

CERN openlab II

Performance Monitoring Results

Andreas Hirstius
Andreas.Hirstius at cern.ch



- The perfmon interface to the Performance Monitoring Units (PMUs) in the processor was used
 - Very detailed information available
 - Very little to no overhead for non-latency bound applications
 - More info:
 - Andrzej presentation from March 2008 (available on openlab web page)
 - <http://perfmon2.sourceforge.net>



CERN
openlab

Performance Monitoring – perfmon/pfmon II

- Pfmon
 - Allows for multiplexing – measuring a large number of performance counter events by switching between different sets – a large number of events can be measured at the same time if the application runs long enough
 - Improvements submitted by Andrzej Nowak include
 - Proper correlation of symbols across `dlopen()/dlclose()` calls
 - Rewrite of symbol table parsing and management
 - And others ...
 - Many more improvements ...
 - Pfmon is now in a state that it can be used on HEP/CERN software ... we have a huge number of shared libs ...

- In connection with the HEPiX benchmark WG
- Running on 3GHz Woodcrests (Xeon 5160)
- Getting low level performance information for
 - SPEC2000
 - int and fp benchmarks together
 - All sub-benchmarks separately
 - SPEC2006
 - int and fp benchmarks together
 - All sub-benchmarks separately
 - All C++ sub-benchmarks together
 - In 32bit and 64bit mode (with CERN settings)
 - Running as single job and with 4 parallel jobs (# of cores)

- Monitoring the jobs running on lxbatch
 - Pfmon running in system-wide mode
 - “overview” over the application-mix on batch/worker nodes
 - Kernel and additional packages have to be “production ready” !!
 - Previous attempt for low level monitoring using “perfctr” resulted in very unstable machines ...
 - The kernel with perfmon patches has to be stable
 - Running pfmon should not introduce instabilities
 - openafs had to rebuild for the new kernel
 - Up to now no instabilities :-)



- Comparing the SPEC performance with and without pfmon

Test	SPECInt2006	SPECfp2006
std. SLC4.6 w/ std. Kernel	34.26	25.6
std. SLC4.6 w/ 2.6.24.2 Kernel & perfmon patches & pfmon <i>inactive</i>	38.16 (111%)	27.77 (108%)
std. SLC4.6 w/ 2.6.24.2 Kernel & perfmon patches & pfmon <i>active</i>	38.23 (112%)	27.78 (108%)

- The performance “impact” of pfmon is
 - ➔ The performance actually **increases !!!**
 - ... mainly due to the newer kernel ...
- Measurements by Alejandro Iribarren

- CPI (Cycles per instruction)
- Load / Store instructions
- x87 instructions
- computational SIMD instructions
- Branch instructions
- Mispredicted branches
- Level2 cache misses (rate)
- Bus utilisation
- Data bus utilisation
- “Bus Not Ready”
- Resource stalls
- Detailed look into the resource stalls:
 - ROB full (Re-Order Buffer full)
 - RS full (Reservation Station full)
 - LD / ST limits
 - Mispredicted branches
 - FP ctrl word

- x87 and computational SIMD instructions
 - They are the “floating point” (FP) instructions
 - x87 instr. are default in 32bit mode
 - SIMD instr. are default in 64bit mode
- Mispredicted branches
 - Percentage of branches that were mispredicted
 - They carry a performance penalty!
- Level2 cache misses (rate)
 - Percentage of requests to the L2 cache that have to go to main memory
 - They carry a significant performance penalty!!!

What do we measure – some details II

- **Bus utilisation**
 - Percentage of bus cycles used to transfer bus transactions of any kind
- **Data bus utilisation**
 - Percentage of bus cycles used to transfer data
- **Bus Not Ready**
 - Percentage of bus cycles during which no new bus transactions can start.
 - A high value indicates a highly loaded and frequently overloaded bus

- SPECint2000 contains ~1% floating point!!
- Only slightly faster in 64bit mode
 - Due to a few sub-benchmarks being extremely slowed down in 64bit mode!!
 - Significantly more L2 cache misses in 64bit mode
 - Slightly more mispredicted branches
 - BNR much higher in 64bit mode
 - Slightly more resource stalls in 64bit mode
 - ... but still significantly fewer actual instructions used !!

