Microsoft Feedback

Burton Smith

Origins

- Intel and Microsoft were worried about client parallelism research.
  - The academic computing community seemed uninterested.
  - US Government agencies were similarly indifferent.
  - Creation of intellectual property was not a primary objective.
- The companies agreed to jointly fund a flagship effort.
  - Proposals were solicited from major US computing departments.
  - We visited four universities and picked Berkeley.
  - Later, we added the University of Illinois at Urbana-Champaign.
- Our intent was to stimulate the whole field.
  - Some said the government could let Microsoft and Intel do it.
  - Still, we have seen some progress within both academia and industry.

Why Berkeley?

- The relocation of faculty and students into a single space
- The precept of using applications to drive the research
- Coverage of the whole stack, clear down to hardware
- A great history of prediction and influence
- A demonstrated enthusiasm for the subject

Achievements

- We have heard about the research achievements already.
  - Our work is barely begun, but begun well.
- The educational achievements are equally impressive.
  - I have never before seen graduate students to match these.
  - They have become agents of change in computing.
- Community outreach has been extraordinary
  - Parallel Programming and Chisel boot camps
  - Patterns evangelism
  - RAMP leadership

Regrets

- Exclusion of client plus cloud as a topic
  - We were worried about pollution by HPC applications.
  - In retrospect, maybe we worried needlessly.
- Insufficient engagement with Microsoft product groups
  - Windows has been an important exception.

A Personal Perspective

- My abiding interest has been general purpose parallel computing.
  - I give a talk with that title about once per decade.
- In 1980, I was certain the need would be obvious by 1990 or so.
  - I was off by some twenty years, thanks to instruction-level parallelism.
  - Some would say the need isn't even obvious today, but I disagree.
- Unfortunately, much of what we do today needs to change.
  - Fortunately, many of us agree about what changes are needed.
  - Maybe, as Max Planck said, "Science progresses one funeral at a time."

Ahead

- ASPIRE at Berkeley
  - Optimal performance and energy efficiency through specialization
  - Supported by DARPA and Semiconductor Research Association
- Mainstream heterogeneity in both client and cloud systems
  - Big/little and CPU/GPU, at least for the moment
  - Lots of operating system work is needed
- A fluid situation in programming languages
  - Parallelism-centric? Browser-centric? Service-centric?

Conclusions

- The Berkeley ParLab has been a tremendous success
- Microsoft and Intel should be proud of their role
- It has started to address the new issues facing our field