KSR INSTITUTE FOR ENGINEERING AND TECHNOLOGY

TIRUCHENGODE -637 215.

DEPARTMENT OF INFORMATION TECHNOLOGY

VOLUME 2
ISSUE 1

August 2014

DIGITIMES

High Performance Computing

KSR INSTITUTE FOR ENGINEERING AND TECHNOLOGY

Vision

To become a globally recognized Institution in Engineering Education, Research and Entrepreneurship.

Mission

M1	Accomplish quality education through improved teaching learning process
M2	Enrich technical skills with state of the art laboratories and facilities
M3	Enhance research and entrepreneurship activities to meet the industrial and societal needs

DEPARTMENT OF INFORMATION TECHNOLOGY

Vision

To produce competent Information Technology Professionals and Entrepreneurs with ethical values to meet the global challenges.

Mission

MD1	Impart quality education with ethical values in Information Technology through improved teaching learning process
MD2	Provide an ambient learning environment using state of the art laboratories and facilities
MD3	Encourage research and entrepreneurship activities to meet the dynamic needs of Information Technology industry and society

Program Educational Objectives (PEOs)

PEO	Key Words	Description
PEO 1	**Core Competency**	Graduates will be successful professionals in career by applying the knowledge of mathematics, science and engineering with appropriate techniques and modern tools.
PEO 2	**Professionalism**	Graduate will exhibit soft skills, professional and ethical values and thrust for continuous learning to maintain professionalism in the IT industries.
PEO 3	**Higher Studies and Entrepreneurship**	Graduates will engage in higher studies and outshine as entrepreneurs through life-long learning which leads to societal benefits.

DIGITIMES

Editorial

We would like to wholeheartedly thank our honorable Chairman, **Lion.Dr.K.S.Rangasamy** and vice chairman **Mr.R.Srinivasan**, and Principal **Dr.M.Venkatesan** for their continuous encouragement and constant support for bringing out the magazine. We profoundly thank our Head of the Department **Dr.P.Meenakshi Devi** for encouraging and motivating us to lead the magazine a successful one right from the beginning. DIGITIMES serves as a platform for updating and enhancing upcoming technologies in Information Technology. We are also grateful to all the contributors and faculty coordinator to bring this magazine.

By,

Editorial Board

CONTENTS

HIGH PERFORMANCE COMPUTING

What is high performance computing?

High Performance Computing most generally refers to the practice of aggregating computing power in a way that delivers much higher performance than one could get out of a typical desktop computer or workstation in order to solve large problems in science, engineering, or business.

It turns out that defining "HPC" is kind of like defining the word

"car" — you probably know what a car is, but I bet you'd be hard pressed to write a concise, simple definition of one that means anything. HPC is actually used in two ways:

It can either mean "high performance computing" or "high performance computer."

It's usually pretty clear from the context which sense is being used.

BY

KEERTHANA R IV Year/IT

HPC Proliferation

"There is an entirely new class of users who do not know they are using HPC. We call it 'implicit HPC.'"
– Raj Hasrah, corporate VP and GM, Intel Corporation

"When someone says HPC it means something really specific to traditional HPC folks; it's tightly coupled, we've got some sort of low latency interconnect, parallel file systems, designed to run high performance, highly scalable custom applications. But today, this has changed. HPC has come to mean pretty much any form of scientific computing and as a result, its breadth has grown in terms of what kind of applications we need to support."
– Gregory Kurtzer, founder, Singularity (HPC container software)

THE ANATOMY OF A HIGH PERFORMANCE COMPUTER

The elements find on desktop — processors, memory, disk, operating system — just more of them. High performance computers of interest to small and medium-sized businesses today are really clusters of computers. Each individual computer in a commonly configured small cluster has between one and four processors, and today's processors typically have from two to four cores. HPC people often refer to the individual computers in a cluster as nodes. A cluster of interest to a small business could have as few as four nodes, or 16 cores. A common cluster size in many businesses is between 16 and 64 nodes, or from 64 to 256 cores.

The point of having a high performance computer is so that the individual nodes can work together to solve a problem larger than any one computer can easily solve. And, just like people, the nodes need to be able to talk to one another in order to work meaningfully together. Of course computers talk to each other over networks, and there are a variety of computer network (or interconnect) options available for business cluster.

