

Efficient Monitoring of Parametric Context-Free Patterns

Patrick O'Neil Meredith and Dongyun Jin and Feng Chen and Grigore Roşu
 University of Illinois at Urbana-Champaign
 {pmeredit, djin3, fengchen, grosu}@cs.uiuc.edu

Abstract— *Recent developments in runtime verification and monitoring show that parametric regular and temporal logic specifications can be efficiently monitored against large programs. However, these logics reduce to ordinary finite automata, limiting their expressivity. For example, neither can specify structured properties that refer to the call stack of the program. While context-free grammars (CFGs) are expressive and well-understood, existing techniques for monitoring CFGs generate large runtime overhead in real-life applications. This paper shows, for the first time, that monitoring parametric CFGs is practical (with overhead on the order of 10% or lower for average cases, several times faster than the state-of-the-art). We present a monitor synthesis algorithm for CFGs based on an LR(1) parsing algorithm, modified with stack cloning to account for good prefix matching. In addition, a logic-independent mechanism is introduced to support matching against the suffixes of execution traces.*

I. INTRODUCTION

Runtime verification (RV) is a relatively new formal analysis approach in which specifications of requirements are given together with the code to check, like in traditional formal verification, but the code is checked against its requirements at runtime, like in testing. A large number of runtime verification approaches and systems, including JPaX [1], TemporalRover [2], JavaMaC [3], Hawk/Eagle [4], Tracematches [5], [6], J-Lo [7], PQL [8], PTQL [9], MOP [10], etc., have been developed recently. In a runtime verification system, monitoring code is generated from the specified properties and integrated with the system to monitor. Therefore, a runtime verification approach consists of at least three interrelated aspects, namely (1) a specification formalism, used to state properties to monitor, (2) a monitor synthesis algorithm, and (3) a program instrumentor. The chosen specification formalism determines the expressivity of the runtime verification approach and/or system.

Monitoring safety properties can be arbitrarily complex [11], but recent developments in runtime verification show that regular and temporal logic formal specifications can be efficiently monitored against large programs. As shown by a series of experiments in the context of Tracematches [6] and JavaMOP [10], parametric regular and temporal logic formal specifications can be monitored against large programs with little runtime overhead: on the order of 10% or lower. However, both regular expressions and temporal logics reduce to ordinary automata when monitored, so they have inherently limited expressivity; more specifically, most runtime verification approaches and systems consider only *flat execution traces*, namely, execution traces without any structure.

Consequently, users of such RV systems are prevented from specifying and checking *structured properties*, referring to the structure of the program. Examples of such structured safety properties include “a resource should be released in the same method which acquired it” or “a resource cannot be accessed if the unsafe method foo is in the current call stack”. A more concrete example is given below.

A. Example

An important and desirable category of properties that cannot be expressed using regular patterns (with or without complement) is one in which pairs of events need match each other, potentially in a nested way. For example, suppose that one prefers to use one’s own locking mechanism for thread synchronization. As usual, for multiple reasons including the allowance of re-entrant synchronized methods (in particular to support recursion), locks are allowed to be acquired and released multiple times by any given thread. However, the lock is effectively released, so that other threads can acquire it, only when the lock releases match the lock acquires. One may want to impose an even stronger locking safety policy, namely that the lock releases should match the lock acquires within the boundaries of each method call. This property is vacuously satisfied when locks are acquired and released in a structured manner using synchronized blocks or methods, like in Java 1.4+, but it may be easily violated when one implements one’s own locking mechanism or uses the semaphores available in Java 5. For example, Figure 1 shows an execution violating this basic safety policy twice (each deeper level symbolizes a nested method invocation): first, the policy is violated when one returns from the last (nested) method invocation, because one does not release the acquired lock; second, the policy is also violated immediately after the return from the last method

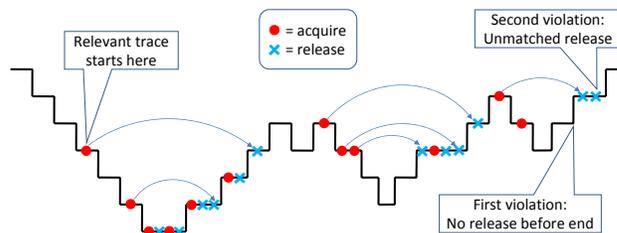


Fig. 1. Example trace for structured acquire and release of locks.

invocation, because the lock is released twice by its caller, but acquired only once.

Supposing that the system is instrumented to emit events begin and end when methods of interest are started and terminated, and events acquire and release when the lock of interest is acquired and released, respectively, then here is a straightforward way to express this safety policy as a context-free grammar:

$$\begin{aligned} S &\rightarrow \epsilon \mid S \text{ acquire } M \text{ release } A \\ M &\rightarrow \epsilon \mid M \text{ begin } M \text{ end} \mid M \text{ acquire } M \text{ release} \\ A &\rightarrow \epsilon \mid A \text{ begin} \mid A \text{ end} \end{aligned}$$

In other words, the pairs of events begin/end and acquire/release should match in a potentially nested way, respectively. S left-recursively refers to itself so that the pattern may be repeated. The non-terminal M stands for “matched” sub-traces, i.e., traces in which all the pairs begin/end and acquire/release are properly matched, and A stands for sequences of (not necessarily matched) begin and end events. A is necessary because the first event seen will be acquire¹ and we do not want to report violations for end’s due to the lack of a matching begin.

Any (finished or unfinished) execution trace that is not a prefix of a word in the language of S in the CFG above is an execution that violates the safety policy. The CFG runtime verification technique presented in this paper and implemented as a logic-plugin in JavaMOP is able to monitor safety properties expressed as CFGs like above². Monitoring-Oriented Programming (MOP) and JavaMOP, the Java implementation of MOP, are discussed in Section III.

```

/*@
scope = global
logic = CFG
SafeLock(Lock l) {
  event acquire<l> : begin(call(* l.acquire()));
  event release<l> : begin(call(* l.release()));
  event begin      : begin(exec(* *.*(..))
                        && !within(Lock));
  event end        : end(exec(* *.*(..))
                        && !within(Lock));

  start symbol : S;
  productions:
  S -> epsilon | S acquire M release A,
  M -> epsilon | M begin M end | M acquire M release,
  A -> epsilon | A begin | A end
}
Violation Handler{
  System.out.println("Unsafe lock operation found!");
}
@*/

```

Fig. 2. JavaMOP specification for the safe lock safety property using the CFG plugin.

Figure 2 shows this safe lock property expressed as a JavaMOP specification, using the CFG logic-plugin discussed in this paper (it is assumed thread local). This specification

¹We wrote the CFG specification this way because we do not want the expense of monitoring `begin` and `end` for *all* methods.

