SOS Workshop Vendor Panel

Richard Kaufmann
High Performance Computing Division

HPTC Is A Dismal Field

- The world really is conspiring against you
  - As soon as you get a better memory system…
    - They double the cores to make sure you don't have enough B/FLOP
- Two developments go against the trend
  - Directly-attached memory systems
    - Alpha EV7, as realized by Opteron
  - Forgetting to break the floating point units ☺
    - 4 F/tick in x86 processors (public as of yesterday)

Q1: LP PFLOP in 2009

- Too conservative – How about 2007?
  - ~One year later: quad core
  - ~3GHz * 4 cores * 4 FLOPS/tick = 50GF socket
  - Four sockets → 200GF Server (peak)
  - Assume ~50% efficiency…
    - 5,000 Servers = 1PF peak; 10K servers for 1PF sustained
    - Not unimaginable IB switch fabric
    - Generally useful system

Q1, cont.

- Accelerators, e.g. Clearspeed
  - Could potentially reduce the number of servers needed for a LP PF
- Leads to Q2: usable FLOPS
  - Accelerators are (for now) ways to pick low-hanging fruit
    - If it solves your problem, it's the best thing in the world
    - If not, it's evil (if only because of envy)
- MTTI issues exacerbated
  - Be mindful if accelerator technology wasn't designed at the start for reliability at a massive scale
    - FPGAs & error corrected links
    - Multimedia accelerators

Q3: The HPC market is growing…

- IDC Revenue Forecast
  - 9.4% CAGR for HPTC (Feb 2006)
  - 3.3% for Commercial (Oct 2005)

Q4: Competitors

- We love competition amongst our suppliers
  - Having three processor lines (Itanium, Opteron, Xeon) is wonderful
    - It's guaranteed that one or two of them will be dogs in any generation
    - And one of them will be the best general purpose processor, period.
- Competition amongst computer vendors is good for the users
  - Emphasis on TCO, usability, manageability

Q4, a better question

- What can the high-end community do to help themselves?
  - **Get your code into ISVs' hands**

YAVOCVCC

- **The Commodity Axis**
  - **2S1U x86 servers connected by**
    - **Ethernet (most)**
    - **Infiniband (some)**
- Outliers
  - "Other processors" in scale-out clusters
  - Large SMPs (e.g. Superdome)
    - Commercially significant for ISV-led accounts
  - Tons of parallel slower, smaller, cooler processors
  - "Low hanging fruit"
    - GPGPUs, FPGAs, ClearSpeed, etc., etc.
  - Traditional Vectors (really: scatter/gather-optimized memory systems)