

Future Architectures for Heterogeneous Computing

- *Heterogeneous Computing*: multiple different Hw. resources (scalar, vector, and multithreaded processors, specialized coprocessors, e.g. FPGAs); Sw. should automatically/dynamically choose the best
- IBM Roadrunner for DoE (2008): hybrid (16,000) AMD Opteron + (16,000) Cell B.E. architecture
- Cray Adaptive Computing: conventional processors + vector processors, multithreaded processors + FPGAs
- SGI Altix systems (now): Intel Itanium 2 + FPGAs module
- ClearSpeed CSX600 coprocessor (now):
 - Commodity system: combination with HP dual Xeon workstation, Sun dual Opteron workstation
 - HPC at Tokio Inst. of Technology: (>10,000) AMD Opteron + CSX600
- Intel hybrid Xeon + Itanium 2 (?)
- Most Pcs: CPU + GPU

Future Software

- In a heterogeneous future, it is likely to be difficult to have “homogeneous” software:
 - Too many different applications
 - Too many different platforms
- OpenMP and MPI the solution or are we going back in time? (Explosion of parallel programming environments by mid-90s)
- We need new tools (compilers, schedulers, profilers etc.); new programming models?
- We need to think also in terms of shared-memory space:
 - Load distribution instead of data partitioning
 - 2-D load distribution (van de Geijn)