The Global Data Intensive Grid Collaboration

Rajkumar Buyya* (Collaboration Coordinator)
+ numerous contributors around the globe**.

* **Gri**d and **D**istributed **S**ystems Laboratory
Dept. of Computer Science and Software Engineering
The University of Melbourne, Australia

** **http://gridbus.cs.mu.oz.au/sc2003/participants.html**

Initial Proposal Authors (Alphabetical Order):
K. Branson (WEHI), R. Buyya (Melbourne), S. Date (Osaka), B. Hughes (Melbourne), Benjamin Khoo (IBM) , R. Moreno-Vozmediano (Madrid), J. Smilie (ANU), S. Venugopal (Melbourne), L. Winton (Melbourne), and J. Yu (Melbourne)

Next Generation Applications (NGA)

- Next generation experiments, simulations, sensors, satellites, even people and businesses are creating a flood of data. They all involve numerous experts/resources from multiple organization in *synthesis, modeling, simulation, analysis,* and interpretation.

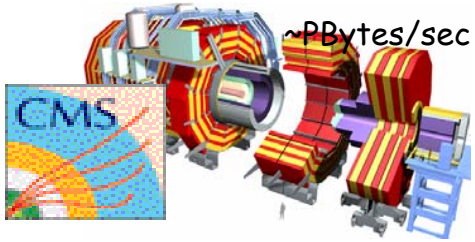

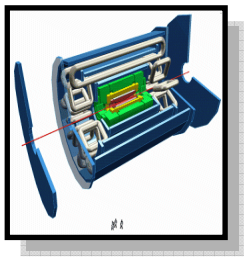

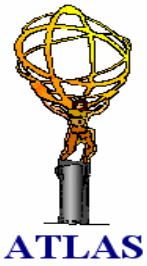

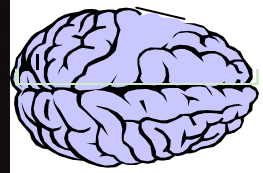

High Energy Physics

Brain Activity Analysis

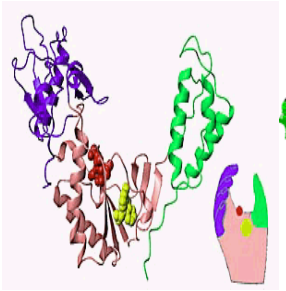

Life Sciences

Digital Biology

Astronomy

Newswire & data mining:
Natural language engineering

Finance: Portfolio analysis

Quantum Chemistry

Internet & Ecommerce

Common Attributes/Needs/Challenges of NGA

- They involve Distributed Entities:
  - Participants/ Organizations
  - Resources
    - Computers
    - Instruments
  - Datasets/Databases
    - Source (e.g., CDB/PDBs)
    - Replication (e.g, HEP Data)
  - Application Components
- Heterogeneous in nature
- Participants require share analysis results of analysis with other collaborators (e.g., HEP)

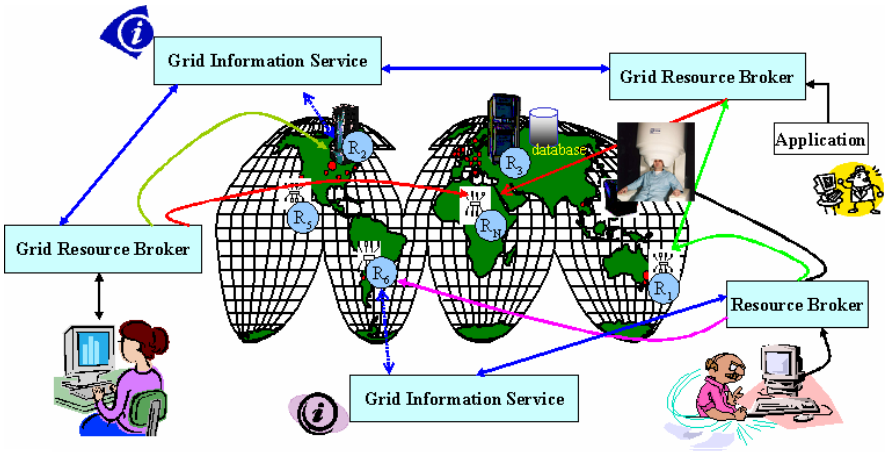

• Grids offer the most promising solution & enable global collaborations.
•The beauty of the grid is that it provides a secure access to a wide range of heterogeneous resources.
•But what does it take to integrate and manage applications across all these resources?

What is "The Global Data Intensive Grid Collaboration" Doing ?

- Assembled several heterogeneous resources, technologies, data-intensive applications of both tightly and loosely coordinated groups and institutions around the world in order to demonstrate both HPC Challenges:
  - Most Data-Intensive Application(s)
  - Most Geographically Distributed Application (s).

