The Abdus Salam International Centre for Theoretical Physics

HPC

"Introduction to High Performance Computing (HPC)"

Clement Onime – onime@ictp.it

Information & Communication Technology Section (ICTS)

International Centre for Theoretical Physics (ICTP)

Outline

- Definition of HPC
- Explore uses and impact of HPC
- Open discussion

Paradigm of Science

- Observe
- Theorize
- Experimentation

Now, less expensive

- Observe
- Theorize
- Computation
- Modeling & simulation

Also, in certain fields, the observe phase is replaced by Simulation. E.g: Study of the early Universe...

Some uses of computers

- Data Collection
- Data analysis
- Visualizations
- Producing reports
- Information exchange
- Communicating

Challenges of growth

- Increasing complexity.
  - Problems have become too hard, expensive, slow, controversial or dangerous.
  - Multiple data collection sites: national or international
  - Bigger data sets (aggregates of sources)
  - teams of researchers with diverse expertise, require many interactions.
  - Solutions are needed now (as fast as possible) to be useful.

In science

- Scientists are challenged to
  - Tackle complex theories numerically
  - Overcome “limitations” in experiments: study (virtual) experiments, where the boundary conditions are greater than what is physically possible.
  - Benchmark and improving correctness of models and theories.

Examples:

- … too HARD
  - e.g. building large wind tunnels
- … too EXPENSIVE
  - building a throw-away passenger jet
  - Simulate lasers behavior
- … too SLOW
  - waiting for climate or galactic evolution
- … too DANGEROUS or CONTROVERSIAL
  - Research on nuclear or radioactive material
  - stem cell research

Solution: Bigger computers?

- Traditional Supercomputers
  - Specially built computers
  - Expensive to create or grow
  - Steep learning curve which cannot be readily reused.
  - Difficult to adapt or re-purpose
- Alternative is bigger computations

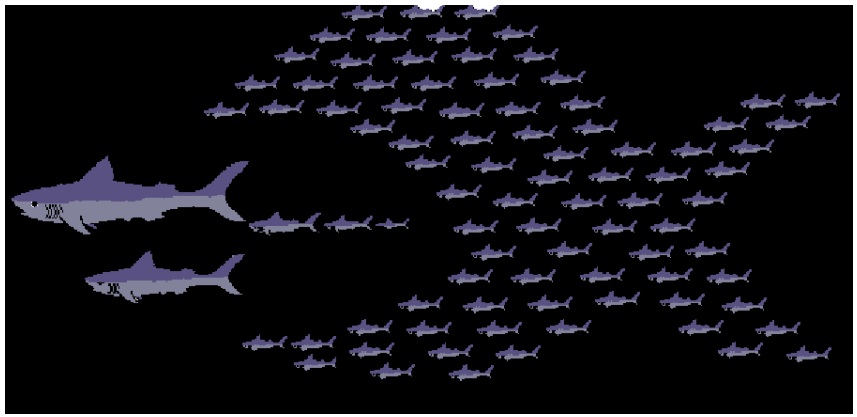

A Definition:

High Performance Computing encompasses a collection of "powerful":

- hardware systems
- software tools
- programming languages
- generic programming approaches

All coordinated together to obtaining faster results to bigger problems.

Some characteristics

- HPC focuses on improving productivity
- HPC can happens on:
  - Single workstation, laptop, smart-phone or multi-core devices such as exotic platforms (Accelerators, GPU, FPGA).
  - HPC become more powerful when computers/devices are clustered together.
  - Linux/AIX/Windows/MacOS

The Opportunity

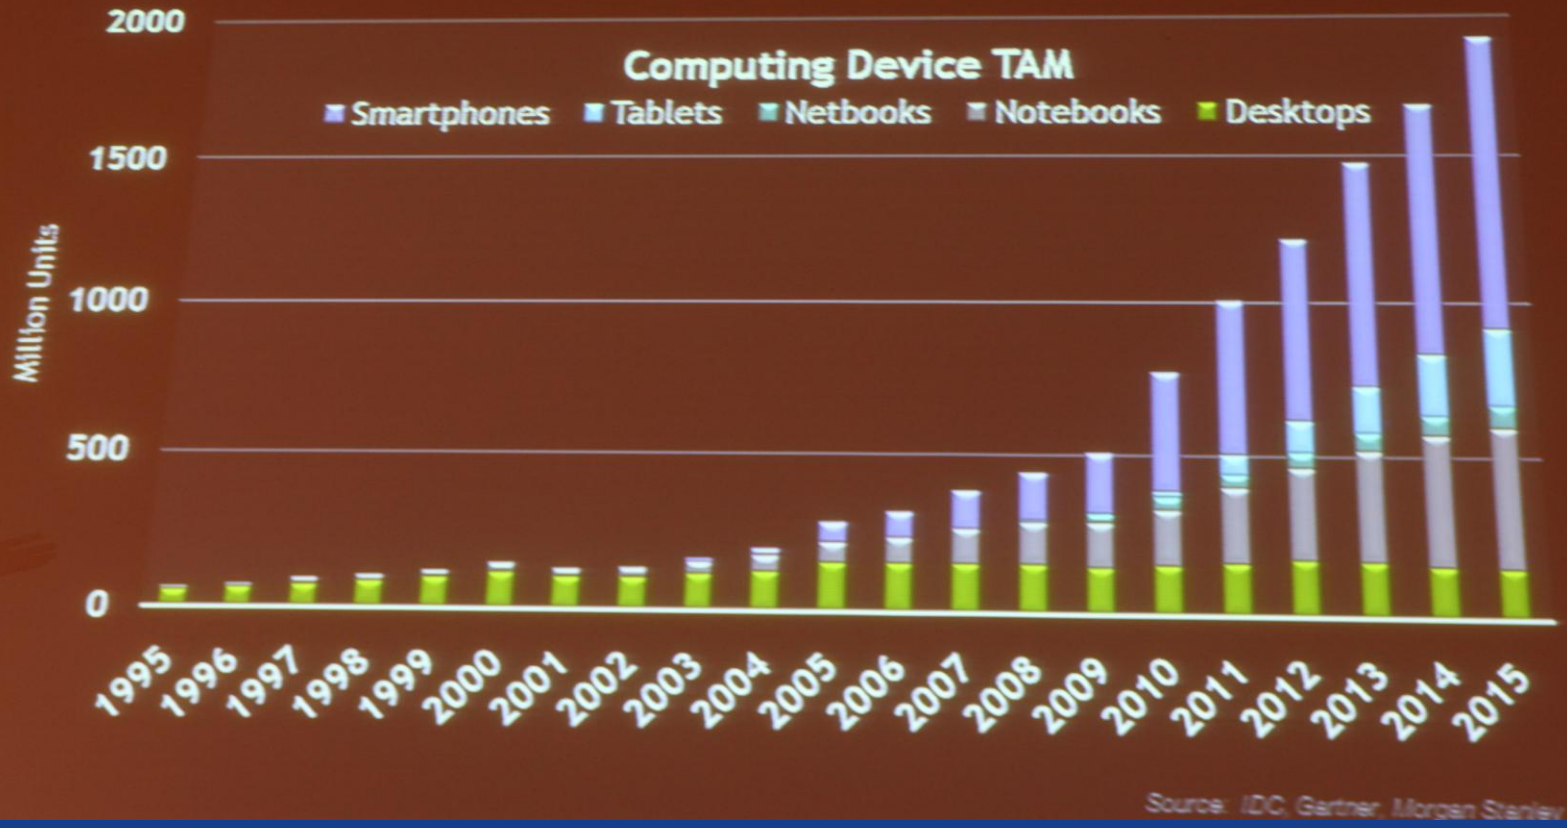

HPC by clusters of nodes (normal PC)

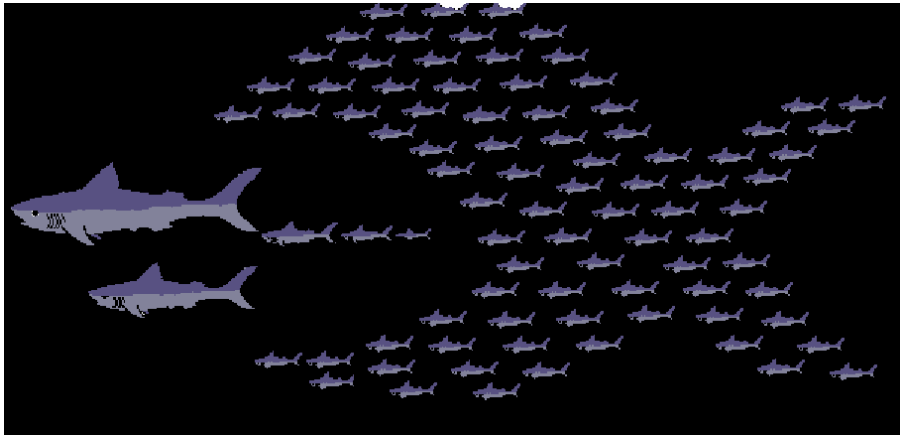

- Many computers (nodes) interconnected by high speed network (infiniband).
- Commodity clusters
- Hybrid platforms that supports both shared & distributed architecture & programming models.

