2015

Overview of Global Supercomputing

Richard S. Segall
Arkansas State University - Jonesboro, rsegall@astate.edu

Neha Gupta
Northeastern University and Osmania University, India

Chapter 1
Overview of Global Supercomputing

Richard S. Segall
Arkansas State University, USA

Neha Gupta
Northeastern University, USA & Osmania University, India

ABSTRACT

In this chapter, a discussion is presented of what a supercomputer really is, as well as of both the top few of the world's fastest supercomputers and the overall top 500 in the world. Discussions are also of cognitive science research using supercomputers for artificial intelligence, architectural classes of supercomputers, and discussion and visualization using tables and graphs of global supercomputing comparisons across different countries. Discussion of supercomputing applications and overview of other book chapters of the entire book are all presented. This chapter serves as an introduction to the entire book and concludes with a summary of the topics of the remaining chapters of this book.

INTRODUCTION

Supercomputers are the fastest computers till date and hence the backbone of Computational Sciences. By processing and generating vast amounts of data with unparalleled speed, they make new developments and research possible. The hardware structure or the architecture of supercomputers determines to a larger extent the efficiency of supercomputing systems. Another important element that is considered is the ability of the compilers to generate efficient code to be executed on a given hardware platform. While the Supercomputers of the 1970s used only a few processors, supercomputers of the 21st century can use over 100,000 processors connected by fast connections.

In 1929, New York World newspaper coins the term "Super Computer" when talking about a giant tabulator custom-built by IBM for Columbia University (Gardner, 2014). In 1966, Seymour Cray developed the world's first "real" supercomputer, the CDC 6600: the first computer specifically designed for science and engineering calculations (Gardner, 2014).

A supercomputer is a computer at the frontlines of current processing capacity and speed of calculations. (Wikipedia, 2014) First introduced in

DOI: 10.4018/978-1-4666-7461-5.ch001

the 1960s, the supercomputers of the 1970s used only few processors, and in the 1990s machines with thousands of processors began to appear.

By the end of the 20th century supercomputers were massively parallel computing systems composed of tens of thousands of processors. In contrast, supercomputers of the 21st century can use over 100,000 processors including those with graphic capabilities. For example, Sequoia, ranked as world's third system in 2103 (Top500, 2013) is a third-generation Blue Gene machine from IBM, and runs on 1.6 million processor cores. It can reach speeds of up to 20 petaflops. A petaflop, equals 1015 operations per second, which means that Sequoia can perform 20 x 1015 operations every second. Sequoia requires 3,000 gallons of water per minute to cool it down. It uses 6 or 7 megawatts on average with peak usage approaching 9 1/2 megawatts. (One megawatt equals 1 million watts), and that's $6 or $7 million a year in power. The 1.6 million cores of supercomputer Sequoia are located on 96 different racks, each of which weigh nearly 5,000 pounds and gives off an average of 100 kilowatts of energy, the amount needed to power about 50 single-family homes (Wagstaff, 2012).

Titan, a Cray XK7 system installed at the U.S. Department of Energy's (DOE) Oak Ridge National Laboratory and previously the No. 1 system, is now ranked No. 2 as of November 2013 (Top500, 2013). Titan achieved 17.59 petaflop/s on the Linpack benchmark using 261,632 of its NVIDIA K20x accelerator cores. Titan is one of the most energy efficient systems on the list, consuming a total of 8.21 MW and delivering 2,143 Mflops/W.

Tianhe-2, a supercomputer developed by China's National University of Defense Technology, is as of November 2013 the world's new No. 1 system (Top500, 2013) with a performance of 33.86 petaflop/s on the Linpack benchmark, according to the 42^{nd} edition of the twice-yearly TOP500 list of the world's most powerful supercomputers. The list was announced November 18, 2013 during the opening session of the 2013 Supercomputing Conference (SC13) in Denver, Colorado USA.

Tianhe-2, or Milky Way-2, was deployed at the National Supercomputer Center in Guangzho, China in 2013. The surprise appearance of Tianhe-2, two years ahead of the expected deployment, marks China's first return to the No. 1 position since November 2010, when Tianhe-1A was the top system. Tianhe-2 has 16,000 nodes, each with two Intel Xeon Ivy Bridge processors and three Xeon Phi processors for a combined total of 3,120,000 computing cores.

The Indian government has stated that it has committed about $940 million to develop what could become the world's fastest supercomputer by 2017, one that would have a performance of1 exaflop, which is about 61 times faster than today's fastest computers (PTI, 2012).

Table 1 lists the top ten supercomputer sites in the world as of November 2013 along with number of cores and other performance measures. The Appendix of this book provides complete information for each of the Top500 supercomputer sites in the world as of November 2013. The authors gratefully acknowledge the permission granted to the co-editors of this book for reprinting this detailed information as one of the Appendices of this book. A more complete discussion of the Tianhe-2 or Milk Way-2, Titan, Sequoia, and other supercomputers are also presented in the following sections of this chapter.

APPLICATIONS OF SUPERCOMPUTERS

Supercomputers are used today for highly-intensive calculation tasks for projects ranging from quantum physics, weather forecasting, molecular modeling, and physical simulations. Supercomputers can be used for simulations of airplanes in wind tunnels, detonations of nuclear weapons, splitting electrons, and helping researchers study how drugs

Table 1. Top 10 Supercomputer Sites in the World as of November 2013 (Source: http://www.top500.org/list/2013/11//#.U3_DUCjRhCg)

Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s	Power (kW)
1	National Super Computer Center in Guangzhou China	**Tianhe-2 (MilkyWay-2)** - TH-IVB-FEP Cluster, Intel Xeon E5-2692 12C 2.200GHz, TH Express-2, Intel Xeon Phi 31S1P NUDT	3,120,000	33,862.7	54,902.4	17,808
2	DOE/SC/Oak Ridge National Laboratory United States	**Titan** - Cray XK7, Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x Cray Inc.	560,640	17,590.0	27,112.5	8,209
3	DOE/NNSA/LLNL United States	**Sequoia** - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom IBM	1,572,864	17,173.2	20,132.7	7,890
4	RIKEN Advanced Institute for Computational Science (AICS) Japan	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect Fujitsu	705,024	10,510.0	11,280.4	12,660
5	DOE/SC/Argonne National Laboratory United States	**Mira** - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM	786,432	8,586.6	10,066.3	3,945
6	Swiss National Supercomputing Centre (CSCS) Switzerland	**Piz Daint** - Cray XC30, Xeon E5-2670 8C 2.600GHz, Aries interconnect, NVIDIA K20x Cray Inc.	115,984	6,271.0	7,788.9	2,325
7	Texas Advanced Computing Center/ Univ. of Texas United States	**Stampede** - PowerEdge C8220, Xeon E5-2680 8C 2.700GHz, Infiniband FDR, Intel Xeon Phi SE10P Dell	462,462	5,168.1	8,520.1	4,510
8	Forschungszentrum Juelich (FZJ) Germany	**JUQUEEN** - BlueGene/Q, Power BQC 16C 1.600GHz, Custom Interconnect IBM	458,752	5,008.9	5,872.0	2,301
9	DOE/NNSA/LLNL United States	**Vulcan** - BlueGene/Q, Power BQC 16C 1.600GHz, Custom Interconnect IBM	393,216	4,293.3	5,033.2	1,972
10	Leibniz Rechenzentrum Germany	**SuperMUC** - iDataPlex DX360M4, Xeon E5-2680 8C 2.70GHz, Infiniband FDR IBM	147, 456	2,897.0	3,185.1	3,423

combat the swine flu virus. Supercomputing can be in the form of grid computing, in which the processing power of a large number of computers is distributed, or in the form of computer clusters, in which a large number of processors are used in close proximity to each other. Jungle computing refers to the use of diverse [vague], distributed and highly non-uniform [vague] high performance computer systems to achieve peak performance (Wikipedia, 2014d).

Application areas for supercomputers also include their use for solving large-scale computational problems in earth and atmospheric sciences, materials sciences and engineering, chemistry, fluid dynamics, physics, design optimization for aerospace, and manufacturing and industrial applications. Chapter by co-editor of this book can be found in Cook (2011) on supercomputers and supercomputing.

Supercomputers have been used in bioinformatics and computational biology, and especially are useable for intensive data sets such as at the genetic level. Previous work of the lead co-editor of this book and of this chapter in this area of supercomputing for bioinformatics and computational biology can be found in publications cited in references of this chapter as (Segall, 2013b), (Segall and Cook, 2014), (Segall and Zhang, 2013) and (Segall et al., 2013, 2011, 2010a, 2010b, 2009), and (Segall, 2013a) that are web links for video posted on the web and YouTube for Invited Plenary Address at 17th World International Institute of Informatics and Systemics (IIS) Conference presented by lead co-editor of this book and co-author of this chapter.

In this Invited Plenary Address (Segall, 2013a) presented discussion of applications of supercomputers and video of the Human Brian Project. The human brain is a complicated machine that neuroscientists continually try to understand, and the Human Brian Project is a new scientific endeavor that hopes to unravel some of these mysteries by creating a highly detailed simulation and working replica of the human brain using a supercomputer. With $1.6 billion in funding and more than 200 researchers, the Human Brain Project is the largest, most ambitious cooperative experiment of its kind. Serious hardware is necessary for a project of this kind — to pack the simulation into a single computer would require a system 1,000 times more powerful than today's supercomputers. The Human Brain Project began in 2012. It will take Europe 10 years to map all of the 100 billion neurons connected by 100,000 billion synapses that make up a human brain (Human Brain Project, 2012) (Wikipedia, 2014c).

COMPUTER ARCHITECTURAL CLASSES

The classification of high performance supercomputers comprises four main architectural classes known as Flynn's taxonomy (Steen, 2005) and (Wikipedia, 2014b).

Single Instruction, Single Data Stream (SISD)

These are the conventional systems that contain one CPU and can accommodate one instruction stream that is executed serially (Steen, 2005). These days, many large mainframes can have more than one CPU but each of these executes instruction streams that are unrelated. Such systems should still be regarded as a couple of SISD machines acting on different data spaces. Examples are HP, IBM, SGI (Steen, 2005).

Single Instruction, Multiple Data Streams (SIMD)

Such systems have a large number of processing units ranging from 1,024 to 16,384 that all may execute the same instruction on different data in lock-step (Steen, 2005). So, a single instruction manipulates many data items in parallel. Examples of SIMD machines in this class are the CPP Gamma II and the Quadrics Apemille which are not marketed currently (Steen, 2005).

Mutiple Instruction, Single Data Stream (MISD)

In these systems, multiple instructions should act on a single stream of data. As yet no practical machine in this class has been constructed nor are such systems easy to conceive (Steen, 2005).

Multiple Instruction, Multiple Data Stream (MIMD)

These execute several instruction streams in parallel on different data. The difference with the multi-processor SISD machines mentioned above lies in the fact that the instructions and data are related because they represent different parts of the same task to be executed (Steen, 2005). So, MIMD machines may run many sub-tasks in parallel in order to shorten the time for the main task to be executed. There is a large variety of MIMD systems. Another important distinction between classes of systems is shown as below (Steen, 2005):

1. Shared–memory systems.
2. Distributed-memory systems.

OVERVIEW OF TIANHE-2 (MILKYWAY-2) SUPERCOMPUTER

This is the first ranked fastest supercomputer in the world as of November 2013.This supercomputer in China has around 100 petaflops system speed. It has a hybrid architecture (Xeon CPU and Xeon Phi). Memory is 1.4 PB in total. Storage system is 256 I/O nodes and 64 storage servers with a total capacity of 12.4 PB. It has 16000 compute nodes in total (Lu, 2014). It has TH Express-2 interconnection network. Cooling type is Close-coupled chilled water cooling and a high cooling capacity of 80kW. Programming languages used for compiling are C/C++/Fortran, OpenMP, OpenMC, MPI/GA, Intel Offload (Lu). The supercomputer also has extended Map Reduce framework on CPU/MIC heterogeneous architecture for big data processing. Automatic fault management is another feature (Lu, 2014) (Wikipedia, 2104g).

