THE EVOLUTION OF ANUPAM SUPERCOMPUTERS

A.G. Apte, R.S. Mundada, B.S. Jagadeesh, K. Rajesh, Kislay Bhatt, D.D. Sonvane, Vaibhav Kumar, Vibhuti Duggal and P.S. Dhekne
Computer Division, BARC

Computers have become ubiquitous in all walks of scientific endeavour. There have been innumerable architectures of computers, aiding equally countless, diverse fields of research, thereby rendering the task of classification of computers itself a nontrivial job. Interestingly, this gamut of computers are classified neither by their explicit features nor by the class of application to which they cater to, but by the way the data and instruction streams flow through them. Specifically, this classification identifies whether there is single stream or multiple streams of data and instructions, which flow through them. Thus, we have Single Instruction Single Data (SISD), Single Instruction Multiple Data (SIMD), Multiple Instruction Single Data (MISD) and Multiple Instruction Multiple Data (MIMD) categories of computers.

Fig. 1: Anupam-Ajeya delivers a sustained performance of 9 teraflops

The conventional sequential machines belong to the SISD category. To speed up computations in SISD machines, one has to either speed up instruction processing (use faster processors) or push data around at a faster rate (use faster system bus, memory, cache etc.) or employ both methods. But their speeds are limited by the hardware technology at that given point of time and the improvements too will depend upon technology development and hence, performance gain will also be incremental in nature. Everything being equal, if one has to derive large performance gains, then one will have to make use of multiple processors, working on multiple sets of data, interconnected through some medium. Machines with specialized functional units, to process multiple chunks of data concurrently (Processor Arrays, Vector Processors) belong to the category of SIMD. In this architecture, the same instruction is executed concurrently over multiple datasets, in lock-steps. This is suitable only for very specific type of applications. MIMD is another interesting architecture wherein, multiple processors execute different instructions on different sets of data independently and yet work in unison, to solve a single problem. This architecture has come to be known as parallel processing. Parallel processing has emerged as the defacto method of achieving performance, that is of greater orders of magnitude than the one that is obtained by single computer at any given point of time. This is the most flexible architecture among the four categories of computers. All modern parallel supercomputers fall in this category. For the sake of comprehensiveness, we can state that there have been only a couple of MISD class of machines and these have remained of academic int erest only.

About fifteen years ago, the development of the Anupam Series of Parallel Supercomputers was started with the goal of achieving computing speeds that were ten times more than the speed of the sequential machines available at that time. BARC has a large pool of Scientists and Engineers working in various aspects of Nuclear Science and Technology and whose computational needs are diverse in nature. To cater to the computational needs of this diverse set of users, it was decided to build a general-purpose parallel computer, rather building many application-specific clusters. To keep the gestation period short, it was decided to build the parallel computer with commercially available off-the-shelf components, with our major contribution being in the areas of system software development, application software development, system engineering, system integration and fine tuning of the system.

The performance of a parallel processing system, essentially depends upon the processor speed and the performance of the interconnect, that glues the participating processors. Earlier versions of ANUPAM systems were based on bus-based (Multibus-II) architecture and i860 processors. This was indeed a simple design, wherein all the processors were latched on to the Multibus-II bus. To initiate the communication, the sender had to gain control of the bus and then send the data to the receiver. Using the ANUPAM systems based on this architecture, a large number of research problems from various domains were resolved (Molecular Dynamical Simulations, Reactor Physics, Theoretical Physics, Computational Chemistry, Computational Fluid Dynamics, Finite Element Analysis to name a few).

As the resolutions of the computational models increased and the models themselves became far more complex, the amount of data that needed to be transferred across also became very huge. This architecture was found wanting as it could not cater to the newer demands. As mentioned previously, in the bus-based architecture, one had to gain control of the bus to perform the data transfer. As only one processor could be the bus master at a given point of time and perform the data transfer, the other processors had to wait to gain control of the bus, even though they had completed their computations and were ready to effect the data transfer. This meant, that even though the computations were done in parallel, the communications remained sequential in nature and hence the turnaround times of the programs did not improve. With the result, these complex, high-resolution models did not scale appreciably.

The issue was to effect the communication that had remained sequential in nature, to parallel. There were two alternatives. The first one was to make use of multiple busses, each of which interconnected a given set of processors. The communication in these multiple busses could go in parallel. The second alternative was to go in for the switch-based architecture wherein, the communications could proceed in parallel, in machines connected to different ports of a switch; the limiting factor being the switching speed of the backbone fabric of the switch. With this, the communication could be carried out in truly parallel manner. This design has come to stay as the typical interconnect model for ANUPAM series of parallel supercomputers. A parallel supercomputer based on Asynchronous Transmission Mode (ATM) interconnect was built in 1997.

