

The CMS Software Performance at the Start of Data Taking

Gabriele Benelli
CERN



Technisches Seminar, DESY Zeuthen, October 21st 2008



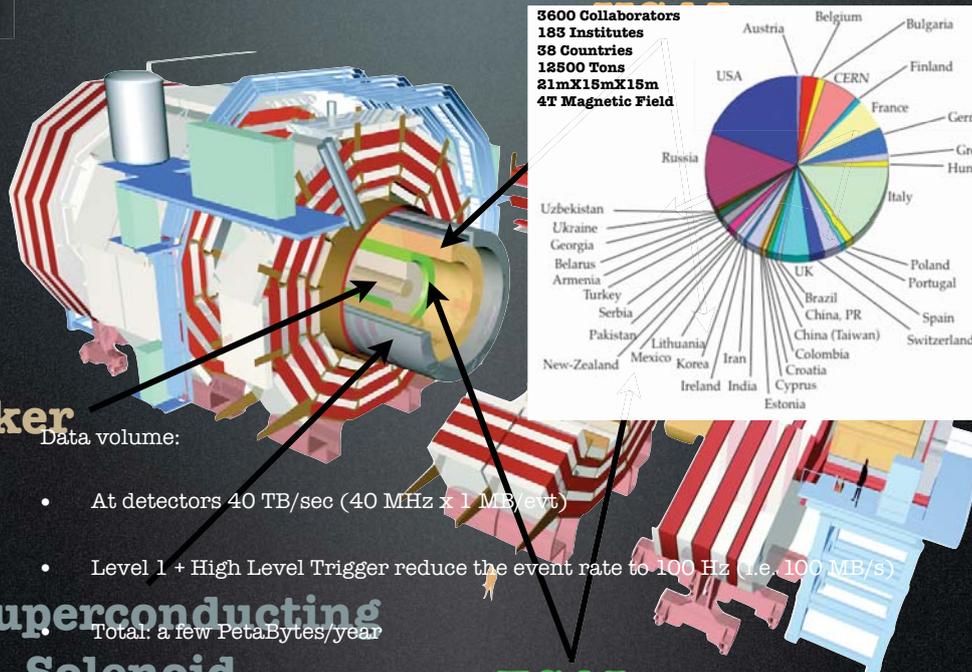
Outline



- The CMS Experiment
- CMS Offline software
- Performance Profiling in CMS
- CMSSW Performance Suite
- Using CMSSW as a benchmarking tool
- Conclusions/Outlook



The CMS experiment



Tracker

Data volume:

- At detectors 40 TB/sec (40 MHz x 1 MB/evt)
- Level 1 + High Level Trigger reduce the event rate to 100 Hz (i.e. 100 MB/s)

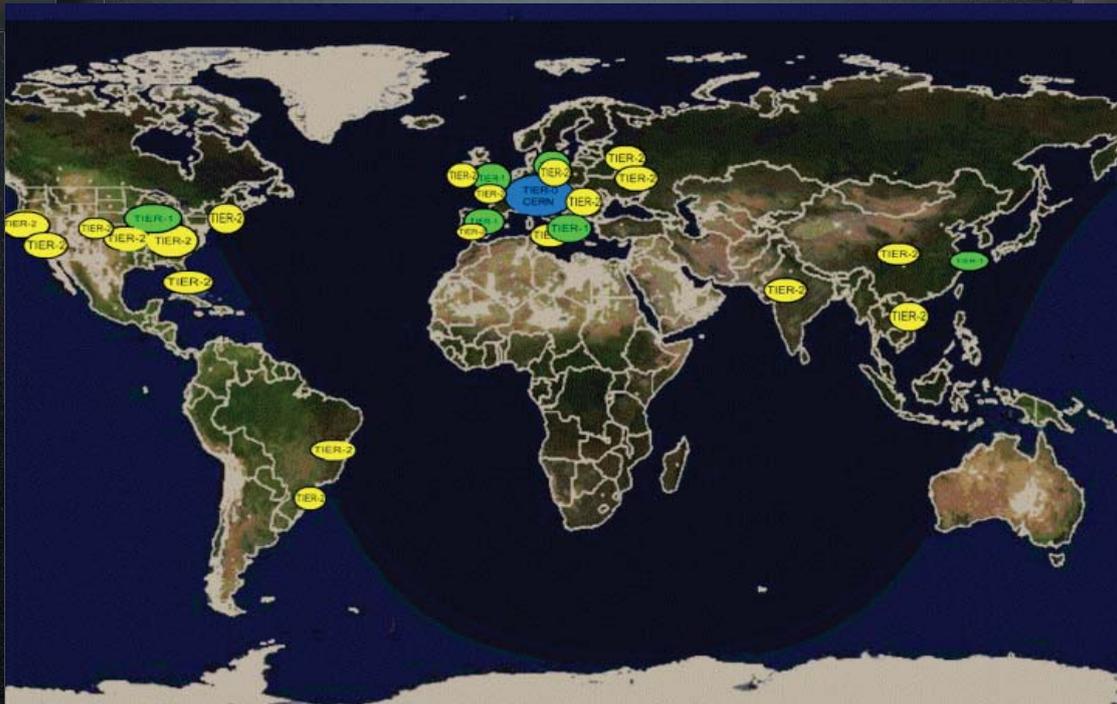
4T Superconducting Solenoid

Total: a few PetaBytes/year

- Offline reconstruction (simulation and analysis) done using the Grid, thousands of computers distributed worldwide, organized in Tier centers



Tiered Computing Model



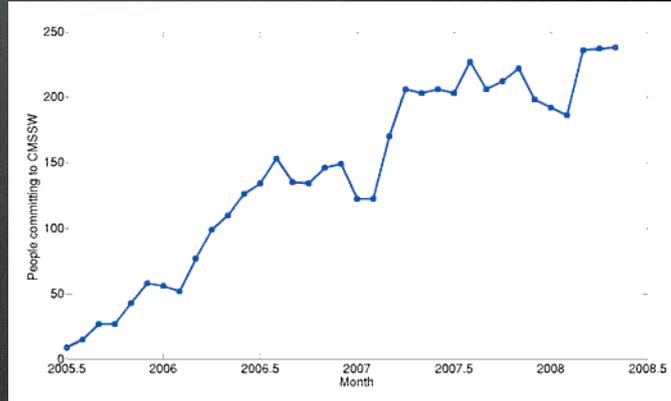


CMS Offline Software



- CMS software framework (CMSSW) consists of approximately:
 - 1100 individual (CVS) packages, organized in 100 Subsystems
 - 2 Millions lines of code (SLOC)
 - 100 external packages (mainly Open Source)
 - 1.5 GB of “data.” packages (mainly for Fast Simulation)

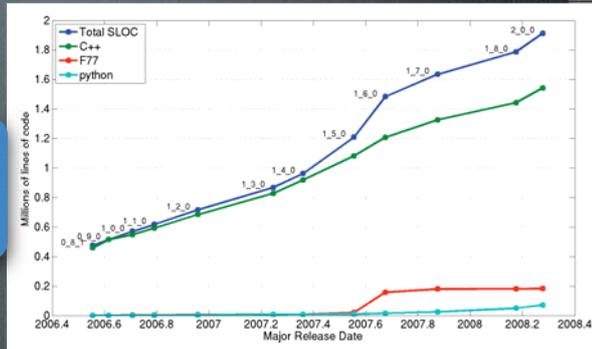
250 active developers
(a lot of work in the
last 3 years!)



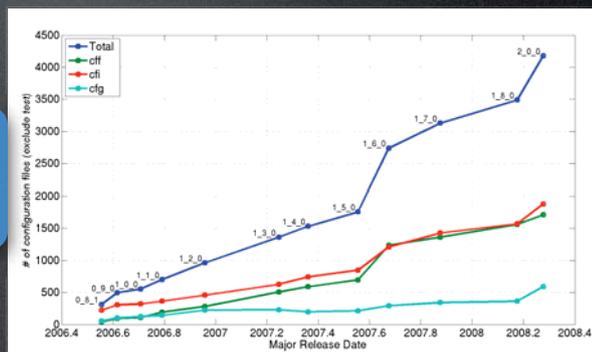
Development metrics



Source Lines Of Code (SLOC)
in C++, Fortran, Python
[F77 contribution due to externally
developed generators]



**Configuration Language
contribution**





CMS Offline Development Model



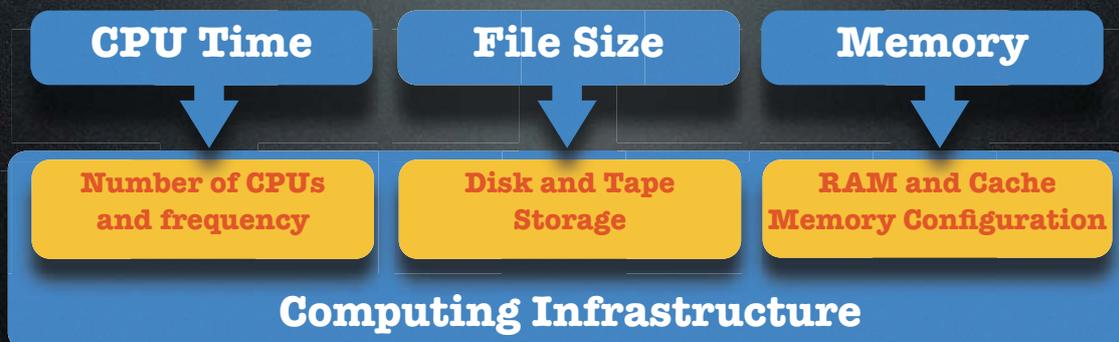
- New tags get queued by developers and need to be blessed (i.e. tested) by the relevant level 2 manager for them to be collected by the Tag Collector
- Integration Builds:
 - Two per day for each release cycle and platform
- Development and Production Releases
 - “open” and “close” phases
- Full set of packages has to build always
- Partial releases done later(FWLite, Online)



Performance Profiling



- The CMS Software Framework (**CMSSW**) rapid evolution requires strict quality assurance and the measurement and monitoring of performance is integral part of the Release Validation (**RelVal**) effort
- The 3 key metrics of software performance:



- The monitoring of these offers guidance to software development and optimization effort