TITAN III SUPERCOMPUTER

Titan is ranked second in the world's Top500 list of supercomputers as of November 2013.Titan was developed as the archetype of a class of machines better suited than a linked workstation for interactive visualization in supercomputing applications (Miranker, 1992).

Architecture

The computational engine of Titan III is a 64 bit vector multiprocessor. Each CPU of the multiprocessor consists of a trio of co-executing processors: a general purpose, RISC-based integer unit or IPU; a low-latency, pipelined floating point unit or FPU; and a multi-pipelined vector floating-point unit or VPU (Miranker, 1992). Each titan integer processor attains 20 Mips average performance; each vector floating-point processor a 32 Mflop peak. It is the only commercially available, desktop supercomputer (Miranker, 1992). It consists of 1 to 4 CPUs, a graphics processor, I/O processor and memory.

The principle features are:

1. **Multiprocessing:** It provides scalability, upgradeability and lower entry price for a machine. Algorithmic and compilation techniques to take advantage of such architectures are well established for the class of computations for which Titan was constructed (Miranker, 1992), (Wikipedia, 2014e).
2. **Vector Processing:** This reflects the nature of the applications targeted for the system. These applications are described as "interactive simulations" requiring both numerical computation and visualization of real-world phenomena (Miranker, 1992), (Wikipedia, 2014i).
3. **Special Graphics Engine:** Titan provides a graphics subsystem with a high resolution, 1,280 X 1,024 pixel display and full color (Miranker, 1992).

IBM's SEQUIOA SUPERCOMPUTER

It is the world's third largest supercomputer after China's Tianhe-2 and US's Titan. It has achieved 16.32 petaflops/s on the Linpack benchmark using 1,572,864 cores (Whittaker, 2012). It will be used to carry out simulations in order to prolong the life of nuclear weapons. The supercomputer is 1.55 times faster than the Japanese Fujitsu model. It also has less than half the number of CPUs than IBM's number crunching behemoth. Sequioa is an IBM/BlueGene Q system which was installed at the Department of Energy's National Nuclear Security Administration at the Lawrence Livermore National Laboratory (LLNL) (Whittaker, 2012).

K COMPUTER

K Computer ranks 4th in the world's Top500 list of supercomputers as of November 2013 (Wikipedia, 2014h). Fujitsu has been actively developing and providing advanced supercomputers for over 30 years since its development of the FACOM 230-75 APU, Japan's first supercomputer in 1977 (Miyazaki, 2012). Their technical expertise has been applied to developing a massively parallel computer system – the K Computer which has been ranked as the top performing supercomputer in the world (Miyazaki, 2012).

The K Computer was developed jointly by RIKEN and Fujitsu as a part of the High Performance Computing Infrastructure (HPCI) initiative. One objective of this project was to achieve a computing performance of 10^{16} FLOPS (Miyazaki, 2012).

Architecture

It consists of an 8-core CPU with a peak performance of 128 GFLOPS called the "SPARC64 VIIIfx". Commercially available DDR3-SDRAM-DIMM is used as the main memory (Miyazaki, 2012). An interconnect architecture called "Tofu" in excess of 80000 nodes was used. The Tofu interconnect constitutes a direct interconnection network that provides scalable connections. The K Computer system consists of 864 compute racks in total. Two adjacent compute racks are connected by a Z-axis cable (Miyazaki, 2012).

MIRA SUPERCOMPUTER

One of the fastest supercomputers, ranked 5th in Top500 list as of November, 2013, Mira, 10-petaflops IBM Blue Gene/Q system is capable of 10 quadrillion calculations per second. With this computing power, what Mira can do in one day, it would take an average personal computer 20 years to achieve. Mira is helping researchers to tackle more complex problems and create more robust models of everything from jet engines to the human body (Argonne National Laboratory, 2014).

It consists of 48 racks 786,432 processors and 768 terabytes of memory. Mira is 20 times faster than Intrepid, its IBM Blue Gene/P predecessor. In addition to being one of the fastest computers, Mira is also among the most energy efficient. By fitting more cores onto a single chip, Mira speeds the communication between cores and saves the energy lost when transporting data across long distances (Argonne National Laboratory, 2014).

Mira's water-cooling system uses copper tubes to pipe cold water directly alongside the chips, saving power by eliminating an extra cooling step (Argonne National Laboratory, 2014).

PIZ DAINT SUPERCOMPUTER

This is ranked 6th in the world's Top500 list as of November 2013 (Wikipedia, 2014h). It is developed by Swiss National Supercomputing Centre. This supercomputer is based on Intel Xeon E5 processors, has 28 cabinets and has been upgraded to a hybrid architecture featuring NVIDIA Tesla K20X graphical processing units (Swiss National Supercomputing Centre, 2014). With a total of 5272 hybrid compute node, it is possible for real simulations to sustain petaflops (10^{15} FLOPS) performance (Swiss National Supercomputing Centre, 2014).

STAMPEDE SUPERCOMPUTER

It is ranked 7th fastest in the world's Top500 list as of November 2013. Texas Advanced Computing Center (TACC) operates many of the most powerful and capable high performance computing systems in the world (Barth, 2013). Stampede is one of the world's most comprehensive systems for the open science community as a part of the National Science Foundation's (NSF) XSEDE program. It is a Dell PowerEdge C8220 cluster with Intel Xeon Phi coprocessors. The scale of Stampede delivers opportunities in computational science and technology research, from highly parallel algorithms to high-throughput computing, from scalable visualization to next generation programming languages (Barth, 2013).

SuperMUC PETASCALE SYSTEM

SuperMUC is the 10th fastest supercomputer in the Top500 list as of November 2013. SuperMUC is the new supercomputer at Leibniz-Rechenzentrum (Leibniz Supercomputing Centre, 2014) in Garching near Munich. With more than 155.000 cores and peak performance of 3 Petaflop/s, superMUC is one of the fastest supercomputers in the world (Leibniz Supercomputing Centre, 2014).

System Overview

There are 155,656 processor cores in 9400 compute nodes with > 300 TB RAM. Infiniband FDR10 interconnect is present. 4 PB of NAS-based permanent disk storage. 10 PB of GPFS-based temporary disk storage. > 30 PB of tape archive capacity (Leibniz Supercomputing Centre, 2014). Powerful visualization and highest energy efficiency.

Energy Efficiency

SuperMUC uses a new form of warm water cooling developed by IBM. Active components like processors and memory are directly cooled with water that can have an inlet temperature of up to 40 degrees Celsius (Leibniz Supercomputing Centre, 2014). The "High Temperature Liquid Cooling" together with very innovative software cuts the energy consumption of the system.

Typically, water used has an inlet temperature of approximately 16 degrees Celsius and after leaving the system, an outlet temperature of approximately 20 degrees Celsius (Leibniz Supercomputing Centre, 2014). To make water with 16 degrees Celsius requires complex and energy-hungry cooling equipment. Also, there is hardly any use of the warmed-up water as it is too cold to be used in any process (Leibniz Supercomputing Centre, 2014).

SuperMUC allows an increased inlet temperature. It is easily possible to provide water having 40 degrees Celsius using simple "free-cooling" equipment as outside temperatures in Germany hardly ever exceed 35 degrees. At the same time, the outlet water can be made quite hot (up to 70 degrees C) and re-used in other technical processes.

System Configuration

SuperMUC consists of 18 Thin Node Islands and one Fat Node Island. Each island contains more than 8,192 cores. All compute nodes within an individual island are connected via a fully Infiniband Network (FDR10 for thin nodes/ QDR for fat nodes) (Leibniz Supercomputing Centre, 2014).

System Software

Suse Linux Enterprise Server. For system management, xCat from IBM is used. For Batch processing, Loadleveler from IBM is used.

Permanent storage for data and programs is provided by a 16-node NAS cluster from Netapp. This cluster has a capacity of 2 petabytes. Data is regularly replicated to a separate 4-node Netapp cluster with another 2 PB of storage for recovery purposes (Leibniz Supercomputing Centre, 2014).

PANGEA SUPERCOMPUTER

Pangea supercomputer is ranked number 14 in the top 500 list of world's fastest supercomputers. Total Exploration production inaugurated its new Pangea supercomputer on March 22, 2013 in France (Total inaugurates its new Pangea supercomputer, 2013). The supercomputer represents an investment of 60 million pounds over four years. It will be used as a tool to assist decision-making in the exploration of complex geological areas and to increase the efficiency of hydrocarbon production complying the safety standards (Total inaugurates its new Pangea supercomputer, 2013).

It is designed by SGI (Silicon Graphics International). The Pangea supercomputer has a computing capacity of 2.3 Pflops. Its unique computing architecture is based on over 110,000 calculation cores, 7 Pb^2 storage capacity and an innovative cooling system (Total, 2013). Requiring 2.8 MW of electric power, the heat generated by this supercomputer is recovered making it possible to heat the Scientific and Technical Centre (Total, 2013).

PLEIADES SUPERCOMPUTER

Pleiades, one of the world's most powerful supercomputers ranked 16th fastest, represents NASA's state-of–the-art technology for meeting the agency's supercomputing requirements (Biswas, 2014). This distributed-memory SGI ICE cluster is connected with InfiniBand in a dual-plane hypercube technology. The system contains the following Intel Xeon Processors: E5-2680v2 (Ivy Bridge), E5-2670 (Sandy Bridge), and X5670 (Westmere) (Biswas, 2014).

System Architecture

Manufacturer is SGI. There are 11,176 nodes (163 racks). 3.59 Pflop/s peak cluster, 1.54 Pflop/s LINPACK rating, 2 racks (64 nodes total) enhanced with NVIDIA graphics processing unit (GPU): 43 teraflops total, Total cores: 184,800 (32,768 additional GPU cores), Total memory: 502 TB (Biswas, 2014).

Interconnects

Internode: InfiniBand with all nodes connected in a partial hypercube topology. There is also Gigabit Ethernet management network (Biswas, 2014).

Storage

SGI Infinite Storage NEXIS 9000 home file system. There is 15 PB of RAID disk storage configured over several cluster-wide Lustre file systems. Operating system is SUSE Linux. Job scheduler is PBS. Compilers are Intel and GNU, C, C++, Fortran (Biswas, 2014).

BLUE WATERS SUPERCOMPUTER

Blue Waters is one of the most powerful supercomputers in the world. It can complete more than 1 quadrillion calculations per second on a sustained basis and more than thirteen times that at peak speed (Kramer, 2014). The peak speed is almost 3 million times faster than the average laptop. The machine architecture balances processing speed with data storage, memory and communication within itself and to the outside world in order to cater to a wide variety of scientific endeavers (Kramer, 2014). Many of the projects require a large portion of the thousands of processors that constitute Blue Waters and would be impossible to run elsewhere. Blue Waters is supported by the National Science Foundation and the University of Illinois (Kramer, 2014).

Three programming environments are available on Blue Waters, namely the Cray Programming Environment, the PGI programming environment and the Gnu programming environment. The environment is managed effectively by the module command (Kramer, 2014).

Compiling and Linking

Multiple compiler suites are available for building and linking application codes. Depending on the type of application, one set of compilers may produce code that performs better than the other (Kramer, 2014). The compilers currently available in this system are Cray (Wikipedia, 2014a), The Portland Group (PGI) and GNU (Kramer, 2014).

System Summary

The Blue Waters system is a Cray XE/XK hybrid machine composed of AMD 6276 “Interlagos” processors (nominal clock speed of at least 2.3 GHz) and NVIDIA GK110 “Kepler” accelerators all connected by the Cray Gemini torus interconnect (Kramer, 2014).

Storage

All file systems on Blue Waters are Lustre based. The home directory has 2.2 PB. Scratch has 22 PB for running batch jobs generating large amounts of data. Projects directory has 2.2 PB. This space is used for sharing frequently used large files within a team.