Software makes the cluster go 'round

Just like desktop or laptop, HPC cluster won't run without software. Two of the most popular choices in HPC are Linux (in all the many varieties) and Windows. Linux currently dominates HPC installations, but this in part due to HPC's legacy in supercomputing, large scale machines, and Unix. Choice of operating system should really be driven by the kinds of applications you need to run on your high performance computer. If using Excel to run option calculations in parallel, want a Windows-based cluster, and so on. In fact the first thing to know when considering buying or building an HPC cluster is what you want to do with it. This sounds obvious, but having a clear sense of what your cluster will, and won't, be required to run will make sure that the decisions you make later are ones you won't regret.

BY

KARTHIKEYAN.G IV Year/IT

Computing is evolving beyond phones, and people are using it in context across many scenarios, be it in their television, be it in their car, be it something they wear on their wrist or even something much more immersive.

-Sundar Pichai

FOUR DIFFERENT MODES OF HIGH PERFORMANCE COMPUTER

1.Dedicated supercomputer (specialised, non-commodity components).

A supercomputer is a computer with a high level of performance compared to a general-purpose computer. Performance of a supercomputer is measured in floating-point operations per second (FLOPS) instead of million instructions per second (MIPS). There are supercomputers which can perform up to nearly a hundred quadrillion FLOPS.All of the world's fastest 500 supercomputers run Linux-based operating systems.

Additional research is being conducted in China, the United States, the European Union, Taiwan and Japan to build even faster, more powerful and more technologically superior exascale supercomputers.

Supercomputers play an important role in the field of computational science, and are used for a wide range of computationally intensive tasks in various fields, including quantum mechanics, weather forecasting, climate research, oil and gas exploration, molecular modeling (computing the structures and properties of chemical compounds, biological macromolecules, polymers, and crystals), and physical simulations (such as simulations of the early moments of the universe, airplane and spacecraft aerodynamics, the detonation of nuclear weapons, and nuclear fusion). Throughout their history, they have been essential in the field of cryptanalysis.

2.Commodity cluster (standard servers with high-speed interconnects).

A **computer cluster** is a set of loosely or tightly connected computers that work together so that, in many respects, they can be viewed as a single system. Unlike grid computers, computer clusters have each node set to perform the same task, controlled and scheduled by software.

The components of a cluster are usually connected to each other through fast local area networks, with each *node* (computer used as a server) running its own instance of an operating system. In most circumstances, all of the nodes use the same hardware [*better source needed*] and the same operating system, although in some setups (e.g. using Open Source Cluster Application Resources (OSCAR)), different operating systems can be used on each computer, or different hardware.

Clusters are usually deployed to improve performance and availability over that of a single computer, while typically being much more cost-effective than single computers of comparable speed or availability.

Computer clusters emerged as a result of convergence of a number of computing trends including the availability of low-cost microprocessors, high-speed networks, and software for high-performance distributed computing.[citation needed] They have a wide range of applicability and deployment, ranging from small business clusters with a handful of nodes to some of the fastest supercomputers in the world such as IBM's Sequoia.

3. HPC cloud computing (compute cycles-as-a-service over the internet).

High-Performance Computing (HPC) in the cloud has reached the mainstream and is currently a hot topic in the research community and the industry. The attractiveness of cloud for HPC is the capability to run large applications on powerful, scalable hardware without needing to actually own or maintain this hardware.

the high computing can be achieved by using cloud or it can be achieve by using a super computer that can cater your application's computing needs.

HPC only indicates the range of computation power required. Whereas cloud computing represents the way of achieving the computational strength.

For example if my program requires 5tera flops of cpu then, i have option to use a single server that can give me 5tflops or i can use a group of five servers having 1tflops that is equivalent to 5tflops anyway.
There are certain domains that requires the very high computation power at certain times, for example if there is a sale going on certain website then, it may attract lakhs of customers and thereby load on server. To cater the need of such requirements ‘cloud bursting’ is used. As soon load on server hit the threshold value another server will launch and that requests can also be handle by that new server.

4.Grid computing

Grid computing is the collection of computer resources from multiple places to reach a common goal. The **grid** can be thought of as a distributed system with non-interactive workloads that involve a large number of files. Grid computing is distinguished from conventional high-performance computing systems such as cluster computing in that grid computers have each node set to perform a different task/application. Grid computers also tend to be more heterogeneous and geographically dispersed (thus not physically coupled) than cluster computers. Although a single grid can be dedicated to a particular application, commonly a grid is used for a variety of purposes. Grids are often constructed with general-purpose grid middleware software libraries.

Grid sizes can be quite large''.

