

EECS 373

Design of Microprocessor-Based Systems

Robert Dick
University of Michigan

Lecture 2: Architecture, Assembly, and ABI

11 September 2017

Many slides from Mark Brehob

R0
R1
R2
R3
R4
R5
R6
R7
R8
R9
R10
R11
R12
R13 (SP)
R14 (LR)
R15 (PC)
xPSR

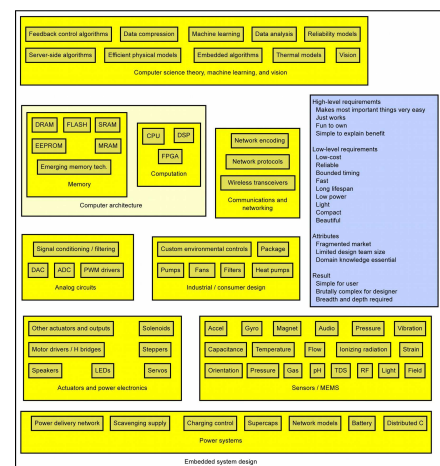
- Course staff
- Implementation technology trends
- Application trends
- Course structure and grading
- Debugging

Outline

- Embedded system
- ISA
- ABI
- Build process



An embedded system



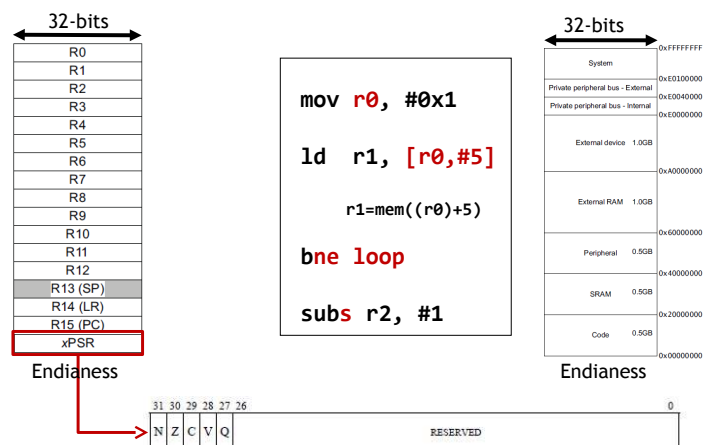
Outline

- Embedded system
- ISA
- ABI
- Build process



Major elements of an Instruction Set Architecture

(registers, memory, word size, endianness, conditions, instructions, addressing modes)





- Little-Endian (default)
 - LSB is at lower address

```

uint8_t a = 1;
uint8_t b = 2;
uint16_t c = 255; // 0x00FF
uint32_t d = 0x12345678;

Memory Offset (LSB) (MSB)
=====
0x0000 01 02 FF 00
0x0004 78 56 34 12

```

- Big-Endian
 - MSB is at lower address

```


uint8_t a = 1;
uint8_t b = 2;
uint16_t c = 255; // 0x00FF
uint32_t d = 0x12345678;

Memory Offset (LSB) (MSB)
=====
0x0000 01 02 00 FF
0x0004 12 34 56 78

```



- Endianness: ordering of bytes within a word
 - Little - increasing numeric significance with increasing memory addresses
 - Big - the opposite, most significant byte first
 - MIPS is big endian, x86 is little endian, ARM supports both



Instruction encoding



- Instructions are encoded in machine language opcodes

Instructions	Register Value	Memory Value
movs r0, #10	001 00 000 00001010 (LSB) (MSB)	(msb) (lsb) 0a 20 00 21
movs r1, #0	001 00 001 00000000	

ARMv7 ARM

Encoding T1 All versions of the Thumb ISA.

MOV<S> <Rd>, #<imm8>

MOV<C> <Rd>, #<imm8>

Outside IT block.

Inside IT block.