SUPERCOMPUTING WITH PARALLELLA

Packing impressive supercomputing power inside a small credit card-sized board running Ubuntu, Adapteva’s $99 ARM-based Parallella system includes the unique Ephiphany numerical accelerator that promises to unleash industrial strength parallel processing on the desktop at a low price (Lucifredi, 2013).

The board is properly structured like a Supercomputer, with a host side powered by a 667 MHz Zynq 7020 ARM A9 System-on-Chip manufactured by Xilinx (Lucifredi, 2013). This is an interesting chip that includes alongside a dual-core ARM v7 CPU a full-fledged programmable logic facility equivalent to an Artix-7 FPGA. The number-crunching side is powered by a 600 MHz, 16-core Adapteva Epiphany-III numerical accelerator, which is replaced by a 64-core version in more expensive configurations (Lucifredi, 2013).

It is built of Linaro 12.10 Ubuntu derivative and hence is extremely standard. The kernel is 3.6.0 with patches additionally and the system runs the SSH daemon. Apt-get has access to the full Linaro 12.10 repositories, so the choice of software is infinite. The board comes loaded with GCC 4.6, Perl 15.4.2 and Python 2.7. Cross-compile is not necessary (Lucifredi, 2013). Man pages are also installed by default. The interesting part is programming the numerical accelerator, binaries for which are created with e-gcc (Lucifredi, 2013). For instance, comparison of performance of the ARM CPU (2 cores) with that of the Epiphany accelerator (16 cores) can be done. The 16-core matrix outpaces the general purpose CPU by 11X (Lucifredi, 2013).

SYMBOLIC SUPERCOMPUTER FOR ARTIFICIAL INTELLIGENCE AND COGNITIVE SCIENCE RESEARCH

Supercomputers are invaluable resources in many areas of science and engineering. However, artificial intelligence and cognitive science research has not benefited from supercomputing because traditional supercomputers are aimed at numerical simulation (Forbus, n.d.). A symbolic supercomputer is built in order to address this. This cluster machine is carefully designed to support the creation and modeling of intelligent systems (Forbus, n.d.). This supercomputer is created in order to support research being conducted on two ONR projects:

- The Artificial Intelligence Project, Qualitative Reasoning for Intelligent Agents.
- The Cognitive Science Project, Analogical Learning and Case-Based Instruction.

The supercomputer is made of Linux Networx Evolocity cluster with 67 nodes and a gigabit switch. Each node has two 3.2 Ghz Pentium Xeon Processors, with 4 GB RAM and 80 GB of disk (Forbus, n.d.). This machine, known as mk2, became operational in 2004.

COMPARISON OF SUPERCOMPUTING ABILITIES ACROSS DIFFERENT COUNTRIES

Visualization of Global Supercomputing Comparisons by Tables and Graphs

This section provides tables and graphs to illustrate the comparisons of supercomputer abilities across the world as of November 2013. The following sections will provide a more detailed description of the supercomputing in the countries of China, India and Japan and those in Europe of Belgium, Bulgaria, Finland, France, Germany, Italy, Norway, Netherlands, Poland, Russia, Slovenia, Spain, Sweden, Switzerland, and United Kingdom.

Table 2 lists each of the countries in the world that have at least one supercomputer and ranks them accordingly to their system share of the top 500. Performance measures of Rmax and Rpeak are also provided as well as number of cores. Figure 1 provides a frequency curve for the count of the number of supercomputers for each of the countries presented in Table 2. This statistic shows the locations of the world's 500 most powerful supercomputers as of November 2013. As of this date, 264 of the world's leading supercomputers were located in the United States.

Figure 2 provides the teraflops per country and log teraflops per country. Log teraflops per country is provided as additional bar graph because of the magnitude of the numbers a more convenient scale is needed to compare these. As Figure 2 illustrates, the United States is the leader with China, Japan, Germany and France lagging behind.

Figure 3 provides a comparison by introducing the factor of average country income by country of the top 500 supercomputers in the world by comparing the megaflops per capita by country. As Figure 3 illustrates the country of Switzerland outpaces all other countries, including the United States, Japan, and Germany in both megaflops and log megaflops per capita.

Figure 4 illustrates the number of people (in millions) per supercomputer by country as of November 2013 in bar graphs, and shows that the India is the leader with 103.06 million because of its population followed by Brazil with 66.72 million and Russia with 28.71 million. Figure 4 also overlays the bar graph with a connected line graph of the number of supercomputers and shows that United States is leader in number of supercomputers of 264 but the lowest in number of people per supercomputer with 1.19 million.

Table 2. Count by country and performance statistics of the 500 most powerful computers in the world as of November 2103 (Source: Top500, http://www.top500.org/statistics/list/#.U30G4SjRhCh)

Countries	Count	System Share (%)	Rmax (GFlops)	Rpeak (GFlops)	Cores
United States	264	52.8	118,261,596	169,499,661	9,837,537
China	63	12.6	48,549,093	89,432,561	4,925,804
Japan	28	5.6	22,472,218	28,925,862	1,558,880
United Kingdom	23	4.6	9,058,329	11,380,215	627,120
France	22	4.4	9,489,912	11,228,571	720,416
Germany	20	4	13,696,834	16,426,807	1,033,252
India	12	2.4	3,040,297	3,812,719	188,252
Canada	10	2	2,077,842	2,627,756	190,752
Korea, South	5	1	1,258,060	1,760,092	154,224
Sweden	5	1	1,067,767	1,297,036	95,680
Russia	5	1	1,846,613	3,242,736	166,432
Australia	5	1	2,180,151	2,635,546	145,036
Italy	5	1	2,665,609	3,212,697	221,120
Switzerland	5	1	7,765,418	9,632,162	253,904
Netherlands	3	0.6	511,071	671,160	47,544
Brazil	3	0.6	626,000	1,182,104	58,880
Norway	3	0.6	735,400	873,164	54,400
Saudi Arabia	3	0.6	1,165,315	1,827,011	128,272
Ireland	2	0.4	268,565	343,310	30,996
Israel	2	0.4	314,056	736,819	35,424
Finland	2	0.4	378,000	436,301	20,976
Hong Kong	2	0.4	352,937	658,368	47,520
Poland	2	0.4	455,909	583,605	41,852
Spain	2	0.4	1,199,031	1,357,824	65,280
Belgium	1	0.2	152,348	175,718	8,448
Austria	1	0.2	152,900	182,829	20,776
Denmark	1	0.2	162,098	183,676	15,672
Taiwan	1	0.2	177,100	231,859	26,244

Table 3 provides the aggregation of the statistics in Table 2 for each geographical region of North America, Western Europe, Northern Europe, Southern Europe, Eastern Europe, Eastern Asia, South-central Asia, Western Asia, Australia and New Zealand, and South America. Table 3 shows a count of 274 supercomputers for North America, and 99 in Eastern Asia, and 52 in Western Europe to account for 85% of the world's supercomputers.

Table 4 provides a count of how many of the top 500 supercomputers use cluster versus massively parallel processors (MPP) and indicates that 84.6% of the world's supercomputers use a cluster architecture.

Figure 1. Locations of the 500 most powerful supercomputers in the world as of November 2013 by country (Source: http://www.statista.com/statistics/264445/number-of-supercomputers-worldwide-by-country/)

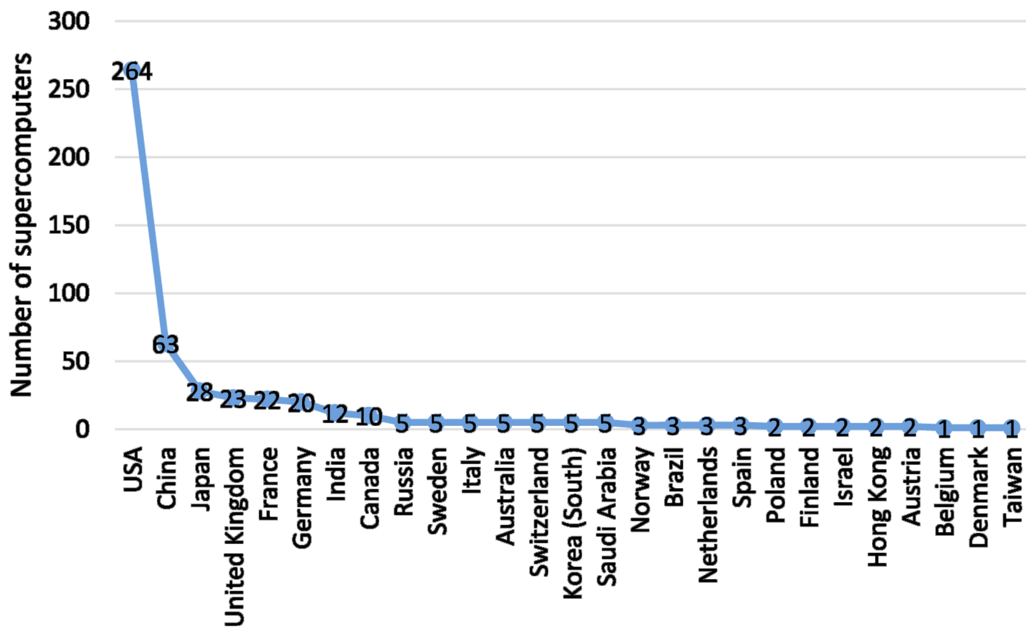

Figure 2. Teraflops and Log Teraflops per country of the world's top 500 supercomputers (Source: top500.org/list/2013/11)

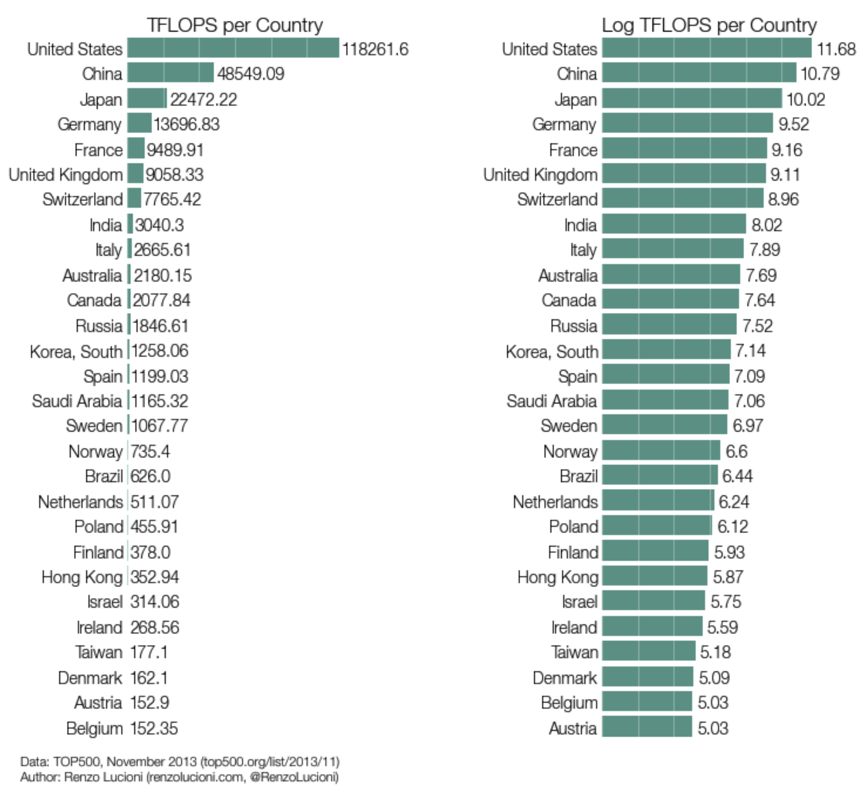

Figure 3. Megaflops and Log Megaflops per capita by country of the world's top 500 supercomputers. (Source: top500.org/list/2013/11)

