Supercomputers: Current Status and Future Trends

Presentation to Melbourne PC Users Group, Inc.

August 2nd, 2016

http://levlafayette.com

0.0 About Your Speaker

0.1 Lev works as the HPC and Training Officer at the University of Melbourne, and do the day-to-day administration of the Spartan HPC/Cloud hybrid and the Edward HPC.

0.2 Prior to that he worked for several years at the Victorian Partnership for Advanced Computing, which he looked after the Tango and Trifid systems and taught researchers at 17 different universities and research institutions. Prior to that he worked for the East Timorese government, and prior to that the Parliament of Victoria.

0.3 He is involved in various community organisations, collect degrees, learns languages, designs games, and writes books for fun.

0.4 You can stalk him at: http://levlafayette.com or https://www.linkedin.com/in/levlafayette

1.0 What is a Supercomputer anyway?

1.1 A supercomputer is a rather nebulous term for any computer system that is at the frontline of current processing capacity.

1.2 The Top500 metric (http://top500.org) measures pure speed of floating point operations with LINPACK. The HPC Challenge uses a variety of metrics (floating point calculation speed, matrix calculations, sustainable memory bandwidth, paired processor communications, random memory updates, discrete Fourier transforms, and communication bandwidth and latency).

1.3 The term supercomputer is not quite the same as "high performance computer", and not quite the same as "cluster computing", and not quite the same as "scientific (or research) computing".

2.0 What are they used for?

2.1 Typically used for complex calculations where many cores are required to operate in a tightly coupled fashion, or for extremely large collections datasets where many cores are required to carry out the analysis simultaneously.

2.1 Scientific uses include; climate modelling and weather forecasting, fluid dynamics, molecular modelling and genomics etc, material sciences, geophysics, astrophyics, particle physics etc, etc.

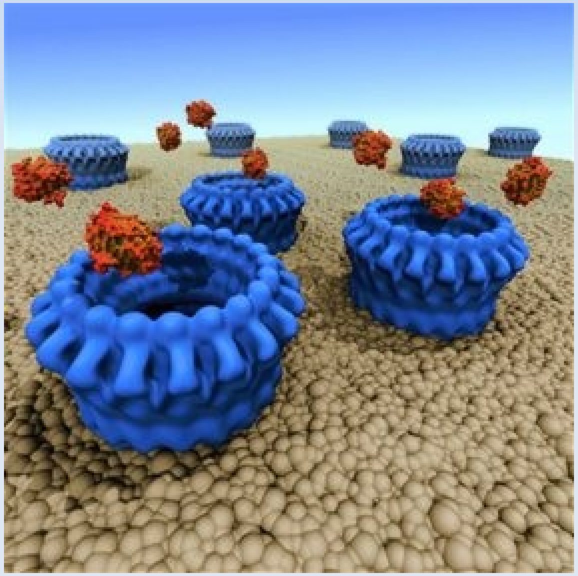

2.3 A local example:
http://www.abc.net.au/science/articles/2010/11/01/3053386.htm
(image of simulated perforin protein by Mike Kuiper at VPAC)

3.0 Historical Examples

3.1 When computers were rare, the proportion of computers that were "supercomputers" was greater. Some major historical examples include:

3.1.1 The CDC 6600 which used germanium-based transistors, and was the fastest computer form 1964-1970. Capable of three megaFLOPS. Successor was the CDC 7600, which could and could deliver about 10 MFLOPS, with a peak of 36 MFLOPS

3.1.2 From 1975 famous Cray 1 was the most notable supercomputer with ICs with two gates per chip and vector processing. First installed at Los Alamos National Laboratory in 1976. Capable of 160 MFLOPS.

3.1.3 Replacements included Cray X-MP, released in 1982, with shared-memory parallel vector processor with better chaining support and multiple memory pipelines, and the Cray-2 liquid cooled computer immersed in a tank of Fluorinert (TM), released in 1985. In the early 1990s, the Fujitsu Numerical Wind Tunnel used 166 processors in a vector parallel archecture to reach 235.8 GFlop/s in November 1993.

