

Undergraduate Topics in Computer Science

Torben Ægidius Mogensen

Introduction to Compiler Design



Undergraduate Topics in Computer Science

Undergraduate Topics in Computer Science (UTiCS) delivers high-quality instructional content for undergraduates studying in all areas of computing and information science. From core foundational and theoretical material to final-year topics and applications, UTiCS books take a fresh, concise, and modern approach and are ideal for self-study or for a one- or two-semester course. The texts are all authored by established experts in their fields, reviewed by an international advisory board, and contain numerous examples and problems. Many include fully worked solutions.

For further volumes:

www.springer.com/series/7592

Torben Ægidius Mogensen

Introduction to Compiler Design

 Springer

Torben Ægidius Mogensen
Department of Computer Science
University of Copenhagen
Copenhagen, Denmark
torbenm@diku.dk
url: <http://www.diku.dk/~torbenm>

Series editor
Ian Mackie

Advisory board
Samson Abramsky, University of Oxford, Oxford, UK
Chris Hankin, Imperial College London, London, UK
Dexter Kozen, Cornell University, Ithaca, USA
Andrew Pitts, University of Cambridge, Cambridge, UK
Hanne Riis Nielson, Technical University of Denmark, Lyngby, Denmark
Steven Skiena, Stony Brook University, Stony Brooks, USA
Iain Stewart, University of Durham, Durham, UK

ISSN 1863-7310
ISBN 978-0-85729-828-7 e-ISBN 978-0-85729-829-4
DOI 10.1007/978-0-85729-829-4
Springer London Dordrecht Heidelberg New York

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

Library of Congress Control Number: 2011933601

© Springer-Verlag London Limited 2011

Apart from any fair dealing for the purposes of research or private study, or criticism or review, as permitted under the Copyright, Designs and Patents Act 1988, this publication may only be reproduced, stored or transmitted, in any form or by any means, with the prior permission in writing of the publishers, or in the case of reprographic reproduction in accordance with the terms of licenses issued by the Copyright Licensing Agency. Enquiries concerning reproduction outside those terms should be sent to the publishers.

The use of registered names, trademarks, etc., in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant laws and regulations and therefore free for general use.

The publisher makes no representation, express or implied, with regard to the accuracy of the information contained in this book and cannot accept any legal responsibility or liability for any errors or omissions that may be made.

Cover design: VTeX UAB, Lithuania

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

“Language is a process of free creation; its laws and principles are fixed, but the manner in which the principles of generation are used is free and infinitely varied. Even the interpretation and use of words involves a process of free creation.”
Noam Chomsky (1928–)

In order to reduce the complexity of designing and building computers, nearly all of these are made to execute relatively simple commands (but do so very quickly). A program for a computer must be built by combining these very simple commands into a program in what is called *machine language*. Since this is a tedious and error-prone process most programming is, instead, done using a high-level *programming language*. This language can be very different from the machine language that the computer can execute, so some means of bridging the gap is required. This is where the *compiler* comes in.

A compiler translates (or *compiles*) a program written in a high-level programming language that is suitable for human programmers into the low-level machine language that is required by computers. During this process, the compiler will also attempt to spot and report obvious programmer mistakes.

Using a high-level language for programming has a large impact on how fast programs can be developed. The main reasons for this are:

- Compared to machine language, the notation used by programming languages is closer to the way humans think about problems.
- The compiler can spot some obvious programming mistakes.
- Programs written in a high-level language tend to be shorter than equivalent programs written in machine language.

Another advantage of using a high-level language is that the same program can be compiled to many different machine languages and, hence, be brought to run on many different machines.

On the other hand, programs that are written in a high-level language and automatically translated to machine language may run somewhat slower than programs that are hand-coded in machine language. Hence, some time-critical programs are still written partly in machine language. A good compiler will, however, be able to

get very close to the speed of hand-written machine code when translating well-structured programs.

The Phases of a Compiler

Since writing a compiler is a nontrivial task, it is a good idea to structure the work. A typical way of doing this is to split the compilation into several phases with well-defined interfaces. Conceptually, these phases operate in sequence (though in practice, they are often interleaved), each phase (except the first) taking the output from the previous phase as its input. It is common to let each phase be handled by a separate module. Some of these modules are written by hand, while others may be generated from specifications. Often, some of the modules can be shared between several compilers.

A common division into phases is described below. In some compilers, the ordering of phases may differ slightly, some phases may be combined or split into several phases or some extra phases may be inserted between those mentioned below.

