The Ulti-Machine

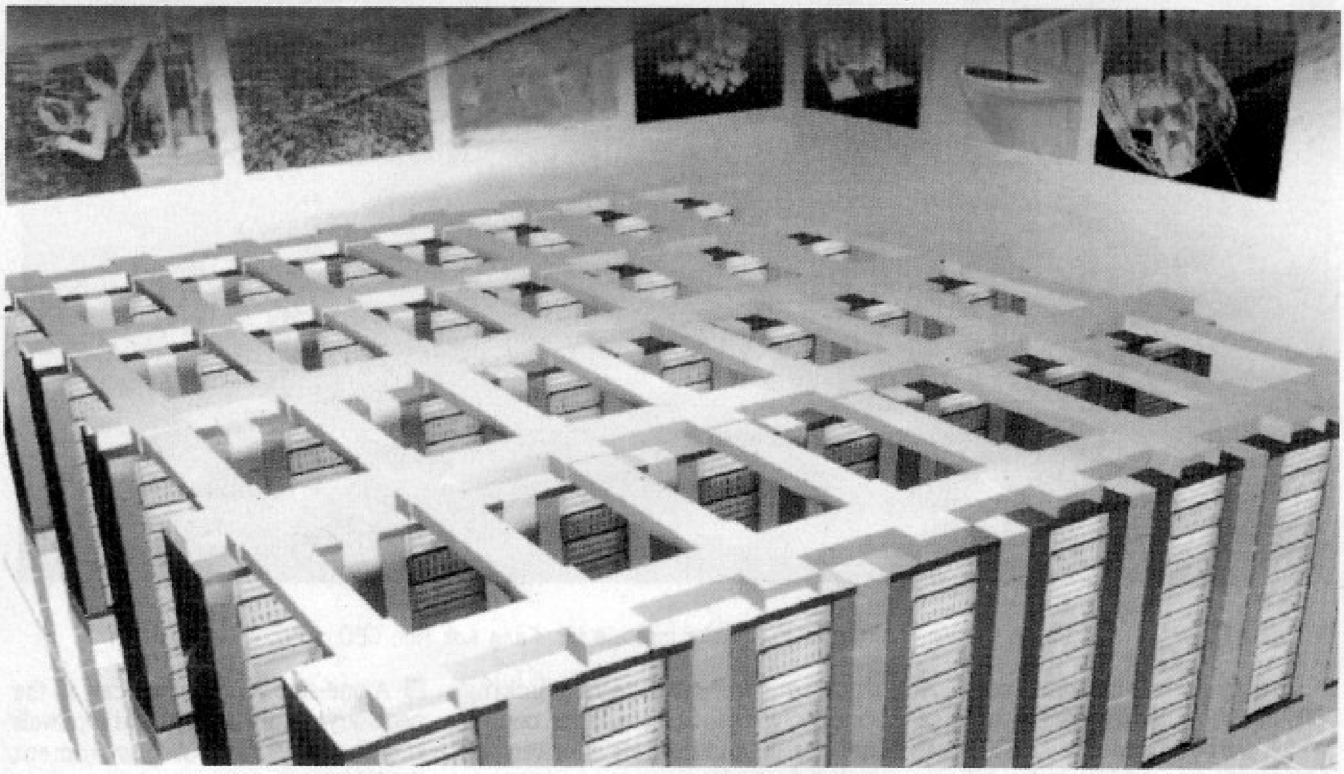

PARAM is a 100 gigaflop machine scalable up to teraflop range. With this, India has become only the second country in Asia and one of the very few countries in the world to have such a machine.

