HPC solutions for Scientific Simulations

Jean-Pierre Panziera –Nov 3, 2010

Bull
Architect of an Open World™

Bull's HPC solutions for Science Simulations

- Bull today, the company, the HPC products
- Large nodes Petaflop
- Toward Exascale: the accelerator's [r]evolution
- Green Power

Bull Group

- **A growing and profitable company**
- **A solid customer base**
  - Public sector, Europe
- **Bull, Architect of an Open World™**
  - Our motto, our heritage, our culture
- **Group commitment to become a leading player in Extreme Computing in Europe**
  - The largest HPC R&D effort in Europe
  - 500 Extreme Computing experts - the largest pool in Europe

REVENUE BREAKDOWN

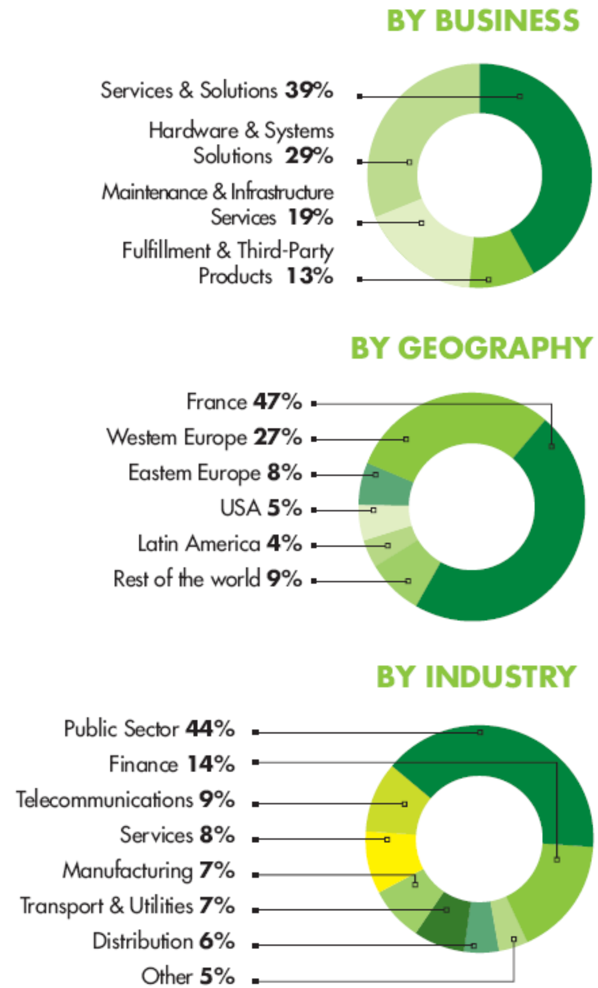

Target markets for Bull in Extreme Computing




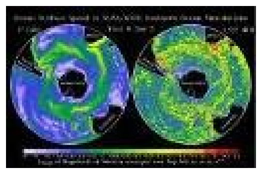

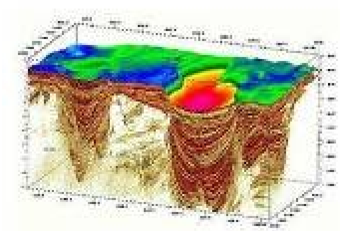



Production HPC

- **Government**
  - Defense
  - Economic Intelligence
  - National Research Centers
  - Weather prediction
  - Climate research, modeling and change
  - Ocean circulation
- **Oil & Gas**
  - Seismic: Imaging, 3D interpretation, Prestack data analysis
  - Reservoir modeling & simulation
  - Geophysics sites Data Center
  - Refinery Data Center
- **Automotive & Aerospace**
  - CAE: Fluid dynamics, Crash simulation
  - EDA: Mechatronics, Simulation & Verification
- **Finance**
  - Derivatives Pricing
  - Risk Analysis
  - Portfolio Optimization

Worldwide references in a variety of sectors

Educ/Research

- ICHEC Irish Centre for High-End Computing
- GENCI
- Universiteit Gent
- www.uvsq.fr
- Heinrich Heine Universität Düsseldorf
- CSIC Consejo Superior de Investigaciones Científicas
- Principe Felipe Centro de Investigacion
- Cardiff University Prifysgol Caerdydd
- Jülich Forschungszentrum
- Universiteit Utrecht
- IRISA
- ENSAM
- RRZN
- Université de Reims Champagne-Ardenne
- CINES
- CCRT

Industry

- Renault
- Pininfarina
- PSA Peugeot Citroën
- Valeo
- Alcan
- BMW
- STMicroelectronics
- Saint-Gobain
- Porsche
- Bosch

Aerospace/Defence

- AWE
- CEA
- DGA
- Dassault Aviation
- Thales
- MBDA Missile Systems
- Alenia Aeronautica
- ONERA
- Safran
- EADS Nucletudes
- Airbus

Other

- atass sports
- EDF
- IFP
- Alstom
- Total
- BR Petrobras
- Meteo France
- kraftway
- Miracle Machines