Lexical analysis This is the initial part of reading and analysing the program text: The text is read and divided into *tokens*, each of which corresponds to a symbol in the programming language, e.g., a variable name, keyword or number.

Syntax analysis This phase takes the list of tokens produced by the lexical analysis and arranges these in a tree-structure (called the *syntax tree*) that reflects the structure of the program. This phase is often called *parsing*.

Type checking This phase analyses the syntax tree to determine if the program violates certain consistency requirements, e.g., if a variable is used but not declared or if it is used in a context that does not make sense given the type of the variable, such as trying to use a boolean value as a function pointer.

Intermediate code generation The program is translated to a simple machine-independent intermediate language.

Register allocation The symbolic variable names used in the intermediate code are translated to numbers, each of which corresponds to a register in the target machine code.

Machine code generation The intermediate language is translated to assembly language (a textual representation of machine code) for a specific machine architecture.

Assembly and linking The assembly-language code is translated into binary representation and addresses of variables, functions, etc., are determined.

The first three phases are collectively called *the frontend* of the compiler and the last three phases are collectively called *the backend*. The middle part of the compiler is in this context only the intermediate code generation, but this often includes various optimisations and transformations on the intermediate code.

Each phase, through checking and transformation, establishes stronger invariants on the things it passes on to the next, so that writing each subsequent phase is easier than if these have to take all the preceding into account. For example, the type

checker can assume absence of syntax errors and the code generation can assume absence of type errors.

Assembly and linking are typically done by programs supplied by the machine or operating system vendor, and are hence not part of the compiler itself, so we will not further discuss these phases in this book.

Interpreters

An *interpreter* is another way of implementing a programming language. Interpretation shares many aspects with compiling. Lexing, parsing and type-checking are in an interpreter done just as in a compiler. But instead of generating code from the syntax tree, the syntax tree is processed directly to evaluate expressions and execute statements, and so on. An interpreter may need to process the same piece of the syntax tree (for example, the body of a loop) many times and, hence, interpretation is typically slower than executing a compiled program. But writing an interpreter is often simpler than writing a compiler and the interpreter is easier to move to a different machine, so for applications where speed is not of essence, interpreters are often used.

Compilation and interpretation may be combined to implement a programming language: The compiler may produce intermediate-level code which is then interpreted rather than compiled to machine code. In some systems, there may even be parts of a program that are compiled to machine code, some parts that are compiled to intermediate code, which is interpreted at runtime while other parts may be kept as a syntax tree and interpreted directly. Each choice is a compromise between speed and space: Compiled code tends to be bigger than intermediate code, which tend to be bigger than syntax, but each step of translation improves running speed.

Using an interpreter is also useful during program development, where it is more important to be able to test a program modification quickly rather than run the program efficiently. And since interpreters do less work on the program before execution starts, they are able to start running the program more quickly. Furthermore, since an interpreter works on a representation that is closer to the source code than is compiled code, error messages can be more precise and informative.

We will discuss interpreters briefly in Chap. 4, but they are not the main focus of this book.

Why Learn About Compilers?

Few people will ever be required to write a compiler for a general-purpose language like C, Java or SML. So why do most computer science institutions offer compiler courses and often make these mandatory?

Some typical reasons are:

- (a) It is considered a topic that you should know in order to be “well-cultured” in computer science.
- (b) A good craftsman should know his tools, and compilers are important tools for programmers and computer scientists.
- (c) The techniques used for constructing a compiler are useful for other purposes as well.
- (d) There is a good chance that a programmer or computer scientist will need to write a compiler or interpreter for a domain-specific language.

The first of these reasons is somewhat dubious, though something can be said for “knowing your roots”, even in such a hastily changing field as computer science.

Reason “b” is more convincing: Understanding how a compiler is built will allow programmers to get an intuition about what their high-level programs will look like when compiled and use this intuition to tune programs for better efficiency. Furthermore, the error reports that compilers provide are often easier to understand when one knows about and understands the different phases of compilation, such as knowing the difference between lexical errors, syntax errors, type errors and so on.

The third reason is also quite valid. In particular, the techniques used for reading (*lexing* and *parsing*) the text of a program and converting this into a form (*abstract syntax*) that is easily manipulated by a computer, can be used to read and manipulate any kind of structured text such as XML documents, address lists, etc.

