School of Design, Engineering \& Computing

BSc (Hons) Computing
BSc (Hons) Software Engineering Management

# ARM: Assembly Language Programming 

Peter Knaggs
AND
Stephen WELSh

August 31, 2004

## Contents

Contents ..... i
List of Programs ..... vii
Preface ..... ix
1 Introduction ..... 1
1.1 The Meaning of Instructions ..... 1
1.1.1 Binary Instructions ..... 1
1.2 A Computer Program ..... 1
1.3 The Binary Programming Problem ..... 2
1.4 Using Octal or Hexadecimal ..... 2
1.5 Instruction Code Mnemonics ..... 3
1.6 The Assembler Program ..... 4
1.6.1 Additional Features of Assemblers ..... 4
1.6.2 Choosing an Assembler ..... 5
1.7 Disadvantages of Assembly Language ..... 5
1.8 High-Level Languages ..... 6
1.8.1 Advantages of High-Level Languages ..... 6
1.8.2 Disadvantages of High-Level Languages ..... 7
1.9 Which Level Should You Use? ..... 8
1.9.1 Applications for Machine Language ..... 8
1.9.2 Applications for Assembly Language ..... 8
1.9.3 Applications for High-Level Language ..... 8
1.9.4 Other Considerations ..... 8
1.10 Why Learn Assembler? ..... 8
2 Assemblers ..... 11
2.1 Fields ..... 11
2.1.1 Delimiters ..... 11
2.1.2 Labels ..... 12
2.2 Operation Codes (Mnemonics) ..... 14
2.3 Directives ..... 14
2.3.1 The DEFINE CONSTANT (Data) Directive ..... 14
2.3.2 The EQUATE Directive ..... 15
2.3.3 The AREA Directive ..... 16
2.3.4 Housekeeping Directives ..... 17
2.3.5 When to Use Labels ..... 17
2.4 Operands and Addresses ..... 17
2.4.1 Decimal Numbers ..... 18
2.4.2 Other Number Systems ..... 18
2.4.3 Names ..... 18
2.4.4 Character Codes ..... 18
2.4.5 Arithmetic and Logical Expressions ..... 18
2.4.6 General Recommendations ..... 19
2.5 Comments ..... 19
2.6 Types of Assemblers ..... 20
2.7 Errors ..... 20
2.8 Loaders ..... 21
3 ARM Architecture ..... 23
3.1 Processor modes ..... 23
3.2 Registers ..... 25
3.2.1 The stack pointer, SP or R13 ..... 26
3.2.2 The Link Register, LR or R14 ..... 27
3.2.3 The program counter, PC or R15 ..... 27
3.2.4 Current Processor Status Registers: CPSR ..... 28
3.3 Flags ..... 28
3.4 Exceptions ..... 29
3.5 Instruction Set ..... 30
3.5.1 Conditional Execution: $\langle c c\rangle$ ..... 31
3.5.2 Data Processing Operands: $\langle o p 1\rangle$ ..... 32
3.5.3 Memory Access Operands: $\langle o p 2\rangle$ ..... 34
4 Instruction Set ..... 37
4.0.4 Branch instructions ..... 38
4.0.5 Data-processing instructions. ..... 38
4.0.6 Status register transfer instructions ..... 39
4.0.7 Load and store instructions ..... 40
4.0.8 Coprocessor instructions ..... 41
4.0.9 Exception-generating instructions ..... 41
4.0.10 Conditional Execution: $\langle c c\rangle$ ..... 42
5 Addressing Modes ..... 45
5.1 Data Processing Operands: $\langle o p 1\rangle$ ..... 45
5.1.1 Unmodified Value ..... 45
5.1.2 Logical Shift Left ..... 45
5.1.3 Logical Shift Right ..... 46
5.1.4 Arithmetic Shift Right ..... 46
5.1.5 Rotate Right ..... 46
5.1.6 Rotate Right Extended ..... 47
5.2 Memory Access Operands: $\langle o p 2\rangle$ ..... 47
5.2.1 Offset Addressing ..... 48
5.2.2 Pre-Index Addressing ..... 49
5.2.3 Post-Index Addressing ..... 49
6 Programs ..... 51
6.1 Example Programs ..... 51
6.1.1 Program Listing Format ..... 51
6.1.2 Guidelines for Examples ..... 51
6.2 Trying the examples ..... 52
6.3 Trying the examples from the command line ..... 53
6.3.1 Setting up TextPad ..... 54
6.4 Program Initialization ..... 55
6.5 Special Conditions ..... 55
6.6 Problems ..... 55
7 Data Movement ..... 57
7.1 Program Examples ..... 57
7.1.1 16-Bit Data Transfer ..... 57
7.1.2 One's Complement ..... 58
7.1.3 32-Bit Addition ..... 59
7.1.4 Shift Left One Bit ..... 60
7.1.5 Byte Disassembly ..... 61
7.1.6 Find Larger of Two Numbers ..... 62
7.1.7 64-Bit Adition ..... 63
7.1.8 Table of Factorials ..... 64
7.2 Problems ..... 65
7.2.1 64-Bit Data Transfer ..... 65
7.2.2 $\quad$ 32-Bit Subtraction ..... 65
7.2.3 Shift Right Three Bits ..... 65
7.2.4 Halfword Assembly ..... 66
7.2.5 Find Smallest of Three Numbers ..... 66
7.2.6 Sum of Squares ..... 66
7.2.7 $\quad$ Shift Left $n$ bits ..... 66
8 Logic ..... 69
9 Program Loops ..... 71
9.1 Program Examples ..... 72
9.1.1 Sum of numbers ..... 72
9.1.2 Number of negative elements ..... 73
9.1.3 Find Maximum Value ..... 75
9.1.4 Normalize A Binary Number ..... 75
9.2 Problems ..... 76
9.2.1 Checksum of data ..... 76
9.2.2 Number of Zero, Positive, and Negative numbers ..... 77
9.2.3 Find Minimum ..... 77
9.2.4 Count 1 Bits ..... 77
9.2.5 Find element with most 1 bits ..... 77
10 Strings ..... 79
10.1 Handling data in ASCII ..... 79
10.2 A string of characters ..... 80
10.2.1 Fixed Length Strings ..... 81
10.2.2 Terminated Strings ..... 81
10.2.3 Counted Strings ..... 82
10.3 International Characters ..... 82
10.4 Program Examples ..... 82
10.4.1 Length of a String of Characters ..... 82
10.4.2 Find First Non-Blank Character ..... 84
10.4.3 Replace Leading Zeros with Blanks ..... 84
10.4.4 Add Even Parity to ASCII Chatacters ..... 85
10.4.5 Pattern Match ..... 86
10.5 Problems ..... 88
10.5.1 Length of a Teletypewriter Message ..... 88
10.5.2 Find Last Non-Blank Character ..... 88
10.5.3 Truncate Decimal String to Integer Form ..... 88
10.5.4 Check Even Parity and ASCII Characters ..... 89
10.5.5 String Comparison ..... 89
11 Code Conversion ..... 91
11.1 Program Examples ..... 91
11.1.1 Hexadecimal to ASCII ..... 91
11.1.2 Decimal to Seven-Segment. ..... 92
11.1.3 ASCII to Decimal ..... 93
11.1.4 Binary-Coded Decimal to Binary ..... 94
11.1.5 Binary Number to ASCII String ..... 95
11.2 Problems ..... 96
11.2.1 ASCII to Hexadecimal ..... 96
11.2.2 Seven-Segment to Decimal. ..... 96
11.2.3 Decimal to ASCII ..... 96
11.2.4 Binary to Binary-Coded-Decimal ..... 97
11.2.5 Packed Binary-Coded-Decimal to Binary String ..... 97
11.2.6 ASCII string to Binary number ..... 97
12 Arithmetic ..... 99
12.1 Program Examples ..... 99
12.1.2 64-Bit Addition ..... 99
12.1.3 Decimal Addition ..... 100
12.1.4 Multiplication. ..... 101
12.1.5 32-Bit Binary Divide ..... 102
12.2 Problems ..... 103
12.2.1 Multiple precision Binary subtraction ..... 103
12.2.2 Decimal Subtraction ..... 103
12.2.3 32-Bit by 32-Bit Multiply ..... 104
13 Tables and Lists ..... 105
13.1 Program Examples ..... 105
13.1.1 Add Entry to List ..... 105
13.1.2 Check an Ordered List ..... 106
13.1.3 Remove an Element from a Queue ..... 107
13.1.4 Sort a List ..... 108
13.1.5 Using an Ordered Jump Table ..... 109
13.2 Problems ..... 109
13.2.1 Remove Entry from List ..... 109
13.2.2 Add Entry to Ordered List ..... 109
13.2.3 Add Element to Queue ..... 109
13.2.4 4-Byte Sort ..... 110
13.2.5 Using a Jump Table with a Key ..... 110
14 The Stack ..... 111
15 Subroutines ..... 113
15.1 Types of Subroutines ..... 113
15.2 Subroutine Documentation ..... 114
15.3 Parameter Passing Techniques ..... 114
15.3.1 Passing Parameters In Registers ..... 114
15.3.2 Passing Parameters In A Parameter Block ..... 115
15.3.3 Passing Parameters On The Stack ..... 115
15.4 Types Of Parameters ..... 116
15.5 Program Examples ..... 116
15.6 Problems ..... 123
15.6.1 ASCII Hex to Binary ..... 123
15.6.2 ASCII Hex String to Binary Word ..... 123
15.6.3 Test for Alphabetic Character ..... 123
15.6.4 Scan to Next Non-alphabetic ..... 123
15.6.5 Check Even Parity ..... 124
15.6.6 Check the Checksum of a String ..... 124
15.6.7 Compare Two Counted Strings ..... 124
16 Interrupts and Exceptions ..... 125
A ARM Instruction Definitions ..... 127
A. 1 ADC: Add with Carry ..... 127
A. 2 ADD: Add ..... 128
A. 3 AND: Bitwise AND ..... 128
A. 4 B, BL: Branch, Branch and Link ..... 129
A. 5 CMP: Compare ..... 129
A. 6 EOR: Exclusive OR ..... 130
A. 7 LDM: Load Multiple ..... 130
A. 8 LDR: Load Register ..... 131
A. 9 LDRB: Load Register Byte ..... 131
A. 10 MOV: Move ..... 131
A. 11 MVN: Move Negative ..... 132
A. 12 ORR: Bitwise OR ..... 132
A. 13 SBC: Subtract with Carry ..... 133
A. 14 STM: Store Multiple ..... 133
A. 15 STR: Store Register ..... 134
A.16 STRB: Store Register Byte ..... 134
A. 17 SUB: Subtract ..... 135
A. 18 SWI: Software Interrupt ..... 135
A. 19 SWP: Swap ..... 136
A. 20 SWPB: Swap Byte ..... 136
B ARM Instruction Summary ..... 139

## List of Programs

7.1 move16.s 16bit data transfer ..... 57
7.2 invert.s Find the one's compliment (inverse) of a number ..... 58
7.3a add.s Add two numbers ..... 59
7.3b add2.s Add two numbers and store the result ..... 59
7.4 shiftleft.s Shift Left one bit ..... 60
7.5 nibble.s Disassemble a byte into its high and low order nibbles ..... 61
7.6 bigger.s Find the larger of two numbers ..... 62
7.7 add64.s 64 bit addition ..... 63
7.8 factorial.s Lookup the factorial from a table by using the address of the memory location ..... 64
8.7a bigger.s Find the larger of two numbers ..... 69
8.7 a add64.s 64 bit addition ..... 69
8.7a factorial.s Lookup the factorial from a table by using the address of the memory location ..... 70
9.1a sum16.s Add a series of 16 bit numbers by using a table address ..... 72
9.1b sum16b.s Add a series of 16 bit numbers by using a table address look-up ..... 72
9.2a countneg.s Scan a series of 32 bit numbers to find how many are negative ..... 73
9.2 b countneg16.s Scan a series of 16 bit numbers to find how many are negative ..... 74
9.3 largest16.s Scan a series of 16 bit numbers to find the largest ..... 75
9.4 normalize.s Normalize a binary number ..... 75
10.1a strlencr.s Find the length of a Carage Return terminated string ..... 82
10.1bstrlen.s Find the length of a null terminated string ..... 83
10.2 skipblanks.s Find first non-blank ..... 84
10.3 padzeros.s Supress leading zeros in a string ..... 84
10.4 setparity.s Set the parity bit on a series of characters store the amended string in Result ..... 85
10.5a cstrcmp.s Compare two counted strings for equality ..... 86
10.5 b strcmp.s Compare null terminated strings for equality assume that we have no $\begin{array}{ll}\text { knowledge of the data structure so we must assess the individual strings } & 87\end{array}$
11.1a nibtohex.s Convert a single hex digit to its ASCII equivalent ..... 91
11.1b wordtohex.s Convert a 32 bit hexadecimal number to an ASCII string and output to the terminal ..... 92
11.2 nibtoseg.s Convert a decimal number to seven segment binary ..... 92
11.3 dectonib.s Convert an ASCII numeric character to decimal ..... 93
11.4a ubcdtohalf.s Convert an unpacked BCD number to binary ..... 94
11.4b ubcdtohalf2.s Convert an unpacked BCD number to binary using MUL ..... 94
11.5 halftobin.s Store a 16bit binary number as an ASCII string of '0's and '1's ..... 95
12.2 add64.s $\quad 64$ Bit Addition ..... 99

| 12.3 addbcd.s | Add two packed BCD numbers to give a packed BCD result . . . . . 100 |
| :---: | :---: |
| 12.4a mul16.s | 16 bit binary multiplication . . . . . . . . . . . . . . . . . . . . . . . . 101 |
| 12.4 b mul32.s | Multiply two 32 bit number to give a 64 bit result (corrupts R0 and R1)101 |
| 12.5 divide.s | Divide a 32 bit binary no by a 16 bit binary no store the quotient and |
|  | remainder there is no 'DIV' instruction in ARM! . . . . . . . . . . . . 102 |
| 13.1a insert.s | Examine a table for a match - store a new entry at the end if no match |
|  | found . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 105 |
| 13.1b insert2.s | Examine a table for a match - store a new entry if no match found |
|  | extends insert.s . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 106 |
| 13.2 search.s | Examine an ordered table for a match . . . . . . . . . . . . . . . . . . 106 |
| 13.3 head.s | Remove the first element of a queue . . . . . . . . . . . . . . . . . . . 107 |
| 13.4 sort.s | Sort a list of values - simple bubble sort . . . . . . . . . . . . . . . . . 108 |
| 15.1a init1.s | Initiate a simple stack . . . . . . . . . . . . . . . . . . . . . . . . . . . 116 |
| 15.1b init2.s | Initiate a simple stack . . . . . . . . . . . . . . . . . . . . . . . . . . . 117 |
| 15.1c init3.s | Initiate a simple stack . . . . . . . . . . . . . . . . . . . . . . . . . . . 117 |
| 15.1d init3a.s | Initiate a simple stack . . . . . . . . . . . . . . . . . . . . . . . . . . . 118 |
| 15.1e byreg.s | A simple subroutine example program passes a variable to the routine |
|  | in a register . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 118 |
| 15.1f bystack.s | A more complex subroutine example program passes variables to the |
|  | routine using the stack . . . . . . . . . . . . . . . . . . . . . . . . . . 119 |
| 15.1g add64.s | A 64 bit addition subroutine . . . . . . . . . . . . . . . . . . . . . . . 121 |
| 15.1hfactorial.s | A subroutine to find the factorial of a number . . . . . . . . . . . . . 122 |

## Preface

Broadly speaking, you can divide the history of computers into four periods: the mainframe, the mini, the microprocessor, and the modern post-microprocessor. The mainframe era was characterized by computers that required large buildings and teams of technicians and operators to keep them going. More often than not, both academics and students had little direct contact with the mainframe-you handed a deck of punched cards to an operator and waited for the output to appear hours later. During the mainfame era, academics concentrated on languages and compilers, algorithms, and operating systems.

The minicomputer era put computers in the hands of students and academics, because university departments could now buy their own minis. As minicomputers were not as complex as mainframes and because students could get direct hands-on experience, many departments of computer science and electronic engineering taught students how to program in the native language of the computer-assembly language. In those days, the mid 1970s, assembly language programming was used to teach both the control of I/O devices, and the writing of programs (i.e., assembly language was taught rather like high level languages). The explosion of computer software had not taken place, and if you wanted software you had to write it yourself.

The late 1970s saw the introduction of the microprocessor. For the first time, each student was able to access a real computer. Unfortunately, microprocessors appeared before the introduction of low-cost memory (both primary and secondary). Students had to program microprocessors in assembly language because the only storage mechanism was often a ROM with just enough capacity to hold a simple single-pass assembler.

The advent of the low-cost microprocessor system (usually on a single board) ensured that virtually every student took a course on assembly language. Even today, most courses in computer science include a module on computer architecture and organization, and teaching students to write programs in assembly language forces them to understand the computer's architecture. However, some computer scientists who had been educated during the mainframe era were unhappy with the microprocessor, because they felt that the 8 -bit microprocessor was a retrograde step-its architecture was far more primitive than the mainframes they had studied in the 1960s.

The 1990s is the post-microprocessor era. Today's personal computers have more power and storage capacity than many of yesterday's mainframes, and they have a range of powerful software tools that were undreamed of in the 1970s. Moreover, the computer science curriculum of the 1990s has exploded. In 1970 a student could be expected to be familiar with all field of computer science. Today, a student can be expected only to browse through the highlights.

The availability of high-performance hardware and the drive to include more and more new material in the curriculum, has put pressure on academics to justify what they teach. In particular, many are questioning the need for courses on assembly language.
If you regard computer science as being primarily concerned with the use of the computer, you can argue that assembly language is an irrelevance. Does the surgeon study metallurgy in order to understand how a scalpel operates? Does the pilot study thermodynamics to understand how a jet engine operates? Does the news reader study electronics to understand how the camera
operates? The answer to all these questions is "no". So why should we inflict assembly language and computer architecture on the student?
First, education is not the same as training. The student of computer science is not simply being trained to use a number of computer packages. A university course leading to a degree should also cover the history and the theoretical basis for the subject. Without a knowledge of computer architecture, the computer scientist cannot understand how computers have developed and what they are capable of.

## Is assembly language today the same as assembly language yesterday?

Two factors have influenced the way in which we teach assembly language - one is the way in which microprocessors have changed, and the other is the use to which assembly language teaching is put. Over the years microprocessors have become more and more complex, with the result that the architecture and assembly language of a modern state-of-the-art microprocessor is radically different to that of an 8 -bit machine of the late 1970s. When we first taught assembly language in the 1970s and early 1980s, we did it to demonstrate how computers operated and to give students hands-on experience of a computer. Since all students either have their own computer or have access to a computer lab, this role of the single-board computer is now obsolete. Moreover, assembly language programming once attempted to ape high-level language programming - students were taught algorithms such as sorting and searching in assembly language, as if assembly language were no more than the (desperately) poor person's C.
The argument for teaching assembly language programming today can be divided into two components: the underpinning of computer architecture and the underpinning of computer software.

Assembly language teaches how a computer works at the machine (i.e., register) level. It is therefore necessary to teach assembly language to all those who might later be involved in computer architecture-either by specifying computers for a particular application, or by designing new architectures. Moreover, the von Neumann machine's sequential nature teaches students the limitation of conventional architectures and, indirectly, leads them on to unconventional architectures (parallel processors, Harvard architectures, data flow computers, and even neural networks).
It is probably in the realm of software that you can most easily build a case for the teaching of assembly language. During a student's career, he or she will encounter a lot of abstract concepts in subjects ranging from programming languages, to operating systems, to real-time programming, to AI. The foundation of many of these concepts lies in assembly language programming and computer architecture. You might even say that assembly language provides bottom-up support for the top-down methodology we teach in high-level languages. Consider some of the following examples (taken from the teaching of Advanced RISC Machines Ltd (ARM) assembly language).

## Data types

Students come across data types in high-level languages and the effects of strong and weak data typing. Teaching an assembly language that can operate on bit, byte, word and long word operands helps students understand data types. Moreover, the ability to perform any type of assembly language operation on any type of data structure demonstrates the need for strong typing.

## Addressing modes

A vital component of assembly language teaching is addressing modes (literal, direct, and indirect). The student learns how pointers function and how pointers are manipulated. This aspect is particularly important if the student is to become a C programmer. Because an assembly language is unencumbered by data types, the students' view of pointers is much simplified by an assembly language. The ARM has complex addressing modes that support direct and indirect addressing, generated jump tables and handling of unknown memory offsets.

## The stack and subroutines

How procedures are called, and parameters passed and returned from procedures. By using an assembly language you can readily teach the passing of parameters by value and by reference. The use of local variables and re-entrant programming can also be taught. This supports the teaching of task switching kernels in both operating systems and real-time programming.

## Recursion

The recursive calling of subroutines often causes a student problems. You can use an assembly language, together with a suitable system with a tracing facility, to demonstrate how recursion operates. The student can actually observe how the stack grows as procedures are called.

## Run-time support for high-level languages

A high-performance processor like the ARM provides facilities that support run-time checking in high-level languages. For example, the programming techniques document lists a series of programs that interface with ' C ' and provide run-time checking for errors such as an attempt to divide a number by zero.

## Protected-mode operation

Members of the ARM family operate in either a priviledge mode or a user mode. The operating system operates in the priviledge mode and all user (applications) programs run in the user mode. This mechanism can be used to construct secure or protected environments in which the effects of an error in one application can be prevented from harming the operating system (or other applications).

## Input-output

Many high-level languages make it difficult to access I/O ports and devices directly. By using an assembly language we can teach students how to write device drivers and how to control interfaces. Most real interfaces are still programmed at the machine level by accessing registers within them.

All these topics can, of course, be taught in the appropriate courses (e.g., high-level languages, operating systems). However, by teaching them in an assembly language course, they pave the way for future studies, and also show the student exactly what is happening within the machine.

## Conclusion

A strong case can be made for the continued teaching of assembly language within the computer science curriculum. However, an assembly language cannot be taught just as if it were another general-purpose programming language as it was once taught ten years ago. Perhaps more than any other component of the computer science curriculum, teaching an assembly language supports a wide range of topics at the heart of computer science. An assembly language should not be used just to illustrate algorithms, but to demonstrate what is actually happening inside the computer.

## 1 Introduction

A computer program is ultimately a series of numbers and therefore has very little meaning to a human being. In this chapter we will discuss the levels of human-like language in which a computer program may be expressed. We will also discuss the reasons for and uses of assembly language.

### 1.1 The Meaning of Instructions

The instruction set of a microprocessor is the set of binary inputs that produce defined actions during an instruction cycle. An instruction set is to a microprocessor what a function table is to a logic device such as a gate, adder, or shift register. Of course, the actions that the microprocessor performs in response to its instruction inputs are far more complex than the actions that logic devices perform in response to their inputs.

### 1.1.1 Binary Instructions

An instruction is a binary digit pattern - it must be available at the data inputs to the microprocessor at the proper time in order to be interpreted as an instruction. For example, when the ARM receives the binary pattern 111000000100 as the input during an instruction fetch operation, the pattern means subtract. Similary the microinstruction 111000001000 means add. Thus the 32 bit pattern 11100000010011101100000000001111 means:
"Subtract R15 from R14 and put the answer in R12."
The microprocessor (like any other computer) only recognises binary patterns as instructions or data; it does not recognise characters or octal, decimal, or hexadecimal numbers.

### 1.2 A Computer Program

A program is a series of instructions that causes a computer to perform a particular task.
Actually, a computer program includes more than instructions, it also contains the data and the memory addresses that the microprocessor needs to accomplish the tasks defined by the instructions. Clearly, if the microprocessor is to perform an addition, it must have two numbers to add and a place to put the result. The computer program must determine the sources of the data and the destination of the result as well as the operation to be performed.
All microprocessors execute instructions sequentially unless an instruction changes the order of execution or halts the processor. That is, the processor gets its next instruction from the next higher memory address unless the current instruction specifically directs it to do otherwise.
Ultimately, every program is a set of binary numbers. For example, this is a snippet of an ARM program that adds the contents of memory locations $8094_{32}$ and $8098_{32}$ and places the result in memory location $809 C_{32}$ :

```
11100101100111110001000000010000
11100101100111110001000000001000
11100000100000010101000000000000
11100101100011110101000000001000
```

This is a machine language, or object, program. If this program were entered into the memory of an ARM-based microcomputer, the microcomputer would be able to execute it directly.

### 1.3 The Binary Programming Problem

There are many difficulties associated with creating programs as object, or binary machine language, programs. These are some of the problems:

- The programs are difficult to understand or debug. (Binary numbers all look the same, particularly after you have looked at them for a few hours.)
- The programs do not describe the task which you want the computer to perform in anything resembling a human-readable format.
- The programs are long and tiresome to write.
- The programmer often makes careless errors that are very difficult to locate and correct.

For example, the following version of the addition object program contains a single bit error. Try to find it:

```
11100101100111110001000000010000
11100101100111110001000000001000
11100000100000010101000000000000
11100110100011110101000000001000
```

Although the computer handles binary numbers with ease, people do not. People find binary programs long, tiresome, confusing, and meaningless. Eventually, a programmer may start remembering some of the binary codes, but such effort should be spent more productively.

### 1.4 Using Octal or Hexadecimal

We can improve the situation somewhat by writing instructions using octal or hexadecimal numbers, rather than binary. We will use hexadecimal numbers because they are shorter, and because they are the standard for the microprocessor industry. Table 1.1 defines the hexadecimal digits and their binary equivalents. The ARM program to add two numbers now becomes:

```
E59F1010
E59f0008
E0815000
E58F5008
```

At the very least, the hexadecimal version is shorter to write and not quite so tiring to examine. Errors are somewhat easier to find in a sequence of hexadecimal digits. The erroneous version of the addition program, in hexadecimal form, becomes:

| Hexadecimal <br> Digit | Binary <br> Equivalent | Decimal <br> Equivalent |
| :---: | :---: | :---: |
| 0 | 0000 | 0 |
| 1 | 0001 | 1 |
| 2 | 0010 | 2 |
| 3 | 0011 | 3 |
| 4 | 0100 | 4 |
| 5 | 0101 | 5 |
| 6 | 0110 | 6 |
| 7 | 0111 | 7 |
| 8 | 1000 | 8 |
| 9 | 1001 | 9 |
| A | 1010 | 10 |
| B | 1011 | 11 |
| C | 1100 | 12 |
| D | 1101 | 13 |
| E | 1110 | 14 |
| F | 1111 | 15 |

Table 1.1: Hexadecimal Conversion Table

$$
\begin{aligned}
& \text { E59F1010 } \\
& \text { E59f0008 } \\
& \text { E0815000 } \\
& \text { E68F5008 }
\end{aligned}
$$

The mistake is far more obvious.
The hexadecimal version of the program is still difficult to read or understand; for example, it does not distinguish operations from data or addresses, nor does the program listing provide any suggestion as to what the program does. What does 3038 or 31 C 0 mean? Memorising a card full of codes is hardly an appetising proposition. Furthermore, the codes will be entirely different for a different microprocessor and the program will require a large amount of documentation.

### 1.5 Instruction Code Mnemonics

An obvious programming improvement is to assign a name to each instruction code. The instruction code name is called a "mnemonic" or memory jogger.

In fact, all microprocessor manufacturers provide a set of mnemonics for the microprocessor instruction set (they cannot remember hexadecimal codes either). You do not have to abide by the manufacturer's mnemonics; there is nothing sacred about them. However, they are standard for a given microprocessor, and therefore understood by all users. These are the instruction codes that you will find in manuals, cards, books, articles, and programs. The problem with selecting instruction mnemonics is that not all instructions have "obvious" names. Some instructions do (for example, ADD, AND, ORR), others have obvious contractions (such as SUB for subtraction, EOR for exclusive-OR), while still others have neither. The result is such mnemonics as BIC, STMIA, and even MRS. Most manufacturers come up with some reasonable names and some hopeless ones. However, users who devise their own mnemonics rarely do much better.

Along with the instruction mnemonics, the manufacturer will usually assign names to the CPU registers. As with the instruction names, some register names are obvious (such as A for Accumulator) while others may have only historical significance. Again, we will use the manufacturer's suggestions simply to promote standardisation.

If we use standard ARM instruction and register mnemonics, as defined by Advanced RISC Machines, our ARM addition program becomes:

```
LDR R1, num1
LDR R0, num2
ADD R5, R1, R0
STR R5, num3
```

The program is still far from obvious, but at least some parts are comprehensible. ADD is a considerable improvement over E59F. The LDR mnemonic does suggest loading data into a register or memory location. We now see that some parts of the program are operations and others are addresses. Such a program is an assembly language program.

### 1.6 The Assembler Program

How do we get the assembly language program into the computer? We have to translate it, either into hexadecimal or into binary numbers. You can translate an assembly language program by hand, instruction by instruction. This is called hand assembly.

The following table illustrates the hand assembly of the addition program:

| Instruction Mnemonic | Register/Memory Location | Hexadecimal Equivalent |
| :---: | :---: | :---: |
| LDR | R1, num1 | E59F1010 |
| LDR | R0, num2 | E59F0008 |
| ADD | R5, R1, R0 | E0815000 |
| STR | R5, num3 | E58F5008 |

Hand assembly is a rote task which is uninteresting, repetitive, and subject to numerous minor errors. Picking the wrong line, transposing digits, omitting instructions, and misreading the codes are only a few of the mistakes that you may make. Most microprocessors complicate the task even further by having instructions with different lengths. Some instructions are one word long while others may be two or three. Some instructions require data in the second and third words; others require memory addresses, register numbers, or who knows what?
Assembly is a rote task that we can assign to the microcomputer. The microcomputer never makes any mistakes when translating codes; it always knows how many words and what format each instruction requires. The program that does this job is an "assembler." The assembler program translates a user program, or "source" program written with mnemonics, into a machine language program, or "object" program, which the microcomputer can execute. The assembler's input is a source program and its output is an object program.
Assemblers have their own rules that you must learn. These include the use of certain markers (such as spaces, commas, semicolons, or colons) in appropriate places, correct spelling, the proper control of information, and perhaps even the correct placement of names and numbers. These rules are usually simple and can be learned quickly.

### 1.6.1 Additional Features of Assemblers

Early assemblers did little more than translate the mnemonic names of instructions and registers into their binary equivalents. However, most assemblers now provide such additional features as:

- Allowing the user to assign names to memory locations, input and output devices, and even sequences of instructions
- Converting data or addresses from various number systems (for example, decimal or hexadecimal) to binary and converting characters into their ASCII or EBCDIC binary codes
- Performing some arithmetic as part of the assembly process
- Telling the loader program where in memory parts of the program or data should be placed
- Allowing the user to assign areas of memory as temporary data storage and to place fixed data in areas of program memory
- Providing the information required to include standard programs from program libraries, or programs written at some other time, in the current program
- Allowing the user to control the format of the program listing and the input and output devices employed


### 1.6.2 Choosing an Assembler

All of these features, of course, involve additional cost and memory. Microcomputers generally have much simpler assemblers than do larger computers, but the tendency is always for the size of assemblers to increase. You will often have a choice of assemblers. The important criterion is not how many off-beat features the assembler has, but rather how convenient it is to use in normal practice.

### 1.7 Disadvantages of Assembly Language

The assembler does not solve all the problems of programming. One problem is the tremendous gap between the microcomputer instruction set and the tasks which the microcomputer is to perform. Computer instructions tend to do things like add the contents of two registers, shift the contents of the Accumulator one bit, or place a new value in the Program Counter. On the other hand, a user generally wants a microcomputer to do something like print a number, look for and react to a particular command from a teletypewriter, or activate a relay at the proper time. An assembly language programmer must translate such tasks into a sequence of simple computer instructions. The translation can be a difficult, time-consuming job.
Furthermore, if you are programming in assembly language, you must have detailed knowledge of the particular microcomputer that you are using. You must know what registers and instructions the microcomputers has, precisely how the instructions affect the various registers, what addressing methods the computer uses, and a mass of other information. None of this information is relevant to the task which the microcomputer must ultimately perform.

In addition, assembly language programs are not portable. Each microcomputer has its own assembly language which reflects its own architecture. An assembly language program written for the ARM will not run on a 486, Pentium, or Z8000 microprocessor. For example, the addition program written for the Z 8000 would be:

$$
\begin{array}{lr}
\text { LD } & \text { RO, } \% 6000 \\
\text { ADD } & \text { RO, } \% 6002 \\
\text { LD } & \% 6004, \text { RO }
\end{array}
$$

The lack of portability not only means that you will not be able to use your assembly language program on a different microcomputer, but also that you will not be able to use any programs that were not specifically written for the microcomputer you are using. This is a particular drawback for new microcomputers, since few assembly language programs exist for them. The result, too frequently, is that you are on your own. If you need a program to perform a particular task, you are not likely to find it in the small program libraries that most manufacturers provide. Nor are you likely to find it in an archive, journal article, or someone's old program File. You will probably have to write it yourself.

### 1.8 High-Level Languages

The solution to many of the difficulties associated with assembly language programs is to use, insted, high-level or procedure-oriented langauges. Such languages allow you to describe tasks in forms that are problem-oriented rather than computer-oriented. Each statement in a high-level language performs a recognisable function; it will generally correspond to many assembly language instruction. A program called a compiler translates the high-level language source program into object code or machine language instructions.

Many different hgih-level languages exist for different types of tasks. If, for exampe, you can express what you want the computer to do in algebraic notation, you can write your FORTRAN (Formula Translation Language), the oldest of the high-level languages. Now, if you want to add two numbers, you just tell the computer:

$$
\text { sum }=\text { num1 }+ \text { num2; }
$$

That is a lot simpler (and shorter) than either the equivalent machine language program or the equivalent assembly language program. Other high-level languages include COBOL (for business applications), BASIC (a cut down version of FORTRAN designed to prototype ideas before codeing them in full), C (a systems-programming language), C++ and JAVA (object-orientated general development languages).

### 1.8.1 Advantages of High-Level Languages

Clearly, high-level languages make program easier and faster to write. A common estimate is that a programmer can write a program about ten times as fast in a high-level langauge as in assembly language. That is just writing the program; it does not include problem definition, program design, debugging testing or documentation, all of which become simpler and faster. The high-level language program is, for instance, partly self-documenting. Even if you do not know FORTRAN, you could probably tell what the statement illustrated above does.

