Provenance-Aware Fault Tolerance for Grid Computing

Professor Jie Xu (*jxu@comp.leeds.ac.uk*)
Director of the WRG e-Science Centre of Excellence
University of Leeds & University of Newcastle upon Tyne, UK

THE WHITE ROSE GRID
e-Science Centre of Excellence

The White Rose Grid Project

- The **three Yorkshire Universities' project** (started in 2001, over £10M investment and research projects) http://www.wrgrid.org.uk/
- **Involves** Leeds (Profs K Brodlie, P M Dew & **J Xu**), York (Prof J Austin), and Sheffield (Profs G Tomlinson & P Fleming); under the guidance of the Chief Executive of WRUC (Dr Julian White – CEO of WRUC)
- White Rose University Consortium – a strategic partnership of the three Universities - http://www.whiterose.ac.uk
- Excellent partnership with **Computing Services** & Comp Science (Dr S Chidlow, C Cartledge, Dr A Turner)
- **Partners:** Esteem Systems in conjunction with Sun Microsystems & Streamline Computing
- Supported by Yorkshire Forward, Y&H Reg Dev Agency

THE WHITE ROSE GRID
e-Science Centre of Excellence

The WRG Architecture

General Purpose HPC node

Computer Science node

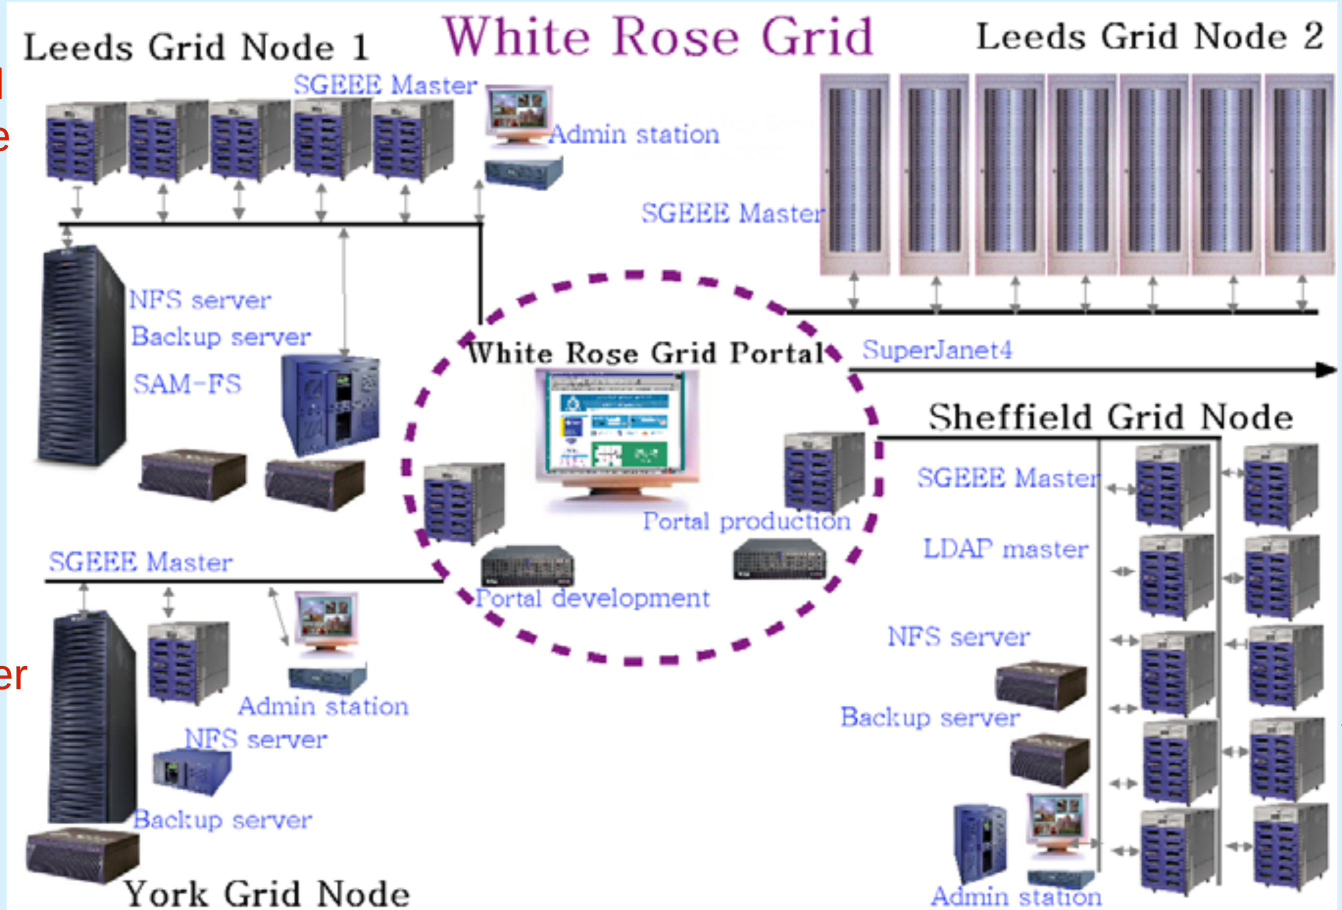

CFD node

Engineering Application node

UK e-Science Centres

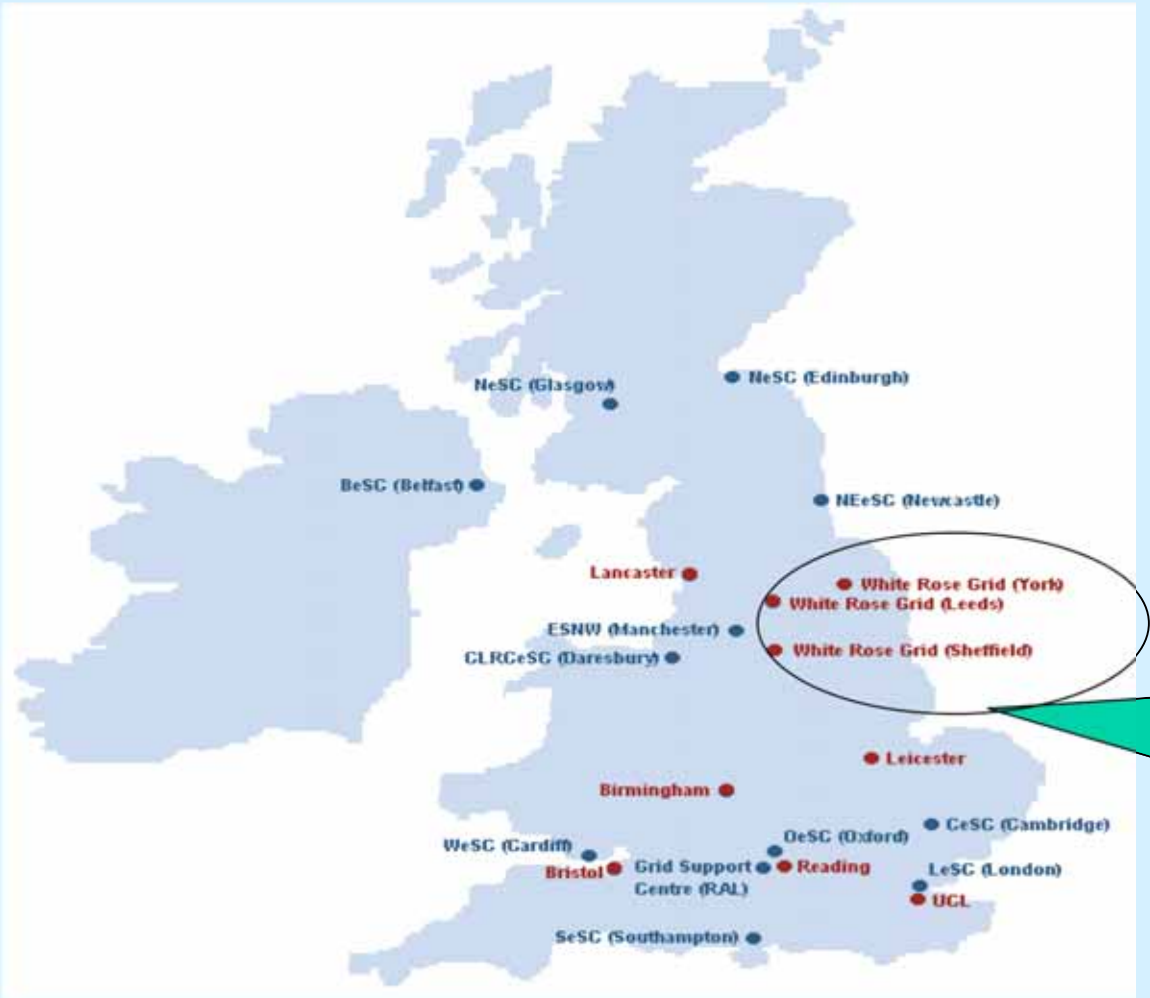

Our Centre:

- To offer focus for a variety of e-science issues and activities in our region
- To develop close links with the UK e-Science CP
- To develop a particular specialism: visualisation, distributed diagnostics and system dependability

The White Rose Grid e-Science Centre of Excellence

UK e-Science Centres (courtesy of NeSC)

