

Compiling with Continuations and LLVM

Kavon Farvardin John Reppy

University of Chicago
Illinois, USA

{kavon, jhr}@cs.uchicago.edu

LLVM is an infrastructure for code generation and low-level optimizations, which has been gaining popularity as a backend for both research and industrial compilers, including many compilers for functional languages. While LLVM provides a relatively easy path to high-quality native code, its design is based on a traditional runtime model which is not well suited to alternative compilation strategies used in high-level language compilers, such as the use of heap-allocated continuation closures.

This paper describes a new LLVM-based backend that supports heap-allocated continuation closures, which enables constant-time `callcc` and very-lightweight multithreading. The backend has been implemented in the Parallel ML compiler, which is part of the Manticore system, but the results should be useful for other compilers, such as Standard ML of New Jersey, that use heap-allocated continuation closures.

1 Introduction

Maintaining a native code generator that targets multiple architectures is a hassle for compiler writers that requires expert knowledge of each new processor’s quirks. Some functional-language compilers avoid this problem by targeting C as a “portable assembly language” [28, 18], but this approach has significant drawbacks in both compile-time and runtime performance. More recently, LLVM [20, 21], which provides a low-level SSA-based representation with many available optimization passes, has emerged as a popular backend for compiler writers. Although designed with imperative and object-oriented languages as its expected clients, LLVM has been used to build backends for functional languages, such as Standard ML [23], SML# [31], Haskell [30], and Erlang [27]. While LLVM addresses the problem of maintaining native code generators and is a better portable assembly language than C code, it still suffers from a bias toward C runtime conventions, which makes it a less than ideal target for a functional-language compiler. Functional language implementations often use specialized register and calling conventions,¹ and require guaranteed tail-call optimization, mechanisms to communicate with the garbage collector, and efficient support for features like first-class continuations.

In this paper, we present our approach to solving the problems of using LLVM as a backend for functional-language implementations. In particular, we show how to use LLVM to support the heap-allocated first-class continuation runtime model [1] used by the SML/NJ system and by the Manticore system [16]. We have integrated our approach into the *Parallel ML* (PML) compiler that is part of the Manticore project [16]. Initial observations suggest that the LLVM backend produces more efficient code relative to the previous MLRisc [17] backend.

The LLVM backends for the Glasgow Haskell Compiler (GHC) [30] and Erlang (ErLLVM) [27] use special language-specific calling conventions added to LLVM that support tail-call optimization (TCO).

¹For example, many implementations dedicate a specific register as an allocation pointer to support efficient open-coding of heap allocation.

The LLVM backend for the MLton SML compiler uses trampolining to avoid the issues with TCO [23]. As far as we know, no one has successfully implemented heap-allocated first-class continuations with LLVM. In the remainder of the paper, we describe a new calling convention that we have added to LLVM, how we handle the GC interface, and how we support capturing continuations to support preemption.

2 Background

In this section, we provide an overview of Manticore’s PML compiler, along with the high-level overview of how the LLVM backend was integrated with the compiler.

2.1 The PML Compiler

The Parallel ML compiler (**pmlc**) that lies at the heart of the Manticore project is implemented by a sequence of translations between intermediate languages (IRs) [14]. There are six distinct IRs in the compiler:

1. Parse tree — the product of the parser.
2. AST — an explicitly-typed abstract-syntax tree representation.
3. BOM — a direct-style normalized λ -calculus.
4. CPS — a continuation-passing-style λ -calculus.
5. CFG — a first-order control-flow-graph representation.
6. MLTree — the expression tree representation used by the MLRISC code generation framework [17].

We support the parallelism and concurrency features of PML by transformations on the AST and BOM representations that introduce explicit continuation binders [15, 16, 4, 25]. The translation from BOM to CPS uses the Danvy-Filinski CPS transformation [7] and we perform several kinds of optimizations on the CPS IR [5, 3]. The conversion from CPS to CFG uses a flat, safe-for-space, closure representation [6]. All closures, including those for return continuations, are immutable and allocated on the heap. The CFG IR produced through closure conversion is still in continuation-passing style, but functions are no longer nested. Finally, we generate an MLTree representation from the CFG and use the MLRISC code generation framework to handle instruction selection and register allocation. In this paper, we describe our experience with replacing the MLRISC code generation framework by LLVM [21] as illustrated in Figure 1.

