Scientific Simulation on Emerging Technologies

Alan O'Cais
Scientific Coordinator
CaSToR^{C}

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

Outline:

- Capacity or capability?
- Capability architecture
  - Accelerators
- Project tools and management

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

Strong/Weak Scalability and HPC Systems

- Strong scaling
  - Same problem size in shorter time
- Weak scaling
  - Larger problem size in same time
- Capacity computing
  - Cycle-scavenging
  - GRID systems
- Capability computing
  - Larger problems, more expensive systems

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

Strong/Weak Scalability and HPC Systems

- Strong scaling
  - Same problem size in shorter time
- Weak scaling
  - Larger problem size in same time
- Capacity computing
  - Cycle-scavenging
  - GRID systems
- Capability computing
  - Larger problems, more expensive systems

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

Architecture: A generic view

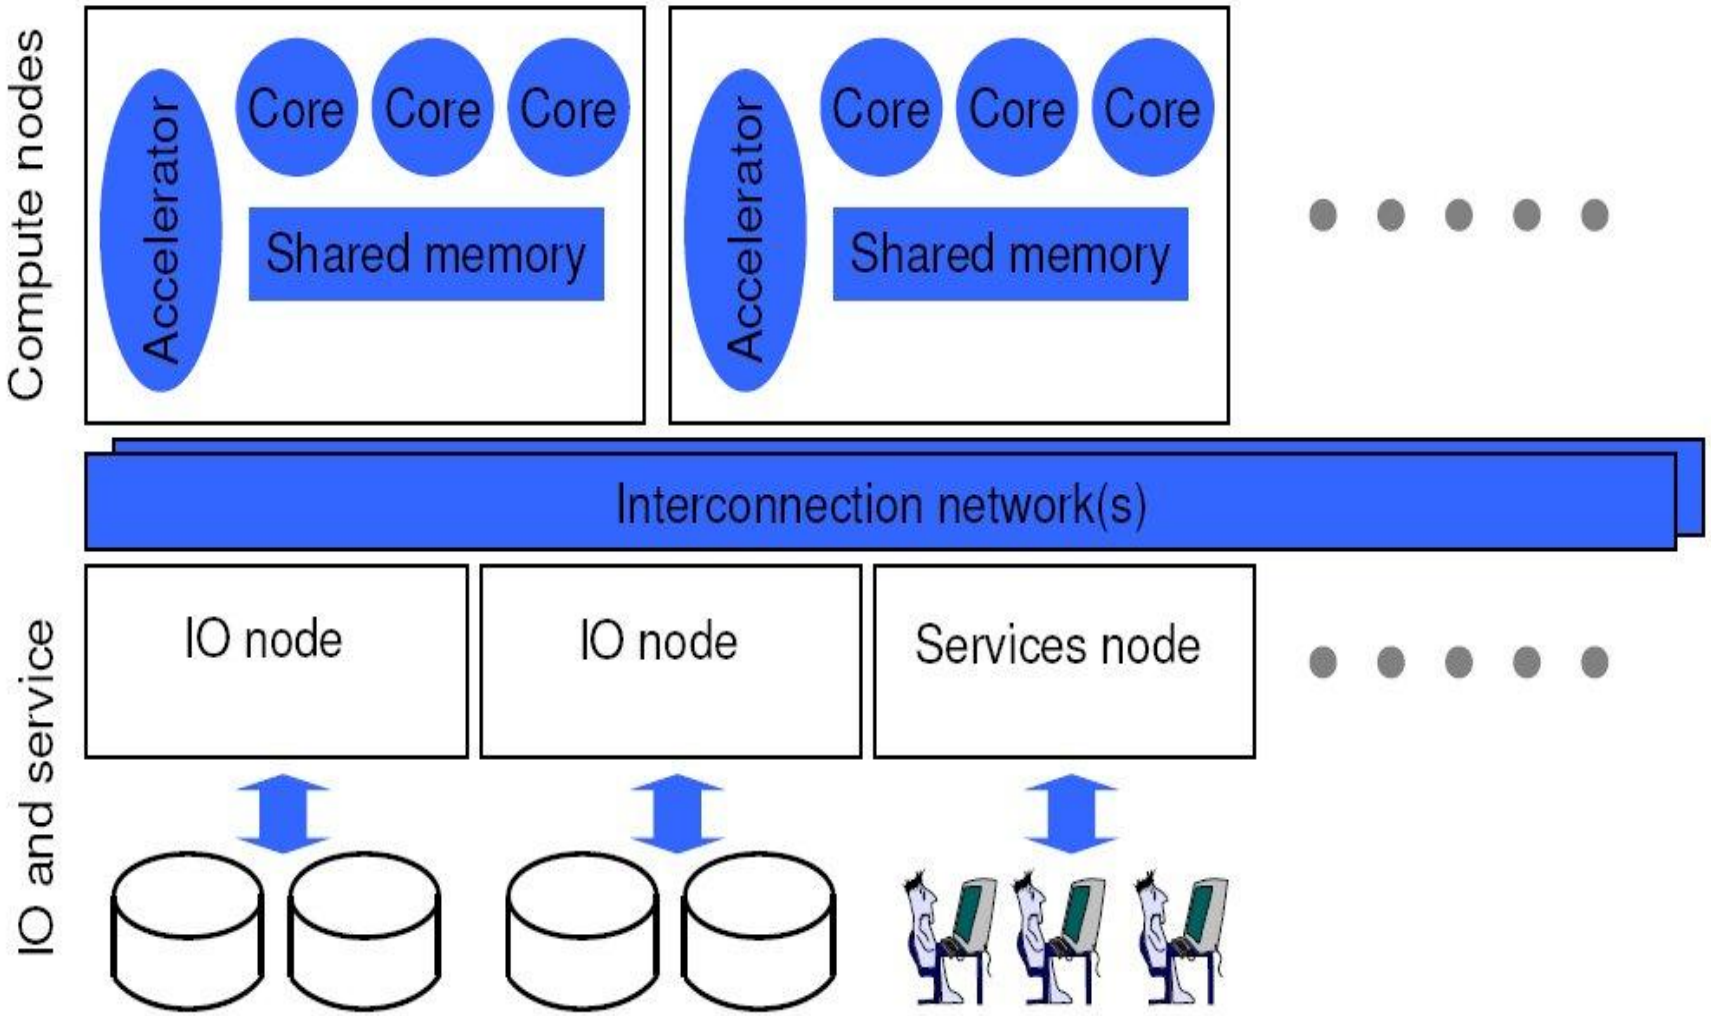

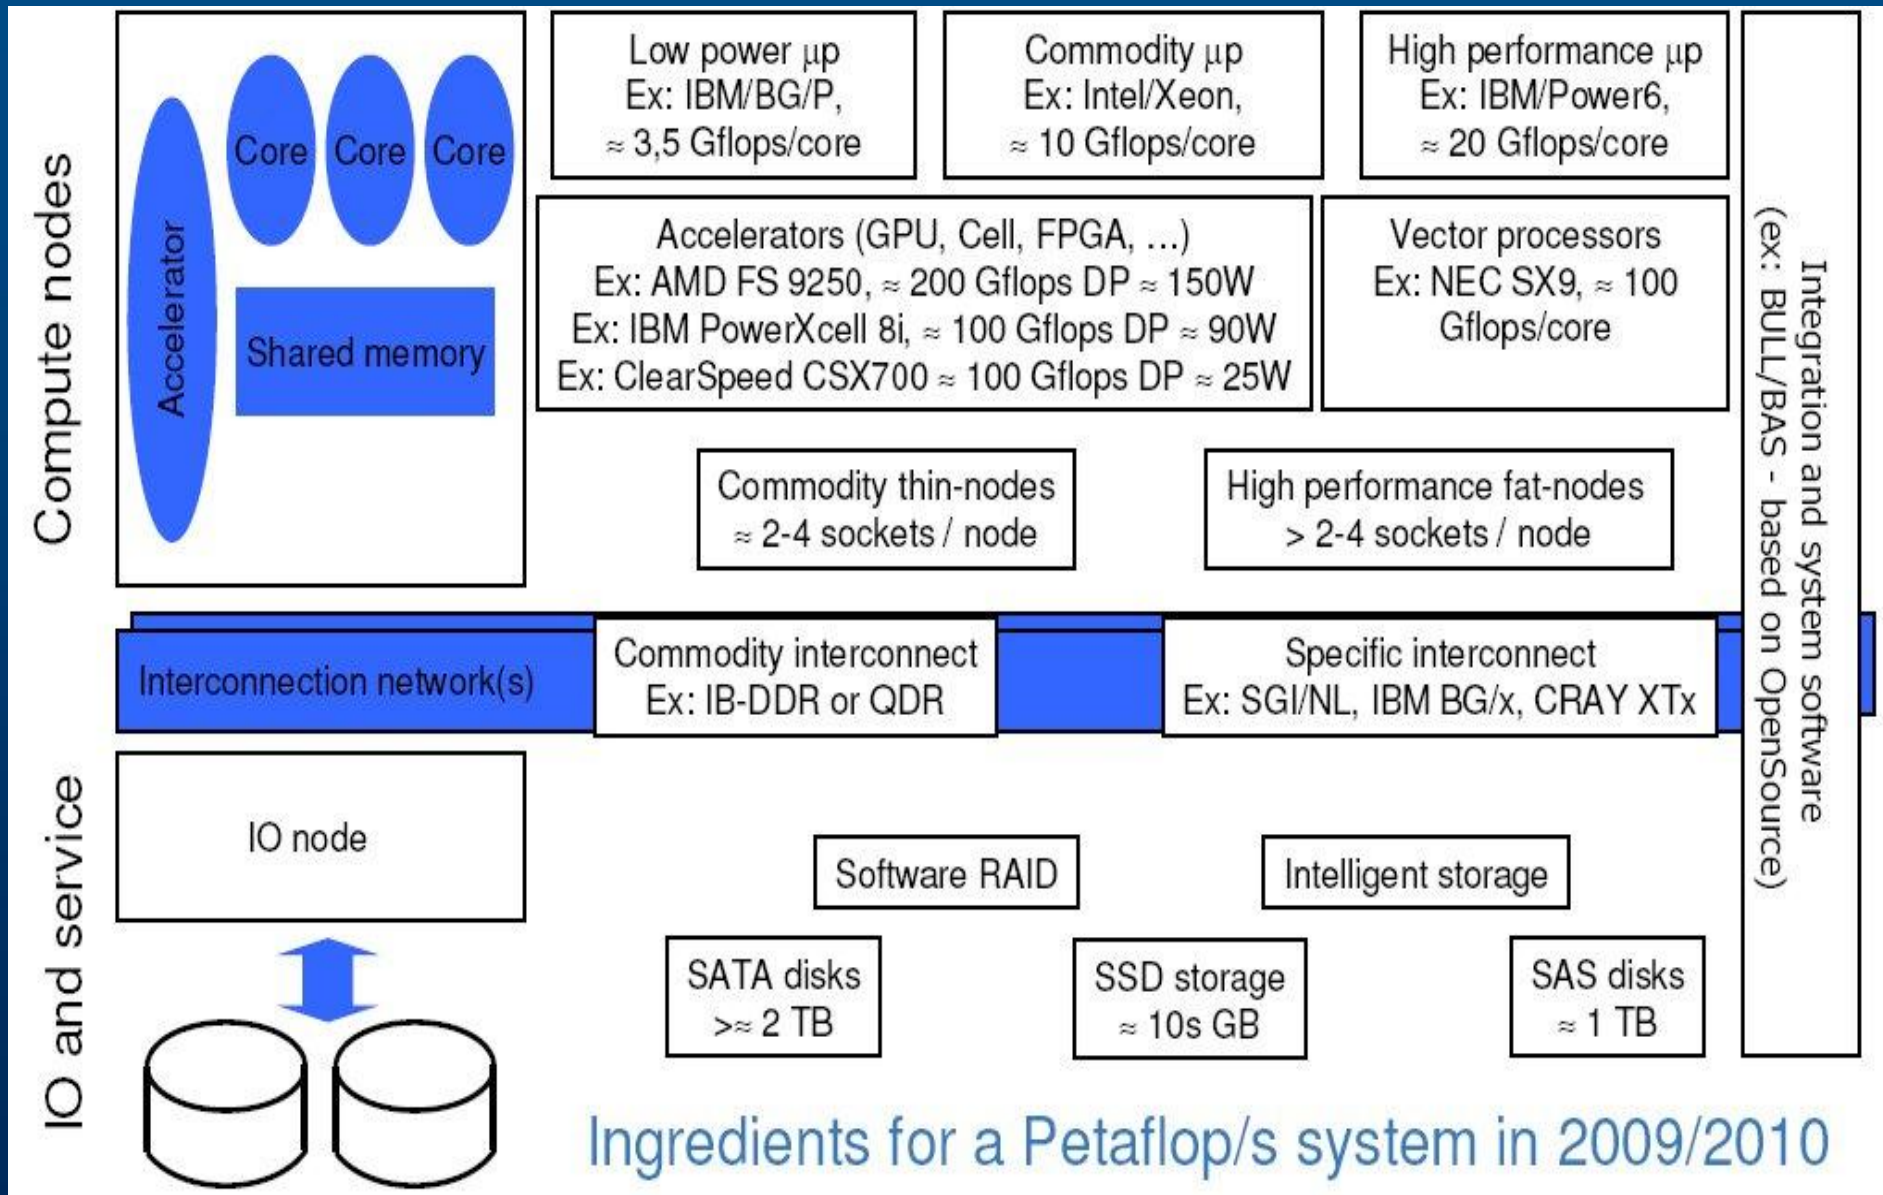

Ingredients for a Petaflop/s system in 2009/2010

Foundation of Cyprus with grant number
ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

Cell

- One PowerPC Processor Element (PPE)
- 8 Synergistic Processor Elements (SPE)
- 204.8 (102.4) GFlops SP (DP)
- ~ 90W