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	1	0	0											
Rd								imm8							

d = UInt(Rd); setFlags = !InITBlock(); imm32 = ZeroExtend(imm8, 32); carry = APSR.C;

Assembly example



```

data:
    .byte 0x12, 20, 0x20, -1

func:
    mov r0, #0
    mov r4, #0
    movw r1, #:lower16:data
    movt r1, #:upper16:data

top:
    ldrb r2, [r1], #1
    add r1, r1, #1
    add r4, r4, r2
    add r0, r0, #1
    cmp r0, #4
    bne top

```

Instructions used



- mov
 - Moves data from register or immediate.
 - Or also from shifted register or immediate!
 - the mov assembly instruction maps to a bunch of different encodings!
 - If immediate it might be a 16-bit or 32-bit instruction.
 - Not all values possible
 - why?
- movw
 - Actually an alias to mov.
 - "w" is "wide"
 - hints at 16-bit immediate.

From the ARMv7-M Architecture Reference Manual (posted on the website under references)

Thumb Instruction Details

A6.7.76 MOV (register)

Move (register) copies a value from a register to the destination register. It can optionally update the condition flags based on the value.

Encoding T1 ARMv6-M, ARMv7-M If <db> and <db> both from R0-R7, otherwise all versions of the Thumb ISA.

MOV<C> <db>, <db> If <db> is the PC, must be outside or last in IT block.

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	1	0	0	0	1	1	0								
D								Rm				Rd			

d = UInt(D:Rd); n = UInt(Rm); setFlags = FALSE; if d == 15 && InITBlock() && !LastInITBlock() then UNPREDICTABLE;

Encoding T2 All versions of the Thumb ISA. (formerly LSL, <db>, <db>, #0) Not permitted inside IT block.

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0	0	0	0	0	0							
Rm								Rd							

d = UInt(Rd); n = UInt(Rm); setFlags = TRUE; if InITBlock() then UNPREDICTABLE;

Encoding T3 ARMv7-M

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	1	0	1	0	1	0	1	0	1	1	1	1	1	1	1
D								Rm				Rd			

d = UInt(Rd); n = UInt(Rm); setFlags = (S == 'I'); if setFlags && (d IN {13,15} || n IN {13,15}) then UNPREDICTABLE; if !setFlags && (d == 32 || n == 35 || (d == 13 && n == 33)) then UNPREDICTABLE;

There are similar entries for move immediate, move shifted (which actually maps to different instructions), etc.

Directives

- `#:lower16:data`
 - What does that do?
 - Why?



A6.7.78 MOVN

Move Top writes an immediate value to the top halfword of the destination register. It does not affect the contents of the bottom halfword.

Encoding T1 ARMv7-M

MOVN<> <Rd>, #<imm16>

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	1	1	1	0	i	1	0	1	1	0	0	imm4	0	imm3	Rd																imm8

$d = \text{UInt}(Rd)$; $\text{imm16} = \text{imm4}::\text{imm3}::\text{imm8}$;
if $d \text{ IN } \{13, 15\}$ then UNPREDICTABLE;

Assembler syntax

MOVN<> <Rd>, #<imm16>

where:

<> See Standard assembler syntax fields on page A6-7.

<Rd> Specifies the destination register.

<imm16> Specifies the immediate value to be written to <Rd>. It must be in the range 0-65535.

Operation

```
if ConditionPassed() then
    EncodingSpecificOperations();
    R[d]<31:16> = imm16;
    // R[d]<15:0> unchanged
```

Exceptions

None.



Loads

- **ldrb** -- Load register byte
 - Note this takes an 8-bit value and moves it into a 32-bit location!
 - Zeros out the top 24 bits.
- **ldrsb** -- Load register signed byte
 - Note this also takes an 8-bit value and moves it into a 32-bit location!
 - Uses sign extension for the top 24 bits.



Addressing modes

- **Offset addressing**
 - Offset is added or subtracted from base register
 - Result used as effective address for memory access
 - [**<Rn>**, **<offset>**]
- **Pre-indexed addressing**
 - Offset is applied to base register
 - Result used as effective address for memory access
 - Result written back into base register
 - [**<Rn>**, **<offset>**]!
- **Post-indexed addressing**
 - The address from the base register is used as the EA
 - The offset is applied to the base and then written back
 - [**<Rn>**], **<offset>**



An ISA defines the hardware/software interface



- A contract between architects and programmers
- Register set
- Instruction set
 - Addressing modes
 - Word size
 - Data formats
 - Operating modes
 - Condition codes
- Calling conventions
 - Really not part of the ISA (usually)
 - Rather part of the ABI
 - But the ISA often provides meaningful support.