2.2 The CFG Representation

The CFG IR is a first-order, machine-like representation whose primary construct is the basic block (Figure 2). Functions consist of a calling convention, a start block, and a list of body blocks. Each block ends in transfer to another block within the function or a tail-call to either a function or continuation. Throws to the continuation of a local control-flow divergence (*i.e.*, a local join point) are identified during closure conversion and are translated as a transfer to a block within the enclosing function.

In preparation for generating native code that uses bump allocation in the heap, we insert heap-limit tests into the CFG representation. We use GC tests as safe-points where preemptive thread switching can occur [26]. Thus, we must guarantee that every loop (even ones that do not allocate) includes at least one heap limit test. To ensure every loop has a limit test, we use a bounded algorithm that

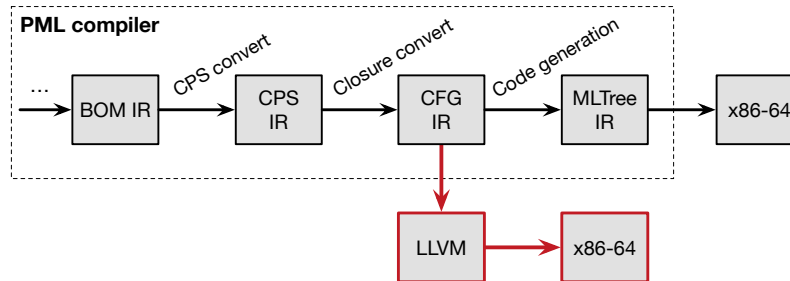


Figure 1: The middle and backends of the PML compiler. The new code-generation path is highlighted in red.

```

1 datatype ty
2   = Raw of machine_ty
3   | Tuple of ty list
4   | Addr of ty
5   | Fun of ...
6   ...
7 type var = id * ty
8 datatype func = FUNC of {
9   entry : convention,
10  start : block,
11  body : block list
12 }
13 and block = BLK of {
14   name : label,
15   args : var list,
16   body : exp list,
17   exit : transfer
18 }
19 and transfer
20 = ApplyFun of ...
21 | ThrowCont of ...
22 | If of ...
23 | HeapCheck of ...
24 ...
  
```

Figure 2: Simplified SML type definitions for the CFG IR.

minimizes the feedback vertex set of the control-flow graph of the program. The limit tests introduce continuation captures that are not present in the original program, which can be used to implement preemptive scheduling (Section 3.5).

2.3 Interfacing with LLVM

There are several ways to generate code for LLVM. One can generate code directly using either the native C++ APIs, which gives one full access to all of the LLVM features, or the more restricted C APIs, which do not support some newer features, such as `musttail` calls. Lastly, one can generate LLVM assembly code and use the `llc` command-line tool to generate assembly code. Since there are no SML bindings for LLVM in the SML/NJ system that we use to develop our code, we built a library for generating LLVM assembly code that is compatible with LLVM 3.8 or newer.

3 Translating to LLVM IR

In this section, we provide complete details of the translation from the CFG IR to LLVM IR, while touching on aspects of our runtime system and LLVM's code generation.

3.1 Conversion to SSA

The LLVM IR is a static-single-assignment (SSA) representation, but it is possible to generate non-SSA code and allow LLVM to do the SSA conversion for you. The approach, which is used by the MLton compiler [23], is to stack allocate variables and emit `load/store` instructions whenever accessing them. The `mem2reg` pass in LLVM will then promote these values to virtual register accesses, using good heuristics to insert ϕ -nodes in the program as-needed [8].

The CFG IR is already in a static single-assignment form, but data-flow joins within a function do not use an explicit ϕ -node to merge their values as in a typical SSA representation. Instead, the CFG IR takes a functional approach by assigning parameters to each block, with the predecessors of a block making up the set of possible values for each parameter [2]. Our library to generate LLVM follows suit by using parameterized blocks and requiring the user to note incoming block transfers. With this information, it is straightforward to emit ϕ -nodes for each block. While this produces a ϕ -node per parameter in each block, the `instcombine` optimization pass in LLVM eliminates redundant nodes that simply rename a value.

