

Programming with Regions in the MLKit

REVISED FOR VERSION 4.7.22

Mads Tofte Lars Birkedal Martin Elsman
Niels Hallenberg Tommy Højfeldt Olesen Peter Sestoft

April 28, 2026



Values and their Representation

<code>int</code>	64 bits, untagged. Unboxed (i.e., not region allocated). One bit is used for tagging when GC is enabled.
<code>real</code>	64 bits, untagged. Boxed (i.e., allocated in region) or unboxed.
<code>string</code>	Unbounded size. Allocated in region.
<code>bool</code>	one 64-bit word. Unboxed.
<code>α list</code>	<code>nil</code> and <code>::</code> cells unboxed (i.e., not region allocated). Auxiliary pairs in one region; elements in zero or more regions. Size of auxiliary pairs: two 64-bit words (three when GC is enabled).
<code>exn</code>	Exception values are boxed and are always stored in a global region.
<code>fn <i>pat</i> \Rightarrow <i>exp</i></code>	An anonymous function is represented by a boxed, untagged closure. Its size is one 64-bit word plus one word for each free variable of the function. Free region variables also count as variables. One extra word is used when GC is enabled.
<code>fun <i>f</i>...</code>	Mutually recursive region-polymorphic functions share the same closure, which is region-allocated, untagged, and whose size (in words) is the number of variables that occur free in the recursive declaration. One extra word is used when GC is enabled.

Regions and their Representation

Finite ($\rho:n$)	Region whose size can be determined at compile time. During compilation, a finite region size is given as a non-negative integer. After multiplicity inference, this integer indicates the number of times a value (of the appropriate type) is written into the region. Later, after physical size inference, the integer indicates the physical region size in words. At runtime, a finite region is allocated on the runtime stack.
Infinite ($\rho:INF$)	All other regions. At runtime, an infinite region consists of a stack allocated region descriptor, which contains pointers to the beginning and the end of a linked list of fixed size region pages.

Storage Modes (only significant for infinite regions)

<code>atbot</code>	Reset region, then store value.
<code>sat</code>	Determine actual storage mode (<code>atbot/atbot</code>) at runtime.
<code>attop</code>	Store at top of region, without destroying any values already in the region.

Contents

I	Overview	17
1	Region-Based Memory Management	19
1.1	Dynamic Memory Management	19
1.2	Checked Deallocation of Memory	20
1.3	Example: the Game of Life	25
1.4	Try it!	32
2	Making Regions Concrete	33
2.1	Finite and Infinite Regions	33
2.2	Runtime Types of Regions	34
2.3	Allocation and Deallocation of Regions	34
2.4	Two Backends	35
2.5	Boxed and Unboxed Values	36
2.6	Intermediate Languages	36
2.7	The Runtime System	38
2.8	Compiling Programs with the MLKit	38
2.9	Compiling with the MLKit Compiler	38
2.10	Running Compiled Programs	39
II	The Language Constructs of SML	41
3	Records and Tuples	43
3.1	Syntax	43
3.2	Example: Basic Record Operations	44
3.3	Region-Annotated Types	44
3.4	Effects and Regions	45
3.5	Runtime Representation	47

4	Basic Values	49
4.1	Integers and Words	49
4.2	Reals	50
4.3	Characters and Strings	51
4.4	Booleans	52
5	Lists	53
5.1	Syntax	53
5.2	Physical Representation	55
5.3	Region-Annotated List Types	55
5.4	Example: Basic List Operations	56
6	First-Order Functions	59
6.1	Region-Polymorphic Functions	59
6.2	Region-Annotated Type Schemes	61
6.3	Endomorphisms and Exomorphisms	62
6.4	Polymorphic Recursion	65
7	Value Declarations	71
7.1	Syntax	71
7.2	Scope Versus Lifetime	72
7.3	Shortening Lifetime	75
8	Static Detection of Space Leaks	77
8.1	Warnings About Space Leaks	78
8.2	Fixing Space Leaks	78
9	References	81
9.1	References in Standard ML	81
9.2	Runtime Representation of References	82
9.3	Region-Annotated Reference Types	83
9.4	Local References	84
9.5	Hints on Programming with References	86
10	Recursive Data Types	87
10.1	Spreading Data Types	87
10.2	Example: Balanced Trees	89
10.3	Unboxing Schemes	91

11 Exceptions	95
11.1 Exception Names	95
11.2 Exception Values	96
11.3 Raising Exceptions	96
11.4 Handling Exceptions	97
11.5 Example: Prudent Use of Exceptions	98
12 Resetting Regions	101
12.1 Storage Modes	102
12.2 Storage Mode Analysis	104
12.3 Calls to Externally Declared Functions	108
12.4 Example: Computing the Length of Lists	109
12.5 Resetting Regions	114
12.6 Example: Improved Mergesort	115
12.7 Example: Scanning Text Files	118
13 Higher-Order Functions	125
13.1 Lambda Abstractions	125
13.2 Region-Annotated Function Types	126
13.3 Arrow Effects	128
13.4 On the Lack of Region Polymorphism	130
13.5 Examples: <code>map</code> and <code>foldl</code>	132
14 The Function Call	137
14.1 Parameter Passing	138
14.2 Tail Calls and Non-Tail Calls	138
14.3 Tail Call to Known Function (<code>jmp</code>)	139
14.4 Non-Tail Call to Known Function (<code>funccall</code>)	140
14.5 Tail Call to Unknown Function (<code>fnjmp</code>)	141
14.6 Non-Tail Call to Unknown Function (<code>fnccall</code>)	142
14.7 Example: Function Composition	142
14.8 Example: <code>foldl</code> Revisited	143
15 ML Basis Files and Modules	147
15.1 ML Basis Files	147
15.2 Structures	151
15.3 Signatures	152
15.4 Functors	153

16	Garbage Collection	157
16.1	Dangling Pointers	157
16.2	Scanning Text Files with Stream IO and Garbage Collection .	159
16.3	Instrumenting the Executable	159
III	System Reference	163
17	Region Profiling	165
17.1	Example: Scanning Text Files Again	168
17.2	Compile-Time Profiling Strategy	174
17.3	The Log File	174
17.4	Runtime Profiling Strategy	176
17.5	Regions Statistics	177
17.6	Processing the Profile Data File	179
17.7	Advanced Graphs with <code>rp2ps</code>	180
18	Controlling MLKit Compilation	185
18.1	Printing of Intermediate Forms	185
18.2	Layout of Intermediate Forms	187
19	Calling C Functions	189
19.1	Declaring Primitives and C Functions	190
19.2	Conversion Macros and Functions	192
19.2.1	Integers	193
19.2.2	Units	193
19.2.3	Reals	193
19.2.4	Booleans	194
19.2.5	Records	195
19.2.6	Strings	195
19.2.7	Lists	196
19.3	Exceptions	198
19.4	Program Points for Profiling	199
19.5	Storage Modes	200
19.6	Endomorphisms by Polymorphism	201
19.7	Compiling and Linking	202
19.8	Dynamic Linking	203
19.9	Auto Conversion	204

19.10Examples	205
20 Summary of Changes	209
20.1 Changes Since Version 4.7.16	209
20.2 Changes Since Version 4.7.2	209
20.3 Changes Since Version 4.6.1	210
20.4 Changes Since Version 4.3.0	211
20.5 Changes Since Version 4	212
20.6 Changes Since Version 3	213
20.7 Changes Since Version 2	214
A Command-Line Options	221

Preface

MLKit is a compiler infrastructure for the Standard ML programming language [MTHM97]. MLKit supports all of Standard ML, including Modules and most parts of the SML Basis Library [GR04]. MLKit features a region-based native backend that generates efficient x64 machine code. This version of the compiler is also named MLKit with Regions. MLKit also features a JavaScript backend, which generates code for execution in web browsers. MLKit with Regions, which this report is about, is intended for the development of stand-alone applications that must be reliable, fast, and space efficient.

There has always been a tension between high-level features in programming languages and the programmer's legitimate need to understand programs at the operational level. Very likely, if a resource conscious programmer is forced to make a choice between the two, he will choose the latter.

MLKit with Regions is the result of a research and development effort, which was initiated at the University of Copenhagen in 1992. The goal of the project has been to develop implementation technology that combines the advantages of using a high-level programming language, in this case Standard ML, with a model of computation that allows programmers to reason about how much space and time their programs use.

In most call-by-value languages, it is not terribly hard to give a model of time usage that is good enough for elementary reasoning.

For space, however, the situation is much less satisfactory. Part of the reason is that many programs must recycle memory while running. For all such programs, the mechanisms that reclaim memory inevitably become part of the reasoning. This is true irrespective of whether memory recycling is done by a stack mechanism or by pointer tracing garbage collection.

In the stack discipline, every point of allocation is matched by a point of deallocation and these points are obvious from the program. By contrast,

garbage collection techniques usually separate allocation, which is done by the programmer, from deallocation, which is done by a garbage collector. The advantage of using reference tracing garbage collection techniques is that they apply to a wide range of high-level concepts now found in programming languages, for example recursive data types, higher-order functions, exceptions, references, and objects. The disadvantage is that it is becoming increasingly difficult for the programmer to reason about lifetimes. Lifetimes may depend on subtle details in the compiler and in the garbage collector. Thus, it is hard to model memory in a way that is useful to programmers. Also, compilers offer little assistance for reasoning about lifetimes.

In this report, we describe how Standard ML can be equipped with a different memory management discipline, namely a *region-based* memory model. Like the stack discipline, the region discipline is, in essence, simple and platform-independent. Unlike the traditional stack discipline, however, the region discipline also applies to recursive data types, references, and higher-order functions, for which one has hitherto mostly used reference tracing garbage collection techniques.

The reader we have in mind is a person with a Computer Science background who is interested in developing reliable and efficient applications written in Standard ML. Also, the report may be of interest to researchers of programming languages, since the MLKit with Regions is a fairly bold exercise in program analysis. We should emphasise, however, that this report is very much intended as a user's guide, not a scientific publication.

This report consists of three parts:

Part I, Overview: This part gives an overview of the ideas that underlie programming with regions in the MLKit.

Part II, Understanding Regions: The second part of the report systematically presents the language constructs of the Standard ML Language, showing for each construct how it can be used when programming with regions.

Part III, System Reference: In this part, we explain how to interact with the system, how to use the region profiler and how to call C functions from the MLKit.

The present report describes the MLKit Version 4.7.22. Besides from featuring bug-fixes, this version of the MLKit extends, and differs from, the MLKit Version 4.7.2 by the following features:

1. Modular storage mode analysis. MLKit now implements a modular version of the storage mode analysis, which allows for external functions to be passed with regions that have storage mode `atbot`. The modular version of the storage mode analysis makes use of local aliasing information for inferring storage modes (Section 12.3).
2. Double-ended bit-stealing. The implementation of algebraic data types now uses an advanced unboxing scheme that allows many data types to be implemented unboxed using the, otherwise non-used, 16 most-significant bits of a pointer or the high unused bits of 8-bit word values [Els24a].
3. Deep-argument flattening. The MLKit optimiser attempts to flatten arguments to functions using a scheme that allows for deep flattening and uncurrying of arguments to functions (also across compilation unit boundaries) [Els25]. With this scheme, if a function takes a tuple of reals and a curried argument of type `real` as arguments, those arguments that are not used in their boxed form will be passed unboxed.
4. Improved abbreviated pretty printing of some intermediate representations.
5. Support for threads and parallelism through a fork-join thread interface [EH23]. The support for threads and parallelism is not documented in this report.
6. Support for explicit programming with regions and effect constraints through an extended syntax (ReML) that allows for annotating expressions and types with region and effect information [Els24b]. The support for programming with explicit regions and effects is not documented in this report.
7. Improved support for auto conversion when interacting with C functions.

Version 4.7.2 of the MLKit extends, and differs from, MLKit Version 4.3.0 by the following features:

1. Type variables are no longer associated with region variables, which simplifies region types significantly and eases the implementation.

2. The machinery for ensuring that no dangling pointers are dereferenced during a garbage collection has been fixed [Els23].
3. A move to machine code generation for the x86_64 architecture. The port also features register allocation for floating point registers, holding intra-procedural values of type `real`.
4. Additional Standard ML Basis Library features, including support for socket programming (e.g., structure `Socket`) and general Unix programming (structure `Unix`).
5. Support for generational garbage collection as a generalisation of the ordinary pointer-tracing garbage collector that can be used as an add-on to region-based memory management [EH20, EH21a].
6. A large number of bug-fixes, performance improvements, and general new features, such as colorful region profile graphs...

Further, Version 4.3.0 extends Version 4 with the following features:

1. Support for compiling ML Basis Files. ML Basis Files allows for expressing source dependencies, exactly (as a directed acyclic graph). ML Basis Files thus provides a mechanism for programming “in the very large”.
2. File-based separate compilation, based on ML Basis Files.
3. An updated Standard ML Basis Library conforming to the specification published in [GR04].
4. Untagged representation of heap-allocated pairs, triples, and Standard ML references, even when garbage collection is enabled.

MLKit Version 4 extends MLKit Version 3 with the following features:

1. Support for pointer tracing garbage collection. Pointer tracing garbage collection works well together with the region memory model. While most deallocations can be efficiently performed by region deallocation, there are some uses of memory for which life time prediction is difficult. In these cases pointer tracing garbage collection does a good job in collaboration with region memory management [Hal99, HET02].

2. An x86 native backend. The backend support has switched from HP PA-RISC to Linux on x86 architectures.
3. A bytecode backend. To improve portability of programs, MLKit Version 4 featured a bytecode backend, which generated code for execution on a stack machine with region primitives. The stack machine closely resembles the stack machine used in the OCaml and Moscow ML compilers.

The MLKit Version 3 extends the MLKit Version 2 with support for the Standard ML Modules language. The MLKit Version 2 is a further development of the MLKit Version 1, which was developed at Edinburgh University and University of Copenhagen [BRTT93]. We hope you will enjoy using the MLKit as much as we have enjoyed developing it. If your experience with the MLKit gives rise to comments and suggestions, specifically with relation to the goals and visions expressed here, please feel free to write. Further information is available at the MLKit Github web site:

<https://github.com/melsman/mlkit>

September, 2001

Mads Tofte, Lars Birkedal, Martin Elsman, Niels Hallenberg,
Tommy Højfeldt Olesen, and Peter Sestoft

Revised 2002, 2004, 2005, 2021, 2025, 2026 by Martin Elsman

Contributions

Many people have contributed to the development of the MLKit, including Peter Bertelsen, Lars Birkedal, Martin Elsmann, Niels Hallenberg, Tommy Højfeldt Olesen, Nick Rothwell, Mads Tofte, David N. Turner, Peter Sestoft, and Carsten Varming.

People who have contributed with bug reports and patches include, but are not limited to (in alphabetical order) Johnny Andersen, Troels Henriksen, Koshy A Joseph, Ken Friis Larsen, Philip Munksgaard, Henning Niss, Daniel Wang, and Stephen Weeks.

License

The MLKit compiler and tools are released under the GNU General Public License:

THIS PROGRAM IS FREE SOFTWARE; YOU CAN REDISTRIBUTE IT AND/OR MODIFY IT UNDER THE TERMS OF THE GNU GENERAL PUBLIC LICENSE AS PUBLISHED BY THE FREE SOFTWARE FOUNDATION; EITHER VERSION 2 OF THE LICENSE, OR (AT YOUR OPTION) ANY LATER VERSION.

THIS PROGRAM IS DISTRIBUTED IN THE HOPE THAT IT WILL BE USEFUL, BUT WITHOUT ANY WARRANTY; WITHOUT EVEN THE IMPLIED WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. SEE THE GNU GENERAL PUBLIC LICENSE FOR MORE DETAILS.

YOU SHOULD HAVE RECEIVED A COPY OF THE GNU GENERAL PUBLIC LICENSE ALONG WITH THIS PROGRAM; IF NOT, WRITE TO THE FREE SOFTWARE FOUNDATION, INC., 675 MASS AVE, CAMBRIDGE, MA 02139, USA.

Parts of the MLKit (the runtime system and the Basis Library) is distributed under the MIT license:

THE MIT LICENSE

PERMISSION IS HEREBY GRANTED, FREE OF CHARGE, TO ANY PERSON OBTAINING A COPY OF THIS SOFTWARE AND ASSOCIATED DOCUMENTATION FILES (THE "SOFTWARE"), TO DEAL IN

THE SOFTWARE WITHOUT RESTRICTION, INCLUDING WITHOUT LIMITATION THE RIGHTS TO USE, COPY, MODIFY, MERGE, PUBLISH, DISTRIBUTE, SUBLICENSE, AND/OR SELL COPIES OF THE SOFTWARE, AND TO PERMIT PERSONS TO WHOM THE SOFTWARE IS FURNISHED TO DO SO, SUBJECT TO THE FOLLOWING CONDITIONS:

THE ABOVE COPYRIGHT NOTICE AND THIS PERMISSION NOTICE SHALL BE INCLUDED IN ALL COPIES OR SUBSTANTIAL PORTIONS OF THE SOFTWARE.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

For details, see the file `doc/license/MLKit-LICENSE` in the source distribution.

Part I

Overview

Chapter 1

Region-Based Memory Management

Region-Based Memory Management is a technique for managing memory for programs that use dynamic data structures, such as lists, trees, pointers, and function closures.

1.1 Dynamic Memory Management

Many programming languages rely on a memory model consisting of a *stack* and a *heap*. Typically, the stack holds temporary values, activation records, arrays, and in general, values whose lifetime is closely connected to procedure activations and whose size can be determined at the latest when creation of the value begins. The heap is what holds all the other values. In particular, the heap holds values whose size can grow dynamically, such as lists and trees. The heap also holds values whose lifetime does not follow procedure activations closely (for example lists and, in functional languages, function closures and suspensions).

The beauty of the stack discipline (apart from the fact that it is often very efficient in practice) is that it couples allocation points and deallocation points in a manner that is intelligible to the programmer. C programmers appreciate that whatever memory is allocated for local variables in a procedure ceases to exist (and take up memory) when the procedure returns. C programmers also know that counting from one to some large number, N , is not best done by making N recursive C procedure calls, because that would

use stack space proportional to N .

By contrast, programmers have much less help when it comes to managing the heap. Two approaches prevail. The first approach is that the programmer manages memory herself, using explicit allocation and deallocation instructions (e.g., `malloc` and `free` in C). For non-trivial programs this can be a very significant burden, because it is, in general, very hard to make sure that none of the values that reside in the memory that one wishes to deallocate are not needed for the rest of the computation. This puts the programmer in a difficult position. If one is too eager to reclaim memory in the heap, the program might crash under some peculiar circumstances, which might be hard to find during debugging. If one is too conservative reclaiming memory, the program might leak space, that is, it might use more memory than expected, perhaps eventually, exhaust the memory of the machine.

The other prevailing approach is to use automatic garbage collection in the heap. Some implementors of some languages even dispense with the stack entirely, relying only on a heap with garbage collection. Garbage collection techniques separate allocation, which is done by the programmer, from deallocation, which is done by the garbage collector. At first, this might seem like the perfect solution: no longer does the programmer have to worry about whether memory that is being reclaimed really is dead, for the garbage collector only reclaims memory that cannot be reached by the rest of the computation. However, reality is less perfect. Garbage collectors are typically based on the idea that if data is reachable via pointers (starting from the stack and other root data) then those data must be kept. Consequently, programs have to be written with care to avoid hanging on to too many pointers. Space conscious programmers (and language implementors) can work their way around these problems, for example by assigning `nil` to pointers that are no longer used. However, such tricks often rely on assumptions about the code that cannot be checked by the compiler and that are likely to be invalidated as the program evolves.

1.2 Checked Deallocation of Memory

Regions offer an alternative to the two approaches to memory management discussed in the previous section. The runtime model is very simple, at least in principle. The store consists of a stack of *regions*, see Figure 1.1. Regions hold values, for example tuples, records, function closures, references, and




Figure 1.1: The store is a stack of regions; every region is depicted by a box in the picture.

values of recursive types (such as lists and trees). All values, except those that fit within one machine word (for example integers), are stored in regions.

The size of a region is not necessarily known when the region is allocated. Thus a region can grow gradually (and many regions can grow at the same time) so one might think of the region stack as a stack of heaps. However, the region stack really is a stack in the sense that (a) if region r_1 is allocated before region r_2 then r_2 is deallocated before r_1 and (b) when a region is deallocated, all the memory occupied by that region is reclaimed in one constant time operation.

Values that reside in one region are often, but not always, of the same type. A region can contain pointers to values that reside in the same region or in other regions. Both forward pointers (i.e., pointers from a region into a region closer to the stack top) and backwards pointers (i.e., pointers to an older region) occur.

As mentioned in the preface, the present version of the MLKit supports reference-tracing garbage collection in combination with region memory management [Hal99]. While most deallocations can be efficiently performed by region deallocation, there are some uses of memory for which it is difficult to predict when memory can be deallocated. In these cases reference-tracing garbage collection does a good job in combination with region deallocation.

In many cases however, one can do just fine without reference-tracing garbage collection. Without reference-tracing garbage collection the region

stack is the only form of memory management provided. Is the region model really general enough to fit a wide variety of computations?

First notice that the pure stack discipline (a stack, but no heap) is a special case of the region stack. Here the size of a region is known at the latest when the region is allocated. Another special case is when one has just one region in the region stack and that region grows dynamically. This case can be thought of as a heap with no garbage collection, which again would not be sufficient.

But when one has many regions, one obtains the possibility of distinguishing between values according to what region they reside in. MLKit has operations for allocating, deallocating, and extending regions. But it also has an explicit operation for resetting an existing region, that is, reclaiming all the memory occupied by the region without eliminating the region from the region stack. This primitive, simple as it is, enables one to cope with most of those situations where lifetimes simply are not nested. Figure 1.2 shows a possible progression of the region stack.

In the MLKit the vast majority of region management is done automatically by the compiler and the runtime system. Indeed, with one exception, source programs are written in Standard ML, with no added syntax or special directives. The exception has to do with resetting of regions. MLKit provides two built-in functions (`resetRegions` and `forceResetting`), which instruct the program to reset regions. Here `resetRegions` is a safe form of resetting where the compiler only inserts region resetting instructions if it can prove that they are safe; it prints thorough explanations of why it thinks resetting might be unsafe otherwise. The function `forceResetting` is for potentially unsafe resetting of regions, which is useful in cases where the programmer jolly well knows that resetting is safe even if the compiler cannot prove it. The function `forceResetting` is the only way we allow users to make decisions that can make the program crash; many programs do not need `forceResetting` and hence cannot crash (unless we have bugs in our system).

All other region directives, including directives for allocation and deallocation of regions, are inferred automatically by the compiler. This happens through a series of fairly complex program analyses and transformations (in the excess of twenty-five passes involving three typed intermediate languages). These analyses are formally defined and the central one, called *region inference*, has been proved correct for a skeletal language. Although the formal rules that govern region inference and the other program analyses

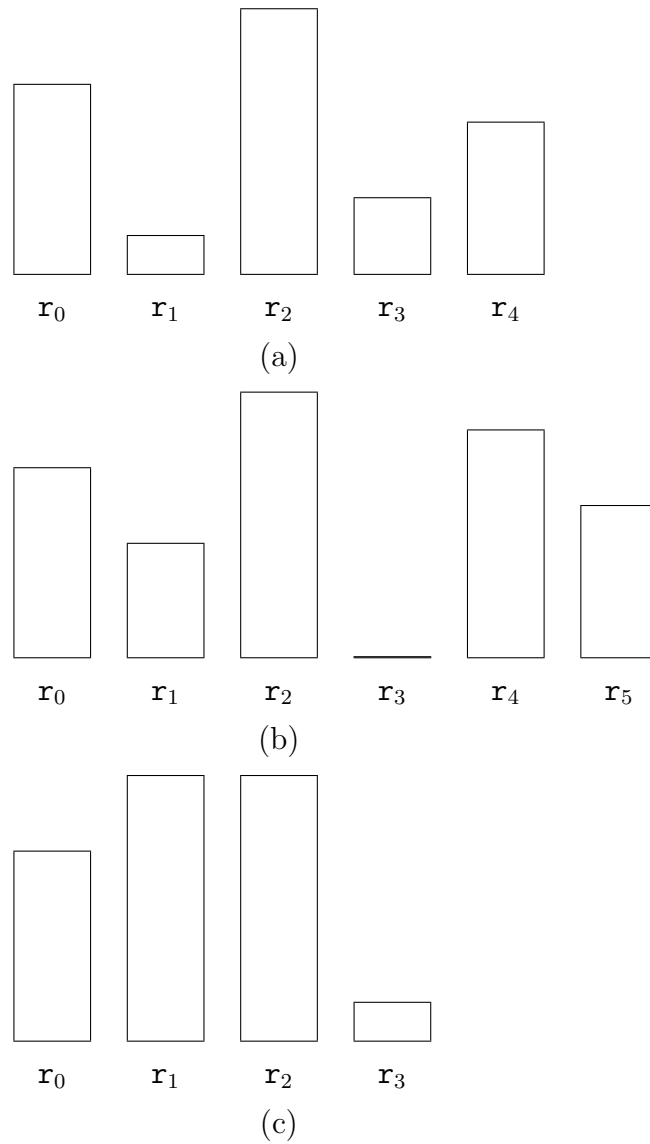


Figure 1.2: Further development of the region stack: (a) after allocation of r_4 ; (b) after growth of r_1 and r_4 , resetting of r_3 and allocation of r_5 ; (c) after popping of r_4 and r_5 but extension of r_1 and r_3 .

are complex, we have on purpose restricted attention to program analyses that we feel capture natural programming intuitions. Moreover, the MLKit implementation is such that, with one exception¹, every region directive takes constant time and constant space to execute. The fact that we avoid interrupting program execution for unbounded lengths of time gives a nice smooth experience when programs are run and should make the scheme attractive for real-time programming.

To help programmers get used to the idea of programming with regions, the MLKit can print region-annotated programs, that is, source programs it has annotated with region directives. Also, it provides a *region profiler* for examining run-time behavior. The region profiler gives a graphical representation of region sizes as a function of time. This tool makes it possible to see what regions use the most space and even to relate memory consumption back to individual allocation points in the (annotated) source program.

To sum up, the key advantages obtained by using regions compared to more traditional memory management schemes are

1. safety of deallocation is checked by the compiler
2. the compiler can in many cases spot potential space leaks
3. region management is under the control of the user, provided one understands the principles of region inference
4. each of the region operations that are inserted use constant time and constant space at runtime
5. it is possible to relate runtime space consumption to allocation points in the source program; we have found region profiling to be a powerful tool for eliminating space leaks

Regions are not a magic wand to solve all memory management problems. Rather, the region scheme encourages a particular discipline of programming. The purpose of this report is to lay out this discipline.

¹The exception has to do with exceptions. When an exception is raised, a search down the stack for a handler takes place; this search is not constant time and it involves popping of regions on the way. However, the number of region operations is bounded by the number of handlers that appear on the stack.

1.3 Example: the Game of Life

To illustrate the general flavor of region-based memory management, let us consider the problem of implementing the game of Life. The game takes place on a board that resembles a chess board, except that the size of the board can grow as the game evolves. Thus every position has eight neighboring positions (perhaps after extension of the board). At any point in time, every position is either *alive* or *dead*. A snapshot of the game consisting of the board together with an indication of which positions are alive is called a *generation*. The rules of the game specify how to progress from one generation to the next. Consider generation n from which we want to create generation $n + 1$ ($n \geq 0$). Let (i, j) be a position on the board, relative to some fixed point $(0, 0)$ in the plane. Assume (i, j) is alive in generation n . Then (i, j) stays alive in generation $n + 1$ if and only if it has two or three live neighbors in generation n . Assume (i, j) is dead at generation n . Then it is born in generation $n + 1$ if and only if it has precisely three live neighbors at generation n . We assume that only finitely many positions are alive initially. An example of two generations of Life is shown below:

```

      0
     0 0
    0 00      0
   00 0 00 0000 0
   00 0 00 0000 0
      0 0      0 0      00
      0      0000      00
          0000
      0
     0000
    00 0 0      0 0
   00 000 0 0 0 0 0
   00 00 0 0      0
      0000 0 0      00
      0      0      00
          0 0
          00

```

To represent the game board, we need a data structure that can grow

```

let val living = alive gen
    fun isalive x = member living x
    fun liveneighbours x = length (filter isalive (neighbours x))
    fun twoorthree n = n=2 orelse n=3
    val survivors = filter (twoorthree o liveneighbours) living
    val newnbrlist = collect (filter (not o isalive) o neighbours)
                          living
    val newborn = occurs3 newnbrlist
in mkgen (cp_list(survivors @ newborn))
end

```

Figure 1.3: An excerpt of (a modified version of) Chris Reade’s Game of Life program.

dynamically (so a two-dimensional array of fixed size is not sufficient). A simple solution is to represent a generation by a list of integer pairs, namely the positions that are alive. Since we want to give all pairs belonging to one generation the same lifetime (in the computer memory, that is!) it is natural to store all the integer pairs belonging to one generation in the same region. Indeed region inference forces this decision upon us, as it happens, since it requires that all elements belonging to the same list lie in the same region. (Different lists can lie in different regions, however.)

Thus, after having built the initial generation, we expect the region stack to look like this

l_n : list of integer pairs representing generation n .

r0

The computation of the next generation involves a considerable amount of list computation. Chris Reade has expressed the key part of the computation as shown in Figure 1.3. Despite the extensive use of higher-order functions

here, there is a great deal of stack structure in this computation. For example, the `survivors` list can be allocated in a local region which can be deallocated after the list has been appended (`@`) to the `newborn` list. The computation of `survivors`, in turn, involves the creation of a closure for (`twoorthree o liveneighbours`) and additional creation of closures as part of the computation of the application of `filter`. Each time `liveneighbours` is called (by `filter`) additional temporary values are created. All of this data should live shorter than `survivors` itself. The details of these lifetimes are determined automatically by the region inference algorithm, which ensures that when the above expression terminates it will simply have created a list containing the live positions of the new generation.

But now we have a design choice. Should we put the new generation in the same region as the previous region or should we arrange that it is put in a separate region? Piling all generations on top of each other in the same region would clearly be a waste of space: only the most recent generation is ever needed. Similarly, giving each generation a separate region on the region stack is no good either, because it would make the stack grow infinitely (although this could be alleviated somewhat by resetting all regions except the topmost one). The solution is simple, however: use two regions, one for the current generation and one for the new generation. When the new generation has been created, reset the region of the old region and copy the contents of the new region into the old region. This effect is achieved by organizing the main loop of the program as follows:

```

1  local
2    fun nthgen (p as(0,g)) = p
3      | nthgen (p as(i,g)) =
4        nthgen (i-1, let val g' = nextgen g
5                      in show g;
6                        resetRegions g;
7                        copy g'
8                      end)
9  in
10   fun iter n = #2(nthgen(n,gun()))
11 end

```

Here `nthgen` is the main loop of the program. It takes a pair as argument; the first component of the pair indicates the number of iterations desired, while the second, `g`, is the current generation. The use of the `as` pattern in line 2 forces the argument and the result of `nthgen` to be in the same regions.

Such a function is called a *region endomorphism*. In line 4, we compute a fresh generation, which lies in fresh regions, as it happens. Having printed the generation (line 5) we then reset the regions containing `g`. The compiler checks that this is safe. Then, in line 7 we copy `g'` and the target of this copy must be the regions of `g`, because `nthgen` is a region endomorphism (see Figure 1.4). All in all, we have achieved that at most two generations are live at the same time (a fact that can be checked by inspecting the region-annotated code, if one feels passionately about it).²

The above device, which we refer to as *double copying*, can be seen as a much expanded version of what is often called “tail recursion optimisation”. In the case of regions, not just the stack space, but also region space, is reused. Indeed, double copying is similar to invoking a copying garbage collector on specific regions that are known not to have live pointers into them. But by doing the copying ourselves, we have full control over when it happens, we know that the cost of copying will be proportional to the size of the generation under consideration and that all other memory management is done automatically by the region mechanism. Because each of the region management directives that the compiler inserts in the code are constant time and space operations, we have now avoided unpredictable interruptions due to memory management. This avoidance of unpredictable interruptions might not be terribly important for the purpose of the game of Life, but if we were writing control software for the ABS brakes of a car, having control over all costs, including memory management, would be crucial!

Region profiles for two hundred generations of `life` starting from the configuration shown earlier appear in Figures 1.5 and 1.6. The highest amount of memory used for regions during the computation is 45KiB. Figure 1.6, which has data collected from 100 snapshots of the computation, clearly shows that most of the 45KiB are reclaimed between every two generations of the game. It turns out that the game essentially stabilises with a small number of live positions on the board after roughly 150 generations. This stabilisation is clearly reflected in the region profile.

Figure 1.5 is from the same computation, but it only includes data from 50 snapshots. In the figures, we can see that the largest region is `r332640`. To find out what this region contains, however, one needs to know about the

²The source file for the life program is `kitdemo/life.sml`. Running programs is described in Section 2.8. When run with `n=10000` under Linux on an x64 box, the memory consumption (resident memory, measured using `top`) quickly reaches 500Kb and stays there for the remaining generations.

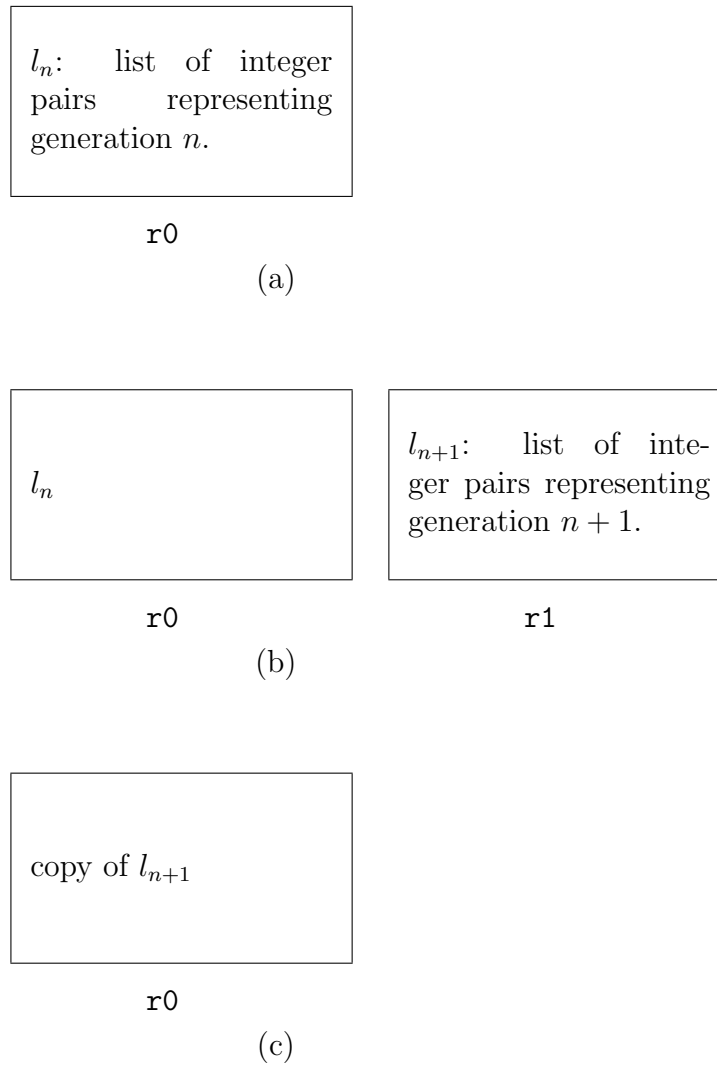


Figure 1.4: Using double-copying in the game of Life: (a) generation number n resides in region $r0$; (b) generation $(n+1)$ has been built in $r1$; (c) region $r0$ has been reset, the new generation copied into $r0$ and $r1$ has been deallocated.




Figure 1.5: A region profile of two hundred generations of the “Game of Life”, showing region sizes as a function of time (50 snapshots).




Figure 1.6: Region profile of two hundred generations of the “Game of Life”, showing region sizes as a function of time (100 snapshots).

methods described in Part II.

1.4 Try it!

This section tells you how to repeat the profiling experiment shown above.

Compile the SML program `kitdemo/life.sml` as follows. First, make a personal copy of the `kit/kitdemo` directory, place yourself in it, and execute the command:³

```
$ mlkit -no_gc -prof life.sml
```

The option `-prof` enables region profiling and the option `-no_gc` disables reference tracing garbage collection. After the MLKit has compiled the program `life.sml`, the executable life program is available as `kitdemo/run`.

Next, you may execute `run`, as follows:

```
$ ./run -microsec 500
```

This command will make a profiling snapshot every 500 microseconds (i.e., every half millisecond). If you are satisfied with less fine-grained information, choose a larger number; it will speed up execution. If you just type

```
./run
```

there will be one snapshot per second.

Finally, you create a Postscript file and convert it into a pdf-file as follows:⁴

```
$ rp2ps -region -name 'Game of life - 100 snapshots' \
        -sampleMax 100
$ ps2pdf region.ps region.pdf
```

The option `-sampleMax N` instructs `rp2ps` to show at most N snapshots (evenly distributed over the duration of the computation).

³We assume that the MLKit compiler command `mlkit` is somehow available through your `PATH` environment variable.

⁴The program `rp2ps` can be found in the `kit/bin` directory.

Chapter 2

Making Regions Concrete

In this chapter, we give a brief overview of how the abstract memory model presented in the last chapter is mapped down to conventional memory. In doing so, we shall introduce notation and concepts that will be used extensively in what follows.

2.1 Finite and Infinite Regions

Not every region has the property that its size is known at compile-time, or even when the region is first allocated at runtime. As we have seen, one typical use of a region is to hold a list, and in general there is no way of knowing how long a given list is going to be.

For efficiency reasons, however, the MLKit distinguishes between two kinds of regions: those regions whose size it can determine at compile-time and those it cannot. These regions are referred to as *finite* and *infinite* regions, respectively.¹ Finite regions are always allocated on the runtime stack. An infinite region is represented as a linked list of fixed-size pages. The runtime system maintains a free list of such pages. An infinite region is represented by a *region descriptor*, which is a record kept on the runtime stack. The region descriptor contains two pointers: one to the first and one to the last region page in the linked list that represents the region. Allocating an infinite region involves getting a page from the free list and pushing a region descriptor onto the runtime stack. Popping a region is done by appending

¹“finite” and “unbounded” would have been better terms, but it is too late to change that.

the region pages of the region and the free list (this is done in constant time) and then popping the region descriptor off the runtime stack.

At runtime, every region is represented by a 64-bit entity, called a *region name*. If the region is finite, the region name is a pointer into the stack, namely to the beginning of the region. If the region is infinite, the region name is a pointer to the region descriptor of the region.

The *multiplicity* of a region is a statically determined upper bound on the number of times a value is put into the region. MLKit operates with three multiplicities: 0, 1 and ∞ , ordered by $0 < 1 < \infty$. Multiplicities annotate binding occurrences of region variables. An expression of the form

```
let region  $\rho$  :  $m$  in  $e$  end
```

where m is a multiplicity, gives rise to an allocation of a region, which is finite if $m < \infty$, and infinite otherwise.

2.2 Runtime Types of Regions

Every region has a runtime type. The following runtime types exist: PAIR_RT, ARRAY_RT, REF_RT, TRIPLE_RT, STRING_RT, TOP_RT, and BOT_RT. Not surprisingly, regions of runtime type STRING_RT contain values of ML type **string**. Regions of runtime type PAIR_RT, ARRAY_RT, REF_RT, and TRIPLE_RT contain pairs, arrays, references, and triples, respectively. Regions with runtime type TOP_RT can contain all other forms of allocated values, that is, constructed values, such as tuples (that are not pairs or triples), records (that do not contain two or three fields), and function closures. Regions of runtime type BOT_RT are not present at runtime and are associated only with so-called *explicit region variables*, which we shall not discuss in this document.

It is often, but not always, the case that all values that reside in the same region have the same type (considered as representations of ML values).

2.3 Allocation and Deallocation of Regions

The analysis that decides when regions should be allocated and deallocated is called *region inference*. Region inference inserts several forms of memory management directives as directives into the program. The target language of region inference is called RegionExp.

In `RegionExp`, region allocation and deallocation are explicit, they are always paired, and they follow the syntactical structure of the source program. If e is an expression in `RegionExp`, then so is

```
let region  $\rho$  in  $e$  end
```

Here ρ is a *region variable*. At runtime, first a region is allocated and bound to ρ . Then e is evaluated, presumably using the region bound to ρ for storing values. Upon reaching `end`, the program pops the region.