²Our CFG plugin actually supports only the LR(1) language; when we use the term context-free we actually mean LR(1), unless explicitly mentioned otherwise.

makes use of JavaMOP’s generic approach to parametric specifications described in [10]: specifications (and events) can be parametric, in this case in the Lock l , and there could be one or more parameters. The logical formalisms in which properties are expressed need not be aware of the parameters; parameters are added automatically and generically by the JavaMOP framework. Only events which can begin a legal trace are monitor creation events, in other words, only the acquire event above can create new monitor instances. The code generated automatically from the JavaMOP specification in Figure 2, following the technique described in the rest of the paper, has more than 700 lines of (human unreadable) code. We ran this property against a hand-crafted program, which generated the trace in Figure 1. Both violations were successfully caught in a single run because the CFG plugin throws away invalid events, allowing it to continue, and catch multiple violations.

The interested reader can synthesize code for the specification in Figure 2 as well as any other, including all those discussed in the remainder of this paper, using JavaMOP’s online interface [12](the above property is called `SafeLock_CFG`).

B. Contributions

Several approaches have been proposed to monitor context-free properties. For example, PQL [8] is based on context-free languages and Hawk/Eagle [4] uses a fix-point logic. These approaches propose rather complex solutions for monitoring CFG patterns, and generate inefficient monitoring code in many cases, thus preventing the practical use of monitoring context-free properties. This paper shows that monitoring (the LR(1) subset of) parametric context-free patterns is *practical*. Our algorithm modifies an LR(1) parsing algorithm to generate monitors instead of parsers for the defined CFG, and is totally different from the CFG monitoring algorithm in the PQL system [8]. We modify the LR(1) parsing algorithm with a *stack cloning* process, which is needed to test for matches of prefix traces without altering the main monitor stack. In addition, generic suffix matching is added to JavaMOP. Suffix matching is, informally, matching against every suffix of a given trace, and is the mode of matching used in Tracematches [6]. A definition of suffix matching can be found in Section IV. We also describe optimizations that make suffix matching in JavaMOP more efficient.

An extensive evaluation of the CFG monitoring algorithm has been carried out, using the same benchmark (DaCapo) and some of the properties used previously for evaluating runtime verification systems [10], [13] (but now expressed as CFGs). Even when monitored using the CFG plugin, however, these regular pattern based specifications still use constant space. We thus perform an evaluation of three strictly context-free properties – which use theoretically unbounded space – to show that, even with such properties, the overhead is reasonable, and to show that monitoring context-free properties is useful. The results of this analysis compare favorably with PQL and Tracematches, two state-of-the-art runtime monitoring systems. One of these properties (ImprovedLeakingSync) is

expressible in neither PQL nor Tracematches, for reasons that will be explained in Section VI. Another of the properties (SafeFileWriter), while expressible in PQL, is not expressible in Tracematches, because Tracematches only has limited ability to express structured properties, rather than the full generality of (deterministic) context-free languages.

Over both the adapted regular properties and the new strictly context-free properties, the overhead of JavaMOP with CFGs is, on average, over 6 times less than Tracematches on properties that Tracematches is able to express, and over 10 times less than PQL on properties that can be expressed in PQL. On all but 9 of the 66 benchmark/property pairs, the overhead is less than 5% in JavaMOP with CFGs.

C. Paper Outline

The remainder of the paper is as follows: Section II illustrates related work. Section III gives a brief overview of MOP and JavaMOP. Section IV describes suffix matching together with its novel, optimized implementation in JavaMOP. Section V explains our CFG monitor synthesis technique in JavaMOP, including considerations for suffix matching. Section VI explains our experimental setup and the results of our experiments. Section VII concludes the paper and describes some future work.

II. RELATED WORK

Many approaches have been proposed to monitor program execution against formally specified properties. Interested readers can refer to [10] for an extensive discussion on existing runtime monitoring approaches. Briefly, all runtime monitoring approaches except MOP [14], [15], [10] have their specification formalisms hardwired and *no two of them* share the same logic. This observation strengthens our belief underlying MOP, that is, there probably is *no silver-bullet* specification formalism for all purposes. Also, most approaches focus on detecting either violations or validations (matches) of the desired property and support only fixed types of monitors, e.g., online monitors that run together with the monitored program or offline monitors that check the logged execution trace.

Specifically, there are four orthogonal attributes of a runtime monitoring system: logic, scope, running mode, and handlers. The logic specifies which formalism is used to specify the property. The scope determines where to check the property; it can be class invariant, global, interface, etc. The running mode denotes where the monitor runs; it can be inline (weaved into the code), online (operating at the same time as the program), outline (receiving events from the program remotely, e.g., over a socket), or offline (checking logged event traces)³. The handlers specify what actions to perform under exceptional conditions; there can be violation and validation handlers. It is worth noting that for many logics, violation and validation are not complementary to each other, i.e., the violation of a formula does not always imply the validation of the negation of the formula.

³Offline implies outline, and inline implies online.

Approach	Logic	Scope	Mode	Handler
JPaX [1]	LTL	class	offline	violation
TemporalRover [2]	MiTL	class	inline	violation
JavaMaC [3]	PastLTL	class	outline	violation
Hawk [4]	Eagle	global	inline	violation
Tracematches [6]	Reg. Exp.	global	inline	validation
J-Lo [7]	LTL	global	inline	violation
Pal [16]	modified Blast	global	inline	validation
PQL [8]	PQL	global	inline	validation
PTQL [9]	SQL	global	outline	validation

TABLE I
RUNTIME VERIFICATION BREAKDOWN.

Most runtime monitoring approaches can be captured by instantiating these attributes, as partly illustrated in Table I. For example, JPaX can be regarded as an approach that uses linear temporal logic (LTL) to specify class-scoped properties, whose monitors work in offline mode and only detect violations.

Of the systems mentioned in Table I, only PQL [8] and Hawk/Eagle [4] can handle arbitrary context-free grammars. Hawk/Eagle adopts a fix-point logic and uses term rewriting during the monitoring, making it rather inefficient. It also has problems with large programs because it does not garbage collect the objects it uses to monitor. In addition, Hawk/Eagle is not publicly available⁴. Because of this, we cannot compare our CFG plugin with Hawk/Eagle. In addition to PQL, we decided, also, to perform comparisons with Tracematches [6], as it is able to monitor a very limited set of context-free properties using compiler-specific support provided by their special AspectJ compiler, ABC [17], and also because it is a very efficient system. Pal [16] is able to monitor properties that take calls and returns into account, giving a limited context-free ability for this one case. Pal is implemented for C, rather than Java, and the implementation is not publicly available.

