MANCHESTER 1824

The University of Manchester

Data Mining on (Cray) Supercomputers

Andrew Jones, Mike Pettipher, Firat Tekiner

Andrew Jones
Manchester Computing

19th May 2005
CUG Summit 2005
Albuquerque, NM

MANCHESTER 1824

The University of Manchester

What are these?


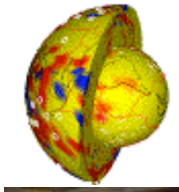




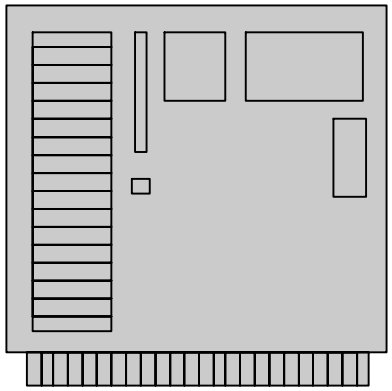
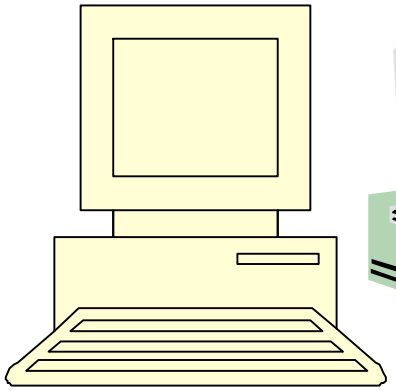


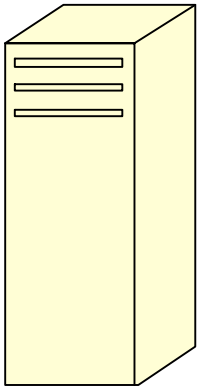
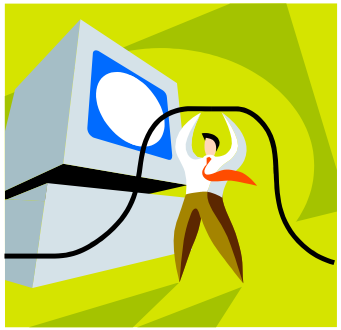
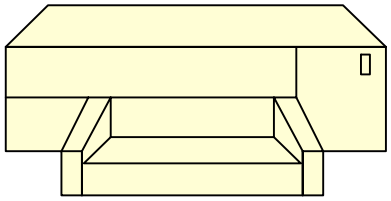

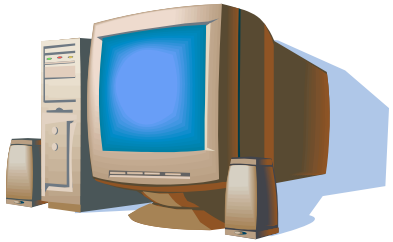

MANCHESTER 1824

The University of Manchester

What are these?


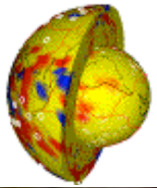



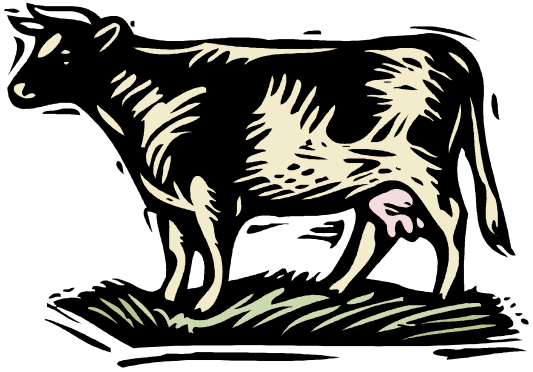

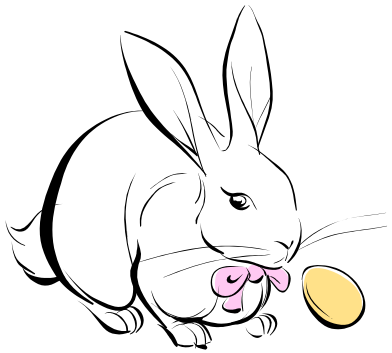

MANCHESTER 1824

The University of Manchester

What are all these data telling us?


