The Performing Arts Data Service

Stephen Malloch, Carola Boehm, Celia Duffy, Catherine Owen:
Performing Arts Data Service, University of Glasgow.
http://pads.ahds.ac.uk
info@pads.ahds.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.

Data relating to Performing Arts is by nature diverse: from text based, to visuals/images, to the intrinsically time-based. Any information system for dealing with this range of material must be able to store complex and composite data, cope with a multitude of single documents, and offer intelligent, user-friendly but controlled access over wide area networks. To be of most use to researchers some means of delivery of data is required as well as effective searching. To this end PADS has acquired two Silicon Graphics Origin 200 servers, one of which will act as a media server streaming audio and video over scalable networks; the other will run an object-orientated database (Hyperwave Information Server) which will store both the non-time-based data and the metadata of the material on the media server. A significant issue facing the PADS is that of streaming audio and video to multiple platforms over varying bandwidths.

This paper will cover the general information systems requirements for complex multimedia data and the web; will describe in detail the hybrid database and media-server system chosen for use at the PADS; and give an overview of current plans for testing video streaming at the PADS in conjunction with the British Film Institute/British Universities Film and Video Council and Joint Information Systems Committee's "Imagination/Universities Network Pilot".

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.

The AHDS will provide 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 will 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 time-based media.

Data relating to Performing Arts is by nature diverse: everything from purely text based (e.g. scripts, stage directions) to visuals/images (musical scores, artwork, photographs) to the intrinsically time-based (recordings of live performances, film, video, radio broadcast and multimedia compositions). Any information system for dealing with this range of material must be able to store complex and composite data, cope with a multitude of single documents, and offer intelligent, user-friendly but controlled access over wide area networks.

This is no small task; however, the University of Glasgow has already proven experience in this area of work, notably through two recent projects. 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[6] for streaming full CD quality music, further developed as part of the **SMaTBaM** project (Serving Massive Time Based Media[7]). The SMaTBaM project researched storing, retrieval and delivery of complex and 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 would be suitable for the PADS, including the means of delivery of time based data as well as the storage and retrieval issues. (Its work contributes significantly to this paper.) This prototype has now 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 complex and composite data.

This paper will cover three principle areas:

- the general information systems requirements for complex multimedia data and the web;
- the hybrid database and media-server system chosen for use at the PADS;
- current plans for audio and video streaming at the PADS, in conjunction with the BFI /BUFVC/JISC's "Imagination Universities Network Pilot."

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

Today, research centres, libraries, universities and the public in general wish to access and use the best information and data possible. This trend is emphasised by a tremendous need for user friendly and flexible systems with advanced information retrieval capacities for research and teaching.

Time-based media - such as that related to the performing arts which includes music, film, video, theatre and dance - have specific requirements which have only been dealt with to a certain extent in past projects. The inherent character of time-based content and the aim to facilitate real-time access in the highest quality possible provides problems that only high performance servers and networks can cope with, and the character of their multimedia content creates archiving problems that cannot be dealt with the traditional relational database model or catalogue systems.

Though much archival material to be stored and distributed will always be simple data types - such as text, image, audio or video - rather than more complex or composite types of data, appropriate mechanisms for searching and standards for exchanging information efficiently are still needed. Furthermore, structures and models are needed to fulfill the needs of the more complex relationships between these types of data.

The use of digital data resources to facilitate research and teaching in the performing arts has to define methods of storing and distributing complex time-based data to be able to serve quality and quantity information across wide area networks.

The important issues here can be split into four areas: the nature of the data; issues relating to metadata; preservation issues; delivery issues; and, interoperability.

2.1 Nature of data in relation to PADS.

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 will be 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* lets other collection holders keep and maintain their collection in their own repository, while access is handled by a central access point.[8]

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.[9]

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.[10]

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'.[11] 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.

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.[12] Relationships can be of numerous variety. For example, five relationships already 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 Bach’s works have a BWV, so each single work inherits the attribute BWV-

verzeichnis-number of the Bach Works Object; or, 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).
- **Web Links** - Web-links can be thought of being a realisation of a certain kind of association. 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.

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:
  1. high performance networks providing high bandwidth and guaranteed quality of service;
  2. client-server software tools to provide the streaming;
  3. 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. [13] 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.[14]

2.2 Metadata – The Dublin Core in relation to Performing Arts.

Regular Web users will be familiar with the limitations of search engines which index text and match hits but are unable to recognise context or intellectual content. Finding author or publisher information can be even more complex, and multimedia objects, which are not recognised by search engines, can be impossible to locate. Ironically, for those who need to use networks for resource location the complexities and promise of the web have encouraged a highly sophisticated series of search strategies. For busy academics or users with less experience, finding the right resource can be a daunting task.

During 1997, the PADS engaged in various activities to investigate and debate how best to facilitate resource discovery in an on-line setting. Specifically, we looked at the metadata standard known as the Dublin Core[15] 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[16], 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[17] 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.

Metadata can be defined as the descriptive information by which users locate re-sources - a sort of electronic catalogue card. The Dublin Core aims to provide users with a way of determining context, subject, intellectual rights and, crucially for the performing arts, the type of resource required. Even by taking a simplified distillation of the Dublin Core elements, for example creator and subject, a user can locate a specific object with accuracy. The ability to specify what type of object, for example a moving image file, a sound clip or an animated gif, facilitates real primary resource location, previously a hit and miss affair.

