

Functional Programming Languages for Verification Tools: A Comparison of ML and HASKELL

Martin Leucker^{1*}, Thomas Noll², Perdita Stevens³, Michael Weber^{2**}

¹ Department of Computer Systems, Uppsala University, Box 337, SE-75105 Uppsala, Sweden
e-mail: leucker@docs.uu.se

² Lehrstuhl für Informatik II, RWTH Aachen, Ahornstr. 55, Aachen, Germany
e-mail: {noll,weber}@i2.informatik.rwth-aachen.de

³ School of Informatics, University of Edinburgh, JCMB, King's Buildings, Mayfield Road, Edinburgh EH9 3JZ, UK
e-mail: Perdita.Stevens@ed.ac.uk

Received: October 23, 2002/ Revised version: date

Abstract. We compare HASKELL with ML as programming languages for verification tools, based on our experience developing the verification platform TRUTH in HASKELL and the Edinburgh Concurrency Workbench (CWB) in ML. We discuss not only technical language features but also the “worlds” of the languages, for example, the availability of compilers, tools, and libraries.

1 Introduction

Concurrent software and hardware systems play an increasing role in today's applications. Due to the large number of states and to the high degree of non-determinism arising from the dynamic behavior of such systems, testing is generally not sufficient to ensure the correctness of their implementation. Formal specification and verification methods are therefore becoming more and more popular, aiming to give rigorous support for the system design and for establishing its correctness properties, respectively (cf. [1] for an overview).

In view of the inherent complexity of formal methods it is desirable to provide the user with tool support. It is even indispensable for the design of safety-critical concurrent systems where an *ad hoc* or conventional software engineering approach is not justifiable. For example, *model checking* is a particularly successful automated approach to verification, in which one employs decision procedures to prove that (a model of) a system has certain properties specified in a suitable logic.

One major concern in the development of model checking tools and other verification tools is *correctness*.

Since a verification tool is employed for verifying hardware, protocols, and software, it would be useless if it were not trustworthy. Thus, a programming language employed for developing a verification tool should support the task of verifying the code, at least in the informal sense of satisfying oneself that the code is correct. In general, *functional languages* are often considered to provide this feature, especially if the language disallows *side effects*.

Considering the broad range of functional languages which have been designed in the past, it is surprising, and in our view unfortunate, that there is little literature comparing different functional languages. The Pseudoknot benchmark paper [3] studies several implementations of functional programming languages (HASKELL and ML among them) with respect to their runtime and memory performance, and [7] compares the module systems of HASKELL and ML. However, the developer trying to choose between the languages needs to be concerned about a much wider class of issues, including both technical language features and “environmental” aspects such as the availability of libraries, documentation, support and multi-platform compilers. We did not find good sources of help for a developer trying to choose between the languages based on a larger collection of relevant aspects like this. By contrast, many comparisons of Java with C++ are readily available.

It is unsurprising, therefore, that developers (those who choose a functional language at all) often choose the language most used in their institution, without seriously considering alternatives. The difficulty of getting information to guide an informed choice may also contribute to developers, whose workplaces do not have a history of functional programming language use, deciding against experimenting with one.

Why is there so little material to help developers make an informed choice? Part of the reason must be that it is very hard to do convincing comparisons of lan-

* supported by European Research Training Network “Games”

** supported by the graduate college “Software für Kommunikationssysteme”

guages without being vulnerable to the criticism that one is not comparing like with like. We think that a fair comparison needs to be based on real experience of people using the languages to build real systems in the same domain; otherwise it is almost impossible to be sure that differences are not due to differences in the domains of application. The systems themselves need to be broadly comparable in size and complexity, need to be more than toys, and should preferably have been developed and maintained over years (since a language that makes development easy might nevertheless encourage the development of code which is unmaintainable). Moreover, the domain should be one where either of the languages is a reasonable choice, and the comparison should be done by people with a reasonably typical level of experience in the languages. A comparison is probably most generally useful to developers when it is done neither by novices in the languages compared nor by people intimately familiar with the compiler internals.

This paper recounts our experiences in using ML and HASKELL for two broadly comparable applications in the domain of verification tools, on which the authors have worked for some years: the Edinburgh Concurrency Workbench (CWB), in ML, and the verification platform TRUTH, in HASKELL. The domain of *verification tools* is eminently suited to the use of a statically typed functional language such as ML and HASKELL, and both languages are popular choices in this domain. All of the authors have accumulated considerable experience with the languages we use, but we are not functional programming researchers.

Thus our primary motivation for writing this paper is that we believe we are in an unusually good position to produce a comparison of HASKELL with ML which may be useful to developers choosing between the language. A secondary motivation is to be helpful to language designers and developers who work to support languages, by providing a record of our experiences, good and bad, with ML and HASKELL.