SPECint2000 – 32bit mode – 4 jobs		SPECint2000 – 64bit mode – 4 jobs	
45201237332100	UNHALTED_CORE_CYCLES	44804927344019	UNHALTED_CORE_CYCLES
45837316931397	INSTRUCTIONS_RETIRED	37967583644997	INSTRUCTIONS_RETIRED
7497308624143	BRANCH_INSTRUCTIONS_RETIRED	6421199461762	BRANCH_INSTRUCTIONS_RETIRED
302799323740	MISPREDICTED_BRANCH_RETIRED	288149260964	MISPREDICTED_BRANCH_RETIRED
19296716652582	INST_RETIRED:LOADS	12159284443871	INST_RETIRED:LOADS
624928	SIMD_COMP_INST_RETIRED:PACKED_SINGLE	413103994039	SIMD_COMP_INST_RETIRED:PACKED_SINGLE
	:SCALAR_SINGLE:PACKED_DOUBLE:SCALAR_DOUBLE		:SCALAR_SINGLE:PACKED_DOUBLE:SCALAR_DOUBLE
9323887361488	INST_RETIRED:STORES	5366523996464	INST_RETIRED:STORES
442567549137	X87_OPS_RETIRED:ANY	57678491535	X87_OPS_RETIRED:ANY
22143424590558	RESOURCE_STALLS:ANY	23933277888657	RESOURCE_STALLS:ANY
378414578559	BUS_TRANS_ANY:ALL_AGENTS	724554943704	BUS_TRANS_ANY:ALL_AGENTS
308895207467	BUS_DRDY_CLOCKS:ALL_AGENTS	574793437730	BUS_DRDY_CLOCKS:ALL_AGENTS
1741997649	BUS_BNR_DRV:ALL_AGENTS	7859133902	BUS_BNR_DRV:ALL_AGENTS
543636967874	LAST_LEVEL_CACHE_REFERENCES	618095730129	LAST_LEVEL_CACHE_REFERENCES
27980144496	LAST_LEVEL_CACHE_MISSES	54049684697	LAST_LEVEL_CACHE_MISSES
5035359994977	CPU_CLK_UNHALTED:BUS	4992386647432	CPU_CLK_UNHALTED:BUS
-----		-----	
Ratios:		Ratios:	
	CPI: 0.9861		CPI: 1.1801
	load instructions %: 42.098%		load instructions %: 32.025%
	store instructions %: 20.341%		store instructions %: 14.134%
	load and store instructions %: 62.440%		load and store instructions %: 46.160%
	resource stalls % (of cycles): 48.989%		resource stalls % (of cycles): 53.417%
	branch instructions %: 16.356%		branch instructions %: 16.912%
	% of branch instr. mispredicted: 4.039%		% of branch instr. mispredicted: 4.487%
	% of 12 loads missed: 5.147%		% of 12 loads missed: 8.745%
	bus utilization %: 15.030%		bus utilization %: 29.026%
	data bus utilization %: 6.135%		data bus utilization %: 11.513%
	bus not ready %: 0.069%		bus not ready %: 0.315%

- SPECint2006 contains ~0.1% floating point ... negligible ...
- Only slightly faster in 64bit mode
 - Due to a few sub-benchmarks being extremely slowed down in 64bit mode!!
 - Slightly more L2 cache misses in 64bit mode, but very high miss rate in general
 - BNR much higher in 64bit mode
 - Slightly more resource stalls in 64bit mode
 - ... again, significantly fewer actual instructions in 64bit mode

SPECint2006 – 32bit mode – 4 jobs		SPECint2006 – 64bit mode – 4 jobs	
461179857748960	UNHALTED_CORE_CYCLES	453153422688509	UNHALTED_CORE_CYCLES
350589714455709	INSTRUCTIONS_RETIRED	303883044666778	INSTRUCTIONS_RETIRED
54906528688928	BRANCH_INSTRUCTIONS_RETIRED	48233916086552	BRANCH_INSTRUCTIONS_RETIRED
2903104696534	MISPREDICTED_BRANCH_RETIRED	2568377286745	MISPREDICTED_BRANCH_RETIRED
162881915117934	INST_RETIRED:LOADS	96299115605813	INST_RETIRED:LOADS
0	SIMD_COMP_INST_RETIRED:PACKED_SINGLE	359749309739	SIMD_COMP_INST_RETIRED:PACKED_SINGLE
	:SCALAR_SINGLE:PACKED_DOUBLE:SCALAR_DOUBLE		:SCALAR_SINGLE:PACKED_DOUBLE:SCALAR_DOUBLE
58657311592597	INST_RETIRED:STORES	33169229301278	INST_RETIRED:STORES
399618515172	X87_OPS_RETIRED:ANY	132684652120	X87_OPS_RETIRED:ANY
265980169455834	RESOURCE_STALLS:ANY	283825996907464	RESOURCE_STALLS:ANY
9030113964523	BUS_TRANS_ANY:ALL_AGENTS	11197303950198	BUS_TRANS_ANY:ALL_AGENTS
7313830415240	BUS_DRDY_CLOCKS:ALL_AGENTS	9007337650460	BUS_DRDY_CLOCKS:ALL_AGENTS
82885334380	BUS_BNR_DRV:ALL_AGENTS	226487203141	BUS_BNR_DRV:ALL_AGENTS
3945120792043	LAST_LEVEL_CACHE_REFERENCES	4818080605870	LAST_LEVEL_CACHE_REFERENCES
478431290431	LAST_LEVEL_CACHE_MISSES	720453702717	LAST_LEVEL_CACHE_MISSES
51581290410280	CPU_CLK_UNHALTED:BUS	50797500445973	CPU_CLK_UNHALTED:BUS
-----		-----	
Ratios:		Ratios:	
	CPI: 1.3154		CPI: 1.4912
	load instructions %: 46.459%		load instructions %: 31.690%
	store instructions %: 16.731%		store instructions %: 10.915%
	load and store instructions %: 63.190%		load and store instructions %: 42.605%
	resource stalls % (of cycles): 57.674%		resource stalls % (of cycles): 62.634%
	branch instructions %: 15.661%		branch instructions %: 15.873%
% of branch instr. mispredicted:	5.287%	% of branch instr. mispredicted:	5.325%
% of 12 loads missed:	12.127%	% of 12 loads missed:	14.953%
bus utilization %:	35.013%	bus utilization %:	44.086%
data bus utilization %:	14.179%	data bus utilization %:	17.732%
bus not ready %:	0.321%	bus not ready %:	0.892%
comp. SIMD inst. ('new FP') %:	0.000%	comp. SIMD inst. ('new FP') %:	0.118%
comp. x87 instr. ('old FP') %:	0.114%	comp. x87 instr. ('old FP') %:	0.044%



