Affordable Supercomputers

Advances in parallel processing make supercomputing technology affordable and accessible to a growing number of users.

By Todd Shumway

Full-scale supercomputers have undoubtedly eased the workload of many physicists, researchers and engineers. The ability of supercomputers to process numbers at very high speeds using a vector architecture, the execution of individual instructions simultaneously rather than one at a time, enables computer problems to be processed at tens of millions of instructions per second (MIPS). In some cases, this speed literally reduces by years the time needed to solve complex computational problems.

These powerful supercomputers, however, are not cheap to own or operate. A Cray 1 supercomputer sells for upward of $12 million and access time can cost up to $7,000 an hour. While these supercomputers are popular in academic circles and large federal research centers, an emerging, more affordable alternative is capturing a wide following of users. Companies offering affordable multiple processor machines to perform the high-speed computing work previously handled only by the likes of those built by Cray Research, Minneapolis, MN, have expanded the supercomputing market with technology known as parallel processing.

Multiple Processors

The term parallel processing is most often applied to multiprocessing, or the employment of multiple processors to execute several instruction streams concurrently. By combining the power of multiple, comparatively-slow processors and dividing a problem between them, the time required to complete an application is much less than it would be for an individual processor.

In most cases, makers of parallel processors have chosen to exploit their own market niche rather than compete directly against true supercomputers. Prices for these minisupercomputers or parallel computers range from $100,000 to $7 million, significantly lower than supercomputers. The parallel processor producers also employ a strategy of marketing more flexible machines that can service a wide variety of engineering and scientific needs.

According to a representative of Evans and Sutherland's Computer Division, Mountain View, CA, makers of the ES-1, "Because of the cost of supercomputers and the difficulties associated with fully utilizing their power, these systems have become essentially 'special purpose' machines." In contrast he said, "The ES-1 was designed from the ground up as a general purpose supercomputer."

Scalar Processors

The ES-1 scalar parallel processor accommodates up to eight processors and operates at 1,600 MIPS and 1,600 million floating point operations per second (Mflops) compared to vectorized supercomputers operating in the billion floating point operation per second (Gflops) range. Because all computations are done separately and sequentially, scalar architecture enables each processor to act either as a separate unit or to tackle a large problem with the other seven processors.

According to Don Oestreicher, applications director of Evans and Sutherland's Computer Division, this architecture gives the ES-1 an advantage over vectorized computers when attempting to solve non-well-structured computational problems. Well-structured computational problems, such as physics problems, are solved most quickly on a vectorized machine, because the problem can be broken up and the computations solved simultaneously. While a scalar parallel processor can solve such problems much more quickly than can a traditional serial computer, such well-structured problems are still most