The rest of this paper is structured as follows. Section 2 discusses the domain of verification tools, and introduces the two systems. Section 3 discusses the class of languages from which ML and HASKELL are drawn, and briefly introduces the two languages for readers who may not be familiar with them. Sections 4 and 5 are the main body of the paper; Section 4 compares HASKELL and ML on the basis of their technical language design features, whereas Section 5 considers the equally important “environment” aspects. Finally, Section 6 concludes.

2 Verification tools

The domain on which we compare ML and HASKELL is that of verification tools. The term “verification tool” covers any tool whose task it is to assist in checking the correctness of some artefact. Usually the artefact

concerned is (an abstraction of) something produced in the software or hardware development process.

We introduce TRUTH and the CWB and briefly summarise their histories, before discussing the characteristic features of verification tools in general.

2.1 The Edinburgh Concurrency Workbench, in ML

Work on the CWB¹ began in 1986. The CWB’s key strength is its breadth: a variety of different verification methods are supported for several different process algebras. In particular, it allows users to:

- define behaviours given either in an extended version of Milner’s CCS (Calculus of Communicating Systems; [5]) or in its synchronous version, SCCS, and to perform various analyses on these behaviours, such as analysing the state space of a given process, or checking various semantic equivalences and pre-orders;
- define propositions in a powerful modal logic and check whether a given process satisfies a property formulated in this logic;
- play Stirling-style model-checking games to understand why a process does or does not satisfy a formula;
- derive automatically logical formulae which distinguish nonequivalent processes;
- interactively simulate the behaviour of an agent, thus guiding it through its state space in a controlled fashion.

One major focus of the CWB was always research; it was a platform which researchers (especially those at Edinburgh) could use to experiment with new relations between processes and new algorithms. In the early years all of these changes were retained in the main tool, even those which had been added experimentally without much consideration for the integrity of the CWB overall. This contributed to the architectural degradation of the CWB and its increasing fragility: an important task faced by Stevens on taking over the maintenance of the CWB in 1994 was to reverse this process. The current version of the CWB consists of around 25kloc in ML, plus several thousand in other languages for various supporting utilities.

The CWB was developed in Standard ML, but variations were long maintained for several major ML compilers because different compilers provided different extensions to the SML’90 standard, and especially because they had different system build facilities. We settled on Standard ML of New Jersey (SML/NJ) because most users of the CWB used that compiler and the effort in maintaining build scripts (the major point of difference) for several compilers did not seem well spent. Perhaps we should once again target Poly/ML², for exam-

¹ <http://www.dcs.ed.ac.uk/home/cwb/>

² <http://www.polym1.org/>

ple, in future. This paper, however, inevitably considers SML/NJ more than any other ML compiler. The contribution made by the language to both the architectural degradation problem and its solution are discussed later.

2.2 The verification platform TRUTH, in HASKELL

In terms of features, TRUTH³ is similar to the CWB. In its current version, it supports tableau-based model checking for the full μ -calculus and game-based model checking for the alternation-free subcalculus. Both operate on finite transition systems, given in terms of CCS processes. The latter can be visualized and simulated in an interactive fashion, to help the user understand TRUTH's answers. Current development activities concentrate on the parallel implementation of model checking on a cluster of workstations and on a specification language compiler generator which, given the definition of a language, automatically generates a corresponding parser and a semantic evaluator.

TRUTH's initial version dates back to 1997, and its development could benefit a lot from the progression made in the design of verification tools over the years. As a consequence, its architecture is quite modular and easy to understand, and a deep change of the module structure was not necessary so far. It now consists of approximately 18kloc in HASKELL. Although there are several HASKELL compilers, TRUTH is written for the Glasgow HASKELL Compiler⁴ (GHC), and since it uses some non-standard HASKELL extensions and libraries only present in the GHC, we have not tried to port it.

Moreover TRUTH employs the parser generator HAPPY⁵ and many of the available HASKELL libraries, and it integrates several standalone systems such as the DAVINCI⁶ graph visualization tool and the GraphViz package⁷. Furthermore it uses existing C and Java libraries to provide functionality such as textual and graphical user interfaces and network communication, comprising approximately 13000 lines of code. It is one of the bigger real-world applications registered in the official HASKELL pages⁸. It is worth mentioning that TRUTH is one of the few tools listed there which was developed *using* but not *for* functional programming.

2.3 Characteristics of verification tools in general

The peculiarities of the verification tool domain from the point of view of software engineering were considered by Stevens in [9]. Here we briefly summarise, and then focus on the implications for language choice.

