

Computational Logical Frameworks and Generic Program Analysis Technologies

José Meseguer and Grigore Roşu

{meseguer,grosu}@cs.uiuc.edu
Department of Computer Science
University of Illinois at Urbana-Champaign

1 Motivation

The technologies developed to solve the verifying compiler grand challenge should be *generic*, that is, not tied to a particular language but widely applicable to many languages. Such technologies should also be *semantics-based*, that is, based on a rigorous formal semantics of the languages.

For this, a *computational logical framework* with efficient executability and a spectrum of *meta-tools* can serve as a basis on which to: (1) define the formal semantics of any programming language; and (2) develop *generic program analysis techniques and tools* that can be instantiated to generate powerful analysis tools for each language of interest.

Not all logical frameworks can serve such purposes well. We first list some specific requirements that we think are important to properly address the grand challenge. Then we present our experience with rewriting logic as supported by the Maude system and its formal tool environment. Finally, we discuss some future directions of research.

2 Logical Framework Requirements

Based on experience, current trends, and the basic requirements of the grand challenge problem, we believe that any logical framework serving as a computational infrastructure for the various technologies for solving the grand challenge should have *at least* the following features:

1. good *data representation* capabilities,
2. support for *concurrency and nondeterminism*,
3. *simplicity* of the formalism,

4. *efficient* implementability, and efficient *meta-tools*,
5. support for *reflection*,
6. support for inductive reasoning, preferably with *initial model* semantics,
7. support for generation of *proof objects*, acting as correctness certificates.

While proponents of a framework may claim that it has all these features, in some cases further analysis can show that it either lacks some of them, or can only “simulate” certain features in a quite artificial way. A good example is the simulation/elimination of concurrency in inherently deterministic formalisms by implementing or defining thread/process scheduling algorithms. Another example might be the claim that the lambda calculus has good data representation capabilities because one can encode numbers as Church numerals.

3 The Rewriting Logic/Maude Experience

At UIUC, together with several students, we are developing semantic definitions of programming languages based on *rewriting logic* (RWL) [27]. Rewriting logic meets the requirements mentioned above, and supports semantic definitions of programming languages that combine *algebraic denotational semantics* and *SOS* semantics in a seamless way [29]. Given a language L , its rewriting logic semantics is a rewrite theory

$$\mathcal{R}_L = (\Sigma_L, E_L, R_L),$$

where Σ_L is a *signature* expressing the syntax of L , E_L is a set of *equations* defining the meaning of the sequential features of L together with that of the various state infrastructure operations, and R_L is a set of *rewrite rules* defining the semantics of the concurrent features of L .

3.1 Maude and its Formal Tools

Rewrite theories are triples (Σ, E, R) , with (Σ, E) an equational theory and R a set of rewrite rules. Intuitively, (Σ, E, R) specifies a computational system in which the *states* are specified as elements of the algebraic data type defined by (Σ, E) , and the system’s *concurrent transitions* are specified by the rewrite rules R . Rewrite theories can be executed in different languages such as CafeOBJ [22], and ELAN [1]. The most general support for the execution of rewrite theories is currently provided by the Maude language [6, 7], in which rewrite theories with very general conditional rules, and whose underlying equational theories can be membership equational theories [28], can be specified and can be executed, provided they satisfy some basic executability requirements. Furthermore, Maude provides very efficient support for rewriting *modulo* any combination of associativity, commutativity, and identity axioms. Since an equational theory (Σ, E) can be regarded as a degenerate rewrite theory of the form (Σ, E, \emptyset) , equational

logic is naturally a sublogic of rewriting logic. In Maude this sublogic is supported by *functional modules* [6], which are theories in membership equational logic. When executed in Maude, the RWL formal semantics \mathcal{R}_L of language L automatically becomes an *efficient interpreter* for L : for example, faster than the Linux bc interpreter, and half the speed of the Scheme interpreter.

