

# 15-411 Compiler Design, Lab 3 (Spring 2025)

Seth, Ben, and co. 🐼

Test Cases Due: 11:59 pm, Sunday, February 23, 2025

Compilers Due: 11:59 pm, Sunday, March 2, 2025

## 1 Introduction

The goal of the lab is to implement a complete compiler for the language L3. This language extends L2 with the ability to define functions and call them. This means you will have to change all phases of the compiler from the second lab. One can write some interesting recursive and iterative functions over integers in this language. Correctness is still paramount, but performance starts to become a bit more of an issue as we run larger and more interesting test cases.

### L3 Global Declarations

In L2, we only had one top-level declaration (`main`), and our program consisted of a single function body. In L3, we allow for multiple top-level/global declarations (`gdecls`). These global declarations should be familiar from C0/C:

- `fdecl`: declares a function (i.e. a function prototype).

$$\langle \text{fdecl} \rangle ::= \langle \text{type} \rangle \textbf{ident} \langle \text{param-list} \rangle ;$$

- `fdefn`: defines a function.

$$\langle \text{fdefn} \rangle ::= \langle \text{type} \rangle \textbf{ident} \langle \text{param-list} \rangle \langle \text{block} \rangle$$

- `typedef`: defines a transparent alias for a type, as in C0/C.

$$\langle \text{typedef} \rangle ::= \textbf{typedef} \langle \text{type} \rangle \textbf{ident} ;$$

An L3 program thus consists of a list of `gdecls`.

## L3 Header Files

Starting in L3 we would like to be able to link programs against a separate runtime (e.g. C0 standard library functions). But your compiler needs to know what external functions are available, and what their interface is. For this purpose we add the notion of “header files”, much like in C0/C.

- L3 programs may be compiled with up to one *header file*
- Header files have the same syntax as program files
- Header files **CANNOT contain** function definitions, only function *declarations* and typedefs.
- Your compiler **cannot** assume that the header file parses and typechecks correctly.
- Your compiler can assume each function declared in the header will be available as an external symbol of the same name, when linking the program together.

The exact constraints on how global declarations interact with those in the L3 source file are complicated. See section 3 for details. Also, the x86 output needs to treat external functions (those declared in the header file) differently from functions declared/defined in the L3 source file. See section 5 for details.

## 2 L3 Syntax

The lexical specification of L3 remains mostly unchanged from that of L2, with the addition of a comma (,) as a lexical token. The syntax of L3 is a superset of L2, as presented in Figure 1. A quick summary of the changes:

- New syntax rules for global declarations, including function declarations and typedefs.
- Function calls as expressions.
- A new statement `assert(e)` with its meaning from C0 and C, aborting the program if *e* evaluates to false.
- The `void` syntactic construct, used in place of a return type for functions that do not return a value. <sup>1</sup>

Ambiguities in the grammar are resolved according to the operator precedence table in Figure 2 and the rule that an `else` provides the alternative for the most recent eligible `if`.

---

<sup>1</sup>Note: Some sources will refer to `void` as a type; this is rather inaccurate in context of C0, since one cannot declare a variable of “type” `void`. In C0, `void` can only be used as a function return type; using `void` anywhere else should be considered a parse error. Moreover, `void` in C0 is semantically different from type theoretical empty type, which is also sometimes called `void`.

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

The precedence of unary and binary operators is given in Figure 2. Non-terminals are in  $\langle \text{angle brackets} \rangle$ . Terminals are in **bold**. The absence of tokens is denoted by  $\epsilon$ .

Figure 1: Grammar of L3

Operator	Associates	Meaning
()	n/a	explicit parentheses
! ~ - ++ --	right	logical not, bitwise not, unary minus, 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 2: Precedence of operators, from highest to lowest

### 3 L3 Statics

The static semantics of L3 is substantially more complex than that of L2. Much of the non-uniform behavior of C has been replicated in L3, especially with respect to name collisions and shadowing. Design your implementation carefully, to ensure it matches the (quite complex) behavior described in this section.