## Machine Independence

High-level languages solve many other problems associated with assembly language programming. The high-level language has its own syntax (usually defined by an international standard). The language does not mention the instruction set, registers, or other features of a particular computer. The compiler takes care of all such details. Programmers can concentrate on their own tasks; they do not need a detailed understanding of the underlying CPU architecture - for that matter, they do not need to know anything about the computer the are programming.

## Portability

Programs written in a high-level language are portable - at least, in theory. They will run on any computer that has a standard compiler for that language.

At the same time, all previous programs written in a high-level language for prior computers and available to you when programming a new computer. This can mean thousands of programs in the case of a common language like C .

### 1.8.2 Disadvantages of High-Level Languages

If all the good things we have said about high-level languages are true - if you can write programs faster and make them portable besides - why bother with assebly languages? Who wants to worry about registers, instruction codes, mnemonics, and all that garbage! As usual, there are disadvantages that balance the advantages.

## Syntax

One obvious problem is that, as with assembly language, you have to learn the "rules" or syntax of any high-level language you want to use. A high-level langauge has a fairly complicated set of rules. You will find that it takes a lot of time just to get a program that is syntactically correct (and even then it probably will not do what you want). A high-level computer language is like a foreign language. If you have talent, you will get used to the rules and be able to turn out programs that the compiler will accept. Still, learning the rules and trying to get the program accepted by the compiler does not contribute directly to doing your job.

## Cost of Compilers

Another obvious problem is that you need a compiler to translate program written in a high-level language into machine language. Compilers are expensive and use a large amount of memory. While most assemblers occupy only a few KBytes of memory, compilers would occupy far larger amounts of memory. A compiler could easily require over four times as much memory as an assembler. So the amount of overhead involved in using the compiler is rather large.

## Adapting Tasks to a Language

Furthermore, only some compilers will make the implementation of your task simpler. Each language has its own target proglem area, for example, FORTRAN is well-suited to problems that can be expressed as algebraic formulas. If however, your problem is controlling a display terminal, editing a string of characters, or monitoring an alarm system, your problem cannot be easily expressed. In fact, formulating the solution in FORTRAN may be more awkward and more difficult than formulating it in assembly language. The answer is, of course, to use a more suitable high-level language. Languages specifically designed for tasks such as those mentioned above do exist - they are called system implementation languages. However, these languages are less widely used.

## Inefficiency

High-level languages do not produce very efficient machine language program. The basic reason for this is that compilation is an automatic process which is riddled with compromises to allow for many ranges of possibilities. The compiler works much like a computerised language translator sometimes the words are right but the sentence structures are awkward. A simpler compiler connot know when a variable is no longer being used and can be discarded, when a register should be used rather than a memory location, or when variables have simple relationships. The experienced programmer can take advantage of shortcuts to shorten execution time or reduce memory usage. A few compiler (known as optimizing cmpilers) can also do this, but such compilers are much larger than regular compilers.

### 1.9 Which Level Should You Use?

Which language level you use depends on your particulr application. Let us briefly note some of the factors which may favor particular levels:

### 1.9.1 Applications for Machine Language

Virtually no one programs in machine language because it wastes human time and is difficult to document. An assembler costs very little and greatly reduces programming time.

### 1.9.2 Applications for Assembly Language

- Limited data processing
- High-volume applications
- Real-Time control applications
- Short to moderate-sized programs
- Application where memory cost is a factor
- Applications involving more input/output or control than computation


### 1.9.3 Applications for High-Level Language

- Long programs
- Low-volume applications
- Programs which are expected to undergo many changes
- Compatibility with similar applications using larger computers
- Applications involing more computation than input/output or control
- Applications where the amout of memory required is already very large
- Availability of a specific program in a high-level language which can be used in the application.


### 1.9.4 Other Considerations

Many other factors are also important, such as the availability of a large computer for use in development, experience with particular languages, and compatibility with other applications.
If hardware will ultimately be the largest cost in your application, or if speed is critical, you should favor assembly language. But be prepared to spend much extra time in software development in exchange for lower memory costs and higher execution speeds. If software will be the largest cost in your application, you should favor a high-level language. But be prepared to spend the extra money required for the supporting hardware and software.
Of course, no one except some theorists will object if you use both assembly and high-level languages. You can write the program originally in a high-level language and then patch some sections in assembly language. However, most users prefer not to do this because it can create havoc in debugging, testing, and documentation.

### 1.10 Why Learn Assembler?

Given the advance of high-level languages, why do you need to learn assembly language programming? The reasons are:

1. Most industrial microcomputer users program in assembly language.
2. Many microcomputer users will continue to program in assembly language since they need the detailed control that it provides.
3. No suitable high-level language has yet become widely available or standardised.
4. Many application require the efficiency of assembly language.
5. An understanding of assembly language can help in evaluating high-level languages.
6. Almost all microcomputer programmers ultimately find that they need some knowledge of assembly language, most often to debug programs, write I/O routines, speed up or shorten critical sections of programs written in high-level languages, utilize or modify operating system functions, and undertand other people's programs.

The rest of these notes will deal exclusively with assembler and assembly language programming.

## 2 Assemblers

This chapter discusses the functions performed by assemblers, beginning with features common to most assemblers and proceeding through more elaborate capabilities such as macros and conditional assembly. You may wish to skim this chapter for the present and return to it when you feel more comfortable with the material.

As we mentioned, today's assemblers do much more than translate assembly language mnemonics into binary codes. But we will describe how an assembler handles the translation of memonics before describing additional assembler features. Finally we will explain how assemblers are used.

### 2.1 Fields

Assembly language instructions (or "statements") are divided into a number of "fields".
The operation code field is the only field which can never he empty; it always contains either an instruction mnemonic or a directive to the assembler, sometimes called a "pseudo-instruction," "pseudo-operation," or "pseudo-op."

The operand or address field may contain an address or data, or it may be blank.
The comment and label fields are optional. A programmer will assign a label to a statement or add a comment as a personal convenience: namely, to make the program easier to read and use.

Of course, the assembler must have some way of telling where one field ends and another begins. Assemblers often require that each field start in a specific column. This is a "fixed format." However, fixed formats are inconvenient when the input medium is paper tape; fixed formats are also a nuisance to programmers. The alternative is a "free format" where the fields may appear anywhere on the line.

### 2.1.1 Delimiters

If the assembler cannot use the position on the line to tell the fields apart, it must use something else. Most assemblers use a special symbol or "delimiter" at the beginning or end of each field.

| Label <br> Field | Operation Code or Mnemonic Field | Operand or Address Field | Comment Field |
| :---: | :---: | :---: | :---: |
| VALUE1 | DCW | 0x201E | ;FIRST VALUE |
| VALUE2 | DCW | 0x0774 | ;SECOND VALUE |
| RESULT | DCW | 1 | ; 16-BIT STORAGE FOR ADDITION RESULT |
| START | MOV | R0, VALUE1 | ;GET FIRST VALUE |
|  | ADD | RO, RO, VALUE2 | ; ADD SECOND VALUE TO FIRST VALUE |
|  | STR | RESULT, RO | ;STORE RESULT OF ADDITION |
| NEXT: | ? | ? | ; NEXT INSTRUCTION |

label $\langle$ whitespace $\rangle$ instruction $\langle$ whitespace $\rangle$; comment

| whitespace | Between label and operation code, between operation code and ad- <br> dress, and before an entry in the comment field |
| :--- | :--- |
| comma | Between operands in the address field |
| asterisk | Before an entire line of comment |
| semicolon | Marks the start of a comment on a line that contains preceding code |

Table 2.1: Standard ARM Assembler Delimiters

The most common delimiter is the space character. Commas, periods, semicolons, colons, slashes, question marks, and other characters that would not otherwise be used in assembly language programs also may serve as delimiters. The general form of layout for the ARM assembler is:

You will have to exercise a little care with delimiters. Some assemblers are fussy about extra spaces or the appearance of delimiters in comments or labels. A well-written assembler will handle these minor problems, but many assemblers are not well-written. Our recommendation is simple: avoid potential problems if you can. The following rules will help:

- Do not use extra spaces, in particular, do not put spaces after commas that separate operands, even though the ARM assembler allows you to do this.
- Do not use delimiter characters in names or labels.
- Include standard delimiters even if your assembler does not require them. Then it will be more likely that your programs are in correct form for another assembler.


### 2.1.2 Labels

The label field is the first field in an assembly language instruction; it may be blank. If a label is present, the assembler defines the label as equivalent to the address into which the first byte of the object code generated for that instruction will be loaded. You may subsequently use the label as an address or as data in another instruction's address field. The assembler will replace the label with the assigned value when creating an object program.

The ARM assembler requires labels to start at the first character of a line. However, some other assemblers also allow you to have the label start anywhere along a line, in which case you must use a colon (: ) as the delimiter to terminate the label field. Colon delimiters are not used by the ARM assembler.

Labels are most frequently used in Branch or SWI instructions. These instructions place a new value in the program counter and so alter the normal sequential execution of instructions. B $150_{16}$ means "place the value $150_{16}$ in the program counter." The next instruction to be executed will be the one in memory location $150_{16}$. The instruction B START means "place the value assigned to the label START in the program counter." The next instruction to be executed will be the on at the address corresponding to the label START. Figure 2.1 contains an example.

Why use a label? Here are some reasons:

- A label makes a program location easier to find and remember.
- The label can easily be moved, if required, to change or correct a program. The assembler will automatically change all instructions that use the label when the program is reassembled.

Assembly language Program

```
START MOV RO, VALUE1
    . (Main Program)
    BAL START
```

When the machine language version of this program is executed, the instruction B START causes the address of the instruction labeled START to be placed in the program counter That instruction will then be executed.

Figure 2.1: Assigning and Using a Label

- The assembler can relocate the whole program by adding a constant (a "relocation constant") to each address in which a label was used. Thus we can move the program to allow for the insertion of other programs or simply to rearrange memory.
- The program is easier to use as a library program; that is, it is easier for someone else to take your program and add it to some totally different program.
- You do not have to figure out memory addresses. Figuring out memory addresses is particularly difficult with microprocessors which have instructions that vary in length.

You should assign a label to any instruction that you might want to refer to later.
The next question is how to choose a label. The assembler often places some restrictions on the number of characters (usually 5 or 6 ), the leading character (often must be a letter), and the trailing characters (often must be letters, numbers, or one of a few special characters). Beyond these restrictions, the choice is up to you.
Our own preference is to use labels that suggest their purpose, i.e., mnemonic labels. Typical examples are ADDW in a routine that adds one word into a sum, SRCHETX in a routine that searches for the ASCII character ETX, or NKEYS for a location in data memory that contains the number of key entries. Meaningful labels are easier to remember and contribute to program documentation. Some programmers use a standard format for labels, such as starting with L0000. These labels are self-sequencing (you can skip a few numbers to permit insertions), but they do not help document the program.

Some label selection rules will keep you out of trouble. We recommend the following:

- Do not use labels that are the same as operation codes or other mnemonics. Most assemblers will not allow this usage; others will, but it is confusing.
- Do not use labels that are longer than the assembler recognises. Assemblers have various rules, and often ignore some of the characters at the end of a long label.
- Avoid special characters (non-alphabetic and non-numeric) and lower-case letters. Some assemblers will not permit them; others allow only certain ones. The simplest practice is to stick to capital letters and numbers.
- Start each label with a letter. Such labels are always acceptable.
- Do not use labels that could be confused with each other. Avoid the letters I, O, and Z and the numbers 0,1 , and 2. Also avoid things like XXXX and XXXXX. Assembly programming is difficult enough without tempting fate or Murphy's Law.
- When you are not sure if a label is legal, do not use it. You will not get any real benefit from discovering exactly what the assembler will accept.

These are recommendations, not rules. You do not have to follow them but don't blame us if you waste time on unnecessary problems.

### 2.2 Operation Codes (Mnemonics)

One main task of the assembler is the translation of mnemonic operation codes into their binary equivalents. The assembler performs this task using a fixed table much as you would if you were doing the assembly by hand.

The assembler must, however, do more than just translate the operation codes. It must also somehow determine how many operands the instruction requires and what type they are. This may be rather complex - some instructions (like a Stop) have no operands, others (like a Jump instruction) have one, while still others (like a transfer between registers or a multiple-bit shift) require two. Some instructions may even allow alternatives; for example, some computers have instructions (like Shift or Clear) which can either apply to a register in the CPU or to a memory location. We will not discuss how the assembler makes these distinctions; we will just note that it must do so.

### 2.3 Directives

Some assembly language instructions are not directly translated into machine language instructions. These instructions are directives to the assembler; they assign the program to certain areas in memory, define symbols, designate areas of memory for data storage, place tables or other fixed data in memory, allow references to other programs, and perform minor housekeeping functions.

To use these assembler directives or pseudo-operations a programmer places the directive's mnemonic in the operation code field, and, if the specified directive requires it, an address or data in the address field.

The most common directives are:

```
DEFINE CONSTANT (Data)
EQUATE (Define)
AREA
DEFINE STORAGE (Reserve)
```

Different assemblers use different names for those operations but their functions are the same. Housekeeping directives include:
END LIST FORMAT TTL PAGE INCLUDE

We will discuss these pseudo-operations briefly, although their functions are usually obvious.

### 2.3.1 The DEFINE CONSTANT (Data) Directive

The DEFINE CONSTANT directive allows the programmer to enter fixed data into program memory. This data may include:

- Names
- Conversion factors
- Messages
- Key identifications
- Commands
- Subroutine addresses
- Tax tables
- Code conversion tables
- Thresholds
- Identification patterns
- Test patterns
- State transition tables
- Lookup tables
- Synchronisation patterns
- Standard forms
- Coefficients for equations
- Masking patterns
- Character generation patterns
- Weighting factors
- Characteristic times or frequencies

The define constant directive treats the data as a permanent part of the program.
The format of a define constant directive is usually quite simple. An instruction like:

```
DZCON DCW 12
```

will place the number 12 in the next available memory location and assign that location the name DZCON. Every DC directive usually has a label, unless it is one of a series. The data and label may take any form that the assembler permits.
More elaborate define constant directives that handle a large amount of data at one time are provided, for example:

```
EMESS DCB 'ERROR'
SQRS DCW 1,4,9,16,25
```

A single directive may fill many bytes of program memory, limited perhaps by the length of a line or by the restrictions of a particular assembler. Of course, you can always overcome any restrictions by following one define constant directive with another:

```
MESSG DCB "NOW IS THE "
    DCB "TIME FOR ALL "
    DCB "GOOD MEN "
    DCB "TO COME TO THE "
    DCB "AID OF THEIR "
    DCB "COUNTRY", 0 ;note the '0' terminating the string
```

Microprocessor assemblers typically have some variations of standard define constant directives. Define Byte or DCB handles 8-bit numbers; Define Word or DCW handles 32-bit numbers or addresses. Other special directives may handle character-coded data. The ARM assembler also defines DCD to (Define Constant Data) which may be used in place of DCW.

### 2.3.2 The EQUATE Directive

The EQUATE directive allows the programmer to equate names with addresses or data. This pseudo-operation is almost always given the mnemonic EQU. The names may refer to device addresses, numeric data, starting addresses, fixed addresses, etc.

The EQUATE directive assigns the numeric value in its operand field to the label in its label field. Here are two examples:

| TTY | EQU | 5 |
| :--- | :--- | :--- |
| LAST | EQU | 5000 |

Most assemblers will allow you to define one label in terms of another, for example:

| LAST | EQU | FINAL |
| :--- | :--- | :--- |
| ST1 | EQU | START+1 |

The label in the operand field must, of course, have been previously defined. Often, the operand field may contain more complex expressions, as we shall see later. Double name assignments (two names for the same data or address) may be useful in patching together programs that use different names for the same variable (or different spellings of what was supposed to be the same name).

Note that an EQU directive does not cause the assembler to place anything in memory. The assembler simply enters an additional name into a table (called a "symbol table") which the assembler maintains.

When do you use a name? The answer is: whenever you have a parameter that you might want to change or that has some meaning besides its ordinary numeric value. We typically assign names to time constants, device addresses, masking patterns, conversion factors, and the like. A name like DELAY, TTY, KBD, KROW, or OPEN not only makes the parameter easier to change, but it also adds to program documentation. We also assign names to memory locations that have special purposes; they may hold data, mark the start of the program, or be available for intermediate storage.

What name do you use? The best rules are much the same as in the case of labels, except that here meaningful names really count. Why not call the teletypewriter TTY instead of X15, a bit time delay BTIME or BTDLY rather than WW, the number of the "GO" key on a keyboard GOKEY rather than HORSE? This advice seems straightforward, but a surprising number of programmers do not follow it.

Where do you place the EQUATE directives? The best place is at the start of the program, under appropriate comment headings such as I/O ADDRESSES, TEMPORARY STORAGE, TIME CONSTANTS, or program locations. This makes the definitions easy to find if you want to change them. Furthermore, another user will be able to look up all the definitions in one centralised place. Clearly this practice improves documentation and makes the program easier to use.

Definitions used only in a specific subroutine should appear at the start of the subroutine.

### 2.3.3 The AREA Directive

The AREA directive allows the programmer to specify the memory locations where programs, subroutines, or data will reside. Programs and data may be located in different areas of memory depending on the memory configuration. Startup routines interrupt service routines, and other required programs may be scattered around memory at fixed or convenient addresses.

The assembler maintains a location counter (comparable to the computer's program counter) which contains the location in memory of the instruction or data item being processed. An area directive causes the assembler to place a new value in the location counter, much as a Jump instruction causes the CPU to place a new value in the program counter. The output from the assembler must not only contain instructions and data, but must also indicate to the loader program where in memory it should place the instructions and data.

Microprocessor programs often contain several AREA statements for the following purposes:

- Reset (startup) address
- Interrupt service addresses
- Trap (software interrupt) addresses
- RAM storage
- Stack
- Main program
- Subroutines
- Input/Output

Still other origin statements may allow room for later insertions, place tables or data in memory, or assign vacant memory space for data buffers. Program and data memory in microcomputers may occupy widely separate addresses to simplify the hardware. Typical origin statements are:

| AREA | RESET |
| :--- | :---: |
| AREA | $\$ 1000$ |
| AREA | INT3 |

The assembler will assume a fake address if the programmer does not put in an AREA statement. The AREA statement at the start of an ARM program is required, and its absence will cause the assembly to fail.

### 2.3.4 Housekeeping Directives

There are various assembler directives that affect the operation of the assembler and its program listing rather than the object program itself. Common directives include:

END, marks the end of the assembly language source program. This must appear in the file or a "missing END directive" error will occur.
INCLUDE will include the contents of a named file into the current file. When the included file has been processed the assembler will continue with the next line in the original file. For example the following line

INCLUDE MATH.S
will include the content of the file math.s at that point of the file.
You should never use a lable with an include directive. Any labels defined in the included file will be defined in the current file, hence an error will be reported if the same label appears in both the source and include file.
An include file may itself include other files, which in turn could include other files, and so on, however, the level of includes the assembler will accept is limited. It is not recommended you go beyond three levels for even the most complex of software.

### 2.3.5 When to Use Labels

Users often wonder if or when they can assign a label to an assembler directive. These are our recommendations:

1. All EQU directives must have labels; they are useless otherwise, since the purpose of an EQU is to define its label.
2. Define Constant and Define Storage directives usually have labels. The label identifies the first memory location used or assigned.
3. Other directives should not have labels.

### 2.4 Operands and Addresses

The assembler allow the programmer a lot of freedom in describing the contents of the operand or address field. But remember that the assembler has built-in names for registers and instructions and may have other built-in names. We will now describe some common options for the operand field.

### 2.4.1 Decimal Numbers

The assembler assume all numbers to be decimal unless they are marked otherwise. So:

```
ADD 100
```

means "add the contents of memory location $100_{10}$ to the contents of the Accumulator."

### 2.4.2 Other Number Systems

The assembler will also accept hexadecimal entries. But you must identify these number systems in some way: for example, by preceding the number with an identifying character.

| $2 \_n n n$ | Binary | Base 2 |
| :--- | :--- | :--- |
| $8 \_n n n$ | Octal | Base 8 |
| $n n n$ | Decimal | Base 10 |
| $0 \times n n n$ | Hexadecimal | Base 16 |

It is good practice to enter numbers in the base in which their meaning is the clearest: that is, decimal constants in decimal; addresses and BCD numbers in hexadecimal; masking patterns or bit outputs in hexadecimal.

### 2.4.3 Names

Names can appear in the operand field; they will be treated as the data that they represent. Remember, however, that there is a difference between operands and addresses. In an ARM assembly language program the sequence:

FIVE EQU 5
ADD R2, \#FIVE
will add the contents of memory location FIVE (not necessarily the number 5) to the contents of data register R2.

### 2.4.4 Character Codes

The assembler allows text to be entered as ASCII strings. Such strings must be surrounded with double quotation marks, unless a single ASCII character is quoted, when single qoutes may be used exactly as in 'C'. We recommend that you use character strings for all text. It improves the clarity and readability of the program.

### 2.4.5 Arithmetic and Logical Expressions

Assemblers permit combinations of the data forms described above, connected by arithmetic, logical, or special operators. These combinations are called expressions. Almost all assemblers allow simple arithmetic expressions such as START+1. Some assemblers also permit multiplication, division, logical functions, shifts, etc. Note that the assembler evaluates expressions at assembly time; if a symbol appears in an expression, the address is used (i.e., the location counter or EQUATE value).

Assemblers vary in what expressions they accept and how they interpret them. Complex expressions make a program difficult to read and understand.

### 2.4.6 General Recommendations

We have made some recommendations during this section but will repeat them and add others here. In general, the user should strive for clarity and simplicity. There is no payoff for being an expert in the intricacies of an assembler or in having the most complex expression on the block. We suggest the following approach:

- Use the clearest number system or character code for data.
- Masks and BCD numbers in decimal, ASCII characters in octal, or ordinary numerical constants in hexadecimal serve no purpose and therefore should not be used.
- Remember to distinguish data from addresses.
- Don't use offsets from the location counter.
- Keep expressions simple and obvious. Don't rely on obscure features of the assembler.


### 2.5 Comments

All assemblers allow you to place comments in a source program. Comments have no effect on the object code, but they help you to read, understand, and document the program. Good commenting is an essential part of writing computer programs, programs without comments are very difficult to understand.

We will discuss commenting along with documentation in a later chapter, but here are some guidelines:

- Use comments to tell what application task the program is performing, not how the microcomputer executes the instructions.
- Comments should say things like "is temperature above limit?", "linefeed to TTY," or "examine load switch."
- Comments should not say things like "add 1 to Accumulator," "jump to Start," or "look at carry." You should describe how the program is affecting the system; internal effects on the CPU should be obvious from the code.
- Keep comments brief and to the point. Details should be available elsewhere in the documentation.
- Comment all key points.
- Do not comment standard instructions or sequences that change counters or pointers; pay special attention to instructions that may not have an obvious meaning.
- Do not use obscure abbreviations.
- Make the comments neat and readable.
- Comment all definitions, describing their purposes. Also mark all tables and data storage areas.
- Comment sections of the program as well as individual instructions.
- Be consistent in your terminology. You can (should) be repetitive, you need not consult a thesaurus.
- Leave yourself notes at points that you find confusing: for example, "remember carry was set by last instruction." If such points get cleared up later in program development, you may drop these comments in the final documentation.

A well-commented program is easy to use. You will recover the time spent in commenting many times over. We will try to show good commenting style in the programming examples, although we often over-comment for instructional purposes.

### 2.6 Types of Assemblers

Although all assemblers perform the same tasks, their implementations vary greatly. We will not try to describe all the existing types of assemblers, we will merely define the terms and indicate some of the choices.

A cross-assembler is an assembler that runs on a computer other than the one for which it assembles object programs. The computer on which the cross-assembler runs is typically a large computer with extensive software support and fast peripherals. The computer for which the cross-assembler assembles programs is typically a micro like the 6809 or MC68000.

When a new microcomputer is introduced, a cross-assembler is often provided to run on existing development systems. For example, ARM provide the 'Armulator' cross-assembler that will run on a PC development system.

A self-assembler or resident assembler is an assembler that runs on the computer for which it assembles programs. The self-assembler will require some memory and peripherals, and it may run quite slowly compared to a cross-assembler.

A macroassembler is an assembler that allows you to define sequences of instructions as macros.
A microassembler is an assembler used to write the microprograms which define the instruction set of a computer. Microprogramming has nothing specifically to do with programming microcomputers, but has to do with the internal operation of the computer.

A meta-assembler is an assembler that can handle many different instruction sets. The user must define the particular instruction set being used.

A one-pass assembler is an assembler that goes through the assembly language program only once. Such an assembler must have some way of resolving forward references, for example, Jump instructions which use labels that have not yet been defined.

A two-pass assembler is an assembler that goes through the assembly language source program twice. The first time the assembler simply collects and defines all the symbols; the second time it replaces the references with the actual definitions. A two-pass assembler has no problems with forward references but may be quite slow if no backup storage (like a floppy disk) is available; then the assembler must physically read the program twice from a slow input medium (like a teletypewriter paper tape reader). Most microprocessor-based assemblers require two passes.

### 2.7 Errors

Assemblers normally provide error messages, often consisting of an error code number. Some typical errors are:

| Undefined name | Often a misspelling or an omitted definition |
| :--- | :--- |
| Illegal character | Such as a 2 in a binary number |
| Illegal format | A wrong delimiter or incorrect operands |
| Invalid expression | for example, two operators in a row |
| Illegal value | Usually the value is too large |
| Missing operand | Pretty self explanatory |
| Double definition | Two different values assigned to one name |
| Illegal label | Such as a label on a pseudo-operation that cannot have one |
| Missing label | Probably a miss spelt lable name |
| Undefined operation code |  |

In interpreting assembler errors, you must remember that the assembler may get on the wrong track if it finds a stray letter, an extra space, or incorrect punctuation. The assembler will then proceed to misinterpret the succeeding instructions and produce meaningless error messages. Always look at the first error very carefully; subsequent ones may depend on it. Caution and consistent adherence to standard formats will eliminate many annoying mistakes.

### 2.8 Loaders

The loader is the program which actually takes the output (object code) from the assembler and places it in memory. Loaders range from the very simple to the very complex. We will describe a few different types.

A bootstrap loader is a program that uses its own first few instructions to load the rest of itself or another loader program into memory. The bootstrap loader may be in ROM, or you may have to enter it into the computer memory using front panel switches. The assembler may place a bootstrap loader at the start of the object program that it produces.

A relocating loader can load programs anywhere in memory. It typically loads each program into the memory space immediately following that used by the previous program. The programs, however, must themselves be capable of being moved around in this way; that is, they must be relocatable. An absolute loader, in contrast, will always place the programs in the same area of memory.
A linking loader loads programs and subroutines that have been assembled separately; it resolves cross-references - that is, instructions in one program that refer to a label in another program. Object programs loaded by a linking loader must be created by an assembler that allows external references. An alternative approach is to separate the linking and loading functions and have the linking performed by a program called a link editor and the loading done by a loader.

## 3 ARM Architecture

This chapter outlines the ARM processor's architecture and describes the syntax rules of the ARM assembler. Later chapters of this book describe the ARM's stack and exception processing system in more detail.

Figure 3.1 on the following page shows the internal structure of the ARM processor. The ARM is a Reduced Instruction Set Computer (RISC) system and includes the attributes typical to that type of system:

- A large array of uniform registers.
- A load/store model of data-processing where operations can only operate on registers and not directly on memory. This requires that all data be loaded into registers before an operation can be preformed, the result can then be used for further processing or stored back into memory.
- A small number of addressing modes with all load/store addresses begin determined from registers and instruction fields only.
- A uniform fixed length instruction (32-bit).

In addition to these traditional features of a RISC system the ARM provides a number of additional features:

- Separate Arithmetic Logic Unit (ALU) and shifter giving additional control over data processing to maximize execution speed.
- Auto-increment and Auto-decrement addressing modes to improve the operation of program loops.
- Conditional execution of instructions to reduce pipeline flushing and thus increase execution speed.


### 3.1 Processor modes

The ARM supports the seven processor modes shown in table 3.1.
Mode changes can be made under software control, or can be caused by external interrupts or exception processing.
Most application programs execute in User mode. While the processor is in User mode, the program being executed is unable to access some protected system resources or to change mode, other than by causing an exception to occur (see 3.4 on page 29). This allows a suitably written operating system to control the use of system resources.


Figure 3.1: ARM Block Diagram

| Processor |  | mode |
| :--- | :--- | :--- |
| User | usr | Nescription |
| FIQ | fiq | Fast Interrupt for high-speed data transfer |
| IRQ | irq | Used for general-purpose interrupt handling |
| Supervisor | svc | A protected mode for the operating system |
| Abort | abt | Implements virtual memory and/or memory protection |
| Undefined | und | Supports software emulation of hardware coprocessors |
| System | sys | Runs privileged operating system tasks |

Table 3.1: ARM processor modes

The modes other than User mode are known as privileged modes. They have full access to system resources and can change mode freely. Five of them are known as exception modes: FIQ (Fast Interrupt), IRQ (Interrupt), Supervisor, Abort, and Undefined. These are entered when specific exceptions occur. Each of them has some additional registers to avoid corrupting User mode state when the exception occurs (see 3.2 for details).

The remaining mode is System mode, it is not entered by any exception and has exactly the same registers available as User mode. However, it is a privileged mode and is therefore not subject to the User mode restrictions. It is intended for use by operating system tasks which need access to system resources, but wish to avoid using the additional registers associated with the exception modes. Avoiding such use ensures that the task state is not corrupted by the occurrence of any exception.

### 3.2 Registers

The ARM has a total of 37 registers. These comprise 30 general purpose registers, 6 status registers and a program counter. Figure 3.2 illustrates the registers of the ARM. Only fifteen of the general purpose registers are available at any one time depending on the processor mode.
There are a standard set of eight general purpose registers that are always available (R0-R7) no matter which mode the processor is in. These registers are truly general-purpose, with no special uses being placed on them by the processors' architecture.
A few registers (R8 - R12) are common to all processor modes with the exception of the fiq mode. This means that to all intent and purpose these are general registers and have no special use. However, when the processor is in the fast interrupt mode these registers and replaced with different set of registers (R8_fiq-R12_fiq). Although the processor does not give any special purpose to these registers they can be used to hold information between fast interrupts. You can consider they to be static registers. The idea is that you can make a fast interrupt even faster by holding information in these registers.
The general purpose registers can be used to handle 8 -bit bytes, 16 -bit half-words 1 or 32 -bit words. When we use a 32 -bit register in a byte instruction only the least significant 8 bits are used. In a half-word instruction only the least significant 16 bits are used. Figure 3.3 demonstrates this.

The remaining registers (R13-R15) are special purpose registers and have very specific roles: R13 is also known as the Stack Pointer, while R14 is known as the Link Register, and R15 is the Program Counter. The "user" (usr) and "System" (sys) modes share the same registers. The exception modes all have their own version of these registers. Making a reference to register R14 will assume you are referring to the register for the current processor mode. If you wish to refer

[^0]| Modes |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Privileged Modes |  |  |  |  |  |  |
|  | Exception Modes |  |  |  |  |  |  |
| User | System | Supervisor | Abort | Undefined | Interrupt | Fast Interrupt |  |
| R0 | R0 | R0 | R0 | R0 | R0 | R0 |  |
| R1 | R1 | R1 | R1 | R1 | R1 | R1 |  |
| R2 | R2 | R2 | R2 | R2 | R2 | R2 |  |
| R3 | R3 | R3 | R3 | R3 | R3 | R3 |  |
| R4 | R4 | R4 | R4 | R4 | R4 | R4 |  |
| R5 | R5 | R5 | R5 | R5 | R5 | R5 |  |
| R6 | R6 | R6 | R6 | R6 | R6 | R6 |  |
| R7 | R7 | R7 | R7 | R7 | R7 | R7 |  |
| R8 | R8 | R8 | R8 | R8 | R8 | R8_fiq |  |
| R9 | R9 | R9 | R9 | R9 | R9 | R9_fiq |  |
| R10 | R10 | R10 | R10 | R10 | R10 | R10_fiq |  |
| R11 | R11 | R11 | R11 | R11 | R11 | R11_fiq |  |
| R12 | R12 | R12 | R12 | R12 | R12 | R12_fiq |  |
| R13 | R13 | R13_svc | R13_abt | R13_und | R13_irq | R13_fiq |  |
| R14 | R14 | R14_svc | R14_abt | R14_und | R14_irq | R14_fiq |  |
| PC | PC | PC | PC | PC | PC | PC |  |


| CPSR | CPSR | CPSR | CPSR | CPSR | CPSR | CPSR |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | SPSR_svc | SPSR_abt | SPSR_und | SPSR_irq | SPSR_fiq |  |
|  |  |  |  |  |  |  |

Figure 3.2: Register Organization


Figure 3.3: Byte/Half Word/Word
to the user mode version of this register you have refer to the R14 usr register. You may only refer to register from other modes when the processor is in one of the privileged modes, i.e., any mode other than user mode.

There are also one or two status registers depending on which mode the processor is in. The Current Processor Status Register (CPSR) holds information about the current status of the processor (including its current mode). In the exception modes there is an additional Saved Processor Status Register (SPSR) which holds information on the processors state before the system changed into this mode, i.e., the processor status just before an exception.

### 3.2.1 The stack pointer, SP or R13

Register R13 is used as a stack pointer and is also known as the SP register. Each exception mode has its own version of R13, which points to a stack dedicated to that exception mode.

The stack is typically used to store temporary values. It is normal to store the contents of any registers a function is going to use on the stack on entry to a subroutine. This leaves the register free for use during the function. The routine can then recover the register values from the stack
on exit from the subroutine. In this way the subroutine can preserve the value of the register and not corrupt the value as would otherwise be the case.

See Chapter 15 for more information on using the stack.