The Members of Collaboration

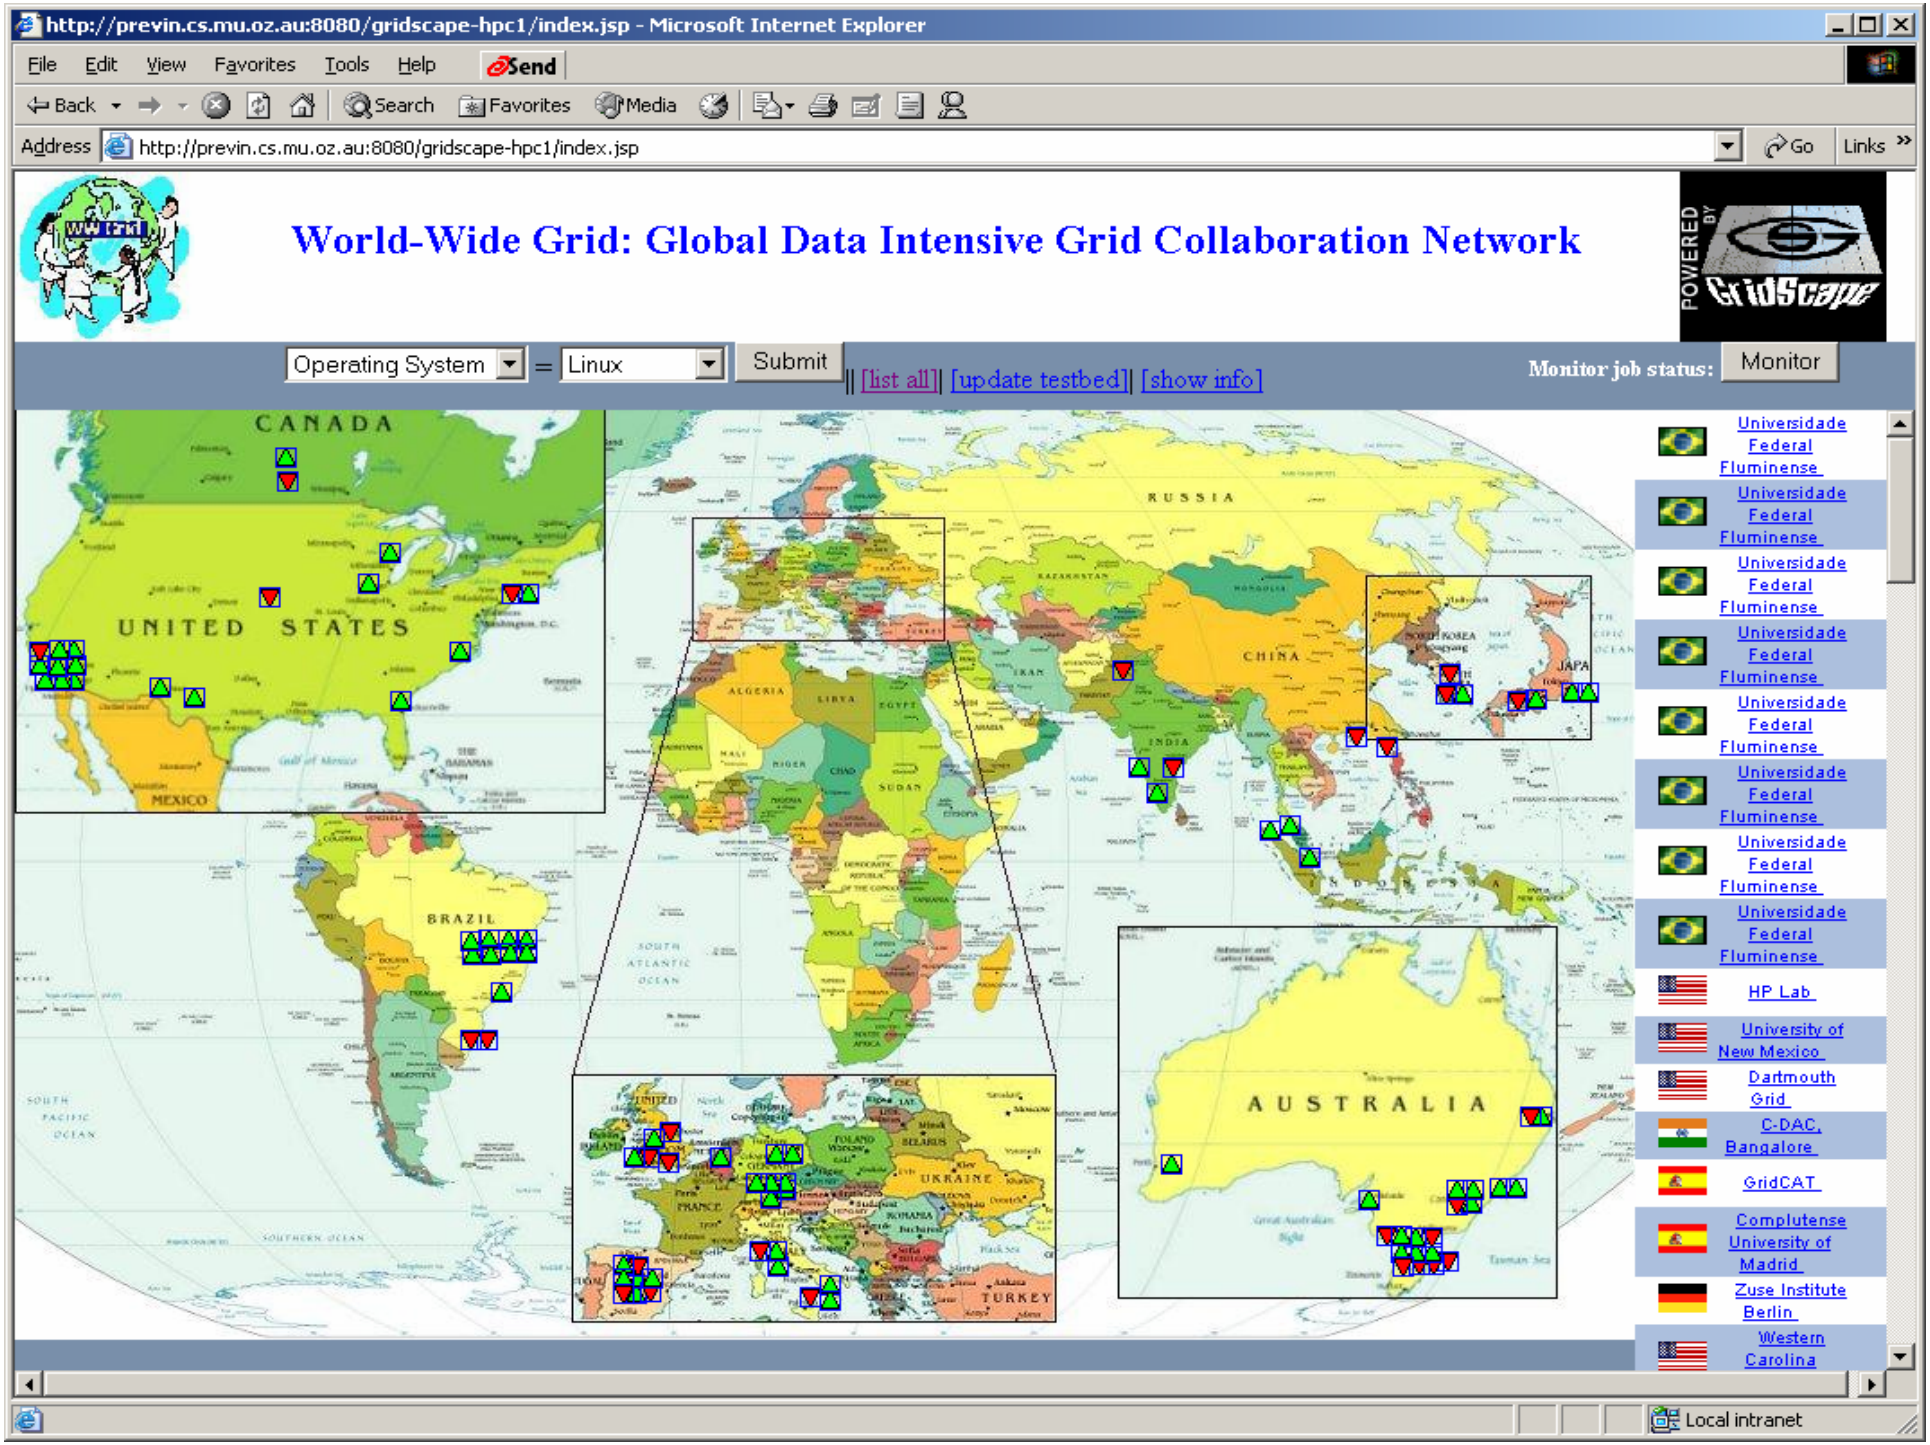
http://previn.cs.mu.oz.au:8080/gridscape-hpc1/index.jsp - Microsoft Internet Explorer
File Edit View Favorites Tools Help Send
Back Search Favorites Media
Address http://previn.cs.mu.oz.au:8080/gridscape-hpc1/index.jsp
Go Links
World-Wide Grid: Global Data Intensive Grid Collaboration Network
POWERED BY GridScape
Operating System = Linux Submit
[list all] [update testbed] [show info]
Monitor job status: Monitor
CANADA
UNITED STATES
MEXICO
RUSSIA
CHINA
INDIA
JAPAN
BRAZIL
AUSTRALIA
TURKEY
UKRAINE
Universidade Federal Fluminense
HP Lab
University of New Mexico
Dartmouth Grid
C-DAC, Bangalore
GridCAT
Complutense University of Madrid
Zuse Institute Berlin
Western Carolina
Local intranet