III. MOP REVISITED

MOP is an extensible runtime verification framework that provides efficient, logic-independent support for parametric specifications. JavaMOP is an implementation of MOP for the Java programming language. By encapsulating our monitor synthesis algorithm for non-parametric CFG patterns in a JavaMOP logic plugin, we have achieved an efficient monitoring tool for universally quantified parametric CFG specifications.

In addition, we have implemented a novel extension of MOP, in JavaMOP, to support suffix matching, also in a logic-independent way. We define suffix matching as matching against every suffix of a given event trace, while total matching, also supported by JavaMOP, attempts to match the entire trace seen at a particular point. It is notable that Tracematches supports *only* suffix matching, and PQL supports a skip semantics more akin to suffix matching than total, so in our experiments we use suffix matching.

⁴[6] makes an argument for the inefficiency of Hawk/Eagle. Since Hawk/Eagle is not publicly available (only its rewrite based algorithm is public [4]), the authors of Hawk/Eagle kindly agreed to monitor some of the simple properties from [13]. We have confirmed the inefficiency claims of [6] with the authors of Hawk/Eagle.

A. MOP in a Nutshell

MOP [14], [15], [10] is a formal framework for software development and analysis, in which the developer specifies desired properties using *definable* specification formalisms, along with code to execute when properties are violated or validated. Monitoring code is then automatically generated from the specified properties and integrated together with the user-provided code into the original system. MOP is a highly extensible and configurable runtime verification framework. The user is allowed to extend the MOP framework with his/her own logics via *logic plugins* which encapsulate the monitor synthesis algorithms. This extensibility of MOP is supported by an especially designed layered architecture [14], which separates monitor generation and monitor integration. By standardizing the protocols between layers, modules can be added and reused easily and independently.

In addition to choosing the formalism to use in the specification, one can also configure the behaviors of the generated monitor through different attributes [15]. Depending upon configuration, the monitors can be separate programs reading events from a log file, from a socket, or from a buffer, or can be inlined within the program at the event observation points; monitors can verify the observed execution trace as a whole or check fragments of the trace. All these configurations are *independent* of the formalism used to specify the property. MOP also provides efficient and logic-independent support for universally quantified parameters in specifications [10], which is useful for specifying properties related to more than one object. This extension allows associating parameters with MOP specifications, and generating efficient monitoring code from parametric specifications with monitor synthesis algorithms for non-parametric specifications. MOP's generic support for universally quantified patterns simplified our CFG plugin's implementation.

The JavaMOP implementation provides several interfaces, including a web-based interface, a command-line interface and an Eclipse-based GUI, providing the developer with different means to manage and process MOP specifications. JavaMOP follows a client-server architecture [18] to flexibly support these various interfaces, as well as for portability reasons. AspectJ [19] is employed for monitor integration: JavaMOP translates outputs of logic-plugins into AspectJ code, which is then merged within the original program by the AspectJ compiler. Five logic-plugins are currently provided with JavaMOP: Java Modeling Language (JML) [20], Extended Regular Expressions (ERE), Past-Time and Future-time Linear Temporal Logics (LTL) (see [18] for more details), and Context-Free Grammar (CFG) that is introduced in this paper. Note that these plugins can be supported by any implementation of MOP.

One might expect some loss of efficiency for MOP's generality of logics. However, the JavaMOP-generated monitors can yield very reasonable runtime overhead in practice, even for properties requiring intensive runtime checking: on the order of 10% or lower, and as efficient as the hand optimized monitoring code in most cases [10].

IV. SUFFIX MATCHING IN JAVAMOP

According to application requirements, one may want to check the desired property against either *the whole execution trace* (at a given program point) or *every suffix of a trace*. Total matching has been adopted by many runtime verification approaches to detect violations of properties, e.g., JPaX [1] and JavaMaC [3]. Suffix matching has been used mainly by monitoring approaches that aim to find validations of properties, e.g., Tracematches [6]. PQL has a skip semantics, wherein a specification is matched against the trace, but events may be skipped. A precise explanation of PQL's semantics is available in [8]. To define suffix and total matching we first must define traces and properties:

Definition 1. Let \mathcal{E} be a set of events. An \mathcal{E} -trace, or simply a trace when \mathcal{E} is understood from context, is any finite sequence of events in \mathcal{E} , that is, an element in \mathcal{E}^* .

In the context of monitoring, an execution trace is a sequence of events observed up to the current moment, thus execution traces are always finite.

Definition 2. An \mathcal{E} -property P , or simply a property, is a pair of disjoint sets (P_+, P_-) in $\mathcal{E}^* \times \mathcal{E}^*$; P_+ is the set of validating traces and P_- is its set of violating traces.

Therefore, our notation of property is quite general; for each particular specification formalism, one need associate an appropriate property to each formula or pattern in that formalism. For example, for a CFG G , we let $P_G = (P_+, P_-)$ be defined as expected: P_+ is $L(G)$ (see Section V-A) and $w \in P_-$ iff w is not the prefix of any $w' \in L(G)$

Definition 3. The total matching semantics of P is a function

$$\llbracket P \rrbracket_{total} : \mathcal{E}^* \rightarrow \{\text{validation}, \text{violation}, ?\}$$

defined as follows for each $w \in \mathcal{E}^*$:

$$\llbracket P \rrbracket_{total}(w) = \begin{cases} \text{validation} & \text{if } w \in P_+, \\ \text{violation} & \text{if } w \in P_-, \\ ? & \text{otherwise} \end{cases}$$

The suffix matching semantics of P is a function

$$\llbracket P \rrbracket_{partial} : \mathcal{E}^* \rightarrow \{\text{validation}, ?\}$$

defined as follows for each $w \in \mathcal{E}^*$:

$$\llbracket P \rrbracket_{partial}(w) = \begin{cases} \text{validation} & \text{if } \exists w_1, w_2 \text{ such that} \\ & w = w_1 w_2 \text{ and} \\ & \llbracket P \rrbracket_{total}(w_2) = \text{validation} \\ ? & \text{otherwise} \end{cases}$$

Because suffix matching semantics is not defined for violation, violation handlers are not permitted in suffix matching JavaMOP properties (an error is reported).