### 3.2.2 The Link Register, LR or R14

Register R14 is also known as the Link Register or LR.
It is used to hold the return address for a subroutine. When a subroutine call is performed via a BL instruction, R14 is set to the address of the next instruction. To return from a subroutine you need to copy the Link Register into the Program Counter. This is typically done in one of the two ways:

- Execute either of these instructions:

$$
\begin{array}{llll}
\text { MOV } \quad \text { PC, LR } & \text { or } & \text { BAL }
\end{array}
$$

- On entry to the subroutine store R14 to the stack with an instruction of the form:

STMFD SP!,\{〈registers $\rangle, \mathrm{LR}\}$
and use a matching instruction to return from the subroutine:

$$
\text { LDMFD SP!,\{〈registers }\rangle, \mathrm{PC}\}
$$

This saves the Link Register on the stack at the start of the subroutine. On exit from the subroutine it collects all the values it placed on the stack, including the return address that was in the Link Register, except it returns this address directly into the Program Counter instead.

See Chapter ?? on page ?? for further details of using the stack, and Chapter 15 on page 113 for further details on using subroutines.

When an exception occurs, the exception mode's version of R14 is set to the address after the instruction which has just been completed. The SPSR is a copy of the CPSR just before the exception occurred. The return from an exception is performed in a similar way to a subroutine return, but using slightly different instructions to ensure full restoration of the state of the program that was being executed when the exception occurred. See 3.4 on page 29 for more details.

### 3.2.3 The program counter, PC or R15

Register R15 holds the Program Counter known as the PC. It is used to identify which instruction is to be preformed next. As the PC holds the address of the next instruction it is often referred to as an instruction pointer. The name "program counter" dates back to the times when program instructions where read in off of punched cards, it refers to the card position within a stack of cards. In spite of its name it does not actually count anything!

## Reading the program counter

When an instruction reads the PC the value returned is the address of the current instruction plus 8 bytes. This is the address of the instruction after the next instruction to be executed ${ }^{2}$

[^1]This way of reading the PC is primarily used for quick, position-independent addressing of nearby instructions and data, including position-independent branching within a program.

An exception to this rule occurs when an STR (Store Register) or STM (Store Multiple Registers) instruction stores R15. The value stored is UNKNOWN and it is best to avoid the use of these instructions that store R15.

## Writing the program counter

When an instruction writes to R15 the normal result is that the value written is treated as an instruction address and the system starts to execute the instruction at that address $s^{3}$,

### 3.2.4 Current Processor Status Registers: CPSR

Rather surprisingly the current processor status register (CPSR) contains the current status of the processor. This includes various condition code flags, interrupt status, processor mode and other status and control information.

The exception modes also have a saved processor status register (SPSR), that is used to preserve the value of the CPSR when the associated exception occurs. Because the User and System modes are not exception modes, there is no SPSR available.

Figure 3.4 shows the format of the CPSR and the SPSR registers.

| 31 | 30 | 29 | 28 | 27 | $\cdots$ | 8 | 7 | 6 | 5 | 4 | $\cdots$ | 0 |
| :---: | :---: | :---: | :---: | ---: | :---: | ---: | ---: | ---: | :---: | :---: | :---: | :---: |
| N | Z | C | V | SBZ | I | F | SBZ | Mode |  |  |  |  |

Figure 3.4: Structure of the Processor Status Registers

The processors' status is split into two distinct parts: the User flags and the Systems Control flags. The upper halfword is accessible in User mode and contains a set of flags which can be used to effect the operation of a program, see section 3.3. The lower halfword contains the System Control information.

Any bit not currently used is reserved for future use and should be zero, and are marked SBZ in the figure. The I and F bits indicate if Interrupts (I) or Fast Interrupts (F) are allowed. The Mode bits indicate which operating mode the processor is in (see 3.1 on page 23).

The system flags can only be altered when the processor is in protected mode. User mode programs can not alter the status register except for the condition code flags.

### 3.3 Flags

The upper four bits of the status register contains a set of four flags, collectively known at the condition code. The condition code flags are:

| Negative | (N) |
| :--- | :--- |
| Zero | (Z) |
| Carry | (C) |
| Overflow | (V) |

[^2]The condition code can be used to control the flow of the program execution. The is often abbreviated to just $\langle c c\rangle$.

N The Negative (sign) flag takes on the value of the most significant bit of a result. Thus when an operation produces a negative result the negative flag is set and a positive result results in a the negative flag being reset. This assumes the values are in standard two's complement form. If the values are unsigned the negative flag can be ignored or used to identify the value of the most significant bit of the result.
Z The Zero flag is set when an operation produces a zero result. It is reset when an operation produces a non-zero result.
C The Carry flag holds the carry from the most significant bit produced by arithmetic operations or shifts. As with most processors, the carry flag is inverted after a subtraction so that the flag acts as a borrow flag after a subtraction.
v The Overflow flag is set when an arithmetic result is greater than can be represented in a register.

Many instructions can modify the flags, these include comparison, arithmetic, logical and move instructions. Most of the instructions have an $S$ qualifier which instructs the processor to set the condition code flags or not.

### 3.4 Exceptions

Exceptions are generated by internal and external sources to cause the processor to handle an event, such as an externally generated interrupt or an attempt to execute an undefined instruction. The ARM supports seven types of exception, and a provides a privileged processing mode for each type. Table 3.2 lists the type of exception and the processor mode associated with it.

When an exception occurs, some of the standard registers are replaced with registers specific to the exception mode. All exception modes have their own Stack Pointer (SP) and Link (LR) registers. The fast interrupt mode has more registers (R8_fiq-R12_fiq) for fast interrupt processing.

| Exception Type | Processor | Mode |
| :--- | :--- | :--- |
| Reset | Supervisor | svc |
| Software Interrupt | Supervisor | svc |
| Undefined Instruction | Undefined | und |
| Prefetch Abort | Abort | abt |
| Data Abort | Abort | abt |
| Interrupt | IRQ | irq |
| Fast Interrupt | FIQ | fiq |

Table 3.2: Exception processing modes

The seven exceptions are:

Reset when the Reset pin is held low, this is normally when the system is first turned on or when the reset button is pressed.

Software Interrupt is generally used to allow user mode programs to call the operating system. The user program executes a software interrupt (SWI, A. 18 on page 135 instruction with a argument which identifies the function the user wishes to preform.

Undefined Instruction is when an attempt is made to preform an undefined instruction. This normally happens when there is a logical error in the program and the processor starts to execute data rather than program code.

Prefetch Abort occurs when the processor attempts to access memory that does not exist.
Data Abort occurs when attempting to access a word on a non-word aligned boundary. The lower two bits of a memory must be zero when accessing a word.

Interrupt occurs when an external device asserts the IRQ (interrupt) pin on the processor. This can be used by external devices to request attention from the processor. An interrupt can not be interrupted with the exception of a fast interrupt.

Fast Interrupt occurs when an external device asserts the FIQ (fast interrupt) pin. This is designed to support data transfer and has sufficient private registers to remove the need for register saving in such applications. A fast interrupt can not be interrupted.

When an exception occurs, the processor halts execution after the current instruction. The state of the processor is preserved in the Saved Processor Status Register (SPSR) so that the original program can be resumed when the exception routine has completed. The address of the instruction the processor was just about to execute is placed into the Link Register of the appropriate processor mode. The processor is now ready to begin execution of the exception handler.

The exception handler are located a pre-defined locations known as exception vectors. It is the responsibility of an operating system to provide suitable exception handling.

### 3.5 Instruction Set

Why are a microprocessor's instructions referred to as an instruction set? Because the microprocessor designer selects the instruction complement with great care; it must be easy to execute complex operations as a sequence of simple events, each of which is represented by one instruction from a well-designed instruction set.

Assembler often frighten users who are new to programming. Yet taken in isolation, the operations involved in the execution of a single instruction are usually easy to follow. Furthermore, you need not attempt to understand all the instructions at once. As you study each of the programs in these notes you will learn about the specific instructions involved.

Table 4.1 lists the instruction mnemonics. This provides a survey of the processors capabilities, and will also be useful when you need a certain kind of operation but are either unsure of the specific mnemonics or not yet familiar with what instructions are available.

See Chapter ?? and Appendix ?? for a detailed description of the individual instructions and chapters 7 through to 15 for a discussion on how to use them.

The ARM instruction set can be divided into six broad classes of instruction.

- Data Movement
- Logical and Bit Manipulation
- Arithmetic
- Memory Access
- Flow Control
- System Control / Privileged

Before we look at each of these groups in a little more detail there are a few ideas which belong to all groups worthy of investigation.

| Operation <br> Mnemonic | Meaning | Operation <br> Mnemonic | Meaning |
| :---: | :--- | :---: | :--- |
| ADC | Add with Carry | MVN | Logical NOT |
| ADD | Add | ORR | Logical OR |
| AND | Logical AND | RSB | Reverse Subtract |
| BAL | Unconditional Branch | RSC | Reverse Subtract with Carry |
| B $\langle c c\rangle$ | Branch on Condition | SBC | Subtract with Carry |
| BIC | Bit Clear | SMLAL | Mult Accum Signed Long |
| BLAL | Unconditional Branch and Link | SMULL | Multiply Signed Long |
| BL〈cc | Conditional Branch and Link | STM | Store Multiple |
| CMP | Compare | STR | Store Register (Word) |
| EOR | Exclusive OR | STRB | Store Register (Byte) |
| LDM | Load Multiple | SUB | Subtract |
| LDR | Load Register (Word) | SWI | Software Interrupt |
| LDRB | Load Register (Byte) | SWP | Swap Word Value |
| MLA | Multiply Accumulate | SWPB | Swap Byte Value |
| MOV | Move | TEQ | Test Equivalence |
| MRS | Load SPSR or CPSR | TST | Test |
| MSR | Store to SPSR or CPSR | UMLAL | Mult Accum Unsigned Long |
| MUL | Multiply | UMULL | Multiply Unsigned Long |

Table 3.3: Instruction Mnemonics

| Mnemonic | Condition | Mnemonic | Condition |
| :---: | :--- | :---: | :--- |
| CS | Carry Set | CC | Carry Clear |
| EQ | Equal (Zero Set) | NE | Not Equal (Zero Clear) |
| VS | Overflow Set | VC | Overflow Clear |
| GT | Greater Than | LT | Less Than |
| GE | Greater Than or Equal | LE | Less Than or Equal |
| PL | Plus (Positive) | MI | Minus (Negative) |
| HI | Higher Than | LO | Lower Than (aka CC) |
| HS | Higher or $S$ ame (aka CS) | LS | Lower or Same |

Table 3.4: $\langle c c\rangle$ (Condition code) Mnemonics

### 3.5.1 Conditional Execution: $\langle c c\rangle$

Almost all ARM instructions contain a condition field which allows it to be executed conditionally dependent on the condition code flags 3.3 on page 28). If the flags indicate that the corresponding condition is true when the instruction starts executing, it executes normally. Otherwise, the instruction does nothing.

Table 4.2 on page 42 shows a list of the condition codes and their mnemonics. To indicate that an instruction is conditional we simply place the mnemonic for the condition code after the mnemonic for the instruction. If no condition code mnemonic is used the instruction will always be executed.

For example the following instruction will move the value of the register R 1 into the R 0 register only when the Carry flag has been set, R0 will remain unaffected if the C flag was clear.
MOVCS R0, R1

Note that the Greater and the Less conditions are for use with signed numbers while the Higher and Lower conditions are for use with unsigned numbers. These condition codes only really make seance after a comparison (CMP) instruction, see A.5 on page 129 .

Most data-processing instructions can also update the condition codes according to their result. Placing an " S " after the mnemonic will cause the flags to be updated. For example there are two versions of the MOV instruction:

MOV RO, \#0 Will move the value 0 into the register R0 without setting the flags.
MOVS RO, \#0 Will do the same, move the value 0 into the register RO, but it will also set the condition code flags accordingly, the Zero flag will be set, the Negative flag will be reset and the Carry and oVerflow flags will not be effected.

If an instruction has this ability we denote it using $\langle S\rangle$ in our description of the instruction. The $\langle S\rangle$ always comes after the $\langle c c\rangle$ (conditional execution) modification if it is given. Thus the full description of the move instruction would be:

$$
\operatorname{MOV}\langle c c\rangle\langle S\rangle \quad \mathrm{Rd}, \quad\langle o p 1\rangle
$$

With all this in mind what does the following code fragment do?

| MOVS | R0, R1 |
| :--- | :--- |
| MOVEQS | R0, |
| M2 |  |
| MOVEQ | R0, |

The first instruction will move R1 into R0 unconditionally, but it will also set the N and Z flags accordingly. Thus the second instruction is only executed if the Z flag is set, i.e., the value of R1 was zero. If the value of R1 was not zero the instruction is skipped. If the second instruction is executed it will copy the value of R2 into RO and it will also set the N and Z flags according to the value of R2. Thus the third instruction is only executed if both R1 and R2 are both zero.

### 3.5.2 Data Processing Operands: $\langle o p 1\rangle$

The majority of the instructions relate to data processing of some form. One of the operands to these instructions is routed through the Barrel Shifter. This means that the operand can be modified before it is used. This can be very useful when dealing with lists, tables and other complex data structures. We denote instructions of this type as taking one of its arguments from $\langle o p 1\rangle$.

An $\langle o p 1\rangle$ argument may come from one of two sources, a constant value or a register, and be modified in five different ways. See Chapter ?? for more detailed information.

## Unmodified Value

You can use a value or a register unmodified by simply giving the value or the register name. For example the following instructions will demonstrate the two methods:

$$
\begin{array}{lll}
\text { MOV } & \text { R0, \#1234 } & \text { Will move the immediate constant value } 1234_{10} \text { into the register R0 } \\
\text { MOV } & \text { R0, R1 } & \text { Will move the value in the register R1 into the register R0 }
\end{array}
$$

## Logical Shift Left

This will take the value of a register and shift the value up, towards the most significant bit, by $n$ bits. The number of bits to shift is specified by either a constant value or another register. The lower bits of the value are replaced with a zero. This is a simple way of performing a multiply by a power of $2\left(\times 2^{n}\right)$.

MOV R0, R1, LSL \#2 R0 will become the value of R1 shifted left by 2 bits. The value of R1 is not changed.

MOV R0, R1, LSL R2 R0 will become the value of R1 shifted left by the number of bits specified in the R2 register. R0 is the only register to change, both R1 and R2 are not effected by this operation.

If the instruction is to set the status register, the carry flag $(C)$ is the last bit that was shifted out of the value.

## Logical Shift Right

Logical Shift Right is very similar to Logical Shift Left except it will shift the value to the right, towards the lest significant bit, by $n$ bits. It will replace the upper bits with zeros, thus providing an efficient unsigned divide by $2^{n}$ function $\left(\left|\div 2^{n}\right|\right)$. The number of bits to shift may be specified by either a constant value or another register.

```
MOV R0, R1, LSR #2 R0 will take on the value of R1 shifted to the right by 2 bits. The
    value of R1 is not changed.
MOV R0, R1, LSR R2 As before R0 will become the value of R1 shifted to the right by the
    number of bits specified in the R2 register. R1 and R2 are not altered
    by this operation.
```

If the instruction is to set the status register, the carry flag (C) is the last bit to be shifted out of the value.

## Arithmetic Shift Right

The Arithmetic Shift Right is rather similar to the Logical Shift Right, but rather than replacing the upper bits with a zero, it maintains the value of the most significant bit. As the most significant bit is used to hold the sign, this means the sign of the value is maintained, thus providing a signed divide by $2^{n}$ operation $\left(\div 2^{n}\right)$.

MOV R0, R1, ASR \#2 Register R0 will become the value of register R1 shifted to the right by 2 bits, with the sign maintained.

MOV R0, R1, ASR R2 Register R0 will become the value of the register R1 shifted to the right by the number of bits specified by the R2 register. R1 and R2 are not altered by this operation.

Given the distinction between the Logical and Arithmetic Shift Right, why is there no Arithmetic Shift Left operation?

As a signed number is stored in two's complement the upper most bits hold the sign of the number. These bits can be considered insignificant unless the number is of a sufficient size to require their use. Thus an Arithmetic Shift Left is not required as the sign is automatically preserved by the Logical Shift.

## Rotate Right

In the Rotate Right operation, the lest significant bit is copied into the carry (C) flag, while the value of the C flag is copied into the most significant bit of the value. In this way none of the bits in the value are lost, but are simply moved from the lower bits to the upper bits of the value.

MOV R0, R1, ROR \#2 This will rotate the value of R1 by two bits. The most significant bit of the resulting value will be the same as the least significant bit of the original value. The second most significant bit will be the same as the Carry flag. In the $S$ version the Carry flag will be set to the second least significant bit of the original value. The value of R1 is not changed by this operation.

MOV R0, R1, ROR R2 Register R0 will become the value of the register R1 rotated to the right by the number of bits specified by the R2 register. R1 and R2 are not altered by this operation.

Why is there no corresponding Rotate Left operation?
An Add With Carry (ADC, A. 1 on page 127) to a zero value provides this service for a single bit. The designers of the instruction set believe that a Rotate Left by more than one bit would never be required, thus they have not provided a ROL function.

## Rotate Right Extended

This is similar to a Rotate Right by one bit. The extended section of the fact that this function moves the value of the Carry (C) flag into the most significant bit of the value, and the least significant bit of the value into the Carry (C) flag. Thus it allows the Carry flag to be propagated though multi-word values, thereby allowing values larger than 32 -bits to be used in calculations.

MOV R0, R1 RRX The register R0 become the same as the value of the register R1 rotated though the carry flag by one bit. The most significant bit of the value becomes the same as the current Carry flag, while the Carry flag will be the same as the least significant bit or R1. The value of R1 will not be changed.

### 3.5.3 Memory Access Operands: $\langle o p 2\rangle$

The memory address used in the memory access instructions may also modified by the barrel shifter. This provides for more advanced access to memory which is particularly useful when dealing with more advanced data structures. It allows pre- and post-increment instructions that update memory pointers as a side effect of the instruction. This makes loops which pass though memory more efficient. We denote instructions of this type as taking one of its arguments from $\langle o p 2\rangle$. For a full discussion of the $\langle o p 2\rangle$ addressing mode we refer the reader to Chapter ?? on page ??.
There are three main methods of specifying a memory address ( $\langle o p 2\rangle$ ), all of which include an offset value of some form. This offset can be specified in one of three ways:

## Constant Value

An immediate constant value can be provided. If no offset is specified an immediate constant value of zero is assumed.

## Register

The offset can be specified by another register. The value of the register is added to the address held in another register to form the final address.

## Scaled

The offset is specified by another register which can be scaled by one of the shift operators used for $\langle o p 1\rangle$. More specifically by the Logical Shift Left (LSL), Logical Shift Right (LSR), Arithmetic Shift Right (ASR), ROtate Right (ROR) or Rotate Right Extended (RRX) shift operators, where the number of bits to shift is specified as a constant value.

## Offset Addressing

In offset addressing the memory address is formed by adding (or subtracting) an offset to or from the value held in a base register.

| LDR | R0, [R1] | Will load the register R0 with the 32 -bit word at the memory address held in the register R1. In this instruction there is no offset specified, so an offset of zero is assumed. The value of R1 is not changed in this instruction. |
| :---: | :---: | :---: |
| LDR | R0, [R1, \#4] | Will load the register R0 with the word at the memory address calculated by adding the constant value 4 to the memory address contained in the R1 register. The register R1 is not changed by this instruction. |
| LDR | R0, [R1, R2] | Loads the register R0 with the value at the memory address calculated by adding the value in the register R 1 to the value held in the register R2. Both R1 and R2 are not altered by this operation. |
| LDR | R0, [R1, R2, LSL \#2] | Will load the register R0 with the 32 -bit value at the memory address calculated by adding the value in the R1 register to the value obtained by shifting the value in R2 left by 2 bits. Both registers, R1 and R2 are not effected by this operation. |

This is particularly useful for indexing into a complex data structure. The start of the data structure is held in a base register, R1 in this case, and the offset to access a particular field within the structure is then added to the base address. Placing the offset in a register allows it to be calculated at run time rather than fixed. This allows for looping though a table.

A scaled value can also be used to access a particular item of a table, where the size of the item is a power of two. For example, to locate item 7 in a table of 32 -bit values we need only shift the index value 6 left by 2 bits $\left(6 \times 2^{2}\right)$ to calculate the value we need to add as an offset to the start of the table held in a register, R1 in our example. Remember that the computer count from zero, thus we use an index value of 6 rather than 7 . A 32-bit number requires 4 bytes of storage which is $2^{2}$, thus we only need a 2 -bit left shift.

## Pre-Index Addressing

In pre-index addressing the memory address if formed in the same way as for offset addressing. The address is not only used to access memory, but the base register is also modified to hold the new value. In the ARM system this is known as a write-back and is denoted by placing a exclamation mark after at the end of the $\langle o p 2\rangle$ code.

Pre-Index address can be particularly useful in a loop as it can be used to automatically increment or decrement a counter or memory pointer.

LDR R0, [R1, \#4]!

LDR R0, [R1, R2]!

LDR R0, [R1, R2, LSL \#2]! First calculates the new address by adding the value in the base address register, R1, to the value obtained by shifting the value in the offset register, R2, left by 2 bits. It will then load the 32 -bit at this address into the destination register, RO. The new address is also written back into the base register, R1. The offset register, R2, will not be effected by this operation.

## Post-Index Addressing

In post-index address the memory address is the base register value. As a side-effect, an offset is added to or subtracted from the base register value and the result is written back to the base register.

Post-index addressing uses the value of the base register without modification. It then applies the modification to the address and writes the new address back into the base register. This can be used to automatically increment or decrement a memory pointer after it has been used, so it is pointing to the next location to be used.

As the instruction must preform a write-back we do not need to include an exclamation mark. Rather we move the closing bracket to include only the base register, as that is the register holding the memory address we are going to access.

LDR R0, [R1], \#4

LDR R0, [R1], R2

Will load the register R0 with the word at the memory address contained in the base register, R1. It will then calculate the new value of R1 by adding the constant value 4 to the current value of R1.

Loads the register RO with the value at the memory address held in the base register, R1. It will then calculate the new value for the base register by adding the value in the offset register, R2, to the current value of the base register. The offset register, R2, is not altered by this operation.

LDR R0, [R1], R2, LSL \#2 First loads the 32-bit value at the memory address contained in the base register, R1, into the destination register, R0. It will then calculate the new value for the base register by adding the current value to the value obtained by shifting the value in the offset register, R2, left by 2 bits. The offset register, R2, will not be effected by this operation.

## 4 Instruction Set

Why are a microprocessor's instructions referred to as an instruction set? Because the microprocessor designer selects the instruction complement with great care; it must be easy to execute complex operations as a sequence of simple events, each of which is represented by one instruction from a well-designed instruction set.

Assembler often frighten users who are new to programming. Yet taken in isolation, the operations involved in the execution of a single instruction are usually easy to follow. Furthermore, you need not attempt to understand all the instructions at once. As you study each of the programs in these notes you will learn about the specific instructions involved.

Table 4.1 lists the instruction mnemonics. This provides a survey of the processors capabilities, and will also be useful when you need a certain kind of operation but are either unsure of the specific mnemonics or not yet familiar with what instructions are available.

The appendix Agives a detailed description of the individual instructions while chapters 7 through to 15 provide a discussion on how to use them.
The ARM instruction set can be divided into six broad classes of instruction.

- Data Movement
- Arithmetic
- Memory Access
- Logical and Bit Manipulation
- Flow Control
- System Control / Privileged

Before we look at each of these groups in a little more detail there are a few ideas which belong to all groups worthy of investigation.

## Important Note:

The ARM instruction set can be divided into six broad classes of instruction:

- Data-processing instructions (Data Movement)
- Branch instructions (Flow Control)
- Status register transfer instructions (Logic/Bit Bashing)
- Load and store instructions (Memory Access)
- Coprocessor instructions (System Control)
- Exception-generating instructions (Privileged)

| Operation <br> Mnemonic | Meaning | Operation <br> Mnemonic | Meaning |
| :--- | :--- | :--- | :--- |
| ADC | Add with Carry | MVN | Logical NOT |
| ADD | Add | ORR | Logical OR |
| AND | Logical AND | RSB | Reverse Subtract |
| BAL | Unconditional Branch | RSC | Reverse Subtract with Carry |
| B $\langle c c\rangle$ | Branch on Condition | SBC | Subtract with Carry |
| BIC | Bit Clear | SMLAL | Mult Accum Signed Long |
| BLAL | Unconditional Branch and Link | SMULL | Multiply Signed Long |
| BL $\langle c c\rangle$ | Conditional Branch and Link | STM | Store Multiple |
| CMP | Compare | STR | Store Register (Word) |
| EOR | Exclusive OR | STRB | Store Register (Byte) |
| LDM | Load Multiple | SUB | Subtract |
| LDR | Load Register (Word) | SWI | Software Interrupt |
| LDRB | Load Register (Byte) | SWP | Swap Word Value |
| MLA | Multiply Accumulate | SWPB | Swap Byte Value |
| MOV | Move | TEQ | Test Equivalence |
| MRS | Load SPSR or CPSR | TST | Test |
| MSR | Store to SPSR or CPSR | UMLAL | Mult Accum Unsigned Long |
| MUL | Multiply | UMULL | Multiply Unsigned Long |

Table 4.1: Instruction Mnemonics

### 4.0.4 Branch instructions

As well as allowing many data-processing or load instructions to change control flow by writing the PC, a standard Branch instruction is provided with a 24 -bit signed offset, allowing forward and backward branches of up to 32 MB .
There is a Branch and Link (BL) option that also preserves the address of the instruction after the branch in R14, the LR. This provides a subroutine call which can be returned from by copying the LR into the PC.

### 4.0.5 Data-processing instructions

The data-processing instructions perform calculations on the general-purpose registers. There are four types of data-processing instructions:

- Arithmetic/logic instructions
- Comparison instructions
- Multiply instructions
- Count Leading Zeros instruction


## Arithmetic/logic instructions

There are twelve arithmetic/logic instructions which share a common instruction format. These perform an arithmetic or logical operation on up to two source operands, and write the result to a destination register. They can also optionally update the condition code flags based on the result.

Of the two source operands:

- one is always a register
- the other has two basic forms:
- an immediate value
- a register value, optionally shifted.

If the operand is a shifted register, the shift amount can be either an immediate value or the value of another register. Four types of shift can be specified. Every arithmetic/logic instruction can therefore perform an arithmetic/logic and a shift operation. As a result, ARM does not have dedicated shift instructions.

Because the Program Counter (PC) is a general-purpose register, arithmetic/logic instructions can write their results directly to the PC. This allows easy implementation of a variety of jump instructions.

## Comparison instructions

There are four comparison instructions which use the same instruction format as the arithmetic/logic instructions. These perform an arithmetic or logical operation on two source operands, but do not write the result to a register. They always update the condition flags based on the result.

The source operands of comparison instructions take the same forms as those of arithmetic/logic instructions, including the ability to incorporate a shift operation.

## Multiply instructions

Multiply instructions come in two classes. Both types multiply two 32-bit register values and store their result:

32-bit result Normal. Stores the 32 -bit result in a register.
64-bit result Long. Stores the 64 -bit result in two separate registers.

Both types of multiply instruction can optionally perform an accumulate operation.

## Count Leading Zeros instruction

The Count Leading Zeros (CLZ) instruction determines the number of zero bits at the most significant end of a register value, up to the first 1 bit. This number is written to the destination register of the CLZ instruction.

### 4.0.6 Status register transfer instructions

The status register transfer instructions transfer the contents of the CPSR or an SPSR to or from a general-purpose register. Writing to the CPSR can:

- set the values of the condition code flags
- set the values of the interrupt enable bits
- set the processor mode


### 4.0.7 Load and store instructions

The following load and store instructions are available:

- Load and Store Register
- Load and Store Multiple registers
- Swap register and memory contents


## Load and Store Register

Load Register instructions can load a 32 -bit word, a 16 -bit halfword or an 8 -bit byte from memory into a register. Byte and halfword loads can be automatically zero-extended or sign-extended as they are loaded.

Store Register instructions can store a 32 -bit word, a 16 -bit halfword or an 8 -bit byte from a register to memory.

Load and Store Register instructions have three primary addressing modes, all of which use a base register and an offset specified by the instruction:

- In offset addressing, the memory address is formed by adding or subtracting an offset to or from the base register value.
- In pre-indexed addressing, the memory address is formed in the same way as for offset addressing. As a side-effect, the memory address is also written back to the base register.
- In post-indexed addressing, the memory address is the base register value. As a side-effect, an offset is added to or subtracted from the base register value and the result is written back to the base register.

In each case, the offset can be either an immediate or the value of an index register. Register-based offsets can also be scaled with shift operations.

As the PC is a general-purpose register, a 32 -bit value can be loaded directly into the PC to perform a jump to any address in the 4GB memory space.

## Load and Store Multiple registers

Load Multiple (LDM) and Store Multiple (STM) instructions perform a block transfer of any number of the general-purpose registers to or from memory. Four addressing modes are provided:

- pre-increment
- post-increment
- pre-decrement
- post-decrement

The base address is specified by a register value, which can be optionally updated after the transfer. As the subroutine return address and PC values are in general-purpose registers, very efficient subroutine entry and exit sequences can be constructed with LDM and STM:

- A single STM instruction at subroutine entry can push register contents and the return address onto the stack, updating the stack pointer in the process.
- A single LDM instruction at subroutine exit can restore register contents from the stack, load the PC with the return address, and update the stack pointer.

LDM and STM instructions also allow very efficient code for block copies and similar data movement algorithms.

## Swap register and memory contents

A swap (SWP) instruction performs the following sequence of operations:

1. It loads a value from a register-specified memory location.
2. It stores the contents of a register to the same memory location.
3. It writes the value loaded in step 1 to a register.

By specifying the same register for steps 2 and 3 , the contents of a memory location and a register are interchanged.

The swap operation performs a special indivisible bus operation that allows atomic update of semaphores. Both 32 -bit word and 8 -bit byte semaphores are supported.

### 4.0.8 Coprocessor instructions

There are three types of coprocessor instructions:

Data-processing instructions These start a coprocessor-specific internal operation.
Data transfer instructions These transfer coprocessor data to or from memory. The address of the transfer is calculated by the ARM processor.

Register transfer instructions These allow a coprocessor value to be transferred to or from an ARM register.

### 4.0.9 Exception-generating instructions

Two types of instruction are designed to cause specific exceptions to occur.

Software interrupt instructions SWI instructions cause a software interrupt exception to occur. These are normally used to make calls to an operating system, to request an OS-defined service. The exception entry caused by a SWI instruction also changes to a privileged processor mode. This allows an unprivileged task to gain access to privileged functions, but only in ways permitted by the OS.

Software breakpoint instructions BKPT instructions cause an abort exception to occur. If suitable debugger software is installed on the abort vector, an abort exception generated in this fashion is treated as a breakpoint. If debug hardware is present in the system, it can instead treat a BKPT instruction directly as a breakpoint, preventing the abort exception from occurring.

| Mnemonic | Condition | Mnemonic | Condition |
| :---: | :--- | :---: | :--- |
| CS | Carry Set | CC | Carry Clear |
| EQ | Equal (Zero Set) | NE | Not Equal (Zero Clear) |
| VS | Overflow Set | VC | Overflow Clear |
| GT | Greater Than | LT | Less Than |
| GE | Greater Than or Equal | LE | Less Than or Equal |
| PL | Plus (Positive) | MI | Minus (Negative) |
| HI | Higher Than | LO | Lower Than (aka CC) |
| HS | Higher or Same (aka CS) | LS | Lower or Same |

Table 4.2: $\langle c c\rangle$ (Condition code) Mnemonics

In addition to the above, the following types of instruction cause an Undefined Instruction exception to occur:

- coprocessor instructions which are not recognized by any hardware coprocessor
- most instruction words that have not yet been allocated a meaning as an ARM instruction.

In each case, this exception is normally used either to generate a suitable error or to initiate software emulation of the instruction.

### 4.0.10 Conditional Execution: $\langle c c\rangle$

Almost all ARM instructions contain a condition field which allows it to be executed conditionally dependent on the condition code flags $(3.3$ on page 28). If the flags indicate that the corresponding condition is true when the instruction starts executing, it executes normally. Otherwise, the instruction does nothing.
Table 4.2 shows a list of the condition codes and their mnemonics. To indicate that an instruction is conditional we simply place the mnemonic for the condition code after the mnemonic for the instruction. If no condition code mnemonic is used the instruction will always be executed.

For example the following instruction will move the value of the register R1 into the R0 register only when the Carry flag has been set, RO will remain unaffected if the C flag was clear.

$$
\text { MOVCS } \quad \text { R0, R1 }
$$

Note that the Greater and the Less conditions are for use with signed numbers while the Higher and Lower conditions are for use with unsigned numbers. These condition codes only really make seance after a comparison (CMP) instruction, see A.5 on page 129 .

Most data-processing instructions can also update the condition codes according to their result. Placing an " $S$ " after the mnemonic will cause the flags to be updated. For example there are two versions of the MOV instruction:
MOV RO, \#O Will move the value 0 into the register R0 without setting the flags.
MOVS RO, \#0 Will do the same, move the value 0 into the register R0, but it will also set the condition code flags accordingly, the Zero flag will be set, the Negative flag will be reset and the Carry and oVerflow flags will not be effected.

If an instruction has this ability we denote it using $\langle S\rangle$ in our description of the instruction. The $\langle S\rangle$ always comes after the $\langle c c\rangle$ (conditional execution) modification if it is given. Thus the full description of the move instruction would be:
$\operatorname{MOV}\langle c c\rangle\langle S\rangle \quad \mathrm{Rd},\langle o p 1\rangle$
With all this in mind what does the following code fragment do?

| MOVS | R0, | R1 |
| :--- | :--- | :--- |
| MOVEQS | R0, | R2 |
| MOVEQ | R0, | R3 |

The first instruction will move R1 into R0 unconditionally, but it will also set the N and Z flags accordingly. Thus the second instruction is only executed if the $Z$ flag is set, i.e., the value of R1 was zero. If the value of R1 was not zero the instruction is skipped. If the second instruction is executed it will copy the value of R2 into RO and it will also set the N and Z flags according to the value of R2. Thus the third instruction is only executed if both R1 and R2 are both zero.

## 5 Addressing Modes

### 5.1 Data Processing Operands: $\langle o p 1\rangle$

The majority of the instructions relate to data processing of some form. One of the operands to these instructions is routed through the Barrel Shifter. This means that the operand can be modified before it is used. This can be very useful when dealing with lists, tables and other complex data structures. We denote instructions of this type as taking one of its arguments from $\langle o p 1\rangle$.