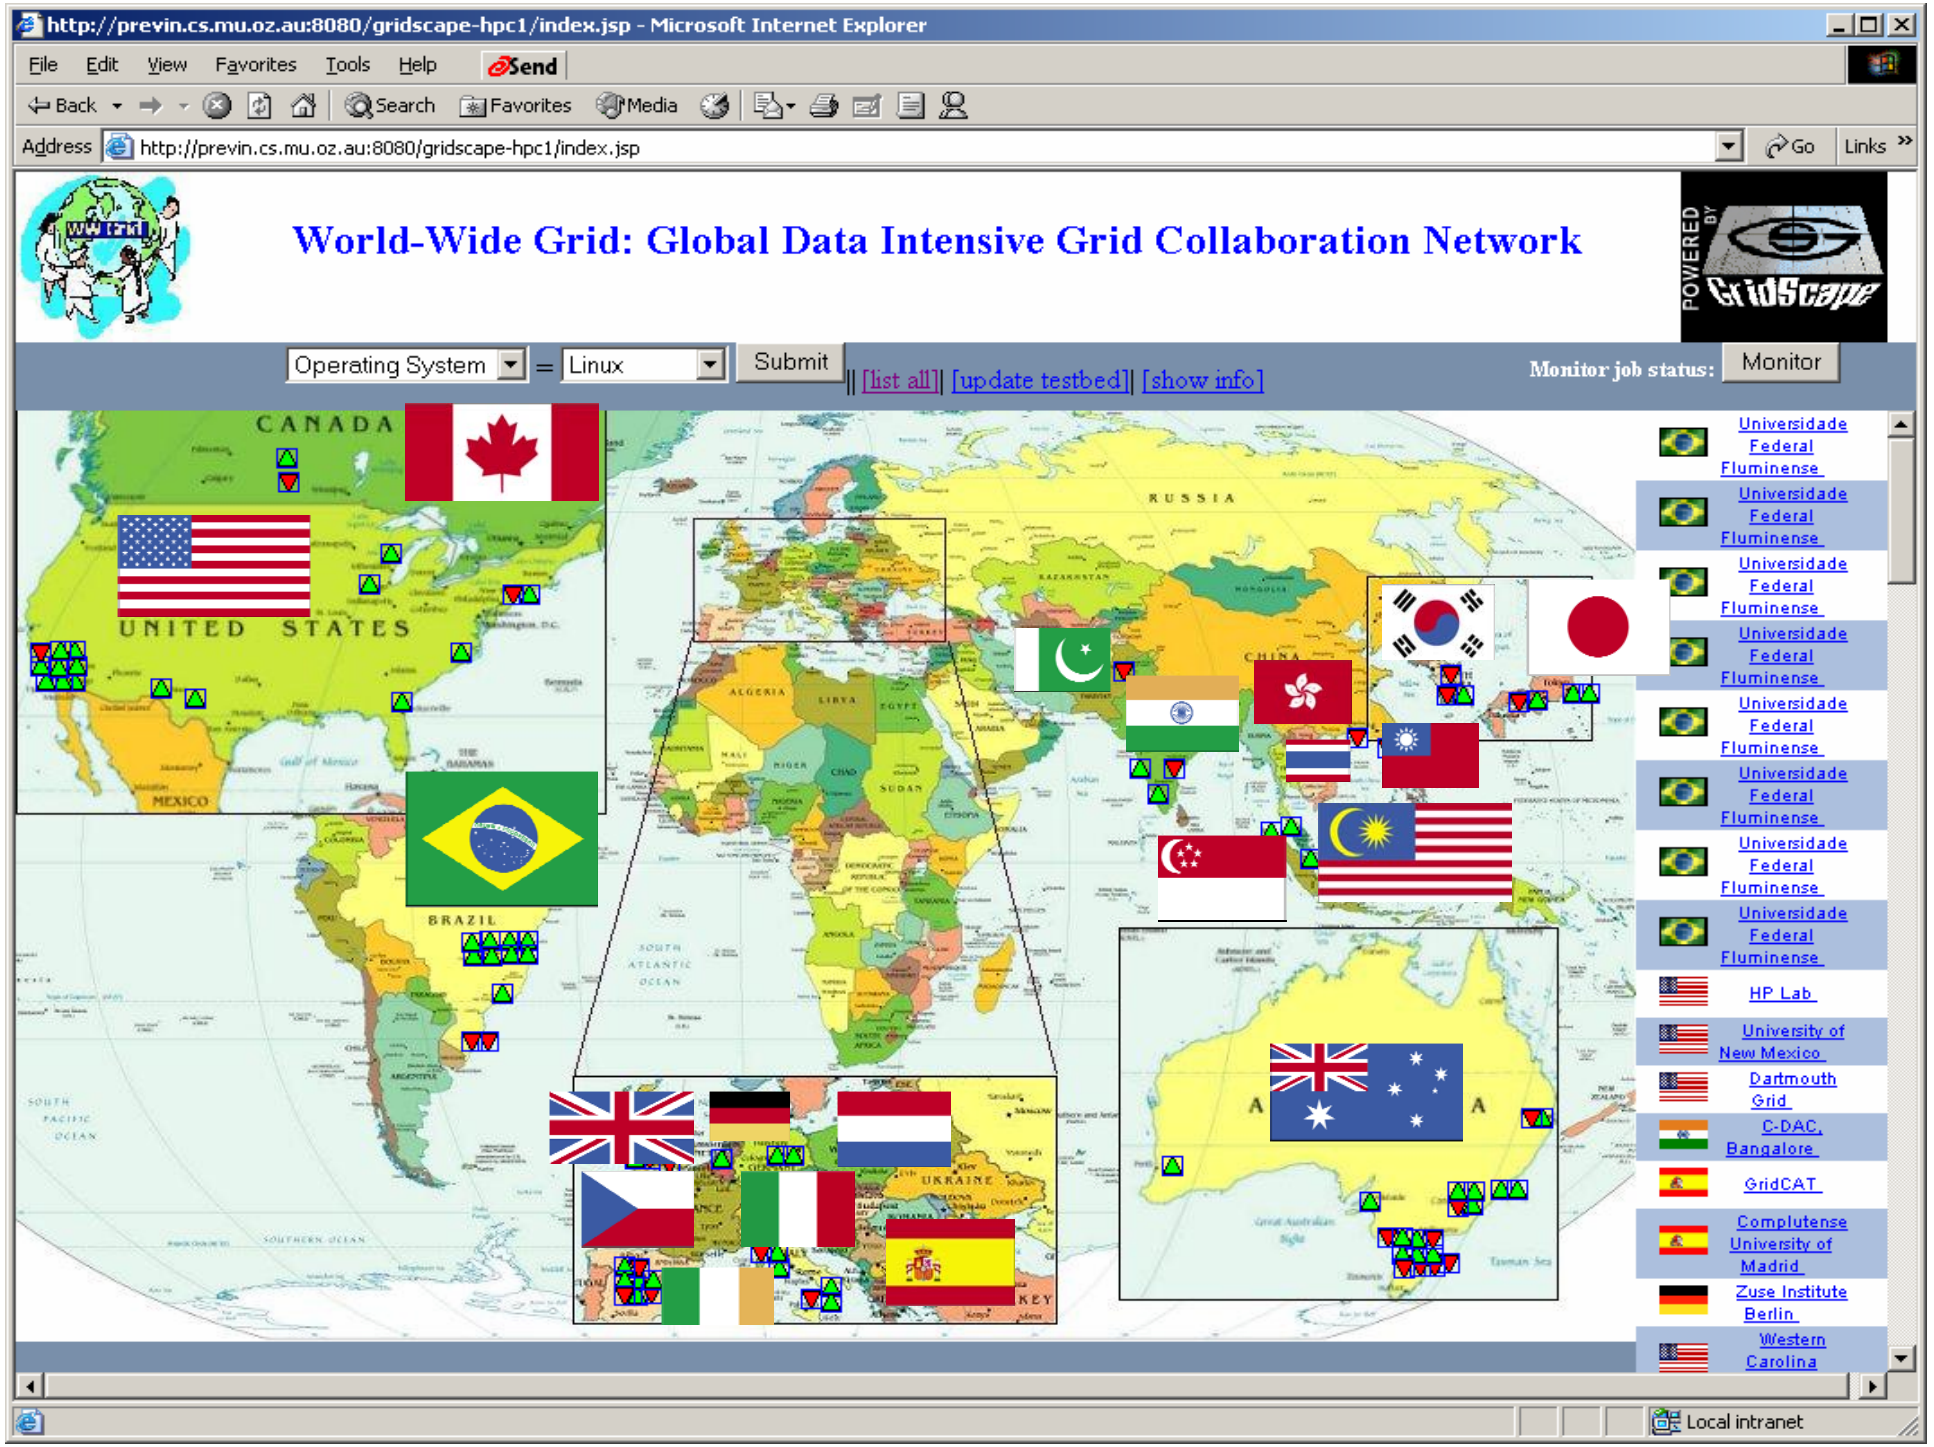
http://previn.cs.mu.oz.au:8080/gridscape-hpc1/index.jsp - Microsoft Internet Explorer
File Edit View Favorites Tools Help Send
Back Search Favorites Media
Address http://previn.cs.mu.oz.au:8080/gridscape-hpc1/index.jsp
Go Links
World-Wide Grid: Global Data Intensive Grid Collaboration Network
POWERED BY GridScape
Operating System = Linux Submit
[list all] [update testbed] [show info]
Monitor job status: Monitor
CANADA
UNITED STATES
MEXICO
BRAZIL
RUSSIA
CHINA
Universidade Federal Fluminense
Universidade Federal Fluminense
Universidade Federal Fluminense
Universidade Federal Fluminense
Universidade Federal Fluminense
Universidade Federal Fluminense
Universidade Federal Fluminense
Universidade Federal Fluminense
HP Lab
University of New Mexico
Dartmouth Grid
C-DAC, Bangalore
GridCAT
Complutense University of Madrid
Zuse Institute Berlin
Western Carolina
Local intranet

Testbed Statistics (Browse the Testbed)

- Grid Nodes: 218 distributed across 62 sites in 21 countries.
  - Laptops, desktop PCs, WS, SMPs, Clusters, supercomputers
  - Total CPUs: 3000+ (~3 TeraFlops)
- CPU Architecture:
  - Intel x86, IA64, AMD, PowerPC, Alpha, MIPS
- Operating Systems:
  - Windows or Unix-variants – Linux, Solaris, AIX, OSF, Irix, HP-UX
- Intranode Network:
  - Ethernet, Fast Ethernet, Gigabit, Myrinet, QsNet, PARAMNet
- Internet/Wide Area Networks
  - GrangeNet, AARNet, ERNet, APAN, TransPAC, & so on.

