

# *SMaTBaM!* - Serving Time-Based Media in the Performing Arts

Carola Boehm, Stephen Arnold  
Department of Music, University of Glasgow  
carola@music.gla.ac.uk, stephen@music.gla.ac.uk

## Abstract

The *SMaTBaM!* project has established a working example of a system to serve massive time-based media in the area of the Performing Arts. It demonstrates the searching and serving of massive multimedia and real-time data (such as musical performances), using a web-based, interactive front-end.

## 1. Introduction and Background

The Performing Arts Data Service (PADS), based at the University of Glasgow, is one five Service Providers appointed by the Arts and Humanities Data Service (AHDS)[1] and funded by the Joint Information Systems Committee (JISC)[2] of the UK's Higher Education Funding Councils. The AHDS's mission is to co-ordinate access to, and facilitate the creation and use of, electronic resources in the arts and humanities by offering a range of services. It will encourage scholarly use of its collections and make information about them available through an on-line catalogue. The AHDS provides a single gateway for arts and humanities scholars wishing to search for data-sets across various discipline areas. The service providers' databases interoperate with other databases within the AHDS and beyond via Z39.50[3], and searching will be available via the Web. In order to achieve meaningful search results, data from all the service providers is indexed with Dublin Core metadata.

The Performing Arts Data Service's role within this framework is to support research and teaching in UK Higher Education by collecting and promoting the use of digital data relating to the performing arts: music, film and video, broadcast arts, theatre and dance. The PADS differs from the other service providers in that it has a particular concern with data consisting of, and representing, time-based media. The results of two recent projects have had a major influence on the system of the PADS as it stands today. The *NetMuse* Project[4] was a project developing web-based music courseware for delivery over the ATM-based Scottish Metropolitan Area Networks (MANs). This included development of a Java-based audio player [Malloch and Pflücke, 1997] for streaming full CD quality music, further developed as part of the *SMaTBaM!* project [Boehm 1997].

The *SMaTBaM!* project researched and implemented a prototype of a system which could provide the technical infrastructure for the PADS, including the means of delivery of time-based data as well as issues concerning the storage and retrieval. This prototype has been scaled up and now forms the basis of the PADS system which consists of two Silicon Graphics (SGI) Origin 200 servers: one is a media server streaming audio and video using SGI MediaBase software; the other runs an object-entented database with a web-gateway (Hyperwave

Information Server) which stores both the non-time-based data and the metadata of the material on the media server. This solution combines the demanding performance of a media server with the advanced database features required for modelling and handling big collections of data. Additional gateways have been set up to secondary remote Hyperwave servers via HGI-CSP, relational database management systems via Net-Dynamics and library databases and catalogues via a newly implemented Hyperwave-Z39.50 gateway. The user therefore accesses information from different locations, different platforms and different kinds of databases, via browseable webpages which guide him through the information space with navigational aids, such as discipline-specific and discipline-independent searching templates, knowledge-domain-related browsing hierarchies, depiction of relationships between resources via metadata and other hyperlinks.

## 2. Information system requirements for Performing Arts data on the Web

As the internet becomes the platform, the browsers become the operating system and applications become services. A digital library project set in the performing arts has to define new methods of storing and distributing time-based data in order to serve information across wide-area networks of high quality and in vast quantity. It also requires solutions of a more philosophical nature, namely how interfaces have to be set up and how information should be represented in order for massive time-based data to be handled as intuitively as possible. Information management services will always have to deal with "the three Is": Information Structure, Information Representation and Information Access.

### 2.1. Nature of the data in the performing arts

A collection dealing with Performing Arts data comprises both secondary resources (materials about the performing arts, moving image and sound-based media) and primary resources (the digitised multimedia objects themselves). As data compression and transmission technologies develop in the future, it is the service's aim to facilitate the real-time access of video clips, sound files, movies, musical performances and multimedia productions - both primary and secondary resources. It is desirable that a collection may be expanded by collections of other service providers holding resources in the same field but at the

same time maintaining a “one-stop shop” in accessing time-based media resources. This distributed resource environment allows the option of other collection holders keeping and maintaining their collection physically in their own repository, while access is handled by a central access point [5].

A performing arts resource collection encompasses a wide range of different disciplines, starting with the disciplines of music and film and stretching further toward dance, theatre and the broadcasting arts [6]. The resources as a whole can be characterised as a) being made out of different types of data, b) containing differing complexities of data, c) possessing different relationships, and d) being time-based in their nature.

