



IHEPCCC/HEPiX Benchmarking WG Status Update

Helge Meinhard / CERN-IT

LCG Management Board 08 July 2008 /
Grid Deployment Board 09 July 2008



Progress since May

- Update to my presentation on 13-May-2008
- All experiments have run a variety of benchmarks on benchmarking nodes
- Simple correlation studies done
- Some anomalies found, in particular in Atlas simulation
- Some work done to understand how random numbers are used
- Some work done to understand impact of HW architecture on random number sequences
- Quantitative estimates for errors done
- Full, statistically consistent least-squares fit in the pipeline, not yet done because of other priorities of the people involved

Options (1)

- Wait for a detailed understanding of the differences in random number usage, and for a full statistical least-squares fit to all data
 - With summer period coming up, no final result before autumn
- Trust the reasonable correlations shown at HEPiX, and pick a (random) choice of one of the SPECcpu2006 benchmarks
 - Not very scientific, not bullet-proof
 - Indications are that it is good enough
 - Can be done now

- Choice of benchmark:
 - SPECint2006
 - SPECfp2006
 - SPECcall_cpp2006
- Choice of benchmarking environment
 - OS
 - Compiler
 - Compilation options
 - Running mode

- SPECint2006 (12 applications)
 - Well established, published values available
 - HEP applications are mostly integer calculations
 - Correlations with experiment applications shown to be fine
- SPECfp2006 (17 applications)
 - Well established, published values available
 - Correlations with experiment applications shown to be fine
- SPECcall_cpp2006 (7 applications)
 - Exactly as easy to run as is SPECint2006 or SPECfp2006
 - No published values (not necessarily a drawback)
 - Takes about 6 h (SPECint2006 or SPECfp2006 are about 24 h)
 - Best modelling of FP contribution to HEP applications
 - Important memory footprint

CERN cluster results (1)


Machine	SPECint2000	SPECint2006	SPECfp2006	SPECcall_cpp 2006
lxbench01	1501	11.06	9.50	20.46
lxbench02	1495	10.09	7.70	19.14
lxbench03	4133	28.76	25.23	55.82
lxbench04	5676	36.77	27.85	70.35
lxbench05	6181	39.39	29.72	75.61
lxbench06	4569	31.44	27.82	60.72
lxbench07	9462	60.89	43.47	113.95
lxbench08	10556	64.78	46.48	120.27

CERN cluster results (2)

Machine	SPECint2006 * 100 / SPECint2000	SPECcall_cpp2006 * 100 / SPECint2000	SPECcall_cpp2006 / SPECint2006
lxbench01	0.74	1.36	1.85
lxbench02	0.67	1.28	1.90
lxbench03	0.70	1.35	1.94
lxbench04	0.65	1.24	1.91
lxbench05	0.64	1.22	1.92
lxbench06	0.69	1.33	1.93
lxbench07	0.64	1.20	1.87
lxbench08	0.61	1.14	1.86

CERN cluster results (3)

Scaled SPEC results



- As discussed often times and as recommended by benchmarking WG:
 1. SL4 x86_64
 2. System compiler gcc 3.4.6
 3. Flags as defined by LCG-SPI (formerly LHC architects' forum)
`-O2 -fPIC -pthread -m32`
 4. Multiple independent parallel runs
- Script available that takes care of steps 2. to 4. and prints one single number – only need SPECcpu2006 and machine installed with SL4 x86_64

- Adopt SPECcall_cpp2006 and environment presented
- CERN will make script available
- Sites to purchase a SPECcpu2006 licence
- Name: Call it HEP-SPEC or something like this?
- Transition of existing capacities: Run old benchmark the old way (whatever that was), and HEP-SPEC on samples of all HW configs in use
- Transition of existing requests/commitments: I presume these were based on some reference machine type; both benchmarks to be run there

Backup / reference slides from previous presentations

- Ran SPECcpu with perfmon (low-level CPU performance counters)
- Perfmon analysis of SPECcpu
 - SI2000 is 1% FP (SI2006: 0.1%)
 - Average bus utilisation low, but Bus Not Ready at percent level
 - L2 cache misses 2x higher on SI2006 than on SI2000
 - SPEC2006 C++: About 10...14% FP
- Perfmon requires a new kernel
 - Can't run on SLC4 2.6.9-xx kernel
 - Even with perfmon enabled, systems running 2.6.24-xx are ~ 10% faster than with 2.6.9-xx ...!