Grid Technologies and Applications

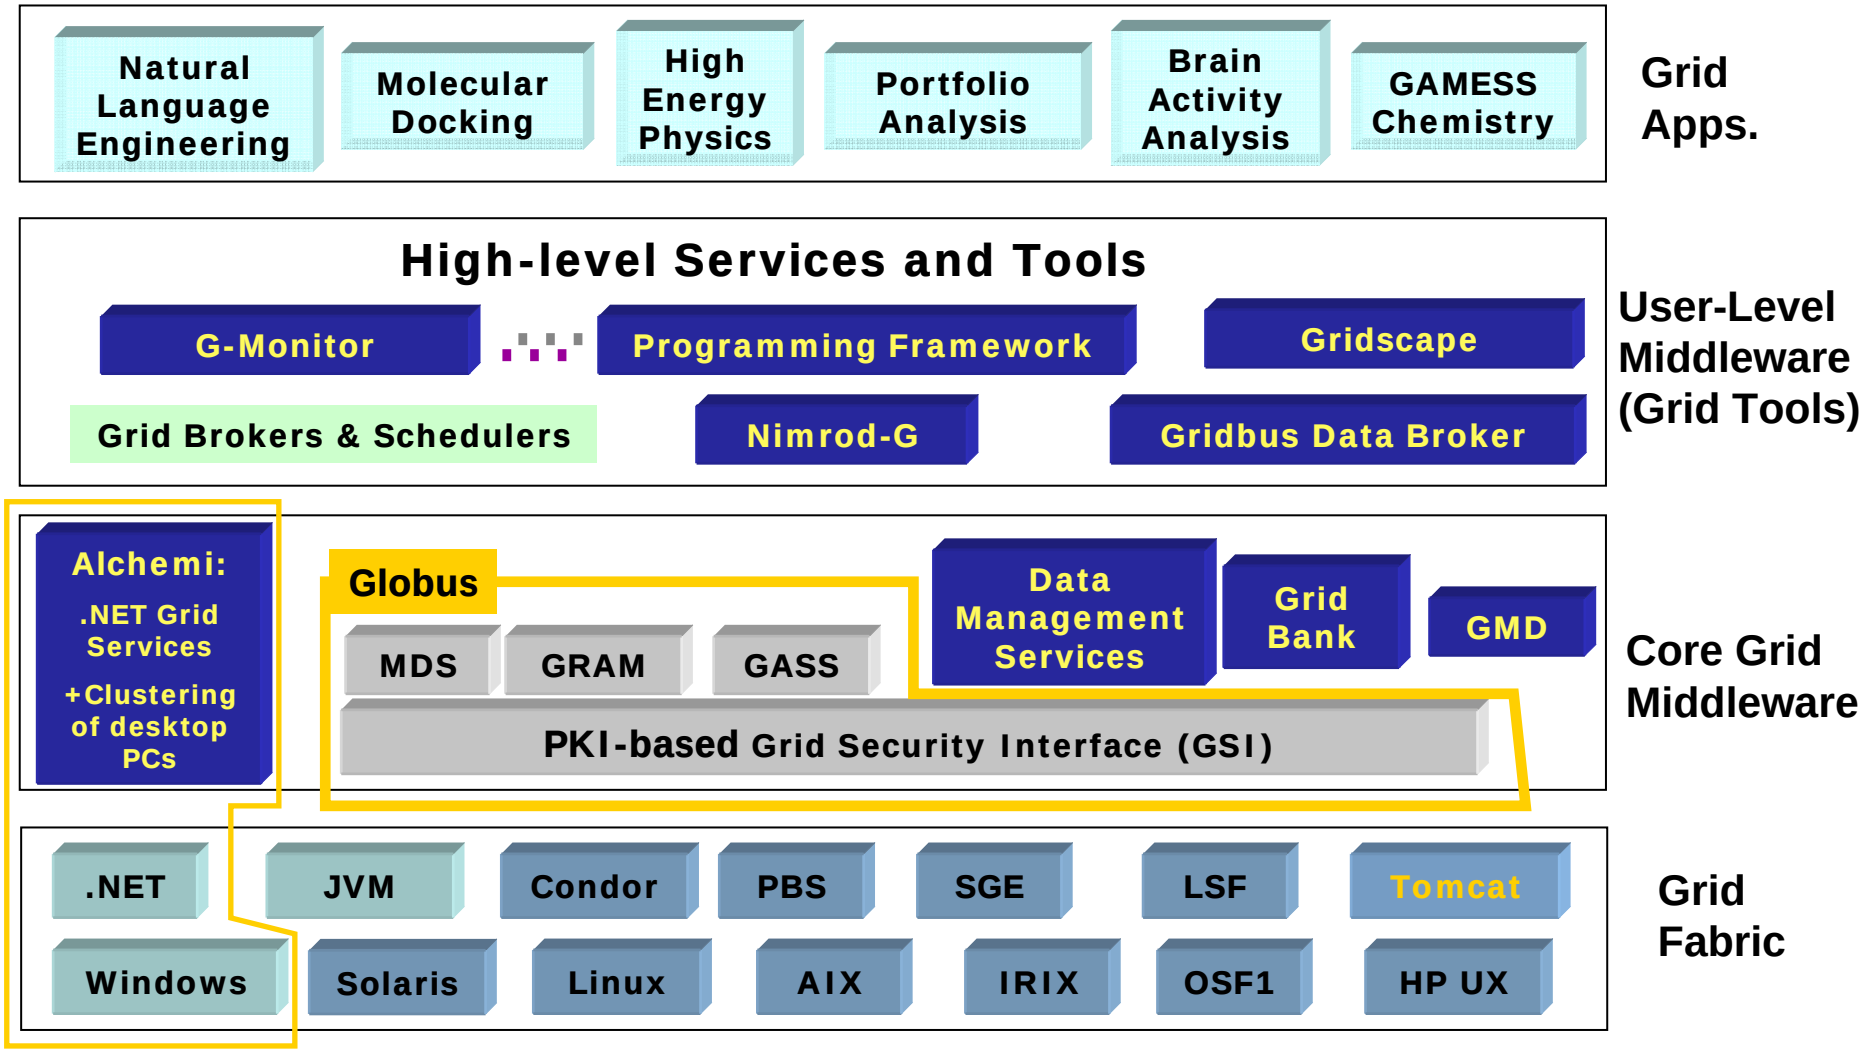

Application Targets

- High Energy Physics – Melbourne School of Physics
  - Belle experiment – CP (charge parity) violation
- Natural Language Engineering – Melbourne School of CS
  - Indexing Newswire Text
- Protein Docking – WEHI for Medical Research, Melbourne
  - Screening molecules to identify their potential as drug candidates
- Portfolio Analysis – UCM, Spain
  - Value at Risk/Investment risk analysis
- Brain Activity Analysis – Osaka University, Japan
  - Identifying symptoms of common disorders through analysis of brain activity patterns.
- Quantum Chemistry - Monash and SDSC effort
  - GAMESS