G.P. Vinayababu

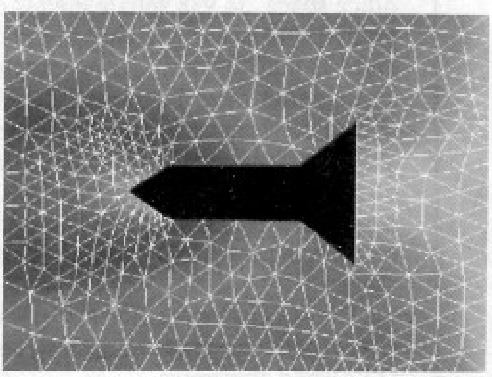

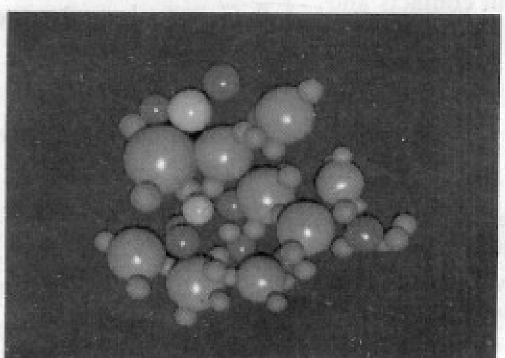

In New York, in the early hours of May 12, 1997, Gary Kasparov, the world's greatest chess player, was fighting a losing battle against a tough opponent. What went against the undisputed world champion and frustrated him to defeat was the brute force of his opponent. Kasparov's challenger here was making moves and countermoves ever so meticulously at speeds unknown to human beings. This challenger was none other than the IBM RS/6000 computer named Deep Blue, specially developed to play chess. The greatest aspect of Deep Blue was its ability to make calculations at the rate of 200 million moves/s, which was no comparison to a maximum of only 2 moves/s which humans can foresee. How could Deep Blue achieve such high speeds of computation? Deep Blue was not an ordinary computer. It was a supercomputer.

Supercomputers

There is no single, clear definition for a supercomputer. Supercomputers are similar to ordinary computers. But the fastest and the most expensive computers available in the market at any given point of time to solve the most difficult problems are called supercomputers. One major consideration which renders any computer a supercomputer is its computing power. "Ten years ago, if a computer crossed 10 megaflops of speed, it was called a supercomputer. But today, only if it calculates at gigaflops, it is described as a supercomputer. After the year 2000, teraflops might be the speed which a computer needs to achieve to be eligible for a supercomputer tag", says Mr. Gangaprasad, Program Co-ordinator, Real Time Systems Group in C-DAC, who is one of the persons involved in the development of software for C-DAC's PARAM series of supercomputers.

Flipflops

Floating point is a datatype in computer programming. It is a number which contains a decimal part. For example, a number like 12.85 has two parts: 12 and 85. And effectively this number will be considered as two separate numbers for calculation by the computer. Multiplication of any two such floating point numbers is called a floating point operation. The speed of a computer is described in terms of the number of floating point operations it can handle in a second. Earlier supercomputers (about a decade ago) were capable of performing megaflops, i.e. 10^6 flops. But today the speed of computers has reached a phenomenal level and hence megaflops are no longer considered supercomputing. Today some of the Pentium PCs itself are capable of achieving close to a megaflop, though it is only a fraction of it. Mainframe computers of the 70s and 80s were considered supercomputers. But today IBM 1401, one of the older mainframes, is slower than the present pentiums. Supercomputers of today are capable of achieving gigaflops.

PARAM 10000 is capable of 100 giga (10^9 flops). However, it is scaleable up to teraflops (10^{12} flops). Future supercomputers would perform at teraflops. A U.S. initiative is developing a 100 teraflop machine which will be released in the year 2004.

PARAM 10000

Compute Node

- ☑ Upto 4 UltraSPARC-II CPUs @ 300 MHz with 2 MB External Cache
- ☑ 512 MB Main Memory (16 Modules of 32 MB-60 ns Access Time) Four Way Interleaved to give Maximum Performance
- ☑ Two 4.2 GB Ultra SCSI Disks
- ☑ 1.44 MB FDD, 12x CD-ROM Drive
- ☑ 10/100 MB Ethernet
- ☑ 2 Serial, 1 Parallel Port
- ☑ 10 PCI Slots
- ☑ PCI Interface to PARAMNet/Myrinet
- ☑ Two 560 W hotswap Power Supplies

