

The Rewriting Logic Semantics Project: A Progress Report

José Meseguer and Grigore Roşu
Department of Computer Science,
University of Illinois at Urbana-Champaign.
{meseguer,grosu}@illinois.edu

Abstract. Rewriting logic is an executable logical framework well suited for the semantic definition of languages. Any such framework has to be judged by its effectiveness to bridge the existing gap between language definitions on the one hand, and language implementations and language analysis tools on the other. We give a progress report on how researchers in the rewriting logic semantics project are narrowing the gap between theory and practice in areas such as: modular semantic definitions of languages; scalability to real languages; support for real time; semantics of software and hardware modeling languages; and semantics-based analysis tools such as static analyzers, model checkers, and program provers.

1 Introduction

The disconnect between theory and practice is one of the worse evils in computer science. Theory disconnected from practice becomes irrelevant; and practice without theory becomes brute-force, costly and ad-hoc engineering. One of the current challenges in formal approaches to language semantics is precisely how to effectively bridge the gap between theory and practice. There are two distinct dimensions to this gap:

- (1) Given a language \mathcal{L} , there is often a substantial gap between: (i) a formal semantics for \mathcal{L} ; (ii) an implementation of \mathcal{L} ; and (iii) analysis tools for \mathcal{L} , including static, dynamic, and deductive tools.
- (2) Even if a formal semantics exists for a programming language \mathcal{L} , there may not be any formal semantics available at the higher level of software designs and models, or at the lower level of hardware.

Regarding (1), a semantics of \mathcal{L} may just be a “paper semantics,” such as some SOS rules on a piece of paper; or it may be a “toy semantics,” not for \mathcal{L} itself, but for a greatly simplified sublanguage. Furthermore, the way a compiler for \mathcal{L} is written may have no connection whatever with a formal semantics for \mathcal{L} , so that different compilers provide different language behaviors. To make things worse, program analysis tools for \mathcal{L} , including tools that supposedly provide some formal analysis, may not be systematically based on a formal semantics either, so that the confidence one can place of the answers from such tools is greatly

diminished. Regarding (2), one big problem is that software modeling notations often lack a formal semantics. A related problem is that this lack of semantics manifests itself as a lack of *analytic power*, that is, as an incapacity to uncover expensive design errors which could have been caught by formal analysis.

We, together with many other colleagues all over the world, have been working for years on the *rewriting logic semantics project* (see [77, 76, 112] for some overview papers at different stages of the project). The goal of this project is to substantially narrow the gap between theory and practice in language specifications, implementations and tools, in both of the above dimensions (1)–(2). In this sense, rewriting logic semantics is a *wide-spectrum framework*, where:

1. The formal semantics of a language \mathcal{L} is used as the *basis* on which both language implementations and language analysis tools are built.
2. The same semantics-based approach is used not just for programming languages, but also for software and hardware modeling languages.

Any attempt to bridge theory and practice cannot be judged by theoretical considerations alone. One has to evaluate the practical effectiveness of the approach in answering questions such as the following:

- *Executability*. Is the semantics executable? How efficiently so? Can semantic definitions be tested to validate their agreement with an informal semantics?
- *Range of Applicability*. Can it be applied to programming languages and to software and hardware modeling languages? Can it naturally support nontrivial features such as concurrency and real time?
- *Scalability*. Can it be used in practice to give full definitions of real languages like Java or C? And of real software and hardware modeling languages?
- *Integrability*. How well can the semantics be integrated with language implementations and language analysis tools? Can it really be used as the *basis* on which such implementations and analysis tools are built?

This paper is a progress report on the efforts by various researchers in the rewriting logic semantics project to positively answer these questions. After summarizing some related work below, we give an overview of rewriting logic semantics in Section 2. Subsequent sections then describe in more detail: (i) modularity of definitions and the support for highly modular definitions provided by the \mathbb{K} framework (Section 3); (ii) semantics of programming languages (Section 4); semantics of real-time language (Section 5); (iv) semantics of software modeling languages (Section 6); (v) semantics of hardware description languages (Section 7); (vi) abstract semantics and static analysis (Section 8); (vii) model checking verification (Section 9); and (viii) deductive verification (Section 10). We finish with some concluding remarks in Section 11.

1.1 Related Work

There is much related work on frameworks for defining programming languages. Without trying to be exhaustive, we mention some of them and point out some relationships to rewriting logic semantics (RLS).

Structural Operational Semantics (SOS). Several variants of structural operational semantics have been proposed. We refer to [112] for an in-depth comparison between SOS and RLS. A key point made in [112], and also made in Section 2.5, is that RLS is a framework supporting many different definitional styles. In particular, it can naturally and faithfully express many different SOS styles such as: small-step SOS [99], big-step SOS [56], MSOS [87], reduction semantics [129], continuation-based semantics [43], and the CHAM [12].

Algebraic denotational semantics. This approach, (see [125, 49, 26, 85] for early papers and [47, 118] for two more recent books), is the special case of RLS where the rewrite theory $\mathcal{R}_{\mathcal{L}}$ defining a language \mathcal{L} is an equational theory. Its main limitation is that it is well suited for giving semantics to *deterministic* languages, but not well suited for concurrent language definitions.

Higher-order approaches. The most classic higher-order approach is *denotational semantics* [109, 110, 108, 86]. Denotational semantics has some similarities with its first-order algebraic cousin mentioned above, since both are based on semantic equations and both are best suited for deterministic languages. Higher-order functional languages or higher-order theorem provers can be used to give an executable semantics to programming languages, including the use of Scheme in [45], the use of ML in [98], and the use of Common LISP within the ACL2 prover in [61]. There is also a body of work on using monads [81, 124, 65] to implement language interpreters in higher-order functional languages; the monadic approach has better modularity characteristics than standard SOS. Some higher-order approaches are based on the use of higher-order abstract syntax (HOAS) [97, 52] and higher-order logical frameworks, such as LF [52] or λ -Prolog [88], to encode programming languages as formal logical systems; for a good example of recent work in this direction see [78] and references there.

Logic-programming-based approaches. Going back to the Centaur project [22, 35], logic programming has been used as a framework for SOS language definitions. Note that λ -Prolog [88] belongs both in this category and in the higher-order one. For a recent textbook giving logic-programming-based language definitions, see [113].

Abstract state machines. Abstract State Machine (ASM) [50] can encode any computation and have a rigorous semantics, so any programming language can be defined as an ASM and thus implicitly be given a semantics. Both big- and small-step ASM semantics have been investigated. The semantics of various programming languages, including Java [114], has been given using ASMs.

Other RLS work. RLS is a collective international project. There is by now a substantial body of work demonstrating the usefulness of this approach, e.g., [23, 120, 117, 115, 72, 119, 31, 104, 122, 42, 40, 55, 25, 73, 77, 30, 28, 41, 34, 106, 1, 116, 36, 107, 58, 54, 46, 39, 5], and we describe some even more recent advances in this paper. A first snapshot of the RLS project was given in [77], a second in [76], and a third in [112], with this paper as the fourth snapshot.

2 Rewriting Logic Semantics in a Nutshell

Before describing in more detail the different advances in the rewriting logic semantics project we give here an overview of it. We begin with a short summary of rewriting logic as a semantic framework for concurrent systems. Then we explain how it can be used to give both an operational and a denotational semantics to a programming language. Thanks to the distinction between equations and rules, this semantics can be given at various levels of abstraction. Furthermore, a wide range of definitional styles can be naturally supported. We explain how rewriting logic semantics has been extended to: (i) real-time languages; (ii) software modeling languages; and (iii) hardware description languages. We finally explain how a rewriting logic semantics can be used for static analysis, and for model checking and deductive verification of programs.

2.1 Rewriting Logic

The goal of rewriting logic [69] is to provide a flexible logical framework to specify concurrent systems. A concurrent system is specified as a *rewrite theory* $\mathcal{R} = (\Sigma, E, R)$, where (Σ, E) is an equational theory, and R is a set of (possibly conditional) rewrite rules. The equational theory (Σ, E) specifies the concurrent system's set of states as an algebraic data type, namely, as the initial algebra of the equational theory (Σ, E) . Concretely, this means that a distributed state is mathematically represented as an E -equivalence class $[t]_E$ of terms built up with the operators declared in Σ , modulo provable equality using the equations E , so that two state representations t and t' describe the *same* state if and only if one can prove the equality $t = t'$ using the equations E .

The rules R specify the system's *local concurrent transitions*. Each rewrite rule in R has the form $t \rightarrow t'$, where t and t' are Σ -terms. The lefthand side t describes a *local firing pattern*, and the righthand side t' describes a corresponding *replacement pattern*. That is, any fragment of a distributed state which is an instance of the firing pattern t can perform a local concurrent transition in which it is replaced by the corresponding instance of the replacement pattern t' . Both t and t' are typically *parametric* patterns, describing not single states, but parametric families of states. The parameters appearing in t and t' are precisely the *mathematical variables* that t and t' have, which can be instantiated to different concrete expressions by a *substitution*, that is, a mapping θ sending each variable x to a term $\theta(x)$. The instance of t by θ is then denoted $\theta(t)$.

The most basic *logical deduction steps* in a rewrite theory $\mathcal{R} = (\Sigma, E, R)$ are precisely atomic concurrent transitions, corresponding to applying a rewrite rule $t \rightarrow t'$ in R to a state fragment which is an instance of the firing pattern t by some substitution θ . That is, up to E -equivalence, the state is of the form $C[\theta(t)]$, where C is the rest of the state not affected by this atomic transition. Then, the resulting state is precisely $C[\theta(t')]$, so that the atomic transition has the form $C[\theta(t)] \rightarrow C[\theta(t')]$. Rewriting is *intrinsically concurrent*, because many other atomic rewrites can potentially take place in the rest of the state C (and in the substitution θ), at the same time that the local atomic transition $\theta(t) \rightarrow \theta(t')$

happens. The rules of deduction of rewriting logic [69, 27] (which in general allow rules in R to be *conditional*) precisely describe all the possible, complex concurrent transitions that a system can perform, so that concurrent computation and logical deduction *coincide*.

2.2 Defining Programming Languages

The flexibility of rewriting logic to naturally express many different models of concurrency can be exploited to give *formal definitions of concurrent programming languages* by specifying the concurrent model of a language \mathcal{L} as a rewrite theory $(\Sigma_{\mathcal{L}}, E_{\mathcal{L}}, R_{\mathcal{L}})$, where: (i) the signature $\Sigma_{\mathcal{L}}$ specifies both the syntax of \mathcal{L} and the types and operators needed to specify semantic entities such as the store, the environment, input-output, and so on; (ii) the equations $E_{\mathcal{L}}$ can be used to give semantic definitions for the *deterministic* features of \mathcal{L} (a sequential language typically has only deterministic features and can be specified just equationally as $(\Sigma_{\mathcal{L}}, E_{\mathcal{L}})$); and (iii) the rewrite rules $R_{\mathcal{L}}$ are used to give semantic definitions for the concurrent features of \mathcal{L} such as, for example, the semantics of threads.

By specifying the rewrite theory $(\Sigma_{\mathcal{L}}, E_{\mathcal{L}}, R_{\mathcal{L}})$ in a rewriting logic language like Maude¹ [32], it becomes not just a mathematical definition but an *executable* one, that is, an *interpreter* for \mathcal{L} . Furthermore, one can leverage Maude’s generic search and LTL model checking features to automatically endow \mathcal{L} with powerful *program analysis capabilities*. For example, Maude’s search command can be used in the module $(\Sigma_{\mathcal{L}}, E_{\mathcal{L}}, R_{\mathcal{L}})$ to detect any violations of invariants, e.g., a deadlock or some other undesired state, of a program in \mathcal{L} . Likewise, for terminating concurrent programs in \mathcal{L} one can model check any desired LTL property. All this can be effectively done not just for toy languages, but for real ones such as Java and the JVM, Scheme, and C (see Section 4 for a discussion of such “real language” applications), and with performance that compares favorably with state-of-the-art model checking tools for real languages.

2.3 Operational vs. Denotational Semantics

A rewrite theory $\mathcal{R} = (\Sigma, E, R)$ has both a *deduction-based operational semantics*, and an *initial model denotational semantics*. Both semantics are defined naturally out of the proof theory of rewriting logic [69, 27]. The deduction-based operational semantics of \mathcal{R} is defined as the collection of *proof terms* [69] of the form $\alpha : t \longrightarrow t'$. A proof term α is an algebraic description of a proof tree proving $\mathcal{R} \vdash t \longrightarrow t'$ by means of the inference rules of rewriting logic. What such proof trees describe are the different *finitary concurrent computations* of the concurrent system axiomatized by \mathcal{R} .

¹ Other rewriting logic languages, such as ELAN or CafeOBJ, can likewise be used. Maude has the advantage of efficiently supporting not only execution, but also LTL model checking verification.

A rewrite theory $\mathcal{R} = (\Sigma, E, R)$ has also a *model-theoretic semantics*, so that the inference rules of rewriting logic are sound and complete with respect to satisfaction in the class of models of \mathcal{R} [69, 27]. Such models are *categories* with a (Σ, E) -algebra structure [69]. These are “true concurrency” denotational models of the concurrent system axiomatized by \mathcal{R} . That is, this model theory gives a precise mathematical answer to the question: when do two descriptions of two concurrent computations denote *the same* concurrent computation? The class of models of a rewrite theory $\mathcal{R} = (\Sigma, E, R)$ has an *initial model* $\mathcal{T}_{\mathcal{R}}$ [69]. The initial model semantics is obtained as a *quotient* of the just-mentioned deduction-based operational semantics, precisely by axiomatizing algebraically when two proof terms $\alpha : t \rightarrow t'$ and $\beta : u \rightarrow u'$ denote the same concurrent computation.