Region inference also decides, for each value-producing expression, into which region (identified by a region variable) the value will be put.

We emphasise that region variables and `region` bindings are not present in source programs. The source language is unadulterated Standard ML, so programs that run on the MLKit should be easy to port to any other Standard ML implementation.

2.4 Two Backends

MLKit provides two different backends, one that generates native code for the x64 architecture (running Linux or macOS), the *native backend* and one that generates JavaScript [Els11].²

Whereas the native x64 backend makes use of regions as its basic memory management discipline, SMLtoJs compiles Standard ML programs to JavaScript and makes use of JavaScript's garbage collection mechanisms for allocating and deallocating memory. In this report, we shall only be concerned about the native backend, for which the linear address space is partitioned into a stack and a heap, which holds region pages, all of the same size.³

For the x64 native backend, programs compile into a sequence of instructions, for example for moving word-size data between two registers or between a register and a memory location. More complex operations, such as function application, are expressed by sequences of more detailed instructions. The

²Previous versions of the MLKit also supported a bytecode backend that allowed for generated bytecode to be executed by a region based abstract machine [EH02] and which was used in the context of a web-server plugin for Standard ML code [EH03, EHV07]; support for this backend has terminated, however.

³With the exception that objects that do not fit into a page are allocated using `malloc` and freed with `free`.

native backend implements Iterated Register Allocation [GA96] for assigning machine registers to temporary variables, using the runtime stack for spilling. Although register allocation as well as other issues, such as the interaction between hardware cache strategies and code selection, are important for generating efficient code on modern architectures, we do not want to go to that level of detail here. Our primary concern is with establishing a model that the user can safely use as a worst-case model of what happens at runtime.

2.5 Boxed and Unboxed Values

As is common with implementations of programming languages, we distinguish between *boxed* and *unboxed* representations of values. An *unboxed* value is one that fits in one machine word (i.e., 64 bits) without being represented as a pointer to allocated data for the value. A *boxed value*, on the other hand, is one that is represented by a word-sized pointer to the value itself, which is stored in one or more regions.

MLKit uses unboxed representation for integers, booleans, words, the unit value, and characters. MLKit uses boxed representation for pairs, records (with at least one element), reals, exception values, function closures, and many constructed values (i.e., data types). Some data types, such as lists and booleans, are represented unboxed, however, which is possible by using the lower and higher bits in pointers to discriminate between the constructed values.

A boxed value may reside in a finite or an infinite region. Unboxed values are not stored in regions, except when they are part of a boxed value. For example, the integer 3 by itself is stored as the (binary representation) of the value 3 in a register or in a machine word. However, the pair (3,4) is represented as a pointer to two consecutive words in a region, the first of which contains the binary representation of 3 and the second of which contains the binary representation of 4.

2.6 Intermediate Languages

The MLKit native compiler compiles Standard ML programs via a sequence of typed intermediate languages into x64 machine code.

The intermediate languages that we shall refer to in the following are (in the order in which they are used in the compilation process):

Lambda: A lambda-calculus like intermediate language. The main difference between the Standard ML Core Language and Lambda is that Lambda only has trivial patterns and allows functions to take multiple arguments.

RegionExp: Same as Lambda, but with explicit region annotations (such as the `region` bindings mentioned in Section 2.3). Region variables have their runtime type (Section 2.2) as an attribute, although, for brevity, the pretty printer omits runtime types when printing expressions, unless instructed otherwise.

MulExp: Same as RegionExp, but now every binding region variable occurrence is also annotated with a multiplicity (Section 2.1) in addition to a runtime type. Again, the default is that the runtime type is not printed. The terms of MulExp are polymorphic in the information that annotate the nodes of the terms. That way, MulExp can be used as a common intermediate language for a number of the internal analyses of the compiler, which add more and more information on the syntax tree. The analysis that computes multiplicities is called the *multiplicity analysis*.

MLKit compiles Standard ML records into Lambda-tuples and it compiles Standard ML match expressions and other constructs containing patterns into simpler Lambda-constructs.

MLKit contains a Lambda optimiser, which will happily rewrite Lambda terms when it is clear that this rewrite results in faster programs (as long as the rewrite cannot lead to increased space usage).

Region inference takes Lambda to be the source language. Region inference happens after the Lambda optimiser has had a go at the Lambda term. Therefore, it was not really true when we said that region inference simply annotates source programs; we ignored the translation from SML to Lambda and the Lambda optimiser. Thus, one has to get used to (mostly minor) differences between the source language and the intermediate languages of the compiler if one wants to read programs in their intermediate forms. Moreover, Modules Language constructs are eliminated during compilation from the intermediate languages (see Chapter 15 for details of compiling with Modules in the MLKit).

When we want to show the result of the analyses, we usually show a `MulExp` expression.

2.7 The Runtime System

The runtime system is written in C. It is small (less than 30Kb of code when compiled). It contains operations for allocating and deallocating regions, extending regions, obtaining more space from the operating system, recording region profiling information, and performing low-level operations for use by the Standard ML Basis Library.

It is possible to call C functions from MLKit code if you use the native backend. MLKit takes care of the memory allocation, by allocating regions for the result of the call before the call and deallocating the regions at some point after the call. The C functions can build ML data structures such as lists through abstract operations provided by the MLKit runtime system. See Chapter 19 for further details.

2.8 Compiling Programs with the MLKit

MLKit is a batch compiler. Thus, executing a program consists of first compiling the program and then running the generated target program. Because the MLKit stores files in the directories where your source files are located, you should make a personal copy of these directories. Before you try any of the examples below, make a personal copy of the `kitdemo` directory, which is part of the distribution, and run the MLKit on your own copy.

2.9 Compiling with the MLKit Compiler

The mechanism the MLKit provides for compiling programs is to give the program source(s) as argument to the MLKit command `mlkit`. Together with the sources, a series of options may be passed to the `mlkit` command. Let us assume that the UNIX command `mlkit` is available on your system.⁴

Compiling an MLB-file (which may list several SML source files) is similar to compiling a single SML source file. However, we shall postpone the in-depth discussion of how to compile MLB-files to Chapter 15.

⁴The `README` file in the distribution tells you how to install the MLKit.

As an example, to compile the file `projection.sml` located in the `kitdemo` directory, first go to this directory and execute the following command:

```
$ mlkit -no_gc projection.sml
```

Execution of this command will result in an executable file `run`, placed in the `kitdemo` directory.

To see some internal representations of the `projection.sml` program, as produced during compilation, try pass the command-line options `-print_types` and `-print_drop_regions_expression` to the `mlkit` command, as follows:

```
$ rm -rf MLB
$ mlkit -no_gc \
    -print_types \
    -print_drop_regions_expression \
    projection.sml
```

Removing the `MLB` directory is necessary to avoid the MLKit to recognise that it can reuse the previous result of compiling the `projection.sml` program. A shorter version of the compilation command is

```
$ mlkit -no_gc -Ptypes -Pdre projection.sml
```

To get more information about which options you can pass to the MLKit at the command-line, try executing `mlkit -help`. The output of executing this command is shown in Appendix A. For instance, you may add the flag `-w 40` to specify that MLKit should pretty-print output using a width of 40 character columns.

2.10 Running Compiled Programs

If no errors were found during compilation, the MLKit produces a *target program* in the form of an executable file, called `run`. MLKit places `run` in the working directory.

Running the target program is done from the UNIX shell by typing

```
$ ./run
```

For small programs, the file will probably be around 250Kb large, even for the trivial examples considered in this chapter. This is because it contains the MLKit runtime system and compiled code for the parts of the SML Basis Library that are needed for linking.

Running the programs presented in this chapter is not particularly exciting, because none of them produce output! However, as an exercise, try compile and execute the `helloworld.sml` program, which, like all other example files in this document, is located in the `kitdemo` directory.

Part II

**The Language Constructs of
SML**

Chapter 3

Records and Tuples

In this chapter we describe construction of records and selection of record components. We also use records to introduce *region-annotated types* and *effects*, which are crucial for understanding when regions are allocated and deallocated.

3.1 Syntax

As part of the SML to Lambda translation, all SML records and SML tuples are compiled into Lambda tuples. The components of Lambda tuples are numbered from left to right, starting from 0. Selection is a primitive operation, both in Lambda and in the other intermediate languages. This primitive is printed using SML notation $\#i$. Components are numbered from 0: the i th components of a tuple of type $\tau_1 * \dots * \tau_n$ is accessed by $\#i$, for $0 \leq i \leq n - 1$.

The tuple constructor in Lambda is written as in SML:

$$(e_1, \dots, e_n)$$

However, the corresponding expression in RegionExp and MulExp takes the form

$$(e_1, \dots, e_n) \text{ at } \rho$$

where ρ is a region variable indicating where the tuple should be put. In the case $n = 0$, the **at** ρ is not printed, because the empty tuple is not allocated; it is just a constant that fits in a register at runtime.

Records are evaluated left to right.

3.2 Example: Basic Record Operations

Consider the source program

```
val xy = ((), ())
val x = #1 xy;
```

Here is the resulting MulExp program:¹

```
let val xy = ((), ()) at r4;
    val x = #0 xy
in {|xy: (_, r4), x: _|}
end
```

There are several things to notice from this example.

1. The MulExp program contains a free region variable, `r4`. Notice that the construction of the pair `xy` has been annotated by “`at r4`”, indicating where the pair should be put;
2. The expression `{|xy: (_, r4), x: _|}` is an example of a *frame expression*. A frame enumerates the components that are exported from a compilation unit. A frame is similar to a record, except that its components are variables, each annotated with a type scheme and a region variable, if the value is boxed. (In records, the components can only have types, not general type schemes.) In the example, the type of the frame is `{|xy: (unit*unit, r4), x: unit|}`. The type shows that, after the program unit has been evaluated, `xy` will reside in `r4`. In the the above example, printing of types was suppressed. Thus types were abbreviated to `_`.

3.3 Region-Annotated Types

ML type inference infers a type for every expression in the program. Region inference extends this idea by inferring for each expression a (*region-annotated*) *type with place*. We use μ to range over types with places

$$\mu ::= (\tau, \rho) \mid \tau$$

¹Program `kitdemo/projection.sm1`. Running programs is described in Section 2.8, including how to print intermediate language representations.

where τ is a *region-annotated type*, which again can contain other region-annotated types with places. The region-annotated type with place of an expression is the ML type of the expression decorated with extra region information; every type constructor that represents boxed values (e.g., pairs and strings) is paired with a region variable, indicating where the value is to be put at runtime. Type constructors that represents unboxed values (e.g., integers and booleans) are not paired with a region.

Here are some examples of region-annotated types with places:

`unit` The type of 0-tuples. Integers, booleans, and 0-tuples are represented unboxed at runtime (rather than being stored in regions), see Section 2.5.

`(string, ρ)` The type of strings in region ρ .

`(int * (string, ρ_1), ρ_2)` The type of pairs in ρ_2 whose first component is an integer and whose second component is a string in region ρ_1 .

One can get the MLKit to print the region-annotated types with places that it infers for binding occurrences of variables. The above example then becomes

```
let val xy:(unit*unit,r4) = (((), ()))at r4;
    val x:unit = #0 xy
in {|xy: (unit*unit,r4), x: unit|}
end
```

3.4 Effects and Regions

We now describe the general principle that the MLKit uses to decide when it is safe to put `region` binding around an expression.

Here is an example of an SML program that first creates a pair and then selects a component of the pair, after which the pair is garbage:²

```
val n = let val pair = if true then (3+4, 4+5)
                    else (4, 5)
in #1 pair
end
```

²Program `kitdemo/elimpair.sml`. To avoid constant case-folding, one need to pass the option `-no_cfold` to the MLKit: `mlkit -no_gc -no_cfold -Pdre elimpair.sml`.

```

let val n =
  let region r9:1;
    val pair = case true of
      true => (3 + 4, 4 + 5)at r9
    | _ => (4, 5)at r9
  in #0 pair
  end
in {n: _}
end

```

Figure 3.1: Region inference decides that the pair is to be allocated in a local, finite region; the region will be deallocated as soon as the pair becomes garbage.

MLKit compiles the declaration into the MulExp program shown in Figure 3.1. The compiler compiles the program as it is, without reducing the conditional to its `then` branch. During evaluation, a region (denoted by `r9`) is introduced before the pair is allocated; it remains on the region stack until the projection of the pair is evaluated, after which the region is deallocated.

The “:1” on the binding occurrences of `r9` is a multiplicity indicating that there is only one store operation into the region. (The multiplicity analysis has discovered that there is at most one store from the `then` branch and at most one store from the `else` branch and that at most one of the branches will be chosen.) Thus, the pair will be allocated in a little region on the runtime stack.

But how does the MLKit know that it is safe to deallocate `r9` when the region goes out of scope?

The answer lies in the fact that the MLKit infers for every expression not just a region-annotated type with place, but also a so-called *effect*. An effect is a finite set of atomic effects. Two forms of atomic effect are `put(ρ)` and `get(ρ)`, where ρ as usual ranges over region variables. The atomic effect `put(ρ)` indicates that a value is being stored in region ρ and `get(ρ)` indicates that a value is being read from region ρ . In our example, the region inference algorithm considers the sub-expression ($e_0 =$)

```

let val pair = case true of
  true => (3 + 4, 4 + 5)at r9

```

```

| _ ⇒ (4, 5)at r9
in #0 pair
end

```

and finds that it has region-annotated type `int` and effect `{put(r9), get(r9)}`.

Whenever a region variable occurs free in the effect of an expression but occurs free neither in the region-annotated type with place of the expression nor in the type of any program variable that occurs free in the expression then that region variable denotes a region that is used only locally within the expression. That this is true is of course far from trivial, but it has been proved for a skeletal version of `RegionExp`. Consequently, when this condition is met, the region inference algorithm wraps a `region` binding of the region variable around that expression.

In our example, there are no free variables in e_0 ; moreover, `r9` occurs in the effect of e_0 but not in the region-annotated type with place of e_0 . Thus, the region inference algorithm inserts a `region` binding of `r9` around e_0 .

3.5 Runtime Representation

A record with 0 components (the value of type `unit`) is represented unboxed. A record with n components ($n \geq 1$) is represented boxed, as a pointer to precisely n words in a region.³ Notice that records are not tagged. Avoiding tags is possible when the reference tracing garbage collector is disabled, because polymorphic equality is compiled into monomorphic equality functions that do not have to examine the type of objects at runtime [Els98].

`Lambda`, `RegionExp`, and `MulExp` allow one to express unboxed tuples, also in the case of function calls and returns. For functions that take a tuple as parameter, the `MLKit` passes the argument tuple unboxed if it can see that the boxed representation of the tuple is not needed by the function. `MLKit` does not at present unbox records returned from functions. See Section 6.1 on page 59 for details about unboxed function arguments.

A tuple is not allocated until its components have been evaluated.

When reference-tracing garbage collection is enabled, records and tuples are tagged (with an extra tag-word) unless they are of size two or three (pairs or triples), in which case they will reside in regions of runtime type

³When garbage collection (GC) is enabled, $n + 1$ words are used to hold a record with n components.

PAIR_RT or TRIPLE_RT, respectively. Such regions are treated specially by the reference-tracing garbage collector.

Chapter 4

Basic Values

In this chapter we describe how basic values such as integers, reals, strings, and booleans are represented in the MLKit. MLKit complies to the Definition of Standard ML (Revised) and to large parts of the Standard ML Basis Library;¹ that is, as a programmer, you can refer to components of the Standard ML Basis Library through the *initial basis*, in which all programs are compiled. Throughout this chapter, we introduce some of the top-level bindings that are provided by the initial basis.

4.1 Integers and Words

Values of type `int` are represented as unboxed 64-bit signed integers. When reference tracing garbage collection is enabled in the MLKit, one bit is used for tagging, thus in this case values of type `int` are really 63-bit signed integers; Chapter 16 describes how to compile programs with garbage collection enabled. The structure `Int` provides many useful operations on integers of type `int`.² MLKit also defines the structures `Int31`, `Int32`, `Int63`, and `Int64` for operations on 31-bit, 32-bit integers, 63-bit, and 64-bit integers. When garbage collection is enabled, values of type `Int64.int` and `Int32.int` (for historical reasons) are represented boxed, whereas values of type `Int31.int` and `Int63.int` are represented unboxed. When garbage collection is enabled, the structure `Int` is identical to the structure `Int63`. When garbage collection

¹See the MLKit web site for a link to the Standard ML Basis Library.

²To see what operations are available in the `Int` structure, consult the file `basis/INTEGER.sml`.

is disabled, the structure `Int` is identical to the structure `Int64`.

The following operations on integers are predefined at top level:

```
infix 4 = <> < > <= >=
infix 6 + -
infix 7 div mod *
val ~ : int → int
val abs : int → int
```

In fact, these operations are overloaded and will work on values of other integer types as well, even on values of type `IntInf.int`, which denotes arbitrarily-sized integers.

Operations on 8-bit, 31-bit, 32-bit, 63-bit, and 64-bit unsigned words are available in the structures `Word8`, `Word31`, `Word32`, `Word63`, and `Word64`.

Similarly as for integers, when garbage collection is enabled, the structure `Word` is identical to the structure `Word63` and values of type `Word64.word` are represented boxed. Contrary, when garbage collection is disabled, the structure `Word` is identical to the structure `Word64` and values of type `Word64.word` are represented unboxed.

4.2 Reals

The initial basis provides the following top-level operations on reals:

```
infix 4 < > <= >=
infix 6 + -
infix 7 * /
val ~ : real → real
val abs : real → real
val real : int → real
val trunc : real → int
val floor : real → int
val ceil : real → int
val round : real → int
```

Values of type `real` are implemented as 64-bit floating point numbers. They are always boxed when they appear in data structures, that is, they are each represented as a pointer to a 64-bit word, which reside in a region of runtime type `TOP_RT`.

A real constant c in the source program is translated into an expression of the form c **at** ρ , where ρ is a region variable, indicating the region into which the real will be stored.

MLKit will do its best to eliminate regions holding reals by unboxing reals and representing them in floating point registers or on the runtime stack. Thus, in internal representations, MLKit distinguishes between (boxed) reals and unboxed 64-bit floating point values (of type **f64**). MLKit supports that values of type **f64** are passed to and returned from functions in floating point registers. Values of type **f64** may also be stored in closures.

The structures **Real**, **IEEEReal**, and **Math** provide other useful operations on reals.³

4.3 Characters and Strings

The initial basis provides the following top-level operations on characters and strings:

```

infix      4 =
infix      6 ^
val ord    : char → int
val chr    : int → char
val str    : char → string
val size   : string → int
val explode : string → char list
val implode : char list → string
val ^      : string * string → string
val concat : string list → string
val substring : string * int * int → string

```

Characters are represented as 64-bit words, although only 8 bits are used to store the character. Characters are always unboxed, also when garbage collection is enabled.

A string is represented by a 64-bit pointer into an infinite region. A string is stored in a region page if it fits in the page; otherwise, it is allocated using **malloc** and linked with the region so that it can be **freed** when the region is deallocated. The internal string representation is completely transparent to the programmer, who does not have to worry about the actual size of

³Consult the files **basis/REAL.sig**, **basis/IEEE_REAL.sig**, and **basis/MATH.sig**.

region pages. Characters of a string take up only 8 bits of memory each. The size of a string is kept along with the characters that make up the string. For compatibility with many C routines, the character sequence is null-terminated.

Calls to `ord`, `chr`, `str`, and `size` take constant time and space. Calls of `explode`, `implode`, `concat`, `substring`, and `^` take time and space proportional to the sum of the size of their input and their output.

The string and character operations can raise exceptions, as detailed in the Standard ML Basis Library documentation.

The structures `Char`, `String`, `CharVector`, `Byte`, and `StringCvt` provide other useful operations on characters and strings.⁴

4.4 Booleans

The boolean values `true` and `false` are represented as 64-bit words, although only one bit is used to denote the value. Booleans are unboxed. The initial basis provides the following top-level operations on booleans:

```
infix 4 =
val not : bool → bool
```

The structure `Bool` provides other useful operations on booleans.⁵

⁴Consult the files `basis/CHAR.sig`, `basis/STRING.sig`, `basis/MONO_VECTOR.sml`, `basis/BYTE.sig`, and `basis/STRING_CVT.sml`.

⁵Consult the file `basis/BOOL.sig`.

Chapter 5

Lists

Section 5.1 gives a summary of the list concept in Standard ML, introduces the notion of the *auxiliary pairs* of a list and presents the syntax of constructors and de-constructors in the intermediate languages. Section 5.3 introduces region-annotated list types and show how they correspond to the layout of lists in memory. Section 5.4 gives a small example.

5.1 Syntax

In Standard ML, all lists are constructed from the two constructors `::` (read: cons) and `nil`. As a shorthand, one can write

$$[exp_1, \dots, exp_n]$$

for

$$exp_1 :: \dots :: exp_n :: \mathbf{nil}$$

which in turn is short for

$$\mathbf{op} :: (exp_1, \dots, \mathbf{op} :: (exp_n, \mathbf{nil}) \dots)$$

where exp ranges over expressions. The type schemes of `nil` and `::` (cons) are

$$\mathbf{nil} \mapsto \forall \alpha. \alpha \mathbf{list} \quad \mathbf{::} \mapsto \forall \alpha. \alpha * \alpha \mathbf{list} \rightarrow \alpha \mathbf{list}$$

Notice that `::` is always applied to a pair. The construction of the pair and the application of `::` should, in principle, not be confused: the pair and the constructed value are in principle separate values inasmuch as they have different type. For example, the declaration

```

val p = (2, nil)
val mylist = (op ::) p
val n = #1 p

```

is legal in Standard ML. We refer to the pairs to which `::` is applied as *auxiliary pairs* (of the list data type).

Decomposition of list values in Standard ML is done by pattern matching. A pattern can extract the pair to which `::` is applied. Pattern matching on pairs can then give access to the components of the pair.

```

val abc = ["a", "b", "c"]
val op :: p = abc (* binds p to the pair ("a", ["b","c"]) *)
val (x::y::_) = abc (* binds x to "a" and y to "b" *)

```

In the last declaration, the pattern `(x::y::_)` is short for the pattern

$$(\text{op} :: (\text{x}, \text{op} :: (\text{y}, _)))$$

which combines decomposition of constructed values with decomposition of pairs.

The intermediate languages Lambda, RegionExp, and MulExp have SML-like constructs for applying constructors, but they decompose constructed values by applying a de-constructor primitive, not by pattern matching.

Lambda, RegionExp, or MulExp	
<code>nil</code>	create <code>nil</code> value
<code>::(e)</code>	create <code>::</code> (cons) value
<code>decon_::(e)</code>	cons decomposition

In Lambda, which has essentially the same type system as SML, `decon_::`, the decomposition function for `::`, has type $\forall\alpha.\alpha \text{ list} \rightarrow \alpha * \alpha \text{ list}$. In addition, Lambda, RegionExp, and MulExp have a simple case construct:

$$\text{case } e \text{ of } :: \Rightarrow e_1 \mid _ \Rightarrow e_2$$

where e must have list type. For convenience, MLKit will layout MulExp `case` constructs of the form

$$\text{case } x \text{ of } :: \Rightarrow \text{let } y = \text{decon_}:: x \text{ in } e_1 \text{ end} \mid _ \Rightarrow e_2$$

as

$$\text{case } x \text{ of } :: y \Rightarrow e_1 \mid _ \Rightarrow e_2$$




Figure 5.1: Layout of the list `["a", "b", "c"] :: ((string, ρ_1), [ρ_2])list` in memory. The auxiliary pairs of the list reside in ρ_2 . Each auxiliary pair takes up two words; the constructors `::` (cons) and `nil` are represented unboxed.

5.2 Physical Representation

The empty list is represented by an odd, unboxed integer. A non-empty list is represented as a pointer to a pair of two words in a region, the first of which contains the head of the list and the second of which contains the representation of the tail of the list. In other words, the physical representation does not distinguish a `::` cell from the auxiliary pair to which `::` is applied. Since `nil` is represented by an odd number and since word addresses are always even, `nil` can be distinguished from the representation of a non-empty list.

As a consequence, there is no cost involved in applying `::` to an auxiliary pair or in applying the decomposition operator `decon_::` to a non-empty list.

5.3 Region-Annotated List Types

In Standard ML, all elements of a given list must have the same type. We extend this constraint to region inference by saying that all element values in the same list must reside in the same region(s) and that all auxiliary pairs of the same list must reside in the same region.

Thus, region inference does not distinguish between a list and its tail. Indeed, a typical use of an infinite region is to hold all the auxiliary pairs of a list. For an example, Figure 5.1 shows how the list `["a", "b", "c"]` is laid out in memory.

In general, the region-annotated type of a list takes the form

$$(\mu, [\rho])\mathbf{list}$$

where μ is the region-annotated type with place of the members of the list and where ρ is the region where the auxiliary pairs of the list are stored. For example, the region-annotated type

$$((\mathbf{string}, \rho_1), [\rho_2])\mathbf{list}$$

classifies lists that have their auxiliary pairs in a region ρ_2 and strings in a region ρ_1 .

Notice that the `list` type constructor is not paired with a region variable. The reason is that the physical representation of lists treats the constructors as unboxed in the sense described in Section 5.2.

Very importantly, not all lists need to live in the same regions. Formally, `nil` and `::` have the following region-annotated type schemes:

$$\begin{aligned} \mathbf{nil} &\mapsto \forall \alpha \rho. (\alpha, [\rho])\mathbf{list} \\ \mathbf{::} &\mapsto \forall \alpha \rho \epsilon. (\alpha * (\alpha, [\rho])\mathbf{list}, \rho) \xrightarrow{\epsilon \emptyset} (\alpha, [\rho])\mathbf{list} \end{aligned}$$

Despite its verbosity, the type scheme for `::` deserves careful study. It is polymorphic not just in types (signified by the bound type variable α) but also in the region signified by the bound region variables ρ , and, in the case of `::`, in the effect signified by the *effect variable* ϵ . The $\epsilon \emptyset$ appearing on the function arrow is called an *arrow effect*. Occurring in a function type, an arrow effect describes the effect of applying the function. In this case, the effect is empty, as only unboxed values are manipulated by `::`. The effect variable ϵ is used for expressing dependencies between effects (examples follow in Chapter 13). Due to the fact that the variables are universally quantified, every occurrence of a list can, potentially, be in its own regions. But notice that the type of `::` forces the element, which is consed onto the list, to be in the same regions as the already existing elements of the list. Similarly, the type forces the auxiliary pairs to be in one region (ρ).

5.4 Example: Basic List Operations

MLKit compiles the program¹

¹Program `kitdemo/onetwothree.sml`.

```

let val it =
  let region r9:INF;
    val l = [1,2,3] at r9
  in case l of
    :: v91 => #0 v91 | _ => raise Bind
  end
in {lit: _}
end

```

Figure 5.2: Example showing construction and de-construction of a small list. Layout of the list `l` is analogous to Figure 5.1. The infinite region `r9` holds the auxiliary pairs of the list. Notice also that `v91` is not bound to an explicit expression `decon_:: l`; instead, the pretty-printing abbreviates such immediate bindings by hoisting the bound variable (here `v91`) into a pattern in the case construct.

```

let val l = [1, 2, 3];
  val (x::_) = l
in x
end

```

into the RegionExp program shown in Figure 5.2.

Chapter 6

First-Order Functions

In this chapter, we shall treat functions that are declared with `fun` and that are first-order (i.e., that neither take functions as arguments nor produce functions as results). Higher-order functions are treated in Chapter 13. Region polymorphism works uniformly over all types; we use lists as an example of the general scheme.

6.1 Region-Polymorphic Functions

It would be a serious limitation if all lists produced by a series of calls to a function were stored in the same region, for then all those lists would have to be kept alive until the last time one of them were used. The solution that the MLKit offers to this problem is *region-polymorphic functions*, that is, functions that are passed regions at runtime.

When one declares a function that, when called, produces a fresh list, then the region inference algorithm will automatically insert extra formal region parameters in the function declaration. At every place one refers to the function, for example because one calls the function, the region inference algorithm inserts actual region parameters that tell the function where to put its result. This is all done automatically; the user does not have to introduce region parameters or pass them as arguments. Even so, it is useful to understand the general principle, so that one can make good use of region polymorphism.

The syntax of a (single) function declaration in MulExp is:

$$\text{fun } f \text{ at } \rho_0 \text{ } [\rho_1, \dots, \rho_k] \text{ } (x_1, \dots, x_n) = e$$

Here ρ_0 denotes the region in which the closure for f is stored, ρ_1, \dots, ρ_k are the *formal region parameters*, x_1, \dots, x_n are value parameters, and e is the body of the function. In general, calls to a function f takes the form

$$f [\rho'_1, \dots, \rho'_k] (e'_1, \dots, e'_n)$$

where $[\rho'_1, \dots, \rho'_k]$ are *actual region parameters* and e'_1, \dots, e'_n are expressions denoting the arguments to the call. Notice that region parameters are enclosed in brackets (`[]`), which should not cause confusion with ML lists. In the special case $k = 0$, no region parameters are passed to the function, and we shall often omit the brackets in this case.

Unlike for Standard ML, functions are allowed to be passed multiple value arguments, without having them passed in a boxed tuple. In the case $n = 1$, we often omit the surrounding brackets (`(...)`).

Different calls of f can use different actual regions; this feature is essential for obtaining good separation of lifetimes. For an example, consider the following program:

```
fun fromto (a, b) = if a>b then []
                  else a :: fromto(a+1, b)
val l = #1(fromto(1,10), fromto(100,110));
```

The corresponding MulExp program is shown in Figure 6.1.

```
let fun fromto at r1 [r13:INF] (v126, v127) =
      case v126 > v127 of
        true => nil
      | _ => :: (v126, fromto[r13] (v126 + 1, v127))at r13;
  val l =
    let region r23:INF, r27:1
        in #0 (fromto[r4] (1, 10), fromto[r23] (100, 110))at r27
        end
  in {|fromto: (_,r1), l: _|}
  end
```

Figure 6.1: The region-annotated version of `fromto` shows that `fromto` is region-polymorphic. (Program: `kitdemo/fromto.sml`, printed by passing the option `-print_drop_regions_expression` (or `-Pdre`) to the MLKit compiler.)

There are several things to notice about the region annotated program. First, notice that the function `fromto` represents its argument `(a, b)` unboxed; the MLKit figures out that the function does not use the boxed representation of the argument and transforms all calls to the function to pass the argument unboxed (on the runtime stack and in registers if possible). This kind of *argument flattening* is supported also for nested tuple arguments and is the mechanism also behind uncurrying and passing floating point arguments in registers [Els25].

Second, notice that `r13` is a formal region parameter of `fromto` and that `r13` is passed along in the recursive call `fromto[r13] (v126 + 1, v127)`. Here the notation `(v126 + 1, v127)` (without an `at` annotation) denotes the passing of the unboxed record to the function `fromto`.

Finally, notice that the regions that hold the two lists generated by this program are distinct. The list that escapes to top level is stored in the global region `r4`, whereas the list that does not escape is stored in the local region `r23`.

6.2 Region-Annotated Type Schemes

A (*region-annotated*) *type scheme* takes the form

$$\sigma ::= \forall \alpha_1 \cdots \alpha_n \rho_1 \cdots \rho_k \epsilon_1 \cdots \epsilon_m . \tau$$

where $\alpha_1, \dots, \alpha_n$ are type variables, ρ_1, \dots, ρ_k are region variables, $\epsilon_1, \dots, \epsilon_m$ are effect variables, and τ is a region-annotated type. The types of `nil` and `::` in Section 5.3 are examples of region-annotated type schemes.

There is a close connection between, on the one hand, the formal and actual region parameters found in `RegionExp` (and `MulExp`) programs, and, on the other hand, the region-annotated type schemes that the region inference algorithm assigns to recursively declared functions. The formal region parameters of a function stem from the bound region variables of the region-annotated type scheme of that function. The actual region parameters which annotate a call of the function are the region variables to which the bound region variables are instantiated at that particular application.

For example, the region-annotated type scheme of `fromto` from Figure 6.1 is

$$\forall \rho_{13} \epsilon . [\text{int}, \text{int}] \xrightarrow{\epsilon . \{\text{put}(\rho_{13})\}} (\text{int}, [\rho_{13}]) \text{list}$$

where we use the syntax $[\tau_1, \dots, \tau_n], n \geq 1$ to denote an unboxed tuple of types τ_1, \dots, τ_n . This syntax is not to be confused with the auxiliary region variables of type constructors (e.g., the list $[\rho_{13}]$ in the region-annotated type scheme of `fromto`.)

At the last call of `fromto` in Figure 6.1, the type scheme is instantiated to the region-annotated type

$$[\text{int}, \text{int}] \xrightarrow{\epsilon'. \{\text{put}(\rho_{23})\}} (\text{int}, [\rho_{23}]) \text{list}$$

The instantiation of bound variables of the type scheme that yields this region-annotated type is

$$\{\rho_{13} \mapsto \rho_{23}, \epsilon \mapsto \epsilon'\}$$

In general, the actual region parameters that annotate a call of a region-polymorphic function are obtained from the range of the substitution by which the type scheme of the function is instantiated at that application.

Region-polymorphic functions also have to be allocated somewhere. Therefore, the region information associated with a region-polymorphic function is a (*region-annotated*) *type scheme with place*, that is, a pair (σ, ρ) . Indeed, every binding of a variable to a boxed value (whether the binding is done by `fun`, `let`, or `fn`) associates a region-annotated type scheme with place to the binding occurrence. (In the case of `let`, the type scheme will have no quantified region and effect variables, however, and in the case of `fn`, the type scheme will have no quantified variables at all.) In the following, when we refer to “the region-annotated type (scheme) with place” of some variable, we mean the region-annotated type (scheme) with place that is associated with the binding occurrence of the variable. The region type scheme should be clearly distinguished from instances of the type scheme, which decorate non-binding occurrences of the variable.

The region-annotated type scheme with place of a variable bound to an unboxed value is always on the form σ (no specified place), where σ is the region-annotated type scheme associated with the variable (see Section 3.3).

6.3 Endomorphisms and Exomorphisms

The `fromto` function from Section 6.2 has the property that it can put its result in regions that are separate from the regions where its argument lies.

This is not surprising, if one looks at the declaration of the function; it creates a brand new list that does not share with the argument (\mathbf{a}, \mathbf{b}) , except for the integers \mathbf{a} and \mathbf{b} , which may end up in the list. The freshness of the generated list is evident from the region type scheme of the function; the region variable in the result type does not appear in the argument type.

Not all region-polymorphic functions create brand new values. Very often, a region-polymorphic function simply adds values to regions that are determined by the argument to the function. A good example is the list append function from the initial basis:¹

```
infixr 5 @
fun [] @ ys = ys
  | (x::xs) @ ys = x :: (xs @ ys)
val l = [1] @ [2,3]
```

Append successively conses the elements of the first list onto the second list. Thus, \mathbf{ys} and $\mathbf{xs} @ \mathbf{ys}$ must be in the same regions. However, the auxiliary pairs of \mathbf{xs} and \mathbf{ys} need not be in the same regions, although the elements of \mathbf{xs} and \mathbf{ys} clearly must be in the same regions, because they end up in the same list. These properties of the append function $@$ are summarised in its inferred region-annotated type scheme:

$$\forall \alpha \rho_{19} \rho_{17} \epsilon. [(\alpha, [\rho_{17}])\text{list}, (\alpha, [\rho_{19}])\text{list}] \xrightarrow{\epsilon. \{\text{get}(\rho_{17}), \text{put}(\rho_{19})\}} (\alpha, [\rho_{19}])\text{list}$$

When one writes a function, it is a good idea to consider whether one wants the function to create values in fresh regions or whether one wants it to add values to existing regions. Adding to existing regions can of course make these regions too large and long-lived, because the entire region will be alive for as long as one of the values in the region may be needed in the future.

The MulExp version of the append function is listed in Figure 6.2. At the application of $@$, the region annotated type scheme for $@$ is instantiated to the region annotated type

$$[(\text{int}, [\rho_{35}])\text{list}, (\text{int}, [\rho_4])\text{list}] \xrightarrow{\epsilon'. \{\text{get}(\rho_{35}), \text{put}(\rho_4)\}} (\text{int}, [\rho_4])\text{list}$$

To avoid passing regions that are never used, the MLKit introduces only formal region variables for those bound region variables in the type scheme for which there appears at least one **put** effect in the type of the function. Reading a value is done simply by following a pointer to the value, irrespective

¹File `kitdemo/append.sml`.

```

let fun @ at r1 [r19:INF] (v129, v130) =
  case v129 of
    nil => v130
  | :: v98 => :: (#0 v98, @[r19] (#1 v98, v130))at r19;
val l =
  let region r35:INF
  in @ [r4] ([1] at r35, [2,3] at r4)
  end
in { |@: (_,r1), l: _ | }
end

```

Figure 6.2: The region-annotated version of `append`.

of what region the value resides in, whereas storing a value in a region uses the name (see Section 2.1) of the region. This omitting of region parameters explains why ρ_{17} does not become a formal region parameter of `@` and why ρ_{35} is not passed to `@` at the call site. This optimisation, which is called *dropping of regions*, is the key reason why the MLKit takes the trouble to distinguish between `put` and `get` effects.

Here are two more examples to highlight the difference between functions that can put values in fresh regions and functions that add values to existing regions:

```

fun cp1 [] = []
  | cp1 (x::xs) = x :: cp1 xs
fun cp2 (l as []) = l
  | cp2 (x::xs) = x :: cp2 xs

```

Here `cp1` can copy the auxiliary pairs of a list into a fresh region, whereas `cp2` always copies the auxiliary pairs of a list into the same region:

$$\begin{aligned}
\text{cp1} &\mapsto \forall \alpha \rho \rho' \epsilon. (\alpha, [\rho]) \text{list} \xrightarrow{\epsilon. \{\text{get}(\rho), \text{put}(\rho')\}} (\alpha, [\rho']) \text{list} \\
\text{cp2} &\mapsto \forall \alpha \rho \epsilon. (\alpha, [\rho]) \text{list} \xrightarrow{\epsilon. \{\text{get}(\rho), \text{put}(\rho)\}} (\alpha, [\rho]) \text{list}
\end{aligned}$$

As we saw in Section 1.3, there are cases where it is useful to copy a list from one region into another region, so as to make it possible to deallocate the old region. This copying can be used as a kind of programmer-controlled garbage collection in cases where garbage has accumulated in the original region.

Because it is often useful to distinguish between functions that can put their result into fresh regions and functions that simply add to regions determined by their value argument, we shall refer informally to the former functions as *region exomorphisms* and the latter as *region endomorphisms*. Notice that these terms do not provide a clear-cut distinction. Often, functions have both an endomorphic and an exomorphic side to them. Also notice that even a region exomorphic function can be forced to act as an endomorphism by the calling context. As an example, consider the expression

```
if true then cp1 1 else 1
```

Because the two branches of the conditional are required to have the same region-annotated type with place, `1` and `cp1 1` are forced to be in the same regions.

6.4 Polymorphic Recursion

A recursive region-polymorphic function

```
fun f at  $\rho_0$  [ $\rho_1, \dots, \rho_k$ ] ( $x_1, \dots, x_n$ ) = e
```

may call itself inside its own body (e) with regions that are different from its own formal region parameter ($[\rho_1, \dots, \rho_k]$). This feature is called *polymorphic recursion in regions*, named after polymorphic recursion, the analogous concept for types. Polymorphic recursion in regions is vital for achieving good memory management in connection with recursion. Unfortunately, it also makes the region inference problem considerably more challenging, but that is a different story [TB98].

We now show a typical use of polymorphic recursion in regions, namely merge sorting of lists. The basic idea of merge sort is simple: first split the input list into two lists l and r of roughly equal length. Then sort l and r recursively and merge the results into a single sorted list. When programming with regions, we need to plan which of these lists we want to reside in the same regions. We do not want to waste space. In particular, if n is the length of the list, it would be quite irresponsible to use $O(n \log n)$ space, say. Let us aim at arranging that the sorting function is a region exomorphism that does not produce any values in its result regions except the sorted list. To sort n elements, we shall need n list cells (to hold the input list) plus roughly

$2 \times (n/2)$ list cells to hold l and r , the two lists that arise from splitting the input list. To sort l recursively, we need space for the two lists obtained by splitting l and so on. The space consumption grows to a maximum of $3n$ list cells (including the n cells to hold the input), before any merging is done. By the time all of l is sorted, that is, just before r is sorted recursively, we have the following lists: the input (n cells), l ($n/2$ cells), l sorted ($n/2$ cells), r ($n/2$ cells). Continuing this way, at the rightmost merge of two lists of length at most one, approximately $4n$ list cells are live. Then a series of final merges occur. Code that uses these ideas is listed in Figure 6.3.²

```

fun cp [] = []
  | cp (x::xs) = x :: cp xs

(* exomorphic merge *)
fun merge (xs, []) : int list = cp xs
  | merge ([], ys) = cp ys
  | merge (l1 as x::xs, l2 as y::ys) =
    if x<y then x :: merge(xs, l2)
    else y :: merge(l1, ys)

(* splitting a list *)
fun split (x::y::zs, l, r) = split(zs, x::l, y::r)
  | split ([x], l, r) = (x::l, r)
  | split ([], l, r) = (l, r)

(* exomorphic merge sort *)
fun msort [] = []
  | msort [x] = [x]
  | msort xs = let val (l, r) = split(xs, [], [])
                in merge(msort l, msort r)
                end
end

```

Figure 6.3: Merge sorting of lists.

