

Microsoft® Research

Faculty Summit

10
YEAR ANNIVERSARY

Energy-efficient computing: Hype or Science?

Chuck Thacker
Technical Fellow
Microsoft Research Silicon Valley

My Position:

- Lots of hype, little science
- But – lots of this hype is real, and there are many challenging engineering problems
- Initially, we focus on data centers:
 - Because we have many and are building even more
 - Because we can optimize the entire center as a *system*
 - Because many of the techniques used will trickle down to desktops and below
- I'll mention several areas: the server computers, power, packaging, cooling, and networking

Server Computers

- Today's CPUs are the result of an evolutionary experiment in extracting highest possible instruction-level parallelism
 - Much of the logic and power in a CPU are there to do this
 - Speculation and prediction improves speed, but failed prediction and speculation wastes energy
- Using lower-powered, simple CPUs can work very well
 - Laptop-class processors are good for I/O intensive workloads
- We need to add Flash and other nonvolatile technologies at the right place in the storage hierarchy
 - Lots of hype here, but also some real progress
- We need to offload computation inappropriate for the "Swiss army knife" CPUs to GPUs, FPGAs

Power and Cooling

- We need to reduce the Joules/operation, not just provide ever-more-exotic cooling technologies
 - Air -> water works well, and is inexpensive
 - It's even better if the water is used once
- Must look at every step in the power distribution chain:
 - AC transformers (99% efficient)
 - AC – DC converters (85-90% efficient)
 - DC – DC converters (90+% efficient)
- We need to do better at energy proportionality
 - Why should an idle server consume 40% of its peak power?

Packaging

- We need packaging optimized for efficient cooling, high density, and low cost
- Shipping containers are a good starting point:
 - Don't need to ship and assemble individual computers
 - Worldwide shipping infrastructure exists
 - Don't need complex buildings
 - They provide a reasonable level of incremental expansion:
 - ~1200 U per 40' container
 - 100 – 200 KW per 40' container
 - May not need to open them during their lifetime
- Server packaging changes when containers are used:
 - Why do servers have fans?
 - What's the steel box for?

Networking

- Networking is the last bastion of non-commodity hardware in a data center; High cost, lots of headaches
- Today's networks have evolved to ship lots of data in the presence of lots of adversity:
 - High error rates, unknown topology, hackers, back-hoes
- Data centers don't have these problems:
 - We know the topology, since we built it
 - The scale is limited to ~100K nodes
 - Security is simpler: No hackers, no back-hoes
- We have an opportunity to rethink data center networking:
 - OpenFlow is a good start (www.openflowswitch.org)
 - More audacious: source-routed, stateless switches/routers

Microsoft[®]

Your potential. Our passion.[™]