

The Performing Arts Data Service

<http://www.pads.ahds.ac.uk>
info@pads.ahds.ac.uk

Carola Boehm, Stephen Malloch, Celia Duffy,
Performing Arts Data Service, University of Glasgow, United Kingdom,
carola@music.gla.ac.uk, S.Malloch@pads.ahds.ac.uk, C.Duffy@pads.ahds.ac.uk

Stephen Arnold, Department of Music,
Tony Pearson, Theatre, Film and Television Studies,
University of Glasgow, United Kingdom,
S.Arnold@music.gla.ac.uk, agp@arts.gla.ac.uk

Abstract: The Performing Arts Data Service (PADS), funded by the Joint Information Systems Committee (JISC) and based at the University of Glasgow, aims 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, broadcast arts, theatre and dance. The PADS is one of 5 service providers of the Arts and Humanities Data Service (AHDS) which will provide a single gateway for arts and humanities scholars wishing to search for datasets across various discipline areas. Data is indexed with Dublin Core metadata, will interoperate with other databases within the AHDS and beyond, and will be available via the Web.

1. Introduction and Background

The Performing Arts Data Service (PADS) is one of a syndicate of five Service Providers appointed by the Arts and Humanities Data Service (AHDS)[1], funded by the Joint Information Systems Committee (JISC)[2] of the UK's Higher Education Funding Councils, and is based at the University of Glasgow. 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 datasets across various discipline areas. Other service providers include: the History Data Service (HDS), The Archaeology Data Service (ADS), the Oxford Text Archive (OTA) and the Visual Arts Data Service (VADS)[3]. The service providers' databases interoperate with other databases within the AHDS and beyond via Z39.50[4], 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, 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 (Scottish Higher Education Funding Council funded) **NetMuse** Project[5] was a project developing web-based music courseware for delivery over the ATM based Scottish Metropolitan Area Networks (MAN's). 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] (Serving Massive Time Based Media[6]). The SMaTBaM project researched storing, retrieval and delivery of time-based media and was designed specifically to benefit the work of the PADS whose remit does not extend to conducting research.

The SMaTBaM project researched and set up a prototype of a system which had to be suitable for the PADS, including the means of delivery of time based data as well as the storage and retrieval issues. This prototype has been scaled up and 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-orientated 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 Netdynamics and library databases and catalogues via a newly implemented Hyperwave-Z39.50 gateway.

Thus the user, although accessing information from different locations, different platforms and different kinds of databases sees one sort of interface, i.e. a browsable webpages guiding him through the information space via navigational aids such as discipline specific and discipline independant searching templates, knowledge domain related, hierarchical browsing hierarchies, depiction of relationships between resources via metadata and other hyperlinks.

The PADS has recently been involved in a major project to provide a digital catalogue records for the Scottish Music Information Centre (SMIC)'s unique collection of scores, manuscripts and recordings, with funding from the Scottish Cultural Resource Access Network (SCRAN). The project "Five Centuries of Scottish Music" aims to digitize 1000 items and metadata of manuscripts, scores, composers' bibliographies, concert programmes and recordings of ten representative scottish composers taken from the last 5 centuries and adding an additional glossary explaining musicological relevant terms mentioned in these records. This project is near its final stage and the collection is planned to be accesable in the near future.

This paper will cover three principle areas:

- the general information systems requirements for serving time-based media over the web;
- the PADS system architecture
- a report on the experiences gained in the first major effort to make a collection of 1000 digital resources available: the SCRAN/SMIC/PADS project.

2. Information Systems Requirements for Performing Arts data on the Web

In a time where 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 to be able to serve quality and quantity information across wide area networks. It also requires solutions of the more philosophical areas of research into how interfaces have to be set up and how information should be represented in order for users to handle vast amounts of data as intuitively and as user friendly as possible. In order to cope with the information increase and user capacity expected information management services will always have to deal with the "three I's": Information Structure, Information Representation and Information Access.

2.1. Nature of the data in the performing arts

A collection dealing with Performing Arts related data consists of both secondary resources i.e. materials about the performing arts, moving image and sound-based media, and primary resources, i.e. 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 is able to 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 keep and maintain their collection physically in their own repository, while access is handled by a central access point.[7]

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.[8] 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.[9]

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'. [10] 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.