²MLB-file `kitdemo/msort.mlb`, file `kitdemo/msort.sml`. To compile the project, go to the `kitdemo` directory and execute `mlkit msort.mlb` from the shell. MLKit places an executable file `run` in the `kitdemo` directory. For an in-depth description of how to compile and run MLB-files and SML-files, see Chapter 15.

The exomorphic merge function is a bit inefficient in that it copies one argument when the other is empty, but the exomorphism ensures that `msort l` and `msort r` are not forced into the same regions. The polymorphic recursion in regions makes it possible for `xs`, `l`, `r`, `msort l`, and `msort r` all to be in distinct regions. For example, in the call `msort l`, the polymorphic recursion makes it possible for `l` to be in a region different from `xs` and it also makes it possible for the result of the call to be in a region different from the result of `msort xs`.

Based on the above analysis we conclude that the space required by `msort xs` is approximately $4nc_1 + c_2 \log_2 n$ plus the extra stack space required for the final merges, where n is the length of `xs`, c_1 is the size of a list cell (2 words in this case) and c_2 is the space on the runtime stack used by one recursive call of `msort` (probably less than 10 words).

Because `merge` is not tail-recursive, a merge requires space both for its two input lists, for its output list, and for temporaries stored on the stack. When one of the lists becomes empty, `merge` calls `cp`, which allocates less for each iterative call than `merge` does. Each return from `merge` allocates a list cell (two words), so the maximum space usage is reached when the last element of the result of the merge is constructed (which happens when the recursion is deepest). Here the space used is (we show $n = 200,000$ list elements as an example)

data	size (words)	$n = 200,000$
input list	$2n$	3,200,000 bytes
l	n	1,600,000 bytes
l sorted	n	1,600,000 bytes
r	n	1,600,000 bytes
r sorted	n	1,600,000 bytes
result list	$2n$	3,200,000 bytes
total in regions	$8n$	12,800,000 bytes

It turns out that each iterative call of `merge` pushes three registers on the stack, so that stack size will be approximately $3 \times 8 \times n$ bytes, which for $n = 200,000$ is 4,800,000 bytes, just before the merged list is constructed, bottom-up. The total space consumption for sorting 200,000 integers should therefore be roughly 14,400,000 bytes (i.e., 12,800,000-3,200,000+4,800,000).

To check the above analysis, we sorted 200,000 integers with the region profiler enabled. As one sees in Figures 6.4 and 6.5, the space usage found by region profiling correspond well to the results of our analysis.

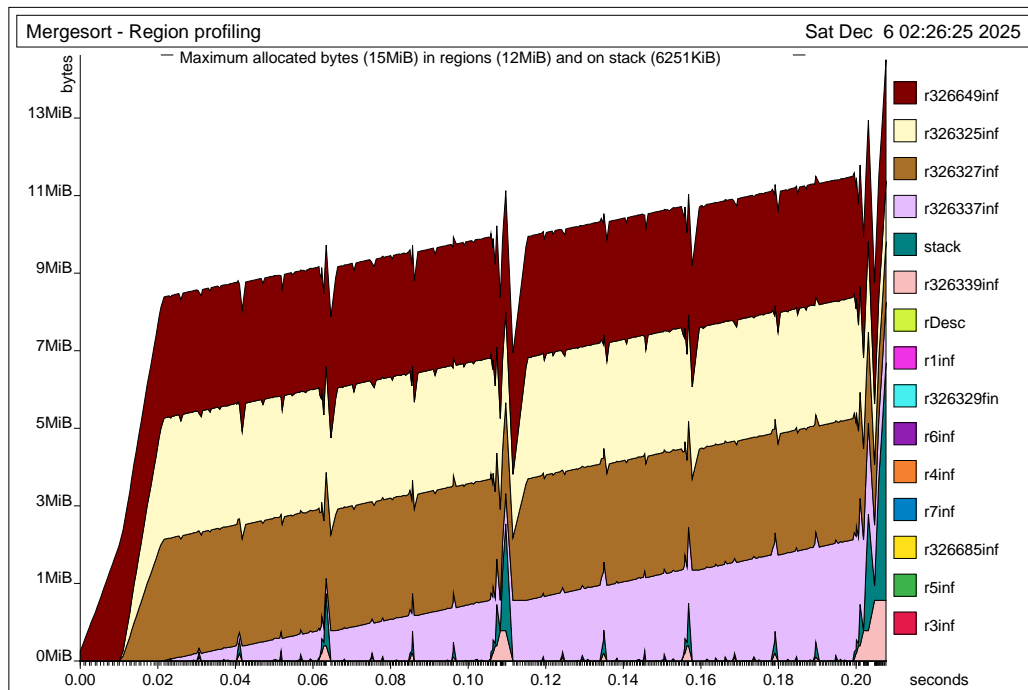


Figure 6.4: Region profiling of `msort` sorting 200,000 integers. The high-level mark denotes the maximum amount of memory allocated in regions and on the stack. Notice that the amount of memory used in regions and the amount of memory used on the stack may not top on the same time, which shows by the high-level mark being lower than the sum of the maximum stack usage and the maximum memory used in regions.

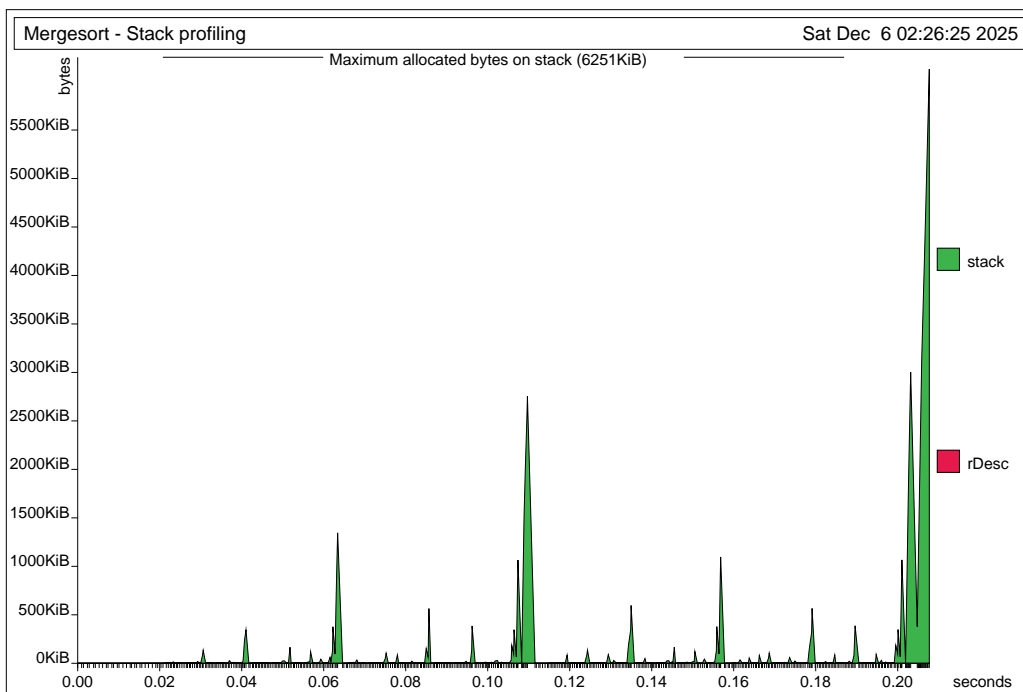


Figure 6.5: Stack profiling of `msort` sorting 200,000 integers.

In Chapter 12, we shall see how one can use resetting of regions to reduce the space usage drastically, to roughly $2nc_1$.

Chapter 7

Value Declarations

Although region inference is based on types and effects, it is also to some extent syntax dependent. That is, two programs that are equivalent in their input-output behavior can easily have very different memory behavior. In this chapter, we discuss how to write declarations so as to obtain good results with region inference. The region inference rules that underlie the MLKit with Regions are related to the scope rules of ML, so we start by a (very informal) summary of the scope rules of ML declarations.

7.1 Syntax

A Standard ML *value declaration* binds a value to a value variable. For example, the result of evaluating the value declaration

```
val x = 3 + 4
```

is the environment $\{x \mapsto 7\}$. More generally, evaluation of a value binding `val id = exp` proceeds as follows. Assume the result of evaluating *exp* is a value, *v*. Then the result of evaluating `val id = exp` is the environment $\{id \mapsto v\}$.

The value declaration is just one form of Core Language declaration (the others being type and exception declarations). We use *dec* to range over declarations. Declarations can be combined in several ways. For example,

$$dec_1; dec_2$$

is a *sequential declaration*. The identifiers declared by this declaration are the identifiers that are declared by *dec*₁ or *dec*₂; moreover, identifiers declared

in dec_1 may be referenced in dec_2 . The semicolon is associative. Thus, in a sequence $dec_1; \dots; dec_n$ of declarations, identifiers declared in dec_i may be referenced in dec_{i+1}, \dots, dec_n ($1 \leq i \leq n$).

The Core Language has two forms of local declarations. The expression

```
let dec in exp end
```

declares identifiers whose scope does not extend beyond exp . Similarly, the declaration

```
local dec1 in dec2 end
```

first declares identifiers (in dec_1) whose scope does not extend beyond dec_2 and then uses these declarations to perform the declarations in dec_2 . An identifier is declared by the entire local construct if and only if it is declared by dec_2 .

7.2 Scope Versus Lifetime

Scope is a syntactic concept: a declaration of an identifier contains a binding occurrence of the identifier; the scope of the declaration is the part of the ensuing program text whose free occurrences of that identifier are bound by that binding occurrence. By contrast, lifetime, as we use the word, is a dynamic concept. A value is “live” if and only if the remainder of the computation uses it (or part of it). The traditional stack discipline couples these two concepts very closely. For example, in the pure stack discipline, the evaluation of

```
let dec in exp end
```

in an environment E proceeds as follows. First evaluate dec to yield an environment, E_1 . Then evaluate exp , in the environment E extended with E_1 , to yield value v . Then v is the result of evaluating the `let` expression in E . In implementation terms: first push an environment E_1 onto the stack, use it to evaluate the expression in the scope of the declaration, and then pop the stack. That this idea works in block-structured languages hinges on a number of carefully made language design decisions. In functional and object-oriented languages, memory cannot be managed that simply. The problem is that while environments can be managed in a stack-like manner, the values in the range of the environment cannot (unless one uses regions, that is). For example consider the ML expression:

```

local
  val private = [2,3,5,7,11,13]
in
  fun smallPrime (n:int) : bool = List.member n private
end

```

Although the scope of the declaration is only the declaration of `smallPrime`, `private` is accessed (at runtime) whenever `smallPrime` is called. Thus, the lifetime of the list of small primes is at least as long as the lifetime of the `smallPrime` function itself.

The region discipline still has a coupling between scope and lifetimes, but, because we want to be able to handle recursive data types and higher-order functions, the coupling is less tight. The ground rule of region inference is that as long as a value variable is in scope, the value bound to it at runtime will remain allocated. More precisely:

Ground Rule: The region rules forbid transforming an expression *exp* into `let region ρ in exp end` if *exp* is in the scope of an identifier that has ρ free in its region-annotated type scheme with place.

For an example, consider

```

let
  val list = [1,2,3]
  val n = length list
  val r = sin(real n)
in
  cos(r)
end

```

At runtime, the list bound to `list` is not used (i.e., it is not live) after its length has been computed; similarly, the value of `n` is not live after it has been converted to a floating point number, and so on. In short, at runtime, we have a sequence of short, non-overlapping lifetimes.

With region inference, however, the list bound to `list` will stay allocated throughout the evaluation of the remainder of the `let` expression.¹

¹One can force deallocation of the list by inserting `val _ = resetRegions(list)` after the declaration of `n`; but, as we shall see, there are less draconian ways of achieving the same result.

For a more interesting example of the consequences of the Ground Rule, consider the following declarations, taken from a program that computes prime numbers using the Sieve of Eratosthenes:

```

fun cp [] = []
  | cp (x::xs) = x :: cp xs

fun sift (n, []) = []
  | sift (n, (x::xs)) = if x mod n = 0 then sift(n,xs)
                       else x::sift(n,xs)

fun sieve (a as ([], p)) = a
  | sieve (x::xs, p) = let val rest = sift(x,xs)
                       in sieve(cp rest,x::p)
                       end

```

Here `sift(n, 1)` produces a list of the numbers from 1 that are not divisible by `n`; `sieve(xs, p)` repeatedly calls `sift`, adding primes to the front of `p`, until the list of numbers remaining in the sieve becomes empty. The programmer has employed the copying technique suggested in Section 1.3 to avoid that the lists that are bound to `rest` during the repeated filtering all are put in the same region. The programmer's intention is that the `cp rest` should overwrite `x::xs` by a copy of `rest`, so that space consumption would be bounded by a constant times the size of the input. But it does not work as intended; because `rest` is in scope at the recursive application of `sieve`, the list that is bound to `rest` will stay allocated for the duration of that call, which is in fact the remainder of the entire computation!

In many cases, the solution is simply to shorten the scope of the declaration. In the above example, a good solution is to move the application of `sieve` outside the `let`:

```

fun sieve (a as ([], p)) = a
  | sieve (x::xs, p) =
      sieve let val rest = sift(x,xs)
            in (cp rest,x::p)
            end

```

That the copying really overwrites the input list relies, in part, on region resetting (Chapter 12). But it also relies on region polymorphism and on the Ground Rule. Rewriting the application of `sieve` ensures that the list bound to `rest` will not live to see the recursive call of `sieve`. Unless forced by context to do otherwise, `sift` will create a list using fresh regions. Because

`cp` is also exomorphic, there will be no sharing between `rest` and the other lists. The region variable that denotes the region that holds the auxiliary pairs of `rest` appears in the effect of the (revised) `let` expression. However, this region variable does not occur free in the region-annotated type scheme with place of any value variable in scope at that point, not even in the region-annotated type scheme with place of `sieve`, which only has the region that contains `sieve` itself free in its region-annotated type scheme with place. Consequently, region inference wraps the `let` expression by a `region` declaration of the region variable in question:

```
fun sieve (a as ([], p)) = a
  | sieve (x::xs, p) =
    sieve let region r10
          val rest = sift[r10](x,xs)
          in (cp rest,x::p)
    end
```

7.3 Shortening Lifetime

Informally, region inference forces the lifetime of an identifier to be at least its scope. Improving memory performance therefore sometimes requires making scopes of identifiers smaller. Useful program transformations include:

Inwards let floating

Transform

```
let val id1 = exp1 val id2 = exp2 in exp end
```

into

```
let val id2 = let val id1 = exp1 in exp2 end in exp end
```

provided `id1` does not occur free in `exp`.

Application extrusion:

Transform

```
let dec in f(exp) end
```

into

```
f let dec in exp end
```

provided f is an identifier that is not declared by dec .

Application extrusion is particularly useful in connection with tail recursion; the reader will see it employed several times in what follows.

Chapter 8

Static Detection of Space Leaks

“Space leak” is the informal term used when a program uses much more memory than one would expect, typically because of memory not being recycled as early as it could (or not at all).

If a region-polymorphic function with region-annotated type scheme σ has a **put** effect on a region variable that is not amongst the bound region variables of σ , then one quite possibly has a space leak; every application of the function may write values into a region that is the same for all calls of the function. For example, consider the source program¹

```
fun g () =  
  let val x = [5,7]  
      fun f y = (if y>3 then x@x else x;  
                5)  
  in  
    f 1; f 4  
  end
```

Here **f** has type `int → int`; yet, when the expression `y>3` evaluates to `true`, an append operation is performed that produces a list in the same region as `x`. The first call of **f** will not cause the append operation to be called, but the second one will. One can say that **f** has a space leak in that it can write values into a more global region, namely a region that is allocated at the beginning of the body of **g**. The sequence of calls to **f** accumulates copies of `x@x` in that region, although none of these lists are accessible anywhere. In this particular case, the values are not even part of the result type of **f**,

¹Program `kitdemo/escape.sml`.

so the writing is a side-effect at the implementation level, even though there are no references in the program.

The region-annotated type scheme inferred for `f` is

$$\forall \epsilon. \text{int} \xrightarrow{\epsilon. \{\text{put}(\rho_{11})\}} \text{int}$$

where the region-annotated type of `x` is

$$(\text{int}, [\rho_{11}])\text{list}$$

Here we see that ρ_{11} is free in the region-annotated type scheme and appears with a `put` effect.

8.1 Warnings About Space Leaks

MLKit can be instrumented to issue a warning each time it meets a function that is declared using `fun` and has a free `put` effect occurring somewhere in its type scheme. The way to tell the MLKit to issue the warnings is by passing the option `-warn_on_escaping_puts` to the MLKit compiler. In practice, this warning mechanism is a valuable device for predicting space leaks. The region-annotated version of our example function `g` is listed in Figure 8.1. During compilation of `g`, the MLKit issues the following warning:²

```
*** Warnings ***
f has a type scheme with escaping put effects on region(s):
r11, which is also free in the type schemes with places of : x
```

We are told that the program might space leak in region `r11`. Looking at the function `f`, we see that this region is an actual region parameter to `@`. It follows that the problem is the call to `@`.

8.2 Fixing Space Leaks

Often one can fix a space leak by delaying the creation of the value that causes the space leak. In the above example, we can move the construction of the list into `f`:³

²To provoke the warning, one has to disable in-lining in the Lambda optimiser; this is done by passing the option `-maximum_inline_size 0` (or `-max_inl_sz 0`) to the MLKit compiler together with the option `-warn_on_escaping_puts`.

³Program `kitdemo/escape1.sml`.

```

let fun g at r1 () =
  let region r11:INF;
      val x = [5,7] at r11;
      region r21:1;
      fun f at r21 y =
        let val _ = case y > 3 of
                    true => @[r11] (x, x)
                    | _ => x
                in 5
                end;
          val _ = f 1
        in f 4
        end
  in {!g: (_,r1)!}
end

```

Figure 8.1: The region-annotated version of `g`.

```

fun g () =
  let fun mk_x () = [5,7]
      fun f y = let val x = mk_x()
                in if y>3 then x@x else x; 5
                end
      in f 1; f 4
  end

```

Of course, this means that the list will be reconstructed upon each application of `f`. Another solution is to move the creation of the list as close to the calls as possible and then pass the list as an extra argument:⁴

```

fun g () =
  let fun f (x,y) = (if y>3 then x@x else x; 5)
      in let val x = [5,7]
          in f(x, 1); f(x, 4)
          end
      end
  end

```

Both solutions stop warnings from being printed, but the second solution is

⁴Program `kitdemo/escape2.sml`.

better than the first: `f` still has a **put** effect on the regions containing `x`, but the difference is that these are now represented by bound region variables in the type scheme of `f`. This quantification has the advantages that (1) allocation of space for the list is delayed until the list is actually used and (2), the list can be deallocated after the calls have been made (whereas in the original version, `x` occurs free in the declaration of `f` and will be kept alive as long as `f` can be called.)

At other times, there is no clean way of avoiding escaping **put** effects. One example is found in the `TextIO` structure of the Basis Library:

```
exception CannotOpen
fun raiseIo fcn nam exn =
  raise IO.Io {function = fcn^"", name = nam^"", cause = exn}

fun openIn (f: string) : instream =
  {ic=prim("openInStream", (f, CannotOpen)),
   name=f} handle exn => raiseIo "openIn" f exn

fun openOut(f: string): ostream =
  {oc=prim("openOutStream", (f, CannotOpen)),
   name=f} handle exn => raiseIo "openOut" f exn
```

As explained in Chapter 11, when a unary exception constructor is applied to a value, both the argument value and the resulting constructed value are forced into a particular global region. Thus, the application

```
IO.Io {function = fcn^"", name = nam^"", cause = exn}
```

has a potential space leak in it; every time we apply the exception constructor, the resulting exception value will be put into a global region. This particular space leak is perhaps not something that would keep one awake at night, because most programs do not make a large number of failed attempts to open files, but it is useful to be warned about this potential problem. Notice, however, that the string arguments to `raiseIo` are copied inside the body of `raiseIo`, so that they are not forced to be placed in the global string region.

Chapter 9

References

Section 9.1 gives a brief summary of references in Standard ML; it may be skipped by readers who know SML. Thereafter, we discuss runtime representation of references and region-annotated reference types.

9.1 References in Standard ML

A reference is a memory address (pointer). Standard ML has three built-in operations on references

<code>ref</code>	$\mapsto \forall \alpha. \alpha \rightarrow \alpha \text{ ref}$	create reference
<code>!</code>	$\mapsto \forall \alpha. \alpha \text{ ref} \rightarrow \alpha$	de-referencing
<code>:=</code>	$\mapsto \forall \alpha. \alpha \text{ ref} * \alpha \rightarrow \text{unit}$	assignment

If the type of a reference r is $\tau \text{ ref}$ then one can store values of type τ (only) at address r . A reference is a value and can therefore be bound to a value identifier by a `val` declaration. While the value stored at a reference may change, the binding between variable and reference does not change. We show an example, because this point can be confusing to programmers who are familiar with mutable variables in languages like C and Pascal:

```
val it =
  let val x : int ref = ref 3
      val y : bool ref = ref true
      val z : int ref = if !y then x else ref 5
  in z := 6; !x
  end
```




Figure 9.1: Creating a reference allocates one word in a region on the region stack. Here, the region is drawn as a finite region, but it could equally well be infinite.

Because `!y` evaluates to true, `z` becomes bound to the same reference (`r`) as `x`. So, the subsequent assignment to `z` changes the contents of the store at address `r` to contain 6. Because `x` and `z` are aliases, the result of the `let` expression is the contents of the store at address `r` (i.e., 6).

9.2 Runtime Representation of References

MLKit translates an SML expression of the form `ref exp` into an expression of the form (assuming `exp` translates into `e`)

$$\text{ref at } \rho \ e$$

which is evaluated as follows. First `e` is evaluated. Assume that this evaluation yields a value `v`. Here `v` may be a boxed or an unboxed value. Next, a 64-bit word is allocated in the region denoted by `ρ`; let `r` be the address of this word. Then `v` is stored at address `r` and `r` is the result of the evaluation.

The situation is depicted in Figure 9.1. The value `v` can be unboxed as shown in Figure 9.2. Or it may be boxed, in which case `v` is an address.

Notice that a reference really is a pointer in the implementation. In particular, a reference is not tagged, so the register allocator may choose to store a particular reference in a register. The content of the reference is also always one word, either an unboxed value (e.g., an integer or a boolean) or a pointer (if the content is boxed). So the content of a reference is not tagged either.

De-referencing a reference `r` is done by reading the content of the memory location `r`. Notice that de-referencing does not require knowledge of what region the word with address `r` resides in.




Figure 9.2: Creating a reference allocates one word in a region on the region stack. Here, the region is drawn as a finite region, but it could equally well be infinite.

Assigning a value v to a reference r simply stores v in the memory at address r . When v is an unboxed value, the assignment can be regarded as copying v into the memory cell r ; otherwise v is a pointer, which the assignment stores in the memory cell r . Either way, assignment is a constant-time operation.

9.3 Region-Annotated Reference Types

The general form of a region-annotated reference type is:

$$(\mu \text{ref}, \rho)$$

Informally, a reference r has this type if it is the address of a word in the region denoted by ρ and, moreover, μ is the region-annotated type with place of the contents of that word. For example, assume ρ is bound to some region name, say `r35`; then the evaluation of the declaration `val x = ref at ρ 3` results in the environment $\{x \mapsto r\}$, where r is the address of a word with contents 3 residing in region `r35`, see Figure 9.2. The type of the variable `x` is `(int ref, ρ)`.

References are treated like all other values by region inference. The region-annotated type schemes given to the three built-in operations are:

$$\begin{aligned} \text{ref} &\mapsto \forall \alpha \rho \epsilon. \alpha \xrightarrow{\epsilon.\{\text{put}(\rho)\}} (\alpha \text{ref}, \rho) \\ ! &\mapsto \forall \alpha \rho \epsilon. (\alpha \text{ref}, \rho) \xrightarrow{\epsilon.\{\text{get}(\rho)\}} \alpha \\ := &\mapsto \forall \alpha \rho \epsilon. [(\alpha \text{ref}, \rho), \alpha] \xrightarrow{\epsilon.\{\text{get}(\rho)\}} \text{unit} \end{aligned}$$

The type scheme for `:=` has in it a `get` effect on the region holding the reference. Although the operator does not actually read the value, the presence of

```

let val it =
  let region r9:INF;
    val x = ref at r9 3;
    region r11:1;
    val y = ref at r11 true;
    val z = case !y of true => x | _ => ref at r9 5;
    val _ = z := 6
  in !x
  end
in {|it: _|}
end

```

Figure 9.3: Region-annotated reference creation.

the value is necessary for it to be updated. Assigning a value v to a reference r does not make a copy of v (unless v is unboxed). Instead, `:=` updates the content of the reference r to point to v .

The advantage of the chosen scheme for handling references is that reference creation, de-referencing, and assignment all are constant-time operations. The disadvantage is that if two values may be assigned to the same reference, then they are forced to be in the same regions (cf. the region-annotated type schemes given above).

If we compile the example from Section 9.1, we get the program shown in Figure 9.3.¹ The region denoted by `r9` contains the memory word whose address is bound to `x` and `z`, and whose contents is first 3, then 6. The region denoted by `r11` contains a single boolean. Also notice that the word containing 5 is designated `r9`, because the `then` and `else` branches must be given the same region-annotated type with place. Finally, notice that all references will be reclaimed automatically at the end of the `let` construct that bind `r9` and `r11`.

9.4 Local References

References that are created locally within a function and that do not escape the function naturally reside in regions that are local to the function body.

¹Program `kitdemo/refs3.sml`.

For example, the declaration:²

```
fun id x = let val r = ref x in ! r end
```

is compiled into

```
let fun id at r1 x =
      let region r11:1 in !(ref at r11 x) end
in  {|id: (_,r1)|}
end
```

Here `r11` will be implemented as one word on the runtime stack. The evaluation of `ref at r11 x` moves the argument `x` to that word on the stack. At the end of the inner `let`-scope, the word (i.e., region `r11`) is popped off the stack.

Now, let us turn to an example of a memory cell whose lifetime extends the scope of its declaration, because it is accessible via a function (in Algol terminology, the reference is an *own variable* of the function.)³

```
local
  val r = ref ([]:string list)
in
  fun memo_id x = (r := x :: !r; x)
end
val y = memo_id "abc"
val z = memo_id "efg";
```

Provided that in-lining by the optimiser is restricted to in-line only those functions that are applied once,⁴ this example compiles into

```
let val r = ref at r6 nil;
      fun memo_id at r1 x =
          let val _ = := (r, :: (x, ! r)at r4) in x end;
      val y = memo_id "abc";
      val z = memo_id "efg"
in  {|memo_id: (_,r1), y: (_,r3), z: (_,r3)|}
end
```

and the MLKit warns us that there is a possible space leak:⁵

²Program `kitdemo/refs1.sml`.

³Program `kitdemo/refs2.sml`.

⁴To restrict the optimiser accordingly, provide the option `-maximum_inline_size 0` (or `-max_inl_sz 0`) to the MLKit compiler.

⁵Warnings are printed only if the option `-warn_on_escaping_puts` is passed to the MLKit compiler along with the option `-max_inl_sz 0`. See Chapter 8.

```
*** Warnings ***  
memo_id has a type scheme with escaping put effects on region(s):  
r4, which is also free in the type schemes with places of : r
```

9.5 Hints on Programming with References

There is no need to shy away from using references when programming with regions. However, one needs to be aware of the restriction that values that may be assigned to the same references are forced to live in the same region, and that this region with all its values will be alive for as long as the reference is live. If the contents type is unboxed (e.g., `int`), there is no problem, for in that case, no region for the contents is allocated. But one should avoid creating long-lived references that are assigned many different large values.

Chapter 10

Recursive Data Types

This chapter describes how the MLKit treats recursive data types. We have already seen how one recursive datatype, namely lists, is handled. This chapter deals with the general case.

10.1 Spreading Data Types

MLKit performs an analysis called “spreading of data types”. Spreading of datatypes analyses `datatype` declarations. This analysis of a `datatype` declaration uses information about the type constructors that appear in the types of the constructors of the data type(s) introduced by the declaration, but it does not use information about the use of the data type.

Spreading determines (a) a so-called arity of every type name that the data type declaration introduces and (b) a region-annotated type scheme for every value constructor introduced by the data type declaration.

In the Definition of Standard ML every type name has an attribute, called its arity [MTHM97, page 15]. The arity of a type name is the number of type arguments it requires. For example, `int` has arity 0 while the type name introduced by the following declaration of binary trees has arity 1:

```
datatype  $\alpha$  tree = Lf | Br of  $\alpha$  *  $\alpha$  tree *  $\alpha$  tree;
```

MLKit extends the notion of arity (in its internal languages) to account for regions and effects. For lists, for example, we need a region for holding the pairs to which `::` is applied. For the type

```
datatype  $\alpha$  foo = A | B of ( $\alpha$  *  $\alpha$ ) * ( $\alpha$  *  $\alpha$ )
```

the type of \mathbf{B} introduces the possibility of three region variables (one for each star). Region variables that are induced by the types of constructors and that do not hold the constructed values themselves are called *auxiliary region variables*. For example, the `list` type

```
datatype  $\alpha$  list = nil | op :: of  $\alpha$  *  $\alpha$  list
```

has one auxiliary region variable, namely the region variable that describes where the pairs of type `α * α list` (i.e., the auxiliary pairs), reside.

Besides auxiliary regions, one sometimes needs auxiliary effects. For an example, consider the type declaration

```
datatype  $V$  = N of int | F of  $V \rightarrow V$ 
```

Here one needs an arrow effect for the function type `$V \rightarrow V$` . We refer to such an arrow effect as an *auxiliary arrow effect* of the data type in question.

We define the (*internal*) *arity* of a type name t to be a triple (n, k, m) of non-negative integers, where n is the usual Standard ML arity of the type name, k is the *region arity* of t , and m is the *effect arity* of t . The region and effect arities indicate the number of auxiliary regions and arrow effects of the data type, respectively.

For efficiency purposes, we have found it prudent to restrict the maximal number of auxiliary regions a data type can have to one for each kind of runtime type of regions and to restrict the maximal number of auxiliary effects to 1. Otherwise, the number of auxiliary regions can grow exponentially in the size of the program:

```
datatype t0 = C
datatype t1 = C1 of t0 * t0
datatype t2 = C2 of t1 * t1
...
```

Here the number of auxiliary region variables would double for each new data type declaration. Furthermore, all type names introduced by a `datatype` declaration are given the same arity (a `datatype` declaration can declare several types simultaneously).

Because of the limit on the number of auxiliary region variables, spreading of data type declarations sometimes unifies two auxiliary region variables that would otherwise be distinct; but it only unifies auxiliary region variables that have the same runtime type. The practical consequence of these restrictions is that applying a constructor to a value v sometimes forces identification of regions of v that hold otherwise unrelated parts of v .

The automatic memory management that we have discussed for lists extends to other recursive data types without problems. For example, binary trees are put into regions and are subsequently deallocated (in a constant time operation) when the region is popped. The next section goes through an example to illustrate the point.

MLKit attempts to use an unboxed representation for value constructors when possible. We have already seen how `cons` (i.e., `::`) and `nil` use an unboxed representation (Chapter 5). MLKit uses an advanced scheme for unboxing that uses both the unused least significant bits and the unused most significant bits of pointer-values for tagging constructors [Els24a]. Section 10.3 documents this scheme and describes how the programmer can learn about MLKit's unboxing decisions. Before presenting unboxing scheme in details, we first give an example of a tree `datatype` declaration for which MLKit will use an unboxed representation that utilises the unused least-significant bits of pointer values for storing constructor tags.

10.2 Example: Balanced Trees

Consider the program in Figure 10.1.¹ We would hope that the balanced tree produced by `balpre` is removed after it has been collapsed into a list by `preord`. And indeed it is. Here is the proof:

```

val it =
  let region r125:INF
  in print
    let region r127:INF
    in implode[r125]
      let region r131:INF
      in preord[r127]
        (let region r135:INF
        in balpre[r131]
          let region r139:1
          in explode[r135] "Greetings from the Kit\n"
          end
        end,
        nil
      )
  )

```

¹MLB-file: `kitdemo/trees.mlb`, file `kitdemo/trees.sml`.

```

datatype  $\alpha$  tree = Lf | Br of  $\alpha$  *  $\alpha$  tree *  $\alpha$  tree

(* preorder traversal of tree *)
fun preord (Lf, xs) = xs
  | preord (Br(x,t1,t2),xs) =
    x::preord(t1,preord(t2,xs))

(* building a balanced binary tree from a list *)
fun balpre [] = Lf
  | balpre (x::xs) =
    let val k = length xs div 2
    in Br(x, balpre(take(xs, k)),
          balpre(drop(xs, k)))
    end

(* preord o balpre is the identity *)
val it = print(implode(preord(balpre(explode
  "Greetings from the MLKit\n"), [])));

```

Figure 10.1: Example showing recycling of memory used for an intermediate data structure. The function `balpre` builds a balanced binary tree from a list and `preord` then flattens the tree to a list (after which the tree is garbage).

```

        end
    end
end

```

The exomorphic behavior of `balpre` causes the tree to be allocated in region `r131`, which is deallocated after the call to `preord`.

This is the kind of certainty about lifetimes we are aiming at. Imagine, for example, that the trees under consideration were terms representing different intermediate forms in a compiler. Then one would like to know that (possibly large) syntax trees are not kept in memory longer than needed.

10.3 Unboxing Schemes

MLKit uses a uniform representation of values, which is important for compiling generic code (e.g., functions) separately from the code that uses it. Still, under these constraints there are many possibilities for securing a compact data representation. MLKit uses tagged regions for some types of values (instead of tagging the values themselves) and, as we have seen for lists and binary trees, datatypes with a single unary constructor that takes boxed arguments are implemented using the lower-bit tags in pointers to discriminate between the constructors.

MLKit does not stop here. It tries hard to represent data type constructors unboxed by utilising (otherwise) non-used bits in pointer-values. On most modern architectures, pointers are aligned to point only at word-addresses, which leaves the three least-significant bits unused and which allows MLKit to represent lists and trees unboxed, as we have seen earlier. Moreover, only the least 48 bits are used for representing pointers; the remaining 16 most-significant bits are unused on modern 64-bit architectures.

When we have a datatype t with u unary constructors and n nullary constructors, if $0 \leq u < 2^{16}$ and if all arguments to the unary constructors are known to be boxed (or not using the most-significant bits), we can represent the type t unboxed, using the high bits for tagging. If $u = 0$, we call the type t an *enum-unboxed* type. We call the type t a *low-unboxed* type if $u = 1$ and the argument to the unary constructor is boxed, and a *high-unboxed* type if $u > 1$ and if all arguments to the unary constructors are known to be boxed. In the worst case, when none of the above conditions are satisfied, MLKit uses a boxed representation of the datatype t . When t is a high-unboxed type, we use the 16 most significant bits to discriminate between the unary and

nullary constructors (for GC safety, MLKit also sets the least significant bit for the nullary constructors). MLKit uses an iterative algorithm to determine the boxity for mutually recursive `datatype` declarations, which also include data types with a single unary constructor. For simplicity of region inference, region inference assumes that, for mutually recursive `datatype` declarations, either all the declared types are unboxed or they are all boxed.

As an example, MLKit implements the following datatype as a high-unboxed type:²

```
datatype t = Leaf of string | Empty | Children of t list
```

MLKit recognises that the arguments to t 's unary constructors are either boxed or use only lower bits for tagging (i.e., `t list`). The higher bits are thus available for discriminating between the three constructor values. We can get MLKit to report on boxing decisions by passing the compile-time flag `-report.boxities`:

```
$ mlkit -no_gc -report_unboxing high-unb.sml | grep Boxities
*** Boxities: [t:hub]
```

Indeed, MLKit reports that it has decided for t to be high-unboxed.

More formally, MLKit distinguishes between the following *boxities* (κ):

κ	::=	<code>lub</code>	low-unboxed
		<code>hub</code>	high-unboxed
		<code>box</code>	boxed
		<code>enum</code>	enumeration (only nullary-constructors)
		<code>single κ</code>	single unary-constructor datatype

When MLKit analyses a set of simultaneously declared data types, it starts with an optimistic guess for each type and checks if the guess is valid according to a series of rules:

1. No unary constructor of a type with boxity `hub` can take a value as argument that uses the upper (i.e., most significant) 16 bits.
2. No type with boxity `lub` can have more than one unary constructor. We admit here that MLKit could relax this condition and allow up to four unary constructors.

²Program `kitdemo/high-unb.sml`.

3. Unary constructors of types with boxity `lub` must take boxed arguments.

If a rule is violated, an attempt is made with less optimistic assumptions.

Many types may be implemented unboxed using this scheme, including data types for representing grammars, datatypes for representing union-find structures, and data types representing finite maps, such as patricia trees [Els24a].

To obtain good boxity decisions, it may sometimes be beneficial for a programmer to make sure that an argument to a constructor is indeed boxed. Consider the following mutually recursive `datatype` declaration:³

```
datatype exp = LET of dec * exp | ADD of exp * exp | INT of int
and dec = VALBIND of string * exp
```

MLKit reports the following boxity decisions:

```
$ mlkit -no_gc -report_boxities expdec.sml | grep Boxities
*** Boxities: [exp:box,dec:box]
```

We see that both `exp` and `dec` are represented boxed, which means that not only will the arguments to `LET`, `ADD`, and `VALBIND` be represented boxed in designated regions, each value of the `exp` and `dec` types will be represented as a pointer to two region-allocated words, one word for storing the constructor tag and one word holding a pointer to the constructor argument (in the case of the `INT` constructor, the argument will be unboxed).

If instead, we arrange that the argument to the `INT` constructor is boxed, the algorithm will determine that `exp` is high-unboxed and `dec` has boxity `single box`.⁴ For arranging that the argument to the `INT` constructor is boxed, we enclose the integer in a singleton record (`INT of {value:int}`), which MLKit will always represent boxed:

```
$ mlkit -no_gc -report_boxities expdec2.sml | grep Boxities
*** Boxities: [exp:hub,dec:single box]
```

We see that now `exp` is represented unboxed and `dec` is represented unboxed and without storing constructor tags at all!

³Program `kitdemo/expdec.sml`.

⁴Program `kitdemo/expdec2.sml`.

Chapter 11

Exceptions

Standard ML exception constructors are introduced by *exception declarations*. The two most basic forms are

```
exception excon
```

and

```
exception excon of ty
```

for introducing nullary and unary exception constructors, respectively.

Exception declarations need not occur at top level. For example, a function body may contain exception declarations.

11.1 Exception Names

Each evaluation of an exception declaration creates a fresh *exception name* and binds it to the exception constructor. This is sometimes referred to as the *generative* nature of Standard ML exceptions.

MLKit implements an exception name as a pointer to a pair consisting of an integer and a string pointer; the string pointer points to the name of the exception, which is a global constant in the target program. The string is used for printing the name of the exception if it ever propagates to top level. The memory cost of creating the pair is, as always with pairs, two words.

11.2 Exception Values

Standard ML has a type `exn` of *exception values*. An exception value is either a *nullary* exception value or a *constructed* exception value. A nullary exception value is a pointer to a word that points to an exception name. A constructed exception value is a pair (en, v) of an exception name en and a value v ; we refer to v as the *argument* of en . This representation of exception values allows for the exception name of an exception value to be fetched in the same way irrespective of whether the exception value is nullary or constructed.

Referring to a nullary exception constructor allocates no memory. By contrast, applying a unary exception constructor to an argument constructs a constructed exception value. The memory cost of such an application is two words for holding the pair (en, v) .

