Knowledge Media Institute

KMi Stadium: Web-based Audio/Visual Interaction As Reusable Organisational Expertise

Marc Eisenstadt, Simon Buckingham Shum& Adam Freeman

KMI-TR-31

September, 1996

**Workshop on Knowledge Media for Improving Organisational Expertise,
1st International Conference on Practical Aspects of Knowledge Management,
Basel, Switzerland, 30-31 October 1996.**

[http://kmi.open.ac.uk/kmi-abstracts/kmi-tr-31-abstract.html]

The Open University

Eisenstadt, M., Buckingham Shum, S. & Freeman, A. (1996). ***KMi Stadium: Web-based Audio/Visual Interaction As Reusable Organisational Expertise.*** Workshop on Knowledge Media for Improving Organisational Expertise, 1st International Conference on Practical Aspects of Knowledge Management, Basel, Switzerland, 30-31 October 1996. [http://kmi.open.ac.uk/kmi-abstracts/kmi-tr-31-abstract.html]

KMi Stadium: Web-based Audio/Visual Interaction As Reusable Organisational Expertise

Marc Eisenstadt, Simon Buckingham Shum & Adam Freeman

Knowledge Media Institute
The Open University
Milton Keynes
MK7 6AA
U.K.

Tel: +44 1908-653800
Fax: +44 1908-653169
Email: {M.Eisenstadt, S.Buckingham.Shum, A.J.Freeman}@open.ac.uk
WWW: http://kmi.open.ac.uk/

Summary

KMi Stadium is a Java-implemented medium for hosting distributed events on a very large scale on the Internet (or an Intranet), allowing thousands of simultaneous participants even over 28.8Kbps dial-up modems. Stadium makes available as a reusable resource audio, coordinated visuals, and secondary resources such as relevant documents, demonstrations and Web sites. Client-based desktop computers and set-top boxes with appropriate browsers can download custom applets which enable the client machines to participate in presentations and other events mediated by a linked-server network. Most of the available bandwidth is dedicated to audio delivery and custom sound effects to help capture the mood of live events. After describing the design philosophy and implementation of Stadium, we consider its niche in the design space of organisational knowledge systems. With very low capture overheads, it enables organisations to make better use of the invaluable resource that can be found in expert speakers' presentations, and makes it easier for any staff member or team to share their expertise in a small, medium or large setting. We briefly consider its possible impact on working practices, both with respect to staff as knowledge consumers, and as knowledge creators.

Introduction

KMi Stadium is an experiment in very large scale telepresence. We are enhancing existing media and developing new media intended to give participants a sense of 'being there' at events of all kinds, including master classes, performances, tutorials, conferences, workshops, ceremonies, parties, jam sessions, recitals, industrial training sessions, university lectures, training on demand, town hall meetings, debates, and so on. Part of our motivation is to help restore the crucial element of human-human communication so often absent in multimedia distance teaching/training environments, and known to be a key ingredient of the UK Open University's success in pioneering open and supported distance learning during the past 26 years.

Experiences with the Virtual Summer School [3,4] convinced us that remote events could be undertaken successfully on a massive scale, provided that we could (i) convey a genuine sense of presence, (ii) harness good quality audio to impart not only the key 'content' of what was taking place but also special crowd moods, and (iii) design a scaleable architecture from the ground up.

For us, ‘being there’ is not primarily about Virtual Reality per se, although VR can certainly help. Rather, it is a question of capturing the right participative aspects of audience presence (such as applauding, laughing, shouting, asking questions, whispering to neighbours) and harnessing those aspects to convey as much of the mood of an event as possible. We are interested in telepresence at both live events and on-demand replays, because we believe that both types of event are enhanced by a sense of the presence of others. We experimented with synchronous computer-mediated conferencing, ‘chat rooms’, avatar-style virtual world environments, and various forms of videoconferencing, and became convinced that we needed something that would facilitate one-to-one, one-to-many, many-to-one, and many-to-many interactions, and help capture a realistic group mood without necessarily allocating all available bandwidth to large numbers of simultaneous video or audio streams. Moreover, we wanted something that could be delivered to desktops in a range of user-selectable bandwidths, from 28.8K modems (our minimal entry point for acceptable streaming audio) up to full 155Mbps ATM networks.

The result of our efforts, KMi Stadium, is not only a software prototype, but also a metaphor for a variety of activities we are undertaking. When we talk about ‘hosting an event in the KMi Stadium’, we mean that we provide the (free) client and (proprietary) server software in which customisable and scaleable events of various kinds can take place, ranging from small-room ‘fireside chats’, to enormous football-stadium-like activities, as appropriate. We use ‘stadium’ as the all-encompassing metaphor, even when it implies a much larger event than we are hosting at any moment, because it conveniently reminds us that scaleability is one of the major challenges we face. Since the ‘veneer’ presented by KMi Stadium is customisable, running small-scale events is easy within the KMi Stadium framework. Large-scale events rely on linked servers distributed across several continents. Our immediate goal is to host an event with 100,000 participants by the end of 1996. The three keys to scaleability are (i) the distributed server environment, (ii) local cacheing of all ‘special effects’ such as laughter, applause and slide shows, and (iii) a simple hierarchy of moderators and meta-moderators to field audience questions and comments.

