Martin Russ, Ian Kegel, and Fred Stentiford

Smart Realisation: Delivering Content Smartly

The Internet has become a medium of interchange and cooperation, as well as distribution of information. But no current vision has taken the creative leap of abstracting the use of the Internet as a tool for distributed communications to its ultimate conclusion. Instead there are a number of initiatives which address specific needs but which fragment rather than consolidate. In order to be truly ready for the demands of users of these initiatives, and the others which will inevitably follow, telcos will need to have a sophisticated and ubiquitous solution in place. Without such a solution, their future role could be critically threatened. This article describes the development of just such a technology: *smart realisation*. It incorporates the essential elements as small parts of a large, coherent and complete solution across the entire communications and multimedia space. This media delivery solution strongly binds the content, communications, and the network, together into a flexible end-to-end process.

Background

The digital manipulation and distribution of media has changed the publishing business. The magazine, television, movie, music and book industries, and many other examples of what were once separate industries, have now been consolidated into large media companies. Portals on the Internet have changed the way that people perceive media and its availability[1]. This has all increased the perceived importance of content; the information which is presented as media.

But content is not everything, regardless of how often we are told that 'content is king'. Human beings continue to communicate using a wide variety of types of medium. The telephone and e-mail business is still many times larger than the movie, games and television businesses combined. However, just as digitisation has aided the merging of the publishing business, it is also capable of subsuming communications functions as well. A converged communications and media world is rather different from the existing partitioned one, and many assumptions will be proved wrong.

A converged communications and media world is rather different from the existing partitioned one, and many assumptions will be proved wrong.

The authors are with BTexact Technologies

Smart realisation: a vision of the next mass-market multimedia form?

As the number of electronic channels to market increases, publishing companies are faced with a problem. They want to deliver the same 'message' to people through web browsers, interactive TV channels, wireless application prototcol (WAP) telephones and other devices. All of these could be in different formats dictated by the characteristics of the user interface in the widest sense. They could use a number of proprietary methods of soliciting feedback, ranging from e-mailed web forms to human beings in telephone call centres. Unstructured, proprietary methods of achieving multimedia delivery, feedback and interaction could well result in expensive chaos, and requiring time-consuming manual reworking of the 'message'.

Smart realisation is a radical solution to this problem. The word 'smart' refers to the pragmatic use of computing technology, while 'realisation' refers to the way that media is put together for presentation to the user. By examining the multimedia publishing chain in detail, and developing ways of making the media content 'smart', then this work has provided British Telecommunications plc (BT) with leading-edge visibility of a technology which could potentially be the next mass-market media form. The value-chain analysis that has been carried out shows that, although the existing roles persist, new relationships and value flows are created by this technique. Developing pathfinder applications which exploit this technique will mean that we are well placed to support and influence this new market as it develops.

Being 'Smart'

So what do we mean by 'smart'? Although the word has been used by many authors in different contexts, we define 'smartness' as simply taking an intelligent approach to the presentation of an electronically-provided

service. Digital content of some description is often the centrepiece of such services, for example an agent searching for Internet information, a personal television service or a multi-player shoot-em-up on a games console. However, a service will inevitably contain transactional elements such as electronic payment, authentication or 'communications' functionality be it text, audio or video.

Fundamental principles

Our approach suggests three fundamental principles which can be applied throughout the service life cycle to enable greater flexibility, efficiency and a range of new service possibilities:

- breakdown into components;
- treat communications and content services as a continuum; and
- make content dynamic and autonomous.

Breakdown into components

Traditionally, services have been built as proprietary systems, published as a complete entity to be received by the consumer. Personalisation or customisation is achieved within the entity, or by selecting a new entity which has been fashioned in a more appropriate form. In the United Kingdom (UK), mobile Internet content has grown in this manner, with providers re-purposing large amounts of their hypertext mark-up language (HTML) content into the appropriate mobile format—usually wireless mark-up language (WML).

A more effective approach can be to break down the service into a series of distinct components which represent the 'essence' of a particular aspect of the service. A 'smart' content component can thus be created by surrounding a fragment of media with a shell of descriptive metadata plus a palette of methods and processes that can be applied to itself and represent its functionality[2,3]. (See Figure 1.)

