

Remote Non-Linear Video Editing

By Robin S. Rowe

Remote non-linear video editing (RNLE) is a natural progression to the next generation of video editing technology. A paradigm shift, editors are no longer chained to the edit suite. RNLE operates a medium or low-bitrate mirror copy of broadcast-quality video clips. With this approach a low-resolution rough cut is created on a PC in a simple cut-and-paste interface. Sometime later, the resulting edit decision list (EDL) is transferred back to the edit suite for the finish cut and committed at high resolution suitable for broadcast as film or HDTV.

Remote non-linear video editing (RNLE) has important implications for the efficiency of television newsrooms and production staffs. Journalists can edit at their desktops freeing up time in the edit suites. Mobile RNLE detaches the RNLE from the network by replicating low-bit rate video clips from the network server onto a laptop.

A director can conveniently create a cut on a laptop while on an airplane en route to a shooting location. The lightweight EDL can be easily e-mailed to an editor in Los Angeles or New York City who has the same footage in high resolution. This enables the director to work from a rough-cut rather than uncut dailies. The director at the site can verify that there is adequate shot coverage and judge the pacing of the show.

The Leitch BrowseCutter is an RNLE implemented in Java via a web browser interface. At NAB'99 an untrained volunteer from the audience operated the BrowseCutter software (running on a laptop) to demonstrate the ease of creating an MPEG-1 rough-cut from raw footage.

Presented at NAB'99 (paper no. 1), in Las Vegas, NV, April 17-22, 1999. Robin S. Rowe is with MovieEditor.com, San Diego, CA 92101. Copyright © 2000 by SMPTE.

BrowseCutter, VR-300, and NewsFlash are trademarks of Leitch Technology Corp. NetShow, Internet Explorer, PowerPoint, and Windows are trademarks of Microsoft Corp. Netscape is a trademark of Netscape Communications Corp. RealVideo is a trademark of RealNetworks, Inc. UNIX is a registered trademark of The Open Group. Panavision is a trademark of Panavision, Inc.

Client/Server Video Editing

Everyone knows that broadcast-quality digital video files are too large to handle over a corporate LAN running at 10 or 100 Mbits/sec Ethernet. That's why television stations use Fibre Channel running at 1 Gbit/sec to connect digital broadcast servers (such as the Leitch VR-300 server and NewsFlash editing system). Video systems with a client/server architecture are a tremendous advantage in the creation of television news programming and play-to-air, because a single high-quality digital video file residing on the server can be shared among multiple users. In a true client/server design, there is no need to wait while a video file copies from an editor to a play-to-air server: it is the same file.

Unfortunately, the expense of high-bandwidth pipes and video mainframes discourages widespread use of client/server video editing. It is simply too costly for most television stations to provide an editing station for every journalist. The result is a traffic jam in the edit suites, as too many users want access to too few machines.

Replication Servers

Ideally, every journalist should have an inexpensive client-based video editing station running on a desktop PC, with video files available from a server using the existing corporate LAN. An adaptation of an advanced technology developed for commercial database servers actually makes this possible. For many years, a technology has existed in database servers called the

replication server, which keeps a copy of the data on a server. If that server fails or loses connectivity to remote users, it would still be possible to continue business as usual by connecting to the alternate database server.

The design of such a system presents a challenge, because when the connection between servers is restored, the data must be automatically resynchronized. A typical example would be a company that has a server on each coast handling sales orders. The servers must track each other so that more items aren't sold than exist in inventory, even if the two servers temporarily lose contact with each other.

The relevancy of replication servers to video mainframes is that most video users don't require the replica to be a perfect copy, except when playing-to-air. On a broadcast server, video is typically encoded at bit rates from 24 to 48 Mbits/sec; on a web server the same video content is encoded in MPEG-1 at 1.1 Mbits/sec (or less). The picture quality of the MPEG-1 video is roughly comparable to VHS SLP-mode—not good enough for broadcast, but quite adequate for shot logging and rough-cuts.

Consider that MPEG-1 can be effectively transmitted over conventional corporate LANs and played back in Microsoft Windows and things become very interesting. All that is required is to keep the broadcast server and the low bitrate replication server in sync; web server recorders do this automatically. Content changes made to the broadcast server will be mirrored in the databases and MPEG-1 files on the web server. Users access the high bitrate version for play-to-air or the low bitrate version for distribution to the desktop.

Java Applets

The trend in desktop interface design is to use Java applets, which have the advantage of running on

Table 1 — Video CODEC Typical Bitrates and File Sizes

CODEC	Selected Bitrate	Typical File Size	% of Total
MPEG full-resolution video	32.00 Mbits/sec	15.00 Gbytes/hr	96.07
MPEG-1 browse video	1.15 Mbits/sec	610.00 Mbytes/hr	03.91
JPEG filmstrip	8.00 Kbits/sec	4.00 Mbytes/hr	0.03
Totals:	14.92 Gbytes/hr	15.61 Gbytes/hr	100.00

almost any platform. Using Microsoft Internet Explorer or Netscape, enables the user to access the replication server without installing additional software on the client machine. Java applets are programs that reside on the server but are executed on the client in the browser. This greatly simplifies things for network systems administrators, because they never have to install new versions of software to the users' PCs. Upgrades are made at the server and everyone is inherently running the same version of the software.