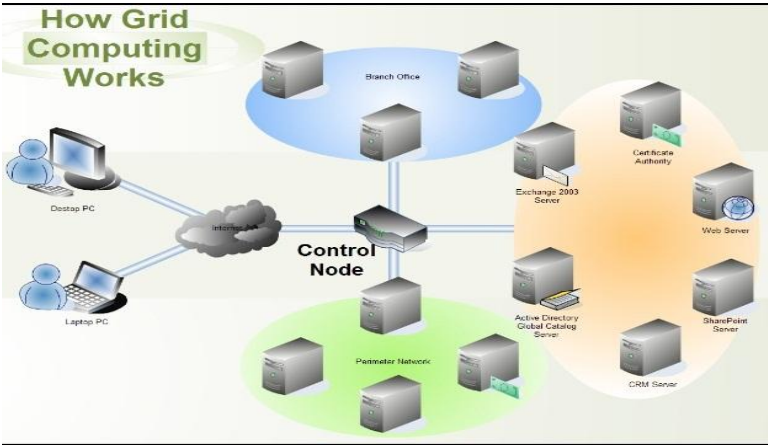

Grids are a form of distributed computing whereby a "**super virtual computer**" is composed of many networked loosely coupled computers acting together to perform large tasks. For certain applications, distributed or grid computing can be seen as a special type of parallel computing that relies on complete computers connected to a computer network (private or public) by a conventional network interface, such as Ethernet. This is in contrast to the traditional notion of a supercomputer, which has many processors connected by a local high-speed computer bus.

BY

GANESAN M II Year/IT

NEED FOR HIGH PERFORMANCE COMPUTING

Case1: Complete a time-consuming operation in less time

I am an automotive engineer – I need to design a new car that consumes less gasoline – I'd rather have the design completed in 6 months than in 2 years – I want to test my design using computer simulations rather than building very expensive prototypes and crashing them

Case 2: Complete an operation under a tight deadline

I work for a weather prediction agency – I am getting input from weather stations/sensors – I'd like to predict tomorrow's forecast today

Case 3: Perform a high number of operations per seconds

I am an engineer at Amazon.com – My Web server gets 1,000 hits per seconds – I'd like my web server and databases to handle 1,000 transactions per seconds so that customers do not experience bad delays

BY

HEMAPRIYA M II Year/IT

WHAT DOES HIGH PERFORMANCE COMPUTING INCLUDE?

• High-performance computing is fast computing

– Computations in parallel over lots of compute elements (CPU, GPU)

– Very fast network to connect between the compute elements

• Hardware – Computer Architecture

• Vector Computers, MPP, SMP, Distributed Systems, Clusters – Network Connections

• InfiniBand, Ethernet, Proprietary (Myrinet, Quadrics, Cray-SeaStar etc.)

• Software – Programming models

• MPI (Message Passing Interface), SHMEM (Shared Memory), PGAS, etc. – Applications

• Open source, commercial

BY

DAVID SHIBIN IV YEAR/IT

Computing is not about computers any more. It is about living.

-Nicholas Negroponte

RISE AND FALL OF HPC COMPUTER ARCHITECTURES

• Vector Computers (VC) - proprietary system

– Provided the breakthrough needed for the emergence of computational science, but they were only a partial answer

• Massively Parallel Processors (MPP)

- proprietary systems – High cost and a low performance/price ratio.

• Symmetric Multiprocessors (SMP)

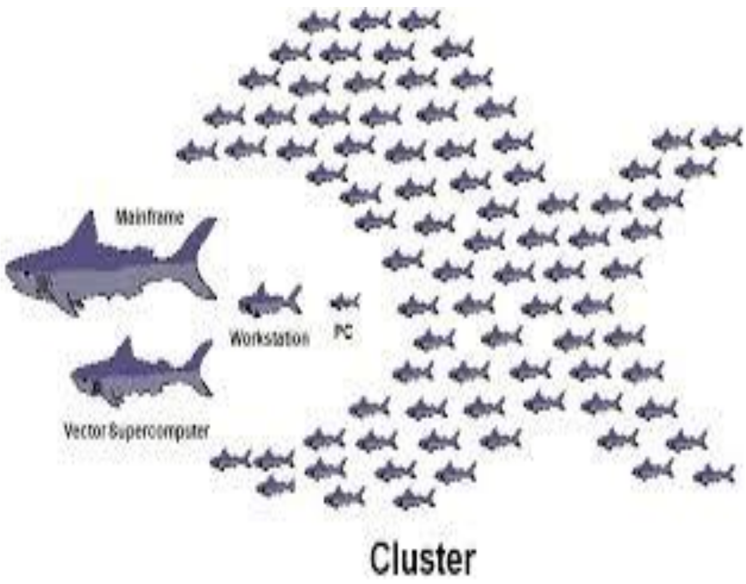

– Suffers from scalability

• Distributed Systems – Difficult to use and hard to extract parallel performance

• Clusters – commodity and highly popular – High Performance Computing - Commodity Supercomputing – High Availability Computing - Mission Critical Applications

HPC CLUSTERS

– Affordable, Efficient and Scalable HPC Solution

Since the 1990s, there has been an increasing trend to move away from expensive /specialized proprietary parallel supercomputers to clusters of computers – Commoditization/standardization are the clustering and interconnect driving forces.

