A Grid-based Distributed Database Solution for Large Astronomy Datasets

Helen X Xiang
Computer Science, University of Hertfordshire, UK
H.Xiang@herts.ac.uk

Abstract

The volume of digital astronomical data is set to expand dramatically over the next ten years, as new satellites, telescopes and instruments come online. For example, both the VISTA [1] and DES [2] programmes will yield databases 20-30 terabytes in size in the coming decade. Storing and accessing such large datasets will be challenging, especially as scientific analysis will require coordinated use of several of these separate multi-terabyte databases, since they will contain complementary data, typically from observations made in different regions of the spectrum.

We used these data to construct a Grid-based SDSS distributed model using Open Grid Services Architecture – Data Access and Integration (OGSA-DAI) [3] and Open Grid Services Architecture – Distributed Query Processing (OGSA-DQP) [4]. Finally we designed and ran a set of distributed queries against the Grid-based distributed SDSS model.

1. Introduction

Scientific data has increased significantly in size in recent years. Advances in the technologies of computing devices and scientific instruments are producing a huge amount of data in science research in almost all disciplines. There is a rising demand for effective ways of housing and managing data.

Dealing with intensive data archives is becoming a challenge for many scientists. The reality is that we have collected an enormous amount of data already, and are still collecting. In observational cosmology, for example, astronomy datasets contain observations of tens of millions of stars and galaxies now, and we expect the catalogues will grow to billions of objects in the next five to ten years. Today, many scientific databases, such as the *Sloan Digital Sky Survey* (SDSS), are terabytes in size. In the next five years, the *Panoramic Survey Telescope-Rapid Response System*

(Pan-STARRS) is expected to produce a science database of more than 100 terabytes [5]. Many more are aiming to generate petabyte volumes of imaging data per year [6], including the *Large Synoptic Survey Telescope* (LSST) project. The LSST project will obtain 30 terabytes of imaging data each night [7].

We believe that better data management is urgently needed in many scientific domains. Grid computing is becoming more and more important in scientific research. In particular the *Open Grid Services Architecture – Data Access and Integration* (OGSA-DAI) project is one of the UK e-Science projects that is working closely with the *Open Grid Forum* (OGF) [8] *Database Access and Integration Services Working Group* (DAIS-WG) [9] to develop Grid middleware to support data access and integration across separate resources in the Grid environment. OGSA-DAI supports a number of *Database Management System* (DBMS)s including Microsoft SQL Server and Oracle, and is therefore an appropriate vehicle to integrate datasets that are distributed over multiple sites. The OGSA-DAI project also provides a useful extension called the *Open Grid Services Architecture – Distributed Query Processing* (OGSA-DQP) that will provide service-based query processing. A distributed query engine can be used to evaluate queries and request data from the appropriate OGSA-DAI data servers.

This paper describes distributed database solution for large astronomy datasets. In particular it explores the relevant Grid technologies to access and integrate data of database systems. We aim to construct a working Grid-based distributed database solution for large astronomy datasets, combining this database solution with OGSA-DAI and OGSA-DQP to provide an open source, distributed database model for accessing and integrating data resources from different *Relational Database Management System* (RDBMS)s at distributed sites. Our initial focus is on the SDSS dataset. Our principal hypothesis is that sooner or later the flood of astronomy data will have to be stored in distributed database systems, and that Grid

IEEE computer society

technologies can contribute to managing the distributed data of astronomy and other sciences.

2. The SDSS Database

The *Sloan Digital Sky Survey* (SDSS) is a project that has built a very detailed digital map of the visible stars and galaxies in the night sky [10]. The SDSS is the first wide-area survey to use electronic light detectors, producing images substantially more sensitive and accurate than earlier surveys, which relied on photographic plates. It is one of the most ambitious astronomical surveys ever attempted. A 2.5 meters SDSS telescope at Apache Point Observatory near Sunspot, New Mexico, in the United States is dedicated full-time to the SDSS, and collects images almost every night (depending on the weather). Once the survey is finished, the SDSS data will map a quarter of the whole sky in detail. The data is made available electronically to the scientific community and the general public, both as images and as digital catalogues of all the objects discovered. With the release of DR6 in June 2007, data for approximately 287 million objects in the sky has been made available so far [11] [12]. The SDSS database mainly consists of photometric objects and all the spectroscopic information. Once entered into the database, data will not change.

