Nils M Holm

PRACTICAL COMPILER CONSTRUCTION

64-Bit Edition

No so-called “AI” was used in the making of this book.

http://t3x.org/reload64/

Contents

Preface **7**
Preface (Second Edition) **9**
Preface (First Edition) **11**
Acknowledgments **15**
Part I — Introduction
Compilers **19**
How Does a Compiler Work? 21
Phases of Compilation 23
A Simplified Model 30
The Peculiarities of C **33**
Rules of the Game **37**
The Source Language 37
The Object Language 40
The Runtime Library 41
Part II — The Tour
Definitions and Data Structures **45**
Definitions 45
Global Data 51
Function Prototypes 57
Utility Functions **61**
Error Handling **65**
Lexical Analysis **69**
A Brief Theory of Scanning 79
The Implementation of the Scanner 83
Symbol Table Management **93**
Syntactic and Semantic Analysis **105**
A Brief Theory of Parsing 106
Mapping Grammars to Parsers *111*

Expression Parsing 117
Constant Expression Parsing 150
Statement Parsing 155
Declaration Parsing 172
Preprocessing 193
Code Generation 201
A Brief Theory of Code Generation 201
The Code Generator 204
Framework 208
Load Operations 210
Binary Operators 211
Unary Operators 219
Jumps and Function Calls 220
Data Definitions 223
Increment Operators 225
Switch Table Generation 230
Store Operations 231
Rvalue Computation 235
Target Description 237
The x86_64 Target 240
Framework 241
Load Operations 242
Stack Operations 245
Binary Operations 245
Unary Operations 249
Increment Operations 250
Jumps and Branches 256
Store Operations 258
Functions and Function Calls 259
Data Definitions 261

The Compiler Controller **263**

Part III — The Runtime Environment

The Runtime Startup Module **279**

The System Calls 283

Using the Host System Library 293

Header Files 296

The System Call Header *296*

The Setjmp Header *297*

The Signal Header *297*

The Runtime Library **299**

Library Initialization 299

Standard I/O 300

The stdio.h Header *301*

Required Stdio Functions *307*

Utility Library 336

The stdlib.h Header *336*

Required Stdlib Functions *337*

String Library 346

The string.h Header *346*

Required String Functions *347*

Character Types 351

The ctype.h Header *351*

The Ctype Functions *351*

The errno.h Header 353

Part IV — Beyond SubC

Code Synthesis **357**

Instruction Queuing 358

CISC versus RISC *368*

Comparisons and Conditional Jumps *369*

Register Allocation 371

Cyclic Register Allocation *377*
Optimization **381**
Peephole Optimization 381
Expression Rewriting 385
Constant Expression Folding *392*
Strength Reduction *395*
Common Subexpression Elimination 400
Emitting Code from an AST 405
Part V — Conclusion
Bootstrapping a Compiler **409**
Design *409*
Implementation *410*
Testing *413*
Having Some Fun 415
Appendix
Where Do We Go from Here? **421**
Piece of Cake 421
This May Hurt a Bit 422
Bring 'Em On! 424
(Sub)C Primer **427**
Data Objects and Declarations 427
Void Pointers *431*
Expressions 431
Pointer Arithmetic *434*
Statements 435
Functions 438
Prototypes and External Declarations 439
Preprocessor 440
Library Functions 442
x86_64 Assembly Primer **443**

Registers 443
Assembler Syntax 445
Assembler Directives 446
Sample Program 447
Addressing Modes 448
Register *448*
Immediate *448*
Memory *448*
Indirect *449*
Instruction Summary 450
Move Instructions *451*
Arithmetic Instructions *452*
Branch and Call Instructions *453*
List of Figures **455**
Bibliography **457**
Index **459**

Preface

In the past years, a few people have asked me to update Practical Compiler Construction and make its compiler target the 64-bit x86 processor instead of the 386. This is what this edition does. The era of 32-bit processors is coming to an end, and many operating system distributions have announced the termination of 32-bit legacy support, which would render the original text an exhibit of rather academic value.

Then, this edition also targets the more popular Linux operating system instead of FreeBSD. Runtime support for FreeBSD is still contained in the SubC source code archive, but it is no longer discussed in the text.

Adapting the text to a 64-bit processor involved much more than just replacing a few processor names and register names. Some assumptions about the target processor had to be abandoned, and then there were lots of places that referred to characteristics of the 386 in one way or another. Finally, the runtime startup module had to be replaced entirely.

The runtime startup module is the part of the compiler that is not only the least portable, it is also kind of fragile, because changes in the underlying operating system can render it non-functional. The text also addresses this problem.

Then, I have worked over some parts of the prose and clarified a few formulations. Nothing far-reaching, but, after having written a few novels in the past, I thought that some parts could need a little bit of work.

Finally, a word about so-called “AI”. No part of this book has been made using LLMs or any other kind of so-called “AI”. Not the code, not the text, and not the illustrations. The entire book is 100% human-made. From a human being for human beings.

Enjoy!

Nils M Holm
October 2026

Preface (Second Edition)

SubC and Practical Compiler Construction were an unexpected success. It is probably still the best-selling one of my books, and people have done interesting things with it, which is probably the greatest praise that the work of a human being can attract. Since the publication of the first edition, the SubC compiler has been ported to Linux, OpenBSD, Mac OS (Darwin), Windows, and DOS. It also supports the ARMv6, 8086, and x86_64 CPUs now, and some of the optimizations explained in this book have been added. The code of the new compiler, however, has evolved substantially and is now quite a bit more complex than the code described initially.

This text still explains the original SubC compiler as it first appeared in the first edition of the book. However, there are some differences. All bugs that have been discovered since the publication of the first edition have been fixed. The dependency on the host system's standard C library has been removed. The system call interface has been made more Unix-like in order to increase the compatibility with existing code.