… and many others

science + computing
ServiWare
Bull

bullx hardware for peta-scalability

bullx supercomputer suite

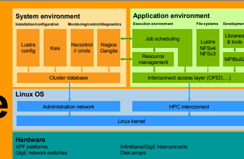

bullx rack-mounted servers

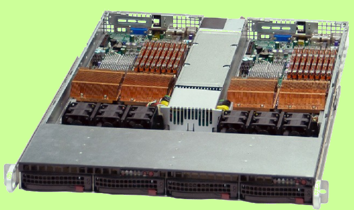

bullx blade system

bullx supernodes

ACCELERATORS

Water cooling

Storage

Architecture

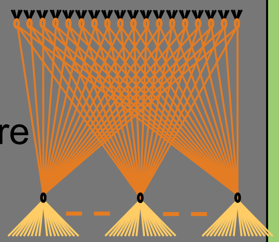

TERA 100

2010: TERA 100

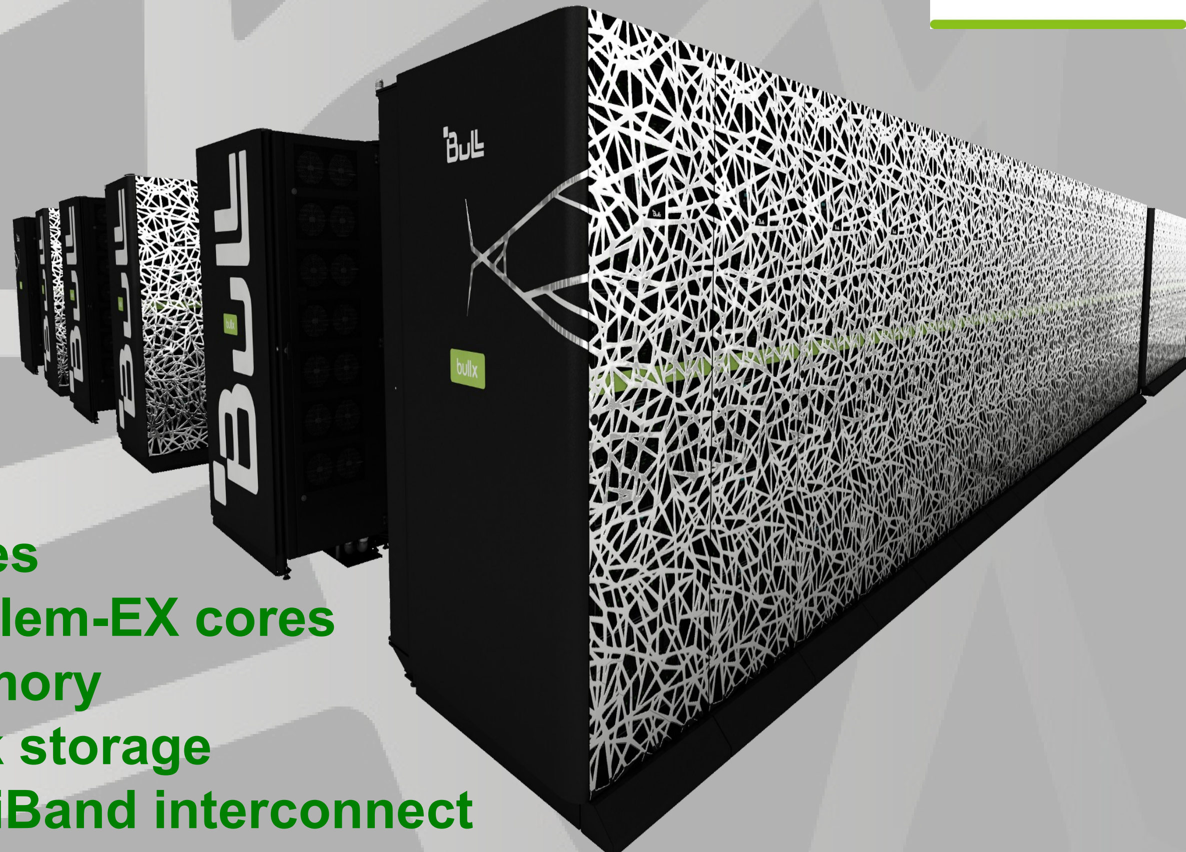

1.25	PFlops
4 300	bullx nodes
140 000	Intel Nehalem-EX cores
300	TB of memory
20	PB of disk storage
	QDR InfiniBand interconnect
500	GB/s bandwidth to the global file system