Reason “d” is becoming more and more important as domain specific languages (DSLs) are gaining in popularity. A DSL is a (typically small) language designed for a narrow class of problems. Examples are data-base query languages, text-formatting languages, scene description languages for ray-tracers and languages for setting up economic simulations. The target language for a compiler for a DSL may be traditional machine code, but it can also be another high-level language for which compilers already exist, a sequence of control signals for a machine, or formatted text and graphics in some printer-control language (e.g. PostScript). Even so, all DSL compilers will share similar front-ends for reading and analysing the program text.

Hence, the methods needed to make a compiler front-end are more widely applicable than the methods needed to make a compiler back-end, but the latter is more important for understanding how a program is executed on a machine.

The Structure of This Book

The first chapters of the book describes the methods and tools required to read program text and convert it into a form suitable for computer manipulation. This process is made in two stages: A lexical analysis stage that basically divides the input text into a list of “words”. This is followed by a syntax analysis (or *parsing*) stage that analyses the way these words form structures and converts the text into a data structure that reflects the textual structure. Lexical analysis is covered in Chap. 1 and syntactical analysis in Chap. 2.

The remainder of the book (Chaps. 3–9) covers the middle part and back-end of interpreters and compilers. Chapter 3 covers how definitions and uses of names (*identifiers*) are connected through *symbol tables*. Chapter 4 shows how you can implement a simple programming language by writing an interpreter and notes that this gives a considerable overhead that can be reduced by doing more things before executing the program, which leads to the following chapters about static type checking (Chap. 5) and compilation (Chaps. 6–9). In Chap. 6, it is shown how expressions and statements can be compiled into an *intermediate language*, a language that is close to machine language but hides machine-specific details. In Chap. 7, it is discussed how the intermediate language can be converted into “real” machine code. Doing this well requires that the registers in the processor are used to store the values of variables, which is achieved by a *register allocation* process, as described in Chap. 8. Up to this point, a “program” has been what corresponds to the body of a single procedure. Procedure calls add some issues, which are discussed in Chap. 9.

The book uses standard set notation and equations over sets. [Appendix](#) contains a short summary of these, which may be helpful to those that need these concepts refreshed.

To the Lecturer

This book was written for use in the introductory compiler course at DIKU, the department of computer science at the University of Copenhagen, Denmark.

At times, standard techniques from compiler construction have been simplified for presentation in this book. In such cases references are made to books or articles where the full version of the techniques can be found.

The book aims at being “language neutral”. This means two things:

- Little detail is given about how the methods in the book can be implemented in any specific language. Rather, the description of the methods is given in the form of algorithm sketches and textual suggestions of how these can be implemented in various types of languages, in particular imperative and functional languages.
- There is no single through-going example of a language to be compiled. Instead, different small (sub-)languages are used in various places to cover exactly the points that the text needs. This is done to avoid drowning in detail, hopefully allowing the readers to “see the wood for the trees”.

Each chapter has a section on further reading, which suggests additional reading material for interested students. Each chapter has a set of exercises. Few of these require access to a computer, but can be solved on paper or black-board. After some of the sections in the book, a few easy exercises are listed as suggested exercises. It is recommended that the student attempts to solve these exercises before continuing reading, as the exercises support understanding of the previous sections.

Teaching with this book can be supplemented with project work, where students write simple compilers. Since the book is language neutral, no specific project is given. Instead, the teacher must choose relevant tools and select a project that fits the level of the students and the time available. Depending on the amount of project work and supplementary material, the book can support course sizes ranging from 5 to 7.5 ECTS points.

Acknowledgements

*“Most people return small favors, acknowledge medium ones
and repay greater ones—with ingratitude.”*
Benjamin Franklin (1705–1790)

The author wishes to thank all people who have been helpful in making this book a reality. This includes the students who have been exposed to earlier versions of the book at the compiler courses “Dat 1E” and “Oversættelse” at DIKU, and who have found numerous typos and other errors in the earlier versions. I would also like to thank the instructors at Dat 1E and Oversættelse, who have pointed out places where things were not as clear as they could be.