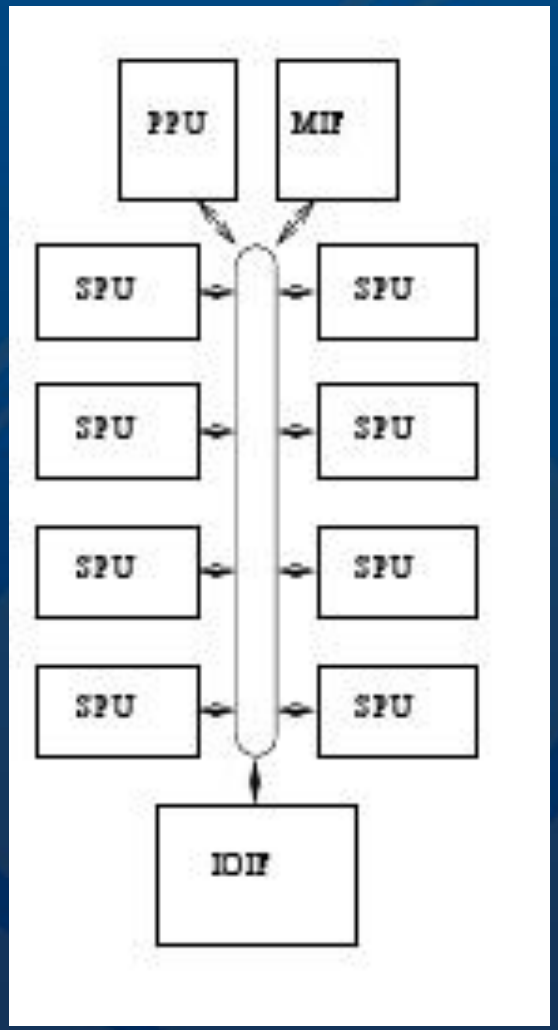

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

GPU

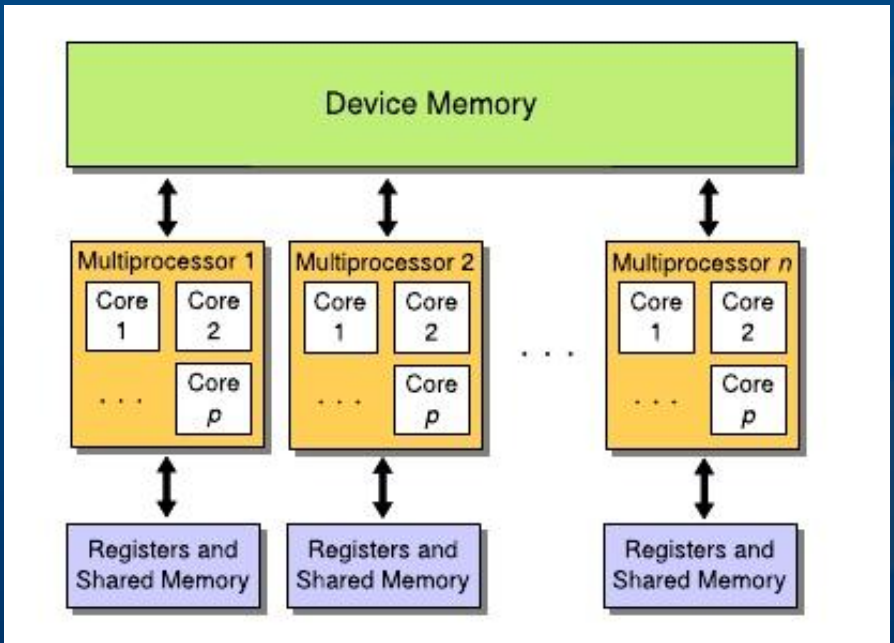

- 30 multiprocessors consisting of a total of 240 cores
- Compute Unified Device Architecture (CUDA)
- ~ 150W

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

FPGA

- Field Programmable Gate Array
- Circuit design can be loaded
- Adapt architecture to program