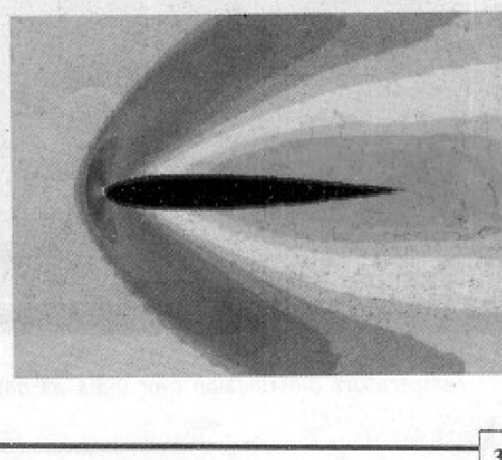

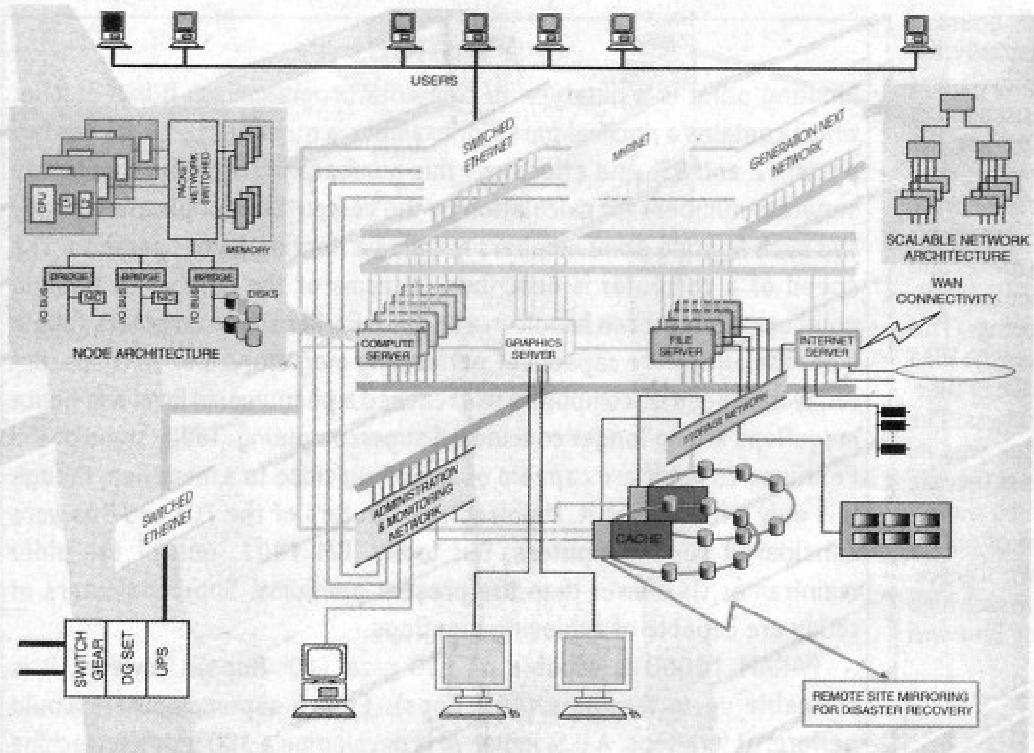

PARAM OPENFRAME HARDWARE ARCHITECTURE

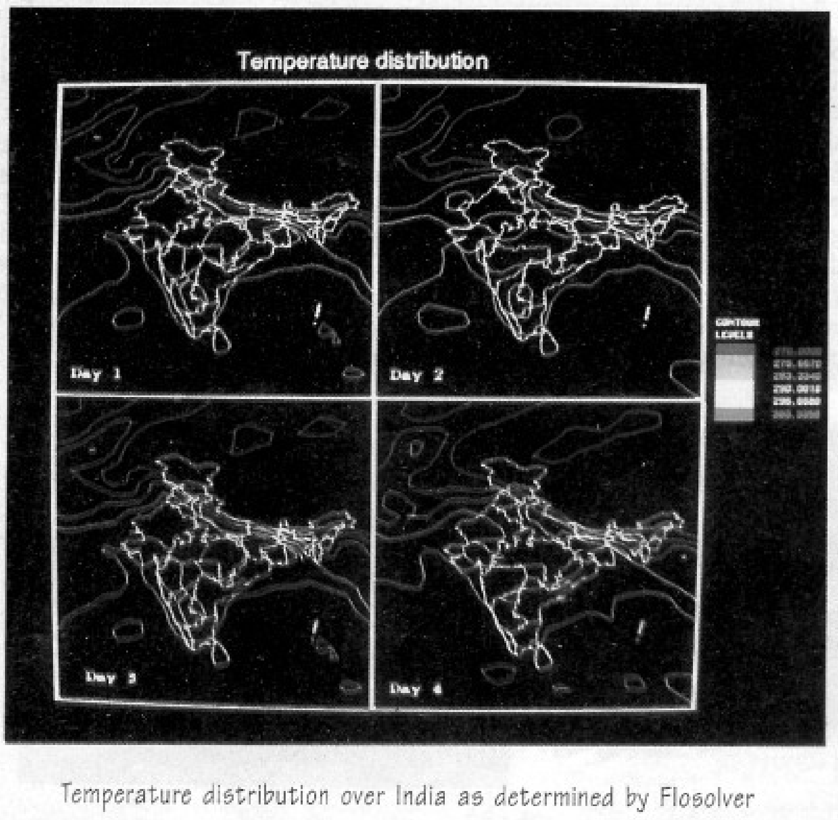

Temperature distribution over India as determined by Flosolver

Chess playing is only one of the areas in which supercomputers are used. In fact, supercomputers have many real-time industrial applications too. Some of the latest technological advances couldn't have been possible without supercomputers, e.g. biochemistry applications. But for the supercomputer, it wouldn't have been possible for biochemists and pharmacologists to develop new drugs for various applications. Drugs are tested under simulated conditions to determine their effectiveness for different combinations. This can optimise the drugs and also save precious lives, eliminating inefficient drug administration. Supercomputers have many other scientific and engineering applications like weather forecasting, remote sensing, seismic data processing, launch vehicle dynamics simulations, medical imaging, etc.

Supercomputers in India

At a time when computer technology in India was still in its infancy and buying advanced computers was considered a luxury, NAL took up the challenge of building the first parallel computing machine in 1986. Prof. R. Narasimha, the then Director of NAL, started the Flosolver project to design, develop, fabricate and use a suitable parallel processing computer for application to fluid dynamical and aerodynamical problems. The CFD researchers at NAL struggled to do worthwhile CFD calculations - often taking weeks to solve a problem which their colleagues in the U.S. could solve in hours! This prompted the development of Flosolver. For all practical purposes, Flosolver can be considered as the front runner for the present-day Indian supercomputers. The fact that Flosolver didn't get attention