HPC clusters / parallel processing

- Hardware
  - Commodity processors give higher performance at lower costs
  - Networking advances give high speed, low latency
  - Easier to integrate into existing networks
  - Costs , Lower initial, running & upgrade
- Versatile
  - Can implement MPP or DSM, Network Ram, Parallel I/O (RAID) & multipath communication.

HPC cluster architecture

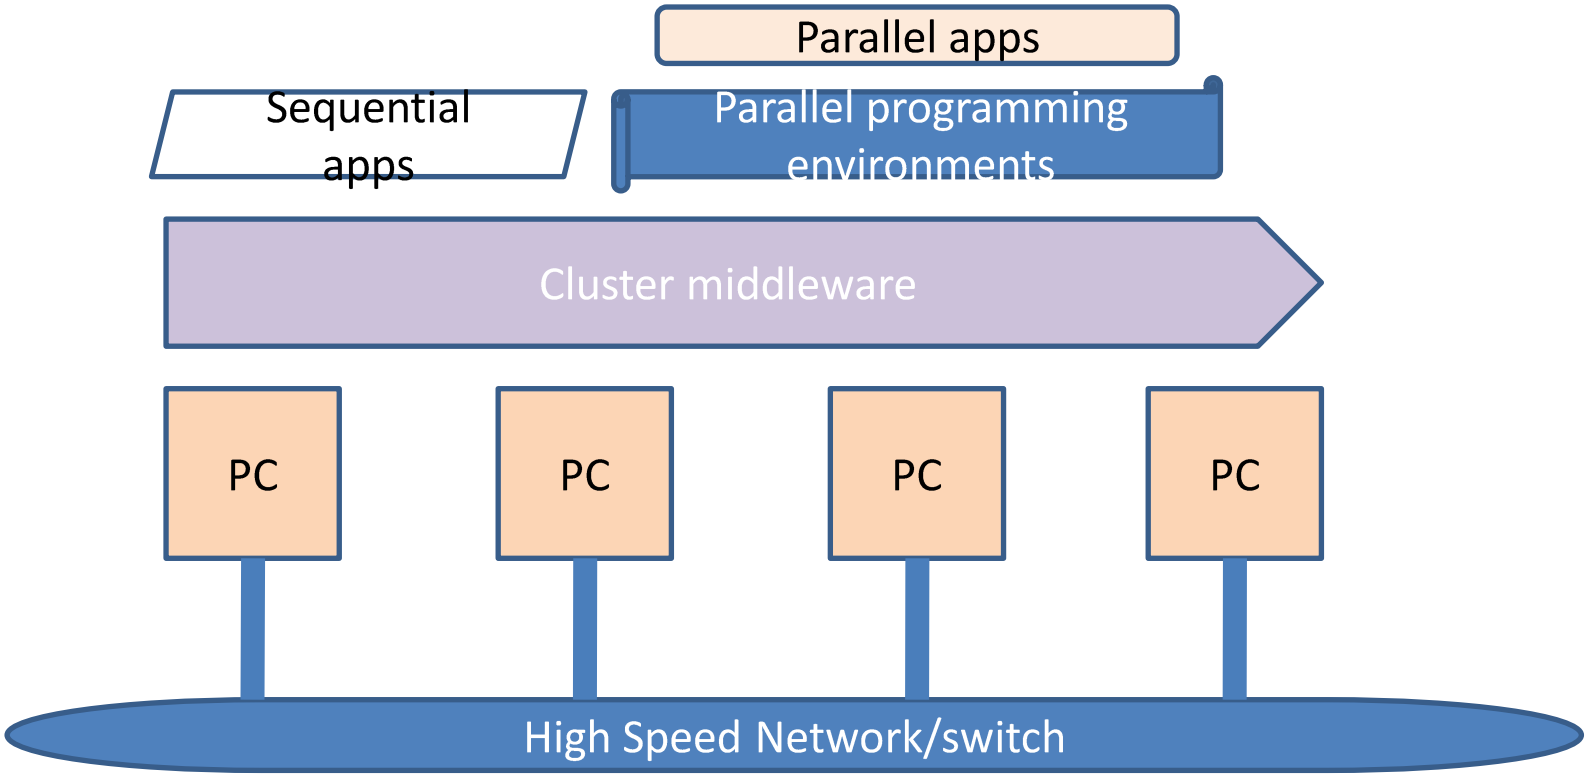

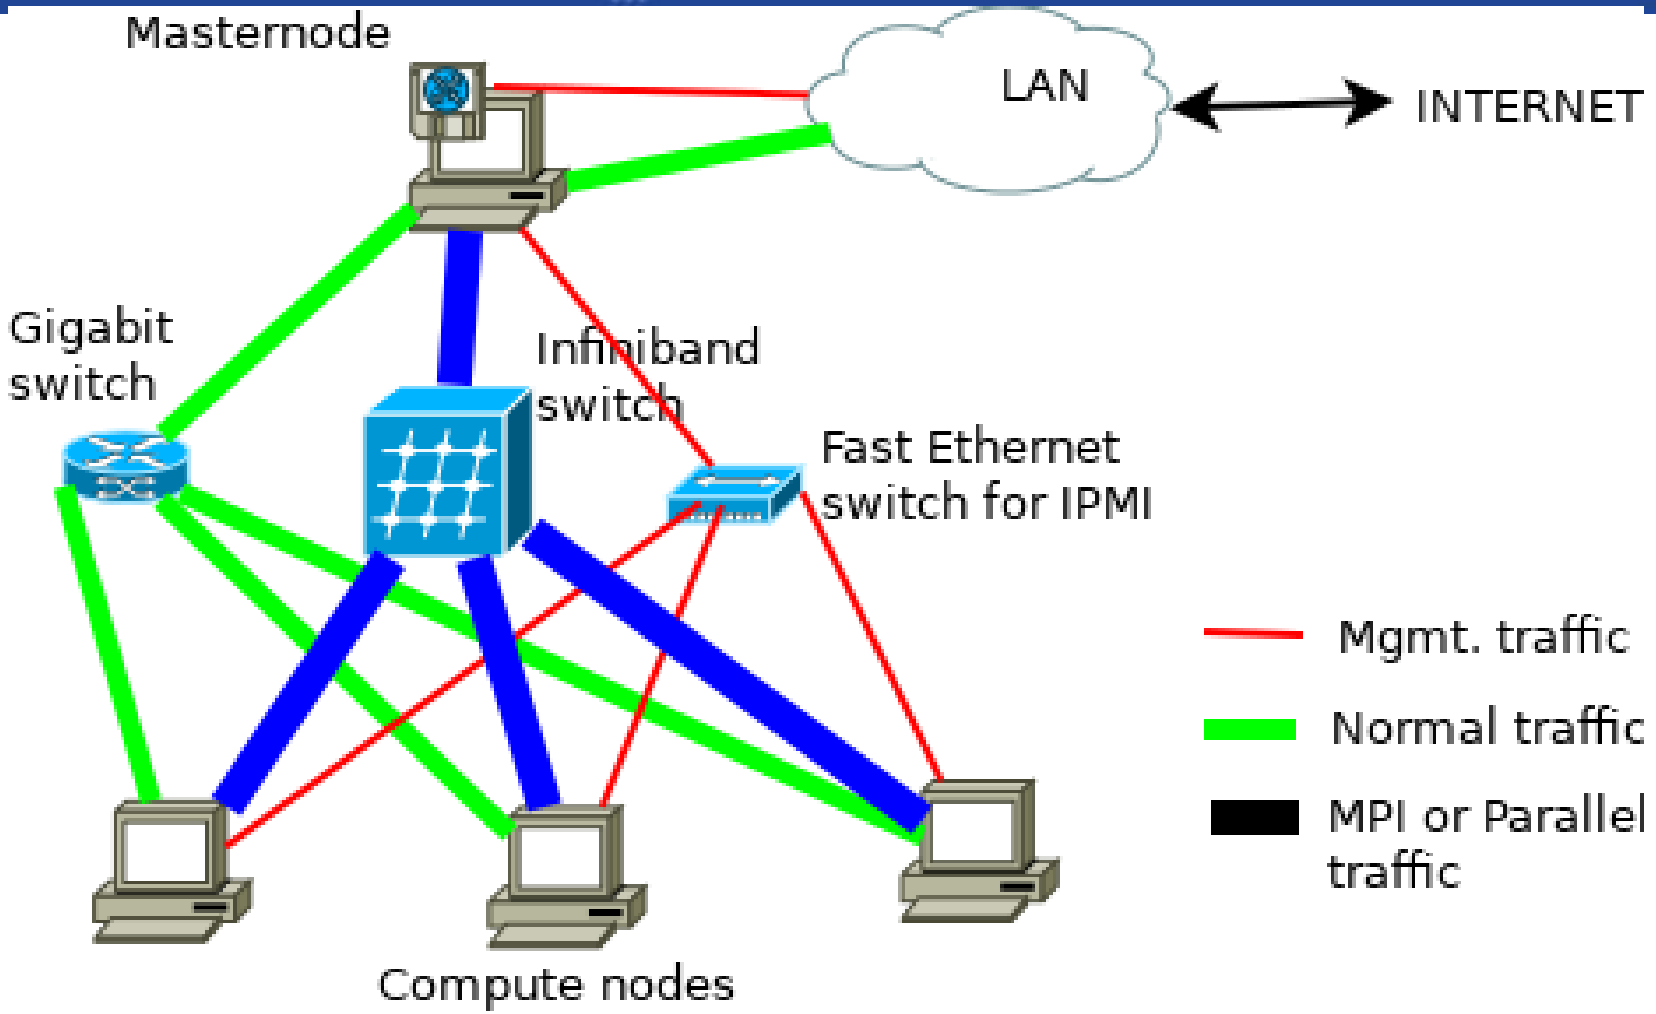
Masternode
LAN
INTERNET
Gigabit switch
Infiniband switch
Fast Ethernet switch for IPMI
Mgmt. traffic
Normal traffic
MPI or Parallel traffic
Compute nodes