How We Do It

The discussion which follows assumes that a main ‘presentation’ is taking place, typically including live audio and a slide show, with the audience responding via applause and laughter, and asking questions. We are most emphatically *not* advocating a ‘chalk and talk’ style of teaching in general, but many of our collaborative partners have requested this highly familiar mode, so this is what we have implemented. Other kinds of event are possible, as described later on.

We decided not to rely upon a high-bandwidth multicast backbone (an excellent vehicle in its own right for suitably-equipped sites), but rather on what is known as a ‘co-channel’ approach involving both local and distant media. In this approach, significant special-effects graphics and sounds are downloaded and cached in advance on each client machine (either over the net for high-bandwidth attendees or by CD-ROM distribution for low-bandwidth attendees). Then, live events and ‘tutor-facilitated on-demand replays’ are run under ‘remote control’, with low-bandwidth audio and mouse events driving synchronous presentations among all clients (clients can of course run their own private on-demand viewings if they don't want the live group atmosphere).

Our Java-implemented clients handle all of the ‘atmosphere’ locally. This includes the look and feel of the Stadium (or auditorium, or classroom, as appropriate), including locally-run animations which depict heads bobbing in the crowd, doors opening, flags waving, and even moving image sequences of the presenter(s). Higher-bandwidth clients, for instance on a corporate intranet, can run live ‘streaming video’, but in anticipation of the worst case (28.8K dial-up modems) the Java client supplies the animated special effects locally, and preserves the valuable bandwidth for the main audio stream.

Significantly, our sound engineers at the BBC Open University Production Centre have created a custom ‘crowd noise synthesizer’, which contains a vast array of crowd noises sampled in different circumstances. For each attendee, an array of buttons at the bottom of the display is used to indicate ‘applaud’, ‘laugh’, ‘boo’, etc. These effects are all available on the local client, so the nearest server can poll all of its clients in a reasonable sampling window, and via a quick communication with the other linked servers instruct all of the clients to play (locally) the appropriate crowd noise(s), in a nearly-faithful replica of how many people are applauding/laughing/booing during that sampling period. Even with the inherent communication delays, a monitoring meter allows both speaker and audience to sense the mood of the moment, given one can

tolerate a 10-to-20-second sampling delay (now the subject of empirical study). We are also experimenting with patching the crowd noise directly into the main server stream, to greatly reduce the sampling delay and help synchronize the presentation with the crowd mood.

To partipate in an event, attendees pass through a simulated entrance lobby in which they can select a cartoon face from a custom 'identi-kit' or 'photo-fit' style tool, or supply their own GIF images. The cartoon face is better for scaleability, because all clients come with the same full identi-kit range, and this is easier to manage remotely across the Internet. During the main presentation, only 'backs of heads' are viewed, and large-audience animation with suitable crowd noise contributes to the overall illusion.

The main presenter can run his or her presentation from a desktop presentation environment, and can also pass control to an audience member, who then appears to join the presenter on 'stage'. The audience member can pose a question via an audio feed (if sitting at a suitably-equipped computer, running Internet Phone, WebTalk, or a similar type of software), and this is then fed back out to the rest of the audience over the Internet using our own custom mixer and our own *RealAudio* server [9]. Peer groups can form spontaneously, and move off to separate 'syndicate rooms' for private discussions. Teletype 'chat room' interactions are trivially supported, but not encouraged.

Asynchronous on-demand replays controlled entirely by client 'pull' (i.e. requiring no mentor to do the driving) are easy, but outside the scope of what we wanted to achieve with KMi Stadium, where synchronicity is crucial to the sense of crowd presence. Unlike commercial 'remote control' software such as Timbuktu and PC-Anywhere, we can guarantee scaleability by distributing only a few bytes (to convey mouse-clicks) over the internet, and allocating the rest of the available bandwidth to streaming audio, using a commercial server provided by Progressive Networks' RealAudio. This technology allows us to distrubute up to 1000 simultaneous audio streams from a suitably-equipped and suitably-licensed server configuration, and through a series of linked servers we can avoid attendees clogging up one site.

The linked servers are themselves implemented entirely in Java, and deal with the client polling, secure access (e.g. only authorised students may be able to attend certain events), dishing out the audio streams, and synchronizing the presenter 'mouse clicks' with the behaviour of each client. The resulting architecture is inherently scaleable, is platform-neutral, responds in nearly-real-time (up to 20 second delay in the worst case), and is deliverable anywhere on the Internet which can guarantee 28.8kbps dial-up access. Each audio stream typically requires 10k of bandwidth, so serving 1000 streams can require 10Mbps in the worst case, and is the current limiting factor. We rely on 'like minded' linked sites that can deliver this capability, and at the moment this is arranged on an ad hoc basis.