With substantial improvement in the performance of PC-based workstations and developments in LAN technologies, building High Performance Computing (HPC) cluster, using these state-of-the-art commodity components, became a very cost-effective approach. We have built quite a large number of clusters using PC-based workstations as compute nodes and commodity Fast and Gigabit Ethernet switches as interconnection networks. Anupam-Ashva, a 64-node system, Anupam-Alpha which replaced CRAY XMP216 at NCMRWF, Mausam Bhavan, NewDelhi, Anupam-Ameya, a 512 processor cluster that delivered 1.73 Teraflops of sustained performance to name a few. Insofar as the interconnects are concerned, the factors that govern their efficacy in parallel processing are the Bandwidth of the point-to-point links and the latency in sending messages from one system to another, in the network.

With ever increasing processor speeds, LAN technologies like Fast and Gigabit Ethernet that are conventionally used as cluster interconnects are fast becoming inadequate, in extracting the full potential of the cluster. To cater to the specific demands of High Performance Computing (HPC) cluster interconnects, there is a growing breed of new interconnect technologies called System Area Networks (SANs). These new technologies incorporate characteristics such as high bandwidth, low latency for communication and scalability to a large number of nodes, that are very essential for HPC cluster interconnects. LAN technologies are aimed primarily at providing connectivity among computers spread across geographical distances of a few kilometers. Whereas SAN technologies are restricted in geographical scope with spanning distances of only up to a few meters, suitable for networks within a system. The various SAN technologies available today include InfiniBand, Myricom Myrinet, Quadrics QsNet, QsNetII and Dolphin SCI. Myrinet and QsNet are proprietary technologies whereas SCI and InfiniBand are industry-standard interconnect architectures. We have used both SCI and InfiniBand SAN technologies as cluster interconnects in ANUPAM systems.

SCI (Scalable Coherent Interface)

This is an interconnect standard (IEEE 1596) for high performance networking which aims to provide high bandwidth, low latency and low cpu overhead for communication operations. An SCI interconnect is defined to be built only from unidirectional point-to-point links between participating nodes. This feature of the SCI links makes it possible to achieve high bandwidths. In contrast to a LAN, the SCI provides hardware-based physical Distributed Shared Memory (DSM), thus exhibiting some characteristics of a Non Uniform Memory Access (NUMA) machine. Because of this architecture, internode communication translates into simple CPU load and stores into DSM segments, which are mapped from remote node memories. Hence, there is no need for a protocol stack, which results in low latencies for communication. The SCI standard, specifies a bandwidth of 1 GB/s but current implementations achieve a link speed of 667 MB/s. SCI clusters can have many topologies such as ring, switch and torus. We have used SCI as an interconnect technology in our 64-node, 128 CPU ANUPAM-ARUNA cluster with an 8x8, 2-D torus topology. This delivered a sustained performance of 365 Gflops.

InfiniBand Architecture

InfiniBand is an industry-standard specification, that defines input/output architecture used to interconnect servers, communication equipment, storage and embedded systems. The InfiniBand specification is defined and maintained by the standards organization InfiniBand Trade Association (IBTA), an industry consortium of hardware and software vendors. When IBTA was formed in 1999, the primary goal of InfiniBand was to replace the PCI bus. Over a period of time, positioning of InfiniBand evolved from being a PCI replacement to being an HPC interconnect for high-speed I/O, networking and inter-process communication.

InfiniBand is a point-to-point, switched fabric, serial I/O interconnect architecture, that can be used for backplane solutions (intra-box) as well as for providing external (inter-box) systems interconnects. It is a low-latency, high-bandwidth interconnect which requires low processing overhead and is ideal for carrying multiple traffic types (clustering, communications, storage, management) over a single connection.

The communication is full duplex and the InfiniBand specification defines 3 link widths, 1x, 4x and 12x. The basic 1x link consists of one high-speed transmission line in each direction (4 wires for bidirectional communication) with a signalling rate of 2.5 Gbps Single Data Rate (SDR). InfiniBand supports double data rate (DDR) and Quadruple Data Rate (QDR) speeds, that is,5 Gbps and 10 Gbps respectively, at the same clock rate. Infiniband uses 8b/10b encoding. This yields a net data rate of 250 MByte/s per direction (SDR). Links with 4x and 12x widths are multiples of 1x link and accomplish raw data rates of 10 Gbps and 30 Gbps (SDR). A raw data rate of 120 Gbps can be accomplished with a 12x QDR link. Table1 shows the raw data rates in InfiniBand for various link widths and signalling rates. Maximum packet payload is 4 KB. As a medium,

Table 1: Raw data rates in InfiniBand

Data Rates	1x	4x	12x
SDR	2.5 Gbps	10 Gbps	30 Gbps
DDR	5 Gbps	20 Gbps	60 Gbps
QDR	10 Gbps	40 Gbps	120 Gbps

InfiniBand defines various copper and fiber-optics cables. A maximum length of 17 meters (SDR, 10 meter for DDR) is specified for copper cable and up to 200 meters for fiber-optic cable.

The latency measured at MPI level is usually 3 to 4 microseconds. An important feature of InfiniBand Architecture that reduces the application level latency and CPU overhead is the support for Remote Direct Memory Access (RDMA). With traditional networks such as Ethernet, the communication between applications is relatively cumbersome. Incoming data is accepted by the network card, processed in the kernel of the operating system and finally delivered to the application. As part of this, data is copied repeatedly from one buffer to the next. Furthermore, several process changes are necessary in the operating system. All this costs CPU cycles and places a load upon the system bus, thus reducing the communication throughput and increasing its latency.