THE WHITE ROSE GRID
e-Science Centre of Excellence

National Grid Service

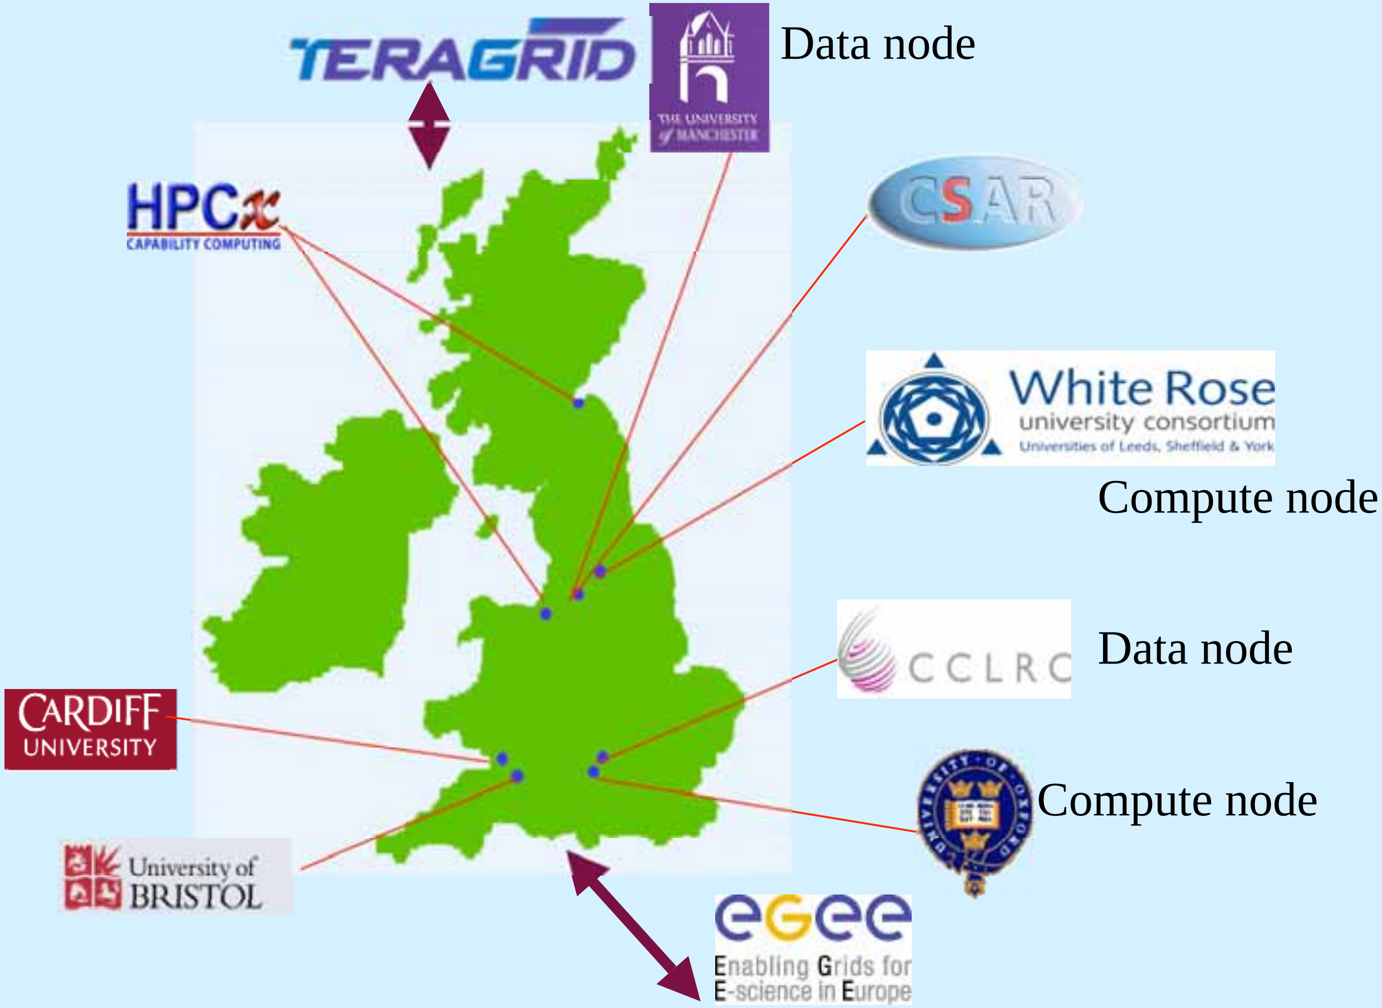

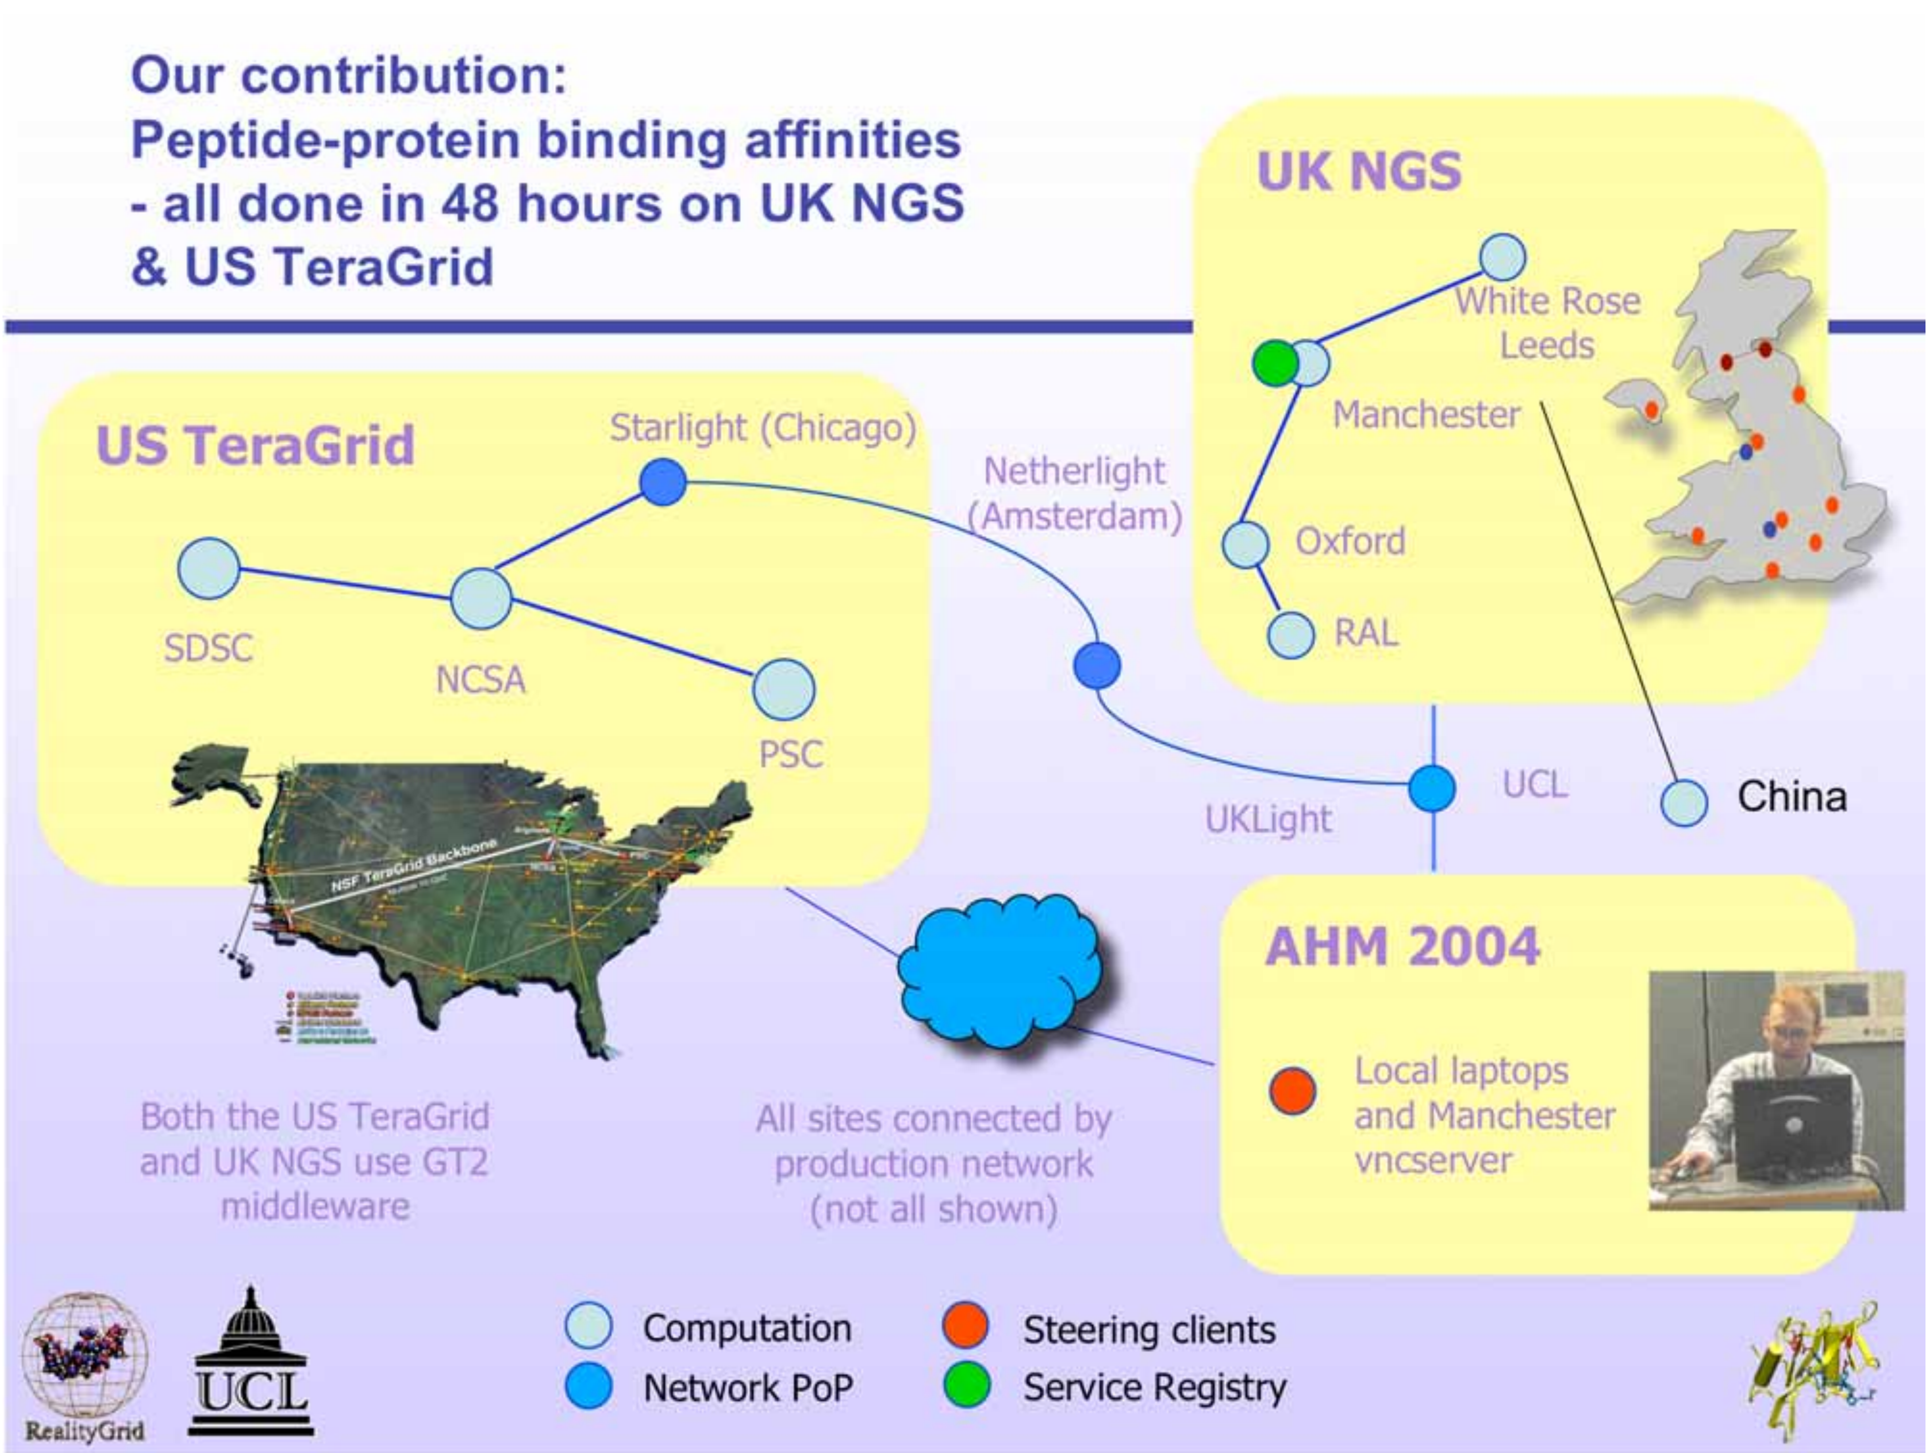
Our contribution:
Peptide-protein binding affinities
- all done in 48 hours on UK NGS
& US TeraGrid
UK NGS
White Rose
Leeds
Manchester
Oxford
RAL
US TeraGrid
Starlight (Chicago)
Netherlight
(Amsterdam)
SDSC
NCSA
PSC
UCL
China
UKLight
NSF TeraGrid Backbone
AHM 2004
Local laptops
and Manchester
vncserver
Both the US TeraGrid
and UK NGS use GT2
middleware
All sites connected by
production network
(not all shown)
Computation
Network PoP
Steering clients
Service Registry
RealityGrid
UCL