InfiniBand and Ethernet are the most used interconnect solutions for HPC systems Cluster Interconnect.

The HPC Advantage: Reduction in Time to Market Source: IDC From concept to engineering, from design to test and manufacturing, from weather prediction to medical discoveries, our day to day life depends more and more on HPC simulations – Safer products, accurate predictions, research, etc. High-performance compute clusters provide the most efficient, flexible, cost effective HPC environment for any HPC simulation

BY

SRIVIDHYA N,III YEAR/IT

HPC UTILIZATION AREAS

What's next?

The rise in cloud computing and multi-core processors has made HPC accessible to everyone.

So, what lies ahead?

Affordable HPC will open more avenues for enterprises and allow them to make use of cutting-edge technologies. Here are the three main areas HPC is helping to improve:

- The Internet of Things (IoT)
- Big data & High Performance Data Analytics (HPDA)
- Machine learning & AI

With these new approaches to HPC, competitive businesses can produce scalable, real-time models at minimal cost. Here's what the future has in store.

A sixth sense: the Internet of Things

IoT isn't a new concept, but with the democratisation of HPC, it's set to take its place at the centre of the business stage.

By 2020, Gartner predicts there will be 20.8 billion smart devices in operation around the globe. To cope with the enormous influx of data from these sensors and appliances, ambitious companies are looking to more robust forms of CPUs.

GPUs

Leading the way is the Graphical Processing Unit (GPU). Although these processors were traditionally designed for use in 3D modelling, they have been adapted to perform other mathematically-intensive tasks, due to their superior number of cores.

This makes them perfect for use in IoT workloads, where vast streams of structured and unstructured data are processed. Since they require less energy to run than the equivalent number of CPUs, they help reduce operational costs.

Industry Insights: Distribution and Customer Spending

HPC is already a big success in logistics, where sensors gather information about distribution performance. The industry demands a high quality of service, with time and cost the most important variables.

Through the use of parallel processing (the ability to distribute a workload across multiple cores), logistics companies are able to collate data such as vehicle faults and weather conditions to reduce idle time.

Similarly, HPC and IoT can be implemented in the banking sector to monitor customer spending habits. Online banking applications collect data over the internet and enable firms to store information on every aspect of a purchase, helping them understand trends in spending.

Companies need to process this data in real-time to provide actionable insight and, without HPC, workload bottlenecks can occur due to overloaded servers.

Analyse this: Big Data & HPDA

With so many connected devices, companies need a way to process petabytes of raw data. HPC is enabling early adopters to cut through the noise Big Data creates and identify those all-important trends and patterns. IDC predicts the global volume of data will hit 44 zettabytes

by 2020, highlighting the need for powerful data analytics engines.

HPDA and the Cloud

High-Performance Data Analytics (HPDA) will be central to querying and classifying new forms of data. Companies can analyse data from IoT devices, CRM systems, stock markets and more with increased speed, performance and scalability.

To achieve this, HPDA uses grid computing to fuel parallel-processing and distributed data storage.

As more and more businesses adopt the cloud, grid computing is becoming virtual, allowing enterprises to run HPDA on-demand. Opting for a cycle-as-a-service platform means you only pay for the computing power you’re using at any given time.

Not only does this reduce CAPEX, it also enables you to scale instantly by building or tearing down virtual infrastructure. Sharing resources in this way makes it easier to store and analyse vast quantities of data that would otherwise require the configuration of expensive hardware.

Industry Insights: Fraud Detection and Personalised Healthcare

While HPDA is still in its infancy, it’s already making waves in the financial services industry. With tighter regulations placed on banks and brokerage firms, finding ways to detect and predict fraud is a crucial part of staying compliant.

Companies such as PayPal are using HPDA to analyse millions of transactions, considering variables such as location and number of transactions in a given timeframe to determine whether they're fraudulent.

Another sector prospering from the use of HPDA is healthcare. Identifying specific compounds in streams of potential variations makes drug discovery simpler.

The combination of Big Data and HPC will allow pharmaceutical companies to administer a more personalised service, using simulated patients as a test-run for new forms of medicine.

