New PC Architectures - Processors, Chipsets, Performance, Bandwidth

1. PC - Schematic overview
2. Chipset schema (Intel 860 example)
3. AMD Athlon, XEON-P4 processor architecture
4. Processor performance
   - SSE(2) instructions
   - Prefetch example
5. Bandwidth considerations
6. Network Interfaces
7. Benchmarks

PC - schematic

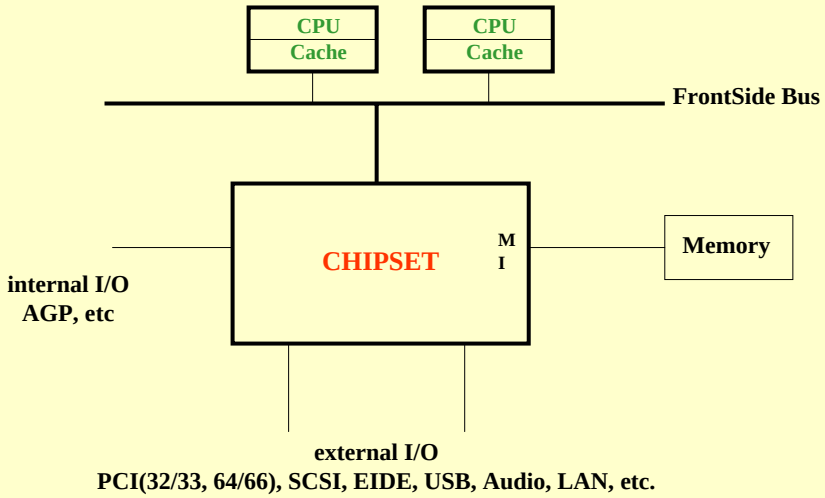

Intel 860 Chipset