### 2.1.1 Different types of data

As with all multimedia related systems, all the “usual” data types are involved from sound, video, text, image and binaries. Storing them in a certain way provides us with a more complex entity of data types: html, sgml, mpeg, wav, gif, jpeg, java, etc. It is certain that these data formats will evolve further in number and content. The use of different formats in a system should therefore be a means but not a solution. In other words, to minimise the danger of storing data in standards that might not be supported in the future, much thought should go into separating the content of a resource from its presentations. To be able to store a resource in the highest quality possible, combined with the ability to convert it into formats suitable for a certain purpose, or added formats in the future, is to provide an open and flexible system with maximum compatibility in the long term.

### 2.1.2 Differing complexity of data

Whereas video and images might be stored largely as single binary data-objects, music, theatre and the broadcasting arts could involve the storing and accessing of highly structured data, presenting complex objects or composite objects. In some cases, it might be hard to distinguish which is the real, the original resource, and which is a composite part of it. If one accepts the fact that the content of a resource might be of complex or composite nature, then the step towards devising a way to store it as such is not far. Technologies are needed that offer the ability to depict, represent, access, store and manipulate complex structures in their complex “Gestalt”. A broadcasting feature, as one resource, might encompass video data, sound data, and text data and still be one work of art.

Future collections may not remain in their binary form and much of our present resources have never been in the “Gestalt” of one entity. Java Applets,

WebObjects and other distributed object environments are already being used by artists to create works made out of many components and having many facades. Also, existing resources, which have been traditionally stored as metadata in catalogues while their real content is being stored as artefacts in shelves, cassettes, or discs, often comprise multiple entities. This suggests that normal library catalogues and conventional relational database management systems are not sufficient: object-oriented or at least object-based information system technologies need to be deployed.

### 2.1.3 Different relationships

Assuming that we have objects stored in a persistent way, the access and search results are influenced by the context these objects are in. The mapping of content and context into a digital world means defining and storing different kinds of relationships between objects. Examples for generic implementations and standard definitions can be found in OMGs Object Request Brokers and their Relationship Service Specification for distributed objects [OMG 1997] or in the Knowledge Interchange Format of the Laboratory for Advanced Information Technology [Finin and Labrou 1997]. Relationships can be of numerous variety. For example, five basic relationships widely used in information systems are:

- a) *Inclusion* - one object is included in another object (e.g. a file in a folder, a certain sound used in a composition, a note in a bar)
- b) *Inheritance* - one object inherits the characteristics of another object (e.g. all service provider users have read rights, these might be inherited down towards the developers of collections, who also have write rights; or, as a third example, all sounds stored at high quality inherit the characteristic of being served out over ATM network only).
- c) *Association* - one object is associated with another object (e.g. Mendelssohn's composition *Fingals Cave* is associated with the geographical rock formation of Staffa. Another example would be that two pages can be associated with each other in form of a sequence. One page should follow the other in a certain context as for instance a book, course, slide show, score etc.).
- d) *Attributes* - an object contains certain attributes, or certain characteristics which describe its state of being or its internal structure (e.g. all objects in the PADS archive have the attribute Dublin Core, where the Dublin Core object itself has 15 further attributes defining the elements of the Dublin Core).
- e) *Web Links* - Web-links can be thought of being a realisation of a certain kind of association in a web environment.

## 2.2 Metadata

Whereas relationships belonging to the categories of inclusion and inheritance are implemented directly into the PADS system, attributes may be thought of as being the metadata of the objects. Unlike a normal library system, in which only metadata about objects is stored and thus not held digitally or stored separately, attributes in the PADS system belong directly to the objects described by them. In order to facilitate an interdisciplinary approach of resource discovery, the PADS, within the wider context of the AHDS, had to seek a standard way of describing resources across different disciplines, across different types of resources and across the different service providers.