"We are on par with European countries" - Sasi Kumar, Director, C-DAC

Sasi Kumar took over in 1993 as the Head of C-DAC's Software Development Centre in Bangalore when the PARAM 10000 project was initiated. In 1994, he became the Director of the Bangalore centre where he is now managing a group of about 120 people. He was instrumental in delivering the software for PARAM 10000 in its present state. Technoworld spoke to him. Excerpts...


Sasi Kumar, Director

Tell us something about PARAM 10000

Param 10000 is a result of our second supercomputing machine project. We have delivered it ahead of time. Our intention was to build a 100 gigaflops machine, which we have achieved through PARAM 10000. The project was started in 1993 and is built on open frame architecture. The same architecture was used in PARAM 9000 but with a different processor. The 10000 uses the Ultraspark-II processor from Sun Microsystems. It has approximately 40 nodes. These are connected by high-speed networks called the system area networks. It is also connected to the fibre optic network of the Poona University campus for other departments to access it. We hope to put it on Internet so that remote users can also log in. The beauty of this system is that it can be used not only as a single parallel machine, but also as individual workstations. It is possible for each and every node to work on different applications or they can all participate in one application. The software for this is developed by this (Bangalore) centre and is called the HPCC software.

What are the applications of PARAM 10000?

It is basically targeted at scientific and engineering applications. Current applications which run on it are seismic data processing, weather forecasting, computational fluid dynamics, medical imaging, finite element analysis and biochemistry.

