

Sustainable Archiving and Storage Management of Audiovisual Digital Assets

M. Addis, R. Beales, R. Lowe, L. Middleton, C. Norlund, and Z. Zlatev

With the advent of end-to-end tapeless production and distribution, the concept of what it means to archive audiovisual content is being challenged. The traditional role of the archive as a repository for material after broadcast is changing as a result of digital file-based technologies and high-speed networking. Rather than being at the end of the production chain, the archive is becoming an integral part of the production process and is being absorbed into wider digital storage environments, including those distributed or used across organizational boundaries. This paper presents some of the work done in the U.K. AVATARm project on service-oriented approaches to digital permanence and preservation of audiovisual content. The main focus is how to specify and govern federated storage services to ensure the long-term safety, security, and accessibility of audiovisual assets in a managed and cost-effective way.

AVATARm is a U.K. collaborative R & D project supported by the Technology Strategy Board in which the IT Innovation Centre, BBC, Xyratex, and Ovation Data Services are developing an innovative approach to large-scale long-term digital archiving within distributed storage infrastructures. This paper presents work from the project on tools to support the planning and management of service-oriented data archiving infrastructures. The tools allow content-centric workflows within an organization to be analyzed in order to profile the generation and consumption of archive assets, including the requirements for safety, security, longevity, and accessibility. These profiles then allow storage provision to be planned in terms of long-term access, ingest, and retention, and technical specifications are created and matched against storage solutions or managed services. The authors also present how service-oriented architectures (SOAs) using automated policies and service level agreements (SLAs) can deliver online archive functions in a managed way within an enterprise, when outsourcing archive hosting, or when collaborating with external organizations.

Motivation and State of the Art

Audiovisual content collections are transforming from archives of analog materials to very large stores of digital data. We are increasingly in an era of direct archive integration into production, distribution, and consumption workflows, with dynamic preservation processes required as a consequence. For example, the BBC Digital Media Initiative project¹ aims to deploy a completely tapeless environment across the entire organization over the next five years. Although there is intense interest in preservation strategies for digital content,²⁻⁵ in general, there is little work on practical implementations tailored for the needs of audiovisual content. For example, the Open Archival Information System (OAIS) Reference Model⁶ defines some of the processes required for long-term preservation and access to information objects, but does not specify how to monitor audiovisual objects or the systems in which they are stored, identify when migration should take place, or determine to what an audiovisual object should migrate. Audiovisual content presents demanding challenges for digital preservation, especially given the

preservation ideal of storing uncompressed content. SD digital video has an uncompressed data rate of about 200 Mb/s, and even when stored with compression (e.g., 50 Mb/s DV), multiple petabytes of storage are required for a typical broadcast archive. HD requires five times as much space. In digital cinema, 4K requires up to 30 times the data rate of SD. This presents a real problem, not least the cost, for which estimates range from “half the price of analog”⁷⁷ to nearly “12 times higher.”⁷⁸

Value chains and business models are also changing and increasingly delivered through multiple service providers or organizations (e.g., outsourced services, federated preservation across organizations, etc.). The economies of scale, power, cooling, and staff costs that can be achieved by organizations such as Google⁹ mean that as network costs continue to fall, in-house solutions will become increasingly expensive compared with outsourced or federated models.

Remote access to archive hosting services has not yet emerged in the broadcast industry, although there are services for remote access to data for distribution, such as VIIA from Ascent Media¹⁰ and data transfer within the enterprise (e.g., DIVAGrid from Front Porch Digital).¹¹ Critical to provision of such services in trusted archive environments is the use of policy-based service governance, which is based on two principles: (1) nonfunctional aspects of a service, including performance should be agreed on in a SLA and (2) the service should be managed, preferably in an automated (self-governing) management environment, so that it conforms to its SLA. Initiatives to standardize how this form of automatable SLA is made and represented include WSLA¹² and WSAgreement.¹³