Verification tools answer precisely defined questions about precisely defined systems. Thus it is comparatively easy to understand what it means for the tool's behaviour to be correct. The downside is that certain classes of bugs are unacceptable in a verification tool; semantic correctness is vital. Thus any language features supporting the development of correct programs are highly desirable.

A further characteristic is that verification tools *tend* to be developed in research environments, where it is easier to be recognised for novel theoretical contributions, or new applications of theory, than for the application of "best practice" in software engineering, which is likely to be discounted because it is not new. Anything that speeds up development is an advantage, as it enables the developers to spend a higher proportion of their time on the work which is most valued. In such environments, it is also difficult to justify spending large amounts of effort on academically uninteresting aspects of the tool, such as a GUI, or on "invisible" areas such as testing(!), documentation, and ensuring portability. Nevertheless, the usability and ultimately success of the tool depends heavily on such aspects. Therefore those planning to develop verification tools will do well to choose a language in which professional results can be achieved with a minimum of effort.

It is perhaps instructive to note that in some cases, the same considerations may apply to those developing languages and their associated tools.

3 The space of programming languages

Clearly HASKELL and ML, the languages of TRUTH and the CWB, have a great deal in common: both are basically functional languages, and both have static type systems which are *strong* in the sense that a well-typed program will be free of certain classes of run-time errors. Moreover, both are minority languages, with their origins in academia. What is the significance of these features for verification tools?

The functional paradigm Essentially, a functional programming language is one in which the natural programming style includes treating functions as first-class concepts. For example, one expects to write *higher order functions*; that is, functions which take other functions as arguments. There is, however, no universally agreed definition of what it is to be a functional programming language, although no reasonable definition would exclude either ML or HASKELL. The difficulty stems from the impure nature of most languages, which stems in turn from the need to permit the use of whichever paradigm is most appropriate for a particular problem. It is possible, for example, to write higher order functions in C; the reason why C is not included in definitions of functional programming language is that this is not the natural, normal way to solve problems in C.

³ <http://www-i2.informatik.rwth-aachen.de/Research/Truth/>

⁴ <http://www.haskell.org/ghc/>

⁵ <http://haskell.cs.yale.edu/happy/>

⁶ <http://www.davinci-presenter.de/>

⁷ <http://www.research.att.com/sw/tools/graphviz/>

⁸ <http://www.haskell.org/practice.html>

The main reason, in our view, for using a functional language for a verification tool is that the paradigm is a good match for the domain, as the most important concepts in the domain tend to be algorithms. It is often claimed that programs written in functional languages are easier to reason about, and hence are more likely to be correct, than those in one of the imperative paradigms (procedural or object oriented). The theoretical concept on which the claim rests is *referential transparency*, essentially the fact that identifiers are used for values, rather than for references whose values may change. Where this property holds, it can indeed facilitate reasoning, at least in small piece of code. However, we have found that in practice, building a verification tool in a way which provides reasonable modularity and efficiency necessitates the use of “impure” features of the languages, so that referential transparency is lost.

Today the most obvious alternative to the functional paradigm for a verification tool writer is the object oriented paradigm. The main argument in favour of the functional paradigm is that the most important concepts in the domain tend to be algorithms, rather than objects. In this respect the verification tool domain differs from most business domains, and the use of a less popular language may be justified. However, as we shall see, being out of the mainstream carries disadvantages sufficient to give one pause.

Static typing In a statically typed language, the compiler carries out certain checks to ensure that the program is free from certain type errors which might otherwise cause incorrect behaviour at runtime. This does not, of course, ensure that the program is free of errors, but it can enable errors to be caught early and easily corrected, thus speeding up the development process. Static typing is often criticised for being inflexible; but when such criticisms are investigated, they turn out to be criticisms of the inflexibility of a particular type system. We argue that a coherent understanding of a solution to a problem includes an understanding of the types of the entities involved; if these fit the type system of the language concerned, it is hard to see how having errors caught by the compiler can fail to be a benefit, although one could still argue about the size of the benefit.

It is clear from the successes achieved by certain groups working with dynamically typed languages such as Erlang (in the functional world) and Smalltalk (in the object oriented world) that it is possible to write complex, correct software without static typing. However, none of the authors would willingly give up the benefits of static typing. We will discuss particular features of the type systems of ML and HASKELL below.

3.1 ML

ML ([6]) is an essentially functional language in the sense discussed above. By “essentially” we mean that

it is not a *pure* functional language: for example, references are permitted. It originated at Edinburgh in the late 1970s. A major rewrite of the language definition resulted in the new definition of Standard ML, sometimes referred to as ML97; this is what we shall mean by ML in the rest of this paper. The clarification is needed because there is a family of ML-like languages, including most notably Caml and its variants such as O’Caml.