Examples of Use

KMi Maven of the Month: This is a 30-minute interview with a leading knowledge media researcher or personality ('maven' = expert or connoisseur), conducted live in the first instance and then available as an on-demand replay, with opportunities for group discussion both during and after the event. These interviews extend the 'talk radio' format (and indeed we refer to the transmission medium for the audio portion of these events as KMi Radio) by adding static and updatable images, slow-frame video, QuickTime movies, audience questions and sound effects to enhance the presentations. Unlike Internet Relay Chat or conference/auditorium events held by the big on-line service providers, we place extensive emphasis on auditory and visual communication, and encourage visual and verbal interactions to take place anywhere in the world. (Figure 1)

Figure 1: KMi Stadium in its non-Java prototype form, being used for a large scale event. "Maven of the Month" is a 30-minute interview with a leading knowledge media researcher or personality, conducted live over the Net in the first instance and then available as an on-demand replay, with opportunities for group discussion both during and after the event. (So far, ten events have been broadcast using this prototype version; the Java enhancements for audience feedback, note-taking and break-out sessions are illustrated in Figure 3).

Professional Update Master Class: For professionals who need fast-turnaround technical updates and career enhancement training, KMi Stadium provides a master class environment in which a renowned subject specialist can provide personalised tuition, even when great distances are involved. This model builds upon the Open University's acknowledged expertise in distance learning methods, and offers a 'mix and match' approach to team up subject experts with customers expressing specific needs. Unlike videoconference tutorials, KMi Stadium master classes facilitate ready 'desktop involvement', serve the needs of self-organising tutorial groups and enable on-demand replays for interested parties anywhere on the Internet. Moreover, scaleability is designed in from the beginning. (Figure 2)