Web service based infrastructures with explicit support for automated service management using automatable SLAs and QoS include FP6 NextGRID,¹⁴ FP6 TrustCOM,¹⁵ and IT Innovation’s GRIA¹⁶ technology. These projects recognize that trust and security (e.g., to support assertions of integrity and authenticity) are equally important in distributed environments (e.g., NextGRID work on interoperation across heterogeneous security environments, including X.509, SAML and Kerberos token exchange). At the same time, the digital library community has been busy creating software frameworks for implementing preservation environments. These include open source solutions, such as DSpace,¹⁷ which provides standard services for ingestion and access and is ported to run on top of SRB for managing distributed data; Fedora,¹⁸ which associates display functions with each data type, allows relationships to be imposed on records, and maps semantic labels on records to an ontology; and simple, off-the-shelf systems such as Greenstone.¹⁹

Commercial digital library systems, such as those from ExLibris,²⁰ are another approach, however; none are designed specifically with the challenges of audiovisual (AV) content in mind. Overall, the work of the digital library community, the broadcast community, the service oriented systems community and the preservation community has largely yet to converge. There is some notable initial work in this direction, however, in the SOA community where a new generation of inter-organization production and post-production infrastructures is emerging through projects such as MUPPITS,²¹ PRISM,²² and BeInGrid,²³ as well as in commercial offerings such as Signiant’s Digital Media Distribution Management Suite.²⁴

In summary, there are clear indicators that digital archiving is changing rapidly in the AV

community and new business models can be anticipated, based on archive service provision. However, the technology state of the art is one of fragmentation in which individual communities, such as the broadcast industry, digital libraries, and SOA, each provide pieces of the puzzle. The challenge is one of integration and adaptation to the specific challenges of AV content.

Approach and Implementation

Our solution is based on three core components. First, GRIA and aggregated storage is used as the basis of secure and managed archive hosting services that operate across administrative domains and can be federated with internal systems. The Open Archival Information System (OAIS) standard is used to specify the interface of these archive services. Second, a multilevel model of archive requirements is used to allow the concerns of the archive manager (assets, users, safety, longevity, value) to be separated from the specifics of a particular technical solution (disk, tape, networking, etc.). This is done through profiles for archive ingest, access, and retention, which specify what goes in and out of an archive. Third, simulation and modeling techniques are being developed to analyze content-centric workflows to determine the workloads these place on an archive and the variations that are likely to occur on a range of timescales.

Storage in AVATAR-m is heterogeneous, reflecting the broad range of storage types that an archive may typically use. The emphasis in the proposed solution is on networked storage, such as spinning disk or media jukeboxes, which may or may not be configured within a storage area network (SAN) or network attached storage (NAS). Additionally, online remote storage provided as a service is also supported to allow archives to make use of third-party storage services such as Amazon. The approach here is to combine these disparate storage types and locations, so that they are aggregated together into a single storage solution, as shown in **Fig. 1**.

Adapters are used for each storage type to which the storage aggregator interfaces, but because most operations are done at the file system level, additional adapters are required only for storage services that offer different application programming interface (APIs). Rather than assigning each asset to a specific tier, available storage locations are ranked dynamically, using a cost function and multi-objective optimization, based on factors such as the current and average read/write rates and availability. The use of the storage is also monitored, ensuring that content that is accessed frequently is made available from higher-ranked (and therefore faster) locations, whereas content accessed often is moved to slower storage. The rules that determine what gets moved can be modified through management policies that can be assigned to specific items or classes of items, such as all files of a certain type, or belonging to a certain user or project. This is similar to hierarchical storage management (HSM) systems, but with the advantage, in this case, of being able to use third-party storage services as well.

Safety of content stored in the system is achieved through a framework for data replication across multiple locations (rules define how many copies to make, and which locations to put these in, with a typical strategy being 3 copies in 3 geographically separate locations). Integrity of content is checked both periodically, e.g., every 6 months, and also when content is accessed or ingested into the system. If corruption is detected then repair takes place by replicating a known good copy from another location. Both CRC and MD5 checksums are used and the frequency of checking and when to repair are configurable. Full details of our approach to data

integrity, including the issues involved and the techniques we use, can be found in other publications from the AVATAR project, notably papers presented at iPRES2008²⁵ and IBC2009²⁶. In this paper we concentrate instead on our service oriented approach to archives as services.

