. Then, it follows the abstract equality comparison algorithm. This algorithm is captured by the `i__abstractEquality` internal function.

The ECMAScript standard notes in the algorithm that *"If $Type(x)$ is $Number$ and $Type(y)$ is $String$, return the result of the comparison $x == ToNumber(y)$ ". Therefore the result of this program is true, as converting the empty string to a number results in the value 0. The equality comparison is the source of much confusion amongst JavaScript developers. It is recommended to use the strict equality operator (`===`) rather than the equality operator (`==`).*

```
(* The setupInitialHeap function creates the initial heap,
   which is established before the execution of any JavaScript program *)
x_0 := "setupInitialHeap"();
x_sc_0 := {{ $lg }};
(* Set this to the global object location *)
x__this := $lg;
(* Get value of first argument *)
x_1_v := "i__getValue"(0.) with elab;
(* Get value of second argument *)
x_2_v := "i__getValue"("") with elab;
(* Perform abstract equality comparison algorithm *)
x_3 := "i__abstractEquality"(x_1_v, x_2_v) with elab;
(* Get the value of the comparison result *)
x_3_v := "i__getValue"(x_3) with elab
xret := x_3_v;
```

JSIL Verify is a semi-automatic verification tool for JSIL. It is based on a symbolic execution engine and an entailment engine. The entailment engine has two components: one for resolving spatial entailments and the other for pure entailments. Spatial entailments are handled by JSIL Verify, whereas the pure ones are delegated to the Z3 SMT solver. When verifying compiled JavaScript code, the JSIL logic specifications for the internal functions are imported before the symbolic execution begins.

Correctness of JaVert. The correctness of JaVert is given by systematic testing of JS-2-JSIL, the soundness result for JSIL Logic and the correctness result for the translation of assertion between JS Logic and JSIL Logic. JS-2-JSIL is tested against the ECMAScript ES6 Test262 test suite, where it passes all 8797 relevant tests. Additionally, the JavaScript internal functions implemented in JSIL are step-by-step faithful to the standard and their specifications are verified by JSIL Verify.

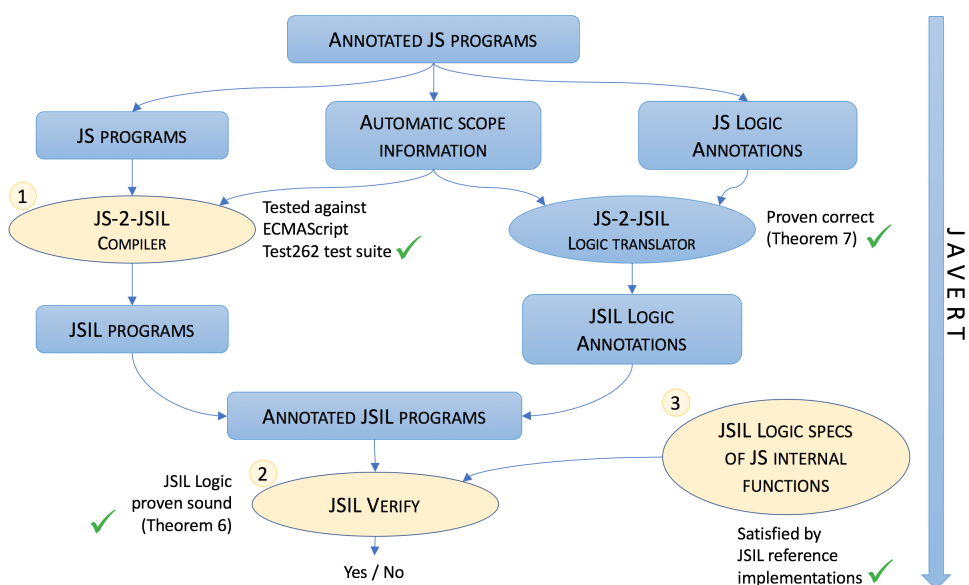


Figure 2.10: JaVerT [35]

Chapter 3

Symbolic Analysis

The symbolic analysis verifies JSIL procedure specifications. It performs this verification by symbolically executing JSIL basic and control flow commands. We take the symbolic execution rules introduced in [35], and modify them, so they closely relate to the implementation. In particular, the revised rules explicitly show the continuation from the precondition to the final state where the execution of the procedure does not fault. In the symbolic analysis we represent the symbolic states as JSIL assertions. The symbolic analysis is split into two sections, first we define the axioms for basic commands (Section 3.1) then we present the proof rules for the control flow commands (Section 3.2). Additionally, in Section 3.2 we present an example of the symbolic analysis. Finally, we establish the soundness of the symbolic analysis formally in Theorem 1 by appealing to the semantics of JSIL commands.

3.1 Axioms for Basic Commands

The axioms for the JSIL basic commands, given in Figure 3.1, have the form $\{P\} \text{bc } \{Q\}$ meaning: when in a state where the JSIL assertion P holds, then executing the basic command bc will not fault and if the execution terminates then it results in a state satisfying JSIL assertion Q . Throughout the rules the notation $e_1 \doteq e_2$ represents $e_1 = e_2 \wedge \text{emp}$.

We briefly explain the non-standard axioms.

- [OBJECT CREATION] When an object is created it has a single `@proto` field with a null value. The `emptyFields(x | @proto)` assertion, states that the object only has the `@proto` field.
- [FIELD DELETION] In the field deletion rule, the object field `@proto` cannot be deleted.
- [MEMBER CHECK] In the member check rule, if the field's value is not `None` (\emptyset) then `x` is `true` otherwise `x` is `false`.

Evaluation. Soundness of the hoare triples for basic commands is established by Lemma 1. The lemma connects the hoare triples to the semantics of basic commands. In order to make this connection abstract heaps are related to concrete heaps using the erasure function, $[\cdot] : \mathcal{H}_{\text{JSIL}}^0 \rightarrow \mathcal{H}_{\text{JSIL}} : [H](l, x) = H(l, x) \stackrel{\text{def}}{\iff} (l, x) \in \text{dom}(H) \wedge H(l, x) \neq \emptyset$.

Lemma 1 (Frame Property and Soundness for Basic Commands). *For all basic commands $\text{bc} \in \text{BCmd}$, abstract heaps $H, \hat{H}_1, \hat{H}_2 \in \mathcal{H}_{\text{JSIL}}^0$, stores $\rho, \rho_f \in \text{Sto}$, logical environments $\epsilon \in \text{Env}$, values*

Basic Commands

$\frac{\text{FIELD ASSIGNMENT}}{\{(e_1, e_2) \mapsto -\} \ [e_1, e_2] := e_3 \ \{(e_1, e_2) \mapsto e_3\}}$	$\frac{\text{FIELD DELETION} \quad P = (e_1, e_2) \mapsto X * X \neq \emptyset * e_2 \neq @proto}{\{P\} \ \text{delete}(e_1, e_2) \ \{(e_1, e_2) \mapsto \emptyset\}}$	
$\frac{\text{FIELD ACCESS} \quad P = (e_1, e_2) \mapsto X * X \neq \emptyset}{\{P\} \ x := [e_1, e_2] \ \{P * x \doteq X\}}$	$\frac{\text{MEMBER CHECK} \quad P = (e_1, e_2) \mapsto X \quad Q = P * x \doteq \text{not}(X = \emptyset)}{\{P\} \ x := \text{hasField}(e_1, e_2) \ \{Q\}}$	
$\frac{\text{VAR ASSIGNMENT}}{\{\text{emp}\} \ x := e \ \{x \doteq e\}}$	$\frac{\text{OBJECT CREATION} \quad Q = (x, @proto) \mapsto \text{null} * \text{emptyFields}(x \mid @proto)}{\{\text{emp}\} \ x := \text{new}() \ \{Q\}}$	$\text{SKIP} \quad \{\text{emp}\} \ \text{skip} \ \{\text{emp}\}$
$\frac{\text{FRAME} \quad \{P\} \ \text{bc} \ \{Q\}}{\{P * R\} \ \text{bc} \ \{Q * R\}}$	$\frac{\text{CONSEQUENCE} \quad \{P'\} \ \text{bc} \ \{Q'\} \quad P \vdash P' \quad Q' \vdash Q}{\{P\} \ \text{bc} \ \{Q\}}$	$\frac{\text{EXISTS} \quad \{P\} \ \text{bc} \ \{Q\}}{\{\exists X. P\} \ \text{bc} \ \{\exists X. Q\}}$

Figure 3.1: Hoare Triples for Basic Commands

$v \in \mathcal{V}_{\text{JSIL}}$, JSIL heaps $\hat{h}_f \in \mathcal{H}_{\text{JSIL}}$, and assertions $P, Q \in \mathcal{AS}_{\text{JSIL}}$, if $\{P\} \text{bc} \{Q\}$, $H, \rho, \epsilon \models P$, and $\llbracket \text{bc} \rrbracket_{[H \uplus \hat{H}_1 \uplus \hat{H}_2], \rho} = (\hat{h}_f, \rho_f, v)$, then there is an abstract heap H_f such that $H_f, \rho_f, v \models Q$, $\hat{h}_f = [H_f \uplus \hat{H}_1 \uplus \hat{H}_2]$ and $\llbracket \text{bc} \rrbracket_{[H \uplus \hat{H}_1], \rho} = ([H_f \uplus \hat{H}_1], \rho_f, v)$.

Proof. For convenience, we name the hypotheses as follows:

- **H1:** $\{P\} \text{bc} \{Q\}$
- **H2:** $H, \rho, \epsilon \models P$
- **H3:** $\llbracket \text{bc} \rrbracket_{[H \uplus \hat{H}_1 \uplus \hat{H}_2], \rho} = (\hat{h}_f, \rho_f, v)$

Our goal is to show that there exists a JSIL abstract heap H_f , such that:

- **G1:** $\llbracket \text{bc} \rrbracket_{[H \uplus \hat{H}_1], \rho} = ([H_f \uplus \hat{H}_1], \rho_f, v)$
- **G2:** $H_f, \rho_f, \epsilon \models Q$.
- **G3:** $\hat{h}_f = [H_f \uplus \hat{H}_1 \uplus \hat{H}_2]$.

We proceed by induction on the derivation of **H1**.

- [SKIP] We have that $\text{bc} = \text{skip}$ and, after applying **H1**, that $P = \text{emp}$ and $Q = \text{emp}$. From the satisfiability of JSIL assertions and **H2**, we obtain that $H = \text{emp}$. From **H3** and the semantics of JSIL basic commands, we obtain that $\hat{h}_f = [\hat{H}_1 \uplus \hat{H}_2]$, $\rho_f = \rho$, and $v = \text{empty}$. We choose $H_f = \text{emp}$, therefore the goals become:

- **G1:** $\llbracket \text{skip} \rrbracket_{[\hat{H}_1], \rho} = ([\hat{H}_1], \rho, \text{empty})$
- **G2:** $\text{emp}, \rho, \epsilon \models \text{emp}$.
- **G3:** $[\hat{H}_1 \uplus \hat{H}_2] = [\text{emp} \uplus \hat{H}_1 \uplus \hat{H}_2]$

and all hold directly from the definitions and hypotheses.

- [FIELD ASSIGNMENT] We have that $bc = [e_1, e_2] := e_3$ and, after applying **H1**, that $P = (e_1, e_2) \mapsto _$ and $Q = (e_1, e_2) \mapsto e_3$. From the satisfiability of JSIL assertions and **H2**, we obtain that $H = ([e_1]_\rho, [e_2]_\rho) \mapsto V$, for some value V , possibly \emptyset . From **H3** and the semantics of JSIL basic commands, we obtain that $\hat{h}_f = \llbracket ([e_1]_\rho, [e_2]_\rho) \mapsto [e_3]_\rho \uplus \hat{H}_1 \uplus \hat{H}_2 \rrbracket$, $\rho_f = \rho$, and $v = [e]_{\rho_3}$. We choose $H_f = ([e_1]_\rho, [e_2]_\rho) \mapsto [e_3]_\rho$, therefore the goals become:

- **G1**: $\llbracket [e_1, e_2] := e_3 \rrbracket_{\llbracket ([e_1]_\rho, [e_2]_\rho) \mapsto V \uplus \hat{H}_1 \rrbracket, \rho} = (\llbracket ([e_1]_\rho, [e_2]_\rho) \mapsto [e_3]_\rho \uplus \hat{H}_1 \rrbracket, \rho, [e]_{\rho_3})$
- **G2**: $([e_1]_\rho, [e_2]_\rho) \mapsto [e_3]_\rho, \rho, \epsilon \models (e_1, e_2) \mapsto e_3$.
- **G3**: $\llbracket ([e_1]_\rho, [e_2]_\rho) \mapsto [e_3]_\rho \uplus \hat{H}_1 \uplus \hat{H}_2 \rrbracket = \llbracket ([e_1]_\rho, [e_2]_\rho) \mapsto [e_3]_\rho \uplus \hat{H}_1 \uplus \hat{H}_2 \rrbracket$.

and all hold directly from the definitions and hypotheses, noting that $[e_3]_\rho \neq \emptyset$.

- [VAR ASSIGNMENT] We have that $bc = x := e$ and, after applying **H1**, that $P = \text{emp}$ and $Q = x \doteq e$. From the satisfiability of JSIL assertions and **H2**, we obtain that $H = \text{emp}$. From **H3** and the semantics of JSIL basic commands, we obtain that $\hat{h}_f = \llbracket \hat{H}_1 \uplus \hat{H}_2 \rrbracket$, $\rho_f = \rho[x \mapsto [e]_\rho]$, and $v = [e]_\rho$. We choose $H_f = \text{emp}$ therefore the goals become:

- **G1**: $\llbracket x := e \rrbracket_{\llbracket \hat{H}_1 \rrbracket, \rho} = (\llbracket \hat{H}_1 \rrbracket, \rho[x \mapsto [e]_\rho], [e]_\rho)$
- **G2**: $\text{emp}, \rho[x \mapsto [e]_\rho], \epsilon \models x \doteq e$
- **G3**: $\llbracket \hat{H}_1 \uplus \hat{H}_2 \rrbracket = \llbracket \text{emp} \uplus \hat{H}_1 \uplus \hat{H}_2 \rrbracket$

and all hold directly from the definitions and hypotheses.

- [OBJECT CREATION] We have that $bc = x := \text{new}()$ and, applying **H1**, that $P = \text{emp}$ and $Q = (x, @proto) \mapsto \text{null} * \text{emptyFields}(x \mid \{\{\@proto\}\})$. From the satisfiability of JSIL assertions and **H2**, we obtain that $H = \text{emp}$. From **H3** and the semantics of JSIL basic commands, we obtain that $\hat{h}_f = (l, @proto) \mapsto \text{null} \uplus \llbracket \hat{H}_1 \uplus \hat{H}_2 \rrbracket$, $\rho_f = \rho[x \mapsto l]$, and $v = l$, for a fresh location l . We know that $l \notin \text{dom}(\hat{H}_1 \uplus \hat{H}_2)$ from **H3**. We choose $H_f = (l, @proto) \mapsto \text{null} \uplus \left(\biguplus_{m \neq @proto} (l, m) \mapsto \emptyset \right)$ (note $\llbracket H_f \rrbracket = (l, @proto) \mapsto \text{null}$). Therefore the goals become:

- **G1**: $\llbracket \text{new}() \rrbracket_{\llbracket \hat{H}_1 \rrbracket, \rho} = (\llbracket (l, @proto) \mapsto \text{null} \uplus H_1 \rrbracket, \rho[x \mapsto l], l)$
- **G2**: $H_f, \rho[x \mapsto l], \epsilon \models (x, @proto) \mapsto \text{null} * \text{emptyFields}(x \mid \{\{\@proto\}\})$
- **G3**: $(l, @proto) \mapsto \text{null} \uplus \llbracket \hat{H}_1 \uplus \hat{H}_2 \rrbracket = \llbracket H_f \uplus \hat{H}_1 \uplus \hat{H}_2 \rrbracket$.

and all hold directly from the definitions and hypotheses, noting that $l \notin \text{dom}(\hat{H}_1 \uplus \hat{H}_2)$.