Many lessons have been learned during the further development of SubC. One is that the C language is really complex, and writing a compiler for even a tiny subset of C is a major undertaking with many small details waiting to be overlooked. This is why the second edition still discusses the original SubC compiler instead of one of its descendants. Another lesson is that emulating **`signal()`** is probably not worth the pain that it inflicts, so it has been removed from this edition. (You can still have it by linking against the host system's C library.)

This book has been available in PDF format only since 2016, because the print master no longer adhered to evolving standards. This was mostly caused by typesetting the book in TeX, which is simply too complicated for my little brain (i.e., this is not a fault of TeX, but a limitation of my understanding of it). So the obvious remedy was to change the typesetting to TROFF, which has never let me down in all the decades of writing. I started this endeavor in

2017, but then lost interest along the way. However, people kept asking about the print edition and so I eventually decided to finish the work on the second edition. It was a much greater effort than originally expected, but I think it was worth it. The second edition is a better book than the first one and, as of now, the book is back in print!

Nils M Holm
May 2022

Preface (First Edition)

A lot of time has passed since the publication of my previous compiler textbook, *Lightweight Compiler Techniques* (LCT), in 1996. I had sold some spiral-bound copies of the book on my own back then and later published a nice soft-cover edition on Lulu.com in 2006. To my surprise, the paperback sold more copies than I had sold when the book was first published, and a few people even asked me if I planned to finish the final chapter on code generation and optimization. However, I never found the motivation to do so. Until now.

Instead of polishing and updating LCT, though, I decided to create an entirely new work that is only loosely based on it. As you probably know, sometimes it is easier to toss all your code to the bin and start from scratch. So while this book is basically an updated, extended, improved, greater, better, funnier edition of LCT, it differs from it in several substantial points:

Instead of the T3X language, which never caught on anyway, it uses the wide-spread C programming language (ANSI C, C89) as both the implementation language and the source language of the compiler being implemented in the course of the text.

It includes much more theory than LCT, thereby shifting the perspective to a more abstract view and making the work more general and less tightly bound to the source language. However, all theory is laid out in simple prose and immediately translated to practical examples, so no prior knowledge in the field is required.

The main part of the book contains the complete code of the compiler, including its code generator, down to the level of the emitted assembly language instructions. It also describes the interface to the operating system and the runtime library—at least those parts that are needed to implement the compiler itself, which is quite a bit!

Then there is a self-contained part dealing with code synthesis and optimization. The approaches discussed here are mostly

taken from LCT, but explained in greater detail and illustrated with lots of diagrams and tables. This part covers topics like abstract program rewriting, register allocation, machine-dependent peephole optimization, and common subexpression elimination.

The final part of the book describes the bootstrapping process that is inherent in the creation of a compiler from the scratch, i.e. it explains where to start and how to proceed when creating a new compiler for a new or an existing programming language.

The intended audience of this book are programmers who are interested in writing a compiler. Familiarity with a programming language is required, working knowledge of the C language is helpful, but a C primer is included in the appendix for the brave of heart. This book is intended as a practical “first contact” to the world of compiler-writing. It covers theory where necessary or helpful, explains the terminology of compilers, and then goes straight to the gory details.

The book is not a classic teaching book. It does not take a break and requires you to pace yourself. Be prepared to read it at least twice: once concentrating on the prose and another time concentrating on the code. The goal of the book is to bring you to a point where you can write your own real-world compiler with all the bells and whistles.

The code spread out in the book is a compiler for a subset of the C programming language as described in the second edition of the book *The C Programming Language* (TCPL2) by Brian W. Kernighan and Dennis M. Ritchie. The compiler targets the plain 386 processor without any extensions. It requires the GNU 386 assembler to compile its output. Its runtime environment has been designed for the FreeBSD[1] operating system, but I suspect that it will also compile on Linux—or should be easy to port, at least.

The complete machine-readable source code to all versions of the compiler can be found at **`http://www.t3x.org`**.

---

[1] See **`http://www.freebsd.org`**.

All the code in this book is in the public domain, so you can do whatever you please with it: compile it, use it, extend it, give it to your friends, sell it, whatever makes you happy.

Nils M Holm
March 2012

Acknowledgments

The following people have contributed to SubC and Practical Compiler Construction by pointing out mistakes, porting SubC to different operating systems and processors, turning it into something else entirely, and giving all kinds of valuable feedback.

Ori Avtalion, Federico Bianchi, Bill Buzbee, Liberto Camus, Paul Edwards, Andy Elvey, Michael Engel, Alexey Frunze, Josh Helzer, Tim Holliefield, Nikolaos Kavvadias, Volkmar Klatt, Jean-Marc Lienher, Romain LWPB, Teran McKinney, Minux, Jens Nyberg, Rick Ohnemus, Kevin T. Ryan, Bakul Shah, Warren Toomey, Quan Tran, Steve Ward.

I apologize in case I have forgotten anyone!

Part I
Introduction

Compilers

Traditionally, most of us probably understand a *compiler* as a stand-alone program that reads source code written in a formal language and generates an executable program for a specific computer system. Indeed, this is how most compilers work, even today. In the past, it has been common to refer to interactive programming language implementations as “interpreters”. Even the distinction between “interpreted” and “compiled” languages has been made. These boundaries are highly artificial, though.

Many languages that have been referred to as interpreted languages in the past have gotten compilers long ago. BASIC and LISP, for instance, both of which were first implemented as “interpreters”, soon evolved into implementations that generated native machine code. Most modern LISP systems compile code to machine code behind the scenes, even in interactive mode. In the age of just-in-time compilation, interactive program execution is no longer an evidence for what we commonly call “interpretation”.

