

BYTE


NOVEMBER 1988

A MCGRAW-HILL PUBLICATION

REVIEWS

Compaq 386s
FlexCache 25386
Mac and PC Transputers
SpinRite
FullWrite
Zortech C++
PC Lint


PRODUCT FOCUS



Steve Jobs' new "machine for the '90s"

The NeXT Computer

- 25-MHz 68030 • Optical Drive
- Math and Digital Signal Processors
- 8 Megabytes of RAM
- Windowing Unix



IN DEPTH

Parallel Processing

PLUS

Scotland's Innovative Rekursiv Chip
PC Backup Power Supplies
Parallelizing Prolog
5 Short Takes



\$3.50 U.S.A./\$4.50 IN CANADA
0360-5280

Parallel Processing

- 275 Side by Side**
by Klaus K. Obermeier
- 287 T800 and Counting**
by Richard M. Stein
- 301 Getting the Job Done**
by David Gelernter
- 311 The Third Dimension**
by Michael J. Little and Jan Grinberg
- 320 Boards and Boxes**

Parallel processing could be described as the ultimate in teamwork. In fact, the kind of teamwork involved is not unlike that found in the football stadium on an autumn Sunday afternoon. The quarterback has his job to do, the center has his, the ends and backs have theirs, and the guards and tackles have theirs. All these jobs are under way at the same time, but they're all different and being done by a different player—parallel processing.

Similarly, when a group of people are raking leaves, different people are doing the same job, at the same time, with the result of significantly cutting down on the time required—also parallel processing. Not all jobs, however, can be done in parallel. That Thanksgiving turkey we look forward to at the end of the month can't be rushed—microwaves aside.

The same basic concepts apply in computing. Multiple processors operating in parallel can perform many, but not all, jobs faster than uniprocessors. A logically sequential program must still run sequentially. However, a modular program, or one that can be made modular, can run different sections on different processors and improve its speed.

Last summer, NASA's Jet Propulsion Laboratory introduced the Mark 3 Hypercube parallel supercomputer. Parallel processing has long been the exclusive realm of very large systems; however, it is now becoming available at the microcomputer level. For example, Zenith has announced the Z-1000 with its parallel 80386s (see Microbytes on page 11), and Cogent has come out with the XTM (see the text box "The Crossbar Connection" on page 278).

This month, we look at the world of parallel processing from the microcomputer view. In "Side by Side," Klaus K. Obermeier looks at the field as a whole:

the appropriate algorithms and applications; the programming languages, including old favorites and new ones with special parallel-processing functionality; and the hardware and operating-system architectures involved.

One particularly applicable piece of hardware is the transputer board. In "T800 and Counting," Richard M. Stein looks at the T800 transputer board from INMOS, discussing both the hardware aspects of the transputer and the related software aspects of the occam language—the two were designed to work together.

Another language designed for parallel processing on the transputer is the University's Linda. In "Getting the Job Done," David Gelernter, one of the language's designers, gives us the inside scoop on the current state of Linda, what it does, how it does it, and its especially parallel features.

Finally, we have an article on a different way of making computers. In "The Third Dimension," Michael J. Little and Jan Grinberg describe the inner workings of Hughes Research Lab's 3-D computer. It's an innately parallel computer built not of chips but of wafers—many wafers. It's a fascinating technology.

While the concept and practice of parallel processing have a history in the large-computer arena, the idea of getting parallel-processing power on a desktop is still very new. The Mark 3 Hypercube is intended for simulations for the Strategic Defense Initiative. Can that kind of power really exist on a desktop?

—Jane Merril
Senior Technical Editor