An $\langle o p 1\rangle$ argument may come from one of two sources, a constant value or a register, and be modified in five different ways. See Chapter ?? for more detailed information.

### 5.1.1 Unmodified Value

You can use a value or a register unmodified by simply giving the value or the register name. For example the following instructions will demonstrate the two methods:

| MOV R0, \#1234 | Will move the immediate constant value $1234_{10}$ into the register R0 |
| :--- | :--- | :--- |
| MOV R0, R1 | Will move the value in the register R1 into the register R0 |

### 5.1.2 Logical Shift Left



This will take the value of a register and shift the value up, towards the most significant bit, by $n$ bits. The number of bits to shift is specified by either a constant value or another register. The lower bits of the value are replaced with a zero. This is a simple way of performing a multiply by a power of $2\left(\times 2^{n}\right)$.

MOV R0, R1, LSL \#2 R0 will become the value of R1 shifted left by 2 bits. The value of R1 is not changed.

MOV R0, R1, LSL R2 R0 will become the value of R1 shifted left by the number of bits specified in the R2 register. R0 is the only register to change, both R1 and R2 are not effected by this operation.

If the instruction is to set the status register, the carry flag $(\mathrm{C})$ is the last bit that was shifted out of the value.

### 5.1.3 Logical Shift Right



Logical Shift Right is very similar to Logical Shift Left except it will shift the value to the right, towards the lest significant bit, by $n$ bits. It will replace the upper bits with zeros, thus providing an efficient unsigned divide by $2^{n}$ function $\left(\left|\div 2^{n}\right|\right)$. The number of bits to shift may be specified by either a constant value or another register.
MOV R0, R1, LSR \#2 R0 will take on the value of R1 shifted to the right by 2 bits. The value of R1 is not changed.

MOV R0, R1, LSR R2 As before R0 will become the value of R1 shifted to the right by the number of bits specified in the R2 register. R1 and R2 are not altered by this operation.

If the instruction is to set the status register, the carry flag ( $C$ ) is the last bit to be shifted out of the value.

### 5.1.4 Arithmetic Shift Right



The Arithmetic Shift Right is rather similar to the Logical Shift Right, but rather than replacing the upper bits with a zero, it maintains the value of the most significant bit. As the most significant bit is used to hold the sign, this means the sign of the value is maintained, thus providing a signed divide by $2^{n}$ operation $\left(\div 2^{n}\right)$.

MOV R0, R1, ASR \#2 Register R0 will become the value of register R1 shifted to the right by 2 bits, with the sign maintained.

MOV R0, R1, ASR R2 Register R0 will become the value of the register R1 shifted to the right by the number of bits specified by the R2 register. R1 and R2 are not altered by this operation.

Given the distinction between the Logical and Arithmetic Shift Right, why is there no Arithmetic Shift Left operation?

As a signed number is stored in two's complement the upper most bits hold the sign of the number. These bits can be considered insignificant unless the number is of a sufficient size to require their use. Thus an Arithmetic Shift Left is not required as the sign is automatically preserved by the Logical Shift.

### 5.1.5 Rotate Right



In the Rotate Right operation, the lest significant bit is copied into the carry (C) flag, while the value of the C flag is copied into the most significant bit of the value. In this way none of the bits in the value are lost, but are simply moved from the lower bits to the upper bits of the value.

MOV R0, R1, ROR \#2 This will rotate the value of R1 by two bits. The most significant bit of the resulting value will be the same as the least significant bit of the original value. The second most significant bit will be the same as the Carry flag. In the $S$ version the Carry flag will be set to the second least significant bit of the original value. The value of R1 is not changed by this operation.

MOV R0, R1, ROR R2 Register R0 will become the value of the register R1 rotated to the right by the number of bits specified by the R2 register. R1 and R2 are not altered by this operation.

Why is there no corresponding Rotate Left operation?
An Add With Carry (ADC, A.1 on page 127) to a zero value provides this service for a single bit. The designers of the instruction set believe that a Rotate Left by more than one bit would never be required, thus they have not provided a ROL function.

### 5.1.6 Rotate Right Extended



This is similar to a Rotate Right by one bit. The extended section of the fact that this function moves the value of the Carry (C) flag into the most significant bit of the value, and the least significant bit of the value into the Carry (C) flag. Thus it allows the Carry flag to be propagated though multi-word values, thereby allowing values larger than 32-bits to be used in calculations.

MOV R0, R1 RRX The register R0 become the same as the value of the register R1 rotated though the carry flag by one bit. The most significant bit of the value becomes the same as the current Carry flag, while the Carry flag will be the same as the least significant bit or R1. The value of R1 will not be changed.

### 5.2 Memory Access Operands: 〈op2〉

The memory address used in the memory access instructions may also modified by the barrel shifter. This provides for more advanced access to memory which is particularly useful when dealing with more advanced data structures. It allows pre- and post-increment instructions that update memory pointers as a side effect of the instruction. This makes loops which pass though memory more efficient. We denote instructions of this type as taking one of its arguments from $\langle o p 2\rangle$. For a full discussion of the $\langle o p 2\rangle$ addressing mode we refer the reader to Chapter ?? on page ??.

There are three main methods of specifying a memory address ( $\langle o p 2\rangle$ ), all of which include an offset value of some form. This offset can be specified in one of three ways:

## Constant Value

An immediate constant value can be provided. If no offset is specified an immediate constant value of zero is assumed.

## Register

The offset can be specified by another register. The value of the register is added to the address held in another register to form the final address.

## Scaled

The offset is specified by another register which can be scaled by one of the shift operators used for $\langle o p 1\rangle$. More specifically by the Logical Shift Left (LSL), Logical Shift Right (LSR), Arithmetic Shift Right (ASR), ROtate Right (ROR) or Rotate Right Extended (RRX) shift operators, where the number of bits to shift is specified as a constant value.

### 5.2.1 Offset Addressing



In offset addressing the memory address is formed by adding (or subtracting) an offset to or from the value held in a base register.

| LDR | R0, [R1] | Will load the register R0 with the 32 -bit word at the memory address held in the register R1. In this instruction there is no offset specified, so an offset of zero is assumed. The value of R1 is not changed in this instruction. |
| :---: | :---: | :---: |
| LDR | R0, [R1, \#4] | Will load the register R0 with the word at the memory address calculated by adding the constant value 4 to the memory address contained in the R1 register. The register R1 is not changed by this instruction. |
| LDR | R0, [R1, R2] | Loads the register R0 with the value at the memory address calculated by adding the value in the register R 1 to the value held in the register R2. Both R1 and R2 are not altered by this operation. |
| LDR | R0, [R1, R2, LSL \#2] | Will load the register R0 with the 32 -bit value at the memory address calculated by adding the value in the R1 register to the value obtained by shifting the value in R2 left by 2 bits. Both registers, R1 and R2 are not effected by this operation. |

This is particularly useful for indexing into a complex data structure. The start of the data structure is held in a base register, R1 in this case, and the offset to access a particular field within the structure is then added to the base address. Placing the offset in a register allows it to be calculated at run time rather than fixed. This allows for looping though a table.

A scaled value can also be used to access a particular item of a table, where the size of the item is a power of two. For example, to locate item 7 in a table of 32 -bit values we need only shift the index value 6 left by 2 bits $\left(6 \times 2^{2}\right)$ to calculate the value we need to add as an offset to the start of the table held in a register, R1 in our example. Remember that the computer count from zero, thus we use an index value of 6 rather than 7 . A 32-bit number requires 4 bytes of storage which is $2^{2}$, thus we only need a 2 -bit left shift.

### 5.2.2 Pre-Index Addressing



In pre-index addressing the memory address if formed in the same way as for offset addressing. The address is not only used to access memory, but the base register is also modified to hold the new value. In the ARM system this is known as a write-back and is denoted by placing a exclamation mark after at the end of the $\langle o p 2\rangle$ code.

Pre-Index address can be particularly useful in a loop as it can be used to automatically increment or decrement a counter or memory pointer.

| LDR | R0, [R1, \#4]! | Will load the register R0 with the word at the memory address calculated by adding the constant value 4 to the memory address contained in the R1 register. The new memory address is placed back into the base register, register R1. |
| :---: | :---: | :---: |
| LDR | R0, [R1, R2]! | Loads the register R0 with the value at the memory address calculated by adding the value in the register R1 to the value held in the register R2. The offset register, R2, is not altered by this operation, the register holding the base address, R1, is modified to hold the new address. |
| LDR | R0, [R1, R2, LSL \#2]! | First calculates the new address by adding the value in the base address register, R1, to the value obtained by shifting the value in the offset register, R2, left by 2 bits. It will then load the 32 -bit at this address into the destination register, R0. The new address is also written back into the base register, R1. The offset register, R2, will not be effected by this operation. |

### 5.2.3 Post-Index Addressing



In post-index address the memory address is the base register value. As a side-effect, an offset is added to or subtracted from the base register value and the result is written back to the base register.

Post-index addressing uses the value of the base register without modification. It then applies the modification to the address and writes the new address back into the base register. This can be used to automatically increment or decrement a memory pointer after it has been used, so it is pointing to the next location to be used.

As the instruction must preform a write-back we do not need to include an exclamation mark. Rather we move the closing bracket to include only the base register, as that is the register holding the memory address we are going to access.

LDR R0, [R1], \#4 Will load the register R0 with the word at the memory address contained in the base register, R1. It will then calculate the new value of R1 by adding the constant value 4 to the current value of R1.

LDR R0, [R1], R2
Loads the register R0 with the value at the memory address held in the base register, R1. It will then calculate the new value for the base register by adding the value in the offset register, R2, to the current value of the base register. The offset register, R2, is not altered by this operation.

LDR R0, [R1], R2, LSL \#2 First loads the 32-bit value at the memory address contained in the base register, R1, into the destination register, R0. It will then calculate the new value for the base register by adding the current value to the value obtained by shifting the value in the offset register, R2, left by 2 bits. The offset register, R2, will not be effected by this operation.

## 6 Programs

The only way to learn assembly language programming is through experience. Throughout the rest of this book each chapter will introduce various aspects of assembly programming. The chapter will start with a general discussion, then move on to a number of example programs which will demonstrate the topic under discussion. The chapter will end with a number of programming problems for you to try.

### 6.1 Example Programs

Each of the program examples contains several parts:

| Title | that describes the general problem |
| :--- | :--- |
| Purpose | statement of purpose that describes the task the program performs <br> and the memory locations used. |
| Problem A sample problem complete with data and results. <br> Algorithm if the program logic is complex. |  |
| Source code for the assembly program. <br> Notes Explanatry notes that discusses the instructions and methods used <br>  in the program. |  |

Each example is written and assembled as a stand-alone program. They can be downloaded from the web sitf ${ }^{1}$

### 6.1.1 Program Listing Format

The examples in the book are the actual source code used to generate the programs. Sometimes you may need to use the listing output of the ARM assembler (the . list file), and in any case you should be aware of the fact that you can generate a listing file. See the section on the ARMulator environment which follows for details of how to generate a .list listing file.

### 6.1.2 Guidelines for Examples

We have used the following guidelines in construction of the examples:

1. Standard ARM assembler notation is used, as summarized in Chapter 2.
2. The forms in which data and addresses appear are selected for clarity rather than for consistency. We use hexadecimal numbers for memory addresses, instruction codes, and BCD data; decimal for numeric constants; binary for logical masks; and ASCII for characters.

[^3]3. Frequently used instructions and programming techniques are emphasized.
4. Examples illustrate tasks that microprocessors perform in communication, instrumentation, computers, business equipment, industrial, and military applications.
5. Detailed comments are included.
6. Simple and clear structures are emphasised, but programs are written as efficiently as possible within this guideline. Notes accompanying programs often describe more efficient procedures.
7. Program are written as an independent procedures or subroutines although no assumptions are made concerning the state of the microprocessor on entry to the procedure.
8. Program end with a SWI \&11 (Software Interrupt) instruction. You may prefer to modify this by replacing the SWI \&11 instruction with an endless loop instruction such as:

## HERE BAL HERE

9. Programs use standard ARM assembler directives. We introduced assembler directives conceptually in Chapter 2. When first examining programming examples, you can ignore the assembler directives if you do not understand them. Assembler directives do not contribute to program logic, which is what you will be trying to understand initially; but they are a necessary part of every assembly language program, so you will have to learn how to use them before you write any executable programs. Including assembler directives in all program examples will help you become familiar with the functions they perform.

### 6.2 Trying the examples

To test one of the example programs, first obtain a copy of the source code. The best way of doing this is to type in the source code presented in this book, as this will help you to understand the code. Alternatively you can download the source from the web site, although you won't gain the same knowledge of the code.

Go to the start menu and call up the "Armulate" program. Next open the source file using the normal "File | Open" menu option. This will open your program source in a separate window within the "Armulate" environment.

The next step is to create a new Project within the environment. Select the "Project" menu option, then "New". Give your project the same name as the source file that you are using (there is no need to use a file extension - it will automatically be saved as a .apj file).

Once you have given the file a name, a further dialog will open as shown in the figure 6.1 on the next page.
Click the "Add" button, and you will again be presented with a file dialog, which will display the source files in the current directory. Select the relevant source file and "OK" the dialog. You will be returned to the previous dialog, but you will see now that your source file is included in the project. "OK" the "Edit Project" dialog, and you will be returned to the Armulate environment, now with two windows open within it, one for the source code and one for the project.
We recommend that you always create a .list listing file for each project that you create. Do this by selecting the "Options" menu with the project window in focus, then the "Assembler" item. This will open the dialog shown in figure 6.2 on the facing page.
Enter -list [yourfilename].list into the "Other" text box and "OK" the dialog.
You have now created your project and are ready to assemble and debug your code.
Additional information on the Armulator is available via the help menu item.


Figure 6.1: New Project Dialog

| Assembler Options | $\times$ |
| :---: | :---: |
| $\Gamma$ Disable source caching (-nocache) | OK |
| $\Gamma$ Ignore 'C' style escape characters (-noesc) | Cancel |
| $\Gamma$ Nowarnings (-nowarn) |  |
| Other |  |
| -list 16datatrans. list |  |

Figure 6.2: Assembler Options Dialog

### 6.3 Trying the examples from the command line

When developing the example programs, we found the "Armulate" environment too clumsy. We used the TextPad editor and assembled the programs from the command line. The Armulate environment provides commands for use from the command line:

## 1. Assembler

The command line assembler is used to create an object file from the program source code. During the development of the add program (program 7.3a) we used the command line:

```
ARMASM -LI -CPU ARM6 -g -list add.list add.s
```


## 2. Linker

It is necessary to position the program at a fixed location in memory. This is done using the linker. In our add example we used the command:

```
ARMLINK -o add add.o
```

Which resolves the relative addresses in the add.o file, producing the add load image.

## 3. Debugger

Finally it is necessary to debug the load image. This can be done in one of two ways, using a command line debugger or the windows debugger. In either case they require a load image (add in our example). To use the command line debugger (known as the source debugger) the following command is used:

## ARMSD add

However, the command driven nature of this system is confusing and hard to use for even the most experienced of developers. Thus we suggest you use the windows based debugger program:

## WINDBG add

Which will provide you with the same debugger you would have seen had you used the Window based Armulate environment.

### 6.3.1 Setting up TextPad

To set up this environment simply download the TextPad editor and the ARM Assembler syntax file. You can download the editor from the download page of the TextPad web sitf ${ }^{2}$,

Download Derek Law's ARM Assembler Syntax Definition file from the TextPad web site. You can find this under the Syntax Definition sub-section of the Add-ons section of the Download page. Unpack the armasm.syn from the arm.zip file into the TextPad Samples directory.
Having installed the Syntax Definitions you should now add a new Document Class to TextPad. Run TextPad and select the New Document Class... wizard from the Configure menu. The wizard will now take you though the following steps:

1. The Document Class requires a name. We have used the name "ARM Assembler".
2. The Class Members, the file name extension to associate with this document class. We associate all .s and .list files with this class: "*.s,*.list"
3. Syntax Highlighting. The next dialog is where we tell TextPad to use syntax highlighting, simply check the Enable Syntax Highlighting box. We now need to tell it which syntax definition file to use. If the armasm.syn file was placed in the Samples directory, this will appear in the drop down list, and should be selected.

While this will create the new document class, you will almost certainly want to change the colour settings for this document class. This class uses the different levels of Keyword colouring for different aspects of the syntax as follows:

$$
\begin{array}{ll}
\text { Keywords 1 } & \text { Instructions } \\
\text { Keywords 2 } & \text { Co-processor and pseudo-instructions } \\
\text { Keywords 3 } & \text { Shift-addresses and logical directives } \\
\text { Keywords 4 } & \text { Registers } \\
\text { Keywords 5 } & \text { Directives } \\
\text { Keywords 6 } & \text { Arguments and built-in names }
\end{array}
$$

You will probably want to set the color (sic) setting for all of these types to the same settings. We have set all but Keywords 2 to the same colour scheme. To alter the color setting you should select the Preferences. . . option from the Configure menu.

In the "Preference" dialog (shown in figure 6.4 on the next page), open the Document Classes section and then your new document class (ARM Assembler). Now you should select the colors section. This will now allow you to change the colours for any of the given color settings.

Finally you may like to consider adding a "File Type Filter" to the "Open File" dialog. This can be done by selecting the File Type Filter entry in the Preference dialog. Simply click on the New button, add the description ("ARM Assembler (*.s, *.list)") and wildcard ("*.s;*.list") details. Finally click on the OK button.
Note the use of a comma to seperate the wildcards in the description, and the use of a semi-colon (without spaces) in the wildcard entry.

[^4]

Figure 6.3: TextPad Colour Preferences Dialog


Figure 6.4: TextPad File Name Filters Preferences Dialog

### 6.4 Program Initialization

All of the programming examples presented in these notes pay particular attention to the correct initialization of constants and operands. Often this requires additional instructions that may appear superfluous, in that they do not contribute directly to the solution of the stated problem. Nevertheless, correct initialization is important in order to ensure the proper execution of the program every time.

We want to stress correct initialization; that is why we are going to emphasize this aspect of problems.

### 6.5 Special Conditions

For the same reasons that we pay particular attention to special conditions that can cause a program to fail. Empty lists and zero indexes are two of the most common circumstances overlooked in sample problems. It is critically important when using microprocessors that you learn with your very first program to anticipate unusual circumstances; they frequently cause your program to fail. You must build in the necessary programming steps to account for these potential problems.

### 6.6 Problems

Each chapter will now end with a number of programming problems for your to try. They have been provided to help you understand the ideas presented in the chapter. You should use the
programming examples as guidelines for solving the problems. Don't forget to run your solutions on the ARMulator to ensure that they are correct.

The following guidelines will help in solving the problems:

1. Comment each program so that others can understand it. The comments can be brief and ungrammatical. They should explain the purpose of a section or instruction in the program, but should not describe the operation of instructions, that description is available in manuals. For example the following line:

ADD R1, R1, \#1
could be given the comment "Add one to R1" or "Increment R1", both of which provide no indication as to why the line is there. They tell us what the instruction is doing, but we can tell that by looking at the instruction itself. We are more interested in why the instruction is there. A comment such as "Increment loop counter" is much more useful as it explains why you are adding one to R1, the loop counter.
You do not have to comment each statement or explain the obvious. You may follow the format of the examples but provide less detail.
2. Emphasise clarity, simplicity, and good structure in programs. While programs should be reasonably efficient, do not worry about saving a single byte of program memory or a few microseconds.
3. Make programs reasonably general. Do not confuse parameters (such as the number of elements in any array) with fixed constants (such as the code for the letter "C").
4. Never assume fixed initial values for parameters.
5. Use assembler notation as shown in the examples and defined in Chapter 2.
6. Use symbolic notation for address and data references. Symbolic notation should also be used even for constants (such as DATA_SELECT instead of 2_00000100). Also use the clearest possible form for data (such as ' C ' instead of $0 \times 43$ ).
7. Use meaningful names for labels and variables, e.g., SUM or CHECK rather than X or $Z$.
8. Execute each program with the emulator. There is no other way of ensuring that your program is correct. We have provided sample data with each problem. Be sure that the program works for special cases.

## 7 Data Movement

This chapter contains some very elementary programs. They will introduce some fundamental features of the ARM. In addition, these programs demonstrate some primitive tasks that are common to assembly language programs for many different applications.

### 7.1 Program Examples

### 7.1.1 16-Bit Data Transfer

Move the contents of one 16 -bit variable Value to another 16 -bit variable Result.
Sample Problems

| Input: | Value $=\mathrm{C} 123$ |
| :--- | :--- | :--- | :--- |
| Output: | Result $=\mathrm{C} 123$ |

```
Program 7.1: move16.s - 16bit data transfer
    ; 16-Bit data transfer
            TTL Ch4ex1 - move16
            AREA Program, CODE, READONLY
            ENTRY
    Main
            LDRB R1, Value ; Load the value to be moved
            STR R1, Result ; Store it back in a different location
            SWI \&11
    Value DCW \&C123 ; Value to be moved
            ALIGN ; Need to do this because working with 16 bit value
    Result DCW 0 ; Storage space
            END
```

This program solves the problem in two simple steps. The first instruction loads data register R1 with the 16 -bit value in location Value. The next instruction saves the 16 -bit contents of data register R1 in location Result.
As a reminder of the necessary elements of an assembler program for the ARMulator, notice that this, and all the other example programs have the following elements. Firstly there must be an ENTRY directive. This tells the assembler where the first executable instruction is located. Next there must be at least one AREA directive, at the start of the program, and there may be other AREA directives to define data storage areas. Finally there must be an END directive, to show where the code ends. The absence of any of these will cause the assembly to fail with an error.

Another limitation to bear in mind is that ARMulator instructions will only deal with BYTE (8 bits) or WORD ( 32 bit ) data sizes. It is possible to declare HALF-WORD ( 16 bit ) variables by the use
of the DCW directive, but it is necessary to ensure consistency of storage of HALF-WORD by the use of the ALIGN directive. You can see the use of this in the first worked example.

In addition, under the RISC architecture of the ARM, it is not possible to directly manipulate data in storage. Even if no actual manipulation of the data is taking place, as in this first example, it is necessary to use the LDR or LDRB and STR or STRB to move data to a different area of memory.

This version of the LDR instruction moves the 32 -bit word contained in memory location Value into a register and then stores it using the STR instruction at the memory location specified by Result.

Notice that, by default, every program is allocated a literal pool (a storage area) after the last executable line. In the case of this, and most of the other programs, we have formalised this by the use of the AREA Data1, DATA directive. Instruction on how to find addresses of variables will be given in the seminars.

### 7.1.2 One's Complement

From the bitwise complement of the contents of the 16 -bit variable Value.
Sample Problems

| Input: | Value $=$ | C123 |
| :--- | :--- | :--- | :--- |
| Output: | Result $=$ | FFFF3EDC |

```
Program 7.2: invert.s - Find the one's compliment (inverse) of a number
    ; Find the one's compliment (inverse) of a number
        TTL Ch4Ex2 - invert
        AREA Program, CODE, READONLY
        ENTRY
    Main
        LDR R1, Value ; Load the number to be complimented
            MVN R1, R1 ; NOT the contents of R1
            STR R1, Result ; Store the result
            SWI &11
    Value DCD &C123 ; Value to be complemented
    Result DCD 0 ; Storage for result
        END
```

This program solves the problem in three steps. The first instruction moves the contents of location Value into data register R1. The next instruction MVN takes the logical complement of data register R1. Finally, in the third instruction the result of the logical complement is stored in Value.

Note that any data register may be referenced in any instruction that uses data registers, but note the use of R15 for the program counter, R14 for the link register and R13 for the stack pointer. Thus, in the LDR instruction we've just illustrated, any of the general purpose registers could have been used.

The LDR and STR instructions in this program, like those in Program 7.1 demonstrate one of the ARM's addressing modes. The data reference to Value as a source operand is an example of immediate addressing. In immediate addressing the offset to the address of the data being referenced (less 8 byes) is contained in the extension word(s) following the operation word of the instruction. As shown in the assembly listing, the offset to the address corresponding to Value is found in the extension word for the LDR and STR instructions.

### 7.1.3 32-Bit Addition

Add the contents of the 32 -bit variable Value1 to the contents of the 32-bit variable Value2 and place the result in the 32 -bit variable Result.

Sample Problems

Input: $\quad$| Value1 |
| :--- |
|  |
|  |
| Value2 |$=37 \mathrm{E} 3 \mathrm{C} 123$

Output: Result $=367402 \mathrm{AA}$

## Program 7.3a: add.s - Add two numbers

; Add two (32-Bit) numbers

AREA Program, CODE, READONLY
ENTRY
Main

| LDR | R1, Value1 | ; Load the first number |
| :--- | :--- | :--- |
| LDR | R2, Value2 | ; Load the second number |
| ADD | R1, R1, R2 | ; ADD them together into $R 1(x=x+y)$ |
| STR | R1, Result | ; Store the result |
| SWI | $\& 11$ |  |
|  |  | ; First value to be added |
| DCD | $\& 37 E 3 C 123$ | ; Second value to be added |
| DCD | $\& 367402 A A$ | Storage for result |
| DCD | 0 |  |
|  |  |  |

The ADD instruction in this program is an example of a three-operand instruction. Unlike the LDR instruction, this instruction's third operand not only represents the instruction's destination but may also be used to calculate the result. The format:

## DESTINATION $\leftarrow$ SOURCE1 operation SOURCE2

is common to many of the instructions.
As with any microprocessor, there are many instruction sequences you can execute which will solve the same problem. Program 7.3b for example, is a modification of Program 7.3a and uses offset addressing instead of immediate addressing.

```
Program 7.3b: add2.s - Add two numbers and store the result
    ; Add two numbers and store the result
            TTL Ch4Ex4 - add2
            AREA Program, CODE, READONLY
            ENTRY
    Main
            LDR RO, =Value1 ; Load the address of first value
            LDR R1, [RO] ; Load what is at that address
            ADD RO, RO, #Ox4 ; Adjust the pointer
            LDR R2, [RO] ; Load what is at the new addr
            ADD R1, R1, R2 ; ADD together
            LDR RO, =Result ; Load the storage address
            STR R1, [RO] ; Store the result
            SWI &11 ; All done
```

```
Value1 DCD &37E3C123 ; First value
Value2 DCD &367402AA ; Second value
Result DCD 0 ; Space to store result
END
```

The ADR pseudo-instruction introduces a new addressing mode - offest addressing, which we have not used previously. Immediate addressing lets you define a data constant and include that constant in the instruction's associated object code. The assembler format identifies immediate addressing with a \# preceding the data constant. The size of the data constant varies depending on the instruction. Immediate addressing is extremely useful when small data constants must be referenced.

The ADR pseudo-instruction could be replaced by the use of the instruction LDR together with the use of the $=$ to indicate that the address of the data should be loaded rather than the data itself.

The second addressing mode - offset addressing - uses immediate addressing to load a pointer to a memory address into one of the general purpose registers.

Program 7.3b also demonstrates the use of base register plus offset addressing. In this example we have performed this operation manually on line 10 (ADD R0, R0, \#0x4), which increments the address stored in RO by 4 bytes or one WORD. There are much simpler and more efficient ways of doing this, such as pre-index or post-index addressing which we will see in later examples.

Another advantage of this addressing mode is its faster execution time as compared to immediate addressing. This improvement occurs because the address extension word(s) does not have to be fetched from memory prior to the actual data reference, after the initial fetch.

A final advantage is the flexibility provided by having R0 hold an address instead of being fixed as part of the instruction. This flexibility allows the same code to be used for more than one address. Thus if you wanted to add the values contained in consecutive variables Value3 and Value4, you could simply change the contents of RO.

### 7.1.4 Shift Left One Bit

Shift the contents of the 16 -bit variable Value to the left one bit. Store the result back in Result.

Sample Problems

| Input: Value | $=4242$ | $\left(0100001001000010_{2}\right)$ |
| :--- | :--- | :--- | :--- |
| Output: | Result $=8484$ | $\left(1000010010000100_{2}\right)$ |

## Program 7.4: shiftleft.s - Shift Left one bit

```
    ; Shift Left one bit
```

        TTL Ch4Ex5 - shiftleft
        AREA Program, CODE, READONLY
        ENTRY
    Main
        LDR R1, Value ; Load the value to be shifted
        MOV R1, R1, LSL \#0x1 ; SHIFT LEFT one bit
        STR R1, Result ; Store the result
        SWI \&11
    Value DCD \& \(\quad\); Value to be shifted
    Result DCD \(0 \quad\); Space to store result
        END
    The MOV instruction is used to perform a logical shift left. Using the operand format of the MOV instruction shown in Program 7.4 a data register can be shifted from 1 to 25 bits on either a byte, word or longword basis. Another form of the LSL operation allows a shift counter to be specified in another data register.

### 7.1.5 Byte Disassembly

Divide the least significant byte of the 8 -bit variable Value into two 4 -bit nibbles and store one nibble in each byte of the 16 -bit variable Result. The low-order four bits of the byte will be stored in the low-order four bits of the least significant byte of Result. The high-order four bits of the byte will be stored in the low-order four bits of the most significant byte of Result.

## Sample Problems

| Input: | Value $=5 \mathrm{~F}$ |
| :--- | :--- |
| Output: | Result $=050 \mathrm{~F}$ |

Program 7.5: nibble.s - Disassemble a byte into its high and low order nibbles
; Disassemble a byte into its high and low order nibbles
TTL Ch4Ex6 - nibble
AREA Program, CODE, READONLY
ENTRY

## Main

LDR R1, Value ; Load the value to be disassembled
LDR R2, Mask ; Load the bitmask
MOV R3, R1, LSR \#0x4 ; Copy just the high order nibble into R3
MOV R3, R3, LSL \#0x8 ; Now left shift it one byte
AND R1, R1, R2 ; AND the original number with the bitmask
ADD R1, R1, R3 ; Add the result of that to
; What we moved into R3
STR R1, Result ; Store the result
SWI \&11
Value DCB ; Value to be shifted
ALIGN ; Keep the memory boundaries
Mask DCW ; Bitmask $=\% 0000000000001111$
Result DCD 0 ; Space to store result
END

This is an example of byte manipulation. The ARM allows most instructions which operate on words also to operate on bytes. Thus, by using the B suffix, all the LDRinstructions in Program 7.5 become LDRB instructions, therefore performing byte operations. The STR instruction must remain, since we are storing a halfword value. If we were only dealing with a one byte result, we could use the STRB byte version of the store instruction.

Remember that the MOV instruction performs register-to-register transfers. This use of the MOV instruction is quite frequent.

Generally, it is more efficient in terms of program memory usage and execution time to minimise references to memory.

### 7.1.6 Find Larger of Two Numbers

Find the larger of two 32 -bit variables Value1 and Value2. Place the result in the variable Result. Assume the values are unsigned.

Sample Problems


```
Program 7.6: bigger.s - Find the larger of two numbers
    ; Find the larger of two numbers
            TTL Ch4Ex7 - bigger
            AREA Program, CODE, READONLY
            ENTRY
    Main
            LDR R1, Value1 ; Load the first value to be compared
            LDR R2, Value2 ; Load the second value to be compared
            CMP R1, R2 ; Compare them
            BHI Done ; If R1 contains the highest
            MOV R1, R2 ; otherwise overwrite R1
    Done
            STR R1, Result ; Store the result
            SWI &11
    Value1 DCD &12345678 ; Value to be compared
    Value2 DCD &87654321 ; Value to be compared
    Result DCD 0 ; Space to store result
            END
```

The Compare instruction, CMP, sets the status register flags as if the destination, R1, were subtracted from the source R2. The order of the operands is the same as the operands in the subtract instruction, SUB.

The conditional transfer instruction BHI transfers control to the statement labeled Done if the unsigned contents of R2 are greater than or equal to the contents of R1. Otherwise, the next instruction (on line 12) is executed. At Done, register R2 will always contain the larger of the two values.

The BHI instruction is one of several conditional branch instructions. To change the program to operate on signed numbers, simply change the BHI to BGE (Branch if Greater than or Equal to):

```
CMP R1, R2
BGE Done
```

You can use the following table 7.1 to use when performing signed and unsigned comparisons.
Note that the same instructions are used for signal and unsigned addition, subtraction, or comparison; however, the comparison operations are different.

The conditional branch instructions are an example of program counter relative addressing. In other words, if the branch condition is satisfied, control will be transfered to an address relative

| Compare Condition | Signed | Unsigned |
| :--- | :---: | :---: |
| greater than or equal | BGE | BHS |
| greater than | BGT | BHI |
| equal | BEQ | BEQ |
| not equal | BNE | BNE |
| less than or equal | BLE | BLS |
| less than | BLT | BLO |

Table 7.1: Signed/Unsigned Comparisons
to the current value of the program counter. Dealing with compares and branches is an important part of programming. Don't confuse the sense of the CMP instruction. After a compare, the relation tested is:

## DESTINATION condition SOURCE

For exampe, if the condition is "less than," then you test for destination less than source. Become familiar with all of the conditions and their meanings. Unsigned compares are very useful when comparing two addresses.

### 7.1.7 64-Bit Adition

Add the contents of two 64-bit variables Value1 and Value2. Store the result in Result.

## Sample Problems

| Input: | Value1 $=12 A 2 E 640$, | F2100123 |  |
| :--- | :--- | :--- | :--- |
|  |  | Value2 | $=001019 B F$, |
|  | 40023F51 |  |  |
| Output: | Result | $=12 B 30000$, | 32124074 |

Program 7.7: add64.s - 64 bit addition
; 64 bit addition
TTL Ch4Ex8 - add64
area Program, CODE, READONLY
ENTRY
Main
LDR RO, =Value1 ; Pointer to first value
LDR R1, [RO] ; Load first part of value1
LDR R2, [RO, \#4] ; Load lower part of value1
LDR RO, =Value2 ; Pointer to second value
LDR R3, [RO] ; Load upper part of value2
LDR R4, [RO, \#4] ; Load lower part of value2
ADDS R6, R2, R4 ; Add lower 4 bytes and set carry flag
ADC R5, R1, R3 ; Add upper 4 bytes including carry
LDR RO, =Result ; Pointer to Result
STR R5, [RO] ; Store upper part of result
STR R6, [R0, \#4] ; Store lower part of result
SWI \&11
Value1 DCD \&12A2E640, \&F2100123 ; Value to be added
Value2 DCD \&001019BF, \&40023F51 ; Value to be added
Result DCD 0 ; Space to store result
END

Here we introduce several important and powerful instructions from the ARM instruction set. As before, at line 8 we use the LDR instruction which causes register RO to hold the starting address of Value1. At line 9 the instruction LDR R1, [RO] fetches the first 4 bytes (32-bits) of the 64 -bit value, starting at the location pointed to by R0 and places them in the R1 register. Line 10 loads the second 4 bytes or the lower half of the 64 -bit value from the memroy address pointed to by R0 plus 4 bytes ([R0, \#4]. Between them R1 and R2 now hold the first 64 -bit value, R1 has the upper half while R2 has the lower half. Lines 11-13 repeat this process for the second 64-bit value, reading it into R3 and R4.