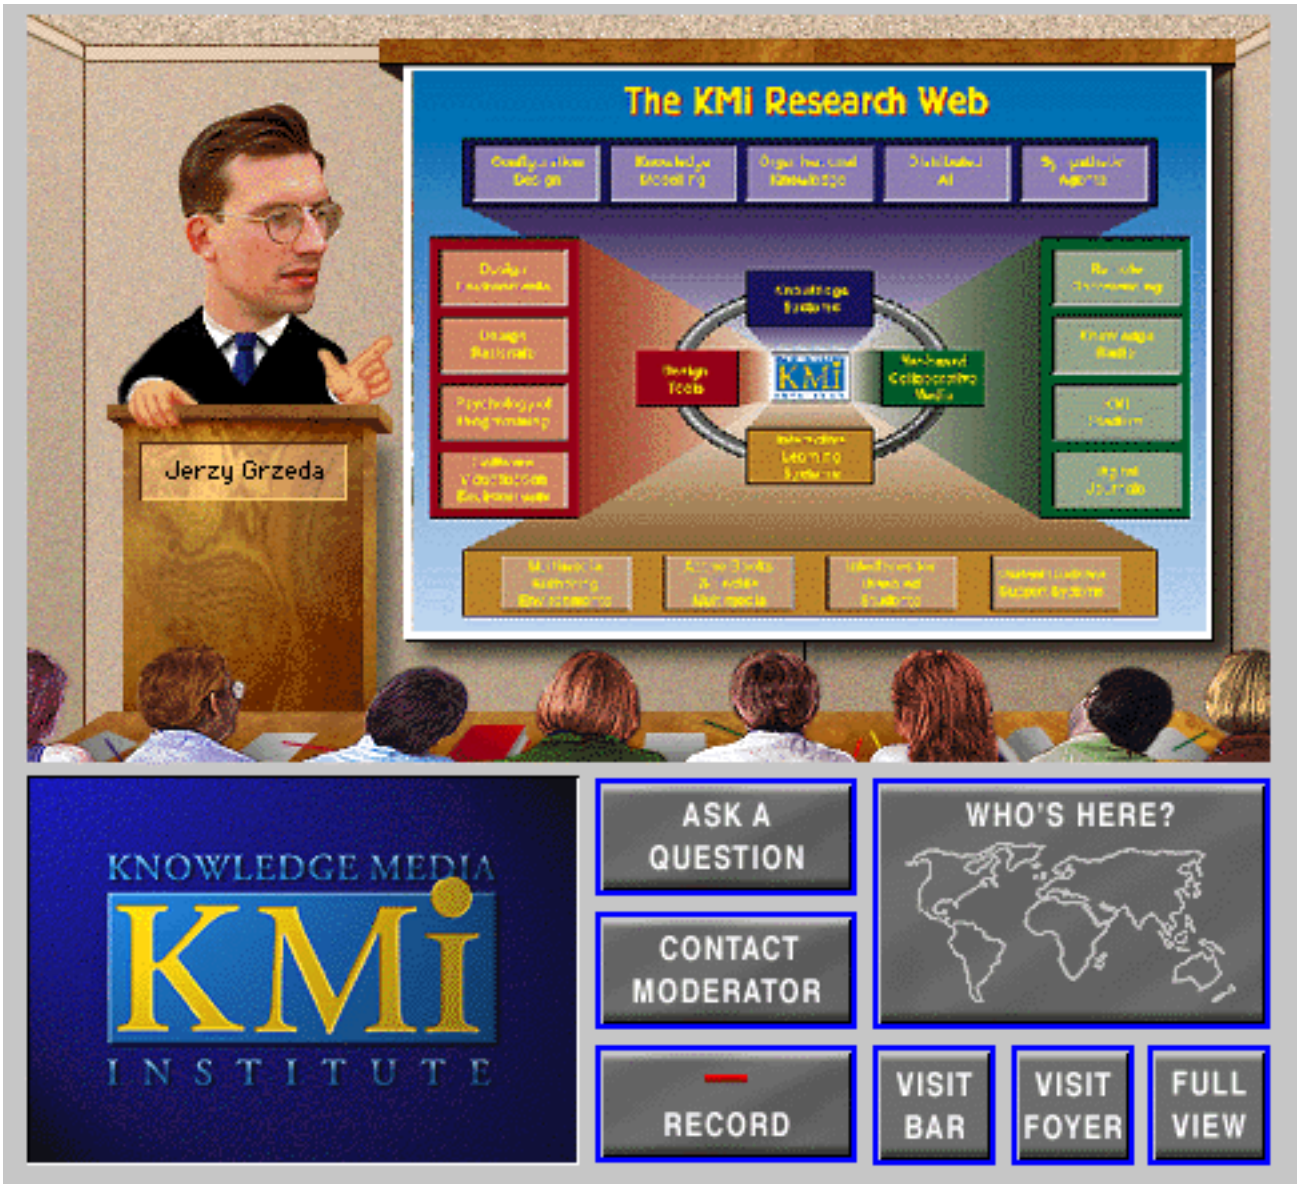

Figure 2: A 'Masterclass' environment for smaller groups, allowing the speaker to present visual aids and interactive demonstrations, which are synchronized with the audio stream live, for subsequent replay.

International symposia: Large international conferences and symposia are a natural activity to be hosted in KMi Stadium. Remote attendees can participate from around the world at their own desktops, and the presenters can benefit from the composite feedback provided by the KMi Stadium Server. Consider a symposium presentation involving 2,000 attendees on-line. If half of those attendees 'applaud' via their personal Stadium client software, then a 'canned applause' accurately representing 1,000 (out of 2,000) people applauding is faithfully reproduced not only via the server, but also via each user's client with a negligible delay. Every client receives synchronised media events as well (including slide shows, audio, movies), and the interaction with the presenter is limited (as in a real symposium) only by the number of questions that can be fielded by the moderator. (Figure 3)

