Introduction to Saturn Assembly Language

F.H. Gilbert

Eric Rechlin

Authors:

Fernandes Henri Gilbert
1 rue des Gâte Ceps
92210 Saint-Cloud
France
andrell@iname.com
ICQ: 12331170

Eric Rechlin
3212 Winnipeg Drive
Bismarck, ND 58501-0453
USA
eric@hpcalc.org
ICQ: 783944

This book, as well as the included examples, is a free work. You can print this for your own purposes or for any other person. This book may be transmitted or reproduced in any form or by any means as long as it is not modified.

In plain English: This is free, so widely distribute it! Just make sure you give credit to us.

Hewlett-Packard is a trademark of the Hewlett-Packard Company.

Edition 1: January 29, 1999

Copyright 1999 F.H. Gilbert and Eric Rechlin
Distribute freely to all your HP48-owning friends.

Preface

Assembly language is very different from most other programming languages. When learning a typical programming language, one of the first lessons will usually explain how to make a simple "Hello World" program. But with assembly language, it's not that simple. First, before any assembly language code can be written, one must understand the concepts of how the processor works and how the operating system works with it. In addition, many commands must be learned in order to do something as simple as printing to the display.

In this book you will have to go through thirty pages of introductory material before even picking up your calculator. But don't worry; you'll get to start programming soon enough! After being introduced to the concepts, assembly language commands, known as instructions, will be learned in small groups, beginning with an extremely simple program.

Because machine code is binary, you will learn assembly language mnemonics that will represent this binary code. This book is aimed at owners of *all* HP48 calculators, even those with only 32 KB of memory, so a small assembler known as HP-ASM will be used. The syntax of this assembler is quite easy, and because the Meta Kernel is compatible with this syntax, one is not stuck with HP-ASM. In addition, the HP syntax for assembly language, used by Jazz and HP Tools, is very similar, so one should be able to easily switch between the two.

We think many more people could learn assembly language. Coders are needed, and the main roadblock is that information about assembly language coding, particularly Saturn assembly, is too hard to get. Most good Saturn assembly language tutorials are in French, and the only English-language one is James Donnelly's "An Introduction to HP 48 System RPL and Assembly Language Programming," and although it is a great book for System RPL, its assembly language section goes by too quickly for most people. With this book, we hope to fill a gap in the documentation that already exists for the HP48.

Though we are not what one would call experienced coders, we would like to share our knowledge. We still consider ourselves newbies as well, so this is not a definitive introduction to assembly language. If you notice any mistakes, please don't flame us, but instead email us at the addresses below. Some, like Jean-Yves Avenard and Cyrille de Brebisson, have many years of coding experience, but don't be intimidated. Everyone is a newbie at some times.

This book goes from the very basics of how processors work, including tutorials covering the binary and hexadecimal bases, to coding complex routines, such as grayscale display routines.

We spent many, many hours on this book, but if you have problems, feel free to email us and we will help as we can and add to the tutorial if needed. This online version of the book is free, so distribute it widely to all HP48 fans!

In addition, we would like to thank the following people for their contributions to this book (in alphabetical order): Jean-Yves Avenard, Cyrille de Brebisson, Chris Griffiths, Eduardo Kalinowski, John H Meyers, Aldiney Oliveira, Mark Power, and Richard Steventon.

Updates to this book are available at <http://www.hpcalc.org>.

Please, please, all that we ask is for you to send feedback to our email addresses:

F.H. Gilbert <andrell@iname.com>
Eric Rechlin <eric@hpcalc.org>

Table of Contents

PART I: HP48 PROGRAMMING CONCEPTS 1

1 INTRODUCTION TO USER RPL AND SYSTEM RPL 1

2 ASSEMBLY LANGUAGE AND MACHINE LANGUAGE 1

3 THE BINARY BASE AND NIBBLES 2

4 CONVERTING FROM BINARY TO DECIMAL 4

5 MOST SIGNIFICANT BIT AND LEAST SIGNIFICANT BIT 5

6 FRACTIONAL NUMBERS 6

7 CALCULATING USING BASE 2 6

7.1 Adding 6
7.2 Subtracting 7
7.3 Negative numbers 7
7.3.1 The signed binary: the first method 7
7.3.2 The complement of 2: the second method 8
7.4 Multiplying 10
7.5 Dividing 11

8 HEXADECIMAL BASE 12

9 CONVERTING FROM HEXADECIMAL TO DECIMAL 12

10 WRITING GROUPS OF FOUR BITS (NIBBLES) USING HEXADECIMAL 13

11 IS ALL THAT NUMBER CRUNCHING NEEDED? 14

12 THE BCD: "BINARY CODED DECIMAL" 15

13 CONVERTING A DECIMAL TO ITS BCD FORM 16

14 USUAL BCD AND "EXTENDED BCD" 16

15 CALCULATING USING BCD 16

15.1 When a nibble overflows 17
15.2 When a carry needs to be moved 17

16 LOGICAL OPERATIONS 17

16.1 Logical OR 18
16.2 Logical AND 19
16.3 Inverting 19
16.4 NOT OR (also called: NOR) 19
16.5 AND NOT (also called: NAND) 19
16.6 EXCLUSIVE OR (also called: XOR) 20

17 REVIEW: NUMERICAL BASES 20

18 REVIEW: CALCULATIONS 20

PART II: PROGRAMMING TOOLS 23

19 AN ASSEMBLER: HP-ASM 23

20 HP-ASM COMMANDS 23

21 ASSEMBLING CODE WITH HP-ASM 24

22 ADDITIONAL HP-ASM FEATURES 25

22.1 The linker 25
22.2 Directives 25
22.3 Macros 25

23 A FAST TEXT EDITOR 25

PART III: THE SATURN PROCESSOR 27

24 SATURN REGISTERS 27

24.1 Processor registers 27
24.1.1 Pointer registers 27
24.1.2 Counting bytes, KB 28
24.1.3 Addressable space 29
24.1.4 Working registers 29
24.1.5 Save registers 29

25 THE FIELDS OF WORKING REGISTERS 30

25.1 Nibble inversion 30
25.2 Available fields 31
25.2.1 The W field 31
25.2.2 The A field 31
25.2.3 The B field 31
25.2.4 The X field 32
25.2.5 The XS field 32
25.2.6 The M field 32
25.2.7 The S field 32
25.2.8 Field summary 32

26 THE P REGISTER AND WP 33

27 CIRCULAR HANDLING OF NUMBERS INSIDE THE PROCESSOR 34

28 RSTK 35

29 FLAGS: ST, HST, AND CARRY 35

29.1 Carry 35
29.2 ST (Status Bits) 36
29.3 HST (Hardware Status Bits) 36

30 OUT AND IN 36

31 HOW THE PROCESSOR WORKS: LIFE OF AN INSTRUCTION 36

32 WHAT IS A PROCESSOR CYCLE? 37

PART IV: SATURN INSTRUCTION SET 39

33 WORKING REGISTER INSTRUCTIONS 39

33.1 Loading a value inside a register 39
33.1.1 LA, LC 40
33.1.2 Setting a register to zero 41

33.1.3 Changing the value of a bit ... 43
33.1.4 Exchanging two registers ... 45
33.1.5 Copying a register into another ... 46
33.2 Working with registers ... 47
33.2.1 Incrementing a register ... 47
33.2.2 Something special with the carry in loops ... 48
33.2.3 Adding two registers ... 49
33.2.4 Adding a register to itself ... 49
33.2.5 Adding a constant to a register ... 49
33.2.6 Decrementing a register ... 51
33.2.7 Subtracting two registers ... 51
33.2.8 Subtracting a constant ... 52
33.2.9 Complement of 2 ... 52
33.2.10 Complement of 1 ... 52
33.2.11 Logical operations ... 53
33.2.12 Shifting one nibble left or right ... 54
33.2.13 Rotating one nibble left or right ... 54
33.2.14 Shifting one bit to the right ... 55

34 SAVE REGISTERS ... 56

34.1 Saving a working register into a save register ... 56
34.2 Saving a working register into a save register within a field ... 56
34.3 Copying a save register to a working register ... 57
34.4 Copying a save register to a working register within a field ... 57
34.5 Exchanging between save and working register ... 57
34.6 Exchange between save/working register with field ... 58

35 POINTERS ... 58

35.1 Giving D0 or D1 a value ... 58
35.2 Adding to or subtracting from D0 or D1 ... 59
35.3 Copying A (or C) to D0 (or D1) ... 59
35.4 Exchanging field A of A/C with D0/D1 ... 60
35.5 Exchanging 4 LSB nibbles A of A/C with D0/D1 ... 64

36 READING MEMORY WITHIN A FIELD ... 64

37 READING FROM MEMORY WITH A VALUE ... 65

38 WRITING TO MEMORY WITHIN A FIELD ... 66

39 WRITING TO MEMORY WITH A VALUE ... 66

40 JUMPING AND TESTS ... 66

41 RELATIVE, ABSOLUTE, AND INDIRECT JUMPS ... 68

41.1 Conditional jumps: GOC and GONC ... 68
41.2 Unconditional jumps: GOTO, GOLONG, GOVLNG ... 69
41.2.1 GOTO ... 69
41.2.2 GOLONG ... 69
41.2.3 GOVLNG ... 69
41.2.4 Unconditional jump to the address in A or C ... 69
41.2.5 Unconditional jump with exchange ... 70
41.2.6 Unconditional and indirect jump ... 70
41.2.7 Saving PC contents ... 70

42 CALLING A SUBROUTINE ... 70

42.1 GOSUB ... 71
42.2 GOSUBL ... 71
42.3 GOSBVL ... 71

42.4 Returning from a subroutine ... 71
42.5 Returning according to the value of the carry bit ... 72

43 THE TESTS ... 72

43.1 Comparing registers to zero ... 72
43.1.1 Equal to zero? ... 72
43.1.2 Different than zero? ... 73
43.2 Equality of two registers ... 73
43.3 Inequality of two registers ... 73
43.4 Less than and greater than ... 73
43.5 Less than or equal to or greater than or equal to ... 74
43.6 Testing a bit ... 74

44 USING REGISTER P ... 75

44.1 P= *n* ... 75
44.2 P=P+1 ... 75
44.3 P=P-1 ... 75
44.4 ?P# *n* ... 75
44.5 ?P= *n* ... 75
44.6 C=P *n* ... 76
44.7 P=C *n* ... 76
44.8 C+P+1 ... 76
44.9 CPEX *n* ... 76

45 THE RSTK STACK ... 76

46 REGISTERS IN AND OUT ... 76

47 STATUS BITS (ST) ... 77

47.1 CLRST ... 77
47.2 C=ST ... 77
47.3 ST=C ... 77
47.4 CSTEX ... 77
47.5 ST=1 *n* ... 78
47.6 ST=0 *n* ... 78
47.7 ?ST=0 *n* and ?ST=1 *n* ... 78

48 HARDWARE STATUS BITS (HST) ... 78

49 SATURN DEC/HEX MODE ... 79

50 INTERRUPTIONS ... 80

51 BUS-RELATED INSTRUCTIONS ... 80

51.1 RESET ... 80
51.2 SREQ? ... 81
51.3 CONFIG ... 81
51.4 UNCNFG ... 81
51.5 C=ID ... 81
51.6 SHUTDN ... 81
51.7 BUSC*x* ... 81

PART V: HP48 OBJECTS ... 83

52 SIMPLE OBJECTS ... 84

52.1 Real number ... 84
52.2 String ... 84
52.3 Global name ... 87
52.3.1 The empty global name and the hidden directory ... 89

52.4 Local name 90
52.5 User binary integer 90
52.6 Graphic object 91
52.7 XLIB name 92
52.8 System binary integer 93
52.9 Long real 93
52.10 Character 93
52.11 Code object 94
52.12 Library data 94
52.13 Extended pointer 94

53 COMPOUND OBJECTS 95

53.1 Complex number 95
53.2 Array 95
53.3 List 97
53.4 RPL program 97
53.5 Algebraic expression 98
53.6 Tagged object 99
53.7 Unit object 99
53.8 Directory object 100
53.8.1 HOME directory 100
53.8.2 Subdirectories 102
53.9 Library 103
53.9.1 Hash table 104
53.9.2 Message table 105
53.9.3 Link table 106
53.10 Backup 106
53.11 Long complex 106
53.12 Linked array 107

54 PRIMITIVE CODE OBJECT 107

PART VI: WRITING PROGRAMS 109

55 DOING LOOPS 109

56 READING FROM THE KEYBOARD 110

56.1 Checking for a single key-press 110
56.2 Checking for key combinations 111
56.3 Reading the entire keyboard 111
56.4 Reading the [ON] key 112

57 MANIPULATING THE STACK 113

57.1 Drop 115
57.2 Swap 115

58 CREATING A PROGRAM FOR CREATING MACROS 116

58.1 →MACRO.S 117
58.2 ROM routines used in this program 119
58.2.1 #5023h, also called ERROR_A 119
58.2.2 #679Bh, also called SAVE_REG 119
58.2.3 #3019h, also called SKIPOB (Skip OBject) 119
58.2.4 #05B7Dh, also called RES_STR (REServe STRing) 119
58.2.5 #0670Ch, also known as COPY DOWN 120
58.2.6 #067D2h, also called GET_REG 120

59 MEMORY MANAGEMENT 121

60 GRAPHICS 121

Part I:
HP48 Programming Concepts

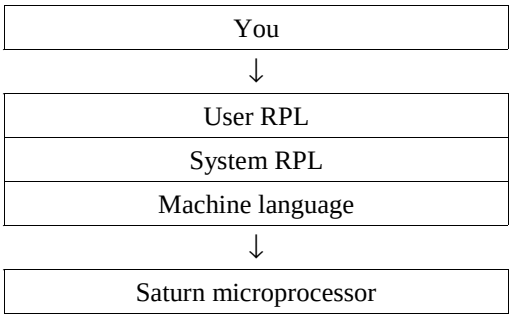
1 Introduction to User RPL and System RPL

The basic user of the HP48 can use User RPL. It's easy to learn, and it's the first language you can use to program your HP.

Each time you use a User RPL command it checks the arguments, so it's a rather safe language. You should not lose any data when using it.

If we look at the HP's internal software, we'll discover that each User RPL command is in fact divided into several simpler commands, each one of which does a single (usually) and simple thing. User RPL is divided internally into elementary commands, which we call System RPL. Each one of these System RPL commands usually ends up calling machine code, which is executed directly by the HP's processor, the Saturn.

To help you understand, look at this:



The HP48 gives you three development languages. User RPL is safe and easy, but the slowest of all. System RPL is faster because no checks are made, so you must take care of all argument at the beginning, but then you no longer need them as you feed data to hungry commands. Finally, we have assembly language, which produces machine language. This is the fastest code, directly executed by the processor, and of course it gives total control over the HP.

If you want to learn System RPL, knowing User RPL will be enough. All you will need is a complete list of all System RPL commands, called entry points, inside the ROM.

Assembly language is not related to User RPL or System RPL, so you're going to learn something really different here. Assembler doesn't have all the commands which are available to the System RPL coder.

It is not a complex language, as its commands are simple and do very elementary things; what **may** seem complicated is that you will need to write a lot of "elementary commands" to do more complicated things than just crashing your HP. ;)

What you'll win is the fastest code around, as well as lack of sleep, but that's being a coder.

2 Assembly language and machine language

The Saturn processor, like any other processor, uses a series of 0's and 1's to work. We are not able, because of our inferior brains*, to understand these series.

* Mr. Avenard is good because he has a Saturn processor surgically implanted inside his brain. Another option is to eat a lot of Saturn processors, like Mika did, but it's quite difficult to eat just one, trust me! :o)

So, we will group these series of 0's and 1's together to create small packets of information. The Saturn processor has been designed to manage this binary information using mostly four-bit packets, each one called a nibble.

We will learn more about both the binary base and the hexadecimal base later, and hexadecimal is the one you will mainly use on computers. Don't worry. :)

So, machine language is a series of binary digits (0 or 1), or bits (the word resulting from the contraction of "binary" and "digit"), which are kept together to form packets of four bits, or nibbles.

Internally, the Saturn processor is only able to do very simple things, like reading some nibbles in memory, writing nibbles, and performing calculations with the nibbles read. Each elementary instruction is a kind of command to us. Assembly language coders use something called mnemonics. What are they?

"Mnemonic" is the name given to each elementary command. It's something very simple, so it reminds us what it's used for (mnemosis → memory, so "mnemonic").

We will write our programs using assembly language, which is composed of mnemonics. There is a mnemonic for every command the Saturn can use. A compiler will translate these mnemonics into binary information, and these bits will be used by the processor to obey our orders!

Remember that:

- Machine language is the *only* language the Saturn processor is able to understand. It's too complicated because it's a series a bits (0 and 1) so we don't use it directly.

- We will then translate those to mnemonics, which are simple words. One exists for each command available on the Saturn processor.

It is important to not mix "assembly language" and "machine language." The latter is the resulting form of the former after a compiler has translated it.

That is why assembly language is called a "low-level" language. In assembly language, each mnemonic corresponds to one instruction. That is why assembly language is **specific** to each processor. When we write assembly language programs, we first write mnemonics in a source file (source because it's the source of our work). Then we'll use a program called an "assembler" or "compiler" that will translate mnemonics to the machine language, a series of 0's and 1's (binary digits), which are kept together using four-bit packets on the HP.

3 The binary base and nibbles

We are going to do some math. Don't be afraid: even if you are bad at math you can do assembly language anyway. :)

As you learned, machine language is composed of series of zeroes and ones. That is because processors work internally with components that let electrons go through or not. In the computer world, 1 usually means an electrical current exists and 0 usually means one does not exist.

So, the Saturn processor keeps the bits as four-bit packets, which are called nibbles. This is something you must *remember*. The traditional spelling is "nybble;" however, the spelling "nibble" has become more common in recent years so we will spell it with an 'i' for the purpose of this book.

Let's start from something we know: the decimal base. We have ten digits so we can easily use the decimal base.

The decimal base is not the only one. Sailors use the sexagesimal base, base 60, because they use hours, minutes, and seconds.

When you type the decimal number 10 on your HP, it displays "10." In fact, the HP internally translates it to bits, and the Saturn processor works with them, thus under the binary base. (I will teach you later that the Saturn processor is able to work with both our decimal base and its binary base, so it's a very good processor isn't it?!)

Decimal base means we use 10 symbols to write numbers:

0 1 2 3 4 5 6 7 8 9

Under the binary base, we only have **two** symbols:

0 1

So, zero in binary base is zero, one is one, and two?

Well, if we use the symbol "2" under the decimal base to represent the number two, we will use "10" under the binary base to represent the number two. In other words, we are going to combine zeroes and ones to represent numbers using the binary base.

The binary base is linked very closely to Boolean algebra, which will be briefly discussed later.

So, here is a small table of decimal numbers and their corresponding form in binary base:

Decimal base	Binary base
0	0
1	1
2	10
3	11
4	100
5	101
6	110
7	111
8	1000
9	1001
10	1010
…	…

You must look at the way the 0's and 1's are combined to form numbers, and then be able to continue. We will discuss this in detail over the next few pages, as it is important to learn. It's not something easy, so don't feel bad if you have trouble understanding.

Now, we are going to kept bits as packets of four bits, because that's the way the Saturn likes.

Four packet bits are called nibbles, so a nibble is composed of four bits. Assembly language coders have some special words when they speak about bits and the different sizes used to group bits.

When four bits are packed together, it is called a "word of four bits." If 32 bits are packed together, it becomes a "word of 32 bits."

How many values can a packet of four bits have? It's easy: we use a power of two to calculate:

$2^4 = 16$

So four bits packed together create a nibble, and:

A nibble can have 16 different values.

Let's calculate another one: how many values can a word of 64 bits have?

$2^{64} = 1.8 * 10^{19}$

A lot. :-)

4 Converting from binary to decimal

We are going to examine two nibbles kept together, which is called a "byte."

4 bits → a nibble
8 bits → a byte

We are going to give each bit a number, from right to left, and will call that the "rank" of a bit. Here is an example: I have two nibbles together (a byte) which are: 01100101, so we have:

rank:	7	6	5	4	3	2	1	0
bits:	0	1	1	0	0	1	0	1

As you see, I have started the "rank" by zero, and then I count from right to left. Each bit has a rank, a specific rank. Each bit has also something special.

The rank is used to know which "weight" the bit has. What does that mean? It means that each bit is, according to its rank, an exponent of base 2. Let me explain; look at this:

rank:	7	6	5	4	3	2	1	0
bits:	0	1	1	0	0	1	0	1
weight:	2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0

What I call weight is the value of the bit in the decimal base. Here is an example: what is the binary value 00100110 in decimal mode? Let's write it:

rank:	7	6	5	4	3	2	1	0
bits:	0	0	1	0	0	1	1	0
weight:	2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0

Now, the "weight" of each bit can be calculated, and we'll find the decimal value! Let's do it:

We are going to multiply each bit by its weight, and we add each together, that is:

$(0 * 2^7) + (0 * 2^6) + (1 * 2^5) + (0 * 2^4) + (0 * 2^3) + (1 * 2^2) + (1 * 2^1) + (0 * 2^0)$

that is: 0 + 0 + 32 + 0 + 0 + 4 + 2 + 0 = 38

Use your HP to check: press [MTH] [BASE] [BIN] and then type the number: #00100110b. Now, press [DEC] and you get:

#38d

It works!! It's cool, isn't it? :-)

Now, you are able to convert a binary number to a decimal one. You write them down, you give each one a rank, from right to left and from zero to the last one. Then, you write the "weight" of each bit, and you can calculate the decimal value.

What is the maximum value of a nibble? Two nibbles?

With two nibbles we have:

rank:	7	6	5	4	3	2	1	0
bits:	1	1	1	1	1	1	1	1
weight:	2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0

That is:

$(1 * 2^7) + (1 * 2^6) + (1 * 2^5) + (1 * 2^4) + (1 * 2^3) + (1 * 2^2) + (1 * 2^1) + (1 * 2^0)$

That is:

$128 + 64 + 32 + 16 + 8 + 4 + 2 + 1 = 255$

So a byte (two nibbles) can go from 0 to 255, that is 256 different values.

Remember!! 0 (zero) is a value too!

And with a nibble we have:

rank:	3	2	1	0
bits:	1	1	1	1
weight:	2^3	2^2	2^1	2^0

That is:

$(1 * 2^3) + (1 * 2^2) + (1 * 2^1) + (1 * 2^0)$

That is:

$8 + 4 + 2 + 1 = 15$

So with a nibble we can go from 0 to 15, that is 16 different values.

To get the decimal form of a binary integer, we multiply each bit's value by $2^{bitrank}$ and we add all values.

5 Most significant bit and least significant bit

This is something important, so don't skip it.

As I have shown you, each bit has a weight. The bit that is the rightmost has a rank of zero, and the bit that is the leftmost has the highest rank (rank n-1 when we have n bits).

This means that each bit has a different weight in the final value.

The bit that is rightmost is called the Least Significant Bit, or LSB.
The bit that is leftmost is called the Most Significant Bit, or MSB.

In the nibble, the LSB and MSB are there:

rank:	3	2	1	0
	MSB			LSB

In the byte, the LSB and MSB are there:

rank: 7 6 5 4 3 2 1 0
MSB LSB

6 Fractional numbers

This is a little trickier.

A fractional number has two parts: the integer part before the fraction mark (a decimal point or comma, depending on region), and the fractional part after the fraction mark.

But…how can we have fractional numbers using bits?

We are going to give each bit a negative exponent weight.

Instead of:

rank: 3 2 1 0
bits: 1 1 1 1
weight: 2^3 2^2 2^1 2^0

we have for fractional numbers:

rank: 3 2 1 0
bits: 1 1 1 1
weight: 2^{-3} 2^{-2} 2^{-1} 2^{-0}

Please note: (2^{-3}) is in fact: ($1 / 2^3$)
(2^{-2}) is in fact: ($1 / 2^2$)

What is the value of 0.0110 in decimal?

rank: 3 2 1 0
bits: 0 1 1 0
weight: 2^{-3} 2^{-2} 2^{-1} 2^{-0}

That is:

$(0 * 2^{-3}) + (1 * 2^{-2}) + (1 * 2^{-1}) + (0 * 2^{-0})$

That is:

$0 + 0.25 + 0.5 + 0 = 0.75$

7 Calculating using base 2

In this section you'll learn how to add, subtract, multiply, and divide two binary values. This is going to help you understand how the Saturn works. :-)

7.1 Adding

You will follow these rules:

0 + 0 = 0 → no carry
0 + 1 = 1 → no carry
1 + 0 = 1 → no carry
1 + 1 = 0 → carry

Here is an example:

```
  0101
+ 1101
 -----
 10010
```

When I have 1 + 1, I put a zero, and I move the carry 1 to the left. I work here from right to left, adding two bits each time.

7.2 Subtracting

You will follow these rules:

0 - 0 = 0
0 - 1 = -1 → negative result
1 - 0 = 1
1 - 1 = 0

So here I need to explain to you how a negative number is represented using the binary base.

7.3 Negative numbers

There are various ways to do negative numbers. I am going to explain two that are important for you to know. You'll use either one or another in your own programs: which one you use will vary if you let the processor do calculations for you, or if you want to make it work the way you want, and rule the beast!!

7.3.1 The signed binary: the first method

We are going to use the most significant bit here (now you know why I told you about the MSB and the LSB!).

When the MSB is zero, the value is *positive*. When the MSB is one, the value is *negative*.

And we use all the other bits available to code the value.

This is called a "signed binary," and it's my favorite one (don't ask why or I'll call Mr. Avenard, and he will eat your brain to feed upon!!)

If we have one nibble, a positive zero is: 0000
a negative zero is: 1000

(yes, there is a +0 and a -0 when using signed binaries!)

If you remember what I told you before, a nibble can go from 0 to 15, thus 16 possible values. Though there are still sixteen possible values, here we cannot use 0 to 15 because we use the MSB to tell if it's positive or negative. So a nibble goes:

Binary	Decimal
from 0000 to 0111 when it is positive	from positive zero to 7 when it is positive

from 1000 to 1111 when it is negative	from negative zero to -7 when it is negative

If we use a byte (two nibbles):

Binary	Decimal
from 00000000 to 01111111 when it is positive	from positive 0 to 127 when it is positive
from 10000000 to 11111111 when it is negative	from negative 0 to -127 when it is negative

Remember!! We have *two* zeroes here: a negative zero and a positive zero.

7.3.2 The complement of 2: the second method

When all bits in a word (whatever its size) are inverted, we get an inverted word. We call that its "complement of 1." The zeroes become ones, and the ones become zeroes. That is not difficult at all.

0101 becomes 1010
1100 becomes 0011
0010 becomes 1101

But the "complement of 2" is the true complement for a word, and it is the complement of 1 plus one; that is:

(complement of 2) = (complement of 1) + 1

Example: what is the "complement of 2" of 01101110 ?

First, we invert all bits, and get: 10010001

Add one and you'll get:

```
   10010001
+         1
   --------
   10010010
```

So the complement of 2 of 01101110 is 10010010

It is more complicated to handle "complement of 2" binary numbers. Please pay attention; it is not complicated, but you must be awake to understand. :)

We will now learn the way the "complement of 2" numbers are encoded. We will use a byte (8 bits) to learn it.

I have a start value of 0000 0000. If I add one, and again and again, it goes up…

But, what will happen if I have 0000 0000 and I remove 1?

I get:

```
0000 0000 - 1 = 1111 1111
```

As you see, when using "complement of 2," the negative values start with 1 in their MSB. If the MSB is 0, the value is positive.

Interesting thing:

Positive values go from 0000 0000 to 0111 1111, or 0 to 127 in decimal.

Negative values go from 1000 0000 to 1111 1111, or -128 to -1 in decimal.

You see?

Note: here we're using 8 bits, but we could use fewer (like 4 with a nibble) or more (like 64…a lot of bits to play with).

We don't get a negative zero and a positive zero here!!

Sniff! Sniff! :'(

(Sniff! = sound of someone close to crying)

But there is something *much more important*!

We cannot do (1000 0000 - 1) !!

Why? We would go from a *negative* value to a *positive* value.

Also, we cannot do (0111 1111 + 1) because we would go from a *positive* value to a *negative* value, and adding +1 should not (unless we redefine mathematics) give a negative value from a positive value.

:-)

You must understand and know the differences between the signed binaries and the "complement of 2." The first one helps one remember and understand the second one. That's what I do, and it's a good way to remember it, as the first one is very easy to remember, with its +0 and –0. :-)

Using the "complement of 2" is a bit more complicated. Here is an example: how do I calculate 8 - 11 (which are decimal values) using 8 bits and "complement of 2"?

8 is 0000 1000
11 is 0000 1011

To subtract 11, I'm going to add its negative value, so I have:

(8 - 11) becomes (8 + -11)

The "complement of 2" of 0000 1011 is inverting all bits and adding one, that is:

0000 1011 —inverting bits→ 1111 0100

Add one to get: 1111 0101

So I now add the two values:

```
   0000 1000
+  1111 0101
   ---------
   1111 1101
```

As you see, the MSB here equals 1, so the result is known to be negative. But, I cannot use that final value directly. I have to calculate its complement of 2, remember it's negative, and add a - sign in front of it.

The complement of 2 of 1111 1101 is: 0000 0011

That is finally: -3

REMEMBER!

When we use complement of 2 to do calculations, like the one I showed above, the final result *is not directly usable*. We have to:

1. calculate the "absolute value"
2. put the sign according to the MSB value

7.4 Multiplying

I don't like this one (because sometimes I lose myself!) but I have to. To do multiplication, we perform addition and shift bits.

When you want to multiply a value, the slowest of the slowest methods is to add the number to itself as long as needed to get the multiplied value. This will eat up the processor's time ("a baaaaaaaaad thing to do," my friend Mr. Avenard would say)

We can do multiplication using bit-shifting and adding values. Here is an example:

```
      0010
    x 0011
      ----
      0010 (that is 0010 * 1)
   + 0010  (that is 0010 * 1)
  + 0000   (well...)
 + 0000
  --------
  00000110
```

Well, it's not very beautiful, but it works. :-)

As you see, I have done four multiplications. As the two left bits are zero, I could simplify it like below:

```
    10
  * 11
  ----
    10
   10
  ----
   110
```

To multiply, I do:

1. I read the multiplier (the second number) from right to left.
2. When I find a bit set to 1, I copy the multiplicand (the first number).
3. Otherwise, if I find a zero, I insert zeroes. That's easy, isn't it? :)
4. Each time I look at a new bit on the second number, I shift the bits to the left (look above, the bits are written one place left at each line).
5. Finally, I add all values like we know to do.