The previous design of JavaMOP supported only total matching. We have implemented a logic-independent extension of JavaMOP to also support suffix matching. This extension is based on the observation that, although total

matching and suffix matching have inherently different semantics, it is not difficult to support suffix matching in a total matching setting, if one maintains *a set of monitor states* during monitoring and *creates a new monitor instance at each event* (this amounts to checking the property on each suffix incrementally). However, the situation becomes more complicated when one wants to develop a logic-independent solution, since different logical formalisms can have different state representations. For example, the monitor state can be an integer when the monitor is based on a state machine, a vector like the past-time LTL monitor, or a stack such as the CFG monitor discussed below. Hence, our solution is to *treat every monitor as a blackbox* without assumptions on its internal state. Also, instead of maintaining a set of monitor states in the monitor, we use a wrapper monitor that keeps a set of total matching monitors as its state for suffix matching. For simplicity, from now on, when we say “monitor” without specific constraints, we mean the monitor generated for total matching. When an event is received, the wrapper monitor for suffix matching operates as follows:

- 1) create a new monitor and add it to the “suffix matching” internal monitor set;
- 2) invoke every monitor in the monitor set to handle the received event;
- 3) if a monitor enters its “violation” state, remove it from the monitor set;
- 4) if a monitor enters its “validation” state, report the validation.

The third step is used to keep the “suffix matching” monitor set small by removing unnecessary monitors. Informally, this implements suffix matching semantics because each total monitor is monitoring a suffix of the current trace and “validation” is only reported if one of the suffixes is valid.

Using our current implementation of suffix matching in JavaMOP, one may further improve the monitoring efficiency if the monitor provides an optional interface, namely, an equals method that compares two monitors with regard to their internal states, and a hashCode method used to reduce the amount of calls to equals. This interface is used to populate a Java HashSet: the combination of the definition of hashCode and equals ensures the monitors in the HashSet are declared duplicates, and removed, based on monitor state rather than memory location. This interface can be easily generated by each JavaMOP logic plugin, because it has full knowledge of the monitor semantics. It is important to note that our approach does *not depend on the underlying specification formalism*.

Moreover, JavaMOP already requires the logic plugin to designate *creation events* that are the starting events of a validating trace in order to avoid unnecessary monitor creation. A new monitor instance need be created only at creation events. This feature is especially useful when combined with suffix matching, which requires creating a new monitor at every event, if no creation events are chosen.

V. CONTEXT-FREE PATTERNS IN JAVAMOP

We support the LR(1) subset of context-free grammars (CFGs). LR(1) can only recognize a subset of the deterministic context-free languages, which are themselves a strict subset of the context-free languages (CFLs) [21], [22]. LR(1), however, is a large subset. Notably, LR(1) is more powerful than the standard mode of Bison [23], which provides LALR⁵. We chose LR rather than LL because LR recognizes a larger number of grammars without translation than LL. We base our implementation on the Knuth algorithm [24] for LR(1) parser table generation as presented in [21]. While the “action” and “goto” tables generated are normal LR(1) “action” and “goto” tables, the algorithm used to parse has been modified to work in the context of monitoring, explained in more detail below.

A. Preliminaries

A CFG G is defined as a tuple of the form, $G = (NT, \Sigma, P, S)$. Σ is the alphabet of the CFG, often referred to as the set of terminals. NT is the set of non-terminals. P is the set of productions, which define what strings non-terminals can derive. $NT \cup \Sigma$ is often called the set of symbols of the CFG. A production is always of the form $A \rightarrow \gamma$, where $A \in NT$ and γ is a string that either consists of symbols, or is the empty string, ϵ , i.e. $\gamma \in (\Sigma \cup NT)^*$. We use the conventional alternation operator, “|”: a production of the form $A \rightarrow \gamma_0 | \gamma_1$ can be equivalently represented as two productions $A \rightarrow \gamma_0$ and $A \rightarrow \gamma_1$. S is the start symbol – that non-terminal from which all strings in the language are derived. For example, $G = (\{A\}, \{a, b\}, P_0, A)$ where $P_0 = \{A \rightarrow aAb | \epsilon\}$ is a simple CFG for the language $\{a^n b^n | n \in \mathbb{N}\}$. The non-terminal A can derive aAb an indeterminate amount of times before deriving ϵ , allowing $a^n b^n$ for any $n \in \mathbb{N}$.

Two important sets are defined for every non-terminal in a grammar: the *first* and *follow* sets. These are used in “action” table construction. The *first* set is also used to decide which terminals in a given grammar define monitor creation events (we shall be more specific about this below). The *follow* set will be useful in illustrating the fundamental challenge of monitoring CFGs. The *first* set of a non-terminal A , denoted $first(A)$, is the set of all terminals t such that the sub-strings which reduce to A may possibly begin with t . The follow set, denoted $follow(A)$, is the set of terminals which follow the strings which reduce to A . These terminals signify a reduction by A , and $follow(A)$, may contain $\$,$ the end of input token.

A reduction is the step whereby a right hand side of a production, γ , is replaced by the left hand side non-terminal in the production. For example, if we have the string $aaabbb$, and we are using our example grammar, G , we can perform a reduction with $\gamma = \epsilon$ resulting in $aaaAbbb$. We can then perform another reduction with $\gamma = aAb$ resulting in $aaAbb$, eventually we reach aAb , which reduces to A (and because A is the start symbol, $aaabbb$ must be in $L(G)$, where $L(G)$ represents the language derived by G).

⁵Bison also has a GLR mode, which can recognize all CFLs.

```

1  globals action_table, goto_table
2  initialize stack.push(initial_state)
3  procedure monitor(event, stack)
4    locals state, state', stack', A
5    state ← stack.top()
6    while (true) {
7      switch (action_table[state,event],action_type) {
8        case shift :
9          state' ← action_table[state, event].next_state
10         if (state' = error) {
11           violation
12           break while
13         }
14         stack.push(state')
15         if (action_table[state', $].action_type = reduce) {
16           stack' ← stack.clone()
17           monitor($, stack')
18         }
19         break while
20       case reduce :
21         stack.pop(action_table[state, event].pop)
22         A ← action_table[state, event].non_terminal
23         state' ← stack.top()
24         stack.push(goto_table[state', A])
25         break switch
26       case validation :
27         validation
28         break while
29     }
30 }

```

Fig. 3. LR(1) Monitoring Algorithm.

B. CFG-plugin Implementation

The CFG plugin allows the user to specify a number of events and a CFG specifying allowable event traces. The events become the terminals of the CFG, i.e., Σ . The translation steps from specification to working Java code gradually transform the specification into AspectJ join points (the events) and aspects (the synthesized monitors), which are then woven into the original application using any off-the-shelf AspectJ compiler. We have implemented the CFG logic plugin using the Maude term rewriting engine [25] and it is available for download at [12]. The operations of the plugin are described in the following sections.

