NATURAL ENVIRONMENT RESEARCH COUNCIL

EPSRC

Introduction to ARCHER

Outline of course

epcc

THE UNIVERSITY OF EDINBURGH

Reusing this material

archer

epcc

THE UNIVERSITY OF EDINBURGH

Course Parameters

- Pre-requisites
  - Familiarity with parallel programming is assumed
- Hands-on practicals form an integral part of the course.
  - Example codes will be supplied, but we hope that many users will work on their own applications in the practical sessions.

Learning Outcomes

- On completion of this course attendees should be able to:
  - Understand the ARCHER hardware environment.
  - Compile and run parallel programs on ARCHER.
  - Exploit ARCHER-specific features of the MPI library in their codes.
  - Port applications to ARCHER.

Structure

- Overview of ARCHER service,
- ARCHER Hardware: arrangement of compute ***nodes***
- ***Sharpen*** practical: check you can log on, compile and run
- Compiling programs and basic job submission/execution
- Explore compilation options using sharpen code
- Node architecture: memory, CPUs, cores, hyperthreads and job execution options
- Explore node architecture using sharpen code
- Overview of MPI library on ARCHER
- Interconnect performance with Intel MPI Benchmark ***IMPI***
- Cray tools available on ARCHER
- Use tools on sharpen, more realistic ***CFD*** or own application

NATURAL ENVIRONMENT RESEARCH COUNCIL

EPSRC

ARCHER Service

Overview and Introduction

epcc

THE UNIVERSITY OF EDINBURGH

ARCHER in a nutshell

- UK National Supercomputing Service
- Cray XC30 Hardware
  - Nodes based on 2×Intel Ivy Bridge 12-core processors
  - 64GB (or 128GB) memory per node
  - 3008 nodes in total (72162 cores)
  - Linked by Cray Aries interconnect (dragonfly topology)
- Cray Application Development Environment
  - Cray, Intel, GNU Compilers
  - Cray Parallel Libraries (MPI, SHMEM, PGAS)
  - DDT Debugger, Cray Performance Analysis Tools

Compared to HECToR (Hardware)

Feature	HECToR	ARCHER
Processors	AMD Interlagos 2.3GHz	Intel Ivy Bridge 2.7GHz
Cores per node	32 (4×8-core NUMA)	24 (2×12-core NUMA)
Memory per node	32 GB (1 GB/core)	64GB (2.66 GB/core) 128GB (5.33 GB/core)
Nodes	2816 (90,112 cores)	3008 (72,192 cores)
Interconnect	Cray Gemini	Cray Aries
Topology	3D Torus	Dragonfly
Post-processing Nodes	(None)	2 Nodes: 48 core SandyBridge 1TB Memory

Compared to HECToR (Software)

- The software environment is very similar
  - Intel Composer replaces PGI Compiler
  - DDT replaces Totalview for debugging
  - Intel MKL replaces ACML library

- If you have your code running on HECToR then it should not be a problem getting it running on ARCHER

Compared to HECToR (PBS)

- ARCHER also uses PBSPro job submission system…
- …but the syntax has changed from HECToR.
- Now ask for number of *nodes* rather than cores.
- e.g. To ask for 64 nodes (64×24 = 1536 cores):

```
#PBS -l select=64
aprun -n 1536 my_app.x
```

- Serial jobs specifier

```
#PBS -l select=serial=true
#PBS -l select=serial=true:ncpus=4
```

Storage

- /home – NFS, not accessible on compute nodes
  - For source code and critical files
  - Backed up
  - > 200 TB total
- /work – Lustre, accessible on all nodes
  - High-performance parallel filesystem
  - Not backed-up
  - > 4PB total
- RDF – GPFS, not accessible on compute nodes
  - Long term data storage

Getting access to ARCHER

- Standard research grant
  - Request Technical Assessment using form on ARCHER website
  - Submit completed TA with notional cost in J-eS
  - Apply for time for maximum of 2 years
- ARCHER Resource Allocation Panel (RAP)
  - Request Technical Assessment using form on ARCHER website
  - Submit completed TA with RAP form
  - Application for computer time only
- Instant Access – Pump-Priming Time
  - Request Technical Assessment using form on ARCHER website
  - Submit completed TA with 2 page description of work
  - Office decision on application

ARCHER Partners

- EPSRC
  - Managing partner on behalf of RCUK
- Cray
  - Hardware provider
- EPCC
  - Service Provision (SP) – Systems, Helpdesk, Administration, Overall Management (also input from STFC Daresbury Laboratory)
  - Computational Science and Engineering (CSE) – In-depth support, training, embedded CSE (eCSE) funding calls
  - Hosting of hardware – datacentre, infrastructure, etc.