REMEMBER:	All you have to do is look at the second number that multiplies, bit by bit from right to left, and do shifts; then, we add all values considering carry when there is one.

The Saturn processor is able to do shifts quickly, like adding values, so it's a good way to multiply values.

Shifting a value one bit to the left is like multiplying by 2. Shifting a value one bit to the right is like dividing by 2.

There is a way to optimize multiplication and division. As we have seen, each 0 found requires shifting one bit to the left. To quickly multiply a value, we find out how many zeroes are inside the number and do as many left bit shifts.

7.5 Dividing

Thanks to Eduardo Kalinowski for this section. I still do not understand this Brazilian style of division, so maybe someone else can help out.

OK. Here we go. First, I'd like to add that if you look at the multiplication you'll see that it is exactly like a multiplication in base 10, but much simpler, because you only have to remember four cases (0x0, 0x1, 1x0, 1x1) rather than the 100 of base 10 multiplication. Take a careful look at the process and note that you do the same thing: multiply the first number by the last digit of the second number. Then you multiply by the digit before that, and so on, and in the end you add everything. Well, division is exactly the same.

Let's practice by dividing #10101b (21) by #111b (7). We should get 11 (3).

First, write the numbers like a normal division. I don't know how this is done in the rest of the world, but here in Brazil we write like this:

```
10101 | 111
-------
```

Then, we see how many times the divisor (111) "fits" in the first digit of the dividend. Since it can fit at most one time (base 2), we only need to check IF it fits. Well, 111 doesn't fit in "1," so we write a zero.

```
.
10101 | 111
---------
0
```

Now, I'm going to take two digits of the dividend (see the dot), and see if it fits. No, it doesn't. Another zero is written:

```
.
10101 | 111
----------
00
```

Now I take the first three digits, and I see it doesn't fit again. So, another zero. I take now four digits (1010), and 111 does fit, once. So, I write 1 in the quotient and subtract to see the partial remainder.

```
10101 | 111
- 111 ----------
---- 0001
11
```

Then, we move a digit down:

```
10101 | 111
- 111| ---------
----v 0001
111
```

And divide this (111) by 111, obtaining 1, and a remainder of zero:

```
10101 | 111
- 111 --------
---- 00011
111
- 111
```

```
---
0
```

Since there are no more digits to move down, we've reached the end!!! The quotient is (ignoring leading zeroes) 11, or 3, as we expected, and the remainder 0.

As you have seen (if you haven't read it yet, *do it*), it is equal, but all operations are done in binary base. I've explained this in as much detail as possible.

8 Hexadecimal base

Unlike the decimal base, which uses base 10, and binary, which uses base 2, the hexadecimal base uses base 16. Why are we going to use base 16? As we've seen, the Saturn prefers to keep the bits in packets of four bits, which we call a nibble.

Earlier, we learned that a nibble can have up to sixteen different values. Here they are in their binary form:

Decimal	Binary	Hexadecimal
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
10	1010	A
11	1011	B
12	1100	C
13	1101	D
14	1110	E
15	1111	F

When we use the decimal base, we have ten symbols and combine them to represent numbers. We have also studied the binary base, which combines zeroes and ones to compose numbers.

As the hexadecimal base has 16 symbols, we have to find 16 symbols; we use the first ten ones from 0 to 9, and then we use the first 6 letters of the alphabet: A, B, C, D, E, F

And we are going to combine those digits and letters to represent numbers. What is cool here is that four bits can be represented by a *single* symbol in the hexadecimal mode. :-)

This base is used in *all* computers I know about, or at least systems using similar processors, as some old processors used the eight symbol octal base instead.

9 Converting from hexadecimal to decimal

When we have a number in its hexadecimal form, we can give its numbers "ranks" like we did for the binary numbers. The exponents are just different (and produce bigger numbers).

Let's look at five nibbles under their hexadecimal form. Why five? Because the Saturn processor uses five nibbles to give each nibble in memory a position, five nibbles are required to have a value that can be used as an address, to get information from memory, or to write memory there.

My hex number is #ABCDE and I want to get its decimal value. I would write:

rank:	4	3	2	1	0
hex:	A	B	C	D	E
weight:	16^4	16^3	16^2	16^1	16^0

It's quite easy: as we are using the hexadecimal base, we are going to give each nibble a weight.

The "weight" of a nibble of rank n is: 16^n

Example: what is the hexadecimal number #1AB0F in decimal?

We have here:

rank:	4	3	2	1	0
hex:	1	A	B	0	F
weight:	16^4	16^3	16^2	16^1	16^0

Now, all we have to do is multiply the nibble by 16^{rank}. Here we go:

Value = $(1 * 16^4) + (A * 16^3) + (B * 16^2) + (0 * 16^0) + (F * 16^0)$

Don't be confused because we have letters here rather than only numbers: because we are using the hexadecimal base, we will often find letters in numbers.

If you prefer, $(B * 16^2)$ is in fact: $(11 * 16^4)$

So we calculate values, and get:

```
(1 * 65536) + (10 * 4096) + (11 * 256) + 0 + 15 = 109 327
```

So #1AB0Fh = 109 327

You can check that using your HP. Type the number 109327 and then type [MTH][BASE][R→B]

R→B is a command that will turn a real to a binary integer. You have to make sure you are using hexadecimal notation; there should be a square dot in the [HEX] menu field. If not, press [HEX] so your binary integers are shown as hexadecimal values.

Try using R→B to turn reals to hexadecimal numbers, and then you can use B→R to turn the binary integer numbers back into reals.

10 Writing groups of four bits (nibbles) using hexadecimal

So, how do we write our nibbles using hexadecimal digits? We know a nibble, with four bits, can have 16 different values, and we have 16 symbols with the hex base, so we just to have to convert the four bits to the hex form.

What if there are more than four bits? Example: I have a 12-bit binary number, and I want to convert it to hex.

The trick is to separate bits into groups of four, from right to left.

For example, let's use the number 0111 1011 0101.

To avoid a complex (and boring as math is) calculation, we won't work with all 12 bits at once, but will instead cut them into packets of four bits. Now, we are going to convert each 4-bit packet to its hexadecimal value:

0111 → 7
1011 → B
0101 → 5

(I gave you a table earlier, but as there aren't many symbols, you can write down numbers from 0 to 15, then their binary form, and then their hex value)

So I know the number:

0111 1011 0101 (binary) is #7B5h

This way we are able to easily convert very big binary numbers.

But...what if we don't have a multiple of four bits? No problem: we add enough zeroes to the left side of the number to get a multiple of four bits.

Example: I want to convert the number 011101 from binary to hexadecimal.

I'm going to separate it into groups of four bits, but obviously, there are not enough numbers:

01 1101

I will add two zeroes to the left-most packet in order to have a multiple of four bits in my number:

01 1101 becomes 0001 1101

Because they are leading zeroes, the value has not changed. :-)

Now I can convert each four-bit packet into its hex form:

0001 → 1
1101 → D

So we have 011101 (binary) → #1Dh

11 Is all that number crunching needed?

Of course, we have an HP, so we could use the available functions to convert from binary to hex in any way we want. But it's important to know *how* bits are encoded and used to represent numbers.

When your knowledge and programming skills increase, you will find some logical functions (described in the part after the next part) *very* useful for checking some values, using what we will call "masks."

A "mask" is a set of bits, which we will then use to check values of one or more bits.

I will tell you how a mask can be useful when we start using the HP's keyboard and reading keys from it.

When I first started learning, I didn't like bits and conversions. You may not care for learning them either, because after all, your HP can do that for you, right?

Wrong! You'll just end up more confused than you would normally have been. You're best off learning how to do it in your head, or you'll get it all wrong.

Go back and reread the previous sections if you are at all confused. Even if you are bad at math you'll be able to understand this.

12 The BCD: "Binary Coded Decimal"

This is the part I disliked most when learning. But it's best to know it, especially if you want to pop reals from the stack.

The Saturn processor works with 4-bit packets called nibbles. We have learned that it stores them using a very cool format called hexadecimal. So when we want the processor to add two values, it will consider those as hexadecimal values: if we ask it to add 9 and 9, the Saturn won't give us a result of "18" but rather "12" because 18 = #12h.

But we humans (or whatever you may be ;) prefer the decimal base, and the Saturn processor has been designed so it can *also* work with decimal numbers.

I told you it's a cool processor, and quite powerful, even if it's small and so hard to swallow (gulp! I ate my second one today).

When the Saturn is asked to do calculations using base 10 numbers, it uses a special mode called BCD.

BCD means Binary Coded Decimal

A nibble, which has four bits, can have 16 values. But when we will use the BCD mode, we will keep our numbers inside our decimal-base 10 form.

Pay attention to the way that BCD codes numbers in the following table. Remain calm, and you'll see it's not so difficult, okay? :-)

Decimal	Nibble (four bits)	BCD
0	0000	0000
1	0001	0001
2	0010	0010
3	0011	0011
4	0100	0100
5	0101	0101
6	0110	0110
7	0111	0111
8	1000	1000
9	1001	1001
10	1010	0001 0000
11	1011	0001 0001
12	1100	0001 0010
13	1101	0001 0011
14	1110	0001 0100
15	1111	0001 0101

As you can see here, there is no change between the nibble/bits value and its value with BCD when under ten.

But, when we reach the value of *ten*, the BCD can no longer increase with a single set of four bits, so we must add another set of four bits. We will start over using the new set of four bits, just like we did on the previous set. Notice how we go from 9 to 10 and how we continue up to 15.

When using BCD, each decimal number uses four bits. We use four bits to count from zero to 9, but "10" is written with two digits, so we have to use eight bits to code the value using BCD.

See, it's not so complicated. It may look strange, but this is not the point we have to pay attention to. :o)

13 Converting a decimal to its BCD form

When you want to convert a decimal number to its BCD form, all you have to do is read the number from left to right and write down *each* number in its binary form:

If I have the number 934, what is its BCD form?

9 → 1001
3 → 0011
4 → 0100

So 934 in BCD form is 1001 0011 0100

See? It's not difficult at all. The BCD looks complex but isn't really.

You must be cautious: the BCD binary form is not the same as the binary form.

The number I used for this example below is 934. Here you have its BCD binary form and its binary form:

BCD:	1001 0011 0100
binary:	0011 1010 0110

Don't mix the two together like I did at first (yep!).

14 Usual BCD and "Extended BCD"

The form of BCD I have explained is the usual one, called the "Compacted BCD." There is another form of BCD called the "Extended BCD."

In the Extended BCD, instead of using four bits per digit, eight bits are used. For example, the number 11 would be coded as follows:

Compacted BCD:	0001 0001
Extended BCD:	00000001 00000001

When using Extended BCD, values can go from 0 to 99 by using one byte (two nibbles) whereas the Compacted BCD can only go from 0 to 9!

So the BCD used by the Saturn processor is not only cool, but clever :-))

(now, look at your HP, and tell it you love it...)

15 Calculating using BCD

Storing numbers is cool, but working with them is even cooler! However, because of the way BCD codes numbers, some nibbles are not allowed.

The Saturn will help us a lot, because not only can it store BCD numbers, but it can also perform calculations with them. All we have to do is tell the processor to be in BCD mode and then use it. It's easy! :-)

There are *two* important things when using BCD: when a nibble gets an overflow, and when a carry appears and has to be moved.

15.1 When a nibble overflows

What will happen if I add 16 twice using BCD?

```
  0001 0110       (16 under BCD form)
+ 0001 0110
  ---------
  0010 1100
```

Problem: we cannot have 1100, as it is not allowed when using BCD.

15.2 When a carry needs to be moved

When we add two 1 bits, we have a carry, and we must move one bit to the left and write a zero. For example, 0001 + 0001 = 0010

But, when using BCD, the carry cannot move to the left as one would like.

16 Logical operations

This is an important part. We will meet operations like OR, AND, and others, like XOR. (Julien Meyer, or SunHP, wrote a prayer to the XOR function, as it is useful in games and graphics, and he prays to it each night. ;)

Processors are made using logical operations. I will only discuss simple ones here, but when using much more complex ones, processors are able to do really complex tasks.

Just think about your COS key. When you press it, the RPL operating system checks if there is an argument, then whether it's a valid one, and finally calls one or more machine language operators, which result in the processor moving bits from places to places, with logical operators working over them. Fascinating! :-)

George Boole (1815-1864) published a book in 1847 called "Analyse mathematique de la logique," or something like "Logic Analysis of Mathematics." The most important part of his work is about *logic*. Some of our words are translated into algebraic components, in order to create rules for the thinking processes of calculations. We call this Boolean algebra.

In Boolean algebra, we have two values: 0 and 1. :-)

0 is said to be "false" (the value "false")
1 is said to be "true"

Because we're speaking of logic, we have to understand the concepts of false or true inside the Boolean algebra. Each value gets one purpose: 0 is used for false, and 1 is used for true.

Electronics extensively use this concept of our beloved Boole. Get a drawing of him, along with Newton and Einstein, and put them on your wall! If you don't, you'll never be a good coder. ;)

So, we're going to use the bits 0 and 1

Zero means “false,” and in electronics, it’s found when there is no voltage.

One means “true,” and it’s found when there is voltage in the circuit.

Because we use bits here, we have two values that we are able to compute:

0 is the “complement” of 1
1 is the “complement” of 0

Because the complement is inverting the bits (I’m speaking of the “complement of 1” also called “restricted complement”. It’s not the “complement of 2” we used to store negative values; take a look at the first part if you are confused).

We have three basic logic operations. From *two* of those we are able to create *all* the other ones that exist (or maybe I should say “those we use and have discovered”).

French name	English name	Symbol used
ET	AND	·
OU	OR	+
NON	NOT	-

When a value has been negated using NOT, a line is drawn over its name, like:

NOT A = $\overline{A}$

And that value is called “A-barre” in French. I think it’s called “A-bar” in English.

16.1 Logical OR

It can also be called “logical adding.”

I am going to show you what is called a “truth table.” Each logical operation has one. On the left part we put letters, like A, B, C and so on, for each bit in input. On the right part, we have one S (the result) or several results, when there are several outputs.

The OR is done between *two* bits, first one is A, second one is B, and in S you have the result of the OR between the two bits.

OR Truth Table:		
A	B	S
0	0	0
0	1	1
1	0	1
1	1	1

The logical OR can also be written using an equation:

S = A + B

Note: the logical OR is *not* the addition we saw when we studied binary numbers.

0101 + 0011 = 1000

0101 OR 0011 = 0111

As you now see, OR result is 1 either if A or B are 1 or if both of them are.

16.2 Logical AND

It is also called the "logical product," but it is NOT multiplication.

The equation is written: $A \cdot B = S$

Here is the truth table for AND:

AND Truth Table:		
A	B	S
0	0	0
0	1	0
1	0	0
1	1	1

Here the result is 1 only if both A and B are 1.

16.3 Inverting

We just invert the bit. The equation is written:

$$S = \overline{A}$$

Notice that there is a bar above the letter.

INVERT Truth Table:	
A	S
0	1
1	0

16.4 NOT OR (also called: NOR)

This is using NOT and OR together. When doing a NOR, output is set to 1 only if all inputs are set to zero.

The equation for NOR is:

$$S = \overline{A+B}$$

NOR Truth Table:		
A	B	S
0	0	1
0	1	0
1	0	0
1	1	0

16.5 AND NOT (also called: NAND)

This uses AND and NOT together. When using NAND, it's the inverse of NOR. In other words, the output is always 1 unless both entries are set to 1.

The equation for NAND is:

$S = \overline{A \cdot B}$

NAND Truth Table:		
A	B	S
0	0	1
0	1	1
1	0	1
1	1	0

16.6 EXCLUSIVE OR (also called: XOR)

This one is extensively used when dealing with graphics. The result is 1 if either A or B is one, but *only* if one of A or B is one. If A and B are *both* set to 1, we get zero.

The equation is:

`A XOR B = (A OR B) - (A AND B)` (note that the minus sign is *not* a logical operation)

XOR Truth Table:		
A	B	S
0	0	0
0	1	1
1	0	1
1	1	0

It is cool because a graphic object is, when in memory, a set of bits. When we want to draw a graphic, we can XOR its bits to the actual graphic on screen. Then, if we do an another XOR at the same position, it makes the graphic object disappear! :-)

We'll see that later, so be patient. :-)

17 Review: Numerical bases

Some people said they either didn't read the part on bases or read it too quickly. I understand, because I did the same when I started learning myself. But, a minimum amount of knowledge is required. To be able to recognize hexadecimal and binary forms is mandatory. When I tell you about masks and how to compose them, both binary and hexadecimal information will be required, because we'll write down bits to create our required mask, convert to hexadecimal, and then load it into C or A for using it to test another register's value.

You don't need to know how to convert values from binary to hexadecimal, and such conversion things, but knowing each bit's weight seems of importance to me.

As a quick review, there are four bits in a nibble and each has a weight. No matter how many bits there are, they are numbered from right to left.

18 Review: Calculations

An important thing here is that we have a limit on how many digits we can use. It's not only because we have 16 nibbles (64-bits) in each working register, but also because we will very often use a field which is smaller than W to do calculations.

Because a byte is two nibbles, the greatest value it can have is #FFh.

The use of negative numbers is not very easy, but complement of 2 is available so we can use it. It's the same for fractional numbers; the first book I used to learn assembly language explained (as far as I remember) that assembly didn't let us fractional numbers. In fact, we can either use negative exponents, or simply use a fixed comma position, and other tricks, like Sophie explained (yes, there are girls that code using assembly language on the HP as well!).

The circular handling of numbers must also be understood. Because each field has a limited number of nibbles, when we add or remove constants the value inside can go through the #0h value, setting the carry flag. This is useful for loops: both wait loops that do nothing and loops that are used to repeat pieces of code.

If the X field is used, there are three nibbles to code values. This means one can use values from #000h to #FFFh. What will happen if one adds #1h to #FFFh? The value goes to #000h and the carry bit is set. It also works with #000h - #1h.

When your values go from one number to another and cross #0h, the carry bit is set to warn you an overflow or underflow occurred.

If you have trouble with shifts, just remember the following:

Shifting a bit to the right divides by 2.
Shifting one nibble to the right divides by 16.

Shifting a bit to the left multiplies by 2.
Shifting one nibble to the left multiplies by 16.

Or even easier:

Divide field f by 2: `CSRB f`
Divide field f by 16: `CSR f`

Multiply field f by 2: `C=C+C f`
Multiply field f by 16: `CSL f`

Here instructions have been given for use with the C register, so if you want to use another register you have to change the first letter. As you can see, to multiply a register (or a part of it) by two, you add it to itself.

Sorry if you didn't understand it the first time. :)

Part II: Programming Tools

19 An assembler: HP-ASM

Before you can start writing code, we must discuss some tools that you'll need. In the beginning, the first coders had to write their assembly programs on paper, then they translated the assembly instructions to their hexadecimal form, and from there they created Code object strings.

Today we have better tools. You write your program in a string, and a compiler called an "assembler" will do the job of checking your code (after all, everyone makes mistakes sometimes) and producing the Code object.

I have tried several assemblers. Jazz is interesting but is too big for most users. I want people using a G to be able to code, and this introduction is aimed at everyone, including those using S and G models. Never let **anyone** tell you an SX or GX is needed to code. They're stupid. ;)

Here I will use the HP-ASM 1.1 syntax. This tool is the G/GX adaptation of ASM Flash from Phong Nguyen, who, according to Jean-Yves, was 8 years old at that time of writing it. For the purpose of this document, HP-ASM is a synonym for ASM Flash.

HP-ASM is really fast. Jazz's ASS command is several times slower, and what takes half a second in HP-ASM needs up to one or several using ASS. It's not very crucial, but impressive.

You may get HP-ASM, for the HP48 G, G+, and GX, from this location:

http://www.hpcalc.org/programming/asm/hpasmgx.zip

ASM Flash, for the HP48 S and SX, is available here:

http://www.hpcalc.org/programming/asm/asmflash.zip

The hpasmgx.zip file is a ZIP file, and inside it there are two versions of HP-ASM for the G series. If you have an S or SX, you should get ASM Flash instead.

HP-ASM is available in both versions 1.0 and 1.1. Version 1.0 is 8606 bytes and works quite well. The other version, version 1.1, is 9124 bytes long. Version 1.1 has some new instructions and perhaps other things.

This library contains an assembler that will create code objects from your sources. In this tutorial version 1.1 commands are described, but 1.0 commands are very similar (→ASM instead of ASM, for example). Because the difference in size between 1.0 and 1.1 is only about 600 bytes, you should get 1.1.

If you have a GX with lots of free uncovered RAM, you can choose Jazz, which has many additional features, such as a disassembler, a debugger, and System RPL support. I won't use its syntax here, as it uses a slightly different syntax. People using G models or S models don't have enough space to use it. Because this introduction is not just for GX people but rather for everyone, I'll use HP-ASM syntax.

HP-ASM syntax is compatible with the Meta Kernel's MASD assembler, so this tutorial is valid for the Meta Kernel as well. In addition, MASD has a few *really* cool features for programmers. If you own the Meta Kernel, read its manual for more information.

20 HP-ASM Commands

Here are the HP-ASM commands in its library menu:

ASM: Takes a string as an argument and returns an assembled object (either a code object or a string with hex codes)

ED: Not an editor, but if you make a mistake, by pressing ED, the source string is edited, and you are moved to the beginning of the line where an error was found. When you use ED, it calls the standard editor, the one given by the HP. ED can be modified so it calls the editor you want if you know a little System RPL, but that's beyond the scope of this book.

OPT: Allows options to be selected. When a menu key is pressed, the menu key contents toggle from one option to another.

CODE/HEXA:	If CODE the assembler produces a Code object; if HEXA it pushes a string that contains the hexadecimal values corresponding to the generated code.
0-15/1-16:	When using 0-15, nibbles are numbered from 0 to 15, otherwise from 1-16. This will affect the use of the P register, like `P= n`, `CPEX`, `C=P`, `P=C`, and the bit mnemonics, like `CBIT`, `ABIT`, and tests involving `P`, and `HS`. With assembly language, we start counting at 0, so 0-15 is strongly preferred, meaning the first nibble of a register is number #0h. Advice: keep it set to 0-15 all the time.
HP/PC:	This affects which operator symbols are used. The HP has extended ASCII characters for comparisons that are not compatible with a PC's extended ASCII. Choosing HP uses these HP characters, and choosing PC uses PC-compatible two-character standard ASCII symbols: #, >= and <=. For example, the "not equal to" sign is "=" with a slash through it when HP is set and the pound sign (#) when PC is set.
JMP/UNJ:	When JMP is chosen, the assembler checks all jumps and calculates them. If you select UNJ, then the jumps to labels are not calculated, and thus not verified. It's a good idea to leave this checked.
ON/OFF:	When compiling a very big program, you can turn the display off and the compilation will run about 13% faster. When OFF is selected, the screen will turn off while assembling; ON keeps it on. It's a good idea to keep the screen off when assembling large sources.
EXIT:	Exits the OPT menu.

21 Assembling code with HP-ASM

The Code objects produced by HP-ASM (or any other assembler, like Jazz, J-ASM, ASM Flash, and MASD) are fast because the HP's Saturn processor directly executes them.

We don't code using bits because that would be too difficult. Assembly language uses mnemonics, with one mnemonic for each instruction available on the processor. An assembler like HP-ASM translates these mnemonics into machine language, which is what is found in hexadecimal form inside the Code objects.

The HP48's RPL cannot control what happens when we use Code objects, so we must code well, or otherwise drain our batteries recovering programs we lost (or for some of us, being happy to have one or more RAM cards).

HP-ASM syntax is not HP's syntax, but rather the one used by French programmers for some time. It's like other tools, such as DEV and TC, so it's a part of history. :-)

You will put your code inside a string. Because the internal HP editor is not very fast (and misuses memory like I have never seen) I recommend that you use an external editor. MiniWriter will work very well on G models. It would be nice if Jean-Yves were to release a MiniWriter that uses the 8-pixel font rather than the minifont. That would be the perfect and smallest editor to be used on S/G for use with ASM-Flash (S models) or HP-ASM (G models).

There must be one instruction per line inside the string that you give to the assembler. You can add comments using the percent symbol (%). Everything that follows a `%` until the carriage return is then considered a comment.

The last symbol in the text file should be the "at" sign (@); the assembler will then know this is the end of the source. This symbol is produced by [ALPHA][Right-Shift][ENTER].

The assembler is case-sensitive, which means that “Label” is different than “LABEL”. When you name your labels, either write them all uppercase or all lowercase. Choose a style of writing, and stick to it to avoid confusion.

To go to the next line, do [Right-Shift][.]. (The state of the Alpha mode doesn’t matter, as it’s the same either way.)

22 Additional HP-ASM features

This is a summarized version of the English documentation included with ASM Flash.

22.1 The linker

This is something really cool. If you divide your sources inside several parts, you can ask the compiler to compile each part and then produce the final Code object. If you have two sources, ONE and TWO, you just put inside of your program:

```
'ONE
'TWO
@
```

The ' symbol is used to place the contents of a file in the HOME directory (or the current directory perhaps?) into your source.

If an error is found, it tells where.

22.2 Directives

A directive is the ‘!’ symbol followed by an OPT value; if you want your code to be compiled with the HP option ON, you just put inside of it:

```
!HP
```

By simply inserting a ‘!’ followed by a name from the OPT menu, you can temporarily turn on or off any of the options from your source code string. It can be done anywhere inside of your source: some parts can use PC comparison operators, and other parts HP ones, if you want.

22.3 Macros

Those are not really macros like coders are used to, but they’re cool anyway. You can compile code and reuse it already compiled, or insert a GROB or any kind of data inside of one of your programs. We will need to write a program to create macro-able objects, and it should be useful when we start the graphics part. See below for that program.

23 A fast text editor

The editor built into the HP is not very good at editing big objects, so you should use a replacement editor to edit your ASM sources.

StringWriter (24 KB) is good but is too big for a G, so you may want to use a smaller one, like MiniWriter (5 KB). Because it is smaller, MiniWriter will easily fit on both G and S models.

If you have a GX with plenty of RAM, StringWriter has many more features, but the font is bigger. Those two editors have a cool option: by pressing [MTH] it will give you a list of all lines that start with a *, and HP-ASM considers those lines as labels. To be able to jump to any label of a source is very, very…mandatory :-)

In addition, TED (12 KB) has many of the same features as StringWriter but it uses the small font.

StringWriter can be downloaded here:

http://www.hpcalc.org/apps/editors/strwrt44.zip

MiniWriter can be downloaded here:

http://www.hpcalc.org/apps/editors/minwrt12.zip

TED can be downloaded here:

http://www.hpcalc.org/apps/editors/ted32.zip

Note: MiniWriter uses a small font, and it needs the UFL. This is a library that contains fonts so each program doesn't have to have its own inside, wasting memory space. MiniWriter does not come with the UFL so it will have to be downloaded separately from:

http://www.hpcalc.org/apps/editors/fonts/ufl102.zip

You'll need to download both MiniWriter *and* the UFL library to your HP, or MiniWriter won't work. You don't need the library called "strwadp.lib." Send the following files to your HP:

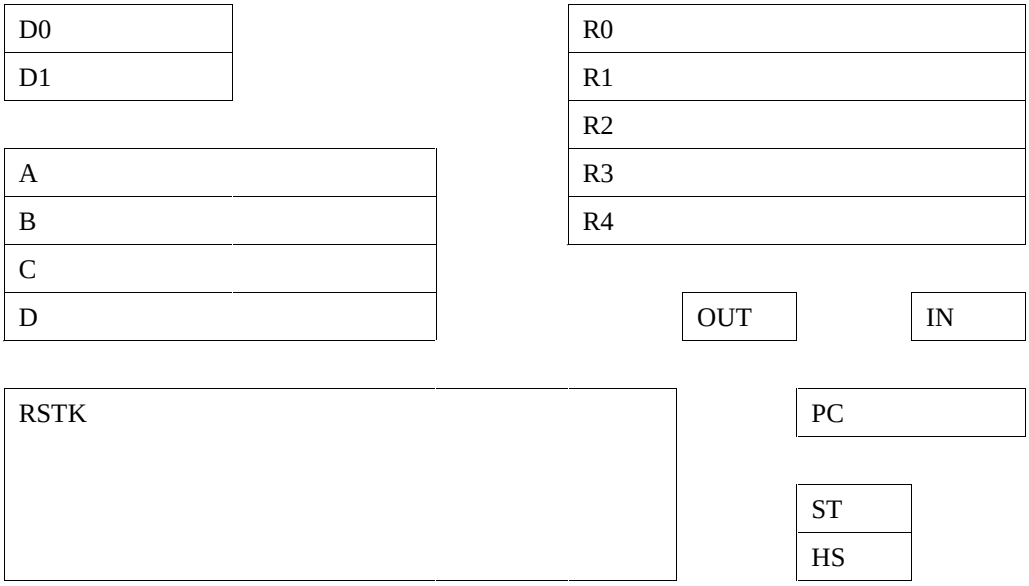
ufl1.lib (found inside the "ufl102.zip" archive)
minwrt12 (found inside the "minwrt12.zip" archive)

I personally prefer StringWriter because the bigger font is easier to read. If you have a G or a S, use MiniWriter. What is important is to have the least memory used. :)

Part III: The Saturn Processor

24 Saturn registers

We can do a kind of drawing of what is inside the Saturn processor:



As you can see above, there are several parts, each inside its own box. Each part has a name, like D0, D1, RSTK, etc.

These are internal components of the Saturn processor, and they are used to make the processor do what you want.

First, where is the processor inside the HP? Not only is it connected to the memory (working upon information stored in memory), but it's connected to the keyboard and the buzzer as well!

