

Curriculum Development in Advanced Computation*

Philip J. Morris, Lyle N. Long, Ali Haghghat, Martin L. Brady
The Pennsylvania State University

Introduction

For U.S. industry to remain competitive in the global market place it must update its approaches to product development. The concept of “concurrent engineering,” in which design and manufacturing procedures are integrated has been shown to lead to decreases in development time, engineering changes, and time to market and to increases in overall quality, productivity, dollar sales and return on assets [1]. Central to the reduced design cost are the use of numerical simulations. This is especially true in the aerospace field. Hoist et. al. [2] argue that the development costs for advanced vehicles are escalating due the constant desire in the commercial field for advancement in aerodynamic, fuel, and propulsion efficiency, weight reduction, flight stability and control, and environmental factors. One way to contain these costs is to place a greater emphasis on computer simulations. However, the simulation of the performance of complex aerospace vehicles is beyond current computing capabilities. This has resulted in the HPCC program: a multi-agency effort to advance the scientific computational capabilities of the U.S. by a factor of 1000 by the end of present decade. The achievement of “scaled teraFLOP performance” must rely on the use of massively parallel computer architecture. The U.S. is no longer the leader in vector computing. The Japanese have drawn even. The conversion of supercomputing from sequential/vector to parallel would be valuable as that is where the U.S. leads; however, the U.S. lead has been jeopardized by our failure to get large parallel machines and parallel software technology into the hands of users. Thus, the training of our university students in the capabilities and use of advanced parallel computers, is an essential component in the establishment of our industry’s competitive edge. In response to this need, faculty in the Colleges of Engineering, Science, and Earth and Mineral Sciences at the Pennsylvania State University have combined to develop a curriculum in advanced computation that emphasizes the capabilities and uses of parallel computers. The goal and method of approach can be summarized as:

Goal: - To train undergraduate and graduate students in advanced computation with an emphasis on the capabilities and uses of parallel computers.

Approach - Through the introduction of a sequence of courses at the senior and introductory graduate level in advanced computation; through the development of software demonstrations for classroom use; and through the use of high-technology classrooms and laboratories.

Course Development

As an initial undertaking we have developed and taught three new lecture courses and have a offered a seminar series in high performance computing. The courses are described here. The course specification

* Supported by NSF Grant EEC-9420592



and numbering reflects the department of the instructor who introduced the course and a cross-listing with High Performance Computing. 400-level courses are senior level courses and the 500-level designation denotes graduate courses.

AERSP/HPCOM 497B, 597G

This course has been taught in a new technology classroom using the latest IBM RS/6000 UNIX workstation and high quality projection equipment. The most recent offering of this course is in new UNIX based classroom established with support from IBM under the Selected University Research (SUR) program. The configuration of these classrooms is described below. A massively parallel CM-5 (at NCSA) and a 62-node IBM SP2 (at Penn State) are integral parts of the course. Numerous programming examples and computer code demonstrations are used. The principal developer of the course is Dr. Lyle Long. The goals of the course are to teach the student to:

1. understand the differences between current parallel computers
 - Intel Paragon, Cray T3D, CM-5, KSR, nCUBE, IBM SP-2, etc.
 - SIMD and MIMD paradigms
 - Message passing and data parallel architectures
 - Parallel Virtual Machine (PVM)
2. effectively map a wide range of algorithms onto parallel computers
3. understand parallel processing performance
4. learn language features for parallel (especially FORTRAN 90)

The students participate in four program development projects. Each project requires the writing of a computer program to solve a particular problem and a report that describes the algorithm and presents the results. The course prerequisites are a first course in computer programming for engineers and scientists and courses in vector calculus, matrices and ordinary and partial differential equations. The course outline contains the following topics:

1. Computer hardware and history
2. Operation of machines
3. Languages
4. Performance
5. Linear algebra
6. Partial differential equations
7. Advanced topics/applications

- Particle methods (DSMC)
- Navier-Stokes



Transposed Thomas algorithm

Boltzmann equation

Molecular dynamics

Most of the material for this course is available on the web at <http://cac.psu.edu/lnl/497/>.