The SDSS database is deployed in a Microsoft SQL Server 2000 instance. The SDSS DR5 database contains 85 user tables, 44 user views, 164 stored procedures and 173 user-defined functions, which enable advanced data searches and query optimisation. Indexing is used automatically in the SQL query optimiser to speed up the performance of queries.

So far there have been six major data releases from the SDSS project. In the four years since the first data release (DR1) up to DR6, data volume has grown from 2.34 terabytes of image data and 0.46 Terabytes of catalogue data to 10.00 and 4.00 terabytes respectively.

3. A Distributed SDSS Database System

The main objective of our project is to distribute the BestDR5 database between two different sites in Portsmouth and Manchester. We migrated the SDSS database from Microsoft SQL Server to Oracle as a preliminary to distributing the dataset – because Oracle can be deployed on more Grid platforms, and because we wanted to experiment with heterogeneous distributed version of the SDSS databases.

After partitioning the SDSS dataset, half of the SDSS dataset was transferred to remote storage on Manchester *National Grid Service* (NGS) site (server:

Vidar) and deployed to a SDSS Oracle database (SAND) there; another half of the SDSS dataset (DR5one) was stored on the SDSS server (Ace) in Portsmouth, in the form of Microsoft SQL Server. The result was a version of the SDSS database distributed over the sites – Manchester and Portsmouth. To our knowledge, this is the first distributed database of its kind in astronomy.

We installed OGSA-DAI WSRF 2.2 instances at both sites – Ace and Vidar – and deployed an OGSA-DAI data service on each host in Apache Tomcat. We then configured the OGSA-DAI data service resources to expose the SDSS databases DR5one and SAND via those OGSA-DAI data services. The resources IDs for the exposed SDSS databases were dr5one_ace and man respectively.

The SDSS BestDR5 database was now distributed across two sites – Ace (SQL Server) at Portsmouth and Vidar (Oracle) at Manchester – and exposed as data resources under the OGSA-DAI middleware.

To be able to query the distributed SDSS BestDR5 database, we installed the OGSA-DQP 3.2 (Tech Preview) toolkit on the distributed sites. On each host, we deployed an OGSA-DQP evaluator service in a Tomcat instance that is separate from the Tomcat used by the OGSA-DAI data service previously deployed.

An OGSA-DQP coordinator can be installed in either of the distributed SDSS hosts on the OGSA-DAI Data Services deployed earlier on, or it can be installed on a separate host. The client application can then interact with the OGSA-DQP coordinator service to create an OGSA-DQP coordinator instance for executing queries. Figure 1 summarizes the architecture of the distributed SDSS DR5 database system developed in this section.

We did a set of experiments we with our modified version of OGSA-DQP. These experiments took the form of a series of queries against the distributed SDSS database, of gradually increasing complexity, culminating in some distributed join queries over the geographically distributed, terabyte sized SDSS dataset. See [13] for details.

In Figure 2 we present a scatter plot of "maximum data volume" against time. Data volume plotted for each point is the maximum of the internal and external volumes[1], so it is an order of magnitude estimate of the

---

[1] *External volume* means the number of "items" returned by the query—rows times columns. *Internal volume* is similar, but it is an estimate of the number of items exchanged *internally* during the execution of the OGSA-DQP query

amount of data exchanged over the network while executing the query.

This is a log-log plot, and the data values span several decades. We cannot trust the detailed values, but we will try to read something from the orders of magnitude. The graph in Figure 2 also shows a "fit" of a linear model to the data, of the form

$$T = m \times V + c \qquad [1]$$

with

$$m \approx 7 \times 10^{-4} \text{ seconds/item}$$
$$c \approx 7 \text{ seconds}$$

Nearly all data points are within a factor of 2 or 3 of the values predicted by this model. So it is a reasonable fit to orders of magnitude.

An important conclusion from this is that the bottleneck in the timings is probably not the network bandwidth – so the bottlenecks are *probably not anything inherent to distributing the database*. Formula 1 implies that processing each item (like an integer or a double precision number) takes almost a millisecond. From a bandwidth point of view we expect it would be possible move data around orders of magnitude faster than that.