What is “program *interpretation*” anyway? It is nothing but the process of performing certain tasks based on the statements written down in a formal program. Does it really matter whether this interpretation is performed by a human being, by another program on the same or a different machine, or by a program that happens to be cast in silicon (a.k.a. a CPU)?

I think it is time to define the terms “compilation” and “interpretation” in a more precise way. In this book, *compilation* will refer to the process of translating a formal language S to a formal language O by a program C. “S” is called the *source language* of the process, “O” is called the *object language*, and “C” the compiler. Interpretation is simply a synonym for program execution, be it by a hardware machine, a virtual machine, or a human being.

Given this definition, even the first BASIC interpreters were in fact compilers *and* interpreters at the same time. The compiler part translated the input language BASIC to some internal form that

was suitable for efficient interpretation. The interpreter part gave meaning to the internal form by "running"—i.e. executing—it. The stages of the process were tightly coupled in this case, so the compiler and interpreter were inseparable parts of the language implementation but, nevertheless, both stages were present.

Another rather useless classification is that of compiled and interpreted languages. Even primordial BASIC is compiled to a token stream and not stored in memory as literal text, and even the output of the most clever and aggressive optimizing compiler is interpreted by a CPU in the end. And then any language that can be interpreted can be compiled as well. If we want to make this differentiation at all, then "interpreted" (by software) versus "compiled" (to native code) is a property of the *implementation* and not one of the language.

Even a code beautifier—a program that improves the formatting of an input program—is a compiler. It just happens to have the same source and object language. The same is true for a code obfuscator, a program that scrambles source code while preserving its meaning.

Compilers come in various flavors, from simple translators that turn their input into a different representation, like early BASIC tokenizers, to full-blown stand-alone optimizing compilers with extensive runtime libraries. However, they all employ the same basic principles.

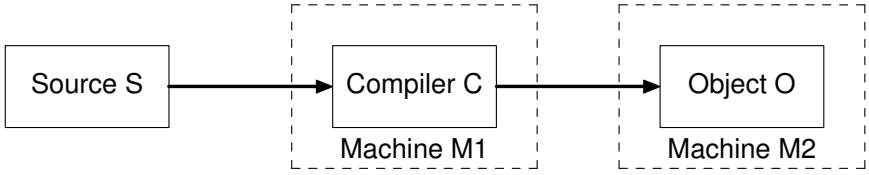

Figure 1: Compiler model

A compiler is a program C written in language L running on machine M_1 that translates source language S to object language O, which may be executable on a machine M_2. S, O, and L may be all different, all the same, or any variation in between. Both M_1 and M_2 may be actual hardware or virtual machines. Here is a non-conclusive list of variations:

- When $S \neq O$, then the compiler is an “ordinary” compiler.
- When $S = L$, then the compiler is *self-hosting*. [1]
- When $M_1 \neq M_2$, the compiler is a *cross-compiler*. [2]
- When $S = O$, the compiler is a code beautifier, obfuscator, etc.
- When O is not a formal language, C is a code analyzer. [3]

In this book, we will concentrate on the first and most common case. We will have a thorough look at a compiler for a subset of the ANSI C programming language (C89) for the 64-bit x86 architecture. We will explore the stages of compilation from source code analysis to native code generation. We will also cover the runtime library and the interface to the operating system (Linux).

The compiler has been designed to be simple and portable. Understanding its principles does not require any prior knowledge other than mastery of C or a comparable programming language. For those who are not familiar with the language, a terse C primer can be found in the appendix (pg. 427). The compiler itself is written in pure ANSI C, so it should compile fine on any platform providing a preexisting C compiler. Porting it to other popular free Unix clones (e.g., the various BSD systems) should be easy. Once bootstrapped, the compiler is capable of compiling itself.

How Does a Compiler Work?

All but the most simple compilers are divided into several *stages* or *phases*, and each of these phases belongs to one of the two principal parts of the compiler called its *front-end* and its *back-end*. The front-end is located on the input side of the compiler. It analyzes the source program and makes sure that it is free of formal errors, such as misspelled keywords, wrong operator symbols, unbalanced parentheses or scope delimiters, etc. This part also collects information about symbols used in the program

[1] It can compile itself.
[2] A compiler that generates code for a CPU and/or platform other than its host CPU and/or platform.
[3] A generator of code metrics, code quality reports, etc.

and associates symbol names with values, addresses, or other entities.

The back-end is on the output side of the compiler. It generates an object program that has the same meaning as the source program read by the front-end. We say that the compiler "preserves the *semantics* of the program being compiled". The semantics of a program is what it "does", while the *syntax* of a program is what it "looks like". A compiler typically reads a program having one syntax and generates a program with a different syntax, but the two programs will always have the same semantics. [1]

Preservation of semantics is the single most important design goal of a compiler. The semantics of a source language is mostly specified semi-formally and sometimes formally. Some formal semantics are even suitable for generating a compiler from them, but this is not (yet) the normal case these days. Languages with formal semantics include, for example, Scheme or Standard ML (SML). Some languages with informal or semi-formal semantics are Common Lisp, Java, C++, and C.

The front-end of a compiler is typically portable and deals with the input language in an abstract way, while the back-end is normally coupled to a specific architecture. When a compiler is intended to be portable among multiple architectures, the front-end and back-end are loosely coupled, for example by passing information from one to the other using a formal protocol or a "glue language". Compilers for single architectures allow for a tighter coupling of the stages and therefore typically generate faster and/or more compact code.

Most phases of compilation belong either to the front-end or the back-end, but some may belong to both of them. The optimization phase, which attempts to simplify the program, is often spread out over both parts, because there are some optimizations that are best applied to an abstract representation of a program, while others can only be performed in the back-end, because they require some knowledge about a specific target. In compiler-

[1] With the sole exception of code analyzers, which are intended to generate human-readable meta-information instead of executable programs.

speak, a *target* is the architecture on which the generated code will run. The SubC compiler discussed in this book will target the x86_64 processor family.