Next, the two low order words, held in R2 and R4 are added, and the result stored in R6.
This is all straightforward, but note now the use of the S suffix to the ADD instruction. This forces the update of the flags as a result of the ADD operation. In other words, if the result of the addition results in a carry, the carry flag bit will be set.

Now the ADC (add with carry) instruction is used to add the two high order words, held in R1 and R3, but taking into account any carry resulting from the previous addition.
Finally, the result is stored using the same technique as we used the load the values (lines 16-18).

### 7.1.8 Table of Factorials

Calculate the factorial of the 8 -bit variable Value from a table of factorials DataTable. Store the result in the 16 -bit variable Result. Assume Value has a value between 0 and 7 .

Sample Problems

| Input: $\quad$ FTABLE | $=0001$ | $(0!$ | $=$ | $\left.1_{10}\right)$ |
| ---: | :--- | :--- | :--- | :--- |
|  | $=0001$ | $(1!$ | $=$ | $\left.1_{10}\right)$ |
|  | $=0002 \quad(2!$ | $=$ | $\left.2_{10}\right)$ |  |
|  | $=0006$ | $(3!$ | $=$ | $\left.6_{10}\right)$ |
|  | $=0018$ | $(4!$ | $=$ | $\left.24_{10}\right)$ |
|  | $=0078$ | $(5!$ | $=$ | $\left.120_{10}\right)$ |
|  | $=02 D 0 \quad(6!$ | $=$ | $\left.720_{10}\right)$ |  |
| Output: $\quad$ Result | $=0078 \quad(5!$ | $=$ | $\left.120_{10}\right)$ |  |

Program 7.8: factorial.s - Lookup the factorial from a table by using the address of the memory location
; Lookup the factorial from a table using the address of the memory location
TTL Ch4Ex9 - factorial
AREA Program, CODE, READONLY
ENTRY
Main
LDR RO, =DataTable ; Load the address of the lookup table
LDR R1, Value ; Offset of value to be looked up
MOV R1, R1, LSL \#0x2 ; Data is declared as 32bit - need
; to quadruple the offset to point at the
; correct memory location
ADD RO, RO, R1 ; RO now contains memory address to store
LDR R2, [RO]
LDR R3, =Result ; The address where we want to store the answer
STR R2, [R3] ; Store the answer
SWI \&11

```
AREA DataTable, DATA
DCD 1 ;0! = 1 ; The data table containing the factorials
    DCD 1 ;1! = 1
    DCD 2 ;2! = 2
    DCD 6 ;3! = 6
    DCD 24 ;4! = 24
    DCD 120 ;5! = 120
    DCD 720 ;6! = 720
    DCD 5040 ;7! = 5040
Value DCB 5
Result DCW 0
    END
```

The approach to this table lookup problem, as implemented in this program, demonstrates the use of offset addressing. The first two LDR instructions, load register RO with the start address of the lookup table ${ }^{1}$ and register R1 contents of Value.

The actual calculation of the entry in the table is determined by the first operand of the R1, R1, LSL \#0x2 instruction. The long word contents of address register R1 are added to the long word contents of data register R0 to form the effective address used to index the table entry. When RO is used in this manner, it is referred to as an index register.

### 7.2 Problems

### 7.2.1 64-Bit Data Transfer

Move the contents of the 64 -bit variable VALUE to the 64 -bit variable RESULT.
Sample Problems

| Input: | VALUE | 3 E 2 A 42 A 1 |
| :--- | :--- | :--- |
|  |  | 21 F 260 A 0 |
| Output: | RESULT | 3 E 2 A 42 A 1 |
|  |  | 21 F 260 A 0 |

### 7.2.2 32-Bit Subtraction

Subtract the contents of the 32 -bit variable VALUE1 from the contents of the 32-bit variable VALUE2 and store the result back in VALUE1.

| Sample Problems |  |  |
| :--- | :--- | :--- |
| Input: | VALUE1 | 12343977 |
|  | VALUE2 | 56782182 |
| Output: | VALUE1 | 4443 E 80 B |

### 7.2.3 Shift Right Three Bits

Shift the contents of the 32 -bit variable VALUE right three bits. Clear the three most significant bit postition.

[^5]Sample Problems
$\begin{array}{lllll} & & \text { Test A } & & \text { Test B } \\ \text { Input: } & \text { VALUE } & \text { 415D7834 } & & 9284 \mathrm{C} 15 \mathrm{D} \\ \text { Output: } & \text { VALUE } & & 082 \mathrm{BAF} 06 & \\ 1250982 \mathrm{~B}\end{array}$

### 7.2.4 Halfword Assembly

Combine the low four bits of each of the four consecutive bytes beginning at LIST into one 16-bit halfword. The value at LIST goes into the most significant nibble of the result. Store the result in the 32 -bit variable RESULT.

Sample Problems

| Input: | LIST | 0 C |
| :--- | :--- | :--- |
|  |  | 02 |
|  |  | 06 |
|  |  | 09 |
| Output: | RESULT | 0000 C 269 |

### 7.2.5 Find Smallest of Three Numbers

The three 32-bit variables VALUE1, VALUE2 and VALUE3, each contain an unsigned number. Store the smallest of these numbers in the 32 -bit variable RESULT.

Sample Problems

| Input: | VALUE1 | 91258465 |
| :--- | :--- | :--- |
|  | VALUE2 | 102 C 2056 |
|  | VALUE3 | 70409254 |
| Output: | RESULT | 102 C 2056 |

### 7.2.6 Sum of Squares

Calculate the squares of the contents of word VALUE1 and word VALUE2 then add them together. Please the result into the word RESULT.

Sample Problems

| Input: | VALUE1 | 00000007 |
| :--- | :--- | :--- |
|  | VALUE2 | 00000032 |
| Output: | RESULT | 000009 F 5 |

That is $7^{2}+50^{2}=49+2500=2549$ (decimal)
or $7^{2}+32^{2}=31+9 C 4=9 F 5$ (hexadecimal)

### 7.2.7 $\quad$ Shift Left $n$ bits

Shift the contents of the word VALUE left. The number of bits to shift is contained in the word COUNT. Assume that the shift count is less than 32. The low-order bits should be cleared.

Sample Problems

| Input: |  | Test A | Test B |
| :---: | :---: | :---: | :---: |
|  | VALUE | 182B | 182B |
|  | COUNT | 0003 | 0020 |
| Output: | value | C158 | 0000 |

In the first case the value is to be shifted left by three bits, while in the second case the same value is to be shifted by thirty two bits.

## Logic

## Program 8.7a: bigger.s - Find the larger of two numbers

; Find the larger of two numbers
TTL Ch4Ex7 - bigger
AREA Program, CODE, READONLY
ENTRY

Main
LDR R1, Value1 ; Load the first value to be compared

LDR R2, Value2 ; Load the second value to be compared
CMP R1, R2 ; Compare them
BHI Done ; If R1 contains the highest
MOV R1, R2 ; otherwise overwrite R1
Done
STR R1, Result ; Store the result
SWI \&11
Value1 DCD \& ; Value to be compared
Value2 DCD \&87654321 ; Value to be compared
Result DCD 0 ; Space to store result

END

## Program 8.7a: add64.s - 64 bit addition

```
    ; 64 bit addition
```

            TTL Ch4Ex8 - add64
            AREA Program, CODE, READONLY
            ENTRY
    Main
    LDR RO, =Value1 ; Pointer to first value
LDR R1, [RO] ; Load first part of value1
LDR R2, [R0, \#4] ; Load lower part of value1
LDR RO, =Value2 ; Pointer to second value
LDR R3, [RO] ; Load upper part of value2
LDR R4, [RO, \#4] ; Load lower part of value2
ADDS R6, R2, R4 ; Add lower 4 bytes and set carry flag
ADC R5, R1, R3 ; Add upper 4 bytes including carry
LDR RO, =Result ; Pointer to Result
STR R5, [RO] ; Store upper part of result
STR R6, [R0, \#4] ; Store lower part of result
SWI \&11
Value1 DCD \&12A2E640, \&F2100123 ; Value to be added
Value2 DCD \&001019BF, \&40023F51 ; Value to be added
Result DCD 0 ; Space to store result

Program 8.7a: factorial.s - Lookup the factorial from a table by using the address of the memory location

```
; Lookup the factorial from a table using the address of the memory location
            TTL Ch4Ex9 - factorial
            AREA Program, CODE, READONLY
            ENTRY
    Main
            LDR RO, =DataTable ; Load the address of the lookup table
            LDR R1, Value ; Offset of value to be looked up
            MOV R1, R1, LSL #0x2 ; Data is declared as 32bit - need
                    ; to quadruple the offset to point at the
                    ; correct memory location
            ADD RO, RO, R1 ; RO now contains memory address to store
            LDR R2, [RO]
            LDR R3, =Result ; The address where we want to store the answer
            STR R2, [R3] ; Store the answer
            SWI &11
            AREA DataTable, DATA
DCD \(1 \quad ; 0!=1 \quad\); The data table containing the factorials
            DCD 1 ; ; ! = 1
            DCD 2 ;2! = 2
            DCD 6 ;3! = 6
            DCD 24 ;4! = 24
            DCD 120 ;5! = 120
            DCD 720 ;6! = 720
            DCD 5040 ;7! = 5040
    Value DCB 5
            ALIGN
    Result DCW O
            END
```


## 9 Program Loops

The program loop is the basic structure that forces the CPU to repeat a sequence of instructions. Loops have four sections:

1. The initialisation section, which establishes the starting values of counters, pointers, and other variables.
2. The processing section, where the actual data manipulation occurs. This is the section that does the work.
3. The loop control section, which updates counters and pointers for the next iteration.
4. The concluding section, that may be needed to analyse and store the results.

The computer performs Sections 1 and 4 only once, while it may perform Sections 2 and 3 many times. Therefore, the execution time of the loop depends mainly on the execution time of Sections 2 and 3. Those sections should execute as quickly as possible, while the execution times of Sections 1 and 4 have less effect on overall program speed.

There are typically two methods of programming a loop, these are the "repeat ... until" loop (Algorithm 9.1a) and the "while" loop (Algorithm 9.1b). The repeat-until loop results in the computer always executing the processing section of the loop at least once. On the other hand, the computer may not execute the processing section of the while loop at all. The repeat-until loop is more natural, but the while loop is often more efficient and eliminates the problem of going through the processing sequence once even where there is no data for it to handle.

The computer can use the loop structure to process large sets of data (usually called "arrays"). The simplest way to use one sequence of instructions to handle an array of data is to have the program increment a register (usually an index register or stack pointer) after each iteration. Then the register will contain the address of the next element in the array when the computer re-

| Algorithm 9.1a |
| :--- |
| Initialisation Section |
| Repeat |
| $\quad$ Processing Section |
| Loop Control Section |
| Until task completed |
| Concluding Section |

## Algorithm 9.1b

Initialisation Section
While task incomplete Processing Section Repeat peats the sequence of instructions. The computer can then handle arrays of any length with a single program.
Register indirect addressing is the key to the processing arrays since it allows you to vary the actual address of the data (the "effective address") by changing the contents of a register. The autoincrementing mode is particularly convenient for processing arrays since it automatically updates the register for the next iteration. No additional instruction is necessary. You can even have an automatic increment by 2 or 4 if the array contains 16 -bit or 32 -bit data or addresses.
Although our examples show the processing of arrays with autoincrementing (adding 1, 2, or 4 after each iteration), the procedure is equally valid with autodecrementing (subtracting 1,2 , or 4 before each iteration). Many programmers find moving backward through an array somewhat awkward
and difficult to follow, but it is more efficient in many situations. The computer obviously does not know backward from forward. The programmer, however, must remember that the processor increments an address register after using it but decrements an address register before using it. This difference affects initialisation as follows:

1. When moving forward through an array (autoincrementing), start the register pointing to the lowest address occupied by the array.
2. When moving backward through an array (autodecrementing), start the register pointing one step ( 1,2 , or 4 ) beyond the highest address occupied by the array.

### 9.1 Program Examples

### 9.1.1 Sum of numbers

## 16-bit

```
Program 9.1a: sum16.s - Add a series of 16 bit numbers by using a table address
    * Add a series of 16 bit numbers by using a table address look-up
        TTL Ch5Ex1
        AREA Program, CODE, READONLY
    ENTRY
    Main
        LDR RO, =Data1 ;load the address of the lookup table
        EOR R1, R1, R1 ;clear R1 to store sum
        LDR R2, Length ;init element count
    Loop
        LDR R3, [RO] ;get the data
        ADD R1, R1, R3 ;add it to r1
        ADD RO, RO, #+4 ;increment pointer
        SUBS R2, R2, #0x1 ;decrement count with zero set
        BNE Loop ;if zero flag is not set, loop
        STR R1, Result ;otherwise done - store result
        SWI &11
        AREA Data1, DATA
    Table DCW &2040
        ;table of values to be added
        ALIGN ;32 bit aligned
        DCW &1C22
        ALIGN
        DCW &0242
        ALIGN
    TablEnd DCD 0
    AREA Data2, DATA
    Length DCW (TablEnd - Table) / 4 ;because we're having to align
    Result DCW 0 %ives the loop count 
    END
```

Program 9.1b: sum16b.s - Add a series of 16 bit numbers by using a table address look-up
1 * Add a series of 16 bit numbers by using a table address look-up

```
    * This example has nothing in the lookup table, and the program handles this
    TTL Ch5Ex2
    AREA Program, CODE, READONLY
    ENTRY
    Main
        LDR RO, =Data1 ;load the address of the lookup table
        EOR R1, R1, R1 ;clear R1 to store sum
        LDR R2, Length ;init element count
        CMP R2, #0
        BEQ Done
    Loop
        LDR R3, [RO]
        ;get the data that RO points to
        ADD R1, R1, R3
        ;add it to r1
        ADD RO, RO, #+4
        ;increment pointer
        SUBS R2, R2, #0x1 ;decrement count with zero set
        BNE Loop ;if zero flag is not set, loop
    Done
        STR R1, Result
        ;otherwise done - store result
        SWI &11
        AREA Data1, DATA
    Table
        ;Table is empty
    TablEnd DCD 0
        AREA Data2, DATA
Length DCW (TablEnd - Table) / 4 ;because we're having to align
ALIGN
;gives the loop count
Result DCW 0
                                    ;storage for result
    END
```


## 32-bit

## 64-bit

### 9.1.2 Number of negative elements


Program 9.2a: countneg.s - Scan a series of 32 bit numbers to find how many are negative
TTL Ch5Ex3
AREA Program, CODE, READONLY
ENTRY
Main
LDR RO, =Data1 ;load the address of the lookup table
; clear R1 to store cou
CMP R2, \#0
R3, [RO] ;get the data
BPL Looptest ;skip next line if +ve or zero
RO, RO, \#+4 ;increment pointer
SUBS R2, R2, \#0x1
; decrement count with zero set

```
BNE Loop
Done
    STR R1, Result
                                    ;otherwise done - store result
    SWI &11
    AREA Data1, DATA
Table DCD &F1522040
    DCD &7F611C22
    DCD &80000242
TablEnd DCD 0
    AREA Data2, DATA
Length DCW (TablEnd - Table) / 4 ;because we're having to align
    ALIGN ;gives the loop count
Result DCW 0 ;storage for result
    END
```

Program 9.2b: countneg16.s - Scan a series of 16 bit numbers to find how many are negative

* Scan a series of 16 bit numbers to find how many are negative

TTL Ch5Ex4
AREA Program, CODE, READONLY
ENTRY
Main
LDR RO, =Data1 ; load the address of the lookup table

EOR R1, R1, R1 ;clear R1 to store count
LDR R2, Length ;init element count
CMP R2, \#O
BEQ Done
Loop
$\begin{array}{ll}\text { LDR } & \text { R3, [RO] } \\ \text { AND } & \text { R3, R3, \#O }\end{array}$
;if table is empty

CMP R3, \#0x8000
; get the data

BEQ Looptest
;bit is 1

ADD R1, R1, \#1
Looptest
RO, RO, \#+4
SUBS R2, R2, \#0x1
;skip next line if zero
;increment -ve number count

BNE Loop
Done
STR R1, Result
SWI \&11

AREA Data1, DATA
Table DCW \&F152 ;table of values to be tested
ALIGN
DCW \&7F61
ALIGN
DCW \&8000
ALIGN
TablEnd DCD 0
AREA Data2, DATA
Length DCW (TablEnd - Table) / 4 ;because we're having to align
ALIGN ;gives the loop count
Result DCW $0 \quad$;storage for result
END
;increment pointer
; decrement count with zero set
;if zero flag is not set, loop
;otherwise done - store result

### 9.1.3 Find Maximum Value

```
Program 9.3: largest16.s - Scan a series of 16 bit numbers to find the largest
    * Scan a series of 16 bit numbers to find the largest
        TTL Ch5Ex5
        AREA Program, CODE, READONLY
        ENTRY
    Main
            EOR R1, R1, R1 ;clear R1 to store largest
            LDR R2, Length ;init element count
            CMP R2, #0
            BEQ Done ;if table is empty
    Loop
            LDR R3, [RO] ;get the data
            CMP R3, R1 ;bit is 1
            BCC Looptest ;skip next line if zero
            MOV R1, R3 ;increment -ve number count
    Looptest
            ADD RO, RO, #+4 ;increment pointer
            SUBS R2, R2, #0x1
                                    ;decrement count with zero set
                                    ;if zero flag is not set, loop
    Done
            STR R1, Result
                                    ;otherwise done - store result
            SWI &11
            AREA Data1, DATA
    Table DCW &A152 ;table of values to be tested
            ALIGN
            DCW &7F61
            ALIGN
            DCW &F123
            ALIGN
            DCW &8000
            ALIGN
    TablEnd DCD
    AREA Data2, DATA
    Length DCW (TablEnd - Table) / 4 ;because we're having to align
    ALIGN ;gives the loop count
    Result DCW 0 ;storage for result
    END
```


### 9.1.4 Normalize A Binary Number

## Program 9.4: normalize.s - Normalize a binary number

* normalize a binary number

TTL Ch5Ex6
AREA Program, CODE, READONLY
ENTRY

Main
LDR RO, =Data1 ;load the address of the lookup table EOR R1, R1, R1 ; clear R1 to store shifts

```
;get the data
;bit is 1
;if table is empty
Loop
            R1, R1, #1
                                    ;increment pointer
            MOVS R3, R3, LSL#0x1 % ; decrement count with zero set 
            MOVS R3, R3, LSL#0x1 % ; decrement count with zero set 
Done
            STR R1, Shifted
                                    ;otherwise done - store result
            STR R3, Normal
            SWI &11
            AREA Data1, DATA
Table
    * DCD &30001000
* DCD &00000000
    DCD &C1234567
    AREA Result, DATA
Number DCD Table
Shifted DCB 0
;storage for shift
Normal DCD 0
;storage for result
```


### 9.2 Problems

### 9.2.1 Checksum of data

Calculate the checksum of a series of 8 -bit numbers. The length of the series is defined by the variable LENGTH. The label START indicates the start of the table. Store the checksum in the variable CHECKSUM. The checksum is formed by adding all the numbers in the list, ignoring the carry over (or overflow).

Note: Checksums are often used to ensure that data has been correctly read. A checksum calculated when reading the data is compared to a checksum that is stored with the data. If the two checksums do not agree, the system will usually indicate an error, or automatically read the data again.

Sample Problem:

| Input: | LENGTH | 00000003 | (Number of items) |
| :---: | :---: | :---: | :---: |
|  | START | 28 | (Start of data table) |
|  |  | 55 |  |
|  |  | 26 |  |
| Output: | CHECKSUM | $28+55+26$ | (Data Checksum) |
|  |  | $=00101000$ (28) |  |
|  |  | + 01010101 (55) |  |
|  |  | $=01111101$ (7D) |  |
|  |  | + 00100110 (26) |  |
|  |  | $=10100011$ (A3) |  |

### 9.2.2 Number of Zero, Positive, and Negative numbers

Determine the number of zero, positive (most significant bit zero, but entire number not zero), and negative (most significant bit set) elements in a series of signed 32-bit numbers. The length of the series is defined by the variable LENGTH and the starting series of numbers start with the START label. Place the number of negative elements in variable NUMNEG, the number of zero elements in variable NUMZERO and the number of positive elements in variable NUMPOS.
Sample Problem:

| Input: | LENGTH | 6 | (Number of items) |
| :---: | :---: | :---: | :---: |
|  | START | 76028326 | (Start of data table - Positive) |
|  |  | 8D489867 | (Negative) |
|  |  | 21202549 | (Positive) |
|  |  | 00000000 | ( Zero) |
|  |  | E605546C | (Negative) |
|  |  | 00000004 | (Positive) |
| Output: | NUMNEG | 2 | (2 negative numbers: 8 D489867 and E605546C) |
|  | NUMZERO | 1 | (1 zero value) |
|  | NUMPOS | 3 | (3 positive numbers: 76028326, 21202549 and 00000004 ) |

### 9.2.3 Find Minimum

Find the smallest element in a series of unsigned bytes. The length of the series is defined by the variable LENGTH with the series starting at the START label. Store the minimum byte value in the NUMMIN variable.

Sample Problem:

| Input: | LENGTH | 5 | ( Number of items) |
| :--- | :--- | :--- | :--- |
|  | START | 65 | (Start of data table) |
|  |  | 79 |  |
|  |  | 15 |  |
|  |  | E 3 |  |
| Output: NUMMIN | 15 | (Smallest of the five) |  |

### 9.2.4 Count 1 Bits

Determine the number of bits which are set in the 32-bit variable NUM, storing the result in the NUMBITS variable.

Sample Problem:

| Input: | NUM | $2866 \mathrm{~B} 794=00111000011001101011011110010100$ |
| :--- | :--- | :--- |
| Output: | NUMBITS | $0 \mathrm{~F}=15$ |

### 9.2.5 Find element with most 1 bits

Determine which element in a series of 32 -bit numbers has the largest number of bits set. The length of the series is defined by the LENGTH variable and the series starts with the START label. Store the value with the most bits set in the NUM variable.

| Input: | LENGTH | 5 | (Number of items) |
| :--- | :--- | :--- | :--- |
|  | START | 205A15E3 | $(00100000010110100001010111010011-13)$ |
|  |  | 256C8700 | $(00100101011011001000011100000000-11)$ |
|  |  | 295468 F 2 | $(00101001010101000110100011110010-14)$ |
|  |  | 29856779 | $(00101001100001010110011101111001-16)$ |
| Output: | NUM | 298567592 A | $(10010001010001110101100100101010-14)$ |
|  | (Number with most 1 - bits $)$ |  |  |

## 10 strings

Microprocessors often handle data which represents printed characters rather than numeric quantities. Not only do keyboards, printers, communications devices, displays, and computer terminals expect or provide character-coded data, but many instruments, test systems, and controllers also require data in this form. ASCII (American Standard Code for Information Interchange) is the most commonly used code, but others exist.

We use the standard seven-bit ASCII character codes, as shown in Table 10.1 the character code occupies the low-order seven bits of the byte, and the most significant bit of the byte holds a 0 or a parity bit.

### 10.1 Handling data in ASCII

Here are some principles to remember in handling ASCII-coded data:

- The codes for the numbers and letters form ordered sequences. Since the ASCII codes for the characters " 0 " through " 9 " are $30_{16}$ through $39_{16}$ you can convert a decimal digit to the equivalent ASCII characters (and ASCII to decimal) by simple adding the ASCII offset: $30_{16}$ $=$ ASCII "0". Since the codes for letters ( $41_{16}$ through $5 \mathrm{~A}_{16}$ and $61_{16}$ through $7 \mathrm{~A}_{16}$ ) are in order, you can alphabetises strings by sorting them according to their numerical values.
- The computer does not distinguish between printing and non-printing characters. Only the I/ 0 devices make that distinction.
- An ASCII I/0 device handles data only in ASCII. For example, if you want an ASCII printer to print the digit " 7 ", you must send it $37_{16}$ as the data; $07_{16}$ will ring the bell. Similarly, if a user presses the " 9 " key on an ASCII keyboard, the input data will be $39_{16} ; 09_{16}$ is the tab key.
- Many ASCII devices do not use the entire character set. For example, devices may ignore many control characters and may not print lower-case letters.
- Despite the definition of the control characters many devices interpret them differently. For example they typically uses control characters in a special way to provide features such as cursor control on a display, and to allow software control of characteristics such as rate of data transmission, print width, and line length.
- Some widely used ASCII control characters are:

| $0 \mathrm{~A}_{16}$ | LF | line feed |
| :---: | :--- | :--- |
| $0 \mathrm{D}_{16}$ | CR | carriage return |
| $08_{16}$ | BS | backspace |
| $7 \mathrm{~F}_{16}$ | DEL | rub out or delete character |


| LSB | MSB |  |  |  |  |  |  |  | Control Characters |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |  |  |  |  |
| 0 | NUL | DLE | SP | 0 | © | P | ' | p | NUL | Null | DLE | Data link escape |
| 1 | SOH | DC1 | ! | 1 | A | Q | a | q | SOH | Start of heading | DC1 | Device control 1 |
| 2 | STX | DC2 | " | 2 | B | R | b | r | STX | Start of text | DC2 | Device control 2 |
| 3 | ETX | DC3 | \# | 3 | C | S | c | s | ETX | End of text | DC3 | Device control 3 |
| 4 | EOT | DC4 | \$ | 4 | D | T | d | t | EOT | End of tx | DC4 | Device control 4 |
| 5 | ENQ | NAK | \% | 5 | E | U | e | u | ENQ | Enquiry | NAK | Negative ack |
| 6 | ACK | SYN | \& | 6 | F | V | f | v | ACK | Acknowledge | SYN | Synchronous idle |
| 7 | BEL | ETB | , | 7 | G | W | g | w | BEL | Bell, or alarm | ETB | End of tx block |
| 8 | BS | CAN | ( | 8 | H | X | h | x | BS | Backspace | CAN | Cancel |
| 9 | HT | EM | ) | 9 | I | Y | i | y | HT | Horizontal tab | EM | End of medium |
| A | LF | SUB | * | : | J | Z | j | z | LF | Line feed | SUB | Substitute |
| B | VT | ESC | + | ; | K | [ | k | \{ | VT | Vertical tab | ESC | Escape |
| C | FF | FS | , | $<$ | L | $\backslash$ | 1 | 1 | FF | Form feed | FS | File separator |
| D | CR | GS | - | = | M | ] | m | \} | CR | Carriage return | GS | Group separator |
| E | SO | RS | - | > | N | - | n | $\sim$ | SO | Shift out | RS | Record separator |
| F | SI | US | / | ? | 0 | - | $\bigcirc$ | DEL | SI | Shift in | US | Unit separator |
|  |  |  |  |  |  |  |  |  | SP | Space | DEL | Delete |

Table 10.1: Hexadecimal ASCII Character Codes

- Each ASCII character occupies eight bits. This allows a large character set but is wasteful when only a few characters are actually being used. If, for example, the data consists entirely of decimal numbers, the ASCII format (allowing one digit per byte) requires twice as much storage, communications capacity, and processing time as the BCD format (allowing two digits per byte).

The assembler includes a feature to make character-coded data easy to handle, single quotation marks around a character indicate the character's ASCII value. For example,

```
MOV R3, #'A'
```

is the same as

$$
\text { MOV } \quad \text { R3, \#0x41 }
$$

The first form is preferable for several reasons. It increases the readability of the instruction, it also avoids errors that may result from looking up a value in a table. The program does not depend on ASCII as the character set, since the assembler handles the conversion using whatever code has been designed for.

### 10.2 A string of characters

Individual characters on there own are not really all that helpful. As humans we need a string of characters in order to form meaningful text. In assembly programming it is normal to have to process one character at a time. However, the assembler does at least allow us to store a string of byes (characters) in a friendly manner with the DCB directive. For example, line 26 of program 10.1a is:

```
DCB "Hello, World", CR
```

which will produce the following binary data:

| Binary: | 48 | 65 | 6 C | 6 C | 6 F | 2 C | 20 | 57 | 6 F | 72 | 6 C | 64 | 0 D |
| ---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Text: | H | e | 1 | 1 | o | , | SP | W | o | r | 1 | d | CR |

Use table 10.1 to check that this is correct. In order to make the program just that little bit more readable, line 5 defines the label CR to have the value for a Carriage Return $\left(0 \mathrm{D}_{16}\right)$.

There are three main methods for handling strings: Fixed Length, Terminated, and Counted. It is normal for a high level language to support just one method. C/C++ and Java all support the use of Zero-Terminated strings, while Pascal and Ada use counted strings. Although it is possible to provide your own support for the alternative string type it is seldom done. A good programmer will use a mix of methods depending of the nature of the strings concerned.

### 10.2.1 Fixed Length Strings

A fixed length string is where the string is of a predefined and fixed size. For example, in a system where it is known that all strings are going to be ten characters in length, we can simply reserve 10 bytes for the string.

This has an immediate advantages in that the management of the strings is simple when compared to the alternative methods. For example we only need one label for an array of strings, and we can calculate the starting position of the $n^{\text {th }}$ string by a simple multiplication.

This advantage is however also a major disadvantage. For example a persons name can be anything from two characters to any number of characters. Although it would be possible to reserve sufficient space for the longest of names this amount of memory would be required for all names, including the two letter ones. This is a significant waist of memory.

It would be possible to reserve just ten characters for each name. When a two letter name appears it would have to be padded out with spaces in order to make the name ten characters in length. When a name longer than ten characters appears it would have to be truncated down to just ten characters thus chopping off part of the name. This requires extra processing and is not entirely friendly to users who happen to have a long name.

When there is little memory and all the strings are known in advance it may be a good idea to use fixed length strings. For example, command driven systems tend to use a fixed length strings for the list of commands.

### 10.2.2 Terminated Strings

A terminated string is one that can be of any length and uses a special character to mark the end of the string, this character is known at the sentinel. For example program 10.1a uses the carriage return as it's sentinel.

Over the years several different sentinels have been used, these include $\$\left(26_{16}\right)$, EOT (End of Text $-04_{16}$ ), CR (Carriage Return $-0 D_{16}$ ), LF (Line Feed $-0 A_{16}$ ) and NUL (No character $00_{16}$ ). Today the most commonly used sentinel is the NUL character, primarily because it is used by $\mathrm{C} / \mathrm{C}++$. The NUL character also has a good feeling about it, as it is represented by the value 0 , has no other meaning and it is easier to detected than any other character. This is frequently referred to as a Null- or Zero-Terminated string or simply as an ASCIIZ string.

The terminated string has the advantage that it can be of any length. Processing the string is fairly simply, you enter into a loop processing each character at a time until you reach the sentinel. The disadvantage is that the sentinel character can not appear in the string. This is another reason why the NUL character is such a good choice for the sentinel.

### 10.2.3 Counted Strings

A counted string is one in which the first one or two byte holds the length of the string in characters. Thus a counted string can be of any number of characters up to the largest unsigned number that can be stored in the first byte/word.

A counted string may appear rather clumsy at first. Having the length of the string as a binary value has a distinct advantage over the terminated string. It allow the use of the counting instructions that have been included in many instruction sets. This means we can ignore the testing for a sentinel character and simply decrement our counter, this is a far faster method of working.

To scan through an array of strings we simply point to the first string, and add the length count to our pointer to obtain the start of the next string. For a terminated string we would have to scan for the sentinel for each string.

There are two disadvantages with the counted string. The string does have a maximum length, 255 characters or 64 K depending on the size of the count value ( 8 - or 16 -bit). Although it is normally felt that 64 K should be sufficient for most strings. The second disadvantage is their perceived complexity. Many people feel that the complexity of the counted string outweighs the speed advantage.

### 10.3 International Characters

As computing expands outside of the English speaking world we have to provide support for languages other than standard American. Many European languages use letters that are not available in standard ASCII, for example: œ, $(\mathbb{E}, \varnothing, \emptyset, æ, \nVdash, \nvdash, €, \mathcal{B}, \mathfrak{i}$, and $i$. This is particularly important when dealing with names: Ångstrøm, Karlstraße or Łukasiewicz.

The ASCII character set is not even capable of handling English correctly. When we borrow a word from another language we also use it's diacritic marks (or accents). For example I would rather see pâté on a menu rather than pate. ASCII does not provide support for such accents.

To overcome this limitation the international community has produced a new character encoding, known as Unicode. In Unicode the character code is two bytes long, the first byte indicates which character set the character comes from, while the second byte indicates the character position within the character set. The traditional ASCII character set is incorporated into Unicode as character set zero. In the revised C standard a new data type of wchar was defined to cater for this new "wide character".

While Unicode is sufficient to represent the characters from most modern languages, it is not sufficient to represent all the written languages of the world, ancient and modern. Hence an extended version, known as Unicode-32 is being developed where the character set is a 23 -bit value (three bytes). Unicode is a subset of Unicode-32, while ASCII is a subset of Unicode.

Although we do not consider Unicode you should be aware of the problem of international character sets and the solution Unicode provides.