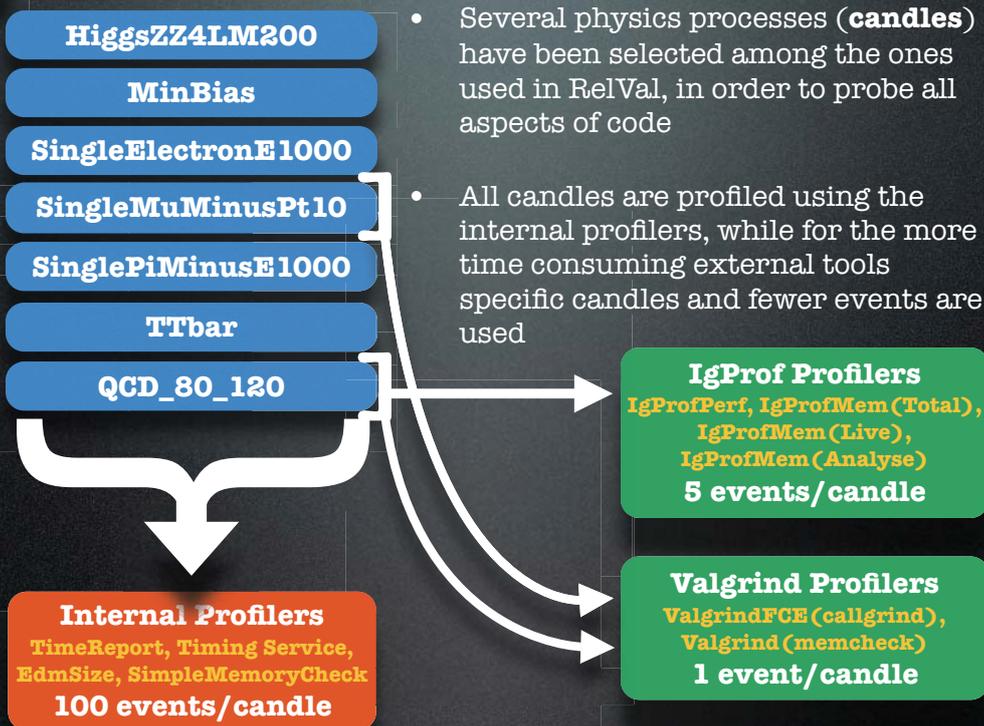
Profilers in CMSSW



- CMSSW contains internal tools to profile itself in terms of CPU time, file size and memory use
- Also external tools (IgProf and Valgrind) are used to gain more information, with significant penalty in terms of execution time
- A Performance Suite of tools has been developed to make use of all these profiling tools, integrating them in the Release Validation, providing a regular measurement of CMSSW performance and allowing regression between software releases



The CMSSW Performance Suite



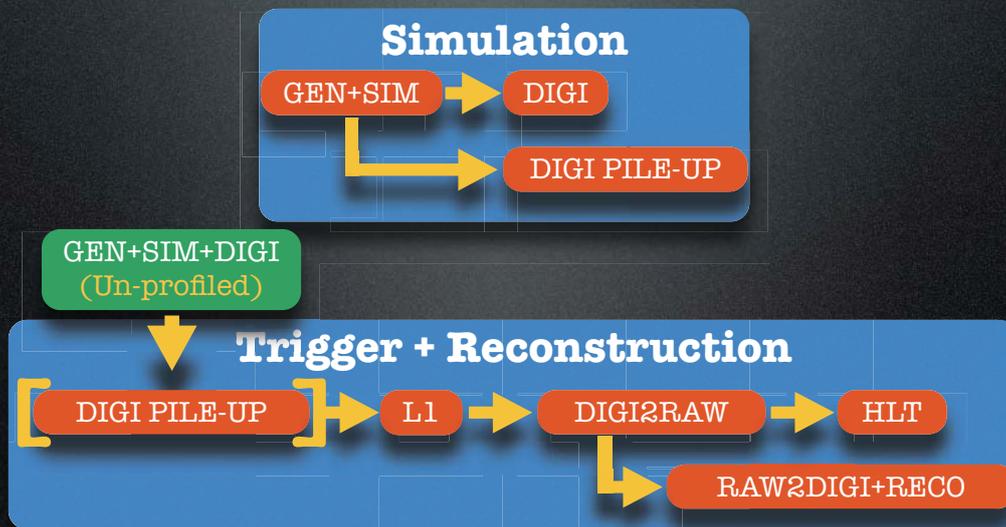
- Several physics processes (**candles**) have been selected among the ones used in RelVal, in order to probe all aspects of code
- All candles are profiled using the internal profilers, while for the more time consuming external tools specific candles and fewer events are used



Profiling CMSSW Steps



- Performance profiling is based on simulated data and it is done for each step individually (output of one step used as input for the next)
- CMSSW pre-release development cycle is 1 week, so based on the different steps execution time, two **24hrs**-workflows have been foreseen, running in parallel on separate machines:



CMSSW Performance Suite



- In order to ensure meaningfulness to regression analysis, the Performance Suite is run on 2 dedicated multicore machines

Simulation Performance Reports for CMSSW_3_0_X_2008-10-10-0200

Simulation Performance Suite run on intel@06.cam.ac.uk at /home/benelli/cmssw_3_0_X_2008-10-10-0200

Results produced on: intel@06.cam.ac.uk

Logfiles available at: /home/benelli/cmssw_3_0_X_2008-10-10-0200

Results for simScoreM2 benchmark (running on the other cores) available at: /home/benelli/cmssw_3_0_X_2008-10-10-0200

Click here to browse the directory containing all results (except the root files)

Results published on Thu Oct 16 13:11:27 2008

Release ROOT file sizes

Table showing current release ROOT file sizes in B/N/G bytes.

Filesize	GEN_SIM	DIGI	RAW2DIGI+RECO	L1	DIGI2RAW	HLT	Total
MinBias	17.6M	37.03M	N/A	N/A	N/A	N/A	54.63M
SingleZL4000	126.33M	137.13M	N/A	N/A	N/A	N/A	263.47M
SingleElectron1000	21.43M	44.85M	N/A	N/A	N/A	N/A	76.28M
SingleMuon4010	1.23M	14.8M	N/A	N/A	N/A	N/A	16.03M
SingleMuon1000	31.28M	46.3M	N/A	N/A	N/A	N/A	77.58M
TBar	107.13M	207.66M	N/A	N/A	N/A	N/A	404.79M
QCD_B0_120	146.47M	159.44M	N/A	N/A	N/A	N/A	305.91M
QCD_B0_120_PileUp	N/A	424.79M	375.25M	423.49M	413.43M	454.48M	2,332M

Release CPU times

Table showing current release CPU times in sec.

CPU Times (s)	GEN_SIM	DIGI	RAW2DIGI+RECO	L1	DIGI2RAW	HLT	Total
MinBias	12.53	2.14	N/A	N/A	N/A	N/A	14.67
SingleZL4000	79.39	2.71	N/A	N/A	N/A	N/A	82.30
SingleElectron1000	112.23	2.11	N/A	N/A	N/A	N/A	114.33
SingleMuon4010	0.45	2.28	N/A	N/A	N/A	N/A	2.54
SingleMuon1000	72.86	2.17	N/A	N/A	N/A	N/A	75.03
TBar	110.75	3.10	N/A	N/A	N/A	N/A	113.86
QCD_B0_120	92.01	2.86	N/A	N/A	N/A	N/A	94.87
QCD_B0_120_PileUp	N/A	16.55	8.72	0.21	0.56	0.62	26.66

MinBias

Logfiles for TimeSize

GenLog
MinBias_Gen_Sim_TrimReport.log
MinBias_Digi_TrimReport.log
MinBias_Reco_TrimReport.log

TimeSize

- TrimReport_Gen_Sim (100 events)
- TrimReport_Digi (100 events)
- TrimReport_Gen_Sim (100 events)
- TrimReport_Digi (100 events)
- SimulTrimReport_Gen_Sim (100 events)
- SimulTrimReport_Digi (100 events)
- GenLog_Gen_Sim (100 events)
- GenLog_Digi (100 events)

- The power-saving BIOS settings and daemons have been disabled and while we run on one core, we run a small cache-contained benchmark on all the other cores, to ensure the same load conditions.
- Once the profiling is done, the suite writes static html reports for each profile and archives a tarball of the working directory on tape
- Finally, all the logs and static html with tables and graphs are published on a web server.



CPU Time Performance



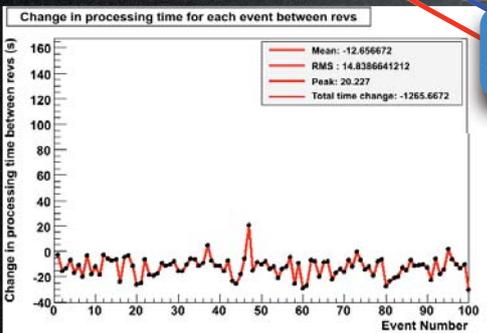
- Sample regression **CMSSW_3_1_X** (G4 9.2 BETA) vs. **CMSSW_2_1_9**

CPU Time regression summary
GEN-SIM, DIGI

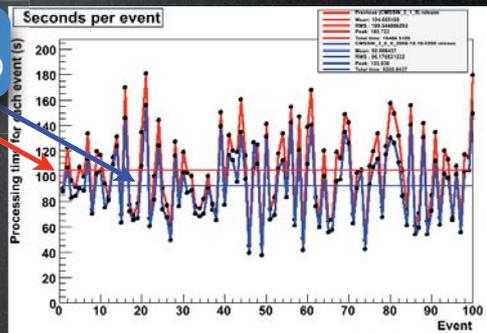
Release CPU times

Table showing previous release CPU times, t1, latest times, t2, and the difference between them Δ in secs.