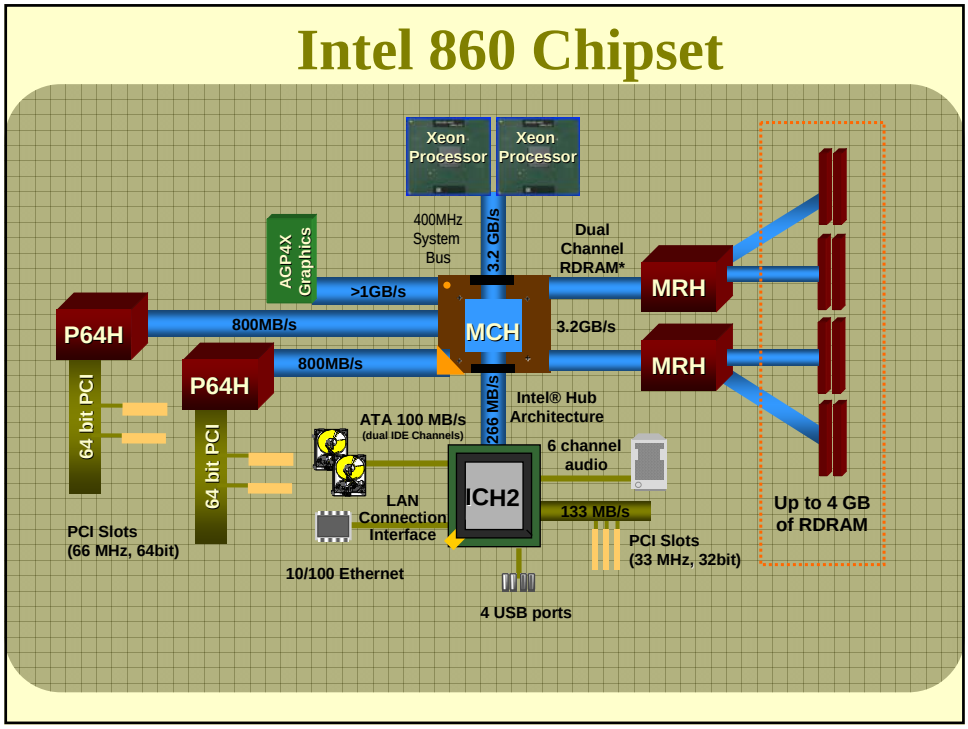

AMD Athlon

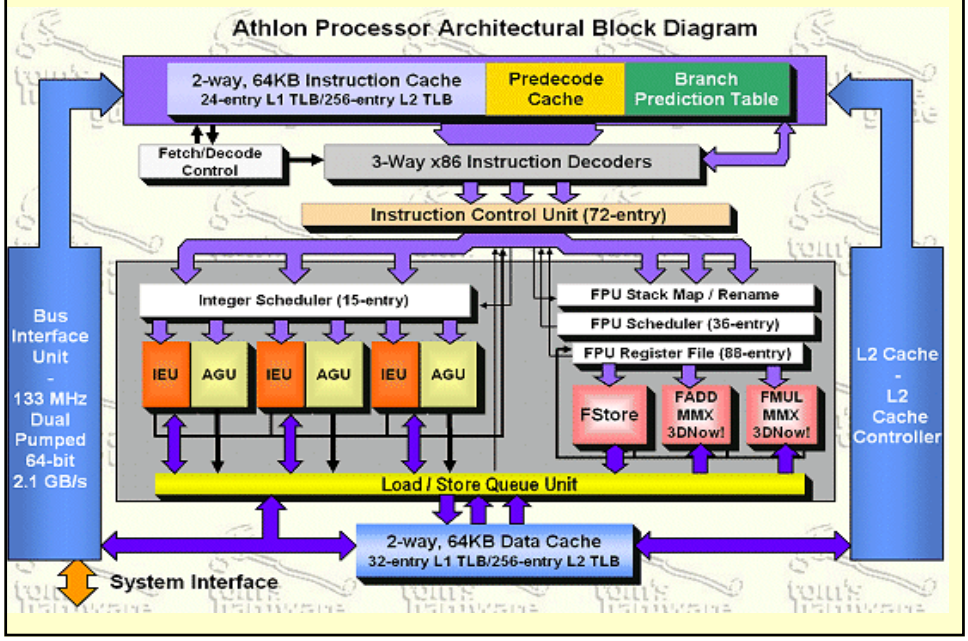

Intel P4

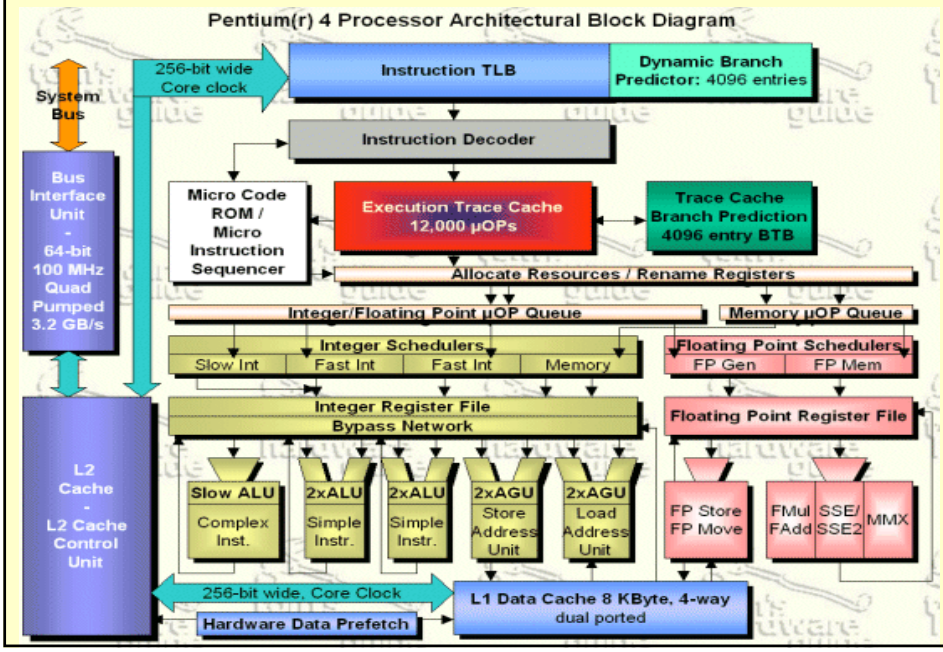

Pentium 4 CPU Core

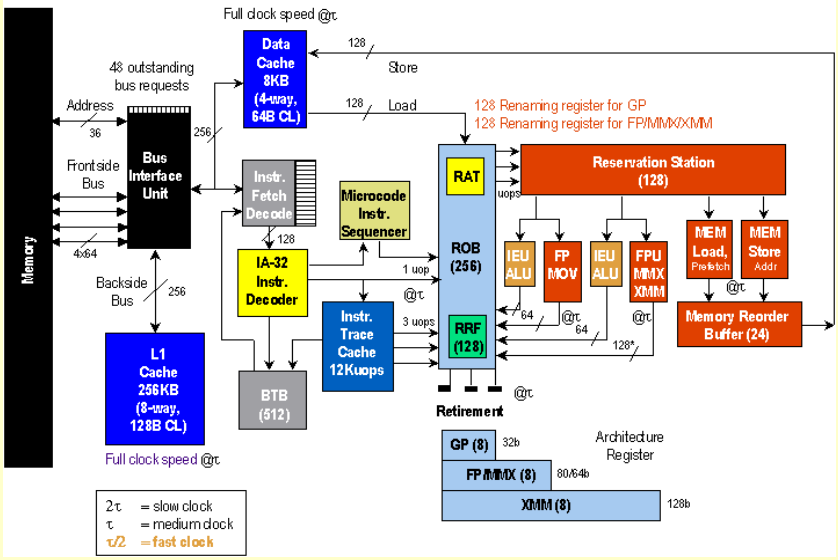

PC Vector instructions

Streaming SIMD extension 1, 2 (SSE1, SSE2), Intel XEON (P4)

144 new instructions, a 4x32-Bit (SSE1) or 2x64-bit (SSE2) SIMD integer arithmetic and 4x32-bit (SSE1) single precision or 2x64-bit (SSE2) double precision SIMD floating point instructions.

3DNow!Professional, AMD Athlon 4

71 new instructions, including SSE1 compatible floating point instructions

PowerPC Altivect

162 SIMD instructions for max. 32-bit floating point arithmetic

Applications : 3D Games (!), Compression (MPEG, JPEG...), Signal processing, ...