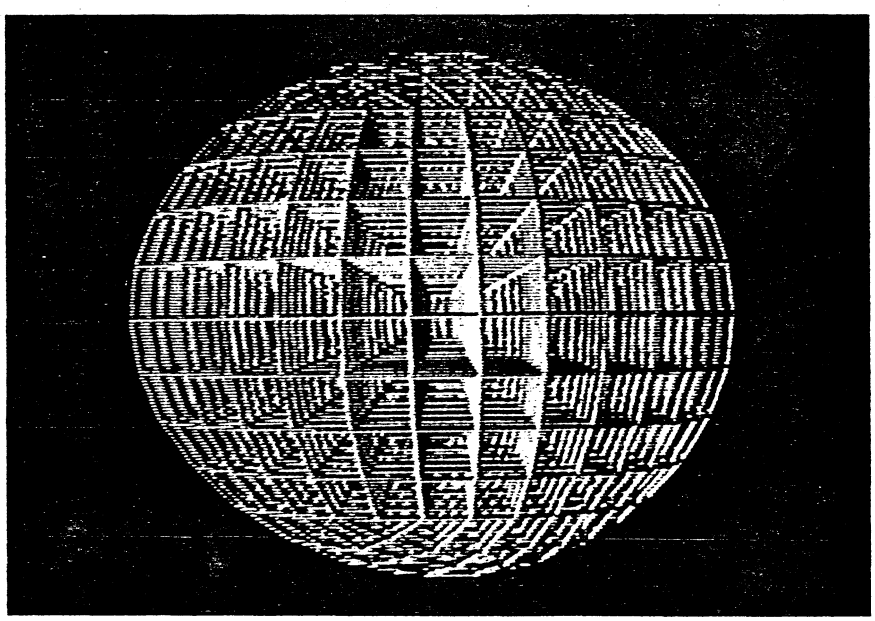

A three-dimensional simulation of a sphere showing both its interior and exterior. Such simulation is used to examine fluid dynamics in aerospace applications. (Photo courtesy of Dr. Daryll Pepper, Advanced Project Research Inc.)

quickly solved on vectorized supercomputers. However, not all engineering or research problems lend themselves to this vectorized approach. Problems that are small or require sequential computation monopolize the vectorized computer's resources even if the problem requires minimal computational power. Oestreicher said this is an inefficient use of expensive resources. Because scalar parallel processors do computations in parallel, a problem can be worked out on one processor while the others can be used to run different problems. Now, as more minisupercomputers with parallel architecture become available, they are being used for doing smaller runs, debug runs and compilations which cannot be vectorized,'' Oestreicher said.

Networking Processors

As computer makers have continued to combine more and more processors, they have also had to find new ways to facilitate interprocessor communication. Initially, communication within a parallel processor was accomplished via a common bus, a single data path connecting a set of processors to one or more memory boards. Many problems resulted from using bus-based communication for parallel processors, however. Because each additional processor competed for the limited space on the bus bandwidth, communication efficiency would decrease as additional processors were added. This competition could tie up data resulting in a bottleneck if a number of processors attempted to simultaneously access the same data. One solution was to create individual cache memories for each processor, but this created problems in maintaining consistent data between memories.

Another solution was the creation of the crossbar switching architecture. In this design, separate buses are established for each global memory module. These buses are connected to all the processors in a grid-like pattern. All the processors can freely access memory without the delays associated with bottlenecks and competition for space on the bus. In crossbar switching, the only time a delay occurs is when more than one processor attempts to access the same memory module. In such a case, an arbiter lets one message pass and temporarily delays the others. However, crossbar switching becomes unmanageable when large numbers of processors are added to the system. For example a crossbar-based computer with 100 processors and 100 memory boards would need a switch with 10,000 wires. This severely limits the number of parallel processors that can use crossbar switching.

Parallel processors are an emerging, affordable alternative that are capturing a new and wide following of users.

A third type of architecture designed to improve data flow between linked processors is the multistage or ''Butterfly'' switch designed by BBN Advanced Computers Inc., Cambridge, MA. The Butterfly switch does not directly connect all the processors and memories. Data is processed through a collection of very large-scale integrated circuit (VLSI) switch nodes that provide paths between processors and memories. According to BBN, the Butterfly switch incorporates techniques used in packet switching. Packet switching is a method of transferring messages between processors by condensing the messages into small, uniform segments. These segments are then passed through a store-and-forward switching network until they reach their destina-