3.2 Type Correspondences

The main base types in the CFG IR, *i.e.*, integers, floats, and addresses, have an obvious corresponding type in LLVM. However, there are two separate difficulties when representing tuples and functions in LLVM.

Ideally, a tuple of values in the CFG IR corresponds to a pointer to a struct in LLVM, providing richer type information for optimization. But, our representation of tuples in memory uses a field alignment convention that deviates from those used in languages such as C. While an LLVM module can specify a `datalayout` string that describes such alignments, this string is only used to inform LLVM's optimizer of what the *code generator* will do. It is currently not possible to override the code generator's memory alignment assumptions for struct fields. Thus, all tuples are represented as a pointer to an 8-byte integer

so that all fields are 8-byte aligned; values are casted as needed for loads and stores involving field pointers.

There are two factors that constrain the type of a CFG function in LLVM. The first is that our runtime system expects certain register conventions for function arguments, so we must order them specially (Section 3.4). The second is that tail calls emitted by our compiler *must* remain as tail calls, so we use the `musttail` call marker (Section 3.3). A `musttail` call requires that the types of the caller and callee match, modulo the type pointed to by a pointer [19].

Since all invocations of CFG functions and continuations are tail calls, and Manticore’s runtime system uses more than one register convention to pass arguments to functions, we must carefully match up the LLVM types used by caller and callee. Every CFG function is given a new type in LLVM by padding the parameter list with 64-bit integers, so that every general-purpose register that could ever be used for passing arguments will be consumed by that type. Functions of this type allow us to skip function casts and implement multiple register conventions without needing more than one LLVM calling convention, since we can reorder parameters so they end up in the right machine register. We pass `undef` values for arguments not used by the callee, and these `undef` values disappear when LLVM generates machine code.

3.3 Proper Tail-Call Optimization

Most functional languages use tail recursion to express loops and, thus, require that tail calls do not grow the stack. In limited circumstances, LLVM can meet this requirement, but even when LLVM is able to apply TCO to a call, there is extra stack-management overhead on function entry and exit [11, 12, 24].

This extra overhead is bad enough for implementations that rely on traditional stacks, but it is prohibitively expensive for implementations that use CPS and heap-allocated continuations [1].

A standard technique to avoid this overhead is to merge mutually tail-recursive functions into a single LLVM function and then use internal jumps instead of tail calls. Unfortunately, this approach does not work for tail calls to unknown functions and incurs substantial compile-time cost. The MLton SML compiler partitions the generated code into multiple large functions (called *chunks*) and uses a trampoline to transfer control between chunks [23]. This approach solves the compile-time issue, but places extra strain on the register allocator because of the nature of the control-flow graphs in the chunks.

Our solution to this problem is to add a new calling convention, which we call Jump-With-Arguments (JWA), to LLVM. This calling convention has the property that it uses all of the available hardware registers and that no registers are preserved by either the caller or callee. Furthermore, the argument registers are exactly the live registers on return (*i.e.*, the call and return have identical register use conventions).² The JWA convention is not specific to any architecture or version of LLVM, though we have currently only implemented it for x86-64 in a fork of LLVM to support the PML compiler.

Second, we mark every function with the `naked` attribute, which tells the code generator to completely skip the emission of a function prologue and epilogue. This attribute *must* be used with care; it was originally designed to support functions consisting entirely of inline assembly that manages the stack explicitly (*e.g.*, interrupt service routines). Using the `naked` attribute means that the generated code is responsible for ensuring that there is sufficient stack space for register spills and any callee-saved registers are preserved.

We use an assembly language shim for switching between the runtime system code (written in C) and the code generated by the PML compiler. This shim code allocates a frame (Figure 3) that is large enough

²These properties are why we need to create a new convention, instead of using an existing convention.