How much did PARAM 10000 cost?

Generally, our computers fall in the range of 20-25% of the cost of any imported supercomputer. With the cost of updating technology increasing, PARAM 10000 may not be as low as 20-25% of the cost of other contemporary machines. But our software cost is nil since we utilise our in-house expertise. For the first project (PARAM 9000) in 1988, Rs. 36 crores was approved and we spent about Rs. 28 crores. The second project was allotted Rs. 48 crores and we have spent all of that.

What are the general applications associated with supercomputers?

Everybody uses supercomputers for the same applications. It is a strategic technology. You can do real things without experimentation by simulation. Things like drug design, weather prediction, launch vehicle dynamics, etc. cannot be done in any other way except this.

Where does PARAM stand in the scheme of things of C-DAC?

PARAM is our flagship. There are many other things we are doing as a fallout of our doing PARAM. The expertise and technology we have gained in building supercomputers since 1988-89 is helping us in developing power monitoring systems and real-time systems which nobody in India is doing. Today the whole northern and eastern grid is monitored by our systems at 1/10th the cost of similar MNC systems. Moreover, our systems are tuned for Indian requirements. We also develop high-speed networks, hospital information systems, provide internet solutions and language technologies. Motorola uses our technology to send pager messages in Indian languages.

What is the next step?

The next step will be to build a third machine. We will continue to develop supercomputer technologies and create test beds for national information infrastructure. These are included in the 9th plan.

Are you in competition with other countries in supercomputing capabilities?

We are definitely on par with European countries. Only Japan in Asia has faster machines than this, either developed in-house or brought from western countries. America has faster supercomputers than ours only because of the amount of money they spend on it. America is developing a 100 teraflop machine for missile simulation tests to be completed in 2004 at a cost of 50 million dollars. ❑

or financial support from the Government to proceed with its intended programs is a different issue. Although Flosolver was not able to compete with the likes of Cray, which is one of the well-known brands of supercomputers, it was instrumental in demolishing the notion that only the West was capable of developing high-performance computers. Till then the U.S. was the only country capable of developing supercomputers. The Cray supercomputers, apart from being exorbitantly priced, were closely monitored by the company itself. After the success of Flosolver, many parallel processing projects were initiated by the government including the PARAM series

UShering in a new era

India has to be eternally grateful to the U.S., in fact, even more than Pakistan, for its present position among the league of nations. The U.S. has now been serving the technological cause of India more than India had hoped for. Every time the U.S. has tried to impose sanctions on the transfer of essential technologies, India has been successful in developing its own indigenous technologies. The most classic case of the U.S. helping India's technological growth has been in high-performance computing. C-DAC, in fact, was set up primarily to counter the denial of the U.S. company Cray to supply supercomputers to India.

But much before C-DAC could be set up, Indian scientists led by Prof. R. Narasimha had begun a parallel computing project in the National Aerospace Laboratories (NAL). Computational fluid dynamics, which is one of the major activities of NAL, Bangalore, requires an enormous amount of calculations which have to be done to determine the pressure distribution of air over an aircraft. It was taking weeks together to do this, which NAL could ill afford. Instead of looking at the foreign supercomputers to solve this problem, Prof. R. Narasimha started a parallel processing computer project. Since the project was aimed at solving problems associated with the flow of air around the aircraft in motion, it was named 'Flosolver'. With just two Intel 8086 processors, Flosolver's MKI version took shape in 1986. "The joy of seeing two processors, for the first time, communicating amongst themselves to solve a complex fluid dynamics problem was simply inexplicable", says Mr. J.P. Singh, one of the senior scientists at the Flosolver unit. The Flosolver proved to be a remarkably faster machine in comparison to other imported versions. Over the years four versions of Flosolver have been released with improvements in processor speed and communicating environment. Flosolver has been successful in performing weather prediction calculations, direct numerical simulation, and hypersonic simulations. Apart from these, Flosolver has handled quantum mechanic projects from Hitachi, Weather code from IBM and Cryptographic code from the Army. In spite of Flosolver's success, NAL as a part of CSIR has not been able to allocate more funds for the Flosolver unit, which is one of the many units of NAL. With better infrastructural support and greater attention, the Flosolver unit under the able leadership of Prof. U.N. Sinha has a great chance to be a parallel supercomputing powerhouse apart from C-DAC in India.