global memory BW **660 TB/s**

Bull
Architect of an Open World

bullx-S6010 Compute module / node

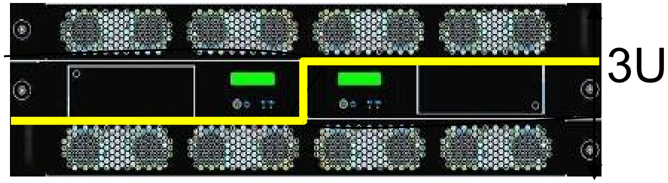

2 Modules (64 cores / 512GB RAM)

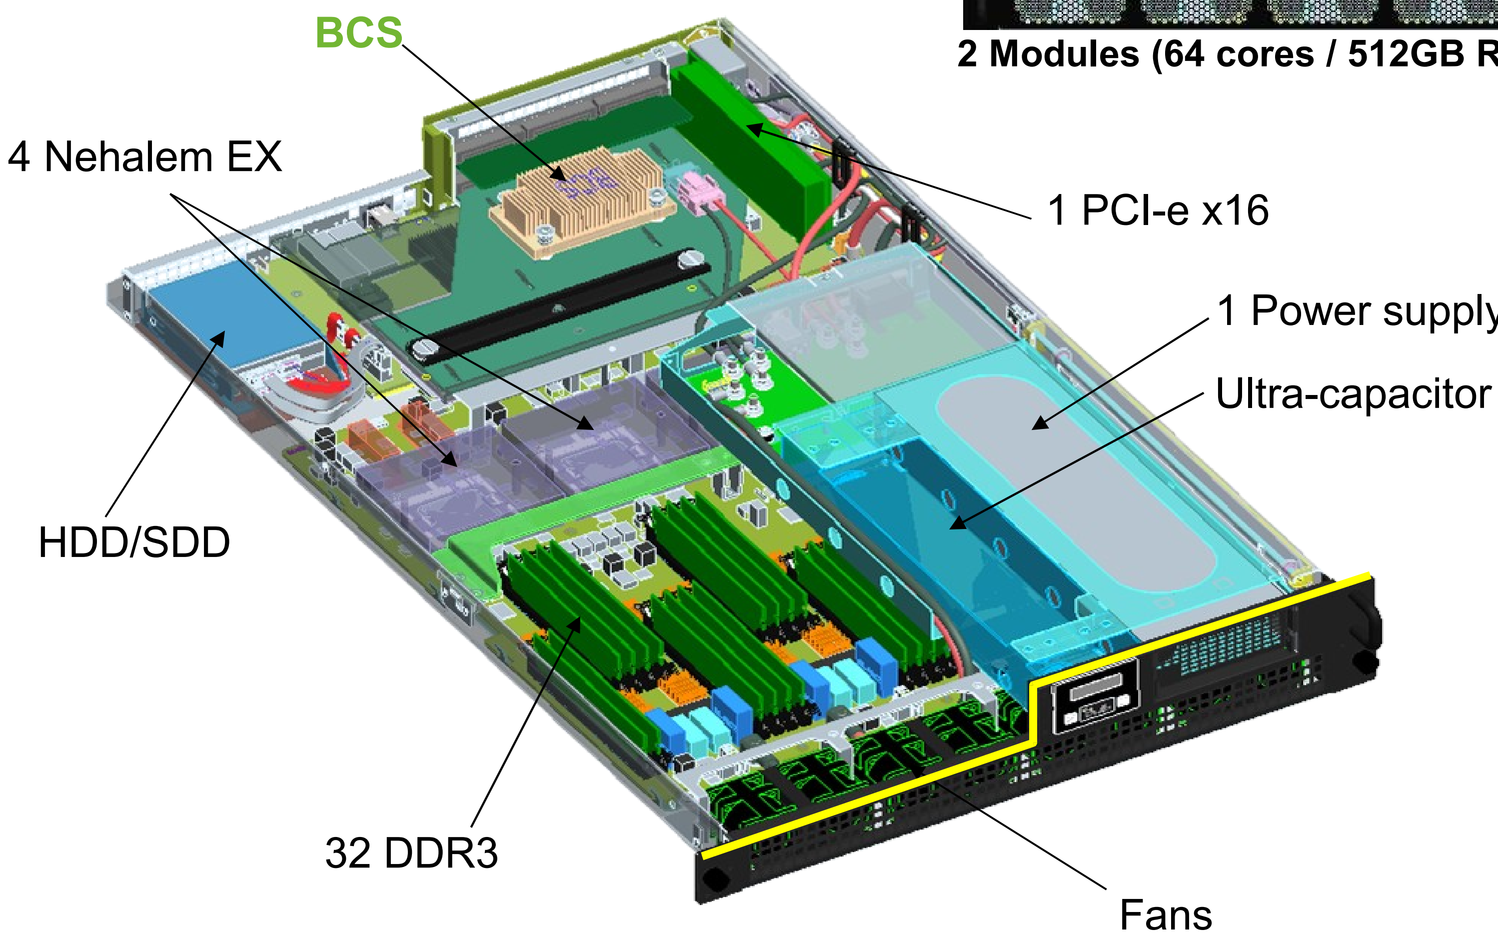

Bullx S6010 – CC-NUMA server

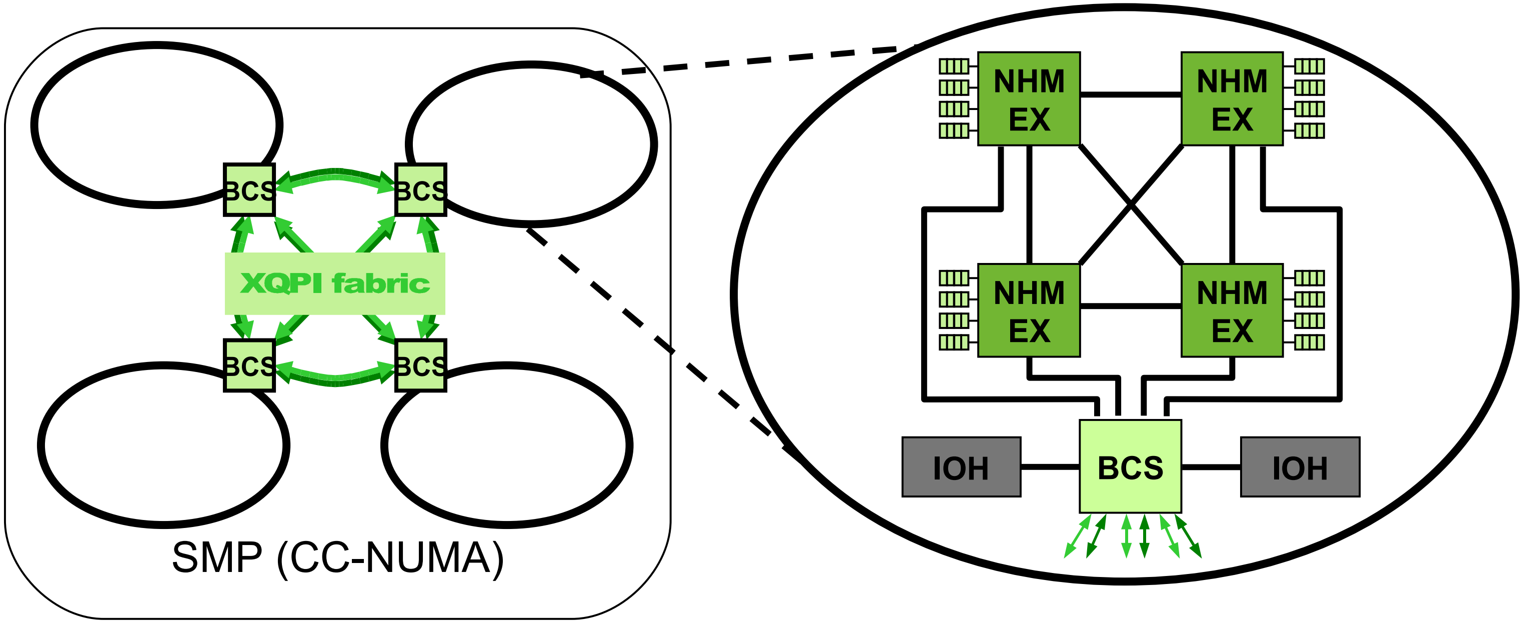

Node maximum configuration :

- 4 modules
- 16 sockets
- 128 cores (Nehalem-EX)
- 128 memory slots (2TB)

Large nodes :

- Large shared memory (pre/post-processing)
- Many more cores (SMP)
- Fewer nodes
- Simpler system administration
- Multi Level Parallelism (MPI/OpenMP)

Continuous growth of HPC systems

- Demand for performance of HPC systems outruns Moore's Law
- CPU performance increases by Moore's Law
- To reach higher system performance, number of CPU's has to increase

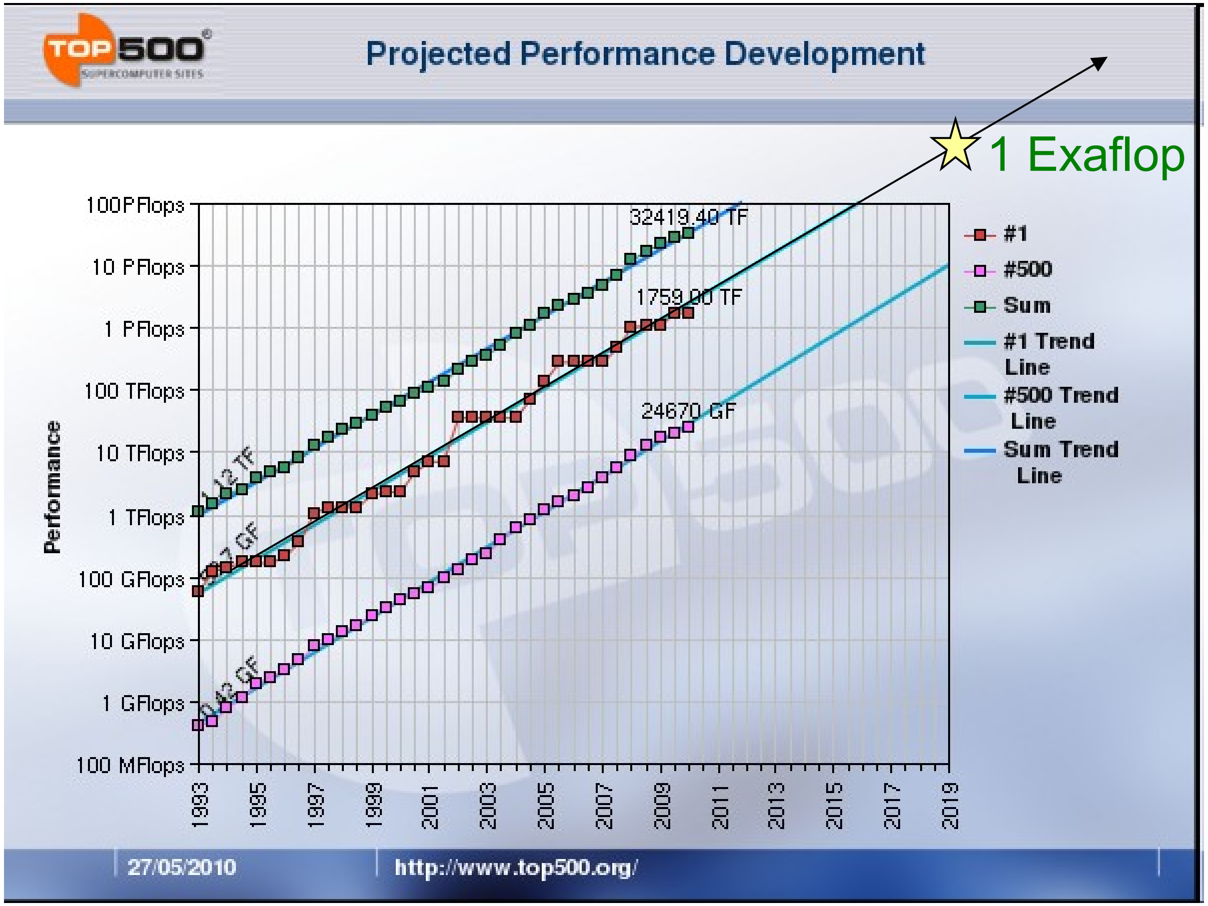

Traditional Sources of Performance Improvement are Flat-Lining

- New Constraints
  - 15 years of *exponential* clock rate growth has ended
- Moore's Law reinterpreted
  - How do we use all of those transistors to keep performance increasing at historical rates?
  - Industry Response : # cores per chip or # flop per cycle double every 18 months or *instead* of clock frequency