- [FIELD DELETION] We have that $bc = \text{delete}(e_1, e_2)$ and, applying **H1**, that $P = (e_1, e_2) \mapsto X * X \neq \emptyset$ and $Q = (e_1, e_2) \mapsto \emptyset$. From the satisfiability of JSIL assertions and **H2**, we obtain that $H = ([e_1]_\rho, [e_2]_\rho) \mapsto \epsilon(X)$, where $[e_2]_\rho \neq @proto$. Note that $\llbracket H \rrbracket = H$ since $X \neq \emptyset$. From **H3** and the semantics of JSIL basic commands, we obtain that $\hat{h}_f = \llbracket \hat{H}_1 \uplus \hat{H}_2 \rrbracket$, $\rho_f = \rho$, and $v = \text{true}$. We choose $H_f = ([e_1]_\rho, [e_2]_\rho) \mapsto \emptyset$ therefore noting that $\llbracket H_f \rrbracket = \text{emp}$ the goals become:

- **G1**: $\llbracket bc \rrbracket_{\llbracket ([e_1]_\rho, [e_2]_\rho) \mapsto \epsilon(X) \uplus \hat{H}_1 \rrbracket, \rho} = (\llbracket \hat{H}_1 \rrbracket, \rho, \text{true})$
- **G2**: $([e_1]_\rho, [e_2]_\rho) \mapsto \emptyset, \rho, \epsilon \models (e_1, e_2) \mapsto \emptyset$
- **G3**: $\llbracket \hat{H}_1 \uplus \hat{H}_2 \rrbracket = \llbracket ([e_1]_\rho, [e_2]_\rho) \mapsto \emptyset \uplus \hat{H}_1 \uplus \hat{H}_2 \rrbracket$.

and all follow directly from the definitions and hypotheses, noting that the disjoint union in **G4** is well-defined due to **H3**.

- [FIELD ACCESS] We have that $\text{bc} = x := [e_1, e_2]$ and, after applying **H1**, that $P = (e_1, e_2) \mapsto X * X \neq \emptyset$ and $Q = (e_1, e_2) \mapsto X * X \neq \emptyset * x \doteq X$. From the satisfiability of JSIL assertions and **H2**, we obtain that $H = (\llbracket e_1 \rrbracket_\rho, \llbracket e_2 \rrbracket_\rho) \mapsto \epsilon(X)$. Note that $\lfloor H \rfloor = H$ since $X \neq \emptyset$. From **H3** and the semantics of JSIL basic commands, we obtain that $\hat{h}_f = \lfloor H \uplus \hat{H}_1 \uplus \hat{H}_2 \rfloor$, $\rho_f = \rho[x \mapsto \epsilon(X)]$, and $v = \epsilon(X)$. We choose $H_f = H = (\llbracket e_1 \rrbracket_\rho, \llbracket e_2 \rrbracket_\rho) \mapsto \epsilon(X)$ therefore the goals become:

- **G1**: $\llbracket x := [e_1, e_2] \rrbracket_{\lfloor (\llbracket e_1 \rrbracket_\rho, \llbracket e_2 \rrbracket_\rho) \mapsto \epsilon(X) \uplus \hat{H}_1 \rfloor, \rho} = (\lfloor (\llbracket e_1 \rrbracket_\rho, \llbracket e_2 \rrbracket_\rho) \mapsto \epsilon(X) \uplus \hat{H}_1 \rfloor, \rho[x \mapsto \epsilon(X)], \epsilon(X))$
- **G2**: $(\llbracket e_1 \rrbracket_\rho, \llbracket e_2 \rrbracket_\rho) \mapsto \epsilon(X), \rho[x \mapsto \epsilon(X)], \epsilon \models (e_1, e_2) \mapsto X * X \neq \emptyset * x \doteq X$.
- **G3**: $\lfloor H \uplus \hat{H}_1 \uplus \hat{H}_2 \rfloor = \lfloor H \uplus \hat{H}_1 \uplus \hat{H}_2 \rfloor$.

and all hold directly from the definitions and hypotheses, noting that $\epsilon(X) \neq \emptyset$.

- [MEMBER CHECK] We have that $\text{bc} = x := \text{hasField}(e_1, e_2)$. After applying **H1**, we obtain $P = (e_1, e_2) \mapsto X$ (**I1**) and $Q = P * x \doteq \text{not}(X = \emptyset)$ (**I2**). Using the satisfiability of JSIL assertions and **H2**, we obtain $H = (\llbracket E_1 \rrbracket_\rho^\epsilon, \llbracket E_2 \rrbracket_\rho^\epsilon) \mapsto \epsilon(X)$ (**I3**).

From **I2** we consider two possible cases:

- $X = \emptyset$ (**I4**)

Due to **I4**, we note $\lfloor H \rfloor = H$. From **H3** and the semantics of JSIL basic commands, we obtain that $\hat{h}_f = \lfloor H \uplus \hat{H}_1 \uplus \hat{H}_2 \rfloor$ and $\rho_f = \rho[x \mapsto \text{true}]$. We choose $H_f = H = (\llbracket e_1 \rrbracket_\rho, \llbracket e_2 \rrbracket_\rho) \mapsto \epsilon(X)$ and therefore the goals become:

- * **G1**: $\llbracket x := \text{hasField}(e_1, e_2) \rrbracket_{\lfloor (\llbracket E_1 \rrbracket_\rho^\epsilon, \llbracket E_2 \rrbracket_\rho^\epsilon) \mapsto \epsilon(X) \uplus \hat{H}_1 \rfloor, \rho} = (\lfloor (\llbracket E_1 \rrbracket_\rho^\epsilon, \llbracket E_2 \rrbracket_\rho^\epsilon) \mapsto \epsilon(X) \uplus \hat{H}_1 \rfloor, \rho[x \mapsto \text{true}], \text{true})$
- * **G2**: $(\llbracket e_1 \rrbracket_\rho, \llbracket e_2 \rrbracket_\rho) \mapsto \epsilon(X), \rho[x \mapsto \text{true}], \epsilon \models P * x \doteq \text{not}(X = \emptyset)$
- * **G3**: $\lfloor H \uplus \hat{H}_1 \uplus \hat{H}_2 \rfloor = \lfloor H \uplus \hat{H}_1 \uplus \hat{H}_2 \rfloor$

and all hold directly from the definitions and hypotheses, noting that $\epsilon(X) \neq \emptyset$.

- $X \neq \emptyset$ (**I5**)

Therefore $H_1 \uplus H_2$ cannot contain the cell $(\llbracket e_1 \rrbracket_\rho, \llbracket e_2 \rrbracket_\rho)$, as the disjoint union $H \uplus \hat{H}_1 \uplus \hat{H}_2$ is well-defined. The goals follow from the $X = \emptyset$ case, with false instead of true.

- [CONSEQUENCE] We have $\{P\} \text{bc} \{Q\}$. From **H1** we obtain $\{P'\} \text{bc} \{Q'\}$ (**I1**), $P \vdash P'$ (**I2**) and $Q' \vdash Q$ (**I3**). From **I2** and **H2** we conclude $H, \rho, \epsilon \models P'$ (**I4**). Applying the inductive hypothesis to **I4**, **I1** and **H3** we obtain $H_f, \rho_f, \epsilon \models Q'$ (**I5**), **G1** and **G3**. Finally, we obtain goal **G2** from **I5** and **I3**.
- [FRAME] We have $\{P * R\} \text{bc} \{Q * R\}$. From **H1** we obtain $\{P\} \text{bc} \{Q\}$ (**I1**). We conclude from the satisfiability of JSIL assertions and **H2**, that there exists a H_1 and H_2 such that $H = H_1 \uplus H_2$ (**I2**), $H_1, \rho, \epsilon \models P$ (**I3**) and $H_2, \rho, \epsilon \models R$ (**I4**). From **I2** we have $\llbracket \text{bc} \rrbracket_{\lfloor H_1 \uplus (H_2 \uplus \hat{H}_1) \uplus \hat{H}_2 \rfloor, \rho} = (\hat{h}_f, \rho_f, v)$ (**I5**). Applying the induction hypothesis to **I1**, **I3** and **I5** we obtain that there exists a JSIL abstract heap H' such that:

- **G1**: $\llbracket \text{bc} \rrbracket_{\lfloor H' \uplus H_2 \uplus \hat{H}_1 \rfloor, \rho} = (\lfloor H' \uplus H_2 \uplus \hat{H}_1 \rfloor, \rho', v)$
- **I6**: $H', \rho', \epsilon \models Q$.
- **G3**: $\hat{h}' = \lfloor H' \uplus H_2 \uplus \hat{H}_1 \uplus \hat{H}_2 \rfloor$.

From **I6**, **I4** and the satisfiability of JSIL assertions, we obtain $H_f, \rho', \epsilon \models Q * R$ (**G2**) where $H_f = H' \uplus H_2$.

- [EXISTS] We have $\{\exists X. P\} \text{ bc } \{\exists X. Q\}$. From **H1** we obtain $\{P\} \text{ bc } \{Q\}$ (**I1**). From the satisfiability of JSIL assertions and **H2**, we obtain $\exists V \in \mathcal{V}_{\text{JSIL}}^L. H, \rho, \epsilon[X \mapsto V] \models P$ (**I2**). Applying the induction hypothesis to **I1**, **I2** and **H3**, we obtain $H_f, \rho_f, \epsilon[X \mapsto V] \models Q$ (**I3**), **G1** and **G3**. Using **I3** and the satisfiability of JSIL assertions we are able to achieve **G2**.

□

3.2 Symbolic Execution Rules for Control Flow Commands

We present symbolic execution rules for the control flow commands in Figure 3.2, these rules allow for verification of procedure specifications.

Procedure specifications have the form $\{P\}_{\text{m}(\bar{x})} \{Q\} \in \text{Spec}$ where m is the procedure name, \bar{x} are the procedure parameters, P is the precondition of the procedure and Q is the postcondition. The pre- and postcondition are JSIL logic assertions. Procedures can return in a normal mode or an error mode, this is denoted by a return flag $fl \in \{\text{nm}, \text{er}\}$. A specification environment, $S : \text{Str} \rightarrow \text{Flag} \rightarrow \text{Spec}$, maps a procedure name and return flag onto its corresponding procedure specification.

The symbolic execution rules take the form $\text{p}, S, \text{m} \vdash_{fl} \{P, j, i\} \rightsquigarrow Q$ where:

- $\text{m} \in \text{Str}$ and $\text{p} \in \text{P}$ respectively denote the JSIL procedure name and program being analysed, and S the specification environment;
- $fl \in \text{Flag}$ denotes the return mode of the procedure currently being executed;
- i denotes the index of the JSIL command to be symbolically executed and j the index of the command that was symbolically executed immediately before i ; and
- P is an assertion describing the precondition of the i th command.

The result of the symbolic execution for control flow commands is an assertion Q which describes the final state reached.

We say that a specification environment S is well-formed if all the specifications in the specification environment are provable using the symbolic execution. Formally, for all fl and m where $S(\text{m}, fl) = \{P\}_{\text{m}(\mathbf{x}_1, \dots, \mathbf{x}_n)} \{Q\}$ then $\text{p}, S, \text{m} \vdash_{fl} \{P, 0, 0\} \rightsquigarrow Q$ and $\text{vars}(P) \cup \text{vars}(Q) \subseteq \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$.

We briefly explain the more complicated rules below.

- [CONDITIONAL GOTO - UNKNOWN] Upon reaching a conditional goto command, where the conditional expression cannot be evaluated to false or true, the proof rule branches down both possibilities.
- [PHI-ASSIGNMENT] We use the notation $i \xrightarrow{k}_{\text{m}} j$ to denote that i is the k -th predecessor of j in procedure m . The phi-assignment rule assigns the k -th element of the list $(\mathbf{x}_1, \dots, \mathbf{x}_n)$ to the variable x , where i is the k -th predecessor of j .
- [RETURN] The terminating commands, return normal and return error, ensure the final symbolic execution state entails the specified postcondition of the procedure.

- [PROCEDURE CALL] Both procedure call rules state that the current state must entail the precondition of the callee and some disjoint frame state P_F . The continuation then combines this frame state with the postcondition of the callee.

Control Flow Commands

BASIC COMMAND			
$p_m(i) = \text{bc}$	$\frac{\text{bc} \in BCmd \quad \{P\} \text{bc} \{Q\} \quad p, S, m \vdash_{fl} \{Q, i, i+1\} \rightsquigarrow Q'}{p, S, m \vdash_{fl} \{P, _, i\} \rightsquigarrow Q'}$		
GOTO			
$p_m(i) = \text{goto } k$	$\frac{p, S, m \vdash_{fl} \{P, i, k\} \rightsquigarrow Q}{p, S, m \vdash_{fl} \{P, _, i\} \rightsquigarrow Q}$		
CONDITIONAL GOTO - TRUE			
$p_m(i) = \text{goto } e \ k \ j$	$\frac{P \vdash \text{true} * e \doteq \text{true} \quad p, S, m \vdash_{fl} \{P, i, k\} \rightsquigarrow Q}{p, S, m \vdash_{fl} \{P, _, i\} \rightsquigarrow Q}$		
CONDITIONAL GOTO - FALSE			
$p_m(i) = \text{goto } e \ k \ j$	$\frac{P \vdash \text{true} * e \doteq \text{false} \quad p, S, m \vdash_{fl} \{P, i, j\} \rightsquigarrow Q}{p, S, m \vdash_{fl} \{P, _, i\} \rightsquigarrow Q}$		
CONDITIONAL GOTO - UNKNOWN			
$p_m(i) = \text{goto } e \ k \ j$	$\frac{P \not\vdash \text{true} * e \doteq \text{false} \quad P \not\vdash \text{true} * e \doteq \text{true} \quad p, S, m \vdash_{fl} \{P * (e \doteq \text{true}), i, k\} \rightsquigarrow Q \quad p, S, m \vdash_{fl} \{P * (e \doteq \text{false}), i, j\} \rightsquigarrow Q}{p, S, m \vdash_{fl} \{P, _, i\} \rightsquigarrow Q}$		
PHI-ASSIGNMENT			
$p_m(i) = x := \phi(x_1, \dots, x_n)$	$\frac{j \xrightarrow{k}_m i \quad p, S, m \vdash_{fl} \{P * x \doteq x_k, i, i+1\} \rightsquigarrow Q}{p, S, m \vdash_{fl} \{P, j, i\} \rightsquigarrow Q}$		
PROCEDURE CALL - NORMAL			
$p_m(i) = x := e_0(e_1, \dots, e_{n_1})$ with k	$\frac{S(m', nm) = \{P'\} \ m'(x_1, \dots, x_{n_2}) \ \{Q' * \text{xret} \doteq e\} \quad e_n = \text{undefined} \mid_{n=n_1+1}^{n_2} \quad P \vdash P_F * P' [e_i/x_i]_{i=1}^{n_2} * e_0 \doteq m' \quad Q'' = P_F * Q' [e_i/x_i]_{i=1}^{n_2} * e_0 \doteq m' * x \doteq e [e_i/x_i]_{i=1}^{n_2} \quad p, S, m \vdash_{fl} \{Q'', i, i+1\} \rightsquigarrow Q}{p, S, m \vdash_{fl} \{P, _, i\} \rightsquigarrow Q}$		
PROCEDURE CALL - ERROR			
$p_m(i) = x := e_0(e_1, \dots, e_{n_1})$ with k	$\frac{S(m', er) = \{P'\} \ m'(x_1, \dots, x_{n_2}) \ \{Q' * \text{xerr} \doteq e\} \quad e_n = \text{undefined} \mid_{n=n_1+1}^{n_2} \quad P \vdash P_F * P' [e_i/x_i]_{i=1}^{n_2} * e_0 \doteq m' \quad Q'' = P_F * Q' [e_i/x_i]_{i=1}^{n_2} * e_0 \doteq m' * x \doteq e [e_i/x_i]_{i=1}^{n_2} \quad p, S, m \vdash_{fl} \{Q'', i, k\} \rightsquigarrow Q}{p, S, m \vdash_{fl} \{P, _, i\} \rightsquigarrow Q}$		
RETURN - NORMAL	$\frac{S(m, nm) = \{P'\} \ m(x_1, \dots, x_{n_2}) \ \{Q'\} \quad P \vdash Q'}{p, S, m \vdash_{nm} \{P, _, i_{nm}\} \rightsquigarrow Q'}$	RETURN - ERROR	$\frac{S(m, er) = \{P'\} \ m(x_1, \dots, x_{n_2}) \ \{Q'\} \quad P \vdash Q'}{p, S, m \vdash_{er} \{P, _, i_{er}\} \rightsquigarrow Q'}$

Figure 3.2: Control Flow Proof Rules

Example. In order to illustrate the symbolic execution proof rules, we present a proof sketch for the procedure swap. We prove the specification:

Precondition:

$$\left\{ \begin{array}{l} \text{cur_node} \neq \text{null} * (\text{cur_node}, \text{"prev"}) \mapsto \#\text{prev} * \#\text{prev} \neq \emptyset \\ * (\text{cur_node}, \text{"next"}) \mapsto \#\text{next} * \#\text{next} \neq \emptyset \end{array} \right\}$$

Postcondition:

$$\left\{ \begin{array}{l} \text{cur_node} \neq \text{null} * (\text{cur_node}, \text{"prev"}) \mapsto \#next * \#prev \neq \emptyset \\ * (\text{cur_node}, \text{"next"}) \mapsto \#prev * \#next \neq \emptyset * \text{ret} = \#next \end{array} \right\}$$

Flag:
Normal

The specification states that the current node cannot be null and it has two none-non fields: previous and next. The procedure swaps these two fields and returns the value of next.

```

proc swap (cur_node) {
  {
    cur_node ≠ null * (cur_node, "prev") ↦ #prev * #prev ≠ ∅
    * (cur_node, "next") ↦ #next * #next ≠ ∅
  }
  goto [cur_node = null] then else;
  then:
    ret := null;
    goto rlab;
  else:
    prev := [cur_node, "prev"];
    {
      cur_node ≠ null * (cur_node, "prev") ↦ #prev * #prev ≠ ∅
      * (cur_node, "next") ↦ #next * #next ≠ ∅ * prev ≐ #prev
    }
    next := [cur_node, "next"];
    {
      cur_node ≠ null * (cur_node, "prev") ↦ #prev * #prev ≠ ∅
      * (cur_node, "next") ↦ #next * #next ≠ ∅ * prev ≐ #prev
      * next ≐ #next
    }
    [cur_node, "prev"] := next;
    {
      cur_node ≠ null * (cur_node, "prev") ↦ #next * #prev ≠ ∅
      * (cur_node, "next") ↦ #next * #next ≠ ∅ * prev ≐ #prev
      * next ≐ #next
    }
    [cur_node, "next"] := prev;
    {
      cur_node ≠ null * (cur_node, "prev") ↦ #next * #prev ≠ ∅
      * (cur_node, "next") ↦ #prev * #next ≠ ∅ * prev ≐ #prev
      * next ≐ #next
    }
    ret := next;
    {
      cur_node ≠ null * (cur_node, "prev") ↦ #next * #prev ≠ ∅
      * (cur_node, "next") ↦ #prev * #next ≠ ∅ * prev ≐ #prev
      * next ≐ #next * ret ≐ #next
    }
  rlab:
    skip
} with {ret: ret, rlab;};

```

Evaluation. The soundness of the symbolic analysis is established in Theorem 1. The proof connects the symbolic execution control flow commands to the JSIL semantics which are given in Appendix A.0.2.