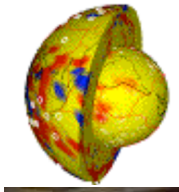





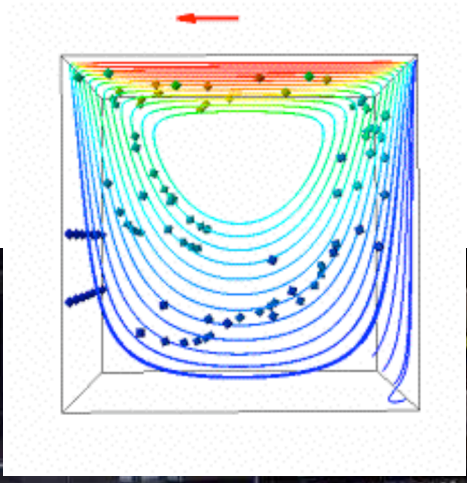
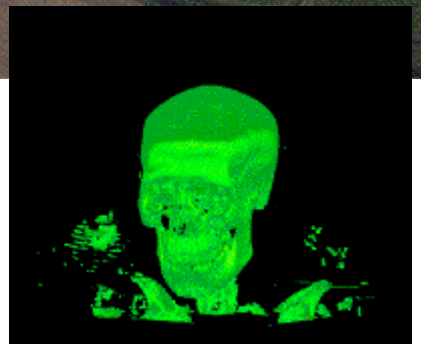
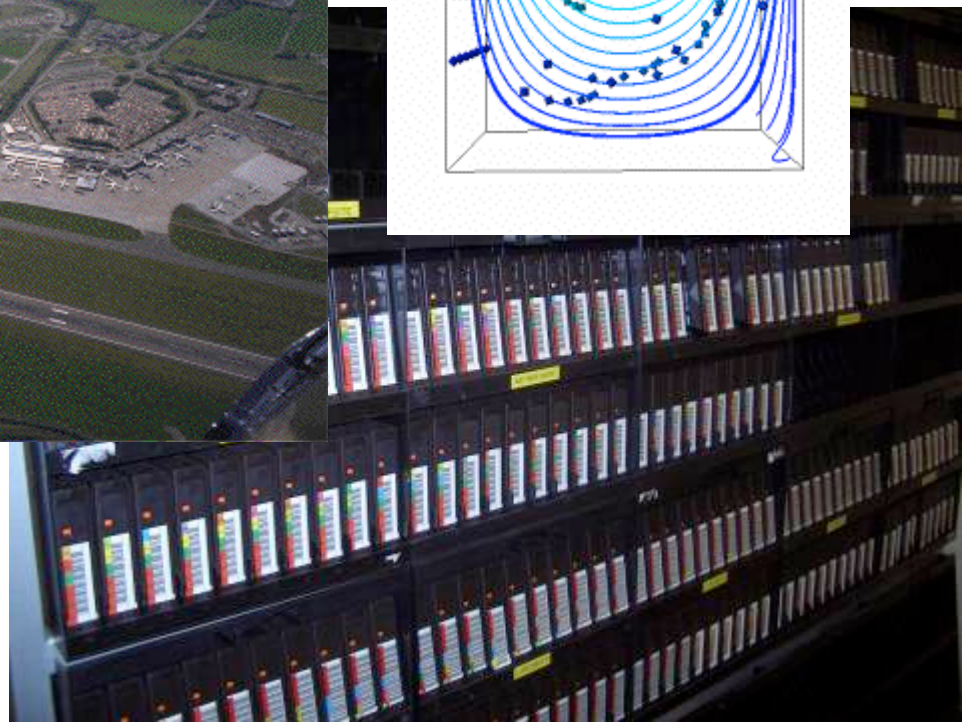

MANCHESTER 1824

The University of Manchester

What are all these data telling us?


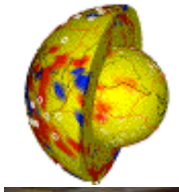



- In this paper, we report on the introduction of a large scale datamining code to the supercomputing environment, and to Cray supercomputers in particular. The codes are taken from a community who are not traditional scientific computing programmers nor use parallel computing and whose codes are subject to rapid development by researchers unfamiliar with the requirements of high end scientific computing. Thus, the challenge is to port the PC written C++ codes to specialist computers such as the X1 and XD1, and extract the enhanced performance, whilst retaining the common code base with the datamining researchers. This investigation is thus an evaluation of the programming environment on the Cray supercomputers as applied to this problem, and of the significant data management issues involved.

MANCHESTER 1824

The University of Manchester

Abstract in bits …


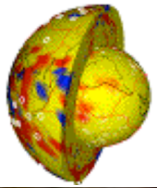



- Introduction of a large scale datamining code to the supercomputing environment
  - and to Cray supercomputers in particular
- Codes
  - from a community who are not traditional scientific computing programmers nor use parallel computing
  - are subject to rapid development by researchers unfamiliar with the requirements of high end scientific computing
- Challenge is to
  - port the PC written C++ codes to specialist computers such as the X1 and XD1 (FPGA?), XT3, MTA?
  - extract the enhanced performance
  - whilst retaining the common code base with the datamining researchers
- evaluation of
  - the programming environment as applied to this problem
  - the significant data management issues involved

MANCHESTER 1824

The University of Manchester

In computer terms …


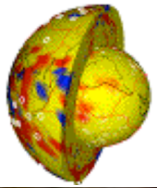



- Take a 50TB collection of data
  - That changes (grows and mutates) daily
  - Residing in multiple geographical and logical locations
    - Hundreds of miles?
- ***Manage*** this data into/though memory
  - Near real time?
- Perform multiple data mining experiments on these data
  - Using a portable, scalable, C++ code
  - In a reasonable timescale – a few minutes?
- Make it all easy to use …