SSE/SSE2 registers

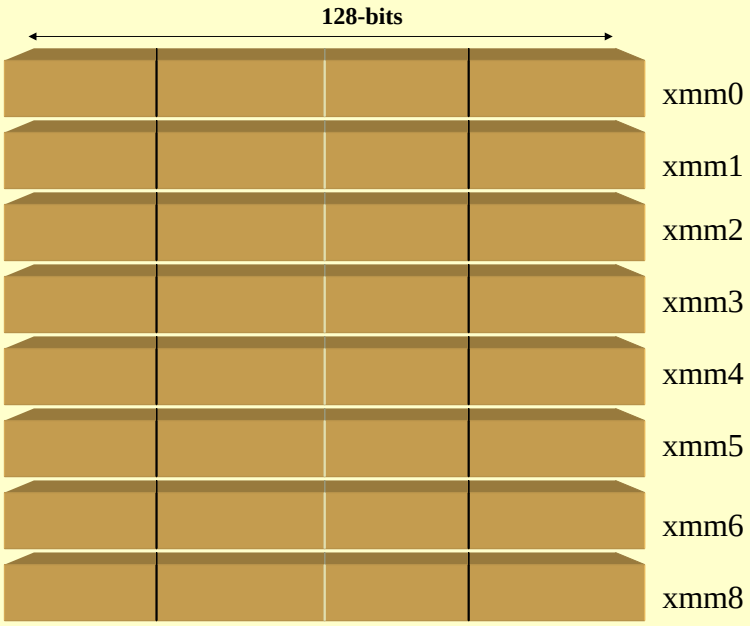

SSE/SSE2 instructions

SSE instruction example: Packed Add, 64 bit

addpd xmm1,xmm2

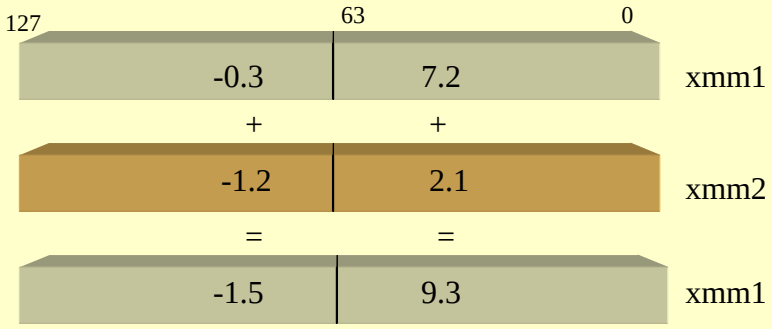

SSE/SSE2 instructions

SSE instruction example: Shuffle, 32 bit

shuffps xmm1, xmm2, 9Ch (10 01 11 00 b)

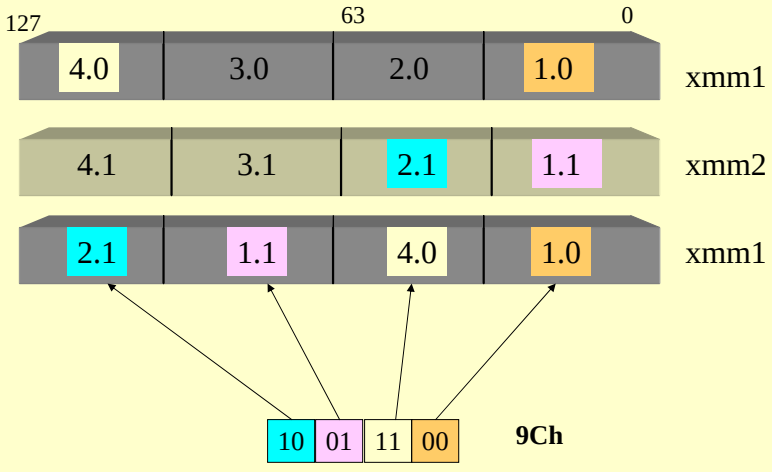

SSE2 prefetch example (M. Lüscher)

Using inlining assembler in GCC:

```
#define _prefetch_su3(addr) \
__asm__ __volatile__ ("prefetcht0 %0 \n\t" \
                      "prefetcht0 %1" \
                      : \
                      : \
                      "m" (*(((char*)((unsigned int)(addr))))), \
                      "m" (*(((char*)((unsigned int)(addr)))+128)))
...
su3 *um;
....
um=&gauge_field[iy][0];
_prefetch_su3(um);
```

SSE(2) support

Kernel	**2.4x or patches for lower versions**
Compiler	**GNU gcc 2.95.2 (no SSE optimization)** **GNU binutils 2.11.x**
	Portland Group compiler, -Mvect=sse
Libraries	**Intel Math Kernel Lib (MKL), JPEG Library, etc.** http://developer.intel.com/software/products/
Applications	**GNU inlining assembler (M. Lüscher)** **GNU C callable NASM assembler (MILC collaboration)** **C callable libraries (e.g. from Intel)**

Double Data Rate (DDR) SDRAM and Rambus

DDR SDRAM (free spec): **Utilizing both rising and falling edges of the clock**
8 Bytes * 2 * clock rate

Clock rate	66	100	133	167	200
Bandwidth GB/s	1.0	1.6	2.1	2.7	3.2

RAMBUS (license): **Direct RDRAM, 16-bit data path, 8-bit control bus, also utilizing both rising and falling edges of the clock,**
800 MHz * 2 Byte = 1.6 GB/s

Dual channel : 3.2 GB/s (XEON, P4)

Todays (02.Oct.2001) Prices for 1 GByte DDR SDRAM: ca. 200,- Euro,
1 GByte PC800 Rambus memory: ca. 530,- Euro
(both 4 x 256 MB modules)

Rambus, DDR SDRAM roadmaps

Rambus:

	PC800	PC1066	PC1200
16-bit	RIMM 1600	RIMM 2100	RIMM 2400
32-bit	RIMM 3200	RIMM 4200	RIMM 4800
64-bit	RIMM 6400	RIMM 8400	RIMM 9600
	now	2002	2005

DDR SDRAM (2002):
400 MHz * 8Byte * 2 ⟶ 6.4 GB/s ??
Double Data Rate-2 SDRAM chip (DDR-2 RAM) ⟶ 4.8 GB/s ??