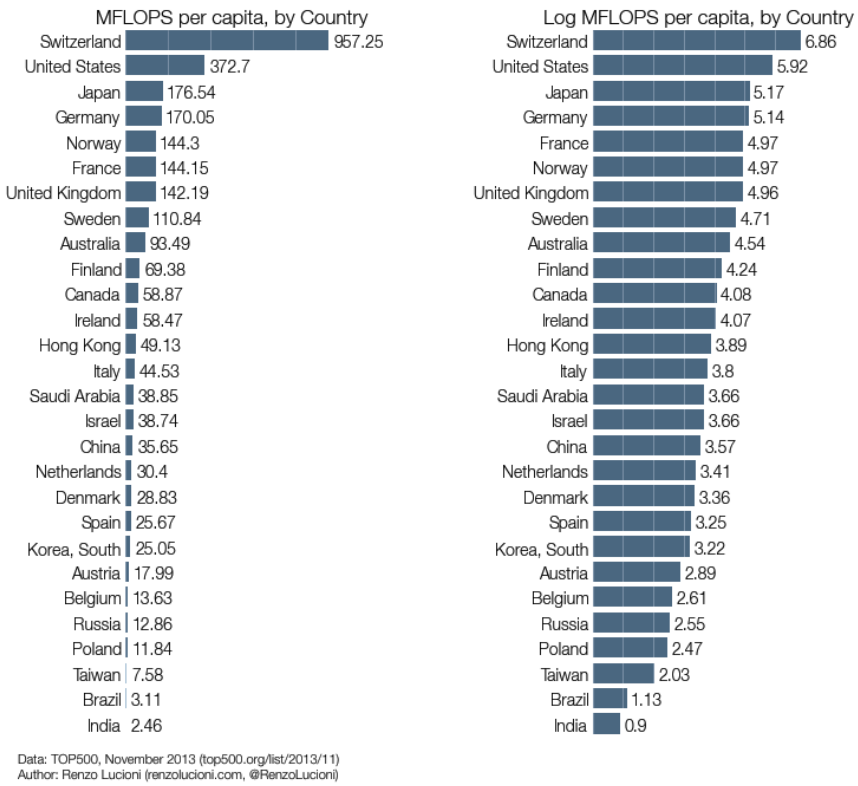

Table 5 shows the counts of operating systems for the top 500 supercomputers in the world.

Table 6 shows the number of supercomputers of the top 500 in the world per processor generation and the system share and other performance measures.

Table 7 shows the counts of the operating system families of Linux, Unix, Mixed and Windows for the Top 500 supercomputers in the world as of November 2013. As Table 6 indicates, 96.4% (482) of the top 500 supercomputers in the world use a Linux Operation System.

Figure 5 shows a bar graph of the aggregated counts for the operating systems of Table 5. The statistics of Figure 5 show a breakdown of the 500 most powerful supercomputers around the world as of November 2013, by operating system family. As of that date the operating system family Linux held a system share of 96.4 percent.

Table 8 shows the numerical counts of the number of supercomputers used for each of the listed application areas as of November 2013. As Table 7 illustrates, 82% (410) of the top 500 supercomputers in the world had a non-specified application area, followed by 11.2% (56) for research, and 1.8% (9) for weather. Figure 6 from Statista shows an aggregated distribution of the 500 most powerful supercomputers as of November 2013 into only 7 categories instead of the 12 categories of Table 7, and thus showing a greater number and percentage (20.6%) of supercomputers being used for research area.

Figure 7 shows the system share of interconnect families used in the 500 most powerful supercom-

Figure 4. Number of people per supercomputer (bar graph) and number of supercomputers by country (line graph) as of November 2013 (Source: http://imgur.com/r/dataisbeautiful/xvsmNbc) also posted at: http://www.reddit.com/r/dataisbeautiful/comments/1tftbr/number_of_people_per_supercomputer_by_country_oc/)

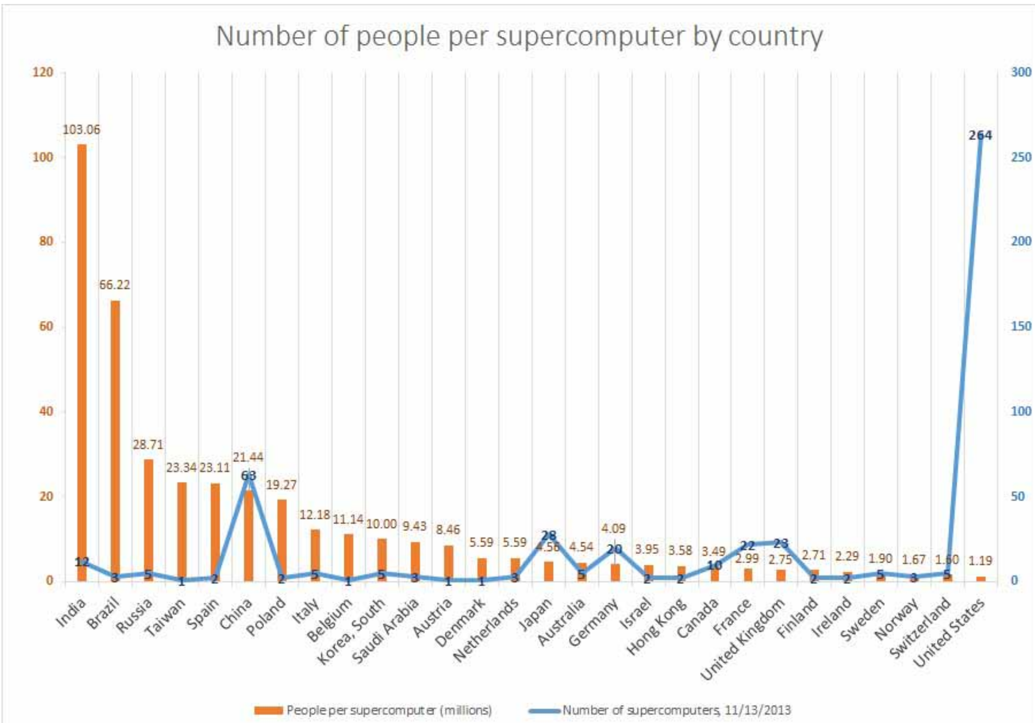

Table 3. Count by geographical region and performance statistics of the 500 most powerful computers in the world as of November 2013 (Source: http://www.top500.org/statistics/list/#.U30G4SjRhCh)

Geographical Region	Count	System Share (%)	Rmax (GFlops)	Rpeak (GFlops)	Cores
North America	274	54.8	120,339,438	172,127,417	10,028,289
Eastern Asia	99	19.8	72,809,408	121,008,742	6,712,672
Western Europe	52	10.4	31,768,483	38,317,247	2,084,340
Northern Europe	36	7.2	11,670,158	14,513,703	844,844
South-central Asia	12	2.4	3,040,297	3,812,719	188,252
Southern Europe	7	1.4	3,864,640	4,570,521	286,400
Eastern Europe	7	1.4	2,302,522	3,826,340	208,284
Western Asia	5	1	1,479,371	2,563,830	163,696
Australia and New Zealand	5	1	2,180,151	2,635,546	145,036
South America	3	0.6	626,000	1,182,104	58,880

Table 4. Architecture classifications of the Top 500 supercomputers in the world as of November 2013 (Source: http://www.top500.org/statistics/list/#.U30G4SjRhCh)

Architecture	Count	System Share (%)	Rmax (GFlops)	Rpeak (GFlops)	Cores
Cluster	423	84.6	156,485,694	244,434,097	13,409,693
MPP	77	15.4	93,594,774	120,124,072	7,311,000

Table 5. Types of operating systems for the Top 500 supercomputer in the world as of November 2013 (Source: http://www.top500.org/statistics/list/#.U30G4SjRhCh)

Operating System	Count	System Share (%)	Rmax (GFlops)	Rpeak (GFlops)	Cores
Linux	414	82.8	158369073	230603624	14277307
Cray Linux Environment	20	4	30911722	43804792	1302984
SUSE Linux Enterprise Server 11	13	2.6	9174795	13081620	432150
CentOS	11	2.2	2685015	3654410	192552
AIX	11	2.2	3496347	4208920	137536
CNK/SLES 9	4	0.8	1184521	1420492	417792
Bullx Linux	4	0.8	1103827	1330204	50960
RHEL 6.2	4	0.8	1738900	2132582	102528
Redhat Enterprise Linux 6	4	0.8	2571639	3388905	321976
bullx SUperCOmputer Suite A.E.2.1	3	0.6	2942070	3583180	165888
Redhat Linux	2	0.4	327834	424760	26636
SLES10 + SGI ProPack 5	2	0.4	398000	439910	38400
Super-UX	1	0.2	122400	131072	1280
Windows Azure	1	0.2	151300	167731	8064
CNL	1	0.2	165600	201216	20960
Windows HPC 2008	1	0.2	180600	233472	30720
Scientific Linux	1	0.2	188725	199680	9600
RHEL 6.1	1	0.2	230600	340915	37056
SUSE Linux	1	0.2	274800	308283	26304
Kylin Linux	1	0.2	33862700	54902400	3,120,0

puters around the world as of November 2013. As of this date Infiniband was the interconnect family used in 41.4 percent of the leading supercomputers.

Figure 8 shows the number of computer cores in the 10 fastest supercomputers in the world (current to November 2009).

1. The “Jaguar” (XT5 HC, Six Core Opteron 2.6 GHz) from Cray is in the Oak Ridge National Laboratory.
2. “Nebulae” (TC3600 Blade, Intel X5650, NVIDIA Tesla C2050 GPU) from Dawning in in the National Supercomputing Centre Shenzhen.

Table 6. Processor generation for the Top 500 supercomputers in the world as of November 2013 (Source: http://www.top500.org/statistics/list/#.U30G4SjRhCh)

Processor Generation	Count	System Share (%)	Rmax (GFlops)	Rpeak (GFlops)	Cores
Intel Xeon E5 (SandyBridge)	307	61.4	87,073,127	136,737,345	6,730,764
Xeon 5600-series (Westmere-EP)	55	11	16,727,544	30,902,871	1,565,654
Intel Xeon E5 (IvyBridge)	34	6.8	45,699,316	70,380,755	3,764,390
Power BQC	24	4.8	46,402,484	54,316,237	4,243,456
Opteron 6100-series “Magny-Cours”	17	3.4	5,295,469	7,163,629	773,640
Opteron 6200 Series “Interlagos”	16	3.2	23,462,905	35,054,646	1,303,280
POWER7	12	2.4	5,011,347	6,153,312	200,896
Xeon 5500-series (Nehalem-EP)	10	2	2,507,531	3,278,734	228,788
Opteron 4100-series “Lisbon”	5	1	691,960	1,057,795	125,928
Xeon 5400-series “Harpertown”	4	0.8	687,264	939,327	81,673
PowerPC 450	4	0.8	1,184,521	1,420,492	417,792
Opteron Quad Core	3	0.6	477,800	596,888	69,904
SPARC64 IXfx	2	0.4	1,209,700	1,317,077	89,088
Opterons 6300 Series (“Abu Dhabi”)	1	0.2	119,300	157,286	16,384
NEC	1	0.2	122,400	131,072	1,280
Xeon 5300-series “Clovertown”	1	0.2	132,800	172,608	14,384
ShenWei	1	0.2	795,900	1,070,160	137,200
Opteron Six Core	1	0.2	919,100	1,173,000	112,800
Xeon 5500-series (Nehalem-EX)	1	0.2	1,050,000	1,254,550	138,368
SPARC64 VIIIfx	1	0.2	10,510,000	11,280,384	705,024

Table 7. Operating system family for the Top 500 supercomputers in the world as of November 2013 (Source: http://www.top500.org/statistics/list/#.U30G4SjRhCh)

Operating System Family	Count	System Share (%)	Rmax (GFlops)	Rpeak (GFlops)	Cores
Linux	482	96.4	244,945,300	358,396,482	20,125,301
Unix	11	2.2	3,496,347	4,208,920	137,536
Mixed	4	0.8	1,184,521	1,420,492	417,792
Windows	2	0.4	331,900	401,203	38,784
BSD Based	1	0.2	122400	131072	1280

Figure 5. Distribution of the 500 most powerful supercomputers worldwide as of November 2013, by operating system family (Source: http://www.statista.com/statistics/249270/distribution-of-leading-supercomputers-worldwide-by-operating-system-family/)