#### Function Signatures

A *function signature* summarizes the argument types of a function, as well as the return type, if the function has one.

$$(\tau_1, \tau_2, \dots, \tau_n) \rightarrow [\tau]$$

In the void return case, we write  $(\tau_1, \dots, \tau_n) \rightarrow \cdot$  to indicate the lack of a return type.

#### Declaration Context $\Delta$

The context  $\Delta$  (*delta*, for *definitions* and *declarations*) will be used to track function declarations and definitions. We have a judgment that  $f$  has a given signature in  $\Delta$ , written:

$$\Delta \vdash f : (\tau_1, \dots, \tau_n) \rightarrow \tau$$

We also have a judgment that  $f$  is defined in  $\Delta$ , written:

$$\Delta \vdash f \text{ defined}$$

We will write  $\Delta, f : (\tau_1, \dots, \tau_n) \rightarrow \tau$  to add a signature to the context, and  $\Delta, f \text{ defined}$  to mark that  $f$  is defined. There are some restrictions on this, detailed in the relevant sections below.

#### Typedef Context $\Omega$

During typechecking we must keep track of which identifiers have been typedef'd, and what their definition is.

For this we introduce  $\Omega$  (*omega* for *om...nevermind*, this one makes no sense).  $\Omega$  is a mapping from typedef'd symbols to their definitions. We have a judgment that, in context  $\Omega$  a type  $t$  “resolves” to  $\tau$ , where  $\tau \in \{\mathbf{int}, \mathbf{bool}\}$ :

$$\Omega \vdash t \rightsquigarrow \tau$$

If  $t$  is **int** or **bool**, then  $t = \tau$ , while if  $t$  is an identifier, we look it up in  $\Omega$ :

$$\frac{\tau \in \{\mathbf{int}, \mathbf{bool}\}}{\Omega \vdash \tau \rightsquigarrow \tau} \qquad \frac{\Omega(t) = \tau}{\Omega \vdash t \rightsquigarrow \tau}$$

Note that typedefs can typedef a symbol to another typedef'd symbol, but for simplicity we will say that  $\Omega$  should map symbols to their fully resolved type (i.e. **int** or **bool**). Assuming  $t$  is a symbol and  $\tau$  is a resolved type, we will write  $\Omega, t \mapsto \tau$  for the context with updated mapping.

## Program Typechecking

An L3 program consists of a list of **gdecls** from a header file, and a list of **gdecls** from an L3 source file. Declarations are processed in the file order, from top down, with header declarations processed before program declarations. Note that you cannot assume the header typechecks; its declarations are subject to the same restrictions as program declarations, with the additional restriction that function *definitions* are not allowed in headers.

A global declaration  $g$  typechecks in context  $\Delta; \Omega$  iff the following judgment is derivable:

$$\Delta; \Omega \vdash g \rightsquigarrow \Delta'; \Omega'$$

The  $\Delta'; \Omega'$  represent the updated contexts, capturing that each **gdecl** adds information to the context.

### Notes:

- The **main()** function is considered implicitly declared with its expected function signature BEFORE the declarations in the L3 source, and can therefore be forward-referenced in source programs.
- Header functions are considered defined before the L3 source.

## Typedefs

Typedefs are relatively straightforward; if valid, the typedef adds a mapping to  $\Omega$ .

$$\frac{t' \notin \text{Dom}(\Delta) \quad t' \notin \text{Dom}(\Omega) \quad \Omega \vdash t \rightsquigarrow \tau}{\Delta; \Omega \vdash \text{typedef } t \ t' \rightsquigarrow \Delta; \Omega, t' \mapsto \tau}$$

The premises require that the symbol  $t'$  is not already declared as a typedef NOR as a function.

Once an identifier is defined as a typedef, **ALL** later identifiers in the program (function names, parameter/local variable names, other typedefs, ...) cannot collide with the typedef'd identifier. We will reiterate this in all relevant rules below.