Of course, the relative balance of media, metadata and methods will depend upon the service the component contributes towards. A video segment which might be used in a personal television service would contain a reference to the media, a detailed description of that media, and perhaps some decoding functionality. The latter could make available a high-definition video stream if being referenced from a television platform, or thumbnail image if being referenced as a search result on a personal digital assistant (PDA). Conversely, a textual 'chat' component would contain very little media content, apart from that pertaining to its presentation, while the functionality element would encapsulate the appropriate messaging interface.

Figure 1 A 'smart' content component

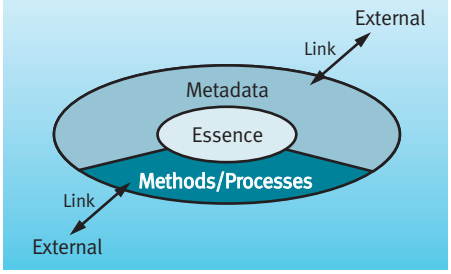

Treat communications and content services as a continuum

Services based around content and communications have historically been kept separate. Digital communication (arguably reaching back to the 'Needle' Telegraph) has been in existence since the 1830s, and telecommunications services have since become a part of everyday life for many people. The possibilities afforded by digital content, however, have only entered the public eye within the last 10 years, largely driven by the success of the Internet. Today, an understanding of the convergence of content and communications challenges the entire telecommunications industry, whose services are rapidly becoming complicated by the need to support a rich 'content experience' for their customers. Recent alliances such as that of America On-Line (AOL) and Time Warner illustrate attempts to come to terms with this challenge.

By allowing communications-related elements to exist as component parts of a service, the user experience will not involve a switch of context[4]. For example, a popular television programme could, when packaged as a service, contain audio communications components. These would enable customers to seamlessly enter a discussion forum, instead of the need to make separate telephone calls to friends and family as is common today.

Make content dynamic and autonomous

One of the hardest tasks facing all content providers is the requirement for continuous churn to maintain the interest of their customers[5]. While broadcast television demands a strict schedule for the customer to obey, the Internet offers no 'rules' for content churn and no guarantee of content quality. The story that many product successes and failures tells is essentially the common-sense view that customers want something in-between these two extremes[6]. Although scheduling is constrictive, a completely free-form world offers too much complexity and too much choice. For example, since approximately the middle of 1999, there have been more pages on the Internet than you could realistically visit in a lifetime. 500 million pages = 1 page viewed per second × 12 hour day × 260 working days × 45 years[7]. Controlling these conflicting demands while maintaining a viable business model is very difficult unless the content itself can do some of the work for you.

One way of making a content component dynamic is to put in place systems which can modify and update its metadata in accordance with external events. Given that metadata is responsible for the selection and presentation of the content component, we can therefore make news articles expire in a timely manner; shopping lists update according to the availability of stock; and make every component record detailed tracking data indicating which customers used it, for how long, and in what context.

The importance of being 'smart'

We live in a world where a number of conventions are accepted because they are familiar. But developments in digital networks and communications have now reached the point where they can challenge these preconceptions and may well replace them with radically new alternatives[8].

One example is the book. Imagine an alternative universe where you are allowed only a very small number of books. These books are kept in small cupboards, each housing a human reader, whose job it is to continually read the book aloud. To access a book, you merely open a cupboard, and listen. If you close the cupboard, then the reading continues, and when you subsequently open the cupboard again, you will have missed a portion of the spoken reading.

You may be smiling at the ludicrous nature of this approach to reading a book, but replace the word 'book' with 'television', and a very different scenario appears. Conventional broadcast television is transmitted continuously, and you need to keep watching in order to ensure contiguous viewing. But because television has always been delivered in this transmission-centric way, then it is acceptable. It is probably a good thing that books were not invented after television! These preconceptions are changing; and we are beginning to see television becoming responsive to audience identity and interaction[9]. It is also worth noting that while books and television programmes are created to be used in a specific way, many people use them in ways which are outside the original intention of the creator. For example, some people habitually read the end of a book first; magazines and newspapers are often read

from the back to the front; and one of the authors of this paper regularly watches movies on videotape in fast-forward mode, and so does not know any quotes, but has a sketchy view of the plot.