MANCHESTER 1824

The University of Manchester

Computing at Manchester


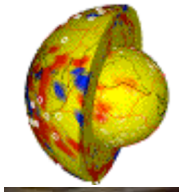



- One of the top research-led universities in the UK
  - Among the world's top 25 universities by 2015
- World leading research pedigree in computing
  - World's first stored program computer in 1948
  - invention of virtual memory
  - first dataflow machine
  - ARM processor
- World class service pedigree in computing
  - several national services
  - ~350 staff
  - ~50 focused on high end – HPC, grid, visualization etc
  - Service, Research and Teaching
  - 23,000+ sq.ft managed machine room space

MANCHESTER 1824

The University of Manchester

Supercomputing at Manchester


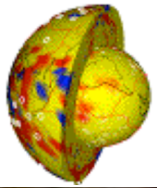



- dominated by the national HPC service (CSAR)
  - 512 processor SGI Altix 3700
  - 512 processor SGI Origin 3800
- and by research group focused facilities, e.g.
  - 146 processor IBM SP
  - 204 processor IBM Beowulf cluster
- no significant central resources
  - 32 processor SGI Onyx300 with 6 IR3 graphics pipes
  - 24 processor Sun E6500 with 30GB shared memory
- "compute-enabled capability science" *not just* "capability compute"
  - make best use of the national HPC resources for capability compute
  - provide central expertise and support, rather than central machines