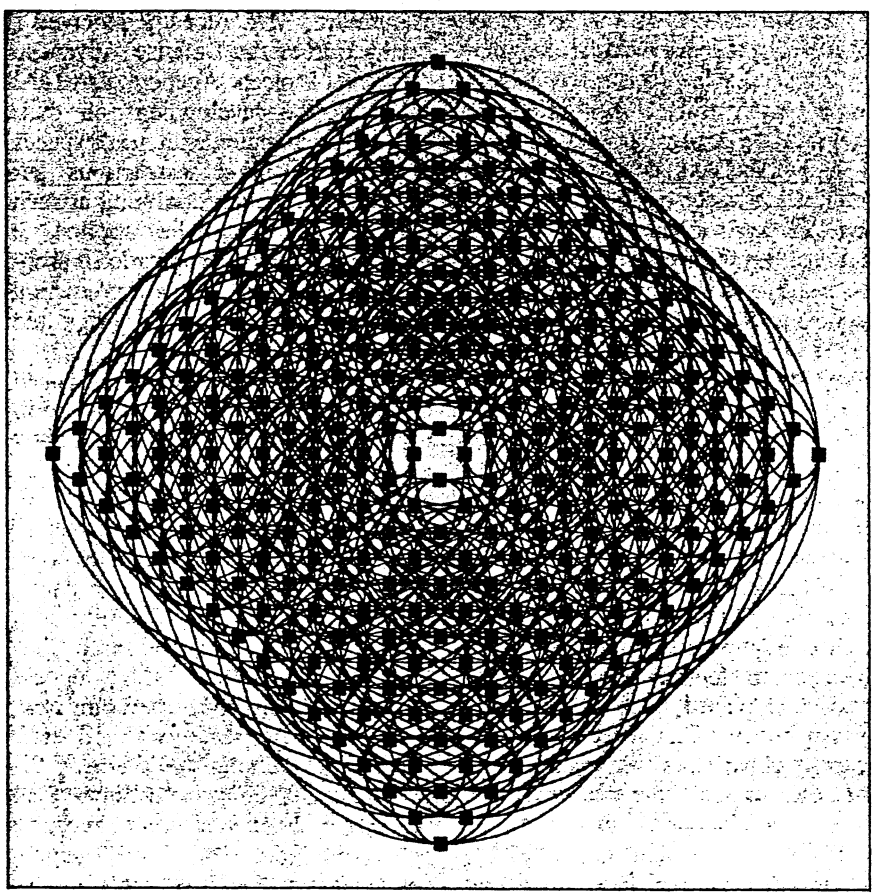

A graphic illustration of a hypercube network providing direct links between an almost infinite number of nodes in a parallel system. (Illustration courtesy of NCube Corp.)

tion. The packets are reassembled into complete messages when they reach their destination. Unlike bus-based designs whose ability to transmit a high volume of data is limited by the bus's bandwidth, the Butterfly switch is able to spread as thickly as required to accommodate heavier flows of information. Also, the number of switching elements within the Butterfly switch is small in comparison to the number of crossover points in crossbar architecture. As a result, it better facilitates future growth than the crossbar or bus-based systems.

Hypercube

Hypercube, which extends communications beyond two dimensions, was created to facilitate data communications between an almost infinite number of processors. Hypercube machines employ thousands of processors and achieve supercomputing performance with a scalar architecture. Essentially, a hypercube is a network of interconnected computers with intertwined data paths connecting every node in the hypercube system to every other node in the system. Communication is permitted between any two non-adjacent nodes via intermediate nodes in a store and forward mode. The intertwined data paths allow the hypercube concept to be extended to hundreds or thousands of individual processors.

Because the hypercube nodes are fully symmetrical, the hypercube network eliminates the competition between simultaneously issued commands attempting to travel on the same data paths. According to Pat Moore, a representative of NCube Corp. Beaverton, OR, maker of the 8192 processor NCube 2, all the processors look the same from any point within the hypercube. "In a hypercube you can lay out the routing to avoid a deadly embrace [a situation in which two or more instructions fight for priority on a data path], because no one node can become a bottleneck," Moore said.

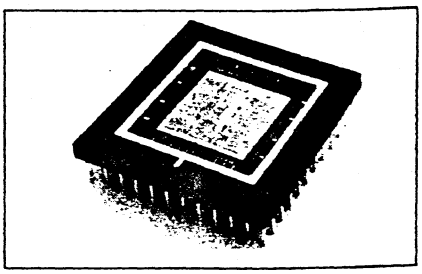

NCube's VLSI processor incorporates a 64-bit computer system into a single computer chip. (Photo courtesy of NCube Corp.)

The hypercube approach extends communications beyond two dimensions.

