ECE 471 – Embedded Systems
Lecture 8

Vince Weaver

https://web.eece.maine.edu/~vweaver

vincent.weaver@maine.edu

21 September 2026

Announcements

- HW#1 was finally graded
- HW#3 was assigned, sorry for the delay
- Most of the issue was making sure you can do it on 64-bit operating system

HW#3 Notes – Getting Weird Errors

- If the code won't assemble with errors about comment char it's often because you are compiling on a non-ARM32 machine
- This will happen if you run "make" on x86
- Be sure you edit the code in the right directory, 32 or 64
- Use `uname -m` to see what kind of machine you are on

HW#3 Notes – Why not a Cross Compiler?

- Some architectures like x86 are strongly backwards compatible, gcc can generate 32 bit code with just `-m32` and generally 32-bit code will run on 64 bit machines (though you might have to install some extra 32-bit libraries first)
- ARM64 vs ARM32 has more differences and PiOS currently doesn't seem to be handling things well
- installing an armhf cross-compiler is how you'd build 32-bit code on 64-bit, but the version in PiOS links wrong (known problem?)
- Running code compiled on actual 32-bit system copied over does not work on Pi5. 16k vs 4k page alignment issue? 32-bit packages installed with apt-get work so

HW#3 Notes – Printing an Integer

- I don't make you do this anymore as part of the assignment but you might be curious about the code
- Writing int to string conversion is a complex task
  There are lots of ways to do it.
- When would you ever need code like this?
  In extreme embedded systems cases you might not have a printf() but still want to debug

HW#3 Notes – Integer to String Algorithm

- Take integer
- Divide by 10, put remainder into array backwards
- Take quotient as next source and repeat until zero
- Also need to convert to ASCII. (by adding `0x30` or '0')

HW3 Notes – ASCII

- American Standard Code for Information Interchange
- Old (late 1960s) standard for mapping text characters to numbers
- 7-bits (top bit either 0 or used for other purposes)
- Below 32 are control chars (like linefeed)
- 32 is space
- 48-57 is 0-1
- 65-90 is A-Z, 97-122 is a-z (bit 5 flipped)

HW3 Notes – Unicode

- what about other languages?
- Unicode, in theory 32 bit should hold all possible
- Windows and Java used 16-bit chars, but turned out not to be enough
- UTF-8 is interesting hack where bottom 127 chars map to ASCII, but when top bit set starts a complicated escape sequence that allows encoding any unicode value in 1 to 5 bytes
- still gives benefit to American English

HW3 Notes – Division if no Divide?

- Original Pi-1B ARM1176 has no divide instruction
- To be backwards compatible even new Pis are compiled w/o divide even though new chips have support
- Various ways in software. Iterative subtraction. Shift and subtract. Newton's method
- For constant values you can instead multiply by the reciprocal
- gcc will do this. It use 32.32 fixed point multiply by 1/10. (0xcccccccd).
- ARM has `umull` instruction that will do a 32x32 multiply and give you the top half of the 64-bit result.

HW#3 Notes – Corner cases in Integer Conver

- Leading zero removal
- Signed numbers (put a '-' in front?)

Why code in Assembly?

- Small binaries
  still useful on small embedded boards
- Improved performance
  skilled programmer can still beat compiler
  harder on modern CPUs where constantly changing underlying architecture makes old optimizations obsolete
  often can optimize the small critical parts of algorithms

ARM32 encoding

```
ADD{S}<c> <Rd>,<Rn>,<Rm>{,<shift>}
```

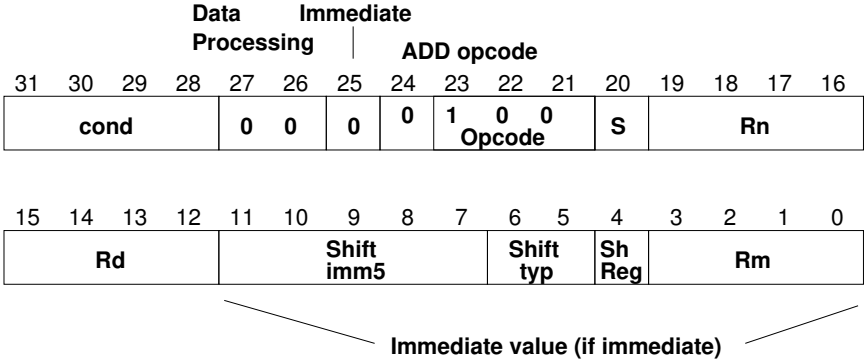

Really Brief Overview of ARM32 Assembler

ARM32 Registers

- Processor has 16 registers r0.. r15
  - SP (r13) (stack pointer) (points to LIFO stack)
  - LR (r14) (link register) when you branch and link to a function, gets the value you are branching from so you can return to it
    how do you handle non-leaf functions? (save on stack usually)
  - PC (r15) (program counter) on original ARM design they were clever and had this as a normal register that you can operate on with regular instructions. In practice not the best ide

ARM32 Instructions

- Instructions have three arguments, destination and two sources
  `add r0,r1,r2` means r0 = r1 + r2
- Assume you have infinite memory, accessible with pointers to bytes.
  With an OS you should allocate it before just using it.