Copenhagen, Denmark

Torben Ægidius Mogensen

Contents

1	Lexical Analysis	1
1.1	Regular Expressions	2
1.1.1	Shorthands	4
1.1.2	Examples	5
1.2	Nondeterministic Finite Automata	6
1.3	Converting a Regular Expression to an NFA	9
1.3.1	Optimisations	9
1.4	Deterministic Finite Automata	11
1.5	Converting an NFA to a DFA	13
1.5.1	Solving Set Equations	13
1.5.2	The Subset Construction	16
1.6	Size Versus Speed	18
1.7	Minimisation of DFAs	19
1.7.1	Example	21
1.7.2	Dead States	23
1.8	Lexers and Lexer Generators	24
1.8.1	Lexer Generators	28
1.9	Properties of Regular Languages	29
1.9.1	Relative Expressive Power	29
1.9.2	Limits to Expressive Power	31
1.9.3	Closure Properties	32
1.10	Further Reading	32
1.11	Exercises	33
	References	37
2	Syntax Analysis	39
2.1	Context-Free Grammars	40
2.1.1	How to Write Context Free Grammars	41
2.2	Derivation	43
2.2.1	Syntax Trees and Ambiguity	45
2.3	Operator Precedence	47

2.3.1	Rewriting Ambiguous Expression Grammars	49
2.4	Other Sources of Ambiguity	51
2.5	Syntax Analysis	52
2.6	Predictive Parsing	52
2.7	<i>Nullable</i> and <i>FIRST</i>	53
2.8	Predictive Parsing Revisited	56
2.9	<i>FOLLOW</i>	58
2.10	A Larger Example	60
2.11	LL(1) Parsing	62
2.11.1	Recursive Descent	63
2.11.2	Table-Driven LL(1) Parsing	64
2.11.3	Conflicts	65
2.12	Rewriting a Grammar for LL(1) Parsing	66
2.12.1	Eliminating Left-Recursion	66
2.12.2	Left-Factorisation	68
2.12.3	Construction of LL(1) Parsers Summarized	69
2.13	SLR Parsing	70
2.14	Constructing SLR Parse Tables	73
2.14.1	Conflicts in SLR Parse-Tables	76
2.15	Using Precedence Rules in LR Parse Tables	77
2.16	Using LR-Parser Generators	79
2.16.1	Declarations and Actions	79
2.16.2	Abstract Syntax	80
2.16.3	Conflict Handling in Parser Generators	83
2.17	Properties of Context-Free Languages	84
2.18	Further Reading	85
2.19	Exercises	85
	References	90
3	Scopes and Symbol Tables	91
3.1	Symbol Tables	92
3.1.1	Implementation of Symbol Tables	92
3.1.2	Simple Persistent Symbol Tables	93
3.1.3	A Simple Imperative Symbol Table	94
3.1.4	Efficiency Issues	95
3.1.5	Shared or Separate Name Spaces	95
3.2	Further Reading	96
3.3	Exercises	96
	References	96
4	Interpretation	97
4.1	The Structure of an Interpreter	98
4.2	A Small Example Language	98
4.3	An Interpreter for the Example Language	99
4.3.1	Evaluating Expressions	100
4.3.2	Interpreting Function Calls	102

4.3.3	Interpreting a Program	102
4.4	Advantages and Disadvantages of Interpretation	102
4.5	Further Reading	105
4.6	Exercises	105
	References	106
5	Type Checking	107
5.1	The Design Space of Types	107
5.2	Attributes	109
5.3	Environments for Type Checking	109
5.4	Type Checking Expressions	110
5.5	Type Checking of Function Declarations	112
5.6	Type Checking a Program	112
5.7	Advanced Type Checking	113
5.8	Further Reading	116
5.9	Exercises	116
	References	117
6	Intermediate-Code Generation	119
6.1	Choosing an Intermediate Language	120
6.2	The Intermediate Language	121
6.3	Syntax-Directed Translation	123
6.4	Generating Code from Expressions	123
	6.4.1 Examples of Translation	127
6.5	Translating Statements	127
6.6	Logical Operators	130
	6.6.1 Sequential Logical Operators	131
6.7	Advanced Control Statements	134
6.8	Translating Structured Data	135
	6.8.1 Floating-Point Values	135
	6.8.2 Arrays	135
	6.8.3 Strings	140
	6.8.4 Records/Structs and Unions	140
6.9	Translating Declarations	141
	6.9.1 Simple Local Declarations	142
	6.9.2 Translation of Function Declarations	142
6.10	Further Reading	143
6.11	Exercises	143
	References	146
7	Machine-Code Generation	147
7.1	Conditional Jumps	147
7.2	Constants	149
7.3	Exploiting Complex Instructions	149
	7.3.1 Two-Address Instructions	153
7.4	Optimisations	153
7.5	Further Reading	155