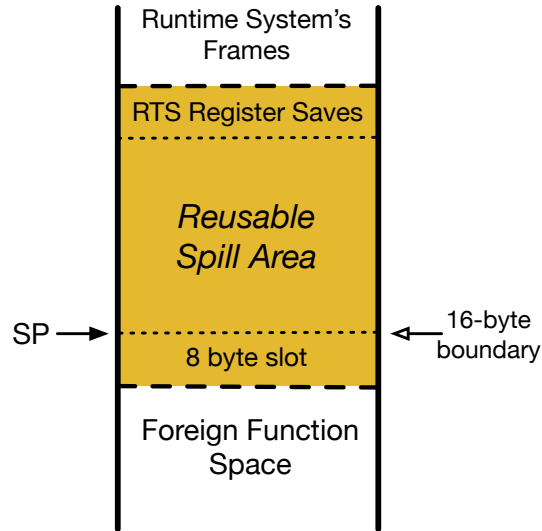


Figure 3: The shaded frame is setup by runtime system and is shared by all PML functions.

to handle the maximum number of register spills. This technique is borrowed from the SML/NJ system; the spill limit is enforced by the compiler limiting the number of live variables at any control point and over-provisioning the spill area.³ In short sequences, LLVM’s code generator may introduce one or two additional spills, but by limiting the number of values early on, we effectively bound the number of register spills. LLVM’s code generator will assign any register spills to frame locations starting from the bottom of the frame, offset from the stack pointer, so the shim preserves all C callee-saved registers at the top of the frame.

3.4 Allocation and Garbage Collection

Our JWA calling convention maps function arguments to hardware registers based on the position of the argument. By using standard positions for special runtime registers, we can effectively pin them to hardware registers (*e.g.*, we always pass the allocation pointer as the first argument). Object allocation then defines new instances of the allocation pointer, which thread the current state of the pointer through the code (recall that LLVM code is in SSA form).

One of the advantages of CPS with heap-allocated continuations is that the interface to garbage collection is very simple. The runtime does not need to scan a stack (since there is no stack) or understand any other properties of the code generator. We did not have to make any modifications to our existing collector to support our LLVM backend.

The compiler is responsible for generating code to check for heap exhaustion and code to invoke the GC when necessary. In Figure 4, we list simplified LLVM code for the heap-limit check (Lines 4–5) and GC invocation. To invoke the GC, we first save the live variables into a new heap object called `roots` using bump-allocation (Lines 8–10) and then do a non-tail JWA call to `@invoke-gc`. When this function returns, we restore the allocation pointer and live variables (Lines 13–15).

We use a non-tail call to `@invoke-gc` for reasons described in below in Section 3.5. We ensure that LLVM does not try to preserve values across the `@invoke-gc` call by taking advantage of the rules about

³As there is only one spill area per hardware thread, allotting a few kilobytes is no problem.

```

1 declare jwa {i64*, i64*} @invoke-gc(i64*, i64*)
2
3 define jwa void @foo(i64 allocPtr_0, ... ) naked {
4 ...
5   if enoughSpace, label continue, label doGC
6
7 doGC:
8   roots_0 = allocPtr_0
9   ; ... save live vals in roots_0 ...
10  allocPtr_1 = getelementptr allocPtr_0, 5 ; bump
11  retV = call jwa {i64*, i64*}
12         @invoke-gc(allocPtr_1, roots_0)
13  allocPtr_2 = extractvalue retV, 0
14  roots_1 = extractvalue retV, 1
15  ; ... restore live vals ...
16  goto label continue
17
18 continue:
19  allocPtr_3 = phi i64* [ allocPtr_0, allocPtr_2 ]
20  liveVal_1 = phi i64* [ ... ]
21 ...

```

Figure 4: An example of a compiler generated safe point for garbage collection.

aliasing. Once the pointer reaching all live values is passed to `@invoke-gc`, which is an external function not visible to LLVM, LLVM must assume that all values may have changed and must use the updated versions from the pointer returned.