Phases of Compilation

All the phases of compilation are outlined in detail in figure 2 (shaded boxes denote processes, clear boxes represent data). The first phase on the input side of the compiler is the *lexical analysis*. This step transforms the source program into a stream of small units called *tokens*. Each token represents a small textual region of the source program, like a *numeric literal*, a *keyword*, an *operator*, etc. This phase also detects and reports input characters that are not part of the source language and sequences of text that cannot be tokenized. [1]

To illustrate the principle of tokenization, here comes a small sample program, in its "natural" form on the left and with blanks between the individual tokens on the right:

```
void printrev(char *s) {
    if (*s) {
        printrev(s+1);
        putchar(*s);
    }
}
```

```
void printrev ( char * s ) {
    if ( * s ) {
        printrev ( s + 1 ) ;
        putchar ( * s ) ;
    }
}
```

The lexical analysis phase is also referred to as the *scanner*. In fact, scanning does a bit more than the above. It does not only split up the source program into manageable parts, but it also categorizes the individual parts. For instance, the scanner output of the above program may look like this (where each symbol denotes a unique small integer):

```
VOID IDENT LPAREN CHAR STAR IDENT RPAREN LBRACE
IF LPAREN STAR IDENT RPAREN LBRACE
IDENT LPAREN IDENT PLUS INTLIT RPAREN SEMI
IDENT LPAREN STAR IDENT RPAREN SEMI
RBRACE RBRACE
```

[1] If such sequences exist in the source language; see pg. 84.

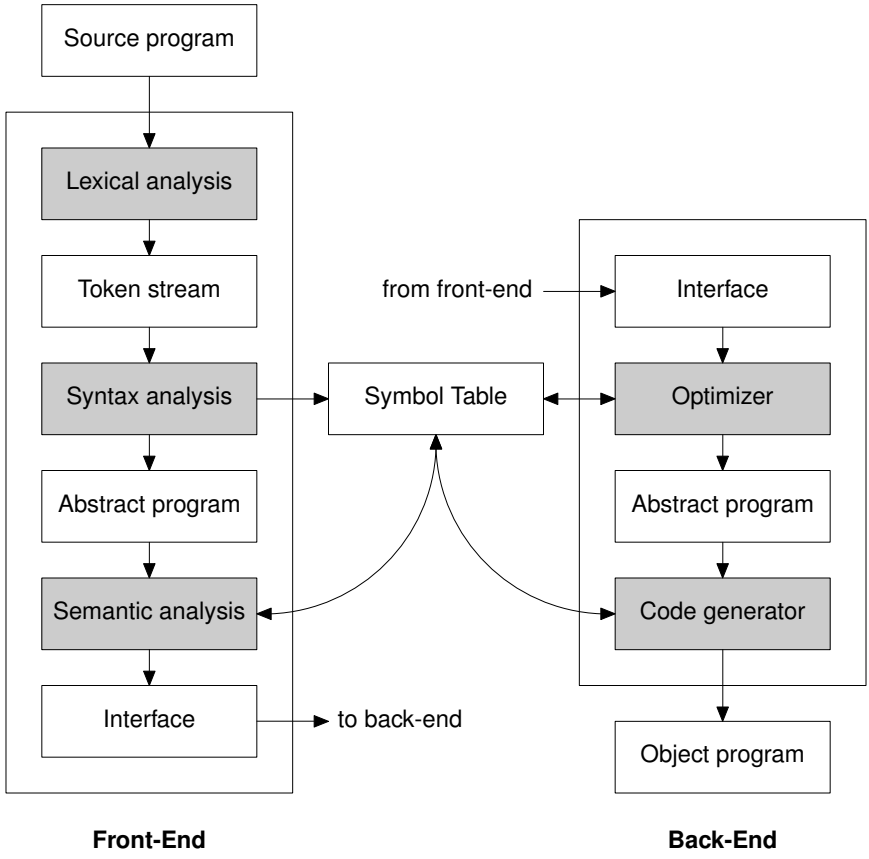

Figure 2: Phases of compilation

Of course this representation loses some information, because all the symbols (like **`s`** or **`putchar`**) are simply referred to as **`IDENT`** (identifiers), and the value of the **`INTLIT`** (integer literal) is lost. This is why the scanner adds *attributes* to some of the tokens, resulting in:

```
VOID IDENT(revstr) LPAREN CHAR STAR IDENT(s) RPAREN LBRACE
IF LPAREN STAR IDENT(s) RPAREN LBRACE
IDENT(revstr) LPAREN IDENT(s) PLUS INTLIT(1) RPAREN SEMI
IDENT(putchar) LPAREN STAR IDENT(s) RPAREN SEMI
RBRACE RBRACE
```

This would be the output of the lexical analysis stage. Some very simple compilers may perform in-place textual comparisons

instead of scanning their input. To find out whether the current input token is the **if** keyword, for instance, they would do something like

```
if (!strncmp(src, "if", 2) && !isalnum(src[2])) {
    /* got an IF */
}
```

This works fine for small languages with few operators and keywords, but does not scale up well, because quite a few comparisons may have to be performed before reaching a procedure that can handle a specific token. Even the SubC language, which is only a subset of C89, knows 77 different token types, so tens of comparisons will have to be made on average before a token can eventually be processed. In larger languages, such as C++ or COBOL, the number of comparisons may easily range in the hundreds and involve longer strings due to longer keywords. This is where tokenization really pays off, because it allows to compare tokens using the **==** operator:

```
if (IF == Token) {
    /* got an IF */
}
```

The next step is to make sure that the sequence of tokens delivered by the scanner forms a valid *sentence* of the source language. Like natural languages, formal languages have grammars that define what a correct sentence looks like. The grammar describes the *syntax* of a language, and the *syntax analyzer* compares the token stream to the rules of a grammar. The syntax analyzer is also called the *parser* and the analysis itself is called “parsing”.