Java has other advantages. Unlike most computer languages, the same software will run on multiple platforms: Windows machines, Macintoshes, and Unix-based workstations. The network administrator installs just one version to support all these platforms. All platforms run the same program and users avoid the headache so prevalent in the Mac/Windows world of one platform supporting a newer version than the other.

File Sizes

Even though MPEG-1 files are relatively smaller than other video files, they are still huge files to routinely move across a corporate LAN. An MPEG-1 file is typically compressed to about 10 Mbytes/min. A 3-min spot is therefore 30 Mbytes in size. In observing how users manipulate video files for editing, it is evident that much of their time is spent scanning for the appropriate footage and marking in/out points. Using a filmstrip representation of the video enables these tasks to be done without transmitting the MPEG-1 files.

In effect, yet another replication database is kept of the video, but this time the JPEG thumbnails sample at one frame every two sec. Each JPEG

still image is about 2K in size. A 3-min spot is therefore roughly 200K in size as a filmstrip, 30 Mbytes as an MPEG-1, and 550 Mbytes as the original high bit rate digital video on the broadcast server.

Keeping three copies of a video obviously takes more storage, but the relative sizes of the copies at 5.5% and 0.04%, respectively, makes that insignificant. In fact, holding the copies after the originals have gone to tape archive is practical.

Architecture

Figure 1 is a diagram of the layout of components. Each browse recorder is connected to a broadcast video server. This browse recorder can march in lockstep with the broadcast server, recording the same clips as they are being ingested. If the web server and broadcast server lose synchronization, the browse recorder can take machine control of the broadcast server to play out the missing segments or edited videos back into the web server.

Browse recorders place the MPEG-1 data and the JPEG stills onto the web server. A database is also created there to provide the data the Java applets display at the client desktops.

A video cataloger is conceptually the reverse of a browse editor in that it takes completed television content and disassembles it into individual stories. For a network news broadcast this enables news analysts to query broadcast television content as conveniently as they would a web site. For example, the user can instantly retrieve a list of the day's broadcast news stories containing the word "president."

HDTV Implications

This replication strategy has important implications for HDTV. Replication databases can mirror video

encoded at any resolution, even film. A browse editor can easily support letterboxing at 16:9 and Panavision aspect ratios. And, NTSC and PAL are both converted into MPEG-1 enabling the same browse-system to support many video standards interchangeably. No matter what the source format, the result is the same international standard MPEG-1 video files that can be played back on any PC.

With HDTV, moving video around the television station is very costly. Once again it requires Fibre Channel connectivity or something similar to have the adequate bandwidth to move such massive files. The expense of routing HDTV video everywhere, in the manner in which stations have become accustomed with analog television, has discouraged some stations from embracing HDTV in the near term. However, with browse editors, the required amount of high-bitrate routing is decreased. MPEG-1 replication copies can travel across existing corporate LAN while keeping the HDTV video mainly in the central rack room. As edits are applied to the MPEG-1 clips, the edit decision list (EDL) is propagated back to the original high-resolution video. The EDL is simply a list of the names of the video clips being assembled with their in/out points and related data: in other words, it's just a text file.

Using an EDL to synchronize edits across video replication servers means that users may detach from the network and carry the browse editor with them on their laptops. This same approach is used in conventional replication databases so that traveling sales reps can take orders without being continually connected to the server. Each day the sales rep dials up the server and reserves the stock he expects to sell, keeping those items in a subset replicated database on his laptop. After taking orders during the day on the laptop, its replication database is relinked temporarily with the main server and the two systems resynchronize. In the same way, an editor can copy the files he expects to edit to his laptop. Later, the extremely lightweight EDL can be e-mailed back to the server from the detached laptop (even using a two-way pager) and applied against the original footage there.