Pfmon – The Results – SPEC2006 C++

CERN
openlab

- The SPEC2006 C++ benchmarks contain 10-14% floating point
- **Slower in 64bit mode !!**
 - Almost all C++ benchmarks are slowed down in 64bit mode!!
 - Slightly more L2 cache misses in 64bit mode, but very high miss rate in general
 - Slightly more mispredicted branches in 64bit mode
 - BNR higher in 64bit mode
 - ~10% more resource stalls in 64bit mode
 - ... again, fewer actual instructions in 64bit mode

SPEC2006 C++ – 32bit mode – 4 jobs		SPEC2006 C++ – 64bit mode – 4 jobs	
196604109151652	UNHALTED_CORE_CYCLES	215947712982588	UNHALTED_CORE_CYCLES
150762226596169	INSTRUCTIONS_RETIRED	119991871388997	INSTRUCTIONS_RETIRED
20597482426950	BRANCH_INSTRUCTIONS_RETIRED	18241941807980	BRANCH_INSTRUCTIONS_RETIRED
764119332313	MISPREDICTED_BRANCH_RETIRED	745321349045	MISPREDICTED_BRANCH_RETIRED
59343184864687	INST_RETIRED:LOADS	36586210736816	INST_RETIRED:LOADS
0	SIMD_COMP_INST_RETIRED:PACKED_SINGLE	16549768252572	SIMD_COMP_INST_RETIRED:PACKED_SINGLE
	:SCALAR_SINGLE:PACKED_DOUBLE:SCALAR_DOUBLE		:SCALAR_SINGLE:PACKED_DOUBLE:SCALAR_DOUBLE
22653050918291	INST_RETIRED:STORES	11817242786897	INST_RETIRED:STORES
16499979718174	X87_OPS_RETIRED:ANY	17200520909	X87_OPS_RETIRED:ANY
121652777590140	RESOURCE_STALLS:ANY	150950098548670	RESOURCE_STALLS:ANY
4058630608266	BUS_TRANS_ANY:ALL_AGENTS	5005296800109	BUS_TRANS_ANY:ALL_AGENTS
3131296782062	BUS_DRDY_CLOCKS:ALL_AGENTS	3821979997184	BUS_DRDY_CLOCKS:ALL_AGENTS
30656103004	BUS_ENR_DRV:ALL_AGENTS	47675704094	BUS_ENR_DRV:ALL_AGENTS
2216257328615	LAST_LEVEL_CACHE_REFERENCES	2767263319645	LAST_LEVEL_CACHE_REFERENCES
307891738955	LAST_LEVEL_CACHE_MISSES	410542553241	LAST_LEVEL_CACHE_MISSES
22121483842719	CPU_CLK_UNHALTED:BUS	24255546086757	CPU_CLK_UNHALTED:BUS
-----		-----	
Ratios:		Ratios:	
	CPI: 1.3041		CPI: 1.7997
	load instructions %: 39.362%		load instructions %: 30.491%
	store instructions %: 15.026%		store instructions %: 9.848%
	load and store instructions %: 54.388%		load and store instructions %: 40.339%
	resource stalls % (of cycles): 61.877%		resource stalls % (of cycles): 69.901%
	branch instructions %: 13.662%		branch instructions %: 15.203%
	% of branch instr. mispredicted: 3.710%		% of branch instr. mispredicted: 4.086%
	% of 12 loads missed: 13.892%		% of 12 loads missed: 14.836%
	bus utilization %: 36.694%		bus utilization %: 41.271%
	data bus utilization %: 14.155%		data bus utilization %: 15.757%
	bus not ready %: 0.277%		bus not ready %: 0.393%
	comp. SIMD inst. ('new FP') %: 0.000%		comp. SIMD inst. ('new FP') %: 13.792%
	comp. x87 instr. ('old FP') %: 10.944%		comp. x87 instr. ('old FP') %: 0.014%