3.5 Preemption and Multithreading

The main motivation for supporting heap-allocated first-class continuations is to enable the efficient implementation of the concurrency mechanisms in the Manticore runtime system [15, 22]. While the mechanisms described in Section 3.3 are sufficient to support the explicit management of continuations, preemptive scheduling requires capturing continuations that are not explicit in the intermediate code. We use the technique developed for supporting asynchronous signals in SML/NJ [26], which limits preemption to *safe points* where all live values are known. Specifically, those places in the code where we perform a heap-limit check serve as safe points.⁴ We store the heap limit pointer in memory, which means that we can set it to zero when we need to force a heap-limit check to fail. The runtime system then constructs a continuation closure, which is passed to the preemption handler where it can be put on scheduling queue *etc.*

This mechanism introduces an additional challenge for our LLVM backend, because the implicit continuations that are captured during preemption do not correspond to LLVM functions and are invoked from unknown locations. For example, consider the heap-limit test in Figure 4. If it is invoked to force a preemption, then a continuation will be created by the runtime system that has Line 13 as its code address (*i.e.*, the return address of the call to `@invoke-gc`). Since Line 13 is not a function entry, there is no way in LLVM to specify a calling convention.

⁴The code generator ensures that even non-allocating loops contain a heap-limit check.

```

1 define jwa void @foo( ... ) naked {
2 ...
3 preempted:
4   env = ; ... save live vars ...
5   closPtr = allocPair (undef, env)
6   ret = call jwa {i64*, i64*}
7         @genLabel(closPtr, @enterRTS)
8   arg1 = extractvalue ret, 0
9   arg2 = extractvalue ret, 1
10 ...
11 }

```

(a) A non-tail call to produce a “block” label.

```

1 ; call convention:
2 ; rsi = closPtr, r11 = @enterRTS
3 genLabel:
4   pop rax          ; pop ret addr
5   mov rax,(rsi) ; finish closure
6   jmp r11

```

(b) Assembly shim to pick up the label from the stack.

Figure 5: An example of generating a first-class label in LLVM.

We are saved, however, by the fact that we can specify our own return convention for structs in LLVM. Normal conventions return a struct using a mix of registers or stack space that does not match up with the way arguments are passed to functions. We setup our JWA convention so that struct field i is returned in the same register that argument i would be passed in during a call. This way, return addresses generated through a non-tail JWA call can be jumped to safely using a standard JWA tail call (which is how we throw to a continuation!).

To help illustrate this point, consider the example in Figure 5. On lines 4–5, we save all live values into a new heap-allocated closure, but the code pointer is left uninitialized. This incomplete closure along with the function we intend to invoke are passed in a non-tail call to `@genLabel`. The key point is that this call will push a return address on the stack that resumes on line 8, which resumes execution by using the values returned. We ensure that there are no values live across the non-tail call, so the stack frame of the caller, `@foo`, can be safely reused. Then, the assembly shim `@genLabel` pops the return address, which has a calling convention that is identical to a call, and places it into the closure before invoking `@enterRTS`.

3.6 LLVM Optimizations

One of the benefits of using LLVM is that it provides a rich set of optimization passes and opportunities to tune the output of its code generator. A particularly beneficial optimization is to guide LLVM’s basic block placement for heap limit tests, since these tests normally fail (*i.e.*, the heap is not exhausted). It is important to keep basic blocks that handle overflow away from the hot path of the function to reduce instruction cache pressure. We use the built-in `@llvm.expect` intrinsic along with the `lowerexpect` pass to add branch probabilities to these tests, which guides the code generator when performing block layout.

We also hand-crafted two sequences of LLVM optimization passes, using trial-and-error guided by