Candles	GEN_SIM			DIGI			RAW2DIGI,RECO			L1			DIGI2RAW			HLT			Total		
	t1	t2	Δ	t1	t2	Δ	t1	t2	Δ	t1	t2	Δ	t1	t2	Δ	t1	t2	Δ	t1	t2	Δ
MinBias	13.44	12.53	-0.91	0.90	2.14	1.24	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	14.34	14.67	0.33
HiggsZZ4LMZ00	88.05	79.59	-8.47	1.45	2.71	1.26	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	89.50	82.30	-7.21
SingleElectronE1000	118.09	112.23	-5.86	0.84	2.11	1.27	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	118.92	114.33	-4.59
SingleMuMinusPt10	0.55	0.45	-0.09	0.87	2.08	1.21	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	1.42	2.54	1.12
SinglePIMinusE1000	81.56	72.86	-8.70	0.96	2.17	1.21	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	82.52	75.03	-7.49
TTbar	123.36	110.75	-12.61	1.86	3.10	1.24	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	125.22	113.86	-11.36
QCD_80_120	104.67	92.01	-12.66	1.62	2.86	1.25	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	106.28	94.87	-11.41
QCD_80_120 PILEUP	N/A	N/A	N/A	8.90	10.55	1.65	32.99	8.72	24.27	0.35	0.35	0.00	0.54	0.56	0.02	0.61	0.62	0.01	43.39	20.80	-22.59



GEN-SIM
QCD_80_120



CPU Time Performance



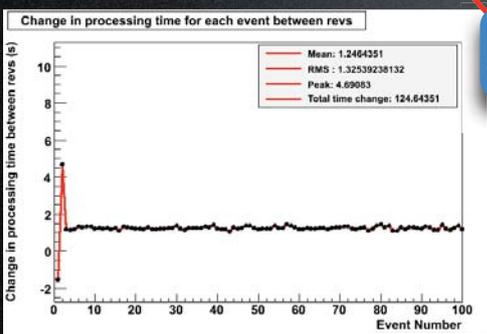
- Sample regression **CMSSW_3_1_X** (G4 9.2 BETA) vs. **CMSSW_2_1_9**

CPU Time regression summary
GEN-SIM, DIGI

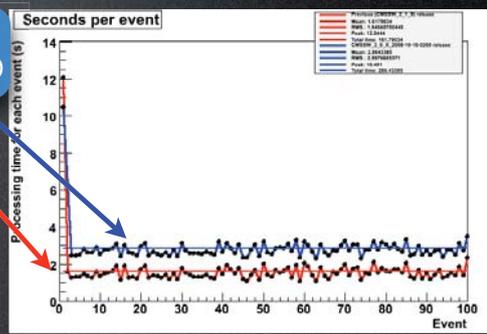
Release CPU times

Table showing previous release CPU times, t1, latest times, t2, and the difference between them Δ in secs.

Candles	GEN_SIM			DIGI			RAW2DIGI,RECO			L1			DIGI2RAW			HLT			Total		
	t1	t2	Δ	t1	t2	Δ	t1	t2	Δ	t1	t2	Δ	t1	t2	Δ	t1	t2	Δ	t1	t2	Δ
MinBias	13.44	12.53	-0.91	0.90	2.14	1.24	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	14.34	14.67	0.33
HiggsZZ4LMZ00	88.05	79.59	-8.47	1.45	2.71	1.26	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	89.50	82.30	-7.21
SingleElectronE1000	118.09	112.23	-5.86	0.84	2.11	1.27	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	118.92	114.33	-4.59
SingleMuMinusPt10	0.55	0.45	-0.09	0.87	2.08	1.21	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	1.42	2.54	1.12
SinglePIMinusE1000	81.56	72.86	-8.70	0.96	2.17	1.21	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	82.52	75.03	-7.49
TTbar	123.36	110.75	-12.61	1.86	3.10	1.24	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	125.22	113.86	-11.36
QCD_80_120	104.67	92.01	-12.66	1.62	2.86	1.25	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	106.28	94.87	-11.41
QCD_80_120 PILEUP	N/A	N/A	N/A	8.90	10.55	1.65	32.99	8.72	24.27	0.35	0.35	0.00	0.54	0.56	0.02	0.61	0.62	0.01	43.39	20.80	-22.59



DIGI
QCD_80_120

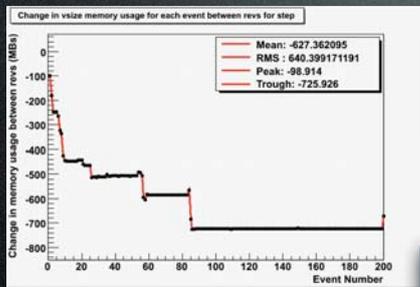




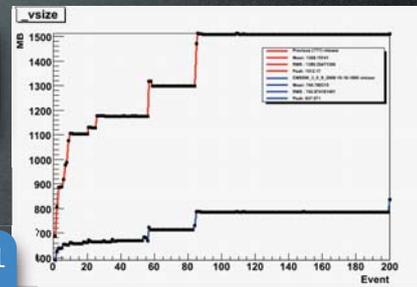
Memory Performance



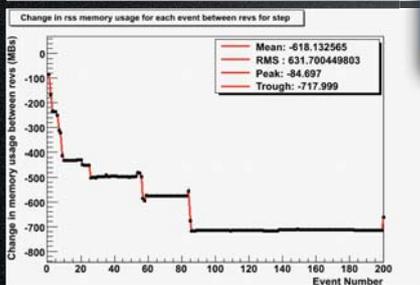
- Virtual memory size (VSIZE) is the parameter constrained in the CMS Computing TDR (1 GB/core).
- The Performance Suite can be used for regression to validate bug fixes with an arbitrary number of events



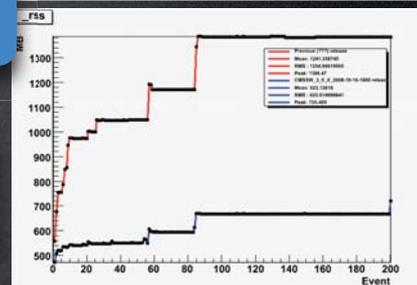
GEN+SIM
TTbar
VSIZE



CMSSW_1_7_1
SimTrack
Fix



GEN+SIM
TTbar
RSS



Gabriele Benelli, CERN

17

DESY Zeuthen, October 21st 2008



Memory Allocation



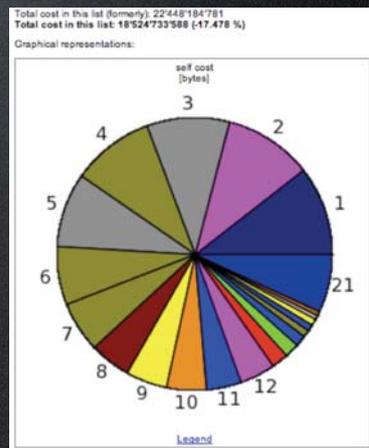
- IgProf can profile memory allocation, preserving all the callstack information for each function, producing very insightful reports:

GEN-SIM
QCD_80_120

Overall
Top 20 functions

	self cost [bytes]	Inclusive cost [bytes]	#times called	name	
-	1966582926 (10.562%)	1966582926	2942436	new[] [more]	1
-	(-19.777%)	(-0.303%)	(-0.808%)	The item appears not to have moved.	2
-	1912750276 (10.325%)	1912750276	15438560	G4NavigationHistory::G4NavigationHistory [more]	3
-	1912370948 (9.784%)	1912370948	8036496	std::vector::M_range_insert [more]	4
-	1788761976 (9.656%)	3700352296	15420361	G4Transportation::PostStepDoIt [more]	5
-	1627101936 (8.783%)	1627101936	30360	deflateInit2 [more]	6
-	1290524084 (6.986%)	1290524084	21998516	std::vector::M_insert_aux [more]	7
-	(-36.821%)	(-2.133%)	(-38.821%)	The item appears not to have moved.	8
-	1105135952 (5.966%)	1162775228	11115510	std::vector::M_insert_aux [more]	9
-	(-50.726%)	(-2.700%)	(-2.886%)	This item has moved up the list by 6 ranks. Former absolute position: 12	10
+ -	904798180 (4.884%)	904811342	8278566	TrackingAction::PreUserTrackingAction [more]	11
-	880728940 (4.754%)	880728940	4937962	std::vector::operator- [more]	12
-	871099568 (4.702%)	871099568	17963026	std::vector::operator- [more]	13
-	(-40.782%)	(-4.851%)	(-40.782%)	The item appears not to have moved.	14
+ -	785969112 (4.133%)	785969112	9725111	std::vector::reserve [more]	15
-	(-124.918%)	(-2.617%)	(-344.896%)	This item has moved up the list by 4 ranks. Former absolute position: 16	16
-	757039920 (4.087%)	757039920	78521	inflateInit2 [more]	17
-	(-20.394%)	(-0.150%)	(-20.394%)	The item appears not to have moved.	18
-	386975704 (2.088%)	386975704	11920	FrontierPayload_create [more]	19
-	301902092 (1.629%)	305365964	1347333	G4Nucleus::InitCandidateVector [more]	20
-	252136430 (1.361%)	630872264	451774	DDPFilteredView::firstChild [more]	21
-	158539784 (0.856%)	158539784	991950	std::vector::M_insert_aux [more]	
-	(-53.533%)	(-0.884%)	(-36.516%)	This item has moved down the list by 9 ranks. Former absolute position: 15	
-	146033008 (0.788%)	146033008	189934	std::vector::reserve [more]	
-	(-4.145%)	(-0.110%)	(-2.923%)	The item appears not to have moved.	
+ -	132805912 (0.717%)	129038960	1184594	G4NucleiModel::generateParticleFate [more]	
-	(-66.108%)	(-0.361%)	(-68.841%)	This element was not in the top 20 before. Former absolute position: 25	
-	96988450 (0.522%)	96988450	4028973	gnu_cxx::new_allocator::allocate [more]	
-	(-94.178%)	(-6.876%)	(-92.839%)	This item has moved down the list by 9 ranks. Former absolute position: 6	
-	89620198 (0.478%)	89620198	207702	xercesc_2_TriMemoryManagerImpl::allocate [more]	
-	1288918011 (6.956%)	n/a	n/a	{others}	

- Besides the memory information, the number of times the functions are called is relevant for CPU time profiling



Gabriele Benelli, CERN

18

DESY Zeuthen, October 21st 2008



Memory Leaks



- Memory errors and leaks are hunted down with Valgrind MemCheck:

