

15-411 Compiler Design, Fall 2015

Lab 4

Rob Simmons, Will Crichton, Grant Della Silva, Matt Bryant, Anshu Bansal

Test Programs Due: 11:00pm, Tuesday, October 27, 2015

Compilers Due: 11:00pm, Tuesday, November 3, 2015

1 Introduction

The goal of the lab is to implement a complete compiler for the language L4. This language extends L3 with pointers, arrays, and structs. With the ability to store global state, you should be able to write a wide variety of interesting programs. As always, correctness is paramount, but you should take care to make sure your compiler runs in reasonable time.

2 L4 Syntax

The lexical specification of L4 is changed by adding '[' , ']', '.', and '->' as lexical tokens; see Figure 1. Whitespace, token delimiting, and comments are unchanged from languages L1, L2, and L3.

The syntax of L4 is defined by the (no longer context-free!) grammar in Figure 2. Ambiguities in this grammar are resolved according to the operator precedence table in Figure 3 and the rule that an `else` provides the alternative for the most recent eligible `if`. Note that according to this precedence table, `*x++` parses as `*(x++)`, which isn't valid syntax. While we could allow `*x++` because `*x` is unambiguously an lvalue and `(*x)++` is unambiguously a statement, we disallow it. This is both to match the precedence rules and also to avoid confusion with the different semantics that statement has in C.

L4 syntax is not context-free

As noted above, the grammar presented for L4 is no longer context free. Consider, for example, the statement

```
foo * bar;
```

If `foo` is a type name, then this is a declaration of a `foo` pointer named `bar`. If, however, `foo` is *not* a type name, then this is a multiplication expression used as a statement.

For those of you using parser combinator libraries, you will be able to backtrack from a parse decision based on whether an identifier is a type name, so this case should not be a problem. However, backtracking raises the specter of a performance bug, so you will need to closely consider performance.

```

ident ::= [A-Za-z_][A-Za-z0-9_]*
num ::= <decnum> | <hexnum>

<decnum> ::= 0 | [1-9][0-9]*
<hexnum> ::= 0[xX][0-9a-fA-F]+

<special characters> ::= ! ~ - + * / % << >>
< > >= <= == != & ^ | && ||
= += -= *= /= %= <<= >>= &= |= ^=
-> . -- ++ ( | ) [ ] , ; ? :

<reserved keywords> ::= struct typedef if else while for continue break
return assert true false NULL alloc alloc_array
int bool void char string

```

Terminals referenced in the grammar are in **bold**. Other classifiers not referenced within the grammar are in *<angle brackets and in italics>*. **ident**, *<decnum>*, and *<hexnum>* are described using regular expressions.

Figure 1: Lexical Tokens

However, those of you using parser generators will have a harder time—the decision whether to shift or reduce could be made well ahead of when an identifier is determined to be a typename or not. Solving this ambiguity is a bit tricky; below, we describe two approaches.

One way to handle this is to perform an ambiguous parse: use one rule to parse both the declaration form and the expression form. Then undo an incorrect decision during elaboration. This approach will almost certainly involve some other adjustment to various pieces of your grammar and lexer. Looking over grammars and parsers from past years, we cannot honestly recommend this technique—results seem to be largely contorted with low confidence in correctness.

Another option is to prevent incorrect decisions from being made. New type identifiers are introduced at the top level (as a *<gdecl>*); you may wish to parse one global declaration at a time, with suitable changes to the start-of-parse and end-of-parse symbols). With this approach, the lexer can produce different tokens for type and non-type identifiers. This was the solution intended by the designers of C, if one considers the relevant footnote in their book.

This solves the parsing problem, but raises another. The lexer performs a lookahead in order to find the longest match. This affects the lexing of a token which is used immediately after it is introduced—consider:

```

typedef int foo;
foo func();

```

In this case, if the parser parses `typedef int foo;` the lexer may already have lexed the `foo` at the beginning of the next line, so be careful! Despite this potential issue, we recommend this approach because the grammar will continue to look natural and follow the actual understanding of the language syntax.

