AIX

IBM POWERparallel Products

AIXpert talked to Rod Cowan, manager of commercial application enablement for the POWERparallel organization. We wanted to learn more about the new SP1 RISC-based parallel processor. He provided a quick summary of the machine's capabilities and commercial applications.

IBM calls your group POWERparallel Systems. What does that encompass?

Cowan: POWERparallel is a business unit that draws upon the unique skills of IBM's Large Scale Computing Division (LSCD) that develops mainframes, the RISC System/6000 Division, and IBM Research. It brings together the following specialties:

- **LSCD**: Years of experience building enterprise-wide solutions, plus mission-critical computing expertise
- **RISC System/6000 Division**: RISC processor technology, understanding how to build machines with compelling price/performance, AIX and open systems knowledge, and ability to get to market quickly
- **Research**: Years of experience in parallel technology and technology for fast processor-to-processor communication

The RISC System/6000 Division has a similar relationship with IBM's personal systems group resulting in the Power Personal Systems Division that is exploiting the PowerPC technology.

We see ourselves at the high end of IBM's "palmtops to teraFLOPs" strategy. That phrase capsulizes our plan to leverage the POWER and PowerPC technology from the smallest machines up to supercomputer-class machines. POWERparallel fits into the high end with our parallel and massively parallel RISC systems.

Rod Cowan

We are using exactly the same hardware and AIX system software that are in the RISC System/6000 to build our parallel machines. RISC System/6000 applications run unchanged on a single node of the Scalable Parallel (SP) system.

Could you describe the SP1 processor?

Cowan: Many people think of the SP as a supercomputer, which it certainly is. At the same time, it also has a place as a commercial machine. We have seen significant interest from commercial customers, and are working quickly to deliver products for decision support and other commercial applications.

The SP was designed in IBM's High-Performance Supercomputing Systems Lab in Kingston, New York. It shipped last September, only 18 months after the lab was formed—a real tribute to the benefits of using "off-the-shelf" RISC System/6000 designs as the building blocks for our parallel machines.

Figure 1 shows an overview of the product. We combine 8 to 64 processors that resemble the RISC System/6000 Model 370 in the SP. The RISC processors are connected by a very high-speed, low-latency, non-blocking switch. If you have a 64-way machine, 32 processors can talk to 32 other processors—all at the same time. We also added hardware and software for availability and system management tailored to the new parallel environment—altogether, delivering one to eight gigaFLOPs of performance.

What application opportunities does the SP provide?

Cowan: The SP was originally intended for the scientific market (universities, petroleum, and so on). Although most of our initial shipments are for those types of applications, the SP is also

Figure 1. The SP POWERparallel system

moving into the commercial world. Small Computer Systems Interface (SCSI) devices can now be attached to each of the nodes, which allows devices such as RAID disks to be accessed by the software. We showed Oracle and Sybase databases running on the SP at UNIX Expo. Databases are critical enablers for commercial applications. You can expect to see several exciting announcements on the availability of major UNIX database products for the platform during the coming year. These will encompass both supercluster and parallel optimized implementations providing solutions to large database problems.

The Oracle database manager actually splits the database across multiple SP nodes. All the disks are available to all the nodes. Multiple queries can be executed on multiple nodes simultaneously. Oracle will have intra-query parallelism later in 1994. This means that a query can be split apart, its functions run on different nodes, and then be brought together to deliver the answer.

We are also working closely with the DB2/6000 team. Their parallel offering will start with early customer studies in 1994.

The availability of parallel databases means that if an application runs on top of one, it can take advantage of the parallelism of the SP without being rewritten.

We are also doing some interesting work with the system as a supercluster, which allows users to take advantage of the client/server environment while moving to parallel applications.

What types of applications will be available for the SP?

Cowan: There are three major areas. The first, called supercluster, consists of large client/server

applications that are pushing the limits of the RISC System/6000. Or perhaps the customer has several servers and would like to consolidate them into one machine with one systems management interface. Client/server applications are easy to move to the SP because they are designed with components that run on different machines. SAP's R/3 business management application is implemented using a client/server architecture and uses the SP as a supercluster. They can take advantage of a database server optimized for parallel processing; other nodes will be used as application servers. The customer can balance the number of application servers to the workload. Terminals are connected to the application servers.

Another interesting area is integrating complex data analysis with mainframes, such as data mining. Customers with this requirement often want to get more information from the vast quantities of data that they have collected. They want to continue to run their transaction processing applications on the mainframe, attach an SP via an ESCON™ channel, and use a parallel database application on the SP to perform data mining that would not otherwise be affordable. A typical application is analyzing trends from data collected by a retailer.

RISC technology t makes up the has been ubling in price/ formance every to 18 months.

The third area is emerging applications. Over the past 10 to 20 years, mainframe price/performance has improved about 15% to 20% annually. But the RISC technology that makes up the SP has been doubling in price/performance every 12 to 18 months. In addition, the SP parallel technology allows that performance to be multiplied by 8 to 64. This means a tremendous amount of computing, memory, and I/O power available at far lower costs than large systems mainframe customers are used to seeing. Applications that stand out are securities portfolio management, analysis of airline routing, retail data analysis, and actuarial or demographic analysis for insurance companies.

How easily can an application running on the RISC System/6000 today be ported to the SP?

Cowan: It depends on how the application is designed. If it is a client/server application divided between database processing and application processing, the SP is ideal. This is because the SP supports communications between processors using TCP/IP over the switch, enabling programs designed to communicate over LANS to get much better communication performance. That becomes transparent to the application. Equally easy to move are applications that run on a parallel enabler such as a database. The most challenging applications to create are those that take advantage of the parallelism of the SP because they need the performance that parallelism can provide.

What is the advantage of the SP1 over multiple RISC System/6000s for client/server applications?

Cowan: There are several advantages. First, the performance of TCP/IP over the switch is faster. Second, centralized systems management allows administrators to treat the SP complex as a single system. For example, a user is added to the system only once and there is a common control point for software upgrades. Finally, parallel applications can be running on the system at the same time.

What can we expect in the next year for the SP?

Cowan: IBM recently announced that it will provide a 512-way version of the SP machine to Cornell University. So you can expect that, in 1994, we will be going beyond the 64-node limit of the first machines.

We will improve the I/O capabilities and switch performance over the next year and also move to the POWER2 technology for the processors. You will also see software enhancements such as system management and additional applications.