- Full least-squares fit awaiting consistent estimation of measurement errors
- Meanwhile, assumed linear dependency and checked Pearson's correlation coefficient

$$r = \frac{1}{n-1} \sum \left(\frac{X_i - \bar{X}}{s_X} \right) \left(\frac{Y_i - \bar{Y}}{s_Y} \right)$$

- For SPECint2006 vs SPECint2000, SPECfp2006 vs SPECint2000: see Alex' talk

Alice results (preliminary)

Benchmark	Test	Exp. Results versus		
		SPECint2000	SPECint2006	SPECfp2006
pp MinBias	GEN+SIM	0.974	0.981	0.980
	DIGI	0.949	0.959	0.979
	RECO	0.956	0.966	0.989
	TOTAL(SUM)	0.965	0.974	0.983
PbPb per2 8.6 - 11.2fm	GEN+SIM	0.976	0.983	0.982
	DIGI	0.754	0.752	0.682
	RECO	0.942	0.949	0.943
	TOTAL(SUM)	0.976	0.983	0.983

CMS results (preliminary) (1)

Benchmark	Test	Exp. Result versus...		
		SPECint2000	SPECint2006	SPECfp2006
HiggsZZ4LM190	GEN+SIM	0.983	0.988	0.986
	DIGI	0.971	0.977	0.974
	RECO	0.979	0.985	0.983
	TOTAL(SUM)	0.982	0.988	0.986
MinBias	GEN+SIM	0.982	0.988	0.986
	DIGI	0.972	0.978	0.973
	RECO	0.970	0.976	0.970
	TOTAL(SUM)	0.981	0.987	0.984

CMS results (preliminary) (2)

QCD_80_120	GEN+SIM	0.980	0.986	0.984
	DIGI	0.973	0.980	0.976
	RECO	0.975	0.981	0.977
	TOTAL(SUM)	0.980	0.986	0.983
SingleElectron E1000	GEN+SIM	0.983	0.989	0.988
	DIGI	0.970	0.976	0.974
	RECO	0.962	0.968	0.960
	TOTAL(SUM)	0.983	0.989	0.987

CMS results (preliminary) (3)

QCD_80_120	GEN+SIM	0.980	0.986	0.984
	DIGI	0.973	0.980	0.976
	RECO	0.975	0.981	0.977
	TOTAL(SUM)	0.980	0.986	0.983
SingleElectronE1000	GEN+SIM	0.983	0.989	0.988
	DIGI	0.970	0.976	0.974
	RECO	0.962	0.968	0.960
	TOTAL(SUM)	0.983	0.989	0.987

CMS results (preliminary) (4)

TTbar	GEN+SIM	0.982	0.987	0.985
	DIGI	0.974	0.980	0.975
	RECO	0.902	0.908	0.891
	TOTAL(SUM)	0.977	0.982	0.978
	Total Total	0.969	0.975	0.970

Atlas results (preliminary)

Machine	ATLAS Generation	ATLAS Simulation	ATLAS Digitization	ATLAS Reconstruction	ATLAS Total
lxbench01	5.291	0.001	0.048	0.055	0.001
lxbench02	5.236	0.001	0.052	0.053	0.001
lxbench03	6.579	0.002	0.076	0.063	0.002
lxbench04	9.434	0.002	0.088	0.109	0.002
lxbench05	10.537	0.003	0.115	0.093	0.003
lxbench06	7.692	0.002	0.088	0.070	0.002
lxbench07	8.333	0.002	0.089	0.099	0.002
SPECint2000	0.645	0.679	0.726	0.691	0.685
SPECint2006	0.651	0.686	0.729	0.706	0.692
SPECfp2006	0.693	0.737	0.760	0.752	0.743