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[18] 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
2.  Creator
3.  Subject
4.  Description
5.  Publisher
6.  Contributor
7.  Date
8.  Type
9.  Format
10. Identifier
11. Source
12. Language
13  Relation
14. Coverage
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 bourne 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.

The Dublin Core often betrays its origins as a tool for text description; moving image resources in particular are generally not amenable to descriptive methods designed for text-based materials (which is why film archives often adopt their own specialised and individualistic procedures and systems for describing resources, rather than try to adapt to the standards used by the book world). For example, the use of Element 5 - `Publisher`, defined in the Dublin Core as *"the entity responsible for making the resource available in its present form"* is far less straightforward than for text-based resources. An archive, an electronic archive (such as the PADS) or a distributor all may provide access to a film; providing access is different from publishing. Production companies and broadcasting organisations each have a different role in the process. Problems were also found with the assignation of authorship to film resources: it is often neither possible nor desirable to assign a principal "author" to a moving image resource. Even if the convention of using the director is adhered to for movie resources, it cannot be consistently applied for TV and other recorded performances. The option of listing those of "secondary" importance implies a hierarchy of artistic effort which is problematic. The definition of Element 14 - `Coverage` which deals with the spatial and temporal characteristics of the resource caused such problems that we recommend that it is not used at all for moving image or sound resources. This element makes sense for an archaeological resource (where, for example, spatial = "Skara Brae" and temporal = "Neolithic"); although a statement about place or provenance is vitally important for moving image resources (for example, users may well want to search country of origin, country of release etc.) users are unlikely to search for this under `Coverage`. The PADS has made recommendations about the use of various elements of the Dublin Core and has proposed a number of qualifying statements[19] which are appropriate to moving image and sound resources.

2.3 Preservation Issues.

Whilst digital resources present some advantage regarding preservation, namely that of loss-less digital transfer between media, they also pose new problems – particularly the rapidly changing world of formats storage media.

As the paper published by the JISC/British Library Workshop of the 27th and 28th November 1995 at the University of Warwick[20] states, there are three types of digital resource preservation: medium preservation, technology preservation and intellectual preservation.

"The problem, and what is new about preservation in the electronic environment, is that electronic information must now be dealt with separately from its medium. This can be illustrated by an analogy …[I]f a book is placed on a closet shelf, and the closet door is closed for 500 years, then at the end of that time one can, broadly speaking, open that door and read the book. With

an electronic resource one does not have that confidence after ten years, and for several reasons." [21]

In the case of having digital resources as the resource itself (not only having records or catalogues describing it) we have, as mentioned above, the content and the representation of that content. The content is the resource itself, the viewers are the means for the user to see or access this resource. An example would be a piece of digitally stored music "viewed" with a Netscape sound player, or a Real-Audio streaming player, or a CD player. The content of the music seemingly does not change. The viewers, or in other words, the representations, do change.

In devising systems in which the rapid changes of technologies will not make the means of viewing information obsolete, we need to implement a separation of content and view as much as possible. The traditional technique of archivists was the "refreshing" of digital information by copying it into a new standard, a new media, or a new format or "migrating" it from one hardware/software configuration to another.[22] Both techniques can be lossy and time consuming. In adapting systems with the separation of content and viewer, combined with the ability of plugging in new viewers, a maximum of independence of technology change is achieved, while the resource is digitally stored in the highest quality possible and remains as that. If compression methods are needed to solve any storage shortages, then a lossless compression method, or a compression method with the least loss, should be used.

In the present era of digital distributed resources, most of the time a storing format that will enable the most efficient delivery is used. This results in using compression methods that would be unacceptable for academic research or cultural heritage preservation. Delivery means and storage means will have to be separated in digital archives and libraries, moving into the era of using digital resource preservation as a means to archive cultural heritage. Even if this seems an unrealistic viewpoint from today's standpoint, we will have to deal in a few centuries with digital artefacts that can no longer be reconstructed back into their original quality because they were originally stored in a compressed, lossy formats. Taking museums and archives for a model, their main aim is to preserve the artefacts as well as possible, and most of the financial expenditure of museums is allocated to this preservation of cultural heritage.

To sum up: in order to guarantee a high quality of digital resource preservation a system design should:

- Separate content from representation;
- Separate storage from delivery, and, if needed, use different storage formats from delivery formats;
- Use lossless compression methods for storing if possible.

The third preservation requirement, specified in the paper published by the JISC/British Library is *intellectual preservation;* this addresses the integrity and authenticity of the information as originally recorded. The JISC/British Library Workshop summed up the changes a digital resource may undergo:

1. Accidental change (data loss during transfer, accidents during updating, saving the wrong version)
2. Intended change / well meaning:
    - New versions or drafts (authorial texts, legislative bills);
    - Structural changes (updating books in print or a telephone directory);
    - Interactive documents, (hypertexts with note-taking capabilities)

3. Intended change / fraud: (political papers, laboratory notebooks, historical rewriting, legal documents, contracts)

This preservation aspect rather addresses security, versioning and copyright issues and has to be handled by a system as such:

- The digital resource in its resource archive must be secure from unwanted changes, i.e. a secure rights administration on the level of user groups and object collections.
- The digital resource may change over time, i.e. a versioning scheme might be required. The digital resource must have a way of identifying its copyright-holders in a secure way, i.e. either the copyright owner information has to be imprinted on the resource itself, or must be attached to it in a secure manner.

