Over the horizon architecture/hardware

But not TOO over the horizon

Chairman's Note: This document was prepared by the "future hardware" working group and was received by the entire workshop in plenary session without modification.

Goal

- To inform Gov initiatives, investment
- To pose the right question
- Does not require consensus

Point of Departure

- For the last ten years, the presumption was that Gov would leverage the advance in desktop microprocessors, and scale them up to provide compute resources for Government mission applications
- Is it time to revisit this assumption?

Comments from Monday session

- End user experience will drive the computer industry
- Killer ap
  - Virtual life
  - 3D holographic telepresence
- Instead of the predicted Uber chip, we currently are seeing domain chips. Specialization versus generalization
  - nVidia – graphics
  - Physics chips

Comments from Monday session

- Can't just ride coat tails of uniprocessors; architectures have to prove themselves
- People will have to write parallel code – what's the programming model that enables that? (multi-cores, etc)
- Reward system is structured to reward incremental improvement
  - Exception: DE Shaw – has 10-15 year horizon for their architecture design
- Have to achieve price-parity for most apps in order for an architecture to succeed

Comments from Monday session

- Question: Why aren't we exploiting commodity embedded processors? Doing it more?
- Question: Will AMD/Intel market drivers (business needs) diverge so much (over next ten years) from what science/gov needs, will it justify investment for government apps (climate change, etc)

Comments from Monday session

- What is the Adv Comp Architecture ten years from now?
  - Beowulfs will take care of themselves
- Observation: 3 classes of arch investment
- 1) clusters. Do nothing
- 2) hybrids (like DSB report)
  - Value added: ex: Red Storm, BGL
- 3) Custom machines
  - Designed for specialized purpose (ex: GRAPE)
- 4) custom designed GP machines
  - X1

Comments from Monday session

- Future trend? May be more embedded memory on chip, which then starts to optimize for density rather than stressing performance
- Question: does it make sense for Government to support long range academic R&D in comp architecture?
- There's a real restriction on type of work funded in academic community (ex: try to help Intel products in next 2-3 years. When Intel decides five years ahead what arch will be) One could argue that for IC technology (silicon, beyond silicon), there really is a problem with too short term academic research focus

Comments from Monday session

- Currently, only one arch design with large team – TRIPS. Other PCA project is MONARCH, an embedded design
- Used to be 5-6 teams/communities
- Lots of arch ideas, but very few explored in the build phase

Tues - Observation

- Revisiting the GP/LD/SPD taxonomy
- GP – generally clusters. Solves a lot of problems. Large customer base. Can take care of itself
- SPDs – totally driven by application requirements. Cost is born by customer, who does this when it's cost effective
- Topical Centers – has been emerging as a strategy since ~2000
  - Red Storm
  - BGL
  - Cyclops
  - Requires large industrial partner (Cray, IBM)

Tues - Observation

- LPDs are happening. Summed over Fed Gov, takes large resource. Software is very immature on systems
- Do we expect continuing creation of LPDs? If so, what can be done to help?
- Would reinvigorating academic arch community help inform LPD design?

Tues – comments

- Sterling – there is a next arch
  - If we call the current thing the MPI architectures (and before that, the vector/SIMD, and before that sequential machines)
  - With possible exception of fine grain, …
  - Implemented with custom design pieces aimed for general purpose

Tues – comments

- Use simple designs that enable complex behavior
- Universal agreement – we need to re-establish sustained funding support for more than one community of comp architectures so they can explore architectures that are more than incremental improvements to the vendors next product

- Try to understand basic requirements of apps (data access patterns)
- Build machine that satisfies those requirements

- Hardware should be able to adapt to computation (virtualization)
- Simple things should be simple. Current hardware makes simple things complex

- There IS no software silver bullet
- You want to be able to express the application as high level as possible
- Hardware and software should minimize the number of hoops the programmer has to jump through
- Software should enable expressability
- Hardware should support the software

Recommendation

- We need to declare sets of requirements
  - Mike's examples of accumulation over tree structures
- Allow freedom of creativity
- Fund it

- The current approach is painful
- We've hit a ceiling
- Can fool around with nVidia, cell, and invest enormous programming effort that we'll get something, but …this cherry picks the problems and lots of other things don't get done
- Need to use different approach to break the ceiling
- We have most of the ideas
- Need gov support to mature the ideas, to cause the transition. The non-linear changes need Gov investment

Tues comments

- Follow through on the things we said we were going to do
  - Original HPCS plan
  - NASA program two years ago (drew proposals, but no funding for them)

Findings

- Arch research pipeline ~ empty
  - DSB and NRC studies
- Barrier to entry for customization is approachable (~$20M)
  - could be lowered by doing tradeoffs and using things like structured ASICs
- Opportunity to exploit specializtion is growing (GPUs, Clearspeed, MTA, etc)

Findings

- Architecture is principle obstruction to better parallel algorithms/software/apps
- There doesn't exist a general parallel arch
  - Can't satisfy all "balances" with 1 arch
- Gov is not funding the creation of building blocks of true parallel systems
  - Locally sequential processors
  - Modest counter example: BG/L barrier

Findings

- Effectively, we haven't explored new paradigms for ten years (since 1994/95)
- Programming models/execution models/architecture models are based on trade-offs that evolve as a function of time and technology
  - Need to explore multiple paradigms again
  - Need to continually explore new ideas and revisit old ones

- Findings

- Exhausted pipelining and ILP
  - Flat lined on conventional arch
  - Freq growth slowing 17%
  - What sustained bulk of growth for 50 years has neared end
  - Multi-core offers more challenges than functions
- Future Challenges/opportunities
  - Concurrency
  - New technology

- Findings

- We anticipate some of the new technologies can contribute to end-user system (5-15 years)

- Findings

- Reliability needs to be architected in to enable successful use of deep submicron VLSI
  - Receiving lip service but little action
- General purpose or at least general components necessary for cost-effective S/W
- Quantum Computing will not address meaningful problems in the next decade

Recommendations

- Money!
- Universal agreement – we need to re-establish sustained funding support for more than one community of comp architectures so they can explore architectures that are more than incremental improvements to the vendors next product
  - *This happens to be NRC/Recommendation #6 and also appears in two recent DSB studies*

Recommendations

- "Gov agencies responsible for supercomputing should underwrite a community effort to develop and maintain a roadmap"
  - *This is NRC Study Recommendation 5*
  - ***We can do this ourselves!***

Open Issues

- Can we get away from using commodity memories?
  - Rambus model of incremental change
  - True renaissance {PIM -> MIND -> Continuum}
  - At what price & volume could one get a custom DRAM?
- Will any of the novel technologies we saw have impact by 2015, 2020?
  - HP – UCLA
  - Nanotube memories
  - MRAM
  - SFQ (memory density is issue)

Participants

- Candy Culhane
- Bob Lucas
- Steve Scott
- Doug Burger
- Peter Zeitzoff
- Thomas Sterling
- Guang Gao
- Larry Bergman
- Mike Merrill