The advantage of the hypercube approach is the vast number of processors that can be coupled to achieve massively parallel computational power. One recognized disadvantage of hypercube architecture is the communication delay caused by the passage of large amounts of input and output (I/O) through intervening processors. The inefficiency caused by the laborious passage of data forces the system to act as a number of loosely coupled microcomputers.

In an attempt to solve this hypercube inefficiency problem, NCube employs VLSI-based parallel processors to reduce the size of their system and maximize its flexibility. By reducing the size of the processors, the travel time between nodes is also reduced. NCube has combined the complete 64-bit computer system with all its components except the memory chips and placed them on one integrated circuit. This reduces the total number of parts per system along with the cost and size of the overall system. By condensing the system it has become more reliable and efficient, Moore said. "We have increased reliability of the chip through the use of custom silicon and have made the packaging much denser. Any processor bigger than five or six feet will start losing performance." he stated.

Parallel Supercomputers

Not all parallel processing vendors have chosen to appeal to a broad range of users. NCube and others have chosen to use scalar processing to compete with vectorized supercomputers, such as Cray, offering similar performance at a reduced cost. Because each processor is made from relatively inexpensive components, the parallel processors can undersell the most expensive vectorized supercomputers. The NCube 2 offers more processing power than any other computer now available, according to the company. It can perform 60 billion instructions per second (BIPS) and 27 billion floating point operations per second (Gflops) and will be sold for approximately $500,000 per system.

Meanwhile, another massively parallel computer under development will reportedly be more powerful than the NCube 2. Thinking Machines Corp., Cambridge, MA, recently received a contract from the Defense Advanced Research Projects Agency (DARPA) to accelerate development on their new supercomputer with peak speeds above one trillion instructions per second. A computer operating at such high speeds would be applied to extremely complex computational problems such as world climate prediction and semiconductor circuit design and testing.

Another parallel processor produced by General Microelectronics Inc., San Diego, CA, claims similar but slightly slower performance than the Cray on a machine that costs $250,000 to $300,000. Allan Simon, chairman of General Microelectronics said his company's configurable architecture parallel processing system (CAPPS) was tested by NASA Ames Research Center, Mountain View, CA, against Cray, Intel and IBM supercomputers by running the same code on each machine and timing the process. Simon said CAPPS proved faster than the IBM and the Intel while finishing behind the Cray. The CAPPS supercomputer was developed with funding from DARPA and Northrop Corp., Los Angeles, CA. According to Simon, the success of CAPPS in the speed

trials was due to a careful tailoring of hardware and software to fit the specific needs of particular applications. CAPPS will be used in several military applications including flight control computers, computational fluid dynamics, structural analysis, computational electromechanics and artificial intelligence.

Software

A lack of application software has deterred the general acceptance of hypercube machines. Engineering efforts have concentrated on creating hardware to allow an increasing number of processors to collaborate. As a result, the evolution of parallel processing hardware has greatly outstripped the evolution of software to run the hypercubed processors. Don Oestreicher of Evans and Sutherland explained, "In order to get the speeds we need in scientific computing, it is clear, from an academic point of view, that we need to increase the amount of parallelism every year. The massively parallel people said, 'We're not going to do it incrementally, we are going to do it right away.' So they jumped way ahead and they have produced technology that we can't produce software for. For today's users it is not a practical solution."

NCube is offering a partial solution to the user overwhelmed by the programming requirements of a massively parallel computer. With the release of the NCube 2, Oracle Corp., Belmont, CA, will offer their industry standard database management system. Also the AT&T standard System v, Unix-based operating system will enable the NCube 2 to run Unix applications and provide compatibility with scientific, engineering and general business programs. However, much of the application software for the massively parallel computers is still being developed on an application-by-application basis. Cross compilers are available to translate software code from one computer to another as are tools for program development, but software development in a parallel environment is lengthy and time consuming. As described by Pat Moore of NCube, the user must initially construct a program designed for either four or eight processors, debug it, run it through the multiple processors, analyze it with the performance monitor and fine tune the configuration. The programmer then reconfigures the program for a parallel processor so it can run on multiple linked computers acting as a single unit.