Recent consumer research[10] has developed the concept of 'media time' for the total percentage of the average day a person spends consuming 'media'. Over the last 10 years, the definition of 'media' has rapidly expanded as new devices, and in particular the Internet, have changed consumption habits. However, the total time has grown very slowly in comparison, indicating significantly greater competition for a share of this precious commodity, which is directly related to revenue opportunities.

Another popular model applied to digital content and services has been the *lean back, lean forward* model. This casts certain modes of consumption, usually related to devices, as passive ('lean back') or interactive ('lean forward'). But these stereotypes are beginning to change. While digital TV content is in its infancy[11], it is clear that interaction with the television is increasing and it has not been difficult to encourage this by means of simple applications such as e-mail. No one device need be the centre of 'lean forward' interactivity any more—consider it more the case that each service will have an imaginary 'interactivity handle', under control of the community of consumers as well as the publisher. Pushing the handle forward may involve joining an audio 'chat room' with the other viewers of your favourite soap opera or game show to air your personal views. Pulling the handle backward may involve leaving an on-line game to view its conclusion as a short movie.

However consumer activity is classified, new digital services are consistently set to challenge the status quo. Being 'smart' suggests adopting a more open state of mind to what is possible.

The Roles of Smart Realisation

By examining the fundamental nature of digital content and services, we have suggested three principles which form the basis of a very different model for how content and communications should be treated (breakdown into components, treat communications and content services as a continuum, make content dynamic and autonomous). We will now focus on three specific parts of the value chain which we have studied in the context of smart realisation. We propose that the exploitation of new ideas in these areas may soon represent a major part of future telecommunications business.

Publishing

In the previous section, it was stated that the traditional approach to the publishing of digital services delivers a single finished article, to all intents and purposes 'frozen' in the form intended by the publisher. So would a change to publishing small well-described content components be an insurmountable barrier to the adoption of smart realisation?

We believe that this need not be the case[12]. Consider a simplified example of a television production house creating reports for a news programme. The report's life begins when the lead is discovered and the structure of the story is drafted out. This will be translated into the appropriate information required for the crew and reporter to go to their chosen locations and record footage and interviews to camera. There may also be a need to record a voice-over back in the studio to cover some of the footage, and images and sound may be taken from library databases. When the editor comes to cut the report together, they will work with the 'shooting script' and database of media clips (which will probably now be transferred into a digital retrieval system) to create a single result which satisfies the producer.

However, what if the process was rewound to the stage just before the final report is committed to tape? The report is composed of a number of individual audio and video components which are well-described in their own right. In fact, the industry tells us that there is probably *too much* metadata available at this point. If the report is authored as a connected set of content components held together by a service 'wrapper', then that individual report could be exploited in a much wider context. For example, the content component containing the script could be combined with a single video frame to generate the web version of the report, while the video itself could be assembled in a variety of different ways according to the preferences of the television network and indeed the consumer. Both the 30-second version and the 3-minute version could be made available simultaneously.

While it is conceivable that all of this could be created from a conventional publishing process, it is important to realise the difference that all of the instances of the report come from a single set of components and a single authoring process. The content has been multi-purposed, and not re-purposed after publishing is complete. The business of post-production re-purposing is present across digital content industries, from mobile Internet content to television. However, there are examples of publishers who are beginning to break the mould. News journalists are embracing content and metadata expressed in the NewsML structure[13], and film director George Lucas, who strongly believes in 'dynamic creation' and 'non-linear filmmaking' is using the same content components for the film *Star Wars: Episode II* as for the video game[14, 15]. The concept of 'purposing', with respect to content creation and manipulation, has several facets. Skinner[16] offers a more detailed breakdown:

- *Recycling*: Content is reused in new, novel or unanticipated ways.
- *Recasting*: Content is provided in a context-aware manner to enrich the user experience.
- *Parallel (or multi-) purposing*: Content is developed with the intention of serving a variety of purposes.
- *Auto-development*: Content is capable of re-composing itself to meet changing circumstances.

Each one of these facets can be exploited as a distinct business advantage for the publisher.

Metadata

The example given above describes a publishing environment where there is an abundance of appropriate metadata. There will be many instances when this is not the case. Of course, generating metadata is expensive if it is done manually and one of the problems being addressed by many organisations is the automatic categorisation of content. In order to respond automatically to requests like: 'I need an image of a landscape with some people in it'[17]; or 'I want a shortened version of this video'[18]; or 'What is strange in this image?'[19]; these all require searchable metadata that adds value and smartness to the content.