2.4 Delivery Issues.

For a number of reasons the traditional 'downloading' method of file transfer (e.g. ftp) is inadequate for moving large amounts of multimedia data around the inter/intranet. Time taken to download, having to download a whole file before listening/watching, reservations of copyright holders concerned about digital copying of material have all led to the concept of "streaming" media.

Digital audio and video data fit well the model of a stream of binary information which is fed to some type of decoder before recreating the original sound/picture. Streaming mechanisms allow the music or video to begin quickly (after a buffer has downloaded), they don't require large amounts of disk space on the client side, and they never result in a pure digital copy of the data being created. The popularity of streaming is clearly visible in the number of RealPlayer, LiquidAudio, Shockwave etc. plugins which have been downloaded from the Internet in recent years.

There are two generic issues relating to delivery of time-based media which give rise to a host of specific questions, namely Quality of Service and bandwidth.

1. **Quality of Service (QoS)** Arguably the most important aspect required when dealing with time based media. If the meaning of the data in question relies on its time-based nature then there must be some way to ensure that meaning is effectively communicated. i.e. data is delivered within the specified time-frame, otherwise the meaning of the data is lost. [23]
2. **Bandwidth.** Perhaps the most obvious feature of multimedia data, especially video, is the sheer quantity of data and hence bandwidth required. Table 1 shows some data rates for selected common digital audio and video formats.[24] As the speed of the existing internet for non-time based use (e.g. general web searching/browsing) currently shows, there is a distinct shortage of bandwidth for current purposes alone, never mind the potentially vast amounts of audio and video data just waiting for the day it can be streamed at high quality.

Standard or format.	Bandwidth Range (bps = bits per sec)
TV and video data	
PAL	400.2 Mbps (580x575px, 50fps)
NTSC	209.5 Mbps (600x485px, 30fps)
ITU-R 601	140-270 Mbps (320x480px)
MPEG-1	1.2–2.0 Mbps (352x240px, 30fps)
MPEG-2	4-60 Mbps 1.5 Mbps (VHS, 352x240px) 5-6Mbps (b'cast, 1440x1152px) 7Mbps (studio, 1920x1080px)
H.320 H.261	RealVideo 1.0 28.8kbps (newscast) 55kbps (full motion) 100kbps ("near" NTSC)

Audio data:-	
ISO audio Layer 1 Layer 2 Layer 3	MPEG Philips DCC 192kbps MUSICAM 128kbps ASPEC 64 kbps.

Table 1 – Data rates for selected media formats.

The other issues relating to delivery – which all must take bandwidth and QoS into account - are network technologies, network protocols, compression/quality, access tools and "network philosophy".

2.4.1 Network technologies.

The most common networks in use are 10Base-T, 100Base-T ethernet, FDDI, token ring or ATM; running over fibre optic or twisted pair cables. Raw bandwidth is a combination of the physical medium and networking technology used. 10Mbps ethernet has been the long established standard for LAN's, 100Mbps ethernet is quite common, and 155Mbps ATM networks are quite widely deployed (e.g. the UK Metropolitan Area Networks).

2.4.2 Network protocols.

Although network protocols are often independent of the underlying technology, the two are not necessarily entirely dissociated. ATM, for example, does not follow the OSI model[25] which makes its applicability to IP traffic more difficult; also, IP switching hardware routes IP packets by examining the IP header deeper down the packet, and by identifying 'flows' of packets.[26]

As for time based protocols, the internet has long used TCP and UDP over IP for most needs: neither of which are at all sympathetic to time-based media. TCP for example is commonly known as a reliable protocol since all packets will eventually arrive, however packets may arrive out of order and dropped packets will need re-transmitted resulting in further delays unacceptable for real time data. New protocols are in development in an attempt to address the needs of time-based data: RTP, RTCP, RSVP.

- **RTP** the Realtime Transport Protocol provides support for applications with real-time properties, including timing reconstruction, loss detection, security and content identification. A supporting protocol **RTCP** – Realtime Transport Control Protocol – is also under development though RTP can be used without RTCP. RTP can be used over a variety of protocols, and often uses UDP. RTP is currently also in experimental use directly over AAL5/ATM.[27]
- **RSVP** is a network control protocol designed to allow Internet applications to obtain QoS measures for their data generally by reserving resources along the data path(s). The RSVP working group of the IETF is developing an Internet standard specification for RSVP which is a component of the future "integrated services" Internet, which will provide both best-effort and realtime qualities of service.[28]

OSI Layer	OSI Layer Name	Examples
4	Transport	TCP, UDP, RTP, RTCP
3	Network	IP, RSVP
2	Data Link	Ethernet 802.1p 802.1Q; Gigabit ethernet 802.3z/ab
1	Physical	10BaseT, FDDI

Table 2 Network technologies and protocols.