GEN-SIM
QCD_80_120

Valgrind MemCheck output (63 leaks)

```

Leak 1: 1,893,816 bytes indirectly lost (record 803 of 811) [href]
operator new(unsigned)
PlaneBuilderFromGeometricDet: plane(GeometricDet const*) const
TrackerGeomBuilderFromGeometricDet: buildPlaneWithMaterial(GeometricDet const*, double) const
TrackerGeomBuilderFromGeometricDet: buildPixel(std::vector<GeometricDet const*, std::allocator<GeometricDet const*> > const&,
TrackerGeometry*, GeometricDetParameters: SubDetector&, std::string const&)
TrackerGeomBuilderFromGeometricDet: build(GeometricDet const*)
TrackerDigiGeometryModule: ddGeometryCallback(IdealGeometryRecord const&)
edm::eventSetup: CallbackProxy<edm::eventSetup: CallbackTrackerDigiGeometryModule, boost::shared_ptr<TrackerGeometry>,
TrackerDigiGeometryRecord, edm::eventSetup: EPFxFunctorDecoratorTrackerDigiGeometryRecord,
edm::eventSetup: DependsOnCallerTrackerDigiGeometryModule, TrackerDigiGeometryRecord, IdealGeometryRecord,
edm::eventSetup: DependsOnDoNothingCallerTrackerDigiGeometryRecord > > >, TrackerDigiGeometryRecord,
boost::shared_ptr<TrackerGeometry> >: make(TrackerDigiGeometryRecord const&, edm::eventSetup: DataKey const&)
edm::eventSetup: DataProxyTemplateTrackerDigiGeometryRecord, TrackerGeometry>: get(TrackerDigiGeometryRecord const&,
edm::eventSetup: DataKey const&) const
void
edm::eventSetup: EventSetupRecordImplementationTrackerDigiGeometryRecord: getImplementationTrackerGeometry(TrackerGeometry
pluginCallIDTrackerStripConnectivityPlugins
edm::eventSetup: ComponentDescription const&*) const
edm::SiPixelDigitizer: produce(edm::Event&, edm::EventSetup const&)
plugin3inTrackerSiPixelDigitizerPlugins

Leak 2: 1,580,208 (104,992 direct, 1,475,216 indirect) bytes definitely lost (record 772 of 811) [href]
operator new(unsigned)
TrackerGeomBuilderFromGeometricDet: buildGeomDet(TrackerGeometry*)
TrackerGeomBuilderFromGeometricDet: build(GeometricDet const*)
TrackerDigiGeometryModule: ddGeometryCallback(IdealGeometryRecord const&)
edm::eventSetup: CallbackProxy<edm::eventSetup: CallbackTrackerDigiGeometryModule, boost::shared_ptr<TrackerGeometry>,
TrackerDigiGeometryRecord, edm::eventSetup: EPFxFunctorDecoratorTrackerDigiGeometryRecord,
edm::eventSetup: DependsOnCallerTrackerDigiGeometryModule, TrackerDigiGeometryRecord, IdealGeometryRecord,
edm::eventSetup: DependsOnDoNothingCallerTrackerDigiGeometryRecord > > >, TrackerDigiGeometryRecord,
boost::shared_ptr<TrackerGeometry> >: make(TrackerDigiGeometryRecord const&, edm::eventSetup: DataKey const&)
edm::eventSetup: DataProxyTemplateTrackerDigiGeometryRecord, TrackerGeometry>: get(TrackerDigiGeometryRecord const&,
edm::eventSetup: DataKey const&) const
void
edm::eventSetup: EventSetupRecordImplementationTrackerDigiGeometryRecord: getImplementationTrackerGeometry(TrackerGeometry
pluginCallIDTrackerStripConnectivityPlugins
edm::eventSetup: ComponentDescription const&*) const
edm::SiPixelDigitizer: produce(edm::Event&, edm::EventSetup const&)
plugin3inTrackerSiPixelDigitizerPlugins
  
```

Leak observed in **SiPixelDigitizer**,
module observed already in CPU Time



CMSSW as

a machine benchmarking tool





CMSSW benchmarking



- In the context of the HEPiX CPU performance working group CMSSW has been used as a tool to do CPU benchmarking
- The working group charge is to:
 - Test the validity of the industry-standard benchmarks (SPEC CPU) when compared with HEP experiments code
 - Provide some recommendation to guide institutional purchases
- CMSSW applications have been used to benchmark a number of machines with different architecture



CMSSW benchmarking



- All tests run using the Performance Suite, in CMSSW_2_0_0_pre5 release (making use of the CPU time profiling capability)
- All seven “candles” (physics processes) were used
- Run 100 events per candle
- Run GEN+SIM, DIGI, RECO steps separately
- Run the 7 candles sequentially on each core
- Four different tests:
 - Loading all cores simultaneously
 - Loading 1, 3 (only for 4 cores machines) and 5 cores (only for 8 cores machines) with our application while running a cpu-intensive, cache-contained benchmarking tool (cmsScimark2) on the other cores



Benchmarked machines



- Used 10 machines with different architectures, frequencies, memory:
 - 7 machines at CERN from the lxbench cluster
 - Cluster TWiki: <https://twiki.cern.ch/twiki/bin/view/FIOgroup/TsiLxbench>
 - 1 machine at DESY Zeuthen (hpbl1)
 - 2 machines at INFN Padua (lxcmsrv7, lxcmsrv8)

	lxbench01	lxbench02	lxbench03	lxbench04	lxbench05	lxbench06	lxbench07	lxcmsrv07	lxcmsrv08	hpbl1
Number of cores	2	2	4	4	4	4	8	8	8	8
Frequency (GHz)	2.8	2.8	2.2	2.66	3.0	2.6	2.33	2.33	2.1	2.83
Cache (could be L2/L3) (MB)	1	2	2	4	4	2	8	12	2	12
Memory (GB)	2	4	2	8	8	8	16	16	16	16
Processor	Nocona	Irvingdale	Opteron 275	Woodcrest	Woodcrest	Opteron 2218 Rev.F	Clovertown	Xeon HarperTown E5410	Opteron Barcelona 2352	Xeon HarperTown E5440
Vendor	Intel	Intel	AMD	Intel	Intel	AMD	Intel	Intel	AMD	Intel



Benchmarking the cluster



- Basically the Performance Suite was submitted on the wanted cores with 100 events and all the internal profilers
- Once the Suite was done running, the logfiles were “harvested” and the timing information from the framework was used to calculate the average for each candle
- The data was collected in a Python dictionary and then tables for publication on the HEPiX Wiki were produced
- Using the same data structure, comparison/analysis plots for the various machines could be generated with Matplotlib



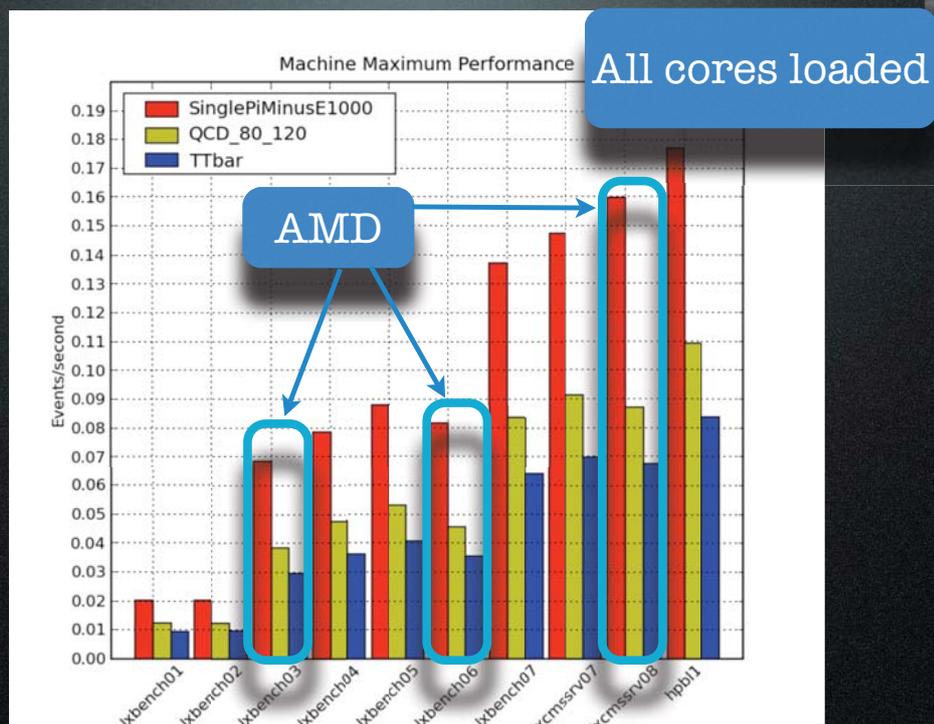
CMSSW Benchmarking



- The result of the benchmarking is seconds/event averaged on the 99 events (skipping the first one to avoid biases due to initialization)
- The results are reported in 3 formats:
 - seconds/event per core
 - events/second per core
 - events/second per machine
- Link: <https://hep.caspr.it/processors/dokuwiki/doku.php?id=benchmarks:cms>