In particular, if a rewrite theory $\mathcal{R}_{\mathcal{L}} = (\Sigma_{\mathcal{L}}, E_{\mathcal{L}}, R_{\mathcal{L}})$ specifies the semantics of a concurrent programming language \mathcal{L} , its denotational semantics is given by the initial model $\mathcal{T}_{\mathcal{R}_{\mathcal{L}}}$, and its operational semantics is given by the proof terms built by the rewriting deduction. As we explain below, many different styles of operational semantics, including various SOS styles, can be naturally obtained as special instances of this general, logic-based operational semantics.

2.4 The Abstraction Dial

Unlike formalisms like SOS, where there is only one type of semantic rule, rewriting logic semantics provides a key distinction between *deterministic rules*, axiomatized by equations, and concurrent and typically *non-deterministic* rules, axiomatized by non-equational rules. More precisely, for the rewriting logic semantics $\mathcal{R}_{\mathcal{L}}$ of a language \mathcal{L} to have good executability properties, we require $\mathcal{R}_{\mathcal{L}}$ to be of the form $\mathcal{R}_{\mathcal{L}} = (\Sigma_{\mathcal{L}}, E_{\mathcal{L}} \cup B_{\mathcal{L}}, R_{\mathcal{L}})$, where: (i) $B_{\mathcal{L}}$ is a collection of *structural axioms*, such as associativity and/or commutativity, and/or identity of certain operators in $\Sigma_{\mathcal{L}}$; (ii) the equations $E_{\mathcal{L}}$ are *confluent modulo* the structural axioms $B_{\mathcal{L}}$; and (iii) the rules $R_{\mathcal{L}}$ are *coherent* with the equations $E_{\mathcal{L}}$ modulo the structural axioms $B_{\mathcal{L}}$ [123]. Conditions (i)–(iii) make $\mathcal{R}_{\mathcal{L}}$ *executable*, so that using a rewriting logic language like Maude we automatically get an interpreter for \mathcal{L} .

As already mentioned, what the equations $E_{\mathcal{L}}$ axiomatize are the *deterministic features* of \mathcal{L} . Instead, the truly concurrent features of \mathcal{L} are axiomatized by the non-equational rules $R_{\mathcal{L}}$. The assumption of determinism is precisely captured by $E_{\mathcal{L}}$ being a set of *confluent equations* (modulo $B_{\mathcal{L}}$), so that their evaluation, if terminating, has a *unique* final result.

All this means that rewriting logic comes with a built-in “abstraction dial.” The least abstract possible position for such a dial is to turn the equations $E_{\mathcal{L}}$ into rules, yielding the theory $(\Sigma_{\mathcal{L}}, B_{\mathcal{L}}, E_{\mathcal{L}} \cup R_{\mathcal{L}})$; this is typically the approach taken by SOS definitions. The specification $\mathcal{R}_{\mathcal{L}} = (\Sigma_{\mathcal{L}}, E_{\mathcal{L}} \cup B_{\mathcal{L}}, R_{\mathcal{L}})$ can already achieve an enormous abstraction, which typically makes the difference between tractable and intractable model checking analysis. The point is that the equations $E_{\mathcal{L}}$ now identify all intermediate execution states obtained by deterministic steps, yielding a typically enormous state space reduction. Sometimes

we may be able to turn the dial to an *even more abstract position* by further decomposing $R_{\mathcal{L}}$ as a disjoint union $R_{\mathcal{L}} = R'_{\mathcal{L}} \cup G_{\mathcal{L}}$, so that the rewrite theory $(\Sigma_{\mathcal{L}}, E_{\mathcal{L}} \cup G_{\mathcal{L}} \cup B_{\mathcal{L}}, R'_{\mathcal{L}})$ still satisfies conditions (i)–(iii). That is, we may be able to identify rules $G_{\mathcal{L}}$ describing concurrent executions which, by being confluent, can be turned into equations. For example, for $\mathcal{L} = Java$, the JavaFAN rewriting logic semantics of Java developed by the late Feng Chen turns the abstraction dial as far as possible, obtaining a set E_{Java} with hundreds of equations, and a set R_{Java} with just 5 rules. This enormous state space reduction is a key reason why the JavaFAN model checker compares favorably with other state-of-the-art Java model checkers [40].

But the abstraction story does not end here. After all, the semantics $(\Sigma_{\mathcal{L}}, E_{\mathcal{L}} \cup G_{\mathcal{L}} \cup B_{\mathcal{L}}, R'_{\mathcal{L}})$ obtained by turning the abstraction dial as much as possible is still a *concrete* semantics. We might call it “the most abstract concrete semantics possible.” For many different static analysis purposes one wants to take a further abstraction step, which further collapses the set of states by defining a suitable *abstract semantics* for a language \mathcal{L} . The point is that, instead of a “concrete semantics” describing the actual execution of programs in \mathcal{L} , one can just as easily define an “abstract semantics” $(\Sigma_{\mathcal{L}}^A, E_{\mathcal{L}}^A, R_{\mathcal{L}}^A)$ describing any desired abstraction A of \mathcal{L} . A good example is type checking, where the values manipulated by the abstract semantics are the types. All this means that many different forms of program analysis, much more scalable than model checking based on a language’s concrete semantics, become available essentially for free by using a tool like Maude to execute and analyze one’s desired abstract semantics $(\Sigma_{\mathcal{L}}^A, E_{\mathcal{L}}^A, R_{\mathcal{L}}^A)$. This is further discussed in Section 8.

2.5 An Ecumenical Movement

For purposes of formally defining the semantics of a programming language, rewriting logic should be viewed not as a competitor to other approaches, but as an “ecumenical movement” providing a framework where many different definitional styles can happily coexist. From its early stages rewriting logic has been recognized as ideally suited for SOS definitions [74, 66], and has been used to give SOS definitions of programming languages in quite different styles, e.g., [119, 25, 121, 122, 40, 42]. What the paper [112] makes explicit is both the wide range of SOS styles supported, and the possibility of defining new styles that may have specific advantages over traditional ones. Indeed, the intrinsic flexibility of rewriting logic means that it does not prescribe a fixed style for giving semantic definitions. Instead, *many different styles* such as, for example, small-step or big-step semantics, reduction semantics, CHAM-style semantics, modular structural operational semantics, or continuation semantics, can all be naturally supported [112]. But not all styles are equally efficient; for example, small-step semantics makes heavy use of conditional rewrite rules, insists on modeling every single computation step as a rule in $R_{\mathcal{L}}$, and is in practice horribly inefficient. Instead, the continuation semantics style described in [112] and used in, e.g., [40] is very efficient. Furthermore, as already mentioned, the distinction between equations and rules provides an “abstraction dial” not available in some definitional styles

but enormously useful for state space reduction purposes. Of particular interest are *modular* definitional styles, which are further discussed in Section 3.

2.6 Defining Real-Time Languages

In rewriting logic, real-time systems are specified with *real-time rewrite theories* [93]. These are just ordinary rewrite theories $\mathcal{R} = (\Sigma, E \cup B, R)$ such that: (i) there is a sort *Time* in Σ such that (Σ, E) contains an algebraic axiomatization of a time data type, where time can be either discrete or continuous; (ii) there is also a sort *GlobalState*, where terms of sort *GlobalState* are pairs (t, r) , with t an “untimed state” (which may however contain time-related quantities such as timers), and r is a term of sort *Time* (that is, the global state is an untimed state plus a global clock); and (iii) the rules R are either: (a) *instantaneous* rules, which do not change the time and only rewrite the discrete part of the state, or (b) *tick* rules, of the form

$$(t, r) \rightarrow (t', r') \text{ if } C$$

where t and t' are term patterns describing untimed states, r and r' are terms of sort *Time*, and C is the rule’s condition. That is, tick rules advance the global clock and also update the untimed state to reflect the passage of time (for example, timers may be decreased, and so on). Real-Time rewrite theories provide a very expressive semantic framework in which many models of real-time systems can be naturally expressed [93]. The Real-Time Maude language [94] is an extension of Maude that supports specification, simulation, and model checking analysis of real-time systems specified as real-time rewrite theories.

How should the formal semantics of a *real-time* programming language be defined? And how can programs in such a language be formally analyzed? The obvious RLS answers are: (i) “with a real-time rewrite theory,” and (ii) “by real-time model checking and/or deductive reasoning based on such a theory.” Of course, the effectiveness of these answers has to be shown in actual languages. This is done in Section 5.

2.7 Defining Modeling Languages

It is well known that the most expensive errors in system development are not coding errors but design errors. Since design errors affect the overall structure of a system and are often discovered quite late in the development cycle, they can be enormously expensive to fix. All this is uncontroversial: there is widely-held agreement that, to develop systems, designs themselves should be made machine-representable, and that tools are needed to keep such designs consistent and to uncover design errors as early as possible. This has led to the development of many software modeling languages.

There are however two main limitations at present. The first is that some of these modeling notations lack a formal semantics: they can and do mean different things to different people. The second is that this lack of semantics

manifests itself at the practical level as a lack of *analytic power*, that is, as an incapacity to uncover expensive design errors which could have been caught by better analysis. It is of course virtually impossible to solve the second problem without solving the first: without a precise mathematical semantics any analytic claims about satisfaction of formal requirements are meaningless.

The practical upshot of all this is that a semantic framework such as rewriting logic can play an important role in: (i) giving a precise semantics to modeling languages; and in (ii) endowing such languages and notations with powerful formal analysis capabilities. Essentially the approach is the same as for programming languages. If, say, \mathcal{M} is a modeling language, then its formal semantics will be a rewrite theory of the form $(\Sigma_{\mathcal{M}}, E_{\mathcal{M}}, R_{\mathcal{M}})$. If the modeling language \mathcal{M} provides enough information about the dynamic behavior of models, the equations $E_{\mathcal{M}}$ and the rules $R_{\mathcal{M}}$ will make \mathcal{M} *executable*, that is, it will be possible to *simulate* models in \mathcal{M} before they are realized by concrete programs, and of course such models thus become amenable to various forms of *formal analysis*. All these ideas are further discussed in Section 6

2.8 Defining Hardware Description Languages

What is hardware? What is software? It depends in part on the level of abstraction chosen, and on specific implementation decisions: a given functionality may sometimes be realized as microcode, other times as code running on an FPGA, and yet other times may be implemented in custom VLSI. All this means that the difference between the semantics of digital hardware in some Hardware Description Language (HDL), and that of a programming language is not an essential one, just one about which level of abstraction is chosen. From the point of view of rewriting logic, both the semantics of an HDL and that of a programming language can be expressed by suitable rewrite theories. We further discuss the rewriting logic semantics of HDLs in Section 7.

2.9 Formal Analysis Methods and Tools

The fact that, under simple conditions, rewriting logic specifications are executable, means that the rewriting logic semantics of a language, whether a programming language, or a modeling language, or an HDL, is *executable* and therefore yields an *interpreter* for the given language when run on a rewriting logic system such as Maude. Since the language in question may not have any other formal semantics, the issue of whether the semantic definitions correctly capture the language's informal semantics is a nontrivial matter; certainly not trivial at all for real languages which may require hundreds of semantic rules. The fact that the semantics is executable is very useful in this regard, since one can *test* the correctness of the definitions by comparing the results from evaluating programs in the interpreter obtained from the rewriting logic semantics and in an actual language implementation. The usefulness of this approach is further discussed for the case of the semantics of C in Section 4.

Once the language specifier is sufficiently convinced that his/her semantic definitions correctly capture the language's informal semantics, various sophisticated forms of program analysis become possible. If some abstract semantics for the language in question has been defined, then the abstract semantic definition can be directly used as an *static analysis tool*. Since various abstract semantics may be defined for diverse analysis purposes, a collection of such tools may be developed. We further discuss this idea in Section 8.

Using a tool like Maude, the concrete rewriting logic semantics of a language becomes not just an interpreter, but also a *model checker* for the language in question. The point is that Maude can model check properties for any user-specified rewrite theory. Specifically, it can perform reachability analysis to detect violations of invariants using its breadth-first search feature; and it can also model check temporal logic properties with its LTL model checker. Such features can then be used to model check programs in the language whose rewriting semantics one has defined, or in an abstraction of it, as explained in Section 9.

Static analysis and model checking do not exhaust the formal analysis possibilities. A language's rewriting logic semantics can also be used as the basis for *deductive reasoning* about programs in such a language. The advantage of directly basing deductive reasoning methods on the semantics is that there is no gap between the operational semantics and the "program logic." This approach has been pioneered by *matching logic* [103, 102], a program verification logic, with substantial advantages over both Hoare logic and separation logic, which uses a language's rewriting logic semantics, including the possibility of using patterns to symbolically characterize sets of states, to mechanize the formal verification of programs, including programs that manipulate complex data structures. More on matching logic and the MatchC tool in Section 10.