Additionally, Real Time Streaming Protocol (RTSP) – an application layer protocol - is under development. RTSP is designed specifically for streaming audio and video and intends to provide a framework for interoperability between streaming client-server applications from different vendors.[29]

Also worth mentioning here is IPv6 (RFC1883) the so-called "next-generation" IP protocol. IPv6 will expand the IP addressing system to allow a much greater number of addresses and will introduce features such as flow labelling, hierarchical routing, and security. When IPv6 is finally implemented it is likely to be of some benefit for time-based media, but this is not imminent.

2.4.3 Compression/quality of media.

Although the argument that "infinite bandwidth" will become available is sometimes used in debates about time-based media over the web, it is fairly safe to assume that millions of streams of uncompressed ITU-601 video will not be zooming round the world for a considerable number of years yet. Compression, as well as realtime network protocols, will therefore play a part in any current media serving solution, and has a direct effect on quality and bandwidth issues. Although audio requires less bandwidth than video, it is perhaps more demanding of compression techniques and network performance for high quality streaming since the slightest glitch is very noticeable, whereas video is more tolerant of occasional dropped frames. MPEG-2 video standard has been designed for broadcasting and should meet all high quality demands; layer 1 MPEG-1 audio – whilst not CD quality - uses compression similar to consumer DCC and MiniDisc and should be adequate for most higher quality audio needs.

Echoing the preservation issue, material can be archived in its highest quality format offline and encoded with a chosen compression format for delivery at a lower bitrate.

2.4.4 'Access tools'

At this point in time access to streaming media is only available via proprietary servers and players which do not interoperate. Some form of interoperation will be useful as (or if) and when media servers (and available bandwidth) become common place. E.g. a single client that could receive streams from Sun Media Centres, SGI MediaBase servers and RealPlayer servers. RTSP based solutions will perhaps change this situation some day.

2.4.5 Network philosophies.

If the hype in the computer media is to be believed, then Gigabit ethernet has already won the battle against ATM (almost without a switch being sold), and with it another victory in the battle between connection-orientated networking versus connection-less networking which has run since the Internet began. This is not too surprising since the design philosophy of the Internet has always been to keep the network simple and puts the complexity into the end systems. ATM goes against this trend since it is an inherently connection-orientated system. ATM has, however, proved itself to be highly successful at reliably and satisfactorily delivering time-based media[30] but it is the applicability of ATM to more conventional internet (IP) traffic which means it is not likely to dominate the IP world of the Web. For example, ATM Switched Virtual Circuits, SVC's, take time to set up and tear down. This is fine if the duration of the connection is much longer than the time taken to set it up, but IP packets are relatively small and self-contained and must find their own way across the internet[31] making IP over ATM a less elegant solution.

Newcomers to the philosophical debate include IPv6, and protocols such as RTP, RTCP, and RSVP. RSVP is interesting in this context since it also (slightly) subverts the purist connection-less paradigm by implying that routers must have some knowledge of state of connections.

These new protocols are yet to prove themselves for the purpose of reliable and satisfactory networked delivery of time-based data, though there will no doubt be much activity in the coming year.

Cable modems are another area where much research and development is going on – since cable is a significant market – and ATM could play an important role here.[32]

2.5 Interoperability.

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.[33] In addition, a large number of music catalogues in a variety of formats has to be also made accessible. [34]

Between library and library-like catalogues, an implementation of the Z39.50 protocol (version 3, 1995) will be 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.[35] 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).[36] 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. [37]

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.

3. PADS Systems Architecture - 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, Z-clients, 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.