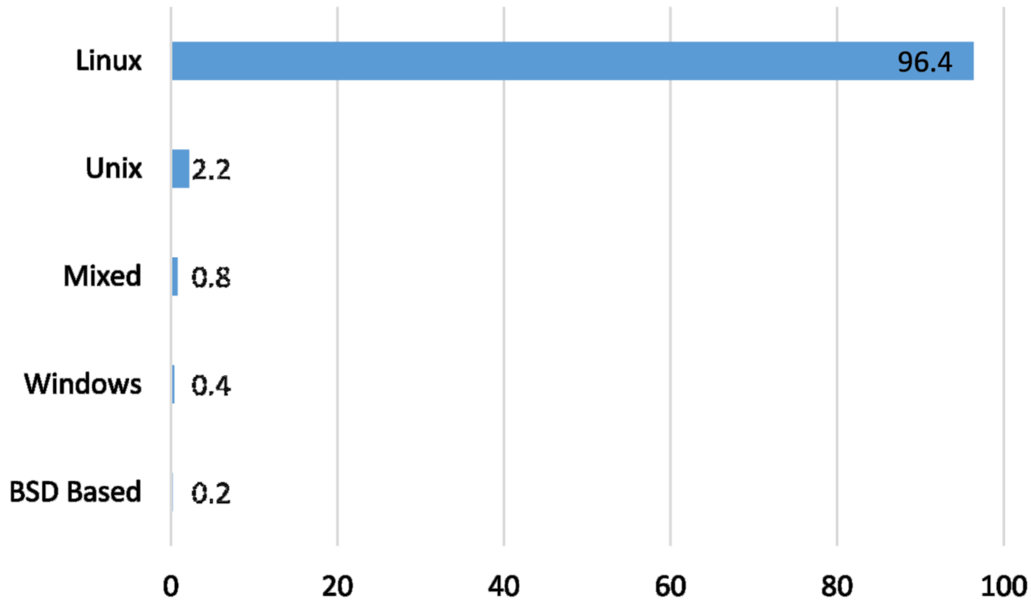

Table 8. Application areas for top 500 supercomputers in the world as of November 2013 (Source: http://www.top500.org/statistics/list/#.U30G4SjRhCh)

Application Area	Count	System Share (%)	Rmax (GFlops)	Rpeak (GFlops)	Cores
Not Specified	410	82	197,710,899	297,627,905	16,061,176
Research	56	11.2	38,576,766	49,804,226	3,589,245
Weather and Climate	9	1.8	3,682,699	4,428,772	295,844
Energy	6	1.2	2,299,298	2,803,536	182,108
Defense	5	1	1,873,434	2,246,243	217,248
Benchmarking	5	1	2,092,440	2,624,667	103,632
Environment	3	0.6	746,307	885,441	43,984
Aerospace	2	0.4	1,785,190	2,403,110	120,384
Web Services	1	0.2	240,090	354,099	17,024
Semiconductor	1	0.2	758,873	933,481	51,392
Software	1	0.2	188,967	209,715	16,384
Finance	1	0.2	125,503	236,974	22,272

Figure 6. Distribution of the 500 most powerful supercomputers as of November 2013, by segment. (Source: http://www.statista.com/statistics/264449/distribution-of-supercomputers-worldwide-by-segment)

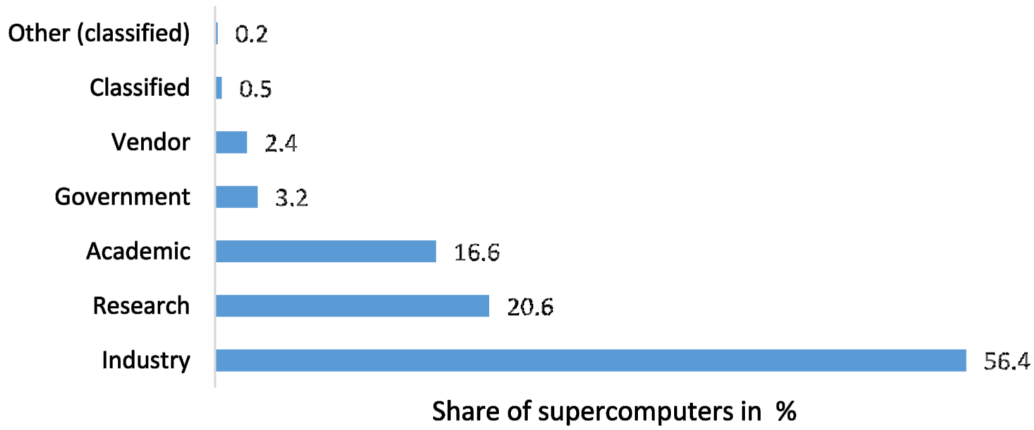

Figure 7. System share of interconnect families used in the most powerful 500 supercomputers worldwide as of November 2013. (Source: http://www.statista.com/statistics/264446/distribution-of-interconnect-families-used-in-supercomputers/)

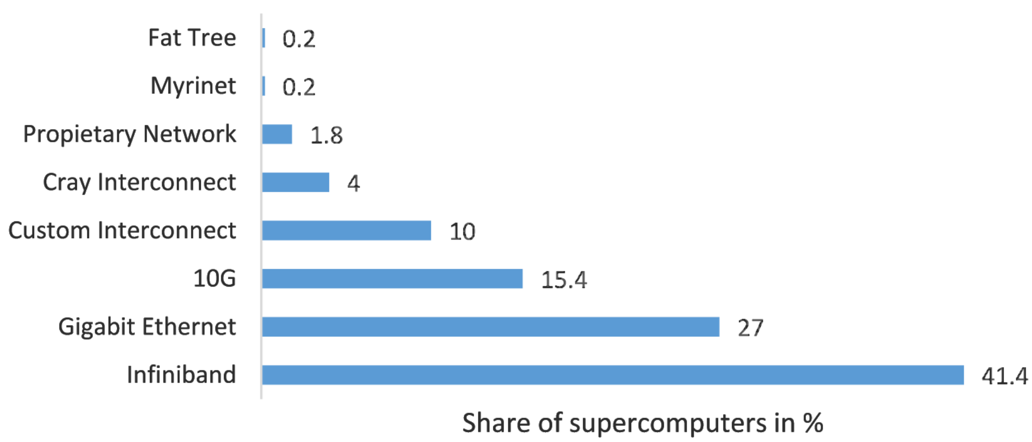

3. The "Roadrunner" (BladeCenter - QS22/ LS21 cluster PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz) from IBM is in the Los Alamos National Laboratory (LANL) in New Mexico.
4. The "Kraken" (Cray XT5 HC, Six Core Xeon 2.36 GHz) from Cray is in the National Institute for Computational Sciences (NICS) at University of Tennessee.
5. "Jugene" (Blue Gene / P Solution) from IBM is at the Research Center Juelich (FZJ).
6. The "Pleiade" (Altix ICE 8200EX/8400EX, Xeon QC HT 3.0/Xeon Westmere 2.93 GHz, Infiniband) from SGI is stored at the National

Figure 8. Number of computer cores in the 10 fastest supercomputers in the world (current to November 2009) (Source: http://www.statista.com/statistics/268280/number-of-computer-cores-in-selected-supercomputers-worldwide/)

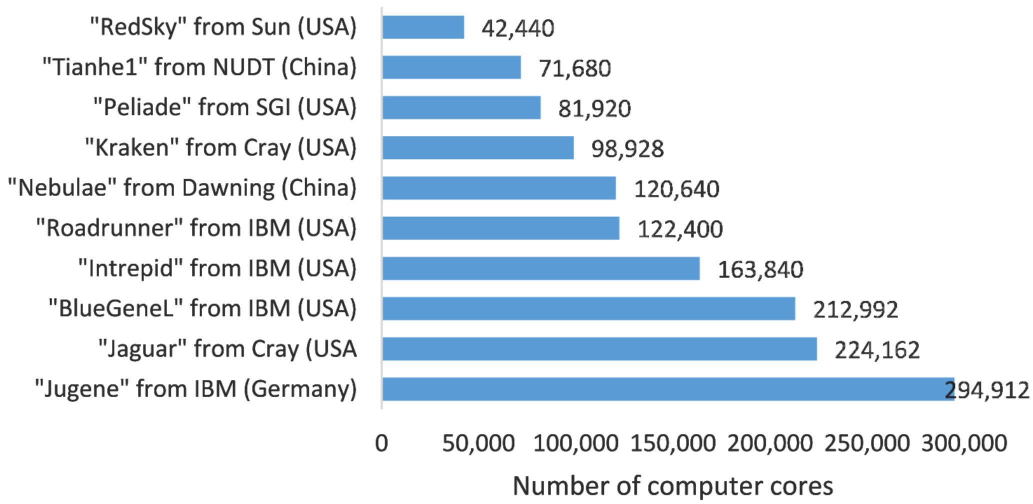

Figure 9. Share of the 500 most powerful supercomputers worldwide as of November 2013, by vendor. (Source: http://www.statista.com/statistics/249268/share-of-leading-supercomputers-worldwide-by-vendor/)

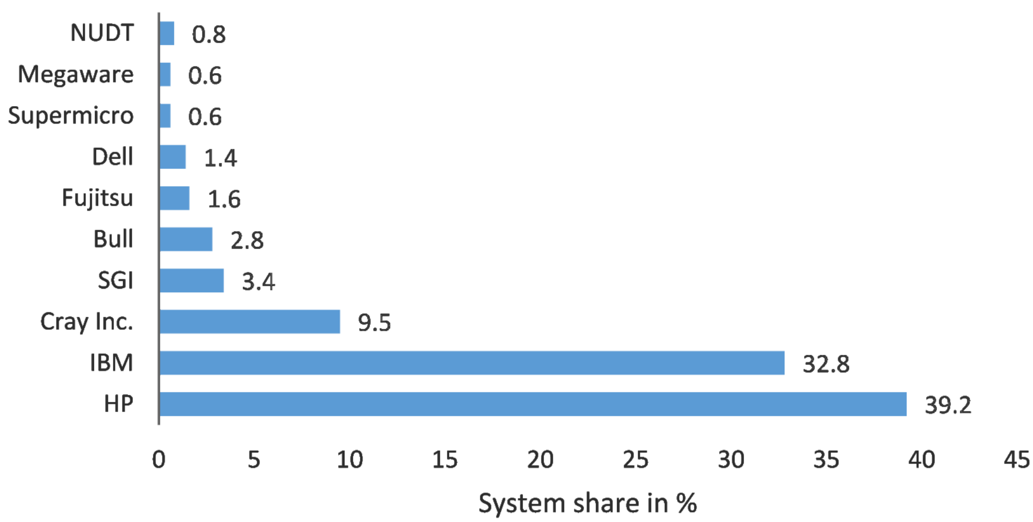

Aeronautic and Space Administration (NASA)/Ames Research Center.

7. The supercomputer from NUDT "Tianhe1" (TH-1-Cluster, Xeon E5540 / E 5450 ATI Radeon HD 4870 2, Infiniband is) is in the National Supercomputer Center in Tianjin.
8. IBM's "BlueGene/L" (eServer Blue Gene Solution) is in the U.S. Department of Energy (DOE) Lawrence Livermore National Laboratory (LLNL).
9. The "Intrepid" Blue (Gene / P Solution) from IBM is in the Argonne National Laboratory.
10. Sun's "RedSky" (SunBlade X6275) is located in Sandia / NRLE.

Figure 9 shows share of the 500 most powerful supercomputers worldwide as of November 2013, by vendor. Figure 9 shows a breakdown of the 500 most powerful supercomputers around the world as of November 2013, by vendor (top 10 vendors only). As of that time IBM had a system share of 32.8 percent of the top 500 supercomputers worldwide.

SUPERCOMPUTING IN CHINA

China has a large number of supercomputing centers which have held world records in speed. The origins of these centers go back to 1989 when the State Planning Commission, the State Science and Technology Commission and the World Bank jointly launched a project to develop networking and supercomputer facilities in China (Supercomputing in China, 2014).