So we have two important parts: first, the processor, which is also called the CPU, meaning Central Processing Unit; and second, the memory.

A simple explanation is that an assembly program reads information from memory (or an input device), modifies it with simple operations (typically logical or algebraic), and then writes back to memory or to an output device.

So the work of the Saturn processor is for it to run instructions, which are programs, working with data in memory.

24.1 Processor registers

Inside the processor there are small areas that, like memory, are able to store information in the form of nibbles. There are three kinds of registers: some are called "pointers," some are simply called "registers," and others are called "save registers."

24.1.1 Pointer registers

The Saturn processor has two pointer registers: D0 and D1.

Each pointer register can store 20 bits. If we divide that by 4 (to get four-bit packets, or nibbles) we get 5. This means that, hardware speaking, the Saturn has two pointer registers, and each can receive *five* nibbles to "point" something in memory.

For example, if I want to read a nibble stored at the address ABCDE, I will put that address into D0 or D1 using this assembly instruction:

```
D0= ABCDE        or      D1= ABCDE
```

Now if I read nibbles, I will read them at the memory address ABCDE.

Thus an address is something that has five nibbles, so the Saturn processor can "point" to any nibble in RAM or ROM from #00000h to #FFFFFh, so we have 1 048 576 nibbles (one nibble being four bits). We will learn more about memory later.

There is another pointer register called PC. PC means "Program Counter." It is a very important register: it contains the address in memory of the current instruction being executed. The processor uses it to know *where* it is, and mostly, where it's going. :-)

24.1.2 Counting bytes, KB

Let's divide that value, 1 048 576, by 1024. Why?

In computers and most devices that have memory, we need to say how much "space" is available.

1. You know that four bits form one nibble.
2. Two nibbles together form one byte

The *byte* is a kind of base unit. As time goes by, our computers get more and more memory, so we use exponents to quantify those increasingly huge numbers, but the principle is the same.

If we multiply 2 by 2, and so on, we get these values:

2^1	2
2^2	4
2^3	8
2^4	16
2^5	32
2^6	64
2^7	128
2^8	256
2^9	512
2^{10}	1024

1024 is a very important number for us, because we use the binary base.

The prefix "kilo" means 1000 times the value it precedes. For example, a kilogram is 1000 grams. However, when referring to memory, it's 1024 rather than 1000!

This is *important*: A kilobyte is *not* 1000 times the byte size but rather 1024 times the byte size

So, we have the "byte" (two nibbles). From the byte we can construct the kilobyte, abbreviated KB, which is 1024 bytes.

When you run BYTES, using [Left-Shift][VAR][BYTES], the HP gives the size in bytes and the CRC of the object on level 1 of stack. We will learn more about the CRC and how to calculate it using assembly language later.

To find out how many KB an object uses, run the BYTES command and then divide by 1024 to get the number of KB.

When you have a 48G with 32 KB, you have 32 * 1024 bytes of RAM, which is 32768 bytes. By multiplying that figure by two we can find how many nibbles of RAM are in a 48G: 65536.

24.1.3 Addressable space

How much memory space is the Saturn able to access? We know the processor has two pointer registers: D0 and D1, and each can have a five-nibble value inside. Five nibbles can have 1 048 576 different values. Because the HP48 works with nibbles, we can access 1 048 576 nibbles, or 524 288 bytes. Divide that figure by 1024, and what do you get? Good! 512 KB.

So the Saturn, at any moment of its execution, can *only* manage 512 KB of memory. (Later we'll learn there are tricks to manage more than 512 KB.)

This is called the addressable space.

24.1.4 Working registers

Those registers are sometimes called calculation registers and are where the processor does its work. If you want to add a value from one area of memory to another, you have to:

- read first and second data (if there's a second one) into a working register
- work with it
- do what you want with resulting data

So if you want to add or subtract values, shift bits, perform logical operations, and so on, you have to get the value into a working register. On the Saturn these working registers are called A, B, C, and D.

So we have four working registers. Each register is 64 bits, or 16 nibbles, wide.

You see, the Saturn processor is able to work with 64 bits in a single step, inside a working register! It's a good processor, even though as time flies by, it won't be such a good one in ten years. :-(

24.1.5 Save registers

There are five save registers: R0, R1, R2, R3, and R4.

Those registers are used to keep data inside the processor, freeing one or more working registers so they can get new data. If you want to free working space, it's better to use a save register by moving nibbles from a working register to a save register than to write the nibbles somewhere else in memory. This is because what is kept *inside* the processor is done *faster* than any exchange with memory outside the processor.

Of course, every save register has the same size as any of the working registers: 64 bits, or 16 nibbles. This means that each save *any* of the working registers.

So now we've learned:

- A, B, C, and D are the working registers; nibbles are worked upon inside of those registers.
- We can use one of the five save registers to save one or several working registers for a short period of time (it's a temporary area).
- D0 and D1 are used to point to any address that the Saturn processor can have access to using five nibbles.
- PC is also a register pointer, but it only points to the current instruction (it's updated by the processor as it runs machine language code).
- When either D0 or D1 "points" to an area of memory, we can read or write nibbles from the address pointed to or from a working register.

25 The fields of working registers

We know working registers can have nibbles inside of them. Their size is 64 bits, or 16 nibbles.

Oftentimes we don't need to work with all 16 nibbles of a register. Perhaps we want to work with only five nibbles, two nibbles, or maybe even only one nibble.

What we call a *field* is a part of a working register. If the register is 16 nibbles wide, we can, for example, call a five-nibble wide section a "field" of the working register.

The Saturn processor already has predefined fields, so it's easier for us to work.

Before we begin learning about fields, I want to tell you about a strange thing with the Saturn processor: nibble inversion. *Stay awake for this part*!! :-)

25.1 Nibble inversion

This is something that will be complicated for some. The Saturn inverts the order of the nibbles when reading from or writing to memory. What I call reading nibbles is moving nibbles from an address in memory to a working register, and writing being moving nibbles from a working register to an address in memory.

Working registers can store up to 16 nibbles. Each nibble inside the register will have a number, as you will learn in the diagrams below. Nibble number zero is the first one (the LSB nibble), and nibble number #Fh is the last one (the MSB nibble). The nibbles are numbered from right to left.

Here are the 16 nibbles of any working register:

F	E	D	C	B	A	9	8	7	6	5	4	3	2	1	0

Note that they are numbered from right to left, so the "start" of the register is the rightmost nibble, and the "end" (nibble #Fh) is the leftmost nibble.

Although *we* (except Arabic or Hebrew friends, and Chinese/Japanese writers) read from left to right, the Saturn starts moving nibbles from the LSB nibbles (nibble number 0, then 1, and so on).

Therefore, the Saturn processor "reads" and "writes" from right to left. For the Saturn, the first nibble is the rightmost.

Let's consider our register C. It's 16 nibbles wide, so we can imagine we have C like below:

F	E	D	C	B	A	9	8	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Now, think of this in memory:

Displacement	Address in memory	Nibbles
0	00100	1B01C
+5	00105	022F0
+10	0010A	00033

If I read 15 nibbles into the C register, starting at address 00100, C will contain the following values:

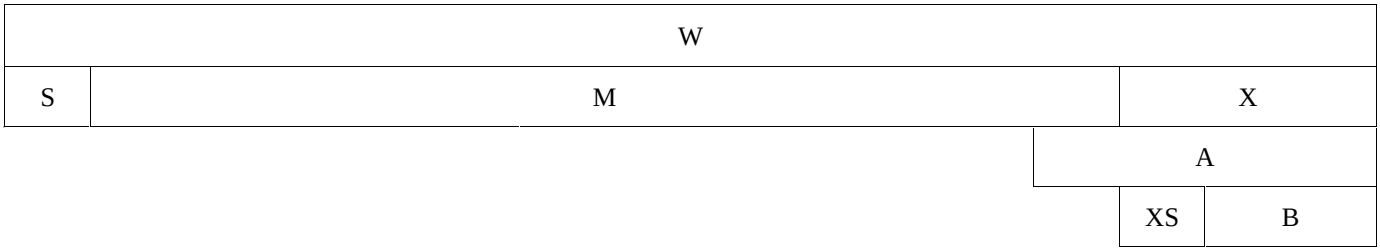
F	E	D	C	B	A	9	8	7	6	5	4	3	2	1	0

0	3	3	0	0	0	0	F	2	2	0	C	1	0	B	1

The nibble number #Fh has not been modified, but you can see that the Saturn loads nibbles from *right* to *left*. If we have 1B01C 022F0 00033 in memory we get 33000 0F220 C10B1 in our working register. (Notice that I removed nibble #Fh, as I am only working with 15 nibbles for this example). Remember that we read the contents of the register from *right* to *left*.

25.2 Available fields

The drawing below is extremely important, so you should *learn* it and be able to reproduce it from memory. I learned it the first time I saw it, following another coder's advice, preventing me from having to search for this information all the time. To make it simple for you: *learn this*! :o)



25.2.1 The W field

This field is easy: it's the *whole* register of 64-bits.

If we give a number to each nibble, we start from 0 and go to 15, because we have 16 nibbles. W is all 16 nibbles:

F	E	D	C	B	A	9	8	7	6	5	4	3	2	1	0

←——————————— W ———————————→

W means **word** (sometimes called "**wide**")

When we want an instruction to work with all 16 nibbles of a working register, we use the W field.

25.2.2 The A field

This is the most used field of all. It's 5 nibbles wide, and of course:

A means **address**

It looks like this:

F	E	D	C	B	A	9	8	7	6	5	4	3	2	1	0

←——— A ———→

Please *remember* that the first nibble is 0, not 1. So when we think of the A field, we have five nibbles: the first one is 0, the second one is 1, the third one is 2, and so on.

Don't forget it's from right to left.

25.2.3 The B field

This is another very useful field: it's the byte one! :-)

As you know, two nibbles form one byte, so the B field is two nibbles wide:

F	E	D	C	B	A	9	8	7	6	5	4	3	2	1	0
														←	B →

Do you still remember the MSB and LSB? We can say that the B field is a part of (or is) the LSB of the working register :-)

25.2.4 The X field

This field is 3 nibbles wide. It's very useful to use with an another component of the processor: OUT.

OUT will be used to check a key or to produce a sound using the HP's buzzer. OUT will get three nibbles of data, so the X field will be used with OUT. Very useful, and for other things too…believe me. ;-)

F	E	D	C	B	A	9	8	7	6	5	4	3	2	1	0
													←—	X	—→

X = eXponent (see M field below)

25.2.5 The XS field

The XS field is the third nibble of the register, that is, the MSB of the X register.

XS = eXponent Sign

25.2.6 The M field

This is the Mantissa field, nibbles 3 to 14

This field comes from the way the HP71 handled numbers with its Saturn processor: S (described below) was the used for the sign, M for the Mantissa, and X for the eXponent.

Now you know why those letters are there. :-)

F	E	D	C	B	A	9	8	7	6	5	4	3	2	1	0
	←————————					M							———————→		

25.2.7 The S field

The S field is the last nibble of the working register, number 15.

25.2.8 Field summary

So, keep this in mind:

- working registers can be divided into fields
- we will often reduce an instruction to a field

- the fields here have *fixed* sizes

Remember, the sizes are:

W	64 bits or 16 nibbles	Whole register
S	4 bits or 1 nibble	Sign
M	48 bits or 12 nibbles	Mantissa
X	12 bits or 3 nibbles	eXponent
A	20 bits or 5 nibbles	Address
XS	4 bits or 1 nibble	eXponent Sign
B	8 bits or 2 nibbles	Byte

26 The P register and WP

We have seven defined fields, which are all described above. There is one field, WP, whose size can be *customized* by the coder, using the P register.

First I will explain P. This register can only store *one* nibble, but it is very special. Its value will go from 0 to F.

1. We can use P as a field. If P contains the value *n*, then the P field is the nibble number *n* of the working register. Example: if P=2, then using the P field reads or writes *one* nibble, nibble number 2 (that is, the *third* nibble of the register, because the first one is 0, second one is 1, etc.)

2. P defines the size of WP. WP is Wide-P. Its size in nibbles goes from the first nibble to nibble *n*. Example: if P=5, then WP is nibbles 0 to 5, thus we can work with 6 nibbles using WP.

WP is useful for reading information with variable sizes, like reading a binary integer.

Remember:

- When you want to read one nibble that is not the first one of the register, we load its position number into P, and then we read it. For example, if we want to read nibble number 8, we will do P=8 and then read one nibble there using the P field (I'll show you how later).
- P defines the size of the WP field, from 0 to *n*, with *n* being the value of P. If you want to read or write *m* nibbles you do P=(*m*-1) and then use WP as the field.

Now, something really, really important about the value of P.

When your program starts, P=0, and when your program ends, you must be sure P is always set to 0. If you modified its value, you'll have to put it back to 0.

Why? When you load data inside a register, P defines where (which nibble) the data will start loading.

There is an instruction to load nibbles inside the C or A registers, but there are none to load nibbles inside B or D. If you want to load nibbles into B or D, you'll first have to load nibbles into C or A and then copy those nibbles into B or D.

If P is set to 0, and you load #A3Bh inside C, you get:

F	E	D	C	B	A	9	8	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	0	0	0	0	0	B	3	A

But if P is set to 2, and you load #A3Bh inside C:

F	E	D	C	B	A	9	8	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	0	0	0	B	3	A	0	0

And if P is set to E, and you load #A3Bh:

F	E	D	C	B	A	9	8	7	6	5	4	3	2	1	0
3	A	0	0	0	0	0	0	0	0	0	0	0	0	0	B

WHAT?

Yes...the 3 nibbles have been loaded, starting *at* nibble E, and because we load 3 nibbles, two are loaded in #Eh and #Fh, and then the last one is loaded at the beginning of the work register.

This introduces:

27 Circular handling of numbers inside the processor

*** IMPORTANT STUFF * READ TWICE IF NEEDED ***

When we work with *fields*, we fix a limit on the number of nibbles that will be used.

The B field works with two nibbles. Every instruction, like adding a value, will work on *only* those two nibbles, and won't affect the remaining 14 nibbles of the register.

If a register contains #FFh in its B field size, what will happen if I, for example, add 1? Remember that we are using the B field so we cannot go from #FFh to #100h, because that would be using 3 nibbles.

What will our lovely Saturn friend do? It does this:

#FFh + 1 = #00h (when working on a B field) *and* something called the "carry" will be set to 1.

The carry is a flag; flags are indicators inside the processor that you can test. They're described later.

If I were working with the A field, and A contained #00000h, removing 1would return #FFFFFh inside the A field, and the carry flag would be set to 1.

If you have #FEAh inside the X field of a register and add #20h, you'll get #00Bh in the X field of the register.

Notice that when you are working on a register and perform an operation, like adding or removing a value, if there is a loss of information because we "go around" inside a field (like #FFh + 1 gives #00h as result when using the B field) the processor sets the carry bit to one to warn you.

There is another thing that can lead to loss information: shifting bits.

I will describe that when I start giving you instructions. I am going to give you a complete description of *every* instruction, how it is coded, how many processor cycles it uses, and so on.

If we do a summary about what we have learned so far about the Saturn processor we have:

- Four working registers: A, B, C, and D
- Five save registers: R0, R1, R2, R3, and R4
- P, which gives you access to *any* nibble of a register, *and* defines the WP field size
- WP, which goes from nibble 0 to nibble (*n*-1) with *n* being the value stored in P
- D0 and D1, which are pointer registers, and allow you to point to a nibble in RAM/ROM (from #00000h to #FFFFFh)
- PC, which contains the current instruction being executed.

Now we must learn about RSTK, ST, HS, OUT, and IN.

Let's go :-)

28 RSTK

RSTK means "Return STacK."

First, what is a "stack?" Hopefully, as HP users, we already use a stack when using our HP. The stack is used *a lot* in computers because of memory structure. This will change when 3D chips are introduced, and when HP will have eaten Intel and Microsoft, but let's be patient until then, and keep that secret. ;)

The RSTK is a LIFO Stack. LIFO means Last In First Out.

Like our HP stack, if you enter a number onto RSTK, the first one to be popped out will be the last one that was put in.

To put it another way, it's like pushing coins into a tube and removing them from the same hole they were put into.

(Another kind of stack is FIFO, First In First Out. This is like a tube where we put coins on one side and we wait for them at the other hole: the first ones pushed into the tube will be the first ones to be recovered.)

This RSTK is like our HP stack. It's limited in size: there are *eight* levels available, and each level has a *fixed* size: 5 nibbles.

This is used by subroutines. When you call a subroutine from an assembly language program, the processor records where we are, stores it into RSTK, and jumps to the subroutine. When the subroutine wants to go back to where it was called from, it pops the "return address" from the return stack and jumps to it. Because it is an address, we only need to have five nibbles per level on the RSTK stack.

So the RSTK contains eight 5-nibble addresses. The interrupt system of the HP uses two levels when called, and sometimes a third level is used, so unless we don't allow *any* interruptions, we will only be allowed to use *five* levels.

RSTK is useful for quickly storing temporary information. We'll discuss this more later.

29 Flags: ST, HST, and Carry

Flags are *special* bits that can have two values: 1 or 0. When a flag is *set*, it has a value of 1. When a flag is *clear*, it has a value of 0.

29.1 Carry

This bit is a flag in itself. Each time an operation (like adding or subtracting, or when we get an overflow* on a register) has a carry, the carry bit is set to tell us that either a carry or an overflow has occurred.

* overflow: I explained it above. When we work with a field of a working register, if adding or removing a value makes us "go around" the carry is set to 1. If field A of C contains #FFFFFh, adding 1 to field A will put #00000h in field A of C, and the carry will be set to 1.

Also, *each* time we perform a test, if the test is *true*, the carry is set (set to 1), and if the test is *false*, the carry is cleared (set to 0).

When we perform a test, the processor will set or clear the carry. Next the carry bit is tested to determine what to do next.

29.2 ST (Status Bits)

The ST flags are composed of 4 nibbles. Therefore, we have 16 bits available, with each one being a flag.

As a coder, you are allowed to use flags 0 to 11. Flags 12 through 15 have special meanings and can be dangerous if misused.

Here are the meanings of bits 12 through 15:

bit 12: Forced wakeup request (overrides DeepSleep)
bit 13: Set if an interrupt has occurred
bit 14: Set if an interrupt is pending
bit 15: Set to 1 if interrupts are enabled or 0 if interrupts are disabled

ST bits 0 through 11 may be given whatever values you want. You can later test them, with each one being a flag with the meaning you give it. :)

The Saturn handles bits 12-15. We'll use bit 15 sometimes, but you should not mess with the others.

29.3 HST (Hardware Status Bits)

This flag has one nibble, or four bits. The Saturn uses these four bits to handle special hardware events. The four bits are:

bit 0: XM (eXternal Module missing)
bit 1: SB (Sticky Bit)
bit 2: SR (Service Request)
bit 3: MP (Module Pulled)

There is only one bit that can be set to 0 by us: XM. We'll use an instruction called RTNSXM (ReTurN and Set XM).

You'll find the SB to be useful: when we shift bits around a register, if a bit goes from the MSB to LSB, or from LSB to MSB, the Saturn will turn the SB to 1 to warn us there has been a loss of information during the bit shift. This will happen if bit 1 of the LSB goes around and gets into the MSB nibble.

30 OUT and IN

The Saturn processor is able to exchange information with the keyboard and produce sounds. Two registers inside the processor are thus used: OUT and IN. They have specific sizes:

OUT is 3 nibbles wide (12 bits)
IN is 4 nibbles wide (16 bits)

When we want to test a key, we'll put a value (from a table of values) into a register like C, and send it to the OUT register. Then electrical power will be given to lines in the keyboard (from 1 to all), and if a key is pressed, the result is put into the IN register. We'll read that value, as it will tell use whether a key was pressed, and if so, which one. :-)

31 How the processor works: life of an instruction

How does the Saturn execute an instruction?

We know the PC register (Program Counter) is used by the processor to know "where" it is. The processor will always update this counter so it knows where to look for the next instruction to execute.

This is the cycle the processor uses to execute an instruction:

1. The PC register tells the processor the location of the next instruction to execute.
2. This address is sent on the bus of the Saturn processor. The bus is a line that links the processor, the RAM, the ROM, the input/output components, etc., together.
3. The Saturn receives the instruction to execute in an internal register. It has to be decoded; this is mandatory, as the processor needs to decode the instruction using specialized circuits.
4. Once the circuits in charge of decoding the instruction have done their work of decoding, they give the instruction to the unit in charge of sequencing commands inside the processor. This unit uses a quartz frequency to execute commands . (The HP48 S/SX have ~2 MHz quartz, and the HP48 G/GX have ~4 MHz quartz.)
5. The unit in charge of sequencing commands sends orders by calling and using the LAU (Logical Arithmetical Unit). This unit is in charge of the logical operations (complementation, logical OR/AND/NOT, etc.) and arithmetic operations (subtracting, adding, etc.).
6. The result is stored in one of the processor's registers. PC is incremented to the next instruction.

And we repeat. :-)

This introduces the concept of:

32 What is a processor cycle?

When we ask the processor to do something, even if it looks simple to us (like adding 1 to a register's value) it is more complex for the processor. The instruction has to be fetched, then decoded, then passed to the unit in charge of executing it, and so on.

So each of these steps is divided into elementary operations, and each one is done in one cycle. Some instructions will require 11 cycles, for example, but others will require more or less.

Each cycle is a "state," or an impulse coming from the processor's quartz.

Consider the HP48G, whose processor runs at about 4 MHz. Each second, four million impulses are generated and each instruction takes a predetermined number of cycles to be completed.

Note that the cycle counts given for instructions in this book are not the exact cycle counts, as it is virtually impossible to get exact counts. That would need an oscilloscope, but these should be close enough. These cycle counts are from the Meta Kernel documentation (thanks, Jean-Yves!) and should be close to the correct figures.

In addition, the cycle counts may vary depending on whether the instruction is at an odd or an even memory address! Some instructions are given with a fractional cycle count. With each of these instructions, if it is at an even address round down (take the FLOOR of the number), and if it is at an odd address round up (take the CEIL of the number). Also, if two counts are given and are separated by a comma, add the FLOOR of the second number if it is reading from an even address or add the CEIL of the second number if it is reading from an odd address.

Part IV: Saturn Instruction Set

33 Working register instructions

33.1 Loading a value inside a register

It is very useful to be able to load a value inside a working register. There are only two instructions available, and only the A and C registers can be loaded with a value.

The instruction LA (Load in A) is used to load a value into the A register, and the instruction LC is used to load a value into the C register.

For example, to load #ABCDE inside the C register, write:

```
LC ABCDE
```

After the code is run, the C register will look like this, with each nibble moved into the register:

F	E	D	C	B	A	9	8	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	0	0	0	E	D	C	B	A

As we have seen, the register is read from right to left, and LC and LA work this way as well.

As the working registers C and A are 16 nibbles wide, we are unable to load more than 16 nibbles into either one.

Let's see what happens. Put an empty string onto the stack, and launch StringWriter or MiniWriter. Write this:

```
LC 1234567890ABCDEF1
@
```

You can get the @ by pressing [ALPHA][Right-shift][ENTER].

HP-ASM will consider the @ as the end of the source file, so anything after that character will be ignored by the compiler.

Now, type ASM (or →ASM if you are using HP-ASM 1.0) to compile the file.

You will get a "Bad Size" error because you tried to load too many nibbles.

On the stack you will have:

```
3:  "LC 12345678...
2:            <0h>
1:  "LC 12345678...
```

The string on level 3 is the original string, level 2 contains a system binary of the line where the error is, and level 1 contains the line that has the error.

Here we can use the ED command of HP-ASM to edit the source, and it will move the editor to the line of the error; but ED will use the HP's internal editor. :P

It is possible to customize ED to use any editor, but that requires a little knowledge of library hacking and System RPL programming. Maybe if there are enough requests for a modified version I'll rewrite ED to use the editor of choice.

Until you customize ED, you'll probably want to drop two levels and edit the string with your favorite when you get an error.

You can use `LC` or `LA` to load from 1 to 16 nibbles inside A or C. There are no instructions to load a value into B or D, so we will have to first load a value into A or C and then move it into B or D.

An important thing to consider when using `LA` or `LC` is the value of P. The value of P, when you will start a program, will be zero. You can change its value, but when you quit, you have to restore it to zero.

If P=0, when you load something into A or C it's loaded from nibble 0 of the register. If P=4, loading will start at nibble 4.

For example, this is what happens if we load ABCDE into C when P=4:

F	E	D	C	B	A	9	8	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	E	D	C	B	A	0	0	0	0

So if you modify P's value (and you will, if you need to read the nibble *n* of a register or use WP) you have to pay attention to `LA` and `LC` instructions.

Also, loading a value overwrites any value previously stored in a register. Suppose C contains:

F	E	D	C	B	A	9	8	7	6	5	4	3	2	1	0
3	A	8	B	B	3	8	2	0	0	0	A	F	0	F	C

After loading three zeros into it, C will contain:

F	E	D	C	B	A	9	8	7	6	5	4	3	2	1	0
3	A	8	B	B	3	8	2	0	0	0	A	F	0	0	0

The nibbles are not "moved" inside the register when you load a value, so any previous value is overwritten.

I have also shown you there is a circular loading of nibbles into registers. If we give P the value of E and load 6 nibbles using `LC ABCDEF` we will get:

F	E	D	C	B	A	9	8	7	6	5	4	3	2	1	0
B	A	0	0	0	0	0	0	0	0	0	0	F	E	D	C

Each time an instruction is listed, five facts are given:
1. The name of the mnemonic (here "`LA`" and "`LC`")
2. How it's coded in memory, in hexadecimal
3. Which fields can be used or are used
4. Whether the carry bit is affected
5. The approximate number of cycles the processor needs to complete the instruction.

33.1.1 LA, LC

Mnemonic	Hex-form	Fields	Carry	Cycles
`LA` *n*`..1`	8082 *n* 1...*n*	according to P value	no	7.5+1.5*n*
`LC` *n*`..1`	3 *n* 1...*n*	according to P value	no	3+1.5*n*

1 is the first nibble being loaded and *n* the last one. You have examples of how nibbles are loaded above. `LA` and `LC` do not modify the carry, even if a circular loading appears.

As you see here, `LC` is faster than `LA`. It's coded with fewer nibbles in memory. It's not very important, unless some part of your code needs to go really fast (like everything linked to display, usually). If you need speed, use `LC` rather than `LA`.

n is the number of nibbles being moved. This piece of code:

```
LC 02A2C
```

will move 5 nibbles into the C register, and needs 3+1.5x5 cycles, or 10.5 cycles.

33.1.2 Setting a register to zero

Here we are going to learn how to set a register, from 1 to 16 nibbles, to zero. The first letter will be the letter of the register, or A, B, C, or D. Then, the equal sign (=) and 0 of course!!

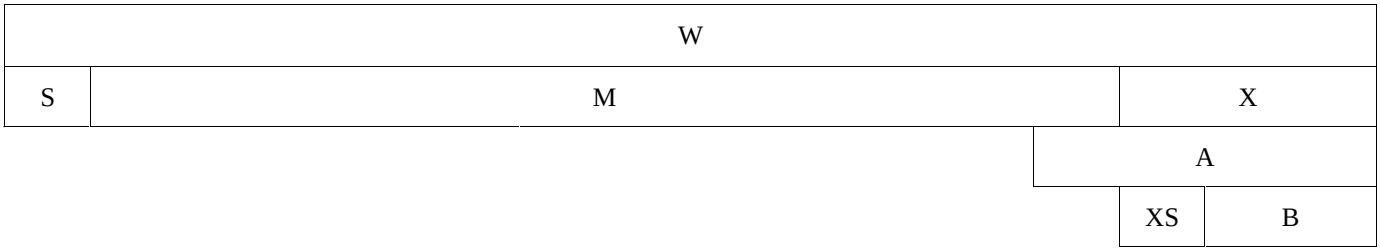
To clear the whole C register (all nibbles), simply write:

```
C=0 W
```

:-)

Remember that W is *word* (or *wide*). It is the field that affects the whole register, or 16 nibbles!

Here are the fields (I described those in part 2):



So if we now want to put all zeroes into the X field of register B, we will use:

```
B=0 X
```

Mnemonic	Hex-form	Fields	Carry	Cycles
`A=0 f`	D0 Ab0	A only P,WP,XS,S,M,B,W	no no	8 4.5+*n*
`B=0 f`	D1 Ab1	A only P,WP,XS,S,M,B,W	no no	8 4.5+*n*
`C=0 f`	D2 Ab2	A only P,WP,XS,S,M,B,W	no no	8 4.5+*n*
`D=0 f`	D3 Ab3	A only P,WP,XS,S,M,B,W	no no	8 4.5+*n*

What is f? It's the field!! :-)

The A field is treated separately than others, because the A field is used so often in the processor that it has specific hex forms.

Also, you will see that the "hex-form" column lists hex numbers, but sometimes, letters. You will use a table to know which letter is to be used, according to the field.

Example: I want to do `A=0 S`. The instruction is Ab0, and by using the table below you can find the value of b:

Field	a values	b values
P	0	8
WP	1	9
XS	2	A
X	3	B
S	4	C
M	5	D
B	6	E
W	7	F

Because there is a "b" you need to use the "b" column. Look for S, and you'll find "C," so you know the instruction to do `A=0 S` is "AC0."

You will not need this information unless you intend to code a disassembler or an assembler. But now if you want to create one you have the info. :-)

As usual, "*n*" is the number of nibbles being cleared to 0.

Let's code! :)

We are going to code something that sets the entire C register to zero and then quits. HP-ASM wants one instruction per line, so each time you write a complete instruction, you will have to move to the next line, like below:

```
C=0 W
A=DAT0 A
D0=D0+ 5
PC=(A)
@
```

I know, there are instructions here you don't know yet, but let's first compile it and run it, okay? If you don't make any mistakes, you'll get a Code object. :-)

Press EVAL, and it will run! The W field of the C register will clear, and then the code will end.

Now, I need to explain the last 3 instructions.

```
A=DAT0 A
```

You know that the Saturn processor has two pointer registers: D0 and DA. (It has PC too, but we won't use it here.) If D0 points to an area somewhere, then DAT0 is the data there. It's the same for D1. If D1 points to #ABCDE, we will use DAT1 to get nibbles from there.