Besides supporting efficient execution, typically in the order of several million rewrites per second, Maude also provides a range of formal tools and algorithms to analyze rewrite theories and verify their properties. These tools can be used almost directly to provide corresponding analysis tools for languages defined as rewrite logic theories. A first very useful formal analysis feature is its *breadth-first search* command. Given an initial state of a system (a term), we can search for all reachable states matching a certain pattern and satisfying an equationally-defined semantic condition P . By making $P = \neg Q$, where Q is an invariant, we get in this way a *semi-decision procedure* for finding failures of invariant safety properties. Note that there is no finite-state assumption involved here: any executable rewrite theory can thus be analyzed. For systems where the set of states reachable from an initial state are finite, Maude also provides a linear time temporal logic (LTL) model checker. Maude's is an explicit-state LTL model checker, with performance comparable to that of the SPIN model checker [24] for the benchmarks that we have analyzed [17, 18].

Reflection is a key feature of rewriting logic, and is efficiently supported in the Maude implementation through its `META-LEVEL` module. One important fruit of this is that it becomes quite easy to build new formal tools and to add them to the Maude environment. Indeed, such tools by their very nature manipulate and analyze rewrite theories. By reflection, a rewrite theory \mathcal{R} becomes a *term* $\overline{\mathcal{R}}$ in the universal theory, which can be efficiently manipulated by the descent functions in the `META-LEVEL` module. As a consequence, Maude formal tools have a reflective design and are built in Maude as suitable extensions of the `META-LEVEL` module. They include the following:

- the Maude Church-Rosser Checker, and Knuth-Bendix and Coherence Completion tools [8, 15, 13, 12]
- the Full Maude module composition tool [11, 16]
- the Maude Predicate Abstraction tool [34]
- the Maude Inductive Theorem Prover (ITP) [5, 8, 9]
- the Real-Time Maude tool [33];
- the Maude Sufficient Completeness Checker (SCC) [23]
- the Maude Termination Tool (MTT) [14].

3.2 Unifying SOS and Equational Semantics

For the most part, equational semantics¹ and SOS have lived separate lives. Pragmatic considerations and differences in taste tend to dictate which framework is adopted in each particular case. For concurrent languages SOS is clearly superior and tends to prevail as the formalism of choice, but for deterministic languages equational approaches are also widely used. Of course there are also practical considerations of tool support for both execution and formal reasoning.

In the end, equational semantics and SOS, although each very valuable in its own way, are “single hammer” approaches. Would it be possible to seamlessly *unify* them within a more flexible and general framework? Could their respective limitations be overcome when they are thus unified? Our proposal is that rewriting logic does indeed provide one such unifying framework. The key to this unification is what we call rewriting logic’s *abstraction knob*. The point is that in equational semantics’ model-theoretic approach entities are *identified by the semantic equations*, and have unique *abstract denotations* in the corresponding models. In our knob metaphor this means that in equational semantics the abstraction knob is *always turned all the way up to its maximum position*. By contrast, one of the key features of SOS is providing a very detailed, step-by-step formal description of a language’s evaluation mechanisms. As a consequence, most entities —except perhaps for built-in data, stores, and environments, which are typically treated on the side— are *primarily syntactic*, and computations are described in full detail. In our metaphor this means that in SOS the abstraction knob is *always turned down to its minimum position*.

How is the unification and corresponding availability of an abstraction knob achieved? Since a rewrite theory is a triple (Σ, E, R) , with (Σ, E) an equational theory with Σ a signature of operations and sorts, and E a set of (possibly conditional) equations, and with R a set of (possibly conditional) rewrite rules, equational semantics is obtained as the special case in which $R = \emptyset$, so we only have the semantic equations E and the abstraction knob is turned up to its maximum position. SOS is obtained as the special case in which $E = \emptyset$, and we only have (possibly conditional) rules R rewriting purely syntactic entities (terms), so that the abstraction knob is turned down to the minimum position.

Rewriting logic’s “abstraction knob” is precisely its crucial distinction between equations E and rules R in a rewrite theory (Σ, E, R) . *States of the computation* are then E -equivalence classes, that is, *abstract elements* in the initial algebra $T_{\Sigma/E}$. A rewrite with a rule in R is understood as a transition $[t] \longrightarrow [t']$ between such abstract states. The knob, however, can be turned up or down. We can turn it *all the way down to its minimum* by converting all equations into rules, transforming (Σ, E, R) into $(\Sigma, \emptyset, R \cup E)$. This gives us the most concrete, SOS-like semantic description possible. Can we turn the knob

¹In equational semantics, formal definitions take the form of *semantic equations*, typically satisfying the *Church-Rosser* property. Both higher-order (denotational semantics) and first-order (algebraic semantics) versions have been shown to be useful formalisms. We use the more neutral term *equational semantics* to emphasize the fact that denotational and algebraic semantics have many common features and can both be viewed as instances of a common equational framework.

