




At the Bleeding Edge: Novel Architectures for High Performance

Richard Kenway


- Edinburgh parallel computing began in 1982



- novel architectures beat Moore's Law

technology push...

- squeeze the physics into the architecture




- forecast behaviour of complex/multi-scale systems
 - predict the future
 - venture beyond direct experience
- confront theory and experiment
 - discover new science
- science is now dependent on simulation
 - Moore's Law is too slow
- drivers
 - better resolution
 - faster solution
 - better models / multidisciplinary science
 - cost-effectiveness
- only architectures that address these drivers will survive



Tier 1 = national




Tier 0 = international



- the definition of capability is application dependent
 - capability vs capacity is a false dichotomy
 - the next phase will integrate HPC and the grid

- every phenomenon can be described by a model that depends on physics over a finite range of scales



- the goal of simulation is to span the range from L_n^{low} to L_n^{high}
 - cost $\propto (L_n^{\text{high}} / L_n^{\text{low}})^{\text{dimension}} + \text{physics} + \text{algorithm}$
 - in practice computer power may be insufficient

(international)
topical machines
+
cool technology



- software persists ... machines are disposable
 - design machines to fit the software/algorithm
 - design software to ride the commodity wave
- the days of general-purpose HPC are numbered



(national)
commodity clusters
+
grid

- “In the Wissenschaftsrat’s opinion, the computer pyramid in Germany ... should be expanded by an additional level of European supercomputers ... more powerful than the national supercomputers” – **Yes!**
- “[these] should be cyclically replaced ... approximately every two to three years ... This means ... the establishment of three European supercomputers” – **No!**
- “It is a condition that the national ... centres retain their function” – **Very unlikely!**
- **topical centres + grid/commodity clusters will destabilise this cosy picture! ... and we need a much stronger case →**

- lack of rigour
 - validate models, software and hardware
 - peer review proposals and papers to higher standards
- lack of a software strategy
 - efficiency + portability + reuse ... people are relatively cheap / skills are scarce
 - open source community codes, standards for interoperability
- lack of scaling
 - cool technologies + cheaper design / fabrication cycles
 - better system integration of commodity / mass market parts
- headline seeking projects
 - set achievable expectations
 - adapt specification to evolving user needs and technological capabilities

- put software first
 - increase investment to improve standards and sustainability
 - improve productivity by re-use ... community codes
 - improve cost-effectiveness by optimisation
- develop topical centres integrated with the grid
 - adapt big projects to user needs
 - opportunity for international cooperation
 - community codes will drive the requirements
 - simulation + data will be the next step
- improve rigour
 - at all levels ... hardware, software, algorithms, models, science

software NOT hardware

science pull NOT technology push