3 Modular Definitions and the \mathbb{K} Framework

One major impediment blocking the broader use of semantic frameworks is the lack of scalability of semantic definitions. Lack of *modularity* is one of the main causes for this lack of scalability. Indeed, in many frameworks one often needs to redefine the semantics of the existing language features in order to include new, unrelated features. For example, in conventional SOS [99] one needs to more than double the number of rules in order to include an abrupt termination construct to a language, because the termination "signal" needs to be propagated through all the language constructs. Mosses' Modular SOS (MSOS) [87] addresses the non-modularity of SOS; it was shown that MSOS can be faithfully represented in rewriting logic, in a way that also preserves its modularity [23, 73, 25, 24, 29]. We here report on the \mathbb{K} framework, developed in parallel with the MSOS approach.

\mathbb{K} [105] is a modular executable semantic framework derived from rewriting logic. It works with terms, but its concurrent semantics is best explained in terms of graph rewriting intuitions [111]. \mathbb{K} was first introduced by the second author in the lecture notes of a programming language design course at the University of Illinois at Urbana-Champaign (UIUC) in Fall 2003 [101], as a means to define

Original language syntax	K Strictness	K Semantics
$AExp ::= Int$		$\langle x \cdot \rangle_k \langle \dots x \mapsto i \cdot \rangle_{state}$
Id		$\frac{}{i}$
$AExp + AExp$	[strict]	$i_1 + i_2 \rightarrow i_1 +_{Int} i_2$
$AExp / AExp$	[strict]	$i_1 / i_2 \rightarrow i_1 /_{Int} i_2$ where $i_2 \neq 0$
$BExp ::= Bool$		
$AExp <= AExp$	[seqstrict]	$i_1 <= i_2 \rightarrow i_1 \leq_{Int} i_2$
not $BExp$	[strict]	not $t \rightarrow \neg_{Bool} t$
$BExp$ and $BExp$	[strict(1)]	$true$ and $b \rightarrow b$ $false$ and $b \rightarrow false$
$Stmt ::= skip$		skip $\rightarrow \cdot$
$Id := AExp$	[strict(2)]	$\langle x := i \cdot \rangle_k \langle \dots x \mapsto \frac{}{i} \cdot \rangle_{state}$
$Stmt ; Stmt$		$s_1 ; s_2 \rightarrow s_1 \frown s_2$
if $BExp$	[strict(1)]	if true then s_1 else $s_2 \rightarrow s_1$
then $Stmt$ else $Stmt$		if false then s_1 else $s_2 \rightarrow s_2$
while $BExp$ do $Stmt$		$\langle \dots \text{while } b \text{ do } s \cdot \rangle_k$ if b then $(s ; \text{while } b \text{ do } s)$ else \cdot
$Pgm ::= \text{var List}\{Id\} ; Stmt$		$\langle \text{var } xl ; s \rangle_k \langle \dots \cdot \rangle_{state}$ $\frac{}{s} \quad xl \mapsto 0$

Fig. 1. \mathbb{K} definition of IMP: syntax (left), annotations (middle) and semantics (right); $x \in Id$, $xl \in \mathbf{List}\{Id\}$, $i, i_1, i_2 \in Int$, $t \in Bool$, $b \in BExp$, $s, s_1, s_2 \in Stmt$

concurrent languages in rewriting logic using Maude. Programming languages, calculi, as well as type systems or formal analyzers can be defined in \mathbb{K} by making use of special, potentially nested *cell* structures, and *rules*. There are two types of \mathbb{K} rules: *computational rules*, which count as computational steps, and *structural rules* (or “half equations”), which do not count as computational steps. The role of the structural rules is to rearrange the term so that the computational rules can apply. \mathbb{K} rules are *unconditional* (they may have side conditions, though), and they are *context-insensitive*.

We introduce \mathbb{K} by means of a simple imperative language, called IMP. In Section 3.2 we extend IMP with dynamic threads into IMP++, and in Section 8.1 we show how one can use \mathbb{K} to define a type checker for IMP++. This language experiment is borrowed from [105], where more details about \mathbb{K} can be found. We also refer the interested reader to <http://k-framework.org> for papers, workshops and an implementation. Our implementation of \mathbb{K} , the \mathbb{K} -Maude tool, consists of a translator to Maude, which is implemented using Perl scripting (about 6,000 lines) and Maude (about 9,000 lines).

3.1 \mathbb{K} Semantics of IMP

Figure 1 shows the complete \mathbb{K} definition of IMP, except for the configuration (explained below). The left column gives the IMP syntax. The middle column

augments it with \mathbb{K} *strictness attributes*, stating the evaluation strategy of some language constructs. Finally, the right column gives the semantic rules.

Language syntax is typically defined in \mathbb{K} using an “algebraic” context-free notation, i.e., one which allows users to make use of list, set, multiset and map structures without defining them. Note, e.g., that we used $\mathbf{List}\{Id\}$ as a non-terminal in the syntax of IMP in Figure 1. System configurations are defined in the same style. Configurations in \mathbb{K} are organized as potentially nested structures of *cells*, which are typically labeled to distinguish them from each other. We use angle brackets as cell wrappers. The \mathbb{K} configuration of IMP can be defined as:

$$Configuration_{IMP} \equiv \langle \langle K \rangle_k \langle \mathbf{Map}\{Id \mapsto Int\} \rangle_{state} \rangle_{\top}$$

In words, IMP configurations consist of a top cell $\langle \dots \rangle_{\top}$ containing two other cells inside: a cell $\langle \dots \rangle_k$ which holds a term of sort K (the computation) and a cell $\langle \dots \rangle_{state}$ which holds a map from variables to integers. As examples of IMP \mathbb{K} configurations, $\langle \langle x := 1; y := x+1 \rangle_k \langle \cdot \rangle_{state} \rangle_{\top}$ is a configuration holding program “ $x := 1; y := x+1$ ” and empty state, $\langle \langle x := 1; y := x+1 \rangle_k \langle x \mapsto 0 \ y \mapsto 1 \rangle_{state} \rangle_{\top}$ is a configuration holding the same program and a state $x \mapsto 0$ and $y \mapsto 1$.

The sort K , for *computational structures* or simply *computations*, has a special meaning in \mathbb{K} . The intuition for terms of sort K is that they have computational contents, such as programs or program fragments have. Technically, computations automatically extend the syntax of the original language (i.e., all syntactic categories are sunk into K) with a list structure with “ \curvearrowright ” (read “followed by”) as binary concatenation of computations and with “.” as the empty computation. For example, the intuition for a computation of the form $T_1 \curvearrowright T_2 \curvearrowright \dots \curvearrowright T_n$ is that the enlisted (computational) tasks should be processed sequentially. Computations give a uniform means to define and handle evaluation contexts and/or continuations as special cases: a computation “ $v \curvearrowright c$ ” can be thought of as “ $c[v]$ ”, that is, evaluation context c applied to v ” or as “passing v to continuation c ”. In fact, \mathbb{K} allows one to define evaluation contexts over the language syntax both directly, like in [105], or indirectly, by means of *strictness attributes* like in the middle column in Figure 1. However, one should be aware that these are nothing but convenient *notations*, which desugar into rules. For example, the evaluation strategies of sum, comparison and conditional in IMP specified by the strictness attributes in Figure 1 can be defined using the following *structural rules* (for diversity, we assume that the sum $+$ evaluates its arguments non-deterministically and the comparison \leq evaluates its arguments sequentially):

$$\begin{aligned} a_1 + a_2 &\rightleftharpoons a_1 \curvearrowright \square + a_2 \\ a_1 + a_2 &\rightleftharpoons a_2 \curvearrowright a_1 + \square \\ a_1 \leq a_2 &\rightleftharpoons a_1 \curvearrowright \square \leq a_2 \\ i_1 \leq a_2 &\rightleftharpoons a_2 \curvearrowright i_1 \leq \square \\ \text{if } b \text{ then } s_1 \text{ else } s_2 &\rightleftharpoons b \curvearrowright \text{if } \square \text{ then } s_1 \text{ else } s_2 \end{aligned}$$

The symbol \rightleftharpoons stands for two structural rules, one left-to-right and another right-to-left. The right-hand sides of the structural rules above contain, besides

the task sequentialization operator \curvearrowright , *freezer* operators containing \square in their names, such as $\square + _$, $_ + \square$, etc. The first rule above says that in any expression of the form $a_1 + a_2$, a_1 can be scheduled for processing while a_2 is being held for future processing. Since these rules are bi-directional, they can be used at will to structurally re-arrange the computations. Thus, when iteratively applied from left-to-right they fulfill the role of *splitting* syntax into an *evaluation context* (the tail of the resulting sequence of computational tasks) and a *redex* (the head of the resulting sequence), and when applied right-to-left they fulfill the role of *plugging* syntax into context. Our current implementation of \mathbb{K} automatically generates rules like the above, plus heuristics to apply them in one direction or the other, from strictness annotations to syntax like in Figure 1 (middle column).

Structural rules like those above decompose and eventually push the tasks that are ready for processing to the top (or the left) of the computation. Semantic rules then tell how to process the atomic tasks. The right column in Figure 1 shows the \mathbb{K} semantic rules of IMP. To explain them, let us first discuss the important notion of a \mathbb{K} rule, which is a strict generalization of the usual notion of a rewrite rule. \mathbb{K} rules explicitly mention the parts of the term that they read, write, or don't care about. The underlined parts are those which are written by the rule; the term underneath the line is the new subterm replacing the one above the line. All writes in a \mathbb{K} rule are applied in *one parallel step*, and, with some reasonable restrictions discussed in [111] that avoid read/write and write/write conflicts, writes in multiple \mathbb{K} rule instances can also apply in parallel. The ellipses “...” represent the volatile part of the term, that is, that part that the current rule does not care about and, consequently, can be concurrently modified by other rules. The operations which are not underlined represent the read-only part of the term: they need to stay unchanged during the application of the rule. For example, consider the assignment rule in Figure 1:

$$\frac{\langle x := i \dots \rangle_{\mathbb{K}}}{\cdot} \langle \dots x \mapsto \frac{_}{i} \dots \rangle_{\text{state}}$$

It says that once the assignment $x := i$ reaches the top of the computation, the value of x in the store is replaced by i and the assignment dissolves; in \mathbb{K} , “...” is a nameless variable of any sort and “.” is the unit (or empty) computation (“.” is a polymorphic unit of all list, set and multiset structures). The rule for variable declarations in Figure 1 (last one) expects an empty state and allocates and initializes with 0 all the declared variables; the dotted or dashed lines signify that the rule is structural, which is discussed next.

\mathbb{K} rules are split in two categories: *computational* and *structural*. Computational rules capture the intuition of computational steps in the execution of the defined system or language, while structural rules capture the intuition of structural rearrangement, rather than computational evolution, of the system. We use dashed or dotted lines in the structural rules. Ordinary rewrite rules are particular \mathbb{K} rules, when the entire term is replaced; in this case, we prefer to use the standard notation $l \rightarrow r$ as syntactic sugar for computational rules and the notation $l \dashrightarrow r$ or $l \dot{\rightarrow} r$ as syntactic sugar for structural rules. Figure 1 shows

Original language syntax	K Strictness	K Semantics
$AExp ::= \dots \mid ++ Id$		$\langle \frac{++ x \ \dots}{i +_{Int} 1} \rangle_k \langle \dots x \mapsto \frac{i}{i +_{Int} 1} \dots \rangle_{state}$
$Stmt ::= \dots$		
$\mid \text{print } AExp$	[strict]	$\langle \frac{\text{print } i \ \dots}{\cdot} \rangle_k \langle \dots \frac{\cdot}{i} \dots \rangle_{output}$
$\mid \text{halt}$		$\langle \frac{\text{halt} \ \dots}{\cdot} \rangle_k$
$\mid \text{spawn } Stmt$		$\langle \frac{\text{spawn } s \ \dots}{\cdot} \rangle_k \frac{\cdot}{\langle s \ \dots \rangle_k}$
$K ::= \dots \mid die$		$\langle die \rangle_k \rightarrow \cdot$

Fig. 2. K definition of IMP++ (extends that in Figure 1, *without changing anything*)

three explicit structural rules (as already discussed, the strictness attributes correspond to implicit ones): $s_1 ; s_2$ is rearranged as $s_1 \curvearrowright s_2$, loops are unrolled when they reach the top of the computation (unconstrained unrolling leads to non-termination), and declared variables are allocated in the state.

\mathbb{K} rewriting is a hybrid between term rewriting and graph rewriting, aimed at keeping the syntactic simplicity of the former and achieving the concurrency semantics benefits of the latter. While rewriting logic can theoretically capture the intended concurrent semantics of graph rewriting [70], the representation in [70] is impractical. The concurrent semantics of \mathbb{K} is given in terms of graph rewriting, taking into account the explicit sharing and variable volatility in \mathbb{K} rules, but avoiding the notational complexity of graph rewriting. However, our current implementation straightforwardly translates \mathbb{K} rules into rewrite rules and then uses Maude for execution and formal analysis. For example, the rule for assignment above gets translated into a rewrite rule of the form:

$$\langle x := i \ \dots \rangle_k \langle \text{before } x \mapsto j \ \text{after} \rangle_{state} \rightarrow \langle \text{rest} \rangle_k \langle \text{before } x \mapsto i \ \text{after} \rangle_{state}$$

Even though our current translation to Maude loses concurrency, a serializability result in [111] connecting \mathbb{K} rewriting and rewriting logic reasoning guarantees the soundness of execution and formal analysis of \mathbb{K} using Maude.