Despite the limitations of the hypercube approach, many applications benefit from the tremendous speed at which the massively parallel machines can operate. Neural networks (computers simulating the functioning of the human brain), image processing and robotics employ the very high-speed parallel machines. Using more than 8,000 processors, a massively parallel computer will be able to employ increasingly sophisticated sensory assimilation that can apply pattern recognition to the process of learning. Compact parallel processors can also be adapted to manipulate a robot. Because these systems have fairly small configurations, several 64-bit processors could be placed in a mobile robot.

Software-Driven Minisupercomputers

While the hypercube approach to supercomputing is the logical extension of a desire for more speed and greater power, a whole class of parallel processors have been designed to be as flexible and adaptable to existing software as possible.

The Alliant FX/2800, from Alliant Computer Systems Corp., Littleton, MA, is advertised as the only supercomputer combining an open industry-standard reduced instruction set computer (RISC) architecture with an open, technical computing software standard. Combining an industry standard RISC architecture with industry standard software means a wide spectrum of users can have supercomputing performance without having to develop new software for most applications.

According to Electronic Trend Publications' *High Performance Computer Market* report, the overall design of parallel processing computer architectures, as well as the design of the system's microprocessor, significantly affect overall system performance. Therefore, the development of RISC architecture has been integral to the development of parallel processors. By increasing the speed at which the most common instructions are processed (researchers found that 20 percent of the microprocessor instructions were used 80 percent of the time) RISC microprocessors registered performance gains of 200 to 300 percent over chips without the ability to streamline instruction processing.

Alliant is attempting to carve a niche for its new FX/2800 by introducing a new solution to the problem of finding software for parallel proc-

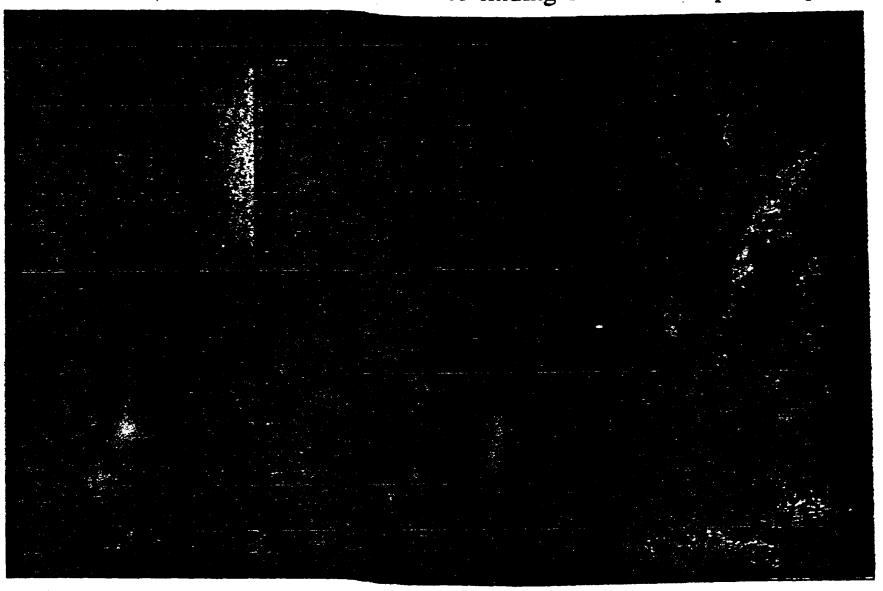

The BBN Advanced Computer's TC2000 is the latest parallel processor by BBN to incorporate the Butterfly switch which provides orderly flow of information between processors and adapts easily to growth. (Photo courtesy of BBN Advanced Computers Inc.)

essing applications. Its solution is an agreement with Intel Corp., Santa Clara, CA, that creates a hardware and software standard for technical computing. The parallel architecture extended (PAX) standard is a set of rules for software and hardware vendors to follow when developing products for i860-based systems, an industry-standard RISC architecture. By making the FX/2800 comply with the PAX standard, any Intel i860 processor-based computer, including the Alliant product, can run the same software. Paul Rubin, a spokesman for Alliant, said the guidelines of the PAX standard assure that any computer complying with the standard will use the same compiler, the same operating system and the same file format for the objects file.

