

Lem: reusable engineering of real-world semantics

Dominic P. Mulligan¹ Scott Owens² Kathryn E. Gray¹ Tom Ridge³ Peter Sewell¹

¹ University of Cambridge ² University of Kent ³ University of Leicester

Abstract

Recent years have seen remarkable successes in *rigorous engineering*: using mathematically rigorous semantic models (not just idealised calculi) of real-world processors, programming languages, protocols, and security mechanisms, for testing, proof, analysis, and design. Building these models is challenging, requiring experimentation, dialogue with vendors or standards bodies, and validation; their scale adds engineering issues akin to those of programming to the task of writing clear and usable mathematics. But language and tool support for specification is lacking. Proof assistants can be used but bring their own difficulties, and a model produced in one, perhaps requiring many person-years effort and maintained over an extended period, cannot be used by those familiar with another.

We introduce Lem, a language for engineering *reusable* large-scale semantic models. The Lem design takes inspiration both from functional programming languages and from proof assistants, and Lem definitions are translatable into OCaml for testing, Coq, HOL4, and Isabelle/HOL for proof, and LaTeX and HTML for presentation. This requires a delicate balance of expressiveness, careful library design, and implementation of transformations – akin to compilation, but subject to the constraint of producing usable and human-readable code for each target. Lem’s effectiveness is demonstrated by its use in practice.

1. Introduction

Recent years have seen a rise in *rigorous engineering*: research projects making essential use of mathematically rigorous semantic models or specifications of key computational abstractions, such as processor architectures, programming languages, protocols, and security enforcement mechanisms. These models are used in many different ways: to elucidate the behaviour of an existing real-world abstraction, as oracles to test implementations against, as the underlying assumptions or goals of verification by mechanised interactive proof, as an explicit basis for program analysis, and as a medium for design. We set the context for our work by first recalling a few representative examples. Experimental semantics research, exploring existing real-world abstractions by a combination of empirical investigation and standards formalisation, has addressed the TCP/IP protocols [6], the sequential and concurrent behaviour of x86, Power, and ARM multiprocessors [9, 10, 30, 31], JavaScript [7, 26], the (sequential) C standard [8], and the concurrency model of C/C++11 [4]; these typically require models that can be used as oracles to decide whether some experimentally observed behaviour is permitted in the model, or to enumerate all the model-permitted behaviour. Verification work using mechanised interactive proof has produced verified compilers [18, 34, 38], verified operating system kernels [16], and verified secure fault isolation [23], each based on rigorous models of the underlying processor or language and of the abstraction the verified system aims to provide. Static and dynamic analysis must either build in implicit assumptions or be explicitly

based on such a model, e.g. for binary analysis w.r.t. processor architectures [19].

This is a remarkable success story, contrasting with the state a decade ago when semantics was largely restricted to small idealised calculi: the above all involve models of (identified aspects of) real-world computer systems. But it raises its own problems. Constructing a substantial model is a major undertaking:

- one needs to understand the system being modelled in detail, often with experimental investigation of a de facto standard that has not been clearly specified even in prose;
- one has to deal with large-scale specifications, more like modest-scale programs (1–10k lines of specification) than the page or two of pencil-and-paper mathematics of a small calculus, with all the engineering and readability issues that entails; and
- models need extensive experimental validation, to establish confidence (in the absence of full formal verification down to the gate level) that they are correct models of reality; they need theoretical validation, e.g. by proofs of sanity properties showing that they are internally consistent; and they need social validation, discussion with the relevant vendors, standards committees, or community to ensure that they capture the right intent (especially for specifications that are looser than any particular implementation).

But language and tool support for such modelling and specification activities is lacking. As we discuss below, there has been no language that is specifically designed for the task, making model development more awkward than it should be, and, more importantly, *reuse* of models in different contexts is challenging and rarely achieved. For example, the above-cited works independently developed no less than six partial models of x86 instruction behaviour and two of JavaScript, and the literature contains yet more. For a small application-specific model this is not a problem, but where model development and validation may take person-years of effort such duplication of effort is not viable. The best one could do at present to make a model reusable in multiple provers would often be ad hoc porting (with scripts and hand-editing), but that is highly error prone, and at odds with the fact that these models must be maintained and developed over time.

Our thesis is that, as the subject matures and rigorous engineering becomes more widespread, the community needs to amortise this effort, establishing a collection of models of the basic abstractions, those processor architectures, programming languages and protocols that are relatively stable interfaces that computer systems depend on. These should be comprehensive and well-tested, and they have to be made available in multiple forms to enable their use for many different purposes by different groups. This should lead to a virtuous circle: the prospect of reuse motivating more complete modelling and validation, and this enabling new research that would be impossible without substantial models.

In this paper we work towards this goal. Our main contribution is the design and implementation of a modelling and specification language, Lem, to support the engineering of reusable large-scale semantic models.

Related work To explain the distinctive features of Lem we first consider the alternatives. In some cases one can use a conventional typed functional programming language (such as Haskell, OCaml, or SML) to express a model as a pure functional reference implementation or test oracle. This gives the advantages of a mature and familiar programming language, but it does not give a basis for proofs about the model, and one often needs more logical expressiveness, especially for loose specifications. In particular, one often needs good support for sets and higher-order logic, inductive relation definitions, and a clearer understanding of when one is in the fragment of the language with a direct mathematical interpretation. For reasoning, one typically turns to a proof assistant such as ACL2, Agda, Coq, HOL4, Isabelle/HOL, Matita, PVS, or Twelf. These provide very powerful proof tooling but are hard to master, and their definition languages have accreted functionality over time rather than being designed top-down as modelling/specification languages; inevitably introducing various idiosyncrasies to the language.

More seriously, the community suffers from *proof-assistant lock-in*: the difficulty in becoming fluent in their use means that very few people can use more than one effectively, and the field is partitioned into schools around each. Indeed, even within some of our own projects we have had to use multiple provers due to differing local expertise. The differences between the tools mean that it is a major and error-prone task to port even the definition of a model from one to another, rarely attempted even where much effort has gone into model development. Sometimes this is for fundamental reasons: for example, definitions which make essential use of the dependent types of Coq may be hard or impossible to practically port to HOL4 or Isabelle/HOL. However, many of the examples cited above are logically undemanding: they have no need of dependent types, the differences between classical and constructive reasoning are not particularly relevant, and there is often little or no object-language variable binding. In such cases, where a model is basically expressible in the intersection of the definition languages of several proof assistants, it should in principle be possible to port definitions; the challenge is one of robustly translating between the source languages, definition styles, and libraries. This is made particularly hard by the sensitivity of proof assistants to whether definitions are *idiomatic*: given two logically equivalent definitions, one may be much more amenable than the other to machine-assisted proof or executable code generation in a particular prover.

Previous work has established connections between different provers at the level of their internal logics [1, 11–14, 37]. These enable results proved in one system to be made available in another, but they do not provide usable *source definitions*. Between provers and programming languages, all the provers mentioned above support some kind of code generation; the other direction is less developed, though Haskell [27] provides a mapping from a fragment of Haskell to Isabelle source.

There are, of course, also many other specification tools, often with extensively engineered support for particular kinds of specification. For some examples in current use (this is by no means exhaustive) we mention ASF+SDF [33], K [28], and Maude [22] (all of which support some form of term rewriting), Ott [32] (for inductive relations over inductive syntax), and PLT Redex [15] (for reduction semantics); other tools target SMT and first-order problems. These all have advantages in their particular domains, but they are tackling rather different problems to Lem (with its focus on specification portable across multiple provers), and they lack the definitional expressiveness of higher-order logic and inductively defined relations. The need for both of these in a range of our large-scale specifications is what motivated us to develop Lem, and the choice of a language expressive enough for them but not so rich (e.g. with general dependent types) that it cannot be translated to multiple targets is central to Lem’s design.

Contribution Lem aims to combine the ease-of-use and uniform language design of programming languages with the logical expressiveness required for specification of the established proof assistants. Most importantly, it aims to support *portable* specifications, that can be used in multiple provers (it is not itself a proof tool). Spelling out our contribution in more detail:

A language of executable mathematics Lem is oriented towards (though not restricted to) executable definitions; the executable fragment of Lem can be translated into OCaml code to use as a test oracle for experimental validation, or for model exploration.

Support for multiple proof-assistant targets Lem definitions can be translated to proof-assistant definitions for Coq, HOL4, and Isabelle/HOL, to support interactive proof. The language design involves a delicate balance of expressiveness: enough for a range of large-scale modelling tasks, but still translatable into usable definitions in the various proof-assistant targets, as idiomatically as we can achieve by automatic translation (reasonably well for HOL4 and Isabelle/HOL; somewhat less so for Coq). These translations are related to, but interestingly different from, conventional programming-language implementation techniques; for example in the translations of equality and pattern matching. Perhaps surprisingly, in some cases it is best to translate into mathematically different code in different targets. The design and implementation choices that make this and the next point possible are described in §3.

Human-readable output It is also important to make the proof assistant generated definitions *human-readable*: Lem preserves the source structure and comments where it can (modulo the tension with generating idiomatic code), and it uses the same machinery to give the user control of layout for generation of production-quality LaTeX that can be used directly in papers and documentation, avoiding the error-prone and tedious manual typesetting of definitions for publication that can be necessary for some proof assistants. Lem can also generate simple HTML.

