

1. Computer Architecture: Computational Features

- CISC/CISC/VLIW
 - x86 RISC in CISC clothes
 - IA64 VLIW (EPIC)
 - Everything else RISC or RISC/VLIW hybrid
- Number of assembly operands:
 - x86 2 operand
 - everything else 3 operand
- Resource limitations
- 32/64 bit native
 - IA64 64 bit, not compatible with x86
 - AMD64 = x86-64 = EM64T = EMT64, is comp with x86
 - * EM64T (Extended Memory 64 Tech) is direct copy of AMD64
- Superscalar PEs
 - PEs include:
 - * Integer units (ALU, shift, etc)
 - * prefetch units
 - * fetch (fp & i) units
 - * various FPUs
 - Use PEs in parallel
 - Make explicit in code with *Superscalar scheduling*
 - Mostly concerned with FPU
- *Hyperthreading*: increases SS exploitation by allowing mult threads to utilize the same pipelines

2. FPU Basics

- Number and type of FPUs
- Real FPUs are pipelined:
 - Latency
 - Repeat rate
- Instruction type(s) supported by FPU:
 - muladd/MACC
 - multiply
 - add
 - multiply or add
- To achieve peak, must use all units:
 - use += for muladd
 - Use pipelined multiply and add for mult and add
 - Respect repeat rate for combined unit
- FDIV/FSQRT often tie up one unit without pipelining, some archs have ded unit

3. FPU Arch Survey

ARCH	Regs	FP Peak	Inst	Lat	rep	# fpu
USIIe	32	2	mul+add	3/3	1/1	2
USIII	32	2	mul+add	4/4	1/1	2
Athlon	8	2	mul+add	4/4	1/1	2
Opteron	8/16	2/2;2/4	mul+add;vmul+vadd	4/4;5/5	1;1	2;2
PIII	8	1/1;0/4	add/mul;vadd+vmul	3/3;4/5	1/2;2/2	1;2
P4	8	1/1;2/4	add/mul;vadd+vmul	5/7;4/6	1/2;2/2	1;2
P4E	8/16	1/1;2/4	add/mul;vadd+vmul	6/8;5/7	1/2;2/2	1;2
Itan2	128	4	MACC	4	1	2
PPC/G4	32	1.5;0/8	MACC;VMACC	3;5	1	1;1
PPC/G5	32	4/4;0/8	MACC;VMACC	6;5	1/1	2;1

- If memory free, what is peak speed for each system:
 - Vec sum, vec scale, dot prod, daxpy?

4. Other PEs of Interest

fetch/store units:

- How many loads can be fetched per cycle?
- How many stores can be retired?
- FPU can be underutilized if fetch/store units inadequate
- How many load/stores needed:
 - vector sum, vec scale, dot prod, daxpy?

Prefetch Units:

- Watch memory access pattern, and attempt to fetch ahead of computation
- What is PFU's ability to detect patterns?
- How many streams can be fetched concurrently?