The ‘Shared Service’ Problem (1)

- A potential approach for achieving fault tolerance in a Grid/Web services environment is to invoke multiple functionally-equivalent services and to act upon the results returned from them, e.g. by comparison or voting.

- A problem for this fault tolerance approach, however, *is* that in most SOA models, the implementational details of a service are hidden from a client of the service.

- The only information available to a client is the service’s interface and – possibly – some QoS metadata.

- This is an issue as services that initially appear disparate may – during the course of their execution – invoke one or more identical, “shared” services.

The 'Shared Service' Problem (2)

- The result is that different services may use the same shared services behind the scenes, which may make common mode failure (CMF) much more likely.

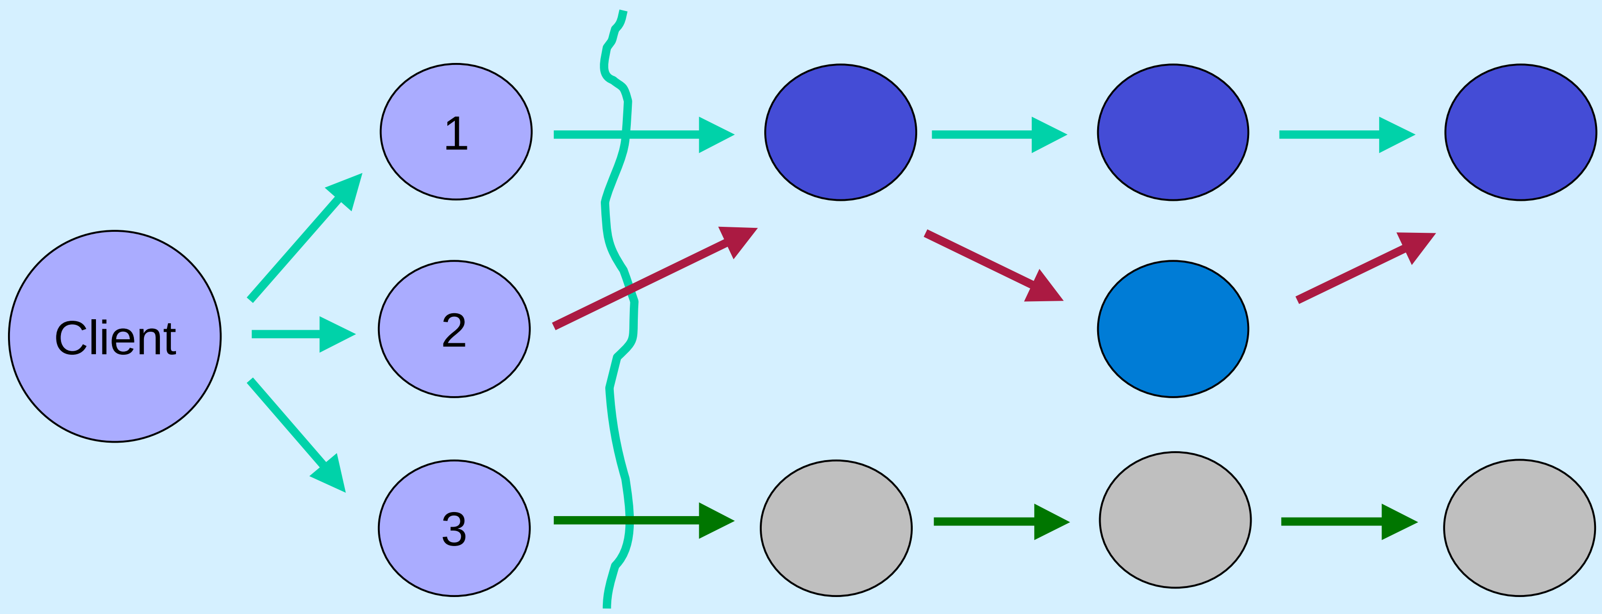

A Solution to This Problem

- One possible way of resolving this problem is to incorporate the technique of **provenance** in the fault tolerance approach used.

- Provenance is the documentation of the process that leads to a result.

- If we assume that data provenance is recorded, it will allow a fault tolerance scheme to build up a “view” of how each result it receives has been constructed.

- By possessing this view, a number of actions can be taken upon the results returned, e.g. weightings can be assigned to each service based upon how closely related it is to another service; services that have many common-dependencies can therefore have less “sway” in the voting algorithm used.

Weighted Voting

- In this “view”, s1 and s2 have 2 common dependencies, whilst s3 has no common dependencies.
- As a trivial example, we could therefore assign weightings of 0.5 to s1 and s2, and 1.0 to s3.
- In this case, should s1 and s2 agree, but s3 disagree with a result, then no overall “trusted” result will emerge.