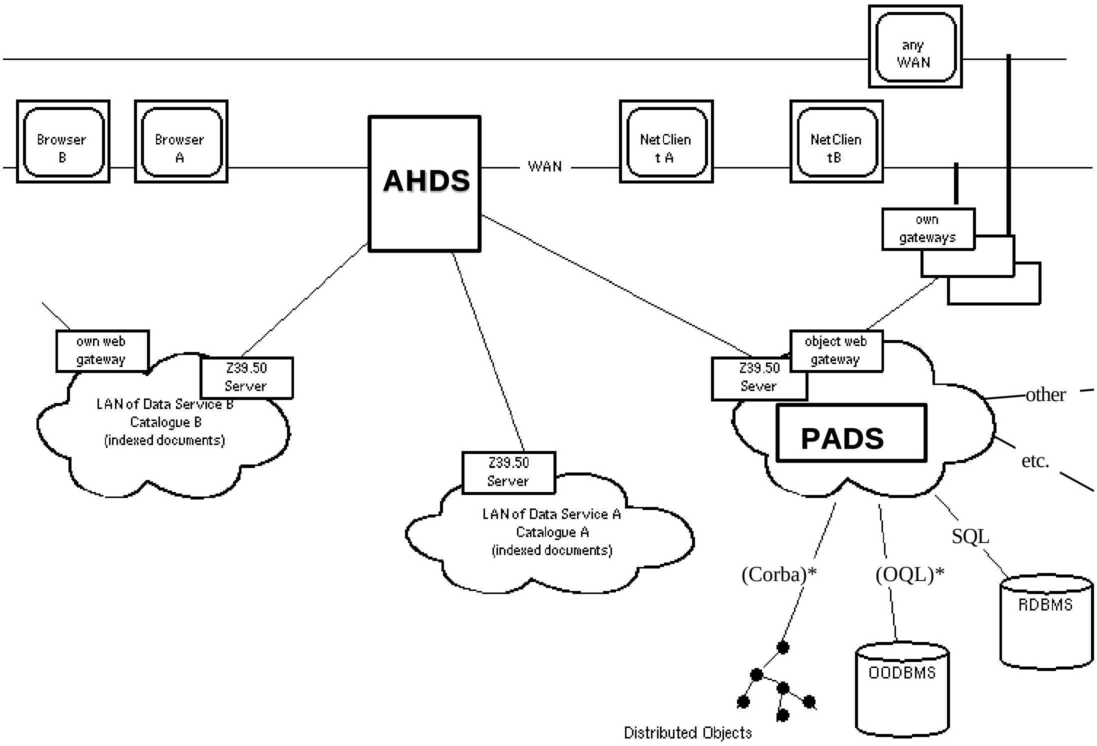

*) not yet implemented interfaces

Greater context of the PADS Service

Figure 1. Context of PADS service.

The central database management system at the PADS is represented and powered by Hyperwave Information Server. The server can be seen as being made up of three different layers: the protocol conversion layer, the session layer, and the database layer. The protocol conversion layer makes the Hyperwave server a multi-protocol server. The WWW gateway transforms HTTP to Hyperwave's client-server protocol (HG-CSP), so that lower layers need to deal only with a single protocol. The session layer communicates with the database layer. An instance of the session layer is created for every client connection, to retain state information and to parallelize client requests. The database layer is where the actual documents, links and meta-information are stored. This layer consists of three modules: the object server, which creates, modifies and deletes objects and their relationships, indexes them for searching, and manages users and access permissions; the full text server, which maintains an inverted index of all text documents for searching; and the document cache server, which stores the local documents of the server, as well as cached documents from remote servers.

The Hyperwave object server (OB) is a multi-user, read-write database. It stores a database of, and relations of, objects. The objects include documents, collections, anchors, users, user groups and server descriptions. The relations define collection membership and links. In fact, the link information of all documents is contained completely in the object server and is also responsible for generating unique local objectIDs (32-bit numbers which are used to identify objects on the server) for new objects. Metadata is stored as attributes of objects. Links and collections are stored as relations between objects. There exists a class hierarchy of objects and attributes are inherited by derived classes.

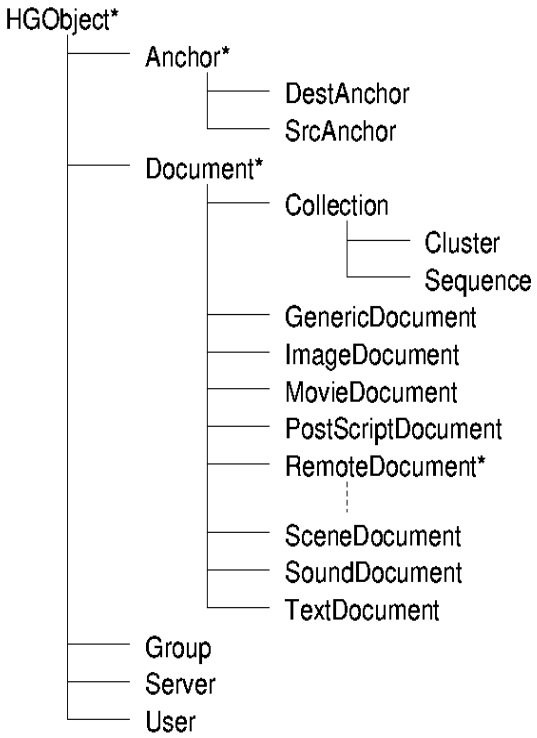

Figure 2. Hyperwave Object Hierarchy - *) abstract base classes, i.e. no instances of such classes can be created

The full text server stores indexes. Unlike in many other systems, indexing is performed immediately. The Verity search engine is a commercial product which has been transparently incorporated into the full text server. Verity uses a so-called universal filter, which recognizes several document types by default. This filter consists of the zone filter, which filters HTML documents, the PDF filter, for Acrobat PDF documents and the Mastersoft Filter Kit V1.5, which supports filtering of WYSIWYG documents. Among the c. 60 document formats supported are WordStar, Microsoft Word, WordPerfect, Microsoft RTF, Lotus 1-2-3, Microsoft Excel, Ami Pro, Microsoft Works, Microsoft Windows Write, MacWrite, Paradox, dBASE, Access and PowerPoint. The document cache server is where the actual documents on a Hyperwave server are kept.[38]

The Hyperwave Information Server can be run in a server pool, enabling multi-server / multi client environments in WAN with its own proprietary connection-oriented hyper-G protocol (or HG-CSP hyper-G Client Server Protocol). Its separate protocol conversion layer enables the use of different and additional protocols, the most important one being the http protocol running the web-gateway. The Hyperwave server with a web-gateway is installed on a Silicon Graphics Origin 200 running Irix 6.4.

During the SMaTBaM project, an external simplified Z39.50 gateway was linked to the server by a perl script, transposing incoming requests into the object query language used by Hyperwave. It will be replaced in mid 1998 by a Z39.50 gateway being implemented especially to conform to the needs of the AHDS-PADS service.