The distinction between nullary and unary exception constructors is important in the MLKit because our region inference analysis takes a simple-minded approach to exceptions:

All exception names and nullary exception values are put into a certain global region and thus never reclaimed automatically. A constructed exception value is put in a region that is live at least as long as the exception constructor is in scope.

We therefore make the following recommendations:

1. Put exception declarations at top level, if possible. That way, the memory required by exception names will be bounded by the program size.
2. Avoid applying unary exception constructors frequently; there is no harm in raising and handling constructed exception values frequently; it is the creation of many different constructed exception values that can lead to space leaks. Nullary constructors may be raised without incurring memory costs.

11.3 Raising Exceptions

An expression of the form

`raise exp`

is evaluated as follows. First *exp*, an expression of type `exn`, is evaluated to an exception value. Then the runtime stack is scanned from top to bottom in search of a handler that can handle the exception. A register points to the top-most exception handler; the exception handlers are linked together as a linked list interspersed with the other contents of the runtime stack. If a matching handler is found, the runtime stack is popped down to the handler. This popping includes popping of regions that lie between that stack top and the handler. Put differently, consider an expression of the form `let region ρ in e end`; if *e* evaluates to an exception packet, then the region bound to ρ is deallocated and the packet is also the result of evaluating the `let region` expression.

We have not attempted to design an analysis that would estimate how far down the stack a given exception value might propagate. Of course, it would not be a very good idea to allocate a constructed exception value in a region that is popped before the exception is handled! This is why we put all exception names in global regions.

11.4 Handling Exceptions

The ML expression form

`exp1 handle match`

is compiled into a `MulExp` expression of the form

```
let region  $\rho$ 
in let f = fn at  $\rho$  match
  in e1 handle f
  end
end
```

where *f* is a fresh variable. So first a handler (expressed as a function) is evaluated and stored in some region ρ . This region will always have multiplicity one and therefore be a finite region which is put on the stack. Then *e*₁, the result of compiling *exp*₁, is evaluated. If *e*₁ terminates with a value, the `let region` construct will take care of deallocating the handler. If *e*₁ terminates with an exception, however, *f* is applied.

```

exception Hd                                (* recommendation 1 *)

fun hd [] = raise Hd
  | hd (x::_) = x

exception Tl

fun tl [] = raise Tl
  | tl (_ ::xs) = xs

exception Error of string

local
  val error_f = Error "f"  (* recommendation 2 *)
in
  fun f l =
    hd(tl(tl l)) handle _ => raise error_f
end

val r = f [1,2,3,4]

```

Figure 11.1: Prudent use of exceptions with MLKit.

Thus the combined cost of raising an exception and searching for the appropriate handler takes time proportional to the depth of the runtime stack in the worst case.

Handling of exceptions is the only operation that takes time that cannot be determined statically, provided one admits arithmetic operations as constant-time operations.

11.5 Example: Prudent Use of Exceptions

Figure 11.1 shows an example of good use of exceptions with MLKit.

The application `Error "f"` has been lifted out from the body of `f`. No matter how many times `f` is applied, it will not create additional exception

values.¹

¹Program `kitdemo/exceptions.sml`.

Chapter 12

Resetting Regions

The idea of region resetting was introduced in Section 1.2.

This chapter gives an informal explanation of the rules that govern resetting. Knowing these rules is useful, irrespective of whether one makes the MLKit decide on region resetting, or prefers to control resetting explicitly in the program.

Resetting only makes sense for infinite regions. Resetting a region is a constant-time operation. Because the same region variable can be bound sometimes to a finite region and sometimes to an infinite region at runtime, resetting a region can involve a test at runtime.

MLKit contains an analysis, called the *storage mode analysis*, which has two purposes:

1. Inserting automatic resetting of infinite regions, when possible.
2. Checking applications of `resetRegions` (and `forceResetting`) so as to report on the safety of the resetting requested by the programmer.

As a matter of design, one might wonder whether it would not be sufficient to rely on the user to indicate where resetting should be done. However, checking whether resetting is safe at a particular point chosen by the user is of course no easier than checking whether resetting is safe at an arbitrary point in the program, so one might as well let the compiler insert region resetting whenever it can prove that it is safe.

In this chapter, we describe the principles that underlie the storage mode analysis. Even if one is willing to insert `resetRegions` and `forceResetting` instructions in the program, one still needs to understand these principles,

so as to be able to act upon the messages that are generated by the system in response to explicit `resetRegions` and `forceResetting` instructions.

12.1 Storage Modes

As we have seen in previous chapters, region inference decorates every allocation point with an annotation of the form `at ρ` , indicating into what region the value should be stored.

Now the basic idea is that storing a value into a region can be done in one of two ways, at runtime. One either stores the value at the *top* of the region, thereby increasing the size of the region; or one stores the value into the *bottom* of the region, by first resetting the region (so that it contains no values) and then storing the value into the region.

The storage mode analysis transforms an allocation point `at ρ` into `attop ρ` when it estimates that ρ contains live values at the allocation point, whereas it transforms it into `atbot ρ` if it can prove that the region will contain no live values at that allocation point. The tokens `attop` and `atbot` are called *storage modes*.

Region polymorphism introduces several interesting problems. Let f be a region-polymorphic function with formal region parameter ρ and consider an allocation point `at ρ` in the body of f . Whether it is safe for f to store the value at bottom in the region depends not only on the body of f but also on the context in which f is called.

For example, consider the compilation unit

```
fun f [] = []
  | f (x::xs) = x+1 :: f xs

val l1 = [1,2,3]
val l2 = if true then f l1 else l1
val x::_ = l1
```

When `f` creates the empty list, it can potentially reset the auxiliary region intended for the auxiliary pairs of the list. In the above program, however, the conditional forces `f l1` and `l2` to be in the same region as `l1`. Because `l1` is live after the application of `f`, this application must not use `atbot` as storage mode. Indeed, even if we removed the last line of the program, the application could still not use `atbot`, because `l1` is exported from the compilation unit and thus potentially used by subsequent compilation units.

By contrast, consider¹

```

fun f [] = []
  | f (x::xs) = x+1 :: f xs

val n = length(let val l1 = [1,2,3]
                in if true then f l1 else l1
                end)

```

When `f` creates the empty list, it is welcome to reset the region that holds `l1`, because by that time, `l1` is no longer needed! (`f` traverses `l1`, but when it reaches the end of the list, `l1` is no longer used.) Indeed, the MLKit will replace the list `[1,2,3]` by `[2,3,4]`. The ability to replace data in regions is crucial in many situations (as we illustrated with the game of Life in Section 1.3).

Because the MLKit allows for separate compilation, it cannot know all the call sites of a region-polymorphic function, when it is declared. Therefore, when considering an allocation point `at` ρ inside the body of some region-polymorphic function f that has ρ as a formal region parameter, one cannot know at compile time whether to use `atop` or `atbot` as storage mode. Instead, the storage mode analysis operates with a third kind of storage mode named `sat`, read: “somewhere at”. Consider an application of f for which ρ is instantiated to some region variable ρ' , say. At runtime, ρ' is bound to some region name (Section 2.1) r' . Then r' is combined with a definite storage mode (i.e., `atop` or `atbot`), to yield r , say, which is then bound to ρ . When r' was originally created (by a `let region` expression), r' was also made to contain an indication of whether it is an infinite region or a finite region.² At runtime, an allocation point `sat` ρ in the body of f will test r to see whether the region is infinite and whether the value should be stored at the top or at the bottom.³

The relevant parts of the result of compiling the last example are shown in Figure 12.1. To see the storage modes, pass the option

```
-print_drop_regions_expression_with_storage_modes
```

¹Program `kitdemo/sma1.sml`.

²On machines that have at least four bytes per word, the two least significant bits of a pointer to a word will always be 00. These two bits hold extra information in the region name. One bit, called the “atbot bit”, holds the current storage mode of the region. Another bit, called the “infinity bit”, indicates whether the region is finite or infinite.

³When ρ has multiplicity infinity, r' must be the name of an infinite region, so the runtime check on whether r has its infinity bit set is omitted.

```

let fun f [r17:INF] var1 =
  case var1 of
  nil => nil
  | :: v94 =>
    let val v95 = #0 v94; val v96 = #1 v94
    in :: (v95 + 1, f[sat r17] v96)attop r17
    end;
  val n =
    let region r31:INF
    in length
      let region r33:INF
      in f[atbot r31] [1,2,3] attop r33
      end
    end
  in {|f: (_,r1), n: _|}
end

```

Figure 12.1: Storage modes inferred by the storage mode analysis for the program `kitdemo/sma1.sml`. Use the flag `-max_inl_sz 0` to avoid inlining of the small function `length`.

to the MLKit compiler (or `-Pdresm`).

12.2 Storage Mode Analysis

For the purpose of the storage mode analysis, actual region parameters to region-polymorphic functions are considered allocation points. Passing a region as an actual argument to a region-polymorphic function involves neither resetting the region nor storing any value in it, but a storage mode has to be determined at that point nonetheless, because it has to be passed into the function together with the region. The storage mode expresses whether, at the call site, there may be any live values in the region after the call. For example, in Figure 12.1, the call to `f` at `(*1*)` passes `r31` with storage mode `atbot` because the only value that exists before the call of `f` and is needed after the call of `f` is `length`, which is declared in a different compilation unit and therefore obviously does not reside in `r31`.

Within every lambda abstraction, the MLKit performs a backwards flow analysis that determines, for every allocation point, a set of *locally live variables*, that is, a set of variables used by the remainder of the computation in the function up to the syntactic end of the function. (This includes variables that appear in function application expressions.) Prior to the computation of locally live variables, a program transformation, called *K-normalisation*, has made sure that every intermediate result that arises during computation becomes bound to a variable. (This happens by introducing extra `let` bindings, when necessary.)⁴ MLKit also computes a set of locally live variables for those allocation points that do not occur inside functions.

We now give an informal explanation of the rules that assign storage modes to allocation points. Let an allocation point

$$\text{at } \rho \tag{12.1}$$

be given. Such an allocation point can be associated with the allocation of a constant value, an immediate closure, a record value, or the allocation of the result of a primitive operation. Another possibility is that the allocation point appears as an argument to a region-polymorphic function. The procedure used for determining storage modes for function arguments depends on whether the function is declared in the same program unit as the function call. The procedure used for calls to externally declared functions is described in Section 12.3.

CASE A: The region ρ is a global region. In this case `atop` is used. There is a deficiency we have to admit here. MLKit only puts `region` bindings around expressions, not around declarations. Thus, if one writes

```
local
  fun f [] = []
    | f (x::xs) = x+1 :: f xs
  val l1 = [1,2,3]
in
  val n = length(if true then f l1 else l1)
end
```

⁴K-normalisation is transparent to users: although the storage mode analysis and all subsequent phases up to code generation operate on K-normal forms, programs are always simplified to eliminate the extra `let`-bindings before they are presented to the user.

at top level, then `l1` is put into a global region, although this decision is really unnecessary. As a consequence, `f` would be called with storage mode `attop` and thus `l1` would not be overwritten.

CASE B: The region variable ρ is not a global region and the allocation point (12.1) occurs inside a lambda abstraction, that is, inside an expression of the form `fn pat \Rightarrow e`. Here we regard every expression of the form

$$\text{let fun } f \ x = e \text{ in } e' \text{ end}$$

as an abbreviation for

$$\text{let val rec } f = \text{fn } x \Rightarrow e \text{ in } e' \text{ end}$$

Then it makes sense to talk about *the smallest enclosing lambda abstraction (of the allocation point)*.

Now there are the following cases:

- B1** ρ is bound outside the smallest enclosing lambda abstraction (and this lambda abstraction is not the right-hand side of a declaration of a region-polymorphic function that has ρ as formal parameter): Use `attop` (see Figure 12.2).
- B2** ρ is bound by a `let region` expression inside the smallest enclosing function: Use `atbot` if no locally live variable at the allocation point has ρ free in its region-annotated type scheme with place (Section 6.2), and use `attop` otherwise (see Figure 12.3).
- B3 (first attempt)** ρ is a formal parameter of a region-polymorphic function whose right-hand side is the smallest enclosing lambda abstraction: Use `sat`, if no locally live variable at the allocation point has ρ free in its region-annotated type scheme with place, and use `attop` otherwise (see Figure 12.4).

The motivation for (B1) is that if ρ is declared non-locally, then we do not attempt to find out whether ρ contains live data (this would require a more sophisticated analysis.)

The intuition behind (B2) is as follows. Region inference makes sure that the region-annotated type of a variable always contains free in it region variables for all the regions that the value bound to the variable needs when used. The lifetime of the region bound to ρ is given by the `let region`

```

let region  $\rho$ 
in ... (fn pat  $\Rightarrow$  ... at  $\rho$  ...)
end

fun f at  $\rho_1$  [ $\rho$ ] x =
  (fn y  $\Rightarrow$  ... at  $\rho$  ...)at  $\rho_2$ 

```

Figure 12.2: Two typical situations where `at ρ` is turned into `attop ρ` by rule B1.

```

(fn pat  $\Rightarrow$  ...
  let region  $\rho$ 
  in ... at  $\rho$  ...
  end)

```

Figure 12.3: The situation considered in B2. If no locally live variable l has ρ occurring in its region-annotated type scheme with place, replace `at ρ` by `atbot ρ` , otherwise by `attop ρ` .

expression, which is in the same function as the allocation point. Thus, if no locally live variable at the allocation point has ρ free in its region-annotated type scheme with place, then ρ really does not contain any live value at that allocation point.

The intuition behind (B3) is the same as behind (B2), but in this case there is a complication: ρ is only a formal parameter so it may be instantiated to different regions; in particular it may be instantiated to a region variable that does occur free in the region-annotated type scheme with place of a

```

fun f at  $\rho_0$  [ $\rho$ , ...] pat =
  ... at  $\rho$  ... l ...

```

Figure 12.4: The situation considered in B3. If no locally live variable l has in its region-annotated type scheme with place a region variable that may be aliased with ρ , replace `at ρ` by `sat ρ` , otherwise by `attop ρ` .

locally live variable at the allocation point. If that happens, rule (B3), as stated, is not sound!

We refer to the phenomenon that two different region variables in the program may denote the same region at runtime as *region aliasing*. To determine whether to use `sat` or `atop` in case (B3), the MLKit builds a *region flow graph* for the entire compilation unit. (This construction happens in a phase prior to the storage mode analysis proper.) The nodes (ranged over by n) of the region flow graph are region variables and effect variables, denoting the region variables and arrow effects that appear in the region-annotated compilation unit. Whenever ρ_1 is a formal region parameter of some function declared in the unit and ρ_2 is a corresponding actual region parameter in the same unit, a directed edge from ρ_1 to ρ_2 is created. Similarly for arrow effects: if $\epsilon_1.\varphi_1$ is a bound arrow effect of a region-polymorphic function declared in the compilation unit and $\epsilon_2.\varphi_2$ is a corresponding actual arrow effect then an edge from ϵ_1 to ϵ_2 is inserted into the graph. Also, edges from ϵ_2 to every region and effect variable occurring in φ_2 are inserted.

Let G be the graph thus constructed. For every node n in the graph, we write $\langle n \rangle$ to denote the set of region variables that can be reached from n , including n itself if n is a region variable. The rule that replaces (B3) is:

B3 ρ is a formal parameter of a region-polymorphic function whose right-hand side is the smallest enclosing lambda abstraction: Use `sat`, if, for every variable l that is locally live at the allocation point and for every region variable or effect variable n that occurs free in the region-annotated type scheme with place of l , it holds that $\langle \rho \rangle \cap \langle n \rangle = \emptyset$; Use `atop` otherwise.

CASE C: ρ is bound by a `let region` expression and the allocation point (12.1) does not occur inside any function abstraction. As in (B2), use `atbot` if no locally live variable at the allocation point has ρ free in its region-annotated type scheme with place, and use `atop` otherwise.

12.3 Calls to Externally Declared Functions

The procedure used in Section 12.2 for determining storage modes for region arguments to a region-polymorphic function assumes that the call appears in the program unit in which the function is declared. For calls to functions in other program units, an additional aliasing check is used for determining if a

region bound by a `let region` construct is passed `atop` or `atbot` or whether a function-bound region is passed `atop` or `sat`. If a region parameter is aliased with any other instantiated region parameter, the storage mode `atop` is used. Otherwise, if it can be determined (using the procedure of Section 12.2) that there are no live values in the region, the storage mode `atbot` or `sat` is used, depending on whether the region is bound by a `let region` construct or whether the region is bound by a region-polymorphic function.

This procedure is used also for determining storage modes for certain primitives like `List.concat` and `^`, for which it will be safe to store the result `atbot` if the region for the result is not aliased with any argument regions and if there are no live values in the result region at the call site.

We refer to this procedure for determining storage modes as being *modular*, as it relies on only local aliasing information for determining the passed storage mode of a region parameter.

12.4 Example: Computing the Length of Lists

We shall now illustrate the storage mode rules of Section 12.2 with some small examples, which also allow us to discuss benefits and drawbacks associated with region resetting.

Consider the functions declared in Figure 12.5;⁵ they implement five different ways of finding the length of a list! The first, `nlength`, is the most straightforward one. It is not tail recursive. Textbooks in functional programming often recommend that functions are written iteratively (i.e., using tail calls) whenever possible. This we have done with `tlength`. Next, `klength` is a version that contains a local region endomorphism `loop` to perform the iteration; `llength` is similar to `klength`, except that the region endomorphism is declared outside `llength`, using a `local` declaration. A region profile resulting from running the program is shown in Figure 12.6. The diagram shows how much space is used in regions (both finite and infinite regions) and on the stack. The `rDesc` band shows how much space is used on the stack for holding region descriptors. The `stack` band shows how much space is used on the stack, including neither finite regions nor region descriptors; the `stack` band mainly consists of registers and return addresses that have been pushed onto the stack.

⁵Program `kitdemo/length.sml`.

```

fun upto n =
  let fun loop (n,acc) = if n=0 then acc
                        else loop(n-1, n::acc)
  in loop(n, [])
  end

fun nlength [] = 0
  | nlength (_::xs) = 1 + nlength xs

fun tlength l =
  let fun tlength' (nil, acc) = acc
      | tlength' (_::xs, acc) = tlength'(xs, acc+1)
  in tlength'(l, 0)
  end

fun klength l =
  let fun loop (p as ([], acc)) = p
      | loop (_::xs, acc) = loop(xs, acc+1)
  in #2(loop(l, 0))
  end

local
  fun llength' (p as ([], acc)) = p
    | llength' (_::xs, acc) = llength'(xs, acc+1)
in
  fun llength l = #2(llength'(l, 0))
end

fun global (p as ([], acc)) = p
  | global (_::xs, acc) = global(xs, acc+1)

fun glength l = #2(global(l, 0))

val k = 5000000
val run = nlength(upto k) + tlength(upto k) + klength(upto k)
         + llength(upto k) + glength(upto k)

```

Figure 12.5: Five different ways of computing the length of lists.

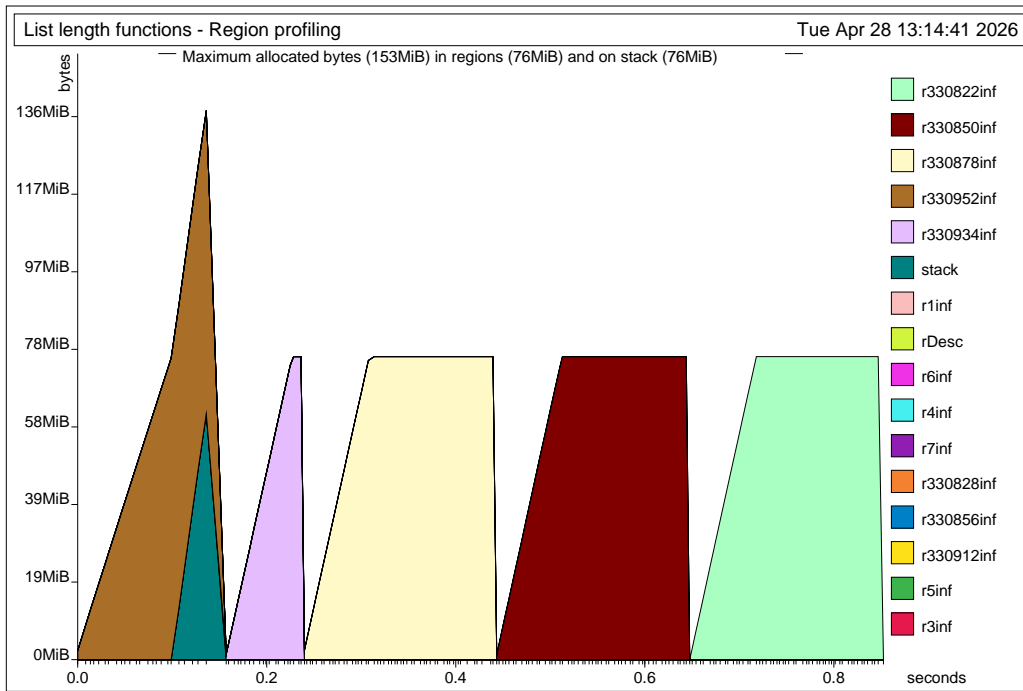


Figure 12.6: Region profiling of five different ways of computing the length of a list, namely, from left to right: `nlength`, `tlength`, `klength`, `llength`, and `glength`.

In Figure 12.6, we clearly see the five phases. In each phase, first a list is built—seen as an almost linear growth in a region; then follows a computation of the length of the list. The space behavior of the five ways of computing the length vary. We shall have more to say about the time behavior in what follows.

As one would expect, `nlength` leads to a peak in stack size; it does not use regions. The peak in stack size is caused by the stacking of a return address.

Next, we see that `tlength` is an improvement over `nlength`, the main reason being that the MLKit has figured out that the argument to `tlength` can be passed unboxed; thus no regions are used to hold the argument pair. However, if we chose to disable the unboxing of arguments that the MLKit performs,⁶ the function would become region-polymorphic and the polymorphic recursion in regions would allow the pair `(xs, acc+1)` to be stored in a region different from the argument pair to `tlength`. In this case, what appeared to be a tail call would in fact not be a tail call, for it would automatically be enclosed in a `let region` construct, introducing a fresh region for each argument pair `(xs, acc+1)`. This region would be finite, so it would be allocated on the stack. Thus, with unboxing of function arguments disabled, we would see a sharp increase in stack size for `tlength`. Although unboxing of function arguments saves us in this situation, we cannot always expect it to do so; if we were to collect boxed data in accumulating parameters to the function and this data is not to be returned by the function, there is a danger that the recursive call would not become a tail call due to the introduction of a `let region` construct being wrapped around the recursive call.⁷

The next function, `klength`, deserves careful study, because it is a prototype of a particular schema that can be used again and again when programming with regions. Iteration is done by a region endomorphism, `loop`, which is declared as a local function to the main function. The use of the same variable `p` on both the left-hand side and the right-hand side of the declaration of `loop` forces `loop` to be a region endomorphism. Because the result of `loop(xs, acc+1)` is also the result of `loop`, the result of `loop(xs, acc+1)` therefore has to be in the same region as `p`; but because `loop` is an endomorphism, `(xs, acc+1)` is forced to be in the same region as `p`. Thus, what appears to

⁶Unboxing of function arguments can be disabled by passing the option `-no_unbox_funargs` to the MLKit compiler.

⁷MLKit features an option `-preserve_tail_calls`, which ensures that no `region` binding is wrapped around a tail-call. This option is enabled by default when garbage collection is enabled.

be a tail call (`loop(xs, acc+1)`) really will be a tail call; in particular, there will be no fresh region for the argument and no growth of the stack.

Better still, we have carefully arranged that memory consumption will be constant throughout the computation of the length of the list. First, the argument to the initial call of `loop` is a pair $(1, 0)$ constructed at that point. Because `loop` is a region endomorphism, the result of `loop(1, 0)` will be in the same region as $(1, 0)$. Moreover, because we then immediately take the second projection of that pair, that region is clearly local to the body of `klength`. Call the region ρ . Because there can be an unbounded number of stores into this region, ρ is classified as infinite by multiplicity inference.

The storage mode passed along with ρ in the initial call `loop(1,0)` is `atbot`, by rule (B2) of Section 12.2. Inside `loop`, the storage mode given to the allocation of `(xs, acc+1)` is `sat`, by rule (B3) of Section 12.2: the only locally live variable at the point where the allocation takes place is `loop`, which we must not destroy before calling! The region that `loop` lies in is clearly different from ρ .

Therefore, every iteration of `loop` resets the infinite region ρ so that it will contain at most one pair. This is seen very clearly in the third hump of Figure 12.6.

Next consider `llength`. The difference from `klength` is that `llength'` is now declared outside `llength`. The storage mode analysis still achieves constant memory usage. As before, we have arranged that iteration is done by a region endomorphism that is initially applied to a freshly constructed pair. This pair can reside in a region that is local to the body of `llength` (once again, the projection `#2(llength'(1, 0))` makes sure that the pair does not escape the body of `llength`). The crucial bit is now what storage mode `llength'` uses when it stores `(xs, acc+1)`. The only locally live variable at that point is `llength'` itself and, as we noted earlier, `length'` lives in a global region, which is clearly different from the region inside `llength` that contains all the pairs. Thus, storage mode `sat` will be used, as desired.

Finally, consider `glength`, which is similar to `llength`, but with the difference that `global` is exported from the compilation unit. Thus, `global` may be called from a different compilation unit. We see, however, that `global` behaves the same as `llength` and `klength` with respect to memory usage. If `global` is called from another program unit then it is determined locally whether the passed region may hold live values at the call site, in which case, the region is passed with storage mode `attop`.

To sum up, here is how one writes a loop without using space proportional

to the number of iterations:

1. The iteration should be done by an uncurried region endomorphic function; we refer (informally) to this function as the *iterator*.
2. The iterator should be tail recursive.
3. Iteration should start from a suitably fresh initial argument; the result of the iteration should be kept clearly separate from the region where the iterator function lies.

Mutual recursion poses no additional complications. All functions in a block of mutually recursive functions are put in the same region.

Finally, the reader may be concerned that the recommended solutions, `klength`, and `llength`, and `glength`, are much slower than the other versions. This is partly an artifact of the profiling software.⁸ To get a better picture of the actual cost of the different versions, we ran the five programs separately (using lists of length 10 million instead of 5 million) on a Mac Book Pro with 16Gb RAM and a 2,7 GHz Quad-Core Intel Core i7 processor.⁹ The results are shown in Figure 12.7. Because `upto` alone takes 0.12 seconds to build the list, the differences in times are clear: the version of the length function that does not use the stack and that takes its argument in registers (i.e., `tlength`) is the fastest. The recommended versions of the length function (i.e., `klength`, `llength`, and `glength`) run as well as the versions that make use of the runtime stack (i.e., `nlength`), but are scalable and follow general useful approaches to writing recursive functions.

12.5 Resetting Regions

It is often the case that there are only a few places in the program where resetting is really essential, for example in some main loop. Therefore, the MLKit provides two operations that the programmer can use to encourage (or force) the MLKit to perform resetting at particular places in the program. The two operations are

`resetRegions` *vid*

⁸When profiling is turned on, every resetting of a region involves resetting of values in the first region page of the region.

⁹Program `kitdemo/length-cmd.sml`. For larger lists, the `nlength` program may cause stack overflow.

Program	upto	nlength	tlength	klength	llength	glength
Time (s)	0.12	0.25	0.15	0.22	0.23	0.26

Figure 12.7: User time in seconds for building a list of 10 million elements and computing its length, using five different length functions. The function `upto` builds the list, but does not compute a length. Times are average over three runs.

and

```
forceResetting vid
```

In both cases, the argument has to be a value identifier. To port programs that contain `resetRegions` and `forceResetting` to other ML systems, simply declare

```
fun resetRegions _ = ()
fun forceResetting _ = ()
```

before compiling the program developed using the MLKit.

Let ρ be a region variable that occurs free in the region-annotated type scheme with place of *vid*. Let m be the storage mode determined for ρ at a program point according to the rules of the previous section. Whether resetting of *vid* at that program point actually takes place at runtime, depends on m and on whether resetting is forced, see Figure 12.8.

12.6 Example: Improved Mergesort

We can now improve on the mergesort algorithm (Section 6.4) by taking storage modes into account. Splitting a list can be done by an iterative region endomorphism that is made local to the sorting function. Also, when the input list has been split, it is no longer needed, so the region it resides in can be reset. Similarly, when the two smaller lists have been sorted (into new regions) the regions of the smaller lists can be reset. These three simple observations lead to the variant of `msort` listed in Figure 12.9.¹⁰

Unfortunately, the storage mode analysis complains:

¹⁰MLB-file: `kitdemo/msortreset1.mlb`, file `kitdemo/msortreset1.sml`.

Does resetting take place at runtime?	<code>resetRegions</code>	<code>forceResetting</code>
$m = \text{atbot}$	yes	yes
$m = \text{sat}$	only if runtime storage mode is <code>atbot</code>	yes*
$m = \text{attop}$	no*	yes*

(*): A compile-time warning is printed in this case.

Figure 12.8: The storage modes that will be used when resetting a region depending on m , the storage mode inferred by the storage mode analysis, and depending on whether the resetting is safe (`resetRegions`) or potentially unsafe (`forceResetting`).

```

*** Warnings ***
resetRegions(xs):
  You have suggested resetting the regions that appear free
  in the type scheme with place of 'xs', i.e., in
  (int, [r147]) list
  (1)
    'r147': there is a conflict with the locally
    live variable
    v187 :(int, [r161]) list
    from which the following region variables can be reached
    in the region flow graph:
      {r161}
    Amongst these, 'r161' can also be reached from 'r147'.
    Thus I have given 'r147' storage mode "attop".

```

There is one complaint concerning the first `resetRegions`, but none concerning the two remaining ones. By inspecting the region-annotated term, one sees that `r147` is a formal parameter of `msort`. Due to the recursive call `msort 1`, the region graph contains an edge from `r147` to `r161`. Thus the analysis decides on `attop`, using rule (B3). This choice shows a weakness in the analysis, for using `sat` would really be sound. (The problem is that, unlike polymorphic recursion, the region flow graph does not distinguish be-

```

local
  fun cp [] = []
    | cp (x::xs) = x :: cp xs

  (* exomorphic merge *)
  fun merge (xs, []) : int list = cp xs
    | merge ([], ys) = cp ys
    | merge (l1 as x::xs, l2 as y::ys) =
      if x<y then x :: merge(xs, l2)
      else y :: merge(l1, ys)

  (* splitting a list *)
  fun split (x::y::zs, l, r) = split(zs, x::l, y::r)
    | split (x::xs, l, r) = (xs, x::l, r)
    | split (p as [], l, r) = p

  (* exomorphic merge sort *)
  fun msort [] = []
    | msort [x] = [x]
    | msort xs = let val (_, l, r) = split(xs, [], [])
                  in resetRegions xs;
                    merge(msort l before resetRegions l,
                          msort r before resetRegions r)
                  end

in
  val runmsort = msort(upto 200000)
  val result = print "Really done\n"
end

```

Figure 12.9: Variant of `msort` that uses `resetRegions` to improve memory usage. MLKit fails to infer that the region holding the argument list `xs` can be reset after `xs` is split.

tween different calls of the same function.) Seeing that this is the problem, we decide to put `forceResetting` to work, see Figure 12.10.¹¹ The region profile of the improved merge sort appears in Figure 12.11. As expected, we have now brought space consumption down from four times to two times the size of the input. Figure 12.11 may be compared to Figure 6.4 on page 68.

12.7 Example: Scanning Text Files

In this section we present a program that can scan a sequence of Standard ML source files and compute what percentage of the source files is made up by comments. Recall that an ML comment begins with the two characters `(*`, ends with `*)`, and that comments may be nested but must be balanced (within each file, we require).

The obvious solution to this problem is to implement an automaton with counters to keep track of the level of nesting of parentheses, number of characters read, and number of characters within comments. This solution provides an interesting test for region inference: although designed with the lambda calculus in mind, does the scheme cope with good old-fashioned state computations?

Let us be ambitious and write a program that only ever holds on to one character at a time when it scans a file. In other words, the aim is to use constant space (i.e., space consumption should be independent of the length of the input file).

To this end, let us arrange to use a region with infinite multiplicity to hold the current input character and then reset that region before we proceed to the next character. The iteration is done by tail recursion, using region endomorphisms to ensure constant space usage.

The bulk of the program appears below.¹² The scanning of a single file is done by `scan`, which contains three mutually recursive region endomorphisms (`count`, `after_lpar`, and `after_star`) written in accordance with the guidelines in Section 12.4.

It turns out that the constraints we have put on ourselves with running in constant space does work well with the combination of how `TextIO` input streams buffers input data. Instead, we shall make use of the lower-level `Posix` operations for opening and reading files.

¹¹MLB-file: `kitdemo/msortreset2.mlb`, file `kitdemo/msortreset2.sml`.

¹²MLB-file: `kitdemo/scan.mlb`, file: `kitdemo/scan.sml`.

```

local
  fun cp [] = []
    | cp (x::xs) = x :: cp xs

  (* exomorphic merge *)
  fun merge (xs, []) : int list = cp xs
    | merge ([], ys) = cp ys
    | merge (l1 as x::xs, l2 as y::ys) =
      if x<y then x :: merge(xs, l2)
      else y :: merge(l1, ys)

  (* splitting a list *)
  fun split (x::y::zs, l, r) = split(zs, x::l, y::r)
    | split (x::xs, l, r) = (xs, x::l, r)
    | split (p as [], l, r) = p

  (* exomorphic merge sort *)
  fun msort [] = []
    | msort [x] = [x]
    | msort xs = let val (_, l, r) = split(xs, [], [])
                  in forceResetting xs;
                    merge(msort l before resetRegions l,
                          msort r before resetRegions r)
                  end

in
  val runmsort = msort(upto 200000)
  val result = print "Really done\n"
end

```

Figure 12.10: Using `forceResetting` to reset regions.

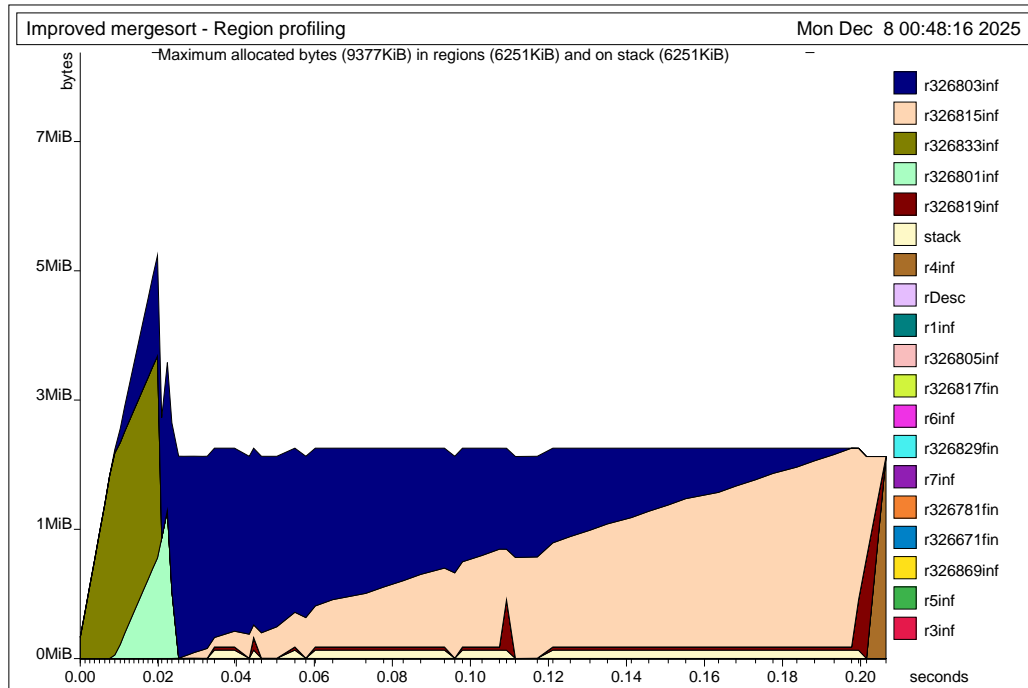


Figure 12.11: Region profiling of the improved mergesort. The upper triangle contains unsorted elements, while the lower triangle contains sorted elements. The program was compiled with profiling enabled and then run with the command `run -microsec 1000`. The Postscript picture `region.ps` was generated with the command `rp2ps -region -sampleMax 200` and then converted into `region.pdf` with `ps2pdf`.

The built-in function `Posix.IO.readVec`, which reads a vector of bytes from an open file descriptor, understands storage modes; if called with storage mode `atbot`, it will reset the region where the byte vector should be put before reading the bytes from the file descriptor. Consequently, at every call of `next`, the “input buffer region” will be reset.