The interactions to and from the aggregated storage occur through GRIA (**Fig. 2**), which is an open-source service-oriented infrastructure (SOI) designed to support B2B collaborations through service provision across organizational boundaries in a secure, interoperable, and flexible manner. GRIA makes use of business models, processes, and semantics to allow service providers and users to discover each other and negotiate terms for access to archive services. Service providers and customers trade resources (applications, data, processing, storage) under the terms of bilateral SLAs, which describes QoS and a promise to provide services, for instance to store and provide access to data for a particular period of time. Different client applications interact with the storage through GRIA, each of which fulfills different users and their roles. The first, the content player application, is the simplest of these, allowing users to access content in the archive but not modify or add to it. The second, the content portal, allows users to ingest new content and the “archive operator” to control and administer the storage access rights and policies. Ingest uses the OAIS model. For example, a provider uploads a submission information package (SIP) to the service provider through a data submission session. The SIP includes the content and preservation information (e.g., the retention schedule). Likewise, content access also uses the OAIS model. For example, a consumer downloads a dissemination information package from an OAIS service provider through a data dissemination session.

Ingest and access is managed according to the SLAs. This is essential for preservation activities using storage services because they need to run efficiently and dependably, to prevent unnecessary risk to the content. This is done through the storage adapters, where instrumented storage and data transfer systems communicate data-centric metrics, for example, I/O (max, min, average), storage usage, frequency of access, latency, etc. Control points allow GRIA to react to this information to manage the services, such as stopping access or upload, throttling bandwidth, or giving different users priority over each. Rules and policies within GRIA encapsulate how to go from the reported metrics through a series of decision points that invoke these control points, such as limiting the volume of content submitted each month according to the agreed terms of the service. WSSecurity is used as the basis of security and GRIA provides both transport (SSL) and message-level security (X509 or SAML). The use of SAML tokens allows GRIA to federate security policies between domains using WS-Federation patterns. This can be integrated with local security management, such as LDAP or Active Directory, at the client and server sides, to allow dynamic and automatic access control between organizations. For example, a content owner could set a policy of who can access their content, including people in other organizations that they trust. This can be dynamically and automatically propagated to the access control mechanisms used for data delivery.

In developing the solution, it was found that there are often significant differences between the parameters with which storage services are defined (storage capacity, access latency, delivery bandwidth etc.) and the level at which archive operators characterize their archive (rates and volumes for ingest and access, retention scheduling to encapsulate value, preservation priorities and asset safety). To address these differences, the storage dashboard tool (**Fig. 3**) allows archive requirements to be specified using parameters (e.g., data volumes and data i/o) that are both application and technology implementation neutral. The tool can be used by an archivist,

external service provider, or in-house IT manager to define SLAs in archivist terms or to interpret resource-level SLAs.

Through a series of screens, the user can specify one or more collections of assets and the associated ingest, access, and retention profiles. For example, a collection might be born digital content of a particular genre or it might be a particular type of analog carrier being migrated into digital form in a preservation project. The ingest profile specifies the rate at which items are put into the archive and can be expressed in various ways, such as items per month or terabytes per year. The access profile specifies how often material is likely to be accessed and can be expressed as an average rate or as a periodic activity. The retention schedule specifies how long each item of content needs to be retained before it is re-appraised and includes an estimate of how much content is likely to be retained after that point. Ingest, access, and retention profiles are aggregated across the collections to define the overall needs of the archive. The workbench allows simple storage solutions to be simulated (e.g., tape libraries) using technology roadmaps (e.g., LTO data tape) to profile investment and migration and find deviations from the archive needs, resulting from device contention during concurrent migration and access.