Major components of a HPC cluster

- PC Computers
  - CPU - Single core, dual core, **six-core**, **quad-core**, SMP or others..
  - Network Interface Card (ethernet, Myrinet or Infiniband)
  - Operating system (thin, fat, micro-kernel, etc)
- High speed network
  - Gigabit, Myrinet, **Infiniband**
- Cluster middleware
  - Resource Manager, shared storage , Parallel FS, Parallel Memory, SSI
- Parallel programming environments
  - **MPI**, PVM, CUDA, debugger & profilers, etc
- Applications
  - Serial, sequential OR **parallel**/distributed

HPC based Supercomputers

- Supercomputers
  - Built using clusters of normal computers
  - In-expensive to create or grow
  - Software is what already know and use.
  - Easy to adapt or repurpose

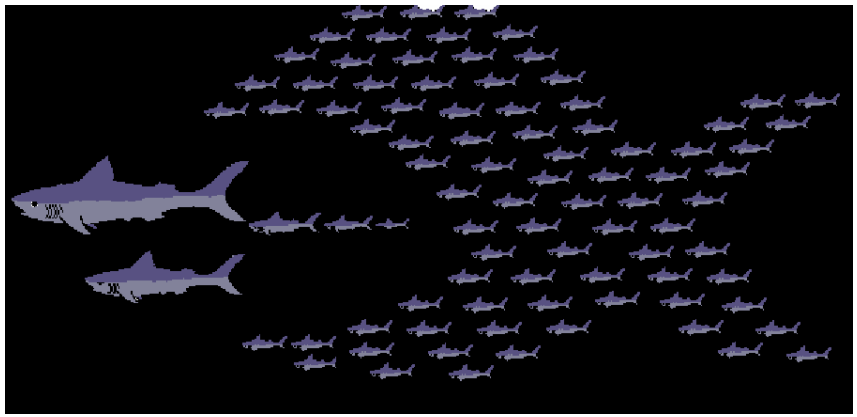

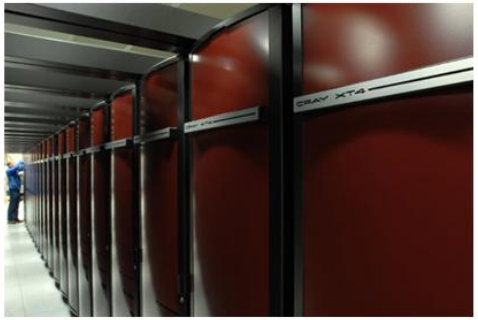

HECToR (XT4 component), U.K

- Cray XT4 supercomputer
  - Quad-core AMD Opteron CPUs
  - 3,072 nodes * 1 CPU * 4 cores
  - 8 GB memory / node
- 12,888 cores and 24.5TB of RAM
- Now only a PART of a bigger HECToR!

JuRoPA in Germany

- Combination of Juropa-JSC & HPC-FF
- Intel Xeon X55xx (Nehalem-EP) quad-core processors
- 3,288 nodes * 2 CPUs * 4 cores
- 24 GB memory / node

- Total of 26,304 cores, 79 TB main memory

- 308 teraflops peak performance

TOP 500

- World wide Ranking of super-computers
  - Twice a year exercise
- See http://www.top500.org/

- Key indicators:
  - Country leaders
  - Supercomputing is now predominantly HPC based and running the Linux operating system

Top 500 supercomputers

(by countries)

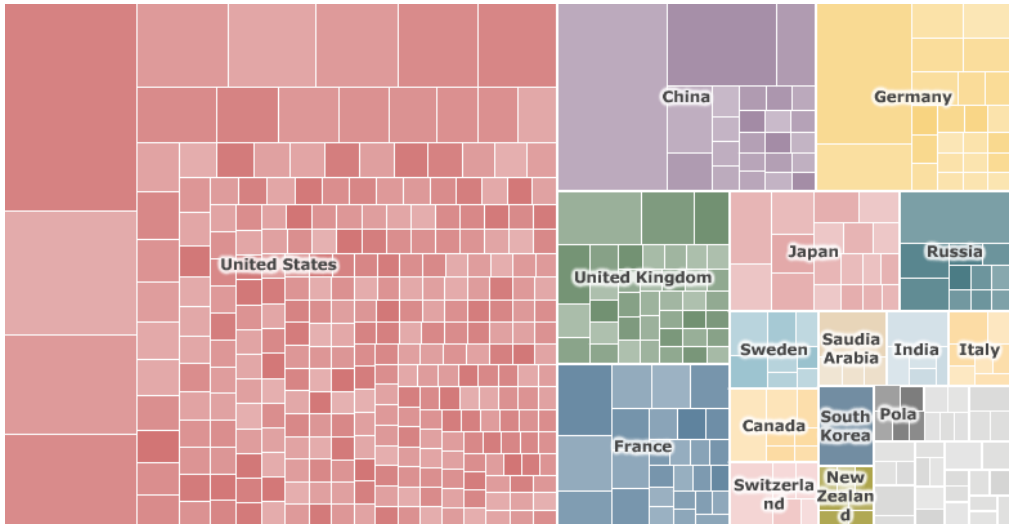

2010

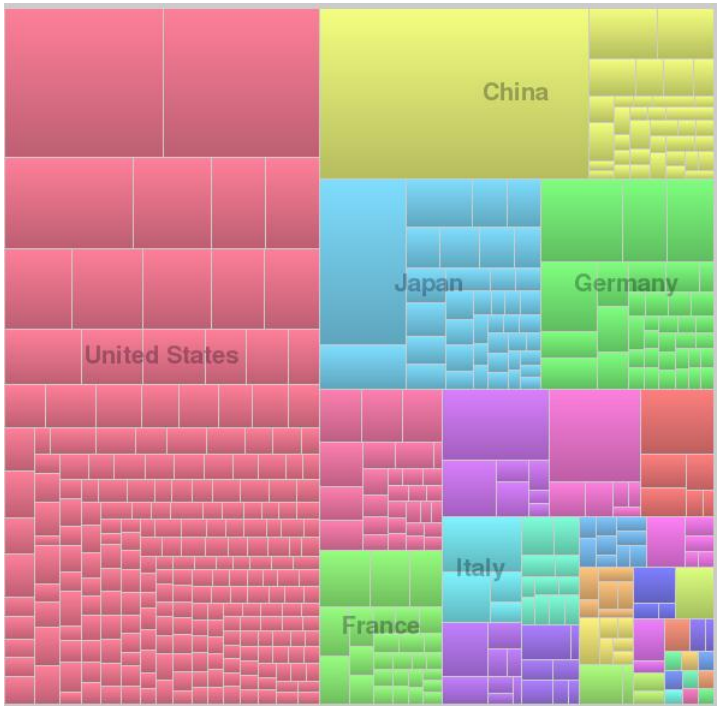

2015

Source: http://www.top500.org/statistics/treemaps/

Top 500 supercomputers

(by Operating Systems)

