An integrated concurrency and core-ISA architectural envelope definition, and test oracle, for IBM POWER multiprocessors

Kathryn E. Gray¹ Gabriel Kerneis¹⁺ Dominic Mulligan¹ Christopher Pulte¹ Susmit Sarkar² Peter Sewell¹

¹ University of Cambridge ¹⁺During work ² University of St Andrews

Typically prose

Typically prose



Typically prose



Sometimes pseudocode

Typically prose



Sometimes pseudocode

B-form

Version 2.06

ranch		I-form	Bran	Branch Conditional		B-form	
4	target_addr	(AA=0 LK=0) (AA=1 LK=0)	bc bca	BO,BI,target_a	addr	(AA=0 LK=0) (AA=1 LK=0)	
a	target_addr target_addr	(AA=0 LK=1) (AA=1 LK=1)	bcl bcla	BO,BI,target_a BO,BI,target_a	addr addr	(AA=0 LK=1) (AA=1 LK=1)	
18	6	AA LK 30 31	0 16	BO BI 6 11	BD	AA LK 30 31	

(if | K=1)

if LK then LR \leftarrow_{iea} CIA + 4 target_addr specifies the branch target address.

If AA=0 then the branch target address is the sum of LIII 0000 sign-extended and the address of this instruction, with the high-order 32 bits of the branch target address set to 0 in 32-bit mode.

If AA=1 then the branch target address is the value LI II 0b00 sign-extended, with the high-order 32 bits of the branch target address set to 0 in 32-bit mode.

If LK=1 then the effective address of the instruction following the Branch instruction is placed into the Link Register.

Special Registers Altered:

if (64-bit mode) then M ← 0 else M ← 32

 $\begin{array}{l} \underset{d \in \mathcal{D}}{\text{cise } A} \leftarrow 12a \quad \text{CR} \leftarrow \text{CRR} - 1 \\ \underset{d \in \mathcal{D}_{Q}}{\text{cise } A} \leftarrow 0 \\ \underset{d \in \mathcal{D}_{Q}}{\text{cond}_{Q}} \land \leftarrow \text{BO}_{Q} \mid ((\text{CR}_{\text{M}_{1} \times 3} \neq 0) \oplus \text{BO}_{3}) \\ \underset{d \in \mathcal{D}_{Q}}{\text{cond}_{Q}} \land \leftarrow \text{cond}_{Q} \land \text{then} \\ \underset{d \in \mathcal{D}_{Q}}{\text{if } At \text{ then NIA}} \leftarrow \underset{d \in \mathcal{D}_{Q}}{\text{cond}_{Q}} \land \text{then} \\ \underset{d \in \mathcal{D}_{Q}}{\text{cise } \text{CR}} \leftarrow \text{CR} + \text{CRS} (\text{BD} \mid \mid \text{ObOO}) \\ \underset{d \in \mathcal{D}_{Q}}{\text{cise } \text{CR}} \land \text{CR} + 4 \\ \end{array}$

if LK then LR \leftarrow_{iea} CIA + 4 BI+32 specifies the Condition Register bit to be tested

The BO field is used to resolve the branch as described in Figure 42. *target_addr* specifies the branch target address.

If AA=0 then the branch target address is the sum of BD II 0b00 sign-extended and the address of this instruction, with the high-order 32 bits of the branch target address set to 0 in 32-bit mode.

If AA=1 then the branch target address is the value BD II 0b00 sign-extended, with the high-order 32 bits of the branch target address set to 0 in 32-bit mode.

If LK=1 then the effective address of the instruction following the *Branch* instruction is placed into the Link Register.

Special Registers Altered: CTR LR

(if BO₂=0) (if LK=1)

Extended Mnemonics:

Examples of extended mnemonics for Branch Conditional:

Exten	ded:	Equiv	Equivalent to:		
blt	target	bc	12,0,target		
bne	cr2,target	bc	4,10,target		
bdnz	target	bc	16.0.target		

But ...