The technical implementation of a system that exploits smart realisation will need to cope with metadata from a wide variety of different content domains. It is reasonable to expect that a British Broadcasting Corporation (BBC) news programme will publish different component metadata to the organisation which has provided a mapping services component allowing the customer to pinpoint the precise location of the reported event. We have identified and demonstrated suitable technologies for translating between these varying metadata schemas. The XSL (extensible stylesheet language) family[20] offers the capability to transform metadata which has been structured in XML (extensible mark-up language). One of two approaches can be taken here.

If a very small number of schemas are in use, a direct mapping between pairs can be most efficient. Alternatively, it is necessary to introduce a semantic repository agent (for example, see Reference 21) whose purpose is to provide transformation services to content components. To enable true multi-purposing of content components, the latter mechanism is preferable, if more vulnerable to errors caused by semantic drift. This condition arises when arbitrary decisions are made regarding the semantic interpretation of an element of metadata and its relation to a different schema. While the results of semantic drift can be minimised by good schema and transformation design, ontological relationships and taxonomic structures can also be established between content components. A detailed treatment of these is beyond the scope of this article, but it is important to realise that both forms of metadata can work together to ensure the right content components are available in the right time and the right place.

The techniques described within this section comprise what we have termed *open publishing*. Figure 2 summarises the key elements that we believe are required to achieve it.

Delivery

In a telecommunications context, the delivery part of the value chain might be assumed to be the most familiar—and therefore the most trivial to implement. But developments from the computing world have had an increasingly significant role to play in the telecommunications space as the Internet protocol (IP) has become ever more ubiquitous. Making the delivery process 'smart' is a prime example of the second principle, where communications and content are linked together.

Traditional media delivery typically involves moving either discrete blocks or continuous streams of content information from one physical location to another. A smart alternative approach would involve several additions:

- 'tokenised' delivery of 'hard' media,
- abstracted vectors used to control stored content components, and
- distributed content hosting.

Taken together with the file and stream delivery, the result can be described as *smart bandwidth*, which Skinner[22] has described as 'the desire to "pack" the bandwidth as effectively as possible, to create the illusion that more information than the bandwidth can carry is being transmitted.'

By *'tokenised' delivery of 'hard' media*, we refer to the common practice of purchasing a licence to personally play a piece of physical media which is normally accompanied by an instance of the media in a readily-playable form. Compact discs (CDs), video home system (VHS) videotapes and digital versatile discs (DVDs) are examples of this type of packaging. Books, magazines, smart cards, and hard disk downloads are also examples of this type of delivery, where a physical object (or in the case of a hard disk download, the actual data bits occupied by the information) actually provides a tangible 'token' for the media. This provides something which can be collected and displayed as well as used to reproduce the media itself. It also provides a single point for the licensing information—for example, a personally-purchased CD normally cannot be played in public unless an additional fee for this use is paid. One interesting comparison is to calculate how long it would take to download a DVD using a low-bandwidth Internet connection: at 128 kbit/s, it would take around a week to download a typical 10 billion byte DVD movie file. The DVD is thus a very-low-cost, fast-access, robust and potentially long-lived alternative to an always-on Internet connection with a latency of days rather than minutes.

Abstracted vectors used to control stored content components refers to utilising storage space in a delivery device as part of the pool of media resources which are then used to generate the final media delivery (see later section on 'Presentation' for the generation part of this process). If the CD, DVD, smart card or hard disk download consists of a number of smart media components, then the 'default' usage might well be to play a number of these components in a specific order. The information required for this fundamental 'built-in' behaviour would be a simple list of the components in order. But by providing additional lists of components, perhaps with additional information about locations inside the components, then the content can be delivered in very different orders. These additional lists, edit decision lists (EDLs) in video parlance, can be supplied either as part of the stored media, or delivered over the network. In addition, if components are not present in the storage device, then they can be delivered across the network connection: via streaming for real-time requirements, or file transfer for non real-time usage. One interpretation of these vectors might be that they are analogous to links or user profiles, but for the internal structure of the media. In this way, the user profiles and vectors can be abstracted into what we term *associative metadata*, and by using these as pointers instead of using names for metadata fields, many of the problems associated with determining explicit naming can be avoided[23, 24, 25]. By applying the concept of abstracted associative metadata to componentised content, a new set of 'services' become possible. These include: sharing vectors between users, so that a user's profile could be made available to other people; a programme where you are aware of which of your friends are watching, and you can then jointly interact, maybe even as part of a virtual audience instead of a closed single person experience.