Supercomputing at Manchester


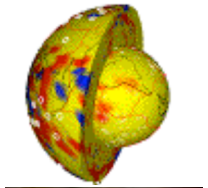



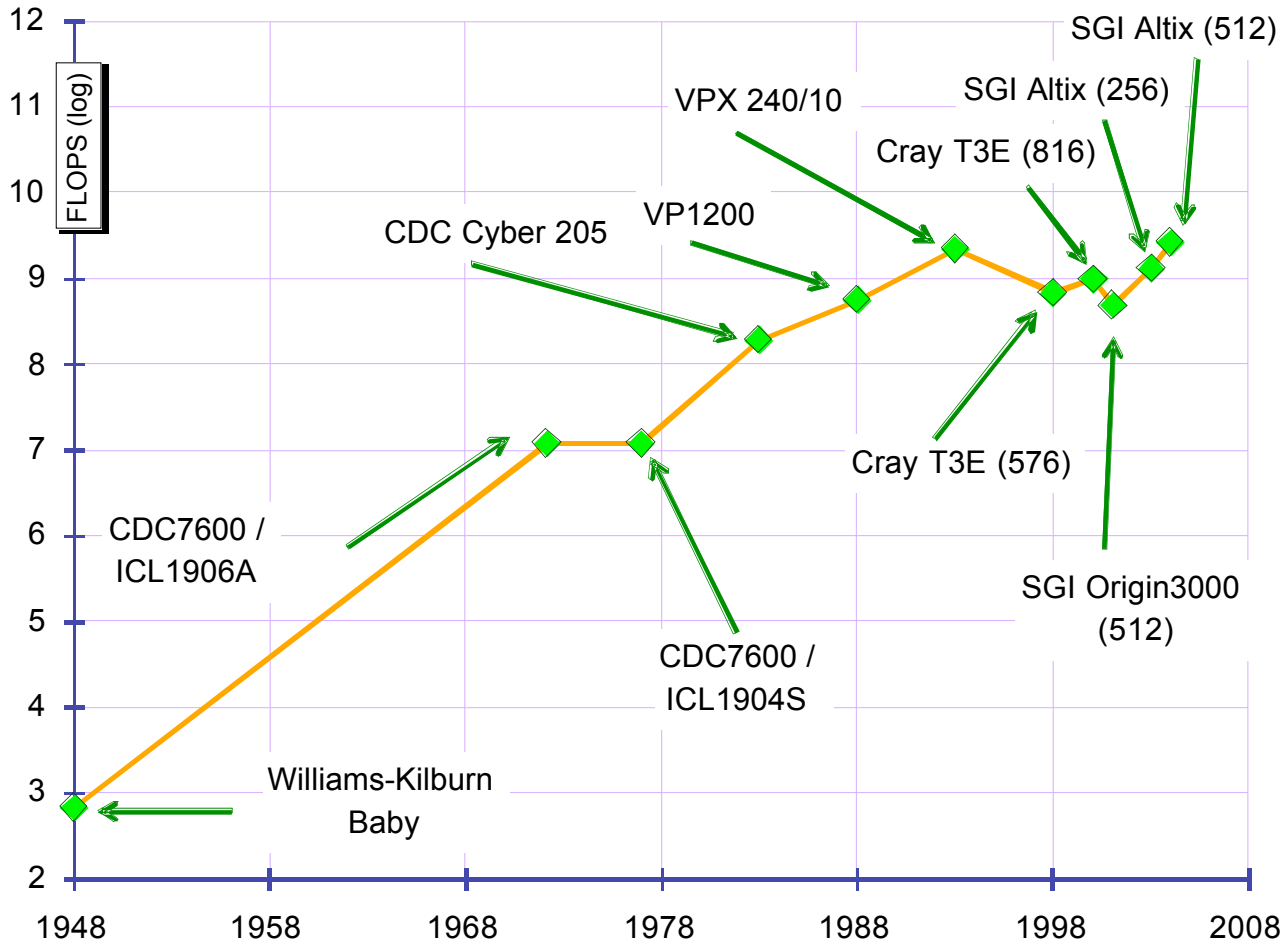

Data at Manchester


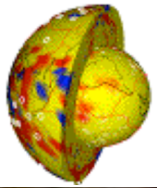



- UK National Grid Service Data Node
- National Centre for e- Social Science (NCeSS);
- National Centre for Text Mining;
- MIMAS national datasets service;
- Supercomputer Data Mining project
- MYGrid
- CSAR national HPC service
- wide data mining research portfolio in 5* computer science dept and other science dept
- one of the four centres of the SuperJANET4

MANCHESTER 1824

The University of Manchester

Vision


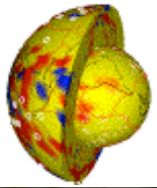



- To co-ordinate and build on the considerable existing expertise, research and services at Manchester to become one of the worlds leading centres for “data”:
  - Data mining
  - Data management
  - Data services
  - HPC + data integration
  - “grid”
  - etc

MANCHESTER 1824

The University of Manchester

Notes


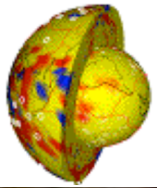



- This is an ambitious goal
- Highly fragmented starting point
  - Including politics!
- Have indication of institutional support

MANCHESTER 1824

The University of Manchester

Some other HPC research at Manchester


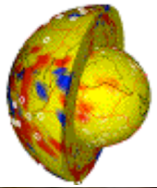



- exploring the role of new HPC technologies such as FPGAs
  - expect to deploy FPGA based HPC system(s) in the next few months
- virtual reality in engineering with ‘real-time’ finite element analysis software
- internationally successful RealityGrid/TeraGyroid project
  - computational steering, inter-system checkpointing and data migration, and remote visualization

MANCHESTER 1824

The University of Manchester

Some near futures


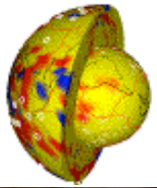



- Current/coming procurements
  - a cluster for high energy physics
  - central HPC/HPV refresh (likely to include FPGA)
  - NW-GRID: $10M total, $3M hardware budget
- ~12 months: a significant university HPC related project
  - Step change in exploitation of HPC
  - Specific novel HPC focus
  - Multi-million $
- HECToR – next UK national HPC service (2007) – $200M?
  - Starting at ~75TF? in 2007, working towards ~400TF? in 2011

MANCHESTER 1824

The University of Manchester

Why data mining?


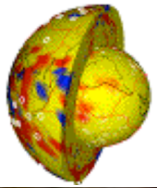



- Huge amounts of data are available to us
  - Unstructured information (text)
  - Semi-structured information (XML+text)
  - Structured information (databases)
  - Raw data
- Created, captured, sought, pushed
- Too much data is worse than not enough
- Data alone is not knowledge

MANCHESTER 1824

The University of Manchester

Data (text) mining: a definition

- Non trivial extraction of implicit, previously unknown, and potentially useful information from mass of (textual) data
- Exploration and analysis of (textual, natural-language) data by automatic and semi automatic means to discover new knowledge and update existing knowledge

MANCHESTER 1824

The University of Manchester

New knowledge

- What is “previously unknown” information, or “new knowledge”?
- Information that not even the creators/authors knew
- Rediscover information that the creators/author encoded in the data

MANCHESTER 1824

The University of Manchester

Step 1: find


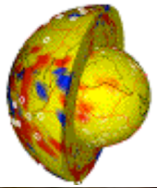



- Information retrieval
- Search
- Gathers, selects, filters data/documents that may prove useful
- Finds what is known
- Google search
- Get all drivers who have had accidents in last 3 years

MANCHESTER 1824

The University of Manchester

Step 2: process


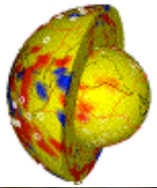



- Information extraction
- Sort & analyse
- Extracts facts & events of interest to user
- Finds only what we are looking for
- Reading the results of a Google search
- Put these drivers into categories – age, city, mileage, etc

MANCHESTER 1824

The University of Manchester

Step 3: mine


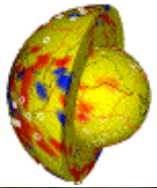



- Data Mining
- Discovers unsuspected associations
- Combines & links facts and events
- Discovers new knowledge, finds new associations
- ?
- What makes a driver a bad insurance risk?

MANCHESTER 1824

The University of Manchester

Beware


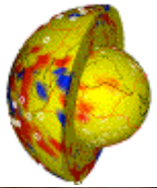



- Information is partial
  - Because of form
  - Because of capture process
  - Because some information is kept confidential
- Information is unreliable
  - Just think of the Web
  - And may contain mistakes
- Drawing conclusions based on statistics is always gambling

Users


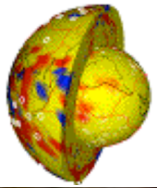



- Data mining = searching + meaning + “hidden” info

- Banks
- Insurance companies
- Supermarkets
- Marketing data
- Police, law enforcement, intelligence etc
- Scientists
- Social scientists
- Web users? (super-Google??)