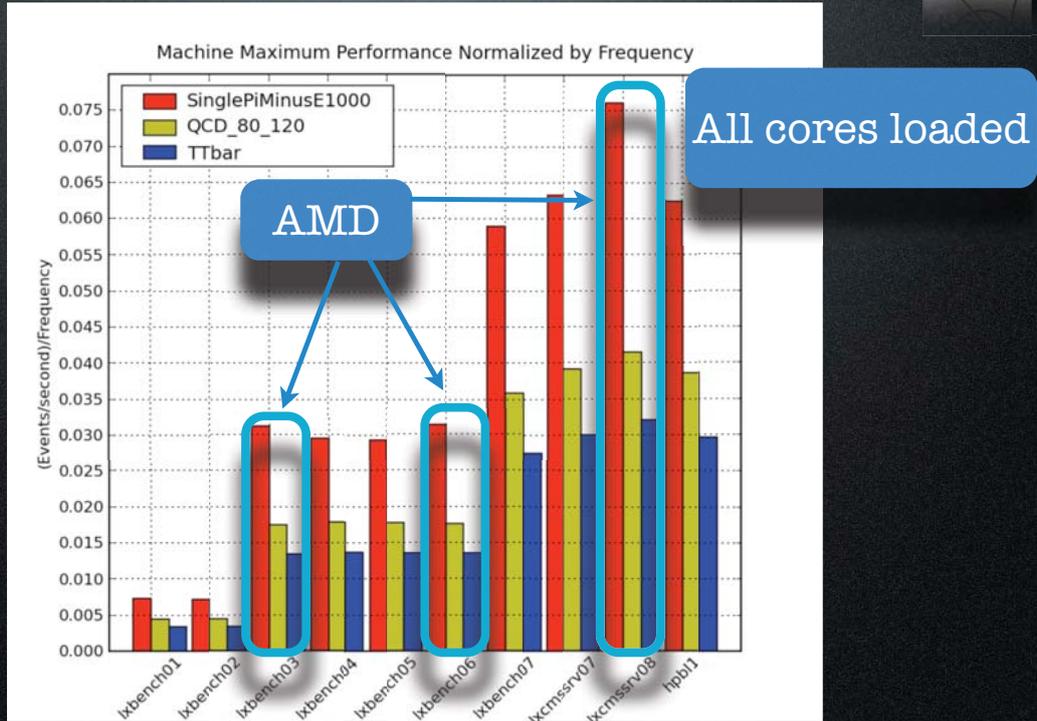


Comparing GEN+SIM

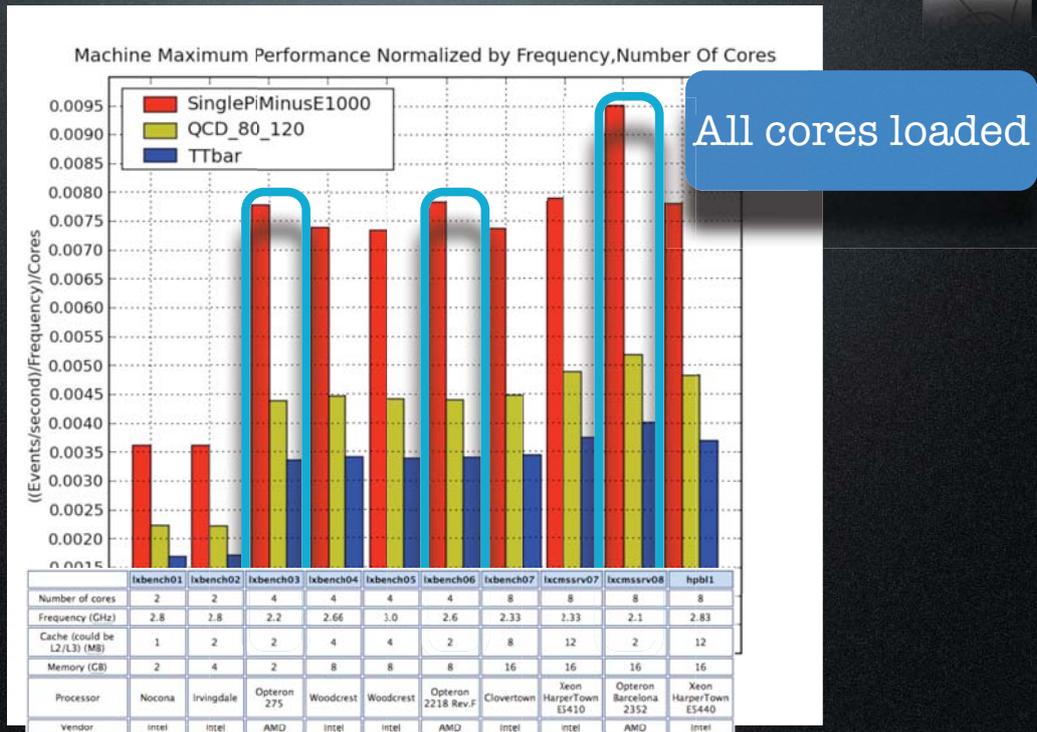




GEN+SIM normalized by frequency

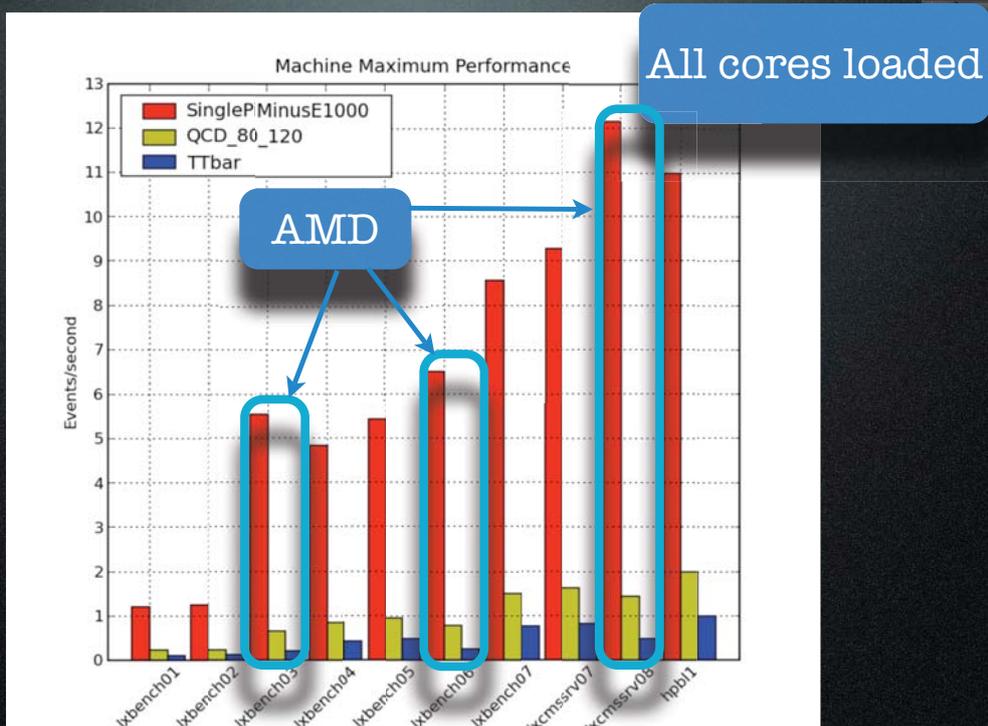


GEN+SIM normalized by frequency and # cores

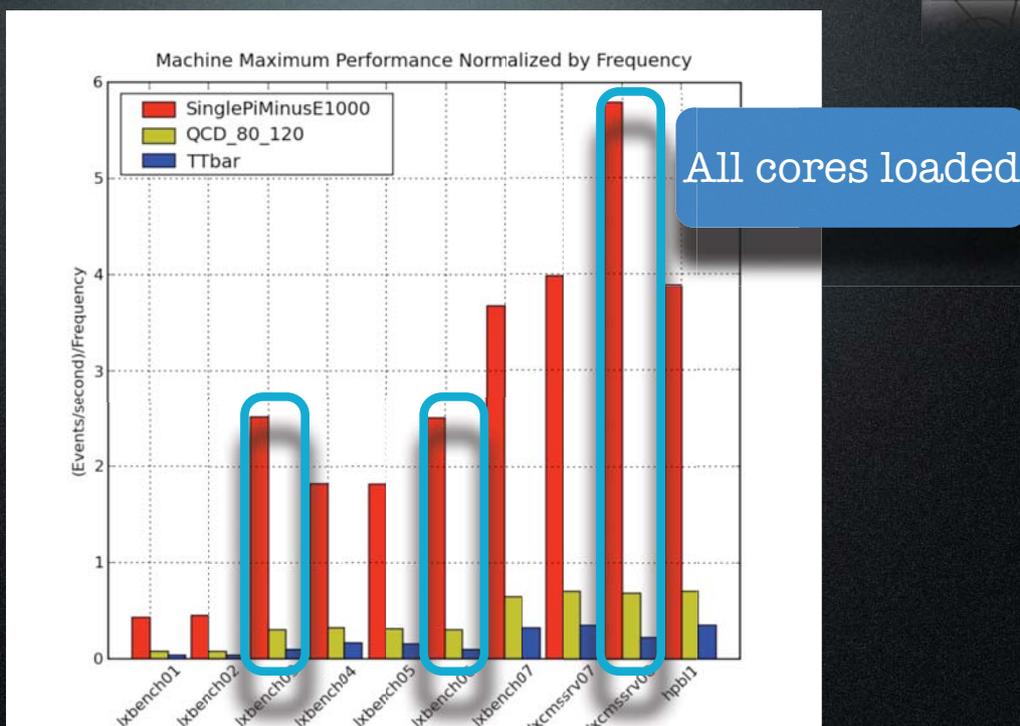




Comparing RECO

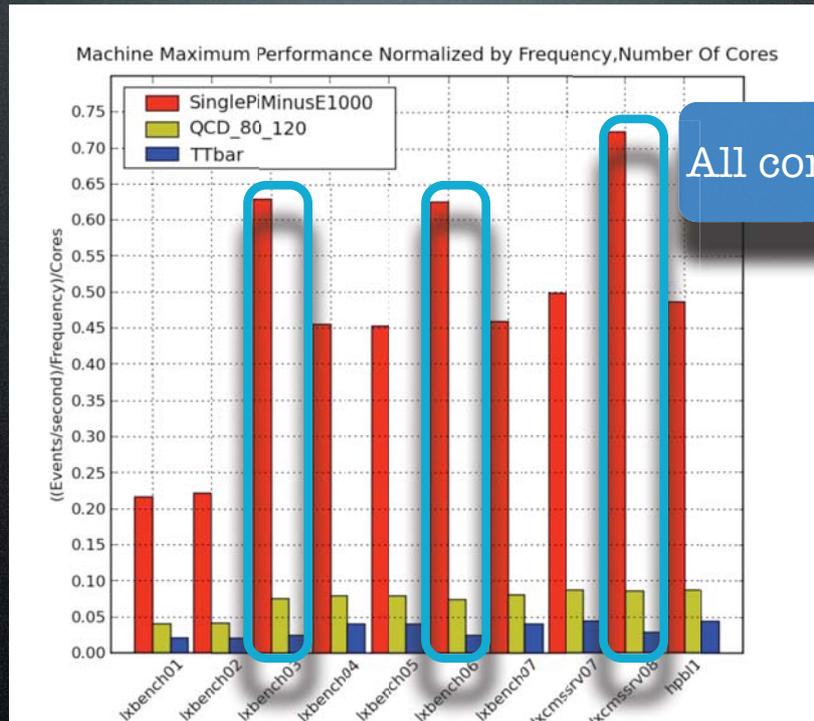


RECO normalized by frequency





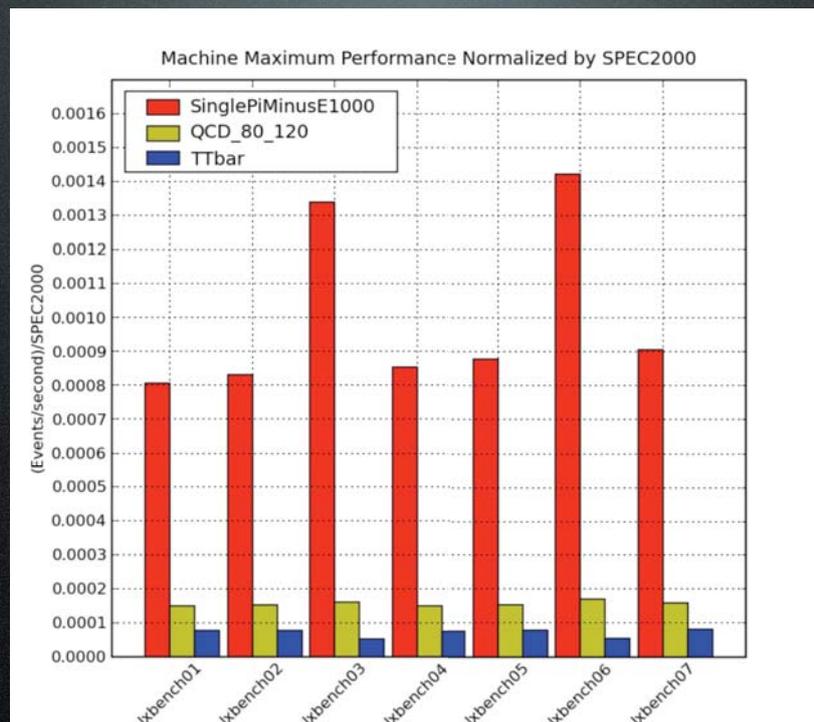
RECO normalized by frequency and # cores