ARM architecture roadmap





- Skim pages 1-84.
- Read pages 85-154.
- Refer to pages 154-end.

A quick comment on the ISA: From: ARMv7-M Architecture Reference Manual



A4.1 About the instruction set

ARMv7-M supports a large number of 32-bit instructions that were introduced as Thumb-2 technology into the Thumb instruction set. Much of the functionality available is identical to the ARM instruction set supported alongside the Thumb instruction set in ARMv6T2 and other ARMv7 profiles. This chapter describes the functionality available in the ARMv7-M Thumb instruction set, and the *Unified Assembler Language* (UAL) that can be assembled to either the Thumb or ARM instruction sets.

Thumb instructions are either 16-bit or 32-bit, and are aligned on a two-byte boundary. 16-bit and 32-bit instructions can be intermixed freely. Many common operations are most efficiently executed using 16-bit instructions. However:

- Most 16-bit instructions can only access eight of the general purpose registers, R0-R7. These are known as the low registers. A small number of 16-bit instructions can access the high registers, R8-R15.
- Many operations that would require two or more 16-bit instructions can be more efficiently executed with a single 32-bit instruction.

The ARM and Thumb instruction sets are designed to *interwork* freely. Because ARMv7-M only supports Thumb instructions, interworking instructions in ARMv7-M must only reference Thumb state execution, see *ARMv7-M and interworking support* for more details.

In addition, see:

- Chapter A5 *Thumb Instruction Set Encoding* for encoding details of the Thumb instruction set
- Chapter A6 *Thumb Instruction Details* for detailed descriptions of the instructions.

ARM Cortex-M3 ISA



Registers



Instruction Set

ADD Rd, Rn, <op2>

Branching
Data processing
Load/Store
Exceptions
Miscellaneous

Register Set

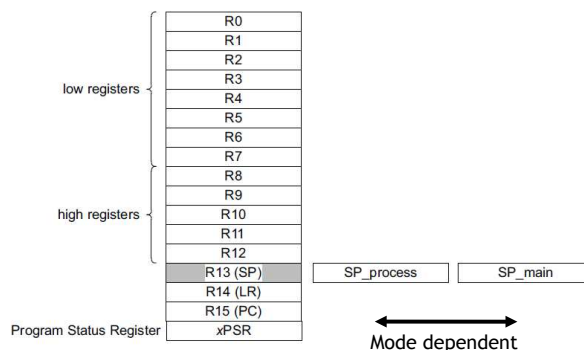
R0
R1
R2
R3
R4
R5
R6
R7
R8
R9
R10
R11
R12
R13 (SP)
R14 (LR)
R15 (PC)
xPSR

32-bits
Endianness

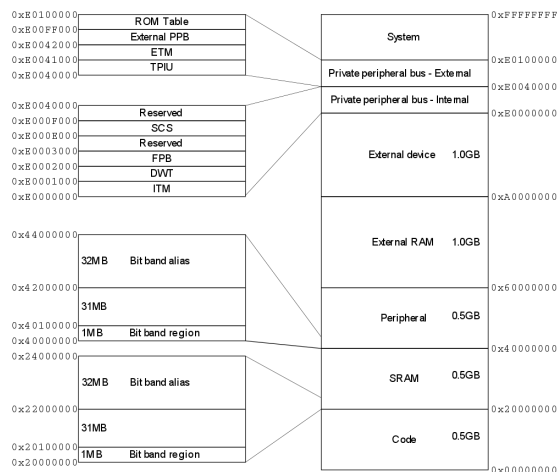
Address Space

System	0xFFFFFFFF
Private peripheral bus - External	0xE0100000
Private peripheral bus - Internal	0xE0040000
External device 1.0GB	0xE0000000
External RAM 1.0GB	0xA0000000
Peripheral 0.5GB	0x60000000
SRAM 0.5GB	0x40000000
Code 0.5GB	0x20000000
	0x00000000

32-bits
Endianness



Address space



Instruction encoding: ADD immediate



Encoding T1 All versions of the Thumb ISA.

ADD <Rd>, <Rn>, #imm3
ADD <Rd>, <Rn>, #imm3

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	1	1	1	0	imm3			Rn	Rd				

Encoding T2 All versions of the Thumb ISA.