Robo-FLOP: Machine Learning & AI

Of all the potential uses of HPC, machine learning and AI are the most fascinating. Cheaper, more robust processing chips have brought software-defined learning to the masses and turned it into an important tool for modern businesses.

In 2015, NVIDIA CEO, Jen Hsun Huang claimed ‘machine learning is high-performance computing's first killer app for consumers.’ And, two years on, he's been proved right.

Machine learning can help both your business and your customers improve the efficiency and performance of their daily operations. Again, GPUs are providing the power to run the complex algorithms needed to deal with the high-value FLOP (floating point operations per second) calculations that underpin intelligent neural networks.

Staying Hyper-Connected

It's the high-speed interconnects between HPC nodes that enable machine learning to function.

Real-time analysis and understanding rely on faster throughput and data sharing capabilities. The best HPC clusters use networks such as 10 Gigabit Ethernet and InfiniBand to increase computation traffic between nodes.

With multiple cores working on the same calculation, it's important that files and data are shared throughout the network as quickly and efficiently as possible.

Industry Insights: Stock Valuation and Mis-sold Insurance

92 percent of businesses believe machine learning will impact them in some capacity. For example, in the stock exchange, traders can use machine learning to forecast trends in stock prices.

HPC improves the results of these predictions by ensuring the calculations are run in near real-time.Insurance will also benefit from deep learning algorithms. Regulatory compliance is at the forefront of industry processes.

As machine learning improves, so will a firm's ability to understand interactions between sales agents and customers. This will give them more control over mis-selling and ensure negligent behaviour doesn't undermine financial success.

Making HPC Work for Your Business

The beauty of HPC is its ability to adapt to a range of different industries and workloads. However, choosing the right datacentre architecture for your business enables you to get more from your HPC cluster.Software-defined infrastructure has made it possible to tailor the number and function of cores available for certain tasks, increasing speed and performance, optimising network connectivity, and reducing the risk of potential bottlenecks.

When investing in HPC technology, a roadmap is key:

- Shop around for the commodity components that best meet specific requirements. While some will be incompatible, choosing resources from a variety of vendors will help you design bespoke infrastructure that maximises ROI.
- Benchmark your applications. Test out actual HPC clusters before you commit to buying. Proof of concept (POC) is the only way to know which configurations work best for your business.
- Work with a partner, not a wholesaler. Purchasing HPC components can be daunting, so pick a company who will work with you to find the perfect software and hardware solutions for your needs.

BY

VIVEK S IV Year/IT

TOP 5 MANUFACTURING APPLICATIONS HPC

The manufacturing segment is one of the largest markets for high performance computing, globally. In fact, the large product manufacturing sub-segment is the biggest vertical in commercial HPC. All leading automotive, aerospace, and heavy equipment manufacturers have employed HPC for decades, using the technology to design and test their products.

The other principle manufacturing sub-segment, consumer manufacturing, is smaller, but represents one of the fastest growing verticals for HPC usage. It employs digital simulations to help design and test a wide array of consumer products such as smartphones, electric drills, laundry detergent, and food packaging. HPC is particularly useful in the fast-moving consumer product space because it enables engineers to iterate quickly through design variations as new products are developed.

Using digital simulations allows manufacturers to reduce costs by replacing costly development of physical models with virtual ones during various stages of the product development workflow. Potential benefits include improved product quality, shorter time to market, and reduced manufacturing costs.

Software application areas encompass two broad categories: structural analysis and fluid dynamics analysis. Fluid dynamics analysis includes applications used to perform computational fluid dynamics (CFD), while structural analysis encompasses applications for analyzing structures, including explicit and implicit finite element analysis (FEA).

In its global survey of HPC manufacturers, Intersect360 Research has found a large number of applications being used across the industry, but a handful (less than 10) are used more frequently. Here are top 5 HPC application packages used by manufacturers:

1. LS-DYNA

Developed by Livermore Software Technology Corporation (LSTC), LS-DYNA is an expansive finite element analysis (FEA) suite used by engineers for automotive, aerospace, construction, military, product manufacturing, and bioengineering applications. Typical use cases include performing virtual crash testing for automobiles, analyzing the deformation profile of plastics and metals, simulating wear and tear of moving parts, and analyzing the explosive characteristics of various weapons. It was used by NASA to simulate the airbags used in the landing of the Mars Pathfinder space probe.

