Music GRID – a Collaborative Virtual Organization for Music Information Retrieval Collaboration and Evaluation

Matthew J. Dovey
Oxford University
13 Banbury Road
Oxford, UK
matthew.dovey@ccc.ox.ac.uk

ABSTRACT

WebServices offer a methodology and technology for large scale distributed systems. GRID offers a methodology for building Virtual Organisations to tackle tasks via distributed technologies such as WebServices. Projects such as OMRAS[1] and Greenstone[2] are establishing modular frameworks for MIR testbeds. A Music GRID could take these frameworks and implement them in a distributed nature where participants make available algorithms for processing formats (e.g. feature extraction), converting formats, searching formats etc. on their own servers as GRID Services. Researchers can then combine these various services with their own approaches in order to perform retrieval experiments and evaluation.

1. OMRAS FRAMEWORK FOR MUSIC INFORMATION RETRIEVAL

1.1 Overview

Within the OMRAS project (Online Music Retrieval And Searching)[1], we have devised a generic framework of how a Music Information Retrieval (MIR) system might look[3]. The UML diagram for this is shown in Figure 1.

The *DocumentRetrievalEngine* implements a particular IR mechanism for music document retrieval (i.e. retrieval of documents which have some musical relevance to the query). Initially musical documents are added to this engine (step 1). Depending on the implementation of this component, at this stage it may build indexes of the musical structures.

A *QueryFormulator* allows an end user to input a particular musical query. Different implementations might include an onscreen piano keyboard, an interface that would take audio input (converting to a score if necessary) or a musical score editor.

The query is then passed onto the *DocumentRetrievalEngine* which builds a ranked list of relevant documents.

The *ResultListManager* handles the display of this list by using an *Output* object for each result. An *Output* object displays the result; different implementations of *Output* objects might include displaying the metadata for the document (such as composer and title), displaying a piano roll version of the score, displaying the musical score playing the audio of the score and so on. The *ResultListManager* handles how the displays are organised: in the case of piano roll displays they may just be displayed as an on-screen list; however, we are also looking at 2- and 3-dimensional techniques for displaying result sets.

In some cases, the *DocumentRetrievalEngine* may not only be able to return relevant documents but also to indicate within a document the relevant passages (typically exact or close variations of the query). However, this is dependent on the techniques used by the retrieval process. In the general case, it may be necessary to apply another search technique on the retrieved documents to locate the relevant passages. This is the role of the *PassageRetrievalEngine* - to locate passages relevant to the query within a document selected by the user.

1.2 WebService Implementation

The OMRAS Framework was originally devised as a set of Interfaces for Java which would enable the various research teams involved with the project to implement the various components they were working on (such as audio to musical score conversion, document retrieval indexing and passage level retrieval techniques) in a way that would allow them to be easily linked together into a single system. Initially it was intended that this hybrid system would run on a single computer.

However the advent of WebServices offers a different approach for implementing the framework which is more suited to the distributed nature of the various teams involved in the OMRAS project. Many of the components in the OMRAS framework can be mapped to WebServices – for example a *QueryFormulator* WebService may take a query and convert it to a suitable form (e.g. audio to score conversion); a *DocumentRetrievalEngine* WebService may then take a suitable query and return a set of relevant documents using a particular algorithm. In this way the various teams within OMRAS could mount the results of their research (be it audio to score conversion, document retrieval algorithms etc.) on their own platforms and these can then be linked together over the internet as examples of complete systems.

2. GRID IR

2.1 GRID

The GRID offers a methodology for building large scale distributed systems in particular for the purposes of performing online experiments and research. A GRID would typically consist of a large number of computing resources (which might be data,

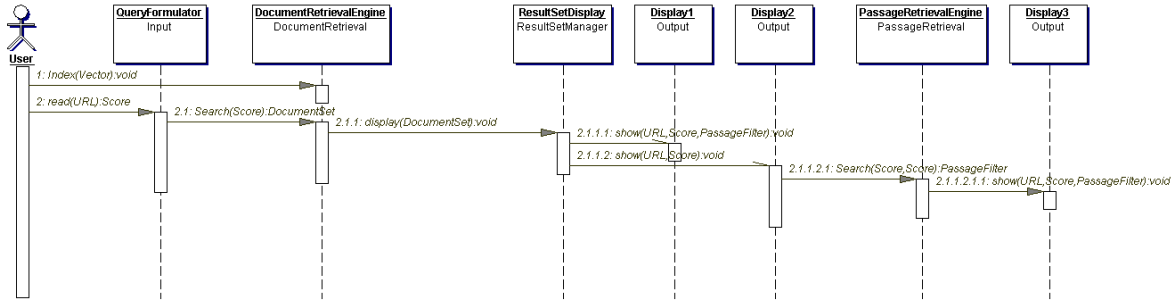

Figure 1 - UML Diagram of the OMRAS Framework

processors, storage, interfaces to remote devices) which can be found and located using directory services. Authentication to use such resources is provided via certificates. A Virtual Organization would bind together a collection of these resources for the purposes of solving or researching a particular problem. So, for example, a group of astrophysicists might build a Virtual Organization which would link together radio telescopes, distributed storage and supercomputers in order to process signals from space. Currently GridServices, namely the use of WebServices to build GRIDs and Virtual Organisations, is being undertaken in the Global Grid Forum[4].