ADD <Rd>, #imm8
ADD <Rd>, #imm8

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	1	1	0	Rdn			imm8						

Encoding T3 ARMv7-M

ADD(S) <Rd>, #imm8, #const>

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	1	1	1	0	1	0	0	0	S	Rn	0	imm3		Rd	imm8

Encoding T4 ARMv7-M

ADD <Rd>, #imm12

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	1	1	1	0	1	0	0	0	0	Rn	0	imm3		Rd	imm8

A6.7.3 ADD (immediate)

This instruction adds an immediate value to a register value, and writes the result to the destination register. It can optionally update the condition flags based on the result.

Encoding T1 All versions of the Thumb ISA.

ADDS <Rd>, <Rn>, #<imm3>

Outside IT block.
Inside IT block.

15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0
0 0 0 1 1 0 imm3 Rn Rd

d = UInt(Rd); n = UInt(Rn); setFlags = !IntTBlock(); imm32 = ZeroExtend(imm3, 32);

Encoding T2 All versions of the Thumb ISA.

ADDS <Rd>, #<imm8>

Outside IT block.
Inside IT block.

15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0
0 0 1 1 0 Rdn imm8

d = UInt(Rdn); n = UInt(Rdn); setFlags = !IntTBlock(); imm32 = ZeroExtend(imm8, 32);

Encoding T3 ARMv7-M

ADD(S) <S>, #<imm8>

15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0
1 1 1 1 0 1 0 1 0 0 0 S Rn 0 imm3 Rd imm8

If Rn == '1111' & S == '1' then SEE ON (immediate);
If Rn == '1101' then SEE ADD (SP plus immediate);
d = UInt(Rd); n = UInt(Rn); setFlags = (S == '1'); imm32 = ThumbExpandImm(i:imm3;imm8);
If d IN {13,15} || n == 15 then UNPREDICTABLE;

Encoding T4 ARMv7-M

ADD(S) <S>, <Rn>, #<imm8>

15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0
1 1 1 1 0 1 1 0 0 0 0 0 Rn 0 imm3 Rd imm8

If Rn == '1111' then SEE ADR;
If Rn == '1101' then SEE ADD (SP plus immediate);
d = UInt(Rd); n = UInt(Rn); setFlags = FALSE; imm32 = ZeroExtend(i:imm3;imm8, 32);
If d IN {13,15} then UNPREDICTABLE;

Branch

Table A4-1 Branch instructions

Instruction	Usage	Range
B on page A6-40	Branch to target address	+/-1 MB
CBNZ, CBZ on page A6-52	Compare and Branch on Nonzero, Compare and Branch on Zero	0-126 B
BL on page A6-49	Call a subroutine	+/-16 MB
BLX (register) on page A6-50	Call a subroutine, optionally change instruction set	Any
BX on page A6-51	Branch to target address, change instruction set	Any
TBB, TBH on page A6-258	Table Branch (byte offsets)	0-510 B
	Table Branch (halfword offsets)	0-131070 B

Data processing instructions

Table A4-2 Standard data-processing instructions

Mnemonic	Instruction	Notes
ADC	Add with Carry	-
ADD	Add	Thumb permits use of a modified immediate constant or a zero-extended 12-bit immediate constant.
ADR	Form PC-relative Address	First operand is the PC. Second operand is an immediate constant. Thumb supports a zero-extended 12-bit immediate constant. Operation is an addition or a subtraction.
AND	Bitwise AND	-
BIC	Bitwise Bit Clear	-
ONM	Compare Negative	Sets flags. Like ADD but with no destination register.
OMP	Compare	Sets flags. Like SUB but with no destination register.
EOR	Bitwise Exclusive OR	-
MV	Copies operand to destination	Has only one operand, with the same options as the second operand in most of these instructions. If the operand is a shifted register, the instruction is an LSL, LSR, ASR, or ROR instruction instead. See <i>Shift instructions</i> on page A4-10 for details. Thumb permits use of a modified immediate constant or a zero-extended 16-bit immediate constant.
MVN	Bitwise NOT	Has only one operand, with the same options as the second operand in most of these instructions.

Many more!