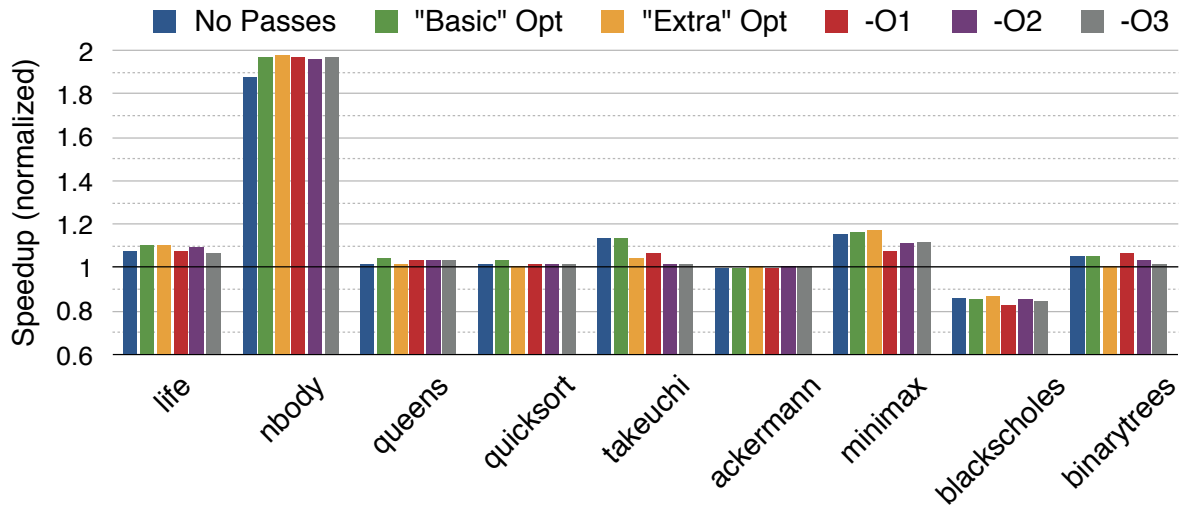


Figure 6: Execution time speedups when compiled with LLVM, normalized to MLRisc. Each bar represents a different set of additional optimizations applied when compiling with LLVM.

intuition,⁵ to simplify the LLVM program output by the PML compiler. Crafting a custom pass sequence is recommended for compiler frontends that target LLVM, as the default “-Ox” passes are tuned for a C/C++ frontend [9]. Our “Basic” optimization will only shrink the size of the program, and consist of a combination of the following passes: `simplifycfg`, `instcombine`, `reassociate`, `constprop`, `early-cse`, `gvn`, and `dce`. “Extra” optimization adds passes to the “Basic” suite that can specifically optimize memory operations, as they are very common: `sink`, `mldst-motion`, and `slp-vectorizer`.

4 Evaluation

We measured the difference in application performance between our two backends to get a sense of how well LLVM can handle the unusual code we generate. Figure 6 summarizes our experiment conducted on a workstation equipped with two E5-2687W processors and 64GB of memory. Speedups reported are relative to the MLRisc backend and were computed using the average of 50 trials. Error bars are omitted as the standard error was less than 1.5%. Each benchmark was compiled with different sets of additional LLVM optimization passes applied before generating assembly (Section 3.6). Optimizations of the form “-Ox” indicate one of the default optimization levels built into LLVM.

The significantly better performance by LLVM over MLRisc on the `nbody` benchmark is owed to poor use of floating-point registers. We have recently identified a mistake in the way that the MLRisc does register allocation for the x86-64, which results in significantly more register shuffling. We have not yet had the opportunity to fix MLRisc, so we do not know how big of a difference this fix will make in the results.

On the other hand, we have not found a reason for the notably worse performance when using LLVM for the parallel `blackscholes` benchmark. When testing this benchmark on a 2013 Macbook Pro, the LLVM-generated version matches or outperforms the MLRisc version. Thus, the `blackscholes` performance may be sensitive to the particular CPU(s) on the machine.

⁵Ideally, we would automate this process [13].

While the takeuchi benchmark only tests the overhead of recursion, its performance takes a hit with aggressive optimization in LLVM because of the `slp-vectorizer` pass. After the pass is applied, the hot path is smaller because of the use of vector instructions to initialize a closure, but the execution cost of such instructions outweighed the size benefit. Overall, LLVM seems to produce better code than MLRisc does, with some programs making significant gains.

5 Related Work

In general, preemption and explicit stack management are treated similarly in the GHC and PML compilers. GHC uses a whole-program CPS transformation to make the call stack explicit in its first-order intermediate representation, regardless of the backend being used. But because of the lack of first-class labels in LLVM, GHC is forced to perform an additional splitting transformation (*i.e.*, “proc-point splitting”) at every return point, such as safe points [29]. A solution similar to the technique described in Section 3.5 can remove the need for a splitting transform when targeting LLVM.