Programming-language engineering The language is designed using best-practice programming-language techniques, taking advantage of the opportunity to do a coherent design without the backwards-compatibility issues faced by proof assistants that have been extended over many years. The syntax and type system of Lem itself are specified using the Ott tool [32], which helped make the design regular (without odd corner cases); it should be easy to use by those familiar with typed functional languages such as OCaml, Haskell, or SML. Lem appears to the user much like a compiler: there is no need to learn a complex interface, and the implementation provides prompt feedback (e.g. for type errors) to the user, so that one can do type-based development and refactoring of specifications in the style of development in a typed programming language.

Library design A specification language needs a good standard library just as much as a programming language does. In §4 we describe the Lem library support and, more challenging, how it has to be related to the differing prover libraries.

Substantial usage Lem has been developed since 2010 and its effectiveness is demonstrated by a number of developments, with both academic and industrial impact. We begin in §2 by recalling those, to explain more clearly what it is (and is not) good for, and discuss one in more detail in §5. Lem and its documentation are available (under a BSD license from a Bitbucket repository) from <http://www.cl.cam.ac.uk/~pes20/lem/>.

Note Some of the motivation behind Lem, and its initial implementation, was initially presented in a short “Rough Diamond” paper [25]. Since then, Lem has been validated by substantial use (§2), and its design and implementation have become significantly more sophisticated (most of §3 and §4).

2. Lem in practice

To explain more concretely the kind of specification work that Lem is aimed at, and to demonstrate its practical effectiveness as a tool for large-scale specification, we describe the main Lem developments produced to date. These underlie multiple academic publications (six in POPL, PLDI, and CAV) and have had industrial impact, in clarifying the IBM Power and ARM concurrency behaviour, on the C/C++11 ISO standard concurrency semantics, and on the compilation scheme of the latter to Power and ARM. For each development we give the number of non-comment lines of specification (LoS) and the Lem targets used, and comment on the mathematical style of the specification. Snapshots of those marked \star are accessible via the URL above.

Sarkar et al. [30] describe an operational model (3008 LoS, \star) for the relaxed-memory behaviour of IBM Power and ARM multiprocessors. The executable Lem-generated OCaml code forms the kernel of the `ppcmem` tool (<http://www.cl.cam.ac.uk/~pes20/ppcmem>) for exhaustive and interactive exploration of the model on examples; that and the generated LaTeX supported experimental validation and extensive discussion with an IBM architect during model development; and the generated Coq has been used for some modest experiments with mechanised proof (unpublished). This model uses a combination of mathematical styles. For specifying the multiprocessor memory subsystem it has an abstract machine built over sets, relations, and lists of memory events, while the thread semantics involves abstract syntax trees of assembly and micro-operation instructions. These each define labelled transition systems (LTSSs), which are combined with an inductive-relation top-level parallel composition.

Batty et al. [4] describe an axiomatic memory model and various sub-models (1517 LoS, \star) for concurrency in the C/C++11 standards. The Lem-generated OCaml and HTML form the kernel of the `cppmem` tool, again with a web interface for interactive and exhaustive exploration of the model on examples (<http://svr-pes20-cppmem.cl.cam.ac.uk/cppmem/>); that and the generated LaTeX supported discussion with the ISO standards committee to develop the model and improve the standard, and `cppmem` has been used by some GCC and Linux developers and within ARM. The generated HOL4 code has been used for mechanised proof of metatheory, and the generated Coq and Isabelle have also been used for proof experiments. The bulk of this development comprises predicates over candidate executions, represented as sets of events with various order relations; we describe it in more detail in § 5.

Batty et al. [3, 29] describe extensions to the above Power and C/C++11 models and correctness proofs for a compilation scheme from C/C++11 concurrency primitives to Power concurrency primitives; these are hand proofs but with lemmas expressed in Lem (931 LoS); this is a useful middle ground between LaTeX and fully mechanised proofs.

Mador-Haim et al. [20] describe an axiomatic memory model for Power and mappings between that and the operational model above (1155 LoS), again with hand proofs; the generated OCaml code was used to test equivalence of the two models on examples and the generated LaTeX to define the model in the paper.

The specification of the de facto standard of the TCP/IP network protocols and the Sockets API by Bishop et al. [6] was originally expressed in HOL4 and has now been ported to Lem (6681 LoS, \star). The style is essentially pure higher-order functional programming, but with sets and maps, a relational monad, inductive relations, and logic. The port to Lem was essentially straightforward except that the original specification made heavy use of HOL4's advanced features for user-defined syntax, while Lem has a more uniform but more restricted syntax; these had to be manually unpicked.

The Ott definition of OCaml_{light} by Owens [24], originally used to generate HOL4, has been adapted to generate Lem (3133 LoS

in Lem, from 4253 lines of Ott source \star). OCaml_{light} uses a typical mathematical style for a typed functional programming language semantics, with inductive relations specifying a type system and small-step operational semantics. Lem's inductive relation support was general enough to support a direct and automatic translation from Ott. This development is of a different character than the others in that here we are translating *into* Lem, showing its potential to be used as a general purpose front-end for domain-specific tools (here Ott) that target multiple provers.

Kumar et al. [17] describe CakeML (<http://cakeml.org>), a mechanically verified ML system above x86-64 machine code. The source language definition (AST, type checking, and small- and big-step operational semantics), compiler, and some additional machinery are written in Lem (4897 LoS); the generated HOL4 code is used for proofs. The operational semantics of CakeML and the low-level CakeML bytecode, and the CakeML type system are given as inductively defined relations. The compiler is specified as a collection of recursive functions, in a typical functional programming style. Furthermore, the semantics and type system rely heavily on helper functions. This is in marked contrast with the OCaml_{light} semantics which were almost entirely relational. It was convenient when doing the proofs for helper (partial) functions (e.g., environment lookup) to be specified as functions rather than relations. Besides relations, functions, and datatype definitions, CakeML uses the Lem list, finite map, and string libraries. The Lem sources are in the compiler, bytecode, semantics, and translator subdirectories of the repository.

Ongoing work describes a model for a substantial part of the C programming language (8274 LoS), largely in a functional style over an inductively defined abstract syntax. This initially used only the OCaml target but has recently been adapted to also generate Coq; the effort involved was nontrivial but remains small (weeks vs years) in comparison to the time required to develop the model.

Together, these demonstrate that Lem is expressive enough for a range of modelling tasks, spanning processor architectures, C and ML-like programming languages, and network protocols, and that it compiles to usable executable code for model exploration and usable definitions in multiple provers for proof. Anecdotally, our experience is that it is reasonably easy to use (analogous to a functional programming language) and that Lem models are malleable: in developing a model, one can largely focus on the domain being modelled rather than issues of expressing the model in Lem, and models can be easily changed as they are developed. Our impression so far is that adding a new target to a specification is usually easy or (at worst) like the last example above, but it is possible that to get a truly idiomatic version in a new target one may need to hand-rewrite substantial parts of the specification in that target, and prove equivalence to the Lem-generated versions. For an intricate specification, embodying much development and validation effort, even that may be preferable to the alternative of hand-rewriting the whole. More experience is needed to see which is normally the case.

Lem is a general-purpose specification language but not, of course, an all-purpose one. It does not aim to support specifications with elaborate dependently typed hierarchies of mathematical structures. It has a straightforward syntax (again similar to that of a functional programming language), without support for rich user-defined syntax. The executable code Lem generates is designed to have a clear relationship to the source and has sufficient performance to support exploration of models on the intricate but small examples that (e.g.) arise as concurrency test cases; it is not aimed at producing performance-optimised code. Lem complements the Ott tool [32]: Ott supports arbitrary context-free user-defined syntax and inductive relations; this is a good fit for high-level programming language and calculus semantics but the lack of the more general


```

typ ::= _ | α | typ1 → typ2 | typ1 × ... × typn
      | id typ1 .. typn | backtick_string typ1 .. typn | (typ)

pat ::= _ | (pat as x) | (pat : typ) | id pat1 .. patn
      | ⟨fpat1; ...; fpatn; ?⟩ | (pat1, ..., patn) | [pat1; ..; patn; ?]
      | (pat) | pat1 :: pat2 | x + num | lit

exp ::= id | backtick_string | fun psexp
      | function |? pexp1 | ... | pexpn end | exp1 exp2 | exp1 ix exp2
      | ⟨fexp1⟩ | ⟨exp with fexp1⟩ | exp.id
      | match exp with |? pexp1 | ... | pexpn end | (exp : typ)
      | let letbind in exp | (exp1, ..., expn) | [exp1; ..; expn; ?]
      | (exp) | begin exp end | if exp1 then exp2 else exp3
      | exp1 :: exp2 | lit | {exp1|exp2}
      | {exp1| forall qbind1 .. qbindn | exp2} | {exp1; ..; expn; ?}
      | q qbind1 ... qbindn.exp | [exp1| forall qbind1 .. qbindn | exp2]
      | do id pat1 ← exp1; .. patn ← expn; in exp end

psexp ::= pat1 ... patn → exp

```

Figure 1. Lem Syntax Excerpts

types, functions and library support of Lem makes it awkward to use for modelling the lower-level systems and languages described above. Lem can serve as an intermediate language for other tools that produce definitions of types, functions, or inductive relations, and we have refactored Ott (which originally produced source code for Coq, HOL4, and Isabelle/HOL directly) to produce Lem definitions, leaving Lem to handle the prover-specific idiosyncrasies.