### 10.4 Program Examples

### 10.4.1 Length of a String of Characters

[^6]```
; Find the length of a CR terminated string
    TTL Ch6Ex1 - strlencr
CR EQU 0xOD
    AREA Program, CODE, READONLY
    ENTRY
Main LDR RO, =Data1 ; Load the address of the lookup table
    EOR R1, R1, R1 ; Clear R1 to store count
Loop LDRB R2, [RO], #1 ; Load the first byte into R2
    CMP R2, #CR ; Is it the terminator ?
    BEQ Done ; Yes => Stop loop
    ADD R1, R1, #1 ; No => Increment count
    BAL Loop ; Read next char
Done STR R1, CharCount ; Store result
    SWI &11
    AREA Data1, DATA
Table
    DCB "Hello, World", CR
    ALIGN
    AREA Result, DATA
CharCount
    DCB 0 ; Storage for count
    END
```

```
Program 10.1b: strlen.s - Find the length of a null terminated string
    ; Find the length of a null terminated string
    TTL Ch6Ex1 - strlen
    AREA Program, CODE, READONLY
    ENTRY
    Main
        LDR RO, =Data1 ; Load the address of the lookup table
        MOV R1, #-1 ; Start count at -1
        ADD R1, R1, #1 ; Increment count
        LDRB R2, [RO], #1 ; Load the first byte into R2
        CMP R2, #0 ; Is it the terminator ?
        BNE Loop ; No => Next char
        STR R1, CharCount ; Store result
        SWI &11
        AREA Data1, DATA
    Table
        DCB "Hello,World", O
        ALIGN
        AREA Result, DATA
    CharCount
        DCB 0 ; Storage for count
    END
```


### 10.4.2 Find First Non-Blank Character

```
Program 10.2: skipblanks.s - Find first non-blank
    * find the length of a string
    TTL Ch6Ex3
    Blank EQU " "
        AREA Program, CODE, READONLY
        ENTRY
    Main
        ADR RO, Data1 ;load the address of the lookup table
        MOV R1, #Blank ;store the blank char in R1
    Loop LDRB R2,[RO], #1 ;load the first byte into R2
        CMP R2, R1 ;is it a blank
        BEQ Loop ;if so loop
        SUB RO, RO, #1 ;otherwise done - adjust pointer
        STR RO, Pointer ;and store it
        SWI &11
        AREA Data1, DATA
    Table
        DCB " 7 "
        ALIGN
        AREA Result, DATA
    Pointer DCD 0 ;storage for count
        ALIGN
        END
```


### 10.4.3 Replace Leading Zeros with Blanks

## Program 10.3: padzeros.s - Supress leading zeros in a string

* supress leading zeros in a string

TTL Ch6Ex4
Blank EQU ,
Zero EQU 'O'
AREA Program, CODE, READONLY
ENTRY
Main

| LDR | RO, =Data1 | ;load the address of the lookup table |
| :--- | :--- | :--- |
| MOV | R1, \#Zero | ;store the zero char in R1 |
| MOV | R3, \#Blank | ;and the blank char in R3 |


|  | SUB | R0, R0, \#1 | ;otherwise adjust the pointer |
| :---: | :---: | :---: | :---: |
|  | STRB | R3, [RO] | ;and store it blank char there |
|  | ADD | RO, RO, \#1 | ;otherwise adjust the pointer |
|  | BAL | Loop | ;and loop |
| Done |  |  |  |
|  | SWI | \&11 | ;all done |
|  | AREA | Data1, DATA |  |
| Table |  |  |  |
|  | DCB | "000007000" |  |
|  | ALIGN |  |  |
|  | AREA | Result, DATA |  |
| Pointer | DCD | 0 | ;storage for count |
|  | ALIGN |  |  |
|  | END |  |  |

### 10.4.4 Add Even Parity to ASCII Chatacters

## Program 10.4: setparity.s - Set the parity bit on a series of characters store the amended string in Result

|  | TTL | Ch6Ex5 |  |
| :---: | :---: | :---: | :---: |
|  | AREA ENTRY | Program, CODE, F |  |
| Main |  |  |  |
|  | LDR | R0, =Data1 | ;load the address of the lookup table |
|  | LDR | R5, =Pointer |  |
|  | LDRB | R1, [RO], \#1 | ;store the string length in R1 |
|  | CMP | R1, \#0 |  |
|  | BEQ | Done | ;nothing to do if zero length |
| MainLoop |  |  |  |
|  | LDRB | R2, [RO], \#1 | ;load the first byte into R2 |
|  | MOV | R6, R2 | ;keep a copy of the original char |
|  | MOV | R2, R2, LSL \#24 | ;shift so that we are dealing with msb |
|  | MOV | R3, \#0 | ;zero the bit counter |
|  | MOV | R4, \#7 | ;init the shift counter |
| ParLoop |  |  |  |
|  | MOVS | R2, R2, LSL \#1 | ; left shift |
|  | BPL | DontAdd | ;if msb is not a one bit, branch |
|  | ADD | R3, R3, \#1 | ;otherwise add to bit count |
| DontAdd |  |  |  |
|  | SUBS | R4, R4, \#1 | ;update shift count |
|  | BNE | ParLoop | ;loop if still bits to check |
|  | TST | R3, \#1 | ; is the parity even |
|  | BEQ | Even | ;if so branch |
|  | ORR | R6, R6, \#0x80 | ;otherwise set the parity bit |
|  | STRB | R6, [R5], \#1 | ;and store the amended char |
|  | BAL | Check |  |
| Even | STRB | R6, [R5], \#1 | ;store the unamended char if even pty |
| Check | SUBS | R1, R1, \#1 | ; decrement the character count |
|  | BNE | MainLoop |  |
| Done | SWI | \&11 |  |

```
    AREA Data1, DATA
    Table DCB 6 ;data table starts with byte length of string
        DCB 0x31
        ;the string
        DCB 0x33
        DCB 0x34
        DCB 0x35
        DCB 0x36
        AREA Result, DATA
        ALIGN
    Pointer DCD 0 ;storage for parity characters
    END
```


### 10.4.5 Pattern Match

## Program 10.5a: cstrcmp.s - Compare two counted strings for equality

* compare two counted strings for equality

TTL Ch6Ex6
AREA Program, CODE, READONLY
ENTRY

Main

| LDR | RO, =Data1 | ;load the address of the lookup table |
| :--- | :--- | :--- |
| LDR | R1, =Data2 |  |
| LDR | R2, Match | ;assume strings not equal - set to -1 |
| LDR | R3, [RO], \#4 | ;store the first string length in R3 |
| LDR | R4, [R1], \#4 | ;store the second string length in R4 |
| CMP | R3, R4 |  |
| BNE | Done | ;if they are different lengths, |
|  |  | ;they can't be equal |
| CMP | R3, \#0 | ;test for zero length if both are |
| BEQ | Same | ;zero length, nothing else to do |

* if we got this far, we now need to check the string char by char

Loop
$\begin{array}{lll}\text { LDRB } & \text { R5, [R0], \#1 } & \text {; character of first string } \\ \text { LDRB } & \text { R6, [R1], \#1 } & \text {; character of second string }\end{array}$
CMP R5, R6 ;are they the same
BNE Done ;if not the strings are different
SUBS R3, R3, \#1 ;use the string length as a counter
BEQ Same ;if we got to the end of the count ;the strings are the same
;not done, loop
Same MOV R2, \#O ;clear the -1 from match ( $0=$ match )
Done STR R2, Mat
SWI \&11

AREA Data1, DATA
Table1 DCD 3 ; data table starts with byte length of string
DCB "CAT" ;the string
Table2 DCD $\begin{aligned} & \text { AREA } \\ & \text { Data2, DATA }\end{aligned} \quad$; data table starts with byte length of string DCB "CAT" ;the string

AREA Result, DATA
;store the result

```
Match
Match DCD &FFFF ;storage for parity characters
    END
```

Program 10.5b: strcmp.s - Compare null terminated strings for equality assume that we have no knowledge of the data structure so we must assess the individual strings

```
; Compare two null terminated strings for equality
    TTL Ch6Ex7
    AREA Program, CODE, READONLY
    ENTRY
Main
        LDR RO, =Data1
        ;load the address of the lookup table
        LDR R1, =Data2
        LDR R2, Match ;assume strings not equal, set to -1
        MOV R3, #0 ;init register
        MOV R4, #0
    Count1
        LDRB R5, [R0], #1 ;load the first byte into R5
        CMP R5, #0 ;is it the terminator
    BEQ Count2 ;if not, Loop
    ADD R3, R3, #1 ;increment count
Count2
    LDRB R5, [R1], #1 ;load the first byte into R5
    CMP R5, #0 ;is it the terminator
    BEQ Next ;if not, Loop
    ADD R4, R4, #1 ;increment count
    BAL Count2
Next CMP R3, R4
    BNE Done
                                    ;if they are different lengths,
                                    ;they can't be equal
    CMP R3,#0 ; test for zero length if both are
    LDR RO, =Data1 ;need to reset the lookup table
    LDR R1, =Data2
    if we got this far, we now need to check the string char by char
Loop
        LDRB R5, [RO], #1 ;character of first string
        LDRB R6, [R1], #1 ;character of second string
        CMP R5, R6 ;are they the same
        BNE Done ;if not the strings are different
        SUBS R3, R3, #1 ;use the string length as a counter
        BEQ Same ;if we got to the end of the count
        BAL Loop ;not done, loop
Same
    MOV R2,#0 ;clear the -1 from match (0 = match)
Done (
    STR R2, Match ;store the result
    SWI &11
    AREA Data1, DATA
Table1 DCB "Hello,World", 0 ;the string
    ALIGN
    AREA Data2, DATA
Table2 DCB "Hello, worl", 0 ; the string
```

```
    AREA Result, DATA
    ALIGN
    Match DCD &FFFF ;flag for match
    END
```


### 10.5 Problems

### 10.5.1 Length of a Teletypewriter Message

Determine the length of an ASCII message. All characters are 7-bit ASCII with MSB $=0$. The string of characters in which the message is embedded has a starting address which is contained in the START variable. The message itself starts with an ASCII STX (Start of Text) character $\left(02_{16}\right)$ and ends with $E T X$ (End of Text) character $\left(03_{16}\right)$. Save the length of the message, the number of characters between the STX and the ETX markers (but not including the markers) in the LENGTH variable.
Sample Problem:

| Input: | START | String | (Location of string) |
| :--- | :--- | :--- | :--- |
|  | String | 02 | (STX - Start Text) |
|  |  | 47 | ("G") |
|  |  | 4 F | ("0") |
| Output: | LENGTH | 03 | (ETX - End Text) |
|  |  | 03 | ("GO") |

### 10.5.2 Find Last Non-Blank Character

Search a string of ASCII characters for the last non-blank character. Starting address of the string is contained in the START variable and the string ends with a carriage return character ( $0 D_{16}$ ). Place the address of the last non-blank character in the POINTER variable.

Sample Problems:

| Input: | Start | Test A | Test B |
| :---: | :---: | :---: | :---: |
|  |  | String | String |
|  | String | 37 ("7") | 41 ("A") |
|  |  | 0D (CR) | 20 (Space) |
|  |  |  | 48 ("H") |
|  |  |  | 41 ("A") |
|  |  |  | 54 ("T") |
|  |  |  | 20 (Space) |
|  |  |  | 20 (Space) |
|  |  |  | 0D (CR) |
| Output: | POINTER | First Char | Fourth Char |

### 10.5.3 Truncate Decimal String to Integer Form

Edit a string of ASCII decimal characters by replacing all digits to the right of the decimal point with ASCII blanks $\left(20_{16}\right)$. The starting address of the string is contained in the START variable and the string is assumed to consist entirely of ASCII-coded decimal digits and a possible decimal point $\left(2 \mathrm{E}_{16}\right)$. The length of the string is stored in the LENGTH variable. If no decimal point appears in the string, assume that the decimal point is at the far right.

Sample Problems:

|  |  | Test A | Test B |
| :---: | :---: | :---: | :---: |
| Input: | Start | String | String |
|  | LENGTH | 4 | 3 |
|  | String | 37 ("7") | 36 ("6") |
|  |  | 2E (".") | 37 ("7") |
|  |  | 38 ("8") | 31 ("1") |
|  |  | 31 ("1") |  |
| Output: |  | 37 ("7") | 36 ("6") |
|  |  | 2E (".") | 37 ("7") |
|  |  | 20 (Space) | 31 ("1") |
|  |  | 20 (Space) |  |

Note that in the second case (Test B) the output is unchaged, as the number is assumed to be "671.".

### 10.5.4 Check Even Parity and ASCII Characters

Cheek for even parity in a string of ASCII characters. A string's starting address is contained in the START variable. The first word of the string is its length which is followed by the string itself. If the parity of all the characters in the string are correct, clear the PARITY variable; otherwise place all ones ( $\mathrm{FFFFFFFF}_{16}$ ) into the variable.

Sample Problems:

|  |  | Test A |  |  |
| :--- | :--- | :--- | :--- | :--- |
| Input: | START | String |  | Test B |
|  | String | 3 |  | 03 |
|  |  | B1 (1011 0001) |  | B1 (1011 0001) |
|  |  | B2 (1011 0010) |  | B6 (1011 0110) |
|  |  | $33(00110011)$ |  | $33(00110011)$ |
| Output: | PARITY | 00000000 (True) |  | FFFFFFFF (False) |

### 10.5.5 String Comparison

Compare two strings of ASCII characters to see which is larger (that is, which follows the other in alphabetical ordering). Both strings have the same length as defined by the LENGTH variable. The strings' starting addresses are defined by the START1 and START2 variables. If the string defined by START1 is greater than or equal to the other string, clear the GREATER variable; otherwise set the variable to all ones ( $\mathrm{FFFFFFFF}_{16}$ ).

Sample Problems:


## 11 Code Conversion

Code conversion is a continual problem in microcomputer applications. Peripherals provide data in ASCII, BCD or various special codes. The microcomputer must convert the data into some standard form for processing. Output devices may require data in ASCII, BCD, seven-segment or other codes. Therefore, the microcomputer must convert the results to the proper form after it completes the processing.

There are several ways to approach code conversion:

1. Some conversions can easily be handled by algorithms involving arithmetic or logical functions. The program may, however, have to handle special cases separately.
2. More complex conversions can be handled with lookup tables. The lookup table method requires little programming and is easy to apply. However the table may occupy a large amount of memory if the range of input values is large.
3. Hardware is readily available for some conversion tasks. Typical examples are decoders for BCD to seven-segment conversion and Universal Asynchronous Receiver/Transmitters (UARTs) for conversion between parallel and serial formats.

In most applications, the program should do as much as possible of the code conversion work. Most code conversions are easy to program and require little execution time.

### 11.1 Program Examples

### 11.1.1 Hexadecimal to ASCII

```
Program 11.1a: nibtohex.s - Convert a single hex digit to its ASCII equivalent
    * convert a single hex digit to its ASCII equivalent
        TTL Ch7Ex1
        AREA Program, CODE, READONLY
        ENTRY
    Main
        LDR RO, Digit ;load the digit
        LDR R1, =Result ;load the address for the result
        CMP RO, #OXA ;is the number < 10 decimal
        BLT Add_0 ;then branch
        ADD RO, RO, #"A"-"O"-OxA ;add offset for 'A' to 'F'
    Add_0
        ADD RO, RO, #"O" ;convert to ASCII
        STR RO, [R1]
        ;store the result
```

```
AREA Data1, DATA
Digit DCD &OC ; the hex digit
Result DRD AREA Data2, DATA 0, % ; storage for result
    END
```

Program 11.1b: wordtohex.s - Convert a 32 bit hexadecimal number to an ASCII string and output to the terminal

* now something a little more adventurous - convert a 32 bit
* hexadecimal number to an ASCII string and output to the terminal

TTL Ch7Ex2
AREA Program, CODE, READONLY
ENTRY
Mask EQU 0x0000000F
start

| LDR | R1, Digit | ;load the digit |
| :--- | :--- | :--- |
| MOV | R4, \#8 | ;init counter |

MainLoop
MOV R3, R1 ;copy original word
MOV R3, R3, LSR R5 ;right shift the correct number of bits
SUB R5, R5, \#4 ;reduce the bit shift
AND R3, R3, \#Mask ; mask out all but the ls nibble
CMP R3, \#0xA ;is the number < 10 decimal
BLT Add_0 ;then branch
ADD R3, R3, \#"A"-"O"-0xA ;add offset for 'A, to 'F,
Add_0 ADD R3, R3, \#"0" ;convert to ASCII
MOV RO, R3 ;prepare to output
SWI \&O ;output to console
SUBS R4, R4, \#1 ; decrement counter
BNE MainLoop
MOV RO, \#\&OD ;add a CR character
SWI \&O ;output it
SWI \&11 ;all done
AREA Data1, DATA
Digit DCD \&DEADBEEF ;the hex word

END

### 11.1.2 Decimal to Seven-Segment

Program 11.2: nibtoseg.s - Convert a decimal number to seven segment binary
convert a decimal number to seven segment binary
TTL Ch7Ex3

AREA Program, CODE, READONLY
ENTRY

| Main |  |  |  |
| :---: | :---: | :---: | :---: |
|  | LDR | R0, =Data1 | ; load the start address of the table |
|  | EOR | R1, R1, R1 | ; clear register for the code |
|  | LDRB | R2, Digit | ;get the digit to encode |
|  | CMP | R2, \#9 | ;is it a valid digit? |
|  | BHI | Done | ; clear the result |
|  | ADD | R0, R0, R2 | ;advance the pointer |
|  | LDRB | R1, [RO] | ;and get the next byte |
| Done |  |  |  |
|  | STR | R1, Result | ;store the result |
|  | SWI | \&11 | ;all done |
|  | AREA | Data1, DATA |  |
| Table | DCB | \&3F | ; the binary conversions table |
|  | DCB | \&06 |  |
|  | DCB | \&5B |  |
|  | DCB | \&4F |  |
|  | DCB | \&66 |  |
|  | DCB | \&6D |  |
|  | DCB | \&7D |  |
|  | DCB | \&07 |  |
|  | DCB | \& 7 F |  |
|  | DCB | \&6F |  |
|  | ALIGN |  |  |
|  | AREA | Data2, DATA |  |
| Digit | $\begin{aligned} & \text { DCB } \\ & \text { ALIGN } \end{aligned}$ | \&05 | ; the number to convert |
|  | AREA | Data3, DATA |  |
| Result | DCD | 0 | ;storage for result |
|  | END |  |  |

### 11.1.3 ASCII to Decimal

Program 11.3: dectonib.s - Convert an ASCII numeric character to decimal

* convert an ASCII numeric character to decimal

TTL Ch7Ex4
AREA Program, CODE, READONLY
ENTRY
Main

| MOV | R1, \#-1 |
| :--- | :--- |
| LDRB | R0, Char |

; set -1 as error flag
RO, Char
;get the character
SUBS RO, RO, \#"O" ;convert and check if character is < 0
BCC Done ;if so do nothing
CMP RO, \#9 ;check if character is > 9
BHI Done ;if so do nothing
MOV R1, RO ;otherwise...
Done
STR R1, Result
;.....store the decimal no
;all done
$\begin{array}{lll} & \text { AREA } & \text { Data1, } \text { DATA } \\ \text { Char } & \text { DCB } & \& 37\end{array}$
;ASCII representation of 7

```
Resul AREA Data2, DATA
Result DCD 0 ;storage for result
    END
```


### 11.1.4 Binary-Coded Decimal to Binary

## Program 11.4a: ubcdtohalf.s - Convert an unpacked BCD number to binary



Program 11.4b: ubcdtohalf2.s - Convert an unpacked BCD number to binary using MUL

* convert an unpacked $B C D$ number to binary using MUL

TTL Ch7Ex6
AREA Program, CODE, READONLY
ENTRY

Main

| LDR | R0, =BCDNum | ;load address of BCD number |
| :--- | :--- | :--- |
| MOV | R5, \#4 | ;init counter |
| MOV | R1, \#0 | ; clear result register |
| MOV | R2, \#0 | ;and final register |
| MOV | R7, \#10 | ;multiplication constant |

```
Loop
    MOV R6, R1
    MUL R1, R6, R7 ;mult by 10
    LDRB R4, [RO], #1 ;load digit and incr address
    ADD R1, R1, R4 ;add the next digit
    SUBS R5, R5, #1 ;decr counter
    BNE Loop ;if count != 0, loop
    STR R1, Result ;store the result
    SWI &11 ;all done
    AREA Data1, DATA
BCDNum DCB &02,&09,&07,&01 ; an unpacked BCD number
    ALIGN
    AREA Data2, DATA
Result DCD 0 ;storage for result
    END
```


### 11.1.5 Binary Number to ASCII String

```
Program 11.5: halftobin.s - Store a 16bit binary number as an ASCII string of '0's and '1's
```




Figure 11.1: Seven-Segment Display

### 11.2 Problems

### 11.2.1 ASCII to Hexadecimal

Convert the contents of the A_DIGIT variable from an ASCII character to a hexadecimal digit and store the result in the H_DIGIT variable. Assume that A_DIGIT contains the ASCII representation of a hexadecimal digit ( 7 bits with $\mathrm{MSB}=0$ ).
Sample Problems:

|  |  | Test A | Test B |
| :---: | :---: | :---: | :---: |
| Input: | A_DIGIT | 43 ("C") | 36 ("6") |
| Output: | H_DIGIT | 0C | 06 |

### 11.2.2 Seven-Segment to Decimal

Convert the contents of the CODE variable from a seven-segment code to a decimal number and store the result in the NUMBER variable. If CODE does not contain a valid seven-segment code, set NUMBER to $\mathrm{FF}_{16}$. Use the seven-segment table given in Figure 11.1 and try to match codes.
Sample Problems:


### 11.2.3 Decimal to ASCII

Convert the contents of the variable DIGIT from decimal digit to an ASCII character and store the result in the variable CHAR. If the number in DIGIT is not a decimal digit, set the contents of CHAR to an ASCII space $\left(20_{16}\right)$.

Sample Problems:

|  |  | Test A |  | Test B |
| :--- | :--- | :--- | :--- | :--- |
| Input: | DIGIT | 07 |  | 55 |
| Output: | CHAR | $37(" 7 ")$ |  | 20 (Space) |

### 11.2.4 Binary to Binary-Coded-Decimal

Convert the contents of the variable NUMBER to four BCD digits in the STRING variable. The 32 -bit number in NUMBER is unsigned and less than 10,000 .

Sample Problem:

| Input: | NUMBER | 1 C 52 | $\left(7250_{10}\right)$ |
| :--- | :--- | :--- | :--- |
| Output: | STRING | 07 | $(" 7 ")$ |
|  |  | 02 | $(" 2 ")$ |
|  |  | 05 | $(" 5 ")$ |
|  |  | 00 | $(" 0 ")$ |

### 11.2.5 Packed Binary-Coded-Decimal to Binary String

Convert the eight digit packed binary-coded-decimal number in the BCDNUM variable into a 32 -bit number in a NUMBER variable.

Sample Problem:
Input: BCDNUM 92529679
Output: NUMBER 0583E40916 (9252967910)

### 11.2.6 ASCII string to Binary number

Convert the eight ASCII characters in the variable STRING to an 8-bit binary number in the variable NUMBER. Clear the byte variable ERROR if all the ASCII characters are either ASCII "1" or ASCII " 0 "; otherwise set ERROR to all ones ( $\mathrm{FF}_{16}$ ).
Sample Problems:

| Input: | STRING | Test A |  | Test B |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | 31 | ("1") | 31 | ("1") |
|  |  | 31 | ("1") | 31 | ("1") |
|  |  | 30 | ("0") | 30 | ("0") |
|  |  | 31 | ("1") | 31 | ("1") |
|  |  | 30 | ("0") | 30 | ("0") |
|  |  | 30 | ("0") | 37 | ("7") |
|  |  | 31 | ("1") | 31 | ("1") |
|  |  | 30 | ("0") | 30 | ("0") |
| Output: | NUMBER | D2 | (1101 0010) | 00 | ( Valid) |
|  | ERROR | 0 | (No Error) | FF | (Error) |

## 12 Arithmetic

Much of the arithmetic in some microprocessor applications consists of multiple-word binary or decimal manipulations. The processor provides for decimal addition and subtraction, but does not provide for decimal multiplication or division, you must implement these operations with sequences of instruction.

Most processors provide for both signed and unsigned binary arithmetic. Signed numbers are represented in two's complement form. This means that the operations of addition and subtraction are the same whether the numbers are signed or unsigned.

Multiple-precision binary arithmetic requires simple repetitions of the basic instructions. The Carry flag transfers information between words. It is set when an addition results in a carry or a subtraction results in a borrow. Add with Carry and Subtract with Carry use this information from the previous arithmetic operation.
Decimal arithmetic is a common enough task for microprocessors that most have special instructions for this purpose. These instructions may either perform decimal operations directly or correct the results of binary operations to the proper decimal form. Decimal arithmetic is essential in such applications as point-of-sale terminals, check processors, order entry systems, and banking terminals.

You can implement decimal multiplication and division as series of additions and subtractions, respectively. Extra storage must be reserved for results, since a multiplication produces a result twice as long as the operands. A division contracts the length of the result. Multiplications and divisions are time-consuming when done in software because of the repeated operations that are necessary.

### 12.1 Program Examples

### 12.1.2 64-Bit Addition

```
Program 12.2: add64.s - 64 Bit Addition
    * 64 bit addition
            TTL 64 bit addition
            AREA Program, CODE, READONLY
            ENTRY
    Main
            LDR RO, =Value1 ; Pointer to first value
            LDR R1, [RO] ; Load first part of value1
            LDR R2, [RO, #4] ; Load lower part of value1
            LDR RO, =Value2 ; Pointer to second value
            LDR R3, [RO] ; Load upper part of value2
            LDR R4, [RO, #4] ; Load lower part of value2
            ADDS R6, R2, R4 ; Add lower 4 bytes and set carry flag
```

```
    ADC R5, R1, R3 ; Add upper 4 bytes including carry
    LDR RO, =Result ; Pointer to Result
    STR R5, [RO] ; Store upper part of result
    STR R6, [RO, #4] ; Store lower part of result
    SWI &11
Value1 DCD &12A2E640, &F2100123 ; Value to be added
Value2 DCD &001019BF, &40023F51 ; Value to be added
Result DCD 0 ; Space to store result
```


### 12.1.3 Decimal Addition

## Program 12.3: addbcd.s - Add two packed BCD numbers to give a packed BCD result

* add two packed BCD numbers to give a packed BCD result

| TTL | Ch8Ex3 |
| :--- | :--- |
| AREA | Program, CODE, READONLY |

ENTRY
Mask EQU 0x0000000F

Main

| LDR | R0, =Result | ; address for storage |
| :--- | :--- | :--- |
| LDR | R1, BCDNum1 | ;load the first BCD number |
| LDR | R2, BCDNum2 | ;and the second |
| LDRB | R8, Length | ; init counter |
| ADD | R0, RO, \#3 | ;adjust for offset |
| MOV | R5, \#0 | ;carry |

## Loop

MOV R3, R1
; copy what is left in the data register ;and the other number
;mask out everything except low order nibble ;mask out everything except low order nibble ;shift the original number one nibble
;shift the original number one nibble
;add the digits
;and the carry
;is it over 10 ?
;if not, reset the carry to 0
;otherwise set the carry
;and subtract 10
; carry reset to 0
Next

```
RCarry2
MOV R5,#O
    MOV R7, R7, LSL #4 ;shift the second digit processed to the left
    ORR R6, R6, R7 ;and OR in the first digit to the ls nibble
    STRB R6, [RO], #-1 ;store the byte, and decrement address
    SUBS R8, R8, #1 ;decrement loop counter
    BNE Loop ;loop while > 0
    SWI &11
    AREA Data1, DATA
Length DCB &04
    ALIGN
BCDNum1 DCB &36, &70, &19, &85 ;an 8 digit packed BCD number
    AREA Data2, DATA
BCDNum2 DCB & 12, &66, &34, &59 ;another }8\mathrm{ digit packed BCD number
    AREA Data3, DATA
Result DCD 0 ;storage for result
    END
```


### 12.1.4 Multiplication

## 16-Bit



## 32-Bit

Program 12.4b: mul32.s - Multiply two 32 bit number to give a 64 bit result (corrupts R0 and

```
R1)
* multiply two 32 bit number to give a 64 bit result
* (corrupts RO and R1)
    TTL Ch8Ex4
    AREA Program, CODE, READONLY
    ENTRY
Main
        LDR RO,Number
        R1, Number2
        LDR R6, =Result ;load the address of result
        MOV R5, RO, LSR #16 ;top half of RO
        MOV R3, R1, LSR #16 ;top half of R1
        BIC RO, RO, R5, LSL #16 ;bottom half of R0
        BIC R1, R1, R3, LSL #16 ;bottom half of R1
        MUL R2, R0, R1 ;partial result
        MUL RO, R3, RO ;partial result
        MUL R1, R5, R1 ;partial result
        MUL R3, R5, R3 ;partial result
        ADDS RO, R1, RO ;add middle parts
        ADDCS R3, R3, #&10000
        ;add in any carry from above
        ADDS R2, R2, RO, LSL #16 ;LSB 32 bits
        ADC R3, R3, RO, LSR #16 ;MSB 32 bits
        STR R2, [R6]
        ADD R6, R6, #4 ;increment pointer
        STR R3, [R6] ;store MSB
        SWI &11 ;all done
    AREA Data1, DATA
    Number2 DCD &ABCDEF01
        ALIGN
        AREA Data2, DATA
    Result DCD 0 ;storage for result
        ALIGN
    END
```


### 12.1.5 32-Bit Binary Divide

Program 12.5: divide.s - Divide a 32 bit binary no by a 16 bit binary no store the quotient and remainder there is no 'DIV' instruction in ARM!


```
    ADD R3, R3, #1 ;add one to quotient
    SUB RO, RO, R1 ;take away the number you first thought of
    B Loop ;and loop
Err
    MOV R3, #OxFFFFFFFF
;error flag (-1)
Done
    STR RO, Remain
    STR R3, Quotient - - 
    SWI &11 ;all done
    AREA Data1, DATA
Number1 DCD &0075CBB1
Number2 DCD &0141
;a 16 bit binary number
;another
    ALIGN
    AREA Data2, DATA
Quotient DCD O
Remain DCD 0
;storage for result
;storage for remainder
```


### 12.2 Problems

### 12.2.1 Multiple precision Binary subtraction

Subtract one multiple-word number from another. The length in words of both numbers is in the LENGTH variable. The numbers themselves are stored (most significant bits First) in the variables NUM1 and NUM2 respectively. Subtract the number in NUM2 from the one in NUM1. Store the difference in NUM1.

Sample Problem:

| Input: | LENGTH | 3 | (Number of words in each number) |
| :---: | :---: | :---: | :---: |
|  | NUM1 | 2F5B8568 <br> 84C32546 <br> 706 C 9567 | (First number is 2F5B856884C32546706C956716) |
|  | NUM2 | 14DF4098 <br> 85B81095 <br> A3BC1284 | (The second number is 14DF409885B81095A3BC128416) |
| Output: | NUM1 | 1 A 7 C 44 CF <br> FF0B14B0 <br> CCB082E3 | (Difference is 1A7C44CFFF0B14B0CCB082E316) |

That is,

- 14DF409885B81095A3BC1284

1A7C44CFFF0B14B0CCB082E3

### 12.2.2 Decimal Subtraction

Subtract one packed decimal (BCD) number from another. The length in bytes of both numbers is in the byte variable LENGTH. The numbers themselves are in the variables NUM1 and NUM2 respectively. Subtract the number contained in NUM2 from the one contained in NUM1. Store the difference in NUM1.

| Sample Problem: |  |  |  |
| :---: | :---: | :---: | :---: |
| Input: | LENGTH | 4 | (Number of bytes in each number) |
|  | nuM1 | 36 | ( The first number is $367019857834_{10}$ ) |
|  |  | 70 |  |
|  |  | 19 |  |
|  |  | 85 |  |
|  |  | 78 |  |
|  |  | 34 |  |
|  | NUM2 | 12 | (The second number is 12663459326910) |
|  |  | 66 |  |
|  |  | 34 |  |
|  |  | 59 |  |
|  |  | 32 |  |
|  |  | 69 |  |
| Output: | nuM1 | 24 | (Difference is 24038526456510 ) |
|  |  | 03 |  |
|  |  | 85 |  |
|  |  | 26 |  |
|  |  | 45 |  |
|  |  | 65 |  |

That is,

$$
\begin{array}{r}
367019857834 \\
-\quad 126634593269 \\
\hline 240385264565
\end{array}
$$

### 12.2.3 32-Bit by 32-Bit Multiply

Multiply the 32-bit value in the NUM1 variable by the value in the NUM2 variable. Use the MULU instruction and place the result in the 64-bit variable PROD1.