All cores loaded

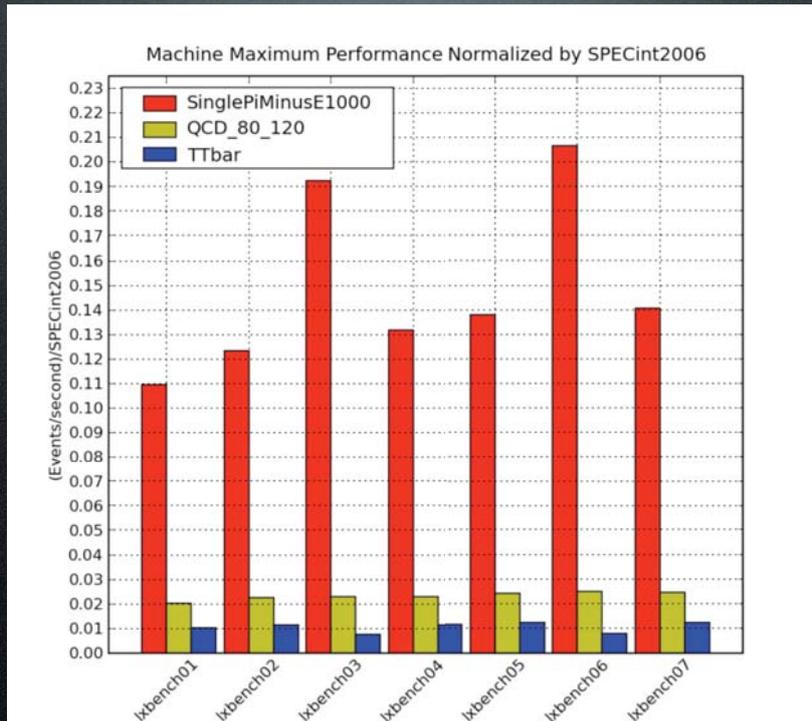


RECO vs SPEC2000

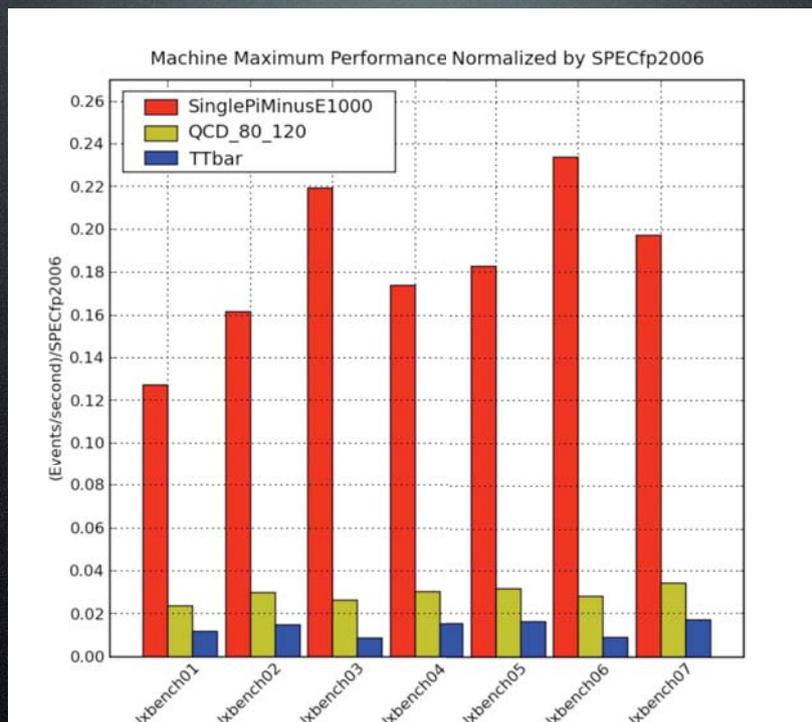




RECO vs SPECint2006

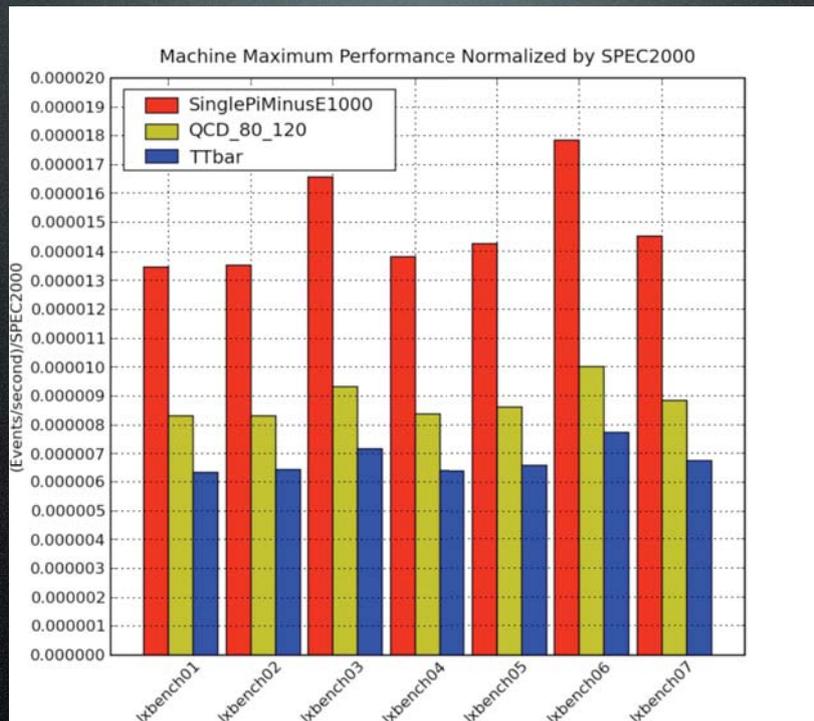


RECO vs SPECfp2006

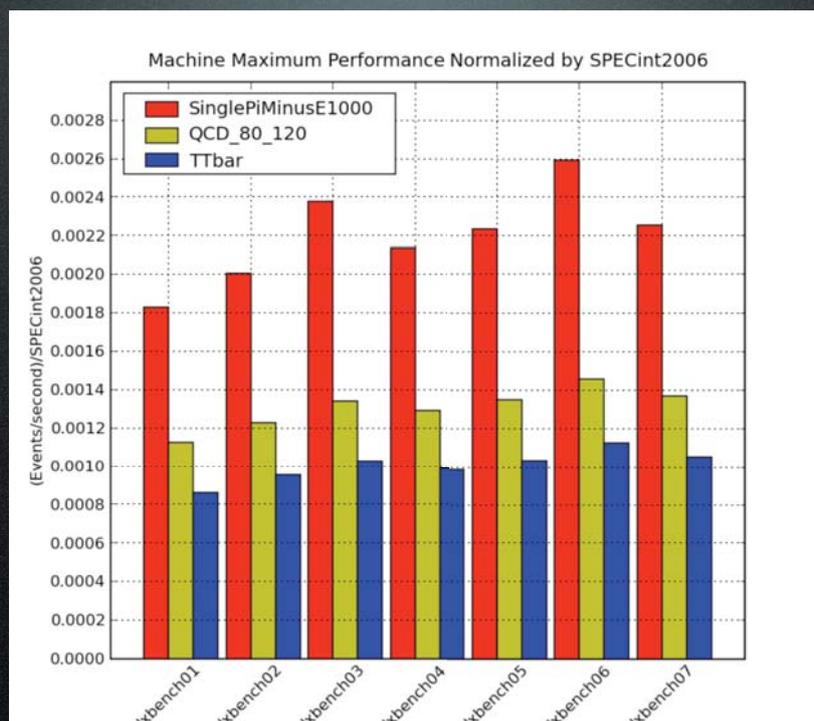




SIM vs SPEC2000

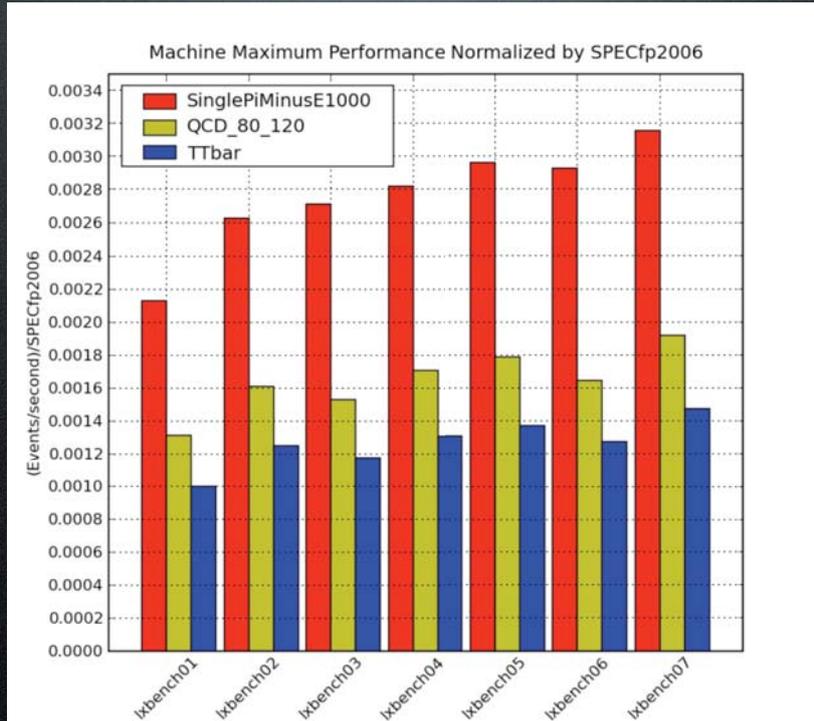


SIM vs SPECint2006

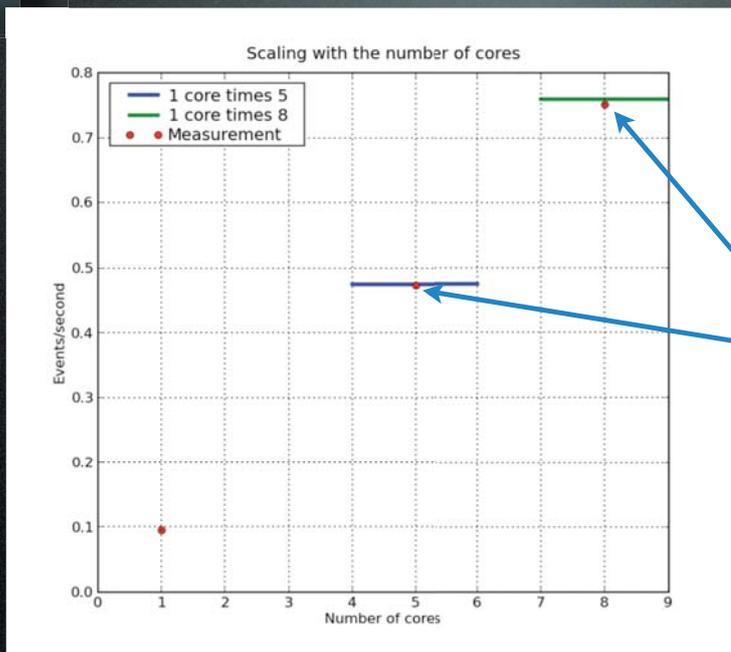




SIM vs SPECfp2006



Scaling with multicores

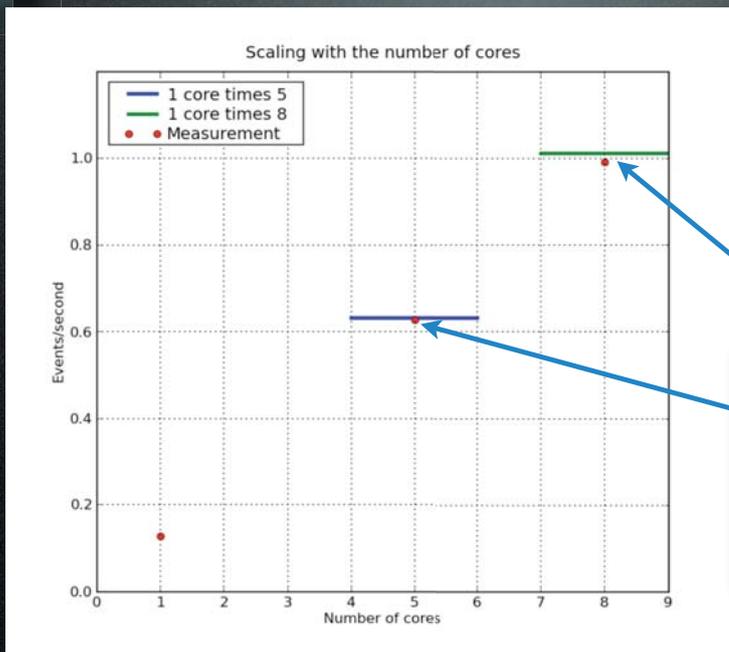


lxbench07
TTbar RECO

Excellent scaling to the last core (cmsScimark2)



Scaling with multicores



hpbl1
TTbar RECO

Excellent
scaling to the
last core
(cmsScimark2)



CMSSW Benchmarking Results



- Observed a different behavior in AMD vs. Intel machines for complex vs. simple events at the RECO step
- Compared CMSSW applications with SPEC benchmarks: differences due to architecture/type of event are larger than the differences between different SPEC benchmarks
- The CMSSW application scales nicely with the current multicore architectures
- A number of open issues from this first experience:
 - Statistical treatment of the data (number of events used, reproducibility, significance of the measurements)
 - Interpretation of the results ("one score" benchmark, weighting of several scores, picking representative candle(s), scores)
 - Graphical/data analysis
- The results of this work, done earlier this year, has inspired the development of a CMSSW benchmark utility that would address these issues and provide the necessary functionality to be used by Tier centers in assessing the CPU performance of machines



CMSSW benchmarking tools



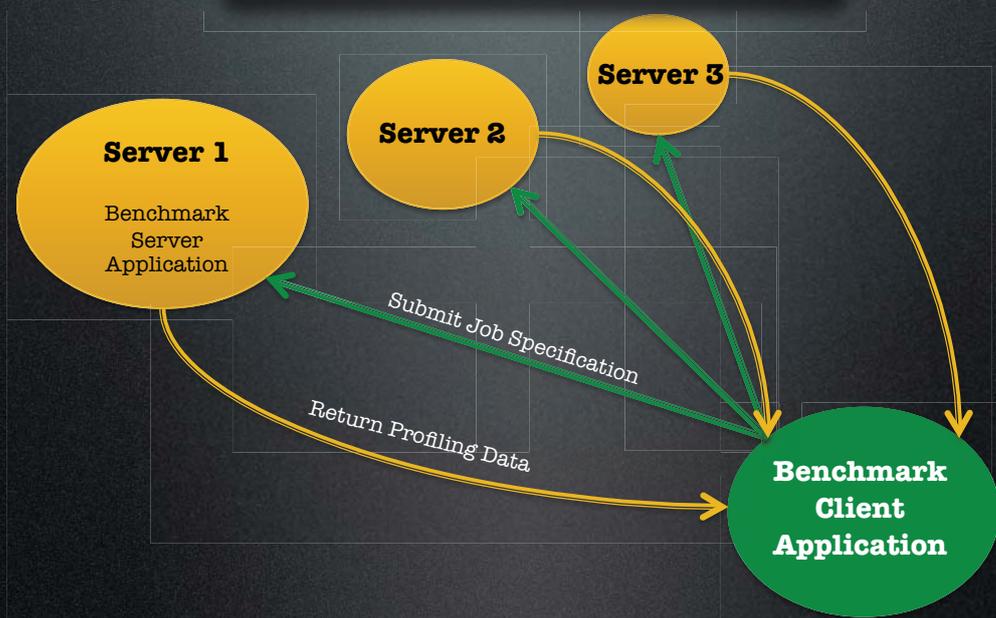
- The idea is to have included in the CMSSW release a “suite” of benchmarking tools
- The basic functionality would be to run a special command of the Performance Suite, then harvest the log files for the CPU timing information
- The command above would return a score (maybe with its composition, in terms of multiple candles, or multiple processing steps, with relative weights, so that a Tier1 vs Tier2 could decide which score is most relevant for their use scenario)
- Since this kind of benchmarking usually involves more than one machine (since one is interested in comparing them), we thought of implementing server-client communication via XML-RPC in Python.
