Parallelism Across the Curriculum

John E. Howland
Department of Computer Science
Trinity University
715 Stadium Drive
San Antonio, Texas 78212-7200
Voice: (210) 999-7364
Fax: (210) 999-7477
E-mail: jhowland@Ariel.CS.Trinity.Edu
Web: http://www.cs.trinity.edu/~jhowland/

Abstract

A sea change within the computing discipline occurred last year which is not widely recognized nor fully understood. Individual processor chips are not going to be faster in the future. The implications of this fact are discussed and changes to the undergraduate computing curriculum are proposed.

Overview of Presentation

- Background of the Problem
- Moore's Law and Changes in Computer Hardware
- Curriculum Issues
- Conclusions

Background

"... the times they are a-changin' ..." [2] : Bob Dylan

- Moore's Law: (April 19, 1965 issue of *Electronics*) ... innovations in technology would allow a doubling of the number of transistors in a given space every year ...
- Moore updated this (1975) to a doubling of the number of transistors every two years to account for the growing complexity of chips
- We have experienced this exponential growth for 40 years

Background (continued)

- Intel Gives Up on Speed Milestone (Wall Street Journal, October 15, 2004)

  > The company said it no longer plans to offer its Pentium 4 chip for desktop computers at a clock speed of four gigahertz, a target that had alread slipped from the end of this year until sometime in 2005. The fastest member of that chip family is now 3.6 gigahertz.

- Intel is now focusing on increasing cache memory and putting multiple processors on a single chip.

- While Moore's law appears to continue to hold, the more dense chips will not include the previous benefit of allowing higher clock speeds due to heat dissapation and current leakage problems.

Background (continued)

- Intel, IBM and other processor manufacturers have announced they too will be following the strategy of providing multiple processor chips which are not necessarily faster than their current chips.
- New Chips Pose a Challenge to Software Makers (Wall Street Journal, April 14, 2005, B3)
- How will our future computing systems change?
- How will our programs go faster in the future?

Background (continued)

- Parallel programming will have to become the standard approach in all of the programs we write (if we wish to go faster).
- New programming languages
- New programming methodoligies
- Changes in CS curricula

Changes to CS Education

- In the past, Moore's law has provided an environment which allows system design to assume that or require that more powerful processors will soon be available to properly execute bloated designs.
- Computer science education has contributed to this trend often by teaching techniques which are not optimal.
- Students learn to rely on ever faster processors to run their non-optimal algorithms and designs.

Computing Across the Curriculum

CS department received input from its advisory board concerning:

- effects of processor speed limits on the computing industry
- importance of parallelism in various sectors of computing

The idea of parallelism across the undergraduate computer science curriculum emerged from these discussions.

It should be noted that Computing Curricula 2001 Computer Science [1] does not address the subject matter of parallel computing being distributed across the undergraduate curriculum.

Computing Across the Curriculum (continued)

Two faculty members of our department have research specialties in parallel computing. As a result, our department already offers two elective courses (Parallel Processing and Advanced Topics in Parallel Computing) which may be taken by junior and senior computer science majors.

Computing Across the Curriculum (continued)

Our department is beginning to experiment with several short modules within several standard courses in our undergraduate curriculum. Many of the modules are short topics with demonstrations which may be covered in one or two lecture periods.

- CS 1

  Simple Threads model to illustrate two tasks computed in parallel using pthreads.

- CS 2

  Several threading examples of removing device latencies in common algorithms using the Java thread library.

Computing Across the Curriculum (continued)

- Data Structures

  Array data structures for parallel algorithms.

- Discrete Structures

  Mathematical models for synchronizing parallel tasks.

Computing Across the Curriculum (continued)

- Functional Languages

  Church-Rosser property of independence of order of evaluation.

  Function application to array/list data structures.

Computing Across the Curriculum (continued)

- Upper Division Elective Courses
  - Computer Graphics
    Parallel processing within modern raster display processors. Techniques of parallel image rendering.
  - Computer Networks
    Multiple network interfaces to increase network transfers. Multiple daemons for handling client-server requests such as serving web pages.
  - Database Systems
    Multiple daemons for handling client-server requests in database applications.

Computing Across the Curriculum (continued)

- Operating Systems
- Simulation Theory
  Application of parallel processing in scientific simulations.
- Artificial Intelligence
  Parallel algorithms for evaluating decision nets.
  Parallel algorithms in multi-agent systems.

Computing Across the Curriculum (continued)

Since two of our computer science faculty have research specialties in parallel computing we have two upper-division courses in parallel computing. These courses provide in-depth treatment of current topics and trends in parallel computing and parallel languages.

- Parallel Processing Fundamentals
- Advanced Topics in Parallel Processing

Laboratory Equipment

Parallel hardware is necessary to provide demonstration of concepts and provide laboratory experience. Fortunately, it is now relatively inexpensive to acquire such equipment as desktop systems with dual core or dual processor systems are now common. We have found that a few such machines running some dialect of Unix (Linux or OS X) are sufficient to provide laboratory experience in all of the above course modules and courses.

Laboratory Equipment (continued)

It is also possible to use lab computers or classroom computers running Linux to provide laboratory experience with parallel clusters using the mpi libraries.

An NSF-CCLI grant proposal is being prepared to fund a cluster of dual processor machines which will be used to support laboratory experiences for our parallelism across the curriculum initiative.

Conclusions

- The computing industry has an insatiable desire for increased performance.
- "Intel Corp. [6] has laid out plans to exploit two trends – parallel computing and a wireless technology known as *WiMax* – that the chip maker believes will drive electronics demand for years."
- "Intel Corp. [4] vowed to deliver a tenfold boost to computer performance in the next three years by putting the equivalent of multiple electronic brains on each chip."

Conclusions (continued)

- With these announcements to the business world, it is clear that parallel computing will be central to the future of computing.
- This future poses a challenge [7] to software makers and computer science education.

Conclusions (continued)

- Motivation: To share preliminary ideas and experience of introducing parallel computing through out the undergraduate computer science curriculum.

- Provoking discussion and interchange concerning changes within the computing discipline which will likely define the future of computing.

References

[1] *Computing Curricula 2001 Computer Science, Final Report*, The Joint Task Force on Computing Curricula, IEEE Computer Society and Association for Computing Machinery, IEEE Computer Society, December 2001.

[2] Bob Dylan, *The Times They Are A-Changin*, CBS Records, February 10, 1964.

[3] Moore, Gordon, “Electronics”, April 19, 1965.

[4] The Wall Street Journal, “Intel Aims for a Tenfold Boost In Chip Performance by 2008”, Page B8, Column 1, December 8, 2004.

[5] The Wall Street Journal, “Intel Gives Up on Speed Milestone”, Page B8, Column 1, October 15, 2004.

[6] The Wall Street Journal, “Intel Plans to Exploit Chip Trends”, Page B3, September 8, 2004.

[7] The Wall Street Journal, “New Chips Pose a Challenge to Software Makers”, Page B3, April 14, 2005.