$\langle \text{program} \rangle$	$::= \epsilon \mid \langle \text{gdecl} \rangle \langle \text{program} \rangle$
$\langle \text{gdecl} \rangle$	$::= \langle \text{fdecl} \rangle \mid \langle \text{fdef} \rangle \mid \langle \text{typedef} \rangle \mid \langle \text{sdecl} \rangle \mid \langle \text{sdef} \rangle$
$\langle \text{fdecl} \rangle$	$::= \langle \text{type} \rangle \mathbf{ident} \langle \text{param-list} \rangle ;$
$\langle \text{fdef} \rangle$	$::= \langle \text{type} \rangle \mathbf{ident} \langle \text{param-list} \rangle \langle \text{block} \rangle$
$\langle \text{param} \rangle$	$::= \langle \text{type} \rangle \mathbf{ident}$
$\langle \text{param-list-follow} \rangle$	$::= \epsilon \mid , \langle \text{param} \rangle \langle \text{param-list-follow} \rangle$
$\langle \text{param-list} \rangle$	$::= () \mid (\langle \text{param} \rangle \langle \text{param-list-follow} \rangle)$
$\langle \text{typedef} \rangle$	$::= \mathbf{typedef} \langle \text{type} \rangle \mathbf{ident} ;$
$\langle \text{sdecl} \rangle$	$::= \mathbf{struct} \mathbf{ident} ;$
$\langle \text{sdef} \rangle$	$::= \mathbf{struct} \mathbf{ident} \{ \langle \text{field-list} \rangle \} ;$
$\langle \text{field} \rangle$	$::= \langle \text{type} \rangle \mathbf{ident} ;$
$\langle \text{field-list} \rangle$	$::= \epsilon \mid \langle \text{field} \rangle \langle \text{field-list} \rangle$
$\langle \text{type} \rangle$	$::= \mathbf{int} \mid \mathbf{bool} \mid \mathbf{ident} \mid \mathbf{void} \mid \langle \text{type} \rangle * \mid \langle \text{type} \rangle [] \mid \mathbf{struct} \mathbf{ident}$
$\langle \text{block} \rangle$	$::= \{ \langle \text{stmts} \rangle \}$
$\langle \text{decl} \rangle$	$::= \langle \text{type} \rangle \mathbf{ident} \mid \langle \text{type} \rangle \mathbf{ident} = \langle \text{exp} \rangle$
$\langle \text{stmts} \rangle$	$::= \epsilon \mid \langle \text{stmt} \rangle \langle \text{stmts} \rangle$
$\langle \text{stmt} \rangle$	$::= \langle \text{simp} \rangle ; \mid \langle \text{control} \rangle \mid \langle \text{block} \rangle$
$\langle \text{simp} \rangle$	$::= \langle \text{lvalue} \rangle \langle \text{asop} \rangle \langle \text{exp} \rangle \mid \langle \text{lvalue} \rangle \langle \text{postop} \rangle \mid \langle \text{decl} \rangle \mid \langle \text{exp} \rangle$
$\langle \text{simpopt} \rangle$	$::= \epsilon \mid \langle \text{simp} \rangle$
$\langle \text{lvalue} \rangle$	$::= \mathbf{ident} \mid \langle \text{lvalue} \rangle . \langle \text{ident} \rangle \mid \langle \text{lvalue} \rangle \rightarrow \langle \text{ident} \rangle$ $\mid * \langle \text{lvalue} \rangle \mid \langle \text{lvalue} \rangle [\langle \text{exp} \rangle] \mid (\langle \text{lvalue} \rangle)$
$\langle \text{elseopt} \rangle$	$::= \epsilon \mid \mathbf{else} \langle \text{stmt} \rangle$
$\langle \text{control} \rangle$	$::= \mathbf{if} (\langle \text{exp} \rangle) \langle \text{stmt} \rangle \langle \text{elseopt} \rangle \mid \mathbf{while} (\langle \text{exp} \rangle) \langle \text{stmt} \rangle$ $\mid \mathbf{for} (\langle \text{simpopt} \rangle ; \langle \text{exp} \rangle ; \langle \text{simpopt} \rangle) \langle \text{stmt} \rangle$ $\mid \mathbf{return} \langle \text{exp} \rangle ; \mid \mathbf{return} ;$ $\mid \mathbf{assert} (\langle \text{exp} \rangle) ;$
$\langle \text{arg-list-follow} \rangle$	$::= \epsilon \mid , \langle \text{exp} \rangle \langle \text{arg-list-follow} \rangle$
$\langle \text{arg-list} \rangle$	$::= () \mid (\langle \text{exp} \rangle \langle \text{arg-list-follow} \rangle)$
$\langle \text{exp} \rangle$	$::= (\langle \text{exp} \rangle) \mid \mathbf{num} \mid \mathbf{true} \mid \mathbf{false} \mid \mathbf{ident} \mid \mathbf{NULL} \mid \langle \text{unop} \rangle \langle \text{exp} \rangle$ $\mid \langle \text{exp} \rangle \langle \text{binop} \rangle \langle \text{exp} \rangle \mid \langle \text{exp} \rangle ? \langle \text{exp} \rangle : \langle \text{exp} \rangle \mid \mathbf{ident} \langle \text{arg-list} \rangle$ $\mid \langle \text{exp} \rangle . \mathbf{ident} \mid \langle \text{exp} \rangle \rightarrow \mathbf{ident} \mid \mathbf{alloc} (\langle \text{type} \rangle) \mid * \langle \text{exp} \rangle$ $\mid \mathbf{alloc_array} (\langle \text{type} \rangle , \langle \text{exp} \rangle) \mid \langle \text{exp} \rangle [\langle \text{exp} \rangle]$
$\langle \text{asop} \rangle$	$::= = \mid += \mid -= \mid *= \mid /= \mid \% = \mid \& = \mid \wedge = \mid \mid = \mid \ll = \mid \gg =$
$\langle \text{binop} \rangle$	$::= + \mid - \mid * \mid / \mid \% \mid < \mid \leq \mid > \mid \geq \mid == \mid !=$ $\mid \&\& \mid \mid \mid \& \mid \wedge \mid \mid \mid \ll \mid \gg$
$\langle \text{unop} \rangle$	$::= ! \mid \sim \mid -$
$\langle \text{postop} \rangle$	$::= ++ \mid --$