3.2 Extending IMP

In this section we highlight the modularity of \mathbb{K} by extending the IMP language in Section 3.1 with variable increment and dynamic threads. Figure 2 shows how the \mathbb{K} semantics of IMP is seamlessly extended into a semantics for IMP++. To accommodate the output, a new cell needs to be added to the configuration:

$$Configuration_{IMP++} \equiv \langle \langle K \rangle_k \langle \text{Map} \{ Id \mapsto Int \} \rangle_{state} \langle \text{List} \{ Int \} \rangle_{output} \rangle_{\top}$$

However, note that none of the existing IMP rules needs to change, because each of them only matches what it needs from the configuration. The construct

`print` is strict and its rule adds the value of its argument to the end of the output buffer (matches and replaces the unit “.” at the end of the buffer). The rule for `halt` dissolves the entire computation, and the rule for `spawn` creates a new $\langle \dots \rangle_k$ cell wrapping the spawned statement. The code in this new cell will be processed concurrently with the other threads. The last rule cools down a terminated thread by simply dissolving it; it is a structural rule since, again, we do not want it to count as a computational step.

4 Programming Language Semantics

Having formal semantics for real programming languages, regardless of the formalism that is being used, is undoubtedly a very important step, useful not only to help us understand those languages better but also to serve as a solid foundation for implementations and for program analysis and verification techniques and tools. Using rewriting logic as a formalism for such semantics has the additional benefit that such techniques and tools can be *directly derived* from the language semantics with minimal effort, as shown throughout this paper.

The rewriting logic semantics technique described in Section 3 has been used to define several programming languages or large fragments of them. Some of these languages serve as models for teaching various language paradigms, which we do not mention here but can be found on webpages for programming language courses at UIUC and can be reached from <http://k-framework.org>, while others are real programming languages, such as C [37], Scheme [67], or Java 1.4 [40, 42]. In this section we only briefly discuss the rewrite logic semantics of C [37], more precisely of the ISO/IEC 9899:1999 (C99) standard, as formalized by Chucky Ellison using the \mathbb{K} framework. This semantics is currently being used by several researchers and research groups, both directly in their tools and indirectly as a basis for understanding (and sometimes criticizing) the C language. This has led to the “C Semantics” Google code project repository at <http://c-semantics.googlecode.com/>.

The C semantics defined by Chucky Ellison defines approximately 120 C syntactic operators and 200 intermediate or auxiliary semantic operators. The definitions of these operators are given by 400 semantic rules and 172 helper rules spread over 2333 lines of code (LOC). However, it takes only 37 of those rules (201 LOC) to cover the behavior of statements, and another 119 for expressions (417 LOC). There are 353 rules for dealing with types, memory, and other necessary mechanisms. Finally, there are about 63 rules for the core of the standard library.

This is the most comprehensive formal semantics of C to date. It is executable and thoroughly tested. All aspects related to the features mentioned below are given a direct semantics. *Expressions*: referencing and dereferencing, casts, array indexing, structure members, arithmetic, bitwise, and logical operators, sizeof, increment and decrement, assignments, sequencing, ternary conditional; *Statements*: for, do-while, while, if, if/else, switch, goto, break, continue, return; *Types and Declarations*: enums, structs, unions, bitfields, initializers, static storage, typedefs; *Values*: regular scalar values (signed/unsigned arithmetic and pointer

types), structs, unions; *Standard Library*: malloc/free, set/longjmp, basic I/O; *Environment*: command line arguments; *Conversions*: (implicit) argument and parameter promotions and arithmetic conversion, and (explicit) casts.

No matter what the intended use is for a formal semantics, such a use is limited if one cannot achieve confidence in its correctness. To achieve this aim, executable semantics has an immense practical advantage over non-executable semantics, because one can simply test it. The C semantics in [37] has been encapsulated inside a drop-in replacement for Gnu's C Compiler (GCC), called "KCC". This allows one to test the semantics as one would test a compiler. Indeed, the C semantics has been successfully run against all the examples in the Kernigham and Ritchie manual that supposedly cover all the features of ANSI C. Moreover, a series of challenging C programs collected from the Internet, such as programs from the Obfuscated C programming competition, totaling more than 10,000 LOC are included in the regression tests of the C semantics, so these are all executed each time the semantics is changed. In addition to the above, the GCC C-torture-test (which contains 715 C programs conforming to the standard semantics of C99) has been executed in the C semantics and its behavior compared to that of GCC itself, as well as to Intel's C Compiler (ICC).

C is so complex that even dedicated and broadly used compilers like GCC or ICC cannot compile and execute all the programs in the GCC torture-test. All in all, considering all the tests that the C semantics has been tested on, the GCC and ICC compilers successfully passed 99% of them, while the C semantics (compiled into Maude using the \mathbb{K} framework tool) passed 96% of them. The C semantics ran over 90% of these programs in under 5 seconds (each). An additional 6% completed in 10 minutes, 1% in 40 minutes, and 2% further in under 2 days. The remaining programs either did not finish because they were computationally very intensive (such as FFTs), or they made use of features which were not yet defined (such as, e.g., unicode characters in strings). While this is not terribly fast performance, especially when compared to compiled C, the reader should keep in mind that this is an interpreter obtained *for free* from a formal semantics and that other existing semantics of C are either "paper" definitions (e.g., [51]), or not executable (e.g., [91]), or very slow (e.g., we were not able to execute factorial of 6 or the 4th Fibonacci's number using the Haskell-based definition in [95, 96]), or covering only a C fragment (e.g., [14]). Moreover, our semantics of C can be used *directly* and *unchanged* for other purposes, such as for model checking (Section 9) and for deductive verification (Section 10).

5 Real-Time Language Semantics

Three real-time programming languages have been given formal semantics as real-time rewrite theories [93] in Real-Time Maude [94]. Using the model checking features of Real-Time Maude it then becomes possible to formally analyze programs in such languages.

In [4], AlTurki et al. present a language for real-time concurrent programming for industrial use in DOCOMO Labs called *L*. The goal of *L* is to serve as

a programming model for higher-level software specifications in SDL or UML. A related goal is to support formal analysis of L programs by both real-time model checking and static analysis, so that software design errors can be caught at design time. The way all this is accomplished is by giving a formal semantics to L in Real-Time Maude, which automatically provides an interpreter and a real-time model checker for L . Static analysis capabilities are added to L by using Maude to define an *abstract semantics* for L in rewriting logic, which is then used as the static analyzer.

The Orc model of real-time concurrent computation [79, 80, 126] has been given semantics in rewriting logic using real-time rewrite theories [5, 6]. Although Orc is a very simple and elegant language, its real-time semantics is quite subtle for two reasons. First, in the evaluation of any Orc expression, internal computation always has higher priority than the handling of external events; this means that, even without modeling time, a vanilla-flavored SOS semantics is not expressive enough to capture these different priorities: two SOS relations are needed [80]. Second, Orc is by design a real-time language, where time is a crucial feature. Using real-time rewrite theories, this double subtlety of the Orc semantics was faithfully captured in [5]; furthermore, this semantics yielded of course an Orc interpreter and a real-time model checker. But Orc is not just a model of computation: it is also a concurrent programming language. This suggested the following challenge question: can a correct-by-construction distributed Orc implementation be derived from its rewriting logic semantics? This question was answered in two stages. Since, as discussed in Section 2.5, a small-step SOS semantics is typically horribly inefficient and it was certainly so in the case of Orc, a much more efficient *reduction semantics* was first defined in [6], and was proved to be bisimilar to the small-step SOS semantics. This semantics provided a much more efficient interpreter and model checker. Furthermore, to explicitly model different Orc clients and various web sites, and their message passing communication, the Orc semantics was seamlessly extended in [6] to a distributed object-based Orc semantics, which modeled what a distributed implementation should look like. The only remaining step was to pass from this model of a distributed implementation to an actual Maude-based distributed real-time implementation. This was accomplished in [7] using three main ideas: (i) the use of sockets in Maude to actually deploy a distributed implementation; (ii) the systematic replacement of logical time by physical time, supported by Ticker objects external to Maude, while retaining the rewriting semantics throughout; and (iii) the experimental estimation of the physical time required for “zero-time” Maude subcomputations, to ensure that the granularity of time ticks is such that all “instantaneous transitions” have already happened before the next tick.

Creol is an object-oriented language supporting concurrent objects which communicate through asynchronous method calls. Its rewriting-logic-based operational semantics was defined in [55] without real-time features. However, to support applications such as sensor systems with wireless communication, where messages expire and may collide with each other, Creol’s design and operational

semantics have been extended in [13] to Timed Creol using rewriting logic. The notion of time used by Timed Creol is described as a “lightweight” one in [13]. Time is discrete and is represented by a time object. This approach does not require a full use of the features in Real-Time Maude (Maude itself is sufficient to define the real-time semantics). The effectiveness of Timed Creol in the modeling and analysis of applications such as sensor networks is illustrated in [13] through a case study.

6 Modeling Language Semantics

Modeling languages are quite useful, but they can be made even more useful by substantially increasing their analytic power through formal analysis, since this can make it possible to catch expensive design errors very early. Formal analysis is impossible or fraudulent without a formal semantics. Early work in developing rewriting-logic-based formal semantics focused on object-oriented design notations and languages [127, 90, 89], and stimulated subsequent work on UML and UML-like notations, e.g., [44, 62, 63, 128, 8, 33, 83, 82, 84].

A more ambitious question is: can we give semantics not just to a single modeling language, but to an entire *modeling framework* where different modeling languages can be defined? This question has been answered positively in [16, 15, 17, 19, 20]. This line of research has led to MOMENT2, an algebraic model management framework and tool written in Maude and developed by Artur Boronat [15]. It permits manipulating software models in the Eclipse Modeling Framework (EMF). It uses OMG standards, such as Meta-Object Facility (MOF), Object Constraint Language (OCL) and Query/View/Transformation (QVT), as a clean interface between rewriting-logic-based formal methods and model-based industrial tools. Specifically, it supports formal analyses based on rewriting logic and graph transformations to endow model-driven software engineering with strong analytic capabilities. MOMENT2 supports not just one fixed modeling language, but any modeling language whose *meta-model* is specified in MOF. In more detail, a modeling language is specified as a pair $(\mathcal{M}, \mathcal{C})$, where \mathcal{M} is its MOF-based metamodel, and \mathcal{C} are the OCL constraints that \mathcal{M} should satisfy. Using rewriting-logic-based reflection and its efficient support in Maude, MOMENT2 provides an *executable algebraic semantics* for such metamodel specifications $(\mathcal{M}, \mathcal{C})$ in the form of a theory $\mathbb{A}(\mathcal{M}, \mathcal{C})$ in membership equational logic (MEL) [71], so that a model M conformant with the metamodel $(\mathcal{M}, \mathcal{C})$ is exactly a term of sort *Model* in $\mathbb{A}(\mathcal{M}, \mathcal{C})$, and so that satisfaction of OCL constraints is also decidable using the algebraic semantics [18, 20].

Due to the executability of MEL specifications in Maude, the realization of MOF metamodels as MEL theories enhances the formalization and prototyping of model-driven development processes, such as: (i) model transformations; (ii) model-driven roundtrip engineering; (iii) model traceability; and (iv) model management. These processes permit, for example, merging models, generating mappings between models, and computing differences between models; they can be used to solve complex scenarios such as the roundtrip problem. In MOMENT2

the formal semantics of *model transformations* is given by rewrite theories specified in a user-friendly QVT-based syntax [17]. Such model transformations can describe the dynamic evolution of systems at the level of their models. Using the search and *LTL* model checking features of Maude, properties about the dynamic evolution of a model M conformant with a metamodel specification $(\mathcal{M}, \mathcal{C})$ can then be formally analyzed by model checking [17]. Real-time modeling languages can likewise be supported and analyzed [21]; this is further discussed below.

6.1 Semantics of Real-Time Modeling Languages

There is strong interest in modeling languages for real-time and embedded systems. The rewriting logic semantics for such modeling languages can be naturally based on real-time rewrite theories. Using a tool like Real-Time Maude, what this means in practice is that such models can then be simulated; and that their formal properties, in particular their safety requirements, can be model checked. Furthermore, the simulations and formal analysis capabilities added to the given modeling language can be offered as “plugins” to already existing modeling tools, so that much of the formal analysis happens “under the hood,” and somebody already familiar with the given modeling notation can perform such formal analysis without having an in-depth understanding of the underlying formalism.

The Ptolemy II modeling language (<http://ptolemy.eecs.berkeley.edu>) supports design and simulation of concurrent, real-time, embedded systems expressed in several models of computation (MoCs), such as state machines, data flow, and discrete-event models, that govern the interaction between concurrent components. A user can visually design and simulate hierarchical models, which may combine different MoCs. Furthermore, Ptolemy II has code generation capabilities to translate models into other modeling or programming languages such as C or Java. Discrete-Event (DE) Models are among the most central in Ptolemy II. Their semantics is defined by the *tagged signal model* [64]. The work by Bae et al. in [11] endows DE models in Ptolemy II with formal analysis capabilities by: (i) defining a semantics for them as real-time rewrite theories; (ii) automating such a formal semantics as a model transformation using Ptolemy II’s code generation features; (iii) providing a Real-Time Maude plugin, so that Ptolemy II users can use an extended GUI to define temporal logic properties of their models in an intuitive syntax and can invoke Real-Time Maude from the GUI to model check their models. This work has been further advanced in [9] to support not just flat DE models, but *hierarchical* ones. That is, above tasks (i)–(iii) have been extended to hierarchical DE models; this extension is non-trivial, because it requires combining synchronous fixpoint computations with hierarchical structure.