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.

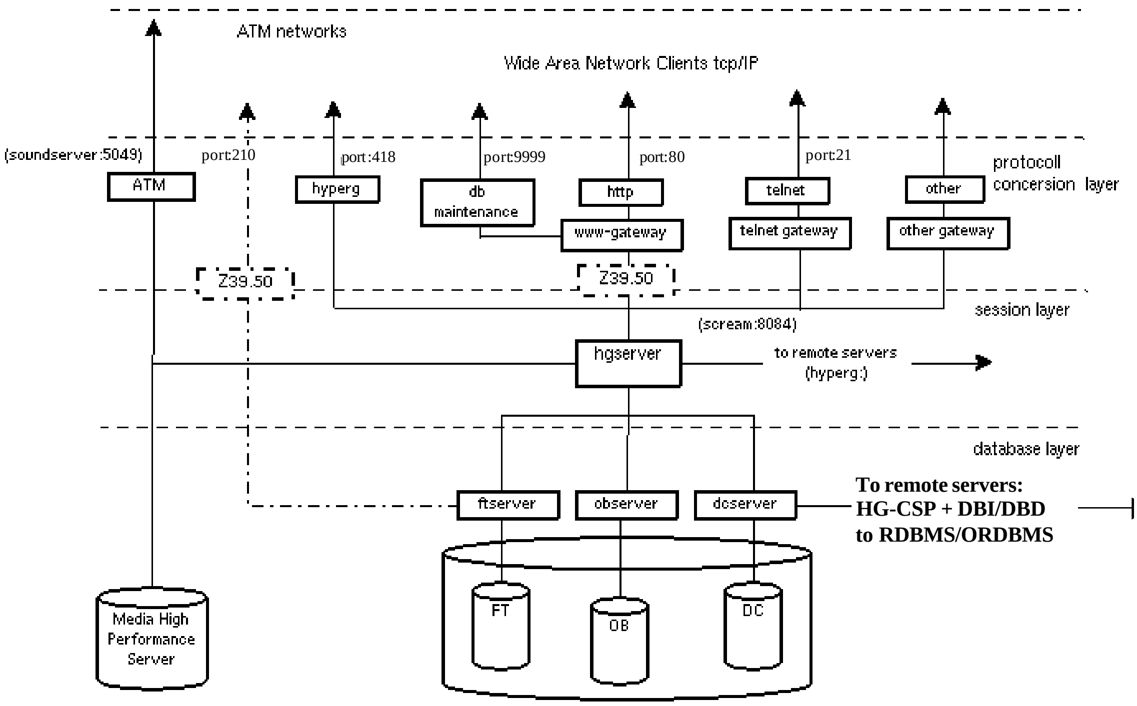

Figure 3 Overview of PADS system architecture

The system can run simultaneously with a normal webserver on the same machine. It is possible to run several different web-gateways with different presentation schemes on different ports and on different machines. The PADS is running two different 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. Also single objects can be linked to certain www-gateway presentation modules (PLACE-templates), and single presentation modules can be programmed with HTML, javascript and a C++-like programming language (PLACE-language). Thus there is a scalable control down to object level of the object presentation over the WWW and it is possible to have different views on one object.

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

Interoperability between the HWIS and remote RDBMS/ORDBMS is supported through a HGI-SQL gateway. For example, to hook up a MS Access database, like a film catalogue, 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 needed is an MS ODBC or SQL server driver.

Another way of linking 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

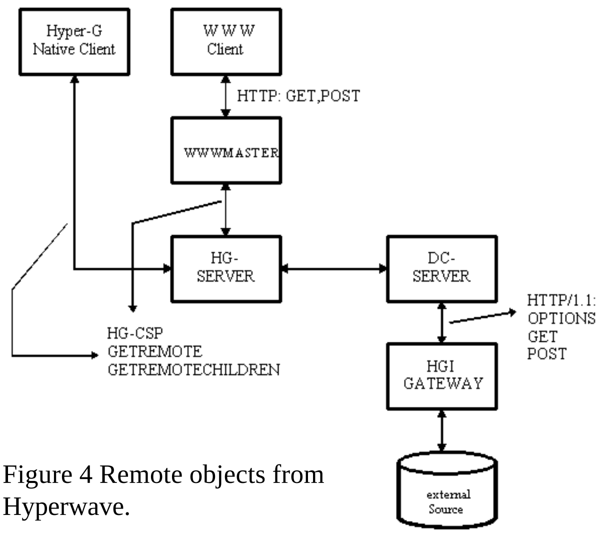

Figure 4 Remote objects from Hyperwave.

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.

It is worth highlighting some of the features of the Mediabase server itself at this point.

Firstly, the software runs on a twin CPU (R10k 195MHz MIPS) machine with 256MB RAM, almost 50GB of disk space and running IRIX 6.4, SGI's 64bit UNIX OS. High performance is maintained through the XFS filesystem with Guaranteed Rate I/O (GRIO) software which uses the realtime section of the disk subsystem to provide sustained quality of service for every individual (realtime) stream. In addition, resource allocation on network interfaces guarantees that all realtimes streams leaving the server have sufficient QoS to maintain the stream satisfactorily.