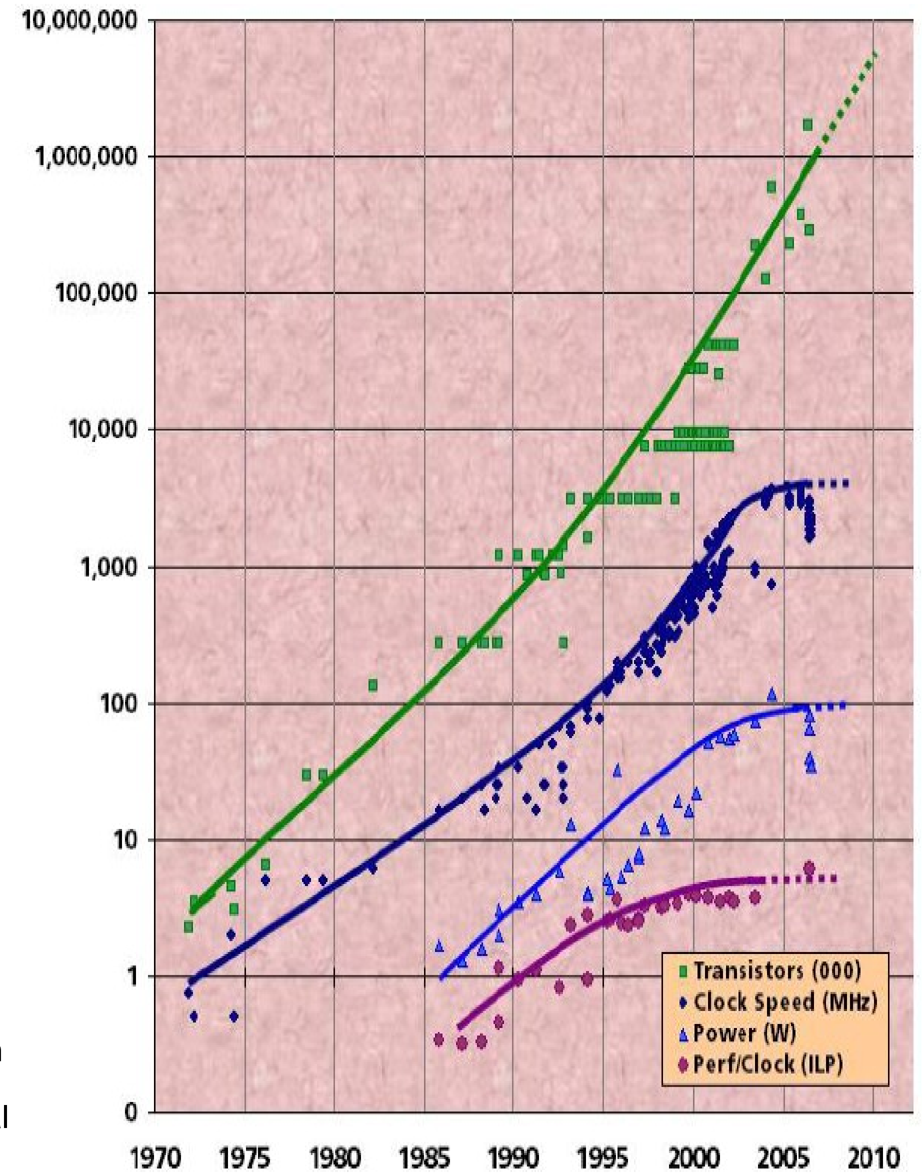

Figure courtesy of Kunle Olukotun, Lance Hammond, Herb Sutter, and Burton Smith

Multi-core CPU architecture evolution

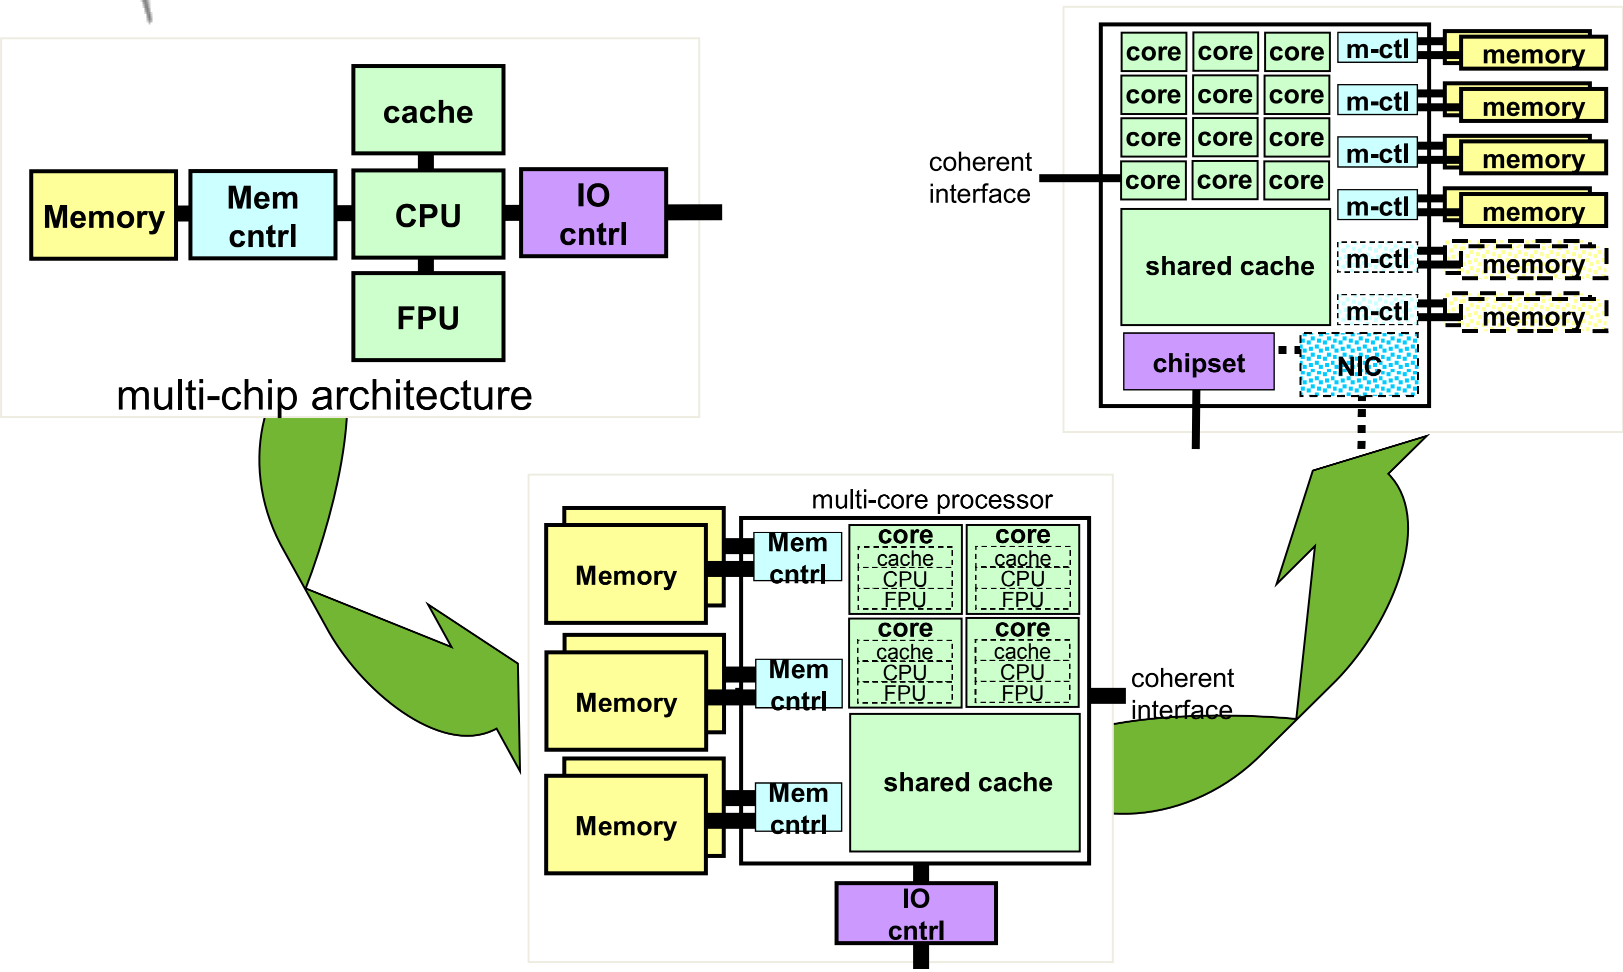

Accelerated HPC has reached the top

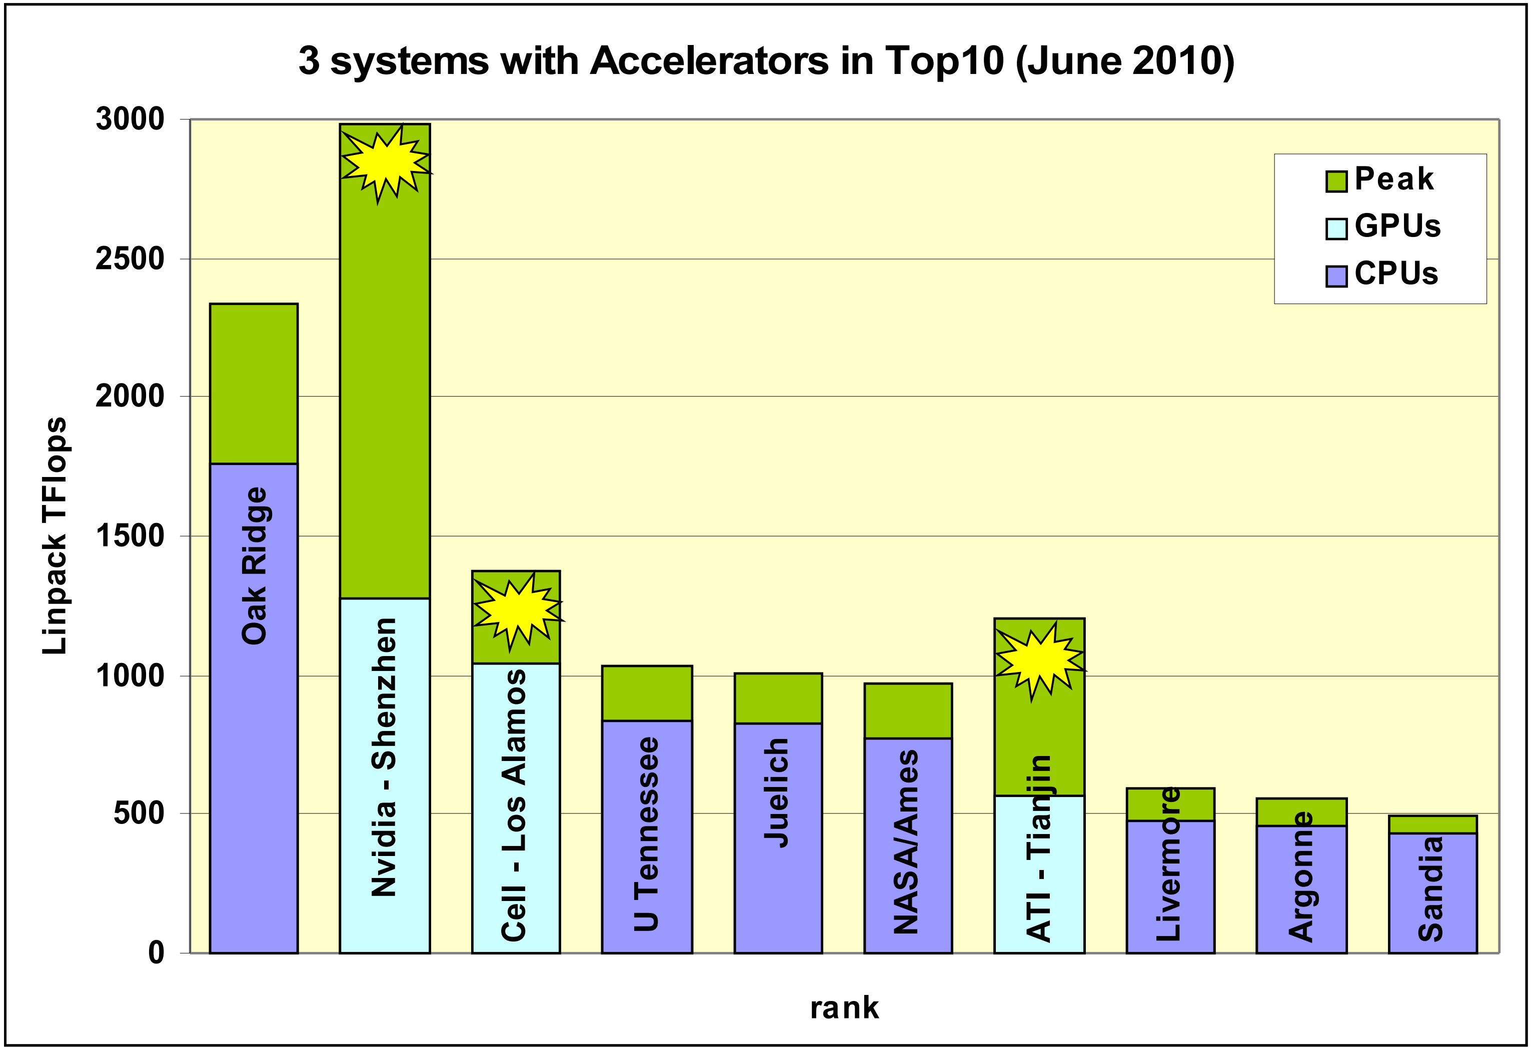

Typical GPU vs bullx blade block diagrams

Best HPC server product or technology

Top 5 new products or technologies to watch

typical GPU system

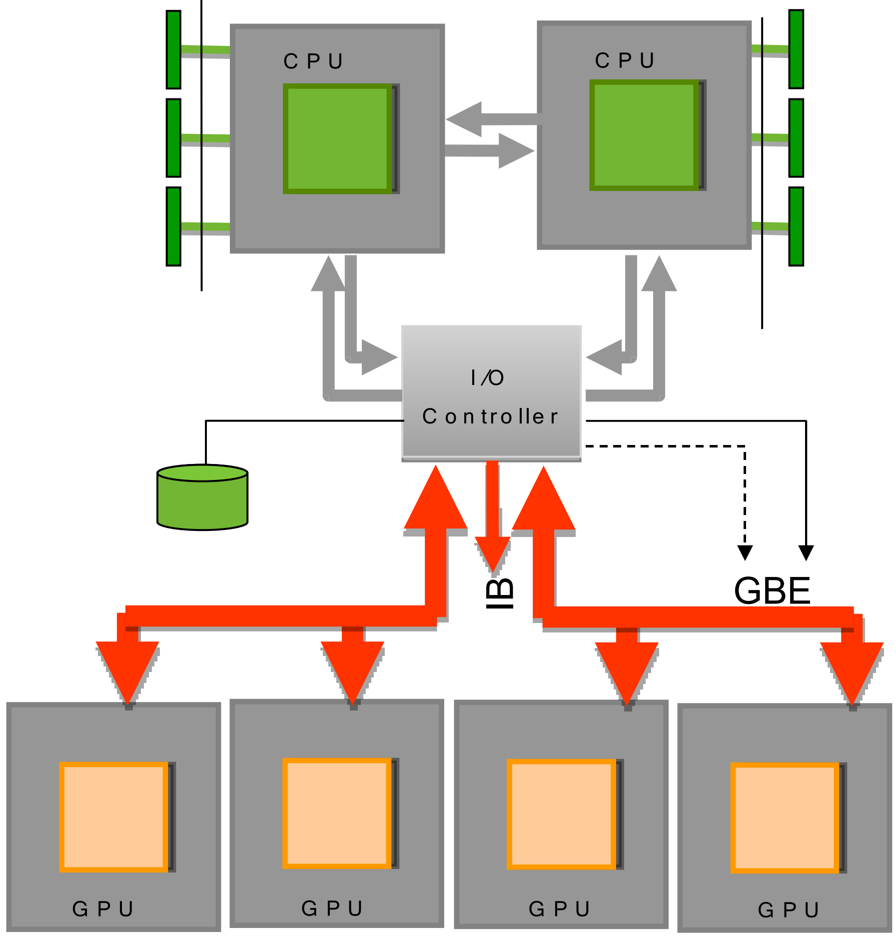

bullx B505 accelerator blade

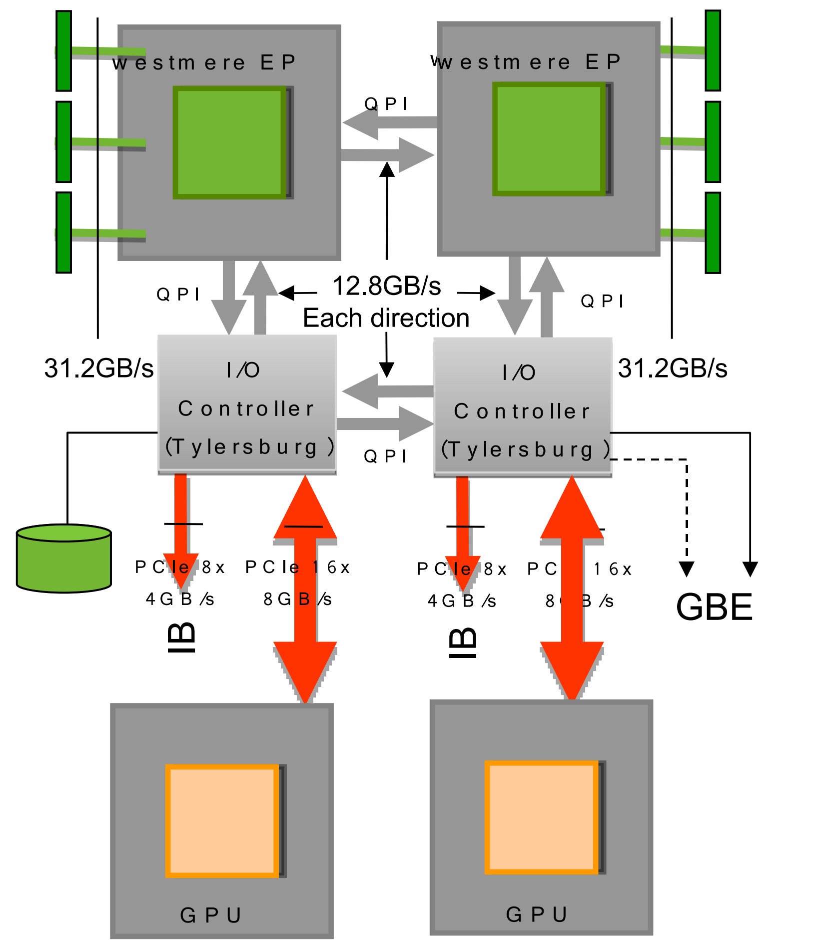

GPU systems bandwidth profiles comparison