3.0 Historical Examples cont..

Image of CDC6600 by Jitze Couperus - Flickr: Supercomputer - The Beginnings, CC BY 2.0, https://commons.wikimedia.org/w/index.php?curid=19382150

4.0 Contemporary Architectures

4.1 In the 1990s supercomputing moved towards massively numbers of processors; with distributed shared memory architectures, message passing, and clusters. Key stystems include the Intel Paragon, which could have up to 4000 processors, and the Intel ASCI Red supercomputer which reached up to 2.3796 TFLOPS.

4.2 The inclusion of GPGPUs in supercomputers in the 21st century was the result of shader and floating point support which assisted in vector and matrix mathematics.

4.3 The IBM Blue gene held the top position in 2004 (70.72 TFLOPS), 2005 (280.6 TFLOPS) and 2007 (478.2 TFLOPS). Notable because of its lower individual processor speed in favour of lower power consumption (more processors per rack), and attention to interconnect innovations (multidimensional torus approaches).

4.4 In recent years the dominance of China in the supercomputing world has become notable with the Tianhe-IA in 2005 (2.566 PFLOPS), the NUDT Tianhe-2 in 2013-2015 (33.86 PFLOPS), and the Sunway TaihuLight in 2016 (93 PFLOPS)

4.0 Contemporary Architectures cont...

4.5 The architecture of supercomputers has been largely the same for over 25 years; a cluster of commodity system units with a high-speed interconnect.