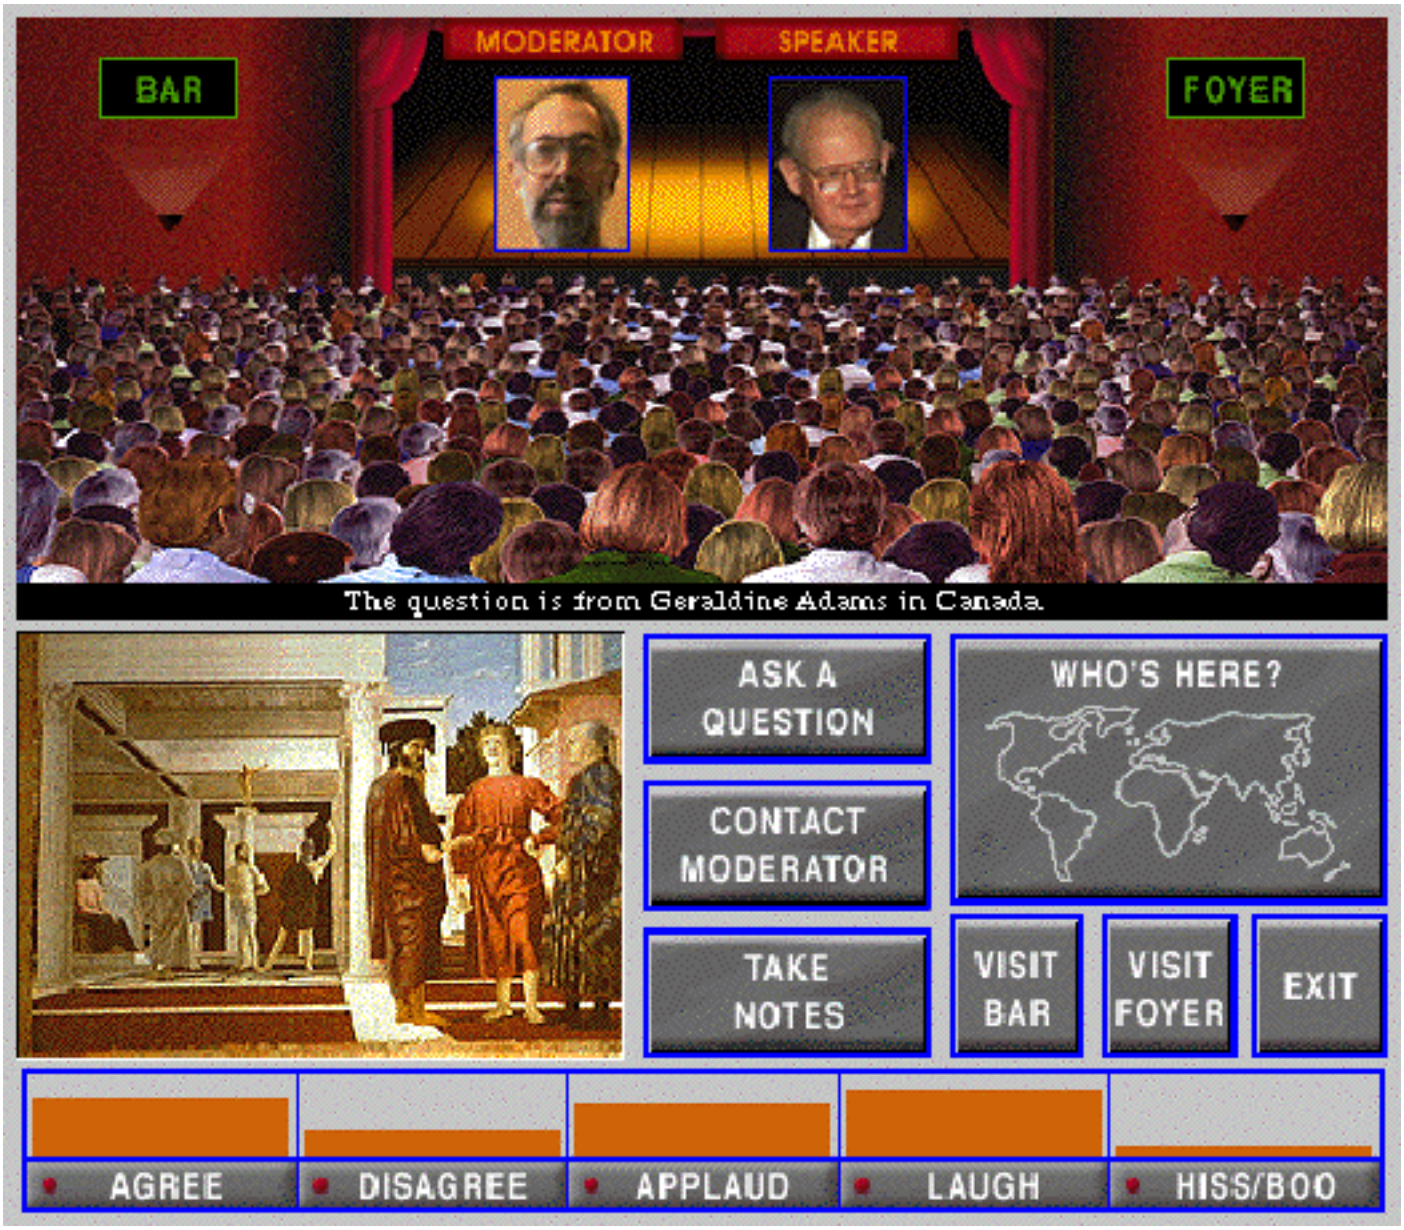

Figure 3: A conference/auditorium environment in Stadium. Whilst listening to the speaker, or to speaker-viewer discussions, the user can view the speaker's visual aids, email a question, hear and contribute 'emotional feedback' (agree/disagree/applaud, etc, with appropriate sound effects.), check who else is attending the event, or go to one of the other 'rooms' (see Figure 4).

Syndicate: Large workshops or conferences often spin off smaller 'syndicates' in which handfuls of participants congregate in separate meeting rooms to focus on specialised issues. The syndicates each select a 'rapporteur' who reports back to the main conference session. KMi Stadium provides arbitrary side-meeting rooms for just the purpose. Self-selecting groups can congregate and engage in smaller sessions, with the rapporteur running a reduced-scale version of the Stadium Server. By relying on just the rapporteurs to report back to main conference session, an important step towards scaleability is achieved. Within each syndicate, multiple audio conversations are allowed, with a designated person authorising turn-taking. (Figure 4)

Figure 4: A meeting room setting, either convened as a standalone event, or breaking out from a main Stadium event. A designated moderator controls turn-taking.