3. Design for portable specification

Aiming to support a range of specification tasks, Lem does not build in any domain-specific assumptions on the form of specification permitted: it is a general language of type, higher-order function, and inductive relation definitions.

The language is intended to be as expressive and straightforward as possible given this generality; it avoids novel or exotic features that would give it a steep learning curve, or render translations into the various targets infeasible. From functional programming languages we take pure higher-order functions, general recursion, recursive algebraic datatypes, records $\langle | \cdot | \rangle$, lists $[\cdot]$, pattern matching, parametric polymorphism, a simple type class mechanism for overloading, and a simple module system. To these we add logical constructs familiar in provers: universal and existential quantification, sets $\{\cdot\}$ (including set comprehensions), relations, finite maps, inductive relation definitions, and lemma statements. The concrete syntax for types and expressions is broadly standard, as one can see in the excerpts in Fig. 1. It largely follows OCaml for the common constructs (the main exceptions being prefix type applications, curried constructors, and records using $\langle | \cdot | \rangle$, to allow $\{\cdot\}$ for sets).

Then there are facilities to let the user tune how Lem definitions are mapped into the various targets (by declaring target representations and controlling notation, renaming, inlining, and type classes), to generate witness types and executable functions from inductive relations, and for assertions; we describe all these below.

Although most of Lem’s features should be unsurprising at first sight, their detailed design must carefully manage tradeoffs between Lem’s expressiveness and usability as specification language, and the need to generate usable code for the various target languages and provers. In this section, we first describe the basic architecture of the Lem implementation, and then describe the design issues and our solutions for each of these seemingly simple features – including polymorphism, equality, partiality, sets, and inductive relations. In the next section (§4), we will describe Lem’s library definition

mechanism, and explain how it connects the Lem standard library to the widely varying standard libraries of the targets.

3.1 System architecture

Lem is written in OCaml, and it follows the architecture of a traditional compiler invoked from the command line, with conventional lexing and parsing of source files into an untyped AST, followed by type inference (in the style of Milner’s Algorithm W) into a typed AST. The OCaml type declarations for the untyped AST are automatically extracted from the formal definition of Lem’s syntax by Ott [32], to help keep the Lem implementation and formal specification in agreement. Its parser and lexer are implemented using `ocamlyacc` and `ocamllex`.

To produce output for a particular target, the typed AST is then transformed, compiling away features that that target does not support (transforming away type classes via dictionary passing, compiling away unsupported pattern matching forms, etc.). Special idiosyncrasies of the target may require additional clean-up (e.g., variable name clashes, extra required parentheses, different infix operator syntax); then the resulting AST is printed in the target’s source syntax. The individual transformations are reused in different combinations for the different targets, as required, and the implementation checks that they preserve typing, greatly easing debugging of Lem itself.

Lem has about 29 000 lines of OCaml code. This compactness makes it easy to understand and adapt. The library is extensive in comparison to the code size: there are about 7700 lines of Lem libraries, together with 1800 lines of OCaml, 1200 lines of Isabelle/HOL, 500 lines of Coq and 300 lines of HOL4.

3.2 Whitespace preservation and refactoring support

Unlike a compiler or proof assistant, Lem’s front end preserves all of the comments, whitespace, and line breaks in the source files. Lem attempts to format its output using the formatting of the input, rather than a pretty printing algorithm, in order to give the user fine-grained control of the output layout. Of course, this is not always possible when the input had to undergo significant transformation, such as the pattern match compilation or dictionary passing translation discussed below. In these cases, we use a standard pretty-printing algorithm for the affected expressions, but can at least keep all of the comments from the input. Crucially, the LaTeX backend does not perform such transformations, and so the user has control over the typesetting of their specifications, including linebreaks and indentation.

Since Lem preserves comments, whitespace and linebreaks, it can reproduce its input exactly. This means that Lem can also easily be used as a powerful refactoring tool for its own input. A special refactoring backend is able to rename functions and types, remove or add function arguments, move definitions to different modules and much more.

Whitespace preservation also means that a standard OCaml profiling tool can be used for coverage analysis of a *specification*, identifying the parts of a Lem specification that are exercised by a particular set of tests run in a semantics exploration tool (with kernel generated from the specification): the generated OCaml is close enough to the Lem source that one can usually easily relate one to the other.

3.3 Polymorphism and dependency

Parametric polymorphism is essential in our specifications – most obviously, for the library functions over lists, sets, and suchlike, and for functions over user-defined polymorphic inductive types. But, as is well-known in the higher-order-logic context, let-polymorphism (the implicit generalisation of types to type-schemes in nested **let** bindings) makes higher-order logic unsound. Accordingly, Lem

supports top-level parametric polymorphism, but type generalisation is restricted to (module) top-level definitions, as in HOL4 and Isabelle/HOL, but diverging from the Hindley-Milner-style polymorphism found in ML-like programming languages. We have not found this to be limiting in practice, and Vytiniotis et al. [35, §4.3] provide empirical evidence that let-polymorphism is rarely used in practical Haskell programming.

More sophisticated type-language features, such as System F-style polymorphism, dependent types, and subtyping, are also not included in Lem, because these would be unduly difficult or impossible to support in many of our chosen targets. However, we do support *ad hoc* polymorphism with type classes.

3.4 Equality and type classes

There are substantial differences in the treatment of equality in our different targets. In the two implementations of higher-order logic, Isabelle/HOL and HOL4, there is a ‘pervasive’ equality constant $=$, at type $\alpha \rightarrow \alpha \rightarrow \text{bool}$. OCaml features a similarly typed equality constant, but it is only usable for non-function types (raising an exception otherwise). Further, this OCaml polymorphic equality is structural, and does not take into account equivalence relations between data types. For abstract types such as sets (implemented in the Lem translation to OCaml as balanced binary trees) one needs to use an equality function specific to sets that compares sets based on their elements, rather than their low-level representation in memory, and that function must have access to the order relation used to build the trees. We could introduce specific equalities at each type, for example setEq at type $\forall \alpha. (\alpha \rightarrow \alpha \rightarrow \text{bool}) \rightarrow \text{set } \alpha \rightarrow \text{set } \alpha \rightarrow \text{bool}$. However, this would force the user to supply the equality function on elements of the set by hand (a task that should be automated), and break the uniformity of the treatment of equality within the language. Lem includes type classes to solve both of these problems.

The Lem Eq type class has the following form:

```
class (Eq  $\alpha$ ) =
  val (=) ['isEqual'] :  $\alpha \rightarrow \alpha \rightarrow \text{bool}$ 
  val (<=) ['isInequal'] :  $\alpha \rightarrow \alpha \rightarrow \text{bool}$ 
end
```

This type class introduces two methods, equality $=$, and inequality $<=$ (together with alphanumeric alternative names). The type class may be instantiated at any type by providing implementations for the equality and inequality methods at that type. For OCaml, Coq, Isabelle/HOL and HOL4, the Lem translation introduces explicit dictionary passing to handle the general case of type classes and their constraints. But there are three situations in which introducing dictionary passing would lead to non-idiomatic code and obstruct use of the extensive proof automation facilities of the provers.

First, for the backends with a general polymorphic boolean equality, Isabelle/HOL and HOL4, we wish to map the Lem overloaded equality constant to that native equality constant. We achieve this with a general inlining method:

```
let inline {hol;isabelle} (=) =
  unsafe_structural_equality
```

This inlining effectively ‘turns off’ the equality type class for HOL4 and Isabelle/HOL. Since all methods are implemented without using the type-class mechanism, the class (and the associated dictionary passing) is not generated for these backends. Here `unsafe_structural_equality`, intended to be used by library authors only, is mapped to the native equality constants in the Isabelle/HOL and HOL4 backends. It is also mapped to the pervasive (but restricted as above, hence the ‘unsafe’) equality function in the OCaml backend.

Secondly, sometimes we want to use a method only for certain backends. With the normal Lem standard library, Lem sets are

represented in Coq and OCaml as ordered, balanced binary trees and therefore an order on their elements needs to be provided. This is achieved via a Lem type class `SetType`. However, the HOL4 and Isabelle/HOL sets do not require an order, so using the type class mechanism naïvely would lead to non-idiomatic HOL4 and Isabelle/HOL code by generating unnecessary dictionary arguments. By restricting class-methods to certain backends, this problem can be solved:

```
class (SetType  $\alpha$ )
  val {ocaml;coq} setElemCompare :  $\alpha \rightarrow \alpha \rightarrow \text{ordering}$ 
end
```

Lem’s type-checker ensures that the method `setElemCompare` is only used in the Coq and OCaml backends, with the type class being safely eliminated for all other backends.

Thirdly, we wish to avoid introducing dictionary passing where all occurrences of a type class can be statically resolved, for example for the Lem type-class `Numerals`, used for overloading numeral syntax for multiple number types. This type class is declared as *inline*. All occurrences must be resolved, and all methods of an inlined type-class are replaced (inlined) with their instantiations.