Theorem 1 (Soundness of Symbolic Analysis). *For any abstract heaps $H, \hat{H}_1, \hat{H}_2 \in \mathcal{H}_{\text{JSIL}}^0$, store $\rho \in \text{Sto}$, logical environment $\epsilon \in \text{Env}$, program $p \in \text{P}$, well-formed specification environment S , JSIL assertions $P, Q \in \mathcal{AS}_{\text{JSIL}}$, procedure name $m \in \text{Str}$ and command labels i and j such that:*

- $H, \rho, \epsilon \models P$ (**H1**),
- $p, S, m \vdash_{fl} \{P, i, j\} \rightsquigarrow Q$ (**H2**),
- $p \vdash \langle [H \uplus \hat{H}_1 \uplus \hat{H}_2], \rho, i, j \rangle \Downarrow_m \langle h_f, \rho_f, o \rangle$ (**H3**)

then there is an abstract heap H_f such that:

- $H_f, \rho_f, \epsilon \models Q$ (**G1**),
- $o = fl\langle v \rangle$ for some value v (**G2**),
- $[H_f \uplus \hat{H}_1 \uplus \hat{H}_2] = h_f$ (**G3**),

$$\bullet \text{ p } \vdash \langle [H \uplus \hat{H}_1], \rho, i, j \rangle \Downarrow_m \langle [H_f \uplus \hat{H}_1], \rho_f, o \rangle \text{ (G4)}$$

Proof. We proceed by induction on the derivation of **H3**.

[PROCEDURE CALL - NORMAL] It follows that $\text{p}_m(j) = \mathbf{x} := \mathbf{e}_0(\mathbf{e}_1, \dots, \mathbf{e}_{n_1})$ with k for a given JSIL variable \mathbf{x} , n JSIL expressions $\mathbf{e}_0, \mathbf{e}_1, \dots, \mathbf{e}_{n_1}$, and index k . We conclude, using **H3** and the semantics of JSIL, that:

$$\begin{aligned} \llbracket \mathbf{e}_0 \rrbracket_\rho &= \mathbf{m}' \text{ (I1)} & \text{p}(\mathbf{m}') &= \text{proc } \mathbf{m}'(y_1, \dots, y_{n_2}) \{ \bar{c} \} \text{ (I2)} & \forall_{1 \leq n \leq n_1} \mathbf{v}_n &= \llbracket \mathbf{e}_n \rrbracket_\rho \text{ (I3)} \\ \forall_{n_1 < n \leq n_2} \mathbf{v}_n &= \text{undefined} \text{ (I4)} & \text{p} \vdash \langle [H \uplus \hat{H}_1 \uplus \hat{H}_2], \emptyset[y_i \mapsto \mathbf{v}_i]_{i=1}^{n_2}, 0, 0 \rangle &\Downarrow_{\mathbf{m}'} \langle \mathbf{h}', \rho', \text{nm}(\mathbf{v}') \rangle \text{ (I5)} \\ \text{p} \vdash \langle \mathbf{h}', \rho[x \mapsto \mathbf{v}'], j, j+1 \rangle &\Downarrow_m \langle \mathbf{h}_f, \rho_f, o \rangle \text{ (I6)} \end{aligned}$$

From **H2**, we conclude that:

$$\begin{aligned} \mathcal{S}(\mathbf{m}'', \text{nm}) &= \{P'\} \mathbf{m}''(\mathbf{x}_1, \dots, \mathbf{x}_{n_3}) \{Q' * \text{xret} \doteq \mathbf{e}\} \text{ (I7)} \\ \mathbf{e}_n &= \text{undefined} \Big|_{n=n_1+1}^{n_3} \text{ (I8)} & P \vdash P_F * P'[e_i/x_i]_{i=1}^{n_3} * \mathbf{e}_0 &\doteq \mathbf{m}'' \text{ (I9)} \\ Q'' &= P_F * Q'[e_i/x_i]_{i=1}^{n_3} * \mathbf{e}_0 \doteq \mathbf{m}'' * \mathbf{x} \doteq \mathbf{e}[e_i/x_i]_{i=1}^{n_3} \text{ (I10)} & \text{p}, \mathcal{S}, \mathbf{m} \vdash_{fl} \{Q'', j, j+1\} &\rightsquigarrow Q \text{ (I11)} \end{aligned}$$

From **H1**, **I1**, **I2**, **I7**, and **I9**, we conclude that $\mathbf{m}'' = \llbracket \mathbf{e}_0 \rrbracket_\rho = \mathbf{m}'$, $n_2 = n_3$, and $(\mathbf{x}_1, \dots, \mathbf{x}_{n_3}) = (y_1, \dots, y_{n_2})$ (**I12**). For convenience, using (**I12**), we rewrite **I7-I10** as follows:

$$\begin{aligned} \mathcal{S}(\mathbf{m}', \text{nm}) &= \{P'\} \mathbf{m}'(y_1, \dots, y_{n_2}) \{Q' * \text{xret} \doteq \mathbf{e}\} \text{ (I13)} & \mathbf{e}_n &= \text{undefined} \Big|_{n=n_1+1}^{n_2} \text{ (I14)} \\ P \vdash P_F * P'[e_i/y_i]_{i=1}^{n_2} * \mathbf{e}_0 &\doteq \mathbf{m}' \text{ (I15)} & Q'' &= P_F * Q'[e_i/y_i]_{i=1}^{n_2} * \mathbf{e}_0 \doteq \mathbf{m}' * \mathbf{x} \doteq \mathbf{e}[e_i/y_i]_{i=1}^{n_2} \text{ (I16)} \end{aligned}$$

Noting that all the specs in \mathcal{S} are well-formed, we conclude, from **I13** that $\text{vars}(P') \cup \text{vars}(Q') \cup \text{vars}(\mathbf{e}) \subseteq \{y_1, \dots, y_{n_2}\}$ (**I17**). From **H1** and **I15**, we conclude that there are two heaps H_1 and H_2 such that: $H = H_1 \uplus H_2$ (**I18**), $H_1, \rho, \epsilon \models P_F$ (**I19**), and $H_2, \rho, \epsilon \models P'[e_i/y_i]_{i=1}^{n_2}$ (**I20**). Applying the Substitution Lemma for Assertions (Lemma 3, given in the Appendix 3) to **I20** and **I17**, we conclude that $H_2, \emptyset[y_i \mapsto \llbracket \mathbf{e}_i \rrbracket_\rho]_{i=1}^{n_2}, \epsilon \models P'$ (**I21**). From **I3**, **I4**, and **I21**, we conclude that:

$$H_2, \emptyset[y_i \mapsto \mathbf{v}_i]_{i=1}^{n_2}, \epsilon \models P' \text{ (I22)}$$

Because the specification environment is well-formed, we conclude that:

$$\text{p}, \mathcal{S}, \mathbf{m}' \vdash_{\text{nm}} \{P', 0, 0\} \rightsquigarrow (Q' * \text{xret} \doteq \mathbf{e}) \text{ (I23)}$$

Applying the induction hypothesis to **I5**, **I18**, **I22**, and **I23**, we conclude that there is an abstract H' such that: $H', \rho', \epsilon \models (Q' * \text{xret} \doteq \mathbf{e})$ (**I24**), $[H' \uplus H_1 \uplus \hat{H}_1 \uplus \hat{H}_2] = \mathbf{h}'$ (**I25**), and $\text{p} \vdash \langle [H_2 \uplus H_1 \uplus \hat{H}_1], \rho, 0, 0 \rangle \Downarrow_{\mathbf{m}'} \langle [H' \uplus H_1 \uplus \hat{H}_1], \rho', \text{nm}(\mathbf{v}') \rangle$ (**I26**). Since we only consider programs in SSA, we conclude, from **I26**, that $\rho' \geq \emptyset[y_i \mapsto \mathbf{v}_i]_{i=1}^{n_2}$ (**I27**). From **I17**, **I24**, and **I27**, we conclude that:

$$H', \emptyset[y_i \mapsto \mathbf{v}_i]_{i=1}^{n_2}, \epsilon \models Q' \text{ (I28)} \quad \llbracket \mathbf{e} \rrbracket_{\emptyset[y_i \mapsto \mathbf{v}_i]_{i=1}^{n_2}}^\epsilon = \rho'(\text{xret}) = \mathbf{v}' \text{ (I29)}$$

Applying the Substitution Lemma to **I28**, we conclude that $H', \rho, \epsilon \models Q'[e_i/y_i]_{i=1}^{n_2}$ (**I30**). Applying the Substitution Lemma for Expressions to **I29**, we conclude that: $\llbracket \mathbf{e}[e_i/y_i]_{i=1}^{n_2} \rrbracket_\rho^\epsilon = \mathbf{v}'$ (**I31**). From **H1**, **I15**, **I16**, **I19**, **I30**, and **I31**, we conclude that: $H_1 \uplus H', \rho, \epsilon \models Q''$ (**I32**). Applying the induction hypothesis to **I6**, **I11**, **I25**, **I32**, we conclude that there is an abstract H_f such that: $H_f, \rho_f, \epsilon \models Q$ (**G1**), $[H_f \uplus \hat{H}_1 \uplus \hat{H}_2] = \mathbf{h}_f$ (**G3**), and $\text{p} \vdash \langle [H_1 \uplus H' \uplus \hat{H}_1], \rho[x \mapsto \mathbf{v}'], j, j+1 \rangle \Downarrow_m \langle [H_f \uplus \hat{H}_1], \rho_f, o \rangle$ (**I33**) and there is a value \mathbf{v} such that $o = fl(\mathbf{v})$ (**G2**). **G4** follows from **I1-I4**, **I26**, and **I33**.

[BASIC COMMAND] It follows that $p_m(j) = \text{bc} \in \text{BCmd}$ (I1). We conclude, using H3 and the semantics of JSIL, that:

$$\llbracket \text{bc} \rrbracket_{\lfloor H \uplus \hat{H}_1 \uplus \hat{H}_2 \rfloor, \rho} = (h', \rho', v) \text{ (I2)} \quad p \vdash \langle h', \rho', j, j+1 \rangle \Downarrow_m \langle h_f, \rho_f, o \rangle \text{ (I3)}$$

From H2 we conclude that:

$$\{P\} \text{bc} \{Q\} \text{ (I4)} \quad p, S, m \vdash_{fl} \{Q, j, j+1\} \rightsquigarrow Q' \text{ (I5)}$$

Applying the Soundness of Basic Commands (Lemma 1) to H1, I2 and I4, then there is an abstract heap H' such that:

$$\llbracket \text{bc} \rrbracket_{\lfloor H \uplus \hat{H}_1 \rfloor, \rho} = (\lfloor H' \uplus \hat{H}_1 \rfloor, \rho', v) \text{ (I6)} \quad H', \rho', \epsilon \models Q \text{ (I7)} \quad h' = \lfloor H' \uplus \hat{H}_1 \uplus \hat{H}_2 \rfloor \text{ (I8)}$$

We then apply the induction hypothesis to I3, I5, I7 so there exists an abstract heap H_f such that:

$$p \vdash \langle \lfloor H' \uplus \hat{H}_1 \rfloor, \rho', j, j+1 \rangle \Downarrow_m \langle \lfloor H_f \uplus \hat{H}_1 \rfloor, \rho_f, o \rangle \text{ (I9)} \quad o = fl(v) \text{ (G2)}$$

$$h_f = \lfloor H_f \uplus \hat{H}_1 \uplus \hat{H}_2 \rfloor \text{ (G3)} \quad H_f, \rho_f, \epsilon \models Q' \text{ (G1)}$$

We are then able to apply the semantics to I6, I9 and I1 in order to obtain $p \vdash \langle \lfloor H \uplus \hat{H}_1 \rfloor, \rho, i, j \rangle \Downarrow_m \langle \lfloor H_f \uplus \hat{H}_1 \rfloor, \rho_f, o \rangle$ (G4)

[GOTO] It follows that $p_m(j) = \text{goto } k$. We conclude, using H3 and the semantics of JSIL, that $p \vdash \langle \lfloor H \uplus \hat{H}_1 \uplus \hat{H}_2 \rfloor, \rho, i, k \rangle \Downarrow_m \langle h_f, \rho_f, o \rangle$ (I1). From H2, we conclude that $p, S, m \vdash_{fl} \{P, i, k\} \rightsquigarrow Q$ (I2).

Applying the induction hypothesis to I1, I2 and H1 gives goals G1-G3 and $p \vdash \langle \lfloor H \uplus \hat{H}_1 \rfloor, \rho, i, k \rangle \Downarrow_m \langle \lfloor H_f \uplus \hat{H}_1 \rfloor, \rho_f, o \rangle$ (I3). From the semantics and I3, goal G4 follows.

[CONDITIONAL GOTO - TRUE] It follows that $p_m(j) = \text{goto } e \ k \ l$. As we are in the conditional goto - true case we also have $\llbracket e \rrbracket_\rho = \text{true}$ (I1). From H2, we consider the three possible cases:

- $P \vdash \text{true} * e \doteq \text{true}$ (I2)

We conclude, using H3 and the semantics of JSIL, that $p \vdash \langle \lfloor H \uplus \hat{H}_1 \uplus \hat{H}_2 \rfloor, \rho, j, k \rangle \Downarrow_m \langle h_f, \rho_f, o \rangle$ (I3)

From H2 and I2 we conclude that $p, S, m \vdash_{fl} \{P, j, k\} \rightsquigarrow Q$ (I4). Applying the induction hypothesis to I1, I4 and H1, gives $p \vdash \langle \lfloor H \uplus \hat{H}_1 \rfloor, \rho, j, k \rangle \Downarrow_m \langle \lfloor H_f \uplus \hat{H}_1 \rfloor, \rho_f, o \rangle$ (I5) and goals G1-G3. Using the semantics and I5 we obtain $p \vdash \langle \lfloor H \uplus \hat{H}_1 \rfloor, \rho, i, j \rangle \Downarrow_m \langle \lfloor H_f \uplus \hat{H}_1 \rfloor, \rho_f, o \rangle$ (G4).

- $P \vdash \text{true} * e \doteq \text{false}$

From H1 we conclude $\llbracket e \rrbracket_\rho = \text{false}$, this is a contradiction of I1.

- $P \not\vdash \text{true} * e \doteq \text{false}$ and $P \not\vdash \text{true} * e \doteq \text{true}$ (I2)

From H2 and I2 we conclude that $p, S, m \vdash_{fl} \{P * (e \doteq \text{true}), j, k\} \rightsquigarrow Q$ (I3) and $p, S, m \vdash_{fl} \{P * (e \doteq \text{false}), i, l\} \rightsquigarrow Q$. Using H1 and I1, then $H, \rho, \epsilon \models P * (e \doteq \text{true})$ follows (I4).

By applying the semantics of JSIL and using H3, we conclude that $p \vdash \langle \lfloor H \uplus \hat{H}_1 \uplus \hat{H}_2 \rfloor, \rho, j, k \rangle \Downarrow_m \langle \lfloor H_f \uplus \hat{H}_1 \rfloor, \rho_f, o \rangle$ (I5) and goals G1-G3. We can then apply the inductive hypothesis to I3, I4 and I5 to obtain $p \vdash \langle \lfloor H \uplus \hat{H}_1 \rfloor, \rho, i, j \rangle \Downarrow_m \langle \lfloor H_f \uplus \hat{H}_1 \rfloor, \rho_f, o \rangle$ (G4).