Most likely these orders of magnitude are accounted for by features in the implementation of OGSA-DQP (which is basically a prototype in its current form). We hope that with improvements to the software it should be possible to gain further orders of magnitude improvement in performance, *making distributed scientific databases a much more realistic proposition*.

4. Conclusion

This paper presents a Grid-based distributed database solution for large astronomy datasets: the SDSS distributed database system, and a working SDSS OGSA-DAI/OGSA-DQP model. We informally analyse the query execution times (Figure 2) and draw some qualitative conclusion based on our experiments. Although we made significant optimisation to OGSA-DQP in the course of our work, the results show that further optimisation will be needed. This could be a good starting point for a new kind of distributed database for astronomical databases.

5. References

[1] Alistair M McPherson, Andrew Born, Will Sutherland, Jim Emerson, Bryan Little, Paul Jeffers, Malcolm Stewart, John Murray, Kim Ward, "VISTA: Project Status", in Ground-based and Airborne Telescopes, edited by Larry M. Stepp, Proc. of SPIE Vol. 6267, 7, 2006

[2] Dark Energy Survey project home page. http://www.darkenergysurvey.com.

[3] M. Antonioletti, M.P. Atkinson, R. Baxter, A. Borley, N.P. Chue Hong, B. Collins, N. Hardman, A. Hume, A. Knox, M. Jackson, A. Krause, S. Laws, J. Magowan, N.W. Paton, D. Pearson, T. Sugden, P. Watson, and M. Westhead. "The Design and Implementation of Grid Database Services in OGSA-DAI", *Concurrency and Computation: Practice and Experience*, Volume 17, Issue 2-4, Pages 357-376, February 2005.

[4] Alpdemir, M. Nedim et al. "OGSA-DQP: A Service for Distributed Querying on the Grid", *Lecture Notes in Computer Science*, Vol. 2992, pp. 858-861, 2004.

[5] N. Kaiser, "The Pan-STARRS Survey Telescope Project", *Advanced Maui Optical and Space Surveillance Technologies Conference*, 2007. http://adsabs.harvard.edu/abs/2007amos.confE...9K

[6] Alexander S. Szalay, Jim Gray, and Jan vandenBerg, "Petabyte scale data mining: Dream or reality?", http://arxiv.org/abs/cs/0208013, 2002.

[7] Z. Ivezic and J.~A. Tyson and R. Allsman and J. Andrew and R. Angel and for the LSST Collaboration, "LSST: from Science Drivers to Reference Design and Anticipated Data Products", 2008.

[8] Open Grid Forum (OGF). http://www.ogf.org/.

[9] Database Access and Integration Services WG (DAIS-WG). https://forge.gridforum.org/projects/dais-wg.

[10] York, D.G. et al., "The Sloan Digital Sky Survey: Technical summary", *Astronomical Journal*, Vol. 120, pp. 1579–1587, 2000.

[11] Alexander S. Szalay, Jim Gray, Ani Thakar, Peter Z. Kunszt, Tanu Malik, Jordan Raddick, Christopher Stoughton, and Jan vandenBerg. "The SDSS SkyServer—public access to the Sloan Digital Sky Server data", *SIGMOD Conference*, pages 570–581, 2002.

[12] The SDSS Data Release 6 (DR6). http://www.sdss.org/dr6/start/aboutdr6.html.