We should accept the fact that our future data might not remain in its 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 the 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, are often not just one entity. In trying to devise resource systems of the next decade, it would be illogical to diminish the resources and their "real-life" manifestation by disregarding their composite character.

It was clear for achieving the above requirements, normal library catalogues and conventional relational database management systems would not be sufficient. Object-oriented or at least object-based information system technologies would have to be employed.

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 OMG's Object Request Brokers and their Relationship Service Specification for distributed objects [OMG CORBA-Relation 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:

- **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)
- **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).
- **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.).

- **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 DublinCore, where the DublinCore object itself has 15 further attributes defining the elements of the Dublin Core, see 2.2 Metadata).
- **Web Links* - Web-links can be thought of being a realisation of a certain kind of association in a web environment. The publication of these resources involves the presenting of one resource via different types of other resources or one resource related to others. For instance, a computer-music piece may exist as a sound file, presenting the first recorded performance, as well as archived as the code of the computer program itself and the secondary information associated with this resource.

Whereas relationships belonging to the categories of inclusion and inheritance are implemented “hard-core” into the PADS system, attributes can be thought of being the metadata of the objects. Unlike normal library system in which only metadata about objects is stored, which are not held digitally or stored separately, in the PADS system attributes belong directly to the objects described by them. In order to facilitate an interdisciplinary approach of resource discovery, the PADS within the greater context of the AHDS searched for 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 investigate and debate how best to facilitate resource discovery in an on-line setting. Specifically, the PADS looked at the metadata standard known as the Dublin Core[11] and how it could be applied as a tool to describe the time-based (sound and image) data resources that are the special responsibility of the PADS. 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[12] with funding from JISC. The aim of the series was to explore how different subject domains both describe and search for electronically held information and to evaluate the usefulness of the Dublin Core as common set of concepts shared across disciplines that may be used in the construction of the AHDS's integrated catalogue.

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. Although the AHDS workshop series[13] and other initiatives from the library and information community have proposed some fairly complex and lengthy qualifiers, and the AHDS has proposed amendments to some of the definitions, the Dublin Core consists of 15 basic elements:

1. Title	6. Contributor	11. Source
2. Creator	7. Date	12. Language
3. Subject	8. Type	13. Relation
4. Description	9. Format	14. Coverage
5. Publisher	10. Identifier	15. Rights

The PADS held two workshops in April-May 1997, inviting participants with a cross-section of expertise and interest in moving image and sound resources from both service provider communities (libraries and archives) and user communities (UK academics in performing arts disciplines). The groups examined the potential use of the Dublin Core for describing time-based resources, tested it against a variety of examples and critically reviewed its application. The findings from the workshops, which have been borne out by subsequent pilot applications to PADS data, were that the Dublin Core could function adequately, but there were some reservations and concerns over certain of its elements. [14] The experience gained was used to develop and implement a Dublin Core Metadata set with a data entry scheme, including additional subelements, which would work for performing arts data resources in general and for the SCRAN/SMIC/PADS project "Five Centuries of Scottish Music" specifically.

2.1.4 Time-Based Data

The common denominator of many prospective resources of the PADS service has the characteristic of being time-based. Storing and accessing time-based media requires special attention in storage and delivery of the objects. Solutions are needed to store information in its inherent complex form on the server side, to transmit these information packages in real-time with high-quality over a wide-area network, and to provide a user

interface able to access and use the resources intelligently. For a high-quality service four types of time-based material, all requiring real-time access, can be identified:

- *large binary data objects*: such as sound or video - streaming binary data combined with using a guaranteed bandwidth to ensure no glitches or breaks. 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.
- *subsets of large binary data objects*: playing just a part of a sound or video
- *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 1998 and Robertson 1998] E.g. 'lip sync' in film and tv, where sound and vision tracks are often recorded on different media.
- *complex objects*: such as MAX music scores, more complex Java applications, or sound-sound combinations require a fast and time-coordinated access of all the composite parts of an object: the synchronisation of multiple, periodic, logically independent streams of arbitrary type.[see Flinn 1995]

In the first two instances, high-performance media servers designed specially for serving time-based media offer many solutions for streaming large amounts of data. Automatic recognition of bandwidth and consequential choice of compression and delivery means ensures that users have the best quality available to them. More difficult are the last two instances. First results to synchronize separate streams on ATM networks have been achieved through the work of George Robertson at the University of Glasgow [Robertson 1997] and are planned to be integrated into the PADS system. Several upcoming standards such as MPEG7 will hopefully further this research and development.