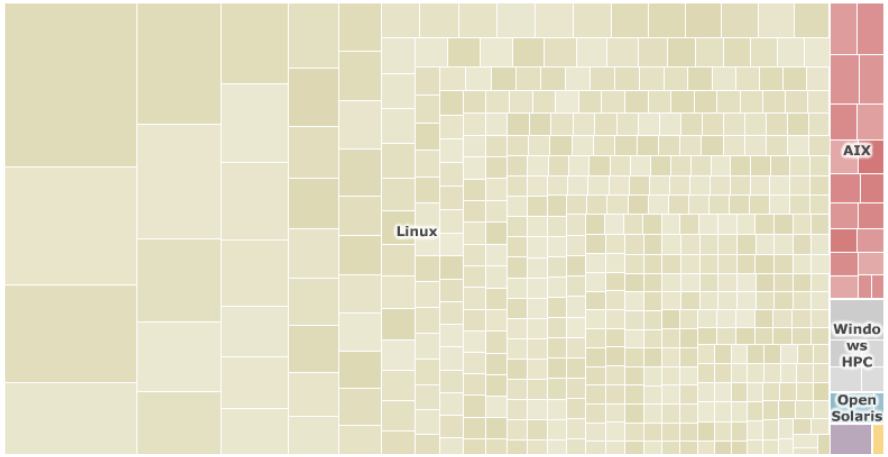

2010

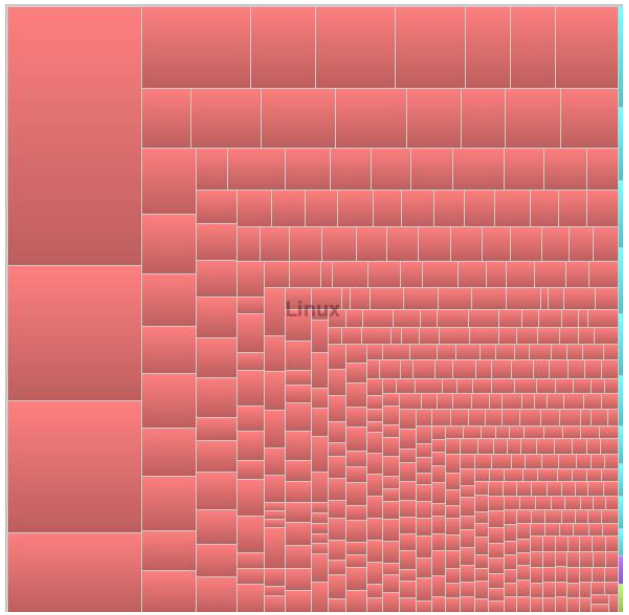

2015

Source: http://www.top500.org/statistics/treemaps/

Top 500 statistics

Application Area System Share from November 2011

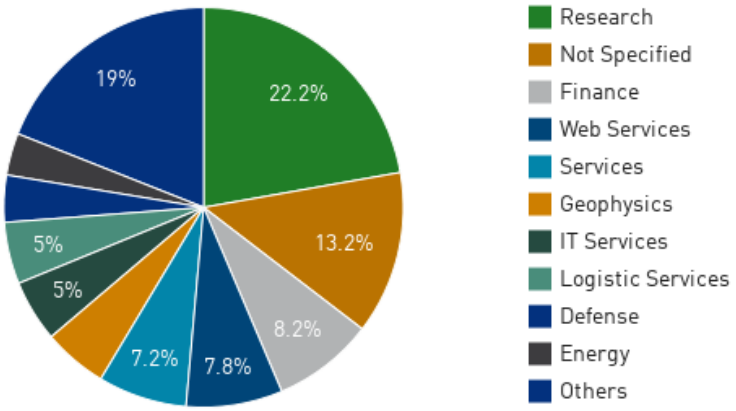

Vendors System Share from June 2015

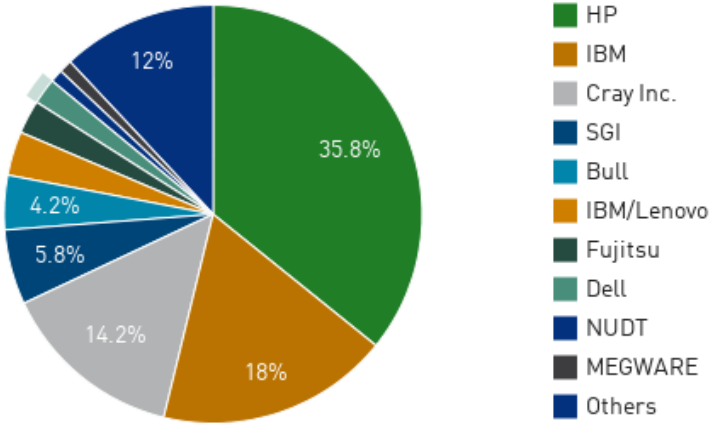

Source: http://www.top500.org/statistics/treemaps/

Example from Material-Science

QUANTUMESPRESSO