ARM32 Instruction Types

- There are many, possibly hundreds
- math: add, sub, mul, div
- bitwise logical: and, orr, eor, bic
- loads/stores: ldr, str
- compare: cmp
- branch: b, beq, bne, bl, blx
- floating point, other extensions

Weird ARM32 Things – Barrel Shifter

- Most ALU instructions you can include a shift/rotate for free
- `add r0,r0,r0 LSL #2`
  Sneaky way to multiply by 5
  `x=x+(x*4)=5x`
- Helps with code density
- LSL, LSR, ASR, ROR, RRX (note, no ASL, why?)

Weird ARM32 Things – Constants

- Loading constant values into registers is hard
- Instructions with immediate values can load any 8-bit value (0..255) rotated by an even amount (4 bits * 2)
- Can load negative with MVN
- Possibly load 16-bit constant with MOV
- Can use the = trick for 32-bit, as you can't fit a 32-bit constant into a 32-bit instruction

Weird ARM32 Things – Addressing modes

- ldr, ldrb, ldrsb, etc
- `ldrb r1,[r2]` – load value in address/pointer in r2
- `ldrb r1,[r2,#20]` – offset r2+ value, useful in structs
- `ldrb r1,[r2,r3]` – add r2+r3 for address
- `ldrb r1,[r2,r3,LSL #2]` – shift before adding, useful array index
- `ldrb r1,[r2,#20]!` – load from r2+20, then update address into r2
- `ldrb r1,[r2],#20` – load from r2, then an and update address into r2

ARM32 Conditional Execution

- Why are branches bad?

```
if (x == 5 )
   a+=2;
else
   b-=2;
```

- Assembly with branches

```
   cmp   r1, #5
   bne   else
   add   r2,r2,#2
   b     done
else:
   sub   r3,r3,#2
done:
```

- Assembly w/o branches

```
   cmp     r1, #5
   addeq   r2,r2,#2
   subne   r3,r3,#2
```

ARM32 Processor Flags

- Zero, oVerflow, Negative, few others
- Can set flags with CMP instruction
- Cat set flags with S in instruction (ADDS)
- Conditional branches use the flags
- Conditional execution uses the flags

ARM32 Assembly – Register moves

- Moving register values around
- `mov r0,r1` – r0 is destination
- `mov r0,#0` – move immediate value
- There are also `msr` and `mrs` to move into special system variables

ARM32 Assembly – Load/Stores

- ARM32 is load/store: must load into register before using
- `ldrb r0,[address]` – load byte into r0 from pointer
- `strb r0,[address]` – store byte r0 to mem at pointer
- can support different widths (ldr, ldrb, ldrh, etc)
- sign vs zero extend (lsrsb)
- Complex addressing modes. register, r1+r2, r1+r2+offset, auto-increment, etc

ARM32 Addressing Modes – Regular

- `ldrb r1, [r2]` @ register
- `ldrb r1, [r2,#20]` @ register/offset
- `ldrb r1, [r2,+r3]` @ register + register
- `ldrb r1, [r2,-r3]` @ register - register
- `ldrb r1, [r2,r3, LSL #2]` @ register +/- register, shift

ARM32 Addressing Modes – Pre-Index

- Calculate address, load, then store back
- `ldrb r1, [r2, #20]!` @ pre-index. Load from r2+20 then write back into r2
- `ldrb r1, [r2, r3]!` @ pre-index. register
- `ldrb r1, [r2, r3, LSL #4]!` @ pre-index. shift

ARM32 Addressing Modes – Post-Index

- load from base, then add in and write new value to base
- `ldrb r1, [r2],#+1` @ post-index. load, then add value to r2
- `ldrb r1, [r2],r3` @ post-index register
- `ldrb r1, [r2],r3, LSL #4` @ post-index shift

ARM32 Assembly – Arithmetic

- add, sub, ...
- `add r0,r1,r2`
- `add r0,r1,#0`
- Barrel shifter allows complex stuff like `add r0,r1,r2 LSL #4` to optionally shift/rotate