- Finally the data from multiple machines would have to be harvested from the logfiles, analyzed and reported in plots and scores



Client/Server Benchmarking



Using XML-RPC Transport





CMSSW benchmarking tool



- Then all any user would have to do, is to follow the APT based installation procedure for CMSSW on each machine intended to be benchmarked.
- On each machine launch an XMLRPC server
- On one machine launch the XMLRPC client, that ships the commands for the Performance Suite to the servers
- Wait for the Suite to finish on each machine (could do multicore running, or multiple runs on the same core(s)) and to report all results in term of a pickled file that contains dictionary data structure
- On this data structure run by default a basic analysis that would produce one score, a table with all results and the relevant comparison plots.
- Up to now the client/server functionality is under testing, the data analysis and score composition requires still some work



Conclusions



- The Performance Suite is an integral part of the CMSSW Release Validation process, providing developers with in-depth performance results and regression information within a 24 hrs cycle
- The Performance Suite can handle a lot of profile data and summarize it making it available and usable by developers,
- Besides its default behavior, it can be tweaked for more statistics, changing tests, single candles, new versions of external tools, new profiling tools, it can also be used as the core of the CMSSW benchmarking utility
- The framework is ready to handle the first collision data . A few improvements are in the works and the performance suite is used in testing and guiding optimization
- A command-line and XML-RPC based client/server CMSSW architecture benchmarking suite is being implemented and will soon ship with the release



Back-up



The CMS Experiment

38 Countries, 183 Institutes, 3000 scientists and engineers (including 400 students)

June 2008

TRIGGER, DATA ACQUISITION & OFFLINE COMPUTING
Austria, Brazil, CERN, Finland, France, Greece, Hungary, Ireland, Italy, Korea, Lithuania, New Zealand, Poland, Portugal, Switzerland, UK, USA

TRACKER
Austria, Belgium, CERN, Finland, France, Germany, Italy, Japan*, Mexico, New Zealand, Switzerland, UK, USA