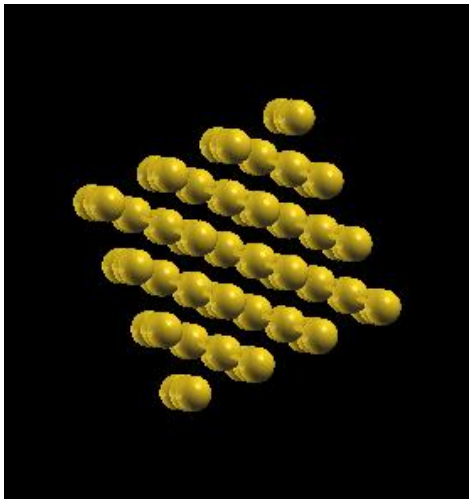

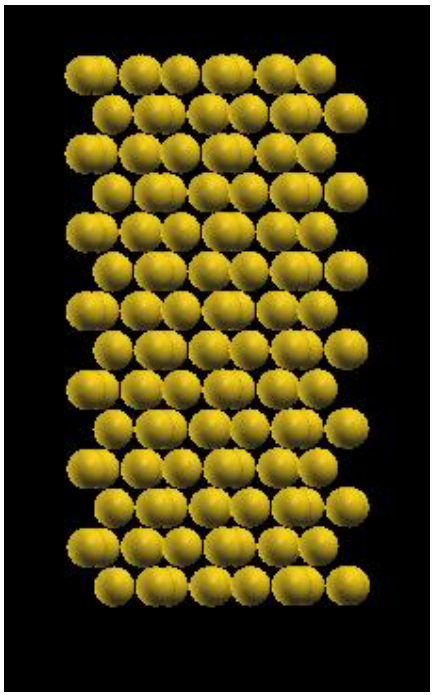

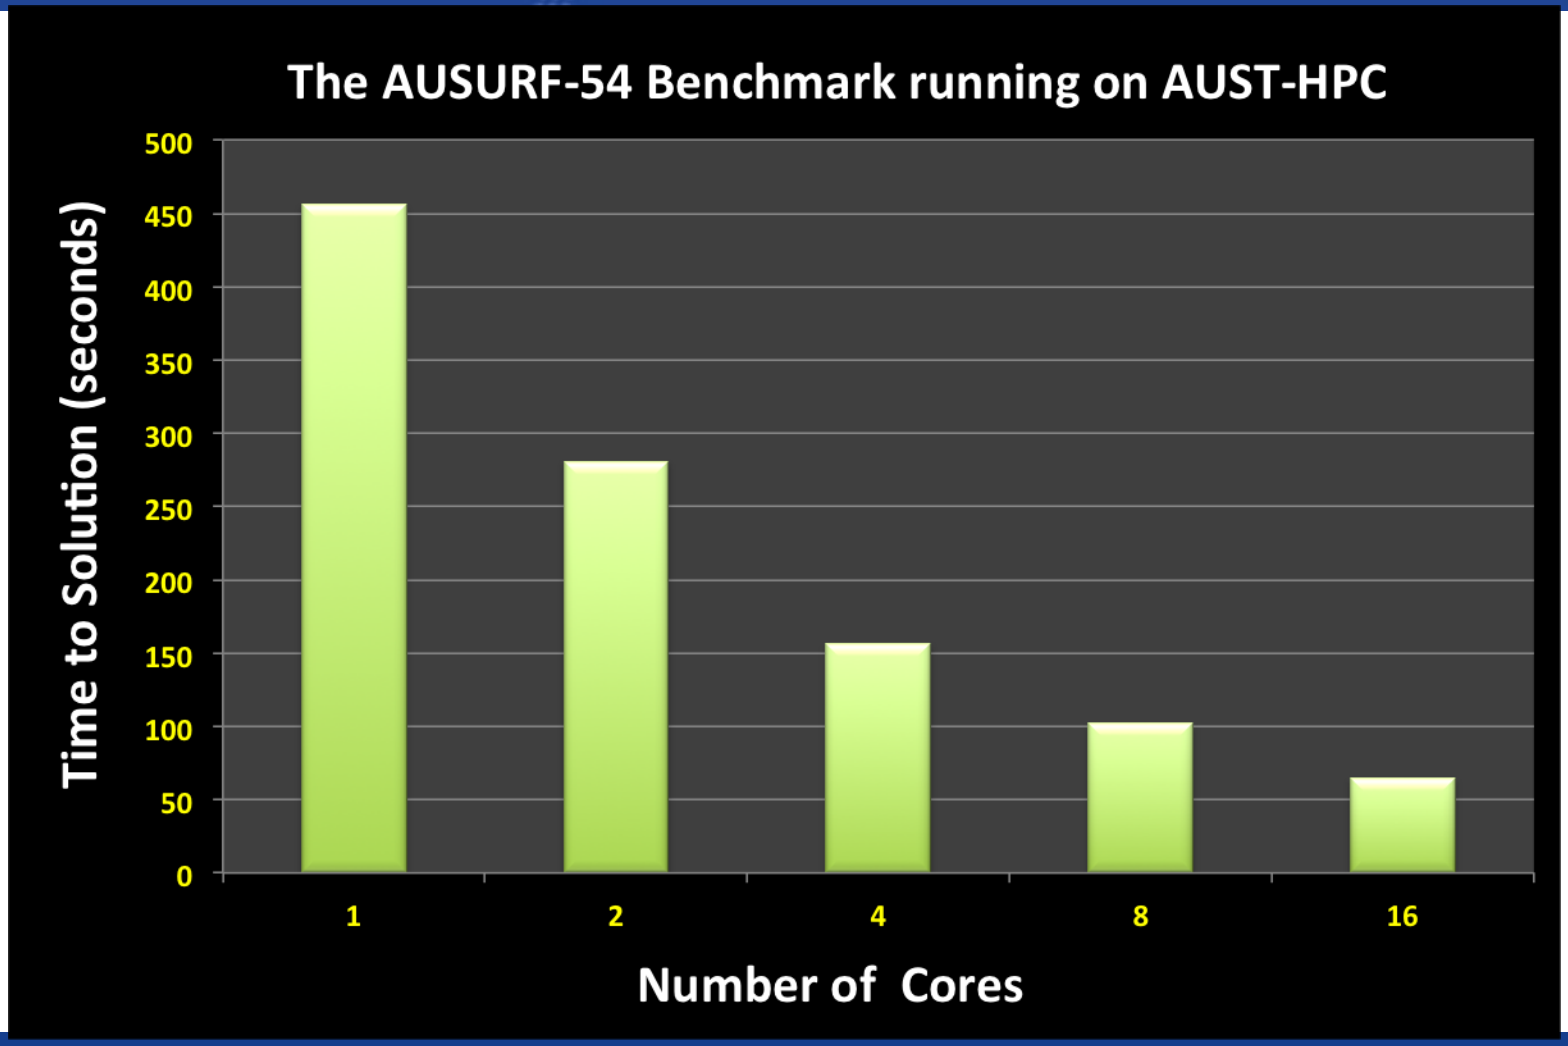
The AUSURF-54 Benchmark running on AUST-HPC
Time to Solution (seconds)
500
450
400
350
300
250
200
150
100
50
0
1
2
4
8
16
Number of Cores

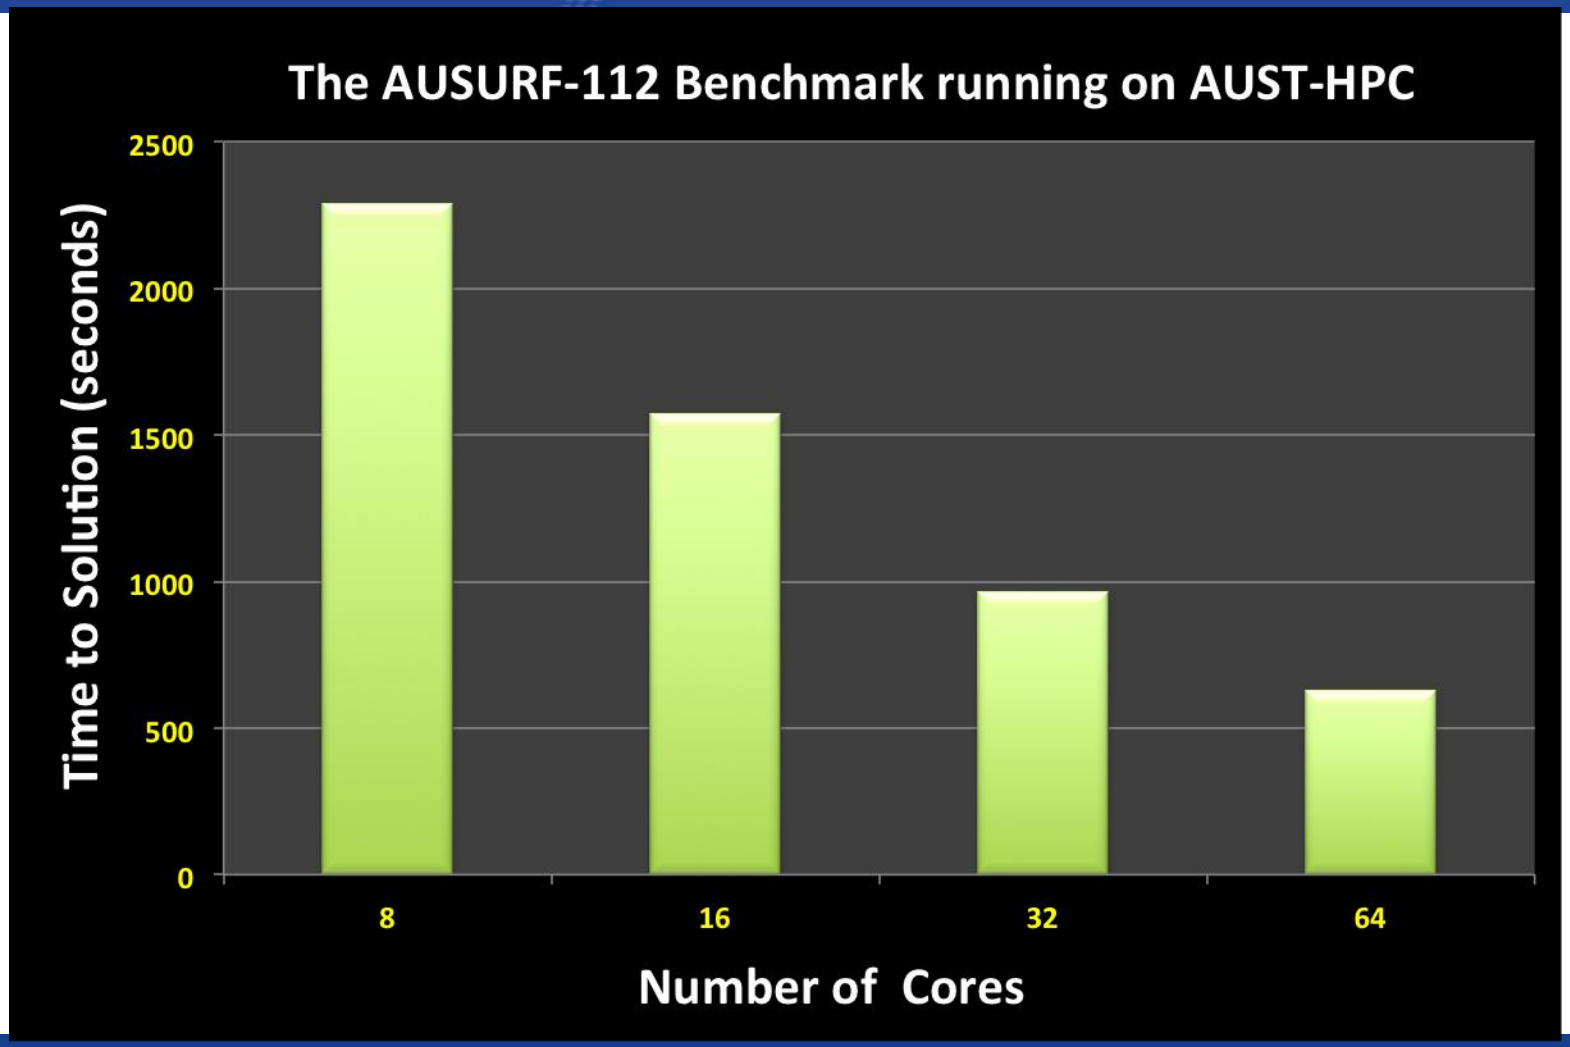
The AUSURF-112 Benchmark running on AUST-HPC
Time to Solution (seconds)
2500
2000
1500
1000
500
0
8
16
32
64
Number of Cores