## Function Declaration

Function declarations insert a new declaration to  $\Delta$ . They have somewhat complex typing constraints:

- The name must not be a typedef
- Parameter names must not be typedefs
- Parameter list should have no duplicate parameter names

- If a function with the same name has been declared, the prior signature must match the new declaration (same argument and return types, though not necessarily the same parameter names).

This last rule is quite important – a program can re-declare any function, even those declared in the header, if the declaration signature matches the previous/existing signature for that function. The constraints are summarized by the following two rules:

$$\begin{array}{c}
\frac{f \notin \text{Dom}(\Delta) \quad f \notin \text{Dom}(\Omega) \quad [\Omega \vdash t \rightsquigarrow \tau] \quad \Omega \vdash t_i \rightsquigarrow \tau_i \quad \forall i \neq j, x_i \neq x_j}{\Delta; \Omega \vdash [t] \ f(t_1 \ x_1, \dots, t_n \ x_n); \rightsquigarrow \Delta, f : (\tau_1, \dots, \tau_n) \rightarrow [\tau]; \Omega} \\
\\
\frac{\Delta \vdash f : (\tau_1, \dots, \tau_n) \rightarrow [\tau] \quad [\Omega \vdash t \rightsquigarrow \tau] \quad \Omega \vdash t_i \rightsquigarrow \tau_i \quad \forall i \neq j, x_i \neq x_j}{\Delta; \Omega \vdash [t] \ f(t_1 \ x_1, \dots, t_n \ x_n); \rightsquigarrow \Delta; \Omega}
\end{array}$$

The square brackets are intended to indicate that these rules apply for functions with no return type; the premises in square brackets simply are ignored in the case of no return type.

## Function Definition

Function definitions simultaneously declare and define a function. The declaration derived from a definition should typecheck with the same restrictions as above. In particular, if a function is already declared, the function signature must match the prior declaration.

A definition also (as the name suggests) defines the function. The function must not have been previously defined, and the function body is a block which must typecheck as a statement returning the expected return type:

$$\frac{\Delta; \Omega \vdash [t] \ f(t_1 \ x_1, \dots, t_n \ x_n); \rightsquigarrow \Delta'; \Omega \quad \Delta \vdash \neg(f \text{ defined}) \quad [\Omega \vdash t \rightsquigarrow \tau] \quad \Omega \vdash t_i \rightsquigarrow \tau_i \quad \Delta'; \Omega; x_1 : \tau_1, \dots, x_n : \tau_n \vdash s : [\tau]}{\Delta; \Omega \vdash [t] \ f(t_1 \ x_1, \dots, t_n \ x_n) \ s \rightsquigarrow \Delta', f \text{ defined}; \Omega}$$

Note the subtleties in the last premise:  $f$  is considered declared in the body  $s$ , thus the updated context  $\Delta'$  is considered when typechecking  $s$ . Also, the initial typing context  $\Gamma$  consists of the parameters of  $f$ . As with the declaration rules, this rule still applies when  $f$  has no return type, with one fewer premise to consider.

In L3, a function is visible after its first declaration, and within its definition. This means a single recursive function does not require a declaration, but two mutually recursive functions require at least one forward declaration.

## Statements & Expressions

The typechecking for statements and expressions is largely the same as in L2. As we will see, the statement judgment needs to be augmented with the  $\Delta, \Omega$  contexts from the global declarations, and the expression judgment needs to be augmented with the  $\Delta$  context.

First, the rule for declarations is now different, because variables may not conflict with typedef'd symbols:

$$\frac{x \notin \text{Dom}(\Omega) \quad x \notin \text{Dom}(\Gamma) \quad \Delta; \Omega; \Gamma, x : \tau \vdash s : [\rho]}{\Delta; \Omega; \Gamma \vdash \text{decl}(\tau, x, s) : [\rho]}$$

There is no restriction on local variables shadowing function names. Note that, in particular, since function parameters are in scope for the entire function, they cannot be shadowed by declarations.