The other important loop in the program is `driver`, a function that repeatedly reads a file name from a given file descriptor, opens the file with that name, and calls `scan` to process the file. Once again, we want to keep at most one file name in memory at a time, so we would like the region containing the file name to be reset upon each iteration. As it turns out, `readWord` will always try to store the string it creates at the bottom of the region in question.

```

local
  structure F = Posix.FileSys
  exception NotBalanced
  fun scan fd : int * int =
    let
      fun next () = Byte.bytesToString(Posix.IO.readVec(fd,1))
      fun up (lev,ins) = if lev > 0 then ins + 1
                        else ins

      (* n   : characters read from 'fd'
         ins : characters inside comments
         lev : current (level) number of unmatched (*
         s   : next input character or empty *) *)

      fun count (p as (n,ins,lev,s)) =
        case s of
          "" => (* end of stream: *) p
        | "(" => after_lpar(n+1,ins,lev,next())
        | "*" => after_star(n+1,up(lev,ins),lev,next())
        | _  => count(n+1,up(lev,ins), lev,next())
      and after_lpar (p as (n,ins,lev,s)) =
        case s of
          "" => p
        | "*" => count(n+1,ins+2, lev+1,next())
        | "(" => after_lpar(n+1,up(lev,ins),lev,next())
        | _  => count(n+1,up(lev,up(lev,ins)),lev,next())

```

```

and after_star (p as (n,ins,lev,s)) =
  case s of
    "" => p
  | ")" => if lev > 0 then
            count(n+1,ins+1,lev-1,next())
          else raise NotBalanced
  | "*" => after_star(n+1,up(lev,ins),lev,next())
  | "(" => after_lpar(n+1,ins,lev,next())
  | _ => count(n+1,up(lev,ins),lev,next())

  val (n,ins,lev,_) = count(0,0,0,next())
  in if lev=0 then (n,ins) else raise NotBalanced
end

fun report_file (filename, n, ins) =
  writeln (filename ^ ": size = " ^ Int.toString n
          ^ " comments: " ^ Int.toString ins ^ " ("
          ^ (Int.toString(percent(ins, n))
            handle _ => "") ^ "%)");

(* scan_file(filename) scans through the file named filename
   returning either SOME(size_in_bytes, size_of_comments)
   or, in case of an error, NONE. In either case a line of
   information is printed. *)

fun scan_file filename : (int*int) option =
  let val fd = F.openf (filename, F.O_RDONLY, F.O.flags[])
  in let val (n, ins) = scan fd
      in Posix.IO.close fd;
        report_file (filename, n, ins);
        SOME (n, ins)
      end handle NotBalanced =>
        (writeln (filename ^ ": not balanced");
         Posix.IO.close fd;
         NONE)
      end handle IO.IOException {name,...} =>
        (writeln (name ^ " failed."); NONE)
  end

fun report_totals (n, ins) =
  writeln ("\nTotal sizes: " ^ Int.toString n

```

```

    ^ " comments: " ^ Int.toString ins
    ^ " (" ^ (Int.toString (percent (ins, n))
      handle _ => "") ^ "%)"

(* main(fd) reads a sequence of filenames from fd, one file
   name pr line (leading spaces are skipped; no spaces
   allowed in file names). Each file is scanned using
   scan_file after which a summary report is printed *)

fun main fd : unit =
  let fun driver (p as (NONE, n, ins)) =
        (report_totals(n, ins); p)
      | driver (p as (SOME filename, n, ins)) =
        driver (case scan_file filename of
                  SOME(n', ins') =>
                    ( resetRegions p
                      ; (readWord fd, n+n', ins+ins')
                    )
                  | NONE => ( resetRegions p
                              ; (readWord fd, n, ins)
                            )
                )
      in driver (readWord fd, 0, 0); ()
    end

in
  val result = main F.stdin
end

```

The program was compiled both with and without profiling turned on. The output from running the program on 14 of the source files for the MLKit is shown here:

```

../src/Parsing/Infixing.sml: size = 32153 comments: 5294 (16%)
../src/Parsing/Parse.sml: size = 5498 comments: 673 (12%)
../src/Parsing/Topdec.grm.sml: size = 212870 comments: 4698 (2%)
../src/Parsing/GRAMMAR_UTILS.sml: size = 5534 comments: 787 (14%)
../src/Parsing/INFIX_STACK.sml: size = 487 comments: 321 (65%)
../src/Parsing/Topdec.lex.sml: size = 49968 comments: 1023 (2%)
../src/Parsing/LEX_BASICS.sml: size = 2048 comments: 1242 (60%)
../src/Parsing/LEX_UTILS.sml: size = 1444 comments: 291 (20%)
../src/Parsing/GrammarUtils.sml: size = 17262 comments: 1820 (10%)

```

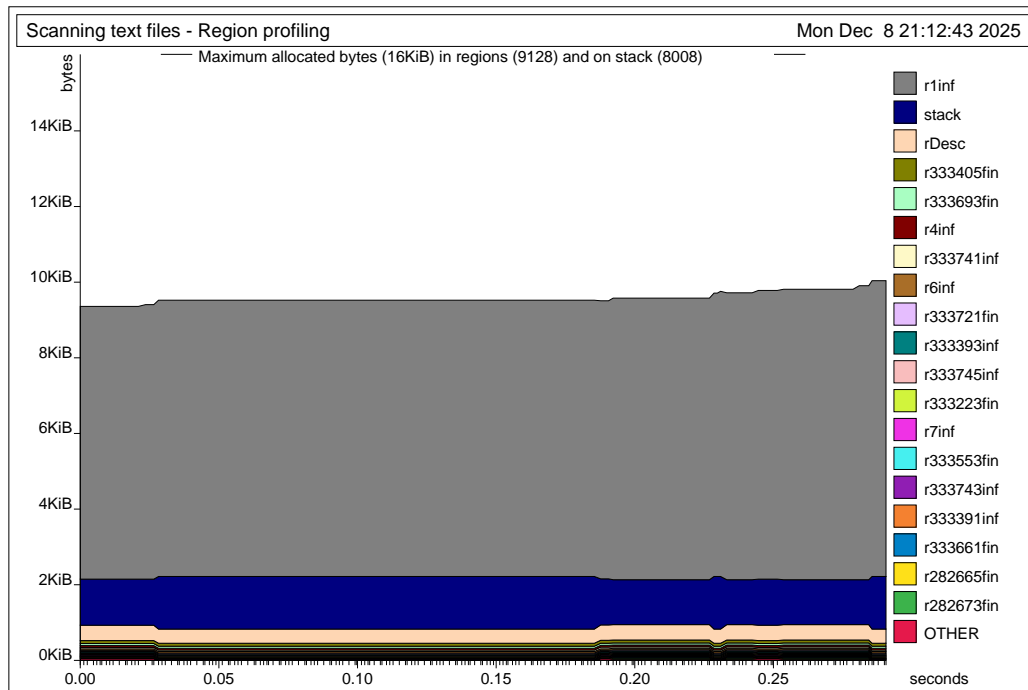


Figure 12.12: Region profile of the comment scanner. The occasional increases in memory use is due to the functions that read a file name from a file descriptor. The program was compiled with profiling enabled, then run with the command `run -notimer 1000 < ../kitdemo/scanfiles`. A Postscript file `region.ps` can be generated with the command `rp2ps -region -sampleMax 200` and converted to `region.pdf` using `ps2pdf region.ps region.pdf`.

```

../src/Parsing/LexBasics.sml: size = 12636 comments: 3079 (24%)
../src/Parsing/MyBase.sml: size = 33725 comments: 11104 (32%)
../src/Parsing/HOOKS.sml: size = 312 comments: 170 (54%)
../src/Parsing/InfixStack.sml: size = 7404 comments: 2972 (40%)
../src/Parsing/LexUtils.sml: size = 8708 comments: 487 (5%)

```

```
Total sizes: 390049 comments: 33961 (8%)
```

A region profile for that run is shown in Figure 12.12. The almost-constant space usage is evident. The occasional disturbances are due to the non-iterative functions that read a file name from input by first reading one line and then extracting the name.

Chapter 13

Higher-Order Functions

13.1 Lambda Abstractions

A *lambda abstraction* in Standard ML is an expression of the form

$$\mathbf{fn} \textit{idpat} \Rightarrow \textit{idexp}$$

where *idpat* is a pattern and *idexp* an expression. Lambda abstractions denote functions. We refer to the *idexp* as the *body* of the function; variable occurrences in *idpat* are binding occurrences; informally, the variables that occur in *idpat* are said to be *lambda-bound* with scope *idexp*.

Lambda abstractions are represented by closures, both in the language definition and in the MLKit. In the MLKit, a closure for a lambda abstraction consists of a code pointer plus one word for each free variable of the lambda abstraction. Closures are not tagged except when garbage collection is enabled, in which case a closure contains one or more words to hold the tag.

At this stage, it will hardly come as a surprise to the reader that closures are stored in regions. Sometimes they reside in finite regions on the stack, other times they live in infinite regions, just like all other boxed values.

Every occurrence of **fn** in the program is considered an allocation point; the region-annotated version of the lambda abstraction is

$$\mathbf{fn} \textit{at } \rho \textit{idpat} \Rightarrow \textit{idexp}$$

Standard ML allows functions to be declared using **val** rather than **fun**, for example,

```
val h = g o f
```

declares the value identifier `h` to be the composition of `g` and `f`. Whereas functions declared with `fun` automatically become region-polymorphic, functions declared with `val` do not in general become region-polymorphic.¹ However, in the special case where the right-hand side of the value declaration is a lambda abstraction, the MLKit automatically converts the declaration into a `fun` declaration, thereby making the function region-polymorphic after all.

ML allows declarations of the form

```
fun f atpat1 atpat2 ... atpatn = exp
```

as a shorthand for

```
fun f atpat1 = fn atpat2 => ... fn atpatn => exp
```

where `atpat` ranges over atomic patterns. Functions declared using this abbreviation are said to be *Curried*.

13.2 Region-Annotated Function Types

The general form of a region-annotated function type is

$$([\mu_1, \dots, \mu_n] \xrightarrow{\epsilon.\varphi} \mu', \rho)$$

where μ_1, \dots, μ_n are the type with places of the arguments, μ' is the type with place of the result, and ρ is the region containing the closure for the function. When a function type has only one argument type, we shall often write it on the form $(\mu \xrightarrow{\epsilon.\varphi} \mu', \rho)$, and so shall the MLKit.

As mentioned in Section 5.3, the unusual looking object $\epsilon.\varphi$ is called an *arrow effect*. Its first component is an effect variable, whose purpose will be explained shortly. The second component is called the *latent effect*, and describes the effect of evaluating the body of the function.

The following example illustrates why latent effects are crucial for knowing the lifetimes of closures.² Consider

¹The reason for this is that the expression on the right-hand side of the value declaration might have an effect (e.g, print something) before returning the function. It would not be correct to suspend this effect by introducing formal region parameters.

²Program `kitdemo/lambda.sml`.

```

let val n =
  let region r9:INF, r19:INF, r25:1;
    val f =
      let val xs = [1,2] at r9
        in fn at r25 ys => length xs + length ys
      end
    in f [7] at r19
  end
in {n: _}
end

```

Figure 13.1: Region-annotated program illustrating that the lifetime of a closure is at least as long as the lifetime of the values that evaluation of the function body will require.

```

val n = let val f = let val xs = [1,2]
              in fn ys => length xs + length ys
            end
        in f [7]
      end

```

Notice that `xs` has to be kept alive for as long as the function (`fn ys => ...`) access `xs`, when called. The region-annotated version of the example appears in Figure 13.1.³

We see that `xs` is put in `r9`, that the function closure for (`fn ys => ...`) is put in `r25` and indeed, `r9` and `r25` have the same lifetime. To understand how the region inference system figured that out, let us consider the effect and the region-annotated types of particular sub-expressions. Looking at the lambda abstraction, it must have a functional type of the form $(\tau \xrightarrow{\epsilon, \varphi} \tau', \mathbf{r25})$ where φ is the effect

$$\{\mathbf{get}(\mathbf{r1}), \mathbf{get}(\mathbf{r9}), \mathbf{get}(\mathbf{r19})\}$$

Notice that `r9` occurs free in the type of the lambda abstraction. But, as pointed out in Section 3.4, the criterion for putting a `region` binding of ρ around an expression e is that ρ occurs free neither in the type with place of e nor in the type scheme with place of any variable in the domain of the

³To see the output programs discussed in this section, enable the flag `-print_drop_regions_expression`.

type environment. The smallest sub-expression of the program for which `r9` does not occur free in the type with place of the expression is the right-hand side of the `val` binding of `n`, for that expression simply has type with place `int`. And at that point, the only region variables that occur free in the type environment are global region variables. Hence the placement of the `region` binding of `r9`.

13.3 Arrow Effects

In a first-order language, effect variables might not be particularly important. But in a higher-order language like ML, effect variables are useful for tracking dependencies between functions. The following example illustrates the point:⁴

```

fun apply f x = f x
val y = apply (fn n => n + 1.0) 5.0
val z = apply (fn m => m) 6

```

Here is the region-annotated type scheme of `apply`:

$$\forall \alpha_0 \alpha_2 \rho_{17} \rho_{15} \epsilon_{16} \epsilon_{20} \epsilon_{18}. (\alpha_0 \xrightarrow{\epsilon_{16}. \emptyset} \alpha_2, \rho_{15}) \xrightarrow{\epsilon_{20}. \{\text{put}(\rho_{17})\}} (\alpha_0 \xrightarrow{\epsilon_{18}. \{\text{get}(\rho_{15}), \epsilon_{16}\}} \alpha_2, \rho_{17})$$

The latent effect associated with ϵ_{20} shows that when `apply` is applied to a function, it may create (in fact: will create) a function closure in ρ_{17} . The latent effect associated with ϵ_{16} is empty, because the declaration of `apply` does not tell us anything about what effect its formal parameter `f` must have. Crucially, however, ϵ_{16} is included as an atomic effect in the latent effect associated with ϵ_{18} ; whenever the body of `apply f` is evaluated, the body of `f` may be (in fact: will be) evaluated.

The polymorphism in effects makes it possible to distinguish between the latent effects of different actual arguments to `apply`. For example, the functions `(fn n => n + 1.0)` and `(fn m => m)` have different latent effects. Let us take the function `(fn n => n + 1.0)` as an example. It has region-annotated type with place

$$((\text{real}, \rho_{21}) \xrightarrow{\epsilon_{26}. \{\text{get}(\rho_{21}), \text{put}(\rho_1)\}} (\text{real}, \rho_1), \rho_{25}) \quad (13.1)$$

⁴Program `kitdemo/apply.sml`.

Here, the effect variable ϵ_{26} and the region variables ρ_{25} and ρ_1 were chosen arbitrarily. (Actually, the region variable ρ_1 denotes the global region for reals.) The region inference algorithm discovers that (13.1) can be derived from the argument type

$$(\alpha_0 \xrightarrow{\epsilon_{16}.\emptyset} \alpha_2, \rho_{15})$$

of the type scheme for `apply` by the instantiating substitution

$$\begin{aligned} S = (& \{\alpha_0 \mapsto (\mathbf{real}, \rho_{21}), \alpha_2 \mapsto (\mathbf{real}, \rho_1)\}, \\ & \{\rho_{17} \mapsto \rho_{27}, \rho_{15} \mapsto \rho_{25}\}, \\ & \{\epsilon_{16} \mapsto \epsilon_{26}.\{\mathbf{get}(\rho_{21}), \mathbf{put}(\rho_1)\}\}, \\ & \epsilon_{20} \mapsto \epsilon_{24}.\{\mathbf{put}(\rho_{27})\}, \\ & \epsilon_{18} \mapsto \epsilon_{22}.\{\epsilon_{26}, \mathbf{put}(\rho_1), \mathbf{get}(\rho_{25})\} \\ & \}) \end{aligned}$$

Formally, a *substitution* is a triple (S^t, S^r, S^e) , where S^t is a finite map from type variables to region-annotated types, S^r is a finite map from region variables to region variables, and S^e is a finite map from effect variables to arrow effects. Let us explain why substitutions map effect variables to arrow effects. One alternative, one might consider, is to let substitutions map effect variables to effect variables. But then substitutions would not be able to account for the idea that effects can grow, when instantiated. In the `apply` example, for instance, the empty effect associated with ϵ_{16} has to grow to $\{\mathbf{get}(\rho_{21}), \mathbf{put}(\rho_1)\}$ at the concrete application of `apply`. Otherwise, as it is easy to demonstrate, the region inference system would become unsound.

Another alternative would be to let substitutions map effect variables to effects. But nor that would work well together with the idea of using substitutions to express growth of effects. For example, when applying the map $\{\epsilon \mapsto \{\mathbf{get}(\rho_0), \mathbf{put}(\rho_2)\}\}$ to the effect $\{\mathbf{get}(\rho_9), \epsilon\}$, say, we would presumably yield the effect $\{\mathbf{get}(\rho_9), \mathbf{get}(\rho_0), \mathbf{put}(\rho_2)\}$ in which the fact that the original effect had to be at least as large as whatever ϵ stands for, is lost. Instead, we define substitution so that applying the effect substitution $\{\epsilon \mapsto \epsilon.\{\mathbf{get}(\rho_2), \mathbf{put}(\rho)\}\}$ to $\{\mathbf{get}(\rho_9), \epsilon\}$ yields $\{\mathbf{get}(\rho_9), \epsilon, \mathbf{get}(\rho_2), \mathbf{put}(\rho)\}$.

We can now give a complete definition of atomic effects. An *atomic effect* is either an effect variable or a term of the form $\mathbf{get}(\rho)$ or $\mathbf{put}(\rho)$, where ρ as usual ranges over region variables. An *effect* is a finite set of atomic effects.

One can get the MLKit to print region-annotated type schemes with places of all binding occurrences of value variables. Also, one can choose to have arrow effects included in the printout by passing the options `print_types`

and `print_effects` to the MLKit compiler. Although passing these options gives very verbose output, it is instructive to look at such a term at least once, to see how arrow effects are instantiated. We show the full output for the `apply` example in Figure 13.2.

In reading the output, it is useful to know that the MLKit represents effects and arrow effects as graphs, the nodes of which are region variables, effect variables, `put`, `get`, or `U` (for “union”; `U` by itself means the empty set). Region variables are leaf nodes. A `put` or `get` node has emanating from it precisely one edge; it leads to the region variable in question. An effect variable node (written `e` followed by a sequence number) is always the handle of an arrow effect; there are edges from the effect variable to the atomic effects of that arrow effect, either directly, or via union nodes or other effect variable nodes. For instance, `e18(U(U,get(r15),e16))` in Figure 13.2 denotes an effect variable with an edge to a union node that has edges to an empty union node, a `get` node, and an effect variable node.

When a term containing arrow effects is printed, shared nodes that have already been printed are marked with a `@`; their children are not printed again. In the figure, the binding occurrence of `apply` has been printed with its region-annotated type scheme. Each non-binding occurrence of `apply` has been printed with four square-bracketed lists. The first list is the actual region arguments; the following three are instantiation lists that show the range of the substitution by which the bound variables of the type scheme was instantiated, in the same order as the bound variables occurred. For example, in the second use of `apply`, `r17` was instantiated to `r45`.

13.4 On the Lack of Region Polymorphism

Unlike identifiers bound by `fun`, lambda-bound function identifiers are never region-polymorphic. So in an expression of the form

$$\mathbf{fn} \mathbf{f} \Rightarrow \dots \mathbf{f} \dots \mathbf{f} \dots$$

all the uses of `f` use the same regions. Indeed, because `f` occurs free in the type environment while region inference analyses the body of the lambda abstraction, none of the regions that appear in the type of `f` will be deallocated inside the body of the lambda abstraction. Also, such a region must be bound outside the lambda abstraction, so any attempt to reset such a region

```

fun apply :all r17,r15,e16,e20,e18,a0,a2.
  (α0-e16→α2,r15) -e20(put(r17))→
  (α0-e18(U(U,get(r15),e16))→α2,r17)
  at r1 [r15:0|r17:1] f = fn at r17 x:α0 ⇒ f x;
val y:(real,r1) =
  let region r21:1, r25:1, r27:1, e26, e22
  in let region e24
    in apply [r27]
      [(real,r21),(real,r1)]
      [r27,r25]
      [e26(get(r21),put(r1)),
       e24(put(r27)),
       e22(e26(get(r21),put(r1)),get(r25))]
    ]
    (fn at r25 n:(real,r21) ⇒
      R64.fromF64 [at r1] (F64.fromR64 n:f64 + 1.0f64)
      :(real,r1)
    )
  end
  5.0
end;
val z:int =
  let region r43:1, r45:1, e46, e42
  in let region e44
    in apply [r45]
      [int,int]
      [r45,r43]
      [e46,e44(put(r45)),e42(e46,get(r43))]
    (fn at r43 m:int ⇒ m)
  end
  6
end

```

Figure 13.2: The instantiation of arrow effects keeps different applications of the same function (here `apply`) apart. The output was obtained by compiling the program `kitdemo/apply.sml` with options `-print_types`, `-print_effects`, `-no_uncurry` and `-max_inl_sz 0`.

inside the body of the abstraction will cause the storage mode analysis to complain (by Rule (B1) of Section 12.2).

Therefore, when a function f is passed as argument to another function g , as in the expression $g(f)$, first regions are allocated for the use of f , then g is called, and finally, the regions are deallocated (provided they are not global regions). Whether the `let region` construct thus introduced encloses the call site immediately, as in

```
let region  $\rho_1, \dots, \rho_n$  in  $g(f)$  end
```

or further out, as in

```
let region  $\rho_1, \dots, \rho_n$  in ...  $g(f)$  ... end
```

depends on the type and effect of the expression $g(f)$ in the usual way: regions can be deallocated when they occur free neither in the type with place of the expression nor in the type environment.

13.5 Examples: map and foldl

Consider the program⁵

```
fun map f [] = []
  | map f (x::xs) = f x :: map f xs

val x = map (fn x => x+1) [7,11]
```

This formulation of `map` is not the most efficient one in the MLKit if it was not because the MLKit optimiser would uncurry the function and furthermore specialise the function for each application point.⁶ However it serves to illustrate the point made in the previous section about allocating regions in connection with higher-order functions. The region-annotated version is listed in Figure 13.3. We see that the region that appears free in the type with place of the successor function (i.e., `r43`) is allocated prior to the call of `map` and that it stays alive throughout the evaluation of the body of `map`. Notice, however, that the closures that are created when `map` is applied do not pile up in `r43`, the region of the successor function. Instead, they are

⁵Program `kitdemo/map.sml`.

⁶The output we present in this section was obtained by passing to the MLKit compiler the options `-maximum_specialise_size 0` and `-no_uncurry`.

```

let fun map at r1 [r39:1, r23:0] var1 =
  fn at r39 var2 ⇒
  case var2 of
    nil ⇒ nil
  | :: v102 ⇒
    let val v103 = #0 v102; val v104 = #1 v102
    in ::
      (var1 v103,
       let region r35:1 in map[r35,r23] var1 v104 end
      )at r23
    end;
  val x =
    let region r43:1, r45:INF, r49:1
    in map[r49,r4] (fn at r43 x ⇒ x + 1) [7,11] at r45
    end
in {|map: (_,r1), x: _|}
end

```

Figure 13.3: Although this version of `map` creates a closure for each list element, the region-polymorphic recursion (of `map`) ensures that that closure is put in a region local to `map`. Thus, these closures do not pile up in `r39`, the region of the initial argument.

put in local regions bound to `r35`, one closure in each region. Also, if we had given some more complicated argument to `map`, the body of that function could include local `region` declarations. For each list element, regions would then be allocated, used, and then deallocated before proceeding to the next list element.

So it might appear that higher-order functions are nothing to worry about when programming with regions. That is not so, however. The limitation that lambda-bound functions are never region-polymorphic can lead to space leaks. Here is an example:

```
fun foldl f acc [] = acc
  | foldl f acc (x::xs) = foldl f (f(x,acc)) xs

val x = foldl (fn (x,acc) => 10*acc+x) 0 [7,2]
```

Because `f` is lambda-bound, all the pairs created by the expression `(x,acc)` will pile up in the same region. The storage mode analysis will infer storage mode `attop` for the allocation of the pair, by rule (B1) of Section 12.2; because `foldl` is curried, there are several lambdas between the formal region parameter of `foldl` that indicates where the pair should be put and the allocation point of the pair.

It does not help to uncurry `foldl` and turn `foldl` into a region endomorphism:

```
fun foldl(p as (f, [], _)) = p
  | foldl(f,x::xs,acc) = foldl(f,xs,f(x,acc))

val x = #3(foldl(fn(x,acc) => 10*acc+x, [7,2], 0))
```

The storage mode analysis will still give `attop` for the allocation of the pair `(x,acc)`, because the region of the pair is free in the region-annotated type of `f`, which is locally live at that point.

What if we require that `f` be curried, so as to avoid the creation of the pair altogether?⁷

```
fun foldl f b xs =
  let fun loop(p as ([], b)) = p
        | loop(x::xs, b) = loop(xs, f x b)
      in #2(loop(xs,b))
  end
```

⁷Program `kitdemo/fold2.sml`.

The region-annotated version of this program appears in Figure 14.4 on page 144. This saves the allocation of a pair inside loop, although the saving is lost if the evaluation of $f\ x$ creates a closure.

In short, folding a function over a list may leak two words of memory for each list element.

Chapter 14

The Function Call

Standard ML allows function applications of the form

$$exp_1 exp_2$$

where exp_1 is the operator and exp_2 is the operand. The syntax for function application is overloaded, in that it is used for three different purposes in ML:

1. Applications of built-in operations such as `+`, `=`, and `:=`.
2. Applications of unary value constructors (including `ref`) and unary exception constructors.
3. Applications of user-defined functions, that is, functions introduced by `fn` or `fun`.

This chapter is about the last kind of function applications; in the following, we use the term function application to stand for applications of user-defined functions only.

Function applications are ubiquitous in Standard ML programs; in particular, iteration is often achieved by function calls. Not surprisingly, careful compilation of function calls is essential for obtaining good performance.

MLKit partitions function calls into four kinds, which are implemented in different ways. At best, a function call is simply realised by a jump in the target code. The resource conscious programmer will want to know the special cases; for example, when doing an iterative computation, it is

important to know whether the space usage is going to be independent of the number of iterations.

MLKit performs a backwards flow analysis, called *call conversion*, to determine what function calls are tail calls and, more generally, what function calls fall into the four special cases. We say that expressions produced by this analysis are *call-explicit*. One can inspect call-explicit programs by passing the option

```
-print_call_explicit_expression
```

to the MLKit compiler, and thus check whether specific function calls in the code turn out as intended. Call-explicit expressions are produced after regions have been dropped (page 64) but before native code generation.

We shall first give a brief description of the parameter passing mechanism in general and then discuss the different kinds of function calls provided, working our way from the most specialised (and most efficient) cases towards the default cases.

14.1 Parameter Passing

Parameters to functions are passed either on the runtime stack or, if possible, in registers. Also region parameters to region-polymorphic functions are passed on the runtime stack or in registers.

14.2 Tail Calls and Non-Tail Calls

A call that is the last action of a function is referred to as a *tail call*. After region inference, the MLKit performs a tail call analysis (in one backwards scan through the program). It is significant that the tail call analysis happens after region inference; as we saw in Section 12.4, a function call that looks like a tail call in the source program may end up as a non-tail call in the region-annotated program, because the function has to return to free memory. The tail call analysis divides function calls into four different kinds of calls:

- jmp:** Tail calls to known functions.
- funcall:** Non-tail calls to known functions.
- fnjmp:** Tail calls to unknown functions.
- fnccall:** Non-tail calls to unknown functions.

In the sections to follow, we describe each of these kinds of calls in detail.

14.3 Tail Call to Known Function (jmp)

A call to a region-polymorphic function (i.e., a known function) takes the form

$$f \ [\rho_1, \dots, \rho_n] \ (e_1, \dots, e_m)$$

where ρ_1, \dots, ρ_n are actual region parameters to the function, f is the name of a region-polymorphic function, and $e_1 \cdots e_m$, $m \geq 1$ are value arguments to the function (we often omit the brackets (\cdots) when $m = 1$.) MLKit turns such a function call into the form

$$\text{jmp } f \ [\rho_1, \dots, \rho_n] \ (e_1, \dots, e_m)$$

if the call appears in a tail-call position, that is, if the call is the last thing the current function needs to do. Because the start address of f is known during compilation (because f is region-polymorphic), such a call is as efficient as an assembly language jump to a constant label (not taking into account the shuffling of arguments needed to match the calling convention for f).

The way to avoid that a **region** binding is wrapped around the function call (and thus causes the call not to be recognised as a tail call) is to turn the calling function into a region endomorphism, when possible.

The following is an example of how one obtains a tail call to a known function:¹

```

local
  fun f'(p as (0,b)) = p
    | f'(n,b) = f'(n-1,n*b)
in
  fun f(a,b) = #2(f'(a,b))
end

```

The call-explicit version of f' appears in Figure 14.1.

There is a more efficient version of the function f that exploits the MLKit's unboxing of function arguments, but in general, one can rely on unboxing to ensure tail-calls only when the elements of the argument tuple themselves are unboxed; otherwise there is a risk that, for each invocation,

¹Program `kitdemo/tail.sml`.

```

fun f' [r15:inf] var2 =
  let val v96 = #0 var2
  in case v96 of
    0 => var2
  | _ => let val v98 = #1 var2
        in jmp f' [sat r15] (v96 - 1, v96 * v98) sat r15
        end
      end
end

```

Figure 14.1: An example where a function call turns into a tail call to a known function.

fresh regions are introduced to hold the arguments to the call, and the call would need to return to deallocate these regions.

MLKit can transform a call into a `jmp` tail call even in the case that the call appears in the body of a `fn` expression. Consider the following two mutually recursive functions `g` and `h`:²

```

fun g (n,b) = h (n-1) b
and h 0 b = b
  | h n b = g(n,n*b)

```

Here `h` calls `g` in a tail position. The call explicit version of the program is listed in Figure 14.2, and indeed, the call to `g` is recognised as a tail call. Also notice that the MLKit does not try to in-line `g` in `h` (or vice-versa), although such an optimisation would certainly improve on the efficiency of the generated code. Another example of a `jmp` tail call is shown in Section 14.8.

The call-explicit program shown in Figure 14.2 was compiled with the flag `-no_abbrev`. Without this flag, the output is abbreviated in several ways. First, `fn`call and `fun`call annotations are not shown and neither are the `at`top `r1` annotations on top-level functions.

14.4 Non-Tail Call to Known Function (`fun`call)

In the case that a call to a known function cannot be turned into a tail call, because the call needs to return to do more work, the call is transformed into

²Program `kitdemo/tail2.sml` compiled with the flags `-no_uncurry` and `-no_abbrev`.

```

let fun g atop r1 (v134, v135) =
  let region r13:3
  in fncall funcall h[atbot r13] (v134 - 1) v135
  end
and h atop r1 [r17:3] var1 =
  fn atop r17 var2 =>
    case var1 of 0 => var2 | _ => jmp g[] (var1, var1 * var2)
in {h: (_,r1), g: (_,r1)|}
end

```

Figure 14.2: A function call can turn into a tail call even in the case that the call appears in the body of a `fn` expression.

$$\text{funccall } f \ [\rho_1, \dots, \rho_n] \ exp$$

where `funccall` is the mnemonic used for non-tail calls to region-polymorphic functions. One example is the call to `h` in Figure 14.2. Here the call to `h` takes a region argument `r13` and an ordinary argument `(v134 - 1)`; the call to `h` returns a closure, which needs to be applied to `v135` before the function `g` can deallocate the region `r13` and return.

This case completes all possible cases of applications of region-polymorphic functions. We now turn to function applications where the operator is not the name of a region-polymorphic function.

14.5 Tail Call to Unknown Function (fnjmp)

Consider the case

$$\text{exp}_1 \ \text{exp}_2$$

where (a) the call is a tail call and (b) `exp1` is not the name of a region-polymorphic function.

Here `exp1` is evaluated to a closure in memory, pointed to by a *standard closure register*. Then `exp2` is evaluated and the result put in a *standard argument register*. The first word in the closure contains the address of the code of the function. This address is fetched into a third register and a jump to the address is made. Because the call is a tail call, it induces no allocation,

neither on the stack nor in regions. It is thus as efficient as an indirect jump in assembly language.

The mnemonic used in call-explicit expressions for this special case is

```
fnjmp exp1 exp2
```

14.6 Non-Tail Call to Unknown Function (fncall)

Consider the case

```
exp1 exp2
```

where (a) the call is not a tail call and (b) *exp*₁ is not the name of a region-polymorphic function.

Applications of this form are implemented as follows. First *exp*₁ is evaluated and the result, a pointer to a closure, is stored in the standard closure register. Then *exp*₂ is evaluated and stored in the standard argument register. Then live registers and a return address are pushed onto the stack and a jump is made to the code address that is stored in the first word of the closure pointed to by the standard closure register. Upon return, registers are restored from the stack.

The mnemonic used in call-explicit expressions for this special case is

```
fncall exp1 exp2
```

14.7 Example: Function Composition

The Standard ML Basis Library declares function composition as follows³

```
fun (f o g) x = f(g x)
```

The resulting call-explicit expression produced by the MLKit when compiled with the flags `-no_uncurry` and `-no_abbrev` is

```
fun o attop r1 [r21:3] (v120, v121) =
  fn attop r21 x => fnjmp v120 (fncall v121 x)
```

Notice that `f o g` first creates a closure in `r21` and then returns. The closure is of size three words and contains a pointer to the code for the function and pointers to the closures for `f` and `g`. When called, the created function first performs a non-tail call of `g` and then a tail call to `f`.

³Program `kitdemo/compose.sml`.

```

fun foldl attop r1 [r41:4, r39:4] f =
  fn attop r41 b =>
  fn attop r39 xs =>
  case xs of
    nil => b
  | :: v100 =>
    let val v101 = #0 v100; val v102 = #1 v100; region r35:4
    in fncall
      let region r37:4
      in fncall
        fncall foldl[atbot r37,atbot r35] f
        (fncall fncall f v101 b)
      end
    v102
  end
end

```

Figure 14.3: The straightforward implementation of `foldl` uses space linear in the length of the list. (Program `kitdemo/fold1.sml` compiled with the flags `-no_uncurry` and `-no_abbrev`.)

14.8 Example: `foldl` Revisited

Consider the following declaration of folding over lists:⁴

```

fun foldl f b xs =
  case xs of
    [] => b
  | x::xs' => foldl f (f x b) xs'

```

The recursive call of `foldl` is a call of a known function, but not a tail call; `foldl` returns a closure, which is subsequently applied to the value of `(f x b)`. This too returns a closure, which in turn is applied to `xs'`. The resulting call-explicit expression is shown in Figure 14.3. Notice that upon each iteration, fresh regions for holding two closures are being allocated for the duration of the recursive call. Thus, space usage is linear in the length of the list (4 words for each list cell, to be precise).

⁴Program `kitdemo/fold1.sml`.

```

fun foldl atop r1 [r63:3, r61:3] f =
  fn atop r63 b =>
  fn atop r61 xs =>
  let region r27:1;
    fun loop atbot r27 [r39:inf] var2 =
      let val v107 = #0 var2
        in case v107 of
          nil => var2
        | :: v109 =>
          let val v110 = #0 v109;
            val v111 = #1 v109;
            val v112 = #1 var2
          in jmp loop[sat r39]
              (v111,
               fncall fncall f v110 v112
              )sat r39
          end
        end;
      region r57:inf
    in #1 (fncall loop[atbot r57] (xs, b)atbot r57)
    end
end

```

Figure 14.4: The result of compiling the efficient version of `foldl` (`kitdemo/fold2.sml`) is an iterative function that avoids argument pairs piling up in one region.

An alternative version of `foldl` assumes that `f` is curried:⁵

```

fun foldl f b xs =
  let fun loop(p as ([], b)) = p
        | loop(x::xs, b) = loop(xs, f x b)
    in
      #2(loop(xs, b))
    end
end

```

It is compiled into the call-explicit expression in Figure 14.4. Here the loop is implemented as a jump and there is no new allocation in each iteration,

⁵Program `kitdemo/fold2.sml`.

except, of course, for the allocation that `f` might make.⁶

As an exercise, consider the following variant of `foldl`, which assumes that `f` takes a pair as an argument:⁷

```
fun foldl' f b xs =
  let fun loop(p as ([], b)) = p
      | loop(x::xs, b) = loop(xs, f(x,b))
  in
    #2(loop(xs, b))
  end
```

Interestingly, this program contains a potential space leak. Can you detect it? If not, the MLKit will tell you when you compile the program if you pass the compiler the option `-warn_on_escaping_puts`.

Finally, we consider the following version of `foldl`, which resembles the implementation of the `List.foldl` function:⁸

```
local
  fun foldl f e [] = e
    | foldl f e (x::xs) = foldl f (f(x,e)) xs
in
  val a = foldl (op +) 0 [0,1,2]
end
```

Figure 14.5 shows the call-explicit version of the program. We see that the function `foldl` is specialised to the supplied function. Notice that MLKit will specialise small functions, such as `List.foldl`, across compilation unit boundaries, which means that MLKit in general generates very efficient code for calls to such functions.

⁶All the allocations made by the calls to `f` (one call for each element of the list) are put in the same regions. If the list is very long or the values produced large, it may be a good idea to copy the final result to separate regions.

⁷Program `kitdemo/fold3.sml`.

⁸Program `kitdemo/fold4.sml`.

```
let val a =
  let region r9:0;
  fun foldl atbot r9 (var2, var3) =
    case var3 of
      nil => var2
    | :: v109 =>
      let val v110 = #0 v109; val v111 = #1 v109
      in jmp foldl (v110 + var2, v111)
      end;
    region r21:inf
  in foldl (0, [0,1,2] atop r21)
  end
in {|a: _|}
end
```

Figure 14.5: The result of applying and specialising the Basis Library version of `List.foldl`.

Chapter 15

ML Basis Files and Modules

In Section 2.8 we described how to compile and run single-file programs. In this chapter, we describe how to program in the large with the MLKit, using Standard ML Modules and the possibility of organising source files in ML Basis Files. MLKit fully supports Standard ML Modules and it has a sophisticated system for avoiding unnecessary recompilation. In the following section, we describe the notion of ML Basis Files. We then turn to show how to program with structures, signatures, and functors. To enable the programmer to write efficient programs using the Modules language, we shall also explain how the MLKit compiles Modules language constructs.

15.1 ML Basis Files

An ML Basis File, in short MLB-file, is a file that lists the SML source files that make up a project or a library. An MLB-file can also *reference* other MLB-files, so one can organise projects in a hierarchical manner. MLB-files are enforced not to be cyclic.

MLB-files have file extension `.mlb`. The content of an MLB-file is a *basis declaration*, for which the grammar is given in Figure 15.1. We assume a denumerable infinite set of *basis identifiers* Bid , ranged over by *bid*. We use *longbid* to range over *long basis identifiers*, that is, non-empty lists of basis identifiers separated by a punctuation letter (`.`). Basis identifiers can be used for giving a name to a group of compilation units and allow for expressing source dependencies, exactly, as a directed acyclic graph, within one MLB-file.

$bdec$	$::=$	$bdec\ bdec$	sequential basis declaration
		ε	empty basis declaration
		local $bdec$ in $bdec$ end	local declaration
		basis $bid = bexp$	basis identifier binding
		open $longbid^*$	opening of a basis
		$atbdec$	
		$path.mlb$	include
$atbdec$	$::=$	$path.sml$	source file
		$path.sig$	source file
$bexp$	$::=$	bas $bdec$ end	basis declaration grouping
		let $bdec$ in $bexp$ end	let expression
		$longbid$	

Figure 15.1: Grammar for MLB-files, i.e., files with extension `.mlb`. For some file extension `.ext`, $path.ext$ denotes either an absolute path or a relative path (relative to the directory in which the MLB-file is located) to a file on the underlying file system.

In an MLB-file, one can reference source files and other MLB-files using absolute or relative paths. Relative paths are relative to the location of the MLB-file. Paths can reference environment variables using the `$(ENVVAR)` notation, where `ENVVAR` is an environment variable.

Until now, we have seen a few examples of MLB-files that reference the Basis Library, using the `$(SML_LIB)` environment variable (see Section 6.4 for such an example). In Section 15.4, we present an example of an MLB-file that reference other MLB-files. In Section 19.7, we shall see an example of how an MLB-file can be compiled and linked with external object files, produced with a C compiler, for instance. MLB-files may contain Standard ML style comments. The declared identifiers of an MLB-file is the union of the identifiers being declared by source files in the MLB-file, excluding source files that are included using `local`. As an example of the use of basis identifiers and `local` to limit what identifiers are declared by an MLB-file, consult the MLB-file `basis/basis.mlb`.

Every source file must contain a Standard ML top-level declaration; the scope of the declaration is all the subsequent source files mentioned in the MLB-file and all other MLB-files that reference this MLB-file. Thus, a source file may depend on source files mentioned earlier in the MLB-file and on other referenced MLB-files. The meaning of an entire MLB-file is the meaning of the top-level declaration that would arise by expanding all referenced MLB-files and then concatenating all the source files listed in the MLB-file (with appropriate renaming of declared identifiers of source files that are included using `local`), in the order they are listed, except that each MLB-file is executed only the first time it is imported.

MLKit has a system for managing compilation and recompilation of MLB-files. The system guarantees that the result of first modifying one or more source files and then using the separate compilation system to rebuild the executable is the same as if all source files were recompiled.

Thus, the separate compilation system is a way of avoiding recompiling parts of a (possibly) long sequence of declarations, while ensuring that the result is always the same as if one had compiled the entire program from scratch. As an example, consider the MLB-file `kitdemo/scan.mlb` for the text scanning example of Section 12.7. It contains the following three lines:

```
$(SML_LIB)/basis/basis.mlb
lib.sml
scan.sml
```

The source files for the project are `lib_posix.sml` and `scan.sml`, which are both located in the directory where `scan.mlb` is located. Whereas each of the source files `lib.sml` and `scan.sml` depends on the Basis Library, the source file `scan.sml` also depends on `lib.sml`.

Compiling an MLB-file is easy; simply give it as an argument to the MLKit executable. When the MLB-file is first compiled, the MLKit detects automatically when a source file has been modified (by checking file modification dates). After a project has been successfully compiled and linked, it can be executed by running the command

```
./run
```

in the working directory.

MLKit compiles each source file of an MLB-file one at a time, in the order mentioned. A source file is compiled under a given set of assumptions, which provides, for instance, region-annotated type schemes with places for free variables of the source file. Also, compilation of a source file gives rise to exported information about declared identifiers. Exported information may occur in assumptions for source files mentioned later in the MLB-file.

There are two rules that govern when a source file is recompiled. A source file is recompiled if either (1) the user has modified the source file or (2) the assumptions under which the source file was previously compiled have changed. To avoid unnecessary recompilation, assumptions for a source file depend on only its free identifiers. Moreover, if a source file has been compiled earlier, the MLKit seeks to *match* the new exported information to the old exported information by renaming generated names to names generated when the source file was first compiled. Matching allows the compiler to use fresh names (stamps) for implementing generative data types, for instance, and still achieve that a source file is not necessarily recompiled even though source files, on which it depends, are modified.

Let us assume that we modify the source file `lib.sml` of the text scanning example, after having compiled the MLB-file `kitdemo/scan.mlb` once. When compiling the MLB-file again, the MLKit checks whether the assumptions under which the source file `scan.sml` was compiled have changed, and if so, recompiles `scan.sml`. Modifying only comments or string constants inside `lib.sml` or extending its set of declared identifiers does not trigger recompilation of `scan.sml`.

Some of the information a source file depends on is the ML type schemes of its free variables. It also depends on, for example, the region-annotated type

schemes with places of its free variables. Thus it can happen that a source file is recompiled even though the ML type assumptions for free variables are unchanged. For instance, the region-annotated type scheme with place for a free variable may have changed, even though the underlying ML type scheme has not.

As an example, consider what happens if we modify the function `readWord` in the source file `lib.sml` so that it puts its result in a global region. This modification will trigger recompilation of the source file `scan.sml`, because the assumptions under which it was previously compiled have changed. Besides changes in region-annotated type schemes with places, changes in multiplicities and in physical sizes of formal region variables of functions may also trigger recompilation.