Load/store instructions

Table A4-10 Load and store instructions

Data type	Load	Store	Load unprivileged	Store unprivileged	Load exclusive	Store exclusive
32-bit word	LDR	STR	LDRT	STRT	LDREX	STREX
16-bit halfword	-	STRH	-	STRHT	-	STREXH
16-bit unsigned halfword	LDRH	-	LDRHT	-	LDREXH	-
16-bit signed halfword	LDRSH	-	LDRSHT	-	-	-
8-bit byte	-	STRB	-	STRBT	-	STREXB
8-bit unsigned byte	LDRB	-	LDRBT	-	LDREXB	-
8-bit signed byte	LDRSB	-	LDRSBT	-	-	-
two 32-bit words	LDRD	STRD	-	-	-	-

Miscellaneous instructions

Table A4-12 Miscellaneous instructions

Instruction	See
Clear Exclusive	CLREX on page A6-56
Debug hint	DBG on page A6-67
Data Memory Barrier	DMB on page A6-68
Data Synchronization Barrier	DSB on page A6-70
Instruction Synchronization Barrier	ISB on page A6-76
If Then (makes following instructions conditional)	IT on page A6-78
No Operation	NOP on page A6-167
Preload Data	PLD, PLDW (immediate) on page A6-176 PLD (register) on page A6-180
Preload Instruction	PLI (immediate, literal) on page A6-182 PLI (register) on page A6-184
Send Event	SEV on page A6-212
Supervisor Call	SVC (formerly SWI) on page A6-252
Wait for Event	WFE on page A6-276
Wait for Interrupt	WFI on page A6-277
Yield	YIELD on page A6-278

Addressing Modes (again)

- Offset Addressing
 - Offset is added or subtracted from base register
 - Result used as effective address for memory access
 - [<Rn>, <offset>]
- Pre-indexed Addressing
 - Offset is applied to base register
 - Result used as effective address for memory access
 - Result written back into base register
 - [<Rn>, <offset>]!
- Post-indexed Addressing
 - The address from the base register is used as the EA
 - The offset is applied to the base and then written back
 - [<Rn>], <offset>

The ARM architecture “books” for this class



The ARM software tools “books” for this class



Outline



- Embedded system
- ISA
- ABI
- Build process

ABI summary



Detailed version

- Pass: r0-r3
- Return: r0 or r0-r1
- Callee saved variables: r4-r8, r11, maybe r9, r10
- Static base: r9
- Stack limit checking: r10
- Veneers, scratch: r12
- Stack pointer: r13
- Link register (function call return address): r14
- Program counter: r15

Simple version

- Callee preserves r4-r11 and r13
- Caller preserves r0-r3

ABI details



1. A subroutine must preserve the contents of the registers r4-r8, r11, maybe r9-r10
2. Arguments are passed through r0 to r3
 - If you need more, we put a pointer into memory in one of the registers.
3. Return value is placed in r0 or r0-r1
4. Allocate space on stack as needed. Use it as needed.
 - Reset stack pointer when done
 - Word align

Outline



- Embedded system
- ISA
- ABI
- Build process

An ARM assembly language program for GNU



```
.equ STACK_TOP, 0x20000800
.text
.syntax unified
.thumb
.global _start
.type start, %function

_start:
.word STACK_TOP, start
start:
movs r0, #10
movs r1, #0
loop:
adds r1, r0
subs r0, #1
bne loop
deadloop:
b deadloop
.end
```

A simple Makefile



```
all:
arm-none-eabi-as -mcpu=cortex-m3 -mthumb example1.s -o example1.o
arm-none-eabi-ld -Ttext 0x0 -o example1.out example1.o
arm-none-eabi-objcopy -Obinary example1.out example.bin
arm-none-eabi-objdump -S example1.out > example1.list
```

An ARM assembly language program for GNU



```
.equ STACK_TOP, 0x20000800
.text
.syntax unified
.thumb
.global _start
.type start, %function

_start:
.word STACK_TOP, start
start:
movs r0, #10
movs r1, #0
loop:
adds r1, r0
subs r0, #1
bne loop
deadloop:
b deadloop
.end
```

Disassembled object code



example1.out: file format elf32-littlearm

Disassembly of section .text:


```
00000000 <_start>:
0: 20000800 .word 0x20000800
4: 00000000 .word 0x00000000

00000008 <start>:
8: 200a      movs r0, #10
a: 2100      movs r1, #0

0000000c <loop>:
c: 1809      adds r1, r1, r0
e: 3801      subs r0, #1
10: d1fc      bne.n c <loop>

00000012 <deadloop>:
12: e7fe      b.n 12 <deadloop>
```

How does an assembly language program get turned into a executable program image?