Remote Direct Memory Access (RDMA) is a communications technique, that allows data to be transmitted from the memory of one computer to the memory of another computer, without passing through either CPUs, without needing extensive buffering and without calling to an operating system kernel. When an application performs an RDMA Write request to the memory on the remote system, the RDMA request is issued from an application running in user space, to the local Network Interface Card (NIC). The "RDMA Engine" on the NIC uses DMA to read data from the user-specified buffer and transmits it as a self-contained message across the network. The "RDMA Engine" on the receiving NIC then uses DMA to place data into the user-specified memory location. There is no intermediary copying and all these operations occur without the involvement of the CPUs reducing the latency and CPU overheads. We are using the InfiniBand technology as the cluster interconnect, in the latest Anupam-Ajeya Supercomputer.

ANUPAM-Ajeya Supercomputer

ANUPAM-Ajeya is the latest in the series of Anupam parallel supercomputers. It deploys 1152 processor cores as its computational workhorses, delivering a sustained performance of 9 teraflops. It is based on the concept of Cluster of Workstations. It is a compact, centralized, homogeneous Linux cluster, comprising of 288 dual-core, dual-processor nodes, interconnected by InfiniBand as well as Gigabit Ethernet networks. A logical view of ANUPAM-Ajeya architecture is shown in Fig. 2.

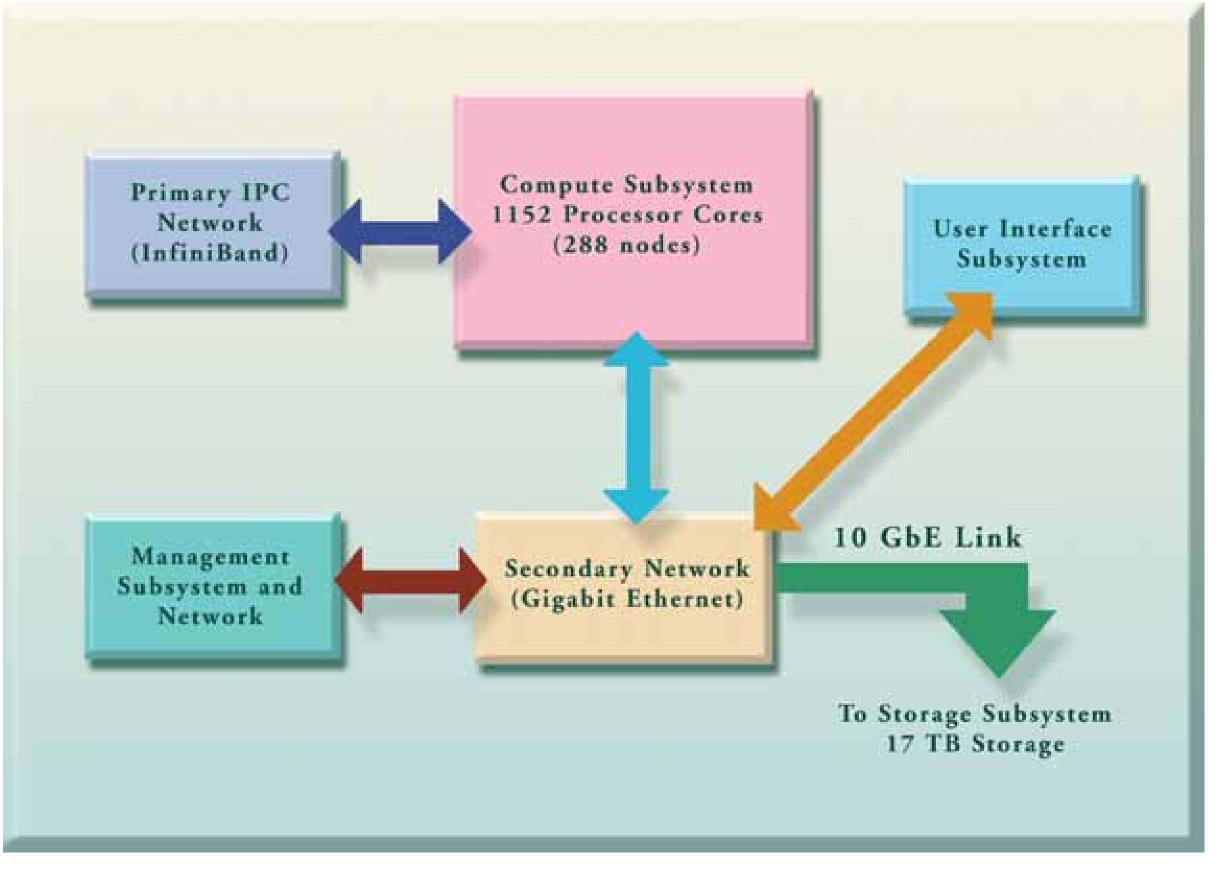

Fig. 2: Block diagram of ANUPAM-Ajeya architecture

The ANUPAM-Ajeya Supercomputer consists of the following subsystems:

Compute Subsystem

The Compute Subsystem is the major subsystem of a Supercomputer, which defines the performance of the system. This subsystem consists of Dual Core Dual Xeon systems with 1U form factor as compute nodes, leading to a total of 1152 CPU Cores. Each node has 4 GB memory, a PCI Express x8 slot and 2 Gigabit Ethernet ports. The Operating System on all the nodes is Scientific Linux 5 and parallel programming environment is provided by MPICH, MPICH2, OPENMPI, PVM and our in-house ANULIB libraries.