Let's say there are 5 nibbles:

Address	Nibbles
#80120h	02A2C00005

If we want to read, say, three nibbles from #80120h, and put them in A, we will code this:

```
D1= 80120
A=DAT1 X
```

First, we make D1 point to #80120h, and then we read three nibbles (field X) into A.

If we want to read 10 nibbles using C and D0, we do this:

```
D0= 80120
C=DAT0 10
```

(you can put a number in the field location, which is cool :-)

Now, let's explain these instructions:

```
A=DAT0 A
D0=D0+ 5
PC=(A)
```

When a program (your program) is given control, some registers and some pointers are used by the HP:

D1 points to the first level of the stack
D0 points to the next object to execute
B points to the return stack of the RPL
D contains the free memory in 5-nibble blocks

When your program starts, those values are used. Sometimes we will have to "save" some register values, because if we lose B or D0, we'll get into trouble (the HP will halt, or if you're unlucky, crash).

D0 points to the next RPL object to execute. When we want to quit, we have to read the address of the next RPL object to run into A, that is:

```
A=DAT0 A
```

So now, A contains the address (5 nibbles read from where D0 points to, A being 5 nibbles wide).

Now that we have read it, we have to move D0 5 nibbles later. It is important to increment D0 to the next object. We are going to execute the current one, so the next instruction that will read from D0 must read from the good zone. :-) That's why we increment D0 by 5 nibbles: we have read 5 nibbles, so we must move 5 nibbles later:

```
D0=D0+ 5
```

The last instruction jumps to that address, directly using the PC pointer. The Saturn processor, when it needs to know "where" to continue, will read from PC. So we make PC point to A, and the execution continues. Remember that the following code is the usual way to "return to RPL":

```
A=DAT0 A
D0=D0+ 5
PC=(A)
```

You will use this to quit your programs most of the time. :-)

Other small pieces of code will be given as we learn more Saturn instructions, like code to drop, swap, and so on.

33.1.3 Changing the value of a bit

We can change the value of any bit inside the first four nibbles of registers A or C (not B or D!). As you have learned, when a bit is equal to 1 it's set and when equal to 0 it's clear.

We will use the `ABIT` (for register A) and `CBIT` (for register C) instructions. When we want to set a bit, we use:

```
ABIT=1 n
CBIT=1 n
```

and to clear a bit:

```
ABIT=0 n
CBIT=0 n
```

where *n* is the number of the bit set to 1 or 0, and *n* may have values from 0 to 15 (number of the bit)

NOTE:	Some people don't like to start counting from 0, and prefer to count from 1. HP-ASM lets you choose. When you press its OPT command, the "B" key lets you choose between 0-15 (the usual way in computers) or 1-16. You can choose what you want. I will always use 0-15 here, so if you choose 1-16, you'll have to adapt your sources.

Here are the mnemonics:

Mnemonic	Hex-form	Fields	Carry	Cycles
`ABIT=0` *n*	8084*n*	bit number *n*	no	7.5
`ABIT=1` *n*	8085*n*	bit number *n*	no	7.5
`CBIT=0` *n*	8088*n*	bit number *n*	no	7.5
`CBIT=1` *n*	8089*n*	bit number *n*	no	7.5

Example: I want to turn off the busy indicator when I launch my code. How would I do this?

First, you have to know that indicators are just bits stored in a certain area of RAM. By giving bits values of 1 we can turn on indicators, and by giving them values of 0 we can turn them off.

I will explain that now, so you can learn about these bits and special values.

At the address #0010Ch there is a nibble, and the value of its bits can be used to turn on or off some indicators:

Bit number	Function
0	busy indicator (1=on, 0=off)
1	transmit indicator (1=on, 0=off)
2	?
3	show indicators? (1=on, 0=off)

So if I want to turn off the busy indicator, I have to read the nibble at #0010C, set its bit number 0 to 0, and then leave.

Let's do it!! :-)

First, we need to use D0 or D1 to point there. But the HP uses D0 and D1! D0 points to the next RPL object to execute, and D1 points to the first level of stack.

So we must "save" the value of D1 (or D0) and restore it before leaving.

I'm going to use C to read the nibble and A to save the value of D1, so I can have D1 point somewhere else.

To move the value of D1, we're going to use an instruction we have not seen yet: `AD1EX`. This instruction moves the 5 nibbles in D1 to A, and the 5 nibbles of A to D1. After using it, the previous value of D1 will be kept in A. This instruction is used to exchange the contents of D1 with A and is explained later in this document. :-)

First, move the value of D1 to A:

```
AD1EX
```

Then, point D1 to #0010Ch, and read one nibble there:

```
D1= 0010C
C=DAT1 P
```

Why P? P is equal to zero, so when using P, it's like using:

```
C=DAT1 0
```

In other words, it reads bit 0 of the memory at the location referred to by D1 and puts it in C. Now, the value of bit 0 in C must be changed. We can use the instruction we have just learned:

```
CBIT=0 0
```

Write it back to memory:

```
DAT1=C P
```

(Instructions for reading and writing data are explained later, but I think you now understand how it works.)

Then give D1 its old value:

```
D1=A
```

and quit to RPL the usual way:

```
A=DAT0 A
D0=D0+ 5
PC=(A)
```

So the code to turn off the busy indicator and return to RPL, is:

```
AD1EX
D1= 0010C
C=DAT1 P
CBIT=0 0
DAT1=C P
D1=A
A=DAT0 A
D0=D0+ 5
PC=(A)
@
```

Of course, it goes very quickly. So, if we want to really see it happen, we have to store the code inside a variable and then wait for a key before we leave.

Compile that code, and store it under the name BOFF (Busy OFF). Next, run the following RPL program:

```
<< BOFF 0 WAIT DROP >>
```

and look at the busy indicator!! :-)

You will have `<< BOFF 0 WAIT DROP >>` on the stack, and the busy indicator will be off. :-)

33.1.4 Exchanging two registers

We can exchange two register values. We can either exchange the whole register or only a smaller field of it.

Mnemonic	Hex-form	Fields	Carry	Cycles
`ABEX f`	DC AbC	A P,WP,XS,S,M,B,W	no no	8 4.5+*n*
`BCEX f`	DD AbD	A P,WP,XS,S,M,B,W	no no	8 4.5+*n*
`ACEX f`	DE	A	no	8

	AbE	P,WP,XS,S,M,B,W	no	4.5+*n*
`CDEX f`	DF AbF	A P,WP,XS,S,M,B,W	no no	8 4.5+*n*

The number *n* is the number of nibbles being exchanged between one register and the other.

> **NOTE:** ACEX is the same thing as CAEX of course. HP-ASM wants you, when you have 2 registers to exchange, to put them in alphabetical order. This means you must write `ACEX f` rather than `CAEX f`, as it will not recognize `CAEX f`.

We know we cannot load a value into D because there is no LD, so we can load nibbles into C or A and then exchange them with C, like:

```
LC 124
CDEX X
```

This exchanges 3 nibbles between C and D, thus "moving" #124h into D. If you want to exchange two registers, you will use the W field, of course. :)

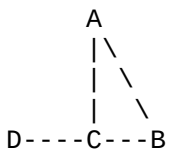
33.1.5 Copying a register into another

Not only can we exchange two registers, but we can also directly overwrite nibbles from one to another one. Here are the instructions to do that:

Mnemonic	Hex-form	Fields	Carry	Cycles
`A=B f`	D4 Ab4	A P,WP,XS,S,M,B,W	no no	8 4.5+*n*
`A=C f`	D8 Ab8	A P,WP,XS,S,M,B,W	no no	8 4.5+*n*
`B=A f`	D5 Ab5	A P,WP,XS,S,M,B,W	no no	8 4.5+*n*
`B=C f`	D9 Ab9	A P,WP,XS,S,M,B,W	no no	8 4.5+*n*
`C=A f`	DA AbA	A P,WP,XS,S,M,B,W	no no	8 4.5+*n*
`C=B f`	D6 Ab6	A P,WP,XS,S,M,B,W	no no	8 4.5+*n*
`C=D f`	D7 Ab7	A P,WP,XS,S,M,B,W	no no	8 4.5+*n*
`D=C f`	DB AdB	A P,WP,XS,S,M,B,W	no no	8 4.5+*n*

That's a lot, isn't it? :-)

If you look closely, you'll see that register D can only be changed with C.



B can be used with A or C, A can be used with B or C, and C can be used with A, B, or D, so D can only be used with C.

33.2 Working with registers

33.2.1 Incrementing a register

This is the same as adding 1 to a register. Here, the field will specify how many nibbles are used. For example, if I have #FFh in C and I do `C=C+1 B` I will get #00h in C and the carry set to 1 (because an overflow occurred). If I use the X field instead, I will get #100h.

Mnemonic	Hex-form	Fields	Carry	Cycles
`A=A+1 f`	E4 Ba4	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`B=B+1 f`	E5 Ba5	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`C=C+1 f`	E6 Ba6	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`D=D+1 f`	E7 Ba7	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*

Example: I want to do a kind of "wait" loop. All I have to do is choose one register, like C. Now, I can choose a field, though not one that's too big, like W. Let's try A. We are going to load the value #FFFFFh into A, subtract one, and loop the subtraction until we reach 0.

First, for the purpose of this exercise it's better to do:

```
C=0 A
C=C-1 A
```

than:

```
LC FFFFF
```

Why?

The second method will need 7 nibbles to be coded (34FFFFF) and 8 cycles

The first method will only need 4 nibbles (D2CE) but more cycles: 14!

As you see, bigger code is usually faster, and if you want to save space, you'll usually need more cycles.

So as we want a small code (although if you need to go fast, choose smart coding, and don't misuse memory) we're going to use:

```
C=0 A
C=C-1 A
```

Now, if we want to loop, we need a *label*. A label is a line which starts with * so the compiler will know we can "branch" to that label.

In this part, we're also going to use a new instruction:

```
GONC
```

`GONC` means: Go if No Carry

Here is the code:

```
C=0 A
```

```
C=C-1 A
*LOOP
C=C-1 A
GONC LOOP
A=DAT0 A
D0=D0+ 5
PC=(A)
@
```

First, we need to have #FFFFFh in C, so we turn C to 0 using field A, and then subtract 1 to get: #FFFFFh.

I also inserted a label called “LOOP.” A label first has a star (*) and then its name. I called it “LOOP” but you can call it what you want.

Inside the loop, all we have to do is remove 1 from C(A)

C(A) means “field A of C”

The instruction GONC will branch, or go to, LOOP until the carry is set to 1. When will it be? When C(A) reaches zero, the next `C=C-1 A` will put C(A) from #00000h to #FFFFFh and the carry will be set to 1. Because the carry is set to 1, `GONC` will no longer jump to LOOP.

Then the program ends.

Try the source above and see. It waits and then quits.

You can use X or B if you want a loop that lasts less time.

We can write it another way: instead of checking whether the carry is set (that is, whether a negative overflow occurred) we could have checked whether C(A) was equal to zero, and loop until it’s equal to zero.

33.2.2 Something special with the carry in loops

When you use the carry in loops, you have to pay attention to the fact that the loop will repeat until an overflow or underflow occurs; in my example, when C(A) reaches zero, it loops again, because the carry is not set to 1.

Take a look at this code:

```
LC F
*LOOP
C=C-1 P
GONC LOOP
```

The LOOP is executed F + 1 times. That’s 16 times, not 15, because when C(P) reaches zero, it loops, and the carry only sets as the underflow occurs, when #0h becomes #Fh.

If we want a piece of code to loop four times, if we use the carry we will then load 3, not 4, inside the register that will be used as a counter:

```
LA 3
*LOOP

(Insert the code that has to be repeated four times here)

A=A-1 P
GONC LOOP
```

If this is difficult for you, ask for help in the newsgroup. Everything must always be clear. If you need help, just ask. Learning well will make you code well. ;)

33.2.3 Adding two registers

Here we add the contents of a register to another one, involving as many nibbles as the field uses. The carry will be affected, of course. For example, adding #Eh to #3h using the P field (with P=0) will set the carry bit.

Mnemonic	Hex-form	Fields	Carry	Cycles
`A=A+B f`	C0 Aa0	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`A=A+C f`	CA AaA	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`B=B+A f`	C8 Aa8	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`B=B+C f`	C1 Aa1	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`C=C+B f`	C2 Aa2	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`C=C+D f`	CB AaB	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`D=D+C f`	C3 Aa3	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*

33.2.4 Adding a register to itself

This will be used to multiply the contents of a register, or, if we use a field other than W, a part of the register. Like before, the carry will be affected if there is an overflow when adding the register to itself on field f:

Mnemonic	Hex-form	Fields	Carry	Cycles
`A=A+A f`	C4 Aa4	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`B=B+B f`	C5 Aa5	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`C=C+C f`	C6 Aa6	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`D=D+D f`	C7 Aa7	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*

33.2.5 Adding a constant to a register

The instructions below let you add a constant to a register's contents using field f. Here we cannot use all fields we would like: we cannot use S, XS, WP or P.

THIS IS BECAUSE OF A BUG IN THE SATURN

The Saturn has a bug: **every** instruction that increments or decrements a register while working with 1 nibble with a constant greater than **one** will turn *all* the bits of the register to 1 if the carry bit is set.

In other words, if the carry bit is set, *all* the bits of the register are set. Because of the bug, the register would then look like this:

F	E	D	C	B	A	9	8	7	6	5	4	3	2	1	0
F	F	F	F	F	F	F	F	F	F	F	F	F	F	F	F

SOLUTION:

If you are working with *one* nibble using field S, XS, WP (with P=0), or P and want to add *more* than 1, you **must** add only 1 as many times as needed.

Example:

```
P=P+2 S
```

will not work! If the carry is set, you will lose the contents of the register!

```
P=P+1 S
P=P+1 S
```

is just fine. :) If the carry is set because of an overflow, everything will be fine.

This also true for:

```
P=P-2 S
```

which will not work either.

```
P=P-1 S
P=P-1 S
```

is okay!

REMEMBER:	When working with a field that deals with *one* nibble, if you want to add or subtract *more* than 1, you'll have to do several +1's or –1's to avoid the carry bug.

When you add 10 to a constant, using `A=A+3 f`, if you want to know how it's coded you have to remove 1 from the constant to find its value in the hex form. Thus, `A=A+3 A` will be coded 818F02 because we have a 2 instead of a 3, because it's c+1. This only matters if you are reading the hex form of instructions that add a constant to a register.

When you code, and write, for example, `A=A+3 A`, the assembler will do all the work, encoding "818F02," so you aren't bothered by that (c+1).

Mnemonic	Hex-form	Fields	Carry	Cycles
`A=A+c f`	818F0(c-1) 818a0(c-1)	A W,M,X,B	yes yes	8+*n* 8+*n*
`B=B+c f`	818F1(c-1) 818a1(c-1)	A W,M,X,B	yes yes	8+*n* 8+*n*
`C=C+c f`	818F2(c-1) 818a2(c-1)	A W,M,X,B	yes yes	8+*n* 8+*n*
`D=D+c f`	818F3(c-1) 818a3(c-1)	A W,M,X,B	yes yes	8+*n* 8+*n*

n = number of nibbles in the object that is being added to.

$2 < c < 16$

If you only want to add 1, it's coded differently than here, but when you will write your programs, that will be the assembler's job. :-)

In the beginning coders had to encode something+1 or something+c with c between 2 or 16. Today, with assemblers, we just add any value from 1 to 16, and the assembler does the work for us.

Here's a cool trick!

You know each nibble has a "weight," so you can alter P and add a value to a specific nibble, doing much more than simply adding a constant c. The nibble #2h (the third nibble in the register) has a weight of 16^2, so each time 1 is added to that nibble, 16^2, or 256, is added to the value.

So, sometimes (but not always) it's better to add a value to nibble *n* of a register rather than adding a big value to the whole register.

33.2.6 Decrementing a register

This is the same as removing one, so it's easy.

Mnemonic	Hex-form	Fields	Carry	Cycles
`A=A-1 f`	CC AaC	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`B=B-1 f`	CD AaD	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`C=C-1 f`	CE AaE	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`D=D-1 f`	CF AaF	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*

The carry bit will be set if we try to remove 1 from any field that is equal to zero.

33.2.7 Subtracting two registers

Well...it's subtracting one register from another.

Mnemonic	Hex-form	Fields	Carry	Cycles
`A=A-B f`	E0 Ba0	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`B=B-C f`	E1 Ba1	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`C=C-A f`	E2 Ba2	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`D=D-C f`	E3 Ba3	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`B=B-A f`	E8 Ba8	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`C=C-B f`	E9 Ba9	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`A=A-C f`	EA BaA	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`C=C-D f`	EB	A	yes	8

	BaB	P,WP,XS,S,M,B,W	yes	4.5+*n*
`A=B-A f`	EC BaC	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`B=C-B f`	ED BaD	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`C=A-C f`	EE BaE	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`D=C-D f`	EF BaF	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*

As always, *n* is the number of nibbles and the carry will be set if there is an underflow.

33.2.8 Subtracting a constant

Mnemonic	Hex-form	Fields	Carry	Cycles
`A=A-c f`	818F8(c-1) 818a8(c-1)	A W,M,X,B	yes yes	8+*n* 8+*n*
`B=B-c f`	818F9(c-1) 818a9(c-1)	A W,M,X,B	yes yes	8+*n* 8+*n*
`C=C-c f`	818FA(c-1) 818aA(c-1)	A W,M,X,B	yes yes	8+*n* 8+*n*
`D=D-c f`	818FB(c-1) 818aB(c-1)	A W,M,X,B	yes yes	8+*n* 8+*n*

We cannot use all fields here because of the carry bug. If you want to subtract c with c > 1, you'll have to use **only 1** as many times as needed to avoid the bug.

You may use the same trick here that was given before, so you can work with a specific nibble by modifying P. By using the weight of each nibble, you can subtract a big value.

33.2.9 Complement of 2

NOTE:	The result you will get using complement of 2 will depend on the *mode* the Saturn is under: decimal or hexadecimal. The instructions to change this mode are explained below. Remember that when your program starts, the processor will *always* be in hexadecimal mode, and when it quits, it *must* also be in hexadecimal mode.

Mnemonic	Hex-form	Fields	Carry	Cycles
`A=-A-1 f`	FC BbC	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`B=-B-1 f`	FD BbD	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`C=-C-1 f`	FE BbE	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
`D=-D-1 f`	FF BbF	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*

If the register (and thus the field) you are working with equals zero, then doing a `r=-r-1 f` (with r being the register) sets the carry.

33.2.10 Complement of 1

As before, the result here depends on the mode the Saturn is running under. All the bits are here inverted:

Mnemonic	Hex-form	Fields	Carry	Cycles
A=-A f	F8 Bb8	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
B=-B f	F9 Bb9	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
C=-C f	FA BbA	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*
D=-D f	FB BbB	A P,WP,XS,S,M,B,W	yes yes	8 4.5+*n*

33.2.11 Logical operations

33.2.11.1 Logical OR (also called “logical sum”)

With logical OR, the “!” symbol is used. The OR is done bit per bit with field f on the register used. The result is placed on the register *to the left* of the equal sign:

Mnemonic	Hex-form	Fields	Carry	Cycles
A=A!B f	0EF8 0Ea8	A P,WP,XS,S,M,B,W	no no	11 6+*n*
A=A!C f	0EFE 0EaE	A P,WP,XS,S,M,B,W	no no	11 6+*n*
B=B!C f	0EF9 0Ea9	A P,WP,XS,S,M,B,W	no no	11 6+*n*
B=B!A f	0EFC 0EaC	A P,WP,XS,S,M,B,W	no no	11 6+*n*
C=C!A f	0EFA 0EaA	A P,WP,XS,S,M,B,W	no no	11 6+*n*
C=C!B f	0EFD 0EaD	A P,WP,XS,S,M,B,W	no no	11 6+*n*
C=C!D f	0EFF 0EaF	A P,WP,XS,S,M,B,W	no no	11 6+*n*
D=D!C f	0EFB 0EaB	A P,WP,XS,S,M,B,W	no no	11 6+*n*

33.2.11.2 Logical AND (also called “logical product”)

The symbol used is “&”.

Mnemonic	Hex-form	Fields	Carry	Cycles
A=A&B f	0EF0 0Ea0	A P,WP,XS,S,M,B,W	no no	11 6+*n*
A=A&C f	0EF6 0Ea6	A P,WP,XS,S,M,B,W	no no	11 6+*n*
B=B&C f	0EF1 0Ea1	A P,WP,XS,S,M,B,W	no no	11 6+*n*
B=B&A f	0EF4 0Ea4	A P,WP,XS,S,M,B,W	no no	11 6+*n*
C=C&A f	0EF2 0Ea2	A P,WP,XS,S,M,B,W	no no	11 6+*n*
C=C&B f	0EF5 0Ea5	A P,WP,XS,S,M,B,W	no no	11 6+*n*
C=C&D f	0EF7	A	no	11

	0Ea7	P,WP,XS,S,M,B,W	no	6+*n*
D=D&C f	0EF3 0Ea3	A P,WP,XS,S,M,B,W	no no	11 6+*n*

33.2.12 Shifting one nibble left or right

There are four instructions to shift one nibble to the left, which multiplies the value by 16, and four others to shift one nibble to the right, which divides the value by 16.

When we shift to the left, all nibbles are moved one nibble to the left and the last nibble (number #Fh) is lost. Nibble number #0h receives a null nibble (value is #0h).

When we shift to the right, all nibbles are moved one nibble to the right and the last nibble (number #Fh) receives a null nibble. The first nibble (number #0h) is lost.

First we write the register's letter, then S (for Shift), and finally the letter of the direction, with R for right and L for left.

The mnemonic to move the C register right is C + S + R, so `CSR f`.

Mnemonic	Hex-form	Fields	Carry	Cycles
ASL f	F0 Bb0	A P,WP,XS,S,M,B,W	no no	8 4.5+*n*
BSL f	F1 Bb1	A P,WP,XS,S,M,B,W	no no	8 4.5+*n*
CSL f	F2 Bb2	A P,WP,XS,S,M,B,W	no no	8 4.5+*n*
DSL f	F3 Bb3	A P,WP,XS,S,M,B,W	no no	8 4.5+*n*
ASR f	F4 Bb4	A P,WP,XS,S,M,B,W	no no	8 4.5+*n*
BSR f	F5 Bb5	A P,WP,XS,S,M,B,W	no no	8 4.5+*n*
CSR f	F6 Bb6	A P,WP,XS,S,M,B,W	no no	8 4.5+*n*
DSR f	F7 Bb7	A P,WP,XS,S,M,B,W	no no	8 4.5+*n*

Here the carry isn't affected, *but* the SB (Sticky Bit) will be. If the nibble lost (on the left or right, depending on the direction we are shifting bits to) has a value other than zero, SB is *set*. It means there has been an important (at least, to be noticed) loss of value.

33.2.13 Rotating one nibble left or right

To rotate a nibble, we first write the letter of the register, then the letter S (for Shift), then the direction (R or L) and finally the letter C (Circling).

We thus have:

Mnemonic	Hex-form	Fields	Carry	Cycles
ASLC	810	All	no	22.5
BSLC	811	All	no	22.5
CSLC	812	All	no	22.5
DSLC	813	All	no	22.5
ASRC	814	All	no	22.5

BSRC	815	All	no	22.5
CSRC	816	All	no	22.5
DSRC	817	All	no	22.5

When we rotate to the right, the contents of nibble number #1h move to nibble #0h and nibble #0h moves to nibble #Fh.

To give you an example, this could be the content of a register before rotating to the right:

F	E	D	C	B	A	9	8	7	6	5	4	3	2	1	0
3	A	8	B	B	3	8	2	0	0	0	A	F	0	F	C

This is what the register would look like after rotating to the right:

F	E	D	C	B	A	9	8	7	6	5	4	3	2	1	0
C	3	A	8	B	B	3	8	2	0	0	0	A	F	0	F

Here, there is no loss of information because nibble #0h has moved to the other side of the register, becoming nibble #Fh.

The Sticky Bit will only be set if the bits of the nibble number #0h (which goes to nibble number #Fh) are null.

33.2.14 Shifting one bit to the right

Shifting one bit to the right is like dividing by 2, as you should know. To form the mnemonics, simply put the letter of the register, then "S" for Shifting, then the direction (R for right) and then B (for Bit).

Mnemonic	Hex-form	Fields	Carry	Cycles
ASRB	81C	All	no	21.5
BSRB	81D	All	no	21.5
CSRB	81E	All	no	21.5
DSRB	81F	All	no	21.5

Here, all 64 bits are moved to the right. Even though the carry bit is not modified, if bit number #0h is lost and was equal to 1, then the Sticky Bit is set.

Note that there are no instructions to shift one bit to the left.

33.2.14.1 Shifting one bit to the right within a field

This is like the previous commands, but here we use a field, so the shifting only occurs inside the field selected. All we have to do is add the letter of the field, like `ASRB A`.

Mnemonic	Hex-form	Fields	Carry	Cycles
ASRB f	819F0	A	no	13.5
	819a0	P,WP,XS,S,M,B	no	8.5+*n*
BSRB f	819F1	A	no	13.5
	819a1	P,WP,XS,S,M,B	no	8.5+*n*
CSRB f	819F2	A	no	13.5
	819a2	P,WP,XS,S,M,B	no	8.5+*n*
DSRB f	819F3	A	no	13.5
	819a3	P,WP,XS,S,M,B	no	8.5+*n*

These instructions divide the value of field f by 2.

The carry is not modified here, but if the bit lost on the right of the field used is equal to 1, then the SB (Sticky Bit) is set. The bit inserted on the left is 0.

Note that there are no instructions to shift one bit in a field to the left either.

34 Save registers

You are going to discover here that registers A and C are really important for the processor. :)

Only A and C can work with the save registers, of which there are five: R0, R1, R2, R3 and R4.

With these mnemonics, the register on the *left* receives the value of the register on the *right* (which is **unmodified**).

34.1 Saving a working register into a save register

Mnemonic	Hex-form	Fields	Carry	Cycles
`R0=A`	100	All	no	20.5
`R1=A`	101	All	no	20.5
`R2=A`	102	All	no	20.5
`R3=A`	103	All	no	20.5
`R4=A`	104	All	no	20.5
`R0=C`	108	All	no	20.5
`R1=C`	109	All	no	20.5
`R2=C`	10A	All	no	20.5
`R3=C`	10B	All	no	20.5
`R4=C`	10C	All	no	20.5

34.2 Saving a working register into a save register within a field

Mnemonic	Hex-form	Fields	Carry	Cycles
`R0=A f`	81AF00	A	no	14
	81Aa00	P,WP,XS,S,M,B,W	no	9+*n*
`R1=A f`	81AF01	A	no	14
	81Aa01	P,WP,XS,S,M,B,W	no	9+*n*
`R2=A f`	81AF02	A	no	14
	81Aa02	P,WP,XS,S,M,B,W	no	9+*n*
`R3=A f`	81AF03	A	no	14
	81Aa03	P,WP,XS,S,M,B,W	no	9+*n*
`R4=A f`	81AF04	A	no	14
	81Aa04	P,WP,XS,S,M,B,W	no	9+*n*
`R0=C f`	81AF08	A	no	14
	81Aa08	P,WP,XS,S,M,B,W	no	9+*n*
`R1=C f`	81AF09	A	no	14
	81Aa09	P,WP,XS,S,M,B,W	no	9+*n*
`R2=C f`	81AF0A	A	no	14
	81Aa0A	P,WP,XS,S,M,B,W	no	9+*n*
`R3=C f`	81AF0B	A	no	14
	81Aa0B	P,WP,XS,S,M,B,W	no	9+*n*
`R4=C f`	81AF0C	A	no	14
	81Aa0C	P,WP,XS,S,M,B,W	no	9+*n*

Example: I want to use D1, but I need to save its value:

```
CD1EX
R0=C A
```

You see?

34.3 Copying a save register to a working register

You cannot work with values stored in the save registers, so here are the instructions to recover information stored on them:

Mnemonic	Hex-form	Fields	Carry	Cycles
A=R0	110	All	no	20.5
A=R1	111	All	no	20.5
A=R2	112	All	no	20.5
A=R3	113	All	no	20.5
A=R4	114	All	no	20.5
C=R0	118	All	no	20.5
C=R1	119	All	no	20.5
C=R2	11A	All	no	20.5
C=R3	11B	All	no	20.5
C=R4	11C	All	no	20.5

34.4 Copying a save register to a working register within a field

Mnemonic	Hex-form	Fields	Carry	Cycles
A=R0 f	81AF10	A	no	14
	81Aa10	P,WP,XS,S,M,B,W	no	9+*n*
A=R1 f	81AF11	A	no	14
	81Aa11	P,WP,XS,S,M,B,W	no	9+*n*
A=R2 f	81AF12	A	no	14
	81Aa12	P,WP,XS,S,M,B,W	no	9+*n*
A=R3 f	81AF13	A	no	14
	81Aa13	P,WP,XS,S,M,B,W	no	9+*n*
A=R4 f	81AF14	A	no	14
	81Aa14	P,WP,XS,S,M,B,W	no	9+*n*
C=R0 f	81AF18	A	no	14
	81Aa18	P,WP,XS,S,M,B,W	no	9+*n*
C=R1 f	81AF19	A	no	14
	81Aa19	P,WP,XS,S,M,B,W	no	9+*n*
C=R2 f	81AF1A	A	no	14
	81Aa1A	P,WP,XS,S,M,B,W	no	9+*n*
C=R3 f	81AF1B	A	no	14
	81Aa1B	P,WP,XS,S,M,B,W	no	9+*n*
C=R4 f	81AF1C	A	no	14
	81Aa1C	P,WP,XS,S,M,B,W	no	9+*n*

34.5 Exchanging between save and working register

Those are really cool. First we write the letter of the working register, then the name of the save register, and finally EX for "EXchange."

