

CS 3330: ISAs and Y86-64

8 September 2016

Changelog

Corrections made in this version not in first posting:

10 Sep 2016: slide 24: jle looks for ZF = 1, not ZF = 0

1

1

ISAs being manufactured today

x86 — dominant in desktops, servers

ARM — dominant in mobile devices

POWER — Wii U, IBM supercomputers

MIPS — common in consumer wifi access points

SPARC — some Oracle servers, Fujitsu supercomputers

z/Architecture — IBM mainframes

Z80 — TI calculators

SHARC — some digital signal processors

Itanium — some HP servers

...

Instruction Set Architecture

the instructions

types of operations?

registers/storage?

how their machine code looks

2

3

ISA variation

instruction set	instr. length	# normal registers	approx. # instrs.
x86-64	1–15 byte	16	1500
Y86-64	1–10 byte	15	18
ARMv7	4 byte*	16	400
POWER8	4 byte	32	1400
MIPS32	4 byte	31	200
Itanium	41 bits*	128	300
Z80	1–4 byte	7	40
VAX	1–14 byte	8	150
z/Architecture	2–6 byte	16	1000

4

Other choices: condition codes?

instead of:

```
testq %r11, %r12
je somewhere
```

could do:

```
/* _B_ranch if _EQ_ual */
beq %r11, %r12, somewhere
```

5

Other choices: addressing modes

ways of specifying **operands**. examples:

x86-64: 10(%r11,%r12,4)

ARM: %r11 << 3 (shift register value by constant)

VAX: ((%r11)) (register value is pointer to pointer)

Other choices: number of operands

add dest, src1, src2
ARM, POWER, MIPS, SPARC, ...

add dest=src1, src2 (or swapped)
x86, AVR, Z80, ...

VAX: both

6

7

CISC and RISC

RISC — Reduced Instruction Set Computer

Reduced from what?

Other choices: instruction complexity

instructions that write multiple values?

x86-64: **push, pop, movs, ...**

more?

8

9

Some VAX instructions

MATCHC *haystackPtr*, *haystackLen*, *needlePtr*, *needleLen*
Find the position of the string in needle within haystack.

POLY *x*, *coefficientsLen*, *coefficientsPtr*
Evaluate the polynomial whose coefficients are pointed to by
coefficientPtr at the value *x*.

EDITPC *sourceLen*, *sourcePtr*, *patternLen*, *patternPtr*
Edit the string pointed to by *sourcePtr* using the pattern
string specified by *patternPtr*.

microcode

MATCHC *haystackPtr*, *haystackLen*, *needlePtr*, *needleLen*
Find the position of the string in needle within haystack.

loop in hardware???

typically: lookup sequence of **microinstructions**

secret simpler instruction set

10

11

Why RISC?

complex instructions were not faster

complex instructions were harder to implement

compilers, not hand-written assembly

Typical RISC ISA properties

fewer, simpler instructions

seperate instructions to access memory

fixed-length instructions

more registers

no “loops” within single instructions

no instructions with two memory operands

few addressing modes

Y86-64 instruction set

based on x86

omits most of the 1000 instructions

leaves

addq jmp pushq

subq jCC popq

andq cmoveq movq (renamed)

xorg call

`nop` `ret`

Y86-64: movq

source | destination

i — immediate
r — register
m — memory

SDmovq

irmovq ~~imnovq~~

rrmovg **rmmovg**

~~mrmoya~~ ~~mrmoya~~

Y86-64: specifying addresses

Valid: `rmmovq %r11, 10(%r12)`

Invalid: ~~`rmmovq %r11, 10(%r12,%r13)`~~

Invalid: ~~`rmmovq %r11, 10(,%r12,4)`~~

Invalid: ~~`rmmovq %r11, 10(%r12,%r13,4)`~~

Y86-64: accessing memory (1)

$r12 \leftarrow \text{memory}[10 + r11] + r12$

Invalid: ~~`addq 10(%r11), %r12`~~

Instead:

`mrmovq 10(%r11), %r11`
/* overwrites %r11 */

`addq %r11, %r12`

Y86-64: accessing memory (2)

$r12 \leftarrow \text{memory}[10 + 8 * r11] + r12$

Invalid: ~~`addq 10(,%r11,8), %r12`~~

Instead:

/* replace %r11 with 8*r11 */
`addq %r11, %r11`
`addq %r11, %r11`
`addq %r11, %r11`

`mrmovq 10(%r11), %r11`
`addq %r11, %r12`

Y86-64 constants (1)

`irmovq $100, %r11`

only instruction with non-address constant operand

Y86-64 constants (2)

$r12 \leftarrow r12 + 1$

Invalid: ~~addq \$1, %r12~~

Instead, need an extra register:

```
irmovq $1, %r11  
addq %r11, %r12
```

Y86-64: operand uniqueness

only one kind of value for each operand

instruction name tells you the kind

(why **movq** was ‘split’ into four names)

20

21

Y86-64: Condition codes

ZF — value was zero?

