Supercomputers: Past, Present, and the Future

by Sumit Narayan

"At each increase of knowledge, as well as on the contrivance of every new tool, human labor becomes abridged."

—Charles Babbage

Computers have become an integral part of everyday life of everyone, more so than Charles Babbage could have foreseen. Computers are more popular and much faster today than they were several years ago. There has been constant improvement in the speed and complexity of microprocessors—the heart of every computing machine. In 1965, Gordon Moore, the co-founder of Intel, predicted that the complexity of integrated circuits would approximately double every year. From single core to multi-core processors, the complexity of these microchips is directly linked to their processing speed, and they have been improving exponentially over time, strictly following Moore's Law. Introduced in 1954, the IBM 704 was the first mass-produced computer with floating-point arithmetic hardware. It was capable of executing 40,000 instructions per second. Today's desktop or laptop computers can perform tens of millions of instructions per second. We have come a long way, driven by dreams of faster calculations, and supercomputers, computers that form the forefront of all computing machines, were designed to fulfill those dreams. From the mysterious Antikythera mechanism calculators to modern-age petaflop supercomputers, we have constantly been searching for ways to make computations faster. Today's supercomputers can do thousands of trillions of floating-point computations per second and have been constantly improving. Applications such as weather prediction or nuclear reaction simulations are comprised of gazillions of operations and may take several days to complete. IBM, Cray, Hitachi, SGI, Fujitsu, and many others have invested millions of dollars and countless hours of work to develop systems to solve these complex problems. In the past, these systems were available only to government-funded national and international research centers; but recent advancement in technology and competition in the technology industry has made these machines cheaper and more valuable for IT organizations.

Past: Early Supercomputers

CDC-6600, a mainframe computer produced by Control Data Corporation (CDC) in 1965 is regarded as the first successful supercomputer. Designed by Seymour Cray and Jim Thornton, the machine was capable of operating at 9 megaflops (MFLOPS)—thousands of times slower than our current desktops. CDC-6600 was also the first machine to introduce separate processors for the handling of housekeeping tasks such as memory access and input/output (I/O). Until then, the central processing unit (CPU) was in charge of performing all operations: computing, memory, and input/output. I/O in those days was usually done using punch cards or standard magnetic tapes and was extremely slow. By providing separate processors for each event, the CPU was responsible only for computations, thus leaving the CPU with fewer instructions. It also resulted in smaller processor size, allowing it to be operated at a higher clock rate. This novel idea, which later came to be known as reduced instruction set computer (RISC), allowed the CPU, peripheral processors (PPs), and I/O units to operate in parallel, thus improving the overall speed and performance of the machine.

Cray's engineering continued at CDC, resulting in CDC-7600, a successor to and ten times faster than CDC-6600. Jim Thornton, on the other hand, became part of a new project, STAR-100, which was designed to operate at 100 MFLOPS. CDC's STAR-100 was released in 1974 and was one of the first machines to use a vector processor for improving math performance. Vector processing design allowed CPUs to perform mathematical operations on multiple data elements simultaneously. The CPU had to decode only a single instruction, set up the hardware and start feeding the data. This technique remained very popular in the scientific community and formed the basis of design for several supercomputers in the 1980s and 1990s. But, STAR-100, although designed to perform at 100 MFLOPS, gave lower than expected numbers in a "real-world" environment because the serialized part of the processing was still slow. Switching from vectors to normal data was still time-consuming, making the real-world performance slower than expected. This conversion theory was developed by Gene Amdahl in early 1967. However, Amdahl's Law was ignored by architects of the STAR-100.

In 1971, Seymour Cray, unable to secure sufficient funds for his project at CDC, left the company to form Cray Research, where he designed Cray-1 (160 MFLOPS). The Cray-1 provided a good balance between scalar and vector performance and also used registers to dramatically improve performance. Registers are small amounts of memory storage available on processors. Their contents can be accessed at much faster speeds compared to external I/O components. However, because they reside on the processor's chip, they are more expensive to manufacture. They also provide less flexibility in terms of size, so Cray's machine could only read small parts of data at a time. The first release of Cray-1 was in 1976, and it dismissed STAR-100 from its top spot as the fastest supercomputer of that time. The first official customer, the National Center for Atmospheric Research (NCAR), paid $8.86 million to own the supercomputer. This machine shaped the computer industry for years to come. Cray-1 was also Cray's first supercomputer to use integrated circuits (ICs).