ARM32 Assembly – Logic

- and, orr, eor
- `and r0,r1,r2`
- `eor r0,r1,#0`
- Barrel shifter too

ARM32 Assembly – Comparison

- `cmp r0,r1` – sets flags
- Same as a subtract but doesn't update destination
- Can do same thing with arithmetic if you add 'S' `adds r0,r1,r2`

ARM32 Assembly – Branches

- Branch based on previous comparison
- beq, blt, bgt, etc
- b – unconditional
- bl – branch and link, calls a function and puts return value in special LR (link register)

ARM32 Assembly – Stack Manipulation

- Old "store multiple" instructions, really powerful, can use any arbitrary reg as stack, arbitrary number of registers to push/pop, can change direction and post or pre-increment
  `ldmia sp!, {r0, r1, r2, r3, ip, pc}^`
- Modern also supports `push {r0, r1}` and `pop {r0,r1}`
- On ARM32 Program Counter (PC) is a regular register. Code will often `push {r0, LR}` at beginning of function to save return, then `pop {r0, PC}` at end which puts LR back into PC to return without an extra `bl LR` instruction

Setting Flags

- `add r1,r2,r3`
- `adds r1,r2,r3` – set condition flag
- `addeqs r1,r2,r3` – set condition flag and prefix compiler and disassembler like `addseq`, GNU as doesn't?

Arithmetic Instructions

Operate on 32-bit integers. Most of these take optional `s` to set status flag

adc	v1	add with carry
add	v1	add
rsb	v1	reverse subtract (immediate - rX)
rsc	v1	reverse subtract with carry
sbc	v1	subtract with carry
sub	v1	subtract

Logic Instructions

and	v1	bitwise and
bfc	??	bitfield clear, clear bits in reg
bfi	??	bitfield insert
bic	v1	bitfield clear: and with negated value
clz	v7	count leading zeros
eor	v1	exclusive or (name shows 6502 heritage)
orn	v6	or not
orr	v1	bitwise or

Register Manipulation

mov, movs	v1	move register
mvn, mvns	v1	move inverted

Loading Constants

- In general you can get a 12-bit immediate which is 8 bits of unsigned and 4-bits of even rotate (rotate by 2*value).

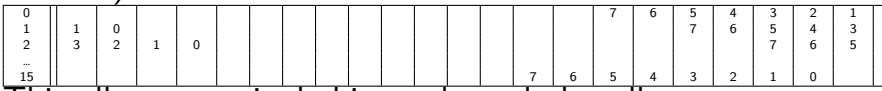

  This allows any single bit mask, and also allows masking of any four sub-bytes.

- You can specify you want the assembler to try to make the immediate for you: `ldr r0,=0xff`
  `ldr r0,=label`
  If it can't make the immediate value, it will store in nearby in a `literal pool` and do a memory read.

Barrel Shift in ALU instructions

If second source is a register, can optionally shift:

- LSL – Logical shift left
- LSR – Logical shift right
- ASR – Arithmetic shift right
- ROR – Rotate Right
- RRX – Rotate Right with Extend
  bit zero into C, C into bit 31 (33-bit rotate)

Barrel Shift Continued

- Why no ASL?
- Adding s `lsls`, `lsrs` puts shifted out bit into C.
- shift pseudo instructions
  `lsr r0, #3` is same as `mov r0,r0 LSR #3`
- For example:
  `add r1, r2, r3, lsr #4`
  `r1 = r2 + (r3>>4)`
- Another example (what does this do):
  `add r1, r2, r2, lsl #2`

Multiply Instructions

Fast multipliers are optional

For 64-bit results,

mla	v2	multiply two registers, add in a third (4 arguments)
mul	v2	multiply two registers, only least sig 32bit saved
smlal	v3M	32x32+64 = 64-bit (result and add source, reg pair rdhi,rdlo)
smull	v3M	32x32 = 64-bit
umlal	v3M	unsigned 32x32+64 = 64-bit
umull	v3M	unsigned 32x32=64-bit

Divide Instructions

- On some machines it's just not there. Original Pi. Why?
- What do you do if you want to divide?
- Shift and subtract (long division)
- Multiply by reciprocal.