Next Steps

In the next phase of the project, to further support the SLA aspects of the project, the plan is to develop a combination of process modeling and statistical techniques to calculate the workloads placed on an archive from the processes that involve the archive, including ingest, access, transcoding, and maintenance (e.g., through migration). This will combine workflow specification languages and enactment engines, queuing theory, and Monte Carlo simulation techniques to analyze the variability of archive workloads and hence the flexibility needed in the systems used to implement the archive. The use of more advanced requirements estimation will form the basis of roundtrip capacity planning, SLA definition, archive service provisioning, and service usage auditing and reporting, including cases in which archive hosting is outsourced. The final objective is to demonstrate a decision support tool (dashboard) for planning, monitoring, and analyzing archiving with wider digital content infrastructures in a way that allows suitability, flexibility, scalability, and cost to be investigated, trade-offs to be explored, and best-fit solutions to be chosen from the perspective of both the consumer and service provider.

Conclusion

A move toward service-oriented and federated archive systems brings with it several challenges. There is a need for archive managers to communicate the requirements of the archive to the technical implementers of archive systems, whether in-house or outsourced. There is also a need for tools supporting capacity planning over long timescales to ensure the IT systems are sufficiently scalable but also planning with fine granularity to ensure systems are robust and flexible to peak loads. Finally, archives, as services, need to be embedded within content-centric environments and deployed across administrative domains with well-defined and automatically managed SLAs and QoS specifications. These challenges are of course interconnected and can be studied by analysis of the processes within and surrounding the archive. This yields an understanding of how embedding archives into AV production environments affects the size, growth, and services required from the archive. Planning technical solutions is currently a skilled, labor-intensive, time-consuming, and error-prone activity. In particular, there are problems of estimating the requirements, such as usage, volumes, and access, and how they vary

over time and with peak loads. Failure to do this can result in overloads at operation time that if not managed, lead to problems, which can put archive content at risk. On the other hand, if this analysis is done effectively, the results provide a basis for policies and SLAs when automating management of the archive.

AVATAR-m addresses these challenges through the use of aggregated and federated storage, an SOI to access and manage this storage, and user interface tools to help with capacity planning and decision support. This allows archive owners to concentrate on the long-term management of their content in a secure, safe, and cost-effective manner.

Acknowledgments

AVATAR-m is a £3M U.K. Technology Strategy Board supported collaborative R&D project that began in November 2006 and runs until 2010. The project involves the BBC, Ovation Data Services, Xyratex, and the IT Innovation Centre. For more information, please contact avatar-m@it-innovation.soton.ac.uk.

References

1. http://backstage.bbc.co.uk/news/archives/2007/09/bbc_dmi_project.html.
2. Digital Preservation Lifecycle Management for Multi-media Collections. Arcot Rajasekar, Reagan Moore, Fran Berman and Brian Schottlaender. Book Chapter. LNCS <http://www.springerlink.com/content/p3m497wt5452012u/>.
3. CLIR pub106 Building a National Strategy for Preservation: Issues in Digital Media Archiving <http://www.clir.org/pubs/abstract/pub106abst.html>.
4. Carl Rauch et al, "Evaluating Preservation Strategies for Audio and Video Files," http://www.ifs.tuwien.ac.at/~andi/publications/pdf/rau_digrep05.pdf.
5. A Miguel Ferreira, Ana Alice Baptista, and José Carlos Ramalho, "Foundation for Automatic Digital Preservation, *Ariadne*, Issue 48, July 2006, <http://www.ariadne.ac.uk/issue48/ferreira-et-al/>.
6. CCSDS 650.0-B-1: Reference Model for an Open Archival Information System (OAIS); *Blue Book*. Issue 1, Jan. 2002. This Recommendation has been adopted as ISO14721:2003. <http://public.ccsds.org/publications/archive/650x0b1.pdf>.
7. David Cavena et al, "Archiving Movies in a Digital World," *Sun Microsystems Report*, Version 2.1, pp. 29, June 8, 2007; David.Cavena@Sun.com; <http://wikis.sun.com/display/SunMediaSpace/2007/11/05/Archiving+Movies+in+a+Digital+World>.
8. "The Digital Dilemma: Strategic Issues in Archiving and Accessing Digital Motion Picture Materials" Academy of Motion Picture Arts & Sciences, 2007; pp .74, available from the academy: <http://www.oscars.org/contact/council.html>.
9. IEEE Spectrum: Google Watches Its Watts <http://spectrum.ieee.org/jul07/5382>.
10. <http://www.ascentmedia.co.uk/>.
11. <http://www.fpdigital.com/>.
12. <http://www.research.ibm.com/wsla/>.
13. <http://www.ogf.org/documents/GFD.107.pdf>.
14. FP6 NextGRID Project. <http://www.nextgrid.org/>.
15. FP6 TrustCOM Project. <http://www.eu-trustcom.com/>.