Cray-1 was succeeded in 1982 by Cray X-MP (800 MFLOPS), the first multiprocessing computer, and in 1985 by Cray-2, the first machine to break the gigaflop barrier at 1.9 GFLOPS. Cray-2 used all IC components instead of individual components and remained the fastest machine until 1987, when ETA Systems, a spin-off from CDC, designed a 10 GFLOPS machine called ETA-10. ETA-10 used fiber optics for communication between processors and I/O devices. ETA later merged back with CDC in 1989. In the meantime, two new companies, Thinking Machines Corporation (1982) and nCUBE (1983), were founded. Both companies specialized in parallel computing architectures. Thinking Machines, started by graduates from the Massachusetts Institute of Technology, produced several supercomputers released as Connection Machines. By 1993, four of the top five fastest supercomputers belonged to Thinking Machines. nCUBE, on the other hand, was started by a group of Intel employees who wanted Intel to enter into parallel computing but couldn't convince the decision-makers to undertake the endeavor. nCUBE released a parallel computer with the same name. In the mid-1990s, the supercomputer market collapsed, and both companies were acquired by bigger players in the business. The crash also forced Cray Research to merge with Silicon Graphics, Inc. (SGI) in 1996.

One of the major companies that has yet to be mentioned is IBM. Although IBM had built several of the fastest computers in the world (for example the IBM 7030), it was not until 1993 that it entered the supercomputer market with IBM SP-1. It was the first member of the IBM's Scalable POWERparallel distributed memory parallel computer, based on RISC System/6000 processing element, which later became known as POWER (Performance Optimization With Enhanced RISC). In a distributed memory system, the memory and address space of each processor in a multi-processor system is local to itself. The data can only be shared between processors using a message passing interface like IBM's message passing library (MPL). IBM continued releasing several successors to IBM-SP, and it faced stiff competition from other players in the market, such as Hitachi and Intel. At the turn of the century, IBM was at the top of the fastest supercomputer list with IBM ASCI White. It had 8,192 processors, 6 TB of memory, 160 TB of storage space, and operated at 7.226 TFLOPS.

Present: Supercomputers Today

In 1993, based on ideas of Hans Meuer, a professor of Computer Science at the University of Mannheim, Germany, project TOP500 was initiated. The aim of this project was to list the 500 most powerful computer systems in the world. The list, compiled biannually, ranks supercomputers based on their performance on the LINPACK benchmark—a linear algebra library for digital computers that tests the floating point computing power of the system. Table 1 shows the list of fastest supercomputers by period since 1993.

After IBM's ASCI White, Earth Simulator developed by NEC in Japan topped the list from 2002 to 2004. It was developed to understand global climate models, and it was capable of operating at over 35 TFLOPS. IBM returned with BlueGene to reposition itself as the leader in building the fastest supercomputer. Several prototypes of BlueGene were announced: BlueGene/L (released March 2005), BlueGene/C (in development), BlueGene/P (released June 2007) and BlueGene/Q (due 2011). BlueGene remained the fastest supercomputer until 2008, when it was replaced by RoadRunner, also designed by IBM. Other powerful supercomputers released during this period include Cray's XT-3 Red Storm, Cray's XT-4 Franklin, Cray's XT-5 Jaguar, Dell's Thunderbird, SGI's Columbia, and HP's Cluster Platform.

According to the list released in November 2008, the top three supercomputers and their average performance are:

1. IBM's RoadRunner at Los Alamos National Laboratory, USA: 1.105 PFLOPS.
2. Cray's Jaguar XT5 at Oak Ridge National Laboratory, USA: 1.059 PFLOPS.
3. SGI's Pleiades Altix ICE 8200EX at NASA/Ames Research Center, USA: 487.01 TFLOPS.

Period	Supercomputer Name	Maker
06/1993–11/1993	CM-5 (Connection Machine)	Thinking Machine Corp.
11/1993–06/1994	Numerical Wind Tunnel	Fujitsu
06/1994–11/1994	Paragon XP/S	Intel
11/1994–06/1996	Numerical Wind Tunnel	Fujitsu
06/1996–11/1996	SR 2201	Hitachi
11/1996–06/1997	CP-PACS	Hitachi
06/1997–11/2000	ASCI Red	Intel
11/2000–06/2002	ASCI White	IBM
06/2002–11/2004	Earth Simulator	NEC
11/2004–06/2008	BlueGene	IBM
06/2008–06/2009	RoadRunner	IBM

Table 1: Fastest Supercomputers (1993–2009).

IBM's RoadRunner

IBM's RoadRunner is a hybrid system. It uses two different processor architectures—dual-core AMD Opteron server processor, based on AMD64 architecture, and IBM's Cell processor, based on POWER architecture. RoadRunner was built by IBM at Los Alamos National Laboratory in the United States. It sports 6,562 Opteron processors, taking care of standard processing, such as file system I/Os as well as 12,240 PowerXCell 8i processors, handling CPU-intensive tasks, such as mathematical calculations. The system boasts 98 TB of memory and 2 PB of external storage. The machine had a peak performance of 1.7 PFLOPS. This design is significantly different from BlueGene systems, which were based on PowerPC processors. The idea behind BlueGene was to trade the speed of processors for lower power consumption. Thus, BlueGene systems had a notably higher amount of processors compared to other supercomputers giving the same performance.