Interconnect Subsystem

The system has two interconnect networks, the primary Inter Process Communication (IPC) network is an InfiniBand SAN (System Area Network) and the secondary one is a Gigabit Ethernet Network. The InfiniBand SAN comprises of a 288 port, 4x DDR (20 Gbps), InfiniBand switch interconnecting the 288 compute nodes using the InfiniBand Host Channel Adapter (HCA) cards (installed on the PCI Express x8 slot on each node) and InfiniBand copper cables. The Gigabit Ethernet Network consists of six 48-port and two 24-port Gigabit Ethernet switches, stacked to form a 336-port switch. All the switches are powered through UPS. This Gigabit Network is connected to the Storage Subsystem using 10 Gbps Ethernet link.

The InfiniBand SAN is used explicitly for inter-process communication, whereas the Gigabit Network is used mainly for accessing the file servers over NFS and for the data traffic pertaining to cluster monitoring and management activities. The Gigabit Network can also be used as a secondary inter-process communication network.

Management Subsystem and Network

Each node in the cluster is equipped with an IPMI (Intelligent Platform Management Interface) Card. IPMI is used for remote management and monitoring of the nodes. Using IPMI, SOL (Serial over LAN) is configured to access the console and BIOS of the nodes over the LAN port, avoiding the wiring for serial ports or KVM switch.

All the Gigabit switches and compute nodes are powered through Power Distribution Units (PDU) in order to control the power to the devices through software, i.e. starting in sequence (in order to prevent sudden surge of current at system startup), selectively powering ON/OFF or resetting the device. These PDUs are connected using a dedicated management network comprising of Fast Ethernet Switches. This network is connected to the secondary network using a Layer 3 switch and VLAN technology.

User Interface Subsystem

This subsystem consists of user terminals to which users can log in directly or through intranet, to compile, submit and monitor their jobs.

Storage Subsystem

All ANUPAM systems (namely, Anupam-Ajeya, Anupam-Ameya and Anupam-Ashva) installed at the Supercomputing Research Facility building share a common storage Subsystem in order to enable users to seamlessly use all the clusters. The Storage Subsystem consists of file servers, backup servers and tape libraries.

