



The future of supercomputing

Greg Astfalk



Prologue



- No inference of future HP products, or lack of, should be made from the material in this talk
- It is completely devoid of vendor-isms

The "c"s



- Future (super)computing is about four "c"s
 - computation
 - communication
 - content (aka data)
 - c (the speed of light "c")
- This talk will touch on each of these

Controversial statement #1



- For the vast majority of scientific computing today the computer industry is providing good solutions
 - Cost, power and volume of a flop is dropping
 - Cost and volume of a byte is dropping
 - Cost of moving a byte is dropping
 - Speed of moving a byte is increasing
 - Absolute delivered performance is generally increasing
- At the high-end
 - Management is growing more complex
 - Reliability is less than desired
 - Efficiencies are too low

Controversial statement #2



- We do not really need more powerful processors
- Need better ways to supply the processors we have
 - Data and instructions
- Need better ways to produce more efficient code

(in)Efficiencies



- There exist three major inefficiencies
 - Instruction-level efficiency
 - Average IPC is ~1.3 out of 4
 - Percentage of processor peak
 - Rule of thumb: 10%
 - Message-passing overheads
 - Large latencies and low bandwidths
 - Decreasing efficiency (or speedup) with increasing scale
- These additively make delivered efficiency quite low

- To a close approximation we know what the processor of +5 years will be
 - It will have a market life of several years
- If it's follow-on were to represent a significant technology shift, it would be in the straw-man or research stages now
- By obvious induction the processors we will have in 10 years are going to be largely evolutionary along the current trajectory

- Herein lies **the** problem
- Memory is not getting faster fast enough
 - Access time has been essentially constant for a decade
 - Processor cycle times are increasing Moore's law-ish
- Latency, measured in processor cycles, is getting worse
- For enough money we can provide sufficient bandwidth for moderately sized systems with contemporary designs

- What holds some promise is the idea of placing processing and memory in very close proximity
- PIM (Processor In Memory) can achieve much better memory characteristics
- PIM is not a panacea however
 - Limited amount of local memory for the “P”
 - Volumes needed to get the cost down
 - Programming/algorithmic changes necessary

- We will always have interconnects since we will always have “parallel” systems
- The latency of inter-system interconnects will not likely get low enough to not be a impediment
- Fully optical interconnects may seem like a solution but it is a complicated problem to utilize optics
 - Just when optics seem better Si and Cu (pleasantly) surprise US

- The relevant way to measure latency is in processor clocks
- In the past the latency problem was “solvable”
- Today the latency problem is not, strictly speaking, solvable
- Consider two length scales
 - Intra- and inter-system
 - see next slide
 - Continental time-of-flight
 - $O(10^2)$ milliseconds

Latency hierarchy



- Consider this approximate latency hierarchy for contemporary architectures
- Disturbing is that these numbers grow worse over time
- More emphasis needed on mitigating the latency effects

Hierarchy	Processor cycles
Register	1
L1 cache	2–4
L2 cache	6–20
L3 cache	15–30
Near memory	80–120
Far memory	100–300
Remote memory	$O(10^2)$
Message passing	$O(10^3)$ – $O(10^4)$

Programming languages



- Current languages are at too low a level of abstraction
- The negative inertia to have the community embrace a new language is enormous
- End-users will adopt a new language only if it offers a large ($>2x$) performance or productivity gain
 - High-end users would be among the most willing to use such a new language

- Over time the centroid of (super)computing is going to be the content (i.e., data)
- The content issues include
 - Storing and moving
 - Gathering and accessing
 - Finding, sharing and securing
 - Rights management
 - Knowledge extraction

Petaflops (controversial statement #3?)



- Seriously contemplating petaflops today is a misuse of people's time and energy
- We hardly know how to efficiently use a O(10) teraflops system today
 - We achieve single-digit percentages of its peak
- With no revolutionary processor or architecture designs being considered what outcome is expected when scaled up 100x?

- The US government is currently reaping the (anti)benefit of not funding architectural research to any significant degree for at least the past decade
- Industry can't do the fundamental research that is necessary to develop **new** computer architectures
- A different funding model is needed to fill the pipeline with ideas and people
 - Large, multi-institution program with guaranteed funding over a long time-horizon
 - DARPA's HPCS and the IHEC program are good ideas



i n v e n t