Pfmon – SPEC2006 example I

- A BAD example
- Much slower in 64bit mode
 - Using MORE instructions in 64bit mode
 - Much higher L2 cache miss rate in 64bit mode
 - More branches in 64bit mode
 - Significantly more mispredicted branches in 64bit mode
 - BNR and bus utilisation actually lower in 64bit mode

SPEC2006 410.bwaves – 32bit mode – 4 jobs		SPEC2006 410.bwaves – 64bit mode – 4 jobs	
66184607716621	UNHALTED_CORE_CYCLES	82646441371942	UNHALTED_CORE_CYCLES
32038993376478	INSTRUCTIONS_RETIRED	43269747823376	INSTRUCTIONS_RETIRED
1045662341511	BRANCH_INSTRUCTIONS_RETIRED	3377043818102	BRANCH_INSTRUCTIONS_RETIRED
5569854319	MISPREDICTED_BRANCH_RETIRED	146650527607	MISPREDICTED_BRANCH_RETIRED
19500786788306	INST_RETIRED:LOADS	14517944486917	INST_RETIRED:LOADS
0	SIMD_COMP_INST_RETIRED:PACKED_SINGLE :SCALAR_SINGLE:PACKED_DOUBLE:SCALAR_DOUBLE	10476785264283	SIMD_COMP_INST_RETIRED:PACKED_SINGLE :SCALAR_SINGLE:PACKED_DOUBLE:SCALAR_DOUBLE
4768056184869	INST_RETIRED:STORES	2998919112401	INST_RETIRED:STORES
6752053870705	X87_OPS_RETIRED:ANY	10959281897	X87_OPS_RETIRED:ANY
55650366867224	RESOURCE_STALLS:ANY	62157492993642	RESOURCE_STALLS:ANY
2795205261519	BUS_TRANS_ANY:ALL_AGENTS	2794173784751	BUS_TRANS_ANY:ALL_AGENTS
1975145132590	BUS_DRDY_CLOCKS:ALL_AGENTS	1974651943707	BUS_DRDY_CLOCKS:ALL_AGENTS
36173999054	BUS_BNR_DRV:ALL_AGENTS	35367284875	BUS_BNR_DRV:ALL_AGENTS
580148778988	LAST_LEVEL_CACHE_REFERENCES	550307895671	LAST_LEVEL_CACHE_REFERENCES
81972164145	LAST_LEVEL_CACHE_MISSES	112587229075	LAST_LEVEL_CACHE_MISSES
7369433865245	CPU_CLK_UNHALTED:BUS	9209403050219	CPU_CLK_UNHALTED:BUS
-----		-----	
Ratios:		Ratios:	
	CPI: 2.0658		CPI: 1.9100
	load instructions %: 60.866%		load instructions %: 33.552%
	store instructions %: 14.882%		store instructions %: 6.931%
	load and store instructions %: 75.748%		load and store instructions %: 40.483%
	resource stalls % (of cycles): 84.084%		resource stalls % (of cycles): 75.209%
	branch instructions %: 3.264%		branch instructions %: 7.805%
	% of branch instr. mispredicted: 0.533%		% of branch instr. mispredicted: 4.343%
	% of 12 loads missed: 14.130%		% of 12 loads missed: 20.459%
	bus utilization %: 75.859%		bus utilization %: 60.681%
	data bus utilization %: 26.802%		data bus utilization %: 21.442%
	bus not ready %: 0.982%		bus not ready %: 0.768%
	comp. SIMD inst. ('new FP') %: 0.000%		comp. SIMD inst. ('new FP') %: 24.213%
	comp. x87 instr. ('old FP') %: 21.074%		comp. x87 instr. ('old FP') %: 0.025%



Pfmon – SPEC2006 example II

- Another BAD example
- Much slower in 64bit mode
 - Using FEWER instructions in 64bit mode
 - L2 cache miss rate only slightly higher in 64bit mode
 - Fewer mispredicted branches in 64bit mode
 - Bus utilisation much higher in 64bit mode
 - BNR and Resource stalls significantly higher in 64bit mode
 - The CPU is waiting for data!!
 - The memory bus is too congested...