Comparison of Indian parallel computing systems

Organisation	No. of processors	Time	Peak speed (Mflops)	Sustained speed (Mflops)
C-DoT	128	3 h 20 min	192	12
NAL	1+4	1 h 40 min (24 h forecast)	400	25
C-DAC IIT-Delhi	1+8	1 h 35 min (24 h forecast)	720	26
BARC	1+3	11 min (1/2 h forecast)	320	-
Cray	2	15 min	460	150

As on 23 July 1994

Supercomputer architecture

Today all supercomputers are parallel computers. A sequential processing machine cannot be a supercomputer. Simply because it takes more time to solve a problem sequentially. We have reached a stage where a single processor cannot achieve gigaflops of speed.

In parallel processing there are two technologies:

a) Shared memory architecture (SMA) and

b) Distributed memory architecture (DMA).

SMA is a highly acceptable standard wherein a common memory is shared by several processors. But there is no growth path in SMA, because more than 64 processors are not possible in SMA architecture as there is a serious limitation on the bus.

DMA defeats this. There is no bus in DMA. The processors are connected by a new communicating paradigm.

Any number of processors can be added. One of the machines working in Russia has 256 processors. The communicating media includes fast ethernet, FDDI, fibre channels, etc. on which a network is built and computers are hooked on to it.

This is the architecture used in PARAM.

from C-DAC. CDAC was set up and the PARAM project was initiated as a reaction to the U.S. denial to deliver high-performance computers to India on the pretext that India would use it for defence purposes.

The other projects initiated during the same period were PACE from Bhabha Atomic Research Centre, Chips from C-DOT (both of which are not functional now), ANUPAM again from BARC and ANURAG from DRDO. Most of these parallel computers were restricted to solving their own in-house problems rather than other data critical applications on a nationwide scale. For example, C-DOT built a parallel computer for radio astronomy, ANURAG was built to make fluid dynamical computations and ANUPAM for nuclear computations.

But the real major effort to address varied national problems as far as data intensive calculations are concerned was made by C-DAC which culminated in the development of the PARAM series of computers.

PARAM - the ultimate

C-DAC advented the open frame architecture with the PARAM series of high-performance computers. Prior to PARAM, we had open system and open architectures. Both of them were clubbed together to deliver open frame architecture in PARAM. Open frame architecture is independent of operating system, language and hardware. It ultimately would provide a heterogeneous kind of environment which provides near idealistic solutions to problems. No single processor can handle all applications efficiently. One processor would be suited to perform one particular job at its best and another processor would perform some other application better. So, with a heterogeneous environment all the jobs can be handled efficiently. PARAM open frame architecture would in fact lead to such a system.

PARAM 10000 series of machines are powered by state-of-the-art and emergent SUN's UltraSparc series of Servers/Workstations configured as Computer nodes, File Servers, Graphics nodes and Internet Server nodes. These nodes are interconnected through PARAMNet high bandwidth, low latency network designed in-house and a choice of other high-performance networks such as Myrinet, Gigabit/Fast Ethernet and ATM.

All these networks can be present concurrently for redundancy and fault tolerance (see box item on Real-time systems). The bandwidth of PARAMNet is 400 + 400 Mbps per port.

Supercomputing technologies

Supercomputers adopt different technologies. There are three different kinds of computing, which can be mixed and matched for special applications.

Distributed computing: Here, the jobs are distributed across several computers and they have a common communicating environment.

Parallel computing: Here, a big job itself is distributed into smaller jobs and a single job is done in parallel.

Meta computing: Several computers are connected through a network. It is not seen as several computers for a user, but as a single computer where all calculations are done. It is like a complex machine with several machines inside. For a user, the entire system is seen as a single machine.