Bandwidth considerations (see Jef Poskanzer http://www.acme.com/build_a_pc/bandwidth.html)

Interface	Width	Frequency	Bytes/Sec	Bits/Sec
Pentium 4 FSB	64bits	100MHz QDR	3.2 GB/s	25.6 Gbps
2-channel PC800 RDRAM	2x16bits	400MHz DDR	3.2 GB/s	25.6 Gbps
PC2100 SDRAM	64bits	133MHz DDR	2.1 GB/s	17 Gbps
EV6 bus (Athlon/Duron FSB)	64bits	100MHz DDR	1.6 GB/s	12.8 Gbps
PC1600 SDRAM	64bits	100MHz DDR	1.6 GB/s	12.8 Gbps
PC800 RDRAM	16bits	400MHz DDR	1.6 GB/s	12.8 Gbps
PC150 SDRAM	64bits	150MHz	1.3 GB/s	10.2 Gbps
133MHz FSB	64bits	133MHz	1.06 GB/s	8.5 Gbps
AGP 4x	32bits	266MHz	1.06 GB/s	8.5 Gbps
100MHz FSB	64bits	100MHz	800 MB/s	6.4 Gbps
PC100 SDRAM	64bits	100MHz	800 MB/s	6.4 Gbps

Memory **Bus** **Spec. I/O**

Bandwidth considerations (cont.)

Interface	Width	Frequency	Bytes/Sec	Bits/Sec
PC66 SDRAM	64bits	66MHz	533 MB/s	4.3 Gbps
fast/wide PCI	64bits	66MHz	533 MB/s	4.3 Gbps
AGP 2x	32bits	133MHz	533 MB/s	4.3 Gbps
Ultra-320 SCSI	16bits	160MHz	320 MB/s	2.6 Gbps
AGP	32bits	66MHz	266 MB/s	2.1 Gbps
Ultra-160 SCSI	16bits	80MHz	160 MB/s	1.3 Gbps
PCI	32bits	33MHz	133 MB/s	1.06 Gbps
ATA/133 disk	8bits	66MHz DDR	133 MB/s	1.06 Gbps
gigabit ethernet	serial	1GHz	125 MB/s	1 Gbps

Memory **Bus** **Spec. I/O** **Network** **Disk**

Bandwidth considerations (cont.)

ATA/100 disk	8bits	50MHz DDR	100 MB/s	800 Mbps
Ultra-2 Wide SCSI	16bits	40MHz	80 MB/s	640 Mbps
OC-12 network	serial	622 MHz	77.7 MB/s	622.08 Mbps
ATA/66 disk	8bits	33MHz DDR	66 MB/s	533 Mbps
USB-2	serial	480 MHz	60 MB/s	480 Mbps
IEEE 1394	serial	400MHz	50 MB/s	400 Mbps
Ultra Wide SCSI	16bits	20MHz	40 MB/s	320 Mbps
ATA/33 disk	8bits	16.6MHz DDR	33 MB/s	266 Mbps
Fast Wide SCSI	16bit	10MHz	20 MB/s	160 Mbps
OC-3 network	serial	155 MHz	19.4 MB/s	155.52 Mbps
100baseT ethernet	serial	100MHz	12.5 MB/s	100 Mbps
T-3 network	serial	45MHz	5.6 MB/s	44.736 Mbps
USB	serial	12MHz	1.5 MB/s	12 Mbps
10baseT ethernet	serial	10MHz	1.25 MB/s	10 Mbps
T-1 network	serial	1.5MHz	193 KB/s	1.544 Mbps

Network Interfaces

	Max.	Sustained
10 GBASE, 10 Gbit Ethernet	1.25 GB/s	?
GSN (&Hippi)	800 MB/s	?
Myrinet	500 MB/s	200...250 MB/s
SCI	500 MB/s (Ring)	ca. 200 MB/s
Gbit Ethernet	125 MB/s	30...80 MB/s
Fast Ethernet	12.5 MB/s	8...10 MB/s

Benchmarks

(Lies, Damned Lies, Benchmarks, Roger Shepherd, Peter Thompson, Inmos Techn. Note 27, Jan. 1988)

Peak performance (never achieved !)