The parser is the heart of the compiler: it pulls the source program in through the scanner and emits an object program through the back-end. The syntax of the source program controls the entire process, which is why this approach is called *syntax-directed translation*.

Most of the error reporting is also done by the parser, because most errors that can be caught by a compiler are syntax errors, i.e.

input sequences that do not match any rule in the grammar of the source language.

When parsing succeeds, i.e. no formal errors are found, the parser transforms the token stream into an even more abstract form. The *abstract program* constructed at this stage may have a variety of forms, but the most common and most practical one is probably the *abstract syntax tree* (*AST*).

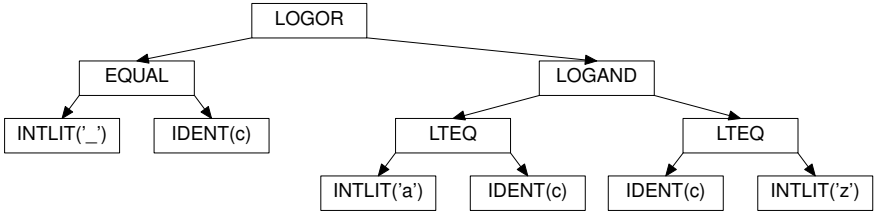

Figure 3: Abstract syntax tree (AST)

ASTs tend to be rather large, which is why we use a smaller example to illustrate them. The C expression

```
'_' == c || 'a' <= c && c <= 'z'
```

would generate the AST displayed in figure 3. (EQUAL denotes the `==` operator, LTEQ is `<=`, LOGAND is `&&`, and LOGOR is `||`). The AST reflects the structure of the input program according to the rules specified in the source language's grammar (we will discuss grammars in great detail at a later point). For example, the AST reflects the fact that the `&&` operator has a higher *precedence* than the `||` operator. [1]

The output of the parser is suitable for more complex operations, because it allows to refer to subsentences by their root nodes. For instance, the subexpression

```
'_' == c
```

of the above expression may be referred to simply by pointing to the EQUAL node in the corresponding AST.

---

[1] When these two operators occur in the same sentence without any explicit grouping (by parentheses), `&&` has to be evaluated before `||`.

The parser also builds the *symbol table*. Symbol tables are used to store the names of identifiers and their associated meta data, like addresses, values, types, *arity,* [1] array sizes, etc. Most compilers for procedural languages employ two symbol tables: one for global symbols and one for local symbols and procedure parameters.

The symbol table is heavily used by all subsequent phases of compilation. The next stage, *semantic analysis*, uses it to look up all kinds of properties of identifiers. It performs tasks such as *type checking* and context checking. Type checking has to be done on the semantic level, because the correctness of a statement like

```
x = 0;
```

cannot be decided without referring to meta information in the symbol table. When **x** is an array or a function, for instance, the above sentence is semantically incorrect. Its syntax is correct, but it makes no sense, because C cannot assign a value to an array or a function. This kind of error cannot be detected on a purely syntactic level, because an indeterminate number of tokens may appear between the definition of **x** and a reference to it.

The reporting of undefined identifiers is also done at this stage, because the semantics of a statement cannot be verified in a static language like C without knowing the properties of an identifier. (In a dynamic language, such tests can sometimes be postponed to the run time of the program.)

Context checking finds out whether a statement appears in a proper context. For instance, a **break** statement is only valid inside of a loop or a **switch** statement and would be flagged otherwise. This kind of analysis could also be performed at the syntax level, but this often leads to cumbersome formal specifications, so, in practice, it is mostly performed at the semantic level.

The *interface* between front-end and back-end can be implemented in a variety of ways. In the simplest case, the parser calls the procedures of the code generator directly. In this case,

[1] The number of arguments of a procedure.

the back-end does not even need access to the symbol table. When the back-end performs more complex operations than just emitting predefined fragments of code, though, the interface must provide procedures for accessing the symbol table from within the back-end and transferring the abstract program to the back-end stages. The most sophisticated approach would be to generate code in a portable abstract *glue language* and pass that program to the back-end. The abstract program would contain all information needed for optimizing the code and synthesizing native machine instructions. No access to the symbol table would be necessary in this case.

No matter in which way the abstract program and the related meta information is transported to the back-end, the next step would be the *optimization* phase. In figure 2, optimization has been placed in the back-end exclusively, but there are some optimizations that can be performed in the front-end as well. When both is possible, an optimization should be performed in the front-end. Only machine-dependent code should be moved to the back-end. A typical platform-neutral optimization is *constant expression folding*, which collapses subtrees of constant factors into single factors. It is illustrated in figure 4.

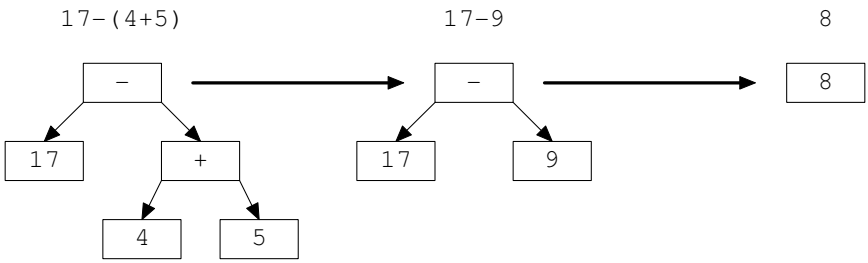

Figure 4: Constant folding

Constant folding is particularly important, because many operations that appear to be simple at the language level may expand to constant expressions when generating an abstract program. For instance, the expression

```
a[0] = 0;
```

(where **a** is an array) may expand to a constant expression adding zero to the address of the array, which can be optimized away. However, adding an arbitrary (non-zero) constant to the address of an array would require to know what the code loading the address will look like, so it must be done in the back-end. Even constant folding may require some knowledge about the target architecture, for example for detecting numeric overflow. In the case of the C language, portable constant folding must be limited to the range $-32768 \cdots 32767$, because this is the minimal numeric range of the **int** type.

