

Virtualisation Case Study: Itanium and vNUMA

Slide 1

Matthew Chapman
University of NSW, National ICT Australia, HP Labs
matthewc@cse.unsw.edu.au



ITANIUM

- High-performance processor architecture
- Also known as IA-64
- Joint venture between HP and Intel
- Not the same instruction set as AMD64/EM64T
- Explicitly Parallel Instruction-Set Computing
- Very good floating-point performance


Slide 2

2000:	Itanium (Merced)	0.18 μm , 733/800Mhz
2002:	Itanium II (McKinley)	0.18 μm , 900Mhz/1Ghz
2003:	Itanium II (Madison)	0.13 μm , 1.3-1.6Ghz
2005:	Montecito	0.09 μm , 1.8-2Ghz

EXPLICITLY PARALLEL INSTRUCTION-SET COMPUTING (EPIC)

- Goal to increase *instructions per cycle*
 - Itanium can have similar performance to x86 at a lower clock speed
- Based on *Very Long Instruction Word (VLIW)*
- Explicit parallelism in instruction set
- Simplified instruction decode and issue
- Scheduling decisions made by compiler

Slide 3



EXAMPLE

Load and add three numbers in assembly code:

Slide 4

```
ld8  r4 = (r1)
ld8  r5 = (r2)
ld8  r6 = (r3)
;;
add  r7 = r4, r5
;;
add  r8 = r6, r7
;;
```

Slide 5

Resulting instructions:

MMI	ld8	r4 = (r1)
	ld8	r5 = (r2)
	nop.i	0
M;MI;	ld8	r6 = (r3)
	::	
	add	r7 = r4, r5 // Arithmetic on M too
	nop.i	0
	::	
M;MI	add	r8 = r6, r7
	::	
	nop.m	0
	nop.i	0

Slide 6

A better way:


MMI;	ld8	r4 = (r1)
	ld8	r5 = (r2)
	nop.i	0
	::	
MI;l;	ld8	r6 = (r3)
	add	r7 = r4, r5
	::	
	add	r8 = r6, r7
	::	

Slide 7

ITANIUM REGISTER SET

Large general register file

→ partially managed by the *register stack engine* (RSE)

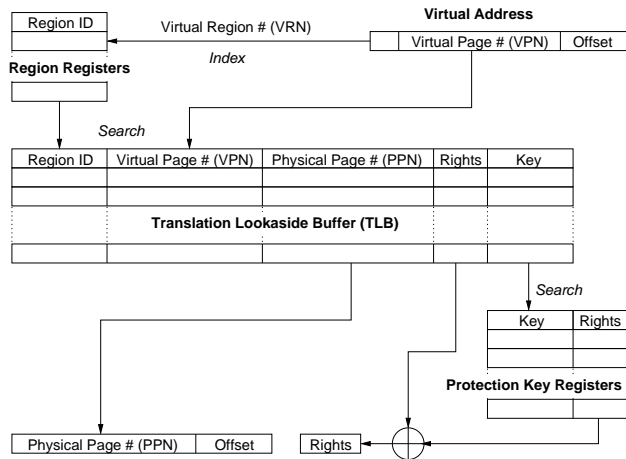


Slide 8

ITANIUM MMU

- 64-bit addressing
- Process address spaces are constructed from 8x61-bit *regions*
- Regions are 61-bit segments of a global 85-bit address space
- Page sizes from 4kB to 4GB
- Protection keys
- Choice of two hardware-walked page table formats

Slide 9



Slide 10

VIRTUALISING THE ITANIUM

- We demote the OS from privilege level 0 (most privileged) to 1
- Which instructions behave differently, and do not trap?
 - *Sensitive instructions*

Slide 11

ITANIUM INSTRUCTION SET

Fixed point arithmetic

add	mix	pavg	pshl	shrp
addp4	mov	pavgsub	pshladd	sub
and	mov	pcmp	pshr	sxt
andcm	mov pr	pmax	pshradd	tbit
cmp	movl	pmin	psub	tnat
cmp4	mux	pmpy	shl	unpack
czx	or	pmpyshr	shladd	xchg
dep	pack	popcnt	shladdp4	xor
extr	padd	psad	shr	zxt

Slide 12

Floating point arithmetic

fabs	fcvt	fnma	fpms	fsetc
fadd	fma	fnmpy	fpneg	fsub
famax	fmax	fnorm	fpnegabs	fswap
famin	fmerge	for	fpnma	fsxt
fand	fmin	fpabs	fpnmpy	fxor
fandcm	fmix	fpack	fprcpa	getf
fchkf	fmpy	fpamax	fprsqarta	mov fr
fclass	fms	fpmerge	frcpa	self
fclrf	fneg	fpmin	frsqarta	xma
fcmp	fnegabs	fpmpy	fselect	xmpy

Slide 13

Branch

br brl brp chk mov br

Register Stack Engine

alloc flushrs mov ar.bspstore
clrrrb invala **mov ar.rsc**^b
cover^a loadrs

^aside-effect at privilege level 0 when psr.ic=0
^bar.rsc contains real privilege level

Slide 14

Memory access

cmpxchg fwb ldfp probe sync
fc^a ld lfetch st
fetchadd ldf mf stf

Other unprivileged

break mov ar mov ip nop srlz
epc **mov cpuid** mov um rum
hint ^b sum

^abypasses protection check at privilege level 0
^breturns cpuid of real processor

Slide 15

Memory management (privileged)

itc mov pkr ptc tak tpa
itr mov rr ptr **thash**^a **ttag**^a

Other privileged

bsw mov dbr mov pmc mov psr rsm
mov cr mov ibr **mov pmd**^b rfi ssm

^anot privileged
^bunprivileged reads return 0 instead of trapping

Slide 16

DEALING WITH NON-VIRTUALISABLE INSTRUCTIONS

- Paravirtualisation
 - Modify guest code
- Static translation (at compile-time)
- Dynamic translation (at runtime)

OTHER ISSUES

- Fixed set of privilege levels (0..3)
 - Guest pl 0..3 now mapped onto 1..3
- Exception handlers live in virtual memory
 - Need to co-ordinate address with guest kernel
- Separate control over instruction/data/register stack translation
 - e.g. data physical, register stack virtual
 - Difficult to replicate in virtual mode
- RSE not completely virtualisable
 - Partially loaded frames cannot exist outside p10

Slide 17

VANDERPOOL TECHNOLOGY FOR ITANIUM (VT-i)

Virtualisation enhancements for Itanium

- Adds new processor operating mode for virtualisation
- Guest OS runs in p10, but with restricted permissions
- Privileged and sensitive instructions trap in this mode
- Some address space reserved for VMM
- Certain issues still problematic (physical mode, RSE)


Slide 18

vNUMA

- Itanium virtual machine monitor developed at UNSW
- Native/standalone/"Type I" VMM
- Good performance
- Supports previrtualised guests for even better performance
- Distributed!

Slide 19

vNUMA DISTRIBUTION



Slide 20

- Run vNUMA on multiple nodes of a cluster
- vNUMA locates and manages CPU/memory/I/O resources
- Presents illusion of a single large NUMA machine to guest OS
- Inter-processor communication sent over Ethernet
- "Physical memory" of virtual machine is distributed using *distributed shared memory* techniques

Slide 21

EXAMPLE

M(p): manager for page p

owner=A



copyset={}

Slide 22

EXAMPLE

M(p): manager for page p

owner=B



fetch

data



+ copyset

Slide 23

Why?

- Single machine easy to use and administer
- Allows distribution of legacy applications
- Cool application of virtualisation!