The progress of supercomputing in China is fluctuating. It was placed 51^{st} in June 2003, then 14^{th} in November 2003, 10^{th} in June 2004, 5^{th} during 2005. By mid-2010, it had reached the second spot and at the end of 2010 the top spot (Supercomputing in China, 2014). To avoid future technology embargo restrictions, the Chinese are developing their own processors like Loongson, a MIPS type processor, etc (Supercomputing in China, 2014).

Supercomputing Centers in China

Tianjin

The National Supercomputing Center in Tianjin is one of the main centers. It houses the Tianhe-1 supercomputer.

Beijing

The Supercomputing Center of the China Academy of Sciences provides academic support functions to the national centers.

Shenzhen

The National Supercomputing Center in Shenzhen houses the second fastest machine in China and the third fastest in the world.

Shanghai

The Shanghai Supercomputer Center operates the Magic Cube supercomputer that runs at 230 teraflops.

Guangzhou

The National Supercomputing Center in Guangzhou operates the top ranked supercomputer in the world Tianhe-2 (MilkyWay-2).

Other centers include Jinan, Hunan, Changsha, etc (Supercomputing in China, 2014).

SUPERCOMPUTING IN EUROPE

Several supercomputing centers are present in Europe and the access to them is coordinated by European initiatives to facilitate high-performance computing. In June 2011, France's Tera 100 was certified the fastest supercomputer in Europe and ranked 9th in the world at that time (Supercomputing in Europe, 2014).

Supercomputing Centers in Europe

Belgium

In 2012, Ghent University in Belgium inaugurated the first Tier 1 supercomputer of the Flemish Supercomputer Center (VSC). The new cluster is ranked 163rd in the world's Top 500 list of supercomputers according to November 2012 (Supercomputing in Europe, 2014).

Bulgaria

The National Center for Supercomputing Applications in Sofia operates an IBM Blue Gene/P supercomputer. The system ranked as 379 in the Top500 list in November 2009 (Supercomputing in Europe, 2014).

Finland

CSC-IT Center for Science is operating a Cray XC30 system with 244 TFlop/s which is to be scaled up to PFlop/s range in 2014 (Supercomputing in Europe, 2014).

France

The CEA operates the Tera 100 machine in the Research and Technology Computing Center in Essonne, France. The Tera 100 has a peak processing speed of 1050 teraflops making it the fastest supercomputer in Europe as of 2011 (Supercomputing in Europe, 2014).

Germany

The three national centers at Garching (LRZ), Juelich (JSC) and Stuttgart (HLRS) together form the Gauss Center for Supercomputing. The Julich Supercomputing Center and the Gauss Center for Supercomputing jointly own the JUGENE computer. JUGENE is based on IBM's BlueGene/P architecture and in June 2011 was ranked the 12th fastest computer in the world by TOP500 (Supercomputing in Europe, 2014).

The Leibniz-Rechenzentrum, a supercomputing center in Munich, houses the SuperMUC system which began operations in 2012 at a processing speed of 3 petaflops. The High Performance Computing Center in Stuttgart fastest computing system is HERMIT with a peak performance of more than 1 petaflops. HERMIT was ranked 12th in the November 2011 TOP500 list (Supercomputing in Europe, 2014).

Italy

The main supercomputing facility is CINECA, a consortium of many universities and research institutions scattered throughout the country. A supercomputer FERMI based on IBM's BlueGene/Q architecture was commissioned in Spring 2012. FERMI is the fastest supercomputer in Italy with a peak performance of 2.1 PFLOPS. Italy also hosts some of the largest nodes of the worldwide LHC Computing Grid (Supercomputing in Europe, 2014).

Netherlands

The European Grid Infrastructure, a distributed computing system is headquartered at the Science Park in Amsterdam (Supercomputing in Europe, 2014).

Norway

The Norwegian University of Science and Technology in Trondheim operates the Vilje supercomputer owned by NTNU and the Norwegian Meteorological Institute (Supercomputing in Europe, 2014).

Poland

The Polish Grid Infrastructure PL-Grid was built between 2009 and 2011 as a nationwide computing infrastructure. The Galera computer cluster at the Gdansk University of Technology was ranked 299th on the TOP500 list in November 2010. The Zeus Computer cluster at the ACK Cyfronet AGH in Krakow was ranked 106th on the TOP500 list in November 2012 (Supercomputing in Europe, 2014).

Russia

In November 2011, the 33,072-processor Lomonosov supercomputer in Moscow was ranked the 18th fastest supercomputer in the world and the third fastest in Europe. In September 2011, T-Platforms stated that it would deliver a water-cooled supercomputer in 2013 (Supercomputing in Europe, 2014).

Slovenia

The Slovenian National Grid Initiative (NGI) provides resources to the European Grid Initiative (EGI). ARNES manages a cluster for testing computing technology where users can also submit jobs. The cluster consists of 2300 cores. Arctur also provides computer resources on its Arctur-1 supercomputer to the Slovenian NGI (Supercomputing in Europe, 2014).

Spain

The Barcelona Supercomputing Center operates the 1 petaflop MareNostrum supercomputer. The Supercomputing and Visualization Center of Madrid at the Technical University of Madrid operates the 72-teraflop Magerit supercomputer. The Spanish Supercomputing Network furthermore provides access to several supercomputers distributed across Spain (Supercomputing in Europe, 2014).

Sweden

The National Supercomputer Center in Sweden operates the Triolith supercomputer which achieved 407.2 Teraflop/s on the Linpack benchmark which was ranked 79 on the November 2013 issue of Top 500 list of the fastest supercomputers in the world (Supercomputing in Europe, 2014).

Switzerland

The Swiss National Supercomputing Centre was founded in 1991 and operated by ETH Zurich. The IBM Aquasar supercomputer became operational at ETH Zurich in 2010. It uses hot water cooling to achieve heat efficiency (Supercomputing in Europe, 2014).

United Kingdom

The EPCC Supercomputer center was established at the University of Edinburgh in 1990. The HECToR project at the University of Edinburgh provides supercomputing facilities using a 360-teraflop Cray XE6 system, the fastest supercomputer in the UK (Supercomputing in Europe, 2014).

SUPERCOMPUTING IN INDIA

India's supercomputer program was started in late 1980s because Cray supercomputers were denied for import due to an arms embargo imposed on India. PARAM 8000 is considered India's first supercomputer. It was built in 1990 by Centre for Development of Advanced Computing (CDAC) and was replicated and installed at ICAD Moscow in 1991 under Russian collaboration (Supercomputing in India, 2014).

As of November 2013, India has 12 systems on the Top500 list ranking 44, 84, 100, 107, 131, 226, 316, 367, 390, 427, 428, 429 (Supercomputing in India, 2014).

Table 9. List of Supercomputers in India and their locations (Source: Supercomputing in India, 2014)

Rank	Site	Name	Rmax (TFlop/s)	Rpeak (TFlop/s)
44	Indian Institute of Tropical Meteorology	iDataPlex DX360M4	719.2	790.7
84	Centre for Development of Advanced Computing	PARAM Yuva - II	388.4	520.4
100	CSIR Centre for Mathematical Modelling and Computer Simulation	Cluster Platform 3000 BL460c Gen8	334.3	362.0
107	National Centre for Medium Range Weather Forecasting	iDataPlex DX360M4	318.4	350.1
131	Indian Institute of Technology Kanpur	Cluster Platform SL230s Gen8	282.6	307.2
226	Vikram Sarabhai Space Centre, ISRO	SAGA - Z24XX/SL390s Cluster	188.7	394.8
316	Manufacturing Company India	Cluster Platform 3000 BL460c Gen8	149.2	175.7
367	IT Services Provider (B)	Cluster Platform 3000 BL460c Gen8	139.2	195.3
291	Computational Research Laboratories	EKA - Cluster Platform 3000 BL460c	132.8	172.6
427	Semiconductor Company (F)	Cluster Platform 3000 BL460c Gen8	129.2	182.0
428	Semiconductor Company (F)	Cluster Platform 3000 BL460c Gen8	129.2	182.0
429	Network Company	Cluster Platform 3000 BL460c Gen8	128.8	179.7

Table 10. Comparison of the supercomputers between different countries as of November 2013 (Source: Supercomputing in India, 2014).

Country	Total Rmax (Gflops)	Number of Computers in TOP500	System Share (%)
India	3,040,297	12	2.4
China	48,549,093	63	12.6
France	9,489,912	22	4.4
Germany	13,696,834	20	4
Japan	22,472,218	28	5.6
Russia	1,846,613	5	1
Poland	455,909	2	0.4
South Korea	1,258,060	5	1
UK	9,058,329	23	4.6
USA	118,261,596	264	52.8
Canada	2,077,842	10	2
Italy	2,665,609	5	1
Australia	2,180,151	5	1

Supercomputing Centers in India

Prithvi

Indian Institute of Tropical Meteorology, Pune, has a machine with a peak of 790.7 teraflop/s, called Prithvi that is used for climate research and forecasting. It is ranked 36th among the world's top 500 supercomputers as of June 2013 list (Supercomputing in India, 2014).

PARAM Yuva II

Unveiled on February 8, 2013, this supercomputer was made by Centre for Development of Advanced Computing in a period of three months at a cost of 3 million US dollars. It performs at a peak of 524 teraflop/s. The supercomputer can deliver sustained performance of 360.8 teraflop/s on Linpack benchmark. In terms of power efficiency, it would have been ranked 33rd in the November 2012 list of Top Green 500 Supercomputers of world (Supercomputing in India, 2014).

Param Yuva II will be used for research in space, bioinformatics, weather forecasting, pharmaceutical development, etc. Educational institutes like Indian Institutes of Technology and National Institutes of Technology can be linked to the computer through the national knowledge network (Supercomputing in India, 2014).

SAGA-220

SAGA-220 built by ISRO, is capable of performing at 220,000 gigaflop/s. It uses about 400 NVIDIA Tesla 2070 GPUs and 400 Intel Quad Core Xeon CPUs (Supercomputing in India, 2014).

EKA

EKA is a supercomputer built by the Computational Research Laboratories with technical assistance and hardware provided by Hewlett-Packard (HP). It is capable of performing at 132 teraflop/s (Supercomputing in India, 2014).

Virgo

IIT, Madras has a 91.1 teraflop/s machine virgo. It is ranked as 364 in the top 500 November 2012 list. It has 292 computer nodes, 2 master nodes, 4 storage nodes and has total computing power of 97 TFlops. According to Linpack performance, Virgo is the fastest cluster in an institution in India. As of 2012, Virgo is at 224th position in the world, 5th ranked energy efficient machine in the world and 1st ranked energy efficient machine in India (Supercomputing in India, 2014).

PARAM Yuva

It belongs to the PARAM series of supercomputer developed by the Centre for Development of Advanced Computing. It is capable of performing at about 54 teraflop/s (Supercomputing in India, 2014).

SUPERCOMPUTING IN JAPAN

Japan operates a number of supercomputing centers which holds world records in speed. The K Computer is the world's fastest as of June 2011 (Supercomputing in Japan, 2014).

Supercomputing Centers in Japan

The GSIC Center at the Tokyo Institute of Technology houses the Tsubame 2.0 supercomputer which has a peak of 2,288 Tflops and in June 2011, ranked fifth in the world (Supercomputing in Japan, 2014).

The RIKEN MDGRAPE-3 for molecular dynamics simulations of proteins is a special purpose petascale supercomputer at the Advanced Center for Computing and Communication (Supercomputing in Japan, 2014).

The next system is Japan Atomic Energy Agency's PRIMERGY BX900 Fujitsu supercomputer. It is ranked as the 38th in the world in 2011 (Supercomputing in Japan, 2014).

DENIGMA is a highly cost and energy efficient computer cluster at the Nagasaki Advanced Computing Center, Nagasaki University. It is used for hierarchical N-body simulations and has a peak performance of 111 TFLOPS with an energy efficiency of 1376 MFLOPS/watt (Supercomputing in Japan, 2014).