15.2 Structures

The support for Modules together with the possibility of dividing top-level declarations into different source files provide a mechanism for programming in the large. In the MLKit, structures exist only at compile time. Thus one need not worry where structures live at runtime.

We illustrate the compile-time nature of structures with the following example. Consider the MLB-file `PolySet.mlb`,¹ which mentions the source files `PolySet.sml`, `INT_SET.sml`, and `IntSet.sml`. The source file `PolySet.sml` contains the following top-level declaration:

```
structure PolySet =
  struct
    type  $\alpha$  set =  $\alpha$  list
    val empty = []
    fun singleton x = [x]
    fun mem(x, []) = false
      | mem(x, y::ys) = x=y orelse mem(x, ys)
    fun union(s1, []) = s1
      | union(s1, x::s2) = if mem(x, s1) then union(s1, s2)
                          else x::union(s1, s2)
  end
```

The code generated by the MLKit for the `PolySet` structure is exactly as if the declarations were written outside of a structure. As a consequence,

¹MLB-file: `kitdemo/PolySet.mlb`.

when you refer to a component of a structure using qualified identifiers (e.g., `PolySet.mem`), no code is generated for fetching the component from the structure. Moreover, when opening a structure, using the `open` declaration, no code is generated for rebinding the identifiers that become visible.

15.3 Signatures

In the MLKit, signature declarations exist only at compile time. That is, a signature declaration does not result in any code being generated. The source file `INT_SET.sml` in the MLB-file `PolySet.mlb`, mentioned earlier, contains the signature declaration

```
signature INT_SET =
  sig
    type  $\alpha$  set
    val empty : int set
    val singleton : int  $\rightarrow$  int set
    val mem : int * int set  $\rightarrow$  bool
    val union : int set * int set  $\rightarrow$  int set
  end
```

Signatures are used in two contexts; for specifying arguments to functors and for providing restricted views of structures using transparent and opaque signature constraints. We defer the discussion of the use of signatures for specifying arguments to functors to Section 15.4.

Transparent signature constraints may both restrict components from a structure and make polymorphic components less polymorphic. Moreover, opaque signature constraints may also make type components of structures abstract. Consider the structure declarations located in the source file `kitdemo/IntSet.sml`:

```
structure IntSet1 : INT_SET = PolySet
structure IntSet2 :> INT_SET = PolySet
```

No code is generated for the structure declarations. Instead, the compiler memorises that if you refer to the long identifier `IntSet1.mem`, for instance, then it is actually `PolySet.mem` that is applied with type instance `int`.

As for the second declaration, opaque signature constraints are eliminated at compile time (after elaboration) and transformed into transparent signature constraints.

15.4 Functors

Functors map structures to structures. MLKit specialises a functor every time it is applied. Thus, types that are abstract for the programmer (inside a functor body) become visible to the compiler. Region-annotated type schemes and other information about identifiers in the actual functor argument are available when the MLKit compiles the functor body.

For practical reasons, it is important that not all functor applications are expanded at once, since this could cause intermediate representations of programs to become as large as (or even much larger than) the entire program. Further, non-restricted in-lining could lead to unnecessary recompilation upon modification of source files. Instead, the largest structure declarations not containing functor applications are compiled into separate chunks of machine object code. Assumptions for compiling these structure declarations are memorised, so that the generated code can be reused upon modification of source files if the assumptions do not change.

Consider the following MLB-file:²

```
$(SML_LIB)/basis/basis.mlb
local utils/utils.mlb
in SET.sml Set.sml SetApp.sml
end
```

The MLB-file reference the MLB-file `utils.mlb` from the `utils` directory.³ This MLB-file provides a structure `ListUtils` that contains the function `pr_list` with type scheme $(\alpha \rightarrow \text{string}) \rightarrow \alpha \text{ list} \rightarrow \text{string}$. The content of the file `Set.sml` is listed in Figure 15.2. It declares the functor `Set`, which takes as arguments the element type for the set, an ordering function on elements, and a function for providing a string representation of elements.

The source file `SetApp.sml` is listed in Figure 15.3. It constructs a structure `IntSet` by applying the functor `Set` to appropriate arguments including an ordering operation on integers and an operation for giving the string representation of an integer. The `IntSet` structure is used for constructing a set `2,5`, which the program prints using the built-in `print` function.

The body of the `Set` functor is instantiated to form the code for the `IntSet` structure. The result of instantiating the `Set` functor is first translated into a

²MLB-file: `kitdemo/Set.mlb`.

³MLB-file: `kitdemo/utils/utils.mlb`.

Lambda program and then translated into a MulExp program. The MulExp call-explicit code for the `mem` function is shown in Figure 15.4.

Notice that the code for the `mem` function holds inlined code for the `lt` function; because the function is sufficiently small, the MLKit propagates its intermediate representation across module boundaries.

```

functor Set (eqtype elem (*total order*)
            val lt : elem * elem → bool
            val pr : elem → string)
: SET where type elem = elem =
struct
  type elem = elem
  type set = elem list
  val empty : set = []
  fun singleton e = [e]
  fun mem x l =
    let fun mem' [] = false
        | mem' (y::ys) = if lt(y,x) then mem' ys
                        else not(lt(x,y))
    in mem' l
    end
  fun union(s1,s2) =
    let fun U (t as ([], [], acc)) = t
        | U ([], y::ys, acc) = U([], ys, y::acc)
        | U (x::xs, [], acc) = U(xs, [], x::acc)
        | U (s1 as x::xs, s2 as y::ys, acc) =
            U(if lt(x,y) then (xs, s2, x::acc)
              else if lt(y,x) then (s1, ys, y::acc)
              else (xs, ys, y::acc))
    in rev(#3(U(s1, s2, [])))
    end
  val pr = fn s ⇒ ListUtils.pr_list pr s
end

```

Figure 15.2: The source file kitdemo/Set.sml.

```

structure IntSet = Set(type elem = int
                    val lt = op <
                    fun pr a = Int.toString a)
open IntSet
val _ = print (pr (union(singleton 2, singleton 5)))

```

Figure 15.3: The source file kitdemo/SetApp.sml.

```
fun mem (x, l) =
  let region r55:1;
      fun mem' atbot r55 var2 =
        case var2 of
          nil => false
        | :: v191 =>
          let val v192 = #0 v191; val v193 = #1 v191
          in case v192 < x of
              true => jmp mem' v193
            | _ =>
              case x < v192 of true => false | _ => true
            end
          end
        in mem' 1
      end
```

Figure 15.4: The MulExp call-explicit code for the `mem` function resulting from instantiating the `Set` functor.

Chapter 16

Garbage Collection

MLKit supports reference tracing garbage collection in combination with the region memory model [Hal99, HET02]. Garbage collection is also possible with region profiling enabled.

The reference-tracing garbage collector is enabled by default and, as we have seen earlier, garbage collection can be disabled by passing the `-no_gc` option to the MLKit compiler at compile time. As we shall see, it is also possible to disable garbage collection at runtime for a program that has been compiled with garbage collection enabled.

MLKit also features generational reference-tracing garbage collection (option `-gengc`), which in some cases is superior to ordinary reference-tracing garbage collection, but which may also cause additional fragmentation [EH21b].

16.1 Dangling Pointers

The region type system supports deallocation of memory that is not accessed in the remainder of the execution of the program. Because of this principle, the execution model may lead to *dangling pointers*, that is, pointers that point into memory that has been discharged. When garbage collection is enabled, the region type system is modified slightly so as to guarantee that no dangling pointers occur during execution [Els03, Els23]. The following example illustrates how the enabling of garbage collection changes the way programs are compiled:

```
val f = let val x = ref (2, [1])
        in fn y => (#1 (!x), y)
```

```

    end
  val r = f 5

```

When garbage collection is disabled, the program is compiled into the following MulExp program:¹

```

val f =
  let region r9:inf;
    val x = ref attop r6 (2, [1] attop r9)attop r4
    in fn attop r1 y => (#0 (!x), y)attop r4
    end
  val r = fncall f 5

```

Notice here that region `r9`, which contains the list `[1]`, is deallocated before the function `f` is applied to the value `5`. If we chose to run this program together with a reference tracing garbage collector, a fatal error could occur: The memory that contains the list `[1]` could be reused for other purposes at the time the garbage collector tries to trace the dangling pointer.

Figure 16.1 shows the MulExp program produced when garbage collection is enabled.² When garbage collection is enabled, the MLKit makes sure

```

val f =
  let val x = ref attop r6 (2, [1] attop r4)attop r4
    in fn attop r1 y => (#0 (!x), y)attop r4
    end
  val r = fncall f 5

```

Figure 16.1: The MulExp program produced when compiling the program `kitdemo/dangling.sml` with garbage collection enabled. To avoid dangling pointers when garbage collection is enabled, all values in the closure for `f` are kept alive as long as the closure itself.

that all values stored in a closure are kept live as long as the closure is live. Assume that the type with place μ of the function associated with the closure is on the form $(\mu_1 \xrightarrow{\epsilon, \varphi} \mu_2, \rho_0)$. MLKit enforces the restriction by requiring

¹Compiled with `mlkit -no_gc -maximum_inline_size 0 -Ppse -w 40 dangling.sml` from within the `kitdemo` directory.

²Compiled with `mlkit -gc -maximum_inline_size 0 -Ppse -w 50 dangling.sml` from within the `kitdemo` directory.

that for each region variable ρ that occur free in the type of free variables of the function (those variables for which values are stored in the closure at runtime), ρ occur free in μ . In the implementation, the requirement may lead to extra `get` effects being added to $\epsilon.\varphi$ when garbage collection is enabled. In the example, an imposed `get` effect on the arrow effect in the type for `f` makes it impossible to wrap a `region` binding around the binding for `f`. (See [TT93, page 50], [Els03], and [Els22, Els23] for more information about this requirement.)

16.2 Scanning Text Files with Stream IO and Garbage Collection

In Section 12.7, we saw how we could use basic non-buffered `Posix` IO operations for scanning text files in constant space. The Standard ML Basis library, however, also features a stream-based `TextIO` structure that supports arbitrary look-ahead, buffering, stream-redirection, and many other features [GR04]. Unfortunately, as mentioned in Section 12.7, this functionality does not work well together with MLKit’s region based memory management, unless combined with a mechanism for dynamically garbage collecting regions. We now consider a modified version of the the text scanning program listed in Section 12.7 that uses the `TextIO.inputN` function for file reading, instead of using `Posix` buffered IO.

Figure 16.2 shows two region profiles for the modified text scanning program.³ We see that the reference-tracing garbage collector periodically cleans up the global regions `r1` and `r4`, which tends to grow steadily without garbage collection.

16.3 Instrumenting the Executable

Executables produced by the MLKit with garbage collection enabled can be instrumented by use of command-line options. For instance, if the MLKit has produced a file `run`, one can pass the option `-verbose_gc` to `run` to enable the printing of garbage collection information at runtime. An overview of available command-line options is shown by passing the option `-help` to the generated executable:

³MLB-file: `scan_stream.mlb`.

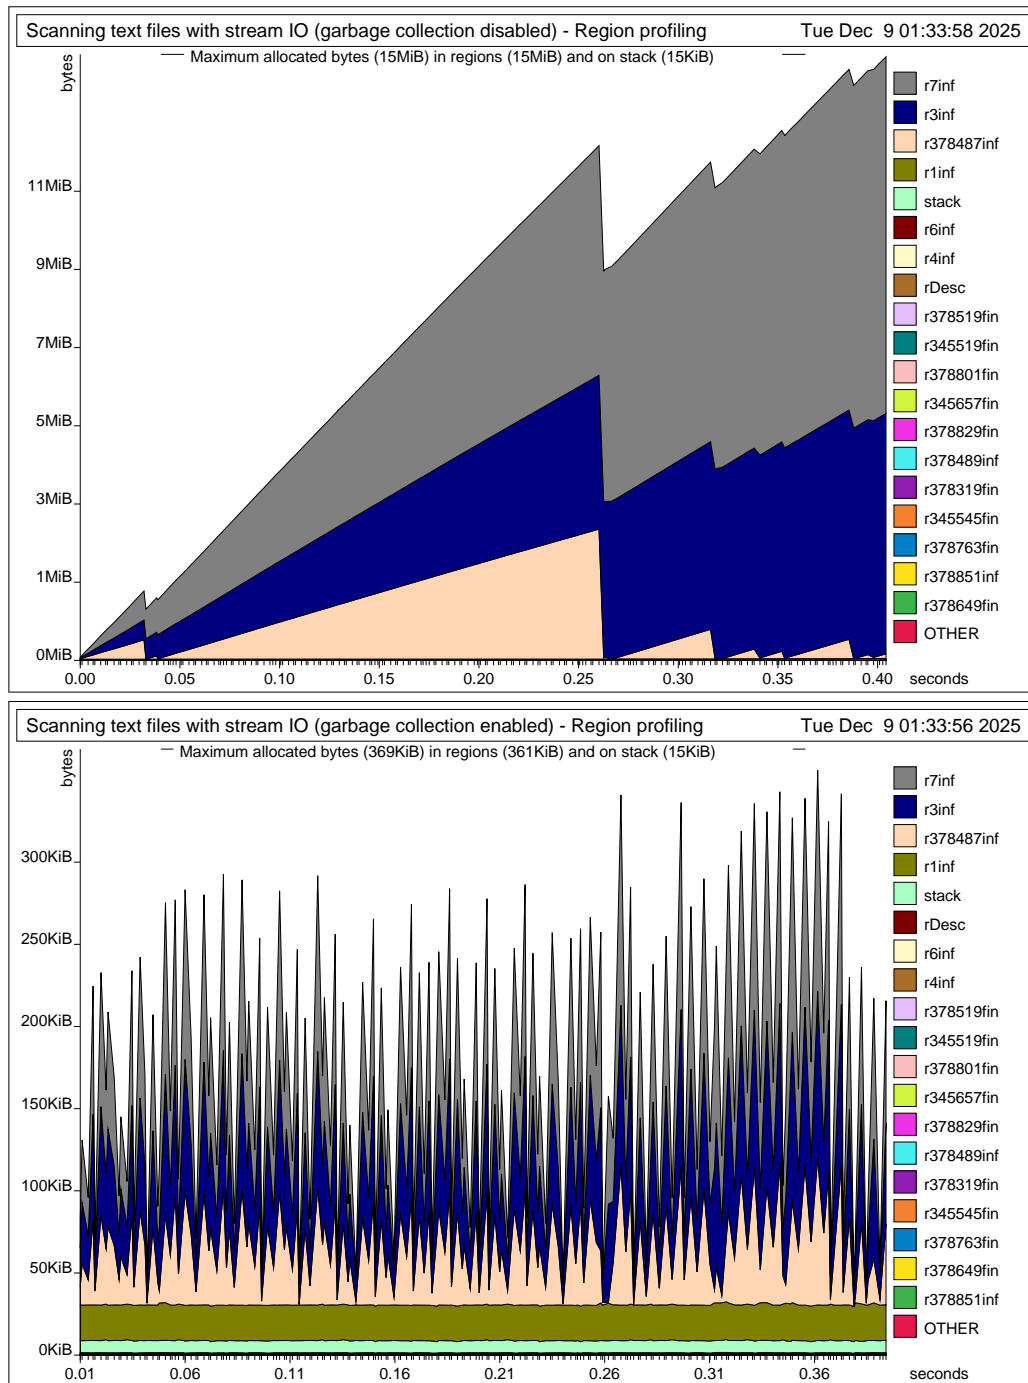


Figure 16.2: Two region profiles of text scanning using buffered stream-based IO operations. The top region profile is with reference-tracing garbage collection disabled and the bottom region profile is with reference-tracing garbage collection enabled.

```
Usage: ./run
       [-help, -h]
       [-disable_gc | -verbose_gc] [-heap_to_live_ratio d]
where
  -help, -h           Print this help screen and exit.

  -disable_gc         Disable garbage collector.
  -verbose_gc         Show info after each collection.
  -heap_to_live_ratio d Use heap to live ratio d, ex. 3.0.
```


Part III
System Reference

Chapter 17

Region Profiling

We have already seen several examples of the use of the profiler. We shall now explain in more detail how to profile programs. For example, we shall see how one can find out precisely what allocation points in the program contribute to allocation in a particular region.

The profiler consists of several tools that can be used to analyse the dynamic memory behavior of a program. First of all, the profiler lets you create graphs of the dynamic memory usage of the program. Three different kinds of graphs may be created:

- A *region profile* is a graph that gives a global view of the memory usage by showing the total number of bytes allocated in regions and on the stack as a function of time. In the graph, regions that arise from the same

region ρ

construct are collected into one colored band, labeled ρ . The region variables that label bands are always global or `region`-bound, never formal region parameters.

- An *object profile* is a graph that, for a particular region, shows the objects allocated in the region, with one colored band for each allocation point in the region-annotated program¹. Each allocation point is annotated with a *program point*, which is a unique number that identifies

¹Every occurrence of an `at` in the region-annotated program is an allocation point.

the allocation.² To inspect region-annotated programs with program points, pass the MLKit compiler the option `-print_program_points` in addition to the option `-print_call_explicit_expression`, say.³

If you have an object profile showing that program point `pp42`, say, contributes with allocation, you can search for `pp42` in the region-annotated program and thus find the construct that caused the allocation.

- A *stack profile* is a graph that shows the stack memory usage, as a function of time.

In addition to the possibility of generating programs with program points, it is also possible, during compilation, to generate a *region flow graph*, which shows how regions may be passed around at runtime when region-polymorphic functions are applied. The region flow graph comes in handy when profiling large programs and when one wants to find out why a formal region variable is instantiated to a certain `region-bound` region variable.

The following example clarifies the use of a region flow graph. Suppose the region profile shows that `r5` is responsible for most of the memory usage. Further, suppose an object profile of `r5` shows that program point `pp345` is responsible for most of the allocation. Searching for `pp345` in the region-annotated program, you may find that the allocation at `pp345` is into some other region variable, `r34`, say. Here `r34` will be a formal region parameter of a region-polymorphic function that at runtime has been instantiated to `r5` by one or more calls of region-polymorphic functions. You can now use the region flow graph to find the cascade of region polymorphic applications that ends up instantiating `r34` to `r5`.

The profiling process is sketched in Figure 17.1.

We will now show an example on how to profile a concrete program that contains a space leak and then show how the profiler can be used to improve the program. We then explain in more detail how to specify the profiling strategies and how the profiles are generated.

²Program points are unique. In particular, for a project with two program units, the program points in the region-annotated programs for the two units will be distinct.

³Program points are annotated during physical size inference.




Figure 17.1: Overview of the profile process. The process sometimes requires the programmer to refine the runtime profiling strategy, or even the compile-time profiling strategy. Dotted boxes represent output from the compiler, from executing the program, and from using the tool `rp2ps`, which generates Postscript graphs from the exported data file.

17.1 Example: Scanning Text Files Again

In this section, we concentrate on the general principles of profiling. As an example, we investigate a revised version of the project `kitdemo/scan.mlb` (see Section 12.7). Instead of asking for a list of input files to scan (as project `scan.mlb` does), the revised version of the scan project asks for only one input file, which it then scans 50 times.⁴

The first thing to do is to get an overview of the memory usage of the program. A region profile of the program gives you just that. See Figure 17.2.

The graph shows that region `r331112` accumulates more memory for each time it scans the file `life.sml`.

To see what happens in region `r331112`, we make an object profile of that region, see Figure 17.3. The object profile shows that program point `pp115` continually allocates memory that is first freed when the program terminates. We now search for `pp115` in the generated log files and in the log file located in the basis library folder `$(SMLLIB)/basis`.⁵ Among other places, we find that the program point `pp115` appears in the following fragment in the file `basis/Int.sml.log`:

```
fun conv attop r1T pp65 [r58078s:inf] (rad, radix, i) =
  case ... of ...
  | _ =>
    let ...
    in String.implode [sat r58078s pp115] chars
    end
```

We also see that this function is called directly by the function `Int.toString` and that the formal region parameter `r59522` is passed directly as a parameter in the call to `conv`:

```
fun toString attop r1T pp775 [r59522s:inf] i =
  let val v1738 = DEC
  in jmp conv[sat r59522s pp776] (v1738, 10, i)
  end
```

⁴Project `kitdemo/scan.rev1.mlb`.

⁵Unfortunately, the MLKit compiler happily resets the program point counter for each compilation of a program unit, thus, there are multiple instances of the program point `pp115`.




Figure 17.2: Memory is accumulated in the top two bands. The global regions `r1` and `r331112` hold the largest amount of memory. The graph was generated by first compiling the `kitdemo/scan_rev1.mlb` project with profiling enabled (option `-prof`). We also add the compiler options `-log_to_file`, `-Ppp`, `-Prfg`, and `-Pcee` for enabling printing of region flow graphs, program points, and call-explicit expressions (redirected to log files). Then by executing `echo life.sml | run -notimer 723` and finally by typing `rp2ps -region -name 'Scanning life.sml 50 times' -sampleMax 200`.




Figure 17.3: There seems to be a space leak at program point pp115. The graph was generated by typing `rp2ps -object r331112`.

So the space leak is caused by the `String.implode` primitive function being called with region `r331112` instantiated for the formal region variable `r59522`.

We now search for `r331112` in file `scan_rev1.sml.log` and find the following fragment of the region flow graph:

```
LETREGION[r331112:inf]
readWord[r330770:inf]  --r330770 atbot→  [*r331112*] ;
toString[r59522:inf]  --r59522 attop→  [*r331112*] ;
```

The fragment is read as follows. The formal region variable `r59522` is instantiated to the `region`-bound region variable `r331112` in a call to `toString`. Moreover, also the formal region variable `r330770` (of function `readWord`) is instantiated to `r331112`. (The asterisks (*) denote that the node has been displayed before.)

Region flow graphs are local to each program fragment in a program. A call to a non-local region-polymorphic function introduces an edge in the region flow graph, but the graph says nothing about in which module the called function is located. Thus, it may be necessary to look in several log files to find the path from a formal region variable to an actual region variable. By inspecting the call-explicit programs found in `basis/Int.sml.log` and `kitdemo/lib.sml.log` one finds that both `toString` and `readWord` eventually call `implode`. However, `readWord` is called only initially, thus, we conclude that the space leak is caused by function `Int.toString` being called with region `r331112` instantiated for the formal region variable `r59522`. Indeed, by inspecting the calls to `toString` in the call-explicit program found in `scan_rev1.sml.log`, we see that `toString` is called with actual region `r331112`.

The `concat` function from the initial basis concatenates a list of strings. But all the strings in the argument list to `concat` are required to be in the same region. Thus, whenever a file is reported (see Figure 17.4), strings created by the `Int.toString` function are put in the region that also holds the file name for the report (which is read using the function `readWord`); and this region is non-local to the `do_it` function, which implements the main loop of the program.

One way of solving the space leak is to make a copy of `filename` at the call to `report_file` in function `scan_file`:

```
fun scan_file filename : (int*int) option =
  let val fd = F.openf (filename, F.O_RDONLY, F.O.flags[])
  in let val (n, ins) = scan fd
     in Posix.IO.close fd;
```

```

fun report_file (filename, n, ins) =
  writeln(concat[filename, ": size = ", Int.toString n,
    " comments: ", Int.toString ins, " (",
    (Int.toString(percent(ins, n))
    handle _ => "-"), "%)"])

fun scan_file filename : (int*int) option =
  let val fd = F.openf (filename, F.O_RDONLY, F.O.flags[])
  in let val (n, ins) = scan fd
    in Posix.IO.close fd;
      report_file (filename, n, ins);
      SOME (n, ins)
    end handle NotBalanced =>
      (writeln (filename ^ ": not balanced");
        Posix.IO.close fd;
        NONE)
  end handle IO.IOException {name,...} =>
    (writeln (name ^ " failed."); NONE)

fun main () : unit =
  case readWord F.stdin of
    SOME filename =>
      let fun do_it 0 = ()
          | do_it n = (scan_file filename; do_it (n-1))
      in do_it 50
      end
    | NONE => ()

```

Figure 17.4: Fragments of `scan_rev1.sml`. All the strings in the argument list to `concat` are put in the same region.




Figure 17.5: There is no space leak: no matter how many times we scan the file, the project will use the same number of words. The graph was generated by executing `echo life.sml | run -notimer 723` and `rp2ps name scan_rev2 -region`.

```

report_file (filename^"", n, ins);
SOME (n, ins)
end handle NotBalanced =>
    (writeln (filename ^ ": not balanced");
     Posix.IO.close fd;
     NONE)
end handle IO.Io {name,...} =>
    (writeln (name ^ " failed."); NONE)

```

Project `kitdemo/scan_rev2.mlb` implements the modification. Figure 17.5 shows a region profile of the `scan_rev2.mlb` project.

17.2 Compile-Time Profiling Strategy

Before compiling a program for the purpose of profiling, one must decide on a *compile-time profiling strategy*; see Figure 17.1. The compile-time profiling strategy directs the embedding of profiling instructions in the generated code and instructs the compiler whether to report a region flow graph.

Region profiling is enabled by passing the option `-region.profiling` (or simply `-prof`) to the MLKit compiler. If you want the MLKit to report region-annotated programs with program points, you should pass the option `-print.all.program.points` (or `-Ppp`) to the MLKit compiler together with one or more of the options `-print.physical.size.inference.expression` (or `-Ppse`) and `-print.call.explicit.expression` (or `-Pcee`).

To make the compiler report a region flow graph, pass the option

```
-print.region.flow.graph
```

(or `-Prfg`) to the MLKit compiler at compile time. The region flow graph is reported in text format.

As a running example, we use the `life` program.⁶ We assume that the options `-prof`, `-Ppp`, `-Prfg`, and `-Pcee` are passed to the MLKit compiler.

By also passing the option `-log.to.file` to the MLKit compiler, the MLKit generates several files, of which we have `life.sml.log` (containing, among other things, the call-explicit region-annotated program with program points and the region flow graph in text layout) and the executable file `run`.

17.3 The Log File

In the file `life.sml.log` you find the call-explicit region-annotated program with program points and the region flow graph in text layout for the `life.sml` source file. The region flow graph is found by searching for the string `REGION FLOW GRAPH FOR PROFILING`. The graph contains the following fragment (modified slightly to fit here):

```
cp_list[r330702:inf]
  --r330702 sat--> nthgen[r332514:inf]
    --r332514 atbot--> LETREGION[r332640:inf] ;
    --r332514 sat--> [*r332514*] ;
```

⁶Program: `kitdemo/life.sml`.

```

--r330702 atbot--> LETREGION[r332002:inf] ;
--r330702 sat--> [*r330754*] ;
--r330702 sat--> [*r330702*] ;

```

The region flow graph is almost equivalent to the graph used by the storage mode analysis (see page 108). In the graph, region variables are nodes and there is an edge between two nodes ρ and ρ' if ρ is a formal region parameter of a function that is applied to actual region parameter ρ' . It follows that **region**-bound region variables are always leaf nodes.

Nodes in the graph are written in square brackets, which are labeled with the token **LETREGION** or the name of the function for which the region variable is a formal parameter. For example, the notation `cp_list[r330702:inf]` identifies the node `r330702`, which is a formal region parameter of the function `cp_list`. An asterisk inside a square bracket means that the node has been written earlier. Only the node identifier (i.e., the region variable) will then be printed. The size of the region is printed after the region variable; we use `inf` for infinite regions and `size` for finite regions of size `size` words.

Edges are written with the *from node* identifier annotated on them. The edge points to the *to node*. The fragment

```

cp_list[r330702:inf]
  --r330702 sat-->    [*r330754*] ;

```

is read: there is an edge from node `r330702` to node `r330754` and node `r330754` has been written earlier. From the cycle in the graph, one can conclude that `cp_list` calls itself recursively; if you look in file `life.sml`, you will find something like

```

fun cp_list [] = []
  | cp_list ((x,y)::rest) =
    let val l = cp_list rest
    in (x,y)::l
    end

```

The region flow graph can get very complicated to read due to mutually recursive functions, which give many edges and cycles. If the graphs get too complicated, you may find help in the *strongly connected component* (scc) version of the graph. The scc graph is found by searching for `[sccNo` in the log file. Each scc is identified by a unique *scc number*. The region variables contained in each scc is annotated on the scc node.

Consider, for example, the following fragment of the scc version of the region flow graph for the `life` program:

```
[sccNo 94: r2768,]  --sccNo 94-->  [sccNo 33: r331990,];
```

Here, we have a scc node (id 94) containing region variable `r2768` and an edge to scc node (id 33) containing region variable `r331990`.

17.4 Runtime Profiling Strategy

When the source program has been compiled and linked, you have an executable file, `run`. Typing `./run` at the command prompt will execute the program with a predefined runtime profiling strategy, which is displayed when the program is run with the `-verbose` option:

```
-----Profiling-Enabled-----
The profile timer (unix virtual timer) is turned on.
A profile tick occurs every 1th second.
Profiling data is exported to file profile.rp.
-----
```

You can change the profiling strategy by passing command line arguments directly to the executable. The second line says that a virtual timer is used. There are three possible timers, each of which can be enabled using one of the following options:⁷

```
-realtime      Real time.
-virtualtime   The execution time for the process.
-profiletime   The execution time for the process together
               with the time used in the operating system
               on behalf of the process.
```

The third line says that a *profile tick* occurs every 1 second. A profile tick is when the program stops normal execution, and memory is traversed to collect profile data. The more often a profile tick occurs the more detailed you profile (and the slower the program will run). The *time slot* (i.e., the time between to succeeding profile ticks) to use is specified by the `-sec n` and `-microsec n` options. A time slot of half a second is specified by `-microsec 500000` and not by `-sec 0.5`.⁸

⁷A complete description can be found in the manual page for `getitimer`.

⁸The lowest possible time slot to use is system dependent.

The fourth line says that the collected profile data is exported to the file `profile.rp`. The default file name setting can be changed with the `-file name` option.

There are several other possible command-line options; use the `-h` option or the `-help` option for details. When garbage collection is enabled, options for controlling garbage collection are also available as command-line options (see Section 16).

17.5 Regions Statistics

If the executable file `run` is executed with the option `-showStat` then *region statistics* is printed just before the program terminates. Region statistics includes information about the use of regions and does not depend on the specifics of the runtime profiling strategy; in fact, region statistics includes only exact, non-sampled values for the program. Assuming that `run` is the executable file generated by compiling the program `life` with profiling enabled, executing `./run -showStat` yields—just before the program terminates—the region statistics shown in Figure 17.6.

The `MALLOC` part of Figure 17.6 shows how memory is allocated from the operating system.

Each infinite region form a linked list of one or more *region pages* whose size is found in the `REGION PAGES` part. The value

```
Max number of allocated pages: 44
```

multiplied by

```
Size of one page: 8176 bytes
```

gives

```
Max space for region pages: 359744 bytes (0.3Mb)
```

In the `INFINITE REGIONS` part, we see the number of calls to infinite region operations such as `allocateRegionInf` and `alloc`. The program allocates 81050 infinite regions and deallocates 81044; the six global regions are not deallocated before the region statistics is printed and the program terminates. The program allocates 1060622 objects in infinite regions. Infinite regions have been reset 85882 times. The `deallocateRegionsUntil` operation is called whenever an exception is raised, thus, we see that no exceptions were raised by the program.

MALLOC

Number of calls to malloc for regions: 1
Alloc. in each malloc call: 819200 bytes
Total allocation by malloc: 819200 bytes (0.8Mb)

REGION PAGES

Size of one page: 8176 bytes
Max number of allocated pages: 44
Number of allocated pages now: 6
Max space for region pages: 359744 bytes (0.3Mb)

INFINITE REGIONS

Size of infinite region descriptor: 32 bytes
Number of calls to allocateRegionInf: 81050
Number of calls to deallocateRegionInf: 81044
Number of calls to alloc: 1060622
Number of calls to resetRegion: 85882
Number of calls to deallocateRegionsUntil: 0

ALLOCATION

Max alloc. space in pages: 121936 bytes (0.1Mb)
incl. prof. info: 243888 bytes (0.2Mb)
Infinite regions utilisation (243888/359744): 68%
Number of allocated large objects: 0

STACK

Number of calls to allocateRegionFin: 358935
Number of calls to deallocateRegionFin: 358935
Max space for finite regions: 144 bytes (0.0Mb)
Max space for region desc: 1280 bytes (0.0Mb)
Max size of stack: 19792 bytes (0.0Mb)
incl. prof. info (stackBot - maxStack): 20232 bytes (0.0Mb)
prof. info: 440 bytes (0.0Mb)
in profile tick: 0 bytes (0.0Mb)

Figure 17.6: Region statistics for the life program.

Because objects allocated in infinite regions are not split across different region pages, it is not always possible to fill out a region page entirely. In the `ALLOCATION` part, the value

```
Infinite regions utilisation (243888/359744): 68%
```

shows memory utilisation for infinite regions at the moment where the program has allocated the largest amount of memory in infinite regions. We also see that the maximum allocated space in region pages is 121936 bytes.

In the `STACK` part, we see that the program allocates and deallocates the same number of finite regions. We also see that the space used for finite regions is 144 bytes and that the total use of stack space is 19792 bytes (excluding space used to hold profiling information). The stack size values

```
prof. info: 440 bytes (0.0Mb)
in profile tick: 0 bytes (0.0Mb)
```

can be used to see if it is necessary to profile with a smaller time slot, which will often lower the difference between the two values.

17.6 Processing the Profile Data File

The profile data-file `profile.rp` can be processed by the graph generator `rp2ps` (read: RegionProfile2Postscript) found in the `bin` directory.⁹ The graph generator is controlled by command line options.

A region profile is produced by typing

```
$ rp2ps -region
```

at the command prompt. The program produces a Postscript file `region.ps` by reading profile information from the profile data file `profile.rp`, see Figure 17.1. A region profile for the `life` program is shown in Figure 1.5 on page 30. The region that occupies the largest area is at the top. If there are more regions than can be shown in different colors, then the smallest regions are collected in an `OTHER` band at the bottom.

Each region is identified with a number that matches a `region`-bound region variable in the region-annotated program. Infinite regions end with `inf` and finite regions end with `fin`. There are also a band named `rDesc` and a band named `stack`. The `rDesc` band shows the memory used on region

⁹The `rp2ps` program is based on a Haskell profiler written by Colin Runciman, David Wakeling and Niklas Røjemo.

descriptors of infinite regions on the stack. The stack band shows stack usage excluding finite regions and region descriptors for infinite regions.

The vertical line marked “Maximum allocated bytes...” in Figure 1.5 is called the *maximum allocation line*; it shows the maximum number of bytes allocated in regions and on the stack when the program was executed. Notice that this maximum is often lower than the sum of the maximum allocated bytes in regions and the maximum allocated bytes on the stack. The space between the maximum allocation line and the top band shows the inaccuracy of the profiling strategy. To decrease the gap, it often helps to use a smaller time slot.

The largest region shown in Figure 1.5 is `r332640`. An object profile of region `r332640` is produced by typing

```
$ rp2ps -object 332640
```

at the command prompt. We obtain the object profile shown in Figure 17.7.

We see that allocation point `pp11` is responsible for the largest amount of allocations in the program. The allocation point may be found in the region-annotated program resulting from compiling the `life` program (remember to enable printing of program points). In general, program points may also stem from the Basis Library (search the `.log` files in the directory `basis`).

The stack profile shown in Figure 17.8 shows memory usage on the stack, excluding space used by finite regions. A stack profile is generated by typing

```
$ rp2ps -stack
```

at the command prompt.

17.7 Advanced Graphs with `rp2ps`

This section gives a quick overview of the more advanced options that can be passed to `rp2ps`. First of all, it is possible to name the profiles with the `-name` option. Comments are inserted on the x-axis with the `-comment` option.

The profile data file may contain a large number of *samples* (the data collected by a profile tick is called a sample). By default, `rp2ps` uses only 64 samples. You can alter the setting with the `-sampleMax` option. The following two possibilities are used to sort out samples:

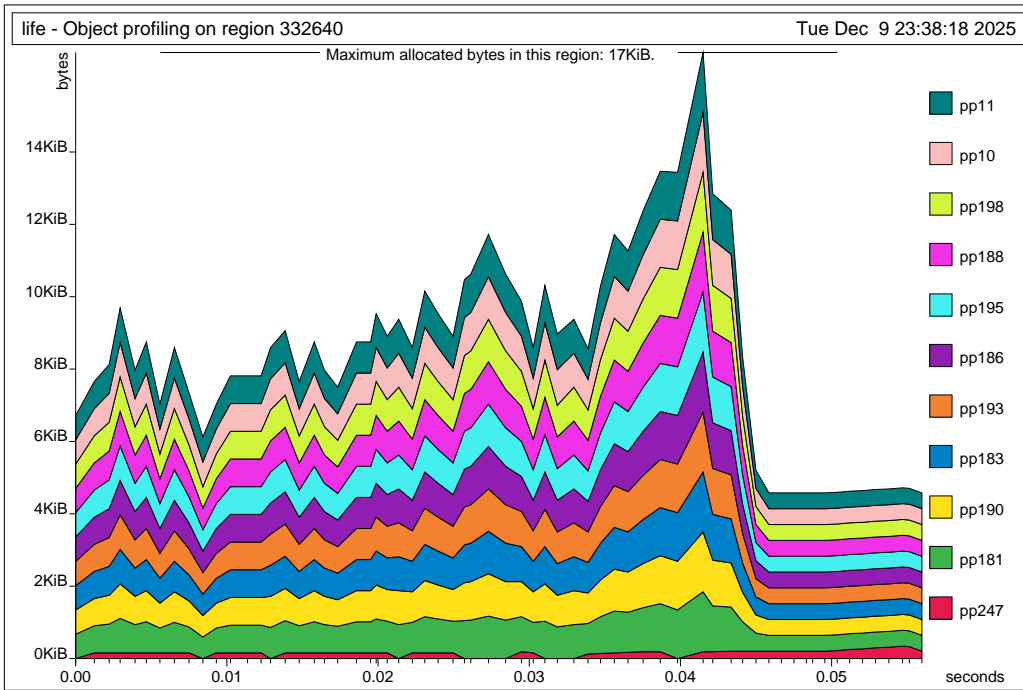


Figure 17.7: The object profile shows all allocation points allocating into region r332640.

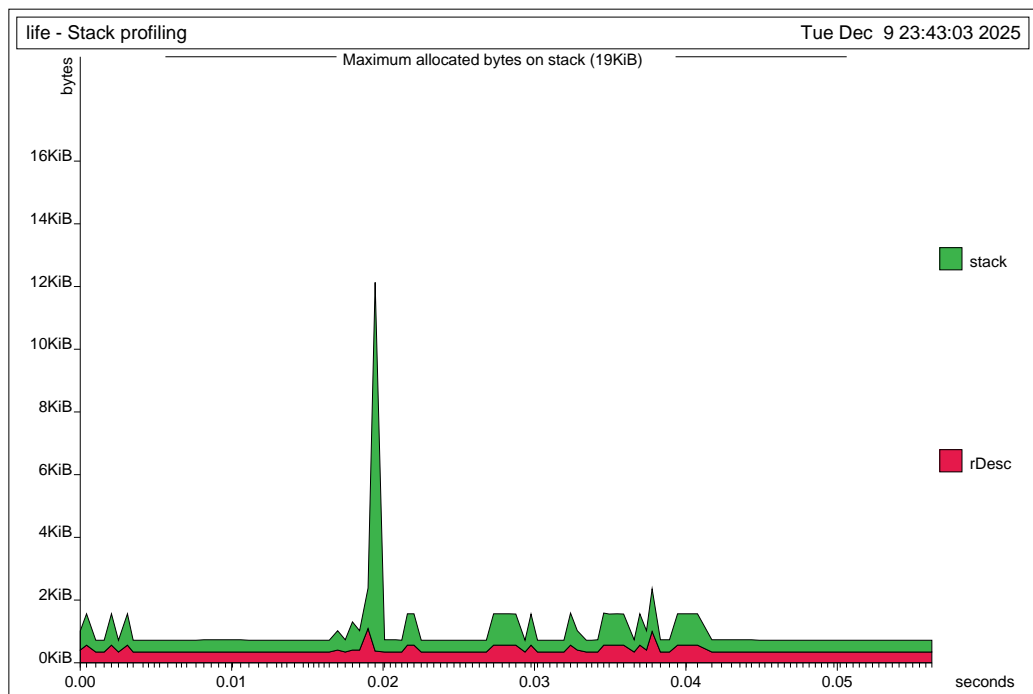


Figure 17.8: Memory usage on the stack excluding space for finite regions.