C. CFG Simplification

The CFG plugin first simplifies the given grammar. The first step of simplification is the removal of non-generating non-terminals (A is non-generating if $\forall s \in \Sigma^*$, s cannot be reduced to A in one or more reductions). The next step in the simplification process is the removal of non-terminals which are unreachable from the start symbol (A is unreachable from the start symbol if there is no string γ that contains A and reduces to the start symbol in any number of steps). The last step removes ϵ -productions from the grammar. After ϵ -productions have been removed from G , resulting in G' , $L(G') = L(G) - \epsilon$.⁶ A monitor matching the empty event trace has little utility, so we feel this is a fair compromise.

⁶The remaining productions are restructured to account for the removal of ϵ -productions without changing the recognized language, other than as mentioned.

D. Tables and Monitoring

After simplification, the tables are generated for a deterministic push-down automaton (DPDA), that is, a deterministic finite automaton with a stack, which is to be used as a monitor. When a new event arrives, the monitoring algorithm must decide how to modify the stack.

The tables are given in a generic intermediate form, which is converted by the Java shell into two Java arrays. The tables are generated by an implementation of the full Knuth algorithm as presented in [21]. This algorithm is known to generate very large tables for complex grammars (such as those for programming language syntax). We believe that monitoring grammars will be relatively small, none-the-less we plan to implement the table compaction optimizations presented in [26] at a future date. Notably, with this optimization, an LALR grammar will create the same sized table as it would with an LALR parser generator, while still generating tables for grammars in LR(1) - LALR(1).

Pseudo-code for our monitoring algorithm is given in Figure 3. The significance of lines 15-18 is explained in the next section. An entry in *action_table* specifies, via the *action_type*, the type of action: *shift*, *reduce*, or *validate*. Each type of action also requires additional information in order for the algorithm to process said action. An entry in *goto_table* simply identifies the next state for the DPDA.

The *shift* action entry contains the next state for the DPDA in the *next_state* field. A *shift* action simply pushes the next state on the stack, if the next state is *not* the error state (lines 10-14). If, however, the table indicates that the next state *is* the error state, the algorithm reports a violation and breaks *without* touching the stack (lines 11-12). This allows the algorithm to continue to find more violations. After a successful *shift* action, the while loop is broken, allowing execution of the monitored program to continue until the next relevant event (line 19).

The *reduce* action is more complicated. The field *non_terminal* describes which non-terminal (A) the production $A \rightarrow \gamma$ reduces to, while the field *pop* denotes how many states to pop from the stack ($|\gamma|$). The reduction proceeds by popping the specified number of states from the stack and consulting *goto_table* to decide the next state (lines 20-25). The state used for indexing *goto_table* is not the current state, but rather the state at the top of the stack after the specified number of states has been popped (line 21). An indeterminate number of reductions can happen in a row, but there *must* be *shift* at the end of the reduction sequence before the algorithm can terminate for a given event. The reductions happen before the *shift* to simulate the look-ahead of one token specified by the 1 in LR(1). If there were no *shift* operation for a given event, it would be effectively *ignored*.

The *validation* action, which directs the DPDA to validate the current input has no special fields, as no more information is necessary (lines 26-28).

E. Monitor Stack Cloning

Note that the LR(1) algorithm assumes that the string of terminals to be evaluated is *completely known ahead of time*.

Thus, it knows where the end of the string (denoted as \$) is. This is important because some reductions happen with the \$ symbol as the look-ahead, and the validation state can *only* be recognized when the next input is \$. When the terminals are events generated dynamically, the end of the “string” (execution trace) cannot be known.

Our implementation of the algorithm attempts to reduce with \$ as the look-ahead after *every* valid *shift* (lines 15-18). The problem with reducing with \$ as the look-ahead where possible is that all state of the current trace evaluation is lost. This means that the monitor can only validate the minimal trace that matches the CFG pattern.

Since we desire to report validations for every sub-trace that matches the pattern⁷, we must *clone* the stack before we perform any reductions with \$ as the look-ahead.

This cloning ensures that the stack is intact for the next, and subsequent events, allowing for multiple validations. For example, consider the language denoted by the regular expression ab^* . While we would suggest using the ERE plugin for such a language, it is a clear example to illustrate the effect of cloning. With no cloning the algorithm would validate for only a . Because it popped during the validation phase, if it sees a b it will report failure. With cloning it will report success for a , and then success again on the input of b , and for any subsequent input of b . An important optimization to cloning is to only clone the stack if there is a reduction with \$ as the look-ahead, rather than blindly for every *shift* operation. This optimization will not help ab^* , but it will help for many other languages. In fact, for $\{a^n b^n | n \in \mathbb{N}\}$, only one clone is necessary no matter how long the input. Any grammar accepting unbounded repetition at the end of the pattern (like ab^*), will require cloning on each input of the repeated character.

F. Correctness of CFG Monitoring Algorithm

We next prove the correctness of our online monitoring algorithm for CFG. It is achieved by showing that our algorithm detects violations and validations of the observed trace in the same way as the Aho, Sethi, Ullman (ASU) parsing algorithm[21], as given in Figure 4.

Theorem: *For every finite prefix of a (possibly infinite) program trace⁸ and a CFG pattern, the MOP algorithm will notify a violation of the pattern if the ASU algorithm would notify a parse failure due to a bad token, and validation if ASU would notify success, given that prefix as total input.*

Proof: First, we construct a new parsing algorithm, as shown in Figure 5. This new algorithm can be proved equivalent to the one in Figure 4 as follows. The major difference between these two algorithms is that the pointer (*ip*) increment is moved to the outer loop in Figure 5. This change does not affect the behavior of the algorithm:

- 1) For a shift action, both algorithms carry out the same operation except that Figure 4 increases the pointer and

⁷This is irrespective of suffix matching which actually generates multiple monitors.

⁸Each prefix is an \mathcal{E} -trace at a given point in a program as per Definition 1.

```

1  globals stack, ip, action_table, goto_table
2  initialize stack.push(initial_state), ip ← 0
3  procedure parse()
4    locals state, state', a, A
5    while (true) {
6      state ← stack.top()
7      a ← get_token_at(ip)
8      switch (at[state, a].action_type) {
9        case shift :
10         state' ← action_table[state, a].next_state
11         if (state' = error) {
12           report_error
13           advance ip
14           continue while
15         }
16         stack.push(state')
17         advance ip
18         continue while
19       case reduce :
20         stack.pop(action_table[state, a].pop)
21         A ← action_table[state, a].non_nonterminal
22         state' ← stack.top()
23         stack.push(goto_table[state', A])
24         continue while
25       case accept :
26         accept
27         return
28     }
29 }

```

Fig. 4. ASU Algorithm.

```

1  globals stack, ip, action_table, goto_table
2  initialize stack.push(initial_state), ip ← 0
3  procedure parse()
4    locals state, state', a, A
5    while (true) {
6      a ← get_token_at(ip)
7      while (true) {
8        state ← stack.top()
9        switch (at[state, a].action_type) {
10         case shift :
11         state' ← action_table[state, a].next_state
12         if (state' = error) {
13           report_error
14           break while
15         }
16         stack.push(state')
17         break while
18       case reduce :
19         stack.pop(action_table[state, a].pop)
20         A ← action_table[state, a].non_nonterminal
21         state' ← stack.top()
22         stack.push(goto_table[state', A])
23         continue while
24       case accept :
25         accept
26         return
27     }
28   }
29   advance ip
30 }

```

Fig. 5. Modified ASU Algorithm.

continues to the next action, while Figure 5 breaks the inner loop, increases the pointer in the outer loop, and then continue to the next action. Both are equivalent.