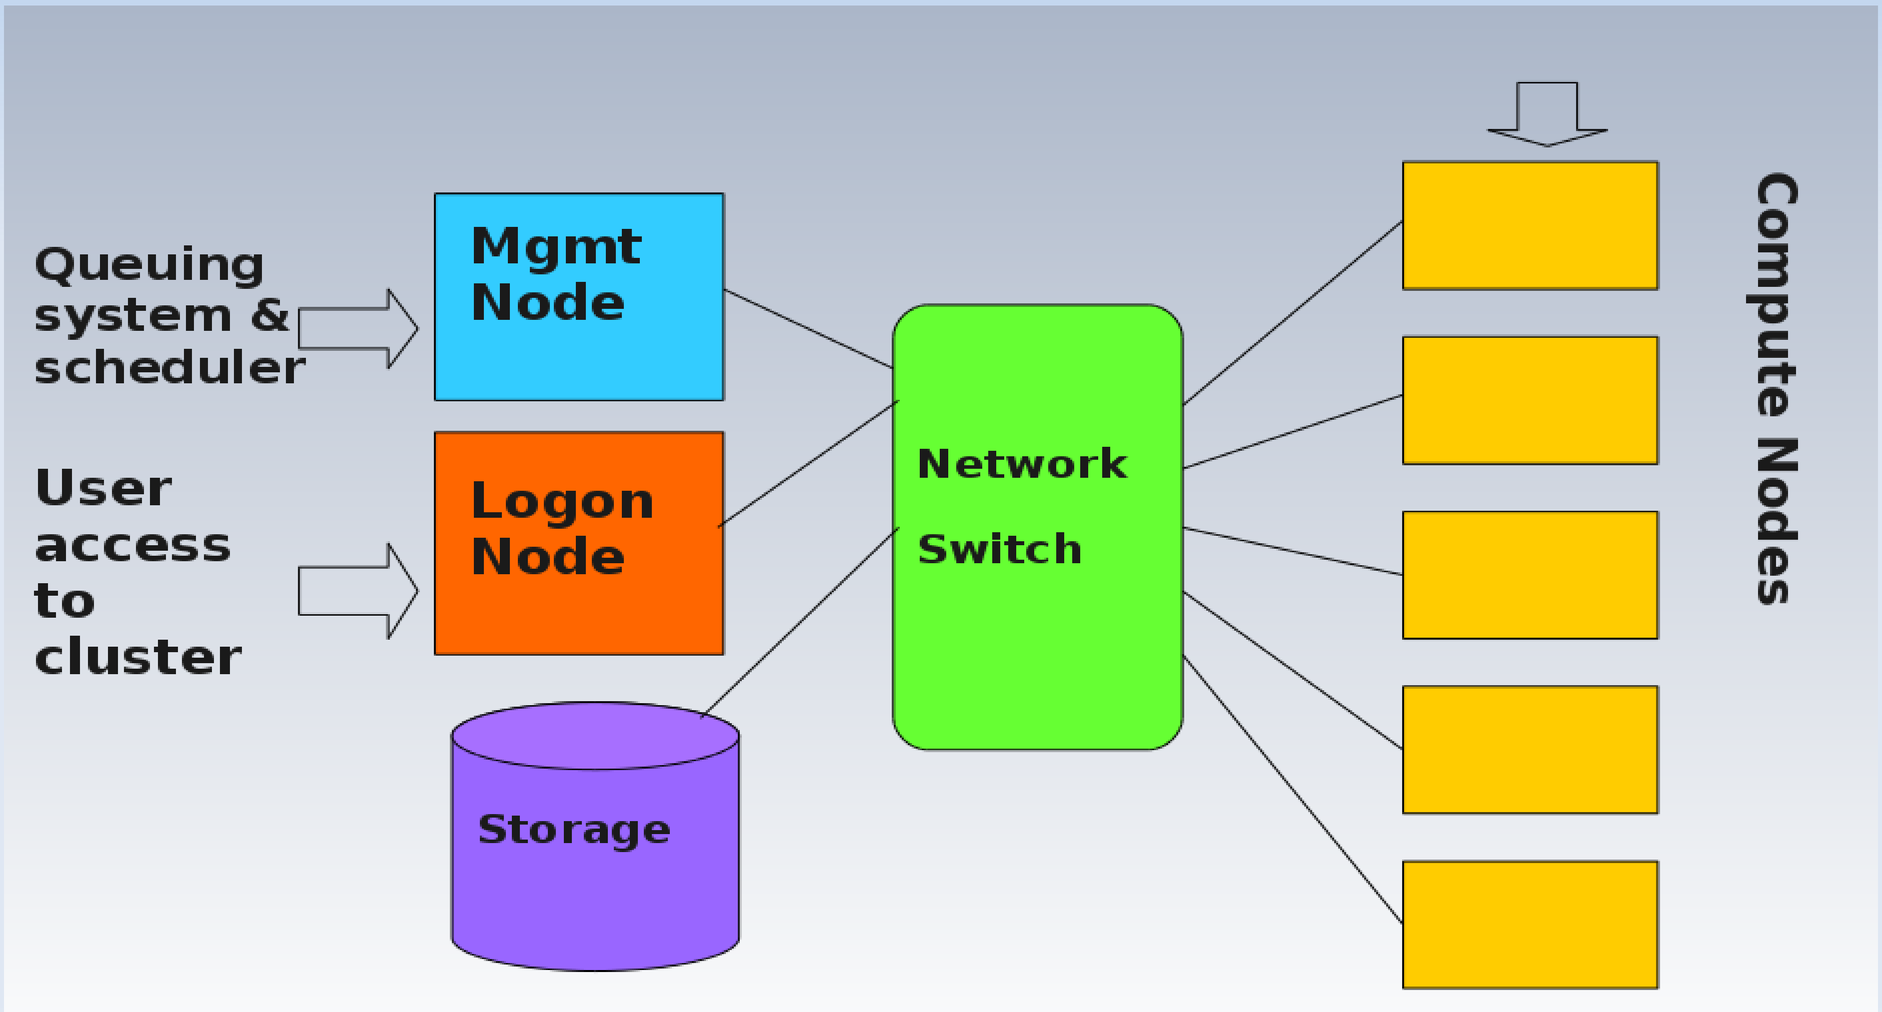

5.0 Operating Systems

5.1 Earliest supercomputers used proprietary operating systems that shipped with the hardware system (Chippewa Operating System for the CDC6000). Increasingly however these became more generic but still proprietary; initial leading example was UNICOS for the Cray-2.

5.2 By the mid-1990s the trend was towards the use of Unix-based systems, reaching a peak in the late 90s where 100% of the Top 500 were *nix-based (99.4% Unix, 0.4% BSD, 0.2% Linux), including Solaris, UNICOS, IRIX, AIX, HP-UX being notable.

5.3 This was quickly replaced by Linux systems; the most recent survey is 66.8% Linux, 0.6% AIX, and the remainder 32.6% "Others" - all being Linux plus customisations (i.e., 99.4% Linux in total).