Comparing this smart component delivery with more conventional re-purposing reveals that whereas re-purposing applies personalisation, customisation and localisation rules at the source of the content (typically an 'edge of network' server), smart delivery applies these rules both at the source and the final delivery client. In addition, smart delivery maintains full control over the final production of the finished media across the entire delivery

Figure 2 An 'open publishing' system

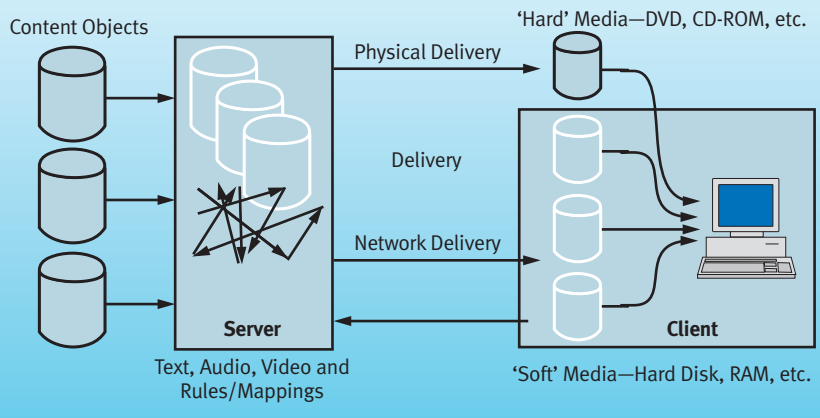

chain, something which normally stops at the point where re-purposing occurs. Control over the complete delivery process is likely to become increasingly important to content producers and publishers as the number of end delivery devices increases, and smart delivery enables this to become an intrinsic part of the delivery chain.

Distributed content hosting extends this concept of being able to 'fly in' components when required, by allowing individual end delivery clients to be more than merely storage devices for local use of the content components. By adopting a peer-to-peer approach, then components stored on one client could be used as a source of that component for other clients. Functionality built into the smartness of the content components deals with the licensing and royalty implications of this retransmission. Venation, a BT spin-out company[26], has already been launched to deliver advanced network capabilities intended to support exactly this type of componentised media delivery.

The end result of these additions is a delivery system:

- where the content need not be delivered as a single file or stream, from a single source;
- where the delivered content is a flexible, personalisable, customisable, localisable, 'unfrozen' set of components of media ready to be 'produced'; and
- where a flexible mix between components both from a physical 'token' and from the network can exist.

Presentation

Digital services today tend to be presented to the customer in a mode which is optimised to the specific application for which they are being used. This is not surprising given the wide variety of devices which mediate these services. Under these circumstances, service providers have found it difficult to offer their products over anything more than a very small range of platforms, each instance of which approximates to a 'frozen' published solution as described earlier. We propose that it is now possible to develop the appropriate software framework to permit digital content and services to be 'synthesised' on any platform which is viable for development. This framework must be highly flexible and sufficiently open to achieve adoption in diverse product environments. By way of analogy, where HTML and hypertext transfer protocol (HTTP) are considered to be the enabling framework which brought about the World Wide Web, we believe that their equivalent across the digital services space could develop into a new media form[8].

Devices

There are presently two specific trends apparent across the range of devices used to deliver digital content and services.

The first is a massive increase in the availability of persistent local data storage. This is most evident close to the TV industry with the development of personal digital recorders (PDRs) such as TiVo, and thus the next generation of home set-top boxes. However, we must not ignore the low cost of high-performance off-the-shelf PCs, and the ubiquity of storage in personal devices—nowadays often enough to hold several audio CDs encoded as MP3.

The second is a commensurate increase in data processing capability. High-performance multimedia processing is now available at very low cost, and the many opportunities this offers have yet to be harnessed. Smart realisation favours this 'fat client' technology because it enables control of the end-to-end process, and the 'synthesis' of personalised, customised and localised services at the point of delivery.