Mnemonic	Hex-form	Fields	Carry	Cycles
AR0EX	120	All	no	20.5
AR1EX	121	All	no	20.5
AR2EX	122	All	no	20.5
AR3EX	123	All	no	20.5

AR4EX	124	All	no	20.5
CR0EX	128	All	no	20.5
CR1EX	129	All	no	20.5
CR2EX	12A	All	no	20.5
CR3EX	12B	All	no	20.5
CR4EX	12C	All	no	20.5

34.6 Exchange between save/working register with field

Mnemonic	Hex-form	Fields	Carry	Cycles
AR0EX f	81AF20 81Aa20	A P,WP,XS,S,M,B,W	no no	14 9+*n*
AR1EX f	81AF21 81Aa21	A P,WP,XS,S,M,B,W	no no	14 9+*n*
AR2EX f	81AF22 81Aa22	A P,WP,XS,S,M,B,W	no no	14 9+*n*
AR3EX f	81AF23 81Aa23	A P,WP,XS,S,M,B,W	no no	14 9+*n*
AR4EX f	81AF24 81Aa24	A P,WP,XS,S,M,B,W	no no	14 9+*n*
CR0EX f	81AF28 81Aa28	A P,WP,XS,S,M,B,W	no no	14 9+*n*
CR1EX f	81AF29 81Aa29	A P,WP,XS,S,M,B,W	no no	14 9+*n*
CR4EX f	81AF2A 81Aa2A	A P,WP,XS,S,M,B,W	no no	14 9+*n*
CR3EX f	81AF2B 81Aa2B	A P,WP,XS,S,M,B,W	no no	14 9+*n*
CR4EX f	81AF2C 81Aa2C	A P,WP,XS,S,M,B,W	no no	14 9+*n*

35 Pointers

D0 and D1 can only contain five nibbles, and this is because all addresses are encoded using five nibbles. They are thus called pointers because according to their value, they "point" a specific nibble in memory.

A and C are also special: they are the only registers that can be used with D0 and D1.

35.1 Giving D0 or D1 a value

We can load 2, 4 or 5 nibbles in D0 or D1. If we only load two, the other three are left untouched; if we load only four, the last one is left untouched.

We will simply write:

```
D0= ..
D1= ..
```

Perhaps it may look strange to only load only two or four nibbles. In fact, sometimes you'll load a five-nibble value, and then only change what needs to be changed.

All you have to write is `D0= ..` (with `..` being 2,4 or 5 nibbles). You can do the same with D1.

The space after the = sign is important! Otherwise, HP-ASM won't recognize it!

The hex forms depend on the number of nibbles we load:

Mnemonic	Hex-form	Fields	Carry	Cycles
`D0= yz`	19zy	First 2 nibbles	no	6
`D0= wxyz`	1Azyxw	First 4 nibbles	no	9
`D0= vwxyz`	1Bzyxwv	First 5 nibbles	no	10.5
`D1= yz`	1Dzy	First 2 nibbles	no	6
`D1= wxyz`	1Ezyxw	First 4 nibbles	no	9
`D1= vwxyz`	1Fzyxwv	First 5 nibbles	no	10.5

As you can see, when you type `D1= 00120`, the Code object will be encoded as "1F02100," so the nibbles to be loaded into D1 are reversed.

I told you that you would love the Saturn…but love or hate it, I'm sure you feel *something*. ;)

35.2 Adding to or subtracting from D0 or D1

We don't always need to load a value into D0 or D1, because we can just move the pointer. Remember this piece of code?

```
A=DAT0 A
D0=D0+ 5
PC=(A)
```

The second line is the instruction used to increment D0; if a minus sign "-" is used, D0 will decrement instead.

Note: don't forget the space between the sign and the value, as it's `D0=D0+ 5` rather than `D0=D0+5.` If you forget the space, the assembler will report an error.

You can add or remove from 1 to 16 to D0 or D1; if you want to remove more, you will have to repeat these instructions as many times as necessary.

Here they are:

Mnemonic	Hex-form	Fields	Carry	Cycles
`D0=D0+` *n*	16(*n*-1)	The entire register	yes	8.5
`D0=D0-` *n*	18(*n*-1)	The entire register	yes	8.5
`D1=D1+` *n*	17(*n*-1)	The entire register	yes	8.5
`D1=D1-` *n*	1C(*n*-1)	The entire register	yes	8.5

n is an integer from 1 through 16. In addition, the carry is affected. For example, if you have #FFFFFh in the pointer register and add a value, the carry is set because there is an overflow. As usual, subtraction can affect the carry as well.

35.3 Copying A (or C) to D0 (or D1)

The A register is useful: it has 5 nibbles, just like D0 and D1! We can make D0 or D1 point to the 5 nibbles inside A or C, like this:

Mnemonic	Hex-form	Fields	Carry	Cycles
`D0=A`	130	A field to D0	no	9.5
`D0=C`	134	A field to D0	no	9.5
`D1=A`	131	A field to D0	no	9.5
`D1=C`	135	A field to D0	no	9.5

It is also possible to only copy four nibbles, leaving the fifth one untouched. In other words, the MSB nibble (MSB = Most Significant Bits) is left in D0 or D1, and four nibbles from A or C overwrite the other four nibbles.

The mnemonics here are similar, with the only difference being an "S" added to the end.

Mnemonic	Hex-form	Fields	Carry	Cycles
D0=AS	138	Four nibbles only	no	8.5
D0=CS	13C	Four nibbles only	no	8.5
D1=AS	139	Four nibbles only	no	8.5
D1=CS	13D	Four nibbles only	no	8.5

35.4 Exchanging field A of A/C with D0/D1

You already know the instructions to exchange between save registers and working registers, but here are some which exchange field A of A or C with D0 or D1.

I like these a lot, because if I want to check for the prologue of an object, I can use one of those exchange registers to modify D0 (or D1) and temporarily save the original value of D0 or D1. First the instructions are listed, and then examples are given.

Mnemonic	Hex-form	Fields	Carry	Cycles
AD0EX	132	Five nibbles	no	9.5
AD1EX	133	Five nibbles	no	9.5
CD0EX	136	Five nibbles	no	9.5
CD1EX	137	Five nibbles	no	9.5

Each object starts with a prologue, which is a five nibble binary number. The stack doesn't contain objects, but rather the addresses of objects.

Why?

As you know, when your program starts, the HP already uses D0, D1, B, and D. D0 points to the address of the next RPL object to use, D1 points to the first level of the stack, D is the free memory in five-nibble blocks (each unit of D is 5 nibbles available) and B points to the RPL return stack.

When we start a program, we can use pointers, but when we quit, we must restore the original values. There are two ways to do that: save all the registers and restore them before we leave, or temporarily save them ourselves.

You can call subroutines found in ROM from your programs, and many are very useful. When you want to save all pointers, you can simply write:

```
GOSBVL 0679B
```

We will learn more about jumps like these later. Once you have called this routine, you can do what you want with D0 or D1. This works because there is a routine at #0679Bh that will save D0, D1, B and D. Later, when we study where those values are saved, I will show you why the HP uses some part of the RAM to store some important values.

Before leaving, we can restore the pointers with this code:

```
GOSBVL 067D2
```

However, if we just want to do something very simple, we may not need to save pointers.

Let's do an example: I want a code program that drops the first level of the stack, but only if it contains a string. First, you have to learn how to do a drop. Next, you must learn how to determine the type of an object on the stack.

We know that D1 points to the first level of the stack. The stack does not contain objects, but instead it contains their addresses. This is better: when you do a SWAP, the addresses on level 1 and 2 are exchanged! The objects don't move, so we can swap the two levels by only moving 10 nibbles. This means it's just as fast when swapping large objects!

This also explains why some strange occurrences. Sometimes when there are two objects on the stack, and one is altered, the second one changes as well!

This happens because when an object is DUPed on the stack, it's the address which is duplicated, not the object. So some commands will modify the first level of stack, and because the two addresses are the same, the object on level 2 is also modified.

This is also why the NEWOB command exists: it duplicates the object and puts its address on the stack.

So, if we have three objects on the stack, the stack will actually contain:

```
D1+15  00000
D1+10  address of third object
D1+5   address of second object
D1     address of first object
```

When our program starts, D1 points to the address of the first object. If we want to have D1 point to the next object's address, we do a:

```
D1=D1+ 5
```

If we use `C=DAT1 f` or `A=DAT1 f` we will read the specified nibbles from the object on level 2.

If we want to do a drop, we can simply move D1 to the object in level 2 and free 5 nibbles of memory (since they are no longer used).

A DROP is then:

```
D1=D1+ 5
D=D+1 A
```

This moves D1 to the next object of the stack and increments D by 1, as it contains the free RAM in five-nibble blocks. We add 1 to it because that frees five nibbles of RAM, since each address on the stack uses five nibbles of memory.

But, we cannot do a drop without checking what's on the stack, as we could get a "Try to recover memory..." Assembly language is cool, but make a mistake and you could corrupt your memory.

So, here is a DROP that does not check if the stack is empty:

```
D1=D1+ 5
D=D+1 A
A=DAT0 A
D0=D0+ 5
PC=(A)
@
```

Don't forget the @, because it tells the assembler where the end of the source is.

Notice that our code moved D1 5 nibbles further (thus, the object on level 2 becomes the object on level 1) and we free 5 nibbles of memory. The three instructions that follow are the usual ones to return to RPL: we read the address of the next RPL object (`A=DAT0 A`), we update D0 (`D0=D0+ 5`), and we continue to the next RPL object (`PC=(A)`).

Compile that code, and store it inside a variable. Then, put 3 or 4 objects on stack, and run it. Each time you run it, the object on level one is dropped. Just make sure not to run it with an empty stack!

Let's do a better example: first we'll check if the stack is empty, and if it's not, we'll do a DROP:

First, we read the address on level 1 of stack. If it's equal to *zero*, the stack is empty. Since the stack contains addresses, we have to check whether 5 nibbles are null there.

We'll read the 5 nibbles, which form the address of the first object of the stack:

```
C=DAT1 A
```

Now C contains the address of the first object on the stack, or #00000h if the stack is empty. So? We need a test, so we'll test whether C is equal to #00000h, using an A field:

```
?C=0 A
```

After this test, we have to tell the assembler what should be done whether the test is true (C=0) or not (C different than 0).

The mnemonic used is "GOYES" followed by a label.

So, write:

```
?C=0 A
GOYES QUIT
...
*QUIT
...
```

If field A of C is equal to zero, we'll jump to the label QUIT. Remember a label is a line that starts with a *.

If C is not equal to zero, the execution continues at the line just after the GOYES.

Here it is:

```
C=DAT1 A      % read what is on level 1
?C=0 A        % if it's equal to zero, we jump to QUIT
GOYES QUIT    % otherwise, we have an object on level 1:
D1=D1+ 5      % we move D1 5 nibbles further
D=D+1 A       % and we free 5 nibbles of memory
*QUIT         % here, we QUIT:
A=DAT0 A      % we read the next RPL object's address
D0=D0+ 5      % update D0
PC=(A)        % and we jump to the object
@             % end of source :-)
```

As you can see, we use the % symbol to add comments.

Programmers use comments to remember what each piece of their code does. It's also very useful when you give code you wrote to someone else. The assembler will ignore everything that follows the % until the line ends.

We could have commented and spaced it like this instead:

```
% check if level 1 is empty

C=DAT1 A
?C=0 A
GOYES QUIT

% DROP
```

```
D1=D1+ 5
D=D+1 A

% we QUIT here

*QUIT
A=DAT0 A
D0=D0+ 5
PC=(A)

@
```

You should put comments in your sources. That will help you know where you are, and when you use old source again, like from several weeks earlier, you don't have to research what each piece of code does.

Now we can write the code we originally wanted: drop the first level of the stack if it's a string.

As before, we first must read the address of the object on the first level of stack. If it's null, we'll quit:

```
C=DAT1 A
?C=0 A
GOYES QUIT
```

Now, we have to ensure that the object on level 1 is a string. Every object on the HP is identified by a 5-nibble prologue at its start, so that we can simply check the prologue to determine the object type.

Here we'll need to use D0 or D1 to point to the object, but we can lose neither D0 nor D1's value, so we can use an exchange instruction:

```
CD1EX
```

Now C contains the D1 value, and D1 contains C. Because C contained the address of the object on level 1 of the stack, D1 now points to the object.

The prologue is the first five nibbles of the object, so we'll use A and read five nibbles:

```
A=DAT1 A
D1=C
```

We have now read the object's prologue and have returned D1 to its previous value, which was saved in C with `CD1EX`.

Here we used `CD1EX` to have D1 point to something else while not losing the original value of D1, as it was stored in C. We used A so C keeps the D1 value.

Now we have to compare A to the string's prologue value. The string prologue is #02A2Ch, and strings always start with these 5 nibbles. (In the next part we will discuss all objects, explaining their prologues and how they're encoded.)

All we have to do here is load #02A2Ch into C and compare A to C using an A field. If A=C, we can drop the object, but otherwise we quit.

Here is the final code to perform our drop while checking for a string on level 1 of the stack:

```
C=DAT1 A        % read level 1's address
?C=0 A          % equal to zero? empty stack --> we quit
GOYES QUIT
CD1EX           % D1 points to the object, C contains old D1
A=DAT1 A        % read the first 5 nibbles, i.e. prologue
D1=C            % restore D1
LC 02A2C        % load #O2A2Ch in C, string's prologue
```

```
?A#C A          % # means: different; check if A not equal to C
GOYES QUIT      % if different, it's not a string --> we quit
D1=D1+ 5        % here we do the DROP, D1+5
D=D+1 A         % and we free 5 nibbles of RAM
*QUIT           % here we quit:
A=DAT0 A
D0=D0+ 5        % usual RPL jump...
PC=(A)
@
```

Try it: it'll work!

If stack is empty, it does nothing. If the stack contains a string, it's dropped, but if it's not a string, nothing is done.

Note: if HP-ASM tells you it doesn't understand "?A#C A" it's because 'HP' is set inside OPT. When HP is set, the not-equal-to symbol is = with a line through it. When PC is set, the not equal to sign is #.

The HP not equal to sign is [ALPHA][Right-Shift][1]. The # is [Right-Shift][Divide key].

Now let's continue with instructions:

35.5 Exchanging 4 LSB nibbles A of A/C with D0/D1

The instructions below exchange the 4 LSB nibbles (the first four nibbles) of working registers A or C with D0 or D1. The MSB nibble is left untouched:

Mnemonic	Hex-form	Fields	Carry	Cycles
AD0XS	13A	Four nibbles only	no	8.5
AD1XS	13B	Four nibbles only	no	8.5
CD0XS	13E	Four nibbles only	no	8.5
CD1XS	13F	Four nibbles only	no	8.5

Those are similar to the previous ones, but here only four nibbles are exchanged.

36 Reading memory within a field

We know how to modify D0 or D1 so it can point to where we want in memory. We also know how to load values into D0 or D1 (using D0/D1= ..). Here, only registers A and C will be used, and we're going to describe the DAT0 and DAT1 instructions we have used.

The following line of code will copy 16 nibbles from the data that D1 points to into C:

```
C=DAT1 W
```

This code copies 12 nibbles, or the mantissa field (M) from working register A to the data pointed to by D0:

```
DAT0=A M
```

Here are the mnemonics:

Mnemonic	Hex-form	Fields	Carry	Cycles
A=DAT0 A	142	A	no	23.5,3.5
A=DAT0 B	14A	B	no	19.5
A=DAT0 f	152a	P,WP,X,XS,S,W,M,B	no	20+*n*,1+*n*/2

A=DAT1 A	143	A	no	23.5,3.5
A=DAT1 B	14B	B	no	19.5
A=DAT1 f	153a	P,WP,X,XS,S,W,M,B	no	20+*n*,1+*n*/2
C=DAT0 A	146	A	no	23.5,3.5
C=DAT0 B	14E	B	no	19.5
C=DAT0 f	156a	P,WP,X,XS,S,W,M,B	no	20+*n*,1+*n*/2
C=DAT1 A	147	A	no	23.5,3.5
C=DAT1 B	14F	B	no	19.5
C=DAT1 f	157a	P,WP,X,XS,S,W,M,B	no	20+*n*,1+*n*/2

The register that receives the nibbles is located to the left of the equal sign (=).

It's not the *value* of the pointer that is moved, but rather the nibbles *where* the pointer D0 or D1 points.

Example: If I want to read 5 nibbles at #00100h, I do:

```
D0= 00100
C=DAT0 A
```

D0 contains #00100h, but C will contain the five nibbles that are at #00100h in memory.

As you see in the table above, fields A and B have special forms in their hexadecimal form, which is the form that is encoded inside the Code object. Reading two nibbles, field B, from D0 can be encoded using one of two different ways in the code object. This means the mnemonic `C=DAT0 B` can be encoded as either 14E or 1566.

Using 14E inside the code is both smaller and faster: 1566 will need 22 cycles but 14E will only need 19.5 cycles. The Saturn processor, when reading or writing data, is somewhat specialized for using A and B fields. In a way, it "prefers" these two fields.

37 Reading from memory with a value

We can also use a number in place of a field, so instead of:

```
C=DAT0 A
```

one can use:

```
C=DAT0 5
```

What is the difference? In the code object it's not encoded the same way. In addition, `C=DAT0 A` needs 23.5 cycles, and `C=DAT0 5` needs 24, so as you see, using A here is better! But, using `C=DAT0 16` is better than using `C=DAT0 W`, as `C=DAT0 16` needs 35 cycles, and `C=DAT0 W` needs 36 cycles. That's only a difference of one cycle, but when used in a loop that is run, for example, 100 times, there are 100 fewer cycles used!

It's a little bit more complicated to give the cycles with the instructions, but when you want the fastest possible routine, you'll use those cycles to choose which instructions to use.

Here they are:

Mnemonic	Hex-form	Fields	Carry	Cycles
A=DAT0 *n*	15A(*n*-1)	n nibbles	no	19+*n*
A=DAT1 *n*	15B(*n*-1)	n nibbles	no	19+*n*
C=DAT0 *n*	15E(*n*-1)	n nibbles	no	19+*n*
C=DAT1 *n*	15F(*n*-1)	n nibbles	no	19+*n*

38 Writing to memory within a field

Here we copy nibbles from a working register (A or C) to the memory:

Mnemonic	Hex-form	Fields	Carry	Cycles
`DAT0=A A`	140	A	no	19.5
`DAT0=A B`	148	B	no	16.5
`DAT0=A f`	150a	P,WP,X,XS,S,W,M,B	no	19+*n*
`DAT1=A A`	141	A	no	19.5
`DAT1=A B`	149	B	no	16.5
`DAT1=A f`	151a	P,WP,X,XS,S,W,M,B	no	19+*n*
`DAT0=C A`	144	A	no	19.5
`DAT0=C B`	14C	B	no	16.5
`DAT0=C f`	154a	P,WP,X,XS,S,W,M,B	no	19+*n*
`DAT1=A A`	145	A	no	19.5
`DAT1=A B`	14D	B	no	16.5
`DAT1=A f`	155a	P,WP,X,XS,S,W,M,B	no	19+*n*

As you can see here, fields A and B have specific and optimized instructions. As you know, A is useful for addresses. B, with two nibbles, is commonly used for bytes. As we'll learn later, characters are also encoded using two nibbles.

39 Writing to memory with a value

Like before, instead of using a field, we give the value we want to read:

Mnemonic	Hex-form	Fields	Carry	Cycles
`DAT0=A` *n*	158(*n*-1)	*n* nibbles	no	18+*n*
`DAT1=A` *n*	159(*n*-1)	*n* nibbles	no	18+*n*
`DAT0=C` *n*	15C(*n*-1)	*n* nibbles	no	18+*n*
`DAT1=A` *n*	15D(*n*-1)	*n* nibbles	no	18+*n*

Suppose you want to write one nibble:

```
DAT0=A P
```

If P is not equal to zero (because you use it for something else), you can write:

```
DAT0=A 1
```

This will write one nibble, and it will be encoded as "1590" inside the code object (159 and 1-1 = 1590)

That's not very difficult now, is it?

40 Jumping and tests

As I've told you (but not really explained) every instruction is located at some address in memory, which means we can do jumps inside our programs. Usually, a program's instructions are executed one after another. But sometimes, after a test or just because we want the code to continue somewhere else, we will want to jump to some other part.

A jump is also called “branching” because we “branch” to some other place in memory.

Register PC is used for that here. You know it contains the next instruction to be executed, so if we modify PC, we can change the flow of the instruction’s execution.

We will discuss three kinds of jumps:

a) The *relative* jump.
   This jumps +n or -n nibbles from our present position. We don’t really know precisely where in memory we are going to jump, but we know how many nibbles before or after it will be.
   This introduces a new word to you: the *offset*. An offset is the numerical value that represents the displacement.
   It’s a *signed* value, which means it can be either *positive* or *negative*. Thus, we can jump either backward or forward.
   For example, if we want to jump 100 nibbles further, the offset is +100.

b) The *absolute* jump
   This jumps to a specific address. As an address is a 5-nibble number, like #ABCDE, an absolute jump means jumping to address ABCDE and continuing execution there.

c) The *indirect* jump
   More complicated (more fun). It’s not jumping to an address, but reading one address, to which we will jump, at another address.
   For example, if I do an *absolute* jump to ABCDE, I will go directly to ABCDE in memory and continue execution there.
   If I do an *indirect* jump, I will read 5 nibbles at #ABCDE and then jump to the address found in those 5 nibbles.

When you put jumps in your code using labels, the assembler will automatically generate the jumps for you. If you jump forward, it directly codes the offset, and if you jump backwards, it codes the complement of 2 of the address to indicate a negative offset, thus jumping backwards in memory.

For example, suppose I have the following jump in my code:

```
?C=0 A
GOYES QUIT
...
*QUIT
```

The assembler will generate the ... code and then calculate the offset so that the `GOYES` jumps to the code that is encoded after the *QUIT line:

Memory address	Encoded in code object	Instructions
#n	8AAzy	?C=0 A
#n + ...	...	…
#n + yz	remaining code	*QUIT label

The assembler will calculate the offset and encode it using two nibbles (yz). The `?C=0 A` will be encoded as 811zy inside the code object.

If the jump is going backwards, the assembler will use the complement of 2 to code the jump value.

Why is this important?

Because the jumps after a test are encoded using *two signed nibbles* inside the Code object, it means sometimes it won’t be able to do the jump because the offset is too big. You’ll have to write something different for your code to work. This will be explained later.

41 Relative, absolute, and indirect jumps

41.1 Conditional jumps: GOC and GONC

A conditional jump is not an automatic jump: the jump will occur according to the value of something else: here, the carry.

Two instructions can be used to jump according to the carry state: `GOC` and `GONC`

`GOC` means Go On Carry
`GONC` means Go On No Carry (or Not Carry)

Example: Suppose I remove 16 from A(A) (field A of working register A). If the carry is set, it means that A was less than 16 (0 to 15). The following code will check if an underflow happens:

```
C=C-16 A
GOC ERROR
...
*ERROR
% underflow occurred!
```

GOC will only jump if the carry bit is set. If it's not, the execution continues just after the GOC ... line.

GONC will only jump if the carry is not set. If it's set, it will not jump.

NOTE: these checks do not change the carry value. They check whether the carry is set but don't change its state. After the jump, the carry is still set until an instruction that modifies it clears it.

Mnemonic	Hex-form	When?	Cycles
`GONC` label	5yz	If carry is clear	4.5 or 12.5
`GOC` label	4yz	If carry is set	4.5 or 12.5

There are two cycle values: 3 cycles are needed if the jump is not made, and 10 cycles are needed if the jump is made.

As you can see here, two nibbles, yz, are used to code the relative jump. This means we have eight bits to code the offset. The value of these two nibbles is *signed* using complement of 2, meaning that GONC and GOC can jump up to 128 nibbles *backwards*, or 127 *forward*.

If the label more than 127 nibbles ahead or 128 nibbles before, the assembler will tell you the jump is too long.

Thus, this:

```
GOC ifCarry
% if no carry here..
```

will become:

```
GONC noCarry
GOTO ifCarry
*noCarry
% we continue here if no carry
```

The `GOTO` uses one more nibble to code jumps for a total of 3 nibbles. If the assembler still says the jump is too long, we will use `GOVLNG`, which can jump to *any* point in memory since it uses 5 nibbles to code the jump address. This works because, as you should remember, all addresses use a maximum of 5 nibbles.

Remember that `GOC` and `GONC` are only able to jump 128 nibbles before and 127 nibbles after the present position.

If you need a bigger jump, you must change the test, using GOTO, GOLONG or GOVLNG (GO Very LoNG) instructions, which are described in the next section.

41.2 Unconditional jumps: GOTO, GOLONG, GOVLNG

Unconditional means the jump will be done without any checking, so one could call it a "blind jump."

A difference between GOTO, GOLONG and GOVLNG is the number of nibbles needed to code jumps: 3, 4 or 5.

41.2.1 GOTO

Here it is:

Mnemonic	Hex-form	When?	Cycles
GOTO label	6xyz	Always (blind jump)	14

GOTO will make the program continue the execution from the LABEL point inside of the source code. It needs 14 cycles to complete. Because 3 nibbles are used to code jumps, it can jump 2048 nibbles before the GOTO position or 2047 nibbles after.

41.2.2 GOLONG

Mnemonic	Hex-form	When?	Cycles
GOLONG label	8Cwxyz	Always (blind jump)	17

This is a bigger jump, using four nibbles to code jumps, so we can jump 32768 nibbles before the GOLONG's position or 32767 nibbles further.

When you code, first try GOTO. If the assembler says the jump is too long, try GOLONG. If it still says the jump is too long, you'll have to use GOVLNG (as this one uses 5 nibbles, it won't complain). This is to use the instructions that take both the fewest cycles and fewest nibbles.

41.2.3 GOVLNG

Mnemonic	Hex-form	When?	Cycles
GOVLNG label	8Dvwxyz	Always (blind jump)	18.5

Here, as all addresses are coded using 5 nibbles, we do a direct jump to any point of memory. In this case it's not an offset encoded, but rather an address.

41.2.4 Unconditional jump to the address in A or C

Two instructions are used here: PC=A and PC=C. We directly modify the PC register and load A or C inside of it. The next instruction will be fetched from the address in A or C.

Mnemonic	Hex-form	When?	Cycles
PC=A	81B2	Field A of register A	26,3.5
PC=C	81B3	Field A of register C	26,3.5

Five nibbles are moved from A or C directly to PC.

41.2.5 Unconditional jump with exchange

Here, instead of copying field A of A or C to PC, we exchange the values. PC makes the flow of instructions continue to the address in field A of A or C, and the previous value of PC is stored inside A or C, because it's an exchange.

Mnemonic	Hex-form	When?	Cycles
APCEX	81B6	Field A of register A	19
CPCEX	81B7	Field A of register C	19

41.2.6 Unconditional and indirect jump

Now for the indirect jump we talked about earlier. We first read five nibbles at the address in A or C and then load them into PC, thus "jumping" there.

Mnemonic	Hex-form	When?	Cycles
PC=(A)	808C	Field A of register A	26,3.5
PC=(C)	808E	Field A of register C	26,3.5

Example: if at address #ABCDE inside A or C there are five nibbles #FGHIJ, then execution continues at #FGHIJ

41.2.7 Saving PC contents

Those two instructions are not used to do jumps, but are useful to save the PC value before a jump. Here they are:

Mnemonic	Hex-form	When?	Cycles
A=PC	81B4	Field A of register A	11
C=PC	84B5	Field A of register C	11

42 Calling a subroutine

When you create a subroutine, it's usually to save space inside of your code, when some piece of code is called a lot. If you need to call a routine (a piece of code) multiple times it's better (if you want to save space) to put it inside a subroutine. It's just a label followed by the code that will be called. When you want, you can return back to where you called the subroutine, and continue running your program.

When jumping to a subroutine, the processor loads the address of the instruction that follows the jump instruction into RSTK (the Return STacK area). Then, the execution continues to where you jumped, until a return instruction is found.

When the return instruction is found, it's loaded into PC, and the program "returns" from the subroutine and execution continues.

The difference here between a usual jump is that an address must be recorded so we can go back.

RSTK is eight levels high, with each level being five nibbles wide, but only five levels can be used. This means you can have five levels of subroutines, or eight if you don't allow *any* interruptions (more on that later).

There are several return instructions available. Some will just return, others will return and modify the carry bit (set it or clear it) and so on.

A subroutine is a *relative* jump, so it's an *offset* which is coded (unless you use GOSBVL, GOSuB Very Long, which encodes a specific five-nibble address).

42.1 GOSUB

Here three nibbles are used to encode the offset of the jump, so we can jump to a subroutine 2048 nibbles before or 2047 nibbles after the GOSUB's position).

Mnemonic	Hex-form	How?	Cycles
GOSUB label	7xyz	1 RSTK level used	15

42.2 GOSUBL

This stands for GOSUB Long. Here we use four nibbles to encode the offset, so we can jump to a subroutine 32768 nibbles before or 32767 nibbles after.

Mnemonic	Hex-form	How?	Cycles
GOSUBL label	8Ewxyz	1 RSTK level used	18

42.3 GOSBVL

Here we directly encode the five nibbles of the address:

Mnemonic	Hex-form	How?	Cycles
GOSUBVL label	8Fvwxyz	1 RSTK level used	19.5

42.4 Returning from a subroutine

There are several instructions to do this:

Mnemonic	Hex-form	How?	Cycles
RTN	01	Usual return	11
RTNSC	02	Sets the carry bit as well	11
RTNCC	03	Clears the carry bit instead	11
RTI	0F	Allows interrupts	11
RTNSXM	00	Sets XM bit	11

RTN simply removes the return address from RSTK and jumps back just after the instruction that called the subroutine.

RTNSC and RTNCC will return and set or clear the carry bit, so it can be used to call a subroutine that will use the carry bit to give you information according to the state of the carry bit.

RTI will allow interrupts again. We'll understand this better when we study how interruptions work, and how they're disallowed or re-allowed by the routines in the HP's ROM (we'll disassemble all that, yes!!)

RTNSXM will return and set the XM (eXternal Module missing) bit.

42.5 Returning according to the value of the carry bit

You can choose whether to return depending on the value of the carry bit. The instructions are:

Mnemonic	Hex-form	When?	Cycles
RTNC	400	If carry is set	4.5 or 12.5
RTNNC	500	If carry is cleared	4.5 or 12.5

RTNC will return only if the carry is set. If it's clear, the testing uses 4.5 cycles; if it's set, the return needs 12.5 cycles.