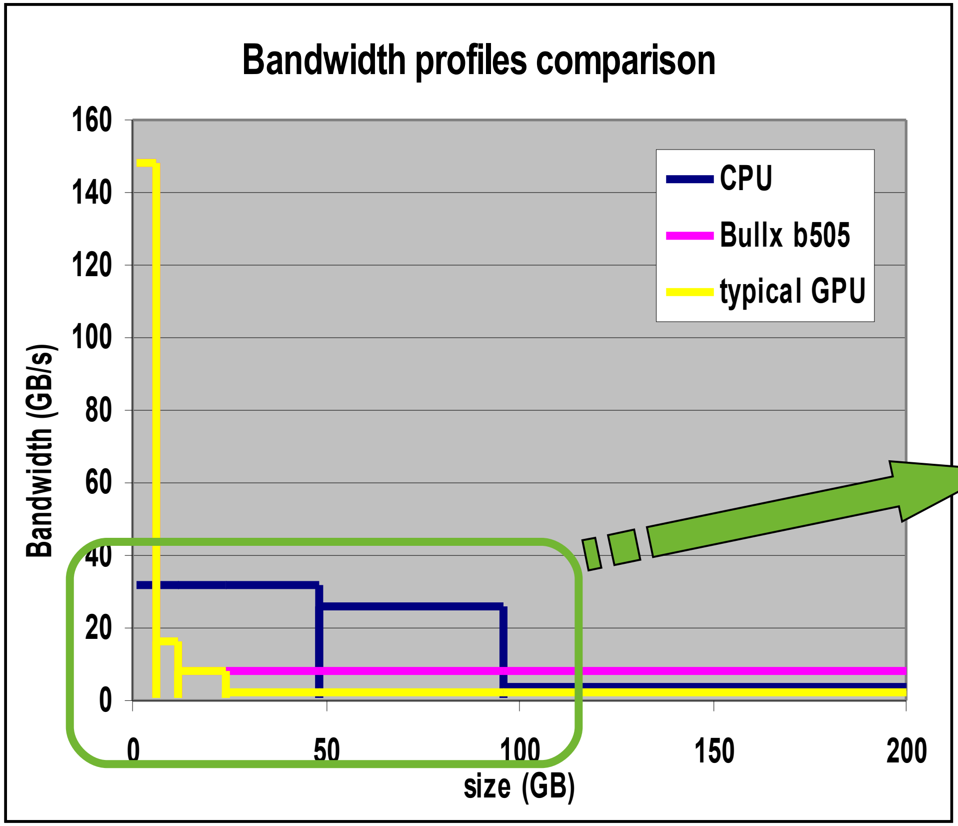

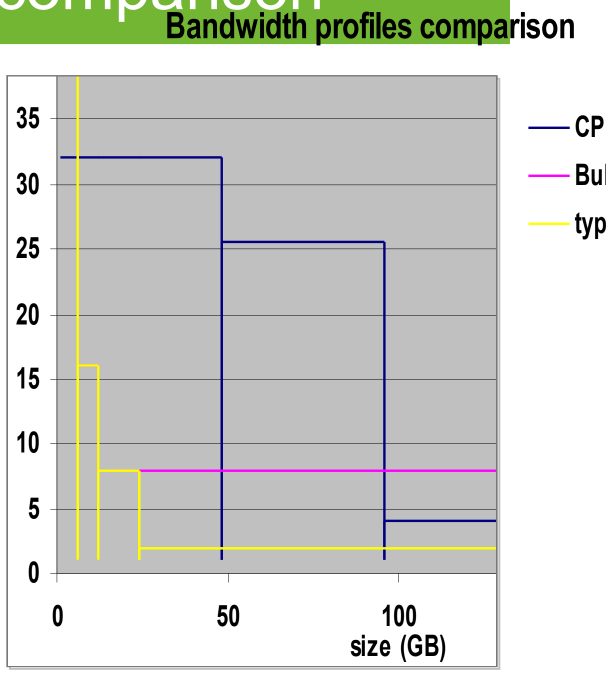

- GPU are powerful compute engines
- A larger bandwidth for communications is required

Programming GPUs

- Until recently GPUs == GOPUs (Graphics Only Processing Unit)
  - very difficult to program: data was mapped as “texture”…OpenGL
  - programming GPUs was the job of computer graphics PhDs
- Cuda was the first environment a (motivated) developer could use
- OpenCL is attempting to standardize the GP-GPU programming … and the CPU as well?
- High level tools are now available (HMPP/CAPS, PGI)

HMPP Example : SGEMM

```
!$HMPP sgemm callsite,args[vin1;vin2;vout;m;n;k2;alpha;beta]
 CALL sgemm_90(n,n,n,alpha,vin1,vin2,beta,vout)
...
!$HMPP sgemm codelet, target=CUDA,
args[vout].io=inout
SUBROUTINE sgemm_90(m,n,k2,alpha,vin1,vin2,beta,vout)
IMPLICIT NONE
Declare all variables ...
!$HMPPCG parallel
DO j=1,n
   !$HMPPCG parallel
   DO i=1,n
      prod = 0.0
      DO k=1,n
         prod = prod + vin1(i,k) * vin2(k,j)
      ENDDO
      vout(i,j) = alpha * prod + beta * vout(i,j) ;
    END DO
END DO
END SUBROUTINE sgemm_90
```