Technically, the revision has been a substantial improvement; but it has led, temporarily at least, to difficulties of tools and libraries not all being updated at once; old ML programs cannot be compiled by new compilers and vice versa.

A variety of compilers is available for Standard ML; by far the most widely used is Standard ML of New Jersey (SML-NJ), and this is the only compiler supported by the CWB.

The definition of Standard ML includes the Standard Basis library, providing such things as string manipulation, operating system interfaces and basic data structures. SML-NJ comes with a more extensive library.

3.2 HASKELL

HASKELL is a purely functional programming language ([8]). The current standard is HASKELL98 which fixes the syntax and semantics as well as a large set of standard libraries.

Until recently the embedding of input and output operations, which have to be considered as side effects, in purely functional programming languages was generally poor. *Monadic I/O* is a very elegant approach to overcome this problem ([11]). HASKELL supports this concept and supplies versatile I/O libraries offering exception handling and file manipulation operations, which were of great help in building a user-friendly and reliable tool.

4 Comparison of language design features

We begin by considering and comparing the more technical aspects of ML and HASKELL, before going on to consider non-technical questions in the next section.

4.1 Typing

As semantic correctness is crucial to any verification tool, it is natural to believe that a strong static type system, enabling a large class of errors to be caught at compile time, is a good thing in a language for verification tools. Our experience supports this; although verification tools have been written in Lisp, for example, we would not like to give up the static typing provided by both ML and HASKELL. To go further, we consider

two related features which HASKELL and ML have in common: parametric polymorphism and type inference.

Extensive type inference is convenient especially in functional programming where identifiers often have complex higher-order types. However, it has serious drawbacks for maintainability of code. The human reader of code needs to understand the types involved, and it is frustrating to know that the compiler has worked out information which is not available to the reader. The natural response is that good programming practice is then to include type annotations; but we have found this hard to put into practice. An annoyance is that the syntax of the languages and the presence of identifiers with complex types sometimes makes this awkward (ML) or impossible (HASKELL). A more serious point is that if type annotations are included which are descriptive enough to be helpful, their presence hinders reuse by parametric polymorphism. There is a tension between trying to enable code reuse on the one hand, and on the other hand trying to make code understandable and trying to maintain appropriate encapsulation barriers. We find that these last two, though different, often go together: one encapsulates the definition of an important type together with appropriate functions for manipulating it, and then uses the new type name in type annotations to elucidate the code. However, in doing so one loses the power of parametric polymorphism for code reuse in clients of this new type, because clients cannot see the structure of the type.

For example, processes in the process calculi we work with can have *restrictions* applied to them. A restriction is conveniently implemented as a list of actions, but certain invariants need to be maintained. If we allow clients to see that a restriction is a list of actions, then when they manipulate processes they can use the standard list functions on the restriction, but we cannot easily enforce the invariants. On the other hand if we use encapsulation to make available only a type *restriction* so that we can enforce the invariants we have to provide all necessary functions for manipulating this type. This is not unreasonable: it is the same work, for example, that we would have to do if we worked in an object oriented language and created a class *Restriction*. However, when we have a variety of slightly different kinds of restriction, we have to implement the manipulating functions afresh every time; to gain encapsulation we have lost parametric polymorphism as a reuse mechanism, and we do not have inheritance available to us as an alternative mechanism. This kind of situation arises very frequently in both the CWB and TRUTH, because we write code to deal with variants of process algebras and logics and with variously processed versions of them. An additional issue in ML is that it is sometimes difficult to decide whether a conceptual type should be implemented at the module level or only at the core level; in the CWB we generally resolve this by using both but not revealing that decision outside the module where it is made, so that, for

example, the signature for processes exports only a type *restriction* whereas an implementation of that signature typically builds a structure *Restriction*, exporting a type from the content of that structure.

Further, we find that the powerful type systems of ML and HASKELL are a mixed blessing, often leading to complex type errors which are understandable only to people who are familiar with the subtleties of the respective type system. Recent work on more informative error messages, such as [4], is to be welcomed, but has yet to make a difference to the compilers. Furthermore, an interactive *type analyzer* would be desirable, a tool which, requested by the user, would visualize the types of certain subexpressions. In the meantime, our advice is that there is little to choose between ML and HASKELL in this respect.

4.2 Strictness vs. laziness

The most obvious difference between ML and HASKELL is that ML is *strict* whilst HASKELL is *lazy*. For discussion of the concepts in general see for example [10]. Basically, laziness means that values are only computed on demand, allowing the implementation of infinite data structures. In contrast, strictness refers to the fact that e.g. the arguments of a function call have to be evaluated before executing the call, no matter whether they are required or not.