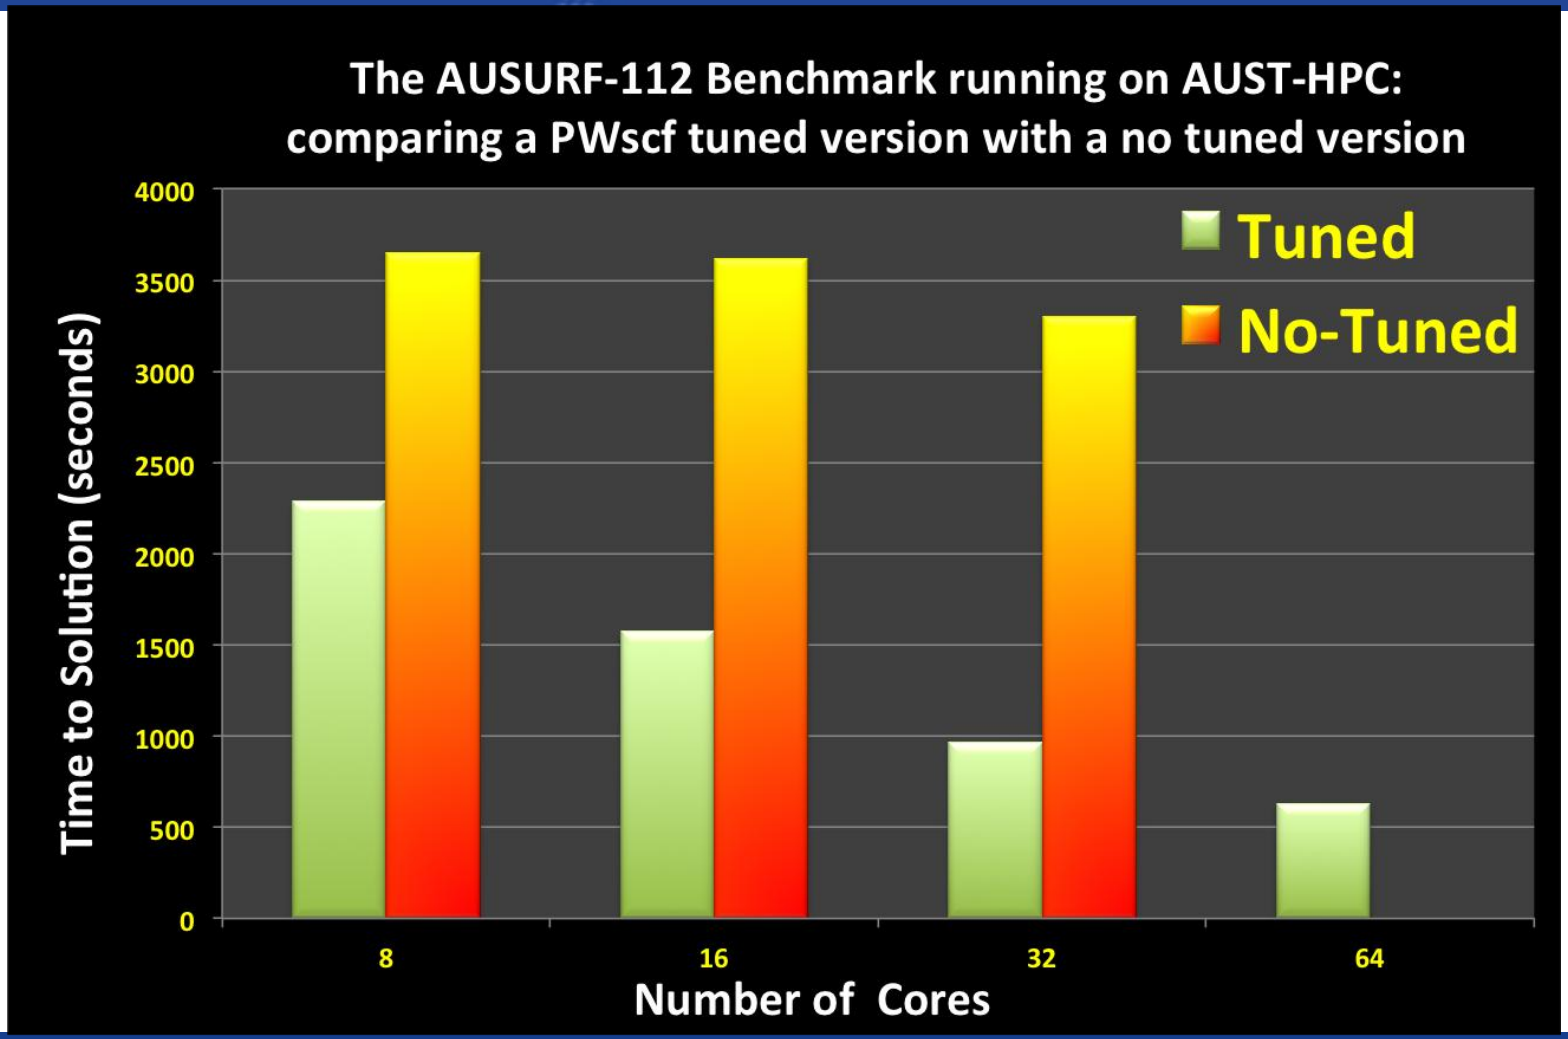
The AUSURF-112 Benchmark running on AUST-HPC:
comparing a PWscf tuned version with a no tuned version
Time to Solution (seconds)
4000
3500
3000
2500
2000
1500
1000
500
0
Tuned
No-Tuned
8
16
32
64
Number of Cores