We also have two (three?) new statements, namely **non-int** returns and asserts. Their rules:

$$\frac{}{\Delta; \Omega; \Gamma \vdash \text{return}; : \cdot} \quad \frac{\Delta; \Gamma \vdash e : \rho}{\Delta; \Omega; \Gamma \vdash \text{return}(e) : \rho} \quad \frac{\Delta; \Gamma \vdash e : \text{bool}}{\Delta; \Omega; \Gamma \vdash \text{assert}(e) : [\rho]}$$

Next, function calls are now a valid expression, so we need a typechecking rule for them. There is a somewhat odd constraint that the function symbol being called must not have been introduced as a variable in scope:

$$\frac{f \notin \text{Dom}(\Gamma) \quad \Delta \vdash f : (\tau_1, \dots, \tau_n) \rightarrow \tau \quad \Delta; \Gamma \vdash e_i : \tau_i}{\Delta; \Gamma \vdash f(e_1, \dots, e_n) : \tau}$$

Note this rule is only applicable to functions with a return type. Void-returning functions cannot appear in arbitrary expressions. Instead, they can only appear as top-level statements. So, we will add a separate statement rule to allow this:

$$\frac{f \notin \text{Dom}(\Gamma) \quad \Delta \vdash f : (\tau_1, \dots, \tau_n) \rightarrow \cdot \quad \Delta; \Gamma \vdash e_i : \tau_i}{\Delta; \Omega; \Gamma \vdash f(e_1, \dots, e_n) : [\rho]}$$

Note that the function name  $f$  must not be in the typing context  $\Gamma$  when the function is called, and variables are introduced to the typing context  $\Gamma$  when they are declared. This means, for example, that `int f = f();` is allowed in L3, but `int f; f = f();` is not. This is a somewhat odd behavior, and it may not be consistent with C.

## Defined & Used Functions

An important check for the safety of your program is ensuring that all functions *used* in the program are *defined* in the program:

- External functions (= at least one declaration in a header file) are considered defined before the start of the L3 source.
- A function is considered used if it is called ANYWHERE in the program (including unreachable/dead code).
- A function need not be defined at the point in code where it is used (see Statements & Expressions subheader), merely by the end of the program.



- `main()` is implicitly considered used in every program, and thus must always be defined (unlike L1/L2, the presence of `main()` is no longer a syntactic restriction but rather a typechecking restriction!)
- A function that is declared, but not used anywhere in the program, DOES NOT need to be defined.<sup>2</sup>

## Return Checking

As in L2, the typechecker must ensure that functions with return types must return a value on every control flow path. However, functions with no return type (i.e. void functions) **do not** need to return on every path. They are assumed to implicitly return at the end of the function body.

## Variable Initialization

As in L2, the typechecker must ensure that all local variables are defined before they are used. Function parameters are considered *defined* at the beginning of the function body (they will have values when the execution of the body of a function commences).

## 4 L3 Dynamic Semantics

The dynamic semantics of L3 directly extends the dynamic semantics from L2.

Function calls  $f(e_1, \dots, e_n)$  are very similar to their counterparts in C with the following significant difference: they must evaluate their arguments from left to right before passing the resulting values to  $f$ . The same is true for other operators that evaluate their arguments. For example, in  $e_1 + e_2$ , we must evaluate  $e_1$  before  $e_2$ . Since L3 has several kinds of effects (arithmetic exception, nontermination, abort, and even output), this specification now becomes significant (when it was not observable in L2).

The new statement `assert( $e$ )` first evaluates  $e$ . If the result is `true`, the assertion succeeds and we just continue with the next statement. If  $e$  is false, it raises the **SIGABRT** exception (6). In your code you can achieve this by calling the function `abort()` which takes no arguments.