In the context of verification tools, laziness seems to be an appealing feature because one might hope to get “for free” certain “on-the-fly” verification techniques that normally have to be worked out in each special case. For example, consider a class of verification questions concerning a system, such as the model-checking problem “does this system satisfy this property”. To answer some questions in the class it will be necessary to calculate the entire state space of the system. For others, only a small part of the state space, perhaps that reachable within three transitions from the starting state, will be relevant. A *global* algorithm is one which always calculates the whole state space; a *local* one does not. Local algorithms are generally harder to design and verify than global ones, and often have poorer worst-case complexity, though in practice they may perform much better. One might hope to be able to get a local algorithm from a global one “for free” using laziness, because the code for calculating certain parts of the state space would simply never be evaluated if its results were not called for. In practice however the TRUTH team found that one has to implement the algorithm generating the state space carefully in order to guarantee the desired behaviour. For example, the use of monads (cf. Section 4.3) or of accumulator techniques can easily destroy the on-the-fly property. Altogether the effort which is required to preserve the locality of a lazily evaluated, global algorithm often corresponds to the design of an algorithm which is local by nature.

Summarizing, lazy evaluation is an attractive feature, but the TRUTH team would have liked a flexible mechanism with which to specify parts of the program which should be evaluated eagerly or lazily.

4.3 Imperative features

Both the TRUTH and the CWB team have found imperative features to be essential. Sometimes the concern is efficiency, but more often it is understandability: many of the algorithms we wish to implement are conceived imperatively and in such cases to implement them functionally makes the implementation more difficult to read and hence more likely to contain errors. Prominent examples of algorithms with an imperative character are graph algorithms, which play an important rôle in tools such as ours which deal with transition systems.

Here ML scores by providing imperative features in the core language in a reasonable and powerful way, although they can be syntactically awkward. I/O is supported by the Standard Basis Library. HASKELL uses *monads* for destructive updates and I/O; they add a restricted form of stateful computation to a pure language, retaining referential transparency ([11]). The disadvantage is that programs become more complicated. Also, if imperative elements of a given application were not taken into account during its design but turn out to be necessary later on, often major parts have to be redesigned or (at least) reimplemented, especially because types change significantly. This is clearly undesirable from a software engineering or economical point of view. Indeed for certain parts of the TRUTH system a redesign turned out to be necessary in the past, mostly in order to implement more efficient versions of algorithms by introducing imperative constructs. The TRUTH team considers this point as one of the biggest drawbacks of the purely functional paradigm as followed by HASKELL.

4.4 Architecture support

The architecture of a system makes a vital contribution to its correctness. We hope to study module systems in this context, building on [7], in future; in this paper we can only indicate the main issues.

A HASKELL *module* defines a collection of values, datatypes, type synonyms, classes, etc., as well as their import/export relationship with other modules. Overloaded functions are provided in a structured way in the form of *type classes*, which can be thought of as families of types (or more generally as families of tuples of types) whose elements are called *instances* of the class. In the instantiation the definitions of the overloaded operations are given.

In Standard ML, *structures* provide the main namespace management mechanism; they may contain substructures, functions, values, types, etc. A structure may

be coded directly, or produced by the application of a *functor*, which may be thought of as a generic or parameterized structure. The programmer may define *signatures* which act as the types for structures; for example, a functor may be defined to take, as argument, any structure matching a given signature. The module system is separate from the core language; one cannot, for example, apply a functor conditionally. Whilst this keeps the language definition clean, in the CWB it has often caused problems leading to code duplication. The changes made to ML in SML'97 are welcome; the elimination of structure sharing and the introduction of “where” clauses have solved several long-standing problems for the CWB.

A basic facility which is desirable in a module system is that it should be possible to define an interface to a module separately from the module itself. This helps developers to understand the system, as they can read interfaces to modules without being distracted by implementation information. We also want to be able to apply the same interface to several modules and to provide several interfaces to the same module. In both the CWB and TRUTH this need arises, for example, because we often work with several variants of a process algebra, logic or algorithm which share an interface. We want the compiler to do the work of making sure that the modules stay consistent, and we want to avoid duplicating code. ML's signatures support this way of working reasonably well, although not without problems. The TRUTH team has found that HASKELL does not support this situation so well: inside the module export list, entities cannot be annotated with types, so a common practice is to add them in comments. However, this is error prone, since there is no way for the compiler to enforce their correctness or check their consistency with the implementation in the module body.

We feel that ML's architectural features are better suited than HASKELL's to our purposes; neither is ideal, however, and this seems an interesting area for future study, especially as we do not think that the class and package systems of C++ or Java would be ideal either.