CRYSTAL ECAL
Belarus, CERN, China, Croatia, Cyprus, France, Italy, Japan*, Portugal, Russia, Serbia, Switzerland, UK, USA

PRESHOWER
Armenia, CERN, Greece, India, Russia, Taiwan

RETURN YOKE
Barrel: Estonia, Germany, Greece, Russia
Endcap: Japan*, USA

SUPERCONDUCTING MAGNET
All countries in CMS contribute to Magnet financing in particular: Finland, France, Italy, Japan*, Korea, Switzerland, USA

FEET
Pakistan, China

FORWARD CALORIMETER
Hungary, Iran, Russia, Turkey, USA

HCAL
Barrel: Bulgaria, India, Spain*, USA
Endcap: Belarus, Bulgaria, Georgia, Russia, Ukraine, Uzbekistan
HO: India

MUON CHAMBERS
Barrel: Austria, Bulgaria, CERN, China, Germany, Hungary, Italy, Spain
Endcap: Belarus, Bulgaria, China, Colombia, Korea, Pakistan, Russia, USA

* Only through industrial contracts

Total weight : 12500 T
Overall diameter : 15.0 m
Overall length : 21.5 m
Magnetic field : 4 Tesla



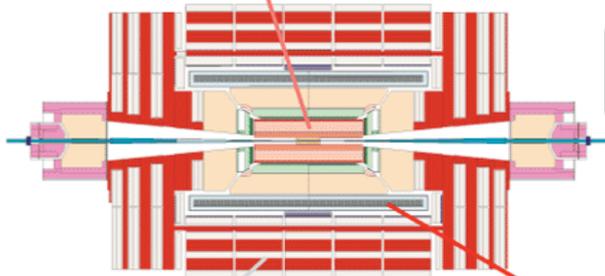
Germany in CMS



6 Institutes
216 Collaborators
As of September 2008

TRACKER
Detector, electronics, mechanics,
support and monitoring
IEH, IEKP, PI3, PI3

HIGH LEVEL TRIGGER
DESY



OFFLINE COMPUTING
IEH, IEKP, PI1, PI2, PI3

MUON Barrel Drifttubes
Construction and assembly,
installation, slow control system,
gas and cooling system
PI2

MAGNET
Procurement
IEH, IEKP, PI1, PI2, PI3

KEY: IEH - Institut fuer Experimentalphysik, Hamburg; IEKP - Institut fur Experimentelle Kernphysik, Karlsruhe;
PI1 - I. Physikalisches Institut, Aachen; PI2 - III. Physikalisches Insitut A, Aachen; PI3 - III. Physikalisches Institut B, Aachen;
DESY - Deutsches Elektronen-Synchrotron, Hamburg



3 slides each (CPU, Size, Memory), maybe 3,24 with an example of use of the tools to investigate the matter/improve (tomorrow submit 2_1_4 to compare with 2_1_10 for SIM and RECO)
CPU:
1-TimevsEvt**+histo for MinBias SIM? (mention possibility to skip first event?) with and without regression
In the same TimeReport breakdown by module (Tabella con tutte le candele)
lgProffPerf con Regression
2-Size Tabella con tutte le sizes to see progressive increase, EdmSize with Regression
3-Memory
SimpleMemcheck:
(Pile-up?), example of bumps and coming down, comparison between successive steps, or Pile-up
lgProfMem con Regression

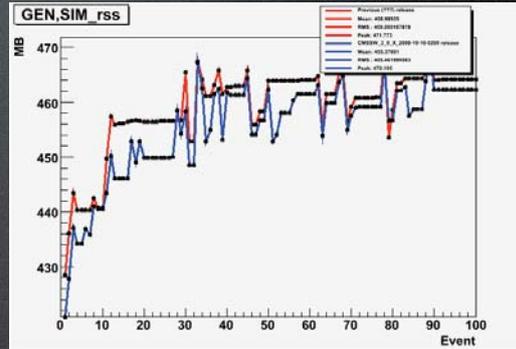
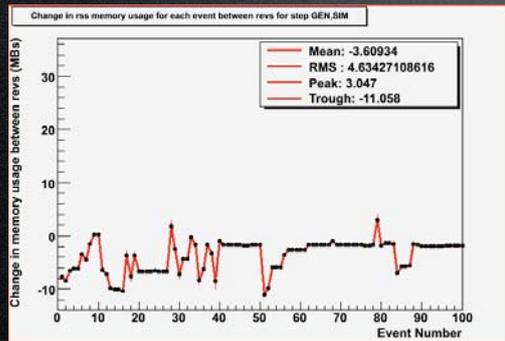
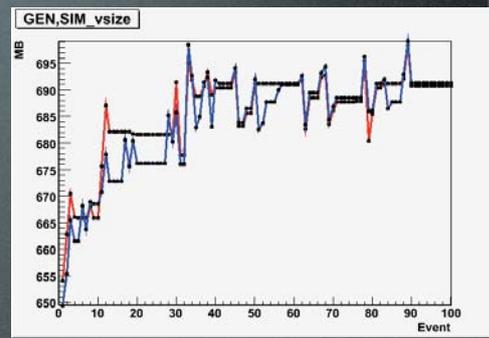
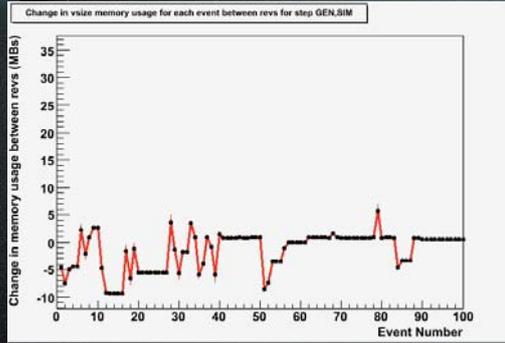
Points to make:
1-Automated procedure to measure and monitor performance
2-Part of Release Validation QA
3-Quick response time/enough detailed information
4-Wide coverage of code, candles, processing steps
5-Robust (based on RelVal samples), options added
6-USed also for other tests
7-Re-used for benchmarking purposes
8-Can include robustness and reproducibility tests
9-All above... but we need to present

Add Graphics with the structure of the Suite itself?
OR Graphic with the issues of power saving paranoia (benchmark before and after on the core), stuff on the other cores

Key point in Results:
1-CPU Time: a table and a couple of plots
2-File Size a regression plot of EdmEventSize or Is table
3-Memory a plot of SimpleMemoryCheck/Total memory estimate from MemTotal (?)



Memory Performance



Gabriele Benelli, CERN

49

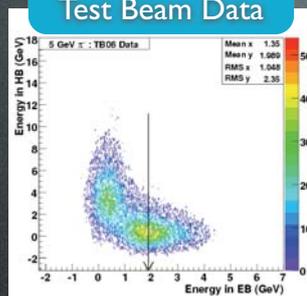
DESY Zeuthen, October 21st 2008



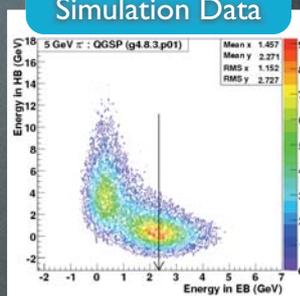
Simulation Test Beam Validation



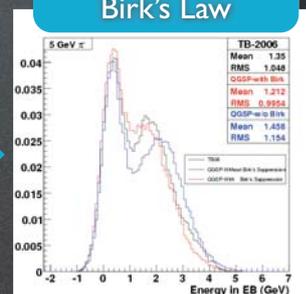
Test Beam Data



Simulation Data



Birk's Law



Scintillator Saturation Effects

Gabriele Benelli, CERN

50

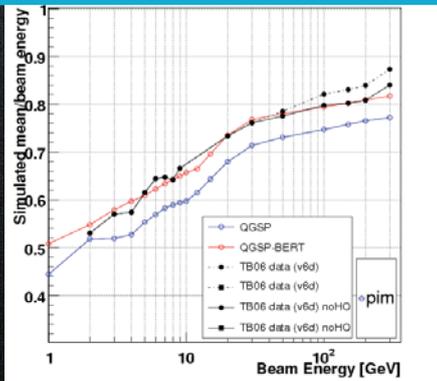
DESY Zeuthen, October 21st 2008



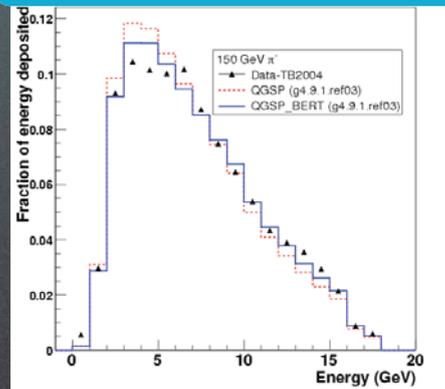
Simulation Test Beam Validation



Mean Energy Response ECAL+HCAL with pion beam



Longitudinal Shower Profile HCAL



Simulation Test Beam Validation



GEANT4 Physics Lists	CPU Time (%)		Event Size (%)	
	MinBias	TTbar	MinBias	TTbar
QGSP_EMV	100	100	100	100
QGSP	116	120	101	103
QGSP_BERT_EMV	141	146	152	177
QGSP_BERT	158	169	152	172



A little bit of history



- Until the last release cycle (21X) CMSSW used a special configuration language to configure the (one and only) cmsRun executable
- The full transition to Python happened with 210 and came with a major improvement in maintainability: cmsDriver.py
- This script is a command-line utility that prepares a full python configuration file, based on a few command-line options, and launches cmsRun on it.
- This is highly configurable and covers most use cases



Details on Python tools



- Graphics to show the cmsDriver, cmsRelvalreport, cmsPerfSuite, cmsRelRegress, cmsPerfRegress, cmsRelvalreportInput ...
- Examples of uses of one and the other (customize fragment and its effect in the log, parsers and plot drawing, regression and publishing on a webserver)
- Issues about the machines used to do the measurements (for CPU only, but also for memory in case of ununderstood crashes)
- Uses of the suite for other aims:
 - tests on external packages (G4 reproducibility, its performance, robustness),
 - use as a machine benchmarking building block (moving to the second part of the talk)



Fast!