AADL (<http://www.aadl.info/>) is a standard for modeling embedded systems that is widely used in avionics and other safety-critical applications. However, AADL lacks a formal semantics, which severely limits both unambiguous communication among model developers and the formal analysis of AADL models. In [92] Ölveczky et al. define a formal object-based real-time concurrent semantics for a behavioral subset of AADL in rewriting logic, which includes the

essential aspects of AADL's behavior annex. Such a semantics is directly executable in Real-Time Maude and provides an AADL simulator and LTL model checking tool called *AADL2Maude*. *AADL2Maude* is integrated with OSATE, so that OSATE's code generation facility is used to automatically transform AADL models into their corresponding Real-Time Maude specifications. Such transformed models can then be executed and model checked by Real-Time Maude. One difficulty with AADL models is that, by being made up of various hierarchical components that communicate asynchronously with each other, their model checking formal analysis can easily experience a combinatorial explosion. However, many such models express designs of distributed embedded systems which, while being asynchronous, should behave in a virtually synchronous way. This suggests the possibility of using the PALS pattern [75], which reduces distributed real-time systems with virtual synchrony to synchronous ones, to pass from simple synchronous systems, which have much smaller state spaces and are much easier to model check, to semantically equivalent asynchronous systems, which often cannot be directly model checked but can be verified indirectly through their synchronous counterparts. This has led to the design of the Synchronous AADL sublanguage in [10], where the user can specify synchronous AADL models by using a sublanguage of AADL with some special keywords. A synchronous rewriting semantics for such models has also been defined in [10]. Using OSATE's code generation facility, synchronous AADL models can be transformed into their corresponding Real-Time Maude specifications in the *SynchAADL2Maude* tool, which is provided as a plugin to OSATE. Likewise, the user can define temporal logic properties of synchronous AADL models based on their features, without requiring knowledge of the underlying formalism, and can model check such models in Real-Time Maude.

A more ambitious goal is to provide a *framework*, where a wide range of real-time Domain-Specific Visual Languages (DSVLs), as well as their dynamic real-time behavior, can be specified with a rigorous semantics. This is precisely the goal of two frameworks and associated tools: (i) the *e-Motions* framework [100]; and (ii) *MOMENT2*'s support for real-time DSVLs [21].

- In *e-Motions*, DSVLs are specified by their corresponding metamodels, and dynamic behavior is specified by rules that define in-place model transformations. But the goals of *e-Motions* do not remain at the syntax/visual level: they also include giving a precise rewriting logic semantics in Real-Time Maude to the different real-time DSVLs that can be defined in *e-Motions*, and to automatically support simulation and formal analysis of models by using the underlying Real-Time Maude engine. The formal semantics translates the metamodel of a DSVL as an object class, the corresponding models as object configurations of that class, and the *e-Motions* rules as rewrite rules. Since all these translations are automatic and define a DSVL's formal semantics, a modeling language designer using *e-Motions* does not have to explicitly define the DSVL's formal semantics: it comes for free, together with the simulation and model checking features, once the DSVL's metamodel and the dynamic behavior rules are specified.

- In [21], the *MOMENT2* framework has been extended to support the formal specification and analysis of real-time model-based systems. This is achieved by means of a collection of built-in timed constructs for defining the timed behavior of such systems. Timed behavior is specified using in-place model transformations. Furthermore, the formal semantics of a *timed behavioral specification* in *MOMENT2* is given by a corresponding real-time rewrite theory. In this way, models can be simulated and model checked using *MOMENT2*'s Maude-based analysis tools. In addition, by using in-place multi-domain model transformations in *MOMENT2*, an existing model-based system can be extended with timed features in a non-intrusive way, in the sense that no modification is needed for the class diagram.

7 Hardware Description Language Semantics

The rewriting logic semantics project has been naturally extended from the level of programming languages to that of *hardware description languages* (HDLs). In this way, hardware designs written in an HDL can be both simulated and analyzed using the executable rewriting semantics of the HDL and tools like ELAN, CafeOBJ, or Maude. The first HDL to be given a rewriting logic semantics in Maude was ABEL [58]; this semantics was used not only for hardware designs, but also for hardware/software co-designs. An important new development has been the use of the rewriting logic semantics of an HDL for *generating sophisticated test inputs for hardware designs*. The point is that random testing can catch a good number of design errors, but uncovering deeper errors after random testing is hard and costly and requires a good understanding of the design to exercise complex computation sequences. The key insight, due to Michael Katelman, is that the rewriting semantics can be used *symbolically* to generate desired test inputs, not on a device's concrete states, but on states that are partly symbolic (contain logical variables) and partly concrete. This symbolic approach, first outlined in [60] and more fully developed in [59], has a number of unique features including: (i) the use of SAT solvers to symbolically solve Boolean constraints; (ii) support for user-guided random generation of partial instantiations; and (iii) a flexible *strategy language*, in which a hardware designer can specify in a declarative, high-level way the kind of test that needs to be generated. The effectiveness of this approach for generating sophisticated tests on real hardware designs has already been demonstrated for medium-sized Verilog designs [59]. The `vlogs1` tool is currently undergoing further enhancements to efficiently handle large designs.

But the value of the rewriting semantics of an HDL is not restricted to testing. For example, the recent Maude-based rewriting logic semantics of Verilog in [68] is arguably the most complete formal semantics to date, both in the sense of covering the largest subset of the language and in its faithful modeling of non-deterministic features. Besides being executable and supporting formal analysis, this semantics has uncovered several nontrivial bugs in various mature Verilog

tools, and can serve as a practical and rigorous standard to ascertain what the correct behavior of such tools should be in complex cases.

A more exotic application of rewriting logic semantics, for which it is ideally suited due to its intrinsically concurrent nature, is that of *asynchronous hardware designs*. These are digital designs which do not have a global clock, so that different gates in a device can fire at different times. Such devices can behave correctly in much harsher environments (e.g., a satellite in outer space) and with much wider ranges of physical operating conditions than clocked devices. Asynchronous designs can be specified with the notation of *production rules*, which roughly speaking describe how each gate behaves when inputs to its wires are available. In [57] a rewriting logic semantics of asynchronous digital devices specified as sets of production rules is given and is realized in Maude. This is the first executable formal semantics of such devices we are aware of. It can be used both for simulation purposes and for model checking verification of small-sized devices (about 100 gates). An interesting challenge is how to scale up model checking for larger devices; this is nontrivial due to the large combinatorial explosion caused by their asynchronous behavior.

8 Abstract vs. Concrete Semantics and Static Analysis

In addition to helping with understanding and experimenting with language designs, a rewriting logic semantics can have several direct uses without having to change the semantics at all. Two such uses of unchanged semantics in the context of program verification are discussed in Sections 9 and 10. Nevertheless, there are program analysis needs where the desired information is not necessarily available in the code itself, or where the desired domain of analysis is not included in, and cannot be obtained from, the concrete domain in which the language semantics operates. In such cases, one can modify the concrete language semantics to operate within a target *abstract domain*. We next first show an overly simplified example, where the concrete semantics of IMP and IMP++ in Sections 3.1 and 3.2 are abstracted into type systems for the defined languages, which yield type checkers when executed. Then we discuss uses of similar but larger scale and more practical abstractions of rewrite logic semantics.

8.1 \mathbb{K} Definition of a Type System for IMP++

The \mathbb{K} semantics of IMP/IMP++ in Sections 3.1 and 3.2 can be used to execute even ill-typed IMP/IMP++ programs, which may be considered undesirable by some language designers. In this section we show how to define a type system for IMP/IMP++ using the very same \mathbb{K} framework. The type system is defined like an (executable) semantics of the language, but one in the more abstract domain of types rather than in the concrete domain of integer and Boolean values.

The typing policy that we want to enforce on IMP/IMP++ programs is easy: all variables in a program have by default integer type and must be declared, arithmetic/Boolean operations are applied only on expressions of corresponding types, etc. Since programs and program fragments are now going to

Original language syntax	K Strictness	K Semantics
$AExp ::= Int$		$i \rightarrow int$
Id		$\langle x \ \dots \rangle_k \langle \dots x \ \dots \rangle_{var}$
$AExp + AExp$	[strict]	int
$AExp / AExp$	[strict]	$int + int \rightarrow int$
$++ Id$		$int / int \rightarrow int$
		$\langle ++ x \ \dots \rangle_k \langle \dots x \ \dots \rangle_{var}$
		int
$BExp ::= AExp <= AExp$	[strict]	$int <= int \rightarrow bool$
not $BExp$	[strict]	not $bool \rightarrow bool$
$BExp$ and $BExp$	[strict]	$bool$ and $bool \rightarrow bool$
$Stmt ::= skip$		skip $\rightarrow stmt$
$Id := AExp$	[strict(2)]	$\langle x := int \ \dots \rangle_k \langle \dots x \ \dots \rangle_{var}$
		$stmt$
$Stmt ; Stmt$	[strict]	$stmt ; stmt \rightarrow stmt$
if $BExp$		
then $Stmt$ else $Stmt$	[strict]	if $bool$ then $stmt$ else $stmt \rightarrow stmt$
while $BExp$ do $Stmt$	[strict]	while $bool$ do $stmt \rightarrow stmt$
print $AExp$	[strict]	print $int \rightarrow stmt$
halt		halt $\rightarrow stmt$
spawn $Stmt$	[strict]	spawn $stmt \rightarrow stmt$
$Pgm ::= var \ List\{Id\} ; Stmt$		$\langle var \ xl ; s \rangle_k \langle \cdot \rangle_{vars}$
		$s \curvearrowright pgm \quad xl$
		$stmt \curvearrowright pgm \rightarrow pgm$

Fig. 3. K type system for IMP++ (and IMP)

be rewritten into their types, we need to add to computations some basic types. Also, in addition to the computation to be typed, configurations must also hold the declared variables. Thus, we define the following (the “...” in the definition of K includes all the default syntax of computations, such as the original language syntax, \curvearrowright , freezers, etc.):

$$K ::= \dots \mid int \mid bool \mid stmt \mid pgm$$

$$Configuration_{IMP++}^{Type} \equiv \langle \langle K \rangle_k \langle List\{Id\} \rangle_{vars} \rangle_{\top}$$

Figure 3 shows the IMP/IMP++ type system as a \mathbb{K} system over such configurations. Constants reduce to their types, and types are straightforwardly propagated through each language construct. Note that almost each language construct is strict now, because we want to type all its arguments in almost all cases in order to apply the typing policy of the construct. Two constructs are exceptional, namely, increment and assignment. The typing policy of these constructs is that they take precisely a variable and not something that types to an integer. If we defined, e.g., the assignment strict and with rule $int := int \rightarrow stmt$, then our type system would allow ill-formed programs like $x+y := 0$. Note how we defined the typing policy of programs $var \ xl ; s$: the declared variables xl are stored into the $\langle \dots \rangle_{vars}$ cell (which is expected to initially be empty) and the statement is scheduled for typing (using a structural rule), placing a “reminder”

in the computation that the *pgm* type is expected; once/if the statement is correctly typed, the type *pgm* is generated.

8.2 Examples of Abstract Rewriting Logic Semantics

We briefly discuss three practical uses of abstract rewriting logic semantics.

C Pluggable Policies. Many programs make implicit assumptions about data. Common examples include assumptions about whether variables have been initialized or can only contain non-null references. Domain-specific examples are also common; a compelling example is units of measurement, used in many scientific computing applications, where different variables and values are assumed to have specific units at specific times/along specific execution paths. These implicit assumptions give rise to implicit domain *policies*, such as requiring assignments to non-null pointers to also be non-null, or requiring two operands in an addition operation to have compatible units of measurement.

Mark Hills et al. [53] propose a framework for *pluggable policies* for C which allows these implicit policies to be made explicit and checked. The core of the framework is a shared annotation engine and parser, allowing annotations in multiple policies to be inserted by developers as comments in C programs, and a shared abstract rewriting logic semantics of C designed as a number of reusable modules that allow for new policies to be quickly developed and plugged in. For instance, a case study for checking non-null references was developed in under two days; another case study for checking units of measurement reuses the shared abstract semantics and only adds domain knowledge [53].

Polymorphic Type Inference. The technique in Section 8.1 for defining type systems using \mathbb{K} is very general and has been used to define more complex type systems, such as higher-order polymorphic ones by Ellison et al. [38]. The \mathbb{K} definition of the type system in [38] is more declarative and thus cleaner and easier to understand than alternative algorithmic definitions. Moreover, the \mathbb{K} definition is formal, so it is amenable for formal reasoning. Interestingly, as shown in [38], the resulting \mathbb{K} definition, when compiled to and executed using Maude, was faster than algorithmic implementations of the same type system found on the internet as teaching material. In fact, experiments in [38] show that it was comparable to state of the art implementations of type inferencers in conventional functional languages! For example, it was only about twice as slow on average than that of OCaml, and had average times comparable, or even better than those of Haskell ghci and SML/NJ.