The Computational Simulation Center, Japan Atomic Energy Agency operates a 1.52 PFLOPS supercomputer in Aomori. The system called as Helios is used for fusion simulation projects (Supercomputing in Japan, 2014).

The University of Tokyo's Information Technology Center in Kashiwa, Chiba began operations of a 1.13 –PFLOPS supercomputer system in April 2012 (Supercomputing in Japan, 2014).

In June 2012, the Numerical Prediction Division, Forecast department of the Japan Meteorological Agency deployed an 847-TFLOPS Hitachi SR16000/M1 supercomputer based on IBM Power 775 in Tokyo (Supercomputing in Japan, 2014).

SUPERCOMPUTING IN PAKISTAN

The fastest supercomputer currently in use in Pakistan is developed and hosted by the National University of Sciences and Technology at its modeling and simulation research centre (Supercomputing in Pakistan, 2014). As of November 2012, there are no supercomputers from Pakistan on the Top500 list (Supercomputing in Pakistan, 2014).

Supercomputing Programs in Pakistan

GIK Institute

HPC platform has been established by Dr. Masroor Hussain, FCSE, GIK Institute. It consists of Dell R815 with 64 CPU cores, 256 GB RAM, 1.8 TB Secondary memory.3 compute nodes are present, Dell Power Connect 8024F layer-3 manageable switch (Supercomputing in Pakistan, 2014). The software consists of Rocks Cluster 6.1 (Emerald Boa) over CentOS has been installed and configured (Supercomputing in Pakistan, 2014).

COMSATS

The COMSATS Institute of Information Technology built a cluster based supercomputer for research purposes in 2004(Supercomputing in Pakistan, 2014). The project was funded by the Higher Education Commission of Pakistan. The Linux-based computing cluster which was tested and configured for optimization achieved a performance of 158 GFLOPS per second (Supercomputing in Pakistan, 2014).

NUST

The National University of Sciences and Technology (NUST) in Islamabad has developed the fastest supercomputing facility in Pakistan till date. The supercomputer has parallel computational abilities and has a performance of 132 teraflops per second making it the fastest GPU parallel computing system in Pakistan (Supercomputing in Pakistan, 2014). The cluster consists of a 66 NODE supercomputer with 30, 992 processor cores, 2 head nodes, 32 dual quad core computer nodes and 32 Nvidia computing processors. Each processor has 960 processor cores, QDR InfiniBand interconnection and 21.6 TB SAN storage (Supercomputing in Pakistan, 2014).

University of Lahore

Installing supercomputer in 2015 (Supercomputing in Pakistan, 2014).

KUST

The cluster is deployed at Kohat University of Science and Technology. Cluster name is KUST-Kohat, number of CPUs are 104, peak performance of 416 GFLOPS (Supercomputing in Pakistan, 2014).

CONCLUSION

This chapter provides an overview of global supercomputing and also as an introduction to the entire book of edited chapters. The subsequent chapters of this book include further discussion of Seymore Cray the father of supercomputing, and the history of supercomputing from past to present to year 2050, and a history of supercomputing and supercomputer centers. The subsequent chapters include detailed studies of applications of supercomputing to nuclear power, nano-materials, steganography, cloud computing, leadership, sequence analysis and genome annotation, population genetics, programming paradigms, modeling of biological systems, renewable energy network design and optimization, data mining, philosophical logic perspective, a philosophy of their language, and current challenges across multiple domains of computational science.

The reader is referred to the Appendix of this book for more information about each of the top 500 supercomputers in the world as of November 2013 that was gratefully provided to us by written permission from Top500.com for educational purposes.

REFERENCES

Argonne National Laboratory. (2014). *Mira: An engine for discovery*. Retrieved from http://www.alcf.anl.gov/mira

Barth, B. (2013). *High performance computing (HPC) systems*. Retrieved from https://www.tacc.utexas.edu/resources/hpc

Biswas, R., & Thigpen, W. (2014). *High-end computing capability*. Retrieved from http://www.nas.nasa.gov/hecc/resources/pleiades.html

Cook, J. S. (2011). Supercomputers and supercomputing. In Q. Zhang, R. S. Segall, & M. Cao (Eds.), *Visual analytics and interactive technologies* (pp. 282–294). IGI Global. doi:10.4018/978-1-60960-102-7.ch017

Forbus, K. (n.d.). *Symbolic supercomputer for artificial intelligence and cognitive science research*. Retrieved from http://www.qrg.northwestern.edu/projects/SymbolicSupercomputing/symb-supercomp_index.html

Gardner, J. (2014). *Introduction to high performance computing (HPC) and the NSF TeraGrid*. Retrieved from http://oldwww.phys.washington.edu/users/gardnerj/tching/IntroToTeraGrid.pdf

Hager, G., & Wellein, G. (2011). Introduction to high performance computing for scientists and engineers. Boca Raton, FL: Academic Press.

Human Brain Project. (2012). Retrieved May 25, 2104 from www.humanbrainproject.eu

Human Brain Project-Video Overview. (n.d.). Retrieved May 25, 2014 from YouTube at: http://www.youtube.com/watch?v=JqMpGrM5ECo

Kondo, J. (Ed.). (1991). *Supercomputing applications, algorithms, and architectures for the future of supercomputing*. Springer-Verlag Tokyo.

Kramer, B., Beldica, C., Gropp, B., & Hwu, W. (2014). *Blue waters sustained petascale computing*. Retrieved May 21, 2014, Retrieved from https://bluewaters.ncsa.illinois.edu/team

Leibniz Supercomputing Centre. (2014). *Leibniz supercomputing centre of the Bavarian academy of sciences and humanities*. Retrieved from http://www.lrz.de/services/compute/supermuc/system-description/

Levesque, J., & Wagenbreth, G. (2011). *High performance computing: Programming and applications*. Chapman & Hall/CRC.

Lu, Y. (2014). *Overview of tianhe-2 (MilkyWay-2) supercomputer*. Retrieved from http://www.slideshare.net/ultrafilter/th2-isc13inspurlyt

Lucifredi, F. (2013). *Supercomputing on the cheap with parallella*. Retrieved from http://programming.oreilly.com/2013/12/supercomputing-on-the-cheap-with-parallella.html

Miranker, G. S. (1992). *Titan III supercomputer architectural overview*. Retrieved from http://link.springer.com/article/10.1007%2FBF02241705

Miyazaki, H., Kusano, Y., Shinjou, N., Shoji, F., Yokokawa, M., & Watanabe, T. (2012). *Overview of the K computer system*. Retrieved from http://www.fujitsu.com/downloads/MAG/vol48-3/paper02.pdf

PTI. (2012, September 17). India plans 61 times faster supercomputer by 2017. *The Times of India.* Retrieved May 25, 2104 from http://articles.timesofindia.indiatimes.com/2012-09-17/hardware/33901529_1_first-supercomputers-petaflop-fastest-supercomputer

Segall, R. S. (2013a). Dimensionalities of computation: From global supercomputing to data, text and web mining. *Invited Plenary Address at International Institute of Informatics and Systemics (IIIS) Conference, 17th World Multi-conference on Systemics, Cybernetics and Informatics (WMSCI 2013).* Retrieved from http://www.iiis.org/ViewVideo2013.asp?id=14

Segall, R. S. (2013b). Computational dimensionalities of global supercomputing. *Journal of Systemics, Cybernetics, and Informatics*, *11*(9), 75–86.

Segall, R. S., & Cook, J. (2014). Data visualization and information quality by supercomputing. In *Proceedings of the Forty-Fifth Meeting of Southwest Decision Sciences Institute (SWDSI).* Dallas, TX: Academic Press.

Segall, R. S., & Zhang, Q. (2013). Information quality and supercomputing. In *Proceedings of the 18th International Conference on Information Quality (ICIQ 2013)* (pp. 400-446). Academic Press.

Segall, R. S., Zhang, Q., & Cook, J. S. (2013). Overview of current research in global supercomputing. In *Proceedings of Forty-Fourth Meeting of Southwest Decision Sciences Institute (SWDSI).* Albuquerque, NM: Academic Press.

Segall, R. S., Zhang, Q., & Pierce, R. (2010a). Data mining supercomputing with SAS JMP® genomics. In *Proceedings of 14th World Multi-Conference on Systemics, Cybernetics and Informatics: WMSCI 2010*. Orlando, FL: Academic Press.

Segall, R. S., Zhang, Q., & Pierce, R. (2011). Data mining supercomputing with SAS JMP® genomics. *Journal of Systemics, Cybernetics and Informatics*, *9*(1), 28–33.

Segall, R. S., Zhang, Q., & Pierce, R. M. (2009). Visualization by supercomputing data mining. In *Proceedings of the 4th INFORMS Workshop on Data Mining and System Informatics*. San Diego, CA: Academic Press.

Segall, R. S., Zhang, Q., & Pierce, R. M. (2010b). Data mining supercomputing with SAS JMP® genomics: Research-in-Progress. In *Proceedings of 2010 Conference on Applied Research in Information Technology*. University of Central Arkansas.

Supercomputing in China. (2014). Retrieved from http://en.wikipedia.org/wiki/Supercomputing_in_China

Supercomputing in Europe. (2014). Retrieved from http://en.wikipedia.org/wiki/Supercomputing_in_Europe

Supercomputing in India. (2014). Retrieved from http://en.wikipedia.org/wiki/Supercomputing_in_India

Supercomputing in Japan. (2014). Retrieved from http://en.wikipedia.org/wiki/Supercomputing_in_Japan

Supercomputing in Pakistan. (2014). Retrieved from http://en.wikipedia.org/wiki/Supercomputing_in_Pakistan

Swiss National Supercomputing Centre. (2014). *Piz daint*. Retrieved from http://www.cscs.ch/computers/piz_daint/index.html

Szczepariski, A. F., Huang, J., Baer, T., Mack, Y. C., & Ahern, S. (2013). Data analysis and visualization in high-performance computing. *Computer*, *46*(May), 84–92. doi:10.1109/MC.2012.192

Top500. (2013). *Supercomputer sites*. Retrieved May 25, 2104 from http://www.top500.org/lists/2103/11

Total, S. A. (2013). *Total inaugurates its new pangea supercomputer, ranking it among the global top ten in terms of computing power*. Retrieved May 22, 2014 from http://total.com/en/media/news/press-releases/20130321-Total-inaugurates-its-new-Pangea-supercomputer-ranking-it-among-the-global-top-ten-in-terms-of-computing-power

van der Steen, A. J. (2005). *Overview of recent supercomputers*. Retrieved from http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.63.3326&rep=rep1&type=pdf

Vetter, J. S. (2013). Contemporary high performance computing: From petascale toward exascale. Chapman & Hall/CRC Computational Science.

