Frame Rate Follies

by Chris & Trish Meyer, CyberMotion

They may be a nuisance, but you need to get them right.

Frame rates are one of the most confusing aspects of dealing with video - particularly if you have to deal with NTSC-related formats. It may be tempting to round off numbers to something that's easier to say or type, or to assume your camera, decks, and software are already handling any sticky issues for you, but this can lead to audio/video synchronization errors as well as duplicated or skipped frames. In this article, we're going to discuss where some of those whacky frame rates came from, and what you can do to protect yourself from errors happening.

Where 29.97 Came From

In the beginning, most film ran at 24 fps (frames per second), while interlaced video ran at a field rate based on the alternating line current frequency in a given country. In Europe, with 50 Hz line current, this meant a field rate of 50 fps giving a frame rate of 25 fps (remembering that there are two interlaced fields per frame). In the US, which has a 60 Hz line current, this meant a field rate of 60 fps and a frame rate of 30 fps.

Then along came color television. It was mandated in the US that color television had to be compatible with black and white television, so that consumers with older TVs could still see something useful when viewing a color broadcast. The solution to this was to continue to broadcast the "luminance" (grayscale, or black and white) information as before, and to tack on additional information that defined the color. This is where the common video "YUV" format came from: Y stands for luminance, while U and V are codings of the color information.

However, there was a concern - untested, we have read - that the additional bandwidth required to send the color information would cause problems with television channels interfering with each, as there may no longer be enough space in between them as they were transmitted over the airwaves. Therefore, to be safe, it was decided to slightly slow down the NTSC frame rate by a ratio of 1001/1000 to approximately 29.97 fps. It was felt that this was close enough to 30 that older receivers should still be able to lock onto the signal, while giving a safety margin between broadcast channels. A precise ratio of 1001/1000 gave a "gear ratio" that was clearly defined and therefore easier to build into hardware.

(By the way, it was later discovered that this additional safety margin wasn't needed. Too bad they weren't able to adequately test for it at the time; we could have avoided this entire mess…)

The Film Case

When transferring 24 fps film to video for broadcast and storage on videotape, two problems had to be overcome: The difference between 24 and 30 fps, and the difference between 30 and 29.97 fps. To solve the first problem, various sequences have been used to spread four frames of film across five frames of video (as 24/30 = 4/5). Initially, every fourth frame of film was repeated to create the fifth frame of video. Since then, various schemes have been employed to spread four frame of film across 10 fields of video, such as the common "3:2 pulldown" system of repeating the first frame of film across three video fields, then repeating the next frame of film across two video fields, continuing this process as long as required.

To solve the 30 versus 29.97 issue, one other step was required: The film had to be slowed down by the same 1001/1000 ratio to approximately 23.976 fps. Since film soundtracks were originally printed optically right alongside the image frames on the same strip of film, slowing down the film reel slowed down the audio by exactly the same amount, so image and sound stayed in sync.

The Digital Age

When desktop video was being born back in the late 80s and early 90s, a lot of the smart computer engineers involved did not have any background in traditional broadcast video. They looked at 29.97, thought it was a pretty inconvenient number, and instead used 30 fps. Indeed, many of the earliest video capture cards tagged the resulting movies as having a frame rate of 30, not 29.97, and many of the earliest stock footage video libraries - including those from Artbeats - also used 30 fps. This error was caught fairly early on, and video hardware and software quickly changed to tagging NTSC content as having a frame rate of 29.97 fps. You will also find that all Artbeats NTSC format stock footage also now uses the rate of 29.97 fps.

Where Things Go Wrong

The difference between 30 and 29.97 is pretty tiny, right? Who