In order that this QoS is not lost as soon as the data reaches the network, MediaBase supports a range of direct network interfaces: 10Base-T, 100Base-T, FDDI, ATM, ISDN; and networking protocols including: native ATM, IP over ATM, RSVP/IP over ATM, RSVP API 4.0, LAN-Emulation over ethernet, token ring and ATM. Which provide varying levels of QoS.

Additionally, MediaBase supports multicast with Time To Live (TTL.) Media formats supported include MPEG-1, MPEG-1 Audio, MPEG-2, and other popular streaming formats such as RealPlayer can be integrated into the system.
Decoding clients for PC's, Mac's and various flavours of UNIX are provided with the system.

A very important feature of Mediabase is automatic scalable streaming. A movie or audio file can be stored in both high quality (MPEG) and low-quality low-bitrate (e.g. RealPlayer) format. By configuring a simple data file on the client side, whenever a client requests a stream the most appropriate version is streamed without any intervention required

Lastly, Mediabase has its own web gateway and relational database management system although these are not extensively used in the PADS environment since Hyperwave performs this function.[39]

A useful diagram showing the Mediabase file, database and web gateway architecture is available on the web.[40]

4. Current plans for audio and video streaming at the PADS.

The question of making available high quality sounds and images relating to performing arts is one the PADS is keen to tackle. While high quality audio and video is desirable, indeed perhaps a necessity, for research of performing arts related data this obviously requires high bandwidth network connections. PADS was always going to have to take a pragmatic view of the level of quality streaming it was able to do.

Although being linked to the Scottish MAN infrastructure via ClydeNet, there are many other groups also using this resource meaning that spare bandwidth is limited.

A more realistic proposal for streaming high quality video is within the campus.

PADS plans have been overtaken somewhat by the selection of University of Glasgow by the BFI/BUFVC/JISC to be a pilot site in their *"Imagination/Universities Network Pilot Project"*.

A consortium from the University of Glasgow led by the PADS and consisting of a range of technical and academic staff (including representatives from the Revelation Project[41]) successfully bid for the opportunity to run trials streaming video to workshops and labs on campus and remote lectures and workshops off campus.

The material to be supplied by the BFI's National Film and Television Archive consists of 30 hours of (MPEG encoded) material in three subject areas: Film and TV, Social History, and Medicine. (10 hours of material in each.)

Final details of the nature of the tests are in the process of being finalised at the time of writing. The intention is to run tests from both PADS' SGI Origin200 MediaServer and Revelation's Sun Media Centre. The tests will stream video to end users over a range of different network technologies and protocols including: contention based 10Mbps ethernet, dedicated 10Mbps ethernet, switched 100Mbps ethernet, native ATM, IP over ATM, LAN-Emulation. Tests will be run over both 'private' temporary dedicated networks on campus, and also into existing 'production' level labs used daily by students – such as the multimedia resource room in the Gilmorehill Centre for the Dept of Theatre, Film and Television Studies.

Comparative trials will be run with and without the use of protocols such as RSVP to see how well they perform. It is also likely that the campus backbone will itself be upgraded to Gigabit ethernet during the life of the project in which case that, too, will be incorporated into the pilot.

The pilot project hopes to gather information which will assist the BFI/BUFVC/JISC in determining what would be required in scaling up such a scheme to allow nationwide delivery of moving image data to Higher Education establishments. Both technical/managerial information is required (network statistics, implications for charging policies, security issues etc.) and information on the pedagogical aspects (how useful is the network as a delivery medium for moving image material).

The pilot is still in the early stages; some preliminary testing has begun and the results of this pilot will be published at a later time. The pilot is due to run until June, with a possible Phase 2 to run until December.

5. CONCLUSION.

This paper has described the general information systems requirements of Performing Arts related data on the web. This complex multimedia data is concerned with metadata, data modelling and preservation as well as the issues of QoS, bandwidth and quality/formats for networked delivery of time-based media. The networking requirements for streaming on the web are far from settled, and it is unlikely for a universal solution to appear in short term.

Acknowledgements

Dr. Stephen Arnold, Co-director of PADS and director of NetMuse and SMaTBaM projects, Dept of Music, University of Glasgow.
SHEFC which funded SMaTBaM project.

Biography of the authors.

Steve Malloch, Systems co-ordinator for the PADS. Worked on SmaTBaM project.
Carola Boehm, Database Systems consultant to the PADS. One of the initiators of the MusicWeb project[42]. Worked on SMaTBaM project.
Celia Duffy, Project Manager of the PADS. Background in music and technology in teaching and learning. Was Project Manager for NetMuse.
Catherine Owen, Collections Co-ordinator for the PADS. Previously librarian at the Scottish Music Information Centre.

References.