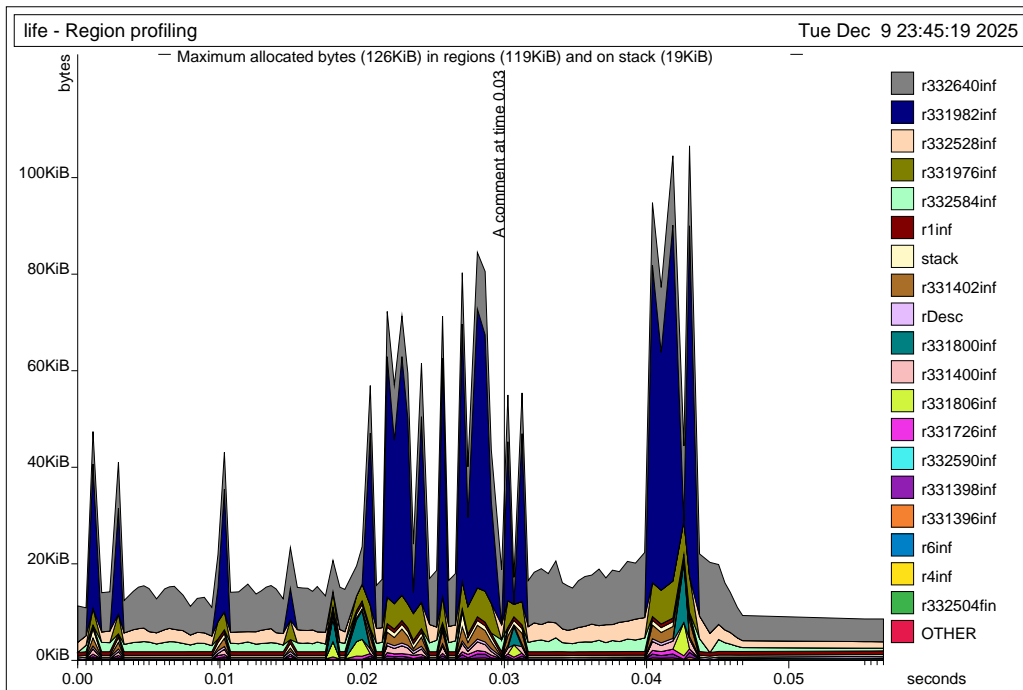


Figure 17.9: It is possible to insert comments in profile graphs.

- sortBySize The n (specified by `-sampleMax`) largest samples are shown.
- sortByTime The n samples shown are equally distributed over time (default).

The `-sortBySize` option is useful if your profiles have a large gap between the top band and the maximum allocation line. If there is a large gap when using option `-sortBySize`, then it may help to profile with a smaller time slot. You can use the `-stat` option to see the number of samples in the profile data file. It is printed as `Number of ticks:`.

Figure 17.9 shows the profile for the following command line:

```
$ rp2ps -region -sampleMax 100 -name life \
    -comment 0.03 "A comment at time 0.03" -sortBySize
```

The graph generator recognises several options that are not mentioned here. Help on these options is obtained by typing `rp2ps -h` or `rp2ps -help`

at the command prompt.

Chapter 18

Controlling MLKit Compilation

We have already described how to compile and run single source files (Section 2.8) and MLB-files (Chapter 15). In the following sections, we give an overview of MLKit options for controlling printing and layout of intermediate forms. One useful command-line option is the `-help` option; Appendix A shows the output of executing `mlkit -help`.

18.1 Printing of Intermediate Forms

A series of options may be used to control printing of intermediate forms during compilation. A summary of the major phases that produce printable intermediate forms is shown in Figure 18.1. The phases are listed in the order they take place in the MLKit.

The optimiser, which rewrites a Lambda program, collects statistics about the optimisation. This statistics is printed if the option

```
-statistics_after_optimisation
```

is provided.

Storage mode analysis (see Chapter 12) results in a MulExp expression, which is printed if the option `-print_storage_mode_expression` is provided. After that, region parameters for which there are only **get** effects on in the type scheme for a region polymorphic function are removed from the MulExp expression (see page 64). To see the resulting expression, turn on `-print_drop_regions_expression` or

```
-print_drop_regions_expression_with_storage_modes
```

Phase	Result	Flag(s) that Print Result
Elaboration	Lambda	(*)
Elim. of Poly. Eq.	Lambda	(*)
Lambda Optimiser	Lambda	-Pole (*)
Spreading	RegionExp	(*)
Region Inference	RegionExp	(*)
Multiplicity Inference	MulExp	(*)
K-normalisation	MulExp	
Storage Mode Analysis	MulExp	-Psme (*)
Dropping of Regions	MulExp	-Pdre (*)
		-Pdresm
Physical Size Inference	MulExp	-Ppse (*)
Call Conversion	MulExp	-Pcee (*)

Figure 18.1: The table shows how different options correspond to printing different intermediate program representations. The option `-debug` causes all intermediate forms marked (*) to be printed. Thus, one can select phases individually or ask to have all intermediate forms printed. The phases that follow K-normalisation all work on K-normal forms, but, for readability, terms are printed as if they had not been normalised.

The latter flag also prints storage modes.

Physical size inference then determines the size in words of finite region variables. For instance, a finite region that will contain a pair will have physical size two words. To see the expression after physical size inference, provide the option `-print_physical_size_inference_expression`. After that, call conversion converts the MulExp expression to a call-explicit expression (see page 138). To see the result, provide the option

```
-print_call_explicit_expression
```

After that, you can inspect the code at different steps of the transformation into machine code by providing different options (use the `-help` option to see which).

18.2 Layout of Intermediate Forms

While the switches described in the previous section concern which intermediate forms to print, the switches described in this section control how the different forms are printed.

The options `-print_types`, `-print_effects`, and `-print_regions` control the printing of region-annotated types, effects, and region allocation points (e.g., `at ρ`). All eight combinations of these three flags are possible, but if `-print_effects` is turned on, it is best also to turn the two others on so that one can see where the effect variables and region variables that appear in arrow effects are bound.

Chapter 19

Calling C Functions

In this chapter, we describe how the MLKit programmer can call C functions from within Standard ML programs. MLKit allows ML values to be passed to C functions, which again may return ML values. Not all ML values are represented as if they were C values. For instance, C strings are null-terminated arrays of characters, whereas ML strings in the MLKit are represented as a linked list of bounded sized character arrays. To allow the programmer to conveniently convert between C values and ML values, the MLKit provides conversion functions and macros for commonly used data structures.

When the MLKit calls a C function, data structures returned by the function are stored in regions that are allocated by the MLKit. For dynamically sized objects of the resulting value, such as strings and lists, regions are allocated by the MLKit and passed to the C function as additional arguments; the C function must then itself allocate space in these regions for the dynamically sized data structures. Moreover, for those parts of the resulting value for which the size can be determined statically, pointers to already allocated space are passed to the C function as additional arguments.

In both cases, the MLKit uses region inference to infer the lifetime of regions that are passed to the C function. The region inference algorithm does not analyse C functions. Instead, the MLKit inspects the ML type provided by the programmer. MLKit assumes that functions with monomorphic types are region exomorphisms; region endomorphic functions may be described using ML polymorphism, see Section 19.6.

For every C function that is called from an ML program, the order of the additional region arguments is uniquely determined by the ML result type of the function. This type must be constructed from lists, records, booleans,

reals, strings, integers, and type variables.

When profiling is enabled, yet another additional argument, a program point, is passed to the C function. This argument provides allocation primitives with information about what points in the program contributes with allocation; see Section 19.4.

Examples of existing libraries that can be accessed from within ML programs include database client drivers (e.g., for Postgresql) and standard UNIX libraries providing functions such as `time`, `cp`, and `fork`. There are limitations to the scheme, however. First, because C and the MLKit do not share value representations, transmitting large data structures between C and ML will often involve significant copying. Second, some C libraries require the user to set up call-back functions to be executed when specific events occur. MLKit only has very limited functionality for letting a C function call ML code.

19.1 Declaring Primitives and C Functions

MLKit conforms in large parts to the Standard ML Basis Library. Part of the functionality found in the basis library is programmed in C and linked to the MLKit runtime system. The declarations in system dependent parts of the library use a special built-in identifier called `prim`, which is declared to have type scheme $\forall\alpha\beta.\text{string} * \alpha \rightarrow \beta$ in the initial basis. A primitive function is then declared by passing its name to `prim`. For example, the declaration

```
fun (s : string) ^ (s' : string) : string =
  prim ("concatStringML", (s, s'))
```

declares string catenation. The argument and result types are explicitly stated so as to give the primitive the correct type scheme. The string constant `"concatStringML"` denotes a C function identifier.¹ For the example declaration, the MLKit generates a call to the C function `concatStringML` with arguments `s` and `s'`. The C function must then of course be present at link-time; if not, the MLKit complains.²

A convenient way to declare a C function is to use the following scheme:

¹Some primitives (e.g., `"="` and `":="`) are recognised and implemented in assembler by the compiler.

²When profiling is enabled, the MLKit automatically appends the extension `Prof` for those functions that take regions (and thus a program point) as argument; see Section 19.4.

```
fun vid ( $x_1:\tau_1, \dots, x_n:\tau_n$ ) :  $\tau = \text{prim}(c\_func, (x_1, \dots, x_n))$ 
```

The result type τ must be of the form

```
 $\tau ::= \alpha \mid \text{int} \mid \text{bool} \mid \text{unit}$   

 $\mid \tau_1 * \dots * \tau_n \mid \tau \text{ list} \mid \text{real} \mid \text{string}$ 
```

If the result type is one of α , `int`, `bool`, or `unit` then the result value can be returned in a single register. Contrary, if the result type represents an allocated value, the C function must be told where to store the value. For any type that is either `real` or a non-empty tuple type, and does not occur in a list type of the result type τ , the MLKit allocates space for the value and passes a pointer to the allocated space as an additional argument to the C function. For any type representing an allocated value that is either `string` or occurs in a list type of the result type τ , the MLKit cannot statically determine the amount of space needed to store the value. Instead, regions are passed to the C function as additional arguments and the C function must then explicitly allocate space in these regions as needed, using a C function provided by the runtime system. The order in which these additional arguments are passed to the C function is determined by a preorder traversal of the result type τ . For a list type, regions are given in the order:

1. Region for auxiliary pairs.
2. Regions for elements (if necessary).

We now give an example to show what extra arguments are passed to a C function, given the result type. In the example, we use the following (optional) naming convention: names of arguments holding addresses of pre-allocated space in regions start with `vAddr`, while names of arguments holding addresses of region descriptors (to be used for allocation in a region) start with `rAddr`.

Example 1 Given the result type `(int * string) list * real`, the following extra arguments are passed to the C function (in order): `vAddrPair`, `rAddrLPairs`, `rAddrEPairs`, `rAddrEStrings` and `vAddrReal`, see Figure 19.1.

Here `vAddrPair` holds an address pointing to preallocated storage in which the tuple of the list and the (pointer to the) real should reside. The argument `rAddrLPairs` holds the region address for the auxiliary pairs of the list. Similarly, the arguments `rAddrEPairs` and `rAddrEStrings` hold region addresses




Figure 19.1: The order of pointers to allocated space and infinite regions is determined from a preorder traversal of the result type `(int * string) list * real`.

for element pairs and strings, respectively. The argument `vAddrReal` holds the address for preallocated storage for the real.

Additional arguments holding pointers to preallocated space and infinite regions are passed to the C function prior to the ML arguments. Consider again the ML declaration

```
fun vid (x1:τ1, ..., xn:τn) : τ = prim(c_func, (x1, ..., xn))
```

The C function `c_func` must then be declared as

```
int c_func (int addr1, ..., int addrm, int x1, ..., int xn)
```

where `addr1, ..., addrm` are pointers to preallocated space and infinite regions as described above.

19.2 Conversion Macros and Functions

The runtime system provides a small set of conversion macros and functions for use by C functions that need to convert between ML values and C values. Using these conversion macros and functions for converting between representations protects you against future changes in the representation of ML

values. The conversion macros and functions are declared in the header files:

```
src/Runtime/Tagging.h
src/Runtime/String.h
src/Runtime/List.h
```

19.2.1 Integers

There are two macros for converting between the ML representation of integers and the C representation of integers:³

```
#define convertIntToC(i)
#define convertIntToML(i)
```

To convert an ML integer `i_ml` to a C integer (type `long int`) `i_c`, write

```
i_c = convertIntToC(i_ml);
```

To convert a C integer (type `long int`) `i_c` to an ML integer `i_ml`, write

```
i_ml = convertIntToML(i_c);
```

The macros demonstrated here are used in the examples 2, 3, and 6 in Section 19.10.

19.2.2 Units

The following constant in the conversion library denotes the ML representation of `()`:

```
#define mlUNIT
```

19.2.3 Reals

An ML real is represented as a pointer into a region containing the real. To convert an ML real to a C real, we dereference the pointer. To convert a C real to an ML real, we update the memory to contain the C real. The following two macros are provided:

```
#define convertRealToC(mlReal)
#define convertRealToML(cReal, mlReal)
```

³These macros are the identity maps when garbage collection is disabled.

Converting an ML real `r_ml` to a C real `r_c` can be done with the first macro:

```
r_c = convertRealToC(r_ml);
```

Converting from a C real to an ML real (being part of the result value of the C function) is done in one or two steps depending on whether the real is part of a list or not. If the real is not in a list the memory containing the real has been allocated before the C call, see Section 19.1:

```
convertRealToML(r_c, r_ml);
```

If the ML real is part of a list element, then space must be allocated for the real before converting it. If `rAddr` identifies a region for the real, you write:

```
allocReal(rAddr, r_ml);
convertRealToML(r_c, r_ml);
```

These macros are used in the examples 3, 6 and 8 in Section 19.10.

19.2.4 Booleans

Four constants provide the values of true and false in ML and in C. These constants are defined by the following macros:⁴

```
#define mlTRUE 3
#define mlFALSE 1
#define cTRUE 1
#define cFALSE 0
```

Two macros are provided for converting booleans:

```
#define convertBoolToC(i)
#define convertBoolToML(i)
```

Converting booleans is similar to converting integers:

```
b_c = convertBoolToC(b_ml);
b_ml = convertBoolToML(b_c);
```

⁴For historical reasons, booleans in the MLKit are tagged even when garbage collection is disabled.

19.2.5 Records

Records are boxed. One macro is provided for storing and retrieving elements:

```
#define elemRecordML(recAddr, offset)
```

An element can be retrieved from a record `rec_ml` by writing

```
e_ml = elemRecordML(rec_ml, offset);
```

where the first element has `offset` 0. An element `e_ml` is stored in an ML record `rec_ml` by writing

```
elemRecordML(rec_ml, offset) = e_ml;
```

Two specialised versions of the `elemRecordML` macro are provided for pairs:

```
#define first(x)
#define second(x)
```

If the record is to be part of a list element then it is necessary to allocate the record before storing into it. This allocation is done with the macro

```
#define allocRecordML(rAddr, size, vAddr)
```

where `rAddr` denotes a region (i.e., a pointer to a region descriptor), `size` is the size of the record (i.e., the number of components), and `vAddr` is a variable in which `allocRecordML` returns a pointer to storage for the record. The record is then stored, component by component, by repeatedly calling `elemRecordML` with the pointer `vAddr` as argument.

The above macros are used in examples 8, 9 and 7 in Section 19.10.

19.2.6 Strings

Strings are boxed and always allocated in infinite regions. It is possible to print an ML string by using the C function

```
void printStringML(String str);
```

Strings are converted from ML to C and vice versa using the two C functions

```
void convertStringToC(String mlStr, char *cStr,
                    size_t cStrLen, int exn);
String convertStringToML(Region rAddr, char *cStr);
```

An ML string `str_ml` is converted to a C string `str_c` in already allocated storage of size `size` bytes by writing

```
convertStringToC(ctx, str_ml, str_c, size, exn);
```

where `exn` is some ML exception value (see Section 19.3) to be raised if the ML string has size greater than `size` and `ctx` is an evaluation context obtained using the `__get_ctx` primitive.

A C string is converted to an ML string in the region denoted by `rAddr` by writing

```
str_ml = convertStringToML(rAddr, str_c);
```

The following macro returns the size of an ML string:

```
size_t sizeStringDefine(String str);
```

These macros are used in the examples 7 and 5 in Section 19.10.

19.2.7 Lists

Lists are always allocated in infinite regions. A list uses, as a minimum, one region for the auxiliary pairs of the list, see Figure 5.1 on page 55.

We shall now show three examples of manipulating lists. The first example traverses a list. Consider the following C function template:

```
void traverse_list(uintptr_t* ls) {
    uintptr_t elemML;
    for ( ; isCONS(ls); ls=t1(ls)) {
        elemML = hd(ls);
        /*do something with the element*/
    }
    return;
}
```

The ML list is passed to the C function in parameter `ls`. The example uses a simple loop to traverse the list. The parameter `ls` points at the first constructor in the list. Each time we have a `CONS` constructor we also have an element, see Figure 5.1. The element can be retrieved with the `hd` macro. One retrieves the tail of the list by using the `t1` macro.

The following four macros are provided in the `src/Runtime/List.h` header file:

```

#define isNIL(x)
#define isCONS(x)
#define hd(x)
#define tl(x)

```

The next example explains how to construct a list backwards. Consider the following C function template:

```

uintptr_t mk_list_backwards(Region pairRho) {
    uintptr_t *resList, *pair;
    makeNIL(resList);
    while (/*more elements*/) {
        ml_elem = ...;
        allocRecordML(pairRho, 2, pair);
        first(pair) = (uintptr_t) ml_elem;
        second(pair) = (uintptr_t) resList;
        makeCONS(pair, resList);
    }
    return (uintptr_t)resList;
}

```

First, we create the NIL constructor, which marks the end of the list. Then, each time we have an element, we allocate a pair. We store the element in the first cell of the pair. A pointer to the list constructed so far is put in the second cell of the pair. (In this release of the MLKit, the `makeCONS` macro simply assigns its second argument the value of its first argument.) In the example, we have assumed that the elements are unboxed, thus, no regions are necessary for the elements.

The last example shows how a list can be constructed forwards. It is more clumsy to construct the list forwards because we have to return a pointer to the first element. Consider the following C function template.

```

uintptr_t mk_list_forwards(Region pairRho) {
    uintptr_t *pair, *cons, *temp_pair, res;

    /* The first element is special because we have to */
    /* return a pointer to it. */
    ml_elem = ...
    allocRecordML(pairRho, 2, pair);
    first(pair) = (uintptr_t) ml_elem;
    makeCONS(pair, cons);
    res = (uintptr_t) cons;
}

```

```

while (/*more elements*/) {
    ml_elem = ...
    allocRecordML(pairRho, 2, temp_pair);
    first(temp_pair) = (uintptr_t) ml_elem;
    makeCONS(temp_pair, cons);
    second(pair) = (uintptr_t) cons;
    pair = temp_pair;
}
makeNIL(cons);
second(pair) = (uintptr_t)cons;
return res;
}

```

We create the `CONS` constructor and pair for the first element and return a pointer to the `CONS` constructor (the pair) as the result. We then construct the rest of the list by constructing a `CONS` constructor and a pair for each element. It is necessary to use a temporary variable for the pair (`temp_pair`) because we have to update the pair for the previous element. The second component of the last pair contains the `NIL` constructor and thus denotes the end of the list.

The two macros `makeCONS` and `makeNIL` are provided in the `List.h` header file:

```

#define makeNIL(rAddr, ptr)
#define makeCONS(rAddr, pair, ptr)

```

19.3 Exceptions

C functions are allowed to raise exceptions and it is possible for the ML code to handle these exceptions. A C function cannot declare exceptions locally, however. For technical reasons, we must first acquire an evaluation context, which can be obtained using the special built-in primitive `__get_ctx`:

```

fun getCtx () : foreignptr =
    prim("__get_ctx", ())

```

Now, as an example, consider the following ML declaration:

```

exception Exn
fun raiseif0 (arg : int) : unit =

```

```
prim("raiseif0", (getCtx(), arg, Exn))
```

If we want the function `raiseif0` to raise the exception value `Exn` if the argument (`arg`) is 0 then we use the function `raise_exn` provided by the runtime system, by including the header file `src/Runtime/Exception.h`. The C function `raiseif0` may be defined thus:

```
void raiseif0(Context ctx, long i_ml, uintptr_t exn) {
    long i_c;
    i_c = convertIntToC(i_ml);
    if (i_c == 0) raise_exn(ctx, exn);
    return;
}
```

Notice that the supplied context is passed to the `raise_exn` function, which uses the context to gain access to the region stack and to the current exception handler.

Notice also that there is no need to make the function `raiseif0` return the value `m1UNIT`; in case the type of the return value is `unit` then the MLKit automatically inserts code for returning the ML value `()` after the call to the C function.

Exceptions are used in examples 6 and 7 in Section 19.10.

19.4 Program Points for Profiling

To support profiling, the programmer must provide special profiling versions of those C functions that allocate space in regions (i.e., that take regions as additional arguments). If profiling is enabled and at least one pointer to a region is passed to the C function then also a program point that represents the call to the C function is passed. The program point is used by the C function when allocating space in regions, as explained in Section 19.4. The program point is passed as the last argument:

```
uintptr_t c_funcProf (Region addr1, ..., Region addrm,
                    uintptr_t x1, ..., uintptr_t xn, long pPoint)
```

No special version of the C function is needed if it does not allocate into infinite regions; in this case, the same C function can be used both when profiling is enabled and disabled.

A program point passed to a C function is an integer; it identifies the allocation point that represents the C call in the program, see Chapter 17.

The runtime system provides special versions of various allocation macros and functions presented earlier in this chapter:

```
#define allocRealProf(realRho, realPtr, pPoint)
#define allocRecordMLProf(rhoRec, ssize, recAddr, pPoint)
String convertStringToMLProf(Region rhoString,
                             char *cStr,
                             long pPoint);
```

Here is the profiling version of the C function `mk_list_backwards`:

```
uintptr_t
mk_list_backwardsProf(Region pairRho, long pPoint) {
    uintptr_t *resList, *pair;
    makeNIL(resList);
    while (/*more elements*/) {
        ml_elem = ...;
        allocRecordMLProf(pairRho, 2, pair, pPoint);
        first(pair) = (uintptr_t) ml_elem;
        second(pair) = (uintptr_t) resList;
        makeCONS(pair, resList);
    }
    return (uintptr_t) resList;
}
```

The example shows that it is not difficult to make the profiling version of a C function; use the `Prof` versions of the macros and use the extra argument `pPoint`, appropriately. The same program point is used for all allocations in the C function, perceiving the C function as one entity.

19.5 Storage Modes

As described in Chapter 12 on page 103, actual region parameters contain a storage mode at runtime, if the region is infinite. A C function may check the storage mode of an infinite region to see whether it is possible to reset the region before allocating space in it. The header file `src/Runtime/Region.h` of the runtime system provides a macro `is_inf_and_atbot`, which can be used to test whether resetting is safe.

The C function `resetRegion`, which is also provided by the runtime system in the header file `src/Runtime/Region.h`, can be used to reset a region. Consider again the `mk_list_backwards` example. If the `atbot` bit of the region for the list is set, then this region can be reset prior to constructing the list:

```
uintptr_t mk_list_backwards(Region pairRho) {
    uintptr_t *resList, *pair;
    if (is_inf_and_atbot(pairRho)) resetRegion(pairRho);
    makeNIL(resList);
    ...
}
```

The C programmer can assume that if a parameter region aliases another parameter region or a region holding an argument, the parameter region is passed with storage mode `attop`.

19.6 Endomorphisms by Polymorphism

Until now, we have seen examples only of C functions that are region exomorphic, that is, functions that, in general, write their result into regions that are different from those in which the arguments reside.

A region endomorphic function has the property that the result of calling the function is stored in the same regions that hold the arguments to the function. Region endomorphic functions are useful when the result of the function shares with parts of the arguments. Consider the C function

```
uintptr_t select_second(uintptr_t pair) {
    return second(pair);
}
```

which selects the second component of `pair` (cast to an integer); the identifier `second` is defined in the header file `Tagging.h` by the macro definition

```
#define second(x) (*((uintptr_t *) (x)+1))
```

Now, for the MLKit to make correct, that is safe, decisions about when to deallocate regions, the endomorphic properties of a C function must be expressed in the region-annotated type scheme for value identifiers to which the C function is bound. The programmer can tell the MLKit about region

endomorphie behavior of a C function by using type variables. For example, here is an ML declaration that binds a value identifier `second` to the C function `select_second`:⁵

```
fun second(pair :  $\alpha * \beta$ ) :  $\beta$  =
  prim("select_second", pair)
```

MLKit associates the following region-annotated type scheme to the value identifier `second`:

$$\forall \alpha_1 \alpha_2 \rho \epsilon. (\alpha_1 * \alpha_2, \rho) \xrightarrow{\epsilon.\{\text{get}(\rho_3)\}} \alpha_2$$

Notice that the region-annotated type scheme expresses the region endomorphic behavior of the C function. That is, for any substitution S , mapping α_1 and α_2 to region types and places, the type $S((\alpha_1 * \alpha_2, \rho) \xrightarrow{\epsilon.\{\text{get}(\rho_3)\}} \alpha_2)$ will express that the result is located in the same regions as (the second projection) of the argument.

19.7 Compiling and Linking

To use a set of C functions in the ML code, one must first compile the C functions into an object file. (Remember to include appropriate header files.)

As an example, the file `kitdemo/libmylib.c` holds a set of example C functions. This file is compiled into an archive (in the form of a single object file) by typing (from the shell)

```
$ gcc -o libmylib.a -c libmylib.c
```

in the `kitdemo` directory. Now, to compile the file to work with profiling, type

```
$ gcc -DPROFILING -o libmylib-p.a -c libmylib.c
```

The MLB-file `mylib.mlb`, which is listed in Figure 19.2, mentions the file `mylib.sml`, which declares a series of ML functions to be used in the file `test_mylib.sml`.

Once the archives have been generated, the appropriate archive can be passed to the `mlkit` compiler, using the options `-libs` and `-libdirs`, as follows:

⁵MLB-file: `kitdemo/select_second.mlb`. The C file `select_second.c` must be compiled (using `gcc`) to form the object file (archive) `libselect_second.a` before the project can be compiled: `mlkit -no_gc -dirlibs "." -libs "m,dl,c,select_second" select_second.mlb`.

```
$(SML_LIB)/basis/basis.mlb
mylib.sml
test_mylib.sml
```

Figure 19.2: Linking with external object files is done by use of the `prim` primitive, which in this case is used in the file `mylib.sml` for declaring a series of ML functions.

```
$ mlkit -no_gc -o mylibtest -libdirs "." \
    -libs "m,c,dl,mylib" mylib.mlb
...
$ mlkit -no_gc -prof -o mylibtest-p -libdirs "." \
    -libs "m,c,dl,mylib-p" mylib.mlb
...
```

To learn more about the options `-libs` and `-libdirs`, type

```
$ mlkit --help
```

on the command line.

You may consult the file `kitdemo/Makefile` to see how one can further automate an appropriate build process.

19.8 Dynamic Linking

MLKit supports dynamic linking at runtime. This is done using the `dlopen` and `dlsym` functions from the MLKit library `basis/dynlink.mlb`. The function `dlopen` opens a given library and the function `dlsym` associates a name with a given function in the library. If the name is already linked, the exception `Fail` is raised.

Using the functions `dlopen` and `dlsym`, as shown in Figure 19.3, you can call a dynamically linked library function using a primitive call to `':`.

If `':` is called with a name that has no association, the exception `Match` is raised.

```

fun isNullFP(s : foreignptr) : bool = prim("__is_null", s)
val b = Dynlib.dlopen (SOME "libcrack.so", Dynlib.NOW, false)
val _ = Dynlib.dlsym ("testdyn", "FascistCheck", b)
fun fascistCheck a : string option =
  let val b : foreignptr =
    prim("@:", ("testdyn", a : string,
               "/usr/lib/cracklib_dict"))
  in if isNullFP b
     then NONE
     else SOME(prim ("fromCtoMLstring", b))
  end

```

Figure 19.3: Dynamic linking of the function `FascistCheck` from the library `libcrack.so`. The ML function `fascistCheck` calls `FascistCheck` with the argument `(a,/usr/lib/cracklib_dict)` and converts the resulting C string into an ML string. This example uses the auto conversion feature as described in the next section.

19.9 Auto Conversion

For C functions that are simple, in a sense that we shall soon define, the MLKit can generate code that automatically converts representations of arguments from ML to C and representations of results from C back to ML.

Auto conversion is enabled by prefixing a `@`-character to the name of the C function, as in the following example:

```

fun power_auto(base : int, n : int) : int =
  prim ("@power_auto", (base, n))

```

The power function may then be implemented in C as follows:

```

long power_auto(long base, long n) {
  long p;
  for (p = 1; n > 0; --n) p = p * base;
  return p;
}

```

No explicit conversion is needed in the C code. The example shown here is example 4 of Section 19.10; it is part of the `mylib.mlb` project.

We define the notion of *base type* as one of `int`, `Int8.int`, `Int16.int`,

`Int31.int`, `Int32.int`, `Int63.int`, `Int64.int`, `word`, `Word8.word`, `Word16.word`, `Word31.word`, `Word32.word`, `Word63.word`, `Word64.word`, `char`, or `bool`.

Auto conversion works when each of the argument types of the ML function is a base type or one of `string`, `CharArray.array`, `Word8Vector.vector`, or `Word8Array.array`. The return type of the ML function must be a base type or the type `unit`.

For the vector and array types, the data pointer is passed as argument to the C function. The C function may then access and mutate the content of the data but should not hold on to the data or store the pointer for later use.

Auto conversion works also when profiling is enabled.

19.10 Examples

Several example C functions are located in the file `kitdemo/libmylib.c`. The MLB-file `kitdemo/mylib.mlb`, which is listed in Figure 19.2, makes use of these functions.

The source file `mylib.sml`, which is part of the `mylib.mlb` project, contains the following ML declarations:

```

fun power(base: int, n: int) : int =
  prim ("power", (base, n))

fun power_auto(base: int, n: int) : int =
  prim ("@power_auto", (base, n))

fun power_real (base: real, n: int) : real =
  prim ("power_real", (base, n))

fun print_string_list (ss: string list) : unit =
  prim ("print_string_list", ss)

exception Power
fun power_exn (base: real, n: int) : real =
  prim ("power_exn", (getCtx(), base, n, Power))

exception DIR
fun dir (directory: string) : string list =
  prim ("dir", (getCtx(), directory, DIR))

```

```

fun real_list () : real list =
  prim ("real_list", ())

fun change_elem (p : int*string) : string*int =
  prim ("change_elem", p)

fun strrev (s:string) : unit =
  prim("@strrev_auto", (size s,s))

```

The C function implementations are summarised below (see the files `libmylib.c` and `mylib.sml` in the `kitdemo` directory for detailed comments.)

Example 2 The `power` function shows how to convert integers with the macros `convertIntToC` and `convertIntToML`.

Example 3 The `power_real` function shows how to convert reals with the macros `convertRealToC` and `convertRealToML`.

Example 4 The `power_auto` function shows the use of auto conversion, which allows for easy linking to certain C functions.

Example 5 The `print_string_list` example shows how to traverse a list of strings. The technique can easily be adopted to other data structures (e.g., to lists of lists of strings).

Example 6 The `power_exn` function shows how an exception can be raised from a C function, as explained in Section 19.3.

Example 7 The `dir` function shows how a list can be constructed backwards. We use the UNIX system calls `opendir` and `readdir` to read the contents of the specified directory.

Notice also that we check the infinite regions for resetting at the start of the C function. The checks should be placed at the start of the function, or else not inserted at all.

If you compare the C functions `dir` and `dirProf` you may see how the function `dirProf` is modified to work with profiling.

Example 8 Function `real_list` constructs a list of reals forwards. The reals are allocated in an infinite region. It may be more convenient to construct

the list backwards in the C function and then apply a list reverse function on the result list in the ML program.

Example 9 Function `change_elem` shows the use of the macro `elemRecordML`. The result type is `string*int`. The function swaps the two elements in the pair. MLKit passes an address to preallocated space for the result pair, and an infinite region for the result string.

At first thought it should be enough to just swap the two arguments, and not copy the string into the string region, that is, one could write the following function:

```
? uintptr_t
? change_elem(uintptr_t newPair, Region stringRho, uintptr_t pair)
  {
?   uintptr_t firstElem_ml, secondElem_ml;
?   firstElem_ml = elemRecordML(pair, 0);
?   secondElem_ml = elemRecordML(pair, 1);
?   elemRecordML(newPair, 0) = secondElem_ml;
?   elemRecordML(newPair, 1) = firstElem_ml;
?   return newPair;
? }

```

This function may work sometimes but it is not safe! Region inference expects the result string to be allocated in `stringRho`, and may therefore deallocate the region containing the argument string, `secondElem_ml`, while the string in the returned pair is still live. The safe version of `change_elem` is found in `libmylib.c`. See Section 19.6 for inspiration to how a safe non-copying swap function can be implemented.

Example 10 This example demonstrates the use of auto conversion with strings. Here is the C code for the function `strrev_auto`, which is referred to by the ML function `strrev`:

```
void
strrev_auto(long len, char* data) {
  long n = len / 2;
  for (long i = 0; i < n; i++) {
    char c = data[i];
    data[i] = data[len-i-1];
    data[len-i-1] = c;
  }
}
```

```
    return;  
}
```

The function reverses the string characters in-place, which is safe even when garbage-collection is enabled. Notice that the size of the string is passed explicitly to the C function as a separate argument. Auto conversion works both with and without profiling and with and without reference-tracing garbage collection.

Chapter 20

Summary of Changes

20.1 Changes Since Version 4.7.16

Here is an overview of the main changes to MLKit since version 4.7.16.

Modular Storage Mode Analysis

MLKit now implements a modular version of the storage mode analysis, which allows for external functions to be passed with regions that have storage mode `atbot`. The modular version of the storage mode analysis makes use of local aliasing information for inferring storage modes. See Section 12.3.

20.2 Changes Since Version 4.7.2

Here is an overview of the main changes to MLKit since version 4.7.2.

Double-Ended Bit-Stealing for Uboxing

MLKit now uses an advanced unboxing scheme that allows many algebraic data types to be implemented unboxed using the, otherwise non-used, 16 most-significant bits of a pointer or the high unused bits of 8-bit word values) [Els24a]. The scheme is described in details in Section 10.3.

Deep-Argument Flattening

The MLKit optimiser now attempts to flatten arguments to functions using a scheme that allows for deep flattening and uncurrying of arguments to functions (also across compilation unit boundaries) [Els25]. With this scheme, if a function takes a tuple of reals and a curried argument of type real as arguments, those arguments that are not used in their boxed form will be passed unboxed.

Support for Threads and Parallelism

MLKit now supports threads and parallelism through a fork-join thread interface [EH23]. The parallelism features are not described in details in this report, however.

Improved Abbreviated Pretty Printing

Intermediate forms in the MLKit (e.g., MulExp programs) now benefit from prettier printing. For example, primitives are now printed using qualified syntax (e.g., `String.concat`).

Explicit Programming with Regions and Effects

MLKit now supports explicit programming with regions and effects. The resulting language and compiler is called ReML and it features techniques for expressing and verifying constraints on regions and effects [Els24b]. ReML is not described in this report.

Improved C FFI Auto Conversion

MLKit now has improved support for interaction with C through an enriched set of supported types, as documented in Section 19.9.

20.3 Changes Since Version 4.6.1

This section provides an overview of the main changes to the MLKit since version 4.6.1.

Simplification of the Region Type System

To simplify the region type system, region variables are no longer associated with type variables and the concept of word regions has been eliminated entirely. Instead of associating type variables with region variables, type substitutions now map type variables to region type and places.

GC Safety Fixed

A problem with garbage-collection safety was fixed and it is now properly ensured that no dangling pointers appear during evaluation, which is a necessary requirement for combining region inference and reference-tracing garbage collection. Details are described in a conference paper [Els23] and a technical report [Els22].

20.4 Changes Since Version 4.3.0

This section provides an overview of the main changes to the MLKit since version 4.3.0.

X64 Backend

The x86 native backend has been replaced with an x64 native backend, which uses the GNU assembler to create native machine code on x86 machines. The new backend also features intra-procedural register allocation for floating-point values. MLKit now also features `Int64.int` as the default integer and `Word64.word` as the default word type (`Int63.int` and `Word63.word` when garbage collection is enabled).

Hosted at Github

MLKit is now hosted at Github. The repository is

`https://github.com/melsman/mlkit`

Bug reports should be filed by submitting a Github issue. Features and bug fixes can be submitted via Github pull requests. Comprehensive tests are executed using Github actions.

JavaScript Backend (SMLtoJs)

SMLtoJs is a stand-alone version of the MLKit that compiles to JavaScript and makes use of MLKit’s front-end and compilation infrastructure (e.g., re-compilation management). It supports large parts of the Standard ML Basis Library. SMLtoJs does not make use of region-based memory management. SMLtoJs has been used for hosting a Standard ML compiler in a web browser [Els11]. To try Standard ML in a browser, visit the site

<https://diku-dk.github.io/sml-ide/>

Generational Garbage Collection

The non-generational pointer-tracing garbage collection technique has been augmented with a generational version (option `-gengc`), which in some cases is superior to ordinary reference-tracing garbage collection, but which may also cause additional fragmentation [EH21b].

Improved Basis Library Coverage

The Standard ML Basis library coverage has been extended to support also the `Unix` structure, socket programming through the `Socket` structure, the `NetHostDb` structure, and the `INetSock` structure. With the move to the x64 architecture, the MLKit now also supports the structures `Word63`, `Word64`, `Int63`, and `Int64`.

20.5 Changes Since Version 4

This section provides an overview of the main changes to the MLKit since version 4, but before version 4.3.0.

Support for Compiling ML Basis Files

ML Basis Files allows for expressing source dependencies, exactly (as a directed acyclic graph). ML Basis Files thus provides a mechanism for programming “in the very large”.

File-based Separate Compilation

MLKit now supports file-based separate compilation, based on dependencies established from ML Basis Files. The compiler serialises symbol table information to disk for each compilation unit, so that this information can be deserialised and used when compiling other compilation units.

Updated Standard ML Basis Library

The MLKit implementation of the Standard ML Basis Library now conforms to the specification published in [GR04].

Untagged Pairs, Triples and References

MLKit now support untagged representations of heap-allocated pairs, triples, and Standard ML references, even when garbage collection is enabled.

20.6 Changes Since Version 3

This section provides an overview of the main changes to the MLKit since version 3, but before version 4.

Garbage Collection

MLKit supports reference tracing garbage collection in combination with the region memory model. Garbage collection is supported only in the native backend version of the MLKit. To enable garbage collection, pass the option `-gc` to the MLKit compiler. Garbage collection is also possible with region profiling enabled. See Chapter 16 for more information about garbage collection with the MLKit.

X86 Backend

The HPPA backend of the MLKit version 3.0 and earlier has been replaced with an x86 native backend, which uses the GNU assembler to create native machine code on x86 machines.

Bytecode Backend

For portability, the MLKit now provides a bytecode backend and a bytecode interpreter. Which backend is used by the MLKit compiler is determined when the MLKit itself is compiled, but it is possible to have both a native version and a bytecode version of the MLKit compiler installed on the same system.

Unboxing of Function Arguments

By default, the MLKit performs a simple local unboxing analysis to figure out if a function taking a tuple as argument can be transformed into a function taking multiple arguments. Only functions that use only the individual elements of the argument tuple undergo transformation. The optimisation can be disabled by passing the option `-no_unbox_function_arguments` to the MLKit compiler.

Removal of Region Vectors

In the MLKit version 3.0 and earlier, actual region parameters were passed to a region polymorphic function in a *region vector*, which itself was allocated in a region. In version 4.0, actual region parameters to region polymorphic functions are passed in registers and on the stack. This simplification improves pretty printing of region annotated terms and on what function calls turn into tail calls (see Section 14.3).

20.7 Changes Since Version 2

This section provides an overview of the main changes to the MLKit since version 2.0 but before version 3.0 of the MLKit.

Modules and Separate Compilation

The most important development since Version 2 is the ability to compile Modules and the discipline of separate compilation. A distinguished feature of the way modules are compiled is that module constructs do not give rise to any code, so there is no runtime overhead in using modules [Els99b, Els99a]. See Chapter 15.

Standard ML Basis Library

MLKit supports a large portion of the Standard ML Basis Library, based on the Moscow ML version of the library. To see exactly what parts of the Standard ML Basis Library are supported, consult the MLB-file `basis.mlb` located in the directory `basis`.