Prefixed instructions

Most instructions can be prefixed with condition codes:

EQ, NE	(equal)	Z==1/Z==0
MI, PL	(minus/plus)	N==1/N==0
HI, LS	(unsigned higher/lower)	C==1&Z==0/C==0|Z==1
GE, LT	(greaterequal/lessthan)	N==V/N!=V
GT, LE	(greaterthan, lessthan)	N==V&Z==0/N!=V|Z==1
CS,HS, CC,LO	(carry set,higher or same/clear)	C==1,C==0
VS, VC	(overflow set / clear)	V==1,V==0
AL	(always)	(this is the default)

Load/Store Instructions

ldr	v1	load register
ldrb	v1	load register byte
ldrd	v5	load double, into consecutive registers (Rd even)
ldrh	v1	load register halfword, zero extends
ldrsb	v1	load register signed byte, sign-extends
ldrsh	v1	load register halfword, sign-extends
str	v1	store register
strb	v1	store byte
strd	v5	store double
strh	v1	store halfword

Addressing Modes

- Regular
  - `ldrb r1, [r2]` @ register
  - `ldrb r1, [r2,#20]` @ register/offset
  - `ldrb r1, [r2,+r3]` @ register + register
  - `ldrb r1, [r2,-r3]` @ register - register
  - `ldrb r1, [r2,r3, LSL #2]` @ register +/- register, shift
- Pre-index. Calculate address, load, then store back
  - `ldrb r1, [r2, #20]!` @ pre-index. Load from r2+20 then write back into r2
  - `ldrb r1, [r2, r3]!` @ pre-index. register
  - `ldrb r1, [r2, r3, LSL #4]!` @ pre-index. shift
- Post-index: load from base, then add in and write new

Why some of these?

- `ldrb r1, [r2,#20]` @ register/offset
  Useful for structs in C (i.e. something.else=4;)
- `ldrb r1, [r2,r3, LSL #2]` @ register +/- register, shift
  Useful for indexing arrays of integers (a[5]=4;)

Comparison Instructions

Updates status flag, no need for s

cmp	v1	compare (subtract but discard result)
cmn	v1	compare negative (add)
teq	v1	tests if two values equal (xor) (preserves carry)
tst	v1	test (and)

Control-Flow Instructions

Can use all of the condition code prefixes.
Branch to a label, which is +/- 32MB from PC

b	v1	branch
bl	v1	branch and link (return value stored in lr)
bx	v4t	branch to offset or reg, possible THUMB switch
blx	v5	branch and link to register, with possible THUMB switch
mov pc,lr	v1	return from a link

Load/Store multiple (stack?)

In general, no interrupt during instruction so long instruction can be bad in embedded
Some of these have been deprecated on newer processors

- ldm – load multiple memory locations into consecutive registers
- stm – store multiple, can be used like a PUSH instruction
- push and pop are thumb equivalent

Load/Store multiple Continued

Can have address mode and ! (update source):

- IA – increment after (start at Rn)
- IB – increment before (start at Rn+4)
- DA – decrement after
- DB – decrement before

Can have empty/full. Full means SP points to a used location, Empty means it is empty:

- FA – Full ascending
- FD – Full descending
- EA – Empty ascending
- ED – Empty descending

Load/Store multiple Continued

Recent machines use the "ARM-Thumb Proc Call Standard" which says a stack is Full/Descending, so use LDMFD/STMFD.
What does `stm SP!, {r0,lr}` then `ldm SP!, {r0,PC,pc}` do?

System Instructions

- svc, swi – software interrupt
  takes immediate, but ignored.
- mrs, msr – copy to/from status register. use to clear interrupts? Can only set flags from userspace
- cdp – perform coprocessor operation
- mrc, mcr – move data to/from coprocessor
- ldc, stc – load/store to coprocessor from memory

Co-processor 15 is the *system control coprocessor* and is used to configure the processor. Co-processor 14 is the debugger 11 is double-precision floating point 10 is single-precision fp as well as VFP/SIMD control 0-7 vendor specific

Other Instructions

- swp – atomic swap value between register and memory (deprecated armv7)
- ldrex/strex – atomic load/store (armv6)
- wfe/sev – armv7 low-power spinlocks
- pli/pld – preload instructions/data
- dmb/dsb – memory barriers

Pseudo-Instructions

adr		add immediate to PC, store address in reg
nop		no-operation

Fancy ARMv6

- mla – multiply/accumulate (armv6)
- mls – multiply and subtract
- pkh – pack halfword (armv6)
- qadd, qsub, etc. – saturating add/sub (armv6)
- rbit – reverse bit order (armv6)
- rbyte – reverse byte order (armv6)
- rev16, revsh – reverse halfwords (armv6)
- sadd16 – do two 16-bit signed adds (armv6)
- sadd8 – do 4 8-bit signed adds (armv6)
- sasx – (armv6)
- sbfx – signed bit field extract (armv6)

Floating Point

ARM floating point varies and is often optional.

- various versions of vector floating point unit
- vfp3 has 16 or 32 64-bit registers
- Advanced SIMD – reuses vfp registers
  Can see as 16 128-bit regs q0-q15 or 32 64-bit d0-d31 and 32 32-bit s0-s31
- SIMD supports integer, also 16-bit?
- Polynomial?
- FPSCR register (flags)