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. To shortly sketch the status-quo situation of using multimedia digital resource collections already available, one can look towards broadcasting stations, music/video archives, record companies and libraries. It must be taken into account that collections are stored in different storage mediums, ranging from simple file systems, to relational database management systems to the growing number of object-oriented database management systems.[15] In addition, a large number of music catalogues in a variety of formats has to be also made accessible.[16]

Between library and library-like catalogues, an implementation of the Z39.50 protocol (version 3, 1995) is sufficient. For interfacing catalogues with relational databases, there will need to be a Z39.50 - SQL interface. There are 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. [17] 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).[18] 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. [19] Assuming a basic interoperability of different collections holding digital, multimedia objects, the underlying transfer protocol will have an influence on the performance, the quality and the representation means of the objects to be delivered. Using a stateless protocol, such as http, means that only one object can be delivered per session. Thus the connection closes after each document is delivered, losing all the information of the former session.

In devising a secure and distributed system, with collections stored in different locations, access handled from a central gateway and user access in the best case being controlled to a point of write, read and execute rights of

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, or/and 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 hooked up to it by Perl Database Modules (DBI:DBD-ODBC, DBI:DBD-Oracle, etc). In cooperation with the AHDS, Index Data (Denmark) and Hyperwave R&D GmbH (Germany) a generic Z39.50-Hyperwave gateway was implemented, translating incoming Z requests into the object query language used by Hyperwave. With this gateway, the requirements of being able to access library catalogues, as well as relational database management systems, as well as other object servers was fulfilled.

3.2. PADS Database and Media-Server System

In the greater context of the AHDS, a central WWW-Z39.50 gateway residing at the AHDS is responsible for incoming requests for multiple-database searches via Z39.50. Z39.50 targets, installed at each of the 5 service providers including the PADS, process the incoming requests and send the result-sets back to the AHDS WWW-Z39.50 central gateway. The PADS also has its own gateways, providing their own specific user community with specialized services relevant for time-based media. Thus gateways to a number of clients is realized, such as www browsers, media-players, Zclients, telnet clients, etc. Direct requests for subject specific searching is possible, with or without the use of Z39.50 protocols. The PADS resources reside on several database management systems, interoperable with each other through hyper-G, CGI and SQL interfaces. Hyperwave Information Servers are responsible for controlling incoming and outgoing requests from clients and interfacing databases. A, for this project implemented generic Z39.50-Hyperwave gateway was linked to the server, transposing incoming requests into the object query language used by Hyperwave. The hyper-G protocol allows proprietary clients (such as Harmony and Mozart) to have session-based protocols with graphical browsing abilities. These graphical representations of relationships of the objects in the DBMS are also realised with the http-Web gateway through Java/JavaScript implementations. The loss of functionality in using http, a stateless protocol, instead of statefull HG-CSP, is compensated by the use of Cookie files or java variables.

As the media server MediaBase on a separate Silicon Graphics Origin 200 is used. Movie players or soundplayers, either from MediaBase or those created by the University of Glasgow are able to make direct connections over the ATM networks. These players are installed on the client machines as browser plugins. Sound and movie objects are put into Hyperwave as objects, calling the functions by scripts to establish the direct ATM connection through Java classes and or C++ classes.

The PADS uses the feature of running several different web-gateways with different presentation schemes on different ports, in different collections and on different machines. The PADS is running two general webgateways at the moment, one serving webpages out for non-java/javascript browsers, the other utilising java and javascript in order to depict dynamically the resources and their relationships with each other. For specific data entry needs another webgateway is running at locations where collections are being archived, digitized and metadata'd. A fourth one is in development, to offer additional functionality needed for usage as a digital library service.