16. GRIA, <http://www.gria.org>.
17. DSpace <http://www.dspace.org/>.
18. Fedora <http://www.fedora.info/>.
19. Greenstone <http://www.greenstone.org/>.
20. ExLibris digital library solution: <http://www.exlibrisgroup.com/>.
21. Multiple User Post-Production IT Services (MUPPITS) <http://www.muppits.org.uk/>.
22. PeRvasive Infrastructure and Services for Media (PRISM).
<http://www.bbc.co.uk/rd/projects/prism/index.shtml>.
23. BeInGrid Movie Post Production Workflow Experiment <http://www.beingrid.eu/be2.html>.
24. <http://www.signiant.com/dmdms>.
25. Matthew Addis and Ant Miller, "The Significance of Storage in the 'Cost of Risk' of Digital Preservation. Richard Wright," Proc. *iPRES 2008: The Fifth International Conference on Preservation of Digital Objects*, London, U.K. 2008.
26. M. Addis, R. Lowe, L. Middleton, N. Salvo, "Reliable Audiovisual Archiving Using Unreliable Storage Technology and Services," *Proc. IBC2009*, Amsterdam, The Netherlands, 2009.

Figures

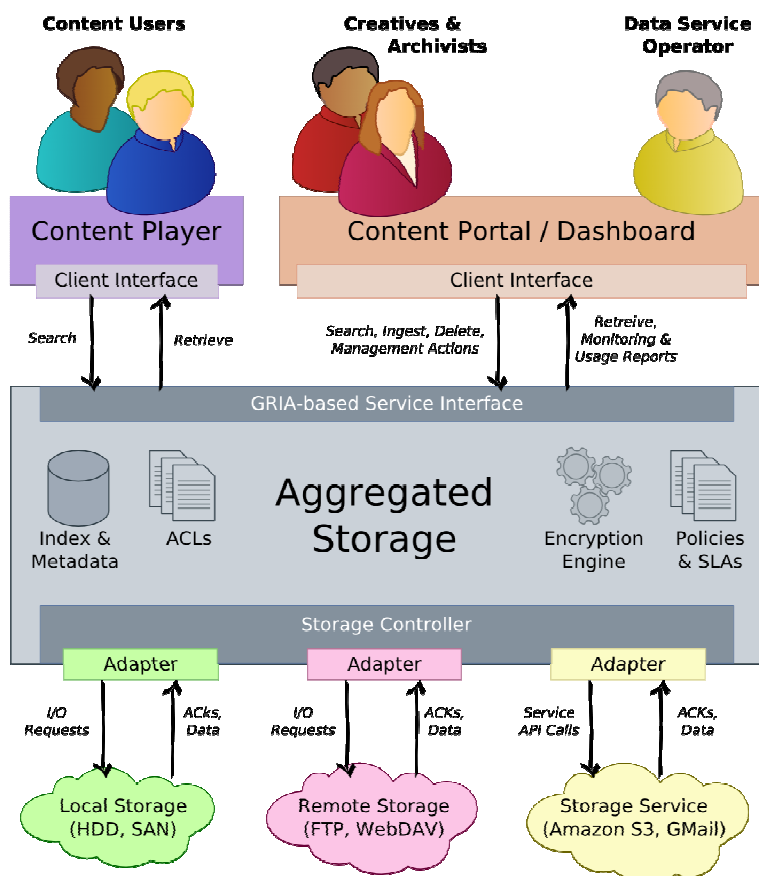


Figure 1. AVATAR-m aggregated storage.