4.5 Exceptions

The CWB used to make heavy use of exceptions as a control flow mechanism. This led to correctness problems because the compiler cannot check whether exceptions are always handled or not. Thus a bug could lead to an exception rising to the top level, where nothing sensible could be done to handle it. To make matters worse in SML'90, although one could write a handler that would catch all exceptions, one could not then tell which exception was actually being handled. Therefore such bugs would result in a message to the user of the CWB along the lines of “Sorry, an ML exception has been raised. This is a bug: please report it.” The exception mechanism has been improved in SML'97 compared

with SML'90: it is now possible to interrogate an exception for its identity, which at least enables the CWB to give a fuller error message which is useful for debugging.

Still, in a language with type safety as a strength, it is a pity to use a programming style in which the programmer cannot be certain that all exceptions are handled. From the point of view of the user of a verification tool, it is not very much better for the application to terminate because of an unhandled exception than it would have been for it to terminate because of a run-time type error. Therefore the CWB now uses exceptions in a more disciplined way which seems to work well. A small number of specified exceptions (corresponding to such things as “error in user input”, “assertion violated”, etc.) are allowed to rise to the top level and are individually handled there. All other exceptions are kept within small pieces of code (e.g., within one ML module) and in each case the programmer verifies by eye that the exception cannot escape. Because the latter is hard work, exceptions are only used where the alternative is really painful.

The TRUTH team found that exceptions interacted badly with laziness: as a rather disturbing effect, partial evaluation enables exceptions to *escape* from an enclosing exception handler. In order to get the exceptions actually raised inside the handler, initially the TRUTH team had to resort to code like `if x==x then x else x` to enforce the evaluation of `x` at the right time. Recently, better ways to trigger full evaluation have been provided (`deepSeq $! x`), but they are non-standard and of course destroy laziness. We think it would be much more natural to avoid situations of this kind by adopting strict rather than lazy evaluation as the standard strategy in the language. Laziness, which is a very costly feature, could then be provided upon request, using annotations of the function and constructor symbols.

We have often seen programming languages compared on the basis of how many lines of code it takes to implement some piece of functionality. We consider this a poor metric. The length of a piece of code is not well correlated either with the time it takes to write it or with the time it takes to understand it; a short piece of code may well be harder to write and to maintain than a longer one. This is why we have not tried to compare ML and HASKELL on this point.

5 Comparison of non-language design features

5.1 The available compilers and their characteristics

There are now three main freely available SML'97 compilers, SML/NJ, Poly/ML and Moscow ML. (Harlequin MLWorks ceased to be available when Harlequin was bought: there was hope that it might become open source, but this now appears unlikely.) SML/NJ can now produce native code for many platforms, which is important for a widely-distributed verification tool. However,

it needs a third-party utility to produce stand-alone applications, and even then there is a problem with running the application from outside its directory. This is hard to explain to users of the CWB and causes embarrassment.

For HASKELL too, three compilers are available, all of them freely: NHC98, HBC, and GHC. For TRUTH, only GHC was considered feature-complete enough; it also provides some extensions to the HASKELL98 language which have proven helpful, for example multi-parameter classes and existential types.

5.2 Libraries and associated tools

Good libraries and tools can help to ensure correctness (e.g. because well-used libraries have been debugged by others) and can cut down development time. We consider and compare what is available for HASKELL and for ML.

General purpose libraries Both HASKELL and ML come along with standard libraries specified alongside the language. SML'97 defines the Standard Basis Library⁹ (ML97SBL); HASKELL98's libraries are described in the Library Report¹⁰ (H98LR). Broadly similar, these provide basic data structures, interface to the operating system etc. Both GHC and SML/NJ ship with some extra, non-standard libraries.

GUI libraries There is an X Window System toolkit, eXene¹¹, written in Concurrent ML, though for a long time this was apparently not usable with SML'97 (because of a signal-handling bug, fixed more recently than the last major CWB changes). Research projects have provided portable GUI library facilities for use with Standard ML, such as `sml.tk`¹². Thus one can implement a GUI in ML; but really good high-level toolkits are still lacking. The CWB has made no serious attempt to do this. For HASKELL too, some bindings for common GUI toolkits are available, but at the time GUI support was added to TRUTH, none of them was regarded as stable or feature-complete enough to be usable for what was planned. In the end, the process simulation GUI for TRUTH was written in Java and was interfaced to the HASKELL part via Unix pipes. (The CWB followed a similar path in a student project, as yet unreleased.)