LS-DYNA originated at Lawrence Livermore National Laboratory (LLNL) in 1976 as a basic FEA program. With the founding of LSTC in 1988 and the subsequent expansion of its capabilities, the software

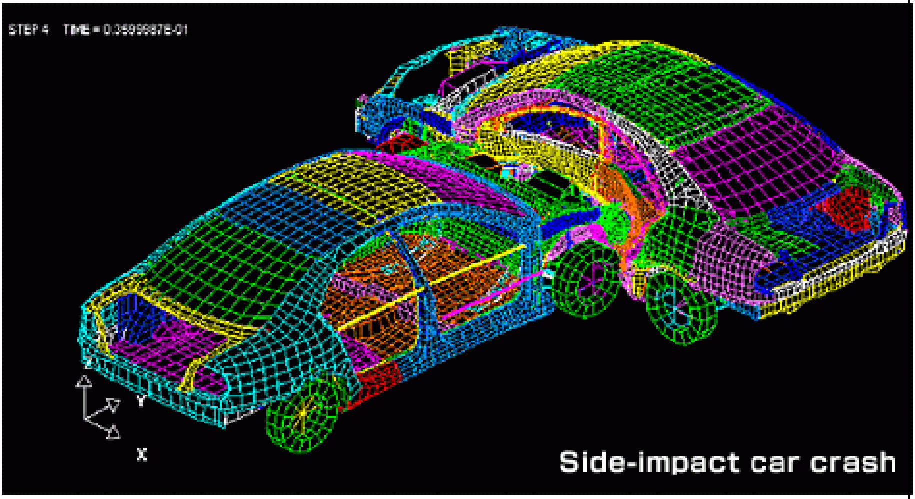

became widely used by commercial manufacturers.

2. Abaqus.

Like LS-DYNA, SIMULIA Abaqus is an FEA suite that has been enhanced to offer full-featured support for engineering simulations and visualizations. The suite encompasses a range of interoperable products, including Standard (general-purpose FEA), CAE (computer-aided engineering), Explicit (advanced FEA), CFD (computational fluid dynamics) and Multiphysics (multiphysics interactions). It also offers an Add-ons product that allows users to integrate applications from third-party tools.

The wide range of capabilities allow Abaqus to support simulations of complex assemblies and heterogeneous materials. It is primary employed by automotive, aerospace, and other industrial manufacturers, and is typically used for designing and testing engines and machinery. The suite is also used for drop testing, crash testing, ballistics, thermal stress analysis, and to assess earthquake resistance of large structures.s LS-DYNA is owned and distributed by Dassault Systèmes

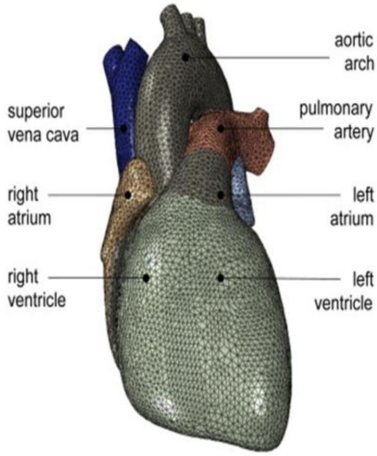

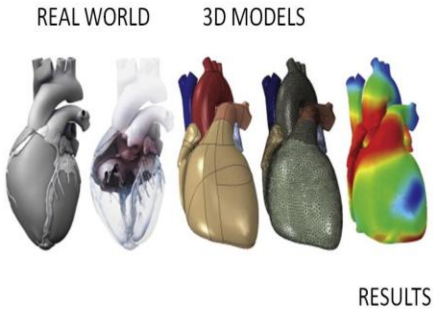

3. ANSYS Fluent.

ANSYS Fluent is one of the most comprehensive computational fluid dynamics (CFD) software packages available today. It allows engineers to model flow, turbulence, heat transfer, and other type of fluid-based reactions for industrial applications. Fluent offers advanced capabilities, such as simulating laminar, turbulent and multiphase flows, as well as interactions with chemical reactions, radiation and particulate dynamics.

Typical manufacturing use cases include simulating airflow over an aircraft wing, performing blood flow analysis for medical devices, studying airflow in heating and cooling systems, and modeling combustion behavior in engines and furnaces. Fluent was used to optimize the aerodynamics of Red Bull’s Formula 1 racecars and to help reduce drag on Speedo’s Olympics-class LZR Racer swimming suit.

4. STAR-CCM+.

CD-adapco. A relatively new offering (2004) from software maker CD-adapco, STAR-CCM+ is a popular CFD package, and like competing products is used to model gas and fluid flow across a wide range of engineering applications. Different solvers are available and are used in different environments, depending on the compressibility of the flow and timescale of the simulation. STAR-CCM+ also supports comprehensive physics modeling for applications that need to incorporate electromagnetics, aeroacoustics, heat transfer, and mechanical stress, to name a few.