In terms of usage, differences are large. To write a program, it is much easier to do it using meta computing since all programming is done on a single machine. In the case of parallel computing, it is not enough to know how to write a program but a clear knowledge of how to distribute the program into split modules and tasks is needed. In distributed computing, it is slightly simplified, but it will waste a lot of power. Of all these, parallel computing is the most popular because of its efficiency.

Real-time systems

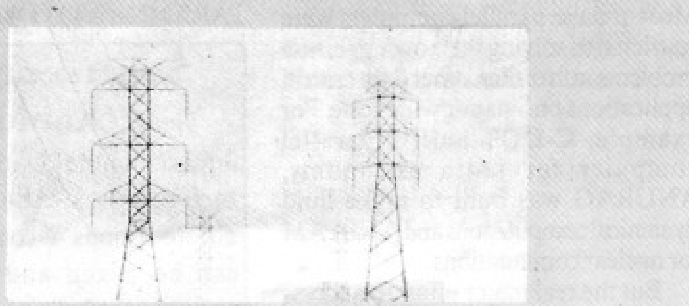

An automated system involves several operations like communications, database handling, number crunching, data acquisition, monitoring, etc. These systems should be operating at 24 hours a day and 365 days a year. In such a scenario, there is no room for failure. One cannot afford to lose precious data. But if the automated systems are capable of tolerating or cushioning such failures, the system would be reliable. C-DAC provides such solutions. It develops 'fault tolerant systems'. The faults which lead to failures are tolerated by a special feature built into the system itself.

There are several mechanisms to tolerate such faults. Some of these mechanisms include building redundant systems, half standby systems, using dual architecture and fault tolerant software. "We have clubbed all these mechanisms and have come up with a fault tolerant system called the UFS - Uninterrupted Functional Service", says Mr. Gangaprasad. C-DAC is in fact working on the next step in fault tolerant systems, i.e. FTS for real-time systems. "We are planning to develop a generic FTS for real-time applications. This is one of the key areas of our development", says Mr. Prasad.

The areas in which real-time systems are behng developed are in water management, power management, transport systems, etc. "In the northern and eastern sectors, our systems are monitoring power distribution and grid management. The whole gamut of activities right from the generation of electricity to the site of its utility is monitored by our systems. Several power generators have to work in synchronisation. This has to be monitored, forecast, managed and economically dispatched. This requires a lot of distributed computing, meta computing and parallel computing", says Mr. Prasad.

The mechanical packaging of PARAM Open Frame is flexible enough to scale from a system having only a few nodes to systems having a very large number of nodes. The PARAM Open Frame cabinet is conceived as a complete, integrated environment for high-performance computing elements. It consists of 19" standard racks, shielded cable trays and accessories for housing all nodes, network switches, associated cables, power sequencers and status displays.

Parallel processing machines, as in the case of the other initiatives, are normally customised to handle specific applications. But PARAM systems are generic machines. A generic system concept would be the concept of the future. For instance, an expert in a steel industry would have a clear knowledge of the nuances of his industry but will have little knowledge of computers. In the same way, a computer expert developing systems for application in a steel industry would have little knowledge about the problems associated with a steel industry. So there is no possibility of building a specific machine for a specific application. We provide the basic things in a generic machine and functional experts have to get back to us with their requirements which we can incorporate" says Mr. Gangaprasad.

What does PARAM mean to India?

PARAM 10000 is a 100 gigaflop machine scaleable to teraflop range. With this, India has become the third country in the world capable of building a teraflop range of machine, the other two being the U.S. and Japan. The major applications associated with PARAM include:

- Weather forecasting and climate modelling
- Remote sensing
- Seismic data processing
- Computational fluid dynamics
- Finite element modelling
- Launch vehicle dynamics simulation
- Medical imaging
- Computational chemistry
- Molecular dynamics
- Biotechnology
- Parallel databases
- Data warehousing

However, the greatest application for which PARAM has made international news is its ability to simulate conditions for nuclear testing which would eliminate the physical testing to ascertain the nuclear capabilities of India. This is a major concern for the U.S.

PARAM is now all set to achieve greater heights. The next five-year plan has already allocated substantial funds for the next project in the PARAM series. PARAM has assured Indians that they neither lack the technology nor the fighting spirit to reach the top of the world. ❑