Example from Climate Modeling and Weather Forecast

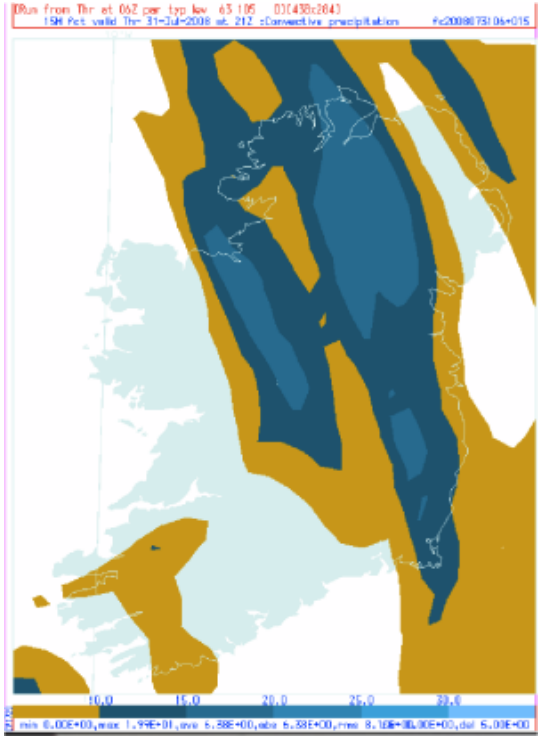

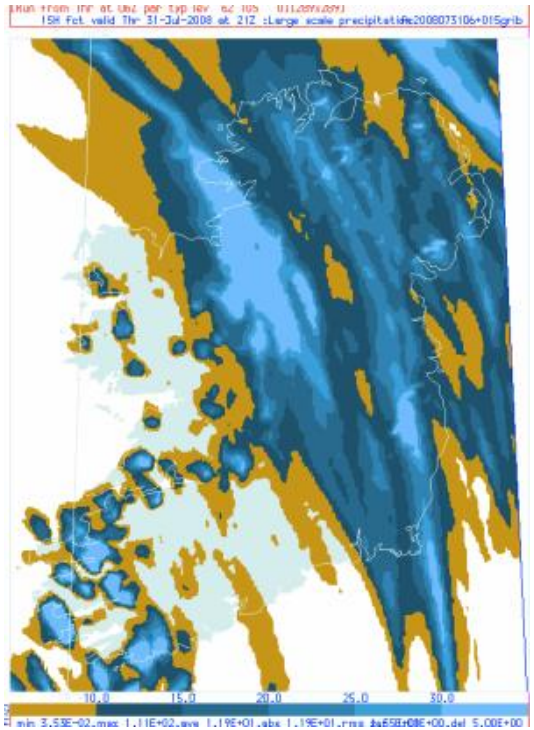

OpenFOAM: An Open Source SW for Engineering Modeling

Oil and gas

- Production facilities for oil & gas
- Pipeline systems
- Tank farms and underground storage facilities
- Refineries and petrochemical plants

Water and environment

- Water supply
- Wastewater treatment & disposal
- Waste treatment & disposal
- Hydropower, dam and river engineering

Energy

- Thermal power plants
- Sea water desalination plants
- Renewable energy
- Climate protection
- Transmission and distribution systems

Civil engineering and infrastructure

- Airports
- Roads
- Railway systems
- Tunnels and caverns
- Buildings and structures
- Alpine engineering

Oil & GAS

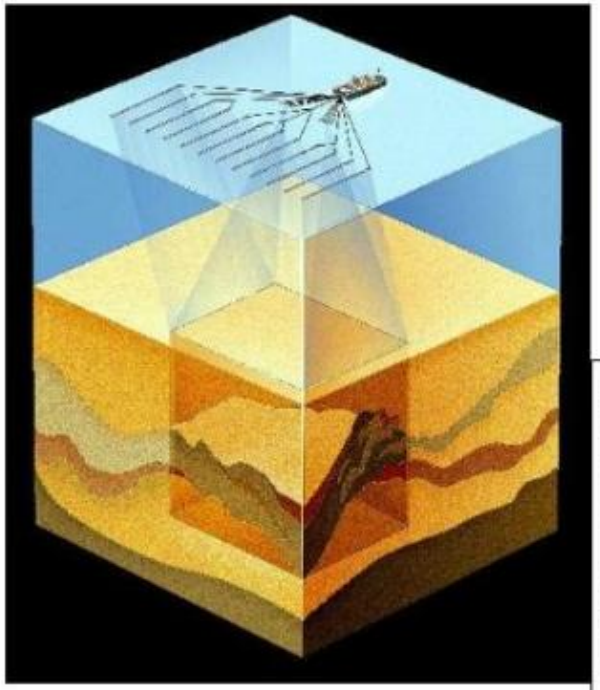

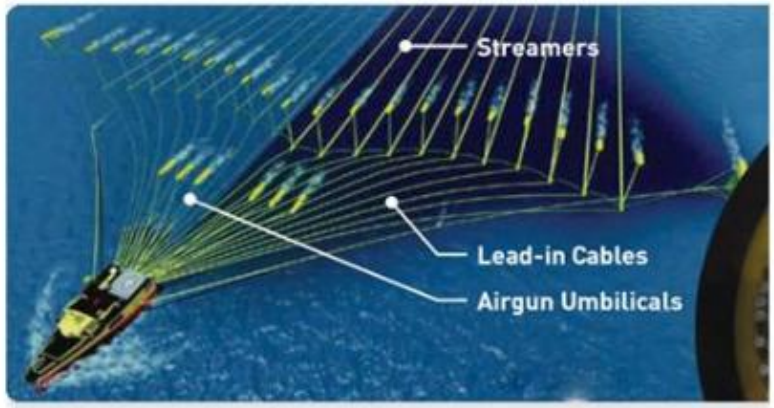

A typical marine seismic acquisition survey - a marine vessel towing energy sources (airguns) and sensors embedded in streamers, sensors may be ~10m apart on 10 or more streamers each of which may extend >10km long.

Shot gather = data recorded by all sensors from a single source detonation, typically there are thousands of sensors and thousands of source locations.

With shots spaced ~20x20m apart and survey sizes of a few thousand sqkm common ... a typical survey could have >50Tb of raw data – ~50,000 shots on a ~20x20m grid over ~2000sqkm, recorded on ~10000 sensors, sampled every 2ms for 6000ms 'listening time'.

Image and pattern analysis

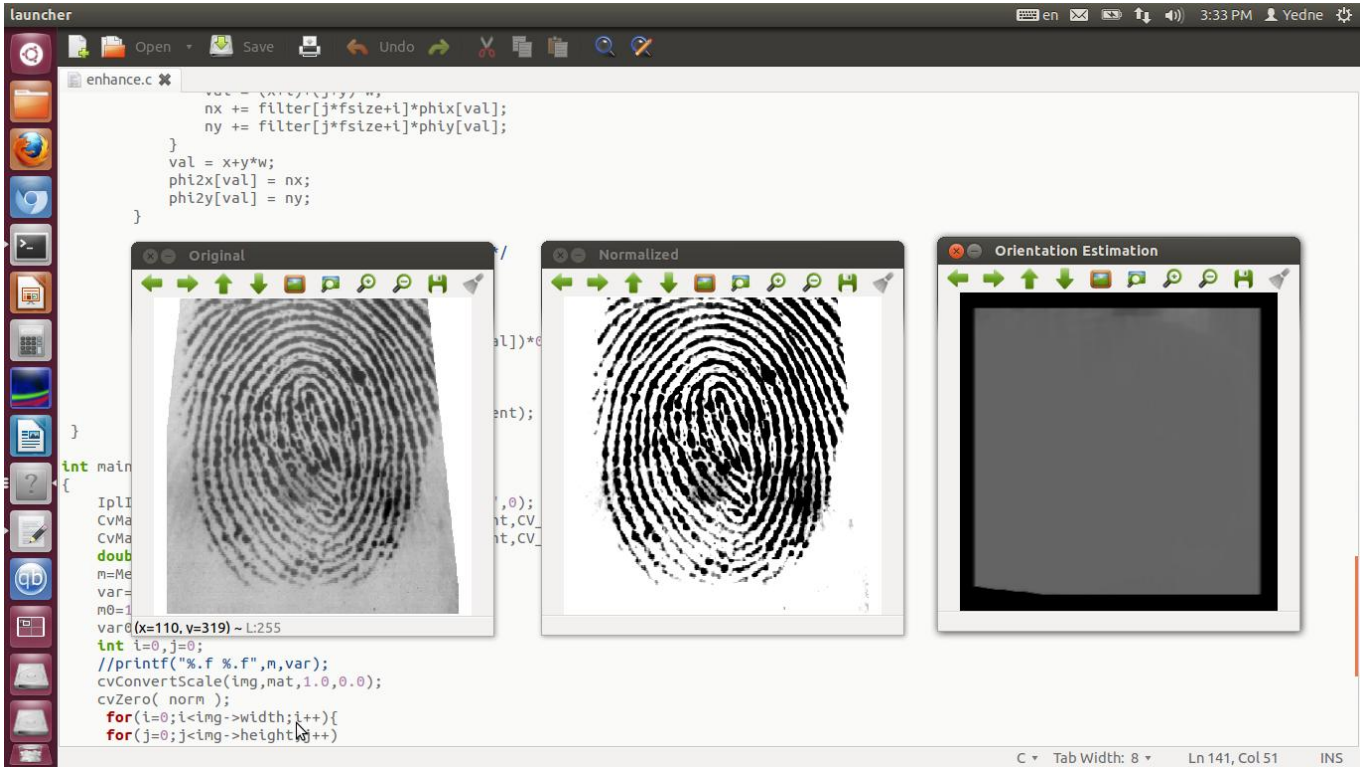

Economic impact

Airlines

System-wide logistics optimizations evaluated on HPC systems save approx. US$100 million per airline per year.

Automotive Design

Major companies use (500+ CPUs) in CAD and CAM for crash testing, structural integrity and aerodynamics saving over US$1 billion per company per year.

