Modelling Summit for Climate Prediction

Trends in High Performance Computing

Walter.Zwieflhofer@ecmwf.int

ECMWF

The world according to Brian

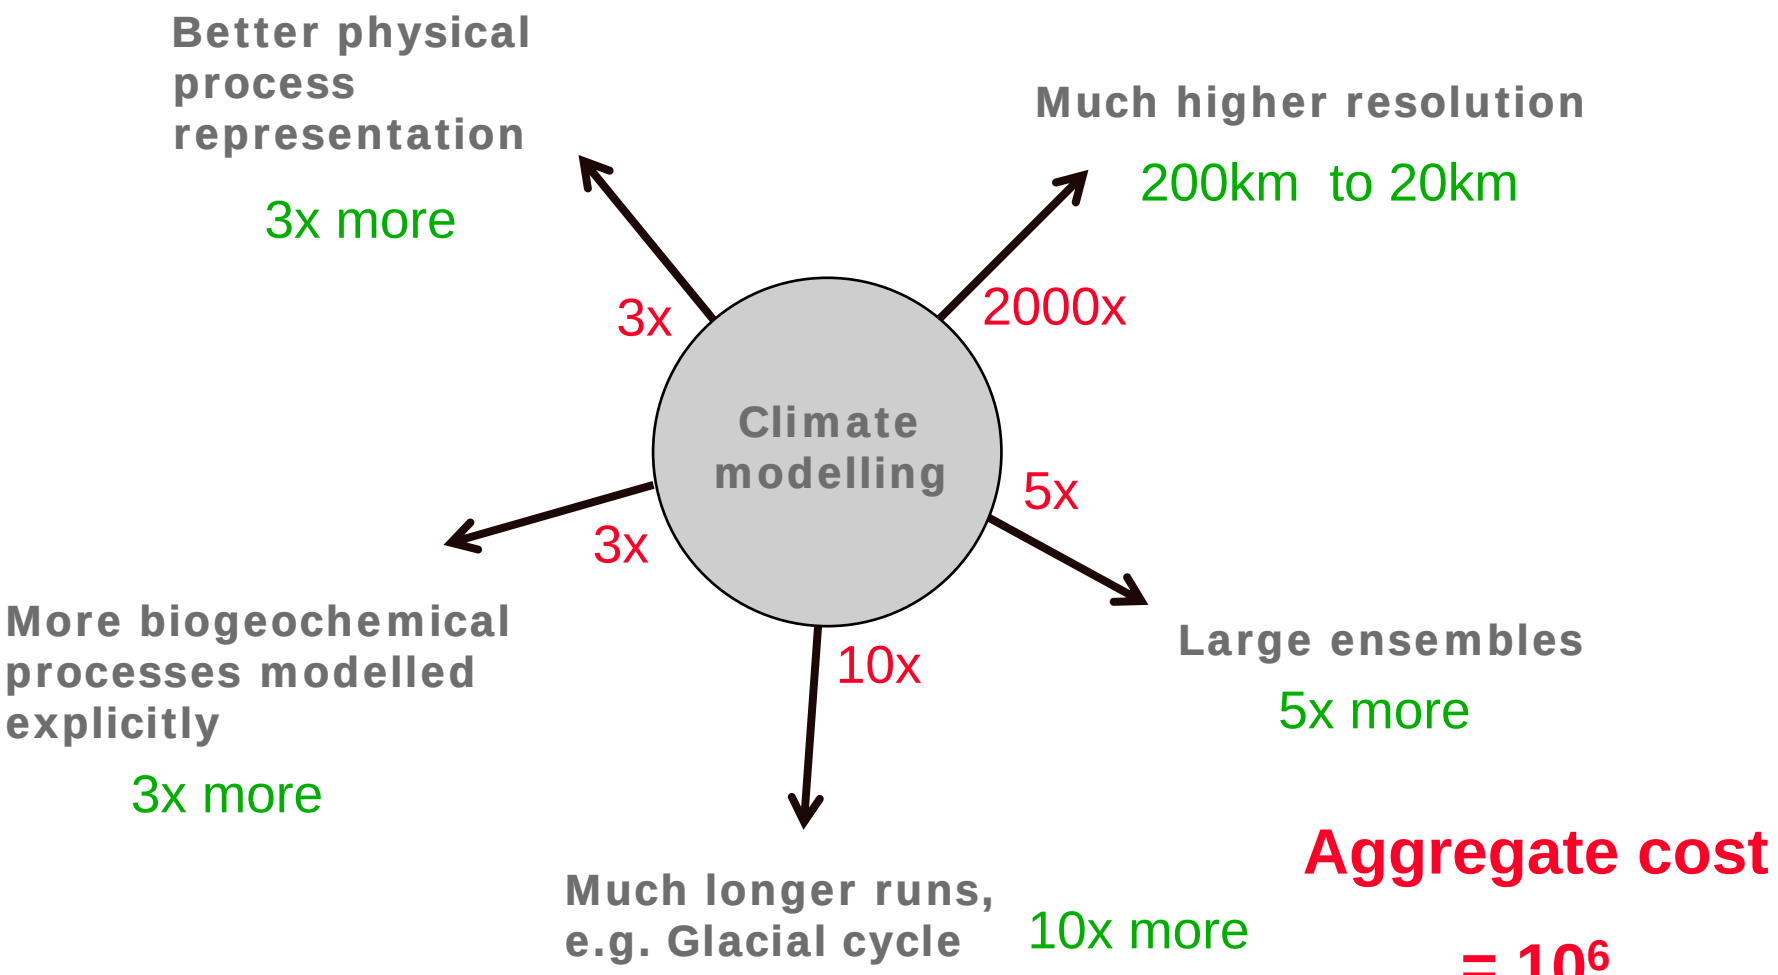

Projected Performance Development

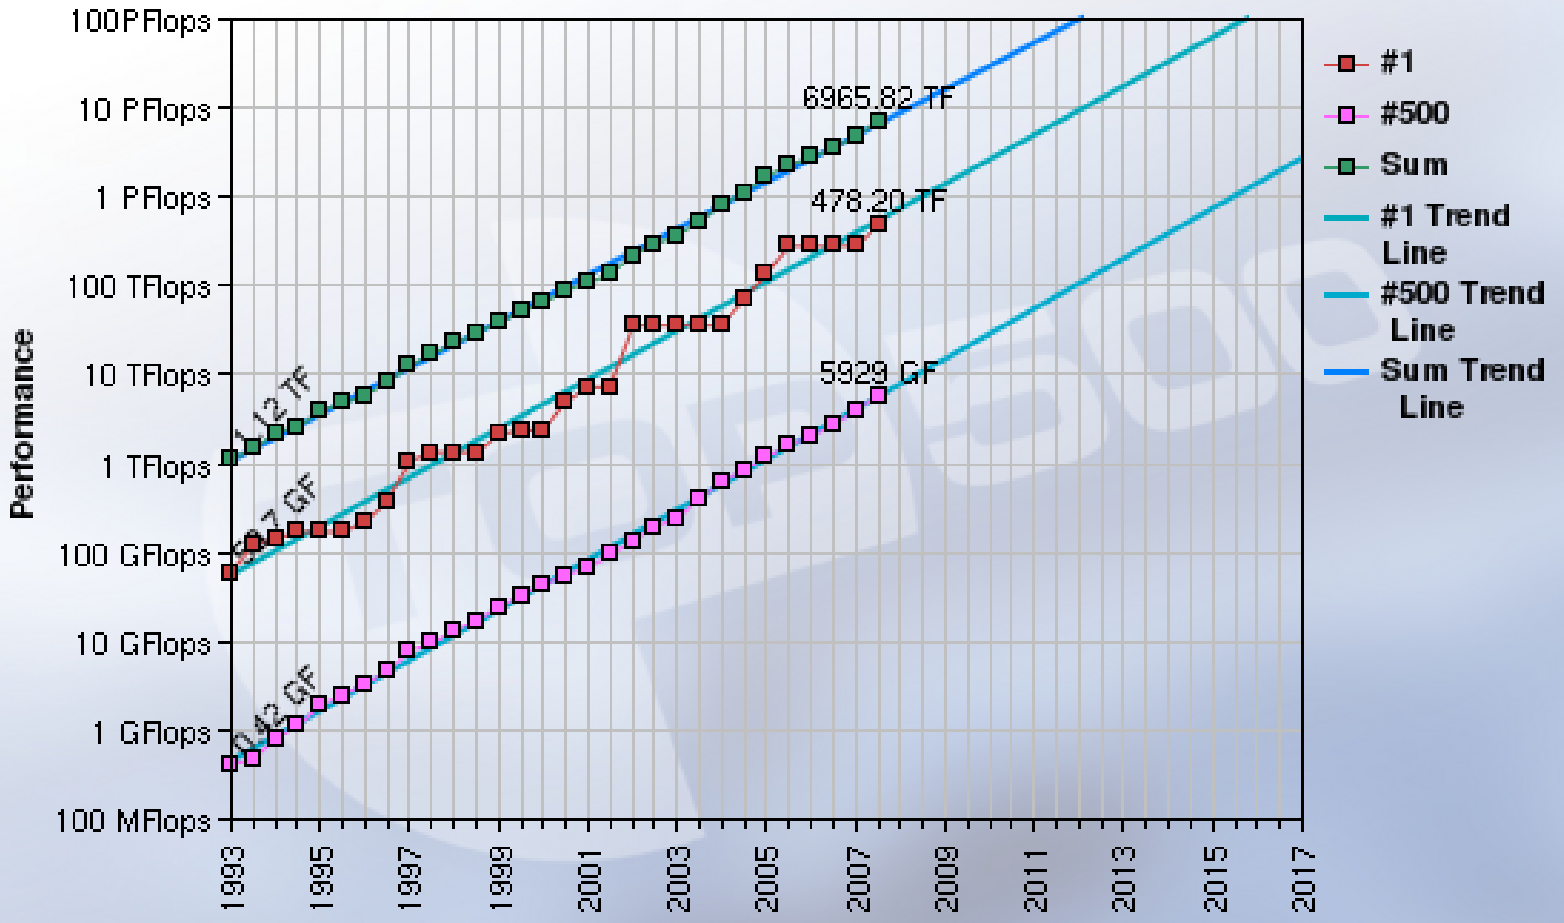

08/11/2007 http://www.top500.org/

How did HPC systems get faster?

