Purdue's Parallel Processing Past & Present

School of Electrical and Computer Engineering
Purdue University
West Lafayette, IN 47907-1285

Parallel processing is the key to dramatic speed increases for supercomputing tasks, but that doesn't mean that it is easy to achieve speedup using parallelism. At the Purdue University School of Electrical and Computer Engineering, we long ago realized that the best speedup will only be obtained through careful design and implementation of the parallel computing system — integrating both hardware and software to maximize system performance. It is perhaps easiest to see the evolution of these ideas through the sampling given in the timeline on the reverse side of this sheet, coupled with the following brief descriptions:

1st Dual Processor VAX 11/780 & Sun4. In 1981, DEC VAX 11/780 machines running UNIX were the mainstream computing support for many schools around the country... but they were not sufficient for the needs of Purdue Electrical Engineering. For this reason, George Goble and Mike Marsh sought a low-cost way to multiply the processing power — which they did by discovering a way to add a second 11/780 processor to a standard uniprocessor system. Further, BSD UNIX was modified to support this dual-processor configuration, thus creating *the first multiprocessor UNIX system* (see `http://ghg.ecn.purdue.edu/vax/paper.html`). Not only were many such systems configured at Purdue, but the hardware and software were widely adopted and adapted. The 4 PE Sun4 system was an early commercial version of SMP UNIX technology. Beyond using such systems, Purdue is also involved with the development of SMP Linux for parallel processing applications (see `http://yara.ecn.purdue.edu/~pplinux/`).

Inmos Transputers. The 16 PE Inmos Transputer system served as a training ground for learning about message-passing MIMD systems. Since then, Purdue has purchased nCUBE, nCUBE2, Intel Paragon XPS, and IBM SP2 message-passing MIMD systems as general-use supercomputers.

NCR GAPP & MasPar MP1. The NCR GAPP is a SIMD system built using an array of 48x48 bit-serial processors equipped with a video input; it was used primarily by Jose Fortes as a target for compiler code scheduling and algorithm mapping research. The GAPP also increased our interest and understanding of massively-parallel SIMD, leading to the purchase of a 16,384 PE MasPar MP1 (the first maximum-configuration system in a US university). The MP1 still serves as a general-use supercomputer, but also was the focus of a wide range of hardware and system software research including Hank Dietz's development of compiler technology that allows a SIMD machine to efficiently execute MIMD programs.

PASM. PASM is the **Pa**rtitionable **S**IMD **M**IMD experimental system designed by H. J. Siegel's research group. There are a number of innovations in this prototype system which was entirely built at Purdue; for example, PASM incorporates a parallel mass storage system and the patented "Extra Stage Cube" network, but perhaps the most significant contribution is the concept of building a machine that can execute in either SIMD or MIMD modes yet uses commodity processors. PASM is capable of being partitioned into up to four submachines, each of which can switch between SIMD and MIMD modes in just a few instructions, thus allowing the mode best matching each portion of a parallel algorithm to be used. This work has since expanded to manage the use of heterogeneous groups of parallel supercomputers.

PAPERS. After building PASM, we realized that the implementation "trick" that allowed the machine to execute both SIMD and MIMD using commodity processors was really a barrier synchronization mechanism that could be invoked at each instruction fetch. This led Hank Dietz and his students to further study the architecture, implementation, and associated compiler technology for barrier synchronization. Major progress was made in developing compiler technology and SIMD/VLIW/MIMD mode emulation, as well as classification of two types of fully partitionable barrier mechanism: SBM and DBM (**S**tatic and **D**ynamic **B**arrier **M**IMD). Limited SBMs were built (e.g., PASM and Thinking Machines CM5), but the complexity of the architecture deterred implementation of a DBM until late in 1993, when we discovered a way to build a DBM using *n* SBM barrier units and a special 1-bit wide multibroadcast network. PAPERS, **P**urdue's **A**dapter for **P**arallel **E**xecution and **R**apid **S**ynchronization, began as a proof-of-concept DBM implementation... but we soon discovered that having a simple stand-alone DBM or SBM unit that could use unmodified PCs or workstations as PEs was a very significant advance toward making parallel processing available to the masses. We also discovered a new communication model in which a single communication operation can compute an *Aggregate Function* of the data *collected from all PEs*, essentially an extension of the concept of the MasPar MP1's "global Or" network. Nearly all aspects of the PAPERS project are available on-line at `http://garage.ecn.purdue.edu/~papers/`.

The above descriptions, and the photos on the reverse side, are only a small sampling of the work and facilities at Purdue's School of Electrical and Computer Engineering. The school has approximately 70 faculty, 1,000 undergraduate students, and 500 graduate students. To make things fit this page, we have omitted mention of the work involving parallelizing compilers, image and speech understanding, and various engineering-oriented applications. For more general information on the school, see: `http://www.ece.purdue.edu/`.