Number of execution units * clock rate

e.g.
MFLOPS : Million Floating point Operations Per Second
MIPS : Million Integer instructions Per Seconds

CERN units : 1 CERN unit = 40 MIPS

Examples of standard (synthetic) Benchmarks

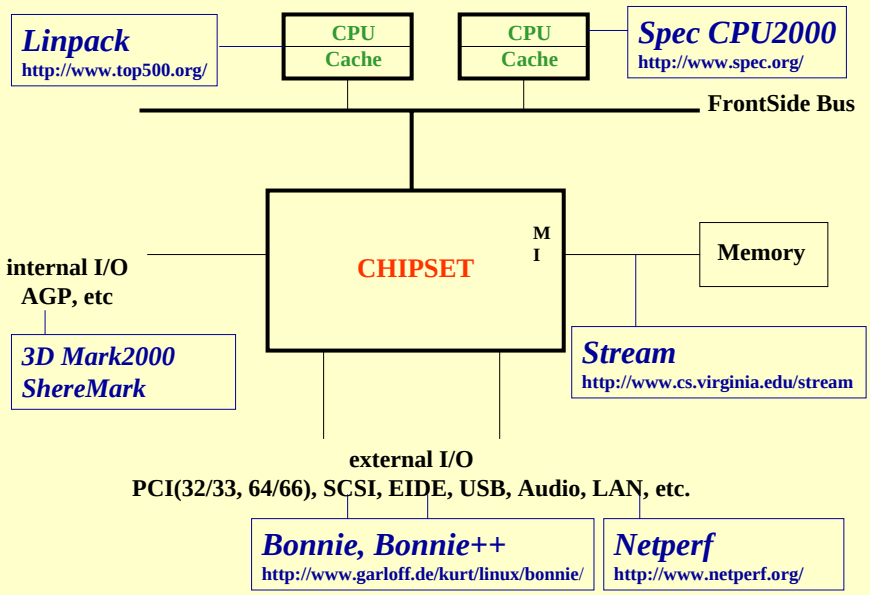

Stream Benchmark for Rambus memory on P4

Streams benchmark, gcc (Don Holmgren, Fermilab, u.a.)

gcc

Function	Rate (MB/s)	RMS time	Min time	Max time
Copy:	1324.0370	0.0492	0.0483	0.0556
a(i) = b(i)				
Scale:	1336.4782	0.0487	0.0479	0.0552
a(i) = q*b(i)				
Add:	1556.6983	0.0621	0.0617	0.0623
a(i) = b(i) + c(i)				
Triad:	1541.3021	0.0627	0.0623	0.0628
a(i) = b(i) + q*c(i)				

Portland Compiler Group build (-Mvect=sse)

Function	Rate (MB/s)	RMS time	Min time	Max time
Copy:	2072.0057	0.0309	0.0309	0.0311
Scale:	1395.3079	0.0463	0.0459	0.0464
Add:	1907.2235	0.0505	0.0503	0.0509
Triad:	1889.2441	0.0509	0.0508	0.0513

(Most of this boost comes from pgcc's use of the SSE prefetch instructions; some benefit comes from moving data via the 128-bit wide SSE registers.)

Application Benchmarks, Stream benchmark

32/64-bit Dirac Kernel, LQCD (Martin Lüscher, CERN, DESY):

P4, 1.4 GHz, 256 MB Rambus

Time per lattice point:

0.926 micro sec (1503 Mflops [32 bit arithmetic])

1.709 micro sec (814 Mflops [64 bit arithmetic])

Amanda Reco:

P4 (1.4 GHz) vs. PIII(800 MHz) 80 ... 100 % improvement

Stream Benchmark:

P4(1.4 GHz, PC800 Rambus) : 1.6 GB/s

PIII (800MHz, PC133 SDRAM) : 400 MB/s

PIII(400 MHz, PC133 SDRAM) : 340 MB/s

Application benchmarks - MILC benchmark (LQCD), non optimized