Town meeting: The emphasis in a town meeting is on the exchange of viewpoints, and it is therefore of crucial importance that KMi Stadium not be merely a 'passive transmission' medium. For this reason, the Stadium software includes several ways for attendees to make their views known, including standard 'voting' buttons and (as described above) the ability to provide feedback via such enhanced audio methods as applause and laughter.

Specialist and celebrity netcasts: A new kind of news 'mediumcast' is emerging, in which KMi Stadium can present events that are of fairly widespread specialised interest, yet not considered to be worthy of airtime on the major television networks. Examples might include the official launch of new charities by ministers or members of the Royal Family, sporting events of interest to expatriates, and public lectures by noted art historians.

20/20 Hindsight: Interested parties can participate in on-demand replays of any of the above activities at a convenient moment, and join a community of others to share and discuss the experience. Moreover, this works using technology that runs on today's Internet infrastructure.

Current Status

We have a fully working Java implementation of multiple servers and multiple clients supporting 'master class' mode with custom slide shows, synchronous audio, and audience sound effects, plus an audio-only 'maven of the month' veneer. We are working with a number of major vendors in order to begin our trials of more widely distributed linked servers and to begin scaleability experiments. We can simulate the linked servers on our SparcCentre-2000 configuration, and are currently undertaking monitoring and perfomance metering.

Improving Organisational Expertise with Stadium

Stadium is a generic technology, a kind of 'shell', which can be turned to many different applications, as indicated above. What applications can we see in the area of organisational knowledge capture and re-use?

As we discuss in another paper at this workshop [1], we need to begin to articulate key dimensions against which to compare organisational knowledge systems, and their associated 'intellectual auditing' methods. There are clearly differences between, for instance, capturing project lessons learned as audio-video records of team members (e.g. [2]), as opposed to developing a knowledge base of corporate knowledge (e.g. [6,8,10]).

We shall use the three questions proposed in [1] as a starting point for articulating Stadium's properties. These questions place great emphasis on the human dimensions of any organisational knowledge media, given that it is an inadequate understanding of these cognitive, group and organisational factors which so often cripple new work technologies. They ask:

1. *What classes of knowledge/expertise are addressed by this approach?*
2. *What representational scheme is proposed, enabling what kinds of analysis and computation, with what justification?*
3. *Who are the stakeholders? How will knowledge encoding and re-use impact their work practices?*

What classes of knowledge/expertise are addressed by this approach?

Stadium occupies a particular niche in the design space of systems for supporting organisational expertise. It makes it possible for an organisation to realistically include interactive AV presentations, or portions of them, as part of its indexible, re-usable intellectual resources.

Video recording a talk is rarely very informative; it is the speaker's *words* and their *coordination with visual aids and demos* which are of interest, and which have most communicative power. In Stadium, an audio/visual presentation can be captured live with minimal effort, and made available on demand over the Web.

Corporate training is an application already described above. The 'Masterclass' version of the user interface provides the environment for employees to tune in together or in their own time to be updated on some aspect of professional knowledge. Rather than merely expecting the trainer/expert to 'impart expertise', Stadium facilitates peer interaction in two ways:

- *Staff interaction and debate:* the tools available to presenters can be made available to audience participants. Thus, when 'control of the floor' is passed from presenter to audience participant, that participant then gains the ability to display graphics, broadcast live audio, and even delegate control onwards to other participants. This gives participants a unique ability to join events on a 'level playing field', and promotes learning as a constructive shared experience, rather than a one-way transfer.
- *'Pansynchronous' interaction:* this is a phrase we have coined to mark the deliberately-blurred distinction between 'asynchronous' (off-line or batch mode) discussions and 'synchronous' (on-line or

live) discussions. In Stadium, participants can observe 'on-demand replays' of pre-recorded events, yet these replays can be restricted to take place only at pre-arranged times, to ensure that a peer group and/or tutor are all on-line at the same time. Thus, the replay is asynchronous with respect to the original 'master presenter', yet synchronous with respect to the tutor and peer group, and all the benefits of shared experience can still be obtained.

As should be clear, an expert's AV presentation is captured and stored as digital audio plus coordinated visuals. In this form of expertise capture, the emphasis is strongly on preserving the richness of the personal, speech and slides. More abstract representations—the summary of the talk, reports by the speaker, related Web-resources—play a secondary role, facilitating access to greater detail as needed.

What representational scheme is proposed, enabling what kinds of analysis and computation, with what justification?

Digitisation allows us to make portions of the AV record quickly accessible, which has always been difficult with analogue AV. Stadium already incorporates a tool which enables the insertion in the audio stream, as it is 'cybercast' live, of triggers linked to slides as the speaker switches visuals. We hope to develop tools for browsing such an audio stream, and to allow the retrospective insertion of keywords into the stream, providing the possibility of indexing the presentation. One can envisage a scenario where various portions from several different talks are retrieved by a system in response to a search for expertise on a particular issue. The only danger in such circumstances is that what the speaker says might be taken out of context. However, just as with a quote from a book, it is the user's responsibility to ensure that they have correctly understood these retrieved 'samples' in context. Given the nature of the representation, little further computational analysis is possible which might merit further concern, from the perspective of a member of staff.