Wagstaff, K. (2012). What exactly is a supercomputer? *Time*, *19*(June). Retrieved from http://techland.time.com/2012/06/19/what-exactly-is-a-supercomputer/

Whittaker, Z. (2012). *U.S. IBM supercomputer is world's fastest: Does it matter?* Retrieved from http://www.zdnet.com/blog/btl/u-s-ibm-supercomputer-is-worlds-fastest-does-it-matter/80122

Wikipedia. (2014a). *Cray*. Retrieved May 25, 2014 from http://en.wikipedia.org/wiki/Cray

Wikipedia. (2014b). *Flynn's taxonomy*. Retrieved May 25, 2014 from http://en.wikipedia.org/wiki/Flynn%27s_taxonomy

Wikipedia. (2014c). *Human Brian project*. Retrieved May 25, 2014 from http//en.wikipedia.org/wiki/Human_Brian_Project

Wikipedia. (2014d). *Jungle computing*. Retrieved May 25, 2014 from http://en.wikipedia.org/wiki/Jungle_computing

Wikipedia. (2014e). *Multiprocessing*. Retrieved May 25, 2104 from http://en.wikipedia.org/wiki/Multiprocessing

Wikipedia. (2014f). *Supercompter*. Retrieved 5/25, 2014 from http://en.wikipedia.org/wiki/Supercomputer

Wikipedia. (2014g). *Tianhe-2, or Milky Way-2*. Retrieved May 25, 2014 from http://en.wikipedia.org/wiki/Tianhe-2

Wikipedia. (2014h). *Top500*. Retrieved 5/25, 2014 from http://en.wikipedia.org/wiki/Top500

Wikipedia. (2014i). *Vector processing*. Retrieved May 25, 2014 from http://en.wikipedia.org/wiki/Vector_processing

ADDITIONAL READING

x58 supercomputer. (2014). Retrieved May 28, 2014 from http://www.asrock.com/mb/Intel/X58%20SuperComputer/

Adiga, N. R., Almasi, G., & Aridor, Y. (2002). An overview of the BlueGene/L supercomputer *Supercomputing, ACM/IEEE 2002 Conference,* Retrieved May 25, 2104 from http://ieeexplore.ieee.org/xpl/login.jsp?tp=&arnumber=1592896&url=http%3A%2F%2Fieeexplore.ieee.org%2Fxpls%2Fabs_all.jsp%3Farnumber%3D1592896. doi:10.1109/SC.2002.10017

Alspector, J., Brenner, A., Leheny, R., & Richmann, J. (2014). China - A new power in supercomputing hardware. Retrieved May 28, 2014 from https://www.ida.org/~/media/Corporate/Files/Publications/IDA_Documents/ITSD/ida-document-ns-d-4857.ashx

Anderson, N., Moody, J., & Lampf, G. (2012). Ciena brings 100G and packet networking innovations to supercomputing 2012. Retrieved May 27, 2014 from http://www.ciena.com/about/newsroom/press-releases/Ciena-Brings-100G-and-Packet-Networking-Innovations-to-Supercomputing-2012.html

Brocade Communication Systems (2013). *Supercomputing 2013 brocade overview*, part I. (Director). (06 December 2013).[Video/DVD]

Brueckner, R. (2011). Video: An overview of CINECA - Italy's largest supercomputer center. Retrieved May 27, 2014, Retrieved fromhttp://insidehpc.com/2011/06/30/video-an-overview-of-cineca-italys-largest-supercomputer-center/

Centre for Development of Advanced Computing. (2014). Architecture overview. Retrieved May 28, 2014 from http://pune.cdac.in/html/npsf/anant/anant1.aspx

Chi, X. (2011). Supercomputing applications in CAS. Retrieved May 28, 2014 from http://www.prace-project.eu/IMG/pdf/chi.pdf

CISL. A computing laboratory. (2014). Retrieved May 27, 2014 from https://www2.cisl.ucar.edu/cisl-computing

Cluster computing overview (n.d.). Retrieved May 28, 2014 from https://www.google.co.in/url?sa=t&rct=j&q=&esrc=s&source=web&cd=72&cad=rja&uact=8&ved=0CDIQFjABOEY&url=http%3A%2F%2Fwww.calstatela.edu%2Ffaculty%2Frpamula%2Fcs370%2Foverview.ppt&ei=cDuFUvzIM_g8AXX3IEQ&usg=AFQjCNE0FzQHKRmhsyPb_LVJSkMsH2GlnQ&sig2=a4qZjekc4AbRmd13lTLXLw&bvm=bv.67720277,d.dGc

Cooling supercomputers using geothermal energy. (2014). Retrieved May 28, 2014 from http://www.csiro.au/Outcomes/Energy/Renewables-and-Smart-Systems/pawsey-geothermal-supercomputing-project.aspx

Defanti, T. (1994). Overview of the I-WAY: Wide area visual supercomputing. Retrieved May 27, 2014 from http://toolkit.globus.org/alliance/publications/papers/iway_overview.pdf

Edelman, A. (2007). The inside story behind interactive supercomputing's star-P platform for high performance computing for MATLAB. Retrieved May 28, 2014 from http://www.cs.berkeley.edu/~yelick/cs267_sp07/lectures/lecture06/lecture06_starp_edelman07.pdf

Ellsworth, M., Goth, G., Zoodsma, R., Arvelo, A., Campbell, L., & Anderl, W. (2012). An overview of the IBM power 775 supercomputer water cooling system. Retrieved May 27, 2014 from http://electronicpackaging.asmedigitalcollection.asme.org/article.aspx?articleid=1410200

Elstner, D. (2014). Supercomputing and E-science, Leibniz-Institut fur Astrophysik Potsdam. Retrieved May 28, 2014 from http://www.aip.de/en/research/research-area-drt/research-groups-and-projects-1/e-science/supercomputing-and-e-science

Feng, W., Feng, X., & Ge, R. (2008). Green supercomputing comes of age. Retrieved May 28, 2014 from http://www.computer.org/csdl/mags/it/2008/01/mit2008010017-abs.html

George, A. (2011). Novo-G: At the forefront of scalable reconfigurable supercomputing. Retrieved May 28, 2014 from http://www.gidel.com/pdf/Novo_G_At_Forefront.pdf

Hurley, J. (2014). CAS- supercomputing -green II, Swinburne University of Technology, Hawthorn, Victoria, Australia. Retrieved May 27, 2014 from http://astronomy.swin.edu.au/supercomputing/green2/

InfoWebLinks – supercomputer (n.d.). Retrieved May 28, 2014 from http://www.infoweblinks.com/content/supercomputers.htm

Intel supercomputing systems division (n.d.). Retrieved May 28, 2014 from http://www.new-npac.org/projects/cdroms/cewes-1999-06-vol1/nhse/hpccsurvey/orgs/intel/intel.html

Introducing titan - advancing the era of accelerated computing. (2012). Retrieved May 27, 2014, from https://www.olcf.ornl.gov/titan/

Lixing, Z. (2009). Supercomputers. Retrieved May 27, 2014 from https://www.google.co.in/url?sa=t&rct=j&q=&esrc=s&source=web&cd=4&cad=rja&uact=8&ved=0CEsQFjAD&url=http%3A%2F%2Fwww.eng.auburn.edu%2F~agrawvd%2FCOURSE%2FE6200_Fall09%2FCLASS_TALKS%2Fsupercomputer.docx&ei=0n6EU6juLsj98QW8ooFg&usg=AFQjCNFX0izRSYgXNGTCM0yEyerjFFdFYg&sig2=uPSOII11gRobzPkkJkIZFA

Mattson, T., & Henry, G. (1998). An overview of the Intel TFLOPS supercomputer. Retrieved May 27, 2014 from http://www.ai.mit.edu/projects/aries/course/notes/ascii_red.pdf

Moscow University supercomputing center (2014). Supercomputing applications. Retrieved May 27, 2014 from http://hpc.msu.ru/?q=node/77

Neeman, H. (2013). Supercomputing in plain English, University of Oklahoma Supercomputing Center for Education & Research, Norman, OK. Retrieved May 27, 2014 from https://www.google.co.in/url?sa=t&rct=j&q=&esrc=s&source=web&cd=8&cad=rja&uact=8&ved=0CGoQFjAH&url=http%3A%2F%2Fwww.oscer.ou.edu%2FWorkshops%2FOverview%2FSiPE_Overview.ppt&ei=0n6EU6juLsj98QW8ooFg&usg=AFQjCNGBGjqFlvc4khw_2ZCdMq5igMCgDA&sig2=zmQ4oqK9PcFx6cCl9ncO2g

New Zealand supercomputing - product specification. (2005). Retrieved May 28, 2014 from http://www.gen-i.co.nz/assets/productspecs/New-Zealand-Super-Computing.pdf

Overview of recent supercomputers. (2012). Retrieved May 27, 2014 from http://www.euroben.nl/reports/web12/overview.php

Overview of supercomputer systems. (2011). Retrieved May 27, 2014 from http://nkl.cc.u-tokyo.ac.jp/13e/01-Intro/ITC.pdf

Oxford Supercomputing Centre. (2012). An introduction to the Oxford supercomputing centre, Oxford University, UK. Retrieved May 28, 2014 from http://www.arc.ox.ac.uk/sites/default/files/uploads/OSC%20Overview%20May%2012.pdf

Pittsburgh supercomputing center (2012). Retrieved May 28, 2014 from http://www.novell.com/docrep/2012/04/pittsburgh_supercomputing_center.pdf

Rao, A. (2012). SeaMicro technology overview. Retrieved May 28, 2014 from http://www.seamicro.com/sites/default/files/SM_TO01_64_v2.5.pdf

Regalado, A. (2014). A $1 billion investment plan for IBM's watson supercomputer. Retrieved May 28, 2014 from http://www.technologyreview.com/news/523411/facing-doubters-ibm-expands-plans-for-watson/

Schroers, W. (2013). Parallel and distributed computing – supercomputing, Field-theory.org, Berlin, Germany. Retrieved May 28, 2014 from http://www.field-theory.org/articles/supercomputing/index.html

Silicon Graphics. (1995). The power challenge technical report. Retrieved May 28, 2014 from http://www.uoks.uj.edu.pl/resources/flugor/POWER/chap1.html

Supercomputer architecture (2014). Retrieved May 26, 2014 from http://en.wikipedia.org/wiki/Supercomputer_architecture

Supercomputers (2014). Retrieved May 28, 2014 from http://www.chilton-computing.org.uk/ccd/supercomputers/overview.htm

Supercomputing architectures: An overview (n.d.). Retrieved May 28, 2014 from http://www.peter-india.net/SuperComputingView.html

Supercomputing in small spaces. (2010). Retrieved May 28, 2014 from http://sss.cs.vt.edu/overview.php

Supercomputing record with bubble collapse simulation. (2013). Retrieved May 28, 2014 from https://www.ethz.ch/en/news-and-events/eth-news/news/2013/11/supercomputing-record.html

Ulmer, D. (2008). The Swiss national supercomputing centre CSCS: An overview. Retrieved May 28, 2014 from http://www.linksceem.eu/ls2/images/stories/DU_Kick-off.pdf

KEY TERMS AND DEFINITIONS

Cray: Cray Inc., an American supercomputer manufacturer based in Seattle, Washington. The company's predecessor, Cray Research, Inc. (CRI), was founded in 1972 by computer designer Seymour Cray (Wikipedia, 2014a).

Flynn's Taxonomy: Classification of computer architectures, proposed by Michael J. Flynn in 1966 (Wikipedia, 2014b).

Human Brain Project: A large scientific research project, directed by the École Polytechnique Fédérale de Lausanne and largely funded by the European Union, which aims to simulate the complete human brain on supercomputers to better understand how it functions. The project is based in Geneva, Switzerland (Wikipedia, 2014c).

Jungle Computing: The use of diverse [vague], distributed and highly non-uniform [vague] high performance computer systems to achieve peak performance (Wikipedia, 2014d).

Multiprocessing: The use of two or more central processing units (CPUs) within a single computer system. The term also refers to the ability of a system to support more than one processor and/or the ability to allocate tasks between them (Wikipedia, 2014e).

Supercomputer: A computer at the frontline of contemporary processing capacity – particularly speed of calculation which can happen at speeds of nanoseconds (Wikipedia, 2014f).

Tianhe-2, or Milky Way-2: Tianhe-2 or TH-2 (Chinese: 天河-2; pinyin: tiānhé-èr; literally: "Heavenriver-2" idiomatically "Milky Way

2") is a 33.86 petaflops supercomputer located in Sun Yat-sen University, Guangzhou, China. It was developed by a team of 1300 scientists and engineers. It is the world's fastest supercomputer according to the TOP500 list for June and November 2013 (Wikipedia, 2014g).

TOP500: The TOP500 project ranks and details the 500 most powerful (non-distributed) computer systems in the world. The project was started in 1993 and publishes an updated list of the supercomputers twice a year (Wikipedia, 2014h).

Vector Processing: A vector processor, or array processor, is a central processing unit (CPU) that implements an instruction set containing instructions that operate on one-dimensional arrays of data called vectors. This is in contrast to a scalar processor, whose instructions operate on single data items (Wikipedia, 2014i).