"Intel is attempting to accomplish in technical computing what they did in the desktop environment. With the X-86 family, Intel, along with IBM and Microsoft, were able to create a standard environment in the PC that was completely upward compatible and software compatible," Rubin said. The PAX standard fills a void in the technical computing industry. "This has been a real problem in the technical computing environment, and in particular in the parallel computing environment because there has been no standard," he added. As a result of the agreement, a wide selection of pre-packaged software will be available for scientific and engineering applications. Programs will work the same on an IBM, a Hewlett-Packard or another i860-compatible computer.

This agreement is representative of a growing trend in the high-performance parallel processing market, as hardware and software vendors team up to make their products more flexible and widely usable. NCube's agreements with Oracle is an example of another such agreement. Further agreements between relational database management software vendors, like Oracle, and other parallel processing vendors are expected according to the *High Performance Computer Market* report.

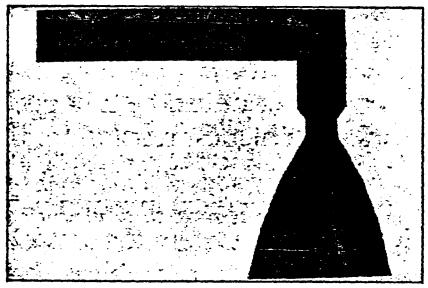

A computer generated image showing the flow of gasses and air from a rocket. Parallel processors are being used to help design the scramjet engines for the proposed National Aerospace Plane. (Photo courtesy of Dr. Daryll Pepper, Advanced Project Research Inc.)

The Alliant provides scientists with a three-dimensional simulation of NASP's rocket-like engines which will propel it into low Earth orbit.

Applications

The bold performance claims made by the designers of parallel processors were supported by Dr. Daryll Pepper, a founder of Advanced Projects Research Inc., Moorpark, CA, who is using the Alliant F/X 2800 in the design of the National Aerospace Plane's (NASP) scramjet engines. The Alliant provides scientists with a three-dimensional simulation of the plane's rocket-like engines which are intended to propel NASP into low Earth orbit. Pepper said that by exploiting the parallelisms of the F/X 2800, he can execute complex physics problems in a 2,000 node parallel processor in significantly less time than a Cray Y-MP would require. "We can throw more physics at the problem and come up with more complete answers when we build a mesh of nodes and operate them in parallel."

The development of parallel processing has provided scientists with opportunities to simulate and test engineering problems that previously would have been prohibitively time-consuming, Pepper explained. Such a project is the design of projectiles to be used in the space-based Brilliant Pebbles missile interceptor system. Pepper said the processors are being used to simulate the detonation of gasses that will propel the projectiles from the lightweight, maneuverable interceptors. Another application for cost-effective, high-speed parallel processors is the development of advanced radar technology. Thermo Electron Corp., Waltham, MA, has developed the Solver, a parallel processor operating peak speeds of 1 billion instructions per second (BIPS) that will deliver detailed pictures of aircraft flying over the horizon and enable pilots to distinguish friendly from hostile aircraft.

By using massively-parallel computer architecture, NASA's Jet Propulsion Laboratory, Pasadena, CA, has designed neural networks that incorporate up to 4,000 processors on a single computer chip measuring eight to nine millimeters square. By incorporating these neural network chips, computational speeds can be increased by several orders of magnitude. Neurocomputers imitate the billions of neurons and trillions of synapses of the human brain by incorporating many processors linked in a parallel architecture. Neurocomputer research money has come from DARPA, the Strategic Defense Initiative Organization, the Office of Naval Research, the Joint Tactical Fusion Program Office and NASA. Applications include: instantaneous map generation for U.S. Army cross-country mobility; graphics for pilots detailing troop positions, weather or air threats; Strategic Defense Initiative (SDI) target location and tracking; interpretation of photographs; and robotic control.

