No escape We have to collaborate !

X. Vigouroux – Director of CEPP

09/23/2015



1 Your business technologists. Powering progress









Atos in the top500







Rank 1 - evolution



6 | 9/23/2015| | © For internal use GBU | Big Data and Security | HPC

Rank 350- evolution



7 | 9/23/2015| | © For internal use GBU | Big Data and Security | HPC

[freq] x [#flop/cycle] x [#cores]

CPU frequency

mean in top500



Frequency and power CONSUMPTION



IVB Voltage-Frequency Shmoo plot



Power vs. Frequency, 1.26 V, variable temperature







[freq] x [#flop/cycle] x [#cores]



Exaflops/s target: 20MW and 50000Mflops/W



Exaflops/s target (list #54) : 50000Mflops/W

Architectures become more complex





A very quick summary



Article suggestion: <u>http://herbsutter.com/welcome-to-the-jungle/</u> Nutrient broth : bouillon de culture

Some people think forstanmers exclaim When doing experiments. But they're way more likely to say... Bollocks!) F*ck oh Shit Stupid piece--of- crap machine! I hate Arse! 44 Science twisteddoodles.com

Impact on Constructors





Adiabatic Cooling





Cooling capacity	108 kW
Power consumption	2,8 kW
Water used (estim. h/year)	472
Water used (estim. m ³)	53
PUE	1.18





10 18 Bull exascale program



- Open platform designed for the long-term
 - Integrate current and future technologies
 - Multiple types of compute nodes
 - (CPUs, GPUs, HPC accelerators)
- Scales up to tens of thousands of nodes
 - Large building blocks to facilitate scaling
 - Embedding the fastest interconnects
- Ultra-energy efficient
 - Enhanced DLC



Integrated

Modulable

coprocessors

Interconnects



High Definition Energy Efficiency Monitoring



Motivations

COMPUTE POWER CONSUMPTION PERFORMANCE (W.h) (flops/s) • Applications are written for • Becomes an increasing budget compute performance Will be charged on end users Is a strict constraint for large Tools are seeking performance systems issues (host spot, idle time) • Unit: flops/Watt • Unit: core.hour

HDEEM MERGES THE BOTH WORLDS Precise metrics in performance AND power consomption



Good optimization needs Goode metrics



...Good metrics

Requires

Accurate Swift Non-intrusive Probes



Sample precision



Measure granularity < 200mW - error < 2%



Sample rate



1000 samples/s at blade level



See the difference !





Tool stack







 \mid 9/23/2015 \mid \mid \odot For internal use GBU \mid Big Data and Security \mid HPC

Interesting paper: "Sesames: a smart-Grid Based Framework for consuming less and better in Extrem-Scale Infrastructures" – M. El Mehdi Diouri, O. Glück, L. Lefèvre



Impact on Constructors





center for excellence in parallel programming

Users questions

- How can I get the best from my computing resource?
 - Do I use the best algorithms, the best solvers, ...?
 - What is the best hardware for my case?
 - Am I using the best runtime conditions?
- How will I get the best from my computing resource?
 - What modification should I anticipate?
 - What are the technology trends and impacts?
 - Will I be able to reach my future goals?
- How I get the best from my computing resource?
 - Discussion with experts, Continuous interactions, Tailored trainings
 - "Brown bag" sessions, Access to computing resources, Creation of funded specific projects



Preparing applications for exascale

- Exascale is disruptive
 - Hardware architecture
 - Resource and data concurrency
 - Energy efficiency



- The Center for excellence in parallel programming reengineers existing applications to make the most of exascale
 - Optimization on extreme scale architectures
 - Resilience
 - Management



Project and Service Map





PhD in Scientific domains

- ► CEPP:
 - fund PhD theses
 - bring the performance expertise to researchers
- Case: "inverse docking"
 - find active sites on thousands of molecules for a given ligand







[Q4/14] Gromacs on K80

► Goal:

- Demonstrate GROMACS improvement on K80
- Demonstrate CEPP expertise to get this performance





[Q4/14] NEMO modernization

► Goal:

modify NEMO to increase performance on Xeon



Make NEMO Xeon Phi ready





H2020

► ESCAPE

Numerical Weather production "dwarfs".

Peter Bauer (ECMWF), Météo France, Météo Swiss, CSC, Cray, Nvidia, Intel

EsiWace Center of Excellence in NWP with DKRZ

 Jetpass with Meteo France. Student hosting (6PM)





Future of CEPP

- Today
 - in Grenoble
 - as many topic as possible
- ► Tomorrow
 - Several CEPP sites
 - On user sites
 - collaborating on specific topics

- Benefit for Atos:
 - relying on local skills
 - starting local projects
- Benefit for « local site »:
 - early access to technologies
 - relying on Grenoble skills

- For instance
 - Large memory
 - Oil and Gas
 - Life Science
 - Chemistry



Questions

Atos, the Atos logo, Atos Consulting, Atos Worldgrid, Worldline, BlueKiwi, Canopy the Open Cloud Company, Yunano, Zero Email, Zero Email Certified and The Zero Email Company are registered trademarks of Atos. July 2014. © 2014 Atos. Confidential information owned by Atos, to be used by the recipient only. This document, or any part of it, may not be reproduced, copied, circulated and/or distributed nor quoted without prior written approval from Atos.



Your business technologists. Powering progress