



IHEPCCC/HEPiX Benchmarking WG Status Report

Helge Meinhard / CERN-IT
HEPiX at CERN
08 May 2008



- Autumn 2006: IHEPCCC chair contacts HEPiX conveners about help on two technical topics
 - File systems
 - (CPU) Benchmarking
- HEPiX meeting at Jefferson Lab (Oct 2006) reacted favourably
- Participation sorted out
 - Chairs: Maslennikov (CASPUR): File systems, HM (CERN): Benchmarking
- File system WG started immediately... not so the Benchmarking WG
 - Mea culpa... and my apologies to everyone



History (2)

- Ad-hoc discussion at HEPiX at DESY (April 2007)
- HEPiX in St Louis (November 2007)
 - Benchmarking track
 - INFN move to SI2006 (Michele Michelotto)
 - Update on GridKA benchmarking (Manfred Alef)
 - Topics around CERN benchmarking (HM)
 - AMD and Intel performance for lattice QCD codes (Don Holmgren)
 - 2 more formal face-to-face meetings of people interested in benchmarking, formal start of WG
- Regular phone conferences
 - About once every two weeks, more frequently when needed

- Used so far (at least in LCG): int_base part of SPEC CPU 2000
 - Inconsistent usage across institutions
 - Measurements under real conditions vs. looking up values
 - Experiment code shown not to scale well with SI2k
 - In particular on multicore processors
 - SI2k memory footprint too small
 - SPEC CPU 2000 phased out in 2007
 - No more support by SPEC
 - No new entries in table of results maintained by SPEC
- New benchmark required
 - Consistent description of experiment requirements, lab commitments and existing resources
 - Tool for procurements
 - Buy performance rather than boxes
- “Natural” candidate: SPEC CPU 2006 (int_base or int_rate part)



WG composition

- Currently on mailing list:
 - Ian Gable (U Victoria)
 - Alex Iribarren (CERN)
 - Helge Meinhard (CERN)
 - Michael Barnes (Jefferson Lab)
 - Sandy Philpott (Jefferson Lab)
 - Manfred Aef (GridKA)
 - Michele Michelotto (INFN Padova)
 - Martin Bly (RAL)
 - Peter Wegner (DESY)
 - Alberto Aimar (CERN)
 - Sverre Jarp, Andreas Hirstius, Andrzej Nowak (all CERN)
 - Atlas: Alessandro di Salvo, Franco Brasolin
 - CMS: Gabriele Benelli et al.
 - Alice: Peter Hristov
 - LHCb: Hubert Degaudenzi



- Focus initially on benchmarking of processing power for worker nodes
- Representative sample of machines needed; centres who can spare a machine (at least temporarily) will announce this to the list
- Environment to be fixed
- Standard set of benchmarks to be run
- Experiments to be invited (i.e. pushed) to run their code
 - Check how well experiments' code scales with industry-standard benchmarks



- Dedicated benchmarking cluster at CERN (see Alex Iribarren's talk)
- DESY Zeuthen:
 - 2 x Harpertown 2.83 GHz
- RAL:
 - 2 x Barcelona – not yet stable
- INFN Padova:
 - 2 x Harpertown 2.33 GHz
 - 2 x Barcelona 2.1 GHz



- SL 4 x86_64, 32-bit applications
 - Some cross-checks with SL 5
- Gcc 3.4.x (system compiler)
 - Cross-checks with gcc 4; will need gcc 4 for SPEC FP 2006
- Compilation options by LCG-SPI (formerly LHC architects' forum)
`-O2 -fPIC -pthread -m32`
- Multiple independent parallel runs
 - Multi-threaded benchmarks (SPECrate) checked as well – although normalisation is the same, values too low by ~5% on fast machines



Standard benchmarks

- **Compulsory:**
 - SPECint_base 2000 (parallel, transition only)
 - SPECint_base 2006 (parallel)
 - SPECfp_base 2006 (parallel, gcc 4)
- **Optional, cross-checks only:**
 - SPECint_rate 2006
 - All above as 64-bit applications
 - SPECint compiled with gcc 4
 - All above under SL5
- **Notes:**
 - “parallel” – as many instances as there are cores run in parallel, results added over all cores
 - Applications are 32 bit unless noted otherwise

- 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)

Exp. Results versus

...