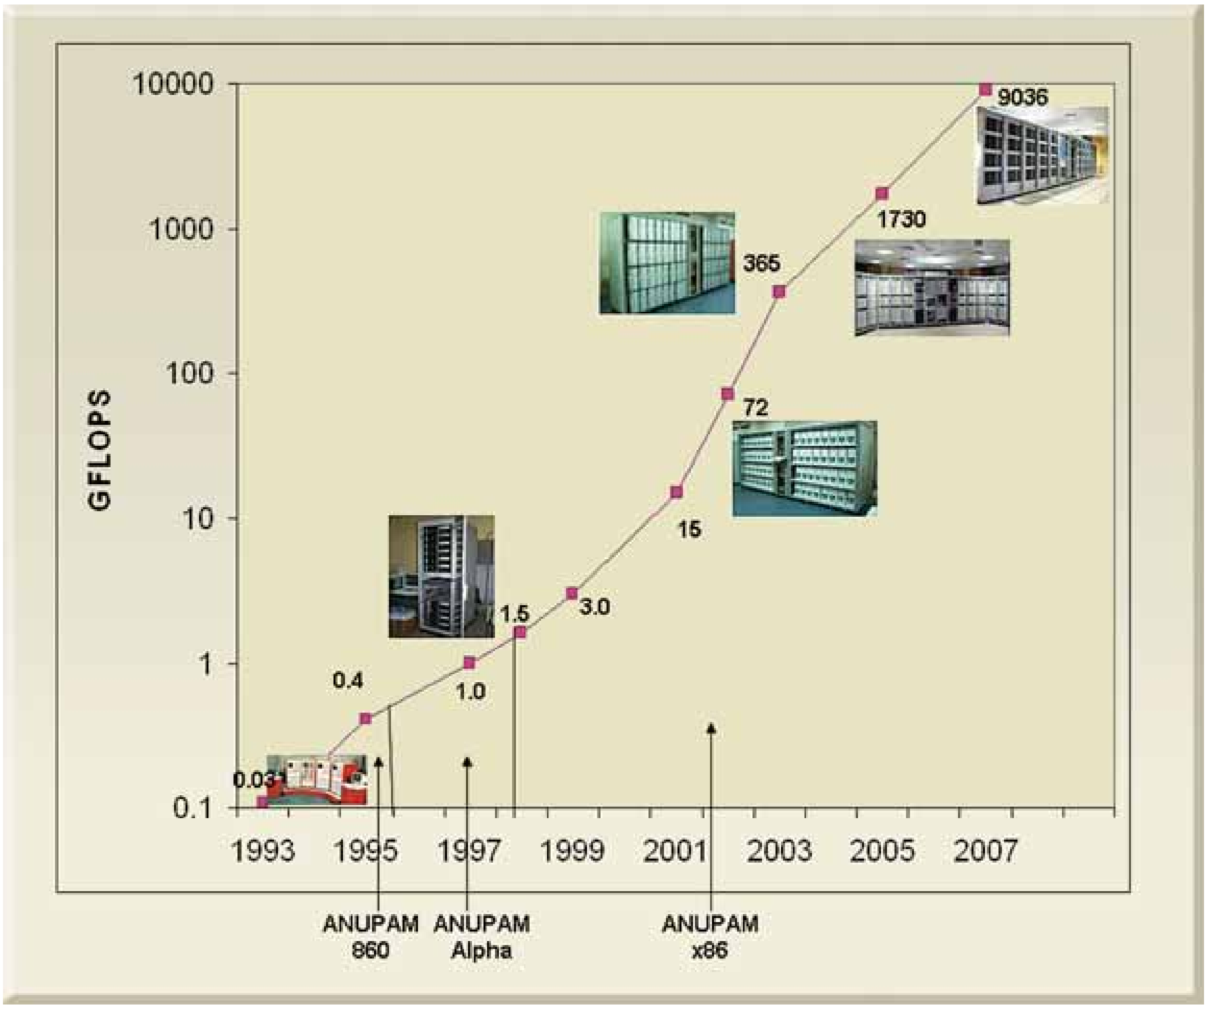

File Servers: There are 12 file servers, which constitute 17 terabytes of storage space. Each server is a 2U rack-mount server, equipped with dual processors.

Backup Servers: There are two backup servers for taking backup of users' data. Each server is a 5U rack-mount server equipped with dual processors and backup devices like DVD writer, DAT and DLT drives. Besides that, each server is connected to a tape library. The backup servers offload the backup load from the main file servers.

Tape Libraries: There are two tape libraries with a storage capacity of 3.2 terabytes each. The backup servers are programmed to take periodic backup of file servers to their local disk and then copy it on to the tape library.

Design challenges for the storage subsystem are performance, reliability and availability. Necessary redundancy is provided to reduce failures and downtime. Moreover, system design ensures minimal effect on services in case of a server failure. RAID is configured on each server to overcome single disk failures. Each server has three Gigabit Ethernet network ports, which have been link-aggregated to increase the availability and throughput three fold. The users are distributed across the 12 file servers so that a single server failure affects only a fraction of users while others can still continue to use the system.

System Layout

The layout of the ANUPAM-Ajeya system is planned in such a way, that it ensures efficient cooling and optimal use of space, without obstructing future repairs and maintenance. The whole system is housed in 12 racks of height 42U, closely aligned in a straight line. Ten racks house compute nodes, each of which contains 28 nodes and four PDUs arranged in 4 groups of 7 nodes each. The nodes within a group are tightly stacked without any gap. There is a 2U gap across the groups to accommodate lengthy cables. We have sealed the gaps on the front side with blanking panels to prevent cold air from escaping through the gaps. Two racks that contain switches, UPS and the 8 service nodes are placed in the centre for symmetry and for reducing the cable lengths.

The Ethernet cables are routed from the top of the racks and the InfiniBand cables that are heavy in weight, are routed from the bottom of the racks. All cables are properly laid through the ducts at the backside corners of the rack, so that, they do not obstruct the airflow. Cold air is provided through grilled tiles in false flooring, right in front of the racks, to ensure proper cooling.

Performance Measurement / Benchmarking

Performance of a Parallel System is measured in terms of GigaFLOPS (GFLOPS) or TeraFLOPS (TFLOPS), where FLOPS stands for FLoating point Operations Per Second. When performance is calculated using the CPU's clock frequency and average Cycle Per Instruction (CPI) rating, it gives *peak performance* of the system. Where as, the *sustained performance* of a system is measured by running benchmarks or real programs on actual machines. In this case, the true performance of a computing system involving the processor, the memory, the peripheral devices and the interconnects as an integrated unit will be measured.

Most of the published GFLOPS and TFLOPS results are based on running LINPACK benchmark code. LINPACK is a general purpose Fortran library for solving dense system of linear equations. We have used a High Performance LINPACK (HPL) benchmark to evaluate the performance of the ANUPAM-Ajeya system. HPL is run for a matrix size of 3,40,000 X 3,40,000 on 1152 cores of Ajeya system and it gave a sustained performance of 9.036 TeraFLOPS. The performances of the various Anupam Systems, developed over the past 15 years, have been depicted in Fig. 3.