Associated tools A debugger is invaluable in program development, especially when experimenting with verification algorithms which may contain bugs. Unfortunately, writing debuggers for functional languages turned out to be harder than for imperative languages like C. This is even more true for a lazily evaluated language like HASKELL, where the inspection of a value

⁹ <http://cm.bell-labs.com/cm/cs/what/smlnj/doc/basis/>

¹⁰ <http://www.haskell.org/definition/>

¹¹ <http://cm.bell-labs.com/cm/cs/who/jhr/sml/eXene/>

¹² http://www.informatik.uni-bremen.de/~cx1/sml_tk/

would sometimes change the evaluation order. Nevertheless there have been made some attempts in this direction, mostly resulting in so-called *tracers* (like Freya, Hood, or Hat), which can record program runs for later analysis. None of them were used during TRUTH development, because they were either not available at that time, or did not support some of the GHC features used in TRUTH. There is no debugger available for SML/NJ. There is, however, a debugger for Poly/ML, which is a welcome development. We have not yet used it.

There is a lexer (ML-Lex) and a parser generator (ML-Yacc) for ML. These were long unavailable in SML'97 versions, but do now seem to work (see below re documentation). The CWB uses ML-Lex but does not use ML-Yacc. At a very early stage a hand-built parser was produced, and by the time the major reengineering work was done on the CWB its syntax (perhaps unfortunately, but understandably) included features which were not supported by ML-Yacc, so that to move to ML-Yacc at that point would have involved a user-visible syntax change. This is an example of the problems that can arise when a suitable third-party component is not available at the right moment; users may not have the option of adopting it later. As mentioned, TRUTH uses the Happy parser generator.

Overall there is little to choose between HASKELL and ML in this category, but both suffer from being minority languages. There are few providers of libraries and tools, and key developers are often more concerned with compilers. This is understandable, but to us libraries and tools are just as important.

5.3 Documentation and other sources of help

Famously, Standard ML has a formal specification [6], but this is impenetrable to most programmers. Fortunately there are also several accessible books and tutorials available. The official specification of HASKELL is given by the HASKELL98 Language Report¹³ which defines the syntax of HASKELL programs and gives an *informal* abstract semantics. For such a technical document it contains much plain text and the general impression of local HASKELL developers was that it is quite readable. On the other hand, as was noted elsewhere¹⁴: “The informal specification in the HASKELL report leaves too much room for confusion and misinterpretation. This leads to genuine discrepancies between implementations, as many subscribers to the HASKELL mailing list will have seen.”

Regarding the compilers’ and associated tools’ documentation, the overall impression of the authors is that GHC’s documentation is slightly better than that of SML/NJ. (The ML-Lex documentation has not been updated for SML'97, for example.) This has not always

been the case, but the GHC developers have improved the documentation quite a lot in the recent past.

In both cases documentation for libraries is patchy, especially in the case of compiler-specific libraries, where it sometimes happens that the programmer must consult the source code to get more information than the signature of a function. The H98LR and the ML97SBL are better documented. From 1997 to 2001 there was no complete and up-to-date documentation of the latter, which was a serious problem. Now, however, an updated web page is available; the documentation will appear in book form [2] in 2003, which is a welcome development. The CWB and TRUTH teams each had the impression initially that the other’s language’s libraries were better documented: this may reflect that one notices faults only on close acquaintance. A plus for HASKELL is that the GHC library documentation has a consistent history of being frequently updated and improved.

Moving from documents to people as sources of help, we have found the newsgroups `comp.lang.ml` and `comp.lang.functional` and the GHC mailing lists to be useful. Naturally it is easier to get help with problems which can be described briefly. When we have needed help with, for example, making architectural decisions, local language experts have been invaluable; this is something that developers should bear in mind.

Last but not least the home pages of Standard ML of New Jersey¹⁵ and of HASKELL¹⁶ provide useful collections of links and references to other resources.

5.4 Foreign function interfaces

We have made no serious attempt to bind C and ML code within the CWB, because the Standard ML foreign function interfaces were perceived (and experienced in a student project) as hard to use and inefficient. Matthias Blume’s new “NLFFI” foreign-function interface may well change the situation.

The foreign function interface in HASKELL has undergone a major redesign and is now quite usable. TRUTH has been extended by a parallel model checking algorithm, which uses the FFI layer to call C functions from the MPICH resp. LAM libraries, both well-known implementations of the *Message Passing Interface* standard. For this application the *marshalling* required to convert between HASKELL and C data formats turned out to be very inefficient, however. Another problem was the instability of the FFI interface at the time the TRUTH team were using it: it changed rapidly between releases of GHC. The TRUTH team made extraordinary use of pre-processor directives and *autoconf* magic in an attempt to allow TRUTH to support many compiler versions, but were eventually forced to give up.