- Require hardware design skills (?)
- Still following Moore's Law
- ~ 25W

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

Success stories in Lattice QCD

- QPACE
  - FPGA for network
  - Cell for computation
  - 25% of peak

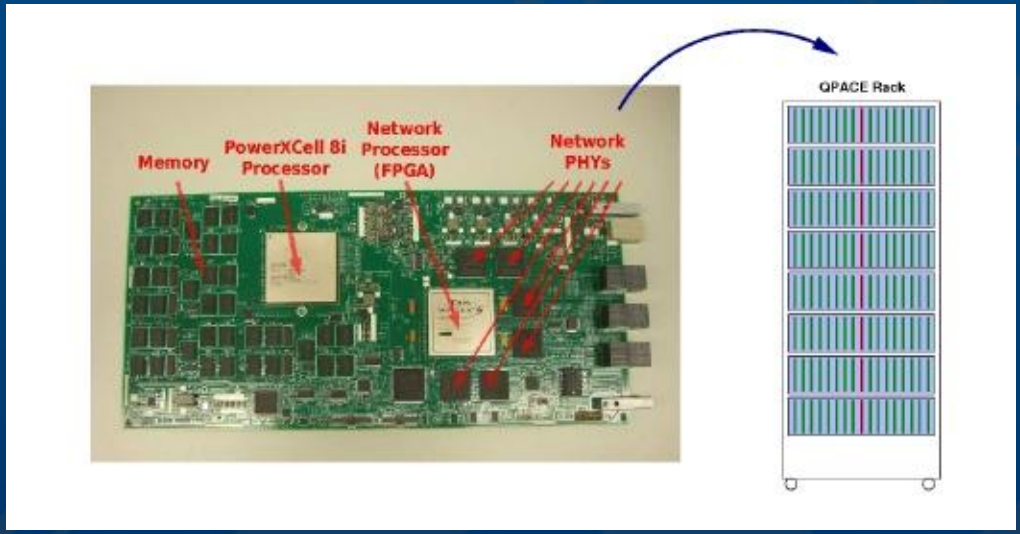

- Boston University (one of many)
  - Data layouts, local storage
  - 80 GFlops sustained (compared to 5GFlops for CPU)
  - $450 per board