The precedence of unary and binary operators is given in Figure 3. Non-terminals are in $\langle \text{angle brackets} \rangle$. Terminals are in **bold**. The absence of tokens is denoted by ϵ .

Figure 2: Grammar of L4

Operator	Associates	Meaning
() [] -> .	n/a	explicit parentheses, array subscript, field dereference, field select
! ~ - * ++ --	right	logical not, bitwise not, unary minus, pointer dereference, increment, decrement
* / %	left	integer times, divide, modulo
+ -	left	integer plus, minus
<< >>	left	(arithmetic) shift left, right
< <= > >=	left	integer comparison
== !=	left	overloaded equality, disequality
&	left	bitwise and
^	left	bitwise exclusive or
	left	bitwise or
&&	left	logical and
	left	logical or
? :	right	conditional expression
= += -= *= /= %=		
&= ^= = <<= >>=	right	assignment operators

Figure 3: Precedence of operators, from highest to lowest

3 L4 Semantics

The static and dynamic semantics for Lab 4 are described in the lecture notes discussing static semantics, dynamic semantics, mutable store, and structs. Some specific points:

- Elaboration will need to deal with the fact that $A[f(x)] += 3$ cannot be elaborated into $\text{assign}(A[f(x)], A[f(x)] + 3)$, because calling $f(x)$ might have an effect, like printing or writing to a pointer.
- You will have to preserve at least some size information during elaboration to generate correct code. We suggest you refer to the current (third) edition of Bryant and O'Hallaron's *Computer Systems: A Programmer's Perspective*, or to the current semester's 15-213 notes. Pay particular attention to the effects of 32 bit operators in the upper 32 bits of the 64 bit registers.
- Struct *definitions* obey scoping rules of the other global declarations, that is, they are available only after their point of definition. However, structs may be *declared* implicitly; see Section 2 of the notes.

- The rules for struct declarations and definitions, mostly inherited from C, are carefully engineered so that it should be possible to compute the size and field offsets of each struct without referring to anything found later in the file. You probably want to store the sizes and field offsets in global tables.
- Like type definitions, struct declarations and definitions can appear in external files.
- We expect your generated code to explicitly capture memory errors, rather than counting on the operating system to notice and raise `SIGSEGV` (11). In order to enforce that, the signal associated with memory errors will be `SIGUSR2` (12). You can raise this signal explicitly with the standard `raise(sig)` function.