The software runs in a client/server model so that engineers can run their applications remotely from their desktop systems on a high performance computing cluster. It provides a modern interface aimed at maximizing ease of use. STAR-CCM+ has been used by Renault to maximize performance on its Formula 1 cars, and to optimize the electrocoat process used for corrosion protection in automobile manufacturing.

5. OpenFOAM.

OpenFOAM is an open source CFD software suite that has been gaining in popularity in the automotive, aerospace, energy, and industrial sectors. Its growing user base of engineers and researchers is a result of its free nature and its easy expandability. In particular, users can create their own customized solvers and interfaces to third party modeling tools. Like other comprehensive CFD suites, it offers extensive support for fluid flow, meshing, turbulence, and heat transfer, while also providing auxiliary functionality in other areas such as electromagnetics, molecular dynamics, combustion and solid dynamics.

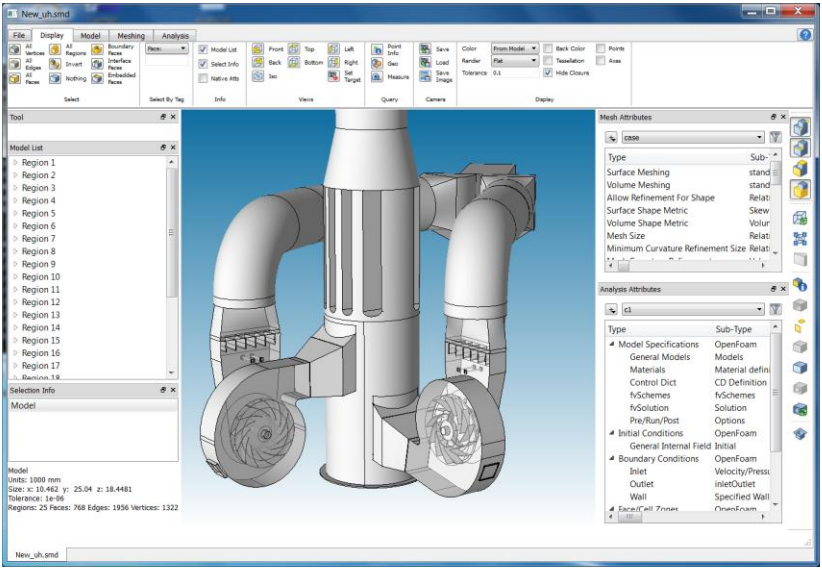

openFOAM code is distributed by the OpenFOAM Foundation, which claims that one million copies of the software have been downloaded since 2005. Support and training is available from ESI Group, which acquired the OpenFOAM trademark when it bought OpenCFD Ltd from SGI in 2012.