Who are the stakeholders? How will knowledge encoding and re-use impact their work practices?

What impact on working practices might a medium such as Stadium lead to? We consider the role of the *knowledge engineer* who has technical responsibility for constructing the corpus of material, the *speakers* who share their expertise, and the *consumers* of that knowledge. We suggest in particular, that Stadium facilitates traditional consumers to become knowledge *creators* as well.

Knowledge engineer

A strength of the Stadium approach is that it has a low 'knowledge capture' overhead, since it simply records audio and visual aids. There is no need for a knowledge engineer in the normal sense, to elicit and structure a more abstract representation of expertise. There is obviously some additional, though relatively small, overhead in the preparation of the final Web resource: checking of the audio-visual coordination; possible digitisation of videos; markup of the Web page linking to secondary resources; indexing of the audio if desired.

Speakers

Speakers must provide their visual aids in digital format, but this is becoming increasingly common with presentation tools like Persuasion and PowerPoint, which can now generate GIF images automatically. Speakers are also used to being audio and video recorded in analogue, but obviously, should be pre-warned that their talk is to be made available over the corporate intranet. Finally, speakers may find that they receive much more exposure, and many more follow-up queries than they would normally have!

Knowledge 'Consumers'

As one would hope, the greatest impact should be on other members of the organisation, especially those distributed geographically. We identify two forms of impact, as *knowledge consumers*, and as *knowledge creators*.

Firstly, as knowledge consumers, staff anywhere will now be able to tune in live, or after the event, and break into discussion groups. What impact might this have on working life? Certainly, we hope that it will make a resource, which was previously too hard to provide or re-use, much more accessible as a source of expertise. We know from empirical evaluation of the Forum presentation/discussion environment [7] that

employees enjoy the convenience of being able to 'drop in' to ongoing events without leaving their offices, and in particular like the ability to 'fast forward' a presentation, break out to check email, and drift back in during the portions of a live session that are expected (because of hopes raised during fast-forward browsing) to be of high interest. Another positive side of Stadium's use is the prospect of enormous savings incurred through reduced travel expenditure, and associated reductions in concomitant fatigue and time lost from work. However, one can also imagine possible negative consequences. For instance, if it became mandatory to connect to a regular presentation, because the management thought that Stadium was effective, staff might resent this. If the number of presentations increased from monthly to weekly because of Stadium, again this might not be appreciated by time-pressured staff. If travel budgets for face-to-face meetings were cut because staff could now meet in Stadium's 'meeting rooms', then again, this might be resented. Media shape work in unpredictable ways, but we are alert to these scenarios, both positive and negative, and will be tracking Stadium's impact as it begins to be used by knowledge organisations.

Secondly, as *knowledge creators*, we are entering an era of 'Training-on-demand' and 'Just-in-time Open Learning' provision, in which we need tools which enable us to create and deliver hot-off-the-press (or hot-out-of-the-mind) materials with unprecedented speed. We are investigating the role of a 'rapid-response network' of tutors who, like emergency medical staff, could be called into action whenever a critical mass of learners/trainees requested a session. We envisage that a medium such as Stadium could empower individuals or teams to deliver quality and up-to-date materials on demand, to more readily share their lessons learned with others, to help them not make the same mistakes, or to build on key strengths. We envisage short presentations in one of Stadium's small-medium settings (e.g. the meeting room, recording studio or lecture theatre), updating colleagues on progress on a project, negotiations over a key contract, or experiences working onsite with a client. We are developing agent-based technology in this area which could assist in automatically scheduling meetings (face-to-face; video conference; Stadium meeting room) to bring together knowledge 'creators' and 'consumers' who expressed particular interest [5].

Finally, the growing wealth of digital resources, and the means for their rapid dissemination, raises the prospect of an 'Educational Utility' in which individuals can mix and match resources, under expert tuition/guidance if desired, to their own tastes and needs. Two of the more challenging aspects of the Educational Utility concept are (i) concerns of quality and (ii) the legal and copyright implications. Because it is trivial for an individual to construct a meta-resource on the Web (a collection of pointers to the resources of others), these two issues are strongly intertwined. For example, someone could launch a new 'Introductory Biology via the Web' course, piecing together segments of (other people's) Web resources. In general, it is safe to assume that the less effort invested in this meta-resource approach, the more likely is that (i) the value-added quality will be low, and (ii) some copyright infringement *may* have occurred along the way. Conversely, care in assembling a good meta-resource, complete with support materials for students (and tutors!), increases the probability of high quality content and sensible rights attribution and clearance. A case in point is the legal dilemma surrounding 'fair use for comment and criticism' of short audio, video, or film clips (either on the Web or elsewhere). In general, clips that are embedded within a scholarly critical review are allowed without unduly restrictive licenses, whereas mere 'samples' or 'clones' are not. This legal framework, it turns out, is a way of ensuring a strong correlation between quality and legal usage (though this is by no means guaranteed!).