Interoperability between the HWIS and remote RDBMS/ORDBMS is supported through a HGI-SQL gateway. For example, to hook up a MS Access database, as the PADS has with a catalogue of a Glasgow -based film association, Hyperwave's own HGI (hyper-G Gateway Interface, using HG-CSP, hyper-G Client Server Protocol) communicates with its SQL gateway, which in turn interfaces to the standard available DBI-DBD perl-gateways. These have RDBMS vendor specific modules, such as DBD-Oracle, DBD-ODBC, DBD-Informix, etc. On the Access side, the only thing required is an MS ODBC or SQL server driver.

Another way used to link external sources to the Hyperwave Information server is by using the HGI (HyperG Gateway Interface) to call upon remote objects using as an additional attribute the protocol used such as http, telnet, Java Objects, CGI programs, etc. Mediabase, the database used by the PADS to store and deliver videos, is hooked up to the HWIS through the HGI. Metadata of videos and high quality soundfiles are stored on the Hyperwave Information Server, as is the protocol used for the delivery, but the BLOBs (Binary Large Objects) of

video or sound themselves are stored on Mediabase. The user thus accesses the videos through the Hyperwave Information Server, which builds up a direct connection between the client of the user and the Mediabase Server.

With this architecture, certain specifications demanded by a service such as the PADS are answered. It is an open but secure system running on a distributed resource environment with full user administration and rights management. It has, among other gateways, a web gateway and can be expanded by future gateways, answering the need of any upcoming specialized clients. It is interoperable with RDBMS and ORDBMS through SQL, and with its own object query language it is theoretically expandable with ODMG's OQL and ODL, linking it to OODBMS and distributed objects. It has powerful indexing capabilities and additional features deriving from the OODBMS world, such as versioning, inheritance of functionality, separation of content and view, and scalability of the delivery as well as the storage system.

References

[AHDS Metadata 1997] Greenstein, D. and Miller, P. (1997). *UKOLN/AHDS Metadata Workshop Series*, King's 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>.

[Flinn 1995] Flinn , Scott, (1995). *Coordinating Heterogeneous Time-Based Media Between Independent Applications*, ACM Multimedia 95 - Electronic Proceedings, San Francisco, California. <http://www.cs.ubc.ca/spider/flinn/publications/mm95/scheduler.html>.

[Kahn and Wilensky, 1997] Kahn & Wilensky (1997). *A Framework for Distributed Digital Object Services*, <http://www.cnri.reston.va.us/home/cstr/arch/k-w.html> , 1998/07/01,

[Lagoze 1995] Lagoze, D.& Shaw, E. & Davis, R. & Krafft, D. (1995) Erin Shaw,(1995) *Dienst/NCSTRL* <http://cs-tr.cs.cornell.edu:80/Dienst/UI/1.0/Display/ncstrl.cornell/TR95-1514>, 1998/07/01.

[Lagoze 1996] Lagoze, C. & Lynch, C. & Daniel, R. (1996) , *The Warwick Framework A Container Architecture for Aggregating Sets of Metadata*.Cornell University. <http://cs-tr.cs.cornell.edu:80/Dienst/UI/1.0/Display/ncstrl.cornell/TR96-1593>., 1998/07/01.

[Malloch and Pflücke, 1997] S. Malloch, S. Arnold, T. Pflücke, "Using Java to stream audio over ATM", Proceedings of the ICMC, Thessaloniki, 1997.

[OMG CORBA-Relation 1997] OMG (1997). *Relationship Service Specification for distributed objects*. OMG <ftp://www.omg.org/pub/docs/formal/97-12-16.pdf>

[Ossenbruggen 1994]. Ossenbruggen, Jacco van & Eliëns, Anton (1994). *Music in Time-Based Hypermedia*, Vrije Universiteit, Amsterdam. <http://www.cs.vu.nl/~jrvosse/Papers/echt94/html/index.html> , 1998/07/01.

[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>, 1998/07/01.

Footnotes

[1] <http://www.ahds.ac.uk/> , 1998/07/01.