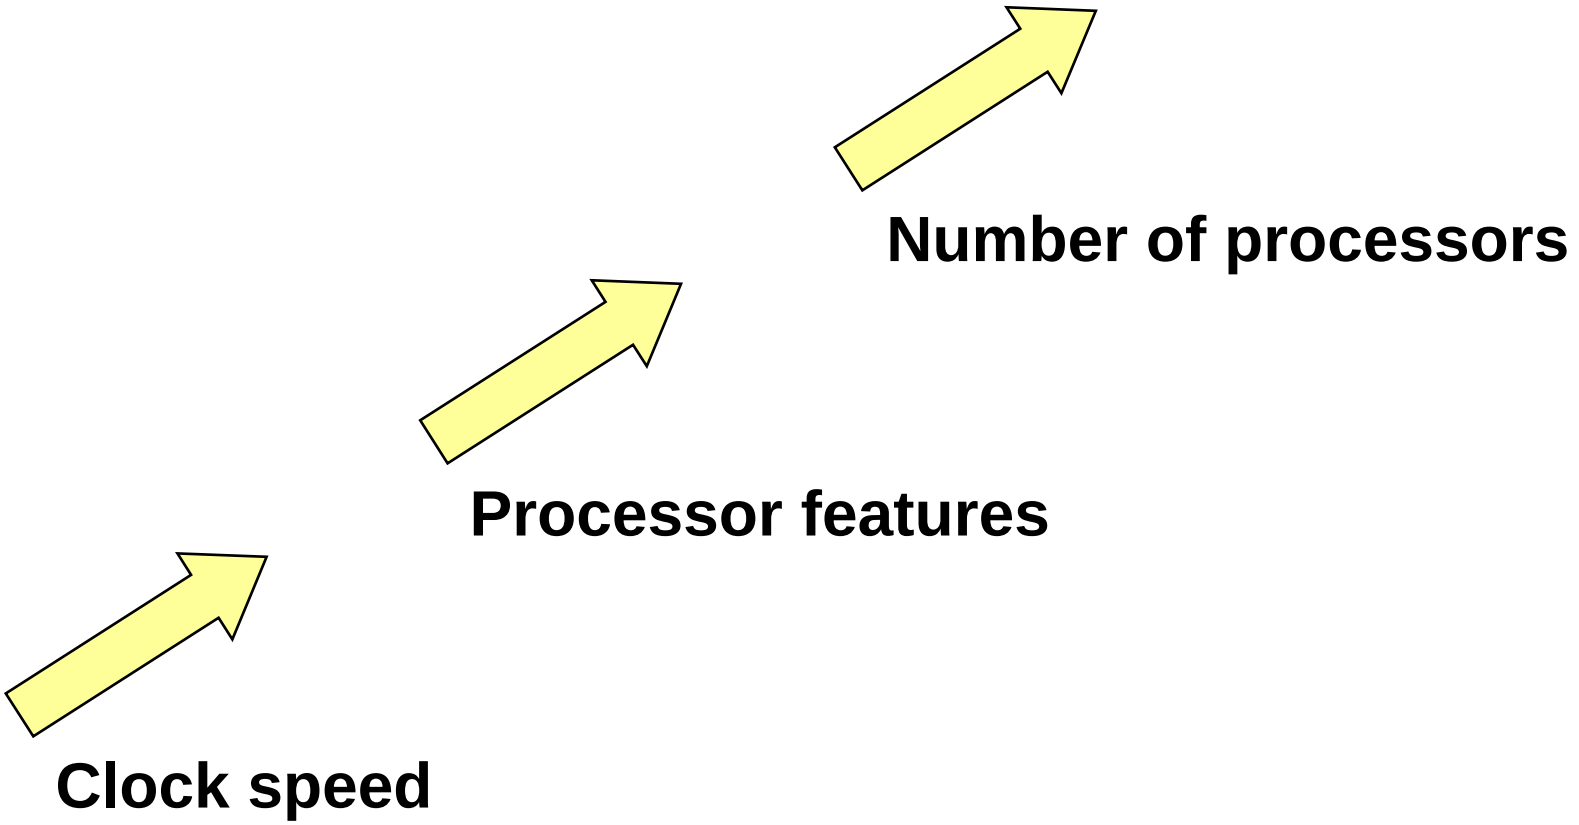

Power became the overriding issue

- **Shrinking of the feature size**
  - **Key to performance improvements**
  - **Processor chip with more than 2 billion transistors by end of 2008**
- **Power density**
  - **Amount of electrical power dissipated per unit area**
  - **Typical chip exceeds that of a hotplate**
- **Why has it been increasing?**
  - **More transistors per area (dynamical load increases)**
  - **Higher frequencies**
  - **Leakage currents (smaller feature sizes)**
- **Vendors have made a huge progress in addressing this**

Outlook .....