What are the real GNU executable names for the ARM?



- Just add the prefix “arm-none-eabi-” prefix
- Assembler (as)
 - arm-none-eabi-as
- Linker (ld)
 - arm-none-eabi-ld
- Object copy (objcopy)
 - arm-none-eabi-objcopy
- Object dump (objdump)
 - arm-none-eabi-objdump
- C Compiler (gcc)
 - arm-none-eabi-gcc
- C++ Compiler (g++)
 - arm-none-eabi-g++

A simple (hardcoded) Makefile example



```
all:
    arm-none-eabi-as -mcpu=cortex-m3 -mthumb example1.s -o example1.o
    arm-none-eabi-ld -Ttext 0x0 -o example1.out example1.o
    arm-none-eabi-objcopy -Obinary example1.out example1.bin
    arm-none-eabi-objdump -S example1.out > example1.lst
```

What information does the disassembled file provide?



```
all:
    arm-none-eabi-as -mcpu=cortex-m3 -mthumb example1.s -o example1.o
    arm-none-eabi-ld -Ttext 0x0 -o example1.out example1.o
    arm-none-eabi-objcopy -Obinary example1.out example1.bin
    arm-none-eabi-objdump -S example1.out > example1.lst
```

```
.equ STACK_TOP, 0x20000800
.text
.syntax unified
.thumb
.global _start
.type start, %function
```

```
_start:
    .word STACK_TOP, start
start:
    movs r0, #10
    movs r1, #0
loop:
    adds r1, r0
    subs r0, #1
    bne loop
deadloop:
    b deadloop
.end
```

```
example1.out:      file format elf32-littlearm
Disassembly of section .text:
```

```
00000000 <start>:
0: 20000800    .word 0x20000800
4: 00000000    .word 0x00000000

00000008 <start>:
8: 200a      movs r0, #10
a: 2100      movs r1, #0

0000000c <loop>:
c: 1809      adds r1, r1, r0
e: 3801      subs r0, #1
10: d1fc     bne.n c, <loop>

00000012 <deadloop>:
12: e7fe     b.n 12 <deadloop>
```

Elements of assembly language program?



```
.equ STACK_TOP, 0x20000800 /* Equates symbol to value */
.text                     /* Tells AS to assemble region */
.syntax unified           /* Means language is ARM UAL */
.thumb                   /* Means ARM ISA is Thumb */
.global _start           /* .global exposes symbol */
/* _start label is the beginning */
/* ...of the program region */
/* Specifies start is a function */
.type start, %function
/* start label is reset handler */

_start:
    .word STACK_TOP, start /* Inserts word 0x20000800 */
/* Inserts word (start) */

start:
    movs r0, #10           /* We've seen the rest ... */
    movs r1, #0

loop:
    adds r1, r0
    subs r0, #1
    bne loop

deadloop:
    b deadloop
.end
```


How are assembly files assembled?



- \$ arm-none-eabi-as
 - Useful options
 - mcpu
 - mthumb
 - o

```
$ arm-none-eabi-as -mcpu=cortex-m3 -mthumb example1.s -o example1.o
```

How does a mixed C/Assembly program get turned into a executable program image?



Register	Synonym	Special	Role in the procedure call standard
r15		PC	The Program Counter.
r14		LR	The Link Register.
r13		SP	The Stack Pointer.
r12		IP	The Intra-Procedure-call scratch register.
r11	v8		Variable-register 8.
r10	v7		Variable-register 7.
r9		v6 SB TR	Platform register. The meaning of this register is defined by the platform standard.
r8	v5		Variable-register 5.
r7	v4		Variable register 4.
r6	v3		Variable register 3.
r5	v2		Variable register 2.
r4	v1		Variable register 1.
r3	a4		Argument / scratch register 4.
r2	a3		Argument / scratch register 3.
r1	a2		Argument / result / scratch register 2.
r0	a1		Argument / result / scratch register 1.



Done.