HPC Challenge Demo Setup

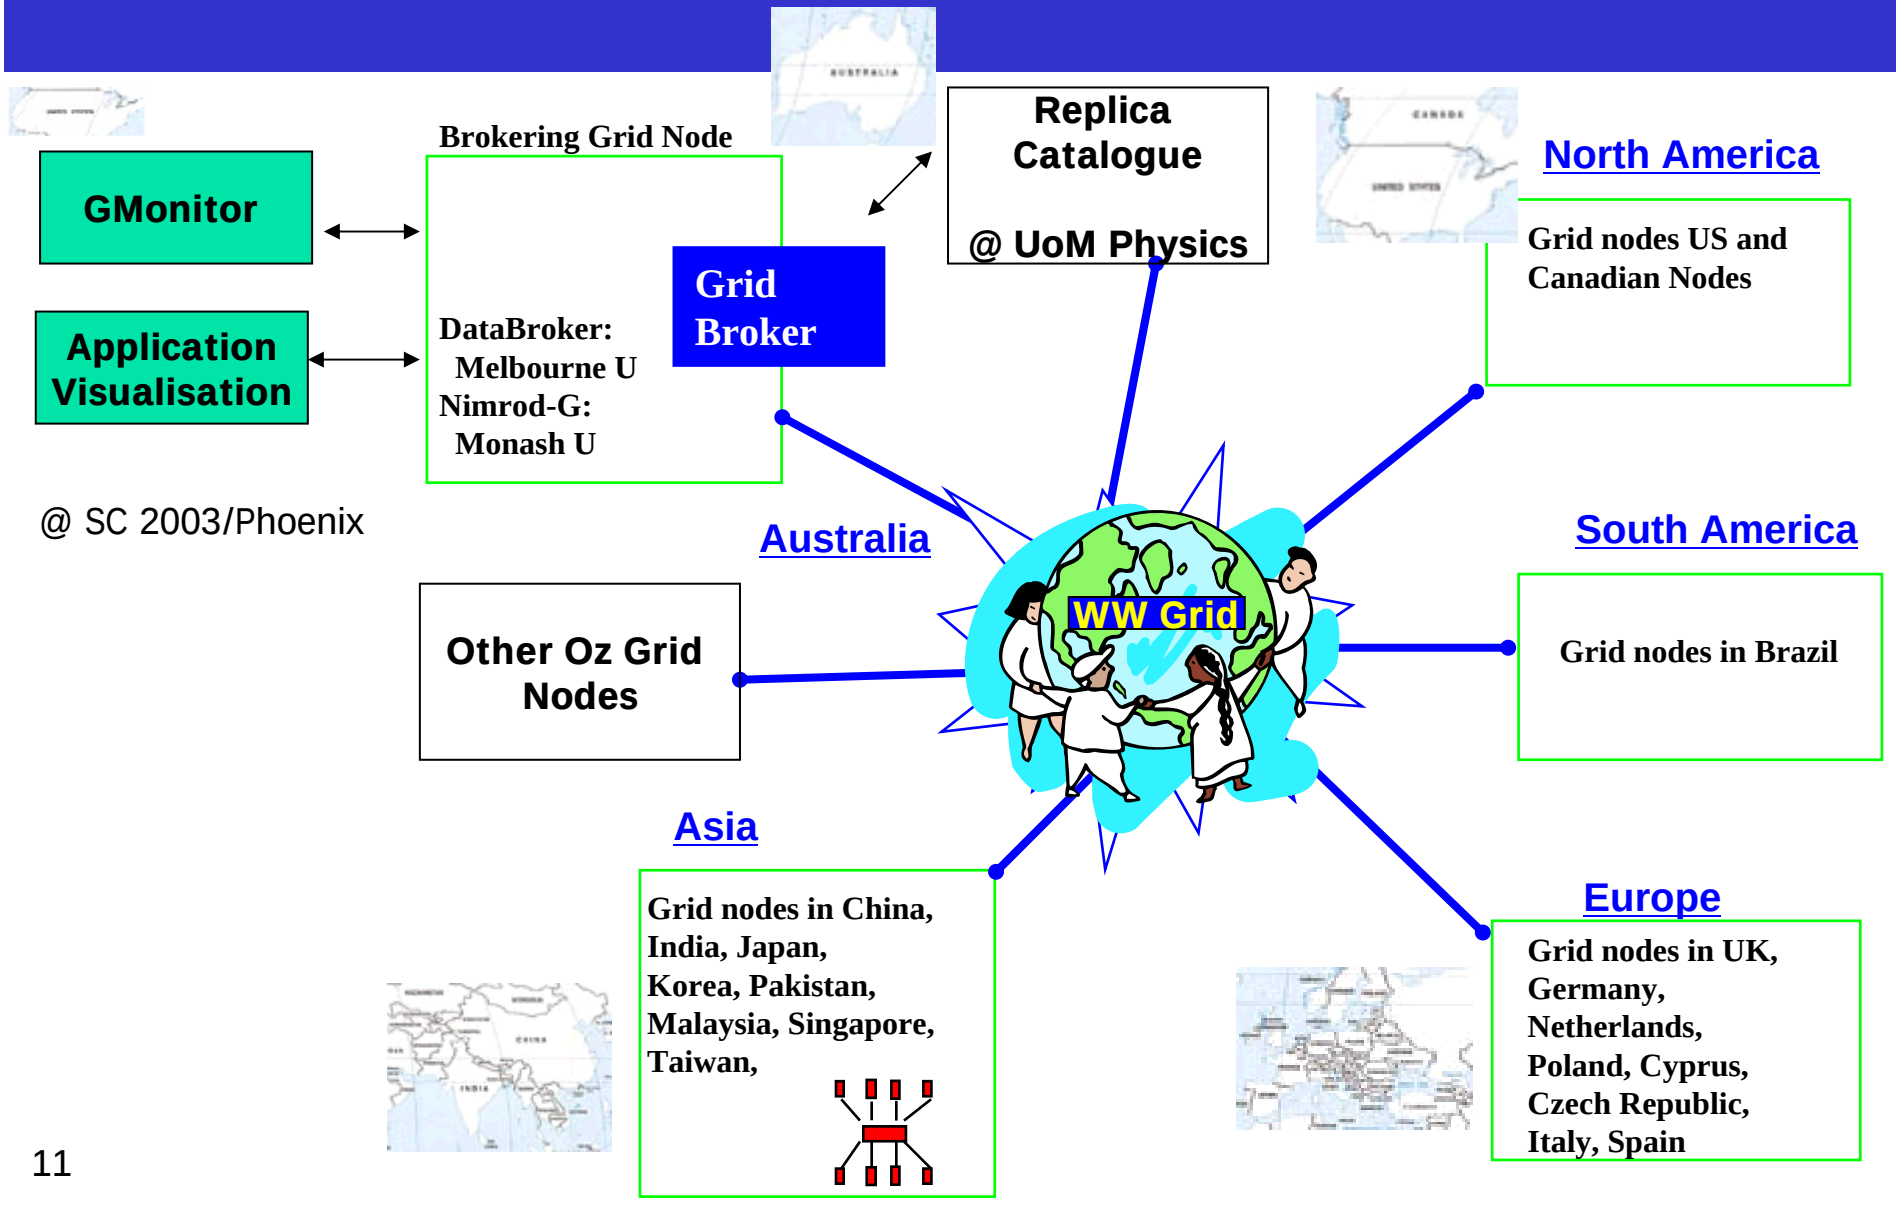

Belle Particle Physics Experiment

- A Running experiment based in KEK B-Factory, Japan
- Investigating fundamental violation of symmetry in nature (Charge Parity) which may help explain the universal matter – antimatter imbalance.
- Collaboration 400 people, 50 institutes
- 100's TB data currently
- UoM School of Physics is an active participant and have led the Grid-enabling of the Belle data analysis framework.