Using type classes for equality in Lem also works around another issue in the Coq backend. In Coq the types `bool` and `Prop` (technically a *sort*) are distinct, with ‘propositional equality’ having type $\forall \alpha. \alpha \rightarrow \alpha \rightarrow \text{Prop}$. It is generally impossible in Coq, without additional axioms, to produce an equality constant of type $\forall \alpha. \alpha \rightarrow \alpha \rightarrow \text{bool}$. Rather, ‘boolean equalities’ are given at specific types, such as `nateq` of type $\mathbb{N} \rightarrow \mathbb{N} \rightarrow \text{bool}$ and `booleq` of type `bool \rightarrow bool \rightarrow bool`, and so on. Lem, following Isabelle/HOL and HOL4, identifies boolean-typed expressions and propositions, or formulae. This poses a problem with extraction to Coq, as innocuous Lem code—e.g. `if 4 = 5 then true else false`—becomes problematic if we try to extract an equality constant in Lem to Coq’s propositional equality constant, as case analysis over `Prop` is not permitted. Boolean equality is therefore needed, and type classes are used to handle this uniformly. We return to `bool` and `Prop` in §3.13.

In principle, it should be straightforward for Lem to automatically generate the appropriate instance declarations for its basic library type classes (`Eq`, `Ord`, `SetType`, and `MapKeyType`) for most user-defined types, analogous to Haskell’s deriving mechanism. Unfortunately that is not currently implemented, and a default instantiation mechanism is used instead. This means the user must take care to give non-default instantiations in some places, e.g., for the OCaml backend, for types which both contain an abstract type and are used as set elements.

Compared to Haskell’s type class system, Lem’s is intentionally less expressive: our goal is to support simple overloaded operators and to solve the above-mentioned issues dealing with the subtle differences between our target systems, not to enable generic and polytypic programming. In particular, Lem does not support constructor classes (type classes with type variables at a higher kind than $*$), default implementations of methods, or backchaining search for instances, nor does it support more recent extensions: multi-parameter type classes, functional dependencies, and so on.

Our implementation already has much of the underlying infrastructure needed for some extensions: instances at compound types, multi-parameter classes, default methods, and other features that would make type classes more convenient to use, and these should be easy to implement in the future. However, constructor classes and the like pose more significant technical challenges, because their dictionary-passing translations are not naturally typeable in the systems of ML-style polymorphism of our OCaml and HOL4 targets and Lem itself. One possible design choice would be to support them, but only at statically known types (i.e., where dictionaries are not required).

3.5 Module system and ‘do’ notation

Lem has a simple module system designed to support the organisation of large-scale specifications into multiple files, and to allow the reuse of specification libraries (including Lem’s standard library itself) across developments. It does not include the *programming-in-the-large* features of advanced PL module/component systems, such as enforced abstraction or parameterisation, because those are primarily useful in code bases that are orders of magnitude larger than even large specifications. Our larger developments each comprise between 10 and 40 non-library Lem files.

The module system is based on a restricted subset of OCaml’s module system: modules contain sequences of definitions, and modules can be defined at the top level (but not inside of expressions). Definitions inside of a module are accessed from outside either through the open declaration, or by explicitly spelling out the module path (with dot notation). There are no signatures or functors.

Lem supports a Haskell-like do notation for specifications that involve monads. Unlike Haskell, Lem’s type class system is not powerful enough to infer which monad is being used (since a monad is a type constructor rather than a type). Instead, each do expression is annotated with a module name that defines the relevant bind and return operations.

3.6 Pattern matching

The usefulness of pattern matching is well-established in functional programming languages, and it is even more valuable in a specification language like Lem because it supports high level, abstract and clear code. But the support for pattern matching varies significantly between the different backends we target. For example, record patterns are supported by OCaml and (recently) Coq, but not by HOL4 or Isabelle/HOL; λ s-patterns are supported by OCaml and Coq, but not by Isabelle/HOL or HOL4; and idiomatic HOL4, Isabelle/HOL and Coq code uses pattern matching on natural numbers (using zero and the successor function as constructors), whereas this is not supported by OCaml. Besides differences in explicit pattern match expressions, there are also differences in how one can use patterns elsewhere. Anonymous functions in OCaml allow arbitrary patterns as arguments, but in HOL4, Isabelle/HOL and Coq only tuples of variables are allowed. The situation is similar for let-expressions and patterns occurring in restricted quantifications. Beyond these local syntactic properties, there are also important semantic differences: Isabelle/HOL supports non-exhaustive pattern matches, but no redundant rows; Coq requires pattern matches to be exhaustive and prohibits redundant rows; OCaml allows both redundant and non-exhaustive pattern matches; and HOL4 allows non-exhaustive pattern matches but redundant rows only at certain places.

We take the opportunity in the Lem language design to provide more general pattern matching, combining the facilities of each target, and compile those general patterns away where necessary. This compilation mostly follows a simple, standard approach, essentially switching on the outermost constructor symbols ([2]) in order to compile pattern matches to decision trees, implemented efficiently following ideas from [21]. But in contrast to the normal PL situation, where one wants to compile away pattern matching altogether, for Lem to produce idiomatic and human-readable code, we have to *preserve* as much of the original structure as possible; Lem needs sophisticated models of the capabilities of each backend in order to compile only the unsupported features and preserve as much of the original structure as possible.

For example, consider the following Lem record pattern match:

```
type t = { f1: nat; f2: bool }

let test_fun x = match x with
| Nothing => 0
| Just { f2 = true } => 1
```

```
| Just { f1 = 0 } => (1 : nat)
| Just { f1 = x } => x + 2
end
```

OCaml supports all these pattern forms, so the resulting OCaml code looks very similar to the input; pattern compilation is not needed. In contrast, HOL4 does not support record patterns, and compiling them away – while preserving as much of the structure as possible – leads to the following HOL4 result:

```
val _ = Define 'test_fun x = case x of
  NONE => 0
| SOME t => (case (t.f2, t.f1) of
  (T,_) => 1
| (_,0) => (1 : num)
| (_,_) => let x = t.f1 in x + 2)';
```

Our implementation also supports a mechanism similar to view patterns [36]. This powerful feature allows users to write even more abstract, higher-level code. For example, consider the type of sets. Functional programmers might be tempted to use the choose function to get the unique element out of a set known by the programmer to be a singleton. However, choose requires the axiom of choice in Coq, and its result is undefined for empty sets and underspecified for sets with more than one element. A solution is using a case-split for sets:

```
let set_case s c_empty c_sing c_else =
  if (null s) then c_empty else
  if (size s = 1) then c_sing (choose s) else
  c_else
```

This set_case can be implemented in all backends and is even executable. Lem’s view-pattern feature allows setting it up together with empty and singleton for pattern matching.

```
declare pattern_match inexhaustive
set 'a = [ empty; singleton ] set_case
```

This setup provides easily readable syntax, as below:

```
let set_test s : nat =
  match (s : set nat) with
  | empty => 0
  | singleton (x + 3) => 2
  | singleton _ => 1
  | _ => 3
end
```

3.7 Partial pattern matches

Programming languages typically permit non-exhaustive or partial pattern matches, with a dynamic exception or error if a match fails at runtime. This supports two use cases: (a) where the programmer knows that the match will never actually fail, because of some invariant (e.g. that a list is nonempty) that may not be expressible in the language’s type system, and (b) where the intended control flow includes paths where an exception is raised and handled.

Proof assistants are more restrictive to retain soundness. For (a), in the HOL4 and Isabelle/HOL logics the Hilbert choice operator lets one construct an arbitrary unknown default value at any type, but in Coq all matches must be total (though the Coq type system can capture complex invariants). None of the three support (b) directly, as that would require a deeply embedded exception monad.

Lem permits partial matches, to support (a). For OCaml they are mapped directly onto similarly partial matches (any Match_failure exceptions can be handled by wrappers around the Lem-generated code, but not within that code). For HOL4 and Isabelle/HOL partial matches are mapped to syntactically partial matches that ultimately use that Hilbert choice operator, which is appropriate in cases where the match can never fail. For Coq, the story is more complex. Given a partial match at a concrete type, e.g.


```
match (m: maybe bool) with
| Just j ⇒ j
end
```

that lacks a case for the `Nothing` constructor of the `maybe` type and whose result type is the concrete `bool`, we generate the following:

```
match (m: maybe bool) with
| Just j ⇒ j
| _ ⇒ bool_default
```

Here `bool_default` is a *default value*. Default values for base types are provided in an external harness file and Lem automatically generates default values for all user-defined types during generation of Coq code. Again this is suitable in cases where the user knows that the result of that branch is irrelevant.

Partial matches with a polymorphic result type are more problematic (though in practice we have not found many cases where they arise), e.g.

```
match (m: maybe  $\beta$ ) with
| Just j ⇒ j
end
```

We could use a type class with a default-value method, at the cost of introducing dictionary passing for that into the generated Coq for any definition which hereditarily involves a partial polymorphic match. Instead, on the assumption that the user would usually prefer to adapt their specification to avoid this, at present we translate into a complete Coq match

```
match (m: maybe  $\beta$ ) with
| Just j ⇒ j
| _ ⇒ DAEMON (* From <position>. *)
end
```

that introduces a placeholder marking a point in the source specification that needs to be addressed, and a warning or error is issued. Here, `DAEMON` is a constant of type $\forall \alpha. \alpha$ —the type of logical falsity. This lets one build the remaining development but its presence as an axiom makes Coq’s logic inconsistent, so one would aim to remove all such usages.