The default library for this lab, `15411-14.h0`, is a modification of the previous library that uses 8-byte floating point values stored in pointers rather than 4-byte floating point values stored in integers. We encourage you to use this library in some of your test cases!

For this lab, you do *not* need to lay out structs in a way that is compatible with C, but you are encouraged to do so. You should respect the machine's alignment requirements so that integers and booleans are aligned at least 0 modulo 4 and addresses at least 0 modulo 8, but beyond that, we will not require strict adherence. One reason for this flexibility is that we allow you to represent structs and arrays containing boolean values however you want. Here is what we will potentially test:

- Integers must be stored in memory as 4 continuous bytes (little-endian, as usual on x86-64)
- Pointers must be stored in memory as 8 continuous bytes (little-endian, as usual on x86-64)
- Structs and arrays which contain only ints (and other structs which contain only ints, and so on) must store the ints continuously, in order, where a , the address of the struct or value of the array, is the address of the first struct field or array element, $a + 4$ is the address of the second struct field or array element, and so on.
- Ditto for structs and arrays which contain only pointers, except that $a + 8$ is the address of the second struct field or array element.

4 Project Requirements

For this project, you are required to hand in test cases and a complete working compiler for L4 that produces correct target programs written in Intel x86-64 assembly language.

We also require that you document your code. Documentation includes both inline documentation and a README document which explains the design decisions underlying the implementation along with the general layout of the sources. If you use publicly available libraries, you are required to indicate their use and source in the README file. If you are unsure whether it is appropriate to use external code, please discuss it with course staff.

When we grade your work, we will use the `gcc` compiler to assemble and link the code you generate into executables using the provided runtime environment on the lab machines.

Your compiler and test programs must be formatted and handed in as specified below. For this project, you must also write and hand in at least 20 test programs, at least two of which must fail to compile, at least two of which must generate a runtime error, and at least two of which must execute correctly and return a value.

Test Files

Test programs should have extension `.14` and start with one of the following lines:

<code>//test return <i>i</i></code>	program must execute correctly and return <i>i</i>
<code>//test div-by-zero</code>	program must compile but raise <code>SIGFPE</code>
<code>//test infloop</code>	program must fail to terminate (see below)
<code>//test abort</code>	program must compile and run but raise <code>SIGABRT</code>
<code>//test memerror</code>	program must compile and run and raise <code>SIGUSR2</code>
<code>//test error</code>	program must fail to compile due to an L3 source error
<code>//test typecheck</code>	program must typecheck correctly (see below)
<code>//test compile</code>	program must typecheck, compile, and link (see below)

followed by the program text. In L4, the exceptions defined are `SIGABRT` (6), `SIGFPE` (8), and `SIGUSR2` (12). Note that infinitely recursing programs might raise `SIGBUS` (7), `SIGSEGV` (11), or `SIGALARM` (14) depending on the system; the autograder will attempt to compensate for this for `infloop` tests.

Since the language now supports function calls, the runtime environment contains external functions providing output capabilities (see the runtime section for details). However, we do not check that the output is correct, merely that correct values are eventually returned from library calls.

If the test program `$test.14` is accompanied by a file `$test.h0` (same base name, but `h0` extension), then we will compile the test treating `$test.h0` as the header file. Otherwise, we will treat `../runtime/15411-14.h0` as the header file for all `14` tests, and we will pass that header file to your compiler with the `-I` argument. The `15411-14.h0` header file describes a library for double-precision floating point arithmetic and printing operations; our testing framework will ignore any output performed from the printing operations. You are strongly encouraged, but not required, to write tests that take advantage of this library.

If your tests use a header file that you wrote, your test *must* start with the line `//test error` or `//test typecheck`. If you include a note that explains what your header file and your test is doing, we might change it to allow the autograder to try to compile (`//test compile`) or run (anything else). Only tests utilizing header files that you wrote should begin with `//test typecheck`.

To register a submission with the autograder, commit your tests to the submission repository, changing only files in the `tests` directory, and make sure the string “GRADE” (all uppercase) appears in the commit message.