[2] <http://www.jisc.ac.uk/> , 1998/07/01.

[3] <http://www.ads.ac.uk/> , <http://www.vads.ac.uk/>, <http://www.ota.ac.uk/>, and <http://www.hds.ac.uk/> 1998/07/01.

[4] <http://lcweb.loc.gov/z3950/> , 1998/07/01.

[5] <http://www.netmuse.gla.ac.uk/> , 1998/07/01.

[6] <http://www.music.gla.ac.uk/HTMLFolder/Research/SMaTBaM.html>, 1998/07/01.

[7] This distribution of information has also implications regarding copyright. I.e. Institutions holding copyright of material will likely want to hold their collection physically on their servers and still be able to offer single user interfaces across remote collections. The National Preservation Office of the National Library of Australia has made this "Distributed responsibility" one of its Statements of Principles of Preservation of and Long-Term Access to Australian Digital Objects. See <http://www.nla.gov.au/nla/staffpaper/preserve.html> , 1998/07/01.

[8] See Categories of time-based Media: <http://www.music.gla.ac.uk/HTMLFolder/Research/smatbam-private/categories.html> , 1998/07/01, in [Boehm 1997 SmatBam].

[9] For an example of the separation of content resource and various representational views, one might think of a picture stored in the highest resolution possible in a central resource archive, and its compression to a lower quality for web use. When using high-speed networks, one would still be able to provide a higher quality resource to appropriate users or maybe an even lower quality one due to any possible copyright restrictions. In the computing world, this separation of content and representation has one of its object-oriented manifestations in the Model-View-Controller paradigm. The model being the content, the data, or a knowledge domain, the view being one possible presentation of it. The controller can be seen as the gadget maintaining the connection between the model and the view. One musical note, for instance, could be depicted in a system by an internal, proprietary data structure. To this note, one or more views can be "plugged in" as for instance a midi representation, a sound representation, and a graphic representation. Devising new views is thus independent of the content. See also [Ossenbruggen 1994].

[10] Elementary or simple objects are objects made out of one entity or one binary (text files, bitmaps, wave format files, midi files). Composite objects consist of a number of elementary or composite objects, for instance a complex/composite music data structure. Complex objects are objects with attributes that change in size.

[11] http://purl.org/metadata/dublin_core, 1998/07/01.

[12] <http://ukoln.ac.uk/> , 1998/07/01.

[13] <http://ahds.ac.uk/public/metadata/discovery.html> , 1998/07/01.

[14] 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 very clear example for 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) Subelements 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 libraries using the Dublin Core face always the question of "Which metadata are we actually using, the one representing the digital object (for instance a 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. For more information about the Dublin Core and its interdisciplinary evaluation at the PADS and AHDS, see [AHDS Metadata 1997].

[15] See examples: Time-Warner Pathfinder Personal Edition, <http://pathfinder.com/@@5cnHOgcAhVYFFeXJ/welcome/> , 1998/07/01.- a personal magazine; the Chicago Tribune's Metromix, <http://www.metromix.com/> , 1998/07/01, EDS, Aniamtion 200 , Liberation (Libraries: Electronic Remote Access to Information Over Networks), <http://www.iicm.edu/liberation>, 1998/07/01.

[16] Whilst in the academic and non-academic library world interoperability has established itself as an important topic, it is not so clear that, mainly commercial, television and broadcasting companies will want their archives to interoperate with those of their competitors. However, given that material tends to decrease in commercial value with time but increase in academic and cultural or heritage value, it is quite possible that their material will end up in such a collection and so issues of interoperability are worth addressing at the outset.

[17] See Proposal for SQL Access in Z39.50: Z39.50/SQL+, http://www.dstc.edu.au/DDU/research_news/reports/zproposal.html, 1998/07/01. It should be noted that although such plans are being discussed elsewhere, the AHDS' plans are limited to procuring specific interfaces between collection holders and Z39.50.

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

[19] See the following research works and projects, which have been influenced largely by projects in cooperation with the Library of Congress: [Kahn and Wilensky, 1997], [Lagoze 1995] and [Lagoze 1996].