Lem itself has no operational semantics, its meaning being defined by the translations to the various targets. It is pure in the sense that it has no I/O or store effects, but in general the use of partial pattern matches and their compilation using exceptions in the OCaml backend can make the OCaml evaluation order observable in generated OCaml code (and there is a similar issue for nontermination). However, if partial matches are only used for (a) above, with missing cases only in places that are unreachable under any evaluation order (and if one uses only Lem functions that are robustly terminating in the same sense), then the target evaluation order should be irrelevant.

3.8 General recursion vs total functions

Functional languages typically allow general recursion, whereas proof assistants generally require some kind of termination proof for all recursive functions. For example, the function `let rec f x = not (f x)` will diverge in OCaml when called, but would introduce an unsoundness to a proof assistant. Defining a recursive function in Coq, HOL4 or Isabelle/HOL therefore requires the user, often assisted by the proof assistant itself but in quite different ways, to supply evidence that the function terminates on all inputs. Lem permits general recursion but with hooks to invoke the backend’s various automatic termination provers, where they exist, e.g.:

```
declare termination_argument my_rec_function = automatic
```

If that does not suffice, for the HOL4 and Isabelle/HOL backends Lem can defer termination proofs, letting the user provide them manually later (see §3.11).

Both facilities are used in CakeML, where simple functions (e.g., `stackshift` in `compiler/toBytecode.lem`) declare an automatic termination argument. More complex ones that HOL cannot automatically prove termination for (e.g., the mutually recursive `exp_to_Cexp` of the same file) are proved in a separate HOL file (e.g. `compiler/compilerTerminationScript.sml`).

In Coq the situation is more complex. Coq can spot that recursive functions exhibiting simple recursion schemes are terminating. If this machinery fails, the user is required to rewrite the function in terms of an accessibility predicate with an explicit well-founded order supplied. This requires more input from the user than the equivalent termination proofs of Isabelle/HOL and HOL4. Lem functions therefore must be written in a style Coq’s termination prover can handle, and if not possible, the user must perform the transformation in Coq. To date, we have come across one function in our Lem developments that require such a rewriting in Coq, but we expect to see more with more extensive use of the Coq backend.

3.9 Per-target representation differences

One might imagine that a portable specification should necessarily map onto mathematically equivalent definitions in the different targets, but this turns out not to be the case: there is a tension between it and the need to generate idiomatic code.

For example, in OCaml the standard type for numbers is `int`, 31- or 63-bit signed integers, while our proof assistant targets use unbounded natural numbers as their standard type. In each target common functions like the list length function use that target’s local standard number type, so either Lem has to add wrappers to such functions or map the same Lem type to different mathematical constructs in different targets.

We give the user the choice, providing a Lem type `nat` which is translated to the standard number type of the targets (for use where the user knows the differences are irrelevant, e.g. when using numbers just as unique identifiers) and a Lem type `natural` which always maps to unbounded naturals. Similarly we provide `int` and `integer`. We provide conversion functions between the different number types; polymorphic numeral constants and polymorphic arithmetic functions make switching between number types easy.

The choice of set representations is more involved. Since Lem is geared towards executability, one might want to model only finite sets, and in practice users are often only concerned with finite structures, e.g. as arising in finite executions of the models described in §2. On the other hand, potentially infinite sets are very common and very useful for specifications and more convenient to work with in some provers. Similar questions occur with quantification. Only bounded quantification is easily executable, but unbounded quantification is often useful for specifications. Lem permits unbounded quantification and infinite sets, but only for the HOL4 and Isabelle/HOL targets; for OCaml and Coq the library currently only provides finite sets and bounded quantification. In future we will provide alternatives as for `nat` and `natural` above.

Having decided on finite or infinite sets, there remains a non-trivial choice of the best target representation. OCaml has a set implementation in its standard library, but with a functorised interface to supply an order on elements; we use a library with a polymorphic interface instead. For Isabelle/HOL and HOL4 the idiomatic potentially infinite sets are used. Isabelle/HOL also has several interesting alternative ones with improved executability, but they are supplied with an order via Isabelle/HOL’s type-class mechanism which would require the user to supply potentially nontrivial proof. For Coq, there is no single idiomatic set library: `FSets` provides a module-oriented implementation, the `Collections` wrapper around this library uses type classes, and one can also represent sets as functions into `bool` or `Prop` (‘ensembles’). They all also require an order, or at least a decidable equality relation, which would have

to be provided via Coq’s type-class mechanism. We add a simple finite-set library, leaving additional Lem libraries to map Lem sets to other Coq representations as future work.

Associative maps have similar issues to sets. Lem translates maps to Lem-specific finite map implementations in OCaml and Coq and into the idiomatic map types for HOL4 and Isabelle/HOL, which are finite maps and infinite maps respectively.

The binary relation type is closely related to sets and therefore also has similar design issues: relations can be seen as sets of pairs or as binary predicates. Both Isabelle/HOL and HOL4 provide dedicated libraries for both representations, and Lem currently maps only to the set representations.

3.10 Naming, notation and namespace issues

The namespaces and the sets of pre-defined identifiers and reserved keywords differ in each backend. For example, `op` is a reserved word in Isabelle/HOL but not in Coq. Without reserving every reserved word and every pre-defined identifier in each of our backends in Lem, we must implement a renaming mechanism to avoid name clashes post extraction.

Lem has a simple model of the namespaces and a list of the reserved words of each target. If a reserved word for the backend in question is encountered, Lem will automatically rename the constant and issue a warning. For example,

```
let op f g = ...
```

is automatically renamed to `op0`, or some other globally fresh name, when extracting to Isabelle/HOL. All occurrences of this constant are also suitably renamed. A warning is issued on the command line to notify users of the renaming. The user can also manually control renaming. For instance, placing

```
rename {isabelle} op = isaop
```

in a Lem source file will rename `op` to `isaop` during extraction to Isabelle/HOL. This mechanism allows us to avoid auto-generated names, making Lem generate stable, predictable output even in the presence of name clashes, and renaming constants to follow the conventions of a particular backend facilitates idiomatic backend code.

The Lem renaming mechanism respects target-specific differences in scoping. For example,

```
let op op = ...
```

features a function `op` taking a parameter called `op`. For some backends, this is problematic, and so the two should be renamed apart there, but not otherwise. In our running example, the function `op` will be renamed to `isaop` and the argument `op` to `op0` to avoid clashes with the Isabelle/HOL keyword. Lem correctly renames the name of the parameter apart from all constants present in the context.

3.11 Assertions, lemmata, and auxiliary outputs

Lem is intended to translate into target output that can be directly used, without manual editing by the user: Lem definitions generate target definitions which can be fully automatically checked. However, Lem can also generate aids for the user in additional *auxiliary* files, intended to be copied and manually edited by the user as necessary. For example, when defining a complicated recursive function in Lem, HOL4 and Isabelle/HOL can leave the termination proof to the user. Lem generates in the auxiliary file a template for its termination proof which the user can flesh out.

Lem also supports *assertions*, *lemmas* and *theorems*. Assertions are executable, for automated testing of simple properties. For OCaml they generate code in an auxiliary file that runs automated unit tests. For the theorem prover backends, they generate

proof obligations, which are attempted to be discharged automatically. Lemmas and theorems are non-executable; they add proof obligations to auxiliary files. Simple, low-level lemmas can also be used for testing: often the resulting proof obligations can be automatically discharged. Exporting complicated lemmas and theorems might be beneficial as well. For example, Isabelle/HOL provides highly automated, powerful tools, and by exporting a lemma to Isabelle/HOL, this powerful machinery is easily accessible even by users not familiar with Isabelle. The following Lem lemma:

```
Lemma unzip_zip:
  ∀l1 l2. unzip (zip l1 l2) = (l1, l2)
```

is translated to the Isabelle/HOL code:

```
Lemma unzip_zip:
  "∀l1 l2. list_unzip (zip l1 l2) = (l1, l2)"
(* try *) by auto
```

The automated proof attempt by the `auto` method fails. If the user then uncomments `try`, various automated methods are run to either prove the lemma or find a counterexample. These methods include running external SMT and first order provers, internal natural deduction tools, and a sophisticated counterexample generator. In this example, Isabelle/HOL quickly finds a counterexample:

```
Nitpick found a counterexample for card 'a = 2 and card 'b = 2:
  Skolem constants: l1 = [a1], l2 = []
```

In general, tools like the counterexample generator Nitpick work well with Lem-generated Isabelle/HOL code, because Lem is tailored toward executability and this executability is (as far as possible) preserved by the translation to Isabelle/HOL. Therefore, non-trivial counterexamples can often be found automatically.

3.12 Inductive relations

Specifications often involve inductively defined relations, such as type systems or evaluation relations defined as the smallest relations satisfying a collection of rules. Lem provides an *inductive relations* mechanism for this purpose.