[CONDITIONAL GOTO - FALSE] Follows from conditional goto true case.

[PHI-ASSIGNMENT] It follows that $p_m(j) = x := \phi(x_1, \dots, x_n)$. We conclude, using **H3** and the semantics of JSIL, that:

$$i \xrightarrow{m}^k j \text{ (I1)} \quad p \vdash \langle [H \uplus \hat{H}_1 \uplus \hat{H}_2], \rho[x \mapsto \rho(x_k)], j, j+1 \rangle \Downarrow_m \langle h_f, \rho_f, o \rangle \text{ (I2)}$$

From **H2** we conclude that $p, S, m \vdash_{fl} \{P * x \doteq x_k, j, j+1\} \rightsquigarrow Q$ **(I3)**. Noting that we assume programs are only written in SSA form, we can obtain from **H1** $H, \rho[x \mapsto \rho(x_k)], \epsilon \models P * x \doteq x_k$ **(I4)**.

Applying the induction hypothesis to **I3**, **I2** and **I4** we obtain goals **G1-G3** and $p \vdash \langle [H \uplus \hat{H}_1], \rho[x \mapsto \rho(x_k)], j, j+1 \rangle \Downarrow_m \langle [H_f \uplus \hat{H}_1], \rho_f, o \rangle$ **(I4)**. We can conclude from the semantics, **I4** and **I1** that $p \vdash \langle [H \uplus \hat{H}_1], \rho, i, j \rangle \Downarrow_m \langle [H_f \uplus \hat{H}_1], \rho_f, o \rangle$ **(G4)**.

[RETURN - NORMAL] It follows that $j = i_{nm}$, $fl = nm$ and $Q = Q'$. From the semantics and **H3** we conclude that $[H \uplus \hat{H}_1 \uplus \hat{H}_2] = h_f$ **(G3)** and $o = nm \langle \rho(x_{ret}) \rangle$ **(G2)**. From the semantics we obtain $\vdash \langle [H \uplus \hat{H}_1], \rho, _, i_{nm} \rangle \Downarrow_m \langle [H \uplus \hat{H}_1], \rho, nm \langle \rho(x_{ret}) \rangle \rangle$ **(G4)**. From **H2**, we obtain $P \vdash Q'$ **(I1)**. We then obtain **G1** from **I1** and **H1**.

[RETURN - ERROR] Follows from the return normal case, with $j = i_{er}$ and $fl = er$. □

Chapter 4

Bi-Abductive Symbolic Analysis

Given a JSIL procedure, the bi-abductive symbolic analysis infers a set of specifications under which the execution of the given procedure does not fault. The analysis presented here is a first step towards a fully-fledged symbolic execution analysis. In particular, we do not provide the means for automatically inferring loop invariants and we cannot deal with recursive procedures. In this section, we first give an intuitive account of the bi-abductive analysis (Section 4.1), then we define the bi-abductive axioms for basic commands (Section 4.2) and the bi-abductive proof rules for control flow commands (Section 4.2), and we conclude with a description of the algorithm used to generate specifications for JSIL programs (Section 4.4).

The bi-abductive analysis builds on the symbolic execution given in Chapter 3. In particular, we establish the soundness of the bi-abductive analysis by appealing to the symbolic execution analysis. In a nutshell, we prove that all the specifications inferred using the bi-abductive symbolic analysis are verifiable using the standard symbolic analysis. This claim is made formally in Theorem 2.

4.1 Bi-Abductive Analysis: High Level Description

Here we give an intuitive account of the bi-abductive symbolic analysis. In contrast to the symbolic analysis, for which the initial precondition holds the entire resource which the program needs in order not to fault, in the bi-abductive symbolic analysis we have to compute the missing resource. Hence, in the bi-abductive symbolic analysis, we associate each program point with:

- an assertion P , describing the state (heap and store) when reaching that program point; and
- an assertion M , describing the resources that need to be added to the initial precondition so that the program can reach the current program point without faulting.

We refer to (P, M) as the bi-abductive symbolic state.

Suppose that (P, M) holds before the bi-abductive symbolic execution of a given command c . After analysing c , we obtain a new bi-abductive symbolic state (P', M') . The assertion P' describes the heap and store after executing c and M' extends M with the resources needed for the execution of c not to fault.

In order to better illustrate this, suppose we use the bi-abductive analysis to find the pre- and postconditions of a program containing the following JSIL code:

```

x := [timeout, interval]
y := convertToSeconds(x)
[timeout, interval] := y

```

Figure 4.1 describes how the bi-abductive symbolic analysis handles this code, comparing it with the normal symbolic execution. In the normal symbolic execution, we start with an assertion P , describing all the resources needed for the execution of these three commands not to fault. After symbolically executing the first command, we obtain P_1 (which holds all the resources for the execution of the last two commands not to fault). After execution the last two commands, we obtain P_2 and P_3 , respectively. In contrast, in the bi-abductive symbolic analysis, we also need to update the assertion describing the missing initial resources. It is important to note that M is monotonically increasing. Therefore, in this example, $M_2 \vdash M_1 * \hat{M}$ where \hat{M} describes the resource required to execute the command $y = \text{convertToSeconds}(x)$. Ultimately, we use the final missing assertion M_3 to generate the precondition for the procedure and we use the final symbolic state P_3 for the postcondition.

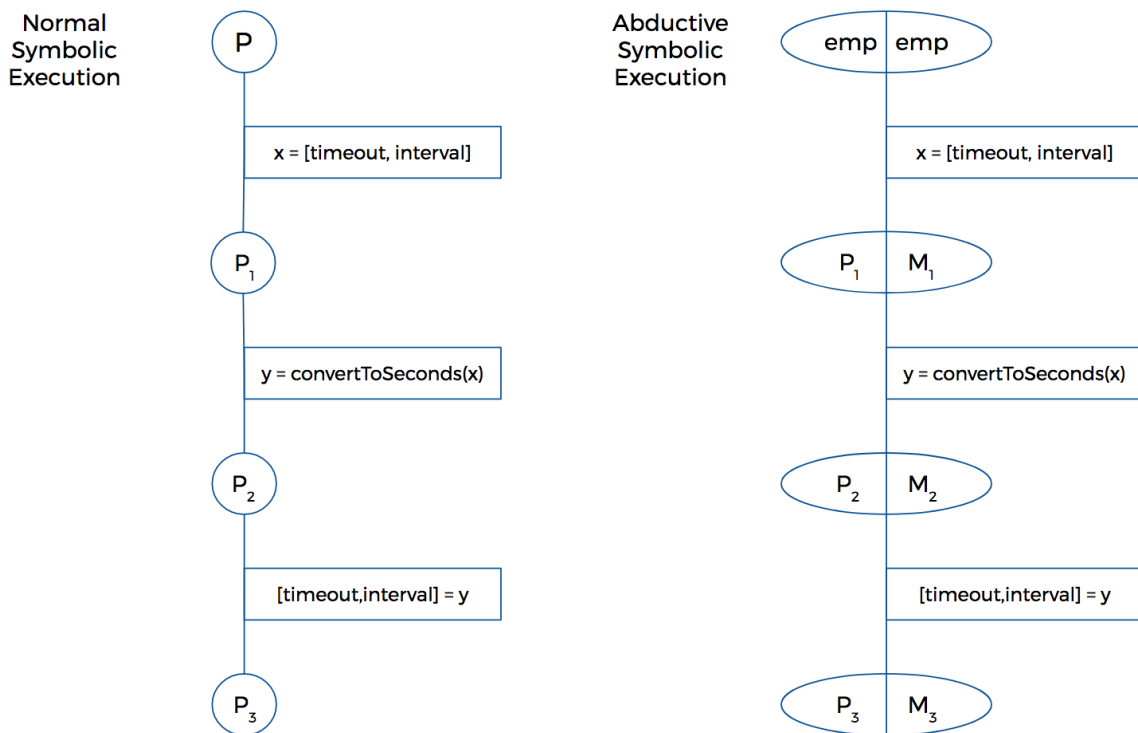


Figure 4.1: Comparison of normal and bi-abductive symbolic execution

In Figure 4.2, we show how the bi-abductive symbolic execution deals with JSIL branching. The execution branches when it is unable to symbolically evaluate the guard of a conditional goto to true or false. When this happens, the bi-abductive symbolic execution branches, obtaining two separate bi-abductive symbolic states: (P_3^T, M_3^T) (for the case in which we assume the guard to evaluate to true) and (P_3^F, M_3^F) (for the false case).

4.2 Bi-Abductive Axioms for Basic Commands

The hoare triples given in Figure 4.3 infer missing resources of the JSIL assertion P needed to execute the command bc . The triples take the form $\{P, M\} bc \{Q, M'\}$, where:

- P is the precondition of bc and M the missing resources computed so far; and

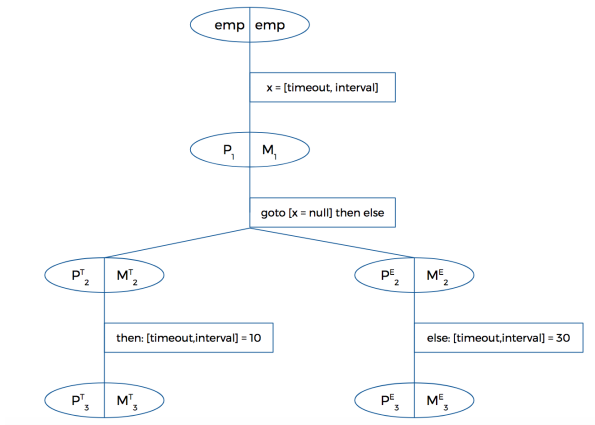


Figure 4.2: Bi-Abductive symbolic execution with branching

- Q is the postcondition of bc and M' the extension of the missing resources for bc not to fault.

If $\{P, M\} bc \{Q, M'\}$ holds, then there is an assertion \hat{M} such that: $M' \vdash M * \hat{M}$ and $\{P * \hat{M}\} bc \{Q\}$ where \hat{M} describes the resources that have to be added to P for the execution of bc not to fault.

Basic Commands

<p style="margin: 0;">FIELD ASSIGNMENT - MISSING</p> $\frac{P \not\vdash \text{true} * (e_1, e_2) \mapsto - \quad Q = P * (e_1, e_2) \mapsto e_3 \quad M' = M * (e_1, e_2) \mapsto -}{\{P, M\} [e_1, e_2] := e_3 \{Q, M'\}}$
<p style="margin: 0;">FIELD DELETION - MISSING</p> $\frac{P \vdash \text{true} * e_2 \not\equiv @proto \quad P \not\vdash \text{true} * (e_1, e_2) \mapsto X * X \not\equiv \emptyset \quad Q = P * (e_1, e_2) \mapsto \emptyset \quad M' = M * (e_1, e_2) \mapsto Y * Y \not\equiv \emptyset}{\{P, M\} \text{delete}(e_1, e_2) \{Q, M'\}}$
<p style="margin: 0;">FIELD ACCESS - MISSING</p> $\frac{P \not\vdash \text{true} * (e_1, e_2) \mapsto X * X \not\equiv \emptyset \quad Q = P * x \doteq X * (e_1, e_2) \mapsto X * X \not\equiv \emptyset \quad M' = M * (e_1, e_2) \mapsto X * X \not\equiv \emptyset}{\{P, M\} x := [e_1, e_2] \{Q, M'\}}$
<p style="margin: 0;">MEMBER CHECK - MISSING</p> $\frac{P \not\vdash \text{true} * (e_1, e_2) \mapsto - \quad Q = P * (e_1, e_2) \mapsto Z * x \doteq \text{not}(Z = \emptyset) \quad M' = M * (e_1, e_2) \mapsto Z}{\{P, M\} x := \text{hasField}(e_1, e_2) \{Q, M'\}}$
<p style="margin: 0;">BASIC COMMAND</p> $\frac{\{P\} bc \{Q\}}{\{P, M\} bc \{Q, M\}}$

Figure 4.3: Bi-Abductive Hoare Triples for Basic Commands

In order to illustrate the mechanics of the bi-abductive Hoare triples for basic commands, we

give the following example:

$$\frac{\begin{array}{l} P = (\text{timeout}, \text{id}) \mapsto 17 \quad P \not\vdash \text{true} * (\text{timeout}, \text{interval}) \mapsto - \\ Q = P * (\text{timeout}, \text{interval}) \mapsto 400 \\ M' = (\text{timeout}, \text{id}) \mapsto - * (\text{timeout}, \text{interval}) \mapsto - \end{array}}{\{(\text{timeout}, \text{id}) \mapsto 17, (\text{timeout}, \text{id}) \mapsto -\} [\text{timeout}, \text{interval}] := 400 \{Q, M'\}} \text{Field Assignment - Missing}$$

We start in a state with the object `timeout` which has one field `id` and we have already inferred a part of the precondition, $(\text{timeout}, \text{id}) \mapsto -$. In order to perform a field assignment operation, we require that the object `timeout` has the field `interval`. Therefore the resulting assertion, $M' = (\text{timeout}, \text{id}) \mapsto - * (\text{timeout}, \text{interval}) \mapsto -$, consists of the previously inferred assertion plus the newly required heap cell. The resulting state after performing the field assignment is $Q = (\text{timeout}, \text{id}) \mapsto 17 * (\text{timeout}, \text{interval}) \mapsto 400$.

Evaluation. Lemma 2 establishes the soundness of the bi-abductive Hoare triples for basic commands. The lemma connects the triples to the non bi-abductive Hoare triples for basic commands.

Lemma 2 (Soundness of Bi-Abductive Hoare Triples for Basic Commands). *For any assertions $P, M, Q, M' \in \mathcal{AS}_{\text{JSIL}}$ and basic commands $\text{bc} \in \text{BCmd}$ such that:*

- **H1:** $\{P, M\} \text{ bc } \{Q, M'\}$

It follows that there exists an $\hat{M} \in \mathcal{AS}_{\text{JSIL}}$ such that:

- **G1:** $M' \vdash M * \hat{M}$
- **G2:** $\{P * \hat{M}\} \text{ bc } \{Q\}$

Proof. We proceed by case analysis on the structure of $\{P, M\} \text{ bc } \{Q, M'\}$.

[FIELD ASSIGNMENT - MISSING] Given $\text{bc} = [e_1, e_2] := e_3$ and $P \not\vdash \text{true} * (e_1, e_2) \mapsto -$ from the bi-abductive field assignment missing basic command rule and **H1** we obtain: $Q = P * (e_1, e_2) \mapsto e_3$ (**I1**) and $M' = M * (e_1, e_2) \mapsto -$ (**I2**). From **I2** it follows that $\hat{M} = (e_1, e_2) \mapsto -$ where $M' \vdash M * \hat{M}$ (**I3**)(**G1**). Using **I3** we obtain $\{P * (e_1, e_2) \mapsto -\} \text{ bc } \{P * (e_1, e_2) \mapsto e_3\}$ (**G2**) as shown in the proof derivation tree:

$$\frac{\frac{\{(\text{e}_1, \text{e}_2) \mapsto -\} [\text{e}_1, \text{e}_2] := \text{e}_3 \{(\text{e}_1, \text{e}_2) \mapsto \text{e}_3\}}{\{P * (\text{e}_1, \text{e}_2) \mapsto -\} [\text{e}_1, \text{e}_2] := \text{e}_3 \{P * (\text{e}_1, \text{e}_2) \mapsto \text{e}_3\}} \text{Field Assignment}}{\text{Frame}}$$

[FIELD DELETION - MISSING] Given $\text{bc} = \text{delete}(e_1, e_2)$ and $P \not\vdash \text{true} * (e_1, e_2) \mapsto X * X \neq \emptyset$ from the bi-abductive field deletion missing basic command rule and **H1** we obtain $P \vdash \text{true} * (e_2 \neq @proto)$ (**I1**), $Q = P * (e_1, e_2) \mapsto \emptyset$ (**I2**) and $M' = M * ((e_1, e_2) \mapsto Y) * (Y \neq \emptyset)$ (**I3**). From **I3** it follows that $\hat{M} = ((e_1, e_2) \mapsto Y) * (Y \neq \emptyset)$ where $M' \vdash M * \hat{M}$ (**I4**)(**G1**). Using **I4** and **I1**, we obtain goal **G2** as shown in the proof derivation tree:

$$\frac{\frac{\{(e_2 \neq @proto) * ((e_1, e_2) \mapsto Y) * (Y \neq \emptyset)\} \text{ delete}(e_1, e_2) \{(e_1, e_2) \mapsto \emptyset\}}{\{P' * (e_2 \neq @proto) * ((e_1, e_2) \mapsto Y) * (Y \neq \emptyset)\} \text{ delete}(e_1, e_2) \{P' * (e_1, e_2) \mapsto \emptyset\}} \text{Field Deletion}}{\text{Frame}}$$

where $P \vdash P' * (e_2 \neq @proto)$

[FIELD ACCESS - MISSING] Given $bc = x := [e_1, e_2]$ and $P \not\vdash true * (e_1, e_2) \mapsto X * X \neq \emptyset$ from the bi-abductive field access missing basic command rule and **H1** we obtain $Q = P * x \doteq X * (e_1, e_2) \mapsto X * X \neq \emptyset$ (**I1**) and $M' = M * (e_1, e_2) \mapsto X * X \neq \emptyset$ (**I2**). From **I2** it follows that $\hat{M} = (e_1, e_2) \mapsto X * X \neq \emptyset$ where $M' \vdash M * \hat{M}$ (**I3**)(**G1**). Using **I3** and **I1**, we obtain goal **G2** as shown in the proof derivation tree:

$$\frac{\frac{\{(e_1, e_2) \mapsto X * X \neq \emptyset\} \quad x := [e_1, e_2] \quad \{x \doteq X * (e_1, e_2) \mapsto X * X \neq \emptyset\}}{\{P * (e_1, e_2) \mapsto X * X \neq \emptyset\} \quad x := [e_1, e_2] \quad \{P * x \doteq X * (e_1, e_2) \mapsto X * X \neq \emptyset\}} \text{Field Access}}{\{P * (e_1, e_2) \mapsto X * X \neq \emptyset\} \quad x := [e_1, e_2] \quad \{P * x \doteq X * (e_1, e_2) \mapsto X * X \neq \emptyset\}} \text{Frame}}$$