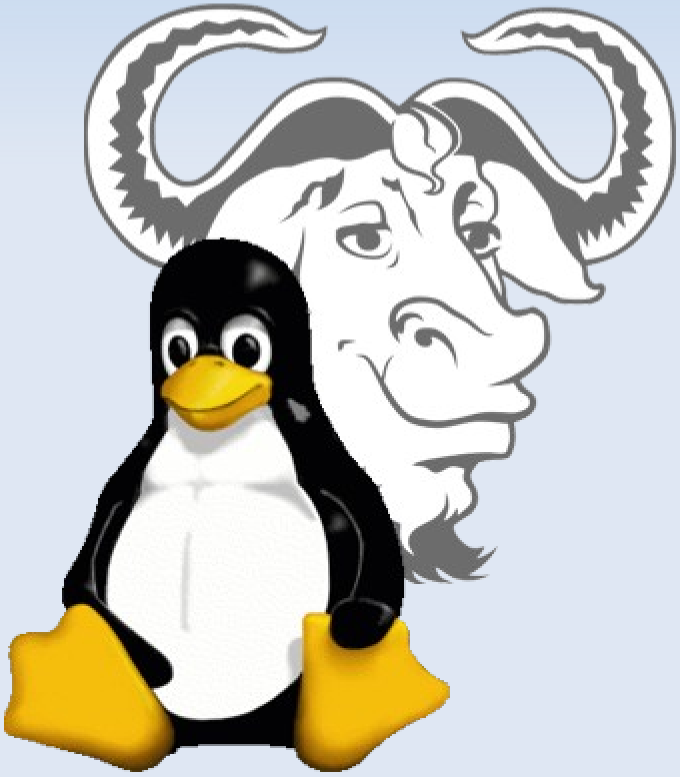

6.0 Parallel Programming

6.1 Historically, software has been written for serial computation (discrete instructions, sequential execution, single processor)

6.2 Parallel programming is the procedures used for simultaneous computation (discrete parts, concurrent executions, multiple processing).

6.3 Flynn's Taxonomy of Computer Architecture into streams of data and instructions (SISD, SIMD, MISD, MIMD). Later three architectures allow for parallel programming.

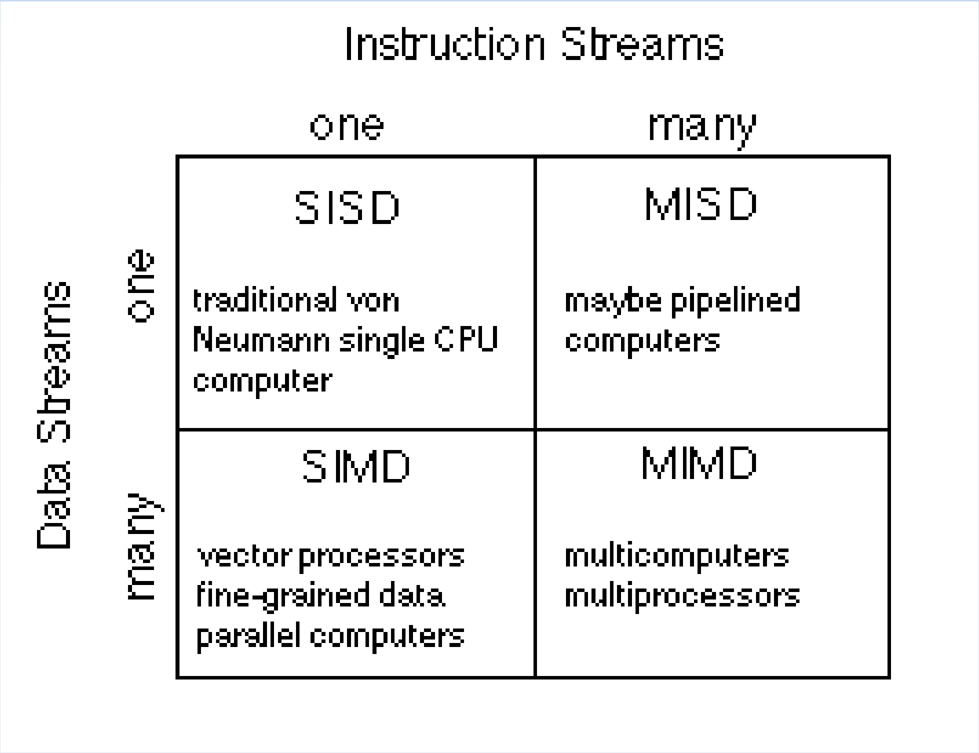

7.0 Data Parallelisation and Task Parallelisation

7.1 Data parallelisation occurs when the same thing is happening many times (e.g., multiple processors performing the same task. Effectively multiple instances of a sequential program (example Octave job array on Edward).