- Not executable test oracles
 - You can't test h/w or s/w against prose



- Not executable test oracles
 - You can't test h/w or s/w against prose
- Not a clear guide to concurrent behaviour
 - Especially for weakly consistent IBM Power and ARM



- Not executable test oracles
 - You can't test h/w or s/w against prose
- Not a clear guide to concurrent behaviour
 - Especially for weakly consistent IBM Power and ARM
- A mass of instruction set detail

We (show how to) make architecture specs that are *real* technical artefacts

• Executable as test oracle

- Executable as test oracle
- Mathematically precise

- Executable as test oracle
- Mathematically precise
- Related to vendor pseudocode and intuition

- Executable as test oracle
- Mathematically precise
- Related to vendor pseudocode and intuition
- Clarify interface between ISA and concurrency

We (show how to) make architecture specs that are *real* technical artefacts

Specifically IBM POWER all non-FP non-vector "user" ISA (153 instructions) and concurrency model

We (show how to) make architecture specs that are *real* technical artefacts

Specifically IBM POWER all non-FP non-vector "user" ISA (153 instructions) and concurrency model

Applicable to ARM as well See Modelling the ARMv8 Architecture, Operationally Concurrency and ISA, POPL16

Emulator

PPCMEM2

Emulator

PPCMEM2

Written in C etc A language with many faults Intermingling of emulation detail & semantics

Emulator

PPCMEM2

Written in C etc A language with many faults Intermingling of emulation detail & semantics

Written in Lem & Sail

Languages for logic, mathematics, and ISAs

Only spec detail Emulation separated

Emulator

PPCMEM2

Running concurrent code: Consider a lock

PPCMEM2

Emulator T1 Lock T1 Set critical section T1 Unlock T2 Lock

. . .

repeat

Running concurrent code: Consider a lock

Emulator PPCMEM2 T1 Lock T1 Set critical section T1 Unlock R3 i8:CBNZ W3, et T2 Lock repeat Running concurrent code: i61:ADD W1, W1, #1 Consider a lock

Beneficiaries

- Compiler writers
- Concurrency primitive implementors
- Security developers
- Hardware developers



Store Word with Update D-form		
stwu RS,D(RA)	<pre>union ast member (bit[5], bit[5], bit[16]) Stwu</pre>	
37RSRAD06111631EA \leftarrow (RA) + EXTS(D)MEM(EA, 4) \leftarrow (RS)MEM(EA, 4) \leftarrow (RS)32:63RA \leftarrow EA	<pre>function clause decode (0b100101 :</pre>	
Let the effective address (EA) be the sum (RA)+ D. (RS) _{32:63} are stored into the word in storage addressed by EA. EA is placed into register RA. If RA=0, the instruction form is invalid. Special Registers Altered: None	<pre>function clause execute (Stwu (RS, RA, D)) = { (bit[64]) EA := 0; EA := GPR[RA] + EXTS(D); GPR[RA] := EA; MEMw(EA,4) := (GPR[RS])[32 63] }</pre>	

Store Word with Update

D-form

stwu RS,D(RA) union 37 RS RA Duction clause 0 6 11 16 31EA \leftarrow (RA) + EXTS(D) Stwu (RS,RA,D)

 $MEM(EA, 4) \leftarrow (RS)_{32:63}$ $RA \leftarrow EA$ (bit[64]) EA

Let the effective address (EA) be the sum (RA)+ D. $(RS)_{32:63}$ are stored into the word in storage addressed by EA.

EA is placed into register RA.

If RA=0, the instruction form is invalid.