[MEMBER CHECK - MISSING] Given $bc = x := hasField(e_1, e_2)$ and $P \not\vdash true * (e_1, e_2) \mapsto -$ from the bi-abductive member check missing basic command rule and **H1** we obtain $Q = P * (e_1, e_2) \mapsto Z * x \doteq not(Z = \emptyset)$ (**I1**) and $M' = M * (e_1, e_2) \mapsto -$ (**I2**). From **I2** it follows that $\hat{M} = (e_1, e_2) \mapsto Z$ where $M' \vdash M * \hat{M}$ (**I3**)(**G1**). Using **I3** and **I1**, we obtain goal **G2** as shown in the proof derivation tree:

$$\frac{\frac{\{(e_1, e_2) \mapsto Z\} \quad x := hasField(e_1, e_2) \quad \{(e_1, e_2) \mapsto Z * x \doteq not(Z = \emptyset)\}}{\{P * (e_1, e_2) \mapsto Z\} \quad x := hasField(e_1, e_2) \quad \{P * (e_1, e_2) \mapsto Z * x \doteq not(Z = \emptyset)\}} \text{Member Check}}{\{P * (e_1, e_2) \mapsto Z\} \quad x := hasField(e_1, e_2) \quad \{P * (e_1, e_2) \mapsto Z * x \doteq not(Z = \emptyset)\}} \text{Frame}}$$

[BASIC COMMAND] Given $\{P, M\} \quad bc \quad \{Q, M\}$ by applying the basic command rule we obtain $\{P\} \quad bc \quad \{Q\}$ (**I1**). Therefore, from **I1**, $\hat{M} = emp$ where $M \vdash M * emp$ (**G1**). By applying the consequences rule we conclude $\{P * emp\} \quad bc \quad \{Q\}$ (**G2**).

□

4.3 Bi-Abductive Proof Rules for Control Flow Commands

The bi-abductive proof rules for the symbolic execution of the control flow commands are given in Figure 4.4. The rules take the form $p, S, m \vdash \{P, M, j, i\} \rightsquigarrow_{ab} \Pi$, where:

- $m \in Str$ and $p \in P$ respectively denote the JSIL procedure name and program being analysed, and S the specification environment;
- i denotes the index of the JSIL command to be symbolically executed and j the index of the command that was symbolically executed immediately before i ; and
- P is an assertion describing the precondition of the i th command, whereas M is an assertion describing the missing resources computed so far.

The result of the bi-abductive symbolic execution for control flow commands is a set Π consisting of triples of the form (P, M, fl) where: P is a possible postcondition, M is the anti-frame that needs be added to the precondition so that the execution of the procedure does not fault, and fl is the return flag for this execution. We briefly explain the more complicated rules below.

- [CONDITIONAL GOTO - UNKNOWN], captures the case when the expression e cannot be evaluated to `false` or `true`. Therefore, the symbolic execution explores both paths, adding the

condition to both the current state and the missing state. However, the condition cannot be added to the missing state if it contains a variable not included in either the collected missing state so far or the procedure arguments. Adding such a variable would entail the generation of a non-well-formed specification. The final result is the union of the results from each path.

- [RETURN] normal and error, these rules return a tuple containing the current state, the collected missing state and the return flag. The return flag is dependant on the index of the final command.
- [PROCEDURE CALL] rule includes the solving of the bi-abductive problem. More concretely, given the current state P and the callee's precondition P' , the system needs to find the assertions P_M and P_F , such that $P * P_M \vdash P' * P_F$. The required missing state is then added to the collected anti-frame and the untouched frame state is combined with the callee's postcondition.

Example. We present a swap procedure example in order to illustrate the bi-abductive proof rules. In particular, if the current node argument is `null` the procedure returns `null`. Otherwise, the procedure swaps the next and previous fields of the current node. In the example we show how the current state is transformed and how the missing state is collected. The result contains two tuples, one with the final state and missing state of the `null` case and the other with the swap case. Both cases have normal return flags.

```

proc swap (cur_node) {
  State: { emp }
  Missing: { emp }
  goto [cur_node = null] then else;
  State: { cur_node ≐ null }
  Missing: { cur_node ≐ null }
  then:   ret := null;
         State: { cur_node ≐ null * ret ≐ null }
         Missing: { cur_node ≐ null }
         goto rlab;
  State: { cur_node ≠ null }
  Missing: { cur_node ≠ null }
  else:   prev := [cur_node, "prev"];
         State: { cur_node ≠ null * (cur_node, "prev") ↦ #prev * #prev ≠ ∅ }
         Missing: { cur_node ≠ null * (cur_node, "prev") ↦ #prev * #prev ≠ ∅ }
         next := [cur_node, "next"];
         State: { cur_node ≠ null * (cur_node, "prev") ↦ #prev * #prev ≠ ∅
                 * prev ≐ #prev * (cur_node, "next") ↦ #next * #next ≠ ∅
                 * next ≐ #next }
         Missing: { cur_node ≠ null * (cur_node, "prev") ↦ #prev * #prev ≠ ∅
                   * (cur_node, "next") ↦ #next * #next ≠ ∅ }
         [cur_node, "prev"] := next;
         State:
         { cur_node ≠ null * (cur_node, "prev") ↦ #next * #prev ≠ ∅
           * prev ≐ #prev * (cur_node, "next") ↦ #next * #next ≠ ∅ *
           next ≐ #next }
         Missing: { cur_node ≠ null * (cur_node, "prev") ↦ #prev * #prev ≠ ∅
                   * (cur_node, "next") ↦ #next * #next ≠ ∅ }
         [cur_node, "next"] := prev;
         State: { cur_node ≠ null * (cur_node, "prev") ↦ #next * #prev ≠ ∅
                 * prev ≐ #prev * (cur_node, "next") ↦ #prev * #next ≠ ∅ *
                 next ≐ #next }

```

```

Missing: { cur_node ≠ null * (cur_node, "prev") ↦ #prev * #prev ≠ ∅
          * (cur_node, "next") ↦ #next * #next ≠ ∅ }
ret := next;
State: { cur_node ≠ null * (cur_node, "prev") ↦ #next * #prev ≠ ∅
        * prev ≐ #prev * (cur_node, "next") ↦ #prev * #next ≠ ∅ *
        next ≐ #next * ret = #next }
Missing: { cur_node ≠ null * (cur_node, "prev") ↦ #prev * #prev ≠ ∅
          * (cur_node, "next") ↦ #next * #next ≠ ∅ }
rlab: skip
} with {ret: ret, rlab;};

```

This bi-abductive symbolic execution generates two specifications, given below.

We give the specification for the case where the current node is `null`:

```

Precondition:
{ cur_node ≐ null }

Postcondition:
{ cur_node ≐ null * ret ≐ null }

Flag:
Normal

```

We give the specification for the case where the current node is not `null`, and therefore the values of previous and next have been swapped.

```

Precondition:
{ cur_node ≠ null * (cur_node, "prev") ↦ #prev * #prev ≠ ∅
  * (cur_node, "next") ↦ #next * #next ≠ ∅ }

Postcondition:
{ cur_node ≠ null * (cur_node, "prev") ↦ #next * #prev ≠ ∅
  * (cur_node, "next") ↦ #prev * #next ≠ ∅ * ret = #next }

Flag:
Normal

```

Evaluation. We establish the soundness of the bi-abductive control flow rules in Theorem 2. We assume the specifications of all callee's a procedure will use are in the specification table and there is no recursion in the program. In this proof we connect the bi-abductive control flow rules to the normal control flow rules. We must also update the specification environment in order to ensure a successful termination of the normal symbolic execution.

Theorem 2 (Soundness of Bi-Abductive Control Flow Proof Rules). *For all programs $p \in \mathcal{P}$, specifications $S \in \text{Spec}$, procedure names $m \in \text{Str}$, assertions $P, M \in \mathcal{AS}_{\text{JSIL}}$ and command labels j and i such that:*

- **H1:** $p, S, m \vdash \{P, M, j, i\} \rightsquigarrow_{ab} \Pi$

It follows that for all $(Q_f, M_f, fl_f) \in \Pi$ there exists a JSIL assertions $\hat{M}'', \hat{P} \in \mathcal{AS}_{\text{JSIL}}$ such that:

- **G1:** $M_f \vdash M * \hat{M}''$
- **G2:** $S' = S[(m, fl_f) \mapsto \{\hat{P}\} m(\bar{x}) \{Q_f\}]$
- **G3:** $p, S', m \vdash_{fl_f} \{P * \hat{M}'', j, i\} \rightsquigarrow Q_f$

Proof. We proceed by induction on the structure of bi-abductive control flow proof rules.

[BASIC COMMAND] Given $p_m(i) = \text{bc}$ where $\text{bc} \in BC\text{cmd}$ (I1), using the control flow rules and H1 we obtain $\{P, M\} \text{bc} \{Q, M'\}$ (I2) and $\text{p}, \text{S}, \text{m} \vdash \{Q, M', i, i+1\} \rightsquigarrow_{ab} \Pi$ (I3). Applying the induction hypothesis to I3, we obtain that there exists some \hat{M}' and P' such that $\text{p}, \text{S}', \text{m} \vdash_{fl_f} \{Q * \hat{M}', i, i+1\} \rightsquigarrow Q_f$ (I4) where $M_f \vdash M' * \hat{M}'$ (I5) and $\text{S}' = \text{S}[(\text{m}, fl_f) \mapsto \{P'\} \text{m}(\bar{x}) \{Q_f\}]$ (G2). From I2 and Lemma 2, it follows that $\{P * \hat{M}\} \text{bc} \{Q\}$ (I6) where $M' \vdash M * \hat{M}$ (I7). Hence using I5 and I7 we are able to obtain, $M_f \vdash M * \hat{M} * \hat{M}'$ where $\hat{M}'' = \hat{M}' * \hat{M}$ (G1). We apply the frame rule to I6, which results in $\{P * \hat{M} * \hat{M}'\} \text{bc} \{Q * \hat{M}'\}$ (I7). Using I1, I3 and I7, we obtain goal G3 as shown in the proof derivation:

$$\frac{\begin{array}{c} p_m(i) = \text{bc} \quad \text{bc} \in BC\text{cmd} \\ \{P * \hat{M}''\} \text{bc} \{Q * \hat{M}'\} \quad \text{p}, \text{S}', \text{m} \vdash_{fl_f} \{Q * \hat{M}', i, i+1\} \rightsquigarrow Q_f \end{array}}{\text{p}, \text{S}', \text{m} \vdash_{fl_f} \{P * \hat{M}'', j, i\} \rightsquigarrow Q_f} \text{Basic Command}$$

[CONDITIONAL GOTO - UNKNOWN] Given $p_m(i) = \text{goto } e \ k \ l$ (I1), $P \not\vdash \text{true} * (e \doteq \text{false})$ (I2) and $P \not\vdash \text{true} * (e \doteq \text{true})$ (I3), we conclude using the control flow rules and H1:

$$\begin{array}{l} \text{p}, \text{S}, \text{m} \vdash \{P * (e \doteq \text{true}), M * (e \doteq \text{true}), i, k\} \rightsquigarrow_{ab} \Pi_t \text{ (I4)} \\ \text{p}, \text{S}, \text{m} \vdash \{P * (e \doteq \text{false}), M * (e \doteq \text{false}), i, l\} \rightsquigarrow_{ab} \Pi_e \text{ (I5)} \\ \text{vars}(e) \subseteq \text{vars}(M) \cup \text{vars}(\text{m}) \text{ (I6)} \quad \Pi = \Pi_t \cup \Pi_e \text{ (I7)} \end{array}$$

From I7 there are two possible cases, either $(Q_f, M_f, fl_f) \in \Pi_t$ or $(Q_f, M_f, fl_f) \in \Pi_e$. We take the case where $(Q_f, M_f, fl_f) \in \Pi_t$ (I8). Applying the induction hypothesis to I4, we obtain that there exists some \hat{M} and P' such that $\text{S}' = \text{S}[(\text{m}, fl_f) \mapsto \{P'\} \text{m}(\bar{x}) \{Q_f\}]$ (I9) and $\text{p}, \text{S}', \text{m} \vdash_{fl_f} \{P * (e \doteq \text{true}) * \hat{M}, i, k\} \rightsquigarrow Q_f$ (I10) where $M_f \vdash M * (e \doteq \text{true}) * \hat{M}$ (I11). Using I1 and I10, we obtain goal G3 as shown in the proof derivation:

$$\frac{\begin{array}{c} p_m(i) = \text{goto } e \ k \ l \quad P * (e \doteq \text{true}) * \hat{M} \vdash \text{true} * (e \doteq \text{true}) \\ \text{p}, \text{S}', \text{m} \vdash_{fl_f} \{P * (e \doteq \text{true}) * \hat{M}, i, k\} \rightsquigarrow Q_f \end{array}}{\text{p}, \text{S}', \text{m} \vdash_{fl_f} \{P * (e \doteq \text{true}) * \hat{M}, j, i\} \rightsquigarrow Q_f} \text{Conditional Goto - True}$$

The case $(Q_f, M_f, fl_f) \in \Pi_e$ follows from the previous case.

[PROCEDURE CALL - NORMAL] Given $p_m(i) = \text{x} := e_0(e_1, \dots, e_{n_1})$ with k (I1), then we obtain using the control flow rules and H1:

$$\begin{array}{l} \text{S}(\text{m}', \text{nm}) = \{P'\} \text{m}'(x_1, \dots, x_{n_2}) \{Q' * \text{xret} \doteq e\} \text{ (I2)} \\ e_n = \text{undefined} \Big|_{n=n_1+1}^{n_2} \text{ (I3)} \\ P * P_M \vdash P_F * P'[\text{e}_i/\text{x}_i]_{i=1}^{n_2} * e_0 \doteq \text{m}' \text{ (I4)} \\ Q'' = P_F * Q'[\text{e}_i/\text{x}_i]_{i=1}^{n_2} * e_0 \doteq \text{m}' * \text{x} \doteq e[\text{e}_i/\text{x}_i]_{i=1}^{n_2} \text{ (I5)} \\ \text{p}, \text{S}, \text{m} \vdash \{Q'', M * P_M, i, i+1\} \rightsquigarrow_{ab} \Pi \text{ (I6)} \end{array}$$

Applying the induction hypothesis to I6, we obtain that there exists some \hat{M} and \hat{P} such that $\text{p}, \text{S}', \text{m} \vdash_{fl_f} \{Q'' * \hat{M}, i, i+1\} \rightsquigarrow Q_f$ (I7) where $M_f \vdash M * P_M * \hat{M}$ (I8)(G1) and $\text{S}' = \text{S}[(\text{m}, fl_f) \mapsto \{\hat{P}\} \text{m}(\bar{x}) \{Q_f\}]$ (G2). From I5, $Q'' * \hat{M} = P_F * Q'[\text{e}_i/\text{x}_i]_{i=1}^{n_2} * e_0 \doteq \text{m}' * \text{x} \doteq e[\text{e}_i/\text{x}_i]_{i=1}^{n_2} * \hat{M}$ (I9). Additionally, from I4 it follows that $P * P_M * \hat{M} \vdash P_F * P'[\text{e}_i/\text{x}_i]_{i=1}^{n_2} * e_0 \doteq \text{m}' * \hat{M}$ (I10).