MANCHESTER 1824

The University of Manchester

Why so much focus on Biosciences?


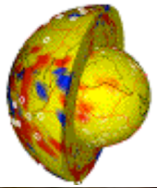



- Journals/databases expanding very fast
  - E.g. MEDLINE: 40,000 new references added per month
- Dynamic nature - new terms (genes, proteins, drugs, chemical compounds) constantly created
- Bio-databases, controlled vocabularies and bio-ontologies encode only small fraction of information
- Bio-informatics is a “hot topic”
  - a current buzz-word for funding proposals etc

MANCHESTER 1824

The University of Manchester

Why supercomputing?


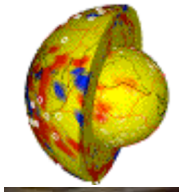



- Supercomputer = Big and/or Fast
- Many terabytes of data
  - IO, data management
- Multiple data sources
  - Geographical
  - Type
  - Owner etc
- Multiple users/queries, mining algorithms
- Time requirements

MANCHESTER 1824

The University of Manchester

Normal people


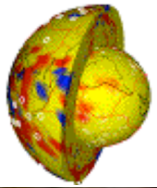



- Most data-miners (developers) are computer scientists
- Most data-miners (end-users) are not computer scientists
- Neither of these are usually supercomputer people
- All have different needs and motivations

MANCHESTER 1824

The University of Manchester

End Users


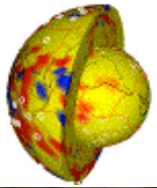



- Special case
  - I'm the only one who needs to do it like this
- In a hurry
  - Deadline in weeks
  - Only just realised I need your help
- One-off
  - It works in my use-case or hack, thus it is good
  - Simplifications, hackery, etc now = trouble later

MANCHESTER 1824

The University of Manchester

Computer Scientists


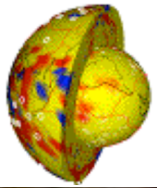



- Lets solve this other problem which isn't your problem but is fun and leads to interesting software
- It works for my toy test set that I synthesised
- Tell us what you want and then go away and we will build it for you
- You can't use it until its finished
- I can understand it, so just need to train you to be just like me!

MANCHESTER 1824

The University of Manchester

Supercomputer experts


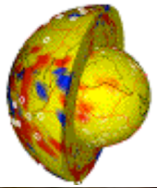



- We're at the high end you know - would you like to come up to our level and join us?
  - We have cryptic & mystic spells for you to learn
  - We promise extreme power and speed
  - But we have a wagon-load of caveats on that …
- What are you doing it that way for?
- Forget all those silly modern fashionable C++, Java, and object-orientated ideas that are easy skills to recruit for
  - use our specialist F77+MPI with customisations for each machine …

MANCHESTER 1824

The University of Manchester

Differing goals


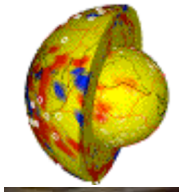



- Users
  - Need to get it done and use it! (simply)
- Computer scientists
  - Want to do it elegantly and all-encompassing
- Supercomputer experts
  - Want to do it bigger and faster

MANCHESTER 1824

The University of Manchester

Goals


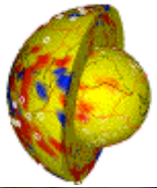



- Take HPC to data-miners, *not* bring data-miners to HPC
  - Our discipline must change and adapt to get best use by new communities
  - Can only expect minimum change on part of new user community
- HPC offers speed
  - Users want speed but ***only*** if its easy
- Our job to bridge gap

MANCHESTER 1824

The University of Manchester

Supercomputer Data Mining Project 1


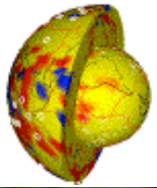



- School of Computer Science, University of the West of England
- School of Computing, University of East Anglia
- Manchester Computing, University of Manchester

- To produce a super computing data mining resource for the UK academic community which utilises machine learning methods and statistical algorithms

General Issues


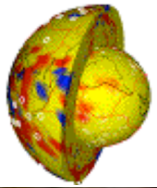



- Quantity of data
- Location
  - where and on what medium does the data reside?
  - given the data quantities, cost will be a major factor
- Speed of access
  - how quickly can the data be transferred to the computer which is going to analyse or extract the data of interest?
- Frequency of access
  - if the same data is to be accessed many times, keep local copies which can be accessed quickly, rather than returning to the original data

Programming Environment


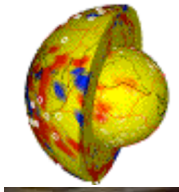



- Aim is to have a portable data mining tool
  - Distributed memory for portability
  - MPI will be the parallel programming paradigm
- Programming language is C++
  - Non-Standard Libraries will also be used (i.e. Boost Libraries)
  - Target compiler will be gcc
- Parallelisation at a high level
  - large number of independent tasks (data miners)
- Initially run many data miners with different parameters/attributes on TB of data

Details and issues - 1


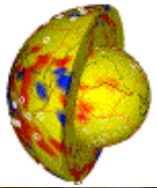



- How, when, where to transfer TBs of data from Norwich and Bristol to Manchester
- Format of data?
  - Flat files – text or binary
  - Structured database files
  - Single file or multiple files (one for each processor => <1/2GB each?)
- How to transfer huge amount of data residing on the disk to the memory?
  - MPI/IO? => binary
  - Must whole dataset fit into memory?