This course has been taught for four years. The enrollment has been typically 25-35 students. They are evenly distributed between undergraduate and graduate students and they come from numerous departments, including aerospace, industrial engineering, chemistry, mechanical engineering, computer science, nuclear engineering, and mathematics.

CSE/HPCOM 597: Parallel Architectures and Commutation

This course was taught for the first time in the Fall semester 1995. The course was developed by Dr. Marty Brady. The goals of the course are to

1. impart a practical knowledge of commercial parallel computer architectures, and to concentrate on the impact of various hardware structures on the efficient design of parallel algorithms.
2. introduce general techniques for implementing parallel algorithms using a variety of practical applications.
3. illustrate, as various structures of parallel computer architectures are introduced, their implication to algorithms and software using basic problems such as matrix multiplication that are often important subroutines in practical engineering/science applications.

The course outline contains the following topics:

1. Fundamentals of parallel computation
2. SIMD architectures and data parallel computing (MasPar MP-1, TMC CM-2)
3. Scalable multiple computer architectures (Intel Paragon, TMC CM-5, nCUBE)
4. Scalable multiprocessor architectures (KSR-1)
5. Other high performance architectures (Cray C-90, networked clusters of workstations)

The students complete three or four small projects on the available parallel machines. The projects coincide with classroom discussion. For instance, when discussing message-passing and networks, the project involves evaluation of a basic algorithm that relies on message-passing, measuring performance and efficiency as a function of problem size. One project is staged as a "contest" based on speedup or efficiency measurements, much like the Gordon Bell competition.

For the first offering of the course 9 students were enrolled (4 Computer Science and Engineering, 2 Electrical Engineering, 1 Aerospace, 1 Meteorology). Students had access to the following facilities at Penn State: 2K node Thinking Machine CM-200; 16 node nCUBE 2; IBM SP-2. In addition, the students were able to submit large jobs to the IBM SP-2 at the Cornell Theory Center.



NUCE/HPCOM 597G: Scientific Algorithms for Parallel Computers

This course discusses mathematical modelling methods used in different scientific and engineering applications. This course has been developed by Dr. Ali Haghghat and is being taught for the first time in Spring 1996. Differential equations and Monte Carlo methods are discussed. A finite-difference discretization for a diffusion equation is derived. This formulation is parallelized via domain decomposition techniques. A simple Monte Carlo formulation is developed and parallelized. Both the finite-difference and Monte Carlo formulations are implemented on two types of parallel architectures: shared memory and distributed memory. Differences and similarities of the two formulations and their suitability and performance are measured.

The course outline contains the following topics:

1. Mathematical modelling of physical systems

- Partial differential equations

- Monte Carlo methods

2. Development of parallel algorithms

- Finite-difference diffusion equation

- simple Monte Carlo particle transport formulation

3. Implementation on parallel computers

- Shared memory (CRAY C90)

- Distributed memory (IBM SP-2)

4. Measurement of parallel performance and issues related to optimization

The course is graded on the basis of homework and projects (80%) and a midterm examination (20%).

Evaluation

Course evaluation materials have been and are being developed for these courses. These include the "one-minute paper" for immediate feedback from students, an "early feedback form" to provide the instructor with information on the effectiveness of their instructional effectiveness early in the semester when small adjustments can be easily made, and a "student survey form" used to determine the students evaluation of the new curriculum. The primary evaluation activities will begin by evaluating the classes as discrete units, focussing on teaching methods and student teacher adjustment to the new curriculum. In the second year and beyond the evaluations will assess student outcomes, including learning and performance. The results of this evaluation process will be reported in more detail at an appropriate later date.

Computer Classroom Facilities

Penn State and IBM recently developed two new classroom/laboratory facilities to promote the teaching of high performance computing. One of these labs has 31 IBM RS/6000 Model 40P computers and a podium/LCD panel. The other room has two IBM RS/6000 SMP machines and 30 IBM X-terminals. We intentionally chose these two different computing configurations to evaluate the effectiveness of each approach. During the first year of service approximately 37 courses have used the rooms and roughly 1000 students have accessed these facilities. The software installed includes: ProEngineer, IDEAS, Patran, MSC/NASTRAN, Tecplot, Matlab, Maple, Mathematical, gnuplot, Fortran 90, and C.