- 2) For reduction, Figure 5 chooses to stay in the inner loop, which is identical to Figure 4, which continues the loop

without increasing the point.

- 3) For acceptance (validation in monitors), both algorithms are identical.

Now we can prove the correctness of the monitoring algorithm in Figure 3 by comparing it with the modified parsing algorithm in Figure 5.

The major difference distinguishing the monitoring algorithm from the parsing algorithm is that the former has to wait for the next event extracted from the execution of the monitored program while the latter can actively retrieve the next token, which is handled in the outer loop in Figure 5. Therefore, we only need prove that the *monitor* procedure in Figure 3 produces the same result as the inner loop in Figure 5, given the same state and event to process.

It is straightforward to compare both pieces of code: the only difference between them is the *stack cloning* (lines 14-17) in Figure 3. It is needed because we wish to continue parsing *after an accept*, and because we can never actually see \$ as an event. We clone the stack after a shift and check for actions with \$ as the input. The only actions possible on this recursive call are reduce and accept on this recursive call, because \$ can never be shifted⁹. Due to this, the recursion is always bounded at depth one. This is the major difference between the MOP and ASU algorithms. Because \$ can never be an event, we must speculatively guess the end of input after every symbol seen. The recursive call must happen iff there is a valid reduction with \$ as the look-ahead. Because the algorithm repeats until a shift action, error, or accept happens, we ensure that the recursive call must happen, if it happens, after the processing of each input¹⁰. Cloning the stack allows us to reduce and accept, while still maintaining the original stack to continue monitoring events as if the end of input had *not* been seen. Thus, this change is equivalent to the ASU algorithm in terms of language recognition because both possibilities (the arrival or non-arrival of \$) are explored. That is, the MOP algorithm will report accept for a given prefix if ASU would, given that prefix as its total input, and it, additionally, retains enough state to continue parsing future (longer) traces. Violation is handled identically in both algorithms. \square

G. Suffix Matching with CFGs

As described in Section IV, several features are needed for monitors to support optimized suffix matching.

The first is identification of monitor creation events¹¹. As already mentioned, monitor creation events are events which, when encountered as the *first* event in a trace would not lead to an immediate failure. For CFGs this would imply an event that can begin a sub-string which reduces to the start symbol. This is the same as the definition of *first* set as given earlier. Thus, the monitor creation events for the CFG plugin are those

⁹This is a property of the CFG parsing table generation algorithm, which we use without proof.

¹⁰Accept need not be considered because it can only happen when the input is \$, which only occurs during a recursive call.

¹¹Though, as mentioned, this is also necessary for total matching.

events which are in $first(S)$, where S is the start symbol for the grammar.

Additionally, it is necessary to define a hash encoding for CFG based monitors because our suffix matching algorithm uses a HashSet to find monitors with *potentially* equivalent states quickly. We decided that two simple defining aspects of CFG based monitors are stack depth and the current state of the monitor (the top of the stack). We chose to xor them together (a broadly used operation for combining two binary quantities into one quantity representative of the two in the same number of bits), because the hashCode method must return an integer. Lastly, we need an equality method (to resolve collisions) defining when two CFG based monitors have *actually* equivalent states. Two CFG monitors can only be equal iff they have the same stack contents. It will be fairly rare for two proper CFG monitors to be equivalent, as they do not have finite state like the other logic plugins of MOP. Thus it is important for failed equality testing to be fast. Because of this, we check to see if two monitors have the same stack depth before beginning element-wise comparisons.

VI. EVALUATION

We evaluated the CFG plugin on the DaCapo benchmark suite [27]. Also, we evaluated PQL and Tracematches against DaCapo, using the same properties, for comparison.

A. Experimental Settings

Our experiments were carried out on a machine with 1.5GB RAM and Pentium 4 2.66GHz processor. The operating system used was Ubuntu Linux 7.10. We used the DaCapo benchmark version 2006-10; it contains eleven open source programs [27]: antlr, bloat, chart, eclipse, fop, hsqldb, jython, luindex, lusearch, pmd, and xalan. The provided default input was used together with the -converge option to execute the benchmark multiple times until the execution time falls within a coefficient of variation of 3%. The average execution time of six iterations after convergence are then used to compute the runtime overhead. Therefore, Table II percentages should be read “ ± 3 ” (meaning negative numbers are possible).

B. Properties

The following general properties borrowed from [13] were checked in the evaluation:

- HashMap: An object’s hash code should not be changed when the object is a key in a HashMap;
- hasNext: For a given iterator, the hasNext() method should be called between all calls to next();
- SafeIterator: Do not update a Collection when using the Iterator interface to iterate its elements.

We also defined three new properties to showcase the power of the CFG plugin; they are all properly context-free:

- ImprovedLeakingSync: The original LeakingSync specified in [13] *only* allows synchronized accesses to synchronized collections. This causes spurious failures because the synchronized methods call the unsynchronized

	HashMap			HasNext			SafeIterator			ImprovedLeakingSync			SafeFileInputStream			SafeFileWriter		
	MOP	PQL	TM	MOP	PQL	TM	MOP	PQL	TM	MOP	PQL	TM	MOP	PQL	TM	MOP	PQL	TM
antlr	3	6	0	1	2	3	2	82	0	1	N/E	N/E	3	113	-1	2	22	N/E
bloat	14	9	-2	1112	5929	2452	627	8694	11258	13	N/E	N/E	1	128	0	27	97	N/E
chart	-1	1	-1	-1	3	0	2	50	11	4	N/E	N/E	0	29	1	0	37	N/E
eclipse	0	1	1	0	2	-1	-2	1	2	1	N/E	N/E	-2.5	3	0	-2	1	N/E
fop	3	2	0	0	2	-1	-1	24	5	1	N/E	N/E	-2	58	-1	-2	47	N/E
hsqldb	0	3	15	0	6	15	0	78	17	1	N/E	N/E	1	280	21	2	95	N/E
ijython	0	23	15	0	0	13	0	12	16	41	N/E	N/E	0	937	12	1	crashes	N/E
luindex	1	8	1	-2	93	2	3	181	9	1	N/E	N/E	-1	233	6	0	33	N/E
lusearch	1	1	8	-1	59	9	4	132	34	2	N/E	N/E	-1	137	7	0	49	N/E
pmd	-1	0	3	191	1870	52	178	1334	175	36	N/E	N/E	-1	547	1	-2	658	N/E
xalan	0	5	1	0	0	2	1	53	10	3	N/E	N/E	-1	90	3	-2.5	164	N/E