When in doubt, it is always better *not* to perform an optimization. It is better to generate inefficient code than to break the preservation of semantics.

The abstract program generated by the optimizer may be an AST or a more specialized representation, like three-address code, which emulates a virtual architecture. The optimizer may also generate trees of actual machine code instructions for the target platform, thereby integrating the stages of back-end optimization and code synthesis.

In the final stage, the *code generator synthesizes* instructions for the target machine and emits them to the output. This is why it is also called the *emitter*. Machine instructions are typically generated from entire subtrees rather than individual nodes of the abstract program because, up to some limit, considering larger subtrees will lead to more efficient code.

For instance, a naive synthesizer may examine each single node of the program

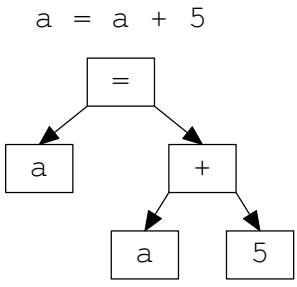

and generate code that works as follows:

- load a into register $r1$;
- load 5 into register $r2$;
- add $r2$ to $r1$;
- store $r1$ in a.

A synthesizer that examines the entire expression before emitting any code could recognize that a constant is being added to a variable and synthesize the single instruction

- add 5 to a.

This process is called “synthesis” because it synthesizes more complex instructions from simpler ones contained in the abstract program. The front-end stages of a compiler decompose complex instructions formulated in the source language into simple instructions of an abstract form. The back-end stages then reconstruct complex instructions from simpler ones, trying to make best use of the features of the target architecture.

The overall task of a compiler may be thought of as *mapping* a source program P_S of a source language S to the most efficient program P_O of an object language O that has the same meaning as P_S.

A Simplified Model

The SubC compiler, which will be discussed in the main part of this book, will use a simplified model of compilation as depicted in figure 5.

Like the full model, this simplified model uses a scanner producing a token stream that is fed to the parser. However, the semantic analysis phase is tightly coupled to the parser in this model, so the compiler performs semantic analysis basically inside of the parser. No abstract program is generated, and the symbol table is also set up in the syntactic analysis phase.

The interface to the back end is a simple procedure call interface, where each procedure in the code generator emits a

parameterized code fragment. The parameters are passed to the code generator via procedure calls. Sometimes these parameters identify slots in the symbol table, allowing the back-end direct access to the attributes of symbols. Because the code generator is only loosely coupled to the rest of the compiler, retargeting the compiler, i.e. porting it to a new target platform, involves just the creation of a new set of procedures that emit the desired object code.

The SubC compiler does not have an optimizer and it often emits inefficient and sometimes even redundant code. Emphasis was put on correctness rather than cleverness. This is the price we pay for a simple and easy-to-comprehend compiler. However, SubC is not *that* bad after all. The entire compiler including its full runtime library bootstraps itself in just about two seconds on a mainstream 750MHz[1] processor, and it is quite sufficient for a lot of every-day programming tasks.

[1] No, this text is not that ancient; the author is just content with old hardware. He has upgraded from 600Mhz to 750MHz since the first edition in 2012, though.
He has upgraded again before writing this edition, to a 1000MHz system, but the compilation time has only decreased slightly.

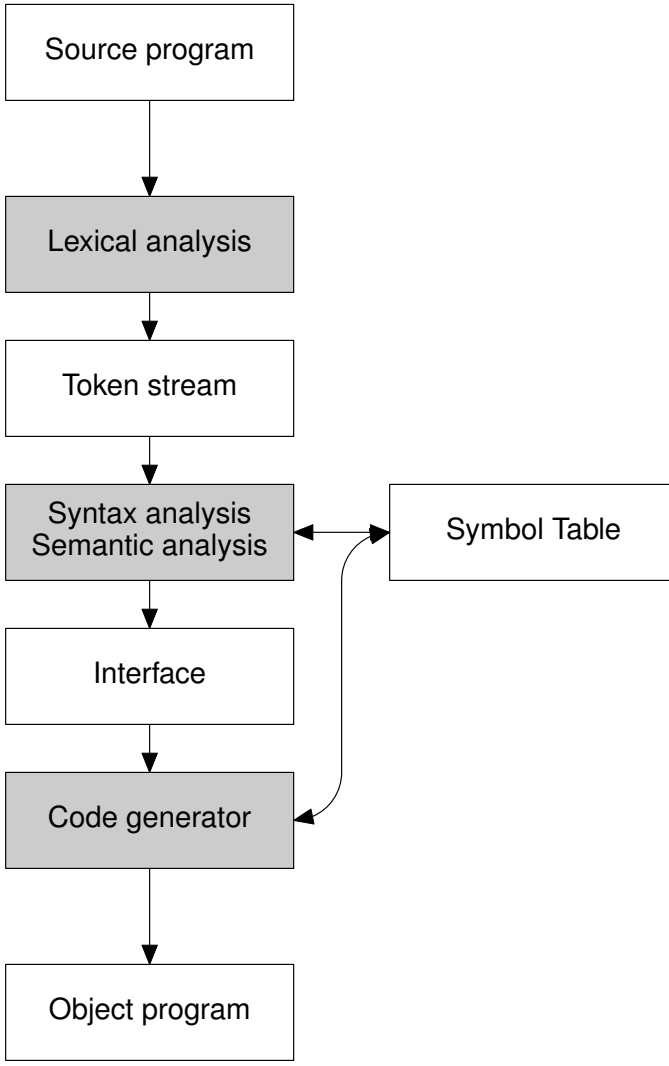

Figure 5: A simplified compiler model

The Peculiarities of C

Viewed superficially, C is a simple, rather low level, block-structured, procedural language. However, it is also quirky as hell. It has quite some hidden complexity in places where you would least expect it. It is not a language that an aspiring compiler-writer would want to implement as their first project. Not without the help of a decent textbook, that is.