BY

OYILA K III Year/IT

CURRENT & FUTURE APPLICATIONS OF HPC

Modelling & Simulation

Prediction of technology breakthroughs

Modelling specific species against climate change

Dynamic longevity prediction

Volcano modelling

Space weather forecasting

Modelling impacts of bio-diversity loss

Power grid simulation & testing

Modelling of organizational behavior

Healthcare & Medicine

Dynamic real-time individual longevity forecasts

Mapping blood flow

Prediction of strokes, brain injury & vascular brain disease

Pandemic modelling

Unravelling protein folding

Curing Alzheimer's disease

Virtual neural circuits

Fintech

Pre-trade risk analysis

Bond pricing

Real-time hedging

Fraud detection

Automated due diligence for M&A

Whole economy simulation

Software & data

Software that writes itself

Holographic data storage

Coding for ultra-low energy use

Data that generates its own models

Engineering, materials & manufacturing

Space station design

Space colony design

Design of new aeronautics materials

Zero gravity manufacturing & design

Security

Recording of every individual human conversation on earth

Modelling of factors likely to lead to a revolution

Deliberate cyber-facilitation of revolutions

Breaking 512-bit encryption ciphers

War forecasting algorhythms

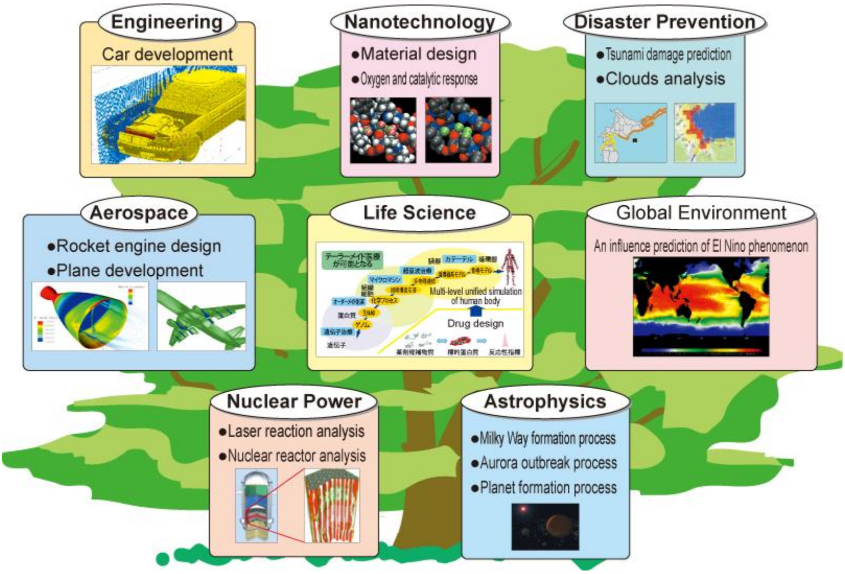

Advantages of HPC

High Performance

HPC systems combines high speed network interconnects, parallel file systems and reliable hardware with clever design to utilize parallel computing on the cluster. This approach enables the system to perform calculations very quickly.

New generation of graphics cards (GPUs, GPGPUs and Co-Processors) pushes HPC system's computing capacity even further. For the supported applications, a GPGPU card can speed up the calculations significantly.

Scalability

Parallel nature of the HPC systems enables them to be scaled very easily and effortlessly. A well designed HPC Cluster can be scaled in a matter of minutes, depending on the number of nodes it contains. This easy scalability gives users a chance to start small and built their cluster up when they feel the need.

Stability

Today most the World’s top HPC Clusters are built on top of Linux operating systems. This high level of utilization of Linux provides HPC systems, a very well designed & tested foundation to lay their cluster design on top. This rock solid foundation when combined with reliable hardware, redundant disks, redundant power supplies and more, delivers a very high uptime and a peace of mind!

Disadvantages of HPC

- **Complex Design**:
  Difficult to implement & build.
- **Troublesome Management**:
  Need of experienced IT staff to manage the systems & keep it running.
- **Expensive**:
  Generally requires more initial investment.

BY

SURIYA V , IV YEAR/IT

Program Outcomes (POs)

PO1	**Engineering Knowledge:** Apply the knowledge of mathematics, science, engineering fundamentals and an engineering specialization to the IT enabled solution of complex engineering problems.
PO2	**Problem Analysis:** Identify, analyze and provide solutions to the problems reaching substantiated IT enabled conclusions.
PO3	**Design/Development of Solutions:** Design solutions for complex engineering problems and design system components or processes that meet the desired needs within realistic constraints.
PO4	**Conduct Investigations of complex problems:** Use research-based knowledge and research methods including design of experiments, analysis and interpretation of data, and synthesis of the information to provide valid conclusions.
PO5	**Modern Tool Usage**: Create, select and apply appropriate techniques, resources and modern engineering and IT tools including prediction and modeling to complex engineering activities with an understanding of the limitations.
PO6	**The Engineer and Society:** Apply reasoning informed by the contextual knowledge to assess societal, health, safety, legal and cultural issues and the consequent responsibilities relevant to the professional engineering practice.
PO7	**Environment and Sustainability:** Understand the impact of the professional engineering solutions in societal and environmental contexts, and demonstrate the knowledge of, and need for sustainable development.
PO8	**Ethics:** Apply ethical principles and commit to professional ethics and responsibilities and norms of engineering practice.
PO9	**Individual and Team Work:** Function effectively as an individual, and as a member or leader in diverse teams, and in multidisciplinary settings.
PO10	**Communication:** Communicate effectively on engineering activities with the engineering community and with society.
PO11	**Project Management and Finance:** Demonstrate knowledge and understanding of the engineering and management principles and apply these to one's own work, as a member and leader in a team, to manage projects and in multidisciplinary environments.
PO12	**Life Long Learning:** Recognize the need for, and have the preparation and ability to engage in independent and life-long learning in the broadest context of technological change.

Program Specific Outcomes(PSOs)

PSO1	**Programming Skill**	Work as Software Engineers for providing solutions to real world problems using programming languages and open source software.
PSO2	**Web Designing Skill**	Ability to use the web designing skill to establish new solutions for the societal needs.


M/S K.S.R
INSTITUE
FOR
ENGINEERING
AND
TECHODOLOGY
Where future begins.