7.2 Task parallelisation occurs when the separate processors engage in different activities but also communicate with other giving data back and forth.

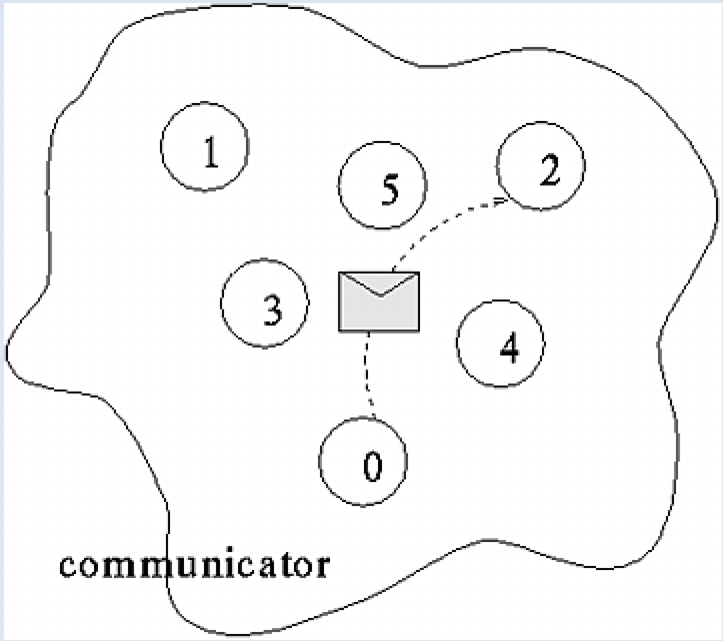

7.3 Usually implemented in software as message passing interface (MPI) for shared and distributed memory systems or as multithreaded fork-join model with multiprocessing (e.g., OpenMP) on shared memory systems (Example Ubiquitin protein simulation using NAMD on Edward).

8.0 Schedulers and Resource Managers

8.1 Supercomputers are invariably shared systems that operate in batch mode. Resource managers track what system resource are being used; this information is fed to a job scheduler which determines when a job will run (depening on the policies employed).

8.2 In early supercomputers these were incorporated into the operating system (Livermore Time Sharing System, Cray Time Sharing System).
8.3 The Portable Batch System was developed by NASA in the early 90s; and released as OpenPBS in 1998. This has two major forks, TORQUE by Adaptive Computing, and PBSPro by Altair. The Oracle Grid Engine (previously Sun Grid Engine) comes form a different code base.

8.4 SLURM (Simple Linux Utility for Resource Management) is a relatively recent initiative from at Lawrence Livermore National Laboratory and others that is now used on 60% of the Top 500 systems.

8.0 Schedulers and Resource Managers cont..

9.0 Future Issues

9.1 After twenty years of emphasis on graphic-user interfaces on the user level, and a more concern on performance rather than usage there is a substantial training gap.

9.2 Data parallel jobs are an increasing need proportionally and relatively easier to program that task parallel jobs. Cloud technologies are a flexible solution to such issues but can have a higher total operations cost, depending on packaging; HPC-Cloud hybrids provide flexibility and performance.

9.3 Open-source has provided a foundation for supercomputing; competitive relations in the development of open-source plus proprietary extensions (in operating system, or network technologies) provides monopolistic advantages for the individual systems but retards aggregate improvement.

9.4 Developments such as massively multicore processors (1,000+) provide enormous computing power but with concurrency issues between the cores and memory (e.g., The Angstrom Project, SGI Altix).

9.5 Extension of grid-like technologies to provide an extended continuum between tightly coupled and loosely coupled clustering (e.g., Plan9 OS)


THANKS FOR WATCHING
& LISTENING PATIENTLY
makeameme.org