“all the way up,” in the sense of converting all rules into equations? Only if the system we are describing is *deterministic* (for example, the semantic definition of a sequential language) is this a sound procedure. In that case, the equational theory $(\Sigma, R \cup E)$ should be Church-Rosser, and we do indeed obtain a most-abstract-possible, purely equational semantics out of the less abstract specification (Σ, E, R) , or even out of the most concrete possible specification $(\Sigma, \emptyset, R \cup E)$. What can we do in general to make a specification *as abstract as possible*? We can identify a subset $R_0 \subseteq R$ such that: (1) $R_0 \cup E$ is Church-Rosser; and (2) R_0 is biggest possible with this property. In actual language specification practice this is not hard to do. Essentially, we can use semantic equations for most of the sequential features of a programming language: only when interactions with memory could lead to nondeterminism (particularly if the language has threads, or they could later be added to the language in an extension) or for intrinsically concurrent features, are rules (as opposed to equations) really needed. In our experience, it is often possible to specify most of the semantic axioms with equations, with relatively few rules needed for truly concurrent or nondeterministic features. For example, the semantics of the JVM described in [21, 19] has about 300 equations and 40 rules; and that of Java described in [19] has about 600 equations but only 15 rules. A semantics for an ML-like language with threads given in [30] has only two rules.

This distinction between equations and rules, besides giving to equational semantics and SOS their due in a way not possible for the other alternative if we were to remain within each of these formalisms, has also important practical consequences for *program analysis*; because it affords a massive *state space reduction* which can make formal analyses such as breadth-first search and model checking enormously more efficient. Because of state-space explosion, such analyses could easily become infeasible if we were to use an SOS-like specification in which all computation steps are described with rules. This capacity of dealing with abstract states is a crucial reason why our generic tools, when instantiated to a given programming language definition, tend to result in program analysis tools of competitive performance. Of course, the price to pay in exchange for abstraction is a *coarser level of granularity* in respect to what aspects of a computation are *observable* at that abstraction level. For example, when analyzing a sequential program using a semantics in which most sequential features have been specified with equations, all sequential subcomputations will be abstracted away, and the analysis will focus on memory and thread interactions. If a finer analysis is needed, we can always obtain it by “turning down the abstraction knob” to the right observability level by *converting some equations into rules*. That is, we can regulate the knob to find for each kind of analysis the best possible balance between abstraction and observability.

3.3 Languages Defined in Rewriting Logic

Many languages have already been given semantics in this way using Maude. The language definitions can then be used as interpreters, and—in conjunction with Maude’s search command and its LTL model checker—to formally analyze

programs in those languages. For example, large fragments of Java and the JVM have been specified in Maude this way, with the Maude rewriting logic semantics being used as the basis of Java and JVM program analysis tools that for some examples outperform well-known Java analysis tools [21, 19]. A similar Maude specification of the semantics of Scheme at UIUC yields an interpreter with .75 the speed of the standard Scheme interpreter on average for the benchmarks tested. The specification of a C-like language and the corresponding formal analyses are discussed in detail in [31]. A semantics of an ML-like language with threads was discussed in detail in [30], a modular rewriting logic semantics of CML has been given in [4], and a definition of the Scheme language has been given in [10]. Other language case studies, all specified in Maude, include: BC [2], CCS [43, 44, 2], CIAO [40], Creol [25], ELOTOS [42], MSR [3, 38], PLAN [39, 40], and the pi-calculus [41]. In fact, the semantics of large fragments of conventional languages are by now routinely developed by UIUC graduate students as course projects [36] in a few weeks, including, besides the languages already mentioned: Beta, Haskell, Lisp, LLVM, Pict, Python, Ruby, and Smalltalk.

3.4 Formal Analysis

Furthermore, Maude’s formal tools, such as its inductive theorem prover, linear temporal logic (LTL) model-checker, and breadth-first search (BFS) capability then become *meta-tools* from which we derive useful program analysis tools for L using \mathcal{R}_L .