Key influencers of both of these trends are a range of 'hybrid' devices whose apparent capabilities place them in potentially more than one marketplace. For example, Sony's PlayStation 2, the ZapMedia Terminal, iDVD, and many more.

Synthesis

The final stage of presentation is the 'production' process. As we have stated, this can now happen much closer to the final end delivery point: typically in the 'fat client', although for some applications, for example, mobile telephony, then an 'edge-of-network' host server/gateway might be used.

The process of turning a collection of components containing content and functionality into a piece of finished media is achieved using what we call a *content synthesiser*. This is essentially nothing more than a way of using production values to present the content. A simple example of a content synthesiser is a World-Wide Web browser—which takes text, pictures, and formatting information and presents it onto the screen as a web page. Some of the text or pictures might already be stored in the browser's local cache, while other fragments of content might be downloaded. Note that for this example the text pictures and formatting information will typically be bound together into a single HTML file, whereas a more generic solution would use either separated elements, or objects with separately accessible elements.

Content synthesis is thus a reversal of the first principle mentioned earlier: content decomposition. Although we have developed a comprehensive content decomposition scheme, we do not believe that a single solution is necessary. Instead, we prefer that an abstracted synthesis engine would provide capability to any scheme, and we have been working with a software vendor on exactly such a generic content synthesis engine.

Market Opportunities

The authors are pursuing a number of exploitation opportunities over the coming year, but these are only pathfinders for a wide range of potential uses for this approach to delivering media solutions.

Some possible end-applications of the technology described here include:

- Providing a flexible solution to bridging the transition from existing server-based networks towards peer-to-peer networks.
- Providing a number of application-area specific interfaces to an open publishing interface, to allow the easy publishing of content databases.
- Providing edge-of-network or host/gateway level access to smart content components for subsequent delivery to clients.
- Providing the mediation and component discovery tools which will assist the dissemination of components around a network, or between peers.

However one looks at the market place, a smart approach has broad applicability; for example:

- *Wireless* A smart realisation solution would encompass both fixed and wireless networks, and could be a key part of a unique selling point for third-generation (3G) mobile.
- *End-to-end IP networks and services* Smart realisation enables any network to be used as part of a total delivery solution. It thus incorporates broadband and narrowband into a coherent complete multimedia delivery system that does not disenfranchise users based on their lack of sufficient bandwidth connection capability.
- *Technology-driven Internet, multimedia or mobile services* Smart realisation brings together all of these areas into one abstracted media delivery solution.
- *Access to content* Smart realisation is designed to make content the key aspect of future media delivery. It removes barriers like operating system, device, platform, etc, and replaces them with an open object-based content synthesis system which delivers the right content, in the right place, at the right time, to any device or platform, without a requirement for re-purposing.

The market sizes for media industries are often over-emphasised by the very self-publicising nature of media itself. The video

games business is now larger than the movie business; for example, with annual video games sales worth nearly $10 billion in the United States in 2000. Radio and TV advertising in the US was worth about $20 billion and $40 billion in 2000 respectively. But the market for communications is about 10 times larger at around $1000 billion in the United States in 2000.

Conclusion

The first pointers towards a demand for componentised content are already appearing, and a smart response is required in order for telcos to be ready to deliver the sort of advanced component-based delivery solutions that we believe will be demanded in the very near future by our customers. BT has been putting in place the underlying framework for making smartness a key part of our future technology strategy for several years, and we are now beginning to exploit this innovative technology commercially.

References

1 DIMITROVA, N.; YU, H.; GALLIANO, F.; KOENEN, R.; ZAKHOR, A.; and BOUMAN, C. Entry into the Content Forest: The Role of Multimedia Portals. *IEEE Multimedia*, July–Sept. 2000, **7**(3), pp. 14–20.

2 CURTIS, K.; FOSTER, P. W.; and STENIFORD, F. Metadata—The Key to Content Management. *3rd IEEE Metadata Conference*, Bethesda, 6–7 April 1999.

3 NACK, F. All Contents Counts: The Future in Digital Media Computing is Meta. *IEEE Multimedia*, July–Sept. 2000, **7**(3), pp. 10–13.