SPEC2006 403.gcc – 32bit mode – 4 jobs		SPEC2006 403.gcc – 64bit mode – 4 jobs	
22932933922318	UNHALTED_CORE_CYCLES	36699844911989	UNHALTED_CORE_CYCLES
16802305051587	INSTRUCTIONS_RETIRED	14818404405994	INSTRUCTIONS_RETIRED
3369056425885	BRANCH_INSTRUCTIONS_RETIRED	3229507742517	BRANCH_INSTRUCTIONS_RETIRED
71024589679	MISPREDICTED_BRANCH_RETIRED	65679138951	MISPREDICTED_BRANCH_RETIRED
6132240496717	INST_RETIRED:LOADS	3006062347368	INST_RETIRED:LOADS
0	SIMD_COMP_INST_RETIRED:PACKED_SINGLE	81295081	SIMD_COMP_INST_RETIRED:PACKED_SINGLE
	:SCALAR_SINGLE:PACKED_DOUBLE:SCALAR_DOUBLE		:SCALAR_SINGLE:PACKED_DOUBLE:SCALAR_DOUBLE
3142263040501	INST_RETIRED:STORES	2621445452291	INST_RETIRED:STORES
98443376	X87_OPS_RETIRED:ANY	24408980456	X87_OPS_RETIRED:ANY
13668930351614	RESOURCE_STALLS:ANY	28348201404117	RESOURCE_STALLS:ANY
477561364166	BUS_TRANS_ANY:ALL_AGENTS	1225261154535	BUS_TRANS_ANY:ALL_AGENTS
345947255307	BUS_DRDY_CLOCKS:ALL_AGENTS	1003956072280	BUS_DRDY_CLOCKS:ALL_AGENTS
2736274388	BUS_BNR_DRV:ALL_AGENTS	120787996557	BUS_BNR_DRV:ALL_AGENTS
201800541169	LAST_LEVEL_CACHE_REFERENCES	250110325508	LAST_LEVEL_CACHE_REFERENCES
36370200772	LAST_LEVEL_CACHE_MISSES	49058830369	LAST_LEVEL_CACHE_MISSES
2565605460197	CPU_CLK_UNHALTED:BUS	4131091711457	CPU_CLK_UNHALTED:BUS
-----		-----	
Ratios:		Ratios:	
	CPI: 1.3649		CPI: 2.4766
	load instructions %: 36.496%		load instructions %: 20.286%
	store instructions %: 18.701%		store instructions %: 17.690%
	load and store instructions %: 55.198%		load and store instructions %: 37.976%
	resource stalls % (of cycles): 59.604%		resource stalls % (of cycles): 77.243%
	branch instructions %: 20.051%		branch instructions %: 21.794%
	% of branch instr. mispredicted: 2.108%		% of branch instr. mispredicted: 2.034%
	% of 12 loads missed: 18.023%		% of 12 loads missed: 19.615%
	bus utilization %: 37.228%		bus utilization %: 59.319%
	data bus utilization %: 13.484%		data bus utilization %: 24.302%
	bus not ready %: 0.213%		bus not ready %: 5.848%
	comp. SIMD inst. ('new FP') %: 0.000%		comp. SIMD inst. ('new FP') %: 0.001%
	comp. x87 instr. ('old FP') %: 0.001%		comp. x87 instr. ('old FP') %: 0.165%



Pfmon – SPEC2006 example III

- A really GOOD example
- Much faster in 64bit mode
 - ~46% fewer “core cycles” !!!
 - Only ~12% fewer instructions