Scalability

MLKit now compiles fairly large programs, including Hafnium's AnnoDomini (58.000 lines of SML) and the MLKit itself (around 80.000 lines).

New Match Compiler

The pattern compiler has been rewritten, based on Sestoft's method [Ses96], which is also the basis of the Moscow ML match compiler.

New StatObject Module

MLKit contains a module, `StatObject`, which implements the semantic objects of the static semantics of the Core. Originally, this was a very clean and very inefficient implementation of the Definition. In version 2 of the MLKit, `StatObject` was replaced by an imperative and efficient, but complicated module. In version 3, `StatObject` uses a clean, efficient and imperative implementation of `StatObject`. This is particularly useful for those who want to reuse the front-end of the MLKit for other purposes.

Unboxed Representation of Lists

List constructors are now represented unboxed, that is, the least significant bits of a list value is used to distinguish between `nil` and a pointer to a pair (`::`) holding the head and the tail of the list. Thus, a list takes up only one region (for the auxiliary pairs) plus any regions for the elements of the list. Consult Chapter 5 for details.

Bibliography

- [BRTT93] Lars Birkedal, Nick Rothwell, Mads Tofte, and David N. Turner. The ML Kit (Version 1). Technical Report DIKU-report 93/14, Department of Computer Science, University of Copenhagen, 1993.
- [BTV96] Lars Birkedal, Mads Tofte, and Magnus Vejlstrup. From region inference to von Neumann machines via region representation inference. In *Proceedings of the 23rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 171–183. ACM Press, January 1996.
- [EH95] Martin Elsmann and Niels Hallenberg. An optimizing backend for the ML Kit using a stack of regions. Student Project 95-7-8, Department of Computer Science, University of Copenhagen (DIKU), July 5 1995.
- [EH02] Martin Elsmann and Niels Hallenberg. A region-based abstract machine for the ML Kit. Technical Report TR-2002-18, Royal Veterinary and Agricultural University of Denmark and IT University of Copenhagen, August 2002. IT University Technical Report Series.
- [EH03] Martin Elsmann and Niels Hallenberg. Web programming with SMLserver. In *Fifth International Symposium on Practical Aspects of Declarative Languages (PADL'03)*. Springer-Verlag, January 2003.
- [EH20] Martin Elsmann and Niels Hallenberg. On the effects of integrating region-based memory management and generational garbage collection in ML. In *Practical Aspects of Declarative Languages:*

- 22nd International Symposium, PADL 2020, New Orleans, LA, USA, January 20–21, 2020, Proceedings*, page 95–112, Berlin, Heidelberg, 2020. Springer-Verlag.
- [EH21a] Martin Elsmann and Niels Hallenberg. Integrating region memory management and tag-free generational garbage collection. *Journal of Functional Programming*, 31, 2021.
- [EH21b] Martin Elsmann and Niels Hallenberg. Integrating region memory management and tag-free generational garbage collection. *Journal of Functional Programming*, 31(E4), Februar 2021.
- [EH23] Martin Elsmann and Troels Henriksen. Parallelism in a region inference context. *Proc. ACM Program. Lang.*, 7(PLDI), June 2023.
- [EHV07] Martin Elsmann, Niels Hallenberg, and Carsten Varming. *SMLserver—A Functional Approach to Web Publishing (Second Edition)*, April 2007. (174 pages).
- [Els98] Martin Elsmann. Polymorphic equality—no tags required. In *Second International Workshop on Types in Compilation*, March 1998.
- [Els99a] Martin Elsmann. *Program Modules, Separate Compilation, and Intermodule Optimisation*. PhD thesis, Department of Computer Science, University of Copenhagen, January 1999.
- [Els99b] Martin Elsmann. Static interpretation of modules. In *Proceedings of the Fourth ACM SIGPLAN International Conference on Functional Programming*, September 1999.
- [Els03] Martin Elsmann. Garbage collection safety for region-based memory management. In *Proceedings of ACM SIGPLAN Workshop on Types in Language Design and Implementation (TLDI’03)*. ACM Press, January 2003.
- [Els11] Martin Elsmann. SMLtoJs: Hosting a Standard ML compiler in a web browser. In *Proceedings of the 1st ACM SIGPLAN International Workshop on Programming Language and Systems Technologies for Internet Clients, PLASTIC ’11*, page 39–48, New York, NY, USA, 2011. Association for Computing Machinery.

- [Els22] Martin Elsman. Garbage-collection safety for region-based type-polymorphic programs. Technical report, Department of Computer Science, University of Copenhagen (DIKU), Denmark, December 2022. DIKU Technical Report.
- [Els23] Martin Elsman. Garbage-collection safety for region-based type-polymorphic programs. *Proc. ACM Program. Lang.*, 7(PLDI), June 2023.
- [Els24a] Martin Elsman. Double-ended bit-stealing for algebraic data types. *Proc. ACM Program. Lang.*, 8(ICFP), August 2024.
- [Els24b] Martin Elsman. Explicit effects and effect constraints in ReML. *Proc. ACM Program. Lang.*, 8(POPL), January 2024.
- [Els25] Martin Elsman. Compositional deep argument flattening. In Informal Proceeding of the 2025 ML Family Workshop (ML ‘25), Singapore, October 2025.
- [GA96] Lal George and Andrew Appel. Iterated register allocation. *ACM Transactions on Programming Languages and Systems*, 18(3):300–324, May 1996.
- [GR04] Emden R. Gansner and John H. Reppy. *The Standard ML Basis Library*. Cambridge University Press, 2004.
- [Hal96] Niels Hallenberg. A Region Profiler for a Standard ML compiler based on Region Inference. Student Project 96-5-7, Department of Computer Science, University of Copenhagen (DIKU), June 14 1996.
- [Hal99] Niels Hallenberg. Combining garbage collection and region inference in the MLKit. Master’s thesis, Department of Computer Science, University of Copenhagen, 1999.
- [HET02] Niels Hallenberg, Martin Elsman, and Mads Tofte. Combining region inference and garbage collection. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI’02)*. ACM Press, June 2002. Berlin, Germany.

- [KO96] Martin Koch and Tommy Højfeldt Olesen. Compiling a higher-order call-by-value functional programming language to a RISC using a stack of regions. Master's thesis, Department of Computer Science, University of Copenhagen, October 1996.
- [MTHM97] Robin Milner, Mads Tofte, Robert Harper, and David MacQueen. *The Definition of Standard ML (Revised)*. MIT Press, 1997.
- [Ses96] Peter Sestoft. *Partial Evaluation*, volume 1110, chapter ML pattern match compilation and partial evaluation, pages 446–464. Springer-Verlag, February 1996.
- [TB98] Mads Tofte and Lars Birkedal. A region inference algorithm. *Transactions on Programming Languages and Systems (TOPLAS)*, 20(4):734–767, July 1998.
- [TBEH04] Mads Tofte, Lars Birkedal, Martin Elsmann, and Niels Hallenberg. A retrospective on region-based memory management. *Higher-Order and Symbolic Computation (HOSC)*, 17(3):245–265, September 2004.
- [TT93] Mads Tofte and Jean-Pierre Talpin. A theory of stack allocation in polymorphically typed languages. Technical Report DIKU-report 93/15, Department of Computer Science, University of Copenhagen, 1993.
- [TT94] Mads Tofte and Jean-Pierre Talpin. Implementing the call-by-value lambda-calculus using a stack of regions. In *Proceedings of the 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 188–201. ACM Press, January 1994.
- [TT97] Mads Tofte and Jean-Pierre Talpin. Region-based memory management. *Information and Computation*, 132(2):109–176, 1997.

Appendix A

Command-Line Options

This appendix shows the output of executing `mlkit -help`, where `mlkit` is the version of the MLKit compiler that uses the x64 native backend.

```
MLKit v4.7.16 ( - ) [X64 Backend]
```

```
Usage: mlkit [OPTION]... [file.sml | file.sig | file.mlb]
```

Options:

```
--version, -v, -V
```

```
    Print version information and exit.
```

```
--help
```

```
    Print extended help information and exit.
```

```
--help S
```

```
    Print help information about an option and exit.
```

```
--man
```

```
    Print man-page and exit.
```

```
--SML_LIB S                (You_did_not_set_path_to_install_dir)
```

```
    Installation directory for the MLKit standard library.
```

```
    For normal execution you should not modify this value.
```

```
    However, if you wish to use the MLKit with an altered
```

```
    runtime system you can update this setting and the
```

```
    system will try to link to a runtime system found in
```

the specified install directory.

- `--aggressive_opt, -aopt` (on)
 Enable aggressive optimisations, including aggressive inlining. These optimisations are not guaranteed to be region safe. Turning off garbage collection automatically turns off this option.
- `--alloc_protect_always` (off)
 Always protect allocation when parallelism is enabled. That is, disregard the result of protection inference.
- `--argobots, -argobots` (off)
 When enabled, executables link with the Argobots lightweight thread library.
- `--assembler S, -as S` (as -q)
 This option specifies the assembler used. On Linux the default is `as --64`. On macOS, the default is `as -q`.
- `--chat, -verbose` (off)
 Print a message for each compilation step in the compiler.
- `--comments_in_asmcode` (off)
 Insert comments in assembler code.
- `--compile_only, -c` (off)
 Compile only. Suppresses generation of executable
- `--compiler_timings, -timings` (off)
 Show compiler timings for each compilation phase.
- `--constant_folding, -cfold` (on)
 Enable constant folding optimisations.
- `--contract` (on)
 Contract is responsible for inlining, specialization, elimination of dead code, and much else (Lambda Expression Optimiser).

- `--contract_regions, -cr` (off)
 When this option is enabled, identically typed regions bound by the same `letregion` construct are unified. Moreover, region parameters to non-exported functions are trimmed whenever possible.
- `--cross_module_opt, -cross_opt` (on)
 Enable cross-module optimisation including inlining of small functions and specialisation of small recursive functions. Which optimisations are performed across modules is controlled by individual optimisation flags.
- `--dangling_pointers, -dangle` (off)
 When this option is disabled, dangling pointers are avoided by forcing values captured in closures to live at-least as long as the closure itself. So as to make garbage collection sound, this option is disabled by default when garbage collection is enabled.
- `--dangling_pointers_statistics` (off)
 When enabled, the compiler prints statistics about the number of times strengthening of the region typing rules (to avoid dangling pointers during evaluation) effects the target program. This flag is useful only when the flag `-gc` or `-no_dangle` is enabled.
- `--debug_compiler, -debug` (off)
 Print intermediate forms of a program during compilation.
- `--debug_constraint_solving, -dcs` (off)
 Debug ReML constraint solving.
- `--debug_linking` (off)
 Debug linking of target code by showing which object files are linked together.

- `--debug_man_enrich` (off)
 Show information about why a program unit need be recompiled. A program unit (or a functor body) is recompiled if either (a) the program unit is modified, or (b) information about an identifier for which the program unit depends upon has changed.
- `--debug_parallelism, -Dpar` (off)
 Debug parallelism, including protection inference.
- `--debug_which_at` (off)
 Debug storage mode analysis.
- `--delete_target_files` (on)
 Delete assembler files produced by the compiler. If you disable this flag, you can inspect the assembler code produced by the compiler.
- `--disable_atbot_analysis` (off)
 Disable storage mode analysis. That is, turn all allocation directives into `attop`.
- `--disable_flow_var` (off)
 Disable optimised compilation of control-flow code, such as conditional expressions.
- `--disable_spurious_type_variables` (off)
 Disable inference of spurious type variables. This option may crash the reference tracing garbage collector as it may result in dangling pointers. This flag is relevant only when garbage collection is enabled.
- `--eliminate_explicit_records` (on)
 Eliminate bindings of explicit records only used for selections. Transform
`let r = (e1,...,en) in ... #i r .. #j r ...`
 into
`let x1=e1 in ... let xn=en in ... xi .. xj ...`
 (Lambda Expression Optimiser).

- `--extra_gc_checks` (off)
Insert check for GC even in functions that do not allocate.
- `--fix_floating` (on)
Float fix-bindings into immediate let-bindings when possible to help tail-calls not be captured in let-region bindings.
- `--garbage_collection, -gc` (on)
Enable garbage collection. When enabled, regions are garbage collected during execution of the program. When garbage collection is enabled, all values are tagged. Due to region inference, for most programs, the garbage collector is invoked less often than for systems based only on garbage collection. When garbage collection is enabled, introduction of dangling pointers are avoided by forcing values captured in closures to live at-least as long as the closure. Moreover, enabling garbage collection implicitly enables the preservation of tail calls (see the option ‘‘preserve_tail_calls’’.)
- `--gdb_support, -g` (off)
When enabled, the compiler passes the option `--g` to ‘as’ (The GNU Assembler) and preserves the generated assembler files (.s files). Passing the `--g` option to ‘as’ makes it possible to step through the generated program using gdb (The GNU Debugger).
- `--generational_garbage_collection, -gengc` (off)
Enable generational garbage collection. Same as option `garbage_collection` except that two generations are used for each region.
- `--high_pointer_tagging` (on)
When enabled, the 16 most-significant bits in pointers may be used for tagging.
- `--import_basislib, -basislib` (on)
Import Basis Library automatically in your projects. If

you wish to make use of the Standard ML Basis Library in your projects, this option should be turned on, unless you wish to import the Basis Library manually in your projects.

`--inline_functors` (off)

Inline functors during static interpretation instead of generating separate target code blocks for functor bodies and arguments. With the flag enabled, performance may increase with the cost of larger (re)compilation times. The flag may be controlled in mlb-files using mlb-annotations.

`--inline_names S`

Names of functions that should always be inlined if possible, no matter the setting of the flag `--maximum_inline_size`.

`--libdirs S`

This option controls where ld looks for archives. The format is a comma-separated list of directories; see the `-libs` entry. The default is the empty list; thus 'ld' will look for libraries in only the system specific default directories. The directories are passed to 'ld' using the `-L` option.

`--libs S` (m,c,dl)

For accessing a foreign function residing in an archive named libNAME.a from Standard ML code (using prim), you need to add 'NAME' to this comma-separated list. Notice that an object file (with extension '.o') is an archive if it is renamed to have extension '.a'. You may need to use the `-libdirs` option for specifying directories for which ld should look for library archives. The libraries are passed to 'ld' using the `-l` option.

`--link_code S`

Link-files to be linked together to form an executable.

`--link_exe S, -ldexe S (gcc -Wl,-ld_classic,-stack_size,0x10000000)`

This option specifies the command used for linking an executable. The standard is to use 'gcc' for linking. When linking with c++ libraries, 'g++' is the linker you want. On Linux the default is 'gcc', whereas on newer macOS systems (Darwin > 23.1), the default is 'gcc -Wl,-ld_classic,-stack_size,0x10000000' and on older macOS systems, the default is 'gcc -Wl,-stack_size,0x10000000'.

`--link_shared S, -ldshared S (gcc)`

This option specifies the command used for linking a shared object file. The standard is to use 'gcc' for linking.

`--link_time_dead_code_elimination, -ltdce (on)`

Link time dead code elimination.

`--load_basis_files S, -load S`

Basis files to be loaded before compilation proper.

`--log_to_file (off)`

Log to files instead of stdout.

`--maximum_inline_size N, -max_inl_sz N (200)`

Functions smaller than this size (counted in abstract syntax tree nodes) are inlined, even if they are used more than once. Functions that are used only once are always inlined. A function declared in one program unit may be inlined when applied in another program unit provided the function is sufficiently closed (refers only to exported identifiers) and provided the option `--cross_opt` is enabled.

`--maximum_specialise_size N, -max_spec_sz N (200)`

Function-parameterised functions smaller than this size

(counted in abstract syntax tree nodes) are specialised if all applications of the function within its own body are applied to its formal function argument, even if they are used more than once. Functions that are used only once are specialised no matter their size. A function declared in one program unit may be specialised when applied in another program unit provided the function is sufficiently closed (refers only to exported identifiers) and provided the option `--cross_opt` is enabled. See also the option `--specialize_recursive_functions`.

- `--messages` (on)
Print messages about reading source files and generating target files.
- `--minimize_fixs` (on)
Minimize fix constructs (Lambda Expression Optimiser).
- `--mlb-subdir S`
For ensuring that the smart recompilation scheme is not reusing target-code compiled with different settings, a string provided with the `mlb-subdir` option can ensure the use of consistently generated code. This option is Useful, in particular, when performing benchmarking.
- `--mlb_path_maps S, -mlb-path-map S`
ML Basis path map files to be used.
- `--namebase S` (dummyBase)
Name base to enforce unique names when compiling `mlb-files`.
- `--no_aggressive_opt, -no_aopt`
Opposite of `--aggressive_opt, -aopt`.
- `--no_constant_folding, -no_cfold`
Opposite of `--constant_folding, -cfold`.
- `--no_contract`

Opposite of `--contract`.

`--no_cross_module_opt`, `-no_cross_opt`
Opposite of `--cross_module_opt`, `-cross_opt`.

`--no_dangling_pointers`, `-no_dangle`
Opposite of `--dangling_pointers`, `-dangle`.

`--no_delete_target_files`
Opposite of `--delete_target_files`.

`--no_eliminate_explicit_records`
Opposite of `--eliminate_explicit_records`.

`--no_fix_floating`
Opposite of `--fix_floating`.

`--no_garbage_collection`, `-no_gc`
Opposite of `--garbage_collection`, `-gc`.

`--no_generational_garbage_collection`, `-no_gengc`
Opposite of `--generational_garbage_collection`, `-gengc`.

`--no_high_pointer_tagging`
Opposite of `--high_pointer_tagging`.

`--no_import_basislib`, `-no_basislib`
Opposite of `--import_basislib`, `-basislib`.

`--no_link_time_dead_code_elimination`, `-no_ltdce`
Opposite of `--link_time_dead_code_elimination`, `-ltdce`.

`--no_messages`
Opposite of `--messages`.

`--no_minimize_fixs`
Opposite of `--minimize_fixs`.

`--no_optimiser`, `-no_opt`
Opposite of `--optimiser`, `-opt`.

`--no_preserve_tail_calls, -no_ptc`
Opposite of `--preserve_tail_calls, -ptc`.

`--no_print_control_abbrev_layout, -no_abbrev`
Opposite of `--print_control_abbrev_layout, -abbrev`.

`--no_print_regions, -no_Pregions`
Opposite of `--print_regions, -Pregions`.

`--no_raggedRight`
Opposite of `--raggedRight`.

`--no_region_inference, -no_ri`
Opposite of `--region_inference, -ri`.

`--no_register_allocation`
Opposite of `--register_allocation`.

`--no_repository, -no_rep`
Opposite of `--repository, -rep`.

`--no_specialize_recursive_functions`
Opposite of `--specialize_recursive_functions`.

`--no_type_check_lambda`
Opposite of `--type_check_lambda`.

`--no_unbox_funargs`
Opposite of `--unbox_funargs`.

`--no_unbox_real_funargs`
Opposite of `--unbox_real_funargs`.

`--no_unbox_reals`
Opposite of `--unbox_reals`.

`--no_uncurrying, -no_uncurry`
Opposite of `--uncurrying, -uncurry`.

- `--objs` (off)
 If enabled, MLKit writes object-file paths to the file `run` or the file specified by `-output`. The path to the runtime system (archive file) is included. The option is best used together with the option `-no_delete_target_files`.
- `--optimiser, -opt` (on)
 Enable optimisation of intermediate language code (Lambda Expressions). Which optimisations are performed is controlled by individual flags. The optimisations include function inlining, function specialisation, fix-minimization, unboxing of function arguments, and elimination of unnecessary record constructions.
- `--output S, -o S` (run)
 The name of the executable file generated by the Kit.
- `--parallel_compilation N, -j N` (1)
 The maximum number of parallel processes used for compilation.
- `--parallelism, -par` (off)
 When enabled, the runtime system supports parallel threads.
- `--parallelism_alloc_unprotected, -par0` (off)
 When enabled, allocation into a region is not guaranteed to be atomic.
- `--preserve_tail_calls, -ptc` (on)
 Avoid the wrapping of `letregion` constructs around tail calls. Turning on garbage collection automatically turns on this option.
- `--pretty_depth N` (5)
 This flag controls the pretty-printing depth of values printed in the REPL. The value must be an integer larger than zero.

- `--pretty_string_size N` (80)
 This flag controls the pretty-printing size of strings printed in the REPL. The value must be an integer larger than zero.
- `--print_K_normal_forms` (off)
 Print Region Expressions in K-Normal Form. Applicable, only after storage mode analysis has been applied.
- `--print_all_program_points, -Ppp` (off)
 Print all program points when printing physical size inference expressions.
- `--print_bit_vectors` (off)
- `--print_calc_offset_program` (off)
- `--print_call_explicit_expression, -Pcee` (off)
 Print Region Expression with call annotations.
- `--print_clos_conv_program, -Pccp` (off)
 Print Region Expression after closure conversion.
- `--print_closed_export_bases, -Pceb` (off)
 Controls printing of closed export bases.
- `--print_constraints` (off)
 Print ReML effect constraints when printing region and effect variables.
- `--print_control_abbrev_layout, -abbrev` (on)
 Abbreviate layout of multiplicity expressions and call-explicit expressions. For instance, do not show at-annotations for top-level functions, do not show at-annotations for immediate constants, do not show region-bindings for zero-size regions that are associated only with immediate constants, do not show empty formal and actual region parameter lists, do not show 'funcall' and 'fncall' annotations, do not show 'id' casts on base values,

do not show unique id for explicit region variables or explicit effect variables.

- `--print_control_prefix_infix` (off)
Prefix infix operations with module identifier to further indicate the type of the operation.
- `--print_drop_regions_expression, -Pdre` (off)
Print Region Expression after dropping word regions and region arguments with only get-effects.
- `--print_drop_regions_expression_with_storage_modes, -Pdresm` (off)
Print Region Expression after dropping word regions and region arguments with only get-effects. Also print atbot and attop annotations resulting from storage mode analysis.
- `--print_effects, -Peffects` (off)
Print effects in region types.
- `--print_excon_name` (off)
Print underlying unique name when printing excons.
- `--print_export_bases, -Peb` (off)
Controls printing of export bases.
- `--print_fetch_and_flush_program` (off)
Print program with instructions for activation record fetching and flushing.
- `--print_linearised_program` (off)
Print a linearised representation of the program unit.
- `--print_lvar_name` (off)
Print underlying unique name when printing lvars.
- `--print_normalized_program` (off)
Print Region Expression after K-normalisation.

- `--print_opt_lambda_expression, -Pole` (off)
Print Lambda Expression after optimisation.
- `--print_physical_size_inference_expression, -Ppse` (off)
Print Region Expression after physical size inference.
- `--print_post_elab_ast, -Ppeast` (off)
Print ast after elaboration.
- `--print_region_flow_graph, -Prfg` (off)
Print a region flow graph for the program fragment and generate a .vcg-file, which can be viewed using the xvcg program.
- `--print_region_inferred_program, -Prip` (off)
Print region-inferred program.
- `--print_region_spreaded_program, -Prsp` (off)
Print region-spreaded program.
- `--print_region_static_env0, -Prse0` (off)
Print imported region static environment prior to region inference.
- `--print_regions, -Pregions` (on)
Print region variables in types and expressions.
- `--print_register_allocated_program` (off)
- `--print_rho_levels` (off)
Print levels of region and effect variables in types and intermediate forms. Levels control quantification of region and effect variables.
- `--print_rho_protection, -Prho_protection` (off)
Print protectedness of region variables if set (P or U).
- `--print_rho_types` (off)
Print region types of region variables in types and intermediate forms. Possible region types are:

p Type of regions containing pairs.
 a Type of regions containing arrays.
 r Type of regions containing references.
 t Type of regions containing triples.
 s Type of regions containing strings.
 T Type of regions containing other than the above
 kinds of values.

--print_simplified_program (off)
 Print simplified program after register
 allocation.

--print_storage_mode_expression, -Psme (off)
 Print Region Expression after storage mode analysis

--print_type_name_stamps, -Ptypestamps (off)
 Print type name stamps and their attributes in types
 and expressions.

--print_types, -Ptypes (off)
 Print types when printing intermediate forms. For Lambda
 Expressions, ordinary ML types are printed, whereas for
 Region Expressions, region types are printed.

--quotation, -quot (off)
 Enable support for quotations and anti-quotations.
 When enabled, the datatype


```

      datatype  $\alpha$  frag = QUOTE of string
                    | ANTIQUOTE  $\alpha$ 
    
```

 is available in the initial environment. Moreover,
 values of this datatype may be constructed using
 the quotation/antiquotation syntax:


```

      val s = "world"
      val a : string frag list = 'hello ^s - goodbye'
    
```

--raggedRight (on)
 Use ragged right margin in pretty-printing of
 expressions and types.

--recompile_basislib, -scratch (off)

Recompile basis library from scratch. This option is useful together with other options that control code generation.

- `--region_inference, -ri` (on)
With this flag disabled, all values are allocated in global regions.
- `--region_profiling, -prof` (off)
Enable region profiling. Object code stemming from compiling a program with region profiling enabled is instrumented with profiling information. When a program compiled with region profiling enabled is run, the program produces a profile file `run.rp`, which can then be read by the profiling tool `rp2ps` that comes with the MLKit to produce profiling graphs of various forms.
- `--regionvar N` (~1)
Uses the provided number as the id of the first generated region variable. When this option is provided together with the `-c` option, a file `f.rv` is written in the `MLB/` directory with two numbers in it: the id for the first region variable generated and the id for the last region variable generated. The number given must be greater than any id for a top-level region/effect variable (>9).
- `--register_allocation` (on)
Perform register allocation. Without register allocation enabled, programs run somewhat slower--but they run and you save about 15 percent on compile time.
- `--reml` (off)
ReML is Standard ML with support for programming with explicit regions, explicit effects, and effect constraints. With ReML, atomic effects also include mutation effects. Whereas ReML includes parallel thread support, currently, ReML does not support integration with reference-tracing garbage collection.

- `--report_boxities` (off)
Report for every datatype declaration the inferred boxity (representation) of its value constructors.
- `--report_file_sig, -sig` (off)
Report signatures for each file read.
- `--repository, -rep` (on)
Use in-memory repository to avoid unnecessary recompilation. This flag should be disabled when compiling mlb-files, which make use of the file system as a repository.
- `--safeLinkTimeElimination` (off)
Treat this module as a library in the sense that the code can be eliminated if it is not used, even if the code has side effects.
- `--specialize_recursive_functions` (on)
Specialise recursive functions. Use the option `maximum_specialise_size` to control which functions are specialised. If this flag is on, functions that are applied only once are specialised, no matter the setting of `--maximum_specialise_size` (Lambda Expression Optimiser).
- `--statistics_after_optimisation, -stats_opt` (off)
Report optimisation statistics after optimisation of Lambda Expression.
- `--statistics_spurious, -stats_spurious` (off)
Report statistics on spurious functions and instantiations of spurious type variable. This flag is relevant only when garbage collection is enabled.
- `--strip` (off)
If enabled, MLKit strips the generated executable.
- `--tag_pairs` (off)
Use a tagged representation of pairs for garbage collection. Garbage collection works fine with a

tag-free representation of pairs, so this option is here for measurement purposes.

- `--tag_values, -tag` (on)
 Enable tagging of values as used when garbage collection is enabled for implementing pointer traversal.
- `--type_check_lambda` (on)
 Type check lambda expression prior to performing region inference. Type checking is very fast and for normal use you should not disable this option. Type checking intermediate forms is very powerful for eliminating bugs in the compiler.
- `--unbox_funargs` (on)
 Unbox arguments to fix-bound functions, for which the argument α' is used only in contexts '#i a'. All call sites are transformed to match the new function (Lambda Expression Optimiser).
- `--unbox_real_funargs` (on)
 Unbox arguments of type real to fix-bound functions, for which the argument α' is used only in contexts that unboxes α' . All call sites are transformed to match the new function. This optimisation has effect only when the flags '-unbox_funargs' and '-unbox_reals' are enabled (Lambda Expression Optimiser).
- `--unbox_reals` (on)
 Unbox real values and computations on real values inside functions. Real values stored in data structures and passed to functions are still boxed.
- `--uncurrying, -uncurry` (on)
 Enable uncurrying of curried functions. The uncurried function takes its arguments unboxed in registers or on the stack. For partial applications and non-application uses of the function, appropriate eta-expansions are applied.

- `--values_64bit` (on)
Support 64-bit values. Should be enabled for backends supporting 64-bit integers and words.
- `--warn_on_escaping_puts` (off)
Enable the compiler to issue a warning whenever a region type scheme contains a put effect on a region that is not quantified.
- `--warn_on_parallel_puts` (off)
Enable the compiler to issue a warning whenever a par-construct is passed functions with intersecting put effects.
- `--warn_spurious` (off)
Warn on the presence of a spurious type variable. This flag is relevant only when garbage collection is enabled.
- `--width N, -w N` (100)
Column width used when pretty printing intermediate code.

Index

- !, 81
- μ , *see* type and place
- *, 50
- +, 50
- , 50
- .mlb, 147
- /, 50
- ::, 53
- :=, 81
- ;, 72
- <, 50
- <=, 50
- <>, 50
- =, 50–52
- >, 50
- >=, 50
- ~, 51
- MLB-file, 147
- ~, 50

- abs, 50
- alignment, 50
- allocation point, 102
- allocReal, 194
- allocRealProf, 200
- allocRecordML, 195
- allocRecordMLProf, 200
- application extrusion, 76
- arguments
 - flattening, 61
 - multiple, 37, 47, 61, 214

- arity, 87
- arrow effect, 56, 126
 - auxiliary, 88
- at, 43, 51, 54, 102
- atbot, 102
- attop, 102, 106
- auto conversion, 11, 204
- auxiliary pairs, 55

- backend
 - bytecode, 13, 35, 214
 - hppa, 213
 - native, 13, 35
 - x64, 211, 221
 - x86, 211, 213
- Basis Library, 49, 213
- batch compilation, 38
- bit stealing, 11, 91
- block structure, 72
- bottom of region, 102
- boxing, 36, 45, 82
- boxity, 92
- Br, 87
- bug report, 211
- bytecode, 13

- C, 19
 - calling, 38, 189
- C examples, 205
- call conversion, 138, 187
- call-back function, 190

- ceil, 50
- cFALSE, 194
- change_elem, 207
- changes
 - since version 2, 214
 - since version 3, 213
 - since version 4, 212
 - since version 4.3.0, 211
 - since version 4.6.1, 210
 - since version 4.7.16, 209
 - since version 4.7.2, 209
- chr, 51
- comment option, 180
- comments
 - in MLB-file, 149
- concat, 51
- convertBoolToC, 194
- convertBoolToML, 194
- convertIntToC, 193
- convertIntToML, 193
- convertRealToC, 193
- convertRealToML, 193
- convertStringToC, 195
- convertStringToML, 195
- convertStringToMLProf, 200
- cp, 66
- cTRUE, 194
- datatype, 87
- declaration
 - local, 72
 - sequential, 71
 - value, 71
- decon, 54
- deep argument flattening, 11, 61
- dir, 206
- div, 50
- dlopen
 - dlopen, 203
- dlsym
 - dlsym, 203
- double copying, 28
- dynamic linking, 203
- effect, 43, 46
 - atomic, 46
 - atomic, definition, 129
 - definition, 129
 - latent, 126
- effect arity, 88
- effect variable, 56, 126
 - bound, 61
- elemRecordML, 195
- endomorphism, *see* region endomorphism
- enum-unboxed, 91
- environment, 71
- eps option, 120
- equality
 - monomorphic, 47
 - polymorphic, 47
- evaluation context, 198
- example programs, *see* kitdemo directory
- exception, 95
 - generative, 95
 - handling, 97
 - raising, 96
- exception, 95
- exception constructor, 95
- exception declaration, 95
- exception name, 95
- exception value, 96
 - constructed, 96
 - nullary, 96
- exn, 96

- exomorphism, *see* region exomorphism
- explicit effects, 11
- explicit regions, 11
- explode, 51
- expression
 - call-explicit, 187
- file option, 177
- first, 195
- floor, 50
- fn, 125
- fnjmp, 141
- foldl, 143
- forceResetting, 22, 101, 115
- frame, 44
- free, 20
- free list, 33
- fromto, 61
- fun, 126
- funcall, 140
- function, 59
 - Curried, 107, 126
 - first-order, 59
 - higher-order, 125
 - region polymorphic, 214
- function arguments
 - multiple, 37, 47, 61, 214
- function call
 - call-explicit, 138
- function type
 - region-annotated, 126
- functor, 153
- garbage collection, 10, 21, 213
 - generational, 212
- get, 46, 64
- Github
 - issue, 211
- hd, 98, 196
- heap, 19, 22
- hello world, 40
 - help option, 177, 183
 - help option to mlkit, 185, 221
- high-unboxed, 91
- implode, 51
- initial basis, 49–52
- Int31 structure, 49
- Int32 structure, 49
- Int63 structure, 49
- Int64 structure, 49
- integer, 49
- is_inf_and_atbot, 200
- isCONS, 196
- isNIL, 196
- iterator, 114
- JavaScript
 - SMLtoJs, 212
- JavaScript backend, 9
- K-normalisation, 105
- kitdemo directory, 40
- Lambda, 37, 43, 185
- lambda abstraction, 125, 126
- Lambda optimiser, 37
- layout control, 187
- length of list, 114
- let, 72
- let floating, 75
- let region, 35, 45, 47, 73, 97, 106
- Lf, 87
- libmylib.c, 205
- Life
 - repository, 211

- game of, 25
- life, 174
- lifetime, 72–75
 - shortening, 75
- list, 53, 215
 - auxiliary pairs, 55
 - region-annotated type, 56
 - tail, 55
- live variable analysis, *see* variable
- local, 72, 109
- low-unboxed, 91

- makeCONS, 198
- makeNIL, 198
- malloc, 20
- matching, 150
- merge sort, 66, 115
- microsec option, 176
- mk_list_backwards, 197
- mk_list_forwards, 197
- ML Basis File, 147
- ML Basis Files, 12, 212
- MLB-file, 38
 - comments in, 149
 - grammar, 147
- MLB-files, 149
- mlFALSE, 194
- MLKit
 - Version 1, 13
 - Version 2, 13
 - Version 3, 13
 - Version 4.7.22, 10
- mlkit
 - executable, 38
- mlTRUE, 194
- mlUNIT, 193
- mod, 50
- msort, 66, 115

- MulExp, 37, 185
- multiple function arguments, 37, 47, 61, 214
- multiplicity, 34
- multiplicity analysis, 37, 46
- mylib.sml, 205

- name option, 180
- nil, 53
- not, 52
- nthgen, 27

- o, 142
- object option, 180
- object profile, 165, 180
- open declaration, 152
- openIn, 80
- openOut, 80
- optimisation
 - statistics, 185
- optimiser, 78
- ord, 51

- pair
 - auxiliary, 54, 88, 102
- parallelism, 11
- path
 - absolute, 149
 - relative, 149
- pattern matching, 54
- physical size inference, 187
- power, 206
- power_auto, 206
- power_exn, 206
- power_real, 206
- prim, 190
- print_all_program_points, 174
- print_call_explicit_expression, 138, 174

- Pdre, 185
- Pdresm, 103
- print_effects, 187
- Ppse, 174, 187
- print_region_flow_graph, 174
- print_regions, 187
- Psme, 185
- print_types, 187
- print_string_list, 206
- Printing of intermediate forms, 185
- printStringML, 195
- profile
 - object, 165
 - region, 165
 - stack, 166
- profile data file, 179
- profile strategy
 - compile-time, 174
 - options, 176
 - runtime, 176
- profile tick, 176
- profiletime option, 176
- profiling
 - time slot, 176
- program point, 165
- program transformation, 75
- projects
 - compiling, 66
 - running, 66
- put**, 46, 63
- put-effect
 - escaping, 78
- rDesc, *see* region descriptor
- real, 50
- real_list, 206
- realtime option, 176
- recompilation, 149
 - cut-off, 150
- record, 43
 - runtime representation of, 47
 - unboxed, 47
- recursion
 - polymorphic, 65
- ref, 81
- reference, 81
 - local, 84
- referencing an MLB-file, 147
- RegionExp, 34, 37
- region, 20
 - auxiliary, 102
 - deallocation, 46, 127
 - dropping of, 64
 - global, 96, 97
 - resetting, 22, 101
- region option, 179
- region**, 46
- region aliasing, 108
- region arity, 88
- region descriptor, 33, 109
- region endomorphism, 28, 65, 109, 112, 114, 189
- region exomorphism, 65, 75, 189
- region flow graph, 108, 166
- region inference, 22
 - ground rule, 73
- region name, 34, 103
- region pages, 33, 177
- region parameter, 61
 - actual, 59, 60
 - formal, 59, 60, 108
- region polymorphism, 59–70, 102, 126, 139
- region profile, 165, 179
- region profiling, 24
- region_profiling, 174

- region size, 21, 33, 187
 - finite, 33
 - infinite, 33
- region stack, 20
- region statistics, 177
- region variable, 35, 43
 - auxiliary, 88
- region vector
 - removed, 214
- region-annotated type, 44
- region-annotated type scheme, 61
 - printing of, 129
- region-annotated type scheme with place,
 - programming, 212
 - 62
- region.ps, 120
- register, 43, 138
 - standard argument, 141, 142
 - standard closure, 141, 142
- ReML, 11
- Repository
 - Github, 211
- resetRegion, 201
- resetRegions, 22
- resetRegions, 114
- round, 50
- rp2ps, 32, 120
- rp2ps options, 179–184
- run, 39, 150
- runtime stack, 33
- runtime system, 38
- runtime type, 34
- sampleMax option, 120, 180
- sat, 106
- scan, 118
- scan_rev1.mlb, 168
- scan_rev2.mlb, 173
- scope rules, 71–76
- sec option, 176
- second, 195
- separate compilation, 12, 213
- Sieve of Eratosthenes, 74
- signature constraint
 - opaque, 152
 - transparent, 152
- signature declaration, 152
- size, 51
- sizeStringDefine, 196
- smallPrime, 73
- socket
 - programming, 212
- sortBySize option, 183
- sortByTime option, 183
- source file, 149
- specialisation
 - functor, 153
- spreading, 87
- stack, 9, 19, 22, 72, 97, 138
 - stack option, 180
- stack band, 109
- stack profile, 166, 180
- standard argument register, 141, 142
- standard closure register, 141, 142
- Standard ML, 9
 - 1997 revision, 49
 - Basis Library, 49, 213
 - Modules, 37, 147
- Standard ML Basis Library, 215
- stat option, 183
- statistics_after_optimisation, 185
- StatObject, 215
- storage mode, 102
- storage mode analysis
 - modular, 11, 109
- str, 51
- String.h, 193

- strongly connected component, 175
- `strrev`, 207
- structure declaration, 151
- substitution, 129
- `substring`, 51
- tagging, 213
- `Tagging.h`, 193
- tail recursion, 76
- target program, 39
- `TextIO`, 80
- timer
 - prof, 176
 - real, 176
 - virtual, 176
- `t1`, 98, 196
- top of region, 102
- `traverse_list`, 196
- tree
 - binary, 87
- `tree`, 87
- `trunc`, 50
- tuple, *see* record
- type
 - region-annotated, 43, 56, 83, 126, 129
- type scheme
 - region-annotated, 61
- type scheme with place
 - region-annotated, 62
- type with place, 44
- unboxing, 11, 89, 91
 - boxity, 92
 - enum, 91
 - high, 91
 - low, 91
- `unit`, 47
- untagging, 213
- `val`, 126
- value
 - boxed, 36, 45
 - unboxed, 36, 45, 91
- value declaration, *see* declaration
- value representation, 213
- variable
 - lambda-bound, 125
 - locally live, 105
 - own, 85
- `-virtualtime` option, 176
- web site, 13
- `Word31` structure, 50
- `Word32` structure, 50
- `Word63` structure, 50
- `Word64` structure, 50
- `Word8` structure, 50

Global Regions

-
- r1 Holds values of type **top**, including records, exceptions, reals, and closures.
 - r2 Holds values of type **bot**. Because no values has type **bot**, this region contains no values. Region variables with region type **bot** are used with so-called explicit regions, which are not covered in this report.
 - r3 Holds values of type **string**.
 - r4 Holds values of type $\tau_1 \times \tau_2$, for any types τ_1 and τ_2 .
 - r5 Holds values of type τ **array** and τ **vector**, for any type τ .
 - r6 Holds values of type τ **ref**, for any type τ .
 - r7 Holds values of $\tau_1 \times \tau_2 \times \tau_3$, for any types τ_1 , τ_2 , and τ_3 .
-