During 1997, the PADS engaged in various activities to determine how best to facilitate resource discovery in an on-line setting. Specifically, the PADS looked at the metadata standard known as the Dublin Core [7] and how it could be applied as a tool to describe resources relating to a performing arts context. The PADS work [PADS Metadata 1997], which formed part of a series of activities in all the arts and humanities discipline areas represented by the AHDS, was conducted under the auspices of the AHDS and the UK Office for Library and Information Networking [8] with funding from JISC. One of the attractions of the Dublin Core metadata set is its simplicity - the Dublin Core was originally intended to be used by non-specialist authors to describe World Wide Web documents. The Dublin Core consists of 15 basic elements (title, creator, subject, description, publisher, contributor, date, type, format, identifier, source, language, relation, coverage, rights) to which the AHDS workshop series [9] and other initiatives from the library and information community have proposed some qualifiers and amendments to some of the definitions. Most of the problems related to the use of the Dublin Core can be put down to its design of describing textual documents on the web. A clear example of this is the author element, which works only for written documents. It does not work for music or many other disciplines, in which there are many more and different creators, which cannot be described and would be even confusing in being described as being the author. Changing this attribute to creator makes it more elegant, but does not solve other problems, as for instance:

- a) Sub-elements or schemes of the Dublin Core are not yet standardized, as are the values in their syntax. (for example date, controlled lists, etc).
- b) Absence of any distinction between the digital item, which is represented in the system, and the physical item it might represent. Thus digital li-

braries using the Dublin Core always face the question of which metadata is actually being used, the one representing the digital object (for instance an image of a manuscript) or the physical object itself (the manuscript). Both sets of information might be important for a specialized user, but the Dublin Core has no option for this possible double existence of the object it is supposed to describe.

- c) Certain basic elements are problematic in their use in an interdisciplinary context, such as subject and coverage [10].

## 2.3 Time-Based Data

The common factor of many prospective Performing Arts resources is their time-based character. Storing and accessing time-based media requires special care in the storage and delivery of the objects. Solutions are needed to store information in its inherently complex form on the server side, to transmit these information packages in real-time and with high-quality over a wide-area network, and to provide a user interface able of accessing and utilizing the resources intelligently. For a high-quality service, four types of time-based material, all requiring real-time access, can be identified:

- a) *large binary data objects*: such as sound or video - streaming binary data combined with using a guaranteed bandwidth to ensure no glitches or breaks. This requires:
  - high performance networks providing high bandwidth and guaranteed quality of service;
  - client-server software tools to provide the streaming;
  - high-performance media servers, and high-end client workstations.
- b) *subsets of large binary data objects*: playing just a part of a sound or video
- c) *two or more parallel large binary objects*: such as synchronisation of multiple audio streams, requires intra-stream and inter-stream synchronisation to maintain the temporal relationship between multiple streams [Robertson 1997 and 1998], e.g. 'lip sync' in film and tv, where sound and vision tracks are often recorded on different media.
- d) *complex objects*: such as MAX music scores, more complex Java applications, or sound-sound combinations require a fast and time-co-ordinated access of all the composite parts of an object: the synchronisation of multiple, periodic, logically independent streams of arbitrary type.

## 3. The PADS system

### 3.1. PADS System Infrastructure

A goal of the PADS service is to provide interoperability with other collection holders by conforming to and implementing relevant standards.

Current usage of multimedia digital resource collections include broadcasting, music/video archives, record companies and libraries. It must be taken into account that collections are stored in different storage media, ranging from simple file systems, through relational database management systems to the growing number of object-oriented database management systems. In addition, a large number of music catalogues in a variety of formats has also to be made accessible. Between library and library-like catalogues, an implementation of the Z39.50 protocol (version 3, 1995) is sufficient. For interfacing catalogues with relational databases, a Z39.50 -SQL interface is required. There are currently very few relational database vendors who have implemented a Z39.50 support, one reason being that their "interoperability protocol" has been SQL, which has been universally accepted and implemented by almost all of the database vendors.

Discussions have already taken place to extend the Z39.50-1995 protocol with SQL [11]. From here, it is logical step and a matter of time to stay interoperable with the present database generation which is based on object-oriented technologies, and has defined an object-query language (OQL) and an object-definition language (ODL)[12]. With the prospective widespread use of digital libraries, object-oriented database management systems will become a major means of storing, accessing and using complex, multimedia data objects. Assuming a basic interoperability of different collections containing digital multimedia objects, the underlying transfer protocol will have an influence on the performance, quality and representation of the objects to be delivered. Using a stateless protocol, such as http, means that only one object can be delivered per session. The connection closes after each document is delivered, thus losing all the information of the former session.