For instance, C has 15 levels of operator precedence, quite a weird declaration syntax, and hairy *pointer semantics*. Quick, what does

```
int *(*foo[17])();
```

declare? And what do the following expressions deliver?

```
a[0]
*a
*a[0]
(*a)[0]
*a++
(*a)++
*++a
```

```
**a++
*(*a)++
(**a)++
++**a
*++*a
**++a
```

Even the SubC compiler discussed in this book will have to get all of the above expressions right! [1]

Most of the complexity of the C compiler is hidden in the expression parser, the part of the parser that analyzes the smallest parts of the language, the formulae used in its statements. It does not only have to support all of the above constructs, it also has to do proper pointer addition and subtraction, even in the **+=** and **−=** operations, etc. There is quite a bit of work hidden in this superficially simple language.

On the other hand, the C standard (“The C Programming Language, 2nd Ed.” (TCPL2) by K&R) gives quite a bit of leeway to the implementor. This is a good thing for a compiler-writer. It

[1] But not the declaration.

allows experienced implementors to choose the most efficient approach and the beginner to use the simplest solution. This is also the primary reason why C is known as a fast and efficient language: it is *underspecified*, i.e., it allows the implementor to interpret the specification quite freely in order to squeeze out a few more optimizations. However, this is often done at the cost of clarity to the user (of the C language). For instance, the statement

```
f(putchar('a'), putchar('b'), putchar('c'));
```

may print any permutation of "a", "b", and "c", because the standard leaves the *order of argument evaluation unspecified*. "Unspecified" means that a compiler can evaluate the arguments to a function in any order it likes. It can even use different orders at different optimization levels or even a different order for exactly the same statement when compiled twice in a row. It could throw a dice. "Unspecified" means: *just do not rely on it*.

It gets even worse when adding in the increment operators. The program fragment

```
x = 0; f(x++, x++);
```

may pass $(0, 0)$ to f or $(0, 1)$, or $(1, 0)$. This is because the standard only says that the post-increment has to take place before the end of the statement. [1] So the compiler may take the first x, increment it, and then take the second x [giving $(0, 1)$]. Or it may evaluate the second argument first [$(1, 0)$]. Or it may decide to increment x by two after calling the function, resulting in $(0, 0)$.

To the compiler writer, of course, such leeway is a great advantage, because it allows them to pick any variant that is convenient to them—and therefore to us! In this text, we will always choose the simplest and most obvious approach in order to keep the compiler source code manageable.

Studying a compiler for a language often changes the way in which we see a language, and mostly for the better. Even if you already have some programming experience in C, this textbook may reveal some of the more subtle details of the language to you.

---

[1] More specifically: before the next "sequence-point", which would be the semicolon (`;`) in this case.

For example:

Why are statements like `i = i++;` not a good idea?

Is `*x[1]` the indirection of an array element or the array element of an indirection?

Why will the declaration

```
extern char *foo[];
```

get you in trouble when the identifier was defined in a different file as

```
char **foo;
```

while at the same time there is absolutely no difference between the meanings of the following two declarations?

```
int main(int argc, char **argv);
int main(int argc, char *argv[]);
```

Answers

In case you want to verify your answers to the questions at the beginning of this chapter:

`int *(*foo[17])();`	An array (size 17) of pointers to functions returning pointer to int
`a[0]`	return first element of `a`
`*a`	return object pointed to by `a`
`*a[0]`	object pointed to by `a[0]`
`(*a)[0]`	return first element of `*a`
`*a++`	return `*a`, then increment `a`
`(*a)++`	return `*a`, then increment `*a`
`*++a`	increment `a`, then return `*a`
`**a++`	return `**a`, then increment `a`
`*(*a)++`	return `**a`, then increment `*a`
`(**a)++`	return `**a`, then increment `**a`
`++**a`	increment `**a`, then return `**a`
`*++*a`	increment `*a`, then return `**a`
`**++a`	increment `a`, then return `**a`

Rules of the Game

The Source Language

The SubC language is an almost *strict subset* of the C programming language as defined in TCPL2 (a.k.a. “C89”, “C90”, or “ANSI C”). This means that the vast majority of the programs that SubC accepts will also be accepted by a C89 compiler (and, by extension, by most modern C compilers).

The SubC language contains only two extensions to the ANSI standard, one unavoidable and one due to the laziness of its author. The unavoidable extension is the addition of the **__argc** keyword, which delivers the number of arguments passed to the current function. It is required to implement variable-argument procedures on top of the left-to-right calling conventions of the SubC compiler. Details will be explained in a later part (pg. 333).

The other extension allows function declarations to mix parameters with and without explicit type information in the same parameter list, e.g.:

```
int f(char *a, b, c);
```

which would be equal to the ANSI C declaration

```
int f(char *a, int b, int c);
```

Such declarations will not be accepted by any C compiler other than SubC. They have been included because the author uses them occasionally in prototyping.

Except for these two extensions, SubC differs from C89 by omission only. The following is a (hopefully!) comprehensive list of restrictions that apply to the SubC language as opposed to C89.

The following keywords are *not* recognized:

auto, **const**, **double**, **float**, **goto**, **long**, **register**, **short**, **signed**, **struct**, **typedef**, **union**, **unsigned**, **volatile**.