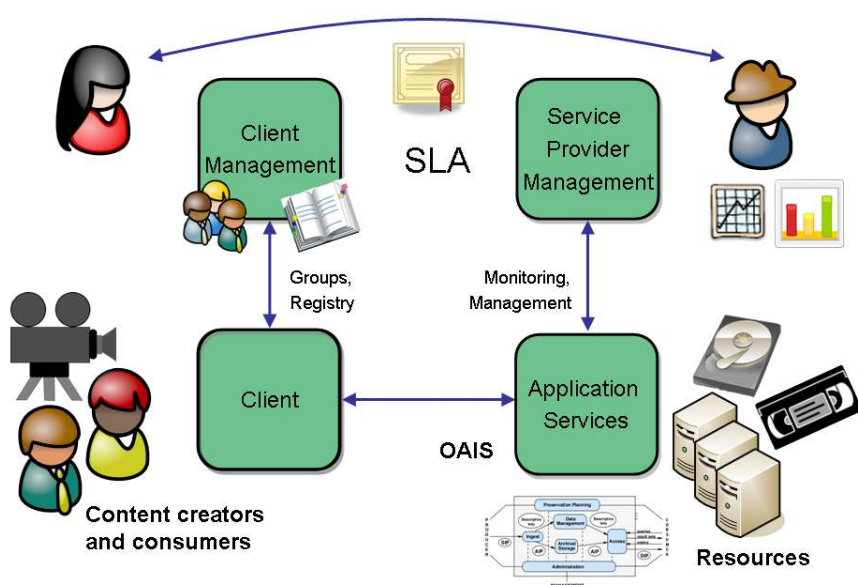


Figure 2. GRIA service-oriented framework as used in AVATAR-m.