Expressions may appear as statements, because we now have the concept of expressions with side-effects. These side-effects are quite simple in L3: arithmetic exceptions, program abort, nontermination, and output done by library functions. An error can also arise when a program runs out of stack space. Different systems handle this differently. The autograder runs on an Ubuntu system that seems to consistently raise **SIGBUS** (7) when it runs out of stack space, but other systems may raise **SIGSEGV** (11). You do not need to handle stack overflows specially.

---

<sup>2</sup>This is largely due to the behavior of C, where the design of the language is supposed to allow separate compilation of individual sections of the program.

## 5 Runtime Environment

Programs compiled with a header file are allowed to assume the declared symbols will be present in the runtime, and that the types declared in the header are accurate to the runtime implementation.

The GNU compiler and linker will be used to link your assembly to the implementations of the 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 in the header file.
- 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.

NOTE: If a header file declares a function with name `_c0_...`, your compiler is allowed to crash after typechecking, since the program may not be compilable in a way that adheres to these conventions.

For Mac users, the function names need to be prefixed with an extra underscore.

- 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 calling conventions. Use the `.globl <label>` directive.

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 docker containers; 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 when calling external functions, though you are free to implement your own representations internally if you wish. 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 or 5.

## 6 Command-Line Interface

Your compiler is also expected to recognize a flag `-t` which, when present on the command line, stops the compiler immediately after typechecking and before the rest of the compiler runs. The exit code of your compiler should indicate success (0) if the code is well-formed, and failure (1) otherwise. If your compiler indicates success when run with `-t`, then it should be able to compile the file without further errors. You are encouraged to implement `-O0` and `-O1` flags, representing different tradeoffs between compilation speed and runtime performance of your compiled program. The autograder will invoke your compiler with whichever flag you have set as the default. You are additionally encouraged to perform analysis on the file you are compiling to determine whether your compiler can feasibly perform more expensive optimizations (such as register allocation) within the specified limit. We recommend falling back to faster, if less optimal, optimization strategies when necessary.

Your compiler should accept an 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/15411-13.h0 $test.13`. Here, `15411-13.h0` is a header file.

If your compiler detects any (compile-time) errors in the source program, it should exit with a non-zero return code. If compilation succeeds and target code is generated, the compiler should then exit with a return code of 0.

## 7 Testing

Test programs have extension `.13` 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 SIGFPE (8)
<code>//test abort</code>	program must compile and run but raise SIGABRT(6)
<code>//test error</code>	program must fail to compile due to an L3 source error
<code>//test typecheck</code>	program must typecheck correctly
<code>//test compile</code>	program must typecheck, compile, and link

If the test program `$test.13` is accompanied by a file `$test.h0` (same base name, but `h0` extension), then we compile the test treating `$test.h0` as the header file. Otherwise, we treat `../runtime/15411-13.h0` as the header file for all `13` tests, and we will pass that header file to your compiler with the `-l` argument. The `15411-13.h0` header file describes a library for floating point arithmetic and printing operations. The implementation of this library can be found in `lab3/runtime/run411.c`. Our testing framework will ignore any output performed from the printing operations.

Tests that use a `.h0` header file might typecheck but fail to link because they refer to functions that aren't provided by the system: the header file can even describe a library that can't possibly be implemented (see `busy.13` and `busy.h0` for an example).

Tests which use a `.h0` header file might also take advantage of functions provided in `libc` or `libgcc` (see `rand.13` and `rand.h0` for an example). But many of these functions can cause the test cases to behave badly. Therefore, tests that have a custom header file *must* start with the line `//test error` or `//test typecheck`, and will not be executed by the autograding harness.

## 8 Submission

For this project, you are expected to hand in a set of 10 test cases and a complete working compiler for L3 that produces correct target programs written in Intel x86-64 assembly language.

A late day is computed based on the active submission on Gradescope. So if your active submission is past **11:59 pm** on the due dates then it will use a late day.