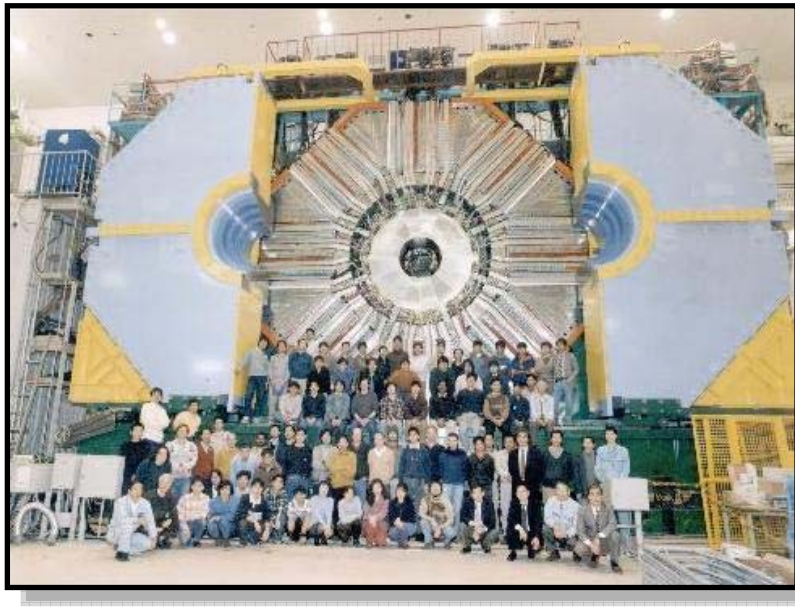

Belle Demo - Simulate specific event of interest B0 → D*-D*+KS

- Generation of Belle data (1,000,000 simulated events):
  - Simulated (or Monte Carlo) data can be generated anywhere, relatively inexpensively
    - Full simulation is very CPU intensive (full physics of interaction, particles, materials, electronics)
  - We need more simulated than real data to help eliminate statistical fluctuations in our efficiency calculations.
- Simulated specific event of interest:
  - Decay Chain: B0 → D*-D*+KS (Particle B0 decays into 3 particles D*, -D*, +KS)

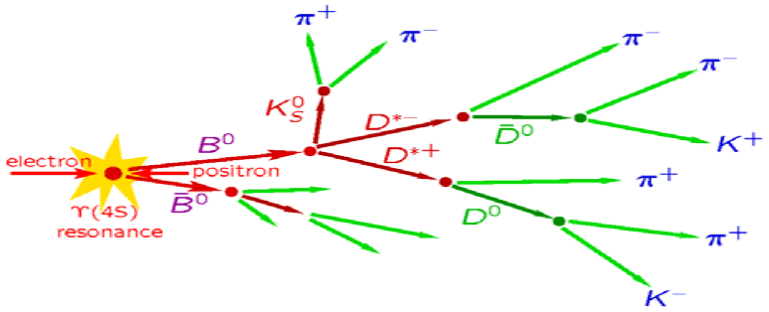

- The data has been made available to the collaboration via global directory structure (Replica Catalog).
- During the analysis, the broker discovers data using Replica Catalog services.

Analysis

- During the demo, we analysed 1,000,000 events using the Grid-enabled BASF (Belle Analysis Software Framework) code .
- The Gridbus broker discovered the catalogued data (lfn:/users/winton/fsimddks/*.mdst) and decomposed them into 100 Grid jobs (each input file size = 3MB) and processed on Belle nodes located in Australia and Japan.
- The broker has optimised the assignment of jobs to Grid nodes to *minimise* both the data transmission time and computation time and finished the analysis in **20** minutes.
- The analysis output histograms has been visualized:

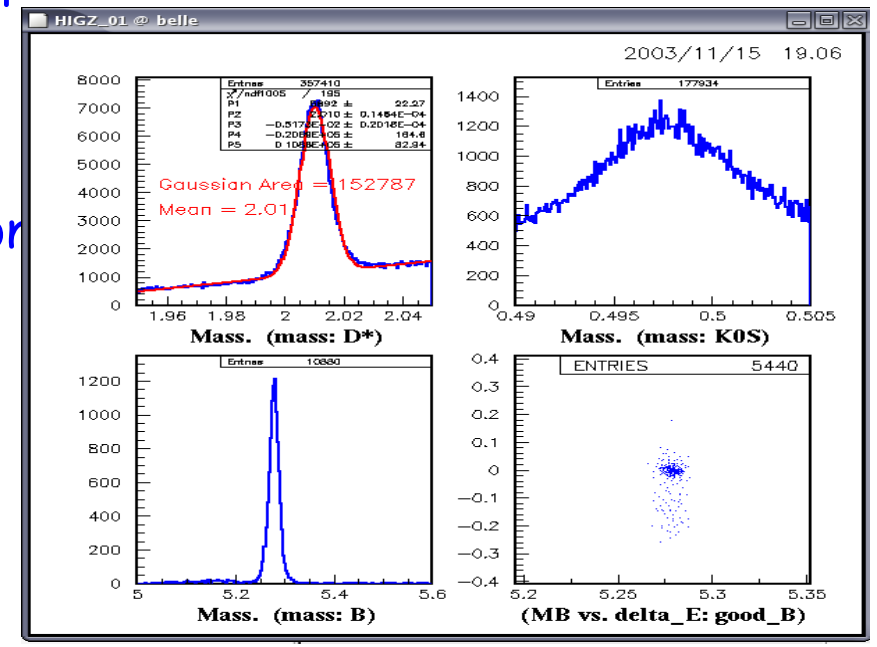

Histogram of an analysis

Indexing Newswire: A Natural Language Engineering Problem

- A newswire service is a dedicated feed of stories from a larger news agency, provided to smaller content aggregators for syndication.
- Essentially a continuous stream of text with little internal structure.
- So, why would we choose to work with such data sources ?
  - Historical enquiry. For example,
    - find all the stories in 1995 about Microsoft and Internet;
    - when was the Bill Clinton and Monica Lewinsky story first exposed.
  - Evaluating how different agencies reported the same event from different perspectives eg US vs European media, New York vs Los Angeles media, television vs cable vs print vs internet.
- The challenge is how do we extract meaningful information from newswire archives efficiently?

Data and Processing

- In this experiment we used samples from the Linguistic Data Consortium's Gigaword Corpus, which is a collection of 4 different newswire sources (Agence France Press English Service, Associated Press Worldstream English Service, New York Times Newswire Service, and Xinhua News Agency over a period of 7 years.
- A typical newswire service generates 15-20Mb per month of raw text.
- We carried two different types of analysis: statistical & indexational. We extracted all the relevant document IDs and headlines for a specific document type to create an index to the archive itself.
- In the demonstration, we used the 1995 collection from Agence France Press (AFP) English Service, which contains about **100Mb** of newswire text.
  - Analysis was carried out on the testbed resources that are conneted by the Australian GrangeNet to minimise the time for input and out data movement and also the processing time.
  - Grid-based analysis was finished in **10** minutes.

Portfolio Analysis on Grid

- Intuitive definition of Value-at-Risk (VaR)
  - Given a trading portfolio, the VaR of the portfolio, provides an answer to the following question:
    - *How much money can I lose over a given time horizon with a given probability ?????*
- Example
  - If the Value-at-Risk of my portfolio is
    - VaR(c=95%,T=10) = 1.0 million dollars
    - c – level of confidence, T is holding period
  - It means:
    - The probability of losing more than 1 million dollars over a holding period of 10 days is lower than 5% (1-c)

Computing VaR: the simulation process

- During the demo, We simulated (Monte-Carlo) N-independent price-paths for the portfolio by using most of the available Grid nodes in the testbed during the demo and finished the analysis within 20 minutes.
- There was significant overlap of Grid nodes during the demo of each application.