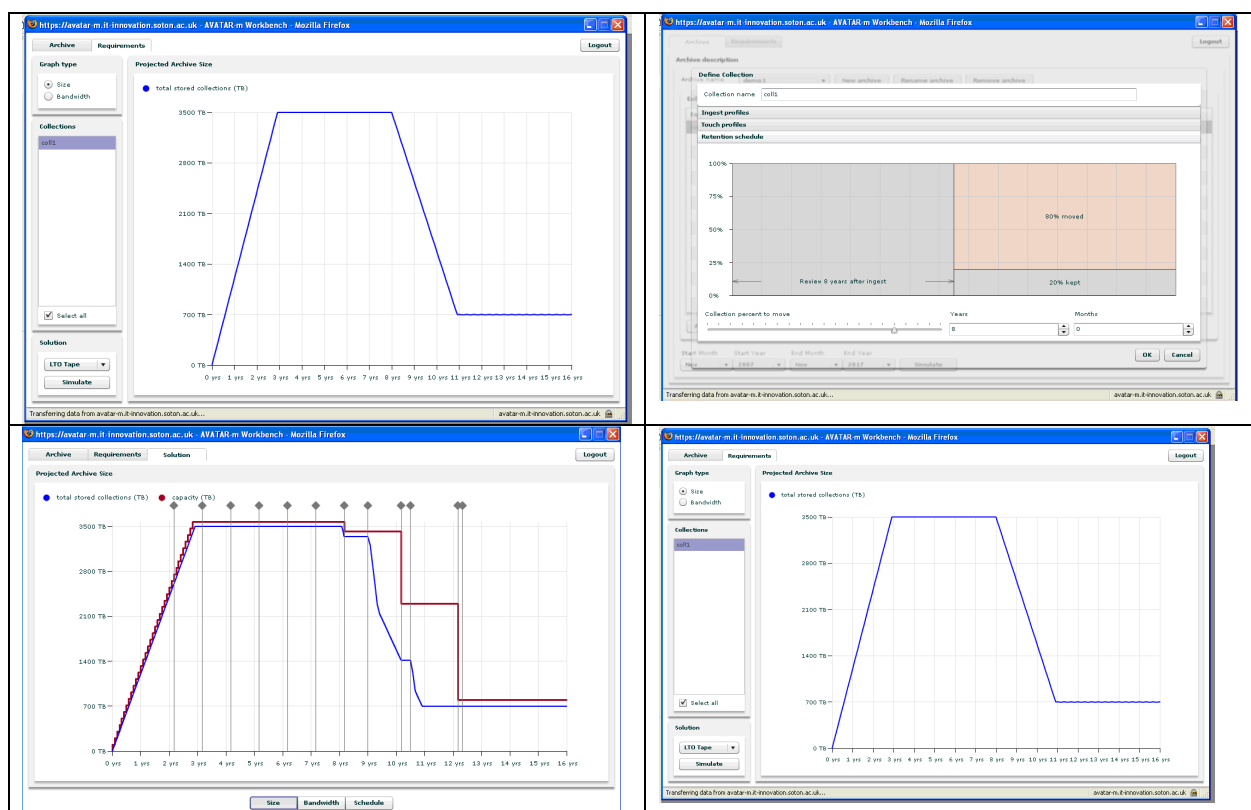


Figure 3. Archive requirement specification using the storage dashboard.

The Authors

Matthew Addis is a research leader and project manager at University of Southampton IT Innovation Centre. Through both national and European collaborative R&D projects, Addis has spent over ten years working with broadcast and film archives to develop strategies, models, tools, and systems for audiovisual restoration, preservation, and digital archiving. Addis currently works on service-oriented approaches to distributed archiving and storage for media-centric applications. His other work includes multimedia digital library technology and developing cost and business models for the preservation and access of cultural-heritage content. Addis has a Ph.D. in physics.

Richard Beales is a senior research engineer at University of Southampton IT Innovation Centre, where he manages collaborative R&D projects involving the film, broadcast, games, and cultural heritage sectors. Beales has a keen interest in media production, semantics, and pervasive computing technologies and is passionate about identifying practical applications of research results from these areas that can deliver commercial benefit to industry. Before joining IT Innovation in 2005, Beales worked as a digital media producer in the advertising industry. He gained a Ph.D. within the Intelligence, Agents & Multimedia group at University of Southampton and holds a B.Sc. in electronic imaging and media communications from University of Bradford.

Richard Lowe is a research engineer at the University of Southampton IT Innovation Centre. Since 2005 Lowe has developed software tools for U.K. and European R&D projects in cultural heritage, digital media, post-production, archiving, and grid- and cloud-based computing applications. Before joining IT Innovation Lowe read for a Masters by Thesis at the University of Hull where he also gained a B.Sc. (Hons) in computer science.

Lee Middleton has a Ph.D. from the University of Auckland, New Zealand, specializing in computer vision where he derived a novel addressing scheme and associated framework for hexagonally sampled images. He moved to the University of Southampton (U.K.) initially to take a research fellow position in biometrics and then to work at the IT Innovation Centre. During this time Middleton has worked on projects that include biometric schemes for gait recognition; a vision based environment for capture of face and gait; human detection and tracking in video; and data replication strategies for long-term storage of audiovisual data. He is a reviewer for several IEEE journals (IEEE Transactions on Image Processing and PAMI) and conferences.

Charlotte Norlund has a background in computer systems engineering and spacecraft technology, initially working as a research engineer at the University of Southampton IT Innovation Centre. During her tenure at the IT Innovation Centre she worked with broadcast and film archives, developing models, tools and systems for digital archiving and also worked with the interoperability of in-situ sensors and sensor networks. Norlund has a keen interest in large-scale modeling and simulation and is currently completing a Ph.D. in astronautics at the University of Southampton, where she is looking at large-scale disaster modeling and evacuation simulation modeling.

Zlatko Zlatev is a research engineer at the University of Southampton IT Innovation Centre, where he is involved in a number of collaborative R&D projects. Zlatev's current work is in the areas of data fusion, statistical learning and systems performance modeling and estimation. He holds an M.Sc. degree in machine learning and data mining from the University of Bristol in the U.K. and also an M.Sc. degree in computing machines and technologies from the Technical University of Varna, Bulgaria.