7.6	Exercises	155
	References	157
8	Register Allocation	159
8.1	Liveness	160
8.2	Liveness Analysis	160
8.3	Interference	164
8.4	Register Allocation by Graph Colouring	165
8.5	Spilling	167
8.6	Heuristics	168
	8.6.1 Removing Redundant Moves	171
	8.6.2 Using Explicit Register Numbers	171
8.7	Further Reading	172
8.8	Exercises	172
	References	174
9	Functions	175
9.1	The Call Stack	175
9.2	Activation Records	176
9.3	Prologues, Epilogues and Call-Sequences	177
9.4	Letting the Callee Save Registers	179
9.5	Caller-Saves Versus Callee-Saves	180
9.6	Using Registers to Pass Parameters	182
9.7	Interaction with the Register Allocator	184
9.8	Local Variables	186
9.9	Accessing Non-local Variables	186
	9.9.1 Global Variables	186
	9.9.2 Call-by-Reference Parameters	187
9.10	Functions as Parameters	188
9.11	Variants	189
	9.11.1 Variable-Sized Frames	189
	9.11.2 Variable Number of Parameters	189
	9.11.3 Direction of Stack-Growth and Position of <i>FP</i>	189
	9.11.4 Register Stacks	190
9.12	Further Reading	190
9.13	Exercises	190
	References	192
Appendix	Set Notation and Concepts	193
A.1	Basic Concepts and Notation	193
	A.1.1 Operations and Predicates	193
	A.1.2 Properties of Set Operations	194
A.2	Set-Builder Notation	195
A.3	Sets of Sets	196
A.4	Set Equations	196

A.4.1 Monotonic Set Functions	197
A.4.2 Distributive Functions	198
A.4.3 Simultaneous Equations	199
Exercises	199
Index	201

List of Figures

Fig. 1.1	Regular expressions	3
Fig. 1.2	Some algebraic properties of regular expressions	5
Fig. 1.3	Example of an NFA	8
Fig. 1.4	Constructing NFA fragments from regular expressions	10
Fig. 1.5	NFA for the regular expression $(a b)^*ac$	11
Fig. 1.6	Optimised NFA construction for regular expression shorthands	11
Fig. 1.7	Optimised NFA for $[0-9]^+$	12
Fig. 1.8	Example of a DFA	12
Fig. 1.9	DFA constructed from the NFA in Fig. 1.5	19
Fig. 1.10	Non-minimal DFA	21
Fig. 1.11	Minimal DFA	22
Fig. 1.12	Combined NFA for several tokens	26
Fig. 1.13	Combined DFA for several tokens	27
Fig. 1.14	A 4-state NFA that gives 15 DFA states	30
Fig. 2.1	From regular expressions to context free grammars	42
Fig. 2.2	Simple expression grammar	42
Fig. 2.3	Simple statement grammar	43
Fig. 2.4	Example grammar	44
Fig. 2.5	Derivation of the string <code>aabbbcc</code> using Grammar 2.4	44
Fig. 2.6	Leftmost derivation of the string <code>aabbbcc</code> using Grammar 2.4	45
Fig. 2.7	Syntax tree for the string <code>aabbbcc</code> using Grammar 2.4	46
Fig. 2.8	Alternative syntax tree for the string <code>aabbbcc</code> using Grammar 2.4	46
Fig. 2.9	Unambiguous version of Grammar 2.4	46
Fig. 2.10	Preferred syntax tree for $2+3*4$ using Grammar 2.2	48
Fig. 2.11	Unambiguous expression grammar	50
Fig. 2.12	Syntax tree for $2+3*4$ using Grammar 2.11	51
Fig. 2.13	Unambiguous grammar for statements	52
Fig. 2.14	Fixed-point iteration for calculation of <i>Nullable</i>	55
Fig. 2.15	Fixed-point iteration for calculation of <i>FIRST</i>	56
Fig. 2.16	Recursive descent parser for Grammar 2.9	64
Fig. 2.17	LL(1) table for Grammar 2.9	65