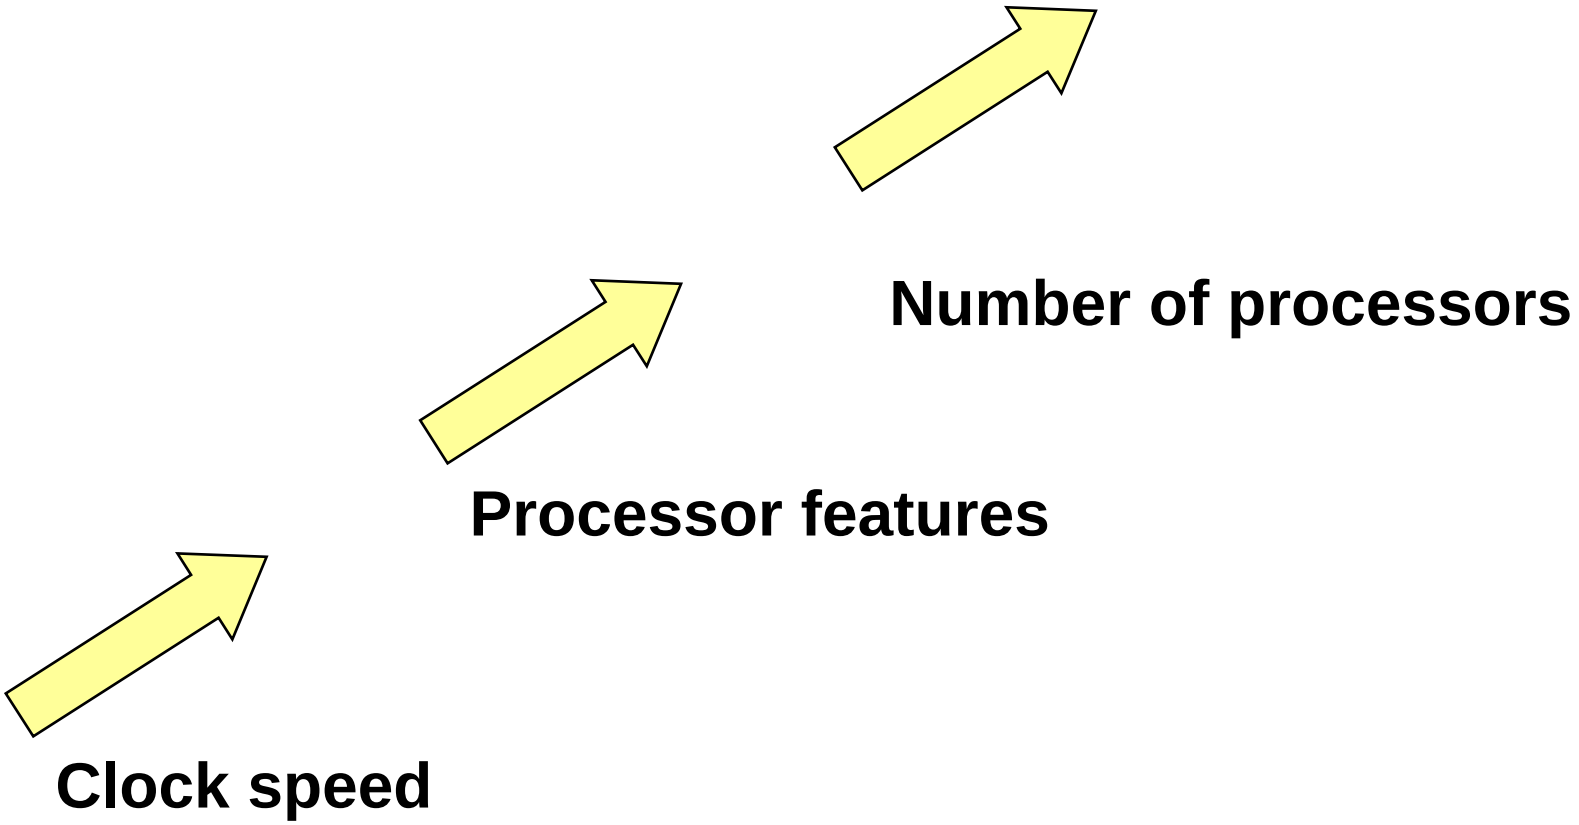

Multi-core

- **IBM POWER4 processor was dual-core already in 2001**
- **Plenty of quad-core chips in consumer desktops (Intel and AMD)**
- **Intel released information about an 80-core prototype**
- **Good application scalability will become much harder to achieve**