L4 is a sophisticated and reasonably expressive language. You should be able to write some very interesting tests, perhaps adapted from the programs and libraries you wrote in the 15-122 course on *Principles of Imperative Computation* that uses C0.

Please do *not* submit test cases which differ in only small ways in terms of the behavior exercised. We would like some fraction of your test programs to compute “interesting” functions on specific values; please briefly describe such examples in a comment in the file. Disallowed are programs which compute Fibonacci numbers, factorials, greatest common divisors, and minor variants thereof. Please use your imagination!

Compiler Files

The files comprising the compiler itself should be collected in a directory `compiler/` which should contain a `Makefile`. **Important:** You should also update the `README` file and insert a description of your code and algorithms used at the beginning of this file. Even though your code will not be read for grading purposes, we may still read it to provide you feedback. The `README` will be crucial information for this purpose.

Issuing the shell command

```
% make
```

should generate the appropriate files so that

```
% bin/c0c <args>
```

will run your L4 compiler. The command

```
% make clean
```

should remove all binaries, heaps, and other generated files.

Runtime Environment

Your compiler should accept a single, optional command line argument `-l` which must be given the name of a file as an argument. For instance, we will be calling your compiler using the following command: `bin/c0c -l ../runtime/15411-14.h0 $test.14`. Here, `15411-14.h0` is the header file mentioned above. You may not assume that the header file parses and typechecks correctly.

The `15411-14.h0` header file describes a library for manipulating double-precision floating-point values. The implementation of this library can be found in `lab3/runtime/run411.c`. You should assume the types of the implementation match the types in the header file.

The GNU compiler and linker will be used to link your assembly to the implementations of the external functions, so you need not worry much about the details of calling to external functions. You should ensure that the code you generate adheres to the C ABI for Linux on x86-64. In order for the linking to work, you must adhere to the following conventions

- External functions must be called as named.
- Non-external functions with name *name* must be called `_c0_name`. This ensures that non-external function names do not accidentally conflict with names from standard library which could cause assembly or linking to fail.
- Non-external functions must be exported from (declared to be *global* in) the assembly file you generate, so that our test harness can call them and verify your adherence to the ABI.

The runtime environment defines a function `main()` which calls a function `_c0_main()` your assembly code should provide and export. Your compiler will be tested in the standard Linux environment on the lab machines; the produced assembly must conform to this environment.

In order to satisfy the ABI of libraries that take or return the type `bool`, you must implement this type such that `false` maps to 0 and `true` maps to 1. For the sake of uniformity, it will still suffice to treat boolean values as 32-bit integers; there may be occasion to reconsider this in Lab 4.

Using the Subversion Repository

Handin and handout of material is via the course subversion repository.

The handout files for this course can be checked out from our subversion repository via

```
% svn checkout https://svn.concert.cs.cmu.edu/15411-f15/groups/<team>
```

where `<team>` is the name of your team. You will find materials for this lab in the `lab4` subdirectory. Or, if you have checked out `15411-f14/groups/<team>` directory before, you can issue the command `svn update` in that directory. *Do not* ever include any compiled files or binaries in the repository!

To register a submission with the autograder, commit your code to the subversion repository, changing only files in the `compiler` directory, and make sure the string “GRADE” (all uppercase) appears in the commit message.

What to Turn In

Hand in:

- At least 20 test cases, at least two of which generate an error, at least two of which raise a runtime exception, and at least two of which return a value. The directory `tests/` should only contain your test files and be submitted via subversion as described above. The server will test your test files and notify you if there is a discrepancy between your answer and the outcome of the reference implementation. You may hand in as many times as you like before the deadline without penalty. If you feel the reference implementation is in error, please notify the instructors. The compiled binary for each test case should run in 2 seconds with the reference compiler on the lab machines; we will use a 5 second limit for testing compilers.

Test cases are due **11:00pm on Tuesday, October 27, 2015**.

- The complete compiler. The directory `compiler/` should contain only the sources for your compiler and be submitted via subversion as described above. The autograder will build your compiler, run it on all extant test files, link the resulting assembly files against our runtime system (if compilation is successful), execute the binaries (each with a 5 second time limit), and finally compare the actual with the expected results. You may hand in as many times as you like before the deadline without penalty.

Compilers are due **11:00pm on Tuesday, November 3, 2015**.