Fig. 2.18	Program for table-driven LL(1) parsing	65
Fig. 2.19	Input and stack during table-driven LL(1) parsing	66
Fig. 2.20	Removing left-recursion from Grammar 2.11	68
Fig. 2.21	Left-factorised grammar for conditionals	69
Fig. 2.22	SLR table for Grammar 2.9	72
Fig. 2.23	Algorithm for SLR parsing	72
Fig. 2.24	Example SLR parsing	73
Fig. 2.25	Example grammar for SLR-table construction	73
Fig. 2.26	NFAs for the productions in Grammar 2.25	74
Fig. 2.27	NFAs for the productions in Grammar 2.25 with epsilon transitions added	75
Fig. 2.28	SLR DFA for Grammar 2.9	75
Fig. 2.29	Summary of SLR parse-table construction	76
Fig. 2.30	Textual representation of NFA states	84
Fig. 4.1	Example language for interpretation	99
Fig. 4.2	Evaluating expressions	101
Fig. 4.3	Evaluating a function call	103
Fig. 4.4	Interpreting a program	104
Fig. 5.1	The design space of types	108
Fig. 5.2	Type checking of expressions	111
Fig. 5.3	Type checking a function declaration	113
Fig. 5.4	Type checking a program	114
Fig. 6.1	The intermediate language	122
Fig. 6.2	A simple expression language	124
Fig. 6.3	Translating an expression	126
Fig. 6.4	Statement language	128
Fig. 6.5	Translation of statements	129
Fig. 6.6	Translation of simple conditions	130
Fig. 6.7	Example language with logical operators	132
Fig. 6.8	Translation of sequential logical operators	133
Fig. 6.9	Translation for one-dimensional arrays	136
Fig. 6.10	A two-dimensional array	138
Fig. 6.11	Translation of multi-dimensional arrays	139
Fig. 6.12	Translation of simple declarations	142
Fig. 7.1	Pattern/replacement pairs for a subset of the MIPS instruction set	152
Fig. 8.1	Gen and kill sets	162
Fig. 8.2	Example program for liveness analysis and register allocation . . .	163
Fig. 8.3	<i>succ</i> , <i>gen</i> and <i>kill</i> for the program in Fig. 8.2	163
Fig. 8.4	Fixed-point iteration for liveness analysis	164
Fig. 8.5	Interference graph for the program in Fig. 8.2	166
Fig. 8.6	Algorithm 8.3 applied to the graph in Fig. 8.5	169
Fig. 8.7	Program from Fig. 8.2 after spilling variable <i>a</i>	169
Fig. 8.8	Interference graph for the program in Fig. 8.7	170
Fig. 8.9	Colouring of the graph in Fig. 8.8	170

Fig. 9.1	Simple activation record layout	177
Fig. 9.2	Prologue for the header $f(p_1, \dots, p_m)$ using the frame layout shown in Fig. 9.1	178
Fig. 9.3	Epilogue for the instruction <code>RETURNresult</code> using the frame layout shown in Fig. 9.1	178
Fig. 9.4	Call sequence for $x := \text{CALL}g(a_1, \dots, a_n)$ using the frame layout shown in Fig. 9.1	179
Fig. 9.5	Activation record layout for callee-saves	180
Fig. 9.6	Prologue for the header $f(p_1, \dots, p_m)$ using callee-saves	180
Fig. 9.7	Epilogue for the instruction <code>RETURNresult</code> using callee-saves	181
Fig. 9.8	Call sequence for $x := \text{CALL}g(a_1, \dots, a_n)$ using callee-saves	181
Fig. 9.9	Possible division of registers for a 16-register architecture	182
Fig. 9.10	Activation record layout for the register division shown in Fig. 9.9	182
Fig. 9.11	Prologue for the header $f(p_1, \dots, p_m)$ using the register division shown in Fig. 9.9	183
Fig. 9.12	Epilogue for the instruction <code>RETURNresult</code> using the register division shown in Fig. 9.9	183
Fig. 9.13	Call sequence for $x := \text{CALL}g(a_1, \dots, a_n)$ using the register division shown in Fig. 9.9	184

Chapter 1

Lexical Analysis

“I am not yet so lost in lexicography as to forget that words are the daughters of earth, and that things are the sons of heaven. Language is only the instrument of science, and words are but the signs of ideas.”
Samuel Johnson (1709–1784)

The word “lexical” in the traditional sense means “pertaining to words”. In terms of programming languages, words are objects like variable names, numbers, keywords etc. Such word-like entities are traditionally called *tokens*.

A *lexical analyser*, also called a *lexer* or *scanner*, will as its input take a string of individual letters and divide this string into tokens. Additionally, it will filter out whatever separates the tokens (the so-called *white-space*), i.e., lay-out characters (spaces, newlines etc.) and comments.