The following Lem definition (generated by Ott from a similar definition expressed over a calculus syntax) captures the reduction relation of the call-by-value λ -calculus:

```
indreln [reduce: term → term → bool]
  ax_app: ∀ x t1 v2.
    (is_val_of_term v2) ⇒
      reduce (T_app (T_lam x t1) v2) (subst v2 x t1)
and
  ctx_app_fun: ∀ t1 t t1'.
    (reduce t1 t1') ⇒
      reduce (T_app t1 t) (T_app t1' t)
and
  ctx_app_arg: ∀ v t1 t1'.
    (is_val_of_term v) ∧ (reduce t1 t1') ⇒
      reduce (T_app v t1) (T_app v t1')
```

Here, `reduce` is introduced as a relation—Lem relations are essentially functions into `bool`—between AST terms. In this case, the `reduce` relation is defined via three clauses. Within a clause, the full power of the Lem language is available, rather than a purely relational subset, as in Prolog. For instance, `is_val_of_term` and `subst` are functions, rather than relations. One may also define mutually recursive inductive relations.

The Isabelle/HOL, HOL4 and Coq backends support inductive relations, and we map Lem inductive relations into the native inductive relations of these backends.

However, inductive relations are not naturally expressible in our OCaml backend. We therefore implement a compilation process, compiling a Lem inductive relation into a function that searches for

derivations. The compilation process is given a *mode* by the user, a description of which components of the relation are to be treated as ‘inputs’ and which are to be treated as ‘outputs’. This compilation scheme is similar to one implemented within the Isabelle/HOL proof assistant, as implemented by Berghofer et al. [5] (there verified within Isabelle/HOL).

We go further than the Isabelle/HOL compilation scheme in automatically generating *witness types*, which encode a derivation tree for a given inductive relation. The functions generated by the compilation scheme can return witnesses, and additional functions check whether an element of the witness type belongs to the relation. These witnesses may also be produced externally, e.g. by a typechecker implementation that one wants to test, using the Lem-generated checker functions, against its definition. The generated types and functions are themselves defined in Lem and thus can be used by later Lem definitions and translated to any of the Lem targets.

For example, the following syntax instructs Lem to generate a reduction function for `reduce`, naming it `onestep`:

```
[reduce: term → term → bool
onestep: input → output ]
```

The mode annotation on `onestep` instructs the compilation machinery to consider the first component of `reduce` as an input, and the second an output. Additional annotations can indicate that the function returns multiple results, that there must be a unique return value, or to generate a witness for the relation. Generated functions that are partial are treated in the same way as partial pattern matches. For non-deterministic rules, the generation searches exhaustively and, according to the annotation, either returns a list of elements in the relation or a single one (or is undefined).

To generate witness types and witness-checking functions, one can write:

```
[reduce : term → term → bool
witness type r_witness; check check_r;]
```

This generates a witness type for the `reduce` relation:

```
type r_witness =
| Ctx_app_arg_witness of term × term × term × r_witness
| Ctx_app_fun_witness of term × term × term × r_witness
| Ax_app_witness of string × term × term
```

Instrumenting an interpreter or type-checker to produce such witnesses should be straightforward. Doing that, and also scaling this code generation up to make it work well on practical examples, is in progress.

3.13 Prop and bool in the Coq backend

As mentioned in our discussion of equality and type classes in Lem, Coq maintains a distinction between a *sort* of propositions, `Prop`, and a *type* of boolean-valued expressions, `bool`. Lem, similarly to Isabelle/HOL and HOL4, collapses these two notions into a single type, `bool`. This mismatch between the languages causes difficulties, most notably in how we handle equality, as we have seen, but also in how we handle inductive relations and lemmata.

Take, as an example, the following Lem inductive relation:

```
indreln [even: nat → bool]
even_zero: true ⇒ even 0
and even_plus: ∀n. even n ⇒ even (n + 2)
```

This is translated to an ordinary Coq inductive type residing in `Prop` (as all inductive types in Coq must reside in a sort):

```
Inductive even: nat → Prop :=
| even_zero: true → even 0
| even_plus: ∀n, even n → even (n + 2).
```

Here the premises of the introduction rules of the inductive relation (for example, `true` in `even_zero`) are of boolean type, whereas they need to inhabit `Prop`. We use a Coq coercion from the Coq `bool` type into `Prop` to circumvent this problem: a function of type `bool → Prop` declared as a coercion automatically lifts a boolean expression into `Prop`. A similar problem occurs with lemma statements, which reside in `bool` in Lem but must reside in `Prop` in Coq.

However, problems persist elsewhere. One cannot perform a case analysis on `Prop` within Coq. However, Lem allows one to perform case analyses with `if`- and `case`-expressions on expressions which reside in `Prop` in the generated code. For example, Lem allows users to perform a case analysis on inductive relations:

```
let odd n = if even n then false else true
```

In the generated Coq code, the term `even n` has type `Prop` when it is expected to have type `bool` due to its position in the `if`-expression. There are four possibilities here:

1. Make this an error in Lem. However for backends like Isabelle/HOL and HOL4 the definition above is completely innocuous and rejecting it would be too restrictive.
2. Make use of Coq’s generalised `if-then-else` notation (supporting any inductive type with exactly two constructors), include the `sumbool` type usually used to capture decidability, and attempt to automatically show decidability for inductive types such as `even` above. The above would then be translated to:

```
let odd n = if even_dec n then false else true
```

where `even_dec` is a decidability theorem of type $\forall n. \{ \text{even } n \} + \{ \neg \text{even } n \}$. However, this approach becomes much more involved if we use inductive relations in more complex ways within Lem, for example, in a list of booleans: `[4 < 5, even 4, false]`, or in some other complex expression, and it is not clear how well it will scale.

3. Enrich Lem’s type system to identify a computational sublanguage, at the cost of significant complication.
4. Admit classical axioms—known to be consistent with Coq’s logic—and collapse `Prop` into `bool`.

We ultimately adopt the last alternative, with a function `bool_of_Prop` of type `Prop → bool`, wrapping this function around any propositional term that is being used in a way where a boolean-typed term is expected. The admission of classical axioms to collapse `Prop` into `bool` is a matter of taste. Some large Coq developments happily assume classical axioms, others stay firmly within the existing constructive logic provided by Coq. We feel, however, that not restricting the Lem source language to accommodate every nuance exhibited by the backends is worth the admission of these axioms.

4. Library mechanisms and design

Library design The Lem distribution supplies a default set of types and functions in its library, focussed on specification. Collections such as lists, sets and maps, basic data types such as disjoint sums, optional types, booleans and tuples, useful combinators on functions, and a library for working with relations are all included.

Specific function names and types exposed to the Lem user by the library are adopted from the Haskell standard library where possible. This is a well-designed library with a focus on purity. We also wish to provide flexibility in the specific choice of *backend* libraries function names and types in the Lem library are mapped to. Often, picking one library over another involves a trade-off between competing factors, with no clear winner. We therefore made it possible for users

to change and extend the library as they see fit, including replacing it wholesale. Lem library files are standard source files, their only distinguishing feature being their inclusion in the Lem distribution. No ‘prelude’ or ‘pervasive’ environment is automatically loaded during Lem compilation.

Lem aims to accommodate both programming languages and proof assistants as backends. Further, we aim to support users who wish to target a subset of these backends, or all of them, favouring neither one nor the other. As a result, the Lem library must be suitably flexible in its design. Our design philosophy in the library is *permit partiality and under-specification, but isolate them*. We bifurcate the library into two sets of modules: the ‘main’ and ‘extra’ modules. The main hierarchy of files contain total, terminating functions that we believe are well-specified enough to be portable across all backends. All other functions are placed in the extra modules. For example, the library file `function.Lem` includes various useful combinators such as `flip` and `const`. The `function_extra.Lem` file, on the other hand, contains the constant `THE` with type $\forall \alpha. (\alpha \rightarrow \text{bool}) \rightarrow \text{maybe } \alpha$, inexpressible in Coq. This design means there is always a conscious decision made on the part of the user to import functionality that assumes choice or exhibits partiality into their development.

Technical mechanisms Proof assistants provide a large body of facts about data types such as lists and numbers. Often, these facts are bundled together into simplification procedures for use in proof automation. In line with our goal of producing idiomatic backend code, we would like to map data types and functions in our Lem source to their corresponding implementations in the proof assistant, rather than generate our own copies, so that those facts and simplification procedures can be used. To this end, Lem features an array of tools for binding Lem functions and types to existing functions and types in the backends. For example

```
declare ocaml target_rep type set = 'Pset.set'
```

declares that Lem sets should be represented in OCaml by the existing OCaml-type `PSet.set`. Similarly, constants and functions can also be mapped to existing target representations:

```
val snoc : ∀α. α → list α → list α
let snoc e l = l ⊕ [e]
declare hol target_rep function snoc = 'SNOC'
let inline {isabelle;coq} snoc e l = l ⊕ [e]
```

Here, we introduce a Lem constant `snoc` and provide a Lem definition for it. For the HOL4 backend, `snoc` is mapped to the `SNOC` function from the HOL4 list-library. Coq and Isabelle/HOL do not provide their own native constants and this operation is expressed idiomatically using list concatenation. Lem’s inlining mechanism replaces all occurrences of the `snoc` constant with the append of a singleton list in the generated Isabelle/HOL and Coq code. Finally, for all other backends that Lem supports (i.e. OCaml in the example), a default implementation of the function is generated. Note no inlining occurs here—the call to `snoc` is preserved, with the resulting OCaml code looking similar to the Lem source code.