GPU accelerators current applications

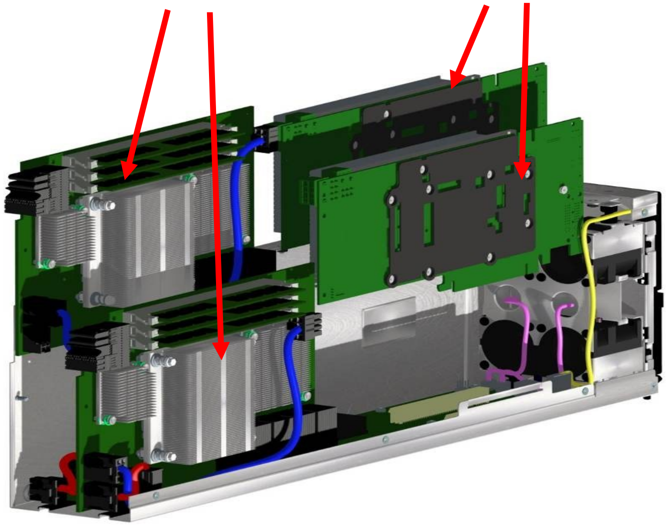

	GPU / CPU ratio
GFlops (DP)	7
Memory BW	4.5
consumption	2
Memory Size	1 / 8

- ***Graphics rendering***
- ***Seismic modeling and imaging***
- ***Molecular Dynamics, Astrophysics***
- ***Financial simulations***
- ***Structure Analysis, Electromagnetism***
- ***Genomics***
- ***Weather/Climate/Oceanography***
- ***… more …***

Successfully Accelerated Applications have:

- **small kernels**
- **moderate size datasets, or good data locality**
- **moderate communications**

Hybrid CPU-GPU architecture [r]evolutions

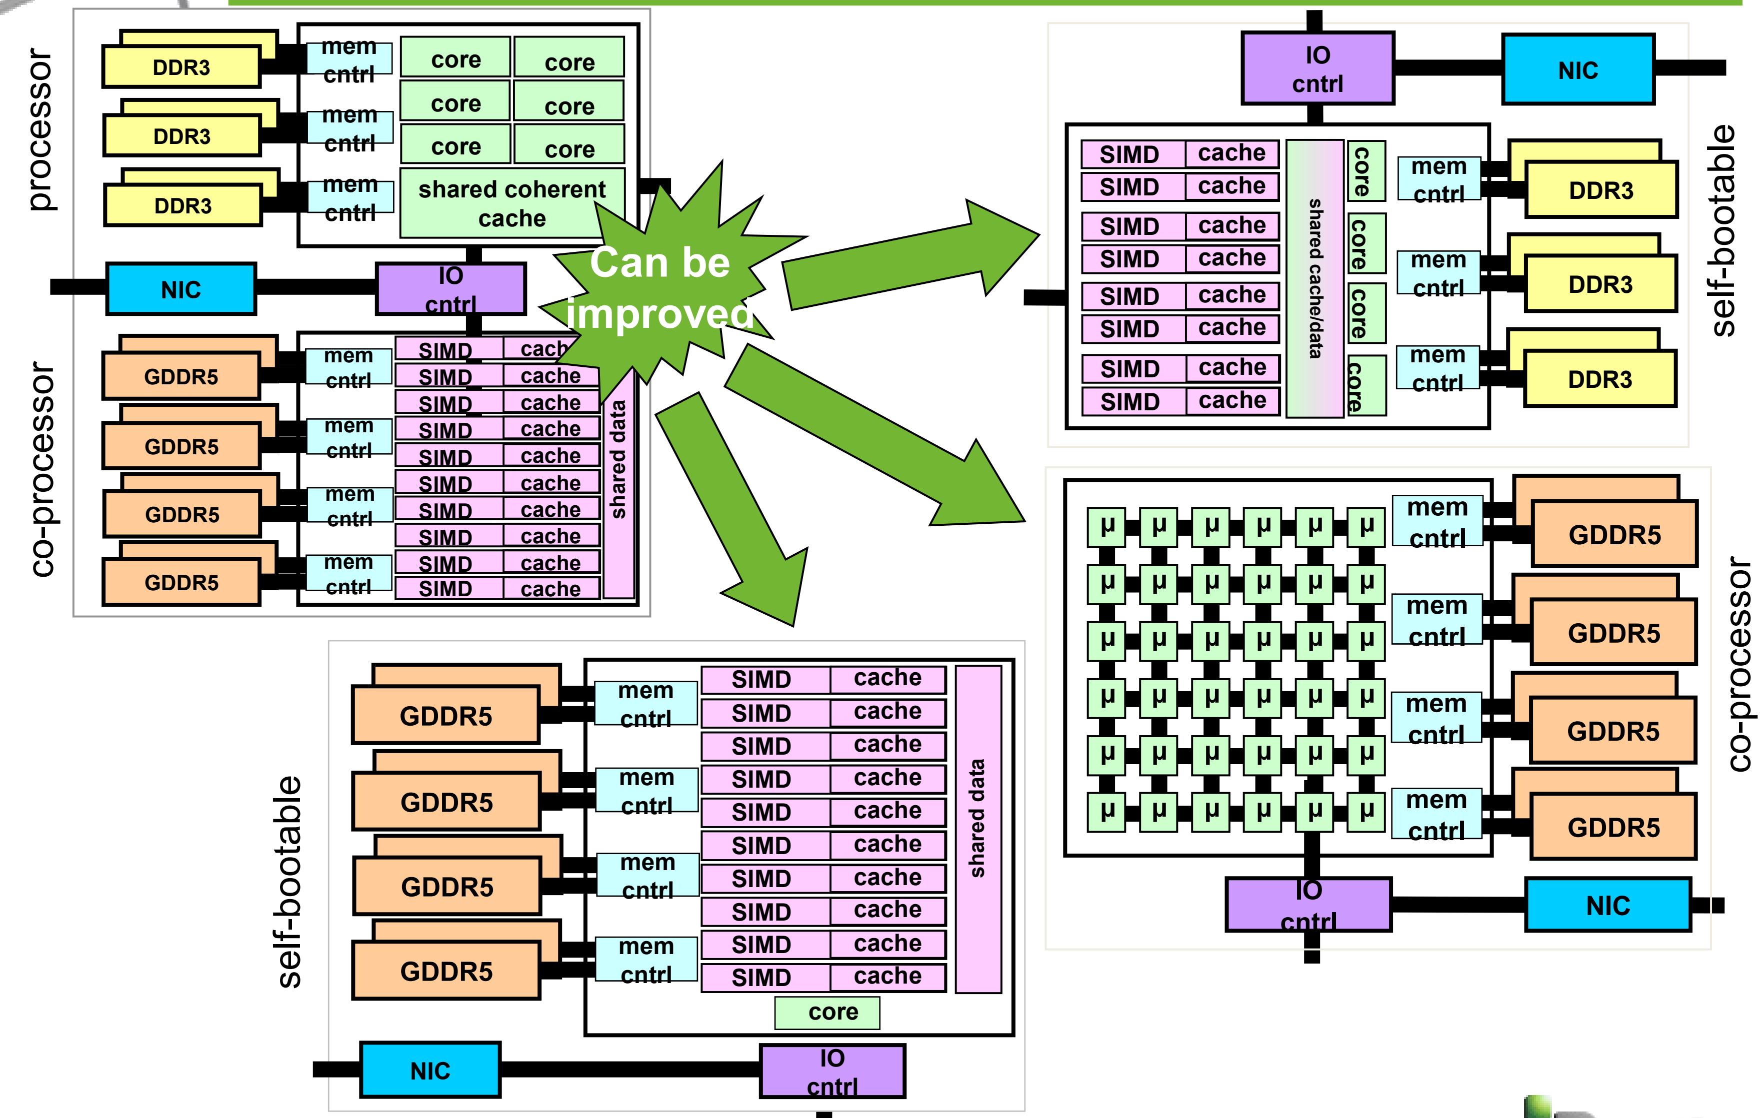

Green Power

<u>Trends:</u>

- HPC Systems are getting ever more powerful
- HPC Systems are consuming more and more
- Energy price is raising fast