What is it used for?

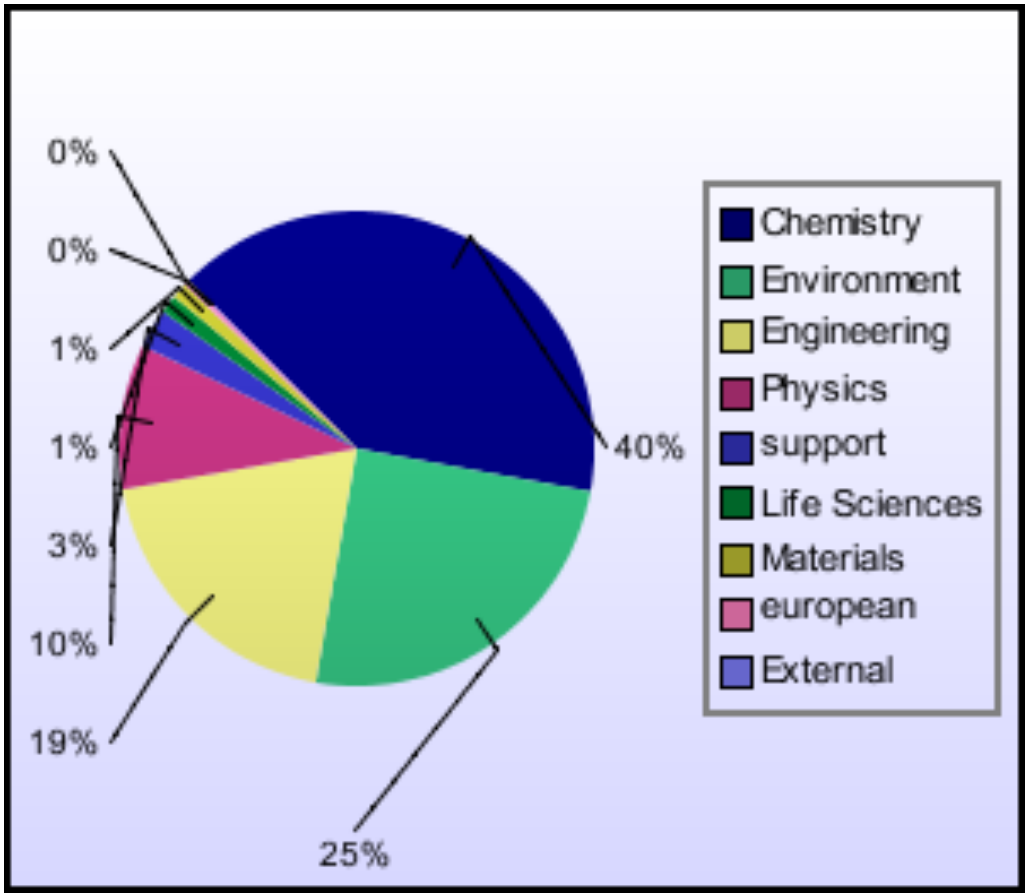

archer

epcc

THE UNIVERSITY OF EDINBURGH

Simulation software

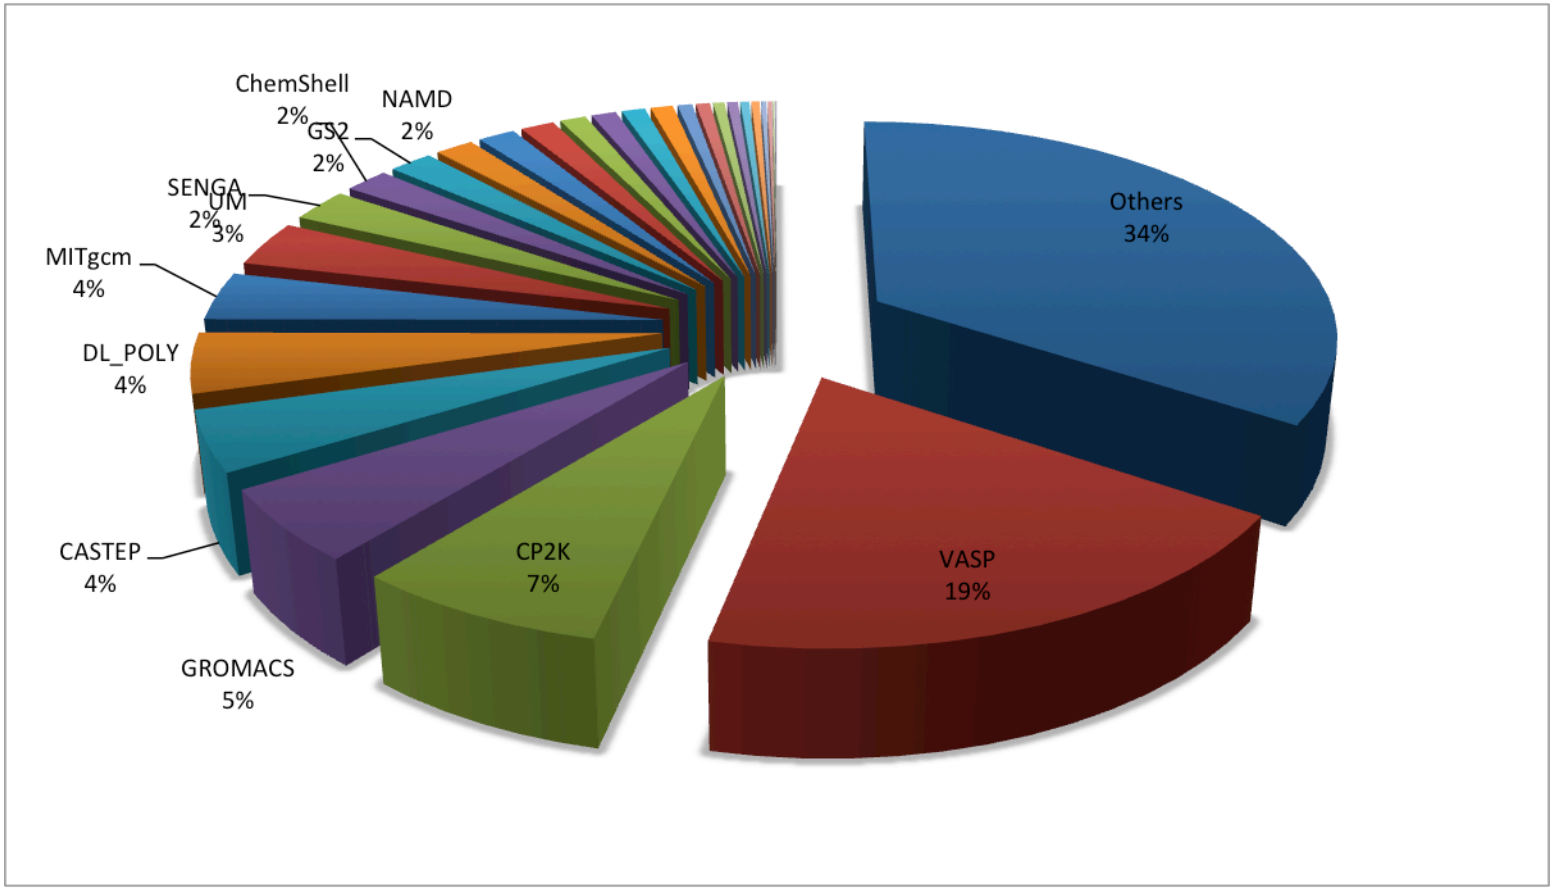

Early Usage

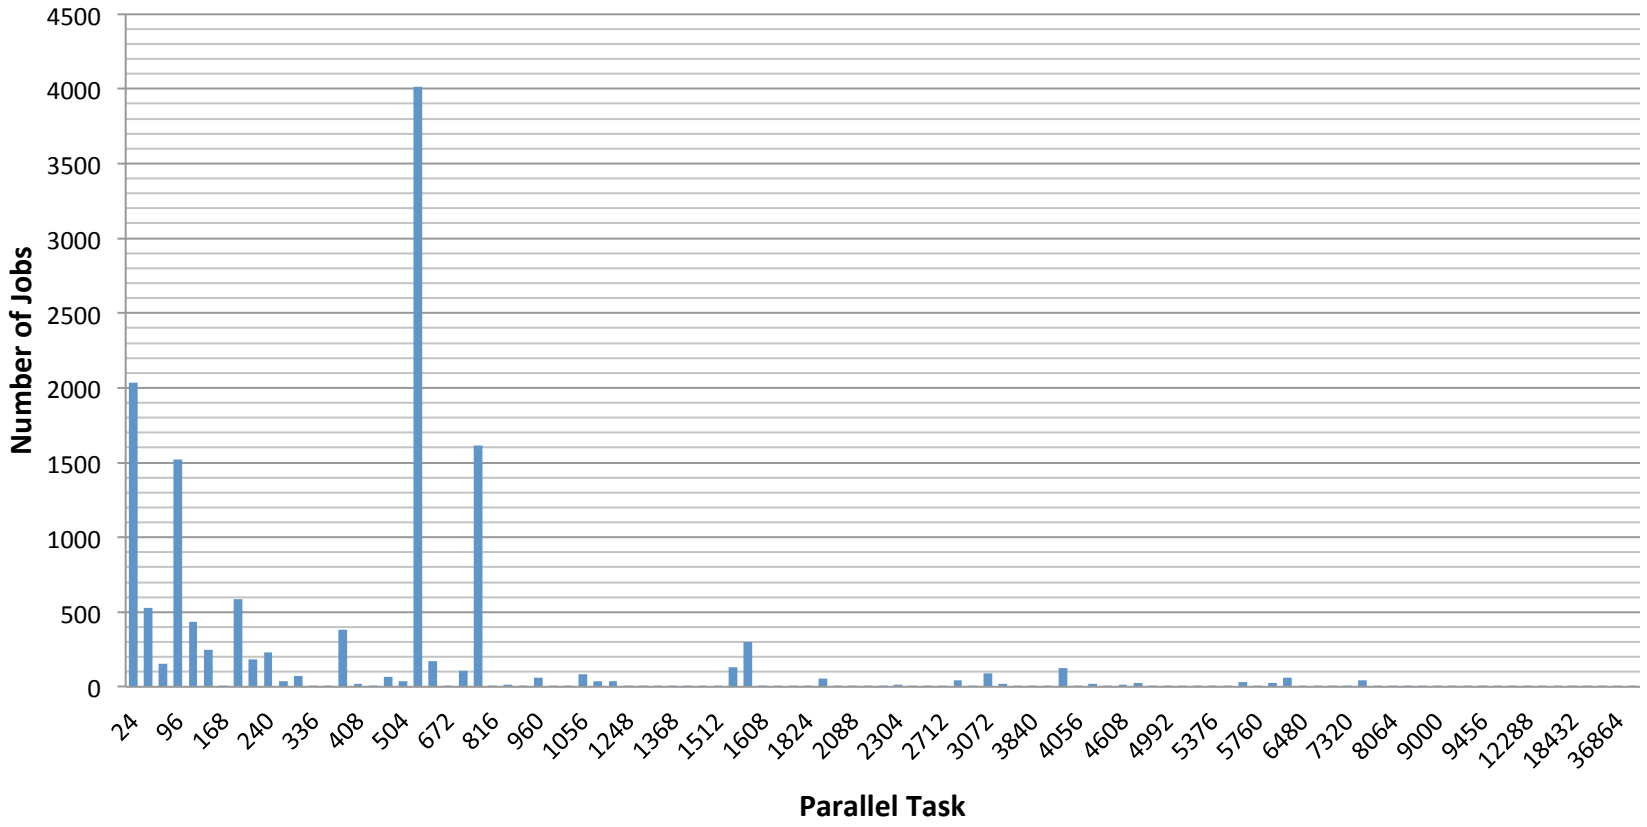

archer