Benchmark	Test	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



LHCb status

- First results received for performance of minimum bias pp on CERN's benchmarking machines
- No fits done yet



Conclusions

- WG has taken long to get started, but is now up and running, and has produced preliminary results
- Momentum is there, interest and (some) person power as well
- Good participation of LHC experiments – thanks!
- Agreement on methodology
- Preliminary results indicate that standard benchmarks are adequate
- Run with perfmon (CPU performance/event counters)
- Some interest in looking at power consumption as well

Alice results (preliminary)

Benchmark	Test	lxbench01	lxbench02	lxbench03	lxbench04	lxbench05	lxbench06	lxbench07	SPECint2000	SPECint2006	SPECfp2006
pp MinBias	GEN+SIM	1.96E-02	1.95E-02	6.72E-02	8.56E-02	9.52E-02	7.80E-02	1.51E-01	0.974	0.981	0.980
	DIGI	3.32E-02	2.68E-02	7.96E-02	9.80E-02	1.00E-01	8.44E-02	1.37E-01	0.949	0.959	0.979
	RECO	2.22E-01	2.48E-01	5.96E-01	7.40E-01	7.36E-01	7.00E-01	9.92E-01	0.956	0.966	0.989
	TOTAL(SUM)	1.17E-02	1.08E-02	3.42E-02	4.32E-02	4.60E-02	3.84E-02	6.70E-02	0.965	0.974	0.983
PbPb per2	GEN+SIM	3.92E-04	3.96E-04	1.32E-03	1.66E-03	1.86E-03	1.54E-03	2.91E-03	0.976	0.983	0.982
	DIGI	1.31E-03	1.26E-03	3.62E-03	4.56E-03	4.92E-03	4.40E-03	7.14E-01	0.754	0.752	0.682
8.6 - 11.2fm	RECO	8.88E-03	9.44E-03	2.51E-02	3.82E-02	4.24E-02	2.79E-02	6.66E-02	0.942	0.949	0.943
	TOTAL(SUM)	2.92E-04	2.92E-04	9.32E-04	1.18E-03	1.31E-03	1.10E-03	2.01E-03	0.976	0.983	0.983

CMS results (preliminary)

Benchmark	Test	Ixbench01	Ixbench02	Ixbench03	Ixbench04	Ixbench05	Ixbench06	Ixbench07	SPECint2000	SPECint2006	SPECfp2006
HiggsZZ4LM190	GEN+SIM	0.012046	0.01217	0.03794	0.046056	0.051428	0.045332	0.080952	0.983	0.988	0.986
	DIGI	0.62131	0.647668	1.97336	2.595716	2.886004	2.293576	4.54804	0.971	0.977	0.974
	RECO	0.211752	0.229542	0.688468	0.842816	0.94384	0.802568	1.48616	0.979	0.985	0.983
	TOTAL(SUM)	0.011192	0.011354	0.035316	0.042944	0.04796	0.042124	0.075496	0.982	0.988	0.986
MinBias	GEN+SIM	0.058672	0.059228	0.182116	0.221584	0.24614	0.21718	0.390584	0.982	0.988	0.986
	DIGI	0.995024	0.943842	2.83286	3.741816	4.20168	3.36984	6.535944	0.972	0.978	0.973
	RECO	0.962	0.985222	2.702704	3.561888	3.95648	3.142184	6.284368	0.970	0.976	0.970
	TOTAL(SUM)	0.052388	0.052746	0.160928	0.197588	0.219612	0.191588	0.348144	0.981	0.987	0.984
QCD_80_120	GEN+SIM	0.012504	0.012446	0.038656	0.04766	0.053276	0.045948	0.083728	0.980	0.986	0.984
	DIGI	0.671818	0.61862	1.929572	2.518892	2.772004	2.27402	4.385968	0.973	0.980	0.976
	RECO	0.21306	0.21732	0.631912	0.803696	0.89666	0.738824	1.41468	0.975	0.981	0.977
	TOTAL(SUM)	0.011608	0.011552	0.035752	0.0442	0.049392	0.042452	0.077656	0.980	0.986	0.983
SingleElectronE1000	GEN+SIM	0.007772	0.007758	0.025188	0.030244	0.03398	0.030076	0.053008	0.983	0.989	0.988
	DIGI	1.017812	0.962928	3.117692	4.20168	4.656576	3.696856	7.285976	0.970	0.976	0.974
	RECO	1.128668	1.142204	2.936856	4.110996	4.5819	3.454232	7.272728	0.962	0.968	0.960
	TOTAL(SUM)	0.007662	0.007644	0.024772	0.029808	0.033488	0.02958	0.052248	0.983	0.989	0.987