Special Registers Altered:

None

Store Word with Update)-form
stwu RS,D(RA)	<pre>union ast member (bit[5], bit[5], bit[16]) Stwu</pre>
37RSRAD061116EA \leftarrow (RA) + EXTS(D)MEM(EA, 4) \leftarrow (RS)RA \leftarrow EA	<pre>31 function clause decode (0b100101 :</pre>
Let the effective address (EA) be the sum (RS) _{32:63} are stored into the word in storage ac by EA.	<pre>(RA)+ D. Idressed function clause execute (Stwu (RS, RA, D)) = { (bit[64]) EA := 0; EA := GPR[RA] + EXTS(D);</pre>
EA is placed into register RA.	GPR[RA] := EA;
If RA=0, the instruction form is invalid.	MEMw(EA, 4) := (GPR[RS])[32 63]
Special Registers Altered: None	ſ



Sail:

for specifying concurrent ISAs

- C-like/ISA Pseudo-code like imperative language with
 - Built-in understanding of registers and memory
 - Type inference, including vector-size checking
- Formal interpreter
 - Executable for sequential or concurrent exploration
 - Analyses instructions for register/memory footprint



ISA + Concurrency Challenges

- No single program point
- No per-thread register state
- Register shadowing effects
- and more

No Per-thread register state

MP+sync+rs		POWER			
Thread	0	Thread 1			
stw r7,0(r1)	# x=1	lwz r5,0(r2)	# r5=y		
sync	# sync	mr r6,r5	# r6=r5		
stw r8,0(r2)	# y=1	lwz r5,0(r1)	# r5=x		
Initial state: 0:r1=x, 0:r2=y, 0:r7=1,					
0:r8=1, 1:r1=x, 1:r2=y, x=0					
Allowed: 1:r6=1, 1:r5=0					

a.001

ELF model

Lem

System State-

Storage subsystem state: writes seen = {W x/8=0, W y/8=0, W x/4=1} coherence $= \{W \ x/8=0 \rightarrow W \ x/4=1\}$ writes past coherence point = {W x/8=0, W y/8=0} events propagated to: Thread 0: [W y/8=0 [0-7], W x/8=0 [0-7], W x/4=1 [0-3], Sync] Thread 1: [W y/8=0 [0-7], W x/8=0 [0-7]] 2 Propagate write to thread: W x/4=1 [0-3] to Thread 1 unacknowledged Sync requests = {Sync } Thread 0 state: unacknowledged Syncs = {Sync } old instructions ioid: 4 loc: 0x00000000000000000 addi RT=1 RA=0 SI=1 ioid: 5 mem writes: [W x/4=1] new instructions ioid: 6 loc: 0x00000000000008 sync L=0 ioid: 7 loc: 0x0000000000000c addi RT=3 RA=0 SI=1 ioid: 8 loc: 0x000000000000000 stw RS=3 RA=4 D=0 instruction kind: write opcode: regs_in: {GPR4, GPR3[32..63]} regs_out: {} ioids_feeding_address: {} nias: {succ} mem writes read from: [] committed mem writes: [] committed barriers: [] reg_reads: [] reg writes: [] committed: false finished: false micro op state: MOS plain b := (GPR[4]); EA := (b + EXTS (64,D)); MEMw(EA,4) := ((GPR[to_num (RS)])[32 .. 63]) Env: D=0b0...0 [0..15], EA=0b0...0 [0..63], RA=0b00100 [0..4], RS=0b00011 [0..4], b=0b0...0 [0..63] 0 (0:8) Register read: GPR4 = y from initialstate of 0-63:yThread 1 state: unacknowledged Syncs = {} 1 (1:2) Fetch from address 0x0000000000000000 lwz RT=1 RA=2 D=0

Enabled transitions:

- <u>0</u>: (0:8) Register read: GPR4 = y from initialstate of 0-63:y
- 1: (1:2) Fetch from address 0x00000000000051000 lwz RT=1 RA=2 D=0
- 2: Propagate write to thread: W x/4=1 [0-3] to Thread 1

Validation

Sequential single instruction

6983 tests of fixed-point user-mode instruction

Concurrent litmus tests

2175 tests run exhaustively including those from prior concurrency models

Conclusions

Combined ISA and concurrency model for IBM POWER

Developed w.r.t existing h/w & in consultation with architects

Usable as reference model for future h/w & s/w

Usable for verification

Relevant for ARM and future models

http://www.cl.cam.ac.uk/~pes20/ppcmem2/