SPEC2006 435.gromacs – 32bit mode – 4 jobs		SPEC2006 435.gromacs – 64bit mode – 4 jobs	
52545097029601	UNHALTED_CORE_CYCLES	28335415184402	UNHALTED_CORE_CYCLES
28239032175254	INSTRUCTIONS_RETIRED	24923444284898	INSTRUCTIONS_RETIRED
786718407198	BRANCH_INSTRUCTIONS_RETIRED	755991835848	BRANCH_INSTRUCTIONS_RETIRED
49833012629	MISPREDICTED_BRANCH_RETIRED	48309357560	MISPREDICTED_BRANCH_RETIRED
11318465152987	INST_RETIRED:LOADS	8722788997385	INST_RETIRED:LOADS
0	SIMD_COMP_INST_RETIRED:PACKED_SINGLE :SCALAR_SINGLE:PACKED_DOUBLE:SCALAR_DOUBLE	9695772998672	SIMD_COMP_INST_RETIRED:PACKED_SINGLE :SCALAR_SINGLE:PACKED_DOUBLE:SCALAR_DOUBLE
5018510676257	INST_RETIRED:STORES	3515929682894	INST_RETIRED:STORES
9717550976444	X87_OPS_RETIRED:ANY	169222816	X87_OPS_RETIRED:ANY
41735531957443	RESOURCE_STALLS:ANY	18613627304402	RESOURCE_STALLS:ANY
74104263298	BUS_TRANS_ANY:ALL_AGENTS	76000740916	BUS_TRANS_ANY:ALL_AGENTS
59386726849	BUS_DRDY_CLOCKS:ALL_AGENTS	61365932872	BUS_DRDY_CLOCKS:ALL_AGENTS
15596654	BUS_BNR_DRV:ALL_AGENTS	32145113	BUS_BNR_DRV:ALL_AGENTS
224259236583	LAST_LEVEL_CACHE_REFERENCES	227650501186	LAST_LEVEL_CACHE_REFERENCES
3380348570	LAST_LEVEL_CACHE_MISSES	3661361080	LAST_LEVEL_CACHE_MISSES
5842738075086	CPU_CLK_UNHALTED:BUS	3152227191304	CPU_CLK_UNHALTED:BUS
-----		-----	
Ratios:		Ratios:	
	CPI: 1.8607		CPI: 1.1369
	load instructions %: 40.081%		load instructions %: 34.998%
	store instructions %: 17.772%		store instructions %: 14.107%
	load and store instructions %: 57.852%		load and store instructions %: 49.105%
	resource stalls % (of cycles): 79.428%		resource stalls % (of cycles): 65.690%
	branch instructions %: 2.786%		branch instructions %: 3.033%
	% of branch instr. mispredicted: 6.334%		% of branch instr. mispredicted: 6.390%
	% of 12 loads missed: 1.507%		% of 12 loads missed: 1.608%
	bus utilization %: 2.537%		bus utilization %: 4.822%
	data bus utilization %: 1.016%		data bus utilization %: 1.947%
	bus not ready %: 0.001%		bus not ready %: 0.002%
	comp. SIMD inst. ('new FP') %: 0.000%		comp. SIMD inst. ('new FP') %: 38.902%
	comp. x87 instr. ('old FP') %: 34.412%		comp. x87 instr. ('old FP') %: 0.001%

- Pfmon in system-wide mode
 - Running on six batch nodes for one month
 - L2 cache misses are missing
 - All “inefficiencies” like waiting for CASTOR are included !!
 - ~10% floating point content
 - ~60% is in 32bit mode
 - ~40% is in 64bit mode

```

71226907483714181 UNHALTED_CORE_CYCLES
60678553830972217 INSTRUCTIONS_RETIRED
8566483902862367 BRANCH_INSTRUCTIONS_RETIRED
235402778733339 MISPREDICTED_BRANCH_RETIRED
20566739654008260 INST_RETIRED:LOADS
28430622301440 LAST_LEVEL_CACHE_MISSES
10318738898916677 INST_RETIRED:STORES
4022701150134275 X87_OPS_RETIRED:ANY
29813618068093779 RESOURCE_STALLS:ANY
772872906674477 BUS_TRANS_ANY:ALL_AGENTS
599509742879852 BUS_DRDY_CLOCKS:ALL_AGENTS
6766408165084 BUS_BNR_DRV:ALL_AGENTS
103029278043584 L2_LINES_IN:ANY
2485005306891005 SIMD_COMP_INST_RETIRED:PACKED_SINGLE:
SCALAR_SINGLE:PACKED_DOUBLE:SCALAR_DOUBLE
9397049467709644 CPU_CLK_UNHALTED:BUS
-----
Ratios:
  CPI (Cycles per instruction): 1.1738399
    load instructions %: 33.895%
    store instructions %: 17.006%
  load and store instructions %: 50.900%
  resource stalls (cycles) %: 41.857%
    branch instructions %: 14.118%
  % branches mispredicted: 2.748%
    bus utilization %: 16.449%
    data bus utilization %: 6.380%
    bus not ready %: 0.144%
  comp. SIMD ("new FP") instr. %: 4.095%
    x87 ("old FP") instr. %: 6.630%

```


- Code seems relatively efficient
 - CPI of 1.17
- Very few resource stalls
- Bus and data bus utilisations are very low
 - Keep in mind, this is the entire system including the “dead time”!
 - The *Bus Not Ready* percentage of 0.144% points to bus congestions under full load !!!!