Dolan *et al.* proposed the SWAPSTACK mechanism for LLVM to enable lightweight context switching [10]. To capture stack-allocated one-shot continuations in LLVM, a mechanism like SWAPSTACK can be used in conjunction with runtime system support. While our focus is on fully general first-class continuations via immutable heap-allocated frames, our mechanism for `@genLabel` (Figure 5) is similar in spirit to SWAPSTACK. A major difference is that we leverage liveness properties of a program in continuation-passing style to correctly implement `@genLabel`. In addition, *explicit* continuation captures in the original program do not need our mechanism at all, thus avoiding the runtime system.

6 Conclusion and Future Work

We have outlined how to extend LLVM to support the heap-allocated first-class continuation runtime model. We are in the process of replacing the MLRisc backend with LLVM using the approach described in this paper. Initial observations suggest that this new LLVM backend produces smaller and more efficient code. We are also hoping to apply these techniques to the SML/NJ system, and integrate a generalized form of this work into LLVM in the future.

Acknowledgments

This material is based upon work supported, in part, by the National Science Foundation under Grant CCF-1010568. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of these organizations or the U.S. Government.

References

- [1] Andrew W. Appel (1992): *Compiling with Continuations*. Cambridge University Press, Cambridge, England.
- [2] Andrew W. Appel (1998): *SSA is Functional Programming*. *SIGPLAN notices* 33(4), pp. 17–20, doi:10.1145/278283.278285.
- [3] Lars Bergstrom, Matthew Fluet, Matthew Le, John Reppy & Nora Sandler (2014): *Practical and Effective Higher-order Optimizations*. In: *ICFP '14*, ACM, New York, NY, pp. 81–93, doi:10.1145/2628136.2628153.