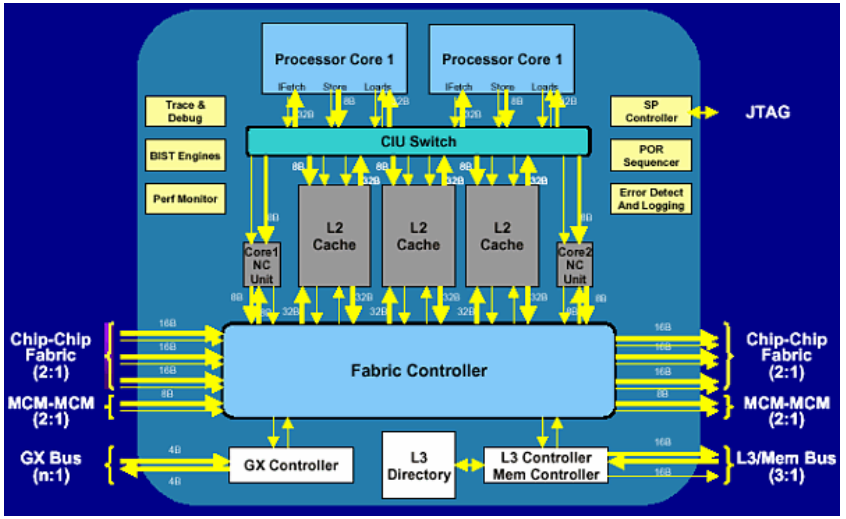

www.ibm.com

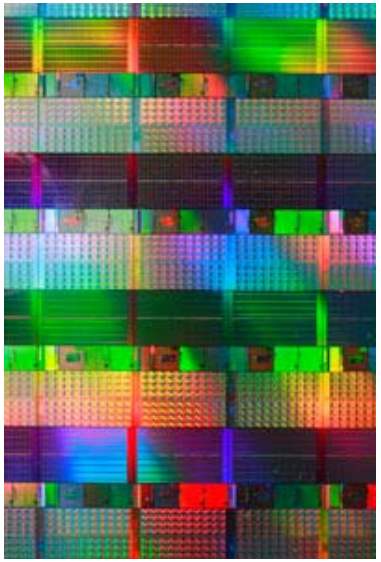

www.intel.com/pressroom

ECMWF

Types of accelerators

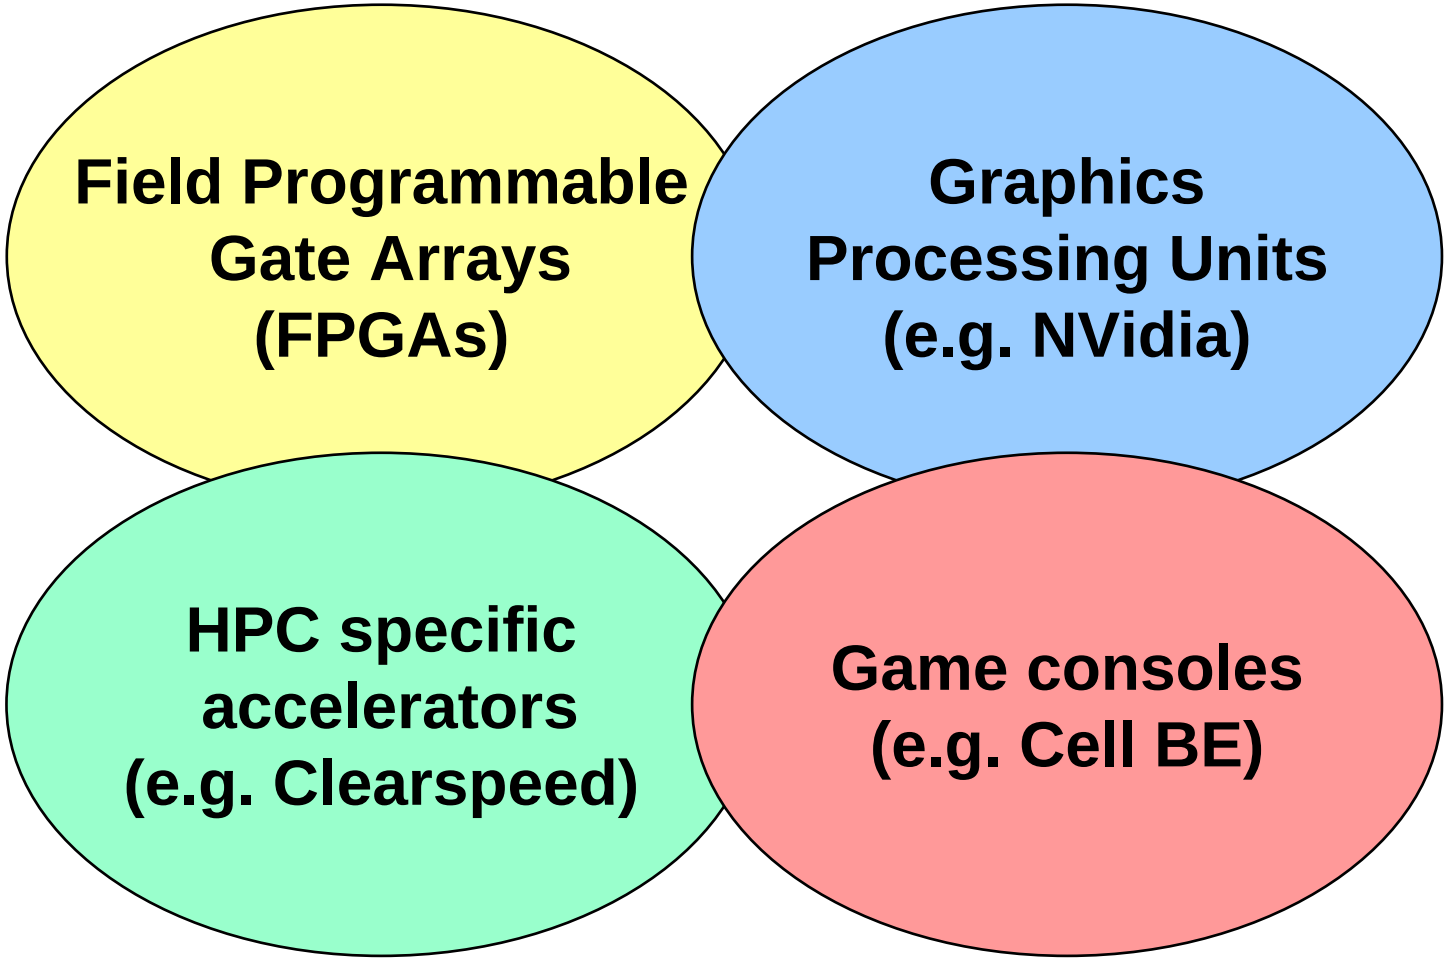

Many improvements are in the pipeline

- **First petaflop system (LINPACK) is close**
- **US High Productivity Computing Systems (HPCS) programme is funding improvements in programming language and tools**
- **Vendors are integrating accelerators and general purpose processors via the memory subsystem**
- **System scalability is being systematically addressed**

- **But application scalability has to be addressed by the user community**
- **Strong cooperation with other user communities might be more beneficial than ever before**

NCAR Aquaplanet experiment (96k cores)