The Lem target-representation and inlining mechanisms are powerful enough to ‘smooth over’ inconsistencies between the backends. For instance, folds over lists display a surprising variety in the order in which arguments are expected. Our mechanisms allow us to provide a consistent interface within Lem for functions such as these whilst mapping to idiomatic backend code:

```
declare hol target_rep function foldr = 'FOLDR'
declare ocaml target_rep function foldr f b l =
  'List.fold_right' f l b
```

Target representations can also declare a constant as infix for certain backends. For example, the set-membership constant is mapped as follows:

```
declare ocaml target_rep function member = 'Pset.mem'
declare hol target_rep function member = infix 'IN'
declare html target_rep function member = infix '&isin;'
```

This function is prefix for OCaml but infix for HOL4. We could also provide additional associativity and binding strength information about infix constants in order to avoid generating superfluous parenthesis. Note that the user can also provide HTML and LaTeX target representations and that target representations are not restricted to valid Lem identifiers.

Unit testing through lemmata The Lem library attempts to clarify the semantics of each function by providing a definition, even if there are target-specific representations for all targets. Moreover, assertions and lemma statements can be used to describe the supposed behaviour of library files. As described in §3.11, assertions are executable tests used for unit testing. They generate executable testing code for OCaml. For the theorem prover backends they generate proof obligations which are (mostly) closed automatically by the prover’s computation mechanisms. Lemmata are non-executable tests. They are ignored by the OCaml backend, while the theorem prover backends generate proof obligations that need to be discharged manually by the user. For the `snoc` example, the library contains the following assertions and lemmata:

```
assert snoc_1 : snoc (2:nat) [] = [2]
assert snoc_2 : snoc (2:nat) [3;4] = [3;4;2]
lemma snoc_length :
  ∀ e l. length (snoc e l) = succ (length l)
```

If both a definition and a target-specific representation are present, Lem automatically generates a lemma that the target representation satisfies the definition. For example, the following lemma is automatically generated for the HOL4 backend:

```
lemma snoc_def_lemma: ∀ l e. (l ⊕ [e]) = (SNOC e l)
```

Whilst we do not aim to completely describe the semantics of every function in the Lem library with assertions and lemmata, we believe this peppering of executable checks and proof obligations provides some assurance that each of the bindings in the respective backends has the intended semantics.

5. Example

In this section we show excerpts from one of our major Lem developments, the C/C++11 axiomatic concurrency model of Batty et al., highlighting especially where any target-specific features of Lem were necessary. The complete specification is around 2500 lines of Lem source (1500 non-comment), defining several related models. The generated OCaml forms the kernel of our `cpmmem` tool, with a web interface (also using the HTML) for interactive and exhaustive exploration of the model on small examples, the generated LaTeX is used in papers and C/C++11 standards committee working notes, and the generated HOL4 code is the basis for mechanised proofs. The generated Coq and Isabelle code has also been used for (relatively minor) proof experiments.

The specification begins by opening standard Lem libraries:

```
open import Pervasives
```

and then proceeds by defining the types of action identifiers, C/C++11 memory orders, and the memory actions of the model. Action identifiers are taken to be strings for convenience in the `cpmmem` user interface; for the meta-theory we just need a type with a decidable equality and infinitely many inhabitants.

```
type aid = string
```

```
type memory_order =
  | NA          | Seq_cst | Relaxed | Release
```

```

| Acquire | Consume | Acq_rel

type action =
| Load of aid × tid × memory_order × location × cvalue
| Store of aid × tid × memory_order × location × cvalue
| Fence of aid × tid × memory_order
| ...

```

In the usual ‘axiomatic memory model’ style, the model is expressed as predicates over a notion of *candidate execution*, comprising a set of actions and various relations over them that describe one execution which might or might not be permitted by the model. Some of those are collected in the following record type:

```

type pre_execution =
<| actions : set (action);
   threads : set (tid);
   lk      : location → location_kind;
   sb      : set (action × action) ;
   asw     : set (action × action) ;
   dd      : set (action × action) ;
|>

```

For the HOL4 meta-theory, we want to consider possibly infinite candidate executions, therefore the representation of Lem sets as the usual idiomatic-HOL4 characteristic functions is appropriate. In the OCaml tool, we will only deal with finite candidate executions, and so the idiomatic, balanced binary tree representation is appropriate there. The sets used here all contain only elements of concrete types (for which the OCaml pervasive comparison will be correct), so no user instantiation of the SetType type class is required.

The specification continues with some routine functions defined by pattern matching, for example:

```

let is_at_non_atomic_location lk a =
  match loc_of a with
  | Just l ⇒ (lk l = Non_Atomic)
  | Nothing ⇒ false
end

```

The bulk of the specification consists of definitions of derived relations and predicates over candidate executions. For example, given a set of actions and a *happens-before* relation *hb*, the following picks out the write-read pairs in *hb* for which the write is a *visible side effect* (in the terms of the C/C++11 standards) for the read. It uses a Lem set comprehension, ranging over the supplied *hb* (which will be known and finite at execution time in the generated OCaml). The body of the comprehension uses standard propositional logic and an existential quantifier, here bounded by the set of actions.

```

let visible_side_effect_set actions hb =
{ (a,b) | forall ((a,b) IN hb) |
  is_write a && is_read b && (loc_of a = loc_of b) &&
  not ( exists (c IN actions). not (c IN {a;b}) &&
    is_write c && (loc_of c = loc_of b) &&
    (a,c) IN hb && (c,b) IN hb) }

```

The Lem-typeset version of this definition can be used simply by including

```
\LEMvisibleSideEffectSet
```

in a LaTeX source file (after including the Lem-generated definitions with `\usepackage{Lem}` and `\include{Cmm-inc}`), to give:

```

let visible_side_effect_set actions hb =
{ (a, b) | ∀ (a, b) ∈ hb |
  is_write a ∧ is_read b ∧ (loc_of a = loc_of b) ∧
  ¬ ( ∃ c ∈ actions. ¬ (c ∈ {a, b}) ∧
    is_write c ∧ (loc_of c = loc_of b) ∧
    (a, c) ∈ hb ∧ (c, b) ∈ hb) }

```

Note the preservation of line breaks and indentation, giving fine control of the typeset layout. The generated HOL4 in this case is

almost a direct transcription of the Lem source, as HOL4 supports similar set comprehension forms. The generated OCaml is somewhat more complex, using Pset library folds, membership tests, and set constructors, all with a suitable comparison function.

These excerpts are representative of typical Lem usage: mathematically straightforward higher-order logic definitions, with the real content in the details of exactly what the defined relations and predicates say. Though this development does not use inductively defined relations or recursively defined data types and functions extensively, others do, in a straightforward manner. In the previous sections we discussed many aspects of the Lem design that make it possible to generate code for all the targets in general, but it is common for only a few of those to be in play at once.

For this development the only target-specific definitions are as follows. First, there are those encapsulated in the Lem libraries used: `Basic_classes`, `Bool`, `Maybe`, `Num`, `Set`, `List`, and `Relation` (this development does not use `Tuple`, `Either`, `Function`, `Map`, `String`, `Word`, or `Sorting`, or any of the ‘extra’ libraries).

We give a general Lem definition for a predicate `strict_total_order_over` but in the HOL4 backend we bind it to the equivalent HOL4 standard library definition:

```

declare hol target_rep function strict_total_order_over
= 'strict_linear_order'

```

For Coq we define one measure function in Lem that is used in a hand proof of the termination of a recursive evaluation function over a type of trees of lists of named predicates:

```

type named_predicate_tree =
| Leaf of (complete_execution → bool)
| Node of list (string × named_predicate_tree)

```

To handle infinitary executions in the meta-theory, the definitions include assumptions that certain relations of a candidate execution, e.g. the coherence order, have finite prefixes. This is given a Lem definition for Isabelle/HOL, mapped to a HOL4 standard library function (in the same manner as above), and mapped to just a constant true for OCaml and Coq, which are both dealing only with the finitary case (using the default Lem set representations in each).

Then there are a number of definitions which are used only in the meta-theory, not in the tool, for which we do not produce OCaml output. This includes the top-level definition of each model (the 15 behaviour predicates), which are parameterised on a thread-local semantics (whereas in the tool we use a fixed thread-local semantics). The meta-theory also uses a receptiveness assumption and 11 other auxiliaries.

The amount of target-specific code needed in other developments is similarly small. For example, in CakeML there are 10 HOL4-specific source lines, introducing 5 HOL4 library functions that are not in the Lem library.

This C/C++11 development is also typical in that it is the product of a considerable investment of effort (multiple person-years by several people, of dialogue with the standards committee, experimentation, and proof) and that it is a specification that we need to maintain over an extended period of time. Manually propagating changes from versions in one target to another would have been prohibitive.

The Lem design is focussed on the reuse of a specification in multiple targets, not on reuse of parts of one specification in another, but the latter is also an important question. In particular, here the metatheory by Batty is over the model as above, but for a different purpose we have recently extended the model, adding new constructors of the action type and corresponding new function and predicate clauses. At present we do this (while avoiding forking the main specification file) by ad hoc means, but ideally one might want a mixin-style module composition.