The main purpose of lexical analysis is to make life easier for the subsequent syntax analysis phase. In theory, the work that is done during lexical analysis can be made an integral part of syntax analysis, and in simple systems this is indeed often done. However, there are reasons for keeping the phases separate:

- **Efficiency:** A lexer may do the simple parts of the work faster than the more general parser can. Furthermore, the size of a system that is split in two may be smaller than a combined system. This may seem paradoxical but, as we shall see, there is a non-linear factor involved which may make a separated system smaller than a combined system.
- **Modularity:** The syntactical description of the language need not be cluttered with small lexical details such as white-space and comments.
- **Tradition:** Languages are often designed with separate lexical and syntactical phases in mind, and the standard documents of such languages typically separate lexical and syntactical elements of the languages.

It is usually not terribly difficult to write a lexer by hand: You first read past initial white-space, then you, in sequence, test to see if the next token is a keyword, a number, a variable or whatnot. However, this is not a very good way of handling

the problem: You may read the same part of the input repeatedly while testing each possible token and in some cases it may not be clear where the next token ends. Furthermore, a handwritten lexer may be complex and difficult to maintain. Hence, lexers are normally constructed by *lexer generators*, which transform human-readable specifications of tokens and white-space into efficient programs.

We will see the same general strategy in the chapter about syntax analysis: Specifications in a well-defined human-readable notation are transformed into efficient programs.

For lexical analysis, specifications are traditionally written using *regular expressions*: An algebraic notation for describing sets of strings. The generated lexers are in a class of extremely simple programs called *finite automata*.

This chapter will describe regular expressions and finite automata, their properties and how regular expressions can be converted to finite automata. Finally, we discuss some practical aspects of lexer generators.

1.1 Regular Expressions

The set of all integer constants or the set of all variable names are sets of strings, where the individual letters are taken from a particular alphabet. Such a set of strings is called a *language*. For integers, the alphabet consists of the digits 0–9 and for variable names the alphabet contains both letters and digits (and perhaps a few other characters, such as underscore).

Given an alphabet, we will describe sets of strings by *regular expressions*, an algebraic notation that is compact and easy for humans to use and understand. The idea is that regular expressions that describe simple sets of strings can be combined to form regular expressions that describe more complex sets of strings.

When talking about regular expressions, we will use the letters (*r*, *s* and *t*) in italics to denote unspecified regular expressions. When letters stand for themselves (i.e., in regular expressions that describe strings that use these letters) we will use typewriter font, e.g., a or b. Hence, when we say, e.g., “The regular expression *s*” we mean the regular expression that describes a single one-letter string “s”, but when we say “The regular expression *s*”, we mean a regular expression of any form which we just happen to call *s*. We use the notation $L(s)$ to denote the language (i.e., set of strings) described by the regular expression *s*. For example, $L(a)$ is the set {“a”}.

Figure 1.1 shows the constructions used to build regular expressions and the languages they describe:

- A single letter describes the language that has the one-letter string consisting of that letter as its only element.
- The symbol ε (the Greek letter *epsilon*) describes the language that consists solely of the empty string. Note that this is not the empty set of strings (see Exercise 1.10).
- $s|t$ (pronounced “*s* or *t*”) describes the union of the languages described by *s* and *t*.

Regular expression	Language (set of strings)	Informal description
a	$\{“a”\}$	The set consisting of the one-letter string “a”.
ε	$\{“”\}$	The set containing the empty string.
$s t$	$L(s) \cup L(t)$	Strings from both languages
st	$\{vw \mid v \in L(s), w \in L(t)\}$	Strings constructed by concatenating a string from the first language with a string from the second language. Note: In set-formulas, “ ” is not a part of a regular expression, but part of the set-builder notation and reads as “where”.
s^*	$\{“”\} \cup \{vw \mid v \in L(s), w \in L(s^*)\}$	Each string in the language is a concatenation of any number of strings in the language of s .