Conclusion

We have described KMi Stadium, a Web-based system for hosting interactive audio-visual presentations and forums. It makes available as a reusable resource audio, coordinated visuals, and secondary resources such as documents, demonstrations and Web sites. We have considered its potential as a medium for capturing and disseminating organisational knowledge, clarifying its niche in relation to other organisational knowledge systems. It enables organisations to make better use of the invaluable resource that can be found in expert speakers' presentations, and makes it easier for any staff member or team to share their expertise in a small, medium or large setting. It has very low capture overheads, and we are interested to assess the impact that it has on work practices. We therefore invite interested organisations who wish to trial Stadium to contact us.

Acknowledgements

This work has been made possible by generous support from Sun Microsystems and Apple Computer.

Stadium graphics are designed by Jon Linney.

'Maven of the Month' and other prototype Stadium events on the Web are made possible and run by Mike Lewis (KMi Webmaster) with assistance from Andrew Rix and Biff Jones (OU Audio-Visual Dept.), Phil Thomas (OU Academic Computing Service), and Andrew Law's team at the BBC's OU Production Centre (including Paul Roberts and Lynton Howell-Hughes).

References

1. Buckingham Shum, S. Negotiating the Construction of Organisational Memory Using Hypermedia Argument Spaces. In *Workshop on Knowledge Media for Improving Organisational Expertise, 1st International Conference on Practical Aspects of Knowledge Management*, Basel, Switzerland (30-31 Oct.), 1996 [http://kmi.open.ac.uk/kmi-abstracts/kmi-tr-30-abstract.html].

2. Carroll, J.M., Alpert, S.R., Karat, J., Deusen, M.S.V. and Rosson, M.B. Raison d'Etre: Capturing Design History and Rationale in Multimedia Narratives. In *Proceedings of ACM CHI'94 Conference on Human Factors in Computing Systems*, ACM Press: New York, 1994, pp. 192-197

3. Eisenstadt, M. The Open University Virtual Summer School. In *Proceedings of ED-MEDIA'94*, Vancouver, BC, Canada (June 1994), 1994, AACE: Association for the Advancement of Computing in Education

4. Eisenstadt, M., Brayshaw, M., Hasemer, T. and Issroff, K. Teaching, Learning and Collaborating at a Virtual Summer School. In *Remote Cooperation: CSCW Issues for Mobile and Tele-Workers*, Dix, A. and Beale, R., (Ed.), Springer: London, 1996 [http://kmi.open.ac.uk/kmi-misc/virtualsummer.html].

5. Eisenstadt, M., Watt, S.N.K. and Zdrahal, Z. A Practical, Modular, Scaleable, Platform-Independent Diary and Workgroup Scheduling Agent. In *Symposium on Intelligent Agents and Collective Intelligence, Exhibition and Symposium on the Interactive and Networked City*, Telepolis, Luxemburg (Nov.), 1995

6. Euzenat, J. Corporate Memory Through Cooperative Creation of Knowledge Bases and Hyperdocuments. In *Proceedings 10th Banff Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Banff, Canada (November, 1996), 1996, SDRG Publications, Dept. Computer Science, U. Calgary, Calgary, Alberta, Canada, T2N 1N4 [ftp://ksi.cpsc.ucalgary.ca/KAW/KAW96/22euzenat.ps.Z].

7. Isaacs, E.A., Morris, T., Rodriguez, T.K. and Tang, J.C. A Comparison of Face-to-Face and Distributed Presentations. In *Proceedings of CHI'95: Conference on Computer-Human Interaction*, Denver, CO, 1994, ACM Press: NY, pp. 354-361 [http://www.sun.com/tech/projects/coco/papers/forum-CHI95.html].

8. KAW'96 "Corporate Memory & Enterprise Modeling" Track. In *Proceedings 10th Banff Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Banff, Canada (November, 1996), 1996, SDRG Publications, Dept. Computer Science, U. Calgary, Calgary, Alberta, Canada, T2N 1N4 [http://ksi.cpsc.ucalgary.ca/KAW/KAW96/KAW96AccTracks.html].

9. Progressive Networks *RealAudio* [http://www.realaudio.com/].

10. van Heijst, G., van der Spek, R. and Kruizinga, E. Organizing Corporate Memories. In *Proceedings 10th Banff Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Banff, Canada (November, 1996), 1996, SDRG Publications, Dept. Computer Science, U. Calgary, Calgary, Alberta, Canada, T2N 1N4 [ftp://ksi.cpsc.ucalgary.ca/KAW/KAW96/25vanheijst.ps.Z].