TABLE II

AVERAGE PERCENT RUNTIME OVERHEAD FOR JAVAMOP CFG (MOP), PQL, AND TRACEMATCHES (TM) (CONVERGENCE WITHIN 3%).

versions. Our improved version allows calls to the unsynchronized methods so long as they happen within synchronized calls.

- **SafeFileInputStream:** SafeFileInputStream is a modification of our SafeLock property from Figure 2. It ensures that a FileInputStream is closed in the same method in which it is created.
- **SafeFileWriter:** SafeFileWriter ensures that all writes to a FileWriter happen between creation and close of the FileWriter, and that the creation and close events are matched pairs.

More properties have been checked in our experiments; we choose the first three regular-language-based properties (HashMap, HasNext, and SafeIterator) to include in this paper because they generate a comparatively larger runtime overhead. We excluded those with little overhead in JavaMOP. For every property, we provide overhead percentages for JavaMOP, as well as PQL and Tracematches where possible. We run the JavaMOP monitors in suffix matching mode; the decentralized indexing of monitors was used in all the experiments (see [10]). We chose the AspectJ compiler 1.5.3 (AJC) in the evaluation to compile the JavaMOP generated monitoring AspectJ code. For Tracematches we used the most recent published version from [28]. All of these properties are available on the JavaMOP Online Interface [12]. Each one is named `Suffix<PropertyName>_CFG`.

C. Results

Table II shows the percent overheads of JavaMOP using the CFG plugin, PQL, and Tracematches. N/E refers to specifications that were not expressible. Negative numbers can be attributed to the 3% noise in the measurements. Tracematches is unable to support ImprovedLeakingSync because the property is truly context-free. PQL is also unable to support it because it requires events corresponding to the beginning and end of synchronized method calls, and PQL can only trigger events on the end of method calls. Tracematches cannot support SafeFileWriter because it is a pure context-free specification. However, Tracematches *can* support SafeFileInputStream because it has the ability to access call stack depth via the `cflowdepth` pointcut term, which is provided only by the ABC compiler for AspectJ.

Over one running of the entire DaCapo benchmark suite, over 355 million events are generated, creating more than 87 million monitors. The average overhead of JavaMOP over the 66 program/property pairs is 34%. There are two considerations here, however: (1) we chose specifically those properties that generated the largest overheads, (2) when the two largest overheads are removed, the average over the remaining 64 pairs drops to a very reasonable 8%. Further, the average JavaMOP overhead for properties expressible in PQL was 39% over 55 pairs, while PQL's overhead on these same properties was 415%. Similarly, for Tracematches expressible properties, JavaMOP's overhead was 48% over 44 pairs, while Tracematches was 322%. Tracematches, PQL, and JavaMOP all feature the same two pairs which have extremely large overhead compared to the median (HasNext and SafeIterator in bloat). When these two pairs are removed from the three averages, the average overhead for JavaMOP with respect to PQL expressible properties is 8%, while PQL still weighs in at 150%. Tracematches is comparable to JavaMOP, with JavaMOP at 9% and Tracematches at 11%. Since Tracematches does not support the full generality of (deterministic) context-free grammars, we view comparable performance to Tracematches as favorable to our approach, especially given that, in the overall data set, our average overhead is over 6 times lower than Tracematches' overhead.

The largest overheads seen, across all three systems, are for SafeIterator and HasNext in bloat. This is due to bloat's extensive use of iterators. bloat is a bytecode optimizer, which uses iterators to process bytecode. PQL and Tracematches perform worse on SafeIterator than they do on HasNext, while our performance is the opposite. The reason for this is that HasNext creates a far larger number of monitors in JavaMOP, because it creates a monitor for every call to next, while SafeIterator only creates monitors on a call to create. The pattern for SafeIterator, however, is more complex. This shows that JavaMOP has, relatively speaking, more overhead in generating and handling the monitor set for suffix matching than it does matching the pattern, while PQL and Tracematches overheads are more affected by the complexity of the pattern. Note that JavaMOP with CFGs far outperforms both PQL and Tracematches on these 2 program/property pairs.

SafeFileInputStream is an interesting case: as a context-free

grammar it is required to match the begin and end of methods. Instrumenting the begin and end of *every* method would be atrociously slow, however. We perform a static analysis which finds those methods in which `FileInputStream`'s are actually used. Then we instrument only those methods for begin and end. Because `Tracematches`, also, is pointcut based, we are able to perform the same optimization for `Tracematches`, so the numbers shown are with the optimization enabled. `PQL` is not pointcut based so the optimization cannot be applied; however, the `PQL` property does not match begins and ends of methods (recall: `PQL` can only match the ends), so this is not an issue. In `PQL`, we specify `SafeFileInputStream` by using an interesting `PQL`-specific feature called `within`. The idea of `within` is that a property matches only *within* a given method or methods matching a particular pattern (in the case of `SafeFileInputStream` we use the pattern `_.*` which specifies all methods of all classes). Additionally, `PQL` will only instrument the same methods that `JavaMOP` and `Tracematches` instrument, because `within` only instruments methods which can generate relevant events.

The memory overhead is reasonable in our experiments: overall, it is 33% on average with a 4% median. Although, when compared with `Tracematches` we still have room for improvement [6], the memory overhead of our approach does not cause unnecessary loss of performance during the evaluation, indicating that it is not a bottleneck for the efficiency of monitoring. More detailed results, including complete numbers for memory usage, generated events, number of monitors, and other interesting statistics, are available at the `JavaMOP` CFG Plugin webpage [12], and in the technical report [29].

VII. CONCLUSIONS AND FUTURE WORK

We implemented a CFG pattern based logic plugin for `JavaMOP` using a modified LR(1) parsing algorithm. Our modification to the algorithm is based on the novel idea of *cloning* the stack in order to “predict” a possible reduction with \$ (end of string) as a look ahead without destroying the state of the monitor. We showed, empirically, that our algorithm is efficient and faster than the state-of-the-art for monitoring CFG properties. We also extended `JavaMOP` with suffix matching, in order to fairly compare `JavaMOP` with `PQL` and `Tracematches`. `Tracematches`, however, cannot handle arbitrary CFG patterns.