¹³ <http://www.haskell.org/onlinereport/>

¹⁴ <http://www.cse.ogi.edu/~mpj/thih/>

¹⁵ <http://cm.bell-labs.com/cm/cs/what/smlnj/>

¹⁶ <http://www.haskell.org/>

5.5 Stability of languages and their implementations

The current stable version of HASKELL is HASKELL98, dating from 1997. This is the fifth major version of the language definition (the next will be *Haskell 2!*). Compilers, of course, provide the effective definition of the language. There have been many changes in what GHC supports (e.g. multi-parameter classes, implicit parameters). Not all changes to extensions have been backwards compatible, which is inconvenient for programmers who need those extensions.

Regarding the stability of HASKELL *implementations*, only GHC has been thoroughly examined, since the other implementations have been ruled out by other issues, as stated before. GHC is under steady development and the quality of released versions differs greatly. Some are quite stable, but for others patchlevel releases have to be made quickly to fix the worst bugs. Unsurprisingly bugs often accompany new features. In fairness, bugs are fixed promptly by the GHC developers once they are reported to the relevant mailing list. However, faced with a show stopper, an application programmer must choose between waiting for an official release of GHC including a fix, or becoming expert in building the compiler itself, which is non-trivial, time consuming, and will not further his/her aims.

Standard ML was subject to a major revision, from ML'90 to ML'97. The SML/NJ compiler has undergone many releases, but now seems fairly stable. As mentioned, tools and libraries tend to lag. The ML2000 project intended to develop a future version of ML. Little has been heard of the project recently and many of its early ideas have been incorporated in O'Caml. We believe that SML'97 will increase, rather than decrease, in stability over the next few years.

5.6 Performance

This is a controversial topic, but it is an important one for developers of verification tools. Both speed and space usage are important, with space usage often being more important, as the amount of memory used by a verification is normally the limiting factor. Our experience with both ML and HASKELL suggests that performance is particularly poor with respect to memory usage. Also, as noted, the key feature of HASKELL, lazy evaluation, comes at a high cost.

6 Conclusion

There have been positive and negative aspects to both our sets of experiences with HASKELL and ML, as there would doubtless have been with whatever language we had chosen. Overall, we consider Standard ML to be a slightly better choice for our kind of application than

HASKELL, more because of a more stable environment of supporting tools than because of language features. Of course, there are many alternatives including other functional languages of which we have less experience; O'Caml might be a strong candidate. However, it turned out in our discussions that none of us were enthusiastic about the idea of using a functional language for a future verification tool, because of their impoverished environments compared with mainstream programming languages. Our conclusion is that if/when we develop new verification tools, these will be written in “a Java-ready subset of C++”. That is, we would prefer to be writing in Java, but at present this appears premature for performance reasons. We would write therefore in C++, trying to avoid using features (such as multiple inheritance, non-virtual functions, “friends”) which cannot be readily translated into Java.

We hope that reporting our experience using major functional languages will help the community to improve such languages and their worlds in future.

References

1. Edmund M. Clarke and Jeanette M. Wing. Formal methods: State of the art and future directions. *ACM Computing Surveys*, 28(4):626–643, December 1996.
2. Emden R. Gansner and John H. Reppy. *The Standard ML Basis Manual*. Cambridge University Press, September 2003. To appear.
3. Pieter H. Hartel et al. Benchmarking implementations of functional languages with ‘Pseudoknot’, a float-intensive benchmark. *Journal of Functional Programming*, 6(4):621–655, July 1996.
4. Bruce McAdam. *Repairing Type Errors in Functional Programs*. PhD thesis, Division of Informatics, University of Edinburgh, 2002.
5. Robin Milner. *Communication and Concurrency*. International Series in Computer Science. Prentice-Hall, 1989.
6. Robin Milner, Mads Tofte, Robert Harper, and David MacQueen. *The Definition of Standard ML (Revised)*. The MIT Press, 1997.
7. Jan Nicklisch and Simon L. Peyton Jones. An Exploration of Modular Programs. In *Glasgow Workshop on Functional Programming*, July 1996.
8. John Peterson, Kevin Hammond, et al. Report on the programming language Haskell, a non-strict purely-functional programming language, version 1.3. Technical report, Yale University, May 1996.
9. Perdita Stevens. A verification tool developer’s vade mecum. *Software Tools for Technology Transfer*, 2(2):89–94, 1999.
10. Philip Wadler. Lazy versus strict. *ACM Computing Surveys*, 28(2):318–320, June 1996.
11. Philip Wadler. How to declare an imperative. *ACM Computing Surveys*, 29(3):240–263, September 1997.