- Use more efficient components
- Avoid wasting energy
- Better system integration
- Better PUE
- Improve Total Cost of Ownership (TCO)

Cooling & Power Usage Effectiveness (PUE)

Air-cooled		Water-cooled doors			Direct-Liquid-cooling	
10(-20) kW/rack		**40 kW/rack**			**70 kW/rack**	
Room	**20°C**	**Room**	**20°C**	**27°C**	**Room**	**27°C**
A/C water	**7-12°C**	**Water**	**7-12°C**	**14-19°C**	**Water**	**ambient θ**
PUE	**1.8-1.9**	**PUE**	**1.6-1.7**	**1.4-1.5**	**PUE**	**1.1-1.2**

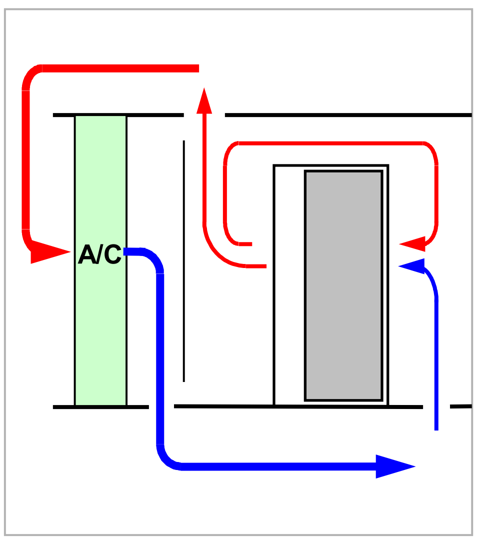

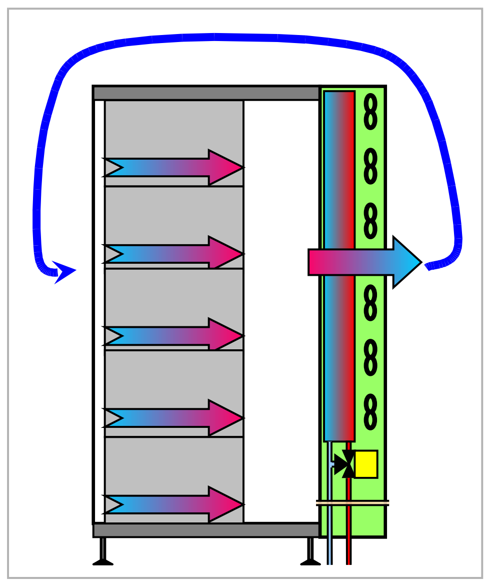

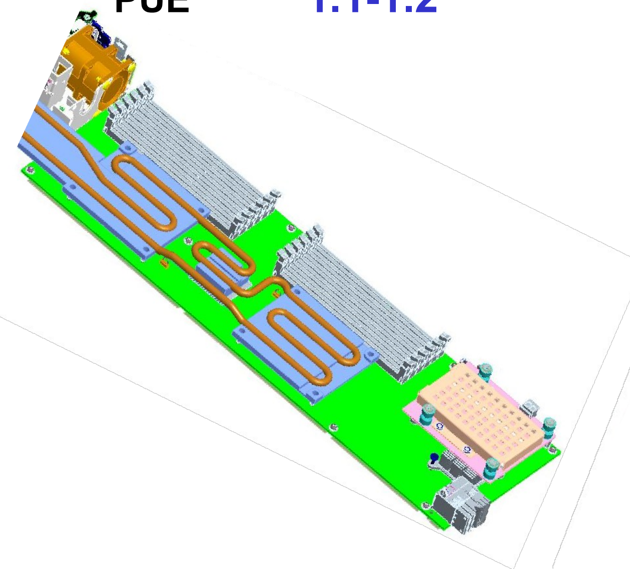

Total Cost of Ownership (TCO)

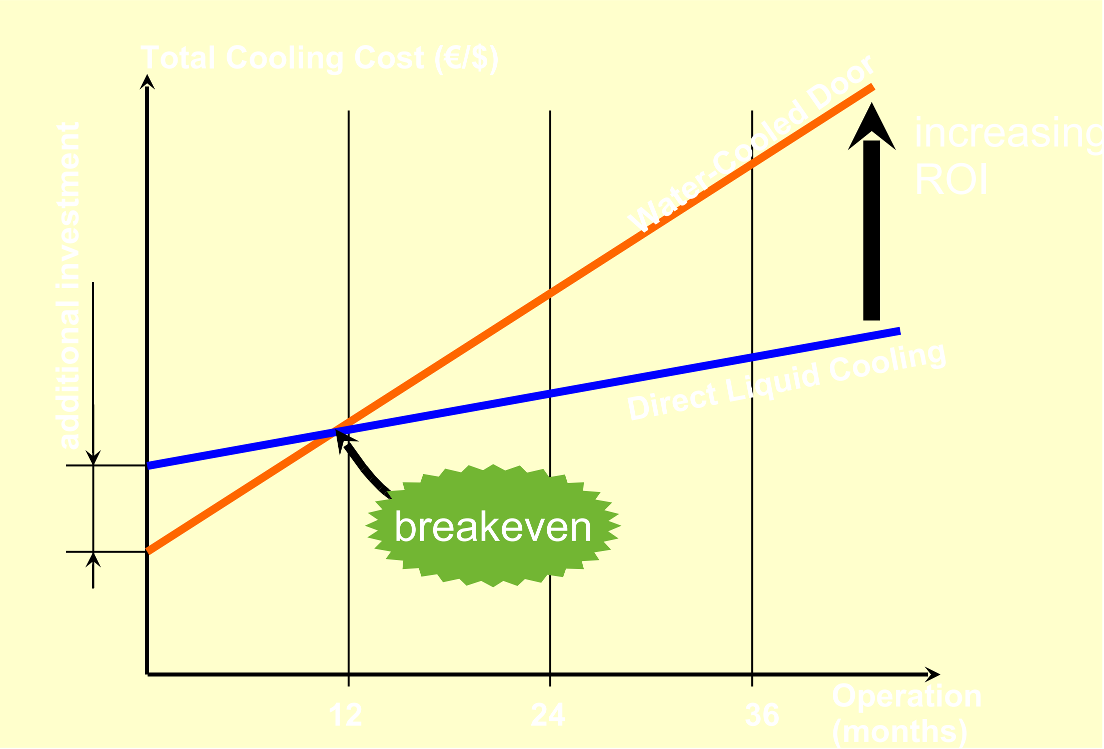

Are you ready to invest to improve
the TCO of your HPC system?

conclusion

- HPC applications requirements keep increasing
- … from Petascale to Exascale during 2010-2020
- Large nodes are an important piece in your workflow
- Today’s GPU accelerators are a stepping stone
- [r]evolution of GPUs / convergence with CPUs
- Greener systems for greener world
- ~~Challenging~~ Exciting times ahead of us

Questions ?

bullx

instruments for innovation

powered by Bull