Details and issues - 2


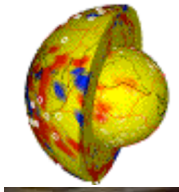



- How to copy/move data from learner to learner ?

  OR ALTERNATIVELY
- How to copy/move learner from processor to processor?
  - 'objects' have to be moved between processors:
    - Learners cannot be re-instantiated at new locations
    - Must serialise objects in order to transfer
- On shared memory systems could leave data and learners in place, and tolerate remote access to data. Is this worth considering as a simple alternative on such systems?

MANCHESTER 1824

The University of Manchester

The project in-summary


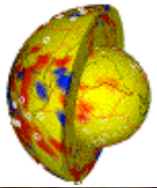



- The objectives are flexible:
  - we wish to provide a tool which is generic and portable, but are not obliged to. If the tool has restricted use (database must fit in memory), we have still met the requirements as in the funding proposal
- Numerous issues in data management and in the use of sophisticated C++ programs to be tackled
- Unclear what the best approach will be, particularly until we have good understanding of the codes
- Plenty of challenges and opportunities!

MANCHESTER 1824

The University of Manchester

Initial work


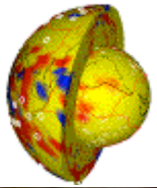



- Bulk of code not yet written
  - only an initial framework
  - Nearest neighbour classifier
    - “what is this”?
- Porting and compiling the base code
  - C++ MPI
  - Minimum changes allowed
  - BOOST C++ libraries are major component
  - Target standard is gcc compiler

MANCHESTER 1824

The University of Manchester

The classifier


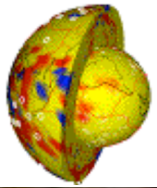



- read in all of the training data
- distribute it amongst the processors
- each processor learns on its data
- read in all of the test data
- loop over records
  - broadcast record to all nodes
  - classify
  - return k class/distance pairs to the master node
  - master then uses merge sort to identify the true k nearest neighbours, which are then used to make the classification

MANCHESTER 1824

The University of Manchester

Current characteristics


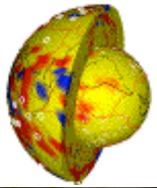



- Repeated cycles of
  - Broadcast small data messages (~100 bytes)
    - Latency would seem to be important
  - Compute/compare
    - Memory performance
    - Few flops?
  - Send results back to master node
- Startup I/O
- Sorting/comparing

MANCHESTER 1824

The University of Manchester

Testing work


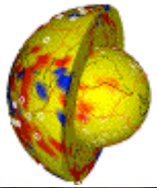



- (C++/MPI) Compilers
- Machines
  - SGI Altix
  - SGI Origin
  - Cray XD1
  - Cray X1
  - Cray XT3
  - [Cray MTA]
  - IBM p690+
- Do some performance tests on resultant code(s)
  - Investigate some performance metrics??
- Investigate some “HPC” options

MANCHESTER 1824

The University of Manchester

Thanks


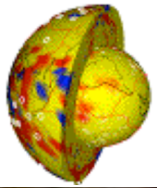



- Special thanks to the folks at PSC for giving me access and support at short notice (< 24hrs) to their XT3
- Also to Cray for access to evaluation platforms
- And to HPCx and CSAR in the UK

MANCHESTER 1824

The University of Manchester

Codes & compilers


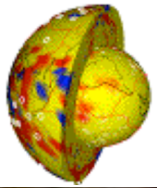



- C++/MPI “Hello World”
- A complex (12k lines) serial C++
- The C++/MPI framework classifier code
  - Very complex C++ style
- C++ can use either
  - C++ interface to MPI
  - C interface to MPI
- Our code (written by computer scientists) uses C++ interface …

Compiler experiences


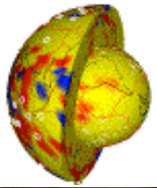



- Easy
  - Few changes to make, good feedback from compilers, didn't take me or machine long, etc
  - Altix, Origin, XD1, XT3
- X1
  - Took longer, but not too difficult
- MTA
  - Ran out of time to start
- P690/Power4+
  - Failed – no support for C++ MPI interface

MANCHESTER 1824

The University of Manchester

System/user experiences


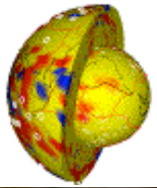



- Have used X1, XT3, XD1 this week
  - Plus Altix, Origin, p690
- Surprising variation between Cray machines, e.g.
  - Getting a simple batch job to run
  - Mpirun, aprun, yod, pbsyod, etc
  - Compiler calls and libraries/includes
- Overall new user experience is good, but not brilliant
  - Certainly no worse than any other new machine I have worked on
- Subjective – only my opinion