We are furthermore developing new *generic program analysis technologies* such as, for example, a *generic partial-order reduction* technique [20] than can apply to any language L with threads, and does not require any changes to an underlying model checker.

Correctness of a compiler can and should mean more than just correctness with respect to functional behavior. Depending upon particular applications of interest, certain important safety policies that transcend the basic semantics of the language under consideration may need to be preserved. For example, in an application referring to physical objects, consistency with respect to units of measurement or coordinate systems needs to be assured. We are also developing *domain-specific certifiers*, which are static analysis tools that check conformance of computation with respect to particular but important domains of interest. For example, we developed RWL-based certifiers for conformance with units of measurement [37], and with coordinate frames [26].

The *cost* of generating tools for a language L this way using its formal semantic definition \mathcal{R}_L is *much lower* (in the order of weeks) than that of building similar language specific analysis tools (man years). For example, it took Feng Chen at UIUC only three weeks to develop the formal semantics of a large subset of Java —including multithreading, inheritance, polymorphism, object references, and dynamic object allocation— as a RWL theory \mathcal{R}_{Java} specified in Maude.

Furthermore, the *formal analysis tools* obtained for free from \mathcal{R}_{Java} and \mathcal{R}_{JVM} are *competitive* for some applications with similar language-specific tools

such a NASA-Ames' Java Path Finder [45] and Stanford's Java model checker [35]. Similarly, our experiments with the generic partial order reduction technique indicate that it can achieve rates of space and time reduction similar to those of language-specific tools such as SPIN [24].

4 Future Directions Related to the Grand Challenge

Our main point has been to emphasize the need for genericity in approaching the grand challenge; otherwise, an answer to the challenge would have a limited applicability to other languages besides those chosen in the challenge project. For this, we have argued that both a computational logical framework in which to give a precise formal semantics to programming languages, and on which to base generic program analysis tools, would be very useful.

We have also summarized our experience so far with one such logical framework, namely rewriting logic, and for applying the Maude RWL language and its generic tools to formally analyze programs in different programming languages. Our results, although encouraging, are very much *work in progress*; we would like to advance in addition the following directions:

1. Modular programming language definitions in the spirit of MSOS [32]. The goal is to build a database of reusable semantic definitions, with the semantics of each language feature defined in a separate module. It should then be possible to define the semantics of a whole language by just combining the modules for the language features, renaming the syntax of each module to the chosen concrete syntax.
2. Developing various language-generic program analysis tools; we have already mentioned the ongoing work on partial order reduction, which should be further advanced; but generic abstraction tools, and also generic tools for static analysis are to other important areas to advance.
3. Language-generic theorem proving environments, based on an axiomatic semantics that uses the language rewriting semantics as its foundation are also an important direction to investigate.
4. Finally, it would be very useful to investigate semantics-preserving translations between languages, and in particular the generation of provably correct compilers from the formal semantics \mathcal{R}_L of a language L .

References

- [1] P. Borovanský, C. Kirchner, H. Kirchner, and P.-E. Moreau. ELAN from a rewriting logic point of view. *Theoretical Computer Science*, 285:155–185, 2002.
- [2] C. Braga and J. Meseguer. Modular rewriting semantics in practice. in Proc. *WRLA'04*, ENTCS.