In devising a secure, distributed system, with collections stored in different locations, access handled from a central gateway and user-access ideally being controlled to the level of write,-read-and-execute rights to single objects and collections, stateless protocols can be a problem. Solutions lie in the underlying existence of user-rights management, such as a database management system able to control the access of many users in dependency of objects or collection of objects, and/or the use of a stateful protocol such as Z39.50 or Hyperwave's HG-CSP.

PADS Hyperwave Information Server is able to handle these security issues, as well as offer an expandable protocol layer. Relational database management systems are interconnected via Perl Database Modules (DBI:DBD-ODBC, DBI:DBD-Oracle, etc). In co-operation with the AHDS, Index Data

(Denmark) and Hyperwave R&D GmbH (Germany), a generic Z39.50-Hyperwave gateway has been implemented, translating incoming Z requests into the object-query language used by Hyperwave. With this gateway, the requirement to access both library catalogues and relational database management systems (as well as other object servers) has been met.

## References

- [AHDS Metadata 1997] Greenstein, D. and Miller, P. (1997). *UKOLN/AHDS Metadata Workshop Series*, Kings College, London (1997) <http://ahds.ac.uk/public/metadata/discovery.html> , 1998/07/01
- [Boehm, SMatBaM 1997] Boehm, C. & Malloch, S. (1997). SmaTBaM, Report on the evaluation process of system needs and demands of serving time-based media in the area of the Performing Arts, Glasgow 1997, <http://www.music.gla.ac.uk/HTMLFolder/Research/SMaTBaM.html>
- [Finin and Labrou 1997] Finin, Labrou & Mayfield,(1997). *Laboratory for Advanced Information Technology: the Knowledge Interchange Format, KIF* , <http://www.cs.umbc.edu/kse/kif/kif101.shtml>
- [Malloch , Pflicke, 1997] S. Malloch, S. Arnold, T. Pflicke, *Using Java to stream audio over ATM*, Proceedings of the ICMC, Thessaloniki, 1997
- [OMG 1997] OMG (1997). *Relationship Service Specification for distributed objects*. OMG <ftp://www.omg.org/pub/docs/formal/97-12-16.pdf>
- [PADS Metadata 1997] Duffy, C. & Owen, C. *UKOLN/AHDS Metadata Workshop: Moving Image Resources and Sound Resources*, (1997) Warwick. <http://www.pads.ahds.ac.uk/padsMetadataWorkshopCollection>
- [Robertson 1997] Robertson, G. (1997). *Sample Rate Synchronization across ATM Network*, ICMC Proceedings of the International Computer Music Association, Thessaloniki.
- [Robertson 1998] Robertson, G. (1998). *MiniMS, Multi-Participant Interactive Music Services*, University of Glasgow. <http://www.music.gla.ac.uk/~george/minims/minis.html>, 1999/05/08.

## Footnotes

- [1] <http://www.ahds.ac.uk/> , 1999/05/08.
- [2] <http://www.jisc.ac.uk> , 1999/05/08.
- [3] <http://lcweb.loc.gov/z3950> , 1999/05/08.
- [4] <http://www.netmuse.gla.ac.uk/> , 1999/05/08.
- [5] This distribution of information has also implications regarding copyright, i.e. institutions holding copyright of material may wish to hold their collection physically on their servers and still be able to offer single-user interfaces across remote collections.
- [6] See Categories of time-based Media: <http://www.music.gla.ac.uk/HTMLFolder/Research/smatbam-private/categories.html> , 1999/05/08, in [Boehm 1997 *SMaTBaM!*].
- [7] [http://purl.org/metadata/dublin\\_core](http://purl.org/metadata/dublin_core), 1998/07/01.
- [8] <http://ukoln.ac.uk/> , 1999/05/08.
- [9] <http://ahds.ac.uk/public/metadata/discovery.html> , 1998/07/01.
- [10] For more information about the Dublin Core and its interdisciplinary evaluation at the PADS and AHDS see [AHDS Metadata 1997].
- [11] See Proposal for SQL Access in Z39.50: Z39.50/SQL+,

[http://www.dstc.edu.au/DDU/research\\_news/reports/zproposal.html](http://www.dstc.edu.au/DDU/research_news/reports/zproposal.html), 1998/07/01. Although such plans are being discussed elsewhere, the AHDS's plans are limited to procuring specific interfaces between collection holders and Z39.50

[12] ODMG 2, <http://www.odmg.org/>, 1998/07/01