6. Conclusion

Lem provides a new alternative for building large-scale semantic models and specifications, combining the uniform language design and ease of use of a good programming language with the definitional expressiveness provided by a theorem prover, and supporting *portable* definitions. It is more general-purpose than existing specification languages like K, Ott, or PLT Redex. Using Lem inevitably imposes some restrictions compared with working natively in a single prover but offers some advantages even in that case, and for modelling/specification exercises where the model creation and validation effort is large, the prospect of portability is compelling. It is demonstrably flexible enough to naturally specify a wide range of large-scale models while also allowing users to use their preferred tools for proof. That said, it is not perfect (of course), and we would like to revisit some aspects of the design with the benefit of hindsight: the type class mechanism, bool vs prop, multiple prover set representations, and inductive relation code generation.

Lem is a higher-order, typed, functional language, as are all of the backends that Lem currently targets. We anticipate that targeting new languages that fall into this pattern, for example Haskell, SML or Matita, would be straightforward.

Lem's syntax and type system are formally defined, but its logical semantics is defined by the translations into the targets. We attempt no formal guarantee that the result of translating into one target has the same mathematical meaning as that of translating into another, and indeed sometimes they intentionally do not (§3.9). Even when intuitively they do, stating and proving that fact would require creating, as a starting point, formal models of the semantics of the various targets, including Coq's underlying type system, Isabelle's datatype package, HOL's inductive relations package, etc. That would be worthwhile and challenging, but ours here is a first, pragmatic, goal: to support the working specifier.

Acknowledgements We thank Thomas Tuerk for his contributions to many aspects of Lem, Thomas Williams for his work on the Lem inductive relations package, Ohad Kammar for his work on model porting, and all the users of Lem for their feedback. We acknowledge funding from EPSRC grants EP/H005633 (Leadership Fellowship, Sewell) and EP/K008528 (REMS Programme Grant).

References

- [1] A. Asperti, C. Sacerdoti Coen, E. Tassi, and S. Zacchiroli. User interaction with the Matita proof assistant. *J. Autom. Reason.*, 2006.
- [2] L. Augustsson. Compiling pattern matching. In *Functional Programming Languages and Computer Architecture*, LNCS 201. 1985.
- [3] M. Batty, K. Memarian, S. Owens, S. Sarkar, and P. Sewell. Clarifying and compiling C/C++ concurrency: from C++11 to POWER. In *Proc. POPL*, 2012.
- [4] M. Batty, S. Owens, S. Sarkar, P. Sewell, and T. Weber. Mathematizing C++ concurrency. In *Proc. POPL*, 2011.
- [5] S. Berghofer, L. Bulwahn, and F. Haftmann. Turning inductive into equational specifications. In *Proc. TPHOLs*, 2009.
- [6] S. Bishop, M. Fairbairn, M. Norrish, P. Sewell, M. Smith, and K. Wansbrough. Rigorous specification and conformance testing techniques for network protocols, as applied to TCP, UDP, and Sockets. In *Proc. SIGCOMM*, 2005.
- [7] M. Bodin, A. Charguéraud, D. Filaretto, P. Gardner, S. Maffei, D. Naudziuniene, A. Schmitt, and G. Smith. A trusted mechanised JavaScript specification. In *Proc. POPL*, 2014.
- [8] C. Ellison and G. Rosu. An executable formal semantics of C with applications. In *Proc. POPL*, 2012.
- [9] A. C. J. Fox and M. O. Myreen. A trustworthy monadic formalization of the ARMv7 instruction set architecture. In *Proc. ITP*, 2010.
- [10] S. Goel, W. A. Hunt Jr., and M. Kaufmann. Abstract Stobjs and their application to ISA modeling. In *Proc. ACL2 Workshop*, 2013.
- [11] M. J. C. Gordon, J. Reynolds, W. A. Hunt Jr., and M. Kaufmann. An integration of HOL and ACL2. In *Proc. FMCAD*, 2006.
- [12] J. Hurd. The OpenTheory standard theory library. In *NASA Formal Methods*, LNCS 6617, 2011.
- [13] C. Kaliszyk and A. Krauss. Scalable LCF-Style proof translation. In *Proc. ITP*, LNCS 7998, 2013.
- [14] C. Keller and B. Werner. Importing HOL Light into Coq. In *Proc. ITP*, LNCS 6172, 2010.
- [15] C. Klein, J. Clements, C. Dimoulas, C. Eastlund, M. Felleisen, M. Flatt, J. A. McCarthy, J. Rafkind, S. Tobin-Hochstadt, and R. B. Findler. Run your research: on the effectiveness of lightweight mechanization. In *Proc. POPL*, 2012.
- [16] G. Klein, K. Elphinstone, G. Heiser, J. Andronick, D. Cock, P. Derrin, D. Elkaduwe, K. Engelhardt, R. Kolanski, M. Norrish, T. Sewell, H. Tuch, and S. Winwood. seL4: formal verification of an OS kernel. In *Proc. SOSP*, 2009.
- [17] R. Kumar, M. O. Myreen, M. Norrish, and S. Owens. CakeML: A Verified Implementation of ML. In *Proc. POPL*, 2014.
- [18] X. Leroy. A formally verified compiler back-end. *Journal of Automated Reasoning*, 43(4):363–446, 2009.
- [19] J. Lim and T. Reps. TSL: A system for generating abstract interpreters and its application to machine-code analysis. *TOPLAS*, 35(1), 2013.
- [20] S. Mador-Haim, L. Maranget, S. Sarkar, K. Memarian, J. Alglave, S. Owens, R. Alur, M. M. K. Martin, P. Sewell, and D. Williams. An axiomatic memory model for POWER multiprocessors. In *CAV*, 2012.
- [21] Luc Maranget. Compiling pattern matching to good decision trees. In *Proc. Workshop on ML*, 2008.
- [22] José Meseguer. Twenty years of rewriting logic. In *Proc. WRLA, WRLA'10*, 2010.
- [23] G. Morrisett, G. Tan, J. Tassarotti, J.-B. Tristan, and E. Gan. RockSalt: better, faster, stronger SFI for the x86. In *Proc. PLDI*, 2012.
- [24] S. Owens. A sound semantics for OCaml light. In *Proc. ESOP, LNCS 4960*, 2008.
- [25] S. Owens, P. Böhm, F. Zappa Nardelli, and P. Sewell. Lem: A lightweight tool for heavyweight semantics. In *Proc. ITP, LNCS 6898*, pages 363–369, 2011. “Rough Diamond” section.
- [26] J. G. Politz, M. J. Carroll, B. S. Lerner, J. Pombrio, and S. Krishnamurthi. A tested semantics for getters, setters, and eval in JavaScript. In *Proc. DSL*, 2012.
- [27] Tobias Rittweiler and Florian Haftmann. Haskell — converting Haskell source files to Isabelle/HOL theories. <http://isabelle.in.tum.de/haskell.html>.
- [28] G. Roşu and T. F. Şerbănuţă. An overview of the K semantic framework. *J. Logic and Algebraic Programming*, 79(6):397–434, 2010.
- [29] S. Sarkar, K. Memarian, S. Owens, M. Batty, P. Sewell, L. Maranget, J. Alglave, and D. Williams. Synchronising C/C++ and POWER. In *Proc. PLDI*, 2012.
- [30] S. Sarkar, P. Sewell, J. Alglave, L. Maranget, and D. Williams. Understanding POWER multiprocessors. In *Proc. PLDI*, 2011.
- [31] P. Sewell, S. Sarkar, S. Owens, F. Zappa Nardelli, and M. O. Myreen. x86-TSO: A rigorous and usable programmer's model for x86 multiprocessors. *C. ACM*, 53(7):89–97, 2010.
- [32] P. Sewell, F. Zappa Nardelli, S. Owens, G. Peskine, T. Ridge, S. Sarkar, and R. Strniša. Ott: Effective tool support for the working semanticist. *J. Funct. Program.*, 20(1):71–122, 2010.
- [33] M. G. J. van den Brand, A. Deursen, J. Heering, H. A. de Jong, M. Jonge, T. Kuipers, P. Klint, L. Moonen, P. A. Olivier, J. Scheerder, J. J. Vinju, E. Visser, and J. Visser. The ASF+SDF meta-environment: A component-based language development environment. In *Proc. LDTA, ENTCS 44*, 2001.
- [34] J. Ševčík, V. Vafeiadis, F. Zappa Nardelli, S. Jagannathan, and P. Sewell. CompCertTSO: A verified compiler for relaxed-memory concurrency. *J. ACM*, 60(3):22:1–22:50, June 2013.
- [35] D. Vytiniotis, S. L. Peyton Jones, T. Schrijvers, and M. Sulzmann. OutsideIn(X) modular type inference with local assumptions. *J. Funct. Program.*, 21(4-5):333–412, 2011.
- [36] P. Wadler. Views: a way for pattern matching to cohabit with data abstraction. In *Proc. POPL*, 1987.
- [37] F. Wiedijk. Encoding the HOL Light logic in Coq, 2007. Note.
- [38] J. Zhao, S. Nagarakatte, M. M. K. Martin, and S. Zdancewic. Formalizing the LLVM intermediate representation for verified program transformations. In *Proc. POPL*, 2012.