Related Activities

The interactions between faculty involved in the CRCD has acted as the stimulus for additional interactions beyond the curriculum area. It rapidly became clear that faculty and students in a broad range of disciplines at Penn State were involved in high performance computing. The topics covered by these researchers not only included the expected areas in science and engineering but also in business and theatre arts. This strong and broad activity in high performance computing applications provided the basis for the establishment, in Fall 1995, of the Institute for High Performance Computing Applications (IHPCA). The primary goal of the IHPCA is

To develop and implement new techniques for the numerical simulation of scientific and engineering problems that require the use of large scale computations with a particular emphasis on both parallel computers and distributed computation.

To achieve this goal the IHPCA is pursuing the following actions:

1. To conduct state-of-the-art research - advance both software and architecture technologies as well as develop innovative numerical solutions to practical problems.
2. To develop coursework in high performance computation - courses will continue to be developed at the senior undergraduate and graduate levels in the uses and applications of parallel and distributed computers. There will also be an emphasis on the use of technology in the classrooms.
3. To establish of a minor in high performance computation - working with other Colleges and the Graduate School a minor in high performance computation is being established. This will be available to students in all Colleges.
4. To acquire high performance computation facilities - the Center will promote and coordinate responses to requests for proposals for scientific computing equipment.
5. To enhance the visibility of high performance computation activities at Penn State - the Center will provide information on high performance computation at Penn State in the form of information brochures and annual reports on activities, and establish an Internet information facility for electronic access.
6. To promote communication between multidisciplinary research groups and to develop a synergistic relationship between groups conducting basic research and those solving engineering problems.
7. To promote and coordinate outreach activities - the Center will participate actively in the education and recruitment of minorities through short course activity for faculty and students at traditionally black colleges. An awareness of opportunities in high performance computation will also be promoted to other women's and minorities groups.

It is clear that the activities of the IHPCA are congruent with those of the CRCD program. At this time there are fifty-two faculty and staff associated with the institute. Further information about the activities of the IHPCA are available on the world wide web at <http://cac.psu.edu/lnl/ihpca/>. (This is a temporary site only).



References

Newsweek, 3157:110-117, April 1990

T.L.. Holst, M. D. Salas, and R. W. Claus. The NASA Computational Aerosciences Program - Toward teraflops computing. AIAA Paper No. 92-0558.

Biographies

Philip J. Morris is the Boeing Professor of Aerospace Engineering at Penn State. He received his B. SC. degree in Aeronautics and Astronautics in 1967, his M. SC. in Advanced Acoustics in 1969, and his Ph.D. in Aeronautics and Astronautics in 1972, all from Southampton University, England. Prior to joining Penn State in 1977 he held positions with the Institute for Aerospace Studies, University of Toronto, and Lockheed Georgia Company. He teaches courses in fluid dynamics, engineering analysis, acoustics, and numerical methods. He is the Director of the Institute for High Performance Computing Applications.

Lyle N. Long is an Associate Professor of Aerospace Engineering at Penn State. He received his B.M.E. in Mechanical Engineering from the University of Minnesota in 1976, his M.S. in Aero. and Astro. Engineering from Stanford University in 1978, and his D. SC. in Aerospace Engineering from George Washington University in 1983. Prior to joining Penn State in 1989 he held positions with Lockheed Aeronautical Systems Co. and has been a visiting scientist at Thinking Machines Corp. He teaches courses in fluid dynamics, parallel processing, and engineering analysis and design. He is the Technical Director of the Institute for High Performance Computing Applications.

Alireza Haghghat is an Associate Professor of Nuclear Engineering at Penn State. He received his B.S. in Physics from Pahlavi University in 1978, his M.S. in Nuclear Engineering from the University of Washington in 1981, and his Ph.D in Nuclear Engineering from the University of Washington in 1986. He teaches courses in parallel computations, particle transport, and modeling of reactor physics.

Martin L. Brady was an Assistant Professor of Computer Science and Engineering at Penn State. He received his B.S. in 1982, his M.S. in 1984, and his Ph.D. in 1987, all in Electrical Engineering at the University of Illinois at Urbana-Champaign. Until 1990, when he joined Penn State, he worked with the parallel processing group of Lockheed R&D Division. He is presently employed by Intel Corporation in Santa Clara, CA.