Performance - 1

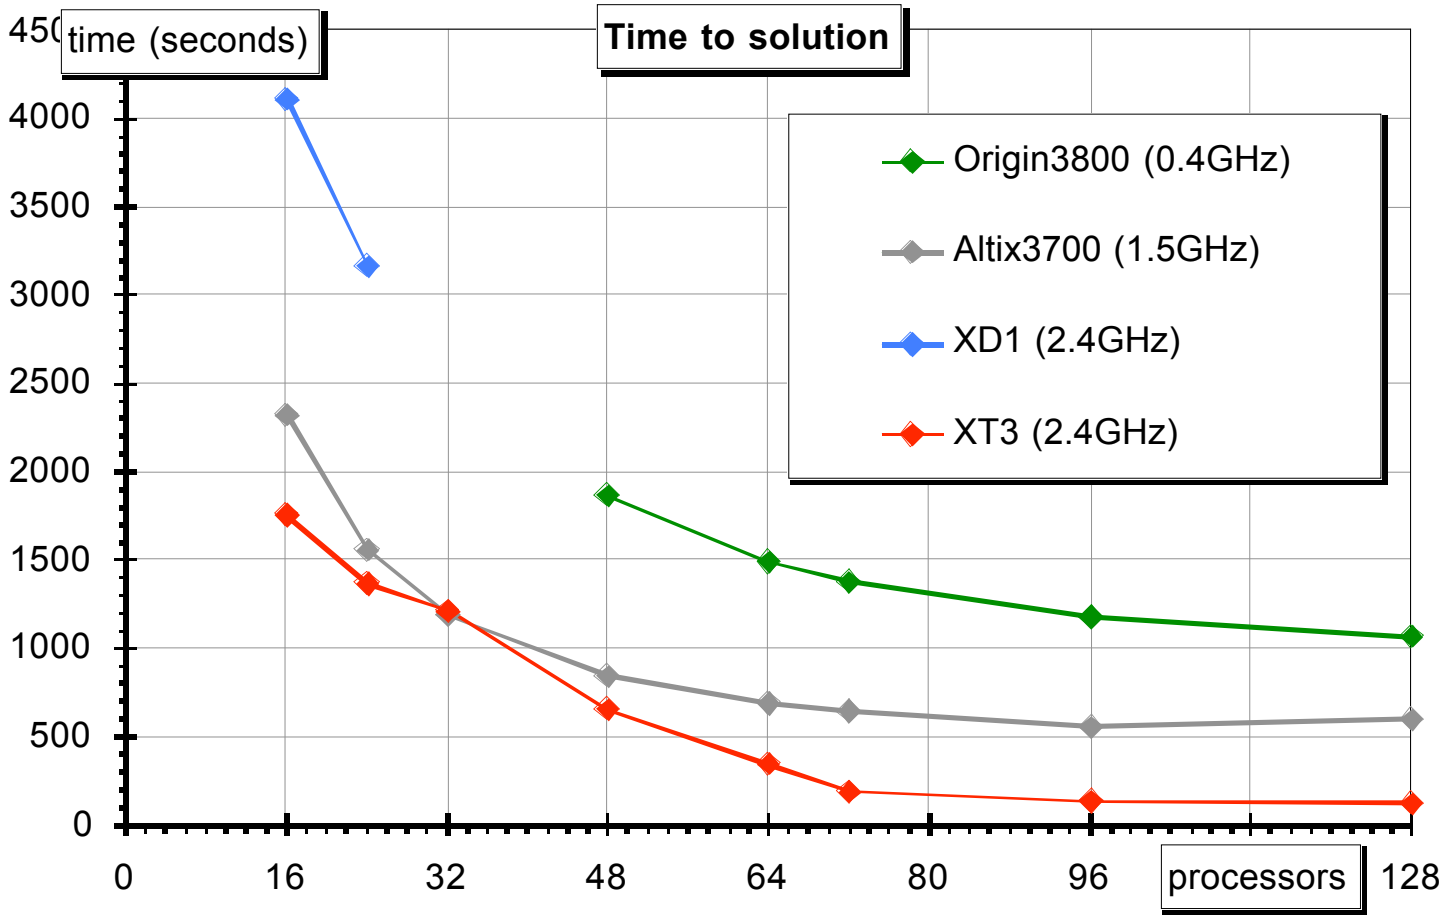

Performance - 2

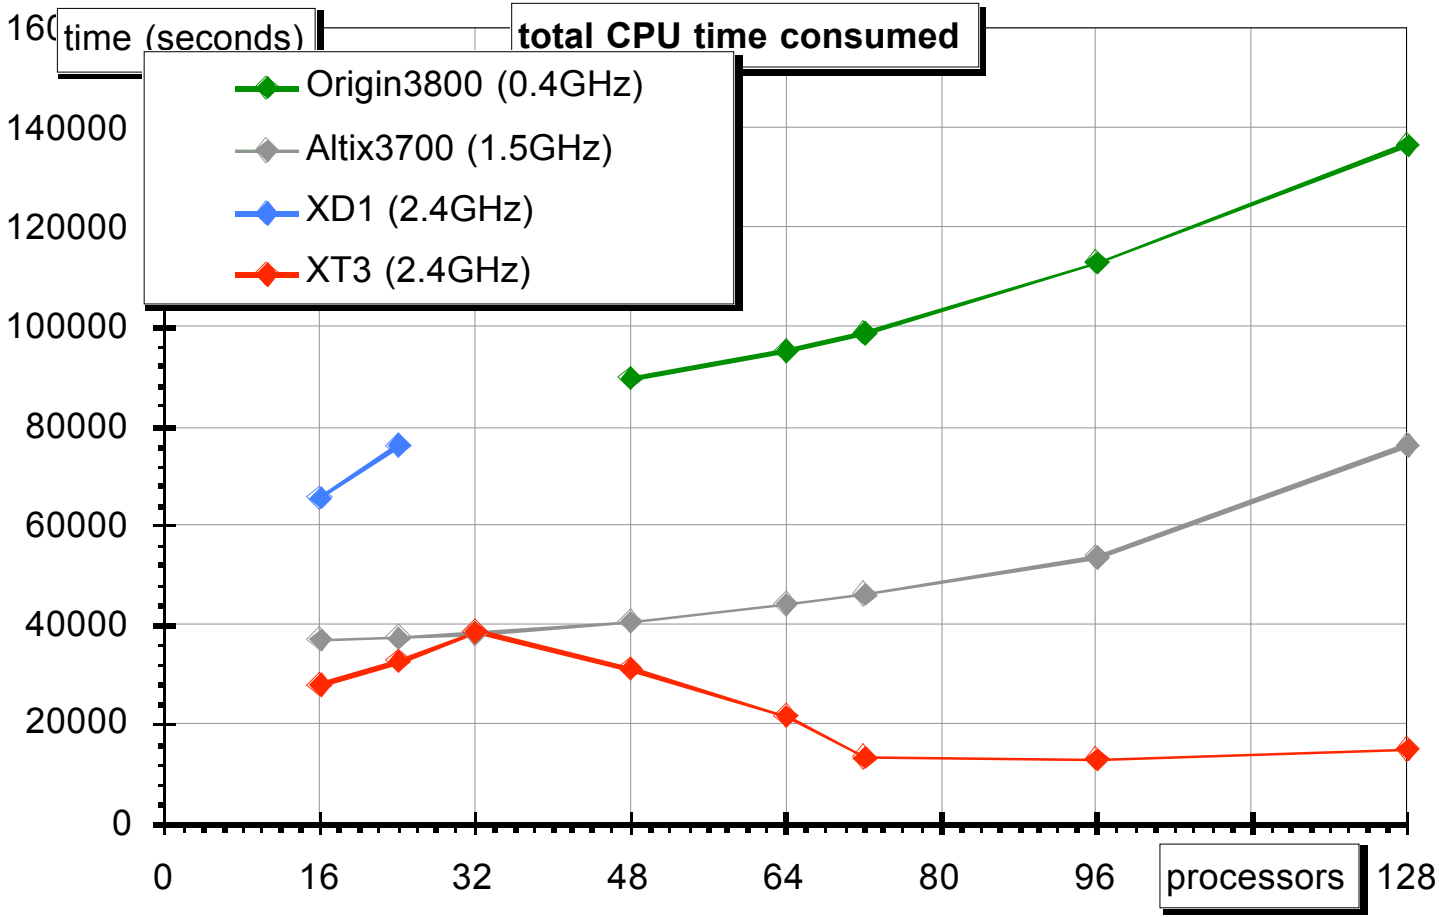

Performance - 3


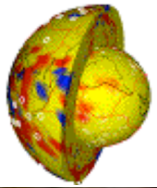



- Performance analysis/evaluation
- Investigate some “HPC options”
  - Write intensive core in UPC?
  - Multi-threading within MPI tasks?
  - Effort to re-write in F90/MPI – and get acceptance for this?
    - (zero chance)
  - What if not FLOPS based?
- Benefit of “HPC options” must be ***overwhelming*** to get acceptance for their use by the code owners
  - 10x faster is not good enough

MANCHESTER 1824

The University of Manchester

Summary


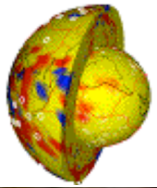



- We expect this project and the overall vision to be challenging
- Computer scientists are funny people
- Data mining has huge potential for new knowledge
- Supercomputing has much to offer data mining
- To get buy-in to new paradigms/technologies requires benefits to be overwhelming ***and*** easy to access/use
- Manchester has massive potential in data intensive computing
  - (Hard) task is to co-ordinate and exploit

MANCHESTER 1824

The University of Manchester

The University of Manchester

World Leading Research & Teaching

www.manchester.ac.uk

Combining the strengths of UMIST and
The Victoria University of Manchester