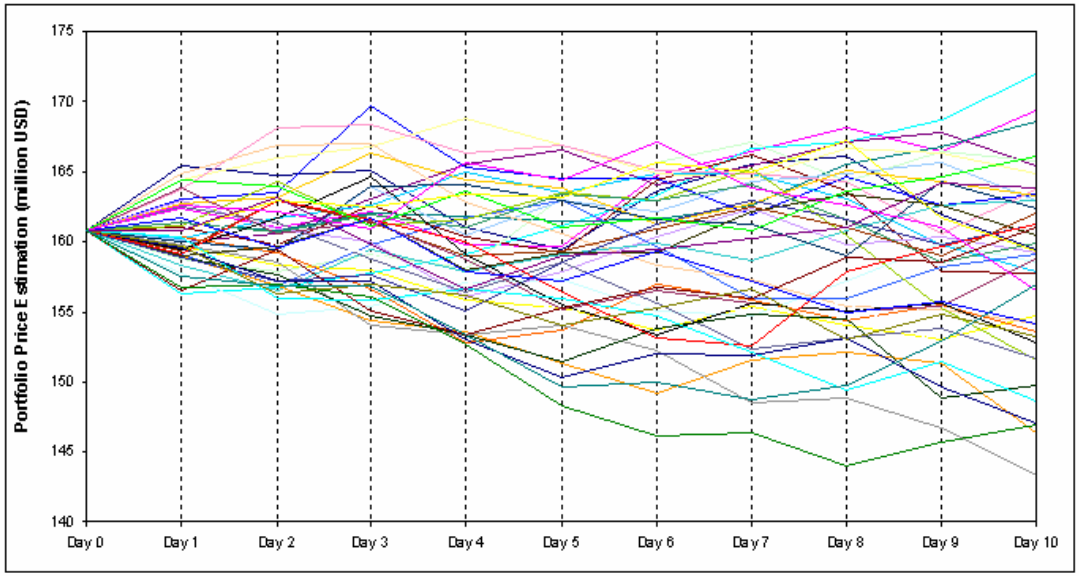

Computing VaR: the output

- Once simulated N independent price paths
  - We obtain a frequency distribution of the N changes in the value of a portfolio
  - The VaR with confidence c can be computed as the (1-c)-percentile of this distribution

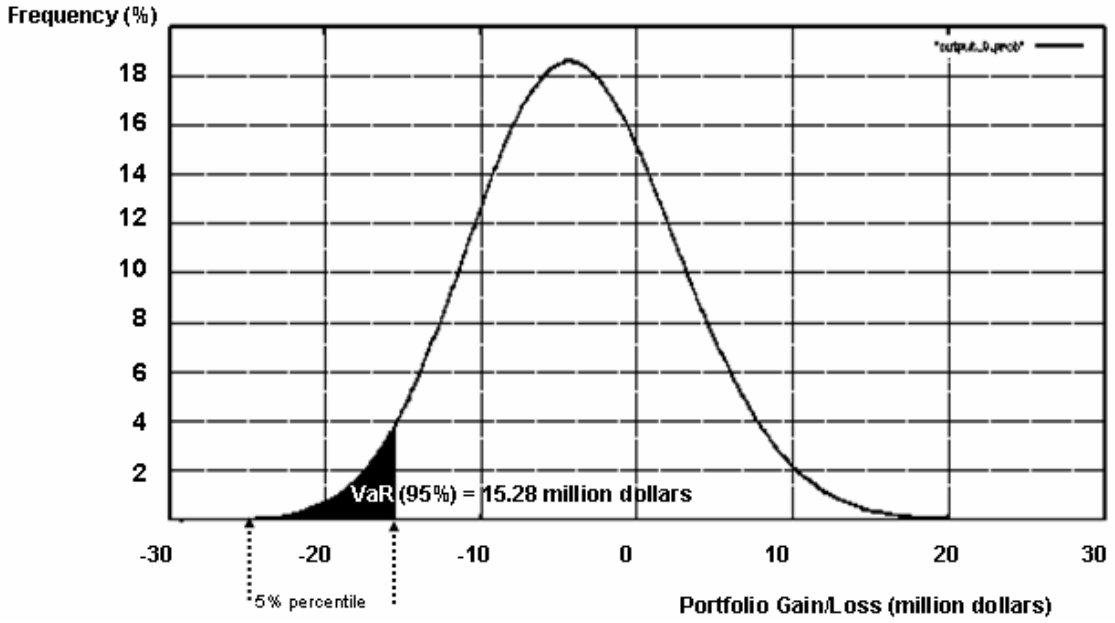

Quantum Chemistry on Grid

- Parameter Scan of an Effective Group Difference Pseudopotential.
- An experiment by:
  - Kim Baldridge and Wibke Sudholt, UCSD
  - David Abramson and Slavisa Garic, Monash
- Using GAMESS (General Atomic and Molecular Electronic Structure System) application and Nimrod-G broker
- A pre-started experiment and continued during the demo and used majority of available Grid nodes.
- Analyzed electrons and positioning of atoms for various scenarios.
- 13,500 jobs (each job took 5-78 minutes) finished in 15 hours.
- Input: 4KB for each job;
- Total output 860MB compressed.