We have also begun work on a logic called `PtCaRet` as another specification formalism to support structured specifications. `PtCaRet` is past time linear temporal logic (PTLTL) with calls and returns. It is a super set of PTLTL that provides operators which apply only on a function/method local level.

Further, we intend to add the `Pager` [26] optimizations to the CFG plugin table creation algorithm, in order to reduce the size of the created action and goto tables.

REFERENCES

- [1] K. Havelund and G. Roşu, “Monitoring Java programs with Java PathExplorer,” in *Runtime Verification(RV'01)*, ser. ENTCS, vol. 55, no. 2, 2001.
- [2] “Temporal Rover,” <http://www.time-rover.com>.
- [3] M. Kim, M. Viswanathan, S. Kannan, I. Lee, and O. Sokolsky, “Java-MaC: A run-time assurance approach for Java programs,” *Formal Methods in System Design*, vol. 24, no. 2, pp. 129–155, 2004.
- [4] M. d’Amorim and K. Havelund, “Event-based runtime verification of Java programs,” *ACM SIGSOFT Software Engineering Notes*, vol. 30, no. 4, pp. 1–7, 2005.
- [5] C. Allan, P. Avgustinov, A. S. Christensen, L. J. Hendren, S. Kuzins, O. Lhoták, O. de Moor, D. Sereni, G. Sittampalam, and J. Tibble, “Adding trace matching with free variables to AspectJ,” in *OOPSLA'05*. ACM, 2005, pp. 345–364.
- [6] P. Avgustinov, J. Tibble, and O. de Moor, “Making trace monitors feasible,” in *OOPSLA'07*. ACM, 2007, pp. 589–608.
- [7] E. Bodden, “J-LO, a tool for runtime-checking temporal assertions,” Master’s thesis, RWTH Aachen University, 2005.
- [8] M. Martin, V. B. Livshits, and M. S. Lam, “Finding application errors and security flaws using PQL: a program query language,” in *OOPSLA'07*. ACM, 2005, pp. 365–383.
- [9] S. Goldsmith, R. O’Callahan, and A. Aiken, “Relational queries over program traces,” in *OOPSLA'05*. ACM, 2005, pp. 385–402.
- [10] F. Chen and G. Roşu, “MOP: An efficient and generic runtime verification framework,” in *OOPSLA'07*. ACM, 2007, pp. 569–588.
- [11] F. B. Schneider, “Enforceable security policies,” *ACM Transactions on Information System Security*, vol. 3, no. 1, pp. 30–50, 2000.
- [12] “CFG Plugin website,” <http://fsl.cs.uiuc.edu/CFGPlugin>.
- [13] E. Bodden, L. Hendren, and O. Lhoták, “A staged static program analysis to improve the performance of runtime monitoring,” in *ECOOP'07*, ser. LNCS, vol. 4609, 2007, pp. 525–549.
- [14] F. Chen and G. Roşu, “Towards monitoring-oriented programming: A paradigm combining specification and implementation,” in *Runtime Verification(RV'03)*, ser. ENTCS, vol. 89, no. 2, 2003.
- [15] F. Chen, M. D’Amorim, and G. Roşu, “A formal monitoring-based framework for software development and analysis,” in *ICFEM'04*, ser. LNCS, vol. 3308, 2004, pp. 357–372.
- [16] S. Chaudhuri and R. Alur, “Instrumenting C programs with nested word monitors,” in *Model Checking Software(SPIN'07)*, ser. LNCS, vol. 4595, 2007, pp. 279–283.
- [17] P. Avgustinov, A. S. Christensen, L. Hendren, S. Kuzins, J. Lhotak, O. Lhotak, O. de Moor, D. Sereni, G. Sittampalam, and J. Tibble, “ABC: an extensible AspectJ compiler,” in *AOSD'05*. ACM, 2005, pp. 87–98.
- [18] F. Chen, M. D’Amorim, and G. Roşu, “Checking and correcting behaviors of Java programs at runtime with JavaMOP,” in *Runtime Verification(RV'06)*, ser. ENTCS, vol. 144, no. 4, 2006, pp. 3–20.
- [19] G. Kiczales, E. Hilsdale, J. Hugunin, M. Kersten, J. Palm, and W. G. Griswold, “An overview of AspectJ,” in *ECOOP'01*, ser. LNCS, vol. 2072, 2001, pp. 327–353.
- [20] G. T. Leavens, K. R. M. Leino, E. Poll, C. Ruby, and B. Jacobs, “JML: notations and tools supporting detailed design in Java,” in *OOPSLA'00*. ACM, 2000, pp. 105–106.
- [21] A. V. Aho, R. Sethi, and J. D. Ullman, *Compilers, principles, techniques, and tools*. Addison-Wesley, 1986, pages 215–246.
- [22] J. E. Hopcroft, R. Motwani, and J. D. Ullman, *Introduction to automata theory, languages, and computation, 2nd edition*. Addison-Wesley, 2001.
- [23] “Bison,” <http://www.gnu.org/software/bison/>.
- [24] D. E. Knuth, “On the translation of languages from left to right,” *Information and Control*, vol. 8, no. 6, pp. 607–639, 1965.
- [25] M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer, and C. Talcott, *All About Maude-A High-Performance Logical Framework: How to Specify, Program, and Verify Systems in Rewriting Logic*. Springer-Verlag New York, Inc., 2007.
- [26] D. Pager, “A practical general method for constructing LR(k) parsers,” *Acta Informatica*, vol. 7, pp. 249–268, 1977.
- [27] S. M. Blackburn, R. Garner, C. Hoffman, A. M. Khan, K. S. McKinley, R. Bentzur, A. Diwan, D. Feinberg, D. Frampton, S. Z. Guyer, M. Hirzel, A. Hosking, M. Jump, H. Lee, J. E. B. Moss, A. Phansalkar, D. Stefanović, T. VanDrunen, D. von Dincklage, and B. Wiedermann, “The DaCapo benchmarks: Java benchmarking development and analysis,” in *OOPSLA'06*. ACM, 2006, pp. 169–190.
- [28] “Tracematches Benchmarks,” <http://abc.comlab.ox.ac.uk/tmahead>.
- [29] P. Meredith, F. Chen, D. Jin, and G. Roşu, “Efficient monitoring of parametric context-free patterns,” Department of Computer Science, University of Illinois at Urbana-Champaign, Tech. Rep. UIUCDCS-R-2008-2954, April 2008.