Security Policy Checking. An elegant application of a programming language's *abstract* rewriting logic semantics to Java code security is presented by Alba-Castro et al. in [2, 3] as part of their rewriting-logic-semantics-based approach to proof carrying code. The key idea is to use an abstract rewriting logic

semantics of Java that correctly approximates security properties such as *noninterference* (that is, the specification of what objects should not have any effects on other objects according to a stated security policy [48]), and *erasure* (a security policy that mandates that secret data should be removed after its intended use). Since the abstract rewriting semantics is finite-state, it supports the automatic creation of certificates for noninterference and erasure properties of Java programs that are independently checkable and small enough to be practical.

9 Model Checking Verification

Once a programming language or system is defined as a rewrite theory, one can use any general-purpose tools and techniques for rewriting logic to obtain tools and techniques specialized for the defined programming language or system. We have reported in the past on the use of Maude's general purpose LTL model checking capabilities to obtain model checkers specialized for various concurrent programming languages, including Java and the JVM (see, e.g., [76, 40, 42]). In this paper we report on some new model checking experiments performed in the context of the C definition discussed in Section 4. We thank Chucky Ellison for extending his C semantics with concurrency primitives and for conducting these experiments. A more detailed presentation of these can be found in [37].

The C semantics in Section 4 can be extended to include semantics for concurrency primitives like “spawn”, “sync”, “lock”, and “unlock”. The former is used to dynamically spawn a new execution thread, “sync” waits for all of the other threads to die before continuing, and “lock” and “unlock” synchronize threads on memory locations (similar to Java locking on references). When formalizing the semantics of C, we did not plan to introduce concurrency. Despite that, as hoped for, the existing rules were left unchanged upon adding configuration support and the semantics of threads.

Dekker's Algorithm We now take a look at the classical Dekker's algorithm, in order to explore thread interleavings.

```

void dekker1(void) {
    flag1 = 1; turn = 2;
    while((flag2 == 1) && (turn == 2)) ;
    critical1();
    flag1 = 0;
}

void dekker2(void) {
    flag2 = 1; turn = 1;
    while((flag1 == 1) && (turn == 1)) ;
    critical2();
    flag2 = 0;
}

```

These two functions get called by the two threads respectively to ensure mutual exclusion of the calls to `criticaln()`. In the program we used for testing, these threads each contain infinite loops while the function `main()` waits on a `sync()`. Thus, the program never terminates.

To test the mutual exclusion property, we model check the following LTL formula: $\Box \neg(\text{enabled}(\text{critical1}) \wedge \text{enabled}(\text{critical2}))$, stating that the two critical sections can never be called at the same time. Applying this formula to our program yields “**result Bool: true**”, in 400ms. If we break the algorithm by changing a `while` to an `if`, the tool instead returns a list of rules, together with the resulting states, that represent a counterexample.

Dining Philosophers Another classic example is the dining philosophers problem.

```

void philosopher( int n ) {
  while(1) {
    // Hungry: obtain chopsticks
    if ( n % 2 == 0 ) { // Even number: Left, then right
      lock(&chopstick[(n+1) % NUM_PHILOSOPHERS]);
      lock(&chopstick[n]);
    } else { // Odd number: Right, then left
      lock(&chopstick[n]);
      lock(&chopstick[(n+1) % NUM_PHILOSOPHERS]);
    }
    // Eating
    // Finished Eating: release chopsticks
    unlock(&chopstick[n]);
    unlock(&chopstick[(n+1) % NUM_PHILOSOPHERS]);
    // Thinking
  }
}

```

The above code shows a solution to the dining philosophers that has even-numbered philosophers picking up their left chopstick first, while odd-numbered philosophers pick up their right chopstick first. This strategy ensures that there is no deadlock. We can use Maude’s search command to verify that there is no deadlock simply by searching for final states. Here are the results:

n	No Deadlock		With Deadlock	
	number of states	time (s)	number of states	time (s)
1	19	0.1	–	–
2	92	0.8	63	0.6
3	987	14.0	490	7.2
4	14610	293.5	5690	119.8
5	288511	8360.3	84369	2376.5

In the “No Deadlock” column we see the results for the code above. We were able to verify that with this algorithm, there were no deadlocks for up to five philosophers. In the “With Deadlock” column, we altered the code so that all philosophers would try to pick up their left chopstick first. For this algorithm, we were able to find counterexamples showing that the program has deadlocks.

While the classic programs above are toy examples, which are far from the complexity of real-life software, we believe that they are sufficient to show that a programming language semantics can be more than a “useless academic intellectual exercise”. The well-known state-space explosion of model checking cannot be avoided, no matter whether one uses a formal semantics of the language or not, but one should note that this is a problem of model checking and not of using a formal semantics for model checking. Also, there are well-known techniques to address the state explosion problem, like partial-order reduction, which can and have also been applied in the context of rewriting logic semantics [41]. And one can use an *abstract semantics* (Section 8) as the basis of the model checker to make it more scalable. The next section shows another use of rewriting logic semantics of programming languages, for deductive program verification.

10 Deductive Verification and Matching Logic

As discussed above, one of the major advantages of giving a rewriting logic semantics to a language is that one can use it not only to obtain a reference implementation of the language, but also to formally analyze programs in the defined language using general-purpose tools developed for rewriting logic, such as Maude's model checker. Moreover, the original rewriting logic semantics of the language is used unchanged for model checking or other similar analyses, which is not only immensely convenient but also offers a high confidence in the results of the analysis (because it excludes the problem of implementing a wrong language semantics in the analyzer). One question, however, still remains unanswered: can we use the language semantics, also unchanged, in a program logic fashion, that is, for deductive verification of programs?

Early work in this direction includes two Hoare logic provers that use directly the rewriting logic semantics of a Pascal like-language and of a fragment of Java and the Maude ITP [34, 107]. Furthermore, the rewriting logic semantics of Java was used in [1] to automatically validate the inference rules of a Java verification tool. In the remainder of this section we report on an alternative approach.

Matching logic [102, 103] is a new program verification logic, which builds upon rewriting logic semantics. Matching logic specifications are constrained symbolic program configurations, called *patterns*, which can be *matched* by concrete configurations. By building upon an executable semantics of the language and allowing specifications to directly refer to the structure of the configuration, matching logic has at least three benefits: (1) one's familiarity with the formalism reduces to one's familiarity with the formal semantics of the language, that is, with the language itself; (2) the verification process proceeds the same way as the program execution, making debugging failed proof attempts manageable because one can always see the "current configuration" and "what went wrong", almost like in a debugger; and (3) nothing is lost in translation, that is, there is no gap between the language definition and its verifier. Moreover, direct access to the structure of the configuration facilitates defining sub-patterns that one may reason about, such as disjoint lists or trees in the heap, as well as supporting framing in various components of the configuration at no additional cost.

To use matching logic for program verification, one must know the structure of the configurations that are used in the executable language semantics. For example, the configuration of some language may contain, besides the code itself, an environment, a heap, stacks, synchronization resources, etc. The configuration of C (see Section 4 and [37]), for example, consists of 75 cells, each containing either other cells or some piece of semantic information. Matching logic specifications, or patterns, allow one to refer directly to the configuration of the program. Moreover, we can use logical variables and thus combine the desired configuration structure with first-order constraints. For example, the pattern

$$\langle \langle \beta, I \rangle_{\text{in}} \langle x \mapsto x, i \mapsto i, n \mapsto n, E \rangle_{\text{env}} \langle \text{list}(x, \alpha), H \rangle_{\text{heap}} C \rangle_{\text{config}} \\ \wedge i \leq n \wedge |\beta| = n - i \wedge A = \text{rev}(\alpha) @ \beta$$

specifies the set of configurations where program variables x , i and n are bound in the environment to some values x , i , and respectively n , such that $i \leq n$, the input buffer contains a sequence β of size $n - i$, and the heap contains a linked list starting with pointer x comprising the sequence of elements α such that the sequence A is the reverse of the sequence α concatenated with β . Here A is a free variable of type sequence of elements. The other variables play the role of cell frames: I is a variable matching the rest of the input cell, E matches the rest of the environment, H the rest of the heap, and C the rest of the configuration. Note that nothing special needs to be done for framing in matching logic (that is, framing is a special case of the more general principle of matching).

A major benefit of matching logic is that it can be used to turn an executable semantics into a program logic without any change to the original semantics. The idea is that the executable semantics can be regarded as a set of rewrite rules between matching logic patterns, and one can use first-order reasoning over patterns to turn the pattern resulting from the application of some rule into a pattern that the next rule expects to match. This way, one can derive rewrite rules from other rewrite rules, using matching logic reasoning as a mechanism to rearrange configurations so that rewrite rules can match and apply.

With the help of Andrei Ştefănescu, we implemented a proof-of-concept matching logic verifier for a fragment of C, called MATCHC, which can be downloaded and executed online at <http://fsl.cs.uiuc.edu/ml>. MATCHC builds upon an executable rewrite-based semantics of this fragment of C, extending it (unchanged) with semantics for pattern specifications. Both the executable semantics and the verifier are implemented using the \mathbb{K} framework (see Section 3).

Figure 4 shows a C program verified using MatchC. The `main()` function reads `n` from the standard input and then calls `readWriteBuffer(n)`. Then `readWriteBuffer(n)` reads from the standard input `n` elements and allocates a linked list putting each element at the top of the list, followed by traversing the linked list and printing each element while deallocating the list nodes. This way, we end up with the reversed sequence of elements printed to the standard output and with the heap unchanged. There are four types of annotations in this program: (1) *assumptions*, which allow one to assume a certain pattern for the remaining program; (2) *assertions*, which generate matching logic proof obligations, namely, that the current pattern implies the asserted pattern; (3) *rules*, which give the claimed \mathbb{K} semantics of the subsequent piece of code; and (4) *invariants*, which are patterns that should hold at each loop iteration.

Some explanations regarding MatchC's notation are necessary. MatchC annotations are introduced like C comments starting with `@`, so they are ignored by C compilers. We use an XML-like notation to specify when cells start and when they end. We use the usual rewriting relation "`=>`" for the in-place rewriting within \mathbb{K} rules. The "`$`" symbol that appears in the computation cell of a rule stands for the subsequent statement (the function body, in our case here). Fourth, to avoid writing quantifiers, variables starting with a question mark are existentially quantified over the pattern. Fifth, we use an underscore in the XML tag to state that the corresponding cell is open in that direction, which can be

```

#include <stdlib.h>
#include <stdio.h>

struct listNode { int val; struct listNode *next; };

void readWriteBuffer(int n)
/*@ rule <k> $ => return; </k> <in> A => epsilon <_/in> <out_> epsilon => rev(A) </out>
   if n = len(A) */
{
  int i; struct listNode *x;
  i = 0; x = 0;
  /*@ inv <in> ?B <_/in> <heap_> list(x)(?A) <_/heap>
     /\ i <= n /\ len(?B) = n - i /\ A = rev(?A) @ ?B */
  while (i < n) {
    struct listNode *y;
    y = x;
    x = (struct listNode*) malloc(sizeof(struct listNode));
    scanf("%d", &(x->val));
    x->next = y;
    i += 1;
  }

  /*@ inv <out_> ?A </out> <heap_> list(x)(?B) <_/heap> /\ A = rev(?A @ ?B)
     while (x) {
       struct listNode *y;
       y = x->next;
       printf("%d ",x->val);
       free(x);
       x = y;
     }
  */
}

void main() {
  int n;
  /*@ assume <in> [5, 1, 2, 3, 4, 5] </in> <out> epsilon </out>
     scanf("%d", &n);
     readWriteBuffer(n);
     /*@ assert <in> epsilon </in> <out> [5, 4, 3, 2, 1] </out>
  */
}

```

Fig. 4. C program making use of the I/O and the heap, verified using MatchC.

regarded as an abbreviation for using a fresh variable; for example, “<in> ?B <_/in>” in the invariant of the first loop abbreviates “<in> ?B, ?E </in>”. Finally, to avoid writing the environment cell all the time, MatchC allows users to refer directly to program variables in patterns; this avoids having to add a binding of the program variable to a logical variable in the environment cell and then using the logical variable throughout the pattern.

The rule giving the semantics of `readWriteBuffer(n)` states that this function returns nothing (“\$ => return;”, that is, its body behaves as if it returns) and takes a sequence A of length n (see the condition “ $n = \text{len}(A)$ ”) from the beginning of the input cell (“<in> A => epsilon <_/in>”) and places it reversed at the end of the output cell (“<out_> epsilon => A </out>”). Since we have a rewrite-based semantics, the fact that no other cells are mentioned implicitly means that *nothing else is modified by this function*, including the heap. The invariant of the first loop is exactly the pattern that we discussed at the beginning of this section. The invariant of the second loop is similar, but dual. We do not show the axiom (matching logic formula) governing the list pattern

in the heap cell; the interested reader can check [102, 103]. Nevertheless, since x is null at the end of the second loop, it follows that the list it points to is empty, so the heap changes by the first loop will be cleaned by the end of the second.

MatchC verifies the program in Figure 4 in about 100 milliseconds:

```
Compiling program ... DONE! [0.311s]
Loading Maude ..... DONE! [0.209s]
Verifying program ... DONE! [0.099s]
Verification succeeded! [82348 rewrites, 4 feasible and 2 infeasible paths]
Output: 5 4 3 2 1
```

We encourage the reader to run MatchC online at <http://fsl.cs.uiuc.edu/ml>.