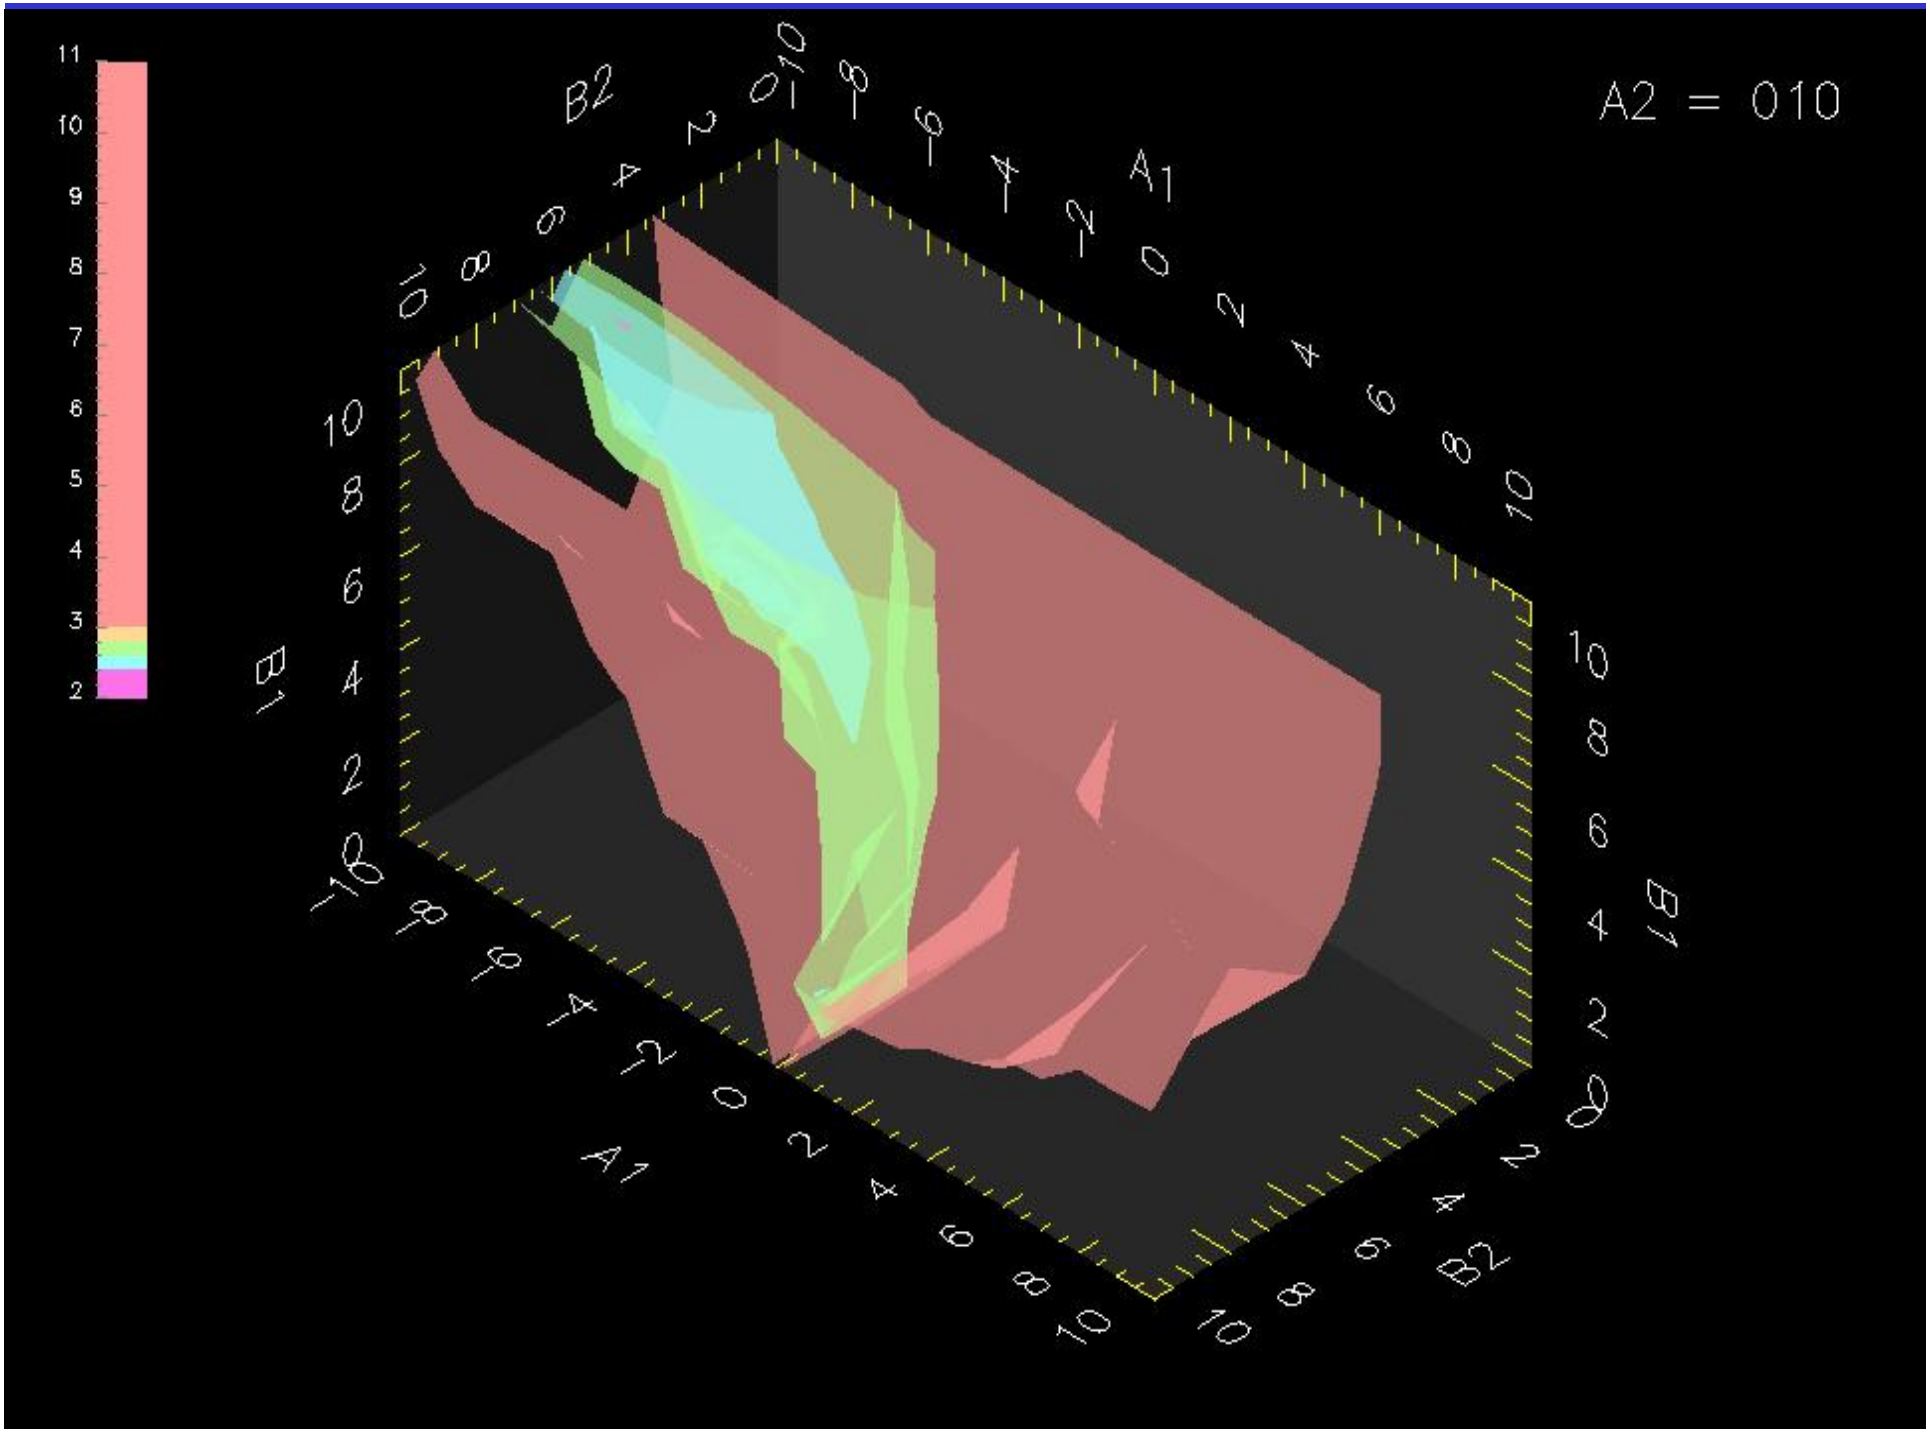
A2 = 010
B2
A1
B1
11
10
9
8
7
6
5
4
3
2

Analysis Summary

Application	Data Size	Processing Time	Nodes
Belle Analysis (HEP)	300 MB input (100 jobs – 3MB each)	30 min.	Australia, Japan
Financial Portfolio Analysis	50 MB output (50 jobs – 1MB each)	20 min.	Global
Newswire Indexing	80 MB input (12 jobs – 7MB each job)	20 min.	GrangeNet, Australia
GAMESS	4KB for each job. Total output: 860MB compressed	Each job took 5-78 minutes. Total 15 hours	Global (130 nodes, 15 sites)

Summary and Conclusion

- The Global Data Intensive Grid Collaboration has successfully put together:
  - 218 heterogeneously Grid nodes distributed across 62 sites in 21 countries around the globe.
  - they are Grid enabled by technologies (Unix and also Windows based Grid technologies),
  - 6 data-intensive applications: HEP, NLE, Docking, Neuroscience, Quantum Chemistry, & Finance
- And demonstrated both HPC Challenges:
  - Most Data-Intensive Application(s)
  - Most Geographically Distributed Application (s).
- It was all possible due to the hard work of numerous volunteers around the world.

Contributing Persons

Akshay Luther
Alexander Reinefeld
Andre Merzky
Andrea Lorenz
Andrew Wendelborn
Arshad Ali
Arun Agarwal
Baden Hughes
Barry Wilkinson
Benjamin Khoo
Christopher Jordan
Colin Enticott
Cory Lueninghoener
Darran Carey
David Abramson
David A. Bader
David Baker
David Glass
Diego Luis Kreutz
Ding Choon-Hoong
Dirk Van Der Knijff
Fabrizio Magugliani
Fang-Pang Lin
Gabriel
Garry Smith
Gee-Bum Koo
Giancarlo Bartoli
Glen Moloney
Gokul Poduval
Grace Foo
Heinz Stockinger
Helmut Heller
Henri Casanova
James E. Dobson
Jem Treadwell
Jia Yu
Jim Hayes
Jim Prewett
John Henriksson
Jon Smillie
Jonathan Giddy
Jose Alcantara
Kashif
Kees Verstoep
Kevin Varvell
Latha Srinivasan
Lluis Ribes
Lyle Winton
Manish Prashar
Markus Buchhorn
Martin Savior
Matthew
Michael Monty
Michal Vocu
Michelle Gower
MohanRam
Nazarul Nasirin
Niall Wilson
Nigel Teow
Oscar Ardaiz
Paolo Trunfio
Paul Coddington
Putchong Uthayopas
R.K.Shyamasundar
Radha Nandakumar
Rafael M-Vozmediano
Rafal Metkowski
Raj Chhabra
Rajalakshmy
Rajiv
Rajiv Ranjan
Rajkumar Buyya
Ricardo
Robert Sturrock
Rodrigo Real
Roy S.C. Ho
S. Anbalagan
Sandeep K. Joshi
Selina Dennis
Sergey
Slavisa Garic
Srikumar
Steven Bird
Steven Melnikoff
Subhek Garg
Subrata Chattopadhyay
Sudarshan
Sugree
Susumu Date
Thomas Hacker
Tony McHale
V.C.V. Rao
Vinod Rebello
Viraj Bhat
Wayne Kelly
Xavier Fernandez
Y.Tanimura
Yeo
Yoshio Tanaka
Yu-Chung Chen

Thanks for your attention!

The Global Data-Intensive Grid Collaboration
htpp://gridbus.cs.mu.oz.au/sc2003/