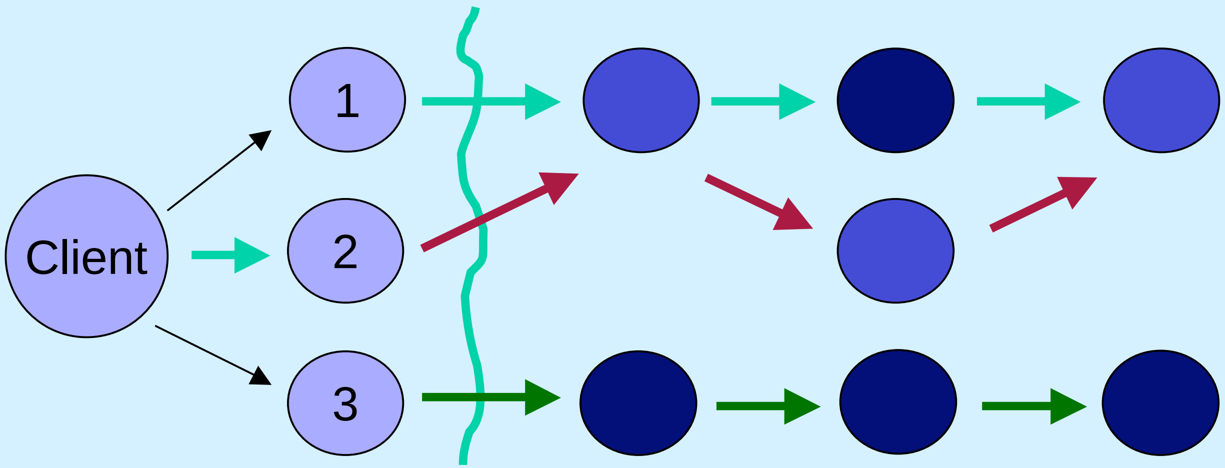

FT-Grid: A framework for achieving fault tolerance

- We have implemented a java-based framework that facilitates the creation of fault tolerance schemes based on diverse services. This is called ***FT-Grid*****.**

- The current implementation consists of both an API allowing developers to easily search for, invoke, and vote on services at run-time, and also a GUI to demonstrate the system.