REMOTE NON-LINEAR VIDEO EDITING

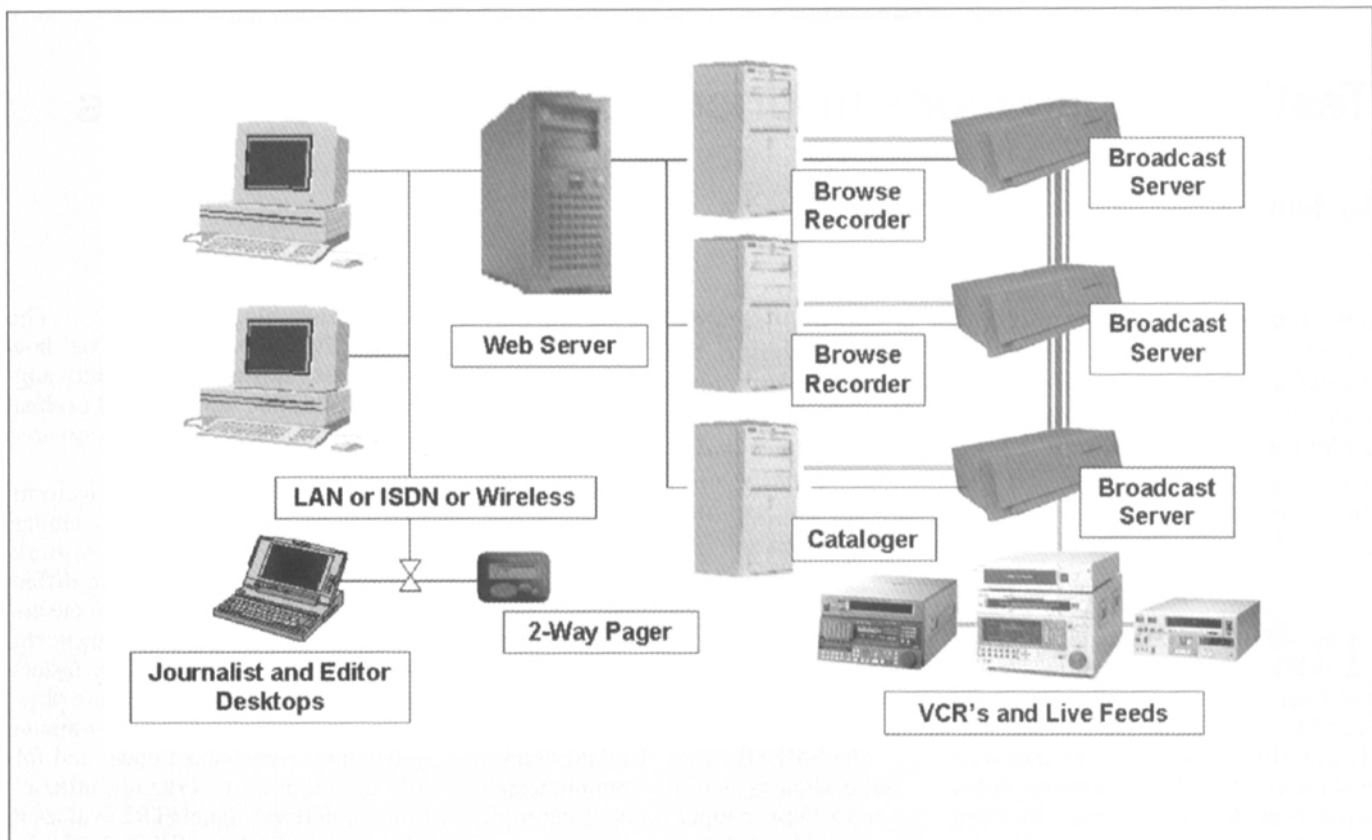


Figure 1. Web-based architecture.

Remote Clients

Handling video content efficiently makes it practical to have the browse-editor user connected by ISDN or cable-modem to a distant browse-server. A television station with a news bureau in another country could access its browse database remotely to extract only the required video clips. Even though the resultant high bitrate video edit would be rather large, it could be cheaper to send electronically by ftp than the alternative of putting a tape in a bag and having a courier deliver it.

By trimming the video duration down to the exact footage needed, it becomes practical to browse content over conventional communications infrastructures. For instance, a television station in Manhattan could connect to its archive in New Jersey and choose just the right clips to retrieve. Even if a courier (sneaker-net) approach is more economical for transporting video in a particular situation, having a lightweight remote browse capability into the archive is a tremendous improvement in selecting footage.

Web Sites

Television stations are embracing web sites as a means of promoting their brand and as a secondary source of income. The act of converting content into a browser-friendly format can help repurpose video content for the Internet.

If placed on an external web server the JPEG filmstrips can be accessed easily over conventional modems. However, MPEG-1 (at CD-ROM quality) is too heavy for such low bitrate connections. An alternative is to reduce the picture quality (and size) by tweaking the encoder settings or use a secondary encoder that converts the video into a format tailored for low bitrate Internet carriage (such as Microsoft NetShow or RealVideo). In fact, the 24-hour news channel Time-Warner New York One is using its browse edit system in this way to export in RealVideo format to its www.nyl.com website.

Conclusion

As this new web-based video editing approach becomes widely avail-

able, it will change the economics of television production. New uses and applications will become possible. For example, it will be simple to burn promotional reels and other video content as MPEG-1 on CD-ROM or MPEG-2 on DVD rather than tape. Browsing and editing video with the convenience of web-based systems will help make video presentations as simple to prepare and as ubiquitous as PowerPoint presentations are today.

THE AUTHOR

Robin Rowe is the CEO/COO of MovieEditor.com. During his tenure as a chief technologist at SAIC, he designed BrowseCutter.

Rowe's broadcast experience includes technical director of broadcast news at WICD-TV, an NBC affiliate in central Illinois. At WMAQ-TV in Chicago, he helped build the robotic studios.