11 Conclusions and Future Work

We have given a progress report on the rewriting logic semantics project. Our main goal has been to show how research in this area is closing the gap between theory and practice by supporting executable semantic definitions that scale up to real languages at the three levels of software modeling languages, programming languages, and HDLs, and with features such as concurrency and real-time semantics. We have also shown how such semantic definitions can be *directly* used as a basis for interpreters and for sophisticated program analysis tools, including static analyzers, model checkers, and program proving tools.

Although reasonably efficient *interpreters* can be currently generated from rewriting logic specifications, one important future challenge is the automatic generation from language definitions of high-performance language implementations that are correct by construction. Another area that should be further developed is that of *meta-reasoning* methods, to prove formal properties not about programs, but about entire language definitions. A third promising future research direction is exploring the systematic interplay between abstract semantics and model checking, as well as the systematic application of state space reduction techniques in the model checking of programs from their rewriting logic language definitions; the overall goal is achieving a high degree of *scalability* in model checking analyses, with a wide spectrum of analysis choices ranging from model checking of programs according to their concrete semantics to various forms of static analysis based on different kinds of abstract semantics.

Acknowledgments. We thank the organizers of FCT 2011 for giving us the opportunity of presenting these ideas, and for their helpful suggestions for improving the exposition. We also thank all the researchers involved in the rewriting logic semantics project for their many contributions, which we have tried to summarize in this paper without any claims of completeness. This research has been partially supported by NSF Grants CNS 08-34709, CCF 09-05584, and CCF-0916893, by NSA contract H98230-10-C-0294, by (Romanian) SMIS-CSNR 602-12516 contract no. 161/15.06.2010, by the Boeing Grant C8088, and by the Samsung SAIT grant 2010-02664.

References

1. W. Ahrendt, A. Roth, and R. Sasse. Automatic validation of transformation rules for Java verification against a rewriting semantics. In *Proc. LPAR 2006*, volume 3835 of *LNCS*, pages 412–426. Springer-Verlag, 2005.
2. M. Alba-Castro, M. Alpuente, and S. Escobar. Abstract certification of global non-interference in rewriting logic. In *Proc. FMCO*, volume 6286 of *LNCS*, pages 105–124. Springer, 2010.
3. M. Alba-Castro, M. Alpuente, and S. Escobar. Approximating non-interference and erasure in rewriting logic. In *Proc. SYNASC*, pages 124–132. IEEE, 2010.
4. M. AlTurki, D. Dhurjati, D. Yu, A. Chander, and H. Inamura. Formal specification and analysis of timing properties in software systems. In *Proc. FASE*, volume 5503 of *LNCS*, pages 262–277. Springer, 2009.
5. M. AlTurki and J. Meseguer. Real-time rewriting semantics of Orc. In *Proc. PPDP, Poland*, pages 131–142. ACM Press, 2007.
6. M. AlTurki and J. Meseguer. Reduction semantics and formal analysis of Orc programs. In *Proc. Workshop on Automated Specification and Verification of Web Systems (WWV'07)*, volume 200(3) of *ENTCS*, pages 25–41. Elsevier, 2008.
7. M. AlTurki and J. Meseguer. Dist-Orc: A rewriting-based distributed implementation of Orc with formal analysis. In *Proc. RTRTS'10*, volume 36 of *Electronic Proceedings in Theoretical Computer Science*, pages 26–45. CoRR, 2010.
8. N. Aoumeur and G. Saake. Integrating and rapid-prototyping UML structural and behavioural diagrams using rewriting logic. In *Proc. CAiSE'02*, volume 2348 of *LNCS*, pages 296–310. Springer, 2002.
9. K. Bae and P. C. Ölveczky. Extending the Real-Time Maude semantics of Ptolemy to hierarchical DE models. In *Proc. RTRTS'10*, volume 36 of *Electronic Proceedings in Theoretical Computer Science*, pages 46–66. CoRR, 2010.
10. K. Bae, P. C. Ölveczky, A. Al-Nayeem, and J. Meseguer. Synchronous AADL and its formal analysis in Real-Time Maude. Technical report, University of Illinois at Urbana-Champaign, 2005. <http://hdl.handle.net/2142/25091>.
11. K. Bae, P. C. Ölveczky, T. H. Feng, and S. Tripakis. Verifying Ptolemy II discrete-event models using Real-Time Maude. In *Proc. of ICFEM'09*, volume 5885 of *LNCS*, pages 717–736. Springer, 2009.
12. G. Berry and G. Boudol. The chemical abstract machine. *Theoretical Computer Science*, 96(1):217–248, 1992.
13. J. Björk, E. B. Johnsen, O. Owe, and R. Schlatte. Lightweight time modeling in timed Creol. In *Proc. RTRTS'10*, volume 36 of *Electronic Proceedings in Theoretical Computer Science*, pages 67–81. CoRR, 2010.
14. S. Blazy and X. Leroy. Mechanized semantics for the Clight subset of the C language. *Journal of Automated Reasoning*, 43(3):263–288, 2009.
15. A. Boronat. *MOMENT: A Formal Framework for Model Management*. PhD thesis, Universitat Politècnica de València, Spain, 2007.
16. A. Boronat, J. A. Carsí, and I. Ramos. Automatic reengineering in MDA using rewriting logic as transformation engine. In *Proc. CSMR'05*, pages 228–231. IEEE, 2005.
17. A. Boronat, R. Heckel, and J. Meseguer. Rewriting logic semantics and verification of model transformations. In *Proc. FASE'09*, volume 5503 of *LNCS*, pages 18–33. Springer, 2009.
18. A. Boronat and J. Meseguer. Algebraic semantics of OCL-constrained metamodel specifications. In *Proc. TOOLS EUROPE'09*, volume 33 of *Lecture Notes in Business Information*, pages 96–115. Springer, 2009.

19. A. Boronat and J. Meseguer. MOMENT2: EMF model transformations in Maude. In A. Vallecillo and G. Sagardui, editors, *Actas de las XIV Jornadas de Ingeniería del Software y Bases de Datos, JISBD 2009, San Sebastián, España, Septiembre 8-11, 2009*, pages 178–179, 2009.
20. A. Boronat and J. Meseguer. An algebraic semantics for MOF. *Formal Aspects of Computing*, 22(3-4):269–296, 2010.
21. A. Boronat and P. C. Ölveczky. Formal real-time model transformations in MOMENT2. In *Proc. FASE'10*, volume 6013 of *LNCS*, pages 29–43. Springer, 2010.
22. P. Borrás, D. Clément, T. Despeyroux, J. Incerpi, G. Kahn, B. Lang, and V. Pascual. CENTAUR: The system. In *Software Development Environments (SDE)*, pages 14–24, 1988.
23. C. Braga. *Rewriting Logic as a Semantic Framework for Modular Structural Operational Semantics*. PhD thesis, Departamento de Informática, Pontifícia Universidade Católica do Rio de Janeiro, Brazil, 2001.
24. C. Braga, E. H. Haesler, J. Meseguer, and P. D. Mosses. Mapping modular SOS to rewriting logic. In *Proc. LOPSTR'02*, LNCS 2664, pages 262–277, 2002.
25. C. Braga and J. Meseguer. Modular rewriting semantics in practice. In *Proc. WRLA'04*, volume 117, pages 393–416. ENTCS, Elsevier, 2004.
26. M. Broy, M. Wirsing, and P. Pepper. On the algebraic definition of programming languages. *ACM TOPLAS*, 9(1):54–99, Jan. 1987.
27. R. Bruni and J. Meseguer. Semantic foundations for generalized rewrite theories. *Theor. Comput. Sci.*, 360(1-3):386–414, 2006.
28. F. Chalub. An implementation of Modular SOS in Maude. Master's thesis, Universidade Federal Fluminense, Niterói, RJ, Brazil, May 2005.
29. F. Chalub and C. Braga. Maude MSOS tool. Universidade Federal Fluminense, www.ic.uff.br/frosario/2o-workshop-vas-novembro-2004.pdf.
30. F. Chalub and C. Braga. A Modular Rewriting Semantics for CML. *Journal of Universal Computer Science*, 10(7):789–807, 2004.
31. F. Chen, G. Roşu, and R. P. Venkatesan. Rule-based analysis of dimensional safety. In *Proc. RTA'03*, volume 2706 of *LNCS*, pages 197–207, 2003.
32. M. Clavel, F. Durán, S. Eker, J. Meseguer, P. Lincoln, N. Martí-Oliet, and C. Talcott. *All About Maude – A High-Performance Logical Framework*. Springer LNCS Vol. 4350, 2007.
33. M. Clavel and M. Egea. ITP/OCL: A rewriting-based validation tool for UML+OCL static class diagrams. In *Proc. AMAST'06*, volume 4019 of *LNCS*, pages 368–373. Springer, 2006.
34. M. Clavel and J. Santa-Cruz. ASIP + ITP: A verification tool based on algebraic semantics. In F. J. López-Fraguas, editor, *Actas de las V Jornadas sobre Programación y Lenguajes, PROLE 2005, Granada, España, Septiembre 14-16, 2005*, pages 149–158. Thomson, 2005.
35. D. Clément, J. Despeyroux, L. Hascoet, and G. Kahn. Natural semantics on the computer. In K. Fuchi and M. Nivat, editors, *Proceedings, France-Japan AI and CS Symposium*, pages 49–89. ICOT, 1986. Also, Information Processing Society of Japan, Technical Memorandum PL-86-6.
36. M. d'Amorim and G. Roşu. An Equational Specification for the Scheme Language. *Journal of Universal Computer Science*, 11(7):1327–1348, 2005. Also Technical Report No. UIUCDCS-R-2005-2567, April 2005.
37. C. Ellison and G. Roşu. A formal semantics of C with applications. Technical Report <http://hdl.handle.net/2142/17414>, University of Illinois, November 2010.

38. C. Ellison, T. F. Şerbănuță, and G. Roşu. A rewriting logic approach to type inference. In *Recent Trends in Algebraic Development Techniques*, volume 5486 of *LNCS*, pages 135–151. Springer, 2009.
39. A. Farzan. *Static and dynamic formal analysis of concurrent systems and languages: a semantics-based approach*. PhD thesis, University of Illinois at Urbana-Champaign, 2007.
40. A. Farzan, F. Cheng, J. Meseguer, and G. Roşu. Formal analysis of Java programs in JavaFAN. In *Proc. CAV'04*, volume 3114 of *LNCS*, 2004.
41. A. Farzan and J. Meseguer. Partial order reduction for rewriting semantics of programming languages. In *Proc. WRLA'06*, pages 61–78. ENTCS 176(4), 2007.
42. A. Farzan, J. Meseguer, and G. Roşu. Formal JVM code analysis in JavaFAN. in *Proc. AMAST'04*, Springer LNCS 3116, 132–147, 2004.
43. M. Felleisen and D. P. Friedman. Control operators, the SECD-machine, and the λ -calculus. In *3rd Working Conference on the Formal Description of Programming Concepts*, pages 193–219, Denmark, Aug. 1986.
44. J. L. Fernández Alemán and J. A. Toval Álvarez. Can intuition become rigorous? Foundations for UML model verification tools. In *Proc. ISSRE'00*, pages 344–355. IEEE, 2000.
45. D. P. Friedman, M. Wand, and C. T. Haynes. *Essentials of Programming Languages*. MIT Press, Cambridge, MA, 2nd edition, 2001.
46. A. Garrido, J. Meseguer, and R. Johnson. Algebraic semantics of the C preprocessor and correctness of its refactorings. Technical Report UIUCDCS-R-2006-2688, CS Dept., University of Illinois at Urbana-Champaign, February 2006.
47. J. Goguen and G. Malcolm. *Algebraic Semantics of Imperative Programs*. MIT Press, 1996.
48. J. Goguen and J. Meseguer. Security policies and security models. In *Proceedings of the 1982 Symposium on Security and Privacy*, pages 11–20. IEEE, 1982.
49. J. A. Goguen and K. Parsaye-Ghomi. Algebraic denotational semantics using parameterized abstract modules. In *Formalizing Programming Concepts*, pages 292–309. Springer-Verlag, 1981. LNCS, Volume 107.
50. Y. Gurevich. Evolving algebras 1993: Lipari Guide. In E. Börger, editor, *Specification and Validation Methods*, pages 9–37. Oxford University Press, 1994.
51. Y. Gurevich and J. K. Huggins. The semantics of the C programming language. In *Computer Science Logic*, volume 702 of *LNCS*, pages 274–308, 1993.
52. R. Harper, F. Honsell, and G. D. Plotkin. A framework for defining logics. *Journal of the ACM*, 40(1):143–184, 1993.
53. M. Hills, F. Chen, and G. Roşu. Pluggable Policies for C. Technical Report UIUCDCS-R-2008-2931, University of Illinois at Urbana-Champaign, 2008.
54. M. Hills, T. F. Şerbănuță, and G. Roşu. A rewrite framework for language definitions and for generation of efficient interpreters. In *Proc. of WRLA'06*, volume 176(4) of *ENTCS*, pages 215–231. Elsevier, 2007.
55. E. B. Johnsen, O. Owe, and E. W. Axelsen. A runtime environment for concurrent objects with asynchronous method calls. In *Proc. WRLA'04*, volume 117. ENTCS, Elsevier, 2004.
56. G. Kahn. Natural semantics. In *Proc. STACS'87*, volume 247 of *LNCS*, pages 22–39. Springer, 1987.
57. M. Katelman, S. Keller, and J. Meseguer. Concurrent rewriting semantics and analysis of asynchronous digital circuits. In *Proc. WRLA'10*, volume 6381 of *LNCS*, pages 140–156. Springer, 2010.