Issues in deploying a cluster in production environment

Building a cluster with novel features, that caters to a given application, is entirely different from building a cluster that will have to go into a production environment. It needs to be run on round-the-clock basis throughout its lifetime, catering to a wide variety of applications. Furthermore, when a cluster is relatively small, (say <32 nodes) its installation, management, monitoring and archival taking corrective action on degradation of environmental conditions and troubleshooting, can be carried out manually. But when the cluster size grows to encompass a large number of components (few hundreds of processors/disk drives, scores of PDUs, scores of Terminal concentrators, scores of switches, fileservers and their associated cooling components), automation becomes a necessity for monitoring their health. The total number of components that needs to be monitored for ANUPAM-AJEYA is about 3000. It goes without saying that other operational tasks like installation, configuration management, job scheduling, accounting of system utilization, archival, power-recycling of given node/ nodes etc. of this large cluster, need to be automated. All these tools have been developed in-house and have been deployed in production clusters.

AnuInstall

One of the important tasks in setting up a cluster, is software installation on nodes and servers and configuring them according to their roles in the system. For a large cluster, having hundreds of nodes, it is a very tedious (and also error-prone) process if one has to install and configure each and every node manually, as it involves several tasks like partitioning of hard disks, selection of appropriate system packages, configuration of network and several post installation tasks like configuration of required services, installation of compilers and other required packages, etc. Therefore, there is a need of an automated installation tool to perform the above tasks in an absolutely error-free manner. For this purpose, *AnuInstall* has been developed.

Various computing elements in Anupam-Ajeya have been categorized into four groups depending on their functionalities namely compute nodes, storage servers, service nodes and backup nodes. ANUPAM-Ajeya has 280 Compute Nodes, 12 Storage Servers, 8 Service Nodes and 2 Backup Servers. Computers belonging to each of these categories would require different configuration detailing. For the automated installation and configuration of these nodes, we have developed AnuInstall tool in-house. This tool uses network installation feature of Linux using PXE LINUX. However, this tool does not keep a separate copy of the kickstart file for each node; instead it generates the kickstart file of the node by a CGI script based on the node's IP address and hostname. The kickstart file is requested using HTTP and the web server runs CGI scripts to generate the kickstart file based on the requesting IP and the configuration is stored in the database. Based on IP address, hostname is resolved and on basis of the hostname, closest matching template of options is selected. List of services to be run is also generated based on the hostname. At the end of the installation, a log file of installation is uploaded on the central server for checking successful installation. An email is also sent to the administrator when installation is completed.

Installation of new nodes is a bit tricky because the mapping of the node's MAC address to its IP address is not known. For this, installation is done in two steps. In the first step, a minimal OS installation is done and hostname to MAC address mapping is registered and in the second step complete OS with full configuration is installed.

Monitoring and Management Tools

For cluster monitoring and management, we have developed the monitoring tool *Anunetra*. Monitoring the nodes, reporting errors, generating alerts and facilitating centralized management are its major functions. In general, it provides centralized interface to perform health check activities in the cluster and other basic management tasks.

Anunetra continuously monitors Ajeya system and reports errors in case of any deviation from normal operating values. Metrics like CPU temperature, fan speed, disk usage, CPU usage, CPU count, CPU speed, swap used, memory free, bytes in/out, packets in/out etc. are under continuous monitoring. The management interface of Anunetra allows the administrator to perform some common tasks, like making the nodes offline/online, network connectivity check, services check, available disk space, etc.

Anunetra uses the publisher and transport mechanism of Ganglia (a public domain utility software). Other modules such as collection, archiving and presentation are designed according to our needs and constraints. Other than common monitoring features, it incorporates facilities such as Reports, Analysis, Alert System, Auto Restart of failed nodes, Resource and Job Information, Management Interface and so on.

Recently, an *Environment Monitoring tool* has been developed, to monitor the temperature and

Fig. 4: Snapshots of various monitoring tools in Anupam Systems

humidity around a large cluster. For the implementation of the system, six temperature and humidity sensors are deployed at various positions of the cluster to monitor temperature and moisture present in the air. These sensors are connected to Environment Management module, which provides sensors' data on the network. The central monitoring program polls the Environment Module Manager using SNMP to collect data for all the sensors. These values are then compared with preset threshold values and SNMP trap is generated. Trap handler, running on a remote machine on operator's desk, switches on an audio hooter to alert the operator on duty. This may be well appreciated when we note, that only a few operators manage a large number of clusters, housed in different buildings on our huge campus, on a round-the-clock basis. These operating personnel who manage multiple sites reside at a central place during silent hours.

Job Management System

To manage jobs on the Ajeya cluster, we have Job Management System, which is based on TORQUE Resource Manager and MAUI Job Scheduler. These tools are complemented by augmenting them with in-house developed codes as per our needs. User jobs have different types of resource requirements namely, some jobs need a large number of processors whereas some saturate at a small number of processors, some jobs run for one or two days whereas some last up to a week, some jobs have to run on some subset of nodes due to software licenses binding and so on. Considering all these requirements, different job-queues are provided with different scheduling and execution policies for each.