An active future for parallel processing seems assured as technology races ahead by leaps and bounds. One possible dark cloud in an otherwise clear sky is the lack of software readily available for the most advanced applications and the arduous effort required

to develop software for such applications. However, the benefits appear to outweigh the time and effort involved. According to Dr. Pepper, "Although a tremendous amount of time and effort is invested in developing software to run applications in parallel, the payoff is high." But he added in a note of caution, "the grief factor is high as well." □

Beginning with this issue, *Defense Electronics* is providing our readers with an easy way to rate their level of interest in our editorial topics. Please take a moment to indicate your opinion of the preceding article by circling the appropriate number on your reader service card.

High Interest	Moderate Interest	Low Interest
162	163	164

Your feedback is extremely useful and important to us, and it is also of interest to your fellow readers. We welcome your further comments and contributions in our "Reader Feedback" department.

There's an epidemic with 27 million victims. And no visible symptoms.

It's an epidemic of people who can't read.
Join the fight against illiteracy by calling toll-free
1-800-228-8813.

Volunteer Against Illiteracy.
The only degree you need is a degree of caring.

Ad Council Coalition for Literacy

MORE EMI/RFI PROTECTION

HIGHER PERFORMANCE
Independent tests show Lindgren's double-electrically-isolated (DEI) design assures higher shielding effectiveness in low-frequency magnetic and high-frequency electric fields.

MORE SHIELDING SYSTEM OPTIONS
We offer a complete line of modular enclosures, cabinets, and custom rooms.

Shielding materials, in solid-sheet or screen, cover almost any shielding requirements.

Our application engineers will help you design enclosures that meet your requirements for windows, filters for power and signal lines, wave guide feedthrus, and lights.

Other options include RF connectors, fiber-optics, special access panels, automatic doors, emergency exits, permanent ramps, and flush floors with low-rise thresholds.

EASY MODULAR ASSEMBLY
Our patented resilient clamping system permits easy assembly. Enclosures can be modified and relocated again and again without compromising performance.

Lindgren also designs and builds a variety of shielding solutions for industrial, medical, military, and scientific applications.

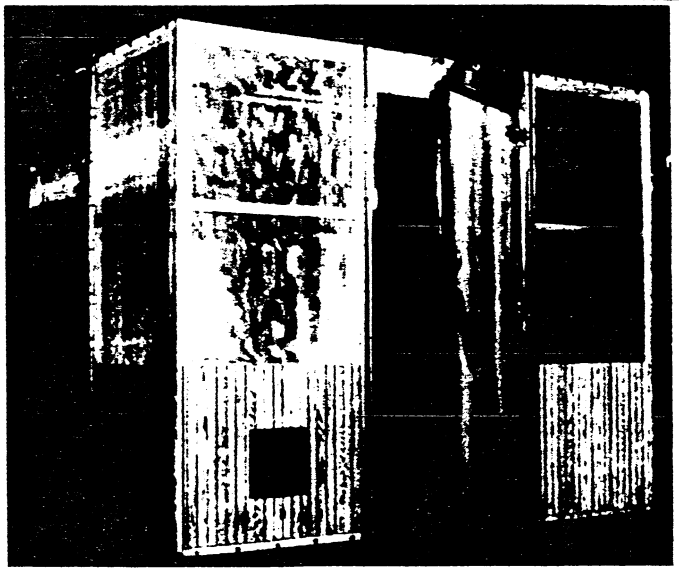

ONLY LINDGREN MODULAR SHIELDING SYSTEMS ARE ASSEMBLED, TESTED, & GUARANTEED TO MEET SPECIFICATIONS BEFORE THEY'RE SHIPPED.

Call us for more information and applications assistance.
708-307-7200

LINDGREN RF ENCLOSURES
400 High Grove Blvd.
Glendale Heights, Illinois 60139
Ph:(708) 307-7200
FAX:(708) 307-7571

Circle Reader Service No. 40