RTNNC will return only if the carry is cleared. 4.5 cycles are needed if the return is not done and 12.5 cycles are used if it returns.

43 The tests

When we do a test, we compare a register to zero or to an another register.

Here you will see two values for the cycles needed. The first one is the cycles used for the test to complete if no jump occurs, and the second is the number of cycles used if the test is completed and the jump is made.

Each test will start with a question mark (?). Next comes a working register (A or C), and then an operator (=, >, <, etc.).

On the line that follows a test, there must be an *action* command. It can be a jump, using GOYES, but it can also be a return from a subroutine if we use RTNYES.

43.1 Comparing registers to zero

43.1.1 Equal to zero?

Here we check whether a working register is equal to zero. The symbol to check equality is =.

Mnemonic	Hex-form	Fields	Carry	Cycles
?A=0 A	8A8 yz	A	yes	13.5/21.5
?A=0 f	9a8 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
?B=0 A	8A9 yz	A	yes	13.5/21.5
?B=0 f	9a9 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
?C=0 A	8AA yz	A	yes	13.5/21.5
?C=0 f	9aA yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
?D=0 A	8AB yz	A	yes	13.5/21.5
?D=0 f	9aB yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*

In the resulting code object, the "yz" is an offset (signed) that is used for the jump if the test is true.

If the test is true then the carry bit is set.

These instructions need 8.5+*n* cycles (*n* = number of nibbles of the field used for comparison) if the test result is false and 16.5+*n* if the test result is true and the jump is made.

You may use "GOYES label" or "RTNYES" after those instructions. GOYES will make the flow of execution jump to the label of your choice inside your code, and RTNYES will return from a subroutine.

Two nibbles are used to code the offset, and it's a signed value, so how far can it jump? Well! Good! 128 nibbles before or 127 nibbles after the jump.

43.1.2 Different than zero?

The symbol for "different" depends on the mode you set in the HP-ASM options: with HP it will be a = with a line through it (not equal to sign) and with "PC" it will be a pound sign (#). I strongly encourage you to use #, so you'll be able to write your code on your PC/Mac/Linux system and send the source to your HP, rather than have to find/replace symbols for another.

Mnemonic	Hex-form	Fields	Carry	Cycles
`?A#0 A`	8AC yz	A	yes	13.5/21.5
`?A#0 f`	9aC yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
`?B#0 A`	8AD yz	A	yes	13.5/21.5
`?B#0 f`	9aD yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
`?C#0 A`	8AE yz	A	yes	13.5/21.5
`?C#0 f`	9aE yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
`?D#0 A`	8AF yz	A	yes	13.5/21.5
`?D#0 f`	9aF yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*

The carry will be set if it tests to be true, otherwise it's clear.

43.2 Equality of two registers

Here the test will be true if equal or false if not equal.

Mnemonic	Hex-form	Fields	Carry	Cycles
`?A=B A`	8A0 yz	A	yes	13.5/21.5
`?A=B f`	9a0 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
`?B=C A`	8A1 yz	A	yes	13.5/21.5
`?B=C f`	9a1 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
`?A=C A`	8A2 yz	A	yes	13.5/21.5
`?A=C f`	9a2 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
`?C=D A`	8A3 yz	A	yes	13.5/21.5
`?C=D f`	9a3 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*

43.3 Inequality of two registers

Mnemonic	Hex-form	Fields	Carry	Cycles
`?A#B A`	8A4 yz	A	yes	13.5/21.5
`?A#B f`	9a4 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
`?B#C A`	8A5 yz	A	yes	13.5/21.5
`?B#C f`	9a5 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
`?A#C A`	8A6 yz	A	yes	13.5/21.5
`?A#C f`	9a6 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
`?C#D A`	8A7 yz	A	yes	13.5/21.5
`?C#D f`	9a7 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*

43.4 Less than and greater than

Mnemonic	Hex-form	Fields	Carry	Cycles

?A>B A	8B0 yz	A	yes	13.5/21.5
?A>B f	9b0 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
?B>C A	8B1 yz	A	yes	13.5/21.5
?B>C f	9b1 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
?C>A A	8B2 yz	A	yes	13.5/21.5
?C>A f	9b2 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
?D>C A	8B3 yz	A	yes	13.5/21.5
?D>C f	9b3 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
?A<B A	8B4 yz	A	yes	13.5/21.5
?A<B f	9b4 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
?B<C A	8B5 yz	A	yes	13.5/21.5
?B<C f	9b5 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
?C<A A	8B6 yz	A	yes	13.5/21.5
?C<A f	9b6 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
?D<C A	8B7 yz	A	yes	13.5/21.5
?D<C f	9b7 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*

43.5 Less than or equal to or greater than or equal to

Mnemonic	Hex-form	Fields	Carry	Cycles
?A>=B A	8B8 yz	A	yes	13.5/21.5
?A>=B f	9b8 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
?B>=C A	8B9 yz	A	yes	13.5/21.5
?B>=C f	9b9 yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
?C>=A A	8BA yz	A	yes	13.5/21.5
?C>=A f	9bA yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
?D>=C A	8BB yz	A	yes	13.5/21.5
?D>=C f	9bB yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
?A<=B A	8BC yz	A	yes	13.5/21.5
?A<=B f	9bC yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
?B<=C A	8BD yz	A	yes	13.5/21.5
?B<=C f	9bD yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
?C<=A A	8BE yz	A	yes	13.5/21.5
?C<=A f	9bE yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*
?D<=C A	8BF yz	A	yes	13.5/21.5
?D<=C f	9bF yz	P,WP,X,XS,S,W,M,B	yes	8.5+*n*/16.5+*n*

As usual, after each of these instructions you can put a GOYES or put a RTNYES. The jump is done is the test is true.

43.6 Testing a bit

Here we have the instructions that can test the value of a bit in C or A. Only the first 16 bits, or four nibbles, of the working registers can be tested.

Twelve cycles are needed to finish the test and continue if it's false, and 21 cycles are needed if it's true and a jump is performed.

Two nibbles are used inside the code (#yz) to encode the relative jump offset, so we can jump 127 nibbles before or 128 nibbles after the test's position.

Mnemonic	Hex-form	Fields	Carry	Cycles
?ABIT=0 *n*	8086n yz	Bit number n	yes	12.5/20.5
?ABIT=1 *n*	8087n yz	Bit number n	yes	12.5/20.5
?CBIT=0 *n*	808An yz	Bit number n	yes	12.5/20.5

?CBIT=1 *n*	808Bn yz	Bit number n	yes	12.5/20.5

44 Using register P

The P register can only contain one nibble. When your program starts, its value will be zero. P not only defines how the loading instructions work (LC and LA), but it also defines the width of the WP register (from nibble 0 to nibble P).

Often, the P register is used as a loop counter, as long there is no LC, LA, or WP-related instruction inside of the loop. Not only are there instructions to give P a value, but there are also instructions to increment or decrement its value.

You can also test the P register, exchange it with the C register, and other things.

The mnemonics are below, followed by an explanation of each:

Mnemonic	Hex-form	Fields	Carry	Cycles
P= *n*	2n	The nibble of P	no	3
P=P+1	0C	The nibble of P	yes	4
P=P-1	0D	The nibble of P	yes	4
?P# *n*	88n yz	The nibble of P	yes	7.5 or 15.5
?P= *n*	89n yz	The nibble of P	yes	7.5 or 15.5
C=P *n*	80Cn	The nibble of P	no	8
P=C *n*	80Dn	The nibble of P	no	8
C+P+1	809	The nibble of P	yes	9.5
CPEX *n*	80Fn	The nibble of P	no	8

44.1 P= *n*

Don't forget the space after the "=" sign!! It gives the value n to P.

44.2 P=P+1

This increments the value of P by 1. NOTE: whichever mode the Saturn is working under, hexadecimal or decimal, the incrementing is done under *hexadecimal* mode. The carry is set if an overflow occurs.

44.3 P=P-1

This decrements the value of P by 1. NOTE: whichever mode the Saturn is working under, hexadecimal or decimal, the incrementing is done under *hexadecimal* mode. The carry is set if an overflow occurs.

44.4 ?P# *n*

This test checks whether P is different than n. Two nibbles inside the code object (signed) will be used to encode the offset of the relative jump. A RTNYES will be coded using #00h; any other value is a signed offset, just like the previous tests. The carry will be set if the test is true.

44.5 ?P= *n*

This test will check if P is equal to n. The carry will be set if the test is true.

44.6 C=P *n*

Gives the value of the P register to nibble number n of working register C. *n* is between 0 and 15, of course.

44.7 P=C *n*

Gives the value of nibble number *n* of working register C to P.

44.8 C+P+1

Adds the value of C(A), P and 1 to field A of C. Strange instruction? You will see it can be very useful. :)

44.9 CPEX *n*

This instruction exchanges nibble number *n* of working register C with the P register.

45 The RSTK stack

RSTK is the return stack where return addresses are stored while subroutines are running. You can also use RSTK to save the working register C for some time, but don't forget to pop the saved value before you quit. It's faster than using a save register, so it can be cool to use.

Mnemonic	Hex-form	Fields	Carry	Cycles
`RSTK=C`	06	Field A of C	no	9
`C=RSTK`	07	Field A of C	no	9

46 Registers IN and OUT

Inputs and outputs done by the processor go through registers IN and OUT. Here the X field will be quite useful because the OUT register is 3 nibbles wide. The IN register is 4 nibbles wide.

We have here four instructions:

Mnemonic	Hex-form	Fields	Carry	Cycles
`OUT=CS`	800	Nibble #0h of C	no	5.5
`OUT=C`	801	Field X of C	no	7.5
`A=IN`	802	Four nibbles of A	no	8.5
`C=IN`	803	Four nibbles of A	no	8.5

ANOTHER BUG

You cannot use `A=IN` or `C=IN` instructions if they are not located at an even address (or odd, I don't remember and that's not very important) because they won't work.

So, instead of writing `A=IN` or `C=IN` in source code, we will always use subroutines in ROM that will do `A=IN` or `C=IN` and then return to our program.

We'll use:

```
GOSBVL 01160 instead of C=IN
GOSBVL 0115A instead of A=IN
```

Those routines contain:

```
C=IN          A=IN
RTN           RTN
```

---

47 Status Bits (ST)

We have many flags we can use in our programs. There are four nibbles inside of ST, that's 4 * 4 = 16 flags; but some are already used by the Saturn. Until you learn how to use them well, we'll avoid playing with bits 12, 13 and 14 (we'll use bit 15 to enable or disable interrupts, but be patient as we'll come to that soon).

Those status bits are like the flags you use in RPL.

Here are the instructions, and explanations follow:

Mnemonic	Hex-form	Fields	Carry	Cycles
`CLRST`	08	First 3 nibble of ST	no	7
`C=ST`	09	X field of C	no	7
`ST=C`	0A	X field of C	no	7
`CSTEX`	0B	X field of C	no	7
`ST=0` *n*	84n	Bit n of ST	no	5.5
`ST=1` *n*	85n	Bit n of ST	no	5.5
`?ST=0` *n*	86n yz	Bit n of ST	yes	8.5 or 16.5
`?ST=1` *n*	87n yz	Bit n of ST	yes	8.5 or 16.5

47.1 CLRST

Clears the first 3 nibbles of ST. The last one is not because it does contain special values for the processor.

47.2 C=ST

Copies 3 nibbles of ST to C.

47.3 ST=C

Copies 3 nibbles from C to ST.

47.4 CSTEX

Exchanges 3 nibbles from C and ST.

47.5 ST=1 *n*

Sets bit n of ST to 1.

47.6 ST=0 *n*

Clears bit n of ST, therefore setting it to 0.

47.7 ?ST=0 *n* and ?ST=1 *n*

Those two instructions are used to check the value of a single bit of the ST register. We can test if it's set or cleared. You can put a `GOYES` or a `RTNYES` with a label after those instructions if you want.

48 Hardware Status Bits (HST)

There are eleven instructions that can be used to affect the HST bits. The useful bits of HST are:

- SR (Service Request)
- SB (Sticky Bit)
- MP (Module Pulled)
- XM (eXternal Module missing)

Mnemonic	Hex-form	Fields	Carry	Cycles
`CLRHST`	82F	Four bits of HST	no	4.5
`XM=0`	821	XM bit	no	4.5
`SB=0`	822	SB bit	no	4.5
`SR=0`	824	SR bit	no	4.5
`MP=0`	828	MP bit	no	4.5
`HST=0` *n*	82n	1 or more bits	no	4.5
`?XM=0`	831	XM bit	yes	7.5 or 15.5
`?SB=0`	832	SB bit	yes	7.5 or 15.5
`?SR=0`	834	SR bit	yes	7.5 or 15.5
`?MP=0`	838	MP bit	yes	7.5 or 15.5
`?HST=0` *n*	83n	1 or more bits	yes	7.5 or 15.5

That's a lot to remember, isn't it?

`CLRHST` will clear the four bits of HST. You can also use other instructions, like `XM=0` or `SB=0` to clear a single bit of HST. There are four bits, thus four instructions to clear the bits.

There is also another way to set some bits of HST to zero using a formula, like this:

value = XM + (2*SB) + (4 * SR) + (8 * MP)

Where 1 means to **clear** a bit and 0 means to set a bit, if you want to have:

XM = 0
SB = 0

SR = 1
MP = 1

You will calculate:

value = 0 + (2 * 0) + (4 * 1) + (8 * 1)
= 0 + 0 + 4 + 8
= 12

So by doing `HS=0 12` one can clear only the SR and MP bits. Note that an `HS=1 n` instruction does not exist.

There are instructions to test each of the four bits available, and the instruction `?HS=0 n` can be used to test whether some (or all) bits of HST are equal to zero, using the same formula as above. In this case, 1 in the formula also means that the test will check whether the bit is **clear**.

After those test instructions, you can put a `GOYES` or `RTNYES` if the test is true.

These instructions need six cycles to perform the test and continue if it's false or 13 cycles is the test is true and the jump is done. It may be a jump to a label or a return using RTNYES if we are inside a subroutine.

For example, let's write some code to return from a subroutine if bit XM is set to 1:

value = 1 + (0 * 2) + (0 * 4) + (0 * 8)
= 1

This means the code can be written like this:

```
?HST=0 1
RTNYES
```

If we want to jump to the label "REACT" if XM=1 and SR=1, we first must calculate the value:

value = 1 + (0 * 2) + (1 * 4) + (0 * 8)
= 1 + 4
= 5

So our code looks like this:

```
?HST=0 5
GOYES REACT
```

To test the SB (Sticky Bit) after a rotation of bits or nibble on a working register we have:

```
?HST=0 2
GOYES RESULT_LOSS
```

etc.

49 Saturn DEC/HEX mode

The Saturn processor can work under two modes: hexadecimal mode (base 16) or decimal mode (base 10). As we studied before, the latter uses compacted BCD.

Remember that with BCD, 4 bits are used for each decimal number to encode. Look back a few dozen pages for more about this.

Mnemonic	Hex-form	Fields	Carry	Cycles
SETHEX	04	Hexadecimal mode	no	4
SETDEC	05	Compact BCD mode	no	4

50 Interruptions

Though only instructions are listed here, we'll discuss interruptions a lot more later.

Mnemonic	Hex-form	Fields	Carry	Cycles
INTOFF	808F	Disable keyboard interrupts	no	7
INTON	8080	Enable keyboard interrupts	no	7
RSI	80810	Reset interrupts	no	8.5

INTOFF and INTON simply disable or enable the interrupts linked to the keyboard. To truly disable all interruptions, we have to set bit 15 of ST to 0 using ST=0 15. In order to allow them again, we must use ST=1 15, INTON and RSI (ST=1 15 should be sufficient, but the HP does much more in its ROM).

51 Bus-related instructions

We will discuss the bus in more depth later, but you can use this section as a reference when learning all about the bus. The bus is linked to the processor, and there is much to learn about modules, memory configurations, and more. That could be a painful part...

The Saturn only has one bus. Usually one might find separate buses, for example, for data and instructions, but because the Saturn only has one, it is "multiplexed." The instructions here will be useful for configuring or deconfiguring a specific module of memory, so some parts are visible and others not.

There are three instructions that don't seem to be used on the HP48: BUSCB, BUSCC and BUSCD.

The mnemonic means: BUS = Bus, C = Command, followed by letter of the command.

Cyrille de Brebisson said in a post that they were not used. I don't know, but he's probably right. I have tried these under several RAM configurations and other circumstances, and they don't seem to be caught when issued on the bus. However, at least one program written by Hewlett-Packard uses one or more of these instructions.

Perhaps HP uses them to send our results to their big supercomputer that rules the world but no one knows about? ;)

Mnemonic	Hex-form	Fields	Carry	Cycles
RESET	80A		no	7.5
SREQ?	80E		no	9.5
CONFIG	805	Field A of C	no	13.5
UNCNFG	804	Field A of C	no	14.5
C=ID	806	Field A of C	no	13.5
SHUTDN	807		no	6.5
BUSCB	8083		no	10
BUSCC	80B		no	8.5
BUSCD	808D		no	10

51.1 RESET

All modules that can be configured are set to the unconfigured state once this command is used. The ROM will not be unconfigured because it has no manager.

51.2 SREQ?

Doesn't seem to be used. Perhaps Jean-Yves could tell us about this little one? ;)

51.3 CONFIG

Used to configure a memory module. This will be explained later, but you are impatient, F.H. Gilbert posted some information about it to comp.sys.hp48.

51.4 UNCNFG

Unconfigures a memory module whose base address is in field A of C.

51.5 C=ID

Copies the ID of the current module of memory to field A of C. Its content will vary depending on the status of the module. This is explained in part discussing "Covered memory."

51.6 SHUTDN

Sets the HP48 to low power usage mode. Only some essential components remain powered, and until the clock reaches zero (when an alarm is due) or the ON key is pressed ("Wake up beloved one!") the HP "sleeps."

51.7 BUSC*x*

These don't seem to be used on the HP48.

Part V: HP48 Objects

Objects are an essential part of the HP48. Everything put on the stack is an object, and they're all encoded the same way: they start with a 5-nibble prologue that identifies the object, and then the object follows in what is called the "body" of the object.

The object is, in a way, a shell, and data is put inside of it.

What is the prologue? It's not only something to identify the object, but in fact it also contains the address of a code routine that is used to execute the object. So the prologue not only identifies something but also executes it.

Some objects, when evaluated, just put themselves on the stack, like strings. If you [EVAL] a string, it remains on the stack. But, if you put a global name onto the stack and [EVAL] it, the contents of the variable that corresponds to the global name is put on the stack. The global name's prologue was evaluated, and it pushed the variable's contents onto the stack.

We are now going to study each object type of the HP48, so this is a rather long part. :-)

Perhaps you know there is a command called TYPE that returns a value, according to the object type that is on the stack. There are 28 TYPE values, but we'll soon see that's not enough. ;)

Of those 28 object types, only some can be used from RPL.

Some objects I will describe are *compound*, like the HOME directory object or the library object (the most interesting of all), and are made up of other objects, but others, like strings, are *simple*.

The following table gives each object type's name and its value given by TYPE:

Object type name	Number	Kind
Real number	0	Simple
Complex number	1	Compound
String	2	Simple
Array	3/4	Compound
List	5	Compound
Global name	6	Simple
Local name	7	Simple
RPL program	8	Compound
Algebraic expression	9	Compound
Binary integer	10	Simple
Graphic object	11	Simple
Tagged object	12	Compound
Unit object	13	Compound
XLIB name	14	Simple
Directory object	15	Compound
Library	16	Compound
Backup	17	Compound
(not used)	18	N/A
(not used)	19	N/A
System binary	20	Simple
Long real	21	Simple
Long complex	22	Compound
Linked array	23	Compound
Character	24	Simple
Code object	25	Simple
Library data	26	Simple

Extended pointer	27	Simple
Reserved 1	27*	N/A
Reserved 2	27*	N/A
Reserved 3	27*	N/A

*These are reserved by HP for future use, if needed.

52 Simple Objects

52.1 Real number

Prologue: 02933 Epilogue: none
Size: 21 nibbles Type: 0

In memory:

Prologue (02933)	5 nibbles
Exponent	3 nibbles
Mantissa	12 nibbles
Sign	1 nibble

This is a typical object type and will be used commonly. The real number is divided into three fields: the mantissa, the sign, and the exponent. The mantissa is the number, the exponent is *n* in 10^n; and the sign is coded using one nibble, where positive is #0h and negative is #9h

The mantissa uses 12 digits, and the number is encoded in memory using compacted BCD (Binary Coded Decimal). The exponent is also encoded using BCD, but in a special way according to its sign.

If the exponent is positive, the 3 nibbles are encoded using BCD.
If the exponent is negative, the exponent is encoded as follows:

```
1000 - ABS(exponent)
```

ABS means the absolute value of the exponent, whose value goes from 0 to 499

For example, let's encode 2.79233×10^{-9}. First, we have the prologue, which is #33920h (remember, it must be reversed). Then, we have the exponent, which here is 1000 - ABS(-9). That's 1000 - 9 = 991, so we have #199 when reversed. The mantissa is 279233, and because it's encoded reversed (like everything else in memory), it's #000000332972h. Finally, the sign is placed: #0h.

In memory we have #33920 199 000000332972 0h (The spaces are for clarity and are not in the real object.)

> **CAUTION!** The Saturn reverses the order of nibbles in memory. This is why each field here is reversed in memory!!

52.2 String

Prologue: 02A2C Epilogue: none
Size: 12 nibbles and up Type: 2

In memory:

Prologue (02A2C)	5 nibbles
Size	5 nibbles
Character 1	2 nibbles
…	
Character *n*	2 nibbles

The size of a string is five nibbles (the five nibbles needed to encode the size) plus two nibbles per character. As with all other objects, the prologue is not considered in the size.

Each character is encoded using 2 nibbles, using ASCII (described below). The size is everything, except the prologue, so it's five nibbles plus 2**n* nibbles (5+2*n*) long.

ASCII is a table used to code characters and stands for American Standard Code for Information Interchange. It uses eight bits (that's one byte or two nibbles) to encode values. Using eight bits lets us have up to 2^8, or 256, different values.

In fact, the original ASCII table only used 7-bits, so only 128 different symbols existed. Today, we use extended ASCII, which uses all eight bits.

The first 32 characters have special meanings, for communicating or printing; these characters, which are called "non-printable characters," include page jumps, tabs, and other codes used between devices to communicate with each other. The first 128 characters are standard, and the second 128 characters vary from country to country and from computer to computer: DOS, Windows, Macintosh, Cyrillic, European, Western, etc.

You can take a look at the HP ASCII table by pressing [Right-Shift][PRG]. At first you will see characters 128 to 191, and you can see 64 characters before using "-64" or 64 characters after using "+64."

Therefore, characters from 0 to 127 are standard, as they were the first 7-bit ASCII standard. Characters from 128 to 255 vary from one platform to another (PC, Mac, etc.) and from one alphabet to another (Roman, Cyrillic, Kanji, etc.).

For example, the string "HP" will be encoded as #02A2Ch for the string prologue plus the body of the objects. Remember we have two characters, occupying four nibbles. Five nibbles are used to encode the size (#9h nibbles), and finally the four nibbles of the characters (#4850). This means that the string object is encoded as follows:

#C2A20 90000 8405h

Note that the two nibbles for each character are reversed, just like anything else.

#48h is #72d, which is the "H" character.
#50h is #80d, which is the "P" character.

We can manipulate strings using assembly language. Let's write a program that changes all letters inside a string to UPPERCASE.

First, take a look at the CHARS table using [Right-Shift][PRG], and notice that there is a fixed distance between the 'A' and 'a' characters: 32.

If you have a lowercase character (in other words, a character whose value is between 97 and 122 inclusive), we can get its corresponding UPPERCASE value by removing 32. That's simple, isn't it? :-)

This example will be interesting: it will include a loop that uses the carry, where you'll see how to calculate the loop value, and several tests. It's a good but not too complex example. :)

(The program to turn your HP into a Star Trek phaser using the incredible boiling power of the IR diodes will come later, but be patient, as it looks like Cyrille doesn't want that secret code to be revealed!)

The code is commented:

```
% -- "Too few arguments" upon empty stack ---

C=DAT1 A                % read first level object's address
?C#0 A                  % if it's not empty (different than zero)...
GOYES STRING?           % check if it's a string
LA 00201                % load #201h into A
*ERROR                  % this label is needed if there isn't
                        % a string on the first level of stack
GOVLNG 05023            % jump to a routine that does error #201h
                        % which is "Too few arguments"

% --- we check if we have a string ---

*STRING?                % label
CD1EX                   % point D1 to the object
A=DAT1 A                % read its prologue
D1=C                    % and restore D1
LC 02A2C                % load string's prologue onto C
?A=C A                  % if the object is a string
GOYES MAIN              % jump to MAIN
LA 00202                % otherwise, we load #202h error code
GOTO ERROR              % which is "Bad argument type" and
                        % jump to ERROR, where the GOVLNG to
                        % the DOERROR routine is in ROM

% --- prepare to loop ---

*MAIN                   % kind of main routine
GOSBVL 0679B            % save registers D1,D0,B,D
C=DAT1 A                % read string's address
D1=C                    % point to it
D1=D1+ 5                % pass the prologue and point to the
                        % size nibbles
C=DAT1 A                % read size of string
D=C A                   % and store it into D

% --- calculate loop value ---

% D contains a number of nibbles.  Since there are 2 nibbles
% per char, I'm going to divide D(A) by two, which will be
% done using DSRB A (shift right to divide by two).

DSRB A                  % divide by two
D=D-1 A                 % and remove one because I'm going to use
                        % a carry loop
GOC END                 % if carry is set, the string is empty, and then
                        % jump to the end and leave
D1=D1+ 3                % move 3 nibbles forward, not 5,
                        % so when LOOP is entered, the D1=D1+ 2
                        % there points to the first char
                        % upon entering the loop, or points to the
                        % next char if it loops :-)

% --- loop ---
```

```
*LOOP
D1=D1+ 2                % point to char
C=DAT1 B                % read it into C
LA 61                   % #61h is #97d
?C<A B                  % if below 97, skip it
GOYES CHECK
LA 7A                   % #7Ah is #122d
?C>A B                  % if above 122, skip it
GOYES CHECK

% here we have a lowercase char, and remove 32 so we get
% an uppercase one:

C=C-16 B
C=C-16 B                % (char value) - 32

DAT1=C B                % write it to the string

*CHECK
D=D-1 A                 % is there an other char?
GONC LOOP               % if so, jump to LOOP

% --- quit here ---

*END

GOSBVL 067D2            % restore registers
GOVLNG 2D564            % and leave :-)

@
```

Now put a string like: “aBcDeFGhij” and run the code. You’ll get “ABCDEFGHIJ”. The Code object is only 84 nibbles long. It’s small, isn’t it? :)

(If you don’t get the same size, something may be wrong)

As practice, write a program to do the opposite: write code that turns to all characters in a string to lowercase.

52.3 Global name

Prologue:	02E48	Epilogue:	none
Size:	7 nibbles and up	Type:	6

In memory:

Prologue (02E48)	5 nibbles
Size	2 nibbles
Character 1	2 nibbles
…	
Character *n*	2 nibbles

Each character is encoded using ASCII, with two nibbles per character. On the stack, global names are stored inside single quotes (' ').

For example, 'HP48' is encoded as #02E48h for the prologue, #04h for the number of characters (up to 255), and finally two nibbles per character; that is:

```
#84E20 40 84054383
```

The HP does not let us create *any* name by default, so how creating "Invalid" Global Names? :)

Our program will take a string as parameter and put the global name on the stack. The main difference between the string and global name objects is the global name can only have up to 255 characters, as its length is encoded using two nibbles, whereas the string encodes the length using five nibbles.

NOTE:	The following program **does not** check the length of the string given in input. Don't give it a string longer than 255 characters.

If you give it an empty string, you'll get an empty global name; otherwise it places a global name with the string's contents on the stack.

```
C=DAT1 A            % read level 1's content
R0=C A              % save into R0 for future use (if needed)
?C#0 A              % if level 1 contains something
GOYES TEST          % check if it's a string
LA 00201            % if not, do error "Too Few Arguments"
*ERROR              % (used to jump here if a string isn't given)
GOVLNG 05023        % this routine makes the error, with the error
                    % code inside Aa
*TEST
CD1EX               % point to the object
A=DAT1 A            % read it's prologue
D1=C                % and restore D1 to first level of stack
LC 02A2C            % 02A2C = string's prologue
?A=C A              % if a string, jump to
GOYES ALLOCATE      % allocate memory
LA 00202            % Otherwise, error #202h, which is
GOTO ERROR          % "Bad Argument Type"
*ALLOCATE
GOSBVL 0679B        % save D0,D1,B,D
C=R0 A              % C contains the address of object in level one
D1=C                % point D1 to object
D1=D1+ 5            % point to string size
C=DAT1 A            % Aa contains string size
C=C-5 A             % remove its length field size
CSRB A              % and divide it per two, now Ca contains
                    % number of characters inside of the string
R1=C A              % save the number of chars into R1a
C=C+7 A             % add 7, with 5 nibbles for global name
                    % prologue and 2 for the length field
A=C A               % put that number into Aa
GOSBVL 039BE        % this routine is called RES_ROOM, it allocates
                    % Aa nibbles and produces an "Insufficient Memory"
                    % error if there is not enough RAM. It also checks
                    % 5 nibbles are free, so we will be able to push
                    % the reserved object onto stack.
CD0EX               % D0 = newly reserved object, now C contains its
                    % address
R2=C A              % save it into R2a (at the end we will push R2a
                    % on stack; that's why we save it)
D0=C                % D0 points to the new reserved memory
```

```
LC 02E48            % 02E48 = global name prologue
DAT0=C A            % write it
D0=D0+ 5            % and move 5 nibbles further
C=R1 B              % recover the size that was calculated
DAT0=C B            % and write it to the global name length field
C=C-1 B             % remove one, if CARRY is set
GOC PUSH            % then string is empty, so push the empty global
                    % name on stack
C=C-1 B             % remove 1 because we're going to use
                    % a carry loop
A=R0 A              % recover string's address
D1=A                % make D1 point to it
D1=D1+ 8            % and move 8 nibbles further. Why not 10?
                    % because inside the loop we have a D1=D1+ 2
                    % so finally get D1=D1+ 10 ! :-)

*LOOP               % in this loop, we copy the chars one at a time
D1=D1+ 2            % move 2 nibbles forward, if we enter the loop,
                    % we get into the first char; otherwise, we move
                    % to the next two chars.
D0=D0+ 2            % move 2 nibbles forward in the new reserved obj
A=DAT1 B            % read one char from the string
DAT0=A B            % and write them in the global name
C=C-1 B             % remove 1 from the counter
GONC LOOP           % and loop until carry is set

*PUSH               % here, push the global name on stack
GOSBVL 067D2        % recover registers D0,D1,B,D
D=D-1 A             % remove 5 nibbles from RAM
D1=D1- 5            % and push one object into the stack
C=R2 A              % recover the global name address
DAT1=C A            % push it on stack
A=DAT0 A            % read next RPL object to execute
D0=D0+ 5            % update the D0 pointer to the next object
PC=(A)              % and jump to continue HP48 program flow...

@                   % *  END OF SOURCE  *
```

Try this with an empty string! If you [EVAL] the resulting empty global name when in HOME, you will get inside the hidden directory. The hidden directory is described next:

52.3.1 The empty global name and the hidden directory

You can create an empty global name, which has a special function on the HP: if you [EVAL] it, you will get inside the hidden directory of HOME. This directory contains three objects by default: Alarms, UserKeys, and UserKeys.CRC. The first one contains a list of alarms, the second one a list with every keyboard assignment, and the last one a four-nibble binary number, which is the CRC for the UserKeys. One rule: *never, ever* change the order of the variables in this directory.