- [3] I. Cervesato and M.-O. Stehr. Representing the msr cryptoprotocol specification language in an extension of rewriting logic with dependent types. In P. Degano, editor, *Proc. Fifth International Workshop on Rewriting Logic and its Applications (WRLA'2004)*. Elsevier ENTCS, 2004. Barcelona, Spain, March 27 - 28, 2004.
- [4] F. Chalub and C. Braga. A Modular Rewriting Semantics for CML. *Journal of Universal Computer Science*, 10(7):789–807, July 2004. http://www.jucs.org/jucs_10_7/a_modular_rewriting_semantics.
- [5] M. Clavel. *Reflection in Rewriting Logic: Metalogical Foundations and Metaprogramming Applications*. CSLI Publications, 2000.
- [6] M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer, and J. Quesada. Maude: specification and programming in rewriting logic. *Theoretical Computer Science*, 285:187–243, 2002.
- [7] M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer, and C. Talcott. Maude 2.0 Manual. June 2003, <http://maude.cs.uiuc.edu>.
- [8] M. Clavel, F. Durán, S. Eker, and J. Meseguer. Building equational proving tools by reflection in rewriting logic. In *CAFE: An Industrial-Strength Algebraic Formal Method*. Elsevier, 2000. <http://maude.cs.uiuc.edu>.
- [9] M. Clavel and M. Palomino. The ITP tool's manual. Universidad Complutense, Madrid, April 2005, <http://maude.sip.ucm.es/itp/>.
- [10] M. d'Amorim and G. Roşu. An Equational Specification for the Scheme Language. In *Proceedings of the 9th Brazilian Symposium on Programming Languages (SBLP'05)*, to appear 2005. Also Technical Report No. UIUCDCS-R-2005-2567, April 2005.
- [11] F. Durán. A reflective module algebra with applications to the Maude language. Ph.D. Thesis, University of Málaga, 1999.
- [12] F. Durán. Coherence checker and completion tools for Maude specifications. Manuscript, Computer Science Laboratory, SRI International, <http://maude.cs.uiuc.edu/papers>, 2000.
- [13] F. Durán. Termination checker and Knuth-Bendix completion tools for Maude equational specifications. Manuscript, Computer Science Laboratory, SRI International, <http://maude.cs.uiuc.edu/papers>, 2000.
- [14] F. Durán, S. Lucas, J. Meseguer, C. Marché, and X. Urbain. Proving termination of membership equational programs. In P. Sestoft and N. Heintze, editors, *Proc. of ACM SIGPLAN 2004 Symposium on Partial Evaluation and Program Manipulation, PEPM'04*, pages 147–158. ACM Press, 2004.
- [15] F. Durán and J. Meseguer. A Church-Rosser checker tool for Maude equational specifications. Manuscript, Computer Science Laboratory, SRI International, <http://maude.cs.uiuc.edu/papers>, 2000.
- [16] F. Durán and J. Meseguer. On parameterized theories and views in Full Maude 2.0. In K. Futatsugi, editor, *Proc. 3rd. Intl. Workshop on Rewriting Logic and its Applications*. ENTCS, Elsevier, 2000.
- [17] S. Eker, J. Meseguer, and A. Sridharanarayanan. The Maude LTL model checker. In F. Gadducci and U. Montanari, editors, *Proc. 4th. Intl. Workshop on Rewriting Logic and its Applications*. ENTCS, Elsevier, 2002.