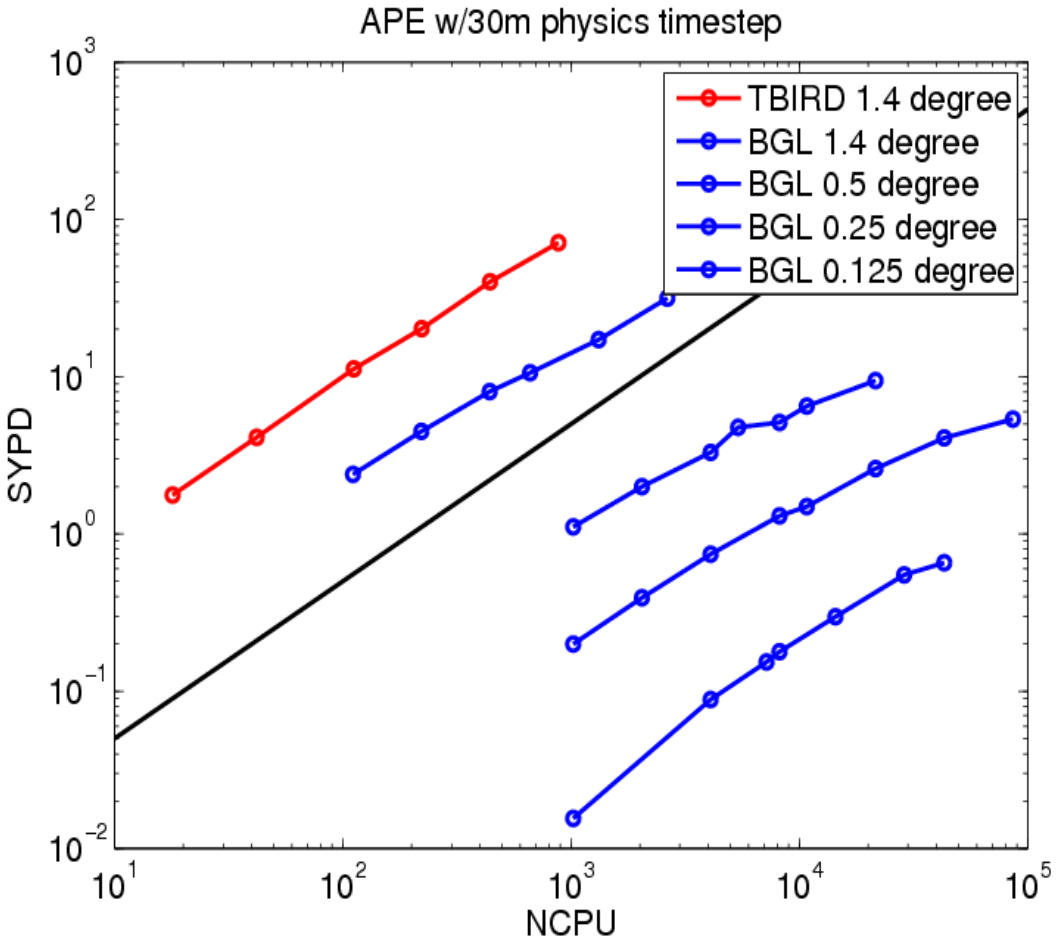

Source: NCAR with thanks to Mark Taylor

Potential for customized climate systems

- **What is the potential for building systems that excel in running earth system codes?**

- **Options for customizing the hardware**
  - **Library call level**
  - **Application code level**
  - **System level**

Library call level