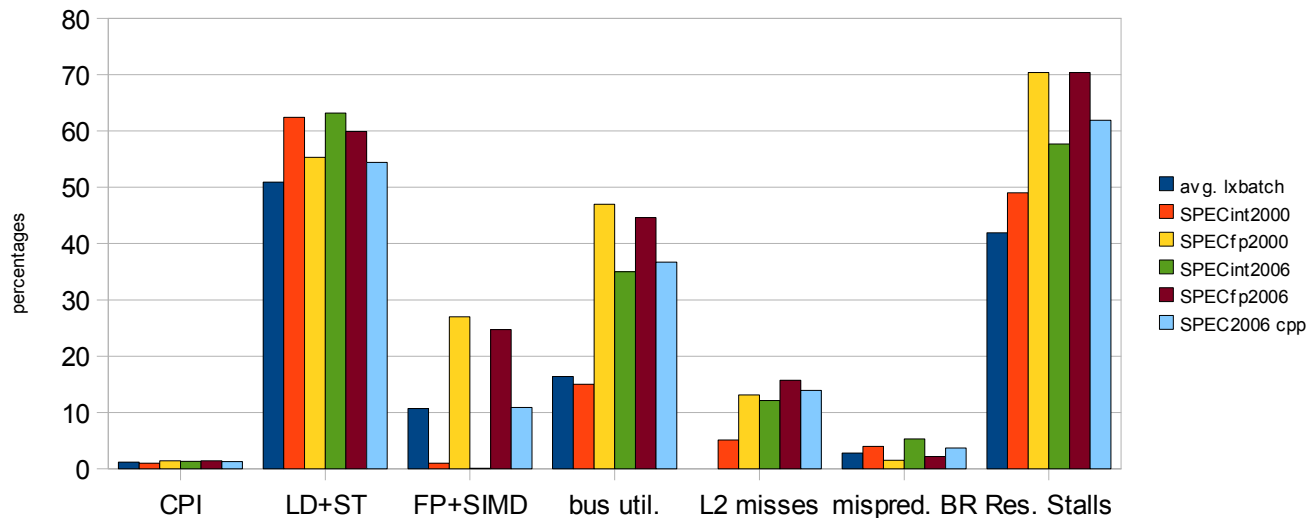
```

71226907483714181 UNHALTED_CORE_CYCLES
60678553830972217 INSTRUCTIONS_RETIRED
8566483902862367 BRANCH_INSTRUCTIONS_RETIRED
235402778733339 MISPREDICTED_BRANCH_RETIRED
20566739654008260 INST_RETIRED:LOADS
28430622301440 LAST_LEVEL_CACHE_MISSES
10318738898916677 INST_RETIRED:STORES
4022701150134275 X87_OPS_RETIRED:ANY
29813618068093779 RESOURCE_STALLS:ANY
772872906674477 BUS_TRANS_ANY:ALL_AGENTS
599509742879852 BUS_DRDY_CLOCKS:ALL_AGENTS
6766408165084 BUS_ENR_DRV:ALL_AGENTS
103029278043584 L2_LINES_IN:ANY
2485005306891005 SIMD_COMP_INST_RETIRED:PACKED_SINGLE:
                                SCALAR_SINGLE:PACKED_DOUBLE:SCALAR_DOUBLE
9397049467709644 CPU_CLK_UNHALTED:BUS
-----
Ratios:
  CPI (Cycles per instruction): 1.1738399
    load instructions %: 33.895%
    store instructions %: 17.006%
  load and store instructions %: 50.900%
    resource stalls (cycles) %: 41.857%
    branch instructions %: 14.118%
    % branches mispredicted: 2.748%
    bus utilization %: 16.449%
    data bus utilization %: 6.380%
    bus not ready %: 0.144%
  comp. SIMD ("new FP") instr. %: 4.095%
    x87 ("old FP") instr. %: 6.630%