Based on demand and available resources, the Scheduling Policy for jobs in Ajeya is optimized. It is designed to avoid resource starvation for large jobs and better utilization of the available resources. Throttle limit for scheduling maximum jobs per user at a time is set and fair-share policy provides a fair portion of the cluster's resources to each user, under heavy load conditions. The Scheduler schedules two jobs per user at a time, but during heavy load, it dynamically decides jobs per user, based on historical data of previously executed jobs.

Accounting Tool

We have developed an Accounting Tool to get information about Ajeya's utilization. One can find system utilization (current as well as previous day's / week's/ month's), users' statistics, current status of the nodes, status of current jobs, etc., using this tool. It has a centralized database that records the information about each and every job submitted to the cluster. Database is populated by data-entry programs, installed on all the compute nodes, which are invoked on submission, startup and termination of jobs. Job parameters such as username, job-name, number of nodes demanded, parallel API used, job's start and end times, are recorded in the accounting database.

To generate meaningful reports out of the collected data, an interactive web interface has been made available, using which, one can extract and display the following information in tabular as well as in graphical form.

- Daily and monthly reports of utilization of individual nodes as well as whole cluster
- Daily and monthly reports of user wise system utilization
- Cluster wise quick reports for getting cluster utilization of the past twelve months in one go
- Average waiting time of a job in a queue.

AnuSakshi

There are a large number of components in Ajeya System like Compute Servers, File Servers, Network Switches, PDUs, UPS, etc. each having a number of subcomponents like Hard disks, Processors, Power supplies, etc. It is difficult to manually keep track of all faults and maintenance operations done on these components / subcomponents. Therefore, AnuSakshi was developed to automatically log all these events and maintain a history of these events along with remarks explaining the events.

AnuSakshi is used to derive the performance of the system and get the availability of the whole system or a component at any given time. It also keeps track of all the faults occurring in any component and the corresponding maintenance operations carried out to rectify them and also the corresponding change of status of components arising or movement of components taking place between BARC and vendors. It also generates a list of the most faulty nodes along with their corresponding number of failures.

Applications

The ANUPAM computers have been extensively used in-house to develop compute-intensive applications such as Ab-Initio Molecular Dynamics, Monte Carlo simulations, Finite Element Analysis, Particle Tracing, Neutron Physics, Computational Fluid Dynamics etc. In fact, Anupam-series of parallel computers have become the main workhorse for the computational work of the scientists and engineers of BARC. In addition to this, thirty-seven supercomputers of

ANUPAM series have been installed at leading R&D and educational institutions in the country including the National Center for Medium Weather Forecasting, New Delhi (currently operating from Noida, UP) and the Aeronautics Development Authority, Bangalore. The R&D work on ANUPAM has found its way into many prestigious publications. Scientists and Engineers of BARC and other institutes who have used ANUPAM for their research work, have published a number of articles in prestigious journals. Fig. 5 depicts a set of applications developed using Anupam.

Acknowledgements

One of the important tasks in commissioning a large parallel computer is the preparation, planning, routing and deployment of a very large number of cables (which easily run into thousands) of various types in a confined space. All these cables were prepared, tested and deployed by our colleagues Mr. Kishor Koli, Mr. Samir Adhikari, Mr. Suresh Vasa and Mr. Manish Duble under strict time schedules. Their efforts deserve special mention. It is our pleasure to acknowledge the unstinting support lent by the Technical Services Division, BARC in providing specialized earthing and meeting our power and air conditioning requirements at short notice.