- **The easiest to implement**
- **Accelerators are invoked via standard library calls**

IFS T799L91 10-day forecast

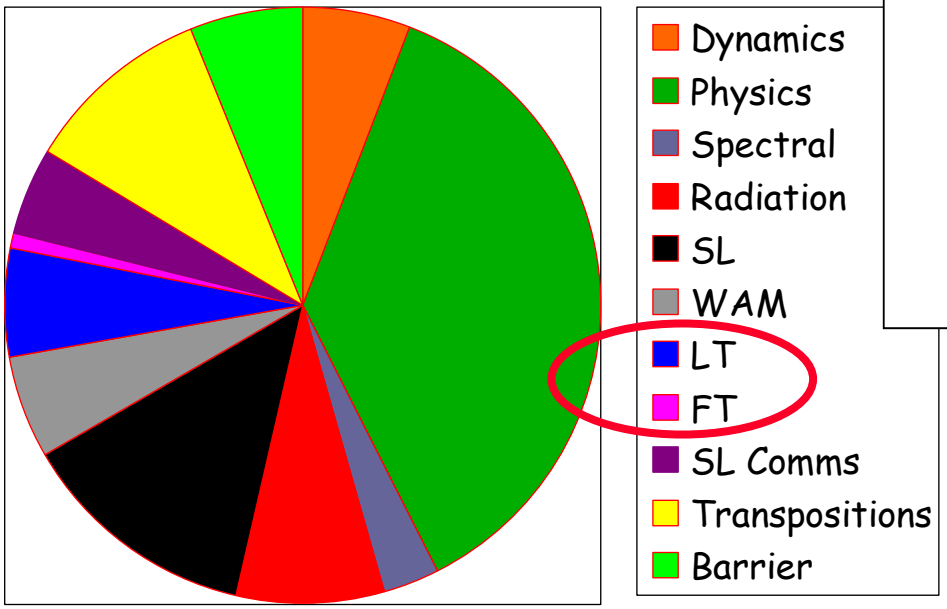

AEFS on the Earth Simulator

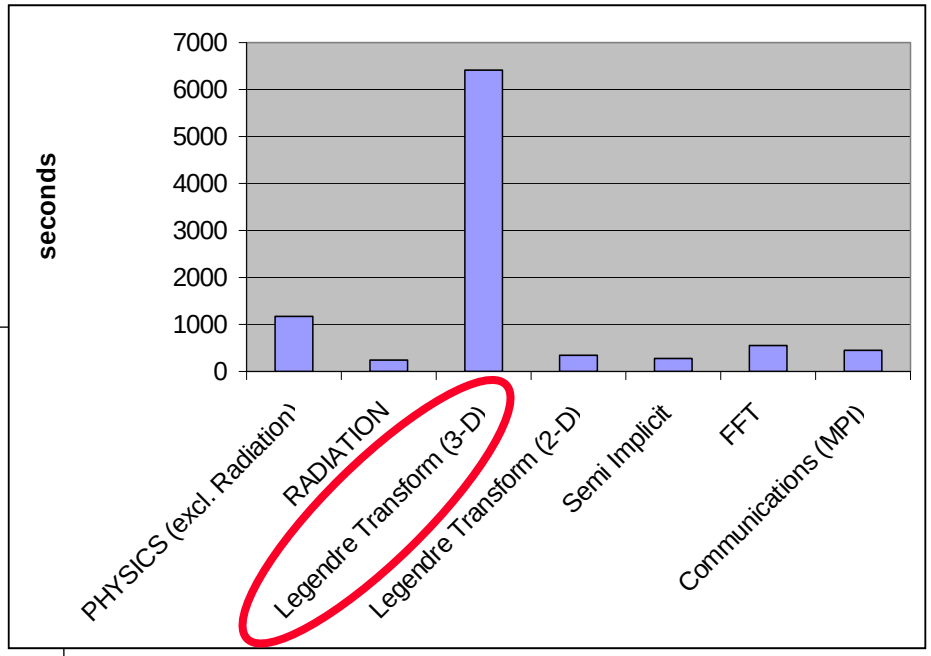

Source: Satoru Shingu et al. "A 26.58 Tflops Global Atmospheric Simulation with the Spectral Transform Method on the Earth Simulator"