Fig. 1.1 Regular expressions

- st (pronounced “ s t ”) describes the concatenation of the languages $L(s)$ and $L(t)$, i.e., the sets of strings obtained by taking a string from $L(s)$ and putting this in front of a string from $L(t)$. For example, if $L(s)$ is $\{“a”, “b”\}$ and $L(t)$ is $\{“c”, “d”\}$, then $L(st)$ is the set $\{“ac”, “ad”, “bc”, “bd”\}$.
- The language for s^* (pronounced “ s star”) is described recursively: It consists of the empty string plus whatever can be obtained by concatenating a string from $L(s)$ to a string from $L(s^*)$. This is equivalent to saying that $L(s^*)$ consists of strings that can be obtained by concatenating zero or more (possibly different) strings from $L(s)$. If, for example, $L(s)$ is $\{“a”, “b”\}$ then $L(s^*)$ is $\{“”, “a”, “b”, “aa”, “ab”, “ba”, “bb”, “aaa”, \dots\}$, i.e., any string (including the empty) that consists entirely of as and bs.

Note that while we use the same notation for concrete strings and regular expressions denoting one-string languages, the context will make it clear which is meant. We will often show strings and sets of strings without using quotation marks, e.g., write $\{a, bb\}$ instead of $\{“a”, “bb”\}$. When doing so, we will use ε to denote the empty string, so the example from $L(s^*)$ above is written as $\{\varepsilon, a, b, aa, ab, ba, bb, aaa, \dots\}$. The letters u, v and w in italics will be used to denote unspecified single strings, i.e., members of some language. As an example, abw denotes any string starting with ab .

Precedence Rules When we combine different constructor symbols, e.g., in the regular expression $a|ab^*$, it is not *a priori* clear how the different subexpressions are grouped. We can use parentheses to make the grouping of symbols explicit such as in $(a|(ab))^*$. Additionally, we use precedence rules, similar to the algebraic convention that $3 + 4 * 5$ means 3 added to the product of 4 and 5 and not multiplying the sum of 3 and 4 by 5. For regular expressions, we use the following conventions: $*$ binds tighter than concatenation, which binds tighter than alternative ($|$). The example $a|ab^*$ from above, hence, is equivalent to $a|(a(b^*))$.

The $|$ operator is associative and commutative (as it corresponds to set union, which has these properties). Concatenation is associative (but obviously not commutative) and distributes over $|$. Figure 1.2 shows these and other algebraic properties of regular expressions, including definitions of some of the shorthands introduced below.

1.1.1 Shorthands

While the constructions in Fig. 1.1 suffice to describe e.g., number strings and variable names, we will often use extra shorthands for convenience. For example, if we want to describe non-negative integer constants, we can do so by saying that it is one or more digits, which is expressed by the regular expression

$$(0|1|2|3|4|5|6|7|8|9)(0|1|2|3|4|5|6|7|8|9)^*$$

The large number of different digits makes this expression rather verbose. It gets even worse when we get to variable names, where we must enumerate all alphabetic letters (in both upper and lower case).

Hence, we introduce a shorthand for sets of letters. Sequences of letters within square brackets represent the set of these letters. For example, we use $[ab01]$ as a shorthand for $a|b|0|1$. Additionally, we can use interval notation to abbreviate $[0123456789]$ to $[0-9]$. We can combine several intervals within one bracket and for example write $[a-zA-Z]$ to denote all alphabetic letters in both lower and upper case.

When using intervals, we must be aware of the ordering for the symbols involved. For the digits and letters used above, there is usually no confusion. However, if we write, e.g., $[0-z]$ it is not immediately clear what is meant. When using such notation in lexer generators, standard ASCII or ISO 8859-1 character sets are usually used, with the hereby implied ordering of symbols. To avoid confusion, we will use the interval notation only for intervals of digits or alphabetic letters.

Getting back to the example of integer constants above, we can now write this much shorter as $[0-9][0-9]^*$.

Since s^* denotes *zero or more* occurrences of s , we needed to write the set of digits twice to describe that *one or more* digits are allowed. Such non-zero repetition is quite common, so we introduce another shorthand, s^+ , to denote one or more

- [*download The Needlecraft Book*](#)
- [**click SPQR: A Roman Miscellany**](#)
- [read Nutrition: An Applied Approach \(3rd Edition\)](#)
- [*Mixed Methods Research: Merging Theory with Practice here*](#)

- <http://kamallubana.com/?library/The-Needlecraft-Book.pdf>
- <http://fitnessfatale.com/freebooks/Chain-of-Blame--How-Wall-Street-Caused-the-Mortgage-and-Credit-Crisis.pdf>
- <http://www.shreesaiexport.com/library/The-Juggler.pdf>
- <http://test1.batsinbelfries.com/ebooks/Mixed-Methods-Research--Merging-Theory-with-Practice.pdf>