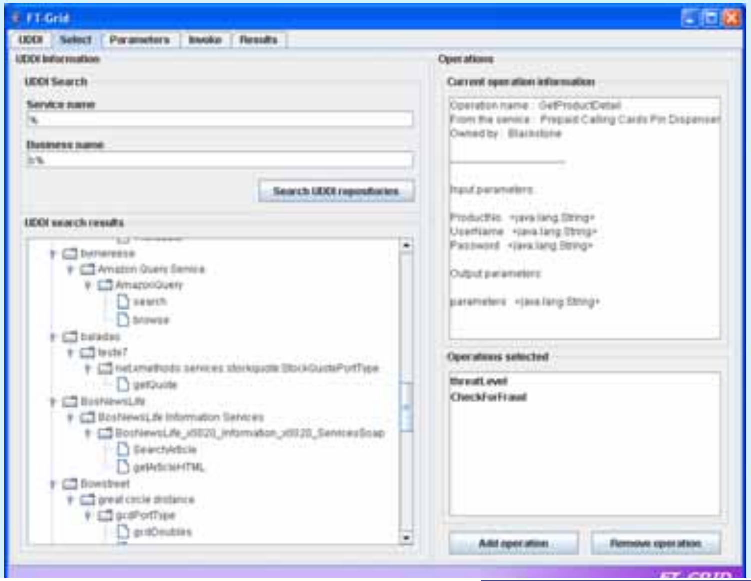

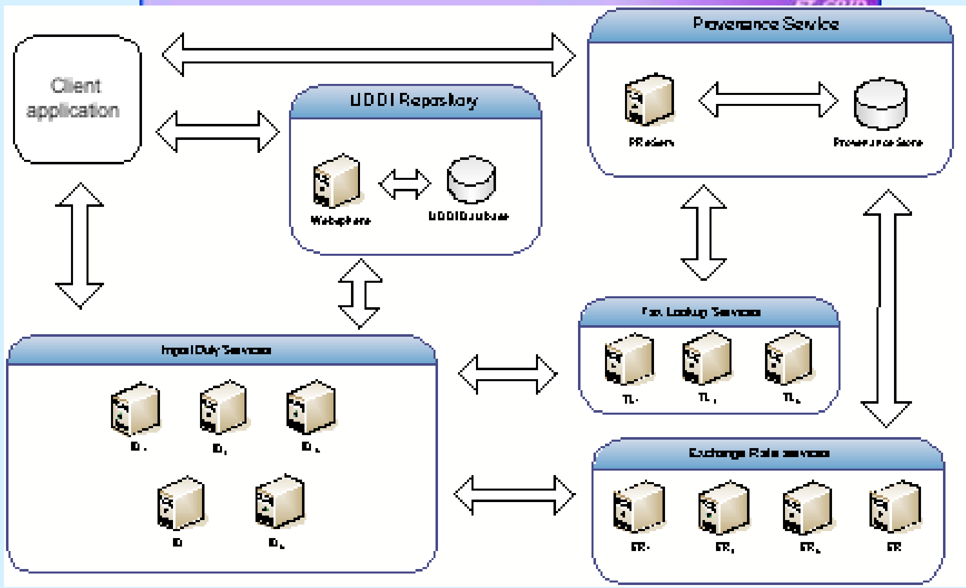

Comparison of Three Schemes

- Using FT-Grid, we developed a system that built up weightings based on the historical results of each service (the frequency with which a service's results agreed with the consensus). 15 Web services were involved and a Grid provenance system, called PASOA (developed at Southampton), was employed.
- We developed three systems in total:
  - A system without fault tolerance
  - A 'traditional' MVS system
  - A provenance-aware MVS system
- The traditional MVS system discarded results from services that had a weighting below a user-specified value, whilst the provenance-aware scheme discarded results where any service in a workflow fell below a user-specified value.
- This experiment yielded a large set of empirical data, and stress-tested both FT-Grid and the underlying infrastructure.

Some Experimental Results

- We performed 3 runs of 1000 tests on each scheme:

	Correct result	No result	CMF
Experiment 1 Run 1	828	172	-
Experiment 1 Run 2	858	142	-
Experiment 1 Run 3	822	178	-
Average	**836**	**164**	**-**
Experiment 2 Run 1	928	9	63
Experiment 2 Run 2	921	14	65
Experiment 2 Run 3	921	7	72
Average	**923.33**	**10**	**66.66**
Experiment 3 Run 1	996	4	0
Experiment 3 Run 2	990	10	0
Experiment 3 Run 3	996	4	0
Average	**994**	**6**	**0**

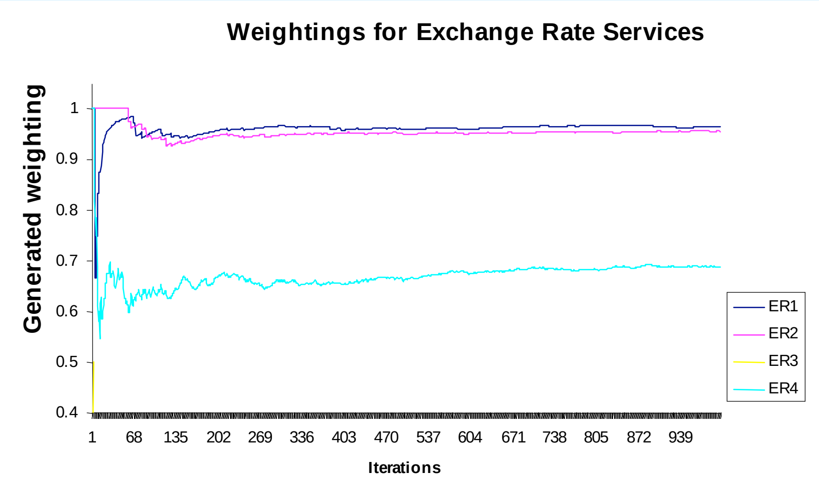

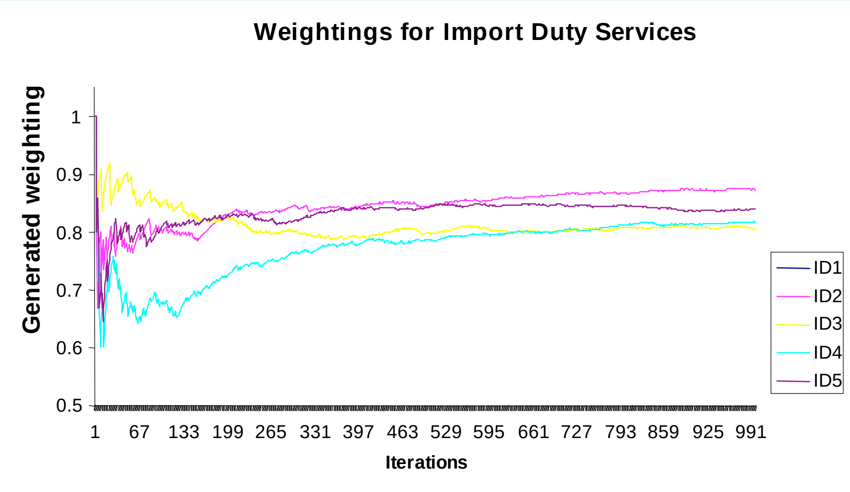

Brief Result Analysis

- The scheme without fault tolerance obtained a correct result in **83.6%** of all tests performed.
- The traditional MVS scheme obtained a correct result in **92.3%** of all tests performed, and a common-mode failure (CMF) occurred in **6.6%** of results.
- The provenance-aware MVS scheme obtained a correct result in **99.4%** of tests performed, and had no CMF.
- These results are encouraging, but it must be remembered that the test scenario was very simple, and in a more realistic environment (with more reliable services), the advantage of the provenance-aware scheme is likely to be reduced.
- We are making progress…

THE WHITE ROSE GRID
e-Science Centre of Excellence

Questions?