```

Pfmon – The Results – Overview

Comparison 32bit



Comparison 64bit

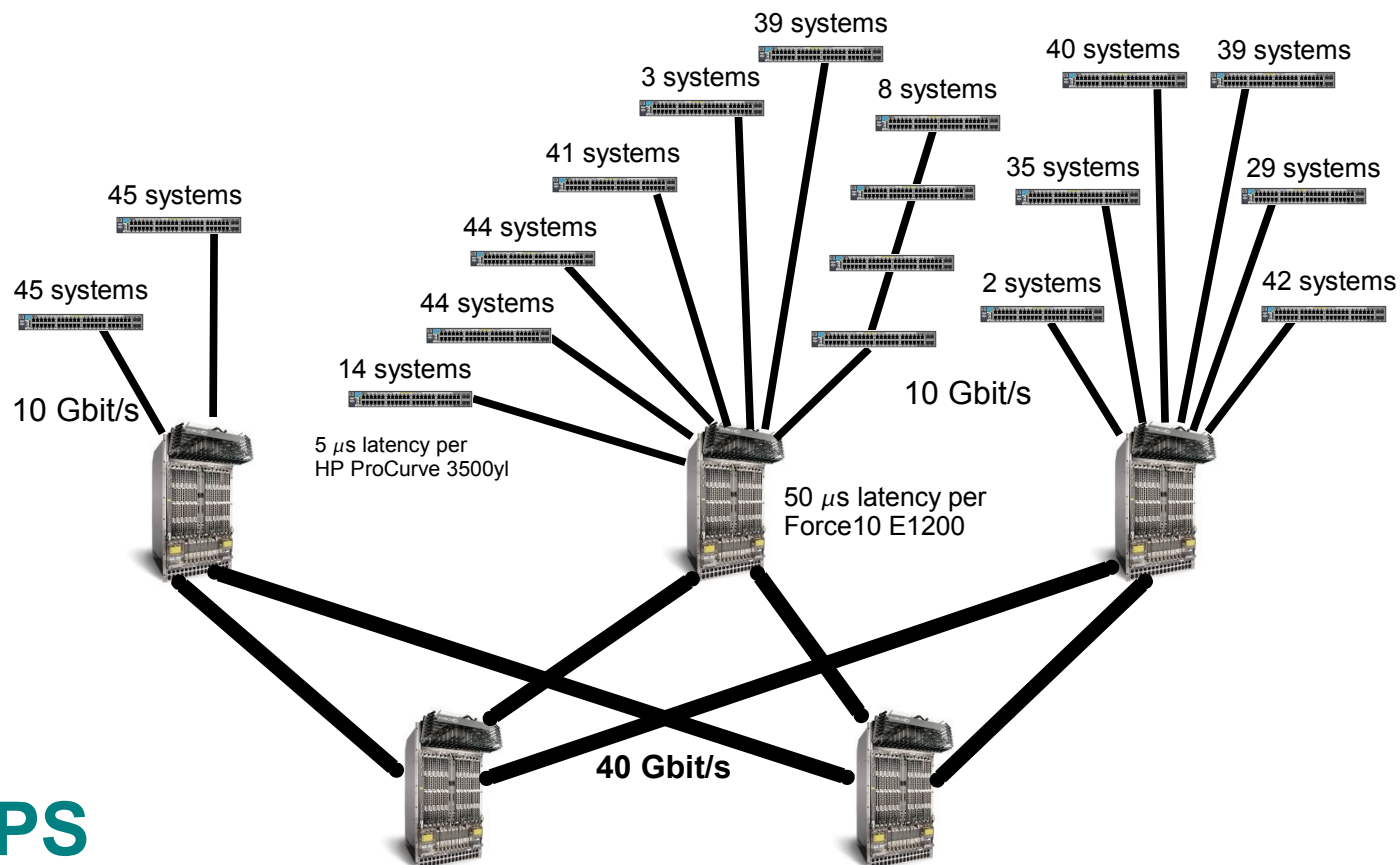


- None of the benchmarks looks like the results from Ixbatch ...
- ... but ...
 - Ixbatch results have to be “cleaned up”
 - only periods with significant activity should be taken into account
 - There are indications (like *BNR*) that the results will look a bit different.
- A new system-wide run is under way right now
 - Includes L2 Cache misses and details about resource stalls

- Batch machines running 2.6.24.x & pfmon are stable !!!!! ... and actually ~10% faster ;-)
- ~250 new machines will go into production soon
 - 50 – 100 machines will be set up to run pfmon in system-wide mode
 - A significant portion of the next batch of new machines might follow in ~3-4 months.
 - Statistics will significantly increase
- Correlation of results with actual load on the machines has to be done
- Using pfmon to take a closer look at software from the experiments
- Paper is being prepared


And now for something completely different

- Another try to get into the TOP500 list
- Using 470 quad-core systems
 - Xeon E5345 CPUs – 2.33GHz
 - 3760 cores total – 35043.2 GFLOPS theoretical max.
 - 16GB RAM per system
 - Only 470 systems instead of the envisioned ~1200 because of delivery problems from our suppliers :-(
 - Gigabit Ethernet interconnect
 - Edge: HP ProCurve 3500yl
 - Core: Force10 E1200
- Worked very closely with Sergey Shalnov (Intel)




19.69 TFLOPS

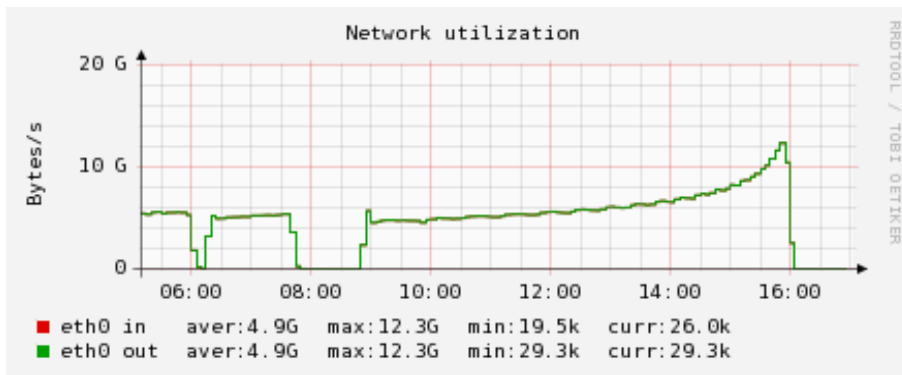
- ~56% of theoretical max. !
- Would be #50 in current list
- Would be the fastest Gigabit Ethernet cluster




CPU Utilization – linpack machines



CPU Utilization – all batchnodes



Network Utilization – linpack machines



Network Utilization – all batchnodes