SF — sign bit was set? i.e. value was negative?

set by all arithmetic

Y86-64: Simple condition codes (1)

If %r9 is -1 and %r10 is 1:

subq %r10, %r9

r9 becomes $-1 - (1) = -2$.

SF = 1 (negative)

ZF = 0 (not zero)

andq %r10, %r10

r10 becomes 1

SF = 0 (non-negative)

ZF = 0 (not zero)

22

23

Y86-64: Using condition codes

	subq SECOND, FIRST (value = FIRST - SECOND) j_ or condition code bit test cmov_	value test
le	SF = 1 or ZF = 1	value ≤ 0
l	SF = 1	value < 0
e	ZF = 1	value = 0
ne	ZF = 0	value $\neq 0$
ge	SF = 0	value ≥ 0
g	SF = 0 and ZF = 0	value > 0

24

Y86-64: Conditionals (1)

~~cmp, test~~

instead: use side effect of normal arithmetic

instead of

`cmpq %r11, %r12
jle somewhere`

maybe:

`subq %r11, %r12
jle`

(but changes %r12)

25

Y86-64: Conditionals (2)

~~cmp, test~~

instead: use side effect of normal arithmetic

instead of

`test %r11, %r11
je somewhere`

use:

`andq %r11, %r11
je somewhere`

(doesn't change %r11)

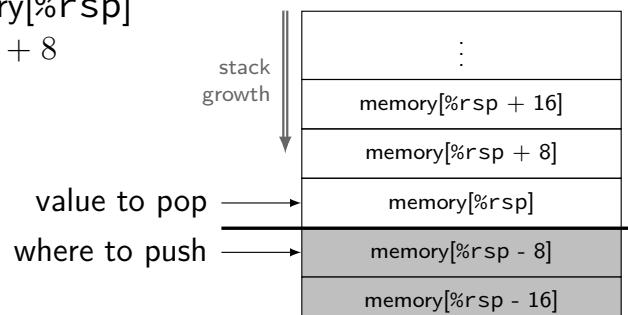
push/pop

`pushq %rbx`

$\%rsp \leftarrow \%rsp - 8$
 $\text{memory}[\%rsp] \leftarrow \%rbx$

`popq %rbx`

$\%rbx \leftarrow \text{memory}[\%rsp]$
 $\%rsp \leftarrow \%rsp + 8$



26

27

Y86-64 state

%rXX — 15 registers

%r15 missing

smaller parts of registers missing

ZF (zero), SF (sign), OF (overflow)

book has OF, we'll not use it

CF (carry) missing

Stat — processor status — halted?

PC — program counter (AKA instruction pointer)

main memory

Typical RISC ISA properties

fewer, simpler instructions

seperate instructions to access memory

fixed-length instructions

more registers

no “loops” within single instructions

no instructions with two memory operands

few addressing modes

Y86-64 instruction formats

byte:	0	1	2	3	4	5	6	7	8	9
halt	0	0								
nop	1	0								
rrmovq/cmovCC rA, rB	2	cc	rA	rB						
irmovq V, rB	3	0	F	rB	V					
rmmovq rA, D(rB)	4	0	rA	rB	D					
mrmovq D(rB), rA	5	0	rA	rB	D					
OPq rA, rB	6	fn	rA	rB						
jCC Dest	7	cc			Dest					
call Dest	8	0			Dest					
ret	9	0								
pushq rA	A	0	rA	F						
popq rA	B	0	rA	F						

Secondary opcodes: cmovcc/jcc

byte:	0	1	2	3	4	5	6	7	8	9
halt	0	0								
nop	1	0								
rrmovq/cmovCC rA, rB	2	cc	rA	rB						
irmovq V, rB	3	0	F	rB	V					
rmmovq rA, D(rB)	4	0	rA	rB	D					
mrmovq D(rB), rA	5	0	rA	rB	D					
OPq rA, rB	6	fn	rA	rB						
jCC Dest	7	cc			Dest					
call Dest	8	0			Dest					
ret	9	0								
pushq rA	A	0	rA	F						
popq rA	B	0	rA	F						

Legend:

- 0 always (jmp/rrmovq)
- 1 le
- 2 l
- 3 e
- 4 ne
- 5 ge
- 6 g

Secondary opcodes: OPq

byte:	0	1	2	3	4	5	6	7	8	9
halt	0	0								
nop	1	0								
rrmovq/cmovCC rA, rB	2	cc	rA	rB						
irmovq V, rB	3	0	F	rB	V	D	D			
rmmovq rA, D(rB)	4	0	rA	rB						
mrmovq D(rB), rA	5	0	rA	rB						
OPq rA, rB	6	fn	rA	rB	0	add	1	sub	2	and
jCC Dest	7	cc								
call Dest	8	0								
ret	9	0								
pushq rA	A	0	rA	F						
popq rA	B	0	rA	F						