epcc

THE UNIVERSITY OF EDINBURGH

Early Usage

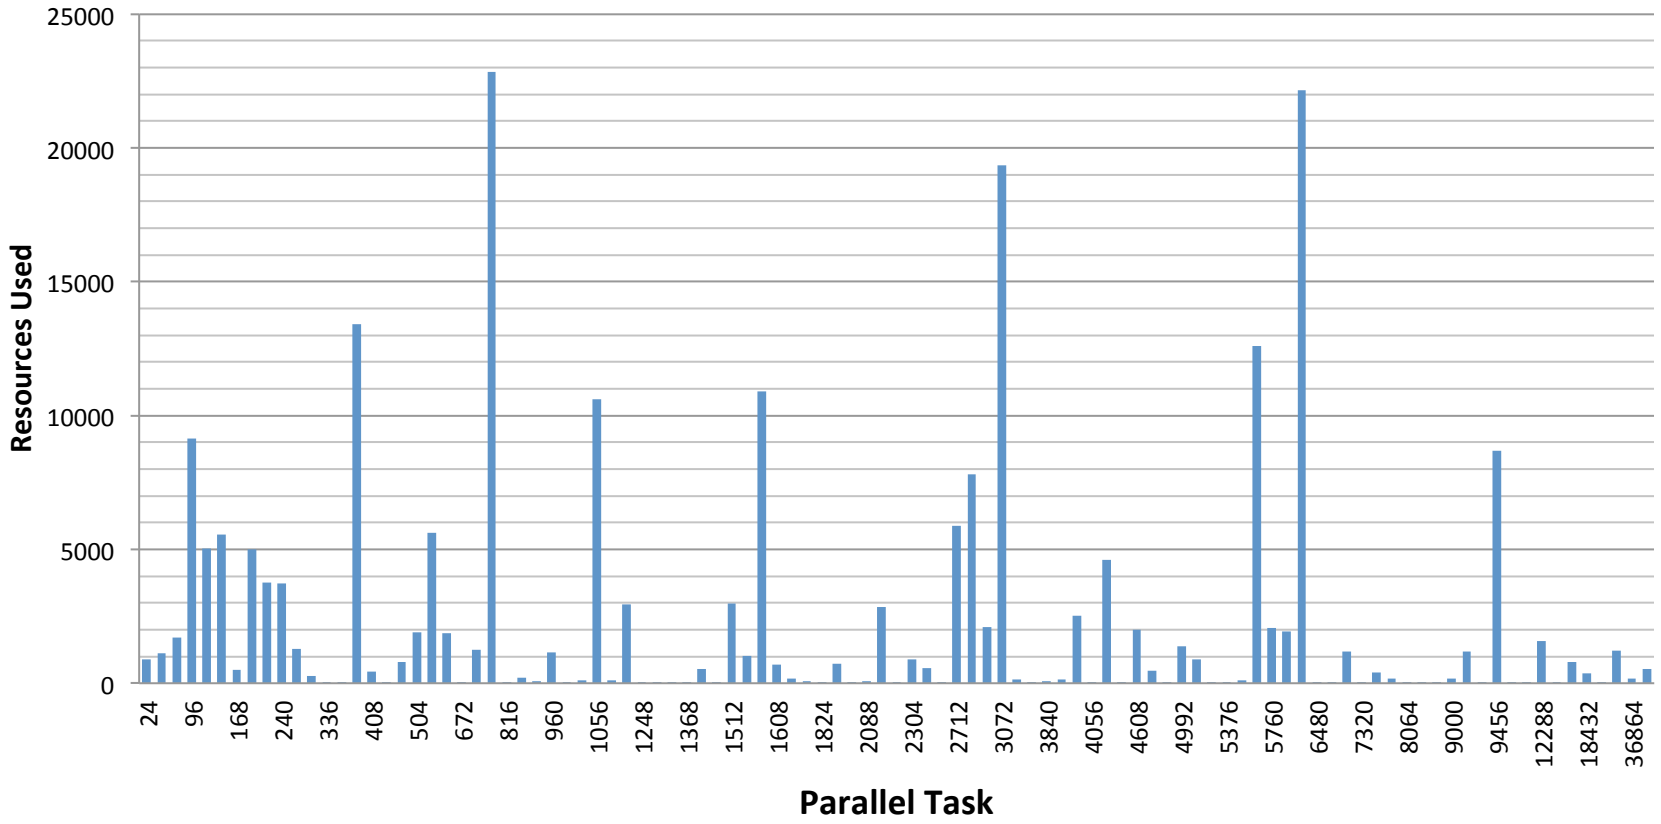

archer

epcc

THE UNIVERSITY OF EDINBURGH

Summary

- ARCHER is a Cray XC30
  - It uses standard Intel processors
    - 2 processors per node, 24 cores per node
    - 64 GB memory of the majority of nodes
    - Nodes similar to many HPC systems
  - Cray ARIES switch
    - High performance, optimised for large jobs
    - Standard usage but can get very good performance
  - Large storage and high performance filesystem
    - 4 PB high performance filesystem
    - 200 TB home space
  - Intel, GNU, and Cray compilers
    - Lots of standard scientific packages, libraries, and software installed