Using **I1,I2,I3,I7,I9** and **I10**, we obtain goal **G3** as shown in the proof derivation:

$$\begin{array}{c}
p_m(i) = x := e_0(e_1, \dots, e_{n_1}) \text{ with } k \\
S(m', nm) = \{P'\} \ m'(x_1, \dots, x_{n_2}) \ \{Q' * \text{xret} \doteq e\} \\
e_n = \text{undefined} \Big|_{n=n_1+1}^{n_2} \\
P * P_M * \hat{M} \vdash P_F * P' [e_i/x_i]_{i=1}^{n_2} * e_0 \doteq m' * \hat{M} \\
Q'' * \hat{M} = P_F * Q' [e_i/x_i]_{i=1}^{n_2} * e_0 \doteq m' * x \doteq e [e_i/x_i]_{i=1}^{n_2} * \hat{M} \\
\frac{p, S', m \vdash_{fl_f} \{Q'' * \hat{M}, i, i+1\} \rightsquigarrow Q_f}{p, S', m \vdash_{fl_f} \{P * P_M * \hat{M}, j, i\} \rightsquigarrow Q_f} \text{Procedure Call - Normal}
\end{array}$$

[PROCEDURE CALL - ERROR] Follows from the procedure call normal case.

[RETURN - NORMAL] Given $i = i_{nm}$, then we obtain using the control flow rules and **H1** that $\Pi = \{(P, M, nm)\}$ (**I1**). Therefore $M_f = M$ and $\hat{M}'' = \text{emp}$, where $M_f \vdash M * \text{emp}$ (**G1**). We set the new specification environment as follows $S' = S[(m, nm) \mapsto \{P'\} \ m(\bar{x}) \ \{P\}]$ (**G2**). Therefore when we lookup the specification in the non bi-abductive normal return rule we get $S'(m, nm) = \{P''\} \ m(x_1, \dots, x_{n_2}) \ \{Q''\}$ (**I2**) where $P'' = P' * M$ and $Q'' = P$ (**I3**). Due to **I3** we obtain $P \vdash Q''$ (**I4**). Using **I2** and **I4**, we obtain goal **G3** as shown in the proof derivation tree:

$$\frac{S'(m, nm) = \{P''\} \ m(x_1, \dots, x_{n_2}) \ \{Q''\} \quad P \vdash Q''}{p, S', m \vdash_{nm} \{P * \text{emp}, j, i_{nm}\} \rightsquigarrow P} \text{Return - Normal}$$

[RETURN - ERROR] Follows from the return normal case.

[ALL OTHER CASES] The phi-assignment, goto, conditional goto true and conditional goto false follow as $\hat{M} = \text{emp}$.

□

4.4 Bi-Abductive Algorithm for JSIL Programs

We provide a procedure level algorithm which converts the result of the symbolic execution proof rules into a specification table. We infer specifications for non recursive procedures. In particular, we start by inferring the specifications of procedures with no procedure calls. We then proceed to infer the specifications of the procedures which only call procedures with already inferred specifications and so on, until all specifications have been inferred. For this algorithm, we assume the procedures in the program p given as input are ordered as such.

Algorithm 1 presents the algorithm to infer specifications of a JSIL program p . The specification environment, S , contains partial specifications. These specifications are optional for each procedure and are not assumed to be correct.

First, the algorithm initialises the new specification environment S' , which will contain the inferred specifications. Next, it loops through all procedures in order. If a partial specification is provided, the initial state is the partial precondition. Otherwise, the initial state is empty. Additionally, the symbolic execution always starts with an empty missing state and the 0-th command.

The algorithm then takes the result of the symbolic execution of the procedure, and loops through all the resulting tuples. It checks if every final state Q_f entails the partial postcondition

if one is provided. Similar to the procedure call rule, this involves solving the bi-abductive problem. More concretely, given a partial postcondition Q and the computed postcondition Q_f , we have to find Q_M such that $Q_f * Q_M \vdash Q * \text{true}$. Intuitively, Q_M describes the missing resources that are required by the provided partial postcondition.

The new specification environment is then extended with the inferred specification. The partial precondition is combined with both the missing resources required for the procedure not to fault as well as the missing resources required by the partial postcondition. This becomes the new precondition. The final state combined with missing resources required by the partial postcondition becomes the new postcondition.

The *InferSpec* function is defined as follows:

$$\Pi := \text{InferSpec}(p, S, m, P) \stackrel{\text{def}}{\Leftrightarrow} p, S, m \vdash \{P, \text{emp}, _, 0\} \rightsquigarrow_{ab} \Pi$$

Algorithm 1 Procedure Level Algorithm

```

1:  $S' := \emptyset$ 
2: for all  $\text{proc } m(\bar{x})\{\bar{c}\} \in p$  do
3:   if  $(m, \_) \in S$  then
4:      $\{P\} m(\bar{x}) \{Q\} = S(m, \_)$ 
5:   else
6:      $P, Q := \text{emp}$ 
7:   end if
8:    $\Pi := \text{InferSpec}(p, S', m, P)$ 
9:   for all  $(Q_f, M_f, fl) \in \Pi$  do
10:    if  $Q_f * Q_M \vdash Q * \text{true}$  then
11:       $S' := S' \cup \{(m, fl) \mapsto \{P * M_f * Q_M\} m(\bar{x}) \{Q_f * Q_M\}\}$ 
12:    end if
13:   end for
14: end for

```

Control Flow Commands

BASIC COMMAND	
$p_m(i) = \text{bc}$	$\frac{\text{bc} \in BCmd \quad \{P, M\} \text{ bc } \{Q, M'\} \quad p, S, m \vdash \{Q, M', i, i + 1\} \rightsquigarrow_{ab} \Pi}{p, S, m \vdash \{P, M, _, i\} \rightsquigarrow_{ab} \Pi}$
GOTO	
$p_m(i) = \text{goto } k$	$\frac{p, S, m \vdash \{P, M, i, k\} \rightsquigarrow_{ab} \Pi}{p, S, m \vdash \{P, M, _, i\} \rightsquigarrow_{ab} \Pi}$
CONDITIONAL GOTO - TRUE	
$p_m(i) = \text{goto } e \ k \ j$	$\frac{P \vdash \text{true} * e \doteq \text{true} \quad p, S, m \vdash \{P, M, i, k\} \rightsquigarrow_{ab} \Pi}{p, S, m \vdash \{P, M, _, i\} \rightsquigarrow_{ab} \Pi}$
CONDITIONAL GOTO - FALSE	
$p_m(i) = \text{goto } e \ k \ j$	$\frac{P \vdash \text{true} * e \doteq \text{false} \quad p, S, m \vdash \{P, M, i, j\} \rightsquigarrow_{ab} \Pi}{p, S, m \vdash \{P, M, _, i\} \rightsquigarrow_{ab} \Pi}$
CONDITIONAL GOTO - UNKNOWN	
$p_m(i) = \text{goto } e \ k \ j$	$\frac{\text{vars}(e) \subseteq \text{vars}(M) \cup \text{vars}(m) \quad P \not\vdash \text{true} * e \doteq \text{false} \quad P \not\vdash \text{true} * e \doteq \text{true} \quad p, S, m \vdash \{P * (e \doteq \text{true}), M * (e \doteq \text{true}), i, k\} \rightsquigarrow_{ab} \Pi \quad p, S, m \vdash \{P * (e \doteq \text{false}), M * (e \doteq \text{false}), i, j\} \rightsquigarrow_{ab} \Pi'}{p, S, m \vdash \{P, M, _, i\} \rightsquigarrow_{ab} \Pi \cup \Pi'}$
PHI-ASSIGNMENT	
$p_m(i) = x := \phi(x_1, \dots, x_n)$	$\frac{j \xrightarrow{k}_m i \quad p, S, m \vdash \{P * x \doteq x_k, M, i, i + 1\} \rightsquigarrow_{ab} \Pi}{p, S, m \vdash \{P, M, j, i\} \rightsquigarrow_{ab} \Pi}$
PROCEDURE CALL - NORMAL	
$p_m(i) = x := e_0(e_1, \dots, e_{n_1})$ with k	$\frac{S(m', nm) = \{P'\} \ m'(x_1, \dots, x_{n_2}) \ \{Q' * \text{xret} \doteq e\} \quad e_n = \text{undefined} \mid_{n=n_1+1}^{n_2} \quad P * P_M \vdash P_F * P'[e_i/x_i \mid_{i=1}^{n_2}] * e_0 \doteq m' \quad Q'' = P_F * Q'[e_i/x_i \mid_{i=1}^{n_2}] * e_0 \doteq m' * x \doteq e[e_i/x_i \mid_{i=1}^{n_2}] \quad p, S, m \vdash \{Q'', M * P_M, i, i + 1\} \rightsquigarrow_{ab} \Pi}{p, S, m \vdash \{P, M, _, i\} \rightsquigarrow_{ab} \Pi}$
PROCEDURE CALL - ERROR	
$p_m(i) = x := e_0(e_1, \dots, e_{n_1})$ with k	$\frac{S(m', er) = \{P'\} \ m'(x_1, \dots, x_{n_2}) \ \{Q' * \text{xerr} \doteq e\} \quad e_n = \text{undefined} \mid_{n=n_1+1}^{n_2} \quad P * P_M \vdash P_F * P'[e_i/x_i \mid_{i=1}^{n_2}] * e_0 \doteq m' \quad Q'' = P_F * Q'[e_i/x_i \mid_{i=1}^{n_2}] * e_0 \doteq m' * x \doteq e[e_i/x_i \mid_{i=1}^{n_2}] \quad p, S, m \vdash \{Q'', M * P_M, i, k\} \rightsquigarrow_{ab} \Pi}{p, S, m \vdash \{P, M, _, i\} \rightsquigarrow_{ab} \Pi}$
RETURN - NORMAL	RETURN - ERROR
$\frac{}{p, S, m \vdash \{P, M, _, i_{nm}\} \rightsquigarrow_{ab} \{(P, M, nm)\}}$	$\frac{}{p, S, m \vdash \{P, M, _, i_{er}\} \rightsquigarrow_{ab} \{(P, M, er)\}}$

Figure 4.4: Abductive Control Flow Command Proof Rules

Chapter 5

Implementation for Bi-Abductive Symbolic Analysis

We give an implementation of the bi-abductive symbolic analysis. Given a JSIL Program, the tool, called JSIL Abduce, infers specifications for the JSIL procedures in the program. We discuss the architecture and limitations of the tool (Section 5.1) and present some of the implementation challenges of the tool (Section 5.2).

5.1 Architecture

JSIL Abduce generates specifications for JSIL procedures, its architecture is given in Figure 5.1. The tool implements the bi-abductive symbolic execution rules and is entirely written in OCaml. JSIL Abduce has two main components: a bi-abductive symbolic execution engine and a structural entailment engine. As input the tool takes JSIL programs and optional partial specifications. The tool verifies each procedure independently, this modular style allows the tool to scale to large programs.

Before symbolic execution begins, a call graph is generated in order to determine the order in which procedures should be analysed. Procedures which call no other procedures are analysed first. Then all procedures which only call procedures with generated specifications are analysed and so on, until all procedures have been analysed. After the symbolic execution, the result of multiple branches are combined. If the tool fails to generate a specification or a branch contains recursion the symbolic execution for that branch ends. Specifications are still generated for the remaining branches.

Symbolic States. The symbolic execution engine transforms symbolic states $(h, \rho, \pi, \Gamma, \delta)$. A symbolic state contains:

- A symbolic heap $h : \mathcal{L} \times \mathcal{X}_{\text{JSIL}} \rightarrow \mathcal{V}_{\text{JSIL}}$ which is a mapping from locations and JSIL variables to JSIL values. Where the locations and JSIL variables represent objects and their fields respectively.
- A variable store $\rho : \mathcal{X}_{\text{JSIL}} \rightarrow \mathcal{V}_{\text{JSIL}}$ which is a mapping from JSIL variables to JSIL values.
- A set of pure formulae $\pi : \mathcal{AS}_{\text{JSIL}}$ which is a set of JSIL assertions.
- A typing environment $\Gamma : \mathcal{X}_{\text{JSIL}} \rightarrow \text{Types}$ which is a map from JSIL variables to JSIL types.

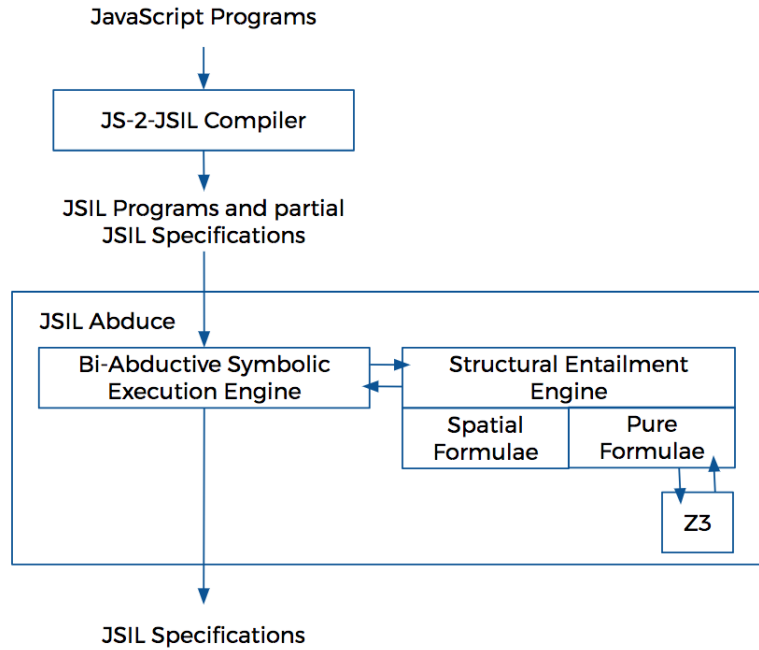


Figure 5.1: Architecture

- A set of predicate assertions δ .

Bi-Abductive Symbolic Execution. The symbolic execution engine implements Algorithm 1, which is explained in Section 4.4. It passes the entailment problem in the procedure call rules and in the final step of the algorithm to the structural entailment engine.

Anti-Frame Inference. We infer missing resources in two stages of the symbolic execution: when calling a procedure and when terminating the symbolic execution.

When the symbolic execution reaches a procedure call command, initially the tool finds all specifications for the procedure being called. The symbolic execution engine then passes the bi-abductive problem $current\ state * missing\ state \vdash precondition * frame$ to the structural entailment engine for each specification. If the current state is able to entail the precondition without the need for missing symbolic state then the symbolic execution transforms the current state into the callee’s postcondition plus any framed state. Otherwise, the tool branches and for each specification the missing footprint is added to anti-frame symbolic state and the current state is transformed into the callee’s postcondition plus any framed state.

When the symbolic execution for a procedure terminates, as outlined in Algorithm 1, we check that the final state entails any partial postcondition. In particular, the symbolic execution engine passes the bi-abductive problem $final\ state * missing\ state \vdash postcondition * frame$ to the structural entailment engine.

Any symbolic state added to the missing footprint must relate to the procedure’s arguments or return value. Otherwise, the generated specification will not be well-formed. Additionally, only type information, heap cells and pure formulae are added to the missing footprint $?M$. Entailments involving pure formulae are checked using the constraint solver Z3, whereas those involving spatial formulae are checked in the structural entailment engine.

5.2 Implementation Challenges

We give examples in order to illustrate some interesting aspects of the JSIL Abduce tool.

Typing and Constraints Complexities. When evaluating expressions the tool infers missing type information and constraints. First, the tool attempts to type an expression with the given type information. If this fails the tool then attempts to infer the types needed to evaluate an expression. This type information is added to the type environment in both the current symbolic state and the missing symbolic state. In addition to type information, the tool infers constraints required to evaluate an expression. For example, when evaluating the expression, `nth(list, 3)` the type `List` of `list` and the constraint that the length of `list` must be greater than 3 is generated.

We give an example to show the additional type information and pure formulae when evaluating expressions. The `getTailElement` procedure is a simple combination of basic commands in order to retrieve the last element of a list. It does not require any objects; instead, it only requires `y` to be `list`. The precondition has the constraint that the list must have at least one element and `y` is of type `List`. These are added when the expressions `l-len(y)` and `l-nth(y, len)` are evaluated.

```
proc getTailElement(y) {
    len := l-len(y) - 1;
    ret := l-nth(y, len);
    rlab: skip
} with {ret: ret, rlab};
```

Generated precondition:

$$\{ (y = \#y) * (!(l-len(\#y) - 1) < 0) * \text{types}(\#y : \$\$list_type) \}$$

Generated postcondition:

$$\left\{ \begin{array}{l} (y = \#y) * (len = (l-len(\#y) - 1)) * (ret = l-nth(\#y, (l-len(\#y) - 1)) * (!(l-len(\#y) - 1) < 0)) \\ * \text{types}(\#y : \$\$list_type, len : \$\$number_type) \end{array} \right\}$$

Return mode:

Normal

Branching. When the tool cannot decide which branch of a `goto` command to explore, it explores both commands. The conditional expression can only be added to the missing symbolic state if the variables are contained within the previously inferred missing state and the parameters. Otherwise the symbolic execution would generate an invalid specification.