- [18] S. Eker, J. Meseguer, and A. Sridharanarayanan. The Maude LTL model checker and its implementation. In *Model Checking Software: Proc. 10th Intl. SPIN Workshop*, volume 2648, pages 230–234. Springer LNCS, 2003.
- [19] A. Farzan, F. Cheng, J. Meseguer, and G. Roşu. Formal analysis of Java programs in JavaFAN. in Proc. CAV’04, Springer LNCS, 2004.
- [20] A. Farzan and J. Meseguer. Partial order reduction for rewriting semantics of programming languages, 2005. Manuscript, submitted for publication.
- [21] A. Farzan, J. Meseguer, and G. Roşu. Formal JVM code analysis in JavaFAN. in Proc. *AMAST’04*, Springer LNCS 3116, 132–147, 2004.
- [22] K. Futatsugi and R. Diaconescu. *CafeOBJ Report*. World Scientific, AMAST Series, 1998.
- [23] J. Hendrix, J. Meseguer, and M. Clavel. A sufficient completeness reasoning tool for partial specifications. In *Proc. RTA’05*, volume 3467 of LNCS, 2005.
- [24] G. Holzmann. *The Spin Model Checker - Primer and Reference Manual*. Addison-Wesley, 2003.
- [25] E. B. Johnsen, O. Owe, and E. W. Axelsen. A runtime environment for concurrent objects with asynchronous method calls. In N. Martí-Oliet, editor, *Proc. 5th. Intl. Workshop on Rewriting Logic and its Applications*. ENTCS, Elsevier, 2004.
- [26] M. Lowry, T. Pressburger, and G. Roşu. Certifying domain-specific policies. In *Proceedings, International Conference on Automated Software Engineering (ASE’01)*, pages 81–90. IEEE, 2001. San Diego, California.
- [27] J. Meseguer. Conditional rewriting logic as a unified model of concurrency. *Theoretical Computer Science*, 96(1):73–155, 1992.
- [28] J. Meseguer. Membership algebra as a logical framework for equational specification. In F. Parisi-Presicce, editor, *Proc. WADT’97*, pages 18–61. Springer LNCS 1376, 1998.
- [29] J. Meseguer and G. Roşu. Rewriting logic semantics: From language specifications to formal analysis tools. In *International Joint Conference on Automated Reasoning (IJCAR’04)*, volume 3097 of *Lecture Notes in Computer Science*, pages 1–44, 2004.
- [30] J. Meseguer and G. Roşu. Rewriting logic semantics: From language specifications to formal analysis tools. In *Proc. Intl. Joint Conf. on Automated Reasoning IJCAR’04, Cork, Ireland, July 2004*, pages 1–44. Springer LNAI 3097, 2004.
- [31] J. Meseguer and G. Roşu. The rewriting logic semantics project. In *Proc. of SOS 2005*, ENTCS, to appear. Elsevier, 2005.
- [32] P. D. Mosses. Foundations of modular SOS. In *Proceedings of MFCS’99, 24th International Symposium on Mathematical Foundations of Computer Science*, pages 70–80. Springer LNCS 1672, 1999.
- [33] P. C. Ölveczky and J. Meseguer. Real-Time Maude 2.1. In N. Martí-Oliet, editor, *Proc. 5th Intl. Workshop on Rewriting Logic and its Applications*. ENTCS, Elsevier, 2004.
- [34] M. Palomino. A predicate abstraction tool for Maude. Documentation and tool available at <http://maude.sip.ucm.es/~miguelpt/bibliography.html>.
- [35] D. Y. W. Park, U. Stern, J. U. Skakkebak, and D. L. Dill. Java model checking. In *ASE’01*, pages 253 – 256, 2000.

- [36] G. Roşu. Programming language classes. Department of Computer Science, University of Illinois at Urbana-Champaign, <http://fsl.cs.uiuc.edu/~grosu/classes/>.
- [37] G. Roşu and F. Chen. Certifying measurement unit safety policy. In *Automated Software Engineering, 2003. Proc. 18th IEEE Intl. Conference*, pages 304–309, 2003.
- [38] M.-O. Stehr, I. Cervesato, and S. Reich. An execution environment for the MSR cryptoprotocol specification language. <http://formal.cs.uiuc.edu/stehr/msr.html>.
- [39] M.-O. Stehr and C. Talcott. PLAN in Maude: Specifying an active network programming language. In F. Gadducci and U. Montanari, editors, *Proc. 4th. Intl. Workshop on Rewriting Logic and its Applications*. ENTCS, Elsevier, 2002.
- [40] M.-O. Stehr and C. L. Talcott. Practical techniques for language design and prototyping. In J. L. Fiadeiro, U. Montanari, and M. Wirsing, editors, *Abstracts Collection of the Dagstuhl Seminar 05081 on Foundations of Global Computing. February 20 – 25, 2005. Schloss Dagstuhl, Wadern, Germany.*, 2005.
- [41] P. Thati, K. Sen, and N. Martí-Oliet. An executable specification of asynchronous Pi-Calculus semantics and may testing in Maude 2.0. In F. Gadducci and U. Montanari, editors, *Proc. 4th. Intl. Workshop on Rewriting Logic and its Applications*. ENTCS, Elsevier, 2002.
- [42] A. Verdejo. *Maude como marco semántico ejecutable*. PhD thesis, Facultad de Informática, Universidad Complutense, Madrid, Spain, 2003.
- [43] A. Verdejo and N. Martí-Oliet. Implementing CCS in Maude. In *Proc. FORTE/PSTV 2000*, pages 351–366. IFIP, vol. 183, 2000.
- [44] A. Verdejo and N. Martí-Oliet. Implementing CCS in Maude 2. In F. Gadducci and U. Montanari, editors, *Proc. 4th. Intl. Workshop on Rewriting Logic and its Applications*. ENTCS, Elsevier, 2002.
- [45] W. Visser, K. Havelund, G. Brat, and S. Park. Java PathFinder - second generation of a Java model checker. In *Proceedings of Post-CAV Workshop on Advances in Verification*, 2000.