Computational and infrastructure support for the biosciences David Abramson

Director, Research Computing Centre Professor of Computer Science University of Queensland



RCC Rationale

Living on the exponential



The future ain't what it used to be









Source: Ed Lazowska

Performance of HPC over the Last 24.5 Years from the Top500



RESEARCH COMPUTING CENTRE

Performance of HPC over the Last 24.5 Years from the Top500



RESEARCH COMPUTING CENTRE

The internet today





THE ARPA NETWORK DEC 1969 4 NODES

2900767	2100	LONDED DP. PROGRAM EDIZ BED BARKER	SK
	22:30	Talked to SRC	este
		thest to thest Lefter inp. program	(sle
		a host dead message	



Memory and Sensors



- > 1972: 1MB, \$1,000,000
- > 1982: 1MB, \$60,000
- > 2005: 1MB, \$0.40 (\$400/GB)
- > 2007: 1MB, \$0.15 (\$145/GB)
- > 2009: 1MB, \$0.015 (\$16/GB)

> Disk capacity,





- > Moore's Law drives sensors as well as processing and memory
 - > LSST will have a 3.2 Gigapixel camera



Source: Ed Lazowska





Science is changing



Emergence of New Problem Solving Methodologies



RESEARCH COMPUTING CENTRE

Computing is increasingly important Recomputing

- > Systems Biology and Health
 - > Human Genome Project, Protein function, Virtual Physiological Human, Blue Brain, ...
- > Engineering
 - > Aerospace, civilian, automotive, domestic, ...
- > Environment
 - > Climate, weather, pollution,
- > Chemistry
 - > Drug design, novel pathways, ...
- > Physics
 - > Particle Physics, Xray treatment, Astrophysics,
- > Business
 - > Manpower planning, Logistics, Resource allocation



Computing is increasingly important



>Systems Biology and Health

> Human Genome Project, Protein function, Virtual Physiological Human, Blue Brain, ...





The Research Computing Centre





Core Technologies

>High Performance Computing
>Data Management
>Scientific Visualization
>Cloud Computing
>Scientific Workflows





SGI LiveArc⁻⁻











UQ HPC Strategy





UQ HPC strategy

- >Persistent services
 - > QRIScloud
- >Very large tightly coupled parallel jobs
 - > support NCG funded research;
 - > support emergent research or require local infrastructure;
- >Loosely-coupled high throughput jobs;
- >Large memory and high input-output (IO) jobs; and
- >Accelerated computing platforms





Persistent Services



A private cloud offering compute, data storage and data collection access to Queensland research and its translation into real-world application









Domain Decomposition in Bio RESEARCH COMPUTING CENTRE Engineering Master Old array

worker 4

worker 1 worker 2 worker 3

rcc.uq.edu.au

THE UNIVERSITY OF QUEENSLAND

Very large parallel jobs



>NCG funded research

>National Merit Allocation Scheme (NCMAS

>NCI

> Raijin: 84,000 cores

> UQ guaranteed 10M additional core hours per yea

> Allocated by RCC and QCIF

>Pawsey

> Magnus: 35,712 cores







Tightly coupled parallel jobs

>Tinaroo

- >Smaller parallel jobs than NCI
- >More dynamic workload than NCI
 - > Access mechanisms are optimised for rapid account creation, which serves post graduate and early career researchers well.
- >Coupled to a significant data holdings
- >~6000 cores
 - >50M core hours per year



RESEARCH

COMPUTING









Loosely-coupled high throughput jobs



>QRIScloud

>National Research Cloud (NeCTAR)

>Research Data Storage Infrastructure (RDSI)

>Awoonga

>KVM based virtualisation

>~ 9M core hours per year





Massive Memory

Put lots of memory on each node

 What is the optimal size?

 Distributed Memory

 Message passing?

 CC-NUMA architecture

 Paying for cache coherence

 Distributed virtual memory

 No free lunch - locality







Massive Memory









Large memory and high input-output (IO) jobs



>FlashLite

- > High throughput solid state disk
- > Large amounts of main memory
- > Software shared memory



Xeon Processor E5-2600 v3 Overview



Large memory and high input-output (IO) jobs



- >~ 70 compute nodes (~1600 cores)
 - > Dual socket Intel E5-2680v3 2.5GHz (Haswell)
 - > 512 GB DDR-2
 - > 4.8 TB NVMe SSD
- >ScaleMP vSMP virtual shared memory
 - > Up to 16TB RAM





Xeon Processor E5-2600 v3 Overview



Accelerated Computing Platforms

- >Distributed Memory
 - >Message passing
 - >MPI
- >Domain decomposition
- >General Purpose Graphics Processing Unit
- >Highly parallel

>Special purpose components for deep learning



First NVIDIA Volta based production system in Asia Pacific, 3rd in the world

A dual parallel filesystem approach

BeeGFS + GPFS, delivering 180GB/sec and

25m IOPS of sustained performance in

RDMA connected nVME flash

rcc.uq.edu.au

workload manager

Wiener

~4.3 petaFLOPS of accelerated Compute capability; UQ's fastest HPC facility

189,440 CUDA cores: 23,680 dedicated hardware tensor





OpenHPC Reference Site



100Gbit/sec EDR non-blocking fabric; 1.9Tbit of combined signaling











Deconvolution

Accelerated EM



Amber protein structure seek acceleration

Deep Learning & Al





UQ Data Strategy









MeDiCI as a parallel file system





DDN SFA12KXE



Parallel file system



Accessing long term collections











Conclusions

>Living on the exponential>Science is changing

>Research Computing Centre

>UQ HPC Strategy

>MeDiCl Data fabric







Subscribe to our mailing list

	* indicates required
mail Address	~~~
	*
irst Name	
.ast Name	
SUBSCRIBE	

Level 5, Axon Building (Bldg 47)

The University of Queensland St Lucia QLD 4072, Australia Office Hours 9am - 5pm, Monday to Friday General enquiries rcc-admin@uq.edu.au Phone: +61 7 336 58350 Technical support rcc-support@uq.edu.au