In order to demonstrate branching within the tool, we give the following example. The procedure simply sets the interval of the timeout object if the interval is not undefined. If the interval is undefined, the timeout object interval is set to the default value 0. As no partial specification is given the tool initiates the symbolic execution of the procedure with the empty symbolic state. As no information is known about the variable `interval` and it is contained within the procedures parameters, the symbolic execution explores both possible paths. A specification is generated for each path.

```
proc getTimeoutInterval (timeout, interval) {
    goto [(interval = undefined)] then1 else1;
    then: [timeout, "interval"] := 0.;
    goto rlab;
    else: [timeout, "interval"] := interval;
    ret := timeout;
    rlab: skip
} with {ret: ret, rlab};
```

The following specification is generated for the undefined case.

Generated precondition:

$$\{ (timeout, "interval") \mapsto - * (interval = undefined) * \text{types}(timeout : \$\$object_type) \}$$

Generated postcondition:

$$\left\{ \begin{array}{l} (timeout, "interval") \mapsto 0 * (interval = undefined) * (ret = timeout) \\ * \text{types}(timeout : \$\$object_type) \end{array} \right\}$$

Return mode:

Normal

The following specification is generated when the given interval is not undefined.

Generated precondition:

$$\{ (timeout, "interval") \mapsto - * (not (interval = undefined)) * \text{types}(timeout : \$\$object_type) \}$$

Generated postcondition:

$$\left\{ \begin{array}{l} (timeout, "interval") \mapsto interval * (not (interval = undefined)) * (ret = timeout) \\ * \text{types}(timeout : \$\$object_type) \end{array} \right\}$$

Return mode:

Normal

Procedure Calls. The following example demonstrates the procedure call rule. First, we infer the specifications of the procedure `getTimeoutInterval` as it calls no other functions. This procedure requires an object `timeout` with a field `interval`. We infer the pre- and postcondition given below. The specification is added to the new specification environment as shown in Algorithm 1 and then the new specification environment is used for symbolic execution of `getTimeoutIntervalInSeconds`.

```
proc getTimeoutInterval (timeout) {
    ret := [timeout, "interval"];
    rlab: skip
} with {ret: ret, rlab};;
```

Generated precondition:

$$\{ (timeout, "interval") \mapsto \#t * (not(\#t = None)) * \text{types}(timeout : \$\$object_type) \}$$

Generated postcondition:

$$\{ (timeout, "interval") \mapsto \#t * (ret = \#t) * (not(\#t = None)) * \text{types}(timeout : \$\$object_type) \}$$

Return mode:

Normal

Next, we symbolically execute `getTimeoutIntervalInSeconds` which calls the procedure `getTimeoutInterval`. The symbolic execution of `getTimeoutIntervalInSeconds` starts with a state and anti-frame emp . First, the field assignment basic command requires the heap cell assertion $(\text{func}, "timeout") \mapsto -$. The assertion is added to the anti-frame. The procedure call rule then solves the problem $\text{emp} * P_M \vdash P_F * P$ where P is the precondition of `getTimeoutInterval`. This is solved by the entailment engine. As none of the precondition P is contained in the current state, P_M is the precondition of `getTimeoutInterval`. This results in $P_F = (\text{func}, "timeout") \mapsto timeout$.

The resulting state is then the adapted postcondition of `getTimeoutInterval` combined with the frame P_F . The resulting anti-frame is $P_M * (\text{func}, "timeout") \mapsto -$. After symbolically executing the final basic command we obtain the specification below.

```
proc getTimeoutIntervalInSeconds (func) {
    t := [func, "timeout"];
    x := "getTimeoutInterval" (t);
    ret := x / 1000;
    rlab: skip
} with {ret: ret, rlab};;
```

Generated precondition:

$$\left\{ \begin{array}{l} (\text{func}, "timeout") \mapsto \#t * (\#t, "interval") \mapsto \#i \\ * \text{types}(\text{func} : object_type, \#t : object_type, \#i : number_type) \end{array} \right\}$$

Generated postcondition:

$$\left\{ \begin{array}{l} (func, "timeout") \mapsto \#t * (\#t, "interval") \mapsto \#i * (x = \#i) * (ret = (\#i/1000)) \\ * (t = \#t) * \text{types}(func : \$\$object_type, ret : \$\$number_type, \#i : \$\$number_type, t : \$\$object_type) \end{array} \right\}$$

Return mode:

Normal

Chapter 6

Evaluation

The evaluation of the tool was twofold; that of the tools ability to generate specifications and that of the quality of the specifications generated. Initially these two qualities were evaluated by a test suite of JSIL programs. The test suite contains a number of representative examples in order to test the correctness of the symbolic execution rules. In total, the test suite has 47 procedures, producing 60 specifications.

This evaluation was extended by inferring the specifications of the internal and built-in functions. We were able to find a bug in the implementation of two of the JavaScript internal functions. We evaluate the tools ability to generate specifications for these functions (Section 6.2). Additionally, we evaluate the quality of the specifications generated, by comparing hand-written specifications to automatically generated specifications. First, we discuss the infrastructure required to carry out this evaluation (Section 6.1).

6.1 Evaluation Infrastructure

The infrastructure required to verify the specifications generated by JSIL Abduce is given in Figure 6.1. The JSIL programs to be evaluated are given to JSIL Abduce which produces a new specification environment contains the JSIL specifications. The JSIL programs and the new specification environment is then passed into JSIL Verify, to verify the generated specifications. JSIL Abduce, generates a call graph before symbolic execution, to calculate the order procedures are analysed. Procedures must be analysed before any procedure which calls it is analysed.

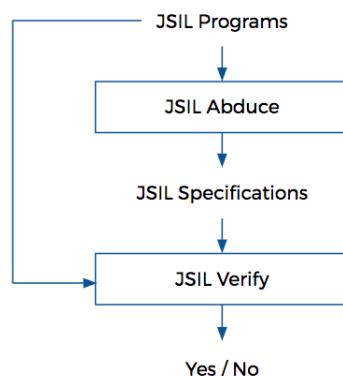


Figure 6.1: Evaluation Infrastructure

Limitations. Currently the basic command `getFields` is not supported by both JSIL Verify and JSIL Abduce, therefore for the evaluation we need to filter out all internal functions containing `getFields`. We perform the filtering while the call graph is generated. Additionally, JSIL Abduce does not support recursive, mutually recursive, or looping branches in procedures. Hence, we detect and halt symbolic execution of these branches, but we continue execution for the remaining branches in the procedure.

6.2 Generating Specifications for the Internal Functions

JSIL implementations for JavaScript internal and built-in functions were written for JSIL Verify. We evaluate the tool's ability to generate specifications for these functions, which involve approximately 3500 lines of JSIL code. JSIL specifications for the internal functions were also written, and can be used to verify JavaScript programs compiled to JSIL. We use these specifications to evaluate the quality of the automatically generated specifications. We give the generated call graph for a subset of these procedures in Figure 6.2.

During the analysis of the internal functions, we found a bug in the implementation of two internal functions, `i_toDataDescriptor` and `i_toAccessorDescriptor`.



Figure 6.2: Call Graph for a subset of the internal functions.

Results. Table 6.2 gives the results of running JSIL Abduce on the internal functions and the separate built-in function libraries. This includes the functions relating to the setup of the JavaScript initial heap contained in the Init library. The number of halted execution branches, relates to the number of branches halted due to recursion. Most of the internal and built-in JSIL implementations contain sizeable procedures with many branches. As a result a significant number of specifications are created per procedure.

Internal and Built-In functions	Lines of code per library	Total number of procedures	Number of procedures containing unsupported operations	Number of halted execution branches	Number of specifications generated	Time taken (in seconds)
Internal	735	30	0	16	287	40.73
Array	467	16	14	8	7	2.92
Boolean	30	4	0	0	28	1.29
Date	94	10	1	231	0	179.17
Errors	87	17	0	107	1	4.37
Functions	47	6	2	2	8	0.45
Global	14	2	0	0	12	0.44
Init	486	6	0	2	11	9.85
Math	73	18	2	0	145	7.27
Number	23	4	2	0	8	0.43
Object	291	19	1	20	31	2.86
String	194	6	3	0	68	5.06
Total	2541	138	25	49	943	254.84

Table 6.1: Results of inferring specifications of internal functions.

6.2.1 Specification Comparison

The semi-automatic JSIL Verify tool requires the JSIL logic specifications of the JavaScript internal functions. These hand-written specifications are verified with respect to the JSIL implementations. We compare these specifications with the automatically generated specifications.

Create Default Object Procedure. The create default object procedure, shown below, creates an object with the default values as defined in the ECMAScript documentation. It has four manually written specifications and four automatically generated specifications.

We evaluate the specifications for the case where both `cl` and `ext` are undefined. Generally the specifications contain similar information about the procedure. The automatically generated specifications contain additional type and spatial information compared to the manually written specifications. This is due to the type inference when evaluating expressions.

```

proc create_default_object (l, pr, cl, ext) {
    goto [cl = undefined] scl text;
    scl:   cl := "Object";
    text:  goto [ext = undefined] sext setall;
    sext:  ext := $$t;

    setall: [l, "@proto"] := pr;
           [l, "@class"] := cl;
           [l, "@extensible"] := ext;

    rlab:  xret := l
}
with
{
    ret:   xret, rlab;
};

```

Hand written specification:

Precondition:

$$\left\{ \begin{array}{l} (l = \#l) * (pr = \#pr) * (cl = \#cl) * (ext = \#ext) * \text{types}(\#l : \$\$object_type) \\ * (\#cl = \text{undefined}) * (\#ext = \text{undefined}) * \text{emptyFields}(\#l :) \end{array} \right\}$$

Postcondition:

$$\left\{ \begin{array}{l} (\#l, "@proto") \mapsto \#pr * (\#l, "@class") \mapsto \text{Object} * (\#l, "@extensible") \mapsto \$\$t \\ * \text{emptyFields}(\#l : "@proto", "@class", "@extensible") * (ret = \#l) \end{array} \right\}$$

Return mode:

Normal

Specification automatically generated by JSIL Abduce:

Precondition:

$$\left\{ \begin{array}{l} (l, "@proto") \mapsto \#s_{-11} * (l, "@class") \mapsto \#s_{-12} * (l, "@extensible") \mapsto \#s_{-13} \\ * (cl = \text{undefined}) * (ext = \text{undefined}) * (pr = \#s_{-14}) * (ext = \text{undefined}) \\ * (cl = \text{undefined}) * \text{types}(l : \$\$object_type, cl : \$\$undefined_type, ext : \$\$undefined_type) \end{array} \right\}$$

Postcondition:

$$\left\{ \begin{array}{l} (l, "@proto") \mapsto \#s_{-14} * (l, "@class") \mapsto \text{"Object"} * (l, "@extensible") \mapsto t \\ * (cl = \text{"Object"}) * (xret = l) * (ext = t) * (pr = \#s_{-14}) \\ * \text{types}(l : \$\$object_type, cl : \$\$string_type, xret : \$\$object_type, ext : \$\$boolean_type) \end{array} \right\}$$

Return mode:

Normal

Error Construct Procedure. The error constructor, shown below, constructs an object with the default fields for an error object. Two specifications were manually written for the procedure whereas seven specifications were automatically generated. We generate more specifications because our tool does not have abstraction. Therefore every time a program branches we generate a new specification. In particular, the hand written specifications generalise the result of calling the procedure `i__toString` which converts `v` to string type. However, the JSIL Abduce tool produces a specification for each possible case of the procedure `i__toString`. The procedure `i__toString` has seven specifications, resulting in significantly more cases.

```

proc Error_construct (xsc, vthis, v) {
    xret := vthis;

    [vthis, "@class"] := "Error";
    [vthis, "@extensible"] := t;

    goto [v = undefined] rlab mess;

    mess:  xerr := "i__toString" (v) with elab;
          [xret, "message"] := {{ "d", xerr, t, f, t }};

    rlab:  skip;
    elab:  skip
}
with
{
    ret:   xret, rlab;
    err:   xerr, elab;
};

```

Is Data Descriptor Procedure. The `isDataDescriptor` procedure, shown below, returns whether a given descriptor is a data descriptor. The five hand written specifications, use the `DataDescriptor` (`d`) and `GenericDescriptor` (`d`) predicates. Predicates can summarise portions of specifications and can lead to more concise specifications. As the tool does not support predicates or abstraction, the specifications generated are less succinct.

```

pred DataDescriptor (d) :
  types (d : $$list_type, #dwrit : $$boolean_type, #denum : $$boolean_type, #
    dconf : $$boolean_type) *
  (! (#dval == $empty)) * (d == {{ "d", #dval, #dwrit, #denum, #dconf }});

pred GenericDescriptor (d) :
  types (d : $$list_type) * (d == {{ "g", #genum, #gconf, #gval, #gwrit, #gget
    , #gset }});

```

With these predicates the hand specification, where the descriptor is a data descriptor, is:

```

Precondition:
{ DataDescriptor(desc) }

Postcondition:
{ ret == t }

Return mode:
Normal

```

Even though the automatically generated specification is correct for this case, it is less concise than the hand-written specifications as the details of the case are not summarised with predicates.

```

Precondition:
{ (desc = #d) * (!(#d = undefined)) * (0. < (l - len#d)) * (l - nth(#d, 0.) = "d") }
  * types(#d : list_type)

Postcondition:
{ (d = l - nth(#d, 0.)) * (desc = #d) * (xret = t) * (!(#d = undefined)) * (0. < (l - len#d)) }
  * (l - nth(#d, 0.) = "d") * types(#d : list_type, desc : list_type, xret : boolean_type)

Return mode:
Normal

```

Setup Initial Heap Procedure. The initial heap setup procedure, is required for any JSIL program compiled from JavaScript. The JavaScript heap requires a significant amount of set up. Therefore the `setupInitialHeap` function contained within the `Init` library consists of 435 lines of code. The automatically generated specifications for this procedure contain approximately 2674 separated assertions in both the precondition and postcondition.

As the first command in a JSIL program compiled from JavaScript is `setupInitialHeap`, unless the initial heap is given in a partial precondition, the entire precondition of `setupInitialHeap` is added to the anti-frame of the main procedure.

Chapter 7

Conclusion

We have progressed from a semi-automatic symbolic analysis to a fully automatic analysis for a subset of JSIL. First, we have successfully modified the symbolic analysis for a more explicit execution. Secondly, we have introduced a bi-abductive symbolic analysis. Finally, we have implemented a bi-abductive tool to infer procedure specifications. We summarise the objectives which have been achieved towards fully automatic verification of JavaScript programs (Section 7.1), and we discuss possible future work (Section 7.2).

7.1 Objectives Achieved

We have succeed in modifying the JSIL logic symbolic rules, in order to describe a symbolic analysis in a style that is closer to the implementation. We have proven this analysis sound with respect to the JSIL semantics.

We have then introduced a bi-abductive symbolic analysis for generating JSIL specifications. This analysis is able to infer the missing resources required to run a given procedure. We have established the soundness of the bi-abductive symbolic analysis in Theorem 2 by appealing to the standard symbolic execution analysis. Additionally, we have presented an algorithm to generate specifications for JSIL programs. This algorithm allows for partial specifications to be given as input to the verification.

Finally, we have developed a tool, JSIL Abduce, which automatically generates specifications for JSIL programs. The tool has generated over 1000 specifications, which have been verified by the semi-automatic tool JSIL Verify. We have evaluated the quality of the generated specifications in Section 6 by comparing them to their handwritten counterparts.

7.2 Future Work

There are two main motivations for generating specifications: to automatically verify properties of JavaScript programs and to support the developer in the verification of functional properties of JavaScript programs. We discuss future work in both of these endeavours. Additionally, future work is required to expand the scope of JSIL programs which the tool can analyse.

Inferring predicate assertions. Currently the tool does not support the inference of logical annotations needed for folding or unfolding user-defined predicates. This is important for partial specifications containing predicate assertions, as well as using the hand-written speci-

fications for the internal functions. For now, due to unsupported operations, in order to infer specifications for JSIL programs compiled from JavaScript, we must use the hand-written internal function specifications. However, issues arise as the hand-written specifications contain a significant number of predicates. Expanding the tool to infer missing annotations needed to fold/unfold predicate assertions, would allow us to generate specifications for a greater range of JSIL programs compiled from JavaScript.

Automatically inferring loop invariants [12]. At present, the tool does not handle looping JSIL procedures. Additionally, the tool does not explore recursive or mutually branches. Future work in automatically generating loop invariants and handling recursion is possible for a pre-established set of abstractions.

Verify properties of JavaScript programs. Future work could explore properties which can be verified by the generated specifications. For example, the specifications are able to record the footprint of each procedure. This can be helpful in detecting confinement violations. For example, detecting whether a program modifies any JavaScript built-in functionality.

Abstraction. As discussed in the evaluation, the quantity of specifications generated for procedures, with multiple branches and procedure calls, is larger than that of the hand-written specifications. This is because the hand-written specifications are able to abstract the result of a procedure call and branching. Expanding the tool to include abstraction would allow us to have the same specification for describing the behaviour of multiple program paths.