2.2 GRID IR

Within the context of the Global Grid Forum, there is currently an activity to establish a GRID Information Retrieval Working Group[5]. The goals of this group are two-fold. Firstly, to look at how GRID applications can benefit from and require information retrieval techniques and technologies. Many GRID projects involve the handling of large amounts of data and hence are in need of sophisticated data handling including data mining, information retrieval and knowledge management. The second goal of this working group is to establish how information retrieval can benefit from GRID techniques and methodologies. This is closely related to the first in that the amount of data many GRID projects are dealing with (typically in the terabyte or petabyte range) is such that multi-processor and distributed processor architectures are needed.

The GridService architectures being developed within this working group are in many cases very similar to the WebService architecture being devised within the OMRAS project. This is not too surprising as the GRID IR architecture is a meant to be a generic IR architecture and the OMRAS architecture could be viewed as a domain-specific implementation of such a generic architecture.

3. MUSIC GRID

3.1 Music GRID Resources

The combination of the work of the GRID IR Working Group and the OMRAS architecture could lead to a basis for an architecture for a Music GRID. Such a GRID would have the following GridServices:

- Conversion GridServices – which allow the conversion of musical data between formats (e.g. MIDI, Humdrum, AIFF, Finale, JPEG image of score etc.)
- Data services – which provide data-sets or evaluation testbeds
- Indexing services – which index a dataset for document level retrieval
- Pattern matching services – which allow the location of a musical pattern (or variation) of the query in a musical document
- Pattern induction services – which locate recurring patterns or themes within a musical document
- Rendering services – which convert musical documents into audio or visual representations

A number of different implementations of each of these GridServices could be made available by different research groups implementing algorithms and approaches developed (or under development) by those groups. Making these available as Web Services would offer a number of advantages:

- Researchers can evaluate different combinations of the various components of a music retrieval system to test for example which IR technique works better with which audio transcription technique
- Researchers can combine their own components with those from other groups i.e. can test the performance of their work without having to implement complete systems
- Researchers would be working with the latest implementations of such services (as these would be running within the developing research group)
- Research groups could make available research that is experimental or in development in a working form without having to release the details of the algorithms involved (the code would be running on a locally controlled machine)

Making these available as GridServices offer a number of additional advantages:

- In addition to allowing data to be passed between different services, the GRID also allows services to run on data in situ. There are some security issues still being worked out within the GRID community.

However, within the context of a Music GRID this would allow indexing algorithms to be tested on large music collections without having to move the collection itself across the Internet (moving the algorithm code rather than the data). This would not only save time and allow those with lower bandwidth to participate but would also reassure copyright holders of the safety of their data (the data would never leave the licensed host machine)

- A Music GRID could take advantage of the large scale storage and processing power of GRID technologies in performing experiments (particularly important for indexing large datasets and digital signal processing applications). The components outlined above may themselves make use of GRID components.
- The GRID would offer the extra security of its certificate system for authentication and authorization (which might help reassure copyright holders of material used for evaluation).

3.2 Music GRID Portals

Within the GRID, a community or Virtual Organization would use a GRID portal[6] to provide a web interface for managing the various resources. Through such a portal, the resources could be found, linked together in an appropriate manner to build a combined service or experiment. Within a Music GRID, such a portal would act as a web based workbench for linking the various components outlined in section 3.1 into a working retrieval system and then entering queries. Such a workbench may also have features such as scripting to allow a series of evaluation experiments to be performed. The use of testing frameworks such as jUnit would also allow automated testing and evaluation of systems.[7]

4. CONCLUSIONS

The combination of WebServices and the GRID (namely GRID Services) could allow new methods of online collaborative work in Music Information Retrieval and could offer a base architecture for evaluation experiments. It offers a number of advantages including real-time collaborative work and offers some solutions to potential problems of protecting musical data licensed for the purposes of evaluation.

5. REFERENCES

[1] OMRAS website. http://www.omras.org

[2] Greenstone website: http://www.nzdl.org/greenstone

[3] Dovey, M et al. Overview of the OMRAS Project: Online Music Retrieval and Searching. Forthcoming edition of the Journal for the American Society of Information Science and Technology

[4] Global Grid Forum. http://www.gridforum.org

[5] GRID IR. http://www.gridir.org

[6] GRIDPortal. http://dast.nlanr.net/Projects/GridPortal/

[7] jUnit. http://www.junit.org