Registers: rA, rB

byte:	0	1	2	3	4	5	6	7	8	9
halt	0	0								
nop	1	0								
rrmovq/cmovCC rA, rB	2	cc	rA	rB	0	%rax	1	%rcx	2	%rdx
irmovq V, rB	3	0	F	rB	2	%rbx	3	%rsp	4	%rbp
rmmovq rA, D(rB)	4	0	rA	rB	3	%rsi	5	%rdi	6	%r8
mrmovq D(rB), rA	5	0	rA	rB	4	%r9	6	%r10	7	%r11
OPq rA, rB	6	fn	rA	rB	5	%r12	7	%r13	8	%r14
jCC Dest	7	cc			6	Dest	7	Dest	8	Dest
call Dest	8	0			7	Dest	8	Dest	9	Dest
ret	9	0			9	Dest	10	Dest	11	Dest
pushq rA	A	0	rA	F	10	Dest	11	Dest	12	Dest
popq rA	B	0	rA	F	11	Dest	12	Dest	13	Dest

Immediates: V, D, Dest

byte:	0	1	2	3	4	5	6	7	8	9
halt	0	0								
nop	1	0								
rrmovq/cmovCC rA, rB	2	cc	rA	rB	V	D	D			
irmovq V, rB	3	0	F	rB						
rmmovq rA, D(rB)	4	0	rA	rB						
mrmovq D(rB), rA	5	0	rA	rB						
OPq rA, rB	6	fn	rA	rB						
jCC Dest	7	cc			Dest					
call Dest	8	0			Dest					
ret	9	0								
pushq rA	A	0	rA	F						
popq rA	B	0	rA	F						

Y86-64 encoding (1)

```
long addOne(long x) {
    return x + 1;
}
```

x86-64:

```
movq %rdi, %rax
addq $1, %rax
ret
```

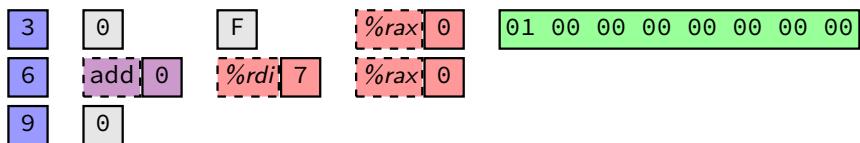
Y86-64:

```
irmovq $1, %rax
addq %rdi, %rax
ret
```

Y86-64 example (2)

addOne:

```
irmovq $1, %rax
addq %rdi, %rax
ret
```



30 F0 01 00 00 00 00 00 00 60 70 90

Y86-64 decoding

20 10 60 20 61 37 72 84 00 00 00 00 00 00 00 20
01 70 68 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

rrmovq %rcx, %rax

- 0 as cc: always
- 1 as reg: %rcx
- 0 as reg: %rax

addq %rdx, %rax

subq %rbx, %rdi

- 0 as fn: add
- 1 as fn: sub

jl 0x84

- 2 as cc: l (less than)
- hex 84 00... as little endian

byte:	0	1	2	3	4	5	6	7	8	9
halt	[0][0]									
nop	[1][0]									
rrmovq/cmovcc rA, rB	[2][cc][rA][rB]									
irmovq V, rB	[3][0][F][rB]	V								
rmmovq rA, D(rB)	[4][0][rA][rB]	D								
mrmovq D(rB), rA	[5][0][rA][rB]	D								
OPq rA, rB	[6][fn][rA][rB]									
jCC Dest	[7][cc]	Dest								
call Dest	[8][0]	Dest								
ret	[9][0]									
pushq rA	[A][0][rA][F]									
popq rA	[B][0][rA][F]									

Y86-64: Convenience for hardware

one nibble to decode instruction size/layout

mostly uniform placement of operands

jumping to zeroes (uninitialized?) by accident halts

no attempt to fit (parts of) multiple instructions in a byte

byte:	0	1	2	3	4	5	6	7	8	9
halt	[0][0]									
nop	[1][0]									
rrmovq/cmovcc rA, rB	[2][cc][rA][rB]									
irmovq V, rB	[3][0][F][rB]	V								
rmmovq rA, D(rB)	[4][0][rA][rB]	D								
mrmovq D(rB), rA	[5][0][rA][rB]	D								
OPq rA, rB	[6][fn][rA][rB]									
jCC Dest	[7][cc]	Dest								
call Dest	[8][0]	Dest								
ret	[9][0]									
pushq rA	[A][0][rA][F]									
popq rA	[B][0][rA][F]									

Y86-64

Y86-64: simplified, more RISC-y version of X86-64

minimal set of arithmetic

only **moves** touch memory

only **jumps**, **calls**, and **moves** take immediates

simple variable-length encoding

next time: implementing with circuits