There are only two data types: the signed **int** and the unsigned **char**; there are also **void** functions and **void** pointers, and there is limited support for **int(*)()** (pointers to functions of type **int**).

No more than two levels of indirection are supported, and arrays are limited to one dimension, i.e. valid declarators are limited to

x, **x[]**, ***x**, ***x[]**, ****x** (and **(*x)()**).

K&R-style function declarations (with parameter declarations between the parameter list and function body) are not accepted.

There are no “register”, “volatile”, or “const” variables. No register allocation takes place, so all variables are implicitly “volatile”.

There is no **typedef**.

There are no unsigned integers and no long integers.

There are no **struct** or **union** data types.

Only **int**, **char**, and array of **int** and **char** types can be initialized in their declarations; pointers can be initialized with 0 (but not with **NULL**).

The **enum** statement does not define a new type, but merely a set of constants.

Local arrays cannot have initializers.

There are no local **extern** qualifiers or **enum** statements.

Local declarations are limited to the beginnings of function bodies, they do not work in other compound statements.

There are no **static** prototypes.

Arguments of prototypes must be named.

There is no **goto**.

There are no parameterized macros.

The **#error**, **#if**, **#line**, and **#pragma** preprocessor commands are not recognized.

The preprocessor does not recognize the **#** and **##** operators.

There may not be any blanks between the **#** character that introduces a preprocessor command and the subsequent command (e.g.: **# define** would not be recognized as a valid command).

Comments in preprocessor commands will cause trouble in many cases and should be avoided.

Preprocessor commands cannot span multiple lines, i.e. a "****" at the end of a line will not continue a preprocessor command.

The **sizeof** operator is limited to types and single identifiers; the operator requires parentheses.

The address of an array must be specified as **&array[0]** instead of **&array** (but just **array** also works).

Subscripting an integer with a pointer (e.g., **1["foo"]**) is not supported.

Function pointers are limited to one single type, **int(*)()**, and they have no argument types.

There is no **assert()** due to the lack of parameterized macros.

The **atexit()** mechanism is limited to one function (this may even be covered by TCPL2).

Environments of **setjmp()** have to be defined as

```
int env[_JMPBUF_SIZ]
```

instead of **jmp_buf env** due to the lack of **typedef**.

FILE is an alias of **int** due to the lack of **typedef** and **struct**.

The **signal()** function returns **int** due to the lack of a more sophisticated type system; the return value should be cast to **int(*)()**.

Most of the time-related functions are missing due to the lack of **struct**; in particular: **asctime()**, **gmtime()**, **localtime()**, **mktime()**, and **strftime()**.

The **clock()** function is missing, because **CLOCKS_PER_SEC** varies among systems.

The **ctime()** function ignores the time zone.

The Object Language

The SubC compiler generates assembly language for the x86_64 processor in "AT&T" syntax, which is predominant on Unix systems. For those familiar with the "Intel" syntax, here are the major differences:

The operands are specified with the *source operand first*, so the instruction **addq %rax,foo** would add the value of **%rax** to **foo**.

The size of an operand is explicitly specified in the name of an instruction. For instance, **movq** moves a quad word (64 bits), **movl** moves a long word (32 bits), **movw** moves a 16-bit word, and **movb** moves a byte.

Immediate operands are prefixed with a "$" sign, e.g.:

```
subq $5,%rax
```

Parentheses are used for indirection and an optional displacement can be specified in front of the parentheses. For example,

```
movq (%rax),%rax
```

loads the value pointed to by **%rax** into **%rax** and

```
incq 16(%rbp)
```

increments the (quad-word) value pointed to by **%rbp+16**.

A star is used for indirect jumps and calls:

```
call *%rbx
```

calls the routine whose address is stored in **%rbx**.

A brief x86_64 assembly language primer can be found in the appendix (pg. 443).

In order to create an executable program, the output of the compiler has to be assembled and linked against the SubC runtime library. The assembler and linker are not discussed in this book and they are not part of the source code presented here.

The SubC compiler controller invokes the host system's assembler and linker in order to generate an executable program.

The Runtime Library

The traditional C runtime environment consists of two parts: the C runtime startup module **`crt0`** and a collection of library functions widely known as **`libc`**. SubC uses this very approach, it just renames **`libc`** to **`libscc`** to avoid naming conflicts. The SubC library is completely independent and does not delegate any calls to the C library of the host system.

The startup module **`crt0`** is the only part of the compiler that is written in assembly language. The primary tasks of the runtime startup module are the initialization of internal library structures and the translation between the calling convention of SubC and the operating system kernel. I.e., this part forms a very thin layer between the SubC library and the operating system.

There are two approaches to implementing the runtime startup module. One would be to just delegate all system calls to the C library of the host operating system. E.g., when a SubC program calls the **`write()`** function of the SubC library, the library just rearranges the parameters and passes them on to the **`write()`** function of the system's **`libc`**. While this approach looks portable, because it just delegates all the nitty-gritty details to an existing library, it is surprisingly quirky. Therefore, another approach is discussed (and, in fact, preferred) in this text. In this approach, the SubC **`crt0`** module will perform all system calls itself. While this approach is more susceptible to changes in the underlying system, it works better (if it works) and produces much leaner output programs.

Except for the startup module, the standard C library[1] for SubC is itself written in SubC. Except for references to system calls, it consists of portable C89 code. Like the compiler, it differs from the full implementation mostly by omission. It also suffers from slow performance due to the lack of parameterized macros (e.g., **`getc`**,

[1] as defined in TCPL2

putc), but instead of tweaking for performance, correctness was considered to be more important, as it should always be done when designing a compiler. Individual design decisions will be discussed in the sections dealing with the C library (pg. 299).