4 RUSS, M. Desktop Conversations—The Future of Multimedia Conferencing. *BT Technology Journal*, Oct. 1997, **15**(4).

5 FINGAR, P. Component-Based Frameworks for E-commerce. *Comm ACM*, Oct. 2000, **43**(10), pp. 61–66.

6 SWIDERSKI, Z. Content Control in Multimedia Applications. *Proc CHI2001*, 31 March–5 April 2001, Seattle, Washington, USA.

7 http://www.clienthelpdesk.com/statistics_research/ web_statistics.html

8 SPOHRER, J.; and STEIN, M. User Experience in the Pervasive Computing Age. *IEEE Multimedia*, Jan.–Mar. 2000, **7**(1), pp. 12–17.

9 BOVE, V. M.; and AGAMANOLIS, S. Responsive Television. *Proc. IBC2000*, Sept. 2000.

10 HENLEY CENTRE, *Media Futures Report 1999*, Section 1: Media Time.

11 TORRIS, T.; and GRIMSDITCH, T. iDTV Needs New Content, Not Web Content. *Forrester Research Inc.*, 8 Feb. 2001.

12 PRICE, M.; and THOMAS, G. 3D Virtual Production and Delivery Using MPEG-4. *Proc. IBC2000*, Sept. 2000.

13 INTERNATIONAL PRESS TELECOMMUNICATIONS COUNCIL. http://www.iptc.org/NMLIntro.htm.

14 Beyond Star Wars. *Wired Magazine*, Feb. 1997.

15 http://www.ifilm.com/db/static_text/0,1699,15700,00.html

16 SKINNER, R. SR—Content, Purpose and Re-Purposing. BT Internal Report by Short Term Fellow from Suffolk College, unpublished.

17 VAILAYA, A.; FIGUEIREDO, M. A. T.; JAIN, A. K.; and ZHANG, H. J. Image Classification for Content-Based Indexing. *IEEE Trans. on Image Processing*, Jan. 2001, **10**(1).

18 XU, L.; ZHU, J.; and STENTIFORD, F. Video Summarisation and Semantic Editing Tools. *Proc SPIE*, Storage and Retrieval for Media Databases 2001, 24–26 Jan. 2001, San Jose, pp. 242–252.

19 STENTIFORD, F. W. M. An Estimator for Visual Attention Through Competitive Novelty with Application to Image Compression. *Picture Coding Symposium 2001*, 25–27 April 2001, Seoul, Korea.

20 WORLD-WIDE WEB CONSORTIUM. http://www.w3.org/Style/XSL/.

21 http://forum.afnor.fr/afnor/WORK/AFNOR/GPN2/ TC154WG1/index.htm.

22 SKINNER, R. Smart Realisation, BT and Motivation. BT Internal Report by Short Term Fellow from Suffolk College, unpublished.

23 http://www.cs.yorku.ca/~peter/Berry.html.

24 RUCKER, R. Infinity and the Mind: The Science and Philosophy of the Infinite.

25 http://www.greylabyrinth.com/Puzzles/answer009.htm.

26 http://www.venation.com.

Biographies

Martin Russ
BTexact Technologies

Martin Russ studied electronics at the University of Manchester Institute of Science & Technology (UMIST). He has worked in television, the music business, computer animation, graphics, and telecommunications research. He has an Open University B.A. degree in Computer Science and several Internet patents. He is currently investigating multimedia content management and delivery systems, as well as learning Japanese in his spare time.
martin.russ@bt.com.

Ian Kegel
BTexact Technologies

Ian Kegel completed a masters degree in Electrical and Information Sciences at the University of Cambridge. He has worked in both the defence and telecommunications industries on projects ranging from radar signal processing to multimedia delivery. His current research is centred on the design and demonstration of novel information management and delivery architectures. In recent years he has contributed to several Internet patents in this area.
ian.c.kegel@bt.com.

Fred Stentiford
BTexact Technologies

Dr Fred Stentiford studied mathematics at St Catharine's College, Cambridge and obtained a Ph.D. in computer aided design at Southampton University. He has worked in a range of research areas including OCR, speech recognition, and machine translation. He now leads a group at BT researching image processing and new multimedia content delivery mechanisms.
fred.stentiford@bt.com.