ECMWF

Application code level

- **To extract the performance gain, code segments need to be carefully mapped into the specialised hardware**
- **Some of the applications that clearly benefit from it are**
  - **Cryptography**
  - **Image processing**
  - **Seismic processing**
- **These applications have very stable algorithms**
- **The code base of earth system models is much more dynamic – and there are often no dominant routines**
- **Accelerators and specialised processors have typically more complex programming models**

WRF12km CONUS Benchmark

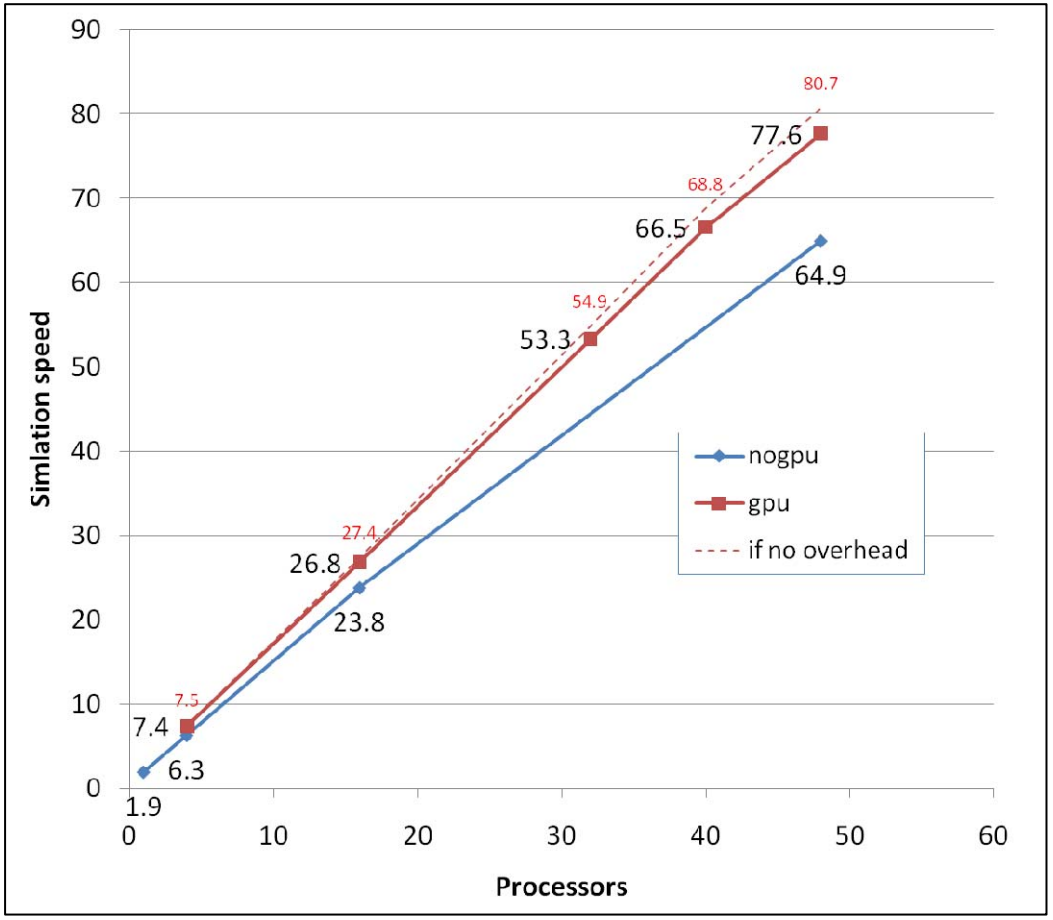

- **16 Dual dual-core 2.4 GHz Opteron nodes, each with 4 NVIDIA 5600 GTX GPUs**
- **Cloud microphysics was off-loaded**
- **Demonstrates that it can be done**

Source: NCAR with thanks to John Michalakes, Wen-mei Hwu, John Stone, and Jeremy Enos

Customization at the system level

- **Important to find the right balance**
  - **Processing capacity**
  - **Memory bandwidth**
  - **Interconnect bandwidth**
- **In particular vector machines have a different balance, but as well a different cost model**
- **Access to memory is getting worse**
  - **Memory is getting further and further away (in cycle times)**
  - **Bandwidth into and out of the chip is limited**
  - **Optical chip-to-chip communication**
  - **Algorithmic enhancements to shift the balance of memory references to local memory would be ideal**

Summary

- **Multi-core is the foreseeable future and the number of cores will grow much faster than ever before**
- **Access to memory will remain the biggest bottleneck**
- **Accelerators and heterogeneous architectures will become more prevalent, but not necessarily the choice for earth system modelling**
- **Efforts put into software engineering and algorithmic improvements must significantly increase – otherwise the next generations of machines will not be used efficiently**