```
Sample Problem:
    Input: NUM1 0024 (The first number is 2468AC16)
        68AC
    NUM2 0328 (The second number is 328108810)
        1088
    Output: PROD1 0000
        72EC (MULU product is 72ECB8C25B6016)
        B8C2
        5B60
    PROD2 0000
        72EC (Shift product is 72ECB8C25B6016)
        B8C2
        5B60
```


## 13 Tables and Lists

Tables and lists are two of the basic data structures used with all computers. We have already seen tables used to perform code conversions and arithmetic. Tables may also be used to identify or respond to commands and instructions, provide access to files or records, define the meaning of keys or switches, and choose among alternate programs. Lists are usually less structured than tables. Lists may record tasks that the processor must perform, messages or data that the processor must record, or conditions that have changed or should be monitored.

### 13.1 Program Examples

### 13.1.1 Add Entry to List

Program 13.1a: insert.s - Examine a table for a match - store a new entry at the end if no match found


```
    AREA Data2, DATA
NewItem DCD &16FA
List DCD Start
    END
```

Program 13.1b: insert2.s - Examine a table for a match - store a new entry if no match found extends insert.s

* examine a table for a match - store a new entry if no match found
* extends Ch9Ex1

| TTL | Ch9Ex2 |
| :--- | :--- |
| AREA | Program, CODE, READONLY |

AREA Program, CODE, READONLY
ENTRY
Main

| LDR | RO, List | ;load the start add |
| :--- | :--- | :--- |
| LDR | $R 1$, NewItem | ;load the new item |

LDR R3, [RO] ; copy the list counter
LDR R2, [RO], \#4 ;init counter and increment pointer
CMP R3, \#0 ;it's an empty list
BEQ Insert ;so store it
LDR R4, [RO], \#4 ;not empty - move to 1st item

## Loop

$\begin{array}{ll}\text { CMP } & \text { R1, R4 } \\ \text { BEQ } & \text { Done }\end{array}$
; does the item match the list?
BEQ Done
;found it - finished
SUBS R2, R2, \#1 ;no - get the next item
LDR R4, [RO], \#4 ;get the next item
BNE Loop
;and loop
SUB RO, RO, \#4 ;adjust the pointer
Insert ADD R3, R3, \#1 ;incr list count
STR R3, Start ;and store it
STR R1, [RO]
Done SWI \&11
;all done
AREA Data1, DATA
Start DCD \&4 ;length of list
DCD \&5376 ;items
$\begin{array}{lr}\text { DCD } & \& 7615 \\ \text { DCD } & \& 138 \mathrm{~A}\end{array}$
DCD \&21DC
;reserve 20 bytes of storage

### 13.1.2 Check an Ordered List

Program 13.2: search.s - Examine an ordered table for a match

* examine an ordered table for a match

TTL Ch9Ex3
AREA Program, CODE, READONLY

```
    ENTRY
Main
    LDR RO, =NewItem ;load the address past the list
    SUB RO, RO, #4
    LDR R1, NewItem
    ;adjust pointer to point at last element of list
    ;load the item to test
    LDR R3, Start ;init counter by reading index from list
    CMP R3, #0
    BEQ Missing
    LDR R4, [RO], #-4
Loop
    CMP R1, R4
    BEQ Done
    BHI Missing
    SUBS R3, R3, #1
    LDR R4, [RO], #-4
    BNE Loop
Missing MOV R3, #OxFFFFFFFF
Done STR R3, Index
    SWI &11
    AREA Data1, DATA
Start DCD &4 ;length of list
    DCD &0000138A ;items
    DCD &000A21DC
    DCD &001F5376
    DCD &09018613
    AREA Data2, DATA
    NewItem DCD &001F5376
    Index DCW O
List DCD Start
    END
```


### 13.1.3 Remove an Element from a Queue

## Program 13.3: head.s - Remove the first element of a queue

|  | TTL | Ch9Ex4 |  |
| :---: | :---: | :---: | :---: |
|  | AREA | Program, COD |  |
|  | Entry |  |  |
| Main |  |  |  |
|  | LDR | R0, Queue | ;load the head of the queue |
|  | STR | R1, Pointer | ;and save it in 'Pointer' |
|  | CMP | R0, \#0 | ;is it NULL? |
|  | BEQ | Done | ;if so, nothing to do |
|  | LDR | R1, [R0] | ; otherwise get the ptr to next |
|  | STR | R1, Queue | ;and make it the start of the queue |
| Done | SWI | \&11 |  |
|  | AREA | Data1, DATA |  |
| Queue | DCD | Item1 | ;pointer to the start of the queue |
| Pointer | DCD | 0 | ;space to save the pointer |
| DArea | \% | 20 | ;space for new entries |

```
    * each item consists of a pointer to the next item, and some data
    Item1 DCD Item2 ;pointer
        DCB 30,20 ;data
    Item2 DCD Item3 ;pointer
        DCB 30, 0xFF ;data
    Item3 DCD 0 ;pointer (NULL)
        DCB 30,&87,&65 ;data
        END
```


### 13.1.4 Sort a List

## Program 13.4: sort.s - Sort a list of values - simple bubble sort

* sort a list of values - simple bubble sort

TTL Ch9Ex5
AREA Program, CODE, READONLY
ENTRY
Main

| LDR | R6, List | ; pointer to start of list |
| :--- | :--- | :--- |
| MOV | RO, \#O | ; clear register |
| LDRB | RO, [R6] | ; get the length of list |

RO, [R6] ;get the length of lis
Sort
ADD $\begin{aligned} & \text { R7, R6, R0 } \\ & \text { MOV } \\ & \text { R1, \#0 }\end{aligned}$
; get address of last element
;zero flag for changes
;move 1 byte up the list each
;iteration
;load the first byte
;and the second
; compare them
;branch if r2 less than r3
otherwise swap the bytes
;like this
STRB R3, [R7]
;flag that changes made
SUB R7, R7, \#1 ;decrement address to check

## NoSwitch

CMP R7, R8 ;have we checked enough bytes?
BHI Next ;if not, do inner loop
CMP R1, \#O ;did we mke changes
BNE Sort ;if so check again - outer loop
Done SWI \&11 ;all done

AREA Data1, DATA
Start DCB 6
DCB \&2A, \&5B, \&60, \&3F, \&D1, \&19
AREA Data2, DATA
List DCD Start
END

### 13.1.5 Using an Ordered Jump Table

### 13.2 Problems

### 13.2.1 Remove Entry from List

Remove the value in the variable ITEM at a list if the value is present. The address of the list is in the LIST variable. The first entry in the list is the number (in words) of elements remaining in the list. Move entries below the one removed up one position and reduce the length of the list by one.

Sample Problems:

|  | Test A |  |  | Test B |  |
| :--- | :--- | :--- | :--- | :--- | :--- |
|  | Input | Output |  | Input | Output |
| ITEM | D010257 |  |  | D0102596 |  |
| LIST | Table |  |  | Table |  |
| Table | 00000004 | No change |  | 00000004 | 0003 |
|  | 2946C121 | since item |  | C1212546 | C1212546 |
|  | 2054A346 | not in list |  | D0102596 | 3A64422B |
|  | 05723A64 |  |  | 3A64422B | 6C20432E |
|  | 12576C20 |  |  | 6C20432E | - |

### 13.2.2 Add Entry to Ordered List

Insert the value in the variable ITEM into an ordered list if it is not already there. The address of the list is in the LIST variable. The first entry in the list is the list's length in words. The list itself consists of unsigned binary numbers in increasing order. Place the new entry in the correct position in the list, adjust the element below it down, and increase the length of the list by one.

Sample Problems

|  | Test A |  |  | Test B |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| ITEM | Input | Output |  | Input | Output |
| LIST | Table |  |  | 7A35B310 |  |
| Table | 00000004 | 0005 |  | 00000005 | No change |
|  | 09250037 | 09250037 |  | 09250037 | since ITEM |
|  | 29567322 | 29567322 |  | 29567322 | already in |
|  | A356A101 | 7A35B310 |  | 7A35B310 | list. |
|  | E235C203 | A356A101 |  | A356A101 |  |
|  | - | E235C203 |  | E235C203 |  |

### 13.2.3 Add Element to Queue

Add the value in the variable ITEM to a queue. The address of the first element in the queue is in the variable QUEUE. Each element in the queue contains a structure of two items (value and next) where next is either the address of the next element in the queue or zero if there is no next element. The new element is placed at the end (tail) of the queue; the new element's address will be in the element that was at the end of the queue. The next entry of the new element will contain zero to indicate that it is now the end of the queue.

Sample Problem:

|  | Input |  | Output |  |
| :--- | :---: | :---: | :---: | :---: |
|  | Value | Next | Value | Next |
| ITEM | 23854760 |  | 23854760 |  |
| QUEUE | 00000001 | item1 | 00000001 | item1 |
| item1 | 00000123 | item2 | 00000123 | item2 |
| item2 | 00123456 | 00000000 | 00123456 | item3 |
| item3 | - | - | $\mathbf{2 3 8 5 4 7 6 0}$ | OOOOOOOO |

### 13.2.4 4-Byte Sort

Sort a list of 4-byte entries into descending order. The first three bytes in each entry are an unsigned key with the first byte being the most significant. The fourth byte is additional information and should not be used to determine the sort order, but should be moved along with its key. The number of entries in the list is defined by the word variable LENGTH. The list itself begins at location LIST.
Sample Problem:
Input Output

| LENGTH | 00000004 |  |  |  |
| :--- | :--- | :--- | :--- | :--- |
| LIST | 41424307 | ("ABC") | 4A4B4C 13 | ("JKL") |
|  | 4A4B4C 13 | ("JKL") | $4 A 4 B 4137$ | ("JKA") |
|  | 4A4B41 37 | ("JKA") | $444 B 413 F$ | ("DKA") |
|  | 444 B 413 F | ("DKA") | 41424337 | ("ABC") |

### 13.2.5 Using a Jump Table with a Key

Using the value in the variable INDEX as a key to a jump table (TABLE). Each entry in the jump table contains a 32 -bit identifier followed by a 32-bit address to which the program should transfer control if the key is equal to that identifier.

Sample Problem:

| INDEX | 00000010 |  |
| :--- | :--- | :--- |
| TABLE | 00000001 | Proc1 |
|  | 00000010 | Proc2 |
|  | 0000001 E | Proc3 |
| Proc1 | NOP |  |
| Proc2 | NOP |  |
| Proc3 | NOP |  |

Control should be transfered to Proc2, the second entry in the table.

14 The Stack

## 15 Subroutines

None of the examples that we have shown thus far is a typical program that would stand by itself. Most real programs perform a series of tasks, many of which may be used a number of times or be common to other programs.

The standard method of producing programs which can be used in this manner is to write subroutines that perform particular tasks. The resulting sequences of instructions can be written once, tested once, and then used repeatedly.

There are special instructions for transferring control to subroutines and restoring control to the main program. We often refer to the special instruction that transfers control to a subroutine as Call, Jump, or Brach to a Subroutine. The special instruction that restores control to the main program is usually called Return.

In the ARM the Branch-and-Link instruction (BL) is used to Branch to a Subroutine. This saves the current value of the program counter (PC or R15) in the Link Register (LR or R14) before placing the starting address of the subroutine in the program counter. The ARM does not have a standard Return from Subroutine instruction like other processors, rather the programmer should copy the value in the Link Register into the Program Counter in order to return to the instruction after the Branch-and-Link instruction. Thus, to return from a subroutine you should the instruction:

MOV
PC, LR
Should the subroutine wish to call another subroutine it will have to save the value of the Link Register before calling the nested subroutine.

### 15.1 Types of Subroutines

Sometimes a subroutine must have special characteristics.

## Relocatable

The code can be placed anywhere in memory. You can use such a subroutine easily, regardless of other programs or the arrangement of the memory. A relocating loader is necessary to place the program in memory properly; the loader will start the program after other programs and will add the starting address or relocation constant to all addresses in the program.

## Position Independent

The code does not require a relocating loader - all program addresses are expressed relative to the program counter's current value. Data addresses are held in-registers at all times. We will discuss the writing of position independent code later in this chapter.

## Reentrant

The subroutine can be interrupted and called by the interrupting program, giving the correct results for both the interrupting and interrupted programs. Reentrant subroutines are required for good for event based systems such as a multitasking operating system (Windows or Unix) and embedded real time environments. It is not difficult to make a subroutine reentrant. The only requirement is that the subroutine uses just registers and the stack for its data storage, and the subroutine is self contained in that it does not use any value defined outside of the routine (global values).

## Recursive

The subroutine can call itself. Such a subroutine clearly must also be reentrant.

### 15.2 Subroutine Documentation

Most programs consist of a main program and several subroutines. This is useful as you can use known prewritten routines when available and you can debug and test the other subroutines properly and remember their exact effects on registers and memory locations.
You should provide sufficient documentation such that users need not examine the subroutine's internal structure. Among necessary specifications are:

- A description of the purpose of the subroutine
- A list of input and output parameters
- Registers and memory locations used
- A sample case, perhaps including a sample calling sequence

The subroutine will be easy to use if you follow these guidelines.

### 15.3 Parameter Passing Techniques

In order to be really useful, a subroutine must be general. For example, a subroutine that can perform only a specialized task, such as looking for a particular letter in an input string of fixed length, will not be very useful. If, on the other hand, the subroutine can look for any letter, in strings of any length, it will be far more helpful.
In order to provide subroutines with this flexibility, it is necessary to provide them with the ability to receive various kinds of information. We call data or addresses that we provide the subroutine parameters. An important part of writing subroutines is providing for transferring the parameters to the subroutine. This process is called Parameter Passing.

There are three general approaches to passing parameters:

1. Place the parameters in registers.
2. Place the parameters in a block of memory.
3. Transfer the parameters and results on the hardware stack.

The registers often provide a fast, convenient way of passing parameters and returning results. The limitations of this method are that it cannot be expanded beyond the number of available registers; it often results in unforeseen side effects; and it lacks generality.

The trade-off here is between fast execution time and a more general approach. Such a trade-off is common in computer applications at all levels. General approaches are easy to learn and consistent; they can be automated through the use of macros. On the other hand, approaches that take advantage of the specific features of a particular task require less time and memory. The choice of one approach over the other depends on your application, but you should take the general approach (saving programming time and simplifying documentation and maintenance) unless time or memory constraints force you to do otherwise.

### 15.3.1 Passing Parameters In Registers

The first and simplest method of passing parameters to a subroutine is via the registers. After calling a subroutine, the calling program can load memory addresses, counters, and other data into registers. For example, suppose a subroutine operates on two data buffers of equal length. The subroutine might specify that the length of the two data buffers be in the register RO while the staring address of the two data buffer are in the registers R1 and R2. The calling program would then call the subroutine as follows:

```
MOV RO, #BufferLen ; Length of Buffer in RO
LDR R1, =BufferA ; Buffer A beginning address in R1
LDR R2, =BufferB ; Buffer B beginning address in R2
BL Subr ; Call subroutine
```

Using this method of parameter passing, the subroutine can simply assume that the parameters are there. Results can also be returned in registers, or the addresses of locations for results can be passed as parameters via the registers. Of course, this technique is limited by the number of registers available.
Processor features such as register indirect addressing, indexed addressing, and the ability to use any register as a stack pointer allow far more powerful and general ways of passing parameters.

### 15.3.2 Passing Parameters In A Parameter Block

Parameters that are to be passed to a subroutine can also be placed into memory in a parameter block. The location of this parameter block can be passed to the subroutine via a register.

```
LDR RO, =Params ; RO Points to Parameter Block
BL Subr ; Call the subroutine
```

If you place the parameter block immediately after the subroutine call the address of the parameter block is automatically place into the Link Register by the Branch and Link instruction. The subroutine must modify the return address in the Link Register in addition to fetching the parameters. Using this technique, our example would be modified as follows:

| BL | Subr |  |
| :--- | :--- | :--- |
| DCD | BufferLen | ;Buffer Length |
| DCD | BufferA | ;Buffer A starting address |
| DCD | BufferB | ;Buffer B starting address |

The subroutine saves' prior contents of CPU registers, then loads parameters and adjusts the return address as follows:

```
Subr LDR RO, [LR], #4 ; Read BuufferLen
    LDR R1, [LR], #4 ; Read address of Buffer A
    LDR R2, [LR], #4 ; Read address of Buffer B
    ; LR points to next instruction
```

The addressing mode [LR], \#4 will read the value at the address pointed to by the Link Register and then move the register on by four bytes. Thus at the end of this sequence the value of LR has been updated to point to the next instruction after the parameter block.
This parameter passing technique has the advantage of being easy to read. It has, however, the disadvantage of requiring parameters to be fixed when the program is written. Passing the address of the parameter block in via a register allows the papa meters to be changed as the program is running.

### 15.3.3 Passing Parameters On The Stack

Another common method of passing parameters to a subroutine is to push the parameters onto the stack. Using this parameter passing technique, the subroutine call illustrated above would occur as follows:

| MOV | RO, \#BufferLen | ; Read Buffer Length |
| :--- | :--- | :--- |
| STR | RO, [SP, \#-4]! | ; Save on the stack |
| LDR | RO, =BufferA | ; Read Address of Buffer A |
| STR RO, [SP, \#-4]! | ; Save on the stack |  |
| LDR | RO, =BufferA | ; Read Address of Buffer B |
| STR | RO, [SP, \#-4]! | ; Save on the stack |
| BL | Subr |  |

The subroutine must begin by loading parameters into CPU registers as follows:


In this approach, all parameters are passed and results are returned on the stack.
The stack grows downward (toward lower addresses). This occurs because elements are pushed onto the stack using the pre-decrement address mode. The use of the pre-decrement mode causes the stack pointer to always contain the address of the last occupied location, rather than the next empty one as on some other microprocessors. This implies that you must initialise the stack pointer to a value higher than the largest address in the stack area.

When passing parameters on the stack, the programmer must implement this approach as follows:

1. Decrement the system stack pointer to make room for parameters on the system stack, and store them using offsets from the stack pointer, or simply push the parameters on the stack.
2. Access the parameters by means of offsets from the system stack pointer.
3. Store the results on the stack by means of offsets from the systems stack pointer.
4. Clean up the stack before or after returning from the subroutine, so that the parameters are removed and the results are handled appropriately.

### 15.4 Types Of Parameters

Regardless of our approach to passing parameters, we can specify the parameters in a variety of ways. For example, we can:

## pass-by-value

Where the actual values are placed in the parameter list. The name comes from the fact that it is only the value of the parameter that is passed into the subroutine rather than the parameter itself. This is the method used by most high level programming languages.

## pass-by-reference

The address of the parameters are placed in the parameter list. The subroutine can access the value directly rather than a copy of the parameter. This is much more dangerous as the subroutine can change a value you don't want it to.

## pass-by-name

Rather than passing either the value or a reference to the value a string containing the name of the parameter is passed. This is used by very high level languages or scripting languages. This is very flexible but rather time consuming as we need to look up the value associated with the variable name every time we wish to access the variable.

### 15.5 Program Examples

```
Program 15.1a: init1.s - Initiate a simple stack
    * initiate a simple stack
        TTL Ch10Ex1
        AREA Program, CODE, READONLY
        ENTRY
    Main
        LDR R1, Value1 ;put some data into registers
        LDR R2, Value2
        LDR R3, Value3
        LDR R4, Value4
        LDR R7, =Data2
        ;load the top of stack
        STMFD R7,{R1 - R4} ;push the data onto the stack
        SWI &11
        ;all done
```

```
AREA Stack1, DATA
Value1 DCD OxFFFF
Value2 DCD 0xDDDD
Value3 DCD 0xAAAA
Value4 DCD 0x3333
Stack % AREA Data2, DATA 40 ;reserve 40 bytes of memory for the stack
StackEnd
    DCD 0
    END
```

Program 15.1b: init2.s - Initiate a simple stack

Program 15.1c: init3.s - Initiate a simple stack


```
\begin{tabular}{|c|c|c|c|}
\hline & SWI & \&11 & ;all done \\
\hline & AREA & Data1, DATA & \\
\hline Value1 & DCD & 0xFFFF & \multirow[t]{4}{*}{;some data to put on stack} \\
\hline Value2 & DCD & \(0 \times\) DDD & \\
\hline Value3 & DCD & OxAAAA & \\
\hline Value4 & DCD & \(0 \times 3333\) & \\
\hline & AREA & Data2, DATA & \\
\hline & - & StackStart & ;reserve 40 bytes of memory for the stack \\
\hline Stack1 & DCD & 0 & \\
\hline
\end{tabular}
    END
```

Program 15.1d: init3a.s - Initiate a simple stack


Program 15.1e: byreg.s - A simple subroutine example program passes a variable to the routine in a register

```
* a simple subroutine example
    * program passes a variable to the routine in a register
        TTL Ch10Ex4
        AREA Program, CODE, READONLY
        ENTRY
    StackStart EQU 0x9000
    Main
        LDRB RO, HDigit ;variable stored to register
        BL Hexdigit ;branch/link
        STRB RO, AChar ;store the result of the subroutine
    SWI &O ;output to console
    SWI &11 ;all done
```

```
* =========================
* Hexdigit subroutine
* =========================
* Purpose
* Hexdigit subroutine converts a Hex digit to an ASCII character
* Initial Condition
* RO contains a value in the range 00 ... OF
* Final Condition
* RO contains ASCII character in the range 'O' ... '9' or 'A' ... 'F'
* 
* Registers changed
* RO only
* Sample case
* Sample case 
* Final condition RO = 36 ('6')
Hexdigit
    CMP RO, #OxA ;is it > 9
    BLE Addz ;if not skip the next
    ADD RO, RO, #"A" - "0" - 0xA ;adjust for A .. F
Addz
    ADD RO, RO, #"O" ;convert to ASCII
    MOV PC, LR ;return from subroutine
    AREA Data1, DATA
HDigit DCB 6 ;digit to convert
AChar DCB 0 ;storage for ASCII character
    END
```

Program 15.1f: bystack.s - A more complex subroutine example program passes variables to the routine using the stack


```
* Second parameter is the address of the string
* Final Condition
* the HEX string occupies 4 bytes beginning with
* the address passed as the second parameter
* Registers changed
* No registers are affected
*
* Sample case
* Initial condition top of stack : 4CDO
* Address of string
* Final condition The string '4,'C',D',O' in ASCII
* occupies memory
Binhex
    MOV R8, R7 ;save stack pointer for later
    STMDA R7, {RO-R6,R14} ;push contents of RO to R6, and LR onto the stack
    MOV R1, #4 ;init counter
    ADD R7, R7, #4 ;adjust pointer
    LDR R2, [R7], #4 ;get the number
    LDR R4, [R7] ;get the address of the string
    ADD R4, R4, #4 ;move past the end of where the string is to be stored
Loop
    MOV RO, R2 ;copy the number
    AND RO, RO, #Mask ;get the low nibble
    BL Hexdigit ;convert to ASCII
    STRB RO, [R4], #-1 ;store it
    MOV R2, R2, LSR #4 ;shift to next nibble
    SUBS R1, R1, #1 ;decr counter
    BNE Loop ;loop while still elements left
    LDMDA R8, {R0-R6,R14} ;restore the registers
    MOV PC, LR ;return from subroutine
* ==========================
* Hexdigit subroutine
* ==========================
* Purpose
* Hexdigit subroutine converts a Hex digit to an ASCII character
* Initial Condition
* RO contains a value in the range 00 ... OF
* Final Condition
* RO contains ASCII character in the range 'O' ... '9' or 'A' ... 'F'
* Registers changed
* RO only
* Sample case
* Initial condition RO = 6
* Final condition RO=36 ('6')
Hexdigit
    CMP RO, #0xA ;is the number 0 ... 9?
    BLE Addz ;if so, branch
    ADD RO, RO, #"A" - "O" - OxA ;adjust for A ... F
Addz
    ADD RO, RO, #"O" ;change to ASCII
    MOV PC, LR ;return from subroutine
    AREA Data1, DATA
    DCD &4CDO
;number to convert
```



Program 15.1h: factorial.s - $A$ subroutine to find the factorial of a number

```
* a subroutine to find the factorial of a number
    TTL Ch10Ex6
    AREA Program, CODE, READONLY
    ENTRY
    Main
        LDR RO, Number ;get number
        BL Factor ;branch/link
        STR RO, FNum ;store the factorial
        SWI &11 ;all done
    * =========================
    * Factor subroutine
    * =========================
    * Purpose
    * Recursively find the factorial of a number
    * Initial Condition
    * RO contains the number to factorial
    * Final Condition
    * RO = factorial of number
    * Registers changed
    * RO and R1 only
    * Sample case
    * Initial condition
    * Number = 5
    * Final condition
    * FNum = 120 = 0x78
    Factor
\begin{tabular}{lll} 
STR & RO, [R12], \#4 & ;push to stack \\
STR & R14, [R12], \#4 & ;push the return address \\
SUBS & RO, RO, \#1 & ;subtract 1 from number
\end{tabular}
    SUBS RO, RO, #1 ;subtract 1 from number
    BNE F_Cont ;not finished
    MOV RO, #1 ;Factorial == 1
    SUB R12, R12, #4 ;adjust stack pointer
    B Return ;done
    F_Cont
        BL Factor ;if not done, call again
    Return
        LDR R14, [R12], #-4 ;return address
        LDR R1, [R12], #-4 ;load to R1 (can't do MUL R0, R0, xxx)
        MUL RO, R1, RO ;multiply the result
        MOV PC, LR ;and return
        AREA Data1, DATA
    Number DCD 5 ;number
    FNum DCD 0 ; ; factorial
```


### 15.6 Problems

Write both a calling program for the sample problem and at least one properly documented subroutine for each problem.

### 15.6.1 ASCII Hex to Binary

Write a subroutine to convert the least significant eight bits in register R0 from the ASCII representation of a hexadecimal digit to the 4 -bit binary representation of the digit. Place the result back into R0.

Sample Problems:

|  |  | Test $A$ | Test $B$ |
| :---: | :---: | :---: | :---: |
| Input: | R0 | 43 'C' | $36^{\prime} 6$ ' |
| Output: | R0 | 0C | 06 |

### 15.6.2 ASCII Hex String to Binary Word

Write a subroutine that takes the address of a string of eight ASCII characters in RO. It should convert the hexadecimal string into a 32 -bit binary number, which it return is RO.
Sample Problem:

| Input: | R0 | String |
| ---: | :--- | :--- |
|  | STRING | $42^{\prime} \mathrm{B}^{\prime}$ |
|  |  | $32^{\prime} 2^{\prime}$ |
|  |  | $46^{\prime} \mathrm{F} \prime$ |
| Output: | R0 | 30 |

### 15.6.3 Test for Alphabetic Character

Write a subroutine that checks the character in register RO to see if it is alphabetic (upper- or lower-case). It should set the Zero flag if the character is alphabetic, and reset the flag if it is not.
Sample Problems:

|  |  | Test $A$ | Test B | Test $C$ |
| :---: | :---: | :---: | :---: | :---: |
| Input: | Ro | 47 'G' | 36 '6' | $6 \mathrm{~A}{ }^{\prime} \mathrm{j}$ ' |
| Output: | Z | FF | 00 | FF |

### 15.6.4 Scan to Next Non-alphabetic

Write a subroutien that takes the address of the start of a text string in register R1 and returns the address of the first non-alphabetic character in the string in register R1. You should consider using the isalpha subroutine you have just define.

Sample Problems:

| Input: | R1 String | Test A | B |
| :---: | :---: | :---: | :---: |
|  |  | String | 6100 |
|  |  | 43 'C' | 32 '2' |
|  |  | 61 'a' | 50 'P' |
|  |  | 74 't' | 49 'I' |
|  |  | 0D CR | 0D CR |
| Output: | R1 | $\begin{aligned} & \text { String }+4 \\ & \text { (CR) } \end{aligned}$ | $\text { String }+0$ <br> (2) |

### 15.6.5 Check Even Parity

Write a subroutine that takes the address of a counted string in the register RO. It should check for an even number of set bits in each character of the string. If all the bytes have an even parity then it should set the Z-flag, if one or more bytes have an odd parity it should clear the Z-flag.

Sample Problems:

|  |  | Test $A$ |  | Test $B$ |
| :---: | :--- | :--- | :--- | :--- |
| Input: | RO | String |  | String |
|  | String | 03 |  | 03 |
|  |  | 47 |  | 47 |
|  |  | AF |  | AF |
| Output: | Z | 18 |  | 19 |
| Onnn | 00 | FF |  |  |

Note that $19_{16}$ is $00011001_{2}$ which has three 1 bits and is thus has an odd parity.

### 15.6.6 Check the Checksum of a String

Write a subroutine to calculate the 8 -bit checksum of the counted string pointed to by the register R0 and compares the calculated checksum with the 8 -bit checksum at the end of the string. It should set the Z-flag if the checksums are equal, and reset the flag if they are not.
Sample Problems:

| Input: | R0 String |  | Test A | Test B |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | String |  | String |  |
|  |  | 03 | (Length) | 03 | (Length) |
|  |  | 41 | ('A') | 61 | ('a') |
|  |  | 42 | ('B') | 62 | ('b') |
|  |  | 43 | ('C') | 63 | ('c') |
|  |  | C6 | (Checksum) | C6 | (Checksum should be $26</$ em>) |
| Output: | Z | Set |  | Clear |  |

### 15.6.7 Compare Two Counted Strings

Write a subroutine to compare two ASCII strings. The first byte in each string is its length. Return the result in the condition codes; i.e., the N -flag will be set if the first string is lexically less than (prior to) the second, the Z-flag will be set if the strings are equal, no flags are set if the second is prior to the first. Note that " ABCD " is lexically greater than " ABC ".

## 16 Interrupts and Exceptions

## A ARM Instruction Definitions

This appendix describes every ARM instruction, in terms of:

## Operation

A Register Transfer Language (RTL) / pseudo-code description of what the instruction does. For details of the register transfer language, see section ?? on page ??.

## Syntax

$<\mathrm{cc}>$ Condition Codes <op1> Data Movement Addressing Modes <op2> Memory Addressing Modes $<$ S $>$ Set Flags bit

## Description

Written description of what the instruction does. This will interpret the formal description given in the operation part. It will also describe any additional notations used in the Syntax part.

## Exceptions

This gives details of which exceptions can occur during the instruction. Prefetch Abort is not listed because it can occur for any instruction.

## Usage

Suggestions and other information relating to how an instruction can be used effectively.

## Condition Codes

Indicates what happens to the CPU Condition Code Flags if the set flags option where to be set.
Notes
Contain any additional explanation that we can not fit into the previous categories.
Appendix B provides a summary of the more common instructions in a more compact manner, using the operation section only.
ADC Add with Carry

| Operation | $\langle c c\rangle: \mathrm{Rd}$ <br> $\langle c c\rangle\langle S\rangle: \mathrm{CPSR} \leftarrow \mathrm{R} n+\langle o p 1\rangle$ <br> Syntax |
| :--- | :--- |
| $\mathrm{ADC}\langle c c\rangle\langle S\rangle \mathrm{Rd} d, \mathrm{Rn},\langle o p 1\rangle$ |  |
| Description | The ADC (Add with Carry) instruction adds the value of $\langle o p 1\rangle$ and the Carry flag to the <br> value of Rn and stores the result in Rd. The condition code flags are optionally updated, |
|  | based on the result. |

```
ADDS R4,RO,R2
ADC R5,R1,R3
```

If the second instruction is changed from:

```
ADC R5,R1,R3
```

to:
ADCS R5,R1,R3
the resulting values of the flags indicate:
$\mathbf{N}$ The 64-bit addition produced a negative result.
C An unsigned overflow occurred.
V A signed overflow occurred.
Z The most significant 32 bits are all zero.
The following instruction produces a single-bit Rotate Left with Extend operation (33-bit rotate through the Carry flag) on R0:
ADCS RO,RO,RO

See Data-processing operands - Rotate right with extend for information on how to perform a similar rotation to the right.

## Condition Codes

The N and Z flags are set according to the result of the addition, and the C and V flags are set according to whether the addition generated a carry (unsigned overflow) and a signed overflow, respectively.
ADD Add

Operation $\langle c c\rangle: \mathrm{Rd} \leftarrow \mathrm{R} n+\langle o p 1\rangle$
$\langle c c\rangle\langle S\rangle: \mathrm{CPSR} \leftarrow \mathrm{ALU}($ Flags $)$
Syntax $\quad \operatorname{ADD}\langle c c\rangle\langle S\rangle \mathrm{Rd}, \mathrm{Rn},\langle o p 1\rangle$
Description Adds the value of $\langle o p 1\rangle$ to the value of register Rn , and stores the result in the destination register Rd. The condition code flags are optionally updated, based on the result.
Usage The ADD instruction is used to add two values together to produce a third.
To increment a register value in $\mathrm{R} x$ use:

$$
\mathrm{ADD} \quad \mathrm{Rx}, \mathrm{Rx}, \# 1
$$

Constant multiplication of $\mathrm{R} x$ by $2^{n}+1$ into Rd can be performed with:
ADD Rd, Rx, Rx, LSL \#n

To form a PC-relative address use:
ADD Rs, PC, \#offset
where the $\langle$ offset $\rangle$ must be the difference between the required address and the address held in the PC, where the PC is the address of the ADD instruction itself plus 8 bytes.

## Condition Codes

The N and Z flags are set according to the result of the addition, and the C and V flags are set according to whether the addition generated a carry (unsigned overflow) and a signed overflow, respectively.

## AND Bitwise AND

Operation $\quad\langle c c\rangle: \mathrm{Rd} \leftarrow \mathrm{R} n \wedge\langle o p 1\rangle$
$\langle c c\rangle\langle S\rangle: \mathrm{CPSR} \leftarrow \mathrm{ALU}($ Flags $)$
Syntax $\quad \operatorname{AND}\langle c c\rangle\langle S\rangle \mathrm{R} d, \mathrm{R} n,\langle o p 1\rangle$
Description The AND instruction performs a bitwise AND of the value of register $R n$ with the value of $\langle o p 1\rangle$, and stores the result in the destination register Rd . The condition code flags are optionally updated, based on the result.

Usage AND is most useful for extracting a field from a register, by ANDing the register with a mask value that has 1 s in the field to be extracted, and 0 s elsewhere.

## Condition Codes

The N and Z flags are set according to the result of the operation, and the C flag is set to the carry output generated by $\langle o p 1\rangle$ (see 5.1 on page 45 The V flag is unaffected.
B, BL Branch, Branch and Link

Operation $\quad\langle c c\rangle\langle L\rangle: \mathrm{LR} \leftarrow \mathrm{PC}+8$
$\langle c c\rangle: \mathrm{PC} \leftarrow \mathrm{PC}+\langle$ offset $\rangle$
Syntax $\quad \mathrm{B}\langle L\rangle\langle c c\rangle\langle$ offs $e t\rangle$
Description The B (Branch) and BL (Branch and Link) instructions cause a branch to a target address, and provide both conditional and unconditional changes to program flow.
The BL (Branch and Link) instruction stores a return address in the link register (LR or R14).
The $\langle$ offset $\rangle$ specifies the target address of the branch. The address of the next instruction is calculated by adding the offset to the program counter (PC) which contains the address of the branch instruction plus 8 .
The branch instructions can specify a branch of approximately $\pm 32 \mathrm{MB}$.
Usage The BL instruction is used to perform a subroutine call. The return from subroutine is achieved by copying the LR to the PC. Typically, this is done by one of the following methods:

- Executing a MOV PC, R14 instruction.
- Storing a group of registers and R14 to the stack on subroutine entry, using an instruction of the form:

STMFD R13!, \{ $\langle$ registers $\rangle$,R14\}
and then restoring the register values and returning with an instruction of the form:

$$
\text { LDMFD R13!, \{〈registers }\rangle, \mathrm{PC}\}
$$

## Condition Codes

The condition codes are not effected by this instruction.
Notes Branching backwards past location zero and forwards over the end of the 32-bit address space is UNPREDICTABLE.
CMP $\quad$ Compare

Operation $\quad\langle c c\rangle: \operatorname{ALU}(0) \leftarrow \mathrm{Rn}-\langle o p 1\rangle$
$\langle c c\rangle:$ CSPR $\leftarrow \mathrm{ALU}($ Flags $)$
Syntax CMP $\langle c c\rangle \operatorname{Rn},\langle o p 1\rangle$
Description The CMP (Compare) instruction compares a register value with another arithmetic value. The condition flags are updated, based on the result of subtracting $\langle o p 1\rangle$ from Rn, so that subsequent instructions can be conditionally executed.

## Condition Codes

The N and Z flags are set according to the result of the subtraction, and the C and V flags are set according to whether the subtraction generated a borrow (unsinged underflow) and a signed overflow, respectively.

| EOR | Exclusive OR |
| :---: | :---: |
| Operation | $\begin{gathered} \langle c c\rangle: \mathrm{Rd} \\ \langle c c\rangle\langle S\rangle: \mathrm{CPSR} \oplus \mathrm{Rn} \oplus\langle\mathrm{op} 1\rangle \\ \leftarrow \mathrm{ALU}(\mathrm{Flags}) \end{gathered}$ |
| Syntax | $\mathrm{EOR}\langle c c\rangle\langle S\rangle \mathrm{Rd}$, Rn , $\langle o p 1\rangle$ |
| Description | The EOR (Exclusive OR) instruction performs a bitwise Exclusive-OR of the value of register $\mathrm{R} n$ with the value of $\langle o p 1\rangle$, and stores the result in the destination register Rd . The condition code flags are optionally updated, based on the result. |
| Usage | EOR can be used to invert selected bits in a register. For each bit, EOR with 1 inverts that bit, and EOR with 0 leaves it unchanged. |
| Condition Codes |  |
|  | The N and Z flags are set according to the result of the operation, and the C flag is set to the carry output bit generated by the shifter. The V flag is unaffected. |

## LDM Load Multiple

Operation if $\langle c c\rangle$
IA: addr $\leftarrow \mathrm{Rn}$
IB: addr $\leftarrow \mathrm{Rn}+4$
DA: addr $\leftarrow \operatorname{Rn}-(\#\langle$ registers $\rangle * 4)+4$
DB: addr $\leftarrow \mathrm{Rn}-(\#\langle$ registers $\rangle * 4)$
for each register Ri in $\langle$ registers $\rangle$
IB: addr $\leftarrow$ addr +4
DB: addr $\leftarrow$ addr - 4
$\mathrm{Ri} \leftarrow \mathrm{M}$ (addr)
IA: addr $\leftarrow$ addr +4
DA: addr $\leftarrow$ addr -1

$$
\langle!\rangle: \operatorname{Rn} \quad \leftarrow \operatorname{addr}
$$

Syntax LDM $\langle c c\rangle\langle$ mode $\rangle \operatorname{Rn}\langle!\rangle,\{\langle$ registers $\rangle\}$
Description The LDM (Load Multiple) instruction is useful for block loads, stack operations and procedure exit sequences. It loads a subset, or possibly all, of the general-purpose registers from sequential memory locations.
The general-purpose registers loaded can include the PC. If they do, the word loaded for the PC is treated as an address and a branch occurs to that address.
The register Rn points to the memory local to load the values from. Each of the registers listed in $\langle$ registers $\rangle$ is loaded in turn, reading each value from the next memory address as directed by $\langle$ mode $\rangle$, one of:

| IB | Increment Before |
| :--- | :--- |
| DB | Decrement Before |
| IA | Increment After |
| DA | Decrement After |

The base register writeback option ( $\langle!\rangle$ ) causes the base register to be modified to hold the address of the final valued loaded.
The register are loaded in sequence, the lowest-numbered register from the lowest memory address, through to the highest-numbered register from the highest memory address.
If the PC (R15) is specified in the register list, the instruction causes a branch to the address loaded into the PC.
Exceptions Data Abort

## Condition Codes

The condition codes are not effected by this instruction.

Notes If the base register Rn is specified in $\langle$ registers $\rangle$, and base register writeback is specified ( $\langle!\rangle$ ), the final value of $R n$ is UNPREDICTABLE.
LDR $\quad$ Load Register

Operation $\quad\langle c c\rangle: \mathrm{Rd} \leftarrow \mathrm{M}(\langle o p 2\rangle)$
Syntax LDR$\langle c c\rangle \operatorname{Rd},\langle o p 2\rangle$
Description The LDR (Load Register) instruction loads a word from the memory address calculated by $\langle o p 1\rangle$ and writes it to register Rd.
If the PC is specified as register Rd, the instruction loads a data word which it treats as an address, then branches to that address.

## Exceptions Data Abort

Usage Using the PC as the base register allows PC-relative addressing, which facilitates positionindependent code. Combined with a suitable addressing mode, LDR allows 32-bit memory data to be loaded into a general-purpose register where its value can be manipulated. If the destination register is the PC, this instruction loads a 32 -bit address from memory and branches to that address.
To synthesize a Branch with Link, precede the LDR instruction with MOV LR, PC.

## Condition Codes

The condition codes are not effected by this instruction.
Notes If $\langle o p 2\rangle$ specifies an address that is not word-aligned, the instruction attempts to load a byte. The result is UNPREDICTABLE and the LDRB instruction should be used.
If $\langle o p 2\rangle$ specifies base register writeback (!), and the same register is specified for Rd and Rn, the results are UNPREDICTABLE.
If the PC (R15) is specified for Rd, the value must be word alligned otherwise the result is UNPREDICTABLE.
LDRB Load Register Byte

Operation $\quad$|  | $\langle c c\rangle: \operatorname{Rd}(7: 0) \leftarrow \mathrm{M}(\langle o p 2\rangle)$ |
| :--- | :--- |
|  | $\langle c c\rangle: \operatorname{Rd}(31: 8) \leftarrow 0$ |

Syntax $\quad \operatorname{LDR}\langle c c\rangle \mathrm{BR} d,\langle o p 2\rangle$
Description The LDRB (Load Register Byte) instruction loads a byte from the memory address calculated by $\langle o p 2\rangle$, zero-extends the byte to a 32 -bit word, and writes the word to register Rd.
Exceptions Data Abort
Usage LDRB allows 8-bit memory data to be loaded into a general-purpose register where it can be manipulated.
Using the PC as the base register allows PC-relative addressing, to facilitate positionindependent code.

## Condition Codes

The condition codes are not effected by this instruction.
Notes If the PC (R15) is specified for Rd, the result is UNPREDICTABLE.
If $\langle o p 2\rangle$ specifies base register writeback (!), and the same register is specified for Rd and Rn, the results are UNPREDICTABLE.
MOV $\quad$ Move $\quad 1$

Operation $\quad\langle c c\rangle: \operatorname{Rd} \quad \leftarrow\langle o p 1\rangle$
$\langle c c\rangle\langle S\rangle: \mathrm{CPSR} \leftarrow \mathrm{ALU}($ Flags $)$

Syntax
$\operatorname{MOV}\langle c c\rangle\langle S\rangle \operatorname{Rd},\langle o p 1\rangle$
Description The MOV (Move) instruction moves the value of $\langle o p 1\rangle$ to the destination register Rd. The condition code flags are optionally updated, based on the result.
Usage MOV is used to:

- Move a value from one register to another.
- Put a constant value into a register.
- Perform a shift without any other arithmetic or logical operation. A left shift by $n$ can be used to multiply by $2^{n}$.
- When the PC is the destination of the instruction, a branch occurs. The instruction:

$$
\text { MOV } \mathrm{PC}, \mathrm{LR}
$$

can therefore be used to return from a subroutine (see instructions $B$, and $B L$ on page 129.

## Condition Codes

The N and Z flags are set according to the value moved (post-shift if a shift is specified), and the C flag is set to the carry output bit generated by the shifter (see 5.1 on page 45 ). The V flag is unaffected.
MVN $\quad$ Move Negative

Operation $\quad\langle c c\rangle: \mathrm{Rd} \quad \leftarrow \overline{\langle o p 1\rangle}$
$\langle c c\rangle\langle S\rangle: \mathrm{CPSR} \leftarrow \mathrm{ALU}($ Flags $)$
Syntax $\quad \operatorname{mVN}\langle c c\rangle\langle S\rangle \operatorname{Rd},\langle o p 1\rangle$
Description The MVN (Move Negative) instruction moves the logical one's complement of the value of $\langle o p 1\rangle$ to the destination register Rd. The condition code flags are optionally updated, based on the result.
Usage MVN is used to:

- Write a negative value into a register.
- Form a bit mask.
- Take the one's complement of a value.


## Condition Codes

The N and Z flags are set according to the result of the operation, and the C flag is set to the carry output bit generated by the shifter (see 5.1 on page 45. The V flag is unaffected.

| ORR | Bitwise OR |
| :---: | :---: |
| Operation | $\begin{array}{cc} \langle c c\rangle: \mathrm{Rd} & \leftarrow \mathrm{R} n \vee\langle o p 1\rangle \\ \langle c c\rangle\langle S\rangle: \mathrm{CPSR} & \leftarrow \mathrm{ALU}(\text { Flags }) \end{array}$ |
| Syntax | $\mathrm{ORR}\langle c c\rangle\langle S\rangle \mathrm{Rd}$, Rn , $\langle o p 1\rangle$ |
| Description | The ORR (Logical OR) instruction performs a bitwise (inclusive) OR of the value of register $\mathrm{R} n$ with the value of $\langle o p 1\rangle$, and stores the result in the destination register Rd . The condition code flags are optionally updated, based on the result. |
| Usage | ORR can be used to set selected bits in a register. For each bit, OR with 1 sets the bit, and OR with 0 leaves it unchanged. |
| Condition Codes |  |
|  | The N and Z flags are set according to the result of the operation, and the C flag is set to the carry output bit generated by the shifter (see 5.1 on page 45 ). The V flag is unaffected. |

## SBC Subtract with Carry

Operation $\quad\langle c c\rangle: \mathrm{Rd} \leftarrow \mathrm{Rn}-\langle o p 1\rangle-\mathrm{NOT}(\operatorname{CPSR}(\mathrm{C}))$
$\langle c c\rangle\langle S\rangle: \mathrm{CPSR} \leftarrow \mathrm{ALU}($ Flags $)$
Syntax $\quad \mathrm{SBC}\langle c c\rangle\langle S\rangle \mathrm{Rd}$, Rn, $\langle o p 1\rangle$
Description The SBC (Subtract with Carry) instruction is used to synthesize multi-word subtraction. SBC subtracts the value of $\langle o p 1\rangle$ and the value of NOT(Carry flag) from the value of register Rn, and stores the result in the destination register Rd. The condition code flags are optionally updated, based on the result.
Usage If register pairs R0,R1 and R2,R3 hold 64-bit values (R0 and R2 hold the least significant words), the following instructions leave the 64-bit difference in R4,R5:

```
SUBS R4,R0,R2
SBC R5,R1,R3
```


## Condition Codes

The N and Z flags are set according to the result of the subtraction, and the C and V flags are set according to whether the subtraction generated a borrow (unsigned underflow) and a signed overflow, respectively.
Notes If $\langle S\rangle$ is specified, the C flag is set to:
0 if no borrow occurs
1 if a borrow does occur
In other words, the C flag is used as a NOT(borrow) flag. This inversion of the borrow condition is usually compensated for by subsequent instructions. For example:

- The SBC and RSC instructions use the C flag as a NOT(borrow) operand, performing a normal subtraction if $\mathrm{C}==1$ and subtracting one more than usual if $\mathrm{C}==0$.
- The HS (unsigned higher or same) and LO (unsigned lower) conditions are equivalent to CS (carry set) and CC (carry clear) respectively.

| STM | Store Multiple |
| :---: | :---: |
| Operation | if $\langle c c\rangle$ |
|  | IA: addr $\leftarrow \mathrm{Rn}$ |
|  | IB: addr $\leftarrow \mathrm{Rn}+4$ |
|  | DA: addr $\leftarrow \mathrm{Rn}-(\#\langle$ registers $\rangle * 4)+4$ |
|  | DB: addr $\leftarrow \mathrm{Rn}-(\#\langle$ registers $\rangle * 4)$ |
|  | for each register Ri in $\langle$ registers $\rangle$ |
|  | IB: addr $\quad \leftarrow$ addr +4 |
|  | DB: addr $\quad \leftarrow$ addr - 4 |
|  | $\mathrm{M}(\mathrm{addr}) \leftarrow \mathrm{Ri}$ |
|  | IA: addr $\leftarrow$ addr +4 |
|  | DA: addr $\quad \leftarrow$ addr - 4 |
|  | $\langle!\rangle: \mathrm{Rn} \quad \leftarrow \mathrm{addr}$ |
| Syntax | $\mathrm{STM}\langle c c\rangle\langle$ mode $\rangle \mathrm{Rn}\langle!\rangle,\{\langle$ registers $\rangle\}$ |

Description The STM (Store Multiple) instruction stores a subset (or possibly all) of the general-purpose registers to sequential memory locations.
The register Rn specifies the base register used to store the registers. Each register given in Rregisters is stored in turn, storing each register in the next memory address as directed by $\langle$ mode $\rangle$, which can be one of:

| IB | Increment Before |
| :--- | :--- |
| DB | Decrement Before |
| IA | Increment After |
| DA | Decrement After |

If the base register writeback option $(\langle!\rangle)$ is specified，the base register（ $\mathrm{R} n$ ）is modified with the new base address．
$\langle r e g i s t e r s\rangle$ is a list of registers，separated by commas and specifies the set of registers to be stored．The registers are stored in sequence，the lowest－numbered register to the lowest memory address，through to the highest－numbered register to the highest memory address． If R15（PC）is specified in $\langle$ registers $\rangle$ ，the value stored is UNKNOWN．

## Exceptions Data Abort

Usage STM is useful as a block store instruction（combined with LDM it allows efficient block copy） and for stack operations．A single STM used in the sequence of a procedure can push the return address and general－purpose register values on to the stack，updating the stack pointer in the process．

## Condition Codes

The condition codes are not effected by this instruction．
Notes If R15（PC）is given as the base register（Rn），the result is UNPREDICTABLE． If $\mathrm{R} n$ is specified as $\langle$ registers $\rangle$ and base register writeback（ $\langle!\rangle$ ）is specified：
－If $\mathrm{R} n$ is the lowest－numbered register specified in $\langle$ registers $\rangle$ ，the original value of $\mathrm{R} n$ is stored．
－Otherwise，the stored value of $\mathrm{R} n$ is UNPREDICTABLE．
The value of $\mathrm{R} n$ should be word alligned．
STR $\quad$ Store Register

Operation $\quad\langle c c\rangle: \mathrm{M}(\langle o p 2\rangle) \leftarrow \mathrm{Rd}$
Syntax $\quad \operatorname{STR}\langle c c\rangle \mathrm{Rd},\langle o p 2\rangle$
Description The STR（Store Register）instruction stores a word from register Rd to the memory address calculated by $\langle o p 2\rangle$ ．

## Exceptions Data Abort

Usage Combined with a suitable addressing mode，STR stores 32－bit data from a general－purpose register into memory．Using the PC as the base register allows PC－relative addressing， which facilitates position－independent code．

## Condition Codes

The condition codes are not effected by this instruction．
Notes Using the PC as the source register（ $R d$ ）will cause an UNKNOWN value to be written． If 〈op2〉 specifies base register writeback（！），and the same register is specified for Rd and Rn ，the results are UNPREDICTABLE．
The address calculated by 〈op2 $\rangle$ must be word－alligned．The result of a store to a non－ word－alligned address is UNPREDICTABLE．
STRB $\quad$ Store Register Byte

Operation $\quad\langle c c\rangle: \mathrm{M}(\langle o p 2\rangle) \leftarrow \operatorname{Rd}(7: 0)$
Syntax $\quad \operatorname{STR}\langle c c\rangle B \operatorname{Rd},\langle o p 2\rangle$
Description The STRB（Store Register Byte）instruction stores a byte from the least significant byte of register Rd to the memory address calculated by $\langle o p 2\rangle$ ．
Exceptions Data Abort

Usage Combined with a suitable addressing mode, STRB writes the least significant byte of a general-purpose register to memory. Using the PC as the base register allows PC-relative addressing, which facilitates position-independent code.

## Condition Codes

The condition codes are not effected by this instruction.
Notes Specifing the PC as the source register (Rd) is UNPREDICTABLE.
If $\langle o p 2\rangle$ specifies base register writeback (!), and the same register is specified for Rd and Rn, the results are UNPREDICTABLE.

## SUB Subtract

Operation $\quad\langle c c\rangle: \mathrm{Rd} \quad \leftarrow \mathrm{Rn}-\langle o p 1\rangle$
$\langle c c\rangle\langle S\rangle: \mathrm{CPSR} \leftarrow \mathrm{ALU}$ (Flags)
Syntax $\quad \operatorname{SUB}\langle c c\rangle\langle S\rangle \operatorname{Rd}, \operatorname{Rn},\langle o p 1\rangle$
Description Subtracts the value of $\langle o p 1\rangle$ from the value of register Rn, and stores the result in the destination register Rd. The condition code flags are optionally updated, based on the result.
Usage SUB is used to subtract one value from another to produce a third. To decrement a register value (in $R x$ ) use:

```
SUBS Rx, Rx, #1
```

SUBS is useful as a loop counter decrement, as the loop branch can test the flags for the appropriate termination condition, without the need for a compare instruction:

```
CMP Rx, #0
```

This both decrements the loop counter in $\mathrm{R} x$ and checks whether it has reached zero.

## Condition Codes

The N and Z flags are set according to the result of the subtraction, and the C and V flags are set according to whether the subtraction generated a borrow (unsigned underflow) and a signed overflow, respectively.
Notes If $\langle S\rangle$ is specified, the C flag is set to:
1 if no borrow occurs
0 if a borrow does occur
In other words, the C flag is used as a NOT(borrow) flag. This inversion of the borrow condition is usually compensated for by subsequent instructions. For example:

- The SBC and RSC instructions use the C flag as a NOT(borrow) operand, performing a normal subtraction if $\mathrm{C}==1$ and subtracting one more than usual if $\mathrm{C}==0$.
- The HS (unsigned higher or same) and LO (unsigned lower) conditions are equivalent to CS (carry set) and CC (carry clear) respectively.
SWI Software Interrupt

| Operation | $\langle c c\rangle:$ R14 svc | $\leftarrow \mathrm{PC}+8$ |
| :--- | :--- | :--- |
|  | $\langle c c\rangle: \operatorname{SPSR}$ svc | $\leftarrow \mathrm{CPSR}$ |
|  | $\langle c c\rangle: \operatorname{CPSR}($ mode $)$ | $\leftarrow$ Supervisor |
|  | $\langle c c\rangle: \operatorname{CPSR}(\mathrm{I})$ | $\leftarrow 1$ (Disable Interrupts) |
|  | $\langle c c\rangle: \mathrm{PC}$ | $\leftarrow 0 \times 00000008$ |

Syntax SWI $\langle c c\rangle\langle$ value $\rangle$
Description Causes a SWI exception (see 3.4 on page 29 .
Exceptions Software interrupt

Usage The SWI instruction is used as an operating system service call. The method used to select which operating system service is required is specified by the operating system, and the SWI exception handler for the operating system determines and provides the requested service. Two typical methods are:

- $\langle v a l u e\rangle$ specifies which service is required, and any parameters needed by the selected service are passed in general-purpose registers.
- $\langle v a l u e\rangle$ is ignored, general-purpose register R 0 is used to select which service is wanted, and any parameters needed by the selected service are passed in other general-purpose registers.


## Condition Codes

The flags will be effected by the operation of the software interrupt. It is not possible to say how they will be effected. The status of the condition code flags is unknown after a software interrupt is UNKNOWN.
SWP Swap

Operation $\quad\langle c c\rangle: \mathrm{ALU}(0) \leftarrow \mathrm{M}(\mathrm{Rn})$
$\langle c c\rangle: \mathrm{M}(\mathrm{R} n) \leftarrow \mathrm{R} m$
$\langle c c\rangle: \mathrm{Rd} \leftarrow \mathrm{ALU}(0)$

## Syntax $\quad \operatorname{SWP}\langle c c\rangle \mathrm{Rd}, \mathrm{Rm},[\mathrm{Rn}]$

Description Swaps a word between registers and memory. SWP loads a word from the memory address given by the value of register $\mathrm{R} n$. The value of register Rm is then stored to the memory address given by the value of $R n$, and the original loaded value is written to register $R d$. If the same register is specified for $R d$ and $R m$, this instruction swaps the value of the register and the value at the memory address.

## Exceptions Data Abort

Usage The SWP instruction can be used to implement semaphores. For sample code, see Semaphore instructions.

## Condition Codes

The condition codes are not effected by this instruction.
Notes If the address contained in Rn is non word-aligned the effect is UNPREDICTABLE.
If the PC is specified as the destination $(R d)$, address ( $R n$ ) or the value ( $R m$ ), the result is UNPREDICTABLE.
If the same register is specified as $\langle R n\rangle$ and $\langle R m\rangle$, or $\langle R n\rangle$ and $\langle R d\rangle$, the result is UNPREDICTABLE.
If a data abort is signaled on either the load access or the store access, the loaded value is not written to $\langle R d\rangle$. If a data abort is signaled on the load access, the store access does not occur.
SWPB $\quad$ Swap Byte

Operation |  | $\langle c c\rangle: \operatorname{ALU}(0) \leftarrow M(\operatorname{Rn})$ |
| :--- | :--- |
|  | $\langle c c\rangle: M(\operatorname{Rn}) \leftarrow \operatorname{Rm}(7: 0)$ |
|  | $\langle c c\rangle: \operatorname{Rd}(7: 0) \leftarrow \operatorname{ALU}(0)$ |

Syntax $\quad \mathrm{SWP}\langle c c\rangle \mathrm{B} d, \mathrm{Rm},[\mathrm{Rn}]$
Description Swaps a byte between registers and memory. SWPB loads a byte from the memory address given by the value of register Rn. The value of the least significant byte of register Rm is stored to the memory address given by Rn, the original loaded value is zero-extended to a 32 -bit word, and the word is written to register Rd. If the same register is specified for Rd and Rm, this instruction swaps the value of the least significant byte of the register and the byte value at the memory address.

Exceptions Data Abort
Usage The SWPB instruction can be used to implement semaphores, in a similar manner to that shown for SWP instructions in Semaphore instructions.

## Condition Codes

The condition codes are not effected by this instruction.
Notes If the PC is specified for Rd, Rn, or Rm, the result is UNPREDICTABLE.
If the same register is specified as $R n$ and $R m$, or $R n$ and $R d$, the result is UNPREDICTABLE.
If a data abort is signaled on either the load access or the store access, the loaded value is not written to $\langle R d\rangle$. If a data abort is signaled on the load access, the store access does not occur.

## B ARM Instruction Summary

| cc: Condition Codes |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | Generic |  | Unsigned |  | Signed |
| CS | Carry Set | HI | Higer Than | GT | Greater Than |
| CC | Carry Clear | HS | Higer or Same | GE | Greater Than or Equal |
| EQ | Equal (Zero Set) | LO | Lower Than | LT | Less Than |
| NE | Not Equal (Zero Clear) | LS | Lower Than or Same | LE | Less Than or Equal |
| VS | Overflow Set |  |  | MI | Minus (Negative) |
| VC | Overflow Clear |  |  | PL | Plus (Positive) |


| op1: Data Access |  |  |
| :---: | :---: | :---: |
| Immediate | \# value ${ }^{\text {l }}$ | $\langle o p 1\rangle \leftarrow \mathrm{IR}($ value) |
| Register | Rm | $\langle o p 1\rangle \leftarrow \mathrm{R} m$ |
| Logical Shift Left Immediate | Rm, LSL \# ${ }^{\text {value }\rangle}$ | $\langle o p 1\rangle \leftarrow \mathrm{Rm} \ll \mathrm{IR}$ (value) |
| Logical Shift Left Register | Rm, LSL Rs | $\langle o p 1\rangle \leftarrow \operatorname{Rm} \ll \operatorname{Rs}(7: 0)$ |
| Logical Shift Right Immediate | Rm, LSR \# $\langle$ value $\rangle$ | $\langle o p 1\rangle \leftarrow \mathrm{R} m \gg \mathrm{IR}$ (value) |
| Logical Shift Right Register | Rm, LSR Rs | $\langle o p 1\rangle \leftarrow \operatorname{Rm} \gg \mathrm{Rs}(7: 0)$ |
| Arithmetic Shift Right Immediate | Rm, ASR \# ${ }^{\text {value }\rangle}$ | $\langle o p 1\rangle \leftarrow \mathrm{R} m+>\mathrm{IR}$ (value) |
| Arithmetic Shift Right Register | Rm, ASR Rs | $\langle o p 1\rangle \leftarrow \mathrm{R} m+>\mathrm{Rs}(7: 0)$ |
| Rotate Right Immediate | Rm, ROR \# ${ }^{\text {value }\rangle}$ | $\langle o p 1\rangle \leftarrow \mathrm{Rm} \ggg$ value $\rangle$ |
| Rotate Right Register | Rm, ROR Rs | $\langle o p 1\rangle \leftarrow \operatorname{Rm} \gg \mathrm{Rs}(4: 0)$ |
| Rotate Right with Extend | Rm, RRX | $\langle o p 1\rangle \leftarrow \mathrm{CPSR}(\mathrm{C})>\mathrm{Rm} \gg \mathrm{CPSR}(\mathrm{C})$ |


| op2: Memory Access |  |  |
| :---: | :---: | :---: |
| Immediate Offset | [Rn, \#土 $\left.{ }^{\text {value }\rangle}\right]$ | $\langle o p 2\rangle \leftarrow \mathrm{R} n+\mathrm{IR}$ (value) |
| Register Offset | [Rn, Rm] | $\langle o p 2\rangle \leftarrow \mathrm{R} n+\mathrm{Rm}$ |
| Scaled Register Offset | [ $\langle$ Rn $\rangle$, Rm, $\langle$ shift $\rangle$ \# $\langle$ value $\rangle$ ] | $\langle o p 2\rangle \leftarrow \mathrm{R} n+(\mathrm{Rm}$ shift $\mathrm{IR}($ value $))$ |
| Immediate Pre-indexed | [Rn, \#土 $\left.{ }^{\text {value }\rangle}\right]$ ! | $\langle o p 2\rangle \leftarrow \mathrm{R} n+\mathrm{IR}$ (value) |
|  |  | $\mathrm{Rn} \leftarrow\langle o p 2\rangle$ |
| Register Pre-indexed | [Rn, Rm]! | $\langle o p 2\rangle \leftarrow \mathrm{R} n+\mathrm{Rm}$ |
|  |  | $\mathrm{Rn} \quad \leftarrow\langle o p 2\rangle$ |
| Scaled Register Pre-indexed |  | $\begin{aligned} \langle o p 2\rangle & \leftarrow \mathrm{Rn}+(\mathrm{Rm} \text { shift IR(value) }) \\ \mathrm{Rn} & \leftarrow\langle o p 2\rangle \end{aligned}$ |
| Immediate Post-indexed | [Rn], \# $\pm\langle$ value $\rangle$ | $\langle o p 2\rangle \leftarrow \mathrm{Rn}$ |
|  |  | $\mathrm{R} n \leftarrow \mathrm{R} n+\mathrm{IR}$ (value) |
| Register Post-indexed | [Rn], Rm | $\langle o p 2\rangle \leftarrow \mathrm{Rn}$ |
|  |  | $\mathrm{R} n \leftarrow \mathrm{R} n+\mathrm{Rm}$ |
| Scaled Register Post-indexed | [Rn], Rm, $\langle$ shift $\rangle$ \# $\langle$ value $\rangle$ | $\langle o p 2\rangle \leftarrow \mathrm{R} n$ |
|  |  | $\mathrm{Rn} \quad \leftarrow \mathrm{R} n+\mathrm{Rm}$ shift IR (value) |

Where $\langle$ shift $\rangle$ is one of: LSL, LSR, ASR, ROR or RRX and has the same effect as for $\langle o p 1\rangle$

| ARM Instructions |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
| Add with Carry | ADC $\langle c c\rangle\langle S\rangle$ | Rd, Rn, $\langle o p 1\rangle$ | $\langle c c\rangle: \mathrm{Rd}$ | $\leftarrow \mathrm{R} n+\langle o p 1\rangle+\operatorname{CPSR}(\mathrm{C})$ |
| Add | $\operatorname{ADD}\langle c c\rangle\langle S\rangle$ | Rd, Rn, $\langle o p 1\rangle$ | $\langle c c\rangle: \mathrm{Rd}$ | $\leftarrow \mathrm{R} n+\langle o p 1\rangle$ |
| Bitwise AND | $\operatorname{AND}\langle c c\rangle\langle S\rangle$ | $\mathrm{Rd}, \mathrm{Rn},\langle o p 1\rangle$ | $\langle c c\rangle: \mathrm{Rd}$ | $\leftarrow \mathrm{R} n \&\langle o p 1\rangle$ |
| Branch | $\mathrm{B}\langle c c\rangle$ | <offset> | $\langle c c\rangle$ : PC | $\leftarrow \mathrm{PC}+\langle$ offset $\rangle$ |
| Branch and Link | BL $\langle c c\rangle$ | <offset> | $\langle c c\rangle$ : LR | $\leftarrow \mathrm{PC}+8$ |
|  |  |  | $\langle c c\rangle$ : PC | $\leftarrow \mathrm{PC}+\langle$ offset $\rangle$ |
| Compare | $\operatorname{CMP}\langle c c\rangle$ | $\mathrm{Rn},\langle o p 1\rangle$ | $\langle c c\rangle$ : CSPR | $\leftarrow(\mathrm{Rn}-\langle o p 1\rangle)$ |
| Exclusive OR | $\mathrm{EOR}\langle c c\rangle\langle S\rangle$ | Rd, Rn, $\langle o p 1\rangle$ | $\langle c c\rangle$ : Rd | $\leftarrow \mathrm{R} n \oplus\langle o p 1\rangle$ |
| Load Register | $\mathrm{LDR}\langle c c\rangle$ | Rd , $\langle o p 2\rangle$ | $\langle c c\rangle: \mathrm{Rd}$ | $\leftarrow \mathrm{M}(\langle o p 2\rangle)$ |
| Load Register Byte | $\mathrm{LDR}\langle c c\rangle$ B | Rd , $\langle o p 2\rangle$ | $\begin{aligned} & \langle c c\rangle: \operatorname{Rd}(7: 0) \\ & \langle c c\rangle: \operatorname{Rd}(31: 8) \end{aligned}$ | $\begin{aligned} & \leftarrow \mathrm{M}(\langle o p 2\rangle) \\ & \leftarrow 0 \end{aligned}$ |
| Move | $\operatorname{MOV}\langle c c\rangle\langle S\rangle$ | Rd , $\langle o p 1\rangle$ | $\langle c c\rangle: \mathrm{Rd}$ | $\leftarrow\langle$ op 1$\rangle$ |
| Move Negative | MVN $\langle c c\rangle\langle S\rangle$ | Rd , $\langle o p 1\rangle$ | $\langle c c\rangle: \mathrm{Rd}$ | $\leftarrow \overline{\langle o p 1\rangle}$ |
| Bitwise OR | $0 \mathrm{RR}\langle c c\rangle\langle S\rangle$ | $\mathrm{Rd}, \mathrm{Rn},\langle o p 1\rangle$ | $\langle c c\rangle: \mathrm{Rd}$ | $\leftarrow \mathrm{Rn} \mid\langle o p 1\rangle$ |
| Subtract with Carry | $\mathrm{SBC}\langle c c\rangle\langle S\rangle$ | Rd , Rn, $\langle o p 1\rangle$ | $\langle c c\rangle: \mathrm{Rd}$ | $\leftarrow \mathrm{R} n-\langle o p 1\rangle-\overline{\operatorname{CPSR}(\mathrm{C})}$ |
| Store Register | $\operatorname{STR}\langle c c\rangle$ | Rd , $\langle o p 2\rangle$ | $\langle c c\rangle: \mathrm{M}(\langle o p 2\rangle)$ | $\leftarrow \mathrm{Rd}$ |
| Store Register Byte | $\operatorname{STR}\langle c c\rangle\langle S\rangle$ | Rd , $\langle o p 2\rangle$ | $\langle c c\rangle: \mathrm{M}(\langle o p 2\rangle)$ | $\leftarrow \operatorname{Rd}(7: 0)$ |
| Subtract | $\operatorname{SUB}\langle c c\rangle\langle S\rangle$ | Rd , Rn, $\langle o p 1\rangle$ | $\langle c c\rangle$ : Rd | $\leftarrow \mathrm{Rn}-\langle o p 1\rangle$ |
| Software Interrupt | SWI $\langle c c\rangle$ | <value〉 |  |  |
| Swap | SWP $\langle c c\rangle$ | Rd, Rm, [Rn] | $\langle c c\rangle: \mathrm{Rd}$ | $\leftarrow \mathrm{M}(\mathrm{Rn})$ |
|  |  |  | $\langle c c\rangle: \mathrm{M}(\mathrm{Rn})$ | $\leftarrow \mathrm{Rm}$ |
| Swap Byte | $\operatorname{SWP}\langle c c\rangle \mathrm{B}$ | Rd , $\mathrm{Rm},[\mathrm{Rn}]$ | $\langle c c\rangle: \operatorname{Rd}(7: 0)$ | $\leftarrow \mathrm{M}(\mathrm{Rn})(7: 0)$ |
|  |  |  | $\langle c c\rangle: \mathrm{M}(\mathrm{R} n)(7: 0) \leftarrow \mathrm{Rm}(7: 0)$ |  |

## Index


Terminated................................. .81


[^0]:    ${ }^{1}$ Although the ARM does allow for Half-Word instructions, the emulator we are using does not.

[^1]:    ${ }^{2}$ This is caused by the processor having already fetched the next instruction from memory while it was deciding what the current instruction was. Thus the PC is still the next instruction to be executed, but that is not the instruction immediately after the current one.

[^2]:    ${ }^{3}$ As the processor has already fetched the instruction after the current instruction it is required to flush the instruction cache and start again. This will cause a short, but not significant, delay.

[^3]:    ${ }^{1}$ http://dec.bournemouth.ac.uk/support/sem/sysarch/examples.zip

[^4]:    ${ }^{2}$ http://www.textpad.com

[^5]:    ${ }^{1}$ Note that we are using a LDR instruction as the data table is sufficently far away from the instruction that an ADR instruction is not valid.

[^6]:    Program 10.1a: strlencr.s - Find the length of a Carage Return terminated string