Cray's Jaguar XT-5

Cray's Jaguar XT-5, which was ranked second in November 2008, is an updated version of Cray's XT-4 supercomputer. It is based on AMD's Opteron quad-core processor. Each Cray XT-5 blade includes four compute nodes for high scalability, and each compute node can be configured with 4–32 GB DDR2 memory. XT-5 blades are interconnected using Cray's SeaStar2+ chips, which provide a very high bi-directional link speed of 9.6 GB/s. The system installed at Oak Ridge National Laboratory in the United States is a combination of both XT-4 and XT-5 machines. In total, the system peaks at 1.6 PFLOPS, consists of 45,376 Opteron processors, has 362 TB of memory and 10 PB of storage space.

Silicon Graphics' Altix

Altix is different from the above two supercomputers in that it is based on Intel processors and is comprised of distributed shared memory machines. The system is installed at the NASA/Ames Research Center/NAS and is nicknamed Pleiades. It consists of 12,800 Intel Xeon processors with 51 TB of RAM and over 1 PB of storage. The system peaks at 608 TFLOPS. The system supplements Columbia, which, with 14,336 cores and 51 TFLOPS, ranked second in 2004, just behind IBM's BlueGene/L. The nodes in Altix are connected using NUMAlink4, developed by SGI, capable of providing bandwidth of up to 6.4 GB/s.

Future: The Next Supercomputers

Normal consumers do not require a supercomputer for their regular computing use. Supercomputers are primarily a necessity of scientists performing mass computing at ultra-high speed. They are used in all plausible domains: space exploration, nuclear energy, climate prediction, environmental simulations, gene technology, math, physics, and many others. While supercomputers excel at highly computationally intensive tasks, they are not the fastest computers on the planet. The human brain controls thousands of human muscles, does audio and visual processing at extremely high speeds, and controls thousands of nerves. It does of these tasks in a fraction of a second and is regarded as the fastest processor in the world. 10 PFLOP is too slow to simulate the whole body, including tissue, blood flow, and movement. The size/performance ratio of the human brain versus that of a supercomputer is beyond comparison. This is an indication that we still have a long way to go.

Construction of supercomputers is a very challenging and expensive task. It may take several years for a supercomputer to move from the laboratory to the market, with costs ranging $150–200 million or more. Most of this work can only be done with the support of government funds and government-funded research centers. Designers of the world's fastest supercomputers, IBM, Cray, SGI, Sun, HP, Hitachi, and many others, are putting forth the effort to create a multi-petaflop machine. IBM is planning a 50 PFLOP machine by the end of 2013, and it is estimated that within the next decade, we will have an exaflop machine. But the

Year	Accomplishments
1962	• Control Data Corporation opened lab; headed by Seymour Cray.
1965	• CDC 6600.
1969	• UNIX operating system developed by group of AT&T employees at Bell Labs.
1972	• Seymour Cray started Cray Research Inc.
1976	• Delivery of Cray-1.
1982	• Delivery of Cray X-MP with two processors.
1985	• Delivery of Cray 2 with four processors; peak performance of 2 GFLOPS. • Delivery of CM-1 by Thinking Machines Corp. • Delivery of iPSC/1 by Intel. • Delivery of Convex C1.
1987	• Delivery of ETA-10. • Delivery of CM-2 by TMC.
1992	• Delivery of Cray-3. • Delivery of CM-5 by TMC. • Production of Paragon/XP series by Intel. • Apple, IBM, and Motorola formed AIM alliance to develop mass marker for POWER processors, resulting in PowerPC.
1993	• TOP500 Project. • Production of IBM SP1.
1997	• Delivery of Intel's ASCI Red, the first TFLOP machine.
1999	• Delivery of IBM's ASCI Blue.
2000	• Delivery of IBM's ASCI White.
2002	• Delivery of NEC's Earth Simulator.
2004	• Delivery of IBM's BlueGene/L. • Delivery of Cray's XT-3.
2006	• Delivery of Cray's XT-4.
2007	• Cell processor released. • Delivery of IBM's BlueGene/P. • Delivery of Cray's XT-5.
2008	• Delivery of IBM's RoadRunner.

Table 2: Chronology of Supercomputers.

process of building faster machines is crippled by input/output units. I/O is not scaling as fast as Moore's Law. Research is being conducted on improving the design and performance of parallel file systems, including introducing solid state drives. Other challenges include the search for experts in computational science, mathematics, and computer science to understand these complex systems and design software for taking advantage of the enormous computing power that these supercomputers provide. As the future unfolds, it will be interesting to see what we will accomplish next.

Acknowledgements

I would like to thank my advisor Dr. John A. Chandy for providing input and sharing his experiences and opinions.

Biography

Sumit Narayan is a PhD candidate at the University of Connecticut, Storrs. He holds a Master's from the University of Connecticut and a Bachelor's of Engineering from University of Madras. His research interests include high-performance computing, parallel file systems, storage system architectures, and I/O subsystems.