58. M. Katelman and J. Meseguer. A rewriting semantics for ABEL with applications to hardware/software co-design and analysis. In *Proc. WRLA'06*, pages 47–60. ENTCS 176(4), Elsevier, 2007.
59. M. Katelman and J. Meseguer. *vlogs1*: A Strategy Language for Simulation-Based Verification of Hardware. In *Proc. HVC'10*, volume 6504 of *LNCS*, pages 129 – 145. Springer Berlin / Heidelberg, 2011.
60. M. Katelman, J. Meseguer, and S. Escobar. Directed-logical testing for functional verification of microprocessors. In *MEMOCODE'08*, pages 89–100. IEEE, 2008.
61. M. Kaufmann, P. Manolios, and J. S. Moore. *Computer-Aided Reasoning: ACL2 Case Studies*. Kluwer Academic Press, 2000.
62. A. Knapp. Generating rewrite theories from UML collaborations. In K. Futatsugi, A. T. Nakagawa, and T. Tamai, editors, *Cafe: An Industrial-Strength Algebraic Formal Method*, pages 97–120. Elsevier, 2000.
63. A. Knapp. *A Formal Approach to Object-Oriented Software Engineering*. Shaker Verlag, Aachen, Germany, 2001. PhD thesis, Institut für Informatik, Universität München, 2000.
64. E. A. Lee. Modeling concurrent real-time processes using discrete events. *Ann. Software Eng.*, 7:25–45, 1999.
65. S. Liang, P. Hudak, and M. Jones. Monad transformers and modular interpreters. In *Proc. POPL'95*, pages 333–343. ACM Press, 1995.
66. N. Martí-Oliet and J. Meseguer. Rewriting logic as a logical and semantic framework. In D. M. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic, Second Edition, Volume 9*, pages 1–87. Kluwer, 2002.
67. P. Meredith, M. Hills, and G. Roşu. A K definition of Scheme. Technical Report UIUCDCS-R-2007-2907, Department of Computer Science, University of Illinois at Urbana-Champaign, 2007.
68. P. Meredith, M. Katelman, J. Meseguer, and G. Roşu. A formal executable semantics of Verilog. In *Proc. MEMOCODE'10*, pages 179–188. IEEE, 2010.
69. J. Meseguer. Conditional rewriting logic as a unified model of concurrency. *Theoretical Computer Science*, 96(1):73–155, 1992.
70. J. Meseguer. Rewriting logic as a semantic framework for concurrency: a progress report. In *Proc. CONCUR'96*, pages 331–372. Springer LNCS 1119, 1996.
71. J. Meseguer. Membership algebra as a logical framework for equational specification. In *Proc. WADT'97*, pages 18–61. LNCS 1376, 1998.
72. J. Meseguer. Software specification and verification in rewriting logic. In M. Broy and M. Pizka, editors, *Models, Algebras, and Logic of Engineering Software, NATO Advanced Study Institute, Marktobendorf, Germany, July 30 – August 11, 2002*, pages 133–193. IOS Press, 2003.
73. J. Meseguer and C. Braga. Modular rewriting semantics of programming languages. in *Proc. AMAST'04*, Springer LNCS 3116, 364–378, 2004.
74. J. Meseguer, K. Futatsugi, and T. Winkler. Using rewriting logic to specify, program, integrate, and reuse open concurrent systems of cooperating agents. In *Proceedings of the 1992 International Symposium on New Models for Software Architecture, Tokyo, Japan, November 1992*, pages 61–106. Research Institute of Software Engineering, 1992.
75. J. Meseguer and P. C. Ölveczky. Formalization and correctness of the PALS architectural pattern for distributed real-time systems. In *Proc. ICFEM'10*, volume 6447 of *LNCS*, pages 303–320, 2010.
76. J. Meseguer and G. Roşu. The rewriting logic semantics project. *Theoretical Computer Science*, 373:213–237, 2007.

77. J. Meseguer and G. Roşu. Rewriting logic semantics: From language specifications to formal analysis tools. In *Proc. IJCAR'04*, volume 3097 of *LNAI*, pages 1–44. Springer, 2004.
78. D. Miller. Representing and reasoning with operational semantics. In *Proc. IJCAR'06*, volume 4130 of *LNCS*, pages 4–20, 2006.
79. J. Misra. Computation orchestration: A basis for wide-area computing. In M. Broy, editor, *Proc. of the NATO Advanced Study Institute, Engineering Theories of Software Intensive Systems Marktoberdorf, Germany, 2004*. NATO ASI Series, 2004.
80. J. Misra and W. R. Cook. Computation orchestration. *Software and System Modeling*, 6(1):83–110, 2007.
81. E. Moggi. An abstract view of programming languages. Technical Report ECS-LFCS-90-113, Edinburgh University, Dept. of Computer Science, June 1989.
82. F. Mokhati and M. Badri. Generating Maude specifications from UML use case diagrams. *Journal of Object Technology*, 8(2):319–336, 2009.
83. F. Mokhati, P. Gagnon, and M. Badri. Verifying UML diagrams with model checking: A rewriting logic based approach. In *Proc. QSIC'07*, pages 356–362. IEEE, 2007.
84. F. Mokhati, B. Sahraoui, S. Bouzaher, and M. T. Kimour. A tool for specifying and validating agents' interaction protocols: From Agent UML to Maude. *Journal of Object Technology*, 9(3):59–77, 2010.
85. P. D. Mosses. Unified algebras and action semantics. In *Proc. STACS'89*, pages 17–35. Springer LNCS 349, 1989.
86. P. D. Mosses. Denotational semantics. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science, Vol. B, Chapter 11*. North-Holland, 1990.
87. P. D. Mosses. Modular structural operational semantics. *J. Log. Algebr. Program.*, 60–61:195–228, 2004.
88. G. Nadathur and D. Miller. An overview of λ Prolog. In K. Bowen and R. Kowalski, editors, *Fifth Int. Joint Conf. and Symp. on Logic Programming*, pages 810–827. The MIT Press, 1988.
89. S. Nakajima. Using algebraic specification techniques in development of object-oriented frameworks. In *Proc. FM'99*, volume 1709 of *LNCS*, pages 1664–1683. Springer, 1999.
90. S. Nakajima and K. Futatsugi. An object-oriented modeling method for algebraic specifications in CafeOBJ. In *Proc. ICSE'97*. ACM, 1997.
91. M. Norrish. C formalised in HOL. Technical Report UCAM-CL-TR-453, University of Cambridge, December 1998.
92. P. C. Ölveczky, A. Boronat, and J. Meseguer. Formal semantics and analysis of behavioral AADL models in Real-Time Maude. In *Proc. FMOODS'10*, volume 6117 of *LNCS*, pages 47–62. Springer, 2010.
93. P. C. Ölveczky and J. Meseguer. Specification of real-time and hybrid systems in rewriting logic. *Theoretical Computer Science*, 285(2):359–405, 2002.
94. P. C. Ölveczky and J. Meseguer. Semantics and pragmatics of Real-Time Maude. *Higher-Order and Symbolic Computation*, 20(1-2):161–196, 2007.
95. N. S. Papaspyrou. *A Formal Semantics for the C Programming Language*. PhD thesis, National Technical University of Athens, February 1998.
96. N. S. Papaspyrou. Denotational semantics of ANSI C. *Computer Standards and Interfaces*, 23(3):169–185, 2001.
97. F. Pfenning and C. Elliott. Higher-order abstract syntax. In *Proc. PLDI'88*, pages 199–208. ACM Press, 1988.

98. B. Pierce. *Types and Programming Languages*. MIT Press, 2002.
99. G. D. Plotkin. A structural approach to operational semantics. *Journal of Logic and Algebraic Programming*, 60-61:17-139, 2004. Previously published as technical report DAIMI FN-19, Aarhus University, 1981.
100. J. E. Rivera, F. Durán, and A. Vallecillo. On the behavioral semantics of real-time domain specific visual languages. In *Proc. WRLA'10*, volume 6381 of *LNCS*, pages 174-190. Springer, 2010.
101. G. Roşu. CS322, Fall 2003 - Programming Language Design: Lecture Notes. Technical Report UIUCDCS-R-2003-2897, University of Illinois at Urbana-Champaign, Dept. of Computer Science, Dec. 2003. Notes of a course taught at UIUC.
102. G. Roşu and A. Ştefănescu. Matching logic: A new program verification approach (nier track). In *Proc. ICSE'11*, 2011.
103. G. Roşu, C. Ellison, and W. Schulte. Matching logic: An alternative to Hoare/Floyd logic. In *Proc. AMAST'10*, pages 142-162. LNCS 6486, 2010.
104. G. Roşu, R. P. Venkatesan, J. Whittle, and L. Leustean. Certifying optimality of state estimation programs. In *Proc. CAV'03*, pages 301-314. Springer, 2003. LNCS 2725.
105. G. Roşu and T. F. Şerbănuţă. An overview of the K semantic framework. *Journal of Logic and Algebraic Programming*, 79(6):397-434, 2010.
106. R. Sasse. Taclets vs. rewriting logic – relating semantics of Java. Master's thesis, Fakultät für Informatik, Universität Karlsruhe, Germany, May 2005. Technical Report in Computing Science No. 2005-16.
107. R. Sasse and J. Meseguer. Java+ITP: A verification tool based on hoare logic and algebraic semantics. In *Proc. WRLA'06*, pages 29-46. ENTCS 176(4), 2007.
108. D. A. Schmidt. *Denotational Semantics – A Methodology for Language Development*. Allyn and Bacon, Boston, MA, 1986.
109. D. Scott. Outline of a mathematical theory of computation. In *Proceedings, Fourth Annual Princeton Conference on Information Sciences and Systems*, pages 169-176. Princeton University, 1970. Also appeared as Technical Monograph PRG 2, Oxford University, Programming Research Group.
110. D. Scott and C. Strachey. Toward a mathematical semantics for computer languages. In *Microwave Research Institute Symposia Series, Vol. 21: Proc. Symp. on Computers and Automata*. Polytechnical Institute of Brooklyn, 1971.
111. T. F. Şerbănuţă. *A Rewriting Approach to Concurrent Programming Language Design and Semantics*. PhD thesis, Department of Computer Science, University of Illinois at Urbana-Champaign, 2010.
112. T. F. Şerbănuţă, G. Roşu, and J. Meseguer. A rewriting logic approach to operational semantics. *Information and Computation*, 207(2):305-340, 2009.
113. K. Slonnegger and B. L. Kurtz. *Formal Syntax and Semantics of Programming Languages*. Addison-Wesley, 1995.
114. R. F. Stärk, J. Schmid, and E. Börger. *Java and the Java Virtual Machine: Definition, Verification, Validation*. Springer, 2001.
115. M.-O. Stehr and C. Talcott. PLAN in Maude: Specifying an active network programming language. In *Proc. WRLA'02*, volume 117. ENTCS, Elsevier, 2002.
116. M.-O. Stehr and C. L. Talcott. Practical techniques for language design and prototyping. In *Dagstuhl Seminar 05081 on Foundations of Global Computing. February 20 - 25, 2005. Schloss Dagstuhl, Wadern, Germany.*, 2005.
117. P. Thati, K. Sen, and N. Martí-Oliet. An executable specification of asynchronous Pi-Calculus semantics and may testing in Maude 2.0. In *Proc. WRLA'02*. ENTCS, Elsevier, 2002.

118. A. van Deursen, J. Heering, and P. Klint. *Language Prototyping: An Algebraic Specification Approach*. World Scientific, 1996.
119. A. Verdejo. *Maude como marco semántico ejecutable*. PhD thesis, Facultad de Informática, Universidad Complutense, Madrid, Spain, 2003.
120. A. Verdejo and N. Martí-Oliet. Implementing CCS in Maude 2. In *Proc. WRLA'02*. ENTCS, Elsevier, 2002.
121. A. Verdejo and N. Martí-Oliet. Two case studies of semantics execution in Maude: CCS and LOTOS. *Formal Methods in System Design*, 27(1-2):113–172, 2005.
122. A. Verdejo and N. Martí-Oliet. Executable structural operational semantics in Maude. *Journal of Logic and Algebraic Programming*, 67(1-2):226–293, 2006.
123. P. Viry. Equational rules for rewriting logic. *Theoretical Computer Science*, 285:487–517, 2002.
124. P. Wadler. The essence of functional programming. In *Proc. POPL '92*, pages 1–14, New York, NY, USA, 1992. ACM Press.
125. M. Wand. First-order identities as a defining language. *Acta Informatica*, 14:337–357, 1980.
126. I. Wehrman, D. Kitchin, W. R. Cook, and J. Misra. A timed semantics of Orc. *Theor. Comput. Sci.*, 402(2-3):234–248, 2008.
127. M. Wirsing and A. Knapp. A formal approach to object-oriented software engineering. In *Proc. WRLA'96*, volume 4 of *ENTCS*, pages 322–360, 1996.
128. M. Wirsing and A. Knapp. A formal approach to object-oriented software engineering. *Theoretical Computer Science*, 285(2):519–560, 2002.
129. A. K. Wright and M. Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, 1994.