[13] H. X. Xiang, *A Grid-based Distributed Database Solution for Large Astronomy Datasets*, University of Portsmouth Ph.D thesis, Portsmouth, UK, February 2008.
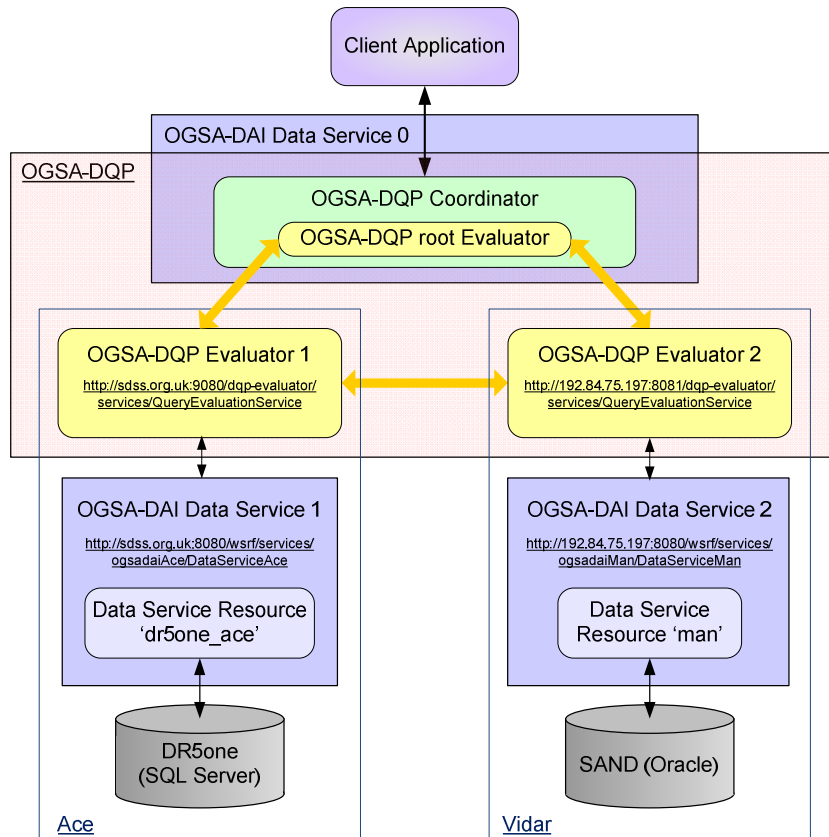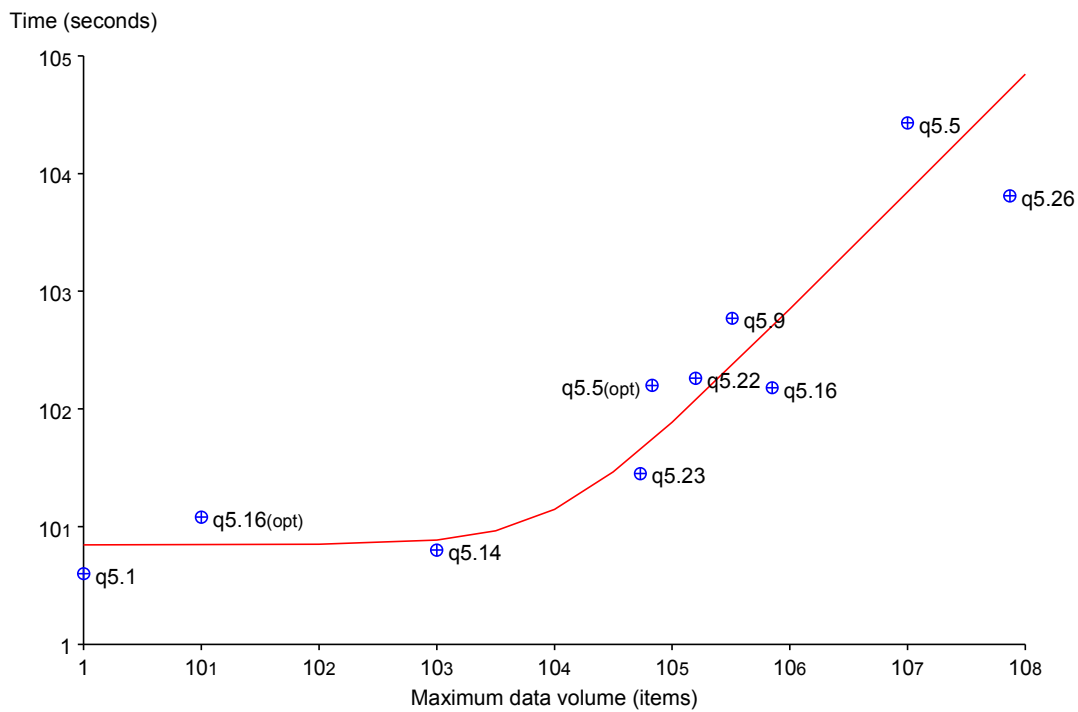
Figure 1: Distributing the SDSS database between two sites: Portsmouth and Manchester

Figure 2: Timing scatter plot