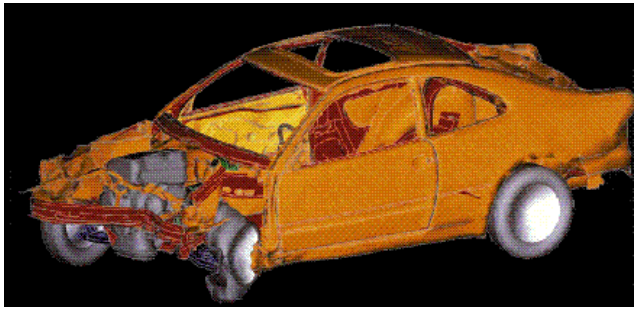

Physics

Detectors at the Large Hadron Collider at CERN, Geneva

- Set to produce several ***Petabytes*** (10^{15} bytes) of data per year
  - a <u>million times</u> the storage capacity of the average desktop computer;
  - accounts for nearly <u>10%</u> of all the information produced by humans each year.

Performing the most rudimentary analysis of the LHC data will require close to 20 TeraFlops (*a trillion floating-point operations per second*)

Animations

- The movie *Shrek 3*, consumed close to 20 million CPU render hours
- Each frame is rendered at DreamWorks Animation, with more than 1,000 Linux desktops and more than 3,000 server CPUs
- Each frame is assigned to a different node of the renderfarm by grid software (using Platform LSF, a commercial Linux package), so that many frames can be output simultaneously.

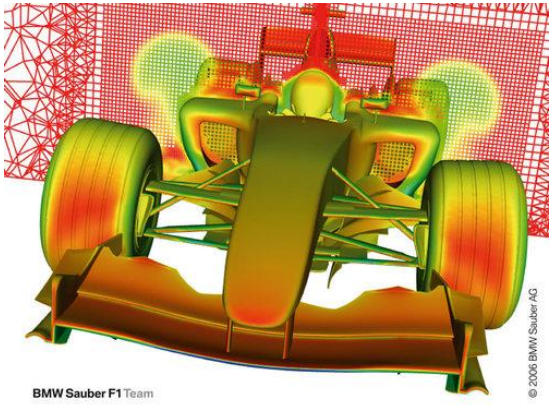

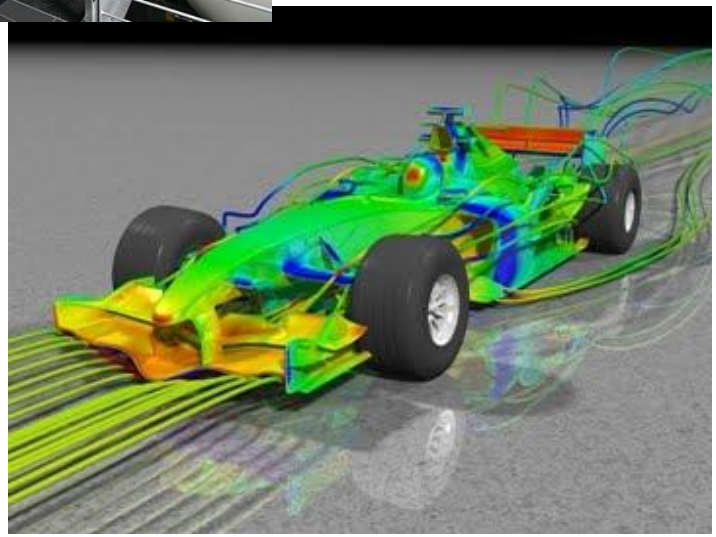

Formula 1 Racing

Swimming

Speedo LZR Racer

- 79 out of 105 world records in 2008.
- CFD technology from ANSYS was used to predict fluid flows around the body.
- Strategically placed polyurethane panels (designed by NASA).

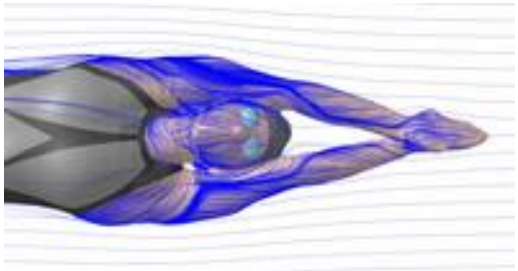

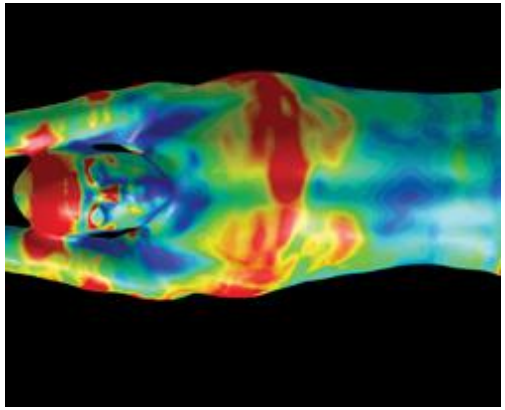

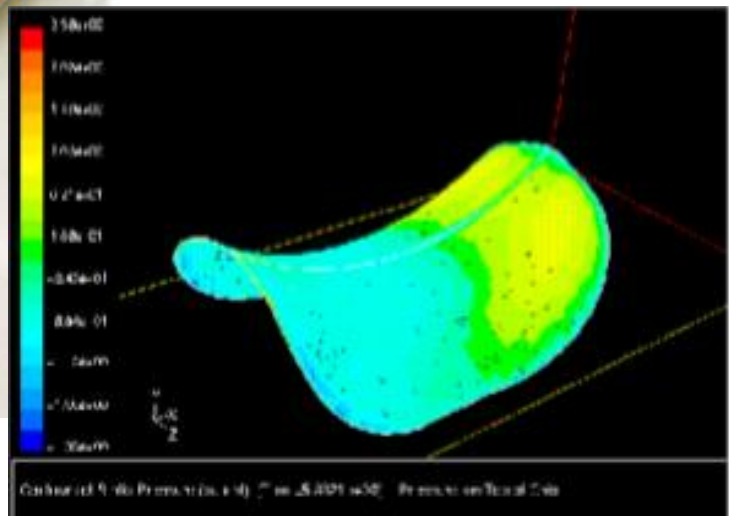

Product Design Procter & Gamble

Packaging & products designed on HPC systems.

Biology

Protein Folding

1 Protein

300 amino acid

~32,000 atoms

1 milisecond

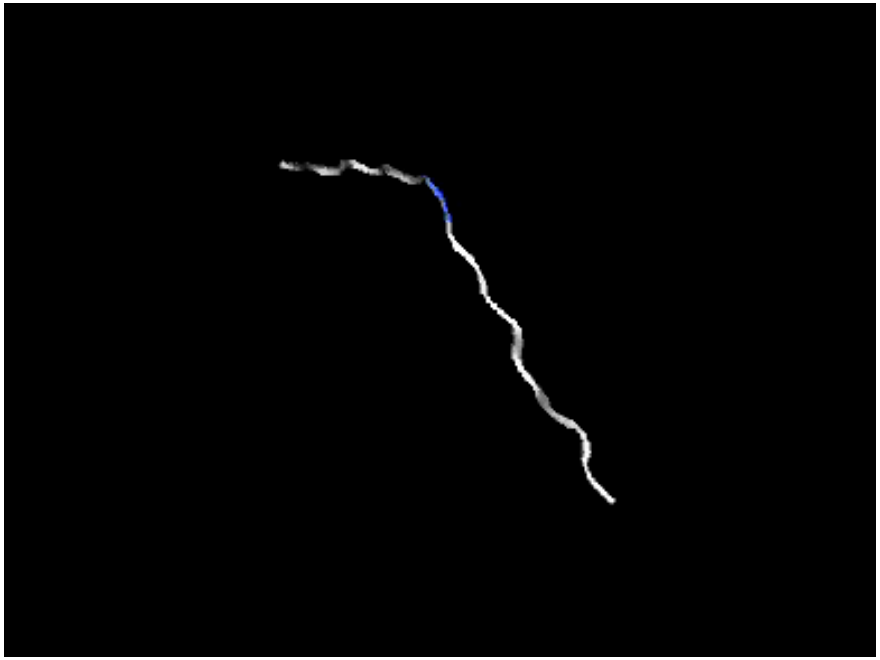

Conclusion

- HPC is for everyone who will like to improve their work as it can lead to faster solutions, better science and informed decisions, more competitive products!
- HPC systems can make research work, smarter, conducive and enable the delivery of world class results faster.
- Commodity based Linux clusters can give super computer class computing for a fraction of the cost and open new areas of development.

Thanks for your attention!!