---

[1] http://www.ahds.ac.uk/
[2]http://www.jisc.ac.uk
[3] http://www.ahds.ac.uk/
[4] http://lcweb.loc.gov/z3950
[5] http://www.netmuse.gla.ac.uk/
[6] S. Malloch, S. Arnold, T. Pflicke, “Using Java to stream audio over ATM”, Proceedings of the ICMC, Thessaloniki, 1997.
[7]http://www.music.gla.ac.uk/HTMLFolder/Research/SMaTBaM.html
[8]The National Preservation Office of the National Library of Australia has made this "Distributed responsibility" one of its Statement of Principles of Preservation of and Long-Term Access to Australian Digital Objects. See http://www.nla.gov.au/npo/natco/princ.html.
[9] See Categories of time-based Media: http://www.music.gla.ac.uk/HTMLFolder/Research/smatbam-private/categories.html.
[10] 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 muiscal 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, a graphic representation. Devising new views is thus independent of the content.
See also Jacco van Ossenbruggen: Music in Time-Based Hypermedia: http://www.cs.vu.nl/~jrvosse/Papers/echt94/htmlindex.html.

[11] 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.

[12] See Relationship Service Specification for distributed objects, (OMG) http://www.omg.org/corbserv/relation.pdf, and also the work done by the Laboratory for Advanced Information Technology: the Knowledge Interchange Format, KIF, http://www.cs.umbc.edu/kse/kif/kif101.shtml.

[13] See, MIniMS, Multi-Participant Interactive Music Services: http://www.music.gla.ac.uk/~george/minims/minis.html. And, G. Robertson, “Sample Rate Synchronization across ATM Network”, Proceedings of the ICMC, Thessaloniki, 1997.

[14] See Scott Flinn: Coordinating Heterogeneous Time-Based Media Between Independent Applications, http://www.cs.ubc.ca/spider/flinn/publications/mm95/scheduler.html.

[15] http://purl.org/metadata/dublin_core

[16] http://pads.ahds.ac.uk/workshops/padswshops.html

[17] http://ukoln.ac.uk/

[18] http://ahds.ac.uk/public/metadata/discovery.html

[19] http://ahds.ac.uk/public/metadata/discovery.html

[20] http://www.ukoln.ac.uk/services/papers/bl/rdr6238/

[21] The Commission on Preservation and Access and The Research Libraries Group, Inc., “The Challenge of Archiving Digital Information:” http://www.rlg.org/ArchTF/tfadi.chelng.htm

[22] See Sites about Digital Research Preservation: AHDS, Digital Research Preservation in the AHDS,http://www.kcl.ac.uk/projects/ahds/preserve.htm; Bibliography of Resources Related to Preservation of Digital Resources: http://www.oclc.org:5046/~weibel/archtf_bib.html; Long Term Preservation Of Electronic Materials: http://ukoln.bath.ac.uk/fresko/warwick/intro.html

[23] Even the argument that “infinite bandwidth” will become available does not contradict this idea, as infinite bandwidth would simply be the *means* of implementing QoS.

[24] M. McCutcheon, M. R. Ito, G. W. Neufeld, “Video and Audio Streams Over an IP/ATM Wide Area Network” UBC Transport Encoded Video over IP/ATM Project, Technical Report 97-3, June 97. http://www.ncstrl.org:3803/.

[25] A. Tanenbaum “Distributed Operating Systems” p44, Prentice Hall International Editions, New Jersey 1995.

[26] For a list of RFC’s and hardware relating to IP switching see Ipsilons web pages: http://www.ipsilon.com/products/specifications.htm#HWPROC

[27] C. Liu, Multimedia Over IP: RSVP, RTP, RTCP, RTSP, http://www.cis.ohio-state.edu/~cliu/

[28] See RFC1633, http://ds.internic.net/rfc/rfc1633.txt

[29] See, http://www.realaudio.com/prognet/rt/.

[30] See, for instance, D. Konstantus, Y. Orlarey et al “Distributed Music Rehearsal” Proceedings of the ICMC, Thessaloniki, 1997 where a conductor in Bonn held a rehearsal with musicians in Geneva connected via ATM audio and video codecs. A much more demanding scenario is hard to imagine: real-time, duplex transmission of audio and video, playing a modern piece of music with unpredictable timing. The fact that this was possible at all

shows what other technologies and protocols have to compete with.

[31] There are various ways of addressing this such as "soft pvc's" and "flow recognition" but they are not ideal solutions. See [24], p15ff.

[32] See [24], Appendix D.

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

[34]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.

[35]See Proposal for SQL Access in Z39.50: Z39.50/SQL+, http://www.dstc.edu.au/DDU/research_news/reports/zproposal.html. 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.

[36] ODMG 2, http://www.odmg.org/

[37] See the following research works and projects, which have been influenced largely by projects in cooperation with the Library of Congress: Kahn/Wilensky, "A Framework for Distributed Digital Object Services, http://www.cnri.reston. va.us/home/cstr/arch/k-w.html, Daniel Lagoze's Dienst/NCSTRL http://www.ncstrl.org/ and "The Warwick Framework A Container Architecture for Aggregating Sets of Metadata." http://cs-tr.cornell.edu:80/Dienst /Repo-sitory/2.0/Body/ncstrl.cornell%2fTR961593/html.

[38]See http://www.hyperwave.de

[39]The database underlying Mediabase is actually a version of Informix; programmer API's are available for the SQL interface.

[40]http://www.css.minsk.by/prod/web/cosmo/media_base_index.html

[41]http://www.revelation.gla.ac.uk/

[42] http://sun1.rrzn.uni-hannover.de/musicweb