- [4] Lars Bergstrom, Matthew Fluet, Mike Rainey, John Reppy, Stephen Rosen & Adam Shaw (2013): *Data-Only Flattening for Nested Data Parallelism*. In: *PPoPP '13*, ACM, New York, NY, pp. 90–106, doi:10.1145/2442516.2442525.
- [5] Lars Bergstrom & John Reppy (2009): *Arity raising in Manticore*. In: *IFL '09*, LNCS, Springer-Verlag, New York, NY, pp. 90–106, doi:10.1007/978-3-642-16478-1_6.
- [6] Luca Cardelli (1983): *The Functional Abstract Machine*. Technical Report TR-107, Bell Laboratories.
- [7] Olivier Danvy & Andrzej Filinski (1992): *Representing Control: A study of the CPS transformation*. *MSCS* 2(4), pp. 361–391, doi:10.1017/S0960129500001535.
- [8] LLVM Developers (2017): *LLVMs Analysis and Transform Passes*. Available at <http://llvm.org/docs/Passes.html>.
- [9] LLVM Developers (2017): *Performance Tips for Frontend Authors*. Available at <https://llvm.org/docs/Frontend/PerformanceTips.html#pass-ordering>.
- [10] Stephen Dolan, Servesh Muralidharan & David Gregg (2013): *Compiler Support for Lightweight Context Switching*. *ACM Trans. Archit. Code Optim.* 9(4), pp. 36:1–36:25, doi:10.1145/2400682.2400695.
- [11] Rafael Ávila de Espíndola (2012): *LLVM Bug 13826 - unreachable prevents tail calls*. Available at https://llvm.org/bugs/show_bug.cgi?id=13826.
- [12] Kavon Farvardin (2015): *LLVM Bug 23766 - musttail calls are not allowed to precede unreachable*. Available at https://llvm.org/bugs/show_bug.cgi?id=23766.
- [13] Kavon Farvardin (2017): *autotune – discover good LLVM passes*. Available at <https://github.com/kavon/autotune>.
- [14] Matthew Fluet, Nic Ford, Mike Rainey, John Reppy, Adam Shaw & Yingqi Xiao (2007): *Status Report: The Manticore Project*. In: *ML '07*, ACM, New York, NY, pp. 15–24, doi:10.1145/1292535.1292539.
- [15] Matthew Fluet, Mike Rainey & John Reppy (2008): *A Scheduling Framework for General-purpose Parallel Languages*. In: *ICFP '08*, ACM, New York, NY, pp. 241–252, doi:10.1145/1411203.1411239.
- [16] Matthew Fluet, Mike Rainey, John Reppy & Adam Shaw (2011): *Implicitly-threaded Parallelism in Manticore*. *JFP* 20(5–6), pp. 537–576, doi:10.1017/S0956796810000201.
- [17] Lal George, Florent Guillaume & John Reppy (1994): *A portable and optimizing back end for the SML/NJ compiler*. In: *CC '94*, LNCS 786, Springer-Verlag, New York, NY, pp. 83–97, doi:10.1007/3-540-57877-3_6.
- [18] Simon Peyton Jones, Norman Ramsey & Fermin Reig (1999): *C--: A Portable Assembly Language that Supports Garbage Collection*. In: *PPDP '99*, Springer-Verlag, New York, NY, pp. 1–28, doi:10.1007/10704567_1.
- [19] Reid Kleckner (2014): *Proposal: Add a guaranteed tail call marker*. Available at <http://lists.llvm.org/pipermail/llvm-dev/2014-April/071704.html>.
- [20] Chris Lattner (2002): *LLVM: An Infrastructure for Multi-Stage Optimization*. Master's thesis, C.S. Dept., UIUC, Urbana, IL.
- [21] Chris Lattner & Vikram Adve (2004): *LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation*. In: *CGO '04*, pp. 75–. Available at <http://dl.acm.org/citation.cfm?id=977395.977673>.
- [22] Matthew Le & Matthew Fluet (2015): *Partial Aborts for Transactions via First-class Continuations*. In: *ICFP '15*, ACM, New York, NY, pp. 230–242, doi:10.1145/2784731.2784736.
- [23] Brian Andrew Leibig (2013): *An LLVM Back-end for MLton*. Master's thesis, Rochester Institute of Technology. Available at https://www.cs.rit.edu/~mtf/student-resources/20124_leibig_msproject.pdf.
- [24] David Majnemer (2015): *LLVM Bug 23470 - Inefficient lowering of 'musttail' call*. Available at https://llvm.org/bugs/show_bug.cgi?id=23470.
- [25] John Reppy, Claudio Russo & Yingqi Xiao (2009): *Parallel Concurrent ML*. In: *ICFP '09*, ACM, New York, NY, pp. 257–268, doi:10.1145/1596550.1596588.

- [26] John H. Reppy (1990): *Asynchronous signals in Standard ML*. Technical Report TR 90-1144, Dept. of CS, Cornell University, Ithaca, NY.
- [27] Konstantinos Sagonas, Chris Stavrakakis & Yiannis Tsiouris (2012): *ErLLVM: An LLVM Backend for Erlang*. In: *ERLANG '12*, ACM, New York, NY, pp. 21–32, doi:10.1145/2364489.2364494.
- [28] David Tarditi, Peter Lee & Anurag Acharya (1992): *No Assembly Required: Compiling Standard ML to C*. *ACM LOPLAS* 1(2), pp. 161–177, doi:10.1145/151333.151343.
- [29] GHC Team (2011): *Work in Progress on the LLVM Backend: Get rid of Proc Point Splitting*. Available at <https://ghc.haskell.org/trac/ghc/wiki/Commentary/Compiler/Backends/LLVM/WIP#GetridofProcPointSplitting>.
- [30] David A. Terei & Manuel M.T. Chakravarty (2010): *An LLVM Backend for GHC*. In: *HASKELL '10*, ACM, New York, NY, pp. 109–120, doi:10.1145/1863523.1863538.
- [31] Katsuhiko Ueno & Atsushi Ohori (2014): *Compiling SML# with LLVM: a Challenge of Implementing ML on a Common Compiler Infrastructure*. In: *Workshop on ML*, pp. 1–2. Available at https://sites.google.com/site/mlworkshope/smlsharp_llvm.pdf.