Bibliography

- [1] *Proceedings of the 3rd International Constraint Solver Competition (2008)*. pas d'éditeur commercial, 2008.
- [2] T. Ball. The concept of dynamic analysis. *SIGSOFT Softw. Eng. Notes*, 24(6):216–234, Oct. 1999.
- [3] C. Barrett, P. Fontaine, and C. Tinelli. The Satisfiability Modulo Theories Library (SMT-LIB). www.SMT-LIB.org, 2016.
- [4] J. Berdine, C. Calcagno, and P. W. Hearn. Smallfoot: Modular automatic assertion checking with separation logic. In *Proceedings of the 4th International Conference on Formal Methods for Components and Objects, FMCO'05*, pages 115–137, Berlin, Heidelberg, 2006. Springer-Verlag.
- [5] J. Berdine, B. Cook, and S. Ishtiaq. Slayer: Memory safety for systems-level code. In *Proceedings of the 23rd International Conference on Computer Aided Verification, CAV'11*, pages 178–183, Berlin, Heidelberg, 2011. Springer-Verlag.
- [6] A. Biere, A. Cimatti, E. M. Clarke, and Y. Zhu. Symbolic model checking without bdds. In *Proceedings of the 5th International Conference on Tools and Algorithms for Construction and Analysis of Systems, TACAS '99*, pages 193–207, London, UK, UK, 1999. Springer-Verlag.
- [7] M. Bodin, A. Chargueraud, D. Filaretti, P. Gardner, S. Maffeis, D. Naudziuniene, A. Schmitt, and G. Smith. A trusted mechanised javascript specification. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14*, pages 87–100, New York, NY, USA, 2014. ACM.
- [8] R. Brummayer and A. Biere. Boolector: An efficient smt solver for bit-vectors and arrays. In *Proceedings of the 15th International Conference on Tools and Algorithms for the Construction and Analysis of Systems: Held As Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009,, TACAS '09*, pages 174–177, Berlin, Heidelberg, 2009. Springer-Verlag.
- [9] J. N. Buxton and B. Randell, editors. *Software Engineering Techniques: Report of a Conference Sponsored by the NATO Science Committee, Rome, Italy, 27-31 Oct. 1969, Brussels, Scientific Affairs Division, NATO*. 1970.
- [10] C. Cadar, D. Dunbar, and D. Engler. Klee: Unassisted and automatic generation of high-coverage tests for complex systems programs. In *Proceedings of the 8th USENIX Conference on Operating Systems Design and Implementation, OSDI'08*, pages 209–224, Berkeley, CA, USA, 2008. USENIX Association.
- [11] C. Calcagno, D. Distefano, J. Dubreil, D. Gabi, P. Hooimeijer, M. Luca, P. O'Hearn, I. Papakonstantinou, J. Purbrick, and D. Rodriguez. *Moving Fast with Software Verification*, pages 3–11. Springer International Publishing, Cham, 2015.

- [12] C. Calcagno, D. Distefano, P. W. O’Hearn, and H. Yang. Compositional shape analysis by means of bi-abduction. *J. ACM*, 58(6):26:1–26:66, Dec. 2011.
- [13] E. Clarke, A. Biere, R. Raimi, and Y. Zhu. Bounded model checking using satisfiability solving. *Form. Methods Syst. Des.*, 19(1):7–34, July 2001.
- [14] E. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, 1999.
- [15] L. De Moura and N. Bjørner. Z3: An efficient smt solver. In *Proceedings of the Theory and Practice of Software, 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems, TACAS’08/ETAPS’08*, pages 337–340, Berlin, Heidelberg, 2008. Springer-Verlag.
- [16] ECMA. 5th Edition of ECMA 262. ECMAScript Language Specification. Technical report, 2011.
- [17] N. Eén and N. Sörensson. *An Extensible SAT-solver*, pages 502–518. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004.
- [18] R. E. Fairley. Tutorial: Static analysis and dynamic testing of computer software. *Computer*, 11(4):14–23, April 1978.
- [19] P. A. Gardner, S. Maffei, and G. D. Smith. Towards a program logic for javascript. In *Proceedings of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL ’12*, pages 31–44, New York, NY, USA, 2012. ACM.
- [20] T. J. Hickey. IAsolver (2010). <http://www.cs.brandeis.edu/~tim/Applets/IAsolver.html>, 2017 (accessed June 15, 2017).
- [21] C. A. R. Hoare. An axiomatic basis for computer programming. *Commun. ACM*, 12(10):576–580, Oct. 1969.
- [22] M. Huth and M. Ryan. *Logic in Computer Science: modelling and reasoning about systems (second edition)*. Cambridge University Press, 2004. Mass-printing license for Indian sub-continent operative, Portuguese language edition and online e-publishing licenses in preparation, 440pp. <http://www.cs.bham.ac.uk/research/lics/>.
- [23] B. Jacobs, J. Smans, P. Philippaerts, F. Vogels, W. Penninckx, and F. Piessens. Verifast: A powerful, sound, predictable, fast verifier for c and java. In *Proceedings of the Third International Conference on NASA Formal Methods, NFM’11*, pages 41–55, Berlin, Heidelberg, 2011. Springer-Verlag.
- [24] G. R. Joshua Bloch. Extra, extra - read all about it: Nearly all binary searches and mergesorts are broken. <https://research.googleblog.com/2006/06/extra-extra-read-all-about-it-nearly.html>, 2017 (accessed June 15, 2017).
- [25] J. C. King. Symbolic execution and program testing. *Commun. ACM*, 19(7):385–394, July 1976.
- [26] D. Kroening and M. Tautschnig. *CBMC – C Bounded Model Checker*, pages 389–391. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014.
- [27] K. R. M. Leino. Dafny: An automatic program verifier for functional correctness. In *Proceedings of the 16th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning, LPAR’10*, pages 348–370, Berlin, Heidelberg, 2010. Springer-Verlag.
- [28] N. G. Leveson. An investigation of the therac-25 accidents. *IEEE Computer*, 26:18–41, 1993.

- [29] S. Maffeis, J. C. Mitchell, and A. Taly. An operational semantics for javascript. In *Proceedings of the 6th Asian Symposium on Programming Languages and Systems, APLAS '08*, pages 307–325, Berlin, Heidelberg, 2008. Springer-Verlag.
- [30] M. J. Marijn Heule and T. Balyo. Sat competition 2017. <http://baldur.iti.kit.edu/sat-competition-2017/index.php>, 2017 (accessed June 15, 2017).
- [31] D. L. Mate Soos and R. Govostes. Stp constraint solver. <http://stp.github.io/>, 2017 (accessed June 15, 2017).
- [32] G. R. Matthias Heizmann and T. Weber. The 12th international satisfiability modulo theories competition. <http://smtcomp.sourceforge.net/2017/>, 2017 (accessed June 15, 2017).
- [33] S. McLaughlin, C. Barrett, and Y. Ge. Cooperating theorem provers: A case study combining hol-light and cvc lite. *Electron. Notes Theor. Comput. Sci.*, 144(2):43–51, Jan. 2006.
- [34] Microsoft. Microsoft security bulletin ms17-010 - critical. <https://technet.microsoft.com/en-us/library/security/ms17-010.aspx>, 2017 (accessed June 15, 2017).
- [35] D. Naudžiūnienė, J. Fragoso Santos, P. Maksimović, T. Wood, and P. Gardner. An Infrastructure for Tractable Verification of JavaScript Programs. Preprint - in submission to ECOOP'17, 2017.
- [36] N. Nethercote and J. Seward. Valgrind: A framework for heavyweight dynamic binary instrumentation. *SIGPLAN Not.*, 42(6):89–100, June 2007.
- [37] C. Peirce, C. Hartshorne, and P. Weiss. *Collected Papers of Charles Sanders Peirce*. Number v. 5-6 in *Collected Papers of Charles Sanders Peirce*. Belknap Press of Harvard University Press, 1974.
- [38] C. S. Păsăreanu and N. Rungta. Symbolic pathfinder: Symbolic execution of java bytecode. In *Proceedings of the IEEE/ACM International Conference on Automated Software Engineering, ASE '10*, pages 179–180, New York, NY, USA, 2010. ACM.
- [39] C. S. Păsăreanu, P. C. Mehltz, D. H. Bushnell, K. Gundy-Burlet, M. Lowry, S. Person, and M. Pape. Combining unit-level symbolic execution and system-level concrete execution for testing nasa software. In *Proceedings of the 2008 International Symposium on Software Testing and Analysis, ISSTA '08*, pages 15–26, New York, NY, USA, 2008. ACM.
- [40] J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Proceedings of the 17th Annual IEEE Symposium on Logic in Computer Science, LICS '02*, pages 55–74, Washington, DC, USA, 2002. IEEE Computer Society.
- [41] U. S. C. E. R. Team. Indicators associated with wannacry ransomware. <https://www.us-cert.gov/ncas/alerts/TA17-132A>, 2017 (accessed June 15, 2017).
- [42] W. Visser, K. Havelund, G. Brat, S. Park, and F. Lerda. Model checking programs. *Automated Software Engg.*, 10(2):203–232, Apr. 2003.
- [43] H. Yang, O. Lee, J. Berdine, C. Calcagno, B. Cook, D. Distefano, and P. O'Hearn. Scalable shape analysis for systems code. In *Proceedings of the 20th International Conference on Computer Aided Verification, CAV '08*, pages 385–398, Berlin, Heidelberg, 2008. Springer-Verlag.

Appendix A

JSIL Semantics and Substitution Lemma

A.0.1 JSIL Substitution Lemma

Lemma 3 (Substitution Lemma, Assertions). *Let* $\text{vars}(P) \subseteq \{x_i \mid_{i=1}^n\}$. *Then:*

$$H, \rho, \epsilon \models P[e_i/x_i \mid_{i=1}^n] \Leftrightarrow H, \emptyset[x_i \mapsto \llbracket e_i \rrbracket_\rho \mid_{i=1}^n], \epsilon \models P$$

A.0.2 JSIL Semantics

Semantics of JSIL Expressions: $\llbracket e \rrbracket_\rho = v$

LITERAL	VARIABLE	UNARY OPERATOR	BINARY OPERATOR
$\llbracket \lambda \rrbracket_\rho \triangleq \lambda$	$\llbracket x \rrbracket_\rho \triangleq \rho(x)$	$\llbracket \ominus e \rrbracket_\rho \triangleq \overline{\llbracket e \rrbracket_\rho}$	$\llbracket e_1 \oplus e_2 \rrbracket_\rho \triangleq \overline{\llbracket e_1 \rrbracket_\rho, \llbracket e_2 \rrbracket_\rho}$
EXPRESSION LIST			
$\llbracket \{e_1, \dots, e_n\} \rrbracket_\rho \triangleq \{\llbracket e_1 \rrbracket_\rho, \dots, \llbracket e_n \rrbracket_\rho\}$			

Semantics of Basic Commands: $\llbracket bc \rrbracket_{h, \rho} = (h', \rho', v)$

<p>SKIP $\llbracket \text{skip} \rrbracket_{h, \rho} \triangleq (h, \rho, \text{empty})$</p>	<p>ASSIGNMENT $\llbracket x := e \rrbracket_{h, \rho} \triangleq (h, \rho', v)$</p>	<p>PROPERTY ACCESS $\llbracket x := [e_1, e_2] \rrbracket_{h, \rho} \triangleq (h, \rho', v)$</p>
<p>PROPERTY ASSIGNMENT $\llbracket [e_1, e_2] := e_3 \rrbracket_{h, \rho} \triangleq (h', \rho, v)$</p>	<p>PROPERTY DELETION $\llbracket \text{delete}(e_1, e_2) \rrbracket_{h, \rho} \triangleq (h', \rho, \text{true})$</p>	
<p>OBJECT CREATION $\llbracket x := \text{new}() \rrbracket_{h, \rho} \triangleq (h', \rho', l)$</p>		
<p>MEMBER CHECK - TRUE $\llbracket x := \text{hasField}(e_1, e_2) \rrbracket_{h, \rho} \triangleq (h, \rho', \text{true})$</p>	<p>MEMBER CHECK - FALSE $\llbracket x := \text{hasField}(e_1, e_2) \rrbracket_{h, \rho} \triangleq (h, \rho', \text{false})$</p>	
<p>GET FIELDS $\llbracket x := \text{getFields}(e) \rrbracket_{h, \rho} \triangleq (h, \rho', v)$</p>		

Semantics of control flow commands: $p \vdash \langle h, \rho, i, j \rangle \Downarrow_m \langle h', \rho', o \rangle$

BASIC COMMAND

$$\frac{p_m(i) = bc \in \text{BCmd} \quad \llbracket bc \rrbracket_{h,\rho} = \langle h', \rho', - \rangle}{p \vdash \langle h', \rho', i, i+1 \rangle \Downarrow_m \langle h'', \rho'', o \rangle}$$

$$p \vdash \langle h, \rho, _, i \rangle \Downarrow_m \langle h'', \rho'', o \rangle$$

COND. GOTO - TRUE

$$\frac{p_m(i) = \text{goto } [e] j, k \quad \llbracket e \rrbracket_\rho = \text{true}}{p \vdash \langle h, \rho, i, j \rangle \Downarrow_m \langle h', \rho', o \rangle}$$

$$p \vdash \langle h, \rho, _, i \rangle \Downarrow_m \langle h', \rho', o \rangle$$

NORMAL RETURN

$$\vdash \langle h, \rho, _, i_{\text{nm}} \rangle \Downarrow_m \langle h, \rho, \text{nm}(\rho(\text{xret})) \rangle$$

PROCEDURE CALL - NORMAL

$$\frac{p_m(i) = x := e(e_1, \dots, e_{n_1}) \text{ with } j \quad \llbracket e \rrbracket_\rho = m' \quad p(m') = \text{proc } m'(y_1, \dots, y_{n_2}) \{ \bar{c} \} \quad \forall_{1 \leq n \leq n_1} v_n = \llbracket e_n \rrbracket_\rho \quad \forall_{n_1 < n \leq n_2} v_n = \text{undefined}}{p \vdash \langle h, \emptyset[y_i \mapsto v_i]_{i=1}^{n_2}, 0, 0 \rangle \Downarrow_{m'} \langle h', \rho', \text{nm}(v) \rangle}$$

$$p \vdash \langle h', \rho[x \mapsto v], i, i+1 \rangle \Downarrow_m \langle h'', \rho'', o \rangle$$

$$p \vdash \langle h, \rho, _, i \rangle \Downarrow_m \langle h'', \rho'', o \rangle$$

PHI-ASSIGNMENT

$$\frac{p_m(j) = x := \phi(x_1, \dots, x_n) \quad i \xrightarrow{k}_m j \quad p \vdash \langle h, \rho[x \mapsto \rho(x_k)], j, j+1 \rangle \Downarrow_m \langle h', \rho', o \rangle}{p \vdash \langle h, \rho, i, j \rangle \Downarrow_m \langle h', \rho', o \rangle}$$

GOTO

$$\frac{p_m(i) = \text{goto } j \quad p \vdash \langle h, \rho, i, j \rangle \Downarrow_m \langle h', \rho', o \rangle}{p \vdash \langle h, \rho, _, i \rangle \Downarrow_m \langle h', \rho', o \rangle}$$

$$p \vdash \langle h, \rho, _, i \rangle \Downarrow_m \langle h', \rho', o \rangle$$

COND. GOTO - FALSE

$$\frac{p_m(i) = \text{goto } [e] j, k \quad \llbracket e \rrbracket_\rho = \text{false}}{p \vdash \langle h, \rho, i, k \rangle \Downarrow_m \langle h', \rho', o \rangle}$$

$$p \vdash \langle h, \rho, _, i \rangle \Downarrow_m \langle h', \rho', o \rangle$$

ERROR RETURN

$$\vdash \langle h, \rho, _, i_{\text{er}} \rangle \Downarrow_m \langle h, \rho, \text{er}(\rho(\text{xerr})) \rangle$$

PROCEDURE CALL - ERROR

$$\frac{p_m(i) = x := e(e_1, \dots, e_{n_1}) \text{ with } j \quad \llbracket e \rrbracket_\rho = m' \quad p(m') = \text{proc } m'(y_1, \dots, y_{n_2}) \{ \bar{c} \} \quad \forall_{1 \leq n \leq n_1} v_n = \llbracket e_n \rrbracket_\rho \quad \forall_{n_1 < n \leq n_2} v_n = \text{undefined}}{p \vdash \langle h, \emptyset[y_i \mapsto v_i]_{i=1}^{n_2}, 0, 0 \rangle \Downarrow_{m'} \langle h', \rho', \text{er}(v) \rangle}$$

$$p \vdash \langle h', \rho[x \mapsto v], i, j \rangle \Downarrow_m \langle h'', \rho'', o \rangle$$

$$p \vdash \langle h, \rho, _, i \rangle \Downarrow_m \langle h'', \rho'', o \rangle$$