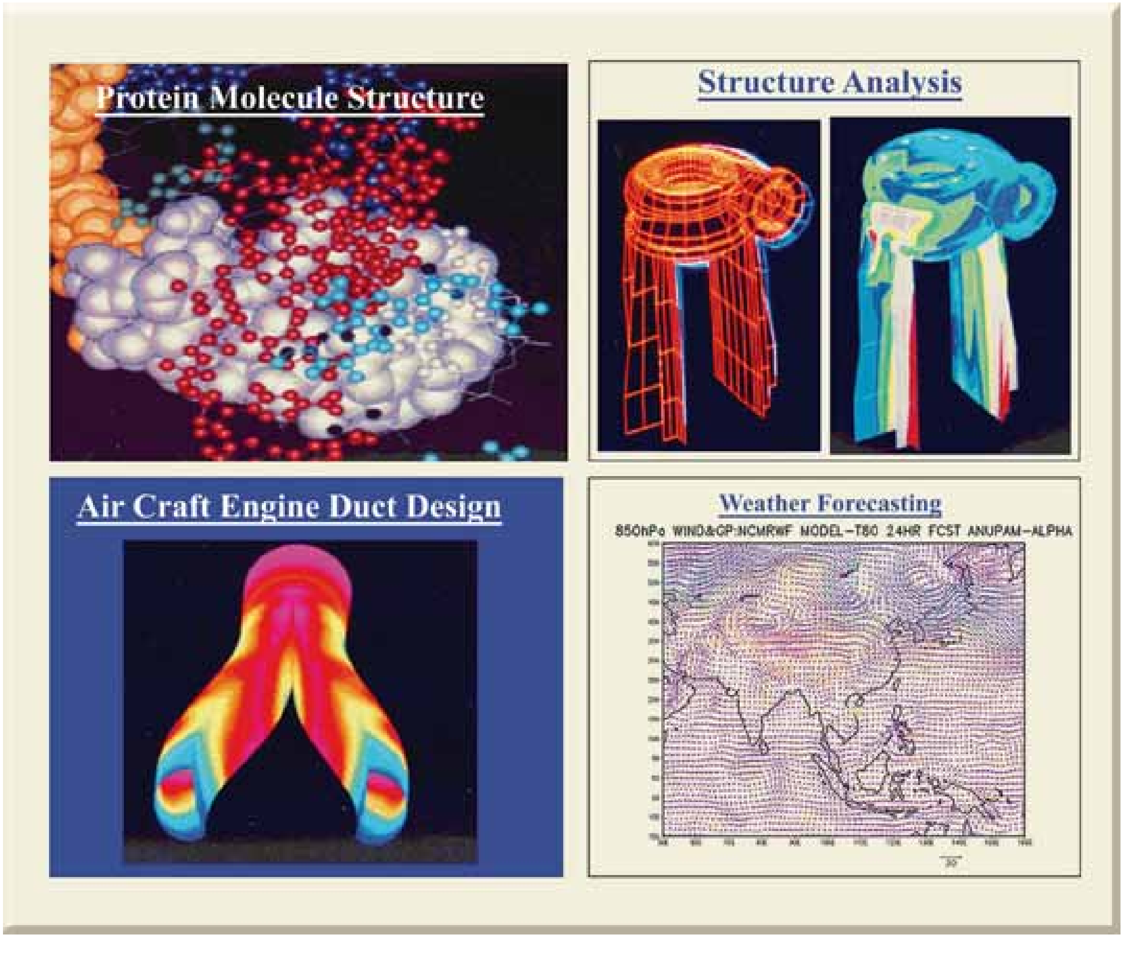

Fig. 5: Representative applications developed using Anupam

There was a presentation on the nuclear symmetry energy in the framework of Brueckner-Hartree-Fock (BHF) formalism, that leads to determination of the neutron skin thickness of ^{208}Pb to be 0.2 fm and the radius of 1.4 solar mass neutron stars as ~21.3 km. There was a review talk on "Quark-Gluon Plasma-Present and Future", covering the frontier area in the field of relativistic heavy ion collisions. The possibility of probing the QCD critical point with an energy scan at RHIC and FAIR facilities was also discussed.

The results from the ongoing international programs associated with STAR, PHENIX, ALICE, CMS, FAIR etc. were presented at the symposium. Many groups from BARC, VECC, Bhubaneswar, Jammu, Jaipur, Delhi, Aligarh etc. are participating in the above programmes.

In the nuclear instrumentation sector, there were talks on a) Advances in gas avalanche detectors and their applications, b) Magnetic separator for light RIB production, c) INGA and NAND instrumentation, d) High energy gamma ray spectrometer and e) A data acquisition system for pelletron-LINAC experiments. The activities and status of various accelerator facilities in the country were updated. The superconducting cyclotron, a major accelerator facility at VECC is expected to be commissioned in 2008. The work on a RIB facility using the 88" variable energy cyclotron at VECC is also in progress. The upgradation of the present pelletron accelerators of both IUAC and BARC/TIFR with the addition of LINAC to boost the energy of the ion beams have made significant progress.

Orientation program

A one day pre-symposium orientation program on **"Introduction to Hadron and Neutrino Physics"** was arranged on 10th December, the preceding day of the main symposium. This was aimed at the students and young researchers to orient themselves through special lectures and intensive interaction with scientists who are experts in the field.

Evening Lectures

There were two popular evening (semi technical) lectures by Prof. V. S. Ramamurthy and Mr. Abasar Beuria. Prof. Ramamurthy (former secretary, DST, Govt. of India, presently DAE Homi Bhaba chair Professor, IUAC-Inter University Accelerator Centre, New Delhi) stressed the fact that **researchers should look for areas where observations and/or theoretical results are apparently anomalous.** He emphasized that giving up such anomalous results which do not confirm to the prevalent scientific wisdom, may actually hide more profound, novel aspects and therefore should be pursued at any cost. To substantiate this point of view, he cited several path-breaking discoveries which were the results of **chasing** such **"anomalous findings".**

Mr. Beuria, IFS, in his evening talk on **"Contemporary world and cultural crisis"** highlighted the onslaught of the currently witnessed phenomenon of globalized market economy on the cultural values, ethnicity and the identity issue of the various communities. He advised the intellectual community to imbibe/ absorb the essence of the brighter and progressive aspects of the different global communities while retaining their own identity, cultural heritage and moral values.

Awards

Following the tradition of the symposium, the best presentations on a) thesis and b) poster were awarded. As a first step, a panel of judges selected the three best theses for oral presentation. One of the three best theses was then selected for the **"C. V. K. Baba Best Thesis Award"**, which was earlier known as "The IPA best thesis award". In the poster category, there were three more awards. A different set of judges chose the three best posters out of the 173 presentations for the **"Best poster award".**