When you are inside a folder inside of HOME, you can use this empty global name to hide any variable following it in the variable list. For example, if you have four variables, and store a null global name after the second one, the two final variables become "invisible," but you can still recall their value or use them!

You probably want to see this in action. :)

First, make sure you are in HOME, by simply pressing [Right-shift][']. Then put #15777h on stack (make sure there's an h) and then type SYSEVAL [ENTER]. Now you have the null global name on stack. Just use [EVAL] to jump into it!

When you are not in HOME, you can use the null name to hide variables. Everything stored *after* the null global name disappears.

Before you forget: if you reorder the invisible directory's contents, you'll get a memory loss. You have been warned twice now.

To exit the invisible directory, do [Left-shift]['] (UPDIR) or [Right-shift]['] (HOME).

Do you want to hide the contents of a directory? Go into the desired directory, type #15777h SYSEVAL [EVAL] to put the null global name on stack, put something like 1 on the stack, and press [SWAP][STO] to store it. We now have a null global name with the value 1 on it. That's not important, because what is important is your variables are no longer visible. To see them again, put the null global name on the stack again and PURGE it.

52.4 Local name

Prologue:	02E48	Epilogue:	none
Size:	7 nibbles and up	Type:	7

In memory:

Prologue (02E6D)	5 nibbles
Size	5 nibbles
Character 1	2 nibbles
…	
Character *n*	2 nibbles

As you can see here, the prologue is the sole difference between a local name and a global name. But, the local name is not available all the time, like the global name.

52.5 User binary integer

There are two kinds of binary integers: a 15-nibble size, and an unlimited (within memory, of course) size. The HP only lets us create the smaller binary integer, but we can create the larger kind manually.

Prologue:	02A4E	Epilogue:	none
Size:	11 nibbles and up	Type:	10

In memory:

Prologue (02A4E)	5 nibbles
Size	5 nibbles
Digit 1	1 nibbles
…	
Digit *n*	1 nibbles

When you enter a number starting with #, the HP will convert it to the current base being used. If you are using the HEX base, if you type # without the last char (h for hex) then the HP will consider it's a hexadecimal number. If you use a letter at

the end (h for hex, d for decimal, b for binary, or o for octal) then the HP will convert it automatically to the current base being used.

Each digit here will use one nibble.

As an example, let's encode #BAh:

The prologue is #02A4Eh, as the preceding table shows. The size is #00015h, as the HP won't allow anything else (see below to get around this). The value is #AB00000000000000h, so we have:

```
#E4A20 51000 AB00000000000000
```

The code below will take a binary number as input, and after running, it will place that binary number with the number of specified digits on the stack, assuming there is enough RAM. You can even create a 200-digit binary number if you want.

This program creates an 85-digit binary number:

```
GOSBVL 0679B        % save registers
LA 0005A            % reserve 90 nibbles
GOSBVL 039BE        % RES_ROOM call
CD0EX               % we save object's address
R0=C A              % into R0a
D0=C                % restore D0
LC 02A4E            % binary number prologue
DAT0=C A            % write it
D0=D0+ 5            % move 5 nibbles further
LC 00055            % size is 90-5=85, #55h
DAT0=C A            % write it to the object's size zone
GOSBVL 067D2        % recover registers
D=D-1 A             % remove 5 nibbles of free ram
D1=D1- 5            % push object on stack
C=R0 A              % its address is inside R0a
DAT1=C A            % we push it
A=DAT0 A            % and here we quit
D0=D0+ 5            % as usual
PC=(A)              %
@
```

WARNING: This program seems to work like it should, and it pushes the binary number on stack. However, if you run MEM, it takes a lot of time, and it really looks like this code somewhat corrupts RAM. I don't know why, but there may be a mistake in this code. Maybe it's not a good idea to create extra-long binary integers. :(

52.6 Graphic object

Prologue:	02B1E	Epilogue:	none
Size:	22 nibbles and up	Type:	11

In memory:

Prologue (02B1E)	5 nibbles
Size	5 nibbles
Number of lines	5 nibbles

Number of rows	5 nibbles
…	Pixels

This is a very commonly used object. A picture, or a grob (Graphic OBject), contains pixels, or more precisely, data that represents pixels. The graphic object first contains a prologue, #02B1Eh, followed by the size of the whole object, but the five nibbles of the prologue are not counted. The next two sets of five nibbles encode the number of lines and rows.

The pixels follow, coded one after another, from the upper-left corner to the lower-right corner. Bytes are used to keep them together. Because the data in a grob is encoded using bytes rather than nibbles, each byte encodes eight bits, which form eight pixels of the grob. A bit set to 1 is a pixel turned on, and a bit set to 0 is a pixel turned off. The word "pixel" is a contraction for "picture element."

IMPORTANT:	In each byte, pixels are encoded using a reversed nibble. Don't forget that pictures must be encoded in a multiple of bytes, or two nibbles. This means that every grob must have an even number of hexadecimal numbers to encode the pixels. Because each nibble encodes four pixels, each grob must be a multiple of eight pixels.

Don't worry if your grob does not have a size that is a multiple of eight pixels. We can simply "pad" the data with zeroes, meaning as many zeroes needed to create a multiple of eight must be added to the end.

For example, suppose I want a pixel off, a pixel on, a pixel off, a pixel on, a pixel off, and then three pixels on, all in one line. The screen would look like this, with 1's being powered pixels and 0's being blank, and the space is added for clarity:

```
0101 0111
```

This has eight pixels, and thus two nibbles, so we don't have to worry about adding trailing zeroes. When we encode this, we must reverse the contents of each nibble, so the pixels must be encoded as `1010 1110`. That's #AEh. After adding the prologue (#02B1Eh), size (#19d, or #13h nibbles), width (#4h), and height (#1h) to the beginning, our graphic is encoded as follows:

```
#E1B20 31000 40000 10000 AE00
```

Note that two trailing zeroes were added, and the addition was reflected in the size of the grob. This was just to make the length #13h rather than #11h so you would remember that the nibbles must be inverted. In addition, this shows that you can have excess data within an image without hurting anything. Of course, if these zeroes were removed the grob would take one byte less memory. A note to Meta Kernel users: the Meta Kernel appears to add a zero-valued byte to the end of a graphic to indicate it is a grayscale grob.

As a refresher, the #13h (#19d) size was calculated by adding five nibbles for the size field, five nibbles for the width field, five nibbles for the height field, and, in this case, four nibbles for the pixel data.

52.7 XLIB name

Prologue:	02E92	Epilogue:	none
Size:	11 nibbles	Type:	14

In memory:

Prologue (02E92)	5 nibbles
Library ID	3 nibbles
Command number	3 nibbles

This is how a command inside a library is referenced on the HP48. Even though you may see the library command name, the HP48 uses this object inside. When you have, for example, DTAG on the stack it's an XLIB object, which the HP48 translates into the command's name when displaying it onscreen. It contains the library number and the command number. When the command has a "user" form (a name you can type) it's displayed, otherwise "XLIB …" appears on the stack.

There is a command you can use to indirectly use XLIB's: the command LIBEVAL. LIBEVAL takes a six-digit binary value (#...) as an argument; the first three digits are the library number and the last three digits are the command number, both in hexadecimal.

52.8 System binary integer

Prologue: 02911 Epilogue: none
Size: 10 nibbles Type: 17

In memory:

Prologue (02911)	5 nibbles
Number	5 nibbles

This is a very simple object. The HP's RPL operating system uses it extensively, and it contains a binary integer encoded using five nibbles (that's 20 bits). The number will be shown according to the active base, **but** it's always encoded in hexadecimal.

52.9 Long real

The long real is also called the "extended real."

Prologue: 02955 Epilogue: none
Size: 26 nibbles Type: 21

In memory:

Prologue (02955)	5 nibbles
Exponent	5 nibbles
Mantissa	15 nibbles
Sign	1 nibble

The long real is more precise than the standard real. The user does not have access to this number, but the HP uses it internally for greater precision during calculations. It's encoded in a way similar to the standard real, but here some fields have greater sizes. BCD is used here, just like with the real object. Refer to the standard real information for details.

52.10 Character

Prologue: 029BF Epilogue: none
Size: 7 nibbles Type: 24

In memory:

Prologue (029BF)	5 nibbles

ASCII code of character	2 nibbles

The character is one of the simplest objects available on the HP48! In the case of this object, the prologue is actually bigger than the data being encoded. Each of these objects contains a character the HP is able to display. Characters 0-128 are the standard ASCII ones, but the HP has some special characters above. Here we have two nibbles per character, so we can code all 256 chars of the extended (8-bit) ASCII table. Because the original ASCII table was only 7-bit, only the first 128 characters are "universal."

52.11 Code object

Prologue: 02DCC Epilogue: none
Size: 10+ nibbles Type: 25

In memory:

Prologue (02DCC)	5 nibbles
Size	5 nibbles
Saturn machine code	varies

This is the kind of object that HP-ASM, ASM Flash, and other code compilers produce.

Because it has the same structure as strings, some people use Code objects to hide strings, or some other object, inside. That's as simple as changing the prologue, which is so easy you know how to do it now, don't you? :o)

52.12 Library data

Prologue: 02B88 Epilogue: none
Size: 10 nibbles and up Type: 26

In memory:

Prologue (02B88)	5 nibbles
Size	5 nibbles
Data	varies

HP has made this object available to programmers so libraries that need to save values can save information inside a separate object. The coder will choose what data and how will be put inside this object, so it's a convenient way to store information, a kind of HP Cookie. ;-)

One common use is to store pieces of Code that are called by other programs; as these pieces of Code don't check the stack, this is a way to prevent the user from "accidentally" running these Code objects. You should be able to easily code a "launcher" for such objects. The goal is to protect your data.

52.13 Extended pointer

Prologue: 02BAA Epilogue: none
Size: 15 nibbles Type: 27

In memory:

Prologue (02BAA)	5 nibbles
Object's address	5 nibbles
Program's address	5 nibbles

As you will learn in the next part, some areas of memory are covered by others, but sometimes we need to point something that is covered. This is the object the HP uses internally for that purpose. There are routines in ROM, which we'll discuss later, that are used to "uncover" the part we want to use. In this case, the object is what we want access to, and the program is the ROM routine that will be associated with the object's address, and will unconfigure and reconfigure memory modules as needed.

53 Compound Objects

53.1 Complex number

Prologue: 02977 Epilogue: none
Size: 37 nibbles Type: 1

In memory:

Prologue (02977)	5 nibbles
Real exponent	3 nibbles
Real mantissa	12 nibbles
Real sign	1 nibble
Imaginary exponent	3 nibbles
Imaginary mantissa	12 nibbles
Imaginary sign	1 nibble

This is the complex number object. In fact, representing a complex number is just keeping two real numbers together. The first one is the real part of the complex, and the second one is the imaginary part, with each number being encoded using 16 nibbles.

For example, let's encode (100,1000).

First comes the prologue, #77920h (reversed, as usual), then the real exponent, #200h, then the mantissa, #000000000001 (for the real part), and then the sign, #0h. Next comes the imaginary part, which is encoded as #300h for the exponent, #000000000001h for the mantissa, and finally the sign, #0h. This means the complex number is stored in memory like this:

#77920 200 000000000001 0 300 000000000001 0

53.2 Array

Prologue: 029E8 Epilogue: none
Size: 27 nibbles and up Type: 3/4

In memory:

Prologue (029E8)	5 nibbles

Size	5 nibbles
Kind of objects	5 nibbles
Number of dimensions	5 nibbles
Dimension 1	5 nibbles
…	
Dimension *n*	5 nibbles
Content of object 1	Size varies
…	
Content of last object	Size varies

The array object can contain any number of objects that are of the same type. When it only has one dimension, it's called a "vector." If there are two dimensions, we have a "matrix." We can also give it more than two dimensions, as five nibbles are available to encode the array's dimensions.

Usually, arrays are used to keep sets of real or complex numbers together. You can store any kind of object you would like using as many dimensions as you want, but the HP will only manage one- or two-dimensional arrays.

The prologue is #029E8h. Within the object, we must first define which objects are going to be kept inside the array object. Then, we have to tell the size of the whole array, but it depends on the kind of objects we're going to collect, and of course, how many of them are to be encoded.

You will see that encoding the length of an object is useful: the HP can use it to "jump" over the object. Most objects will have their size inside of them, but some won't, so the "epilogue" will be very useful to find the end of these objects.

What is interesting here is that even if we are collecting ten reals, there are not ten real prologues stored into the array: their prologue is only stored once. This is why the kind of objects is encoded using five nibbles and also why only **one** kind of object can be stored inside an array. Once we have set the kind of objects inside the array, we encode the number of dimensions using five nibbles: #00001h for vectors or #00002h for matrices. After the dimensions, we must tell how many objects are in *each* dimension. If we have a matrix, we will encode the number of lines and columns.

Once we have done all that, we can start encoding the objects inside the array. The way values are stored will vary according to the kind of objects stored. The prologues are not needed because we have already recorded it.

An important thing is that no matter what the size of the array is or what the objects are, if some parts are "empty," they still use memory.

For example, let's create an array of strings, with some names inside of it.

```
[ HP48S HP48G ]
```

First, we have the array prologue: #029E8h. As usual, it will be encoded reversed in memory. The total size is #00032h (you will learn why later; this takes five nibbles to encode). Next we must tell what kind of objects are included–strings–so we use #02A2Ch (five nibbles). There is only one dimension so we must write #00001h (another five nibbles), and there are two elements so we must write #00002h (five more nibbles).

Finally we encode the two strings. Since there are two nibbles per character, each string takes ten nibbles plus five more nibbles in front of it to tell its length. Because each string takes 15 nibbles, including the length field, the length is #0000Fh, reversed to show #F0000h. This means a total of 50 nibbles (hence the #32h length, as #50d = #32h), plus the prologue are needed to encode the object, like this:

```
#8E920 23000 C2A20 10000 20000 F00008405438335 F00008405438374
```

The HP48's stack doesn't show arrays of strings, so it will simply say "Array of String." If you are running the Meta Kernel, however, it will show the true array.

53.3 List

Prologue: 02A74 Epilogue: 0312B
Size: 10 nibbles and up Type: 5

In memory:

Prologue (02A74)	5 nibbles
Object 1	Size varies
…	
Last object	Size varies
Epilogue (0312B)	5 nibbles

The list object is interesting: it can contain any kind of object, and as many objects as RAM allows. This is done by using a prologue to mark the beginning and an epilogue to mark the end, with each object (including its prologue) between them. This is often found in ROM :-)

An empty list will be { } and 10 nibbles long: there is a prologue and an epilogue but nothing between.

For example, let's encode the list `{ 10.2 "dX" }`.

I have chosen two kinds of objects we have already seen: a real number and a string. It's quite easy: first encode each object, and then put the two between the list's prologue and epilogue:

10.2 is `#33920 100 000000000201 0`
"dX" is `#C2A20 90000 4685`

This is how the list would be encoded:

```
#47A20 3392010000000000002010 C2A20900004685 B2130
```

53.4 RPL program

Prologue: 02D9D Epilogue: none
Size: 20 nibbles and up Type: 7

In memory:

Prologue (02D9D)	5 nibbles
<< User RPL start (2361E)	5 nibbles
…	
>> User RPL end (23639)	5 nibbles
Epilogue (0312B)	5 nibbles

When an RPL program is created, this is the object produced. If you look closely, you'll see that it's similar to a list, containing a prologue and an epilogue. Inside the RPL program, objects are found, one after another, but each object may be an address rather than the object itself. For example, suppose you create this RPL program:

```
<<  OFF  >>
```

Instead of finding the code that turns the HP48 off inside this RPL object, there is only an address that points to the code. The OFF code is found at that address. When the HP displays the RPL program to you, it decodes each address into the corresponding name, if it has one.

The RPL program `<< OFF >>` is coded this way in memory:

```
D9D20
  E1632          % <<
    E13A1
  93632          % >>
B2130
```

Here, the nibbles are reversed, as the Saturn always does when dealing with memory. The first and last five nibbles are the RPL object: #02D9Dh is the prologue and #0312Bh is the epilogue.

Inside the object itself there is the `<<` (being #2361Eh) and the `>>` (#23639h), and between those delimiters is an address, #1A31Eh, which is the OFF command.

You can verify it: type #1A31Eh, SYSEVAL, [ENTER], and your HP will turn off. :-)

Every time you put a command inside a RPL program, the HP puts the address of the command.

Here the `<<` and `>>` create a block inside the RPL object. We can even put these "blocks" inside of each other.

These `<<` and `>>` are not really required in RPL programs you write, but the prologue and epilogue are required to have a valid, and thus executable, RPL program. Even if you do `<< << OFF >> >>` the HP codes something different: the `<<` and `>>` are there, but to make sure the `<< OFF >>` inside is not executed, something is inserted just before and just after.

This book is not aimed to be a tutorial on System RPL so this will not be discussed. If you would like to learn more about the secrets of RPL, there are many excellent tutorials. Eduardo Kalinowski's "Programming in System RPL," Hewlett Packard's RPLMAN.DOC, and James Donnelly's "An Introduction to HP 48 System RPL and Assembly Language Programming" are recommended. The first two are available at http://www.hpcalc.org and the last can be purchased from Armstrong Publishing Company (http://www.peak.org/~jimd/Armstrong/).

53.5 Algebraic expression

This object type is also called a "symbolic."

Prologue:	02AB8	Epilogue:	0312B
Size:	10 nibbles and up	Type:	9

In memory:

Prologue (02AB8)	5 nibbles
Object 1	Size varies
…	
Object *n*	Size varies
Epilogue (0312B)	5 nibbles

This is the usual algebraic expression used in RPL. Both a prologue and an epilogue are included. The content of the object is encoded using RPL-notation; although on the stack it does not look like RPL, **it is**. As you know, because it is encoded

like RPL objects, there is no "+" sign (or whatever the desired command is) inside this object, but each mathematical symbol is coded using its address in ROM using five-nibble addresses.

For example, let's encode `'A+2'`:

The prologue is 021B8, as shown in the table above. Next, we must encode the global name A, which is encoded as 02E48 as prologue, 01 for the size (two nibbles used), and 41 for the ASCII code of the character 'A'. If this is confusing, look at the section on encoding global objects. Next comes the "+" command, encoded using its address in ROM, #2A2Deh, and the number 2, encoded using its ROM address, #1AB67h. Finally, we find the epilogue in the table above, #0312Bh. This means that the complete object, as reversed in ROM, is:

```
#8BA20 84E201014 ED2A2 76BA1 B2130
```

53.6 Tagged object

Prologue:	02AFC	Epilogue:	none
Size:	varies	Type:	12

In memory:

Prologue (02AFC)	5 nibbles
Tag size	2 nibbles
First tag character	2 nibbles
…	
*n*th tag character	2 nibbles
Object	varies

A tagged object is simply an object with a tag to label it. A label of up to 256 characters is created and the object is put just after. The whole thing is kept inside the "tagged" object type.

All we do to create a tagged object is start with the tag prologue, next put the size of the string which tags the object, then put the characters of the tag, and finally encode the object to be tagged as you normally would. The size of the string used to tag is encoded using two nibbles, and thus it is limited to 256 characters.

As an example, let's tag the value 10 using "Ten". It will be encoded like this:

```
#CFA20 30 4556E6 33920 100 000000000001 0
```

Notice that this begins with the tag prologue, #02AFCh, the number of characters used by the tag (#03h, but reversed of course) follows, and then the ASCII values of the characters forming the tag are encoded (Ten becomes #6E6554h, and is then reversed) of the chars. Finally, the real number 10 is encoded as you learned before.

53.7 Unit object

Prologue:	02ADA	Epilogue:	0312B
Size:	varies	Type:	13

In memory:

Prologue (02ADA)	5 nibbles
Object	varies

Unit	varies
Epilogue (0312B)	5 nibbles

This object is defined with a prologue and epilogue, and between them we find two values. The first one is the object that will receive the unit; it can be a real value, such as the 2 in `'2 m^2'`. The second one is the object used to encode the unit. Symbol like + or * are encoded using their address in ROM, and characters encode the various letter symbols.

For example, let's encode 300 000 000 meters per second. That's slightly over the light of speed, which is about 299792458 meters per second according to the HP48's constant library.

First, put the prologue, which is #02ADAh. Next, put the object, which is the real number 300 000 000 and is encoded as `#33920 800 000000000003 0`.

Finally the unit is encoded. This is important: we must use RPN notation!! This shouldn't be a surprise to you, though, as you know that the HP48 is an RPL-based calculator.

Using RPN, we see that the unit "meters per second" is then: `m s /`

We have two characters, m and s. That's two strings, so by applying our knowledge of encoding strings they are:

```
#C2A20 70000 D6
#C2A20 70000 37
```

The division sign (/) comes next, to divide m by s, and then a multiplication sign (*) to multiply the number by the unit.

The epilogue is inserted at the end, so we finally get:

`#ADA20`	(prologue)
`#339200800000000000030`	(the number 300000000)
`#C2A2C70000D6`	(m)
`#C2A2C7000037`	(s)
`#86B01`	(/)
`#68B01`	(*)
`#B2130`	(epilogue)

53.8 Directory object

We have two kinds of directories: the HOME directory, and every other directory. HOME is often called the "root directory."

Libraries are "attached" to directories. HOME can have many libraries attached, but each subdirectory can only have *one* library attached to it. Two descriptions follow: one for the HOME directory and the other for subdirectories. This is a very complex object, so pay attention!

53.8.1 HOME directory

Prologue:	02A96	Epilogue:	none
Size:	57 nibbles and up	Type:	15

In memory:

Prologue (02A96)	5 nibbles
Number of attached	3 nibbles

libraries	
Number of first attached library	3 nibbles
Address of first library hash table	5 nibbles
Address of first library message table	5 nibbles
…	
Number of last attached library	3 nibbles
Address of last library hash table	5 nibbles
Address of last library mess table	5 nibbles
Offset of last object	5 nibbles
00000	5 null nibbles
Number of characters of object name 1	2 nibbles
Character number 1 of object 1	2 nibbles
…	
Character number *n* of object 1	2 nibbles
Object 1	varies
Size of object 1 zone	5 nibbles
…	
Number of chars of last object name	2 nibbles
Character number 1 of last object name	2 nibbles
…	
Character number *n* of last object name	2 nibbles
Last object	varies

So, first we have the HOME prologue. It's followed by a three-nibble value, which tells how many libraries are currently attached to HOME (in other words, the libraries you have inside your HP). Even if you have not installed any libraries, the HP48 has libraries built into ROM that are attached to HOME! Three very well known libraries are number 2, which contains ASR, RL, RLB, etc., number 171, which contains XVOL, YVOL, VERSION, etc., and number 1792, which contains IF, THEN, ELSE, etc. Although your HP may appear to be empty, these three libraries are always attached.

Next we will find 13 nibbles for each library that is attached. The first three nibbles contain the number of the library, the next five contain the address of the hash table of the library (this will be described later, but it is quite complex), and the last five contain the address of the library message table.

The hash table will be used to quickly access commands and the message table, if it exists, to look for custom error messages inside the lib. It's not the best optimization available, but an interesting one nevertheless.

After that, we have a five-nibble address, which is an offset to the *last* object inside the directory. This is very useful, and very, very clever.

Then, the objects inside the directory are encoded, one after another. For each object, the length of its name is listed first, encoded using 2 nibbles, as the length of a global name is two nibbles, and then each character of the name is given. If you are used to programming in Pascal, this is similar, as it uses the same variable naming convention. The Pascal language first codes the size and then the data; on the other hand, the C language encodes the data with a null value following it and no size field. Here, it's very useful to have the length before.

After the last character of the object's name, we have the object, and after the object, five nibbles give us the *whole* space used by object 1, from the two nibbles that encode its name length to the end of the object.

For each object, it's the same: first shown is the length of the name, then the characters that make up the name, then the object, and finally the space the whole thing uses.

NOTE 1:	If the hash table or the message table of one of the libraries is located in covered memory that cannot be accessed directly, then an extended pointer is found there. The extended pointer object, as well as covered memory, will be explained later. If a library doesn't have a message table, you will find #00000h.

NOTE 2: These numbers are available for library ID numbers:

Number's range	Hewlett-Packard's recommended usage:
#000h to #100h	HP libraries in ROM (RESERVED)
#101h to #200h	HP libraries in RAM (RESERVED)
#201h to #300h	Non-HP libraries (RESERVED)
#301h to #6FFh	User libraries (FREE)
#700h to #7FFh	HP own use (RESERVED)

As you see, of the three libraries attached to the HP, { #700h #002h #0ABh }, the first one is in the area reserved for HP's own use (presumably in application cards) and the two others in the area reserved for libraries resident in ROM.

53.8.2 Subdirectories

In memory:

Prologue (02A96)	5 nibbles
Number library attached to directory	3 nibbles
Offset of last object	5 nibbles
00000	5 null nibbles
Number of characters of object name 1	2 nibbles
Character number 1 of object 1	2 nibbles
…	
Character number *n* of object 1	2 nibbles
Object 1	varies
Size of object 1 zone	5 nibbles
…	

Number of chars of last object name	2 nibbles
Character number 1 of last object name	2 nibbles
…	
Character number *n* of last object name	2 nibbles
Last object	varies
Size of last object zone	5 nibbles

As you can see, a subdirectory can only have *one* library attached to it. If no libraries are attached, then #7FFh will be used as the number of the attached library.

NOTE:	An offset is contained inside the directory object so you can quickly and easily find the last object inside one directory object. The first object has *five* null nibbles before it, and you also will find a *five* nibble value after each object which gives you the whole length used by the object, though the last object does not have that information. We also can find the number of characters of each object's encoded name.

53.9 Library

Caution! This is a very complex object. :)

Prologue:	02B40	Epilogue:	none
Size:	varies	Type:	16

In memory:

Prologue (02B40)	5 nibbles
Size (CRC not included)	5 nibbles
Library name's size	2 nibbles
Library name's first character	5 nibbles
…	
Library name's last character	2 nibbles
Library name's size	2 nibbles
Library number	3 nibbles
Hash table offset	5 nibbles
Message table offset	5 nibbles
Link table offset	5 nibbles
Config object offset	5 nibbles
Message table array	varies
Link table	varies
XLIB 1 (kind)	1 or 3 nibbles
Library number	3 nibbles

Command number	3 nibbles
XLIB 1 objects	varies
XLIB 2 (kind)	1 or 3 nibbles
Library number	3 nibbles
Command number	3 nibbles
XLIB 2 objects	varies
…	
XLIB *n* (kind)	1 or 3 nibbles
Library number	3 nibbles
Command number	3 nibbles
XLIB *n* objects	varies
Masked object 1	varies
…	
Last masked object	varies
Config object	varies
CRC	4 nibbles

Not only is the library structure somewhat complex, but it varies too. Four objects inside the library object have a very specific structure: the hash table, the message table, the link table and the config object.

As you have seen, there are four offsets inside the library with this very purpose.

If the library has no name, its first length field will be zero, and there will **not** be a second length field. If there is a name, we find its length, then the characters, then again the length. A library without a name will not appear when you will do [Right-shift][2].

The hash table is used to increase the access speed of library commands. The message table contains all custom error messages of the library and the link table references all objects of the library. Offsets are used inside the library, and if an offset is #00000h then the object we are looking for does not exist.

We will discuss the four object types below. But first there is something else about the library that's important: there are two types of objects, those that are visible, called XLIB's, and those which are masked. For each *visible* object of the library, you will find:

- its kind (encoded using 1 to 3 nibbles)
- the library number (3 nibbles)
- the command number (3 nibbles)

The CRC is calculated using a hardware CRC circuit inside the HP. The CRC of an object is calculated from the size field (the prologue is not included) to the end of the object. The CRC is *not* a part of the object that is CRC'd. Then, the CRC is added to the library. The calculated size of the library contains neither the prologue nibbles nor the CRC nibbles.

53.9.1 Hash table

The purpose of this table is to speed up access to library commands. It's very simple, so the gain is not monstrous, but it's a good idea. Commands are distributed according to their name's length, from 1 to 16. A command can, however, have more than 16 characters, but then it will be grouped with the 16-character long commands.