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

Evolution of Accelerators

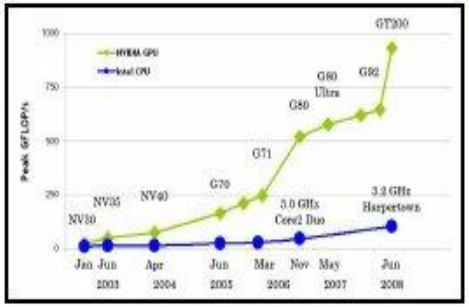

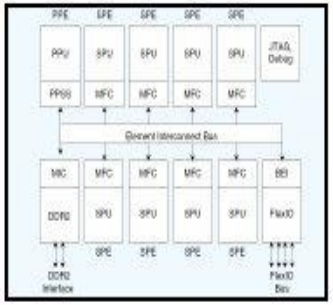

IBM PowerXCell 8i
>0.1 Tflops DP

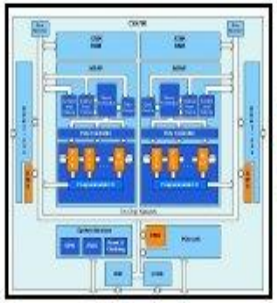

ClearSpeed CSX700
0,1 Tflops DP

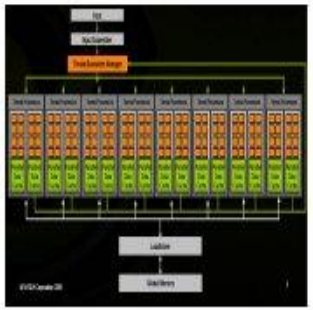

Nvidia Tesla T10
1 Tflops SP >0.1 Tflops DP
AMD FireStream 9250
1 Tflops SP, 0.2 Tflops DP

Programmability

Outstanding performance/price and performance/electricity ratio for well suited and programmed applications

GPU
- Evolving toward general-purpose computing
- Addressing the HPC market (data-parallel programming)

FPGA
- Less flexible but best performance/watt

Fully programmable

Partially programmable

Fixed function

GPU
FPGA

Performance

Based on an Intel presentation at SIGGRAPH 2008

Beginning the Project

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

Beginning the Project

- Understand the system
  - Heterogeneity
  - Memory (at all levels)
  - Network

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

Beginning the Project

- Understand the system
  - Heterogeneity
  - Memory (at all levels)
  - Network
- Use all the available tools
  - Ease of development
  - Ease of management

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

Creating the Project

- Language
  - Programming (PGAS languages, hybrid programming)
  - Accelerator (CUDA,RapidMind,OpenCL,…)
- Algorithm
  - Software technology doubles speed every 6 years
- Libraries
  - BLAS, LAPACK, FFTW…

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

Controlling the project

- Version control
  - CVS, SVN
- Debugging
  - Alinea DDT, TotalView
- Compiling
  - PGI, Pathscale, GCC, Intel, IBM,…
- Profiling
  - Gprof, Vtune, Intel Trace Analyzer,…
- Documentation

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

Implementing the Project

- GNU/Linux
- Scripting
  - Perl, Python, RUBY,…
- Job submission and queue management
- Fault tolerance and checkpointing
- Visualization
  - Parallel I/O
  - Paraview, OpenDX, AVS,…

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

Conclusions

- Technology is advancing

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

Conclusions

- Technology is advancing and changing

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

Conclusions

- Technology is advancing and changing
- With access to facilities, cutting-edge research is achievable anywhere

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

Conclusions

- Technology is advancing and changing
- With access to facilities, cutting-edge research is achievable anywhere
- Must attempt to understand architecture
- Must attempt to use available tools

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

Conclusions

- Technology is advancing and changing
- With access to facilities, cutting-edge research is achievable anywhere
- Must attempt to understand architecture
- Must attempt to use available tools
- Must exhibit capacity to exploit capability

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09

The End

Ίδρυμα Προώθησης Έρευνας
Acknowledgements:
This work was supported by the Research Promotion Foundation of Cyprus with grant number ΠΡΟΣΕΛΚΥΣΗ/ΠΡΟΝΕ0308/09