You may turn in code and have it autograded as many times as you like, without penalty. In fact, we encourage you to hand in to verify that the autograder agrees with the driver results that you use for development, and also as insurance against a last-minute rush. The submission with the highest grade will count. You will submit:

**Before Sunday, February 23th, 11:59 pm** Ten (10) test cases, at least 1 for each of the first 5 categories mentioned in section 7. All of your test cases should utilize new language features introduced in L3. You should submit to **Test 3** assessment on Gradescope. The directory tests should only contain your test files. Please name your test files with the extension `.l3` prefixed with your team name in lower case.

**Before Sunday, March 2nd, 11:59 pm** The complete compiler. You will submit to the **Lab 3** assessment on Gradescope. The directory `compiler/lab3` should contain only the sources for your compiler. The autograder will build your compiler, run it on all existing test files, link the resulting assembly files against our runtime system (if compilation is successful), execute the binaries, and finally compare the actual with the expected results.

The files comprising the compiler itself should be collected in a directory `compiler/` which should contain a `Makefile`.

Issuing the shell command

```
% make lab3
```

from within the `compiler` directory should generate the appropriate files so that

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

will run your L3 compiler. The command

```
% make clean
```

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

**Important:** You should also update the `README` file to include

1. a description of your compiler's structure.
2. major design decisions or specific algorithms utilized.
3. any publicly available external libraries you used.

Even though we will not be grading you on your code quality, we will be reading your code for the code review (see below), and to provide you feedback. The README will be crucial for us to understand what is going on in your compiler.

You are permitted to make use of external, publicly available libraries in your compiler, in accordance with the course policy. If you are unsure whether it is appropriate to use an external library, please discuss it with course staff.

## Autograded Scoring

You may turn in code and have it autograded as many times as you like, without penalty. In fact, we encourage you to hand in to verify that the autograder agrees with the driver results that you use for development, and also as insurance against a last-minute rush. The submission with the highest grade will count.

Your score for each submission is computed as follows:

$$\begin{aligned}\text{subtotal} &= 25 * \frac{\text{passed basic tests}}{\text{total basic tests}} + 65 * \frac{\text{passed large \& only tests}}{\text{total large \& only tests}} \\ \text{compiler} &= \text{subtotal} - (1 \text{ point per compiler failure}) - (0.1 \text{ points per executable timeout})\end{aligned}$$

Your total score for the lab is computed as follows:

$$\begin{aligned}\text{test cases} &= 10 * \frac{\min(\text{valid test cases}, 10)}{10} \\ \text{total} &= \text{test cases} + \text{compiler}\end{aligned}$$

## Code Review

We will be holding code review sessions during the week of March 10. The purpose of the code review is to make sure that both partners understand the whole compiler, from the lexing and parsing all the way to x86-64 generation. We will ask questions pertaining to any part of the compiler, and both partners are expected to be able to answer them. The instructors will be looking through your code and READMEs before meeting with you, so make sure that the READMEs are updated to reflect the organization of your code.

**We will post sign ups for slots on Piazza as soon as we finalize the schedules.** Be sure to check Piazza frequently for updates.

## 9 Notes and Hints

### Elaboration

We *highly* recommend using an explicit elaboration pass, or even multiple elaboration passes, rather than transforming source code on the fly during parsing. A well-designed elaboration pass can significantly streamline your compiler and reduce the amount of work it takes to implement later phases of your compiler.

### Static Checking

Follow the rules for static analysis closely. You should take some care to produce useful error messages when typechecking; this phase is notoriously prone to small or subtle bugs in L3, and good error messages can be a life saver when debugging your typechecker.

### Compiling Functions

It is not a strict requirement, but we recommend compiling functions completely independently from each other, taking care to respect the calling conventions and making no other assumptions. Interprocedural program analysis and optimization is difficult and, if you do it at all, is better left to a later lab.

### Calling Conventions

Your code must strictly adhere to the x86-64 calling conventions. Please refer to the course webpage for resources on the ABI and calling conventions. A particular point of note: `%rsp` must be 16-byte aligned. GCC often ignores this rule when it isn't actively using floating point numbers, because the requirement is only consequential when the floating point stack is in use.