Here is the memory scheme of the hash table:

Prologue:	02A4E	Epilogue:	none
Size:	varies	Type:	none

Category 1: commands whose names are 1 character long
…
Category *n*: commands whose name are *n* characters long

In memory:

Prologue (02A4E)	5 nibbles
Size of hash table	5 nibbles
First category offset	5 nibbles
…	
Last category offset	5 nibbles
Size of the names list	2 nibbles
Size of one-character command names	2 nibbles
First one-character command	2 nibbles
Number of command	2 nibbles
Second one-character command	2 nibbles
Number of command	2 nibbles
…	
*n*th one-character command	2 nibbles
Size of 2 character command names	2 nibbles
First two-character command	4 nibbles
Number of command	2 nibbles
…	
Offset to command 1	5 nibbles
…	
Offset to last command	5 nibbles

Here, a command that has no visible name is not listed in the hash table. Each library command has a name and number inside the library. This information is ordered in the hash table. The command number follows the names: just after the first one-character command we find its number, and so on. First, we order by length of name, and after each name we put the command number as inside the library, but only for visible commands. Next, we find the hash table offsets, of which there are two kinds: one to find a command according to its length and another to find the command using the offsets found at the end of the hash table.

53.9.2 Message table

The message table can have two forms: it can either be a single-dimensional array or an *indexed* single-dimensional array.

If we find an array (prologue 029E8) we will find the number of messages inside the library, followed by all messages, each one starting with its length (Pascal string encoding), and then each ASCII byte of each character.

If we find an indexed array (prologue 02A0A) we will find the number of messages and then a list of offsets to each message. Its length (still Pascal string encoding) precedes each message and is followed with one ASCII byte for each character of the message.

53.9.3 Link table

The link table is a simple binary integer. It encodes the addresses of the library's objects inside. You can divide this binary integer into chunks of five nibbles. The first one is the prologue (02A4E) and is followed by the length, which, as usual, does not include the prologue's length. After that, every five nibbles is an address to an object. Here we just have a table of addresses, one for each object inside the library.

53.10 Backup

Prologue: 02B62 Epilogue: none
Size: varies Type: 17

In memory:

Prologue (02B62)	5 nibbles
Size	5 nibbles
Number of characters in object's name	2 nibbles
First character of name	2 nibbles
…	
Last character of name	2 nibbles
Object	varies
System binary prologue (02911)	5 nibbles
0	1 nibble
CRC	4 nibbles

Several objects can be put inside of a backup object. The HP will only put one object and will follow it with a hidden system binary whose first nibble is 0 along with four other nibbles containing the CRC of the object. The CRC is used to ensure that the object has not been corrupted. The first of the five nibbles of the system binary used will always be zero because the CRC is encoded using only four nibbles.

53.11 Long complex

Prologue: 0299D Epilogue: none
Size: 47 nibbles Type: 22

In memory:

Prologue (0299D)	5 nibbles
Real exponent	5 nibbles

Real mantissa	15 nibbles
Real sign	1 nibble
Imaginary exponent	5 nibbles
Imaginary mantissa	15 nibbles
Imaginary sign	1 nibble

The long complex is like two extended reals together, with the first one being the real part and the second one being the imaginary part. This object is used internally by the HP for greater precision.

53.12 Linked array

Prologue: 02A0A Epilogue: none
Size: 32+ nibbles Type: 23

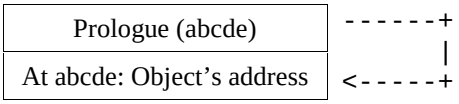
In memory:

Prologue (02A0A)	5 nibbles
Size	5 nibbles
Size of objects	5 nibbles
Number of dimensions	5 nibbles
Dimension 1	5 nibbles
…	
Dimension *n*	5 nibbles
Pointer to first object	5 nibbles
…	
Pointer to last object	5 nibbles
Contents of first object	Size varies
…	
Contents of last object	Size varies

The linked array is a special kind of array. There are pointers inside a linked array, with each one being a shortcut to each object being encoded. If you have a big array that contains a lot of identical values, this kind of array will save a lot of memory space. For example, a pointer is 5 nibbles long, while a real is 16, so if you have a 100x100 array with five different real number values you will save a lot of space by using a linked array.

54 Primitive Code Object

Those don't really exist as objects, as you cannot have them on the stack except through a pointer. The prologue of those PCO's points to its address+5, for example:



As you can see, a PCO is defined as above: it's an address where five nibbles are found, located at the PCO address plus 5.

There you will find machine language Code, without any prologue or indication of size. This is used a lot by the HP; for example, all stack commands like SWAP, DUP, DROP, etc. are encoded using PCO's.

Someone who codes using System RPL may make calls to these PCO objects in ROM, called "primitives," though that doesn't mean the mathematical primitive function!.

You can call a PCO using the SYSEVAL command. Put a binary integer number on the stack and type SYSEVAL [ENTER].

Part VI: Writing Programs

55 Doing loops

There are several ways to do loops, serving various purposes. A loop can be used to "wait" or slow down a program, because sometimes it will go too fast, or one can just repeat a set of instructions several times.

When you want to do a loop, something must be used to keep a value. It can be a register like A, B, C, D, or even P.

Here, we're going to use the register C as loop counter:

```
LC FF
*LOOP
C=C-1 B
GONC LOOP
A=DAT1 A
D0=D0+ 5
PC=(A)
@
```

This is a simple program. The number FF is loaded and the B field is used to loop. Each time the program loops, 1 is removed from the counter register, using the B field. The loop continues until the carry is set.

An important thing to remember when using the carry is that the loop is done 1 more time than the counter's value. Why?

FF was loaded into field B. 1 will be removed each loop, but when zero is read, the carry is still not set, so it loops again. The next time, #00h is in field B and removing 1 sets the carry, so the loop stops. Therefore, the loop above is executed #FFh + 1 times, so #100h times.

When the carry is used to do loops in your program, don't forget to remove one unit from the counter value before starting. Because the loop will be done one more time, it will loop the proper amount.

Here is an example of using P as counter. We are going to create a program that does a little "click" on the HP's buzzer.

To turn on the buzzer, we use the OUT register of the Saturn processor. We send #800h to OUT, and the buzzer turns on. If we want to turn it off, we send #000h to the OUT.

Here is the example program that clicks:

```
LC 800
OUT=C
*LOOP
P=P+1
GONC LOOP
C=0 X
OUT=C
A=DAT0 A
D0=D0+ 5
PC=(A)
@
```

Compile it, store it inside a variable, put the back of your HP close to your ear, and press the key several times to run the program. You will hear a small click. :-)

By varying the time and turning off and on the speaker, you can produce sounds. There is a routine in ROM that produces sound of the frequency (in hertz) and duration (in milliseconds) that you give it.

56 Reading from the keyboard

When we want to read a key, we have to consider that keys are internally located along lines. We're going to have current sent over some lines, and we will look at the output to check if a key was pressed. Pressing a key "closes" a circuit, which we can detect.

This introduces the concept of "couples." We are going to first send a specific value to OUT, then read IN, and finally check whether a key was pressed by comparing the IN with what should be found according to the OUT sent.

Below is the keyboard layout of the HP. Below each key name are two numbers: one on the left, and one on the right. The one on the left will be put into OUT, and if they key is pressed we will find the corresponding value on the right in the IN register.

A 002/0010	B 100/0010	C 100/0008	D 100/0004	E 100/0002	F 100/0001
MTH 004/0010	PRG 080/0010	CST 080/0008	VAR 080/0004	(UP) 080/0002	NXT 080/0001
' 001/0010	STO 040/0010	EVAL 040/0008	(LEFT) 040/0004	(DOWN) 040/0002	(RIGHT) 040/0001
SIN 008/0010	COS 020/0010	TAN 020/0008	(SQRT) 020/0004	(POWER) 020/0002	(INV) 020/0001
ENTER 010/0010		+/- 010/0008	EEX 010/0004	DEL 010/0002	(BACKSPACE) 010/0001
(ALPHA) 008/0020	7 008/0008	8 008/0004	9 008/0002	(DIVIDE) 008/0001	
(LEFT-SHIFT) 004/0020	4 004/0008	5 004/0004	6 004/0002	(MULTIPLY) 004/0001	
(RIGHT-SHIFT) 002/0020	1 002/0008	2 002/0004	3 002/0002	(SUBTRACT) 002/0001	
ON /8000	0 001/0008	. ¼	SPC 001/0002	(ADD) 001/0001	

56.1 Checking for a single key-press

Make a program loop until the [SPC] key is pressed:

```
INTOFF          % turn off keyboard interrupts
*LOOP           % LOOP starts here
LC 001          % we load #001h onto C(A)
OUT=C           % send it so SPC line receives current
GOSBVL 01160    % we do C=IN (see below)
?CBIT=0 1       % does C contains #0002h ?
GOYES LOOP      % no? jump to LOOP
INTON           % allow keyboard interrupts
A=DAT0 A        % return to RPL
D0=D0+ 5
PC=(A)
@
```

A few notes here. First, INTOFF and INTON have been used. This is to tell the HP we're going to handle the keyboard in our code. It is important to use INTON at the end of the source. Then, if you look at the couple of values for the [SPC] key you will find that it's 001/0002.

One could 0002 into A, and then compare register C with register A. But, if one only considers the first nibble, #0002h in hex is the same as #0010b in binary. This means all one has to do is check whether bit 1 is equal to zero. If so, the key has not been pressed, and it will loop again.

Remember that the first bit is number zero and the second bit is number one.

There are others ways to check for a key. Next we'll use one that uses a logical operator, and the concept of a "mask."

56.2 Checking for key combinations

Masks are a useful concept when it comes to detecting multiple key-presses. Let's write a program that will quit if [1] and [9] are pressed at the same time.

The output for 1 and 9 keys is #002h (#0010b) and #008h (#1000b). Performing an OR of the two values gives #Ah (#1010b).

You can make this process simpler by using your calculator! Simply put #2h and #8h on the stack and press [MTH][F][NXT][A][B] (or simply type OR). This performs an OR on the first two levels of stack

NOTE:	You must be in HEX mode to get the value you will put inside your program. HEX mode is [MTH][F][A].

The IN values will be found using OR. The couples are #8h and #2h, so #8h OR #2h = #Ah

Source is:

```
INTOFF
*LOOP
LC 00A
OUT=C
GOSBVL 01160
LA 000A
?C#A A
GOYES LOOP
INTON
A=DAT0 A
D0=D0+ 5
PC=(A)
@
```

56.3 Reading the entire keyboard

And now, what if we want to check if any key has been pressed? We can simply do an OR of all output masks, and then check whether the value read from the keyboard is different than zero.

But, the [ON] key is special: it's directly wired to a pin on the Saturn processor. To test whether [ON] has been pressed, we must do a `C=IN` and then check whether bit 15 is null. If it is null, [ON] was pressed.

The OR mask of all output masks is #1h OR #2h OR #4h OR #8h OR #10h OR #20h OR #40h OR #80h OR #100h. That returns #1FFh.

The source to test whether any key has been pressed follows:

```
INTOFF
*LOOP
LC 1FF
GOSBVL 01EEC    % does C=OUT and then C=IN and returns
?C=0 A
GOYES LOOP
INTON
A=DAT0 A
D0=D0+ 5
PC=(A)
@
```

There is one problem with this: as soon as we launch it, it quits. We have to flush the keyboard buffer before we can start the loop. This piece of code will flush the keyboard buffer:

```
D1= 0669
C=0 A
DAT1=C A
```

This routine is extracted from the routine at #00D57h in ROM, which contains the code to flush the keyboard.

If we use D1, we must save its value before, of course. :)

We get:

```
INTOFF          % we turn off keyboard interrupts
*LOOP
LC 1FF          % mask for all keys
GOSBVL 01EEC    % does C=OUT then C=IN and returns :-)
?C=0 A          % no key pressed?
GOYES LOOP      % we loop
INTON
LC 00669        % point to buffer
CD1EX           % old D1 saved in C
A=0 A
DAT1=A A        % we empty it
D1=C            % restore D1
A=DAT0 A        % and we quit
D0=D0+ 5
PC=(A)
@
```

It goes so fast that you will need something like the following RPL program, where WAITK is the variable the code will be inside:

```
<< 0 WAIT DROP WAITK >>
```

Put this program on the stack, press EVAL quickly, the busy indicator should turn off, and the RPL program will remain on stack. Press any key and it will quit.

Code executes very quickly, as you can see.

56.4 Reading the [ON] key

The [ON] key is special: because it is directly wired to a pin of the processor; if we want to test it, INTOFF is not enough. We have to disable all interrupts, so need to set flag 15 of ST to zero. As we'll see later, in the disassembled code of the

interrupt handler, if this bit is set to zero, no interrupt is done. (In fact it's a bit more complicated, but let's keep that for later. Interrupts are complex, so we'll leave that for the last part.)

Code to check for [ON]:

```
ST=0 15          % disable all interrupts
*LOOP
GOSBVL 01160     % we directly do C=IN
?CBIT=0 15       % [ON] pressed?
GOYES LOOP       % if not, we loop
ST=1 15          % we allow all interrupts
A=DAT0 A         % and quit
D0=D0+ 5
PC=(A)
@
```

> **CAUTION!** Don't make mistakes with `ST=0 15`, because once it's performed, you can not use ON-C or ON-A-F. If your program traps itself inside an infinite loop, you will have to do a hard reset using the hole below your HP, and you may lose parts of memory.

We have to use `ST=0 15` and `ST=1 15` because otherwise we cannot check for [ON] being pressed.

There is an another way to check for multiple key presses: check each key, one after another. Machine language goes so fast that it will detect multiple-key presses: the program below loops until [A] and [F] are pressed together.

First, we wait for [A], and as soon it's detected, we check for [F]. This program goes so fast that it will detect both keys being pressed:

```
INTOFF           % turn off interrupts
*LOOP
LC 002           % out for [A]
GOSBVL 01EEC     % does OUT=C and C=IN
LA 0010          % in for [A]
?C#A A           % if [A] not pressed
GOYES LOOP       % we loop
LC 100           % out for [F]
GOSBVL 01EEC     % we do OUT=C and C=IN
?CBIT=0 0        % bit 0 = 0 ?
GOYES LOOP       % then wait for [A], [F] being not pressed
INTON            % restore interrupts
LC 00669         % point to buffer
CD1EX            % old D1 saved in C
A=0 A
DAT1=A A         % we empty it
D1=C             % restore D1
A=DAT0 A         % and we quit
D0=D0+ 5
PC=(A)
@
```

57 Manipulating the stack

Previously you learned that the stack ends with four null nibbles, #00000h. Though this is partially true, the *full* truth is a bit more complicated. :-)

Internally, the HP calculates the stack's depth with the following code:

```
LC 806FD        % 806FD on G/GX and 7057E on S/SX
CD1EX
A=DAT1 A
C=A-C A
GOSUBL 06A8E    % a routine that divides C(A) by 5
C=C-1 A         % we remove 1 and we have stack's depth
```

Here, D1 points to the first level of stack. The `CD1EX` moves the address in D1 to C and makes D1 point to #806FDh. At that address, using five nibbles, the address of the command line's start is stored. Next, that address is loaded inside A, and D1 is restored. Then, we calculate:

stack size * 5 = (D1 address) - (Command line address)

So we need to divide the result, C(A), by five. There is a routine in ROM to do this, located at address #06A8Eh, and it uses registers A, C, and D.

Once divided by five, remove one, and we have the stack's depth. :-)

Why do we do all that? Is it not enough just to check for the #00000h on stack? Even though it's true that the end of stack contains #00000h, there can also be the address #00000h on stack, so the stack could look like this:

```
4:
3:              4
2:    <External>  ← this is #00000h !!
1:              3
```

So we can have two (or more) #00000h on stack. This is why, if we really want to know what is the stack's depth, we have to calculate it, and not just search for the first #00000h on stack.

If you'd like, you can make your HP show this stack display. Put three numbers on the stack, and then compile this code:

```
D1=D1+ 5        % we move to level 2
C=0 A
DAT1=C A        % and set it to #00000h
D1=D1-5         % we come back to level 1
GOVLNG 2D564    % this does A=DAT0 A, D0=D0+ 5 and PC=(A)
```

Store this Code inside a variable, and run it over the 3 level stack you have. What do you see?

You have on level two an <External>, which is in fact #00000h.

Press [DEL] to clear the whole stack. Don't [EVAL] the <External>, as that would make your HP jump to #00000h, causing a warmstart. Your HP will only restart, which is nothing dangerous, but don't play with fire. :)

Here is the code that calculate the real depth of the stack:

```
CD1EX           % we push D1 into C
D1= 806FD       % and make D1 point to #806FDh
A=DAT1 A        % we read command's line start
C=A-C A         % and calculate D1 address minus command line address
GOSBVL 06A8E    % divide it per 5
C=C-1 A         % and remove one
```

C contains the stack's real depth, and this routine uses registers A, B, C, and D1.

57.1 Drop

Coding a DROP is very easy. Here, we don't need to know the stack's real depth. We're just going to read the first's level value. If it's zero, then the stack "should" be empty. If it's not equal to zero, we drop one level.

You know that D1 points to the stack's address. Each time we increment D1 we reduce the stack's size, and each time we decrement it we make the stack bigger. The stack only contains addresses, each one being 5 nibbles wide, so if we read from the stack, we'll use field A.

If we increment D1 by 5, we do a drop. The code below does it:

```
C=DAT1 A
?C=0 A
GOYES QUIT
D1=D1+ 5
D=D+1 A
*QUIT
GOVLNG 2D564
@
```

Here, we start by reading the contents of the first level. If it's null, we quit. At 2D564 is the following code:

```
A=DAT0 A
D0=D0+ 5
PC=(A)
```

The GOSBVL 2D564 uses 7 nibbles in memory, and the usual instructions would need 11 nibbles. Here we save 4 nibbles by using this jump rather than using three instructions.

57.2 Swap

Since the stack only contains addresses, doing a swap is simply exchanging the addresses of level one and level two. First we must ensure that the stack is two levels deep, and if it is, we will exchange the two addresses.

We could do:

```
C=DAT1 A        % read first level's address
?C=0 A          % if it's null, we quit
GOYES QUIT
D1=D1+ 5        % we point to level 2
A=DAT1 A        % and read it into A
D1=D1- 5        % we come back to level 1
?A=0 A          % if A empty, no level 2
GOYES QUIT      % so we quit
DAT1=A A        % otherwise, we write A to DAT1 which moves level 2's address to level 1
D1=D1+ 5        % we point to level 2
DAT1=C A        % and write C there, which moves level 1's address to level 2
D1=D1- 5        % we make D1 come back to level 1
*QUIT
GOVLNG 2D564    % and we quit  :-)
@
```

Try it. If there is an empty stack or only one object it does nothing, but if there are two objects they are swapped! :-)

How does the HP perform a SWAP?

When we call SWAP, an RPL program is called. It contains two calls: the first one checks whether two object's addresses are on stack, and if not, an error is done. Otherwise, the SWAP code is called, and it does a SWAP without checking, because the other piece of code already checked.

It's better because this way, the HP is able to handle a #00000h on the stack, swapping it with another object. Our code will recognize any #00000h as the end of stack, so it will not do anything if the stack contains a #00000h.

Of course, the "usual" HP user should **never** have an <External> on the stack. If they do have one, chances are it's not the #00000h one, whose evaluation does a warmstart. But, Externals are extensively used by System RPL, the basis of User RPL, and it must be able to manage SWAP or any other stack function with Externals.

58 Creating a program for creating macros

You know that you can put parts of your programs into several files, and use the ' symbol to create links. But, as you learned in the HP-ASM documentation, you can also create macros. Remember that a macro is something precompiled that can be included inside any Code object.

For example, suppose we want to write a program that displays "YES" and "NO" choices on the menu bar, and returns <1h> if YES is pressed or <0h> if NO is pressed. We could create a grob, with "YES" and "NO" words inside, display it on the menu area from our program, and wait for a key to be pressed!

So, if we want to put a grob inside of our code, we don't really need to put the whole object inside. We can just put the data we need. ASM Flash and HP-ASM are able to include a precompiled object anywhere in one of our programs. So we are going to write a →MACRO program that does two things:

- If a Code object is given to it, it removes the prologue and size fields and puts a string with the code inside on the stack
- If a GROB object is given to it, remove the prologue, the size field and the 10 nibbles that follow, that are used to code the number of lines and columns.

For example, instead of writing this at the end of a program all the time:

```
A=DAT0 A
D0=D0+ 5
PC=(A)
```

we could compile it and use →MACRO on it. Then, we call it 'RPL' and store it in the Macro directory of HOME. Each time we want to use it in a program, type the following and ASM Flash/HP ASM will put the contents of RPL there:

```
...
/RPL
...
```

This is VERY cool to avoid remembering the addresses of routines in ROM: instead of remembering #2D564h, which contains the RPL exit, we simply code this and use →MACRO on it:

```
GOVLNG 2D564
@
```

Store this as 'RPL' and /RPL can be used in source code. This can be done with any piece of code. :)

The commented source is in the next section. It has been tested, so it should work.

Some assemblers, however, have names built-in for some addresses and functions. For example, the Meta Kernel has a command built-in for exiting to RPL.

58.1 →MACRO.S

```
!0-15
!PC

C=DAT1 A            % read first level's content
?C#0 A              % if it's not null, we start
GOYES BEGIN
LA 00201            % otherwise do an error; Aa contains error code
GOVLNG 05023        % and 05023 does the error

*BEGIN              % here we begin
R1=C A              % save first level's address into R1a
GOSBVL 0679B        % save registers
C=R1 A              % put level one address into C
B=C A               % and save it into B (see why below)
D0=C                % D0 points to object's first nibble
GOSBVL 03019        % call SKIP so D0 points to the first nibble
                    % after the object
CD0EX               % C contains that address
C=C-B A             % calculate (end address - start address)
                    % so C = length of object
R2=C A              % R2a contains object's length (in nibbles)
C=R1 A              % get back object's address saved in R1a
D0=C                % and make D0 point to it
A=DAT0 A            % read its prologue
LC 02DCC            % 02DCC = Code prologue
?A#C A              % if there isn't a code object,
GOYES GROB?         % check if there's a grob
% Code
C=R1 A              % C = address of object
C=C+10 A            % skip the code prologue and size (10 nibbles)
R1=C A              % and save it back to R1a
C=R2 A              % here C = length to allocate for string
C=C-10 A            % remove 10 (since we have 10 fewer nibbles to copy)
R2=C A              % and save it back to R2A

*ALLOCATE
GOSBVL 05B7D        % reserves a string Ca long
CD0EX               % C = address of first nibble of the newly reserved
                    % string
D1=C                % D1 point to it (because the COPY routine I'm going
                    % to use copies from D0 to D1..)
C=R1 A              % C is the "corrected" object's address
                    % corrected = prologue and non-used fields skipped
D0=C                % make D0 point to it
C=R2 A              % C = number of nibbles to copy (R2 = the length
                    % calculated using SKIPOB)
GOSBVL 0670C        % COPY DOWN routine (explained below, with others)
GOSBVL 067D2        % recover D0,D1,B and D registers
C=R0 A              % R0 contains reserved string's address
D=D-1 A             % remove 5 nibbles of the free RAM
```

```
GOC EXIT            % if CARRY = not enough memory to push object
                    % on the stack
D1=D1- 5            % add one object to the stack: D1-5
DAT1=C A            % write object's address (that is: pushes object
                    % on stack)
*EXIT               % here we find the usual return to
GOVLNG 2D564        % RPL

*GROB?              % check if there's a GROB as the object
LC 02B1E            % 02B1E = GROB prologue
?A=C A              % if a GROB, update R1a and R2a
GOYES OK
GOVLNG 05143        % otherwise, 05143 makes us recover the saved
                    % registers, and quits (using the usual RPL return)
*OK                 % here there's a GROB, so:
C=R1 A              % get the start of GROB's object
C=C+16 A            % we add 20 to it, skipping the prologue, the size,
C=C+4 A             % and the ten nibbles of number of lines and columns
R1=C A              % wrote back to R1a
C=R2 A              % C = calculated length
C=C-16 A            % remove 20 from it
C=C-4 A
R2=C A              % and wrote back to R2a
GOTO ALLOCATE       % we reserve, copy and quit

@
```

In this program, R1a and R2a are used. (R1a means field A of R1; this notation comes from Cyrille de Brebinson). R0 isn't used because the routine used to reserve a string sends the address of the reserved string back into R0.

First, the first level of the stack is read. If it's empty, #00201h is loaded into Aa. The routine at #05023h does an error, using the error code inside of Aa (all routines used are described just below the explanations about that code).

If it's not empty, the object's length is calculated. I have used R1a to keep the first object's address, because I'll need it later when I do the copy from the object to my reserved string. I also use the SKIPOB routine located at #03019h, which will be described below.

Once done, D0 points to the first nibble just after the object. An exchange instruction is then used, and then the length of the object is calculated using this formula: length = (end address - start address). R2a is used to keep the object length, which is needed to copy the object later.

Then, I get the object's pointer, and check if its prologue is a Code prologue. If so, I update R1a and R2a, by adding the number of nibbles skipped to R1a (10 for Code, 10 for GROB) and then removing that number from R2a, the number of nibbles to copy. When you want to modify →MACRO so it can handle the object kind you want, just add tests and update R1a and R2a the same way.

Once it's done, we recover the string's length and reserve it using a ROM subprogram found at #05B7Dh, which allocates Ca nibbles in a string object. When it comes back, R0a contains the address of the string, which will be put on the stack later, and D0 points inside the string.

Then, all I do is use the COPY DOWN copy routine in ROM (see below for all routines used) but it copies from D0 to D1. When I return from reserving the string, D0 points to the string. I make D0 point to the object and D1 to the reserved string, and then put the number of nibbles I want to copy into Ca.

Once the copy is done, the object is put on the stack. A GOC is used so if there are not five nibbles available, the object isn't pushed and it quits. The next garbage collection would then remove the string that was reserved.

58.2 ROM routines used in this program

58.2.1 #5023h, also called ERROR_A

Put the usual error code into Aa, and call this routine using a GOVLNG. You **must not** have registers saved when you call it, so if registers need to be recovered, do it first and then call this routine.

58.2.2 #679Bh, also called SAVE_REG

Let's disassemble it and see how it works. Here is the commented source of SAVE_REG:

```
CD0EX          % move D0 address to C
DO= 8072F      % and make D0 point to #8072Fh (see below)
DAT0=C A       % write D0 to it (we save D0)
D0= 806F8      % make D0 point to #806F8
CD1EX          % move D1 address to C
D1=C           % and write it there
DAT0=C A       % so we save D1
D0= 806F3      % make D0 point to #806F3h
C=B A          % move Ba to Ca
DAT0=C A       % and save Ba to #806F3h
D0= 807ED      % make D0 point to #807EDh
C=D A          % move Da to Ca
DAT0=C A       % and save Da to #807EDh
RTNCC          % return AND clear the CARRY flag
```

Notice that the register D0 is lost when we return, but D1, B, and D still have their values. You can see that the HP saves D0, D1, B and D to some addresses in memory: those are reserved just for that. Later in this document, I will describe an area of memory called the "System area." There, we have space reserved so we can save the registers. When we recover the registers (see below) the HP reads where we saved. This is why we can only save registers one time; old values are lost for new ones.

58.2.3 #3019h, also called SKIPOB (Skip OBject)

This one is a little more complicated.

I have myself limited knowledge of RPL, so I don't know how it works exactly.

Before calling this routine, make D0 point to its first nibble. When it returns, D0 points to the first nibble after the object. This routine uses D0, A and C, so if one of those registers contains something you really need, save it somewhere! :)

58.2.4 #05B7Dh, also called RES_STR (REServe STRing)

This ROM routine uses D0, D1, A, B, C, D, R0.

Before calling it, put the number of chars you want reserved into Ca; if there is not memory, an "Insufficient Memory" error is done (registers should be saved before calling this routine, as any error recovers them).

If there is enough space in the temporary memory area, the subroutine returns, and D0 points *inside* the string on its first nibble, and its address is stored in R0a. When you quit the program, you will put that address onto the stack to put the string on it.

58.2.5 #0670Ch, also known as COPY DOWN

I have used a call to this subroutine because memory-handling routines are very well written in the ROM. There are two routines to copy data in memory: COPY DOWN and COPY UP.

What's the difference? When you use COPY DOWN, Ca nibbles are copied from D0 to D1, moving D0 and D1 with "+*n*" (*n* a value the routine sets), and COPY UP copies from D0 to D1, but from the *end* of the zone being copied (D0 and D1 are decremented during the copy).

I use COPY DOWN since I copy from an object to another, from the beginning of each to the end. If I copied backwards in memory, I would have used COPY UP.

When you call COPY DOWN or COPY UP, D0 must point to where data will be read, and D1 to where it will be written. Ca is used to say how many nibbles we want to have copied.

Those two routines (COPY DOWN at #0670Ch and COPY UP at #066B9h) use D0,D1,A and C.

58.2.6 #067D2h, also called GET_REG

This routine restores the previous saved registers; in ROM, we find this:

```
D0= 807ED    % where D is saved
C=DAT0 A
D=C A
DO= 806F3    % where B is saved
C=DAT0 A
B=C A
DO= 806F8    % where D1 is saved
C=DAT0 A
D1=C
D0= 8072F    % where D0 is saved
C=DAT0 A
D0=C
RTNCC
```

Of course, we can call an address *after* 067D2. What does this mean? Easy :-)

If we want to recover D0 or B and D0, we will call the routine, but not at its beginning! This also works for all routines: you can choose from *where* you start it, as long as it finishes with a `RTN`.

Here is the same code as before, but with addresses:

```
#067D2h     D0= 807ED    % where D is saved
#067D9h     C=DAT0 A
#067DCh     D=C A
#067DEh     DO= 806F3    % where B is saved
#067E5h     C=DAT0 A
#067E8      B=C A
#067EAh     DO= 806F8    % where D1 is saved
#067F1h     C=DAT0 A
#067F4h     D1=C
#067F7h     D0= 8072F    % where D0 is saved
#067FEh     C=DAT0 A
#06801h     D0=C
#06804H     RTNCC
```

If you just want to recover D0, simply call `GOSBVL 067F7`.

59 Memory management

Sorry, this section has yet to be written. Would anyone like to help?

60 Graphics

Sorry, this section hasn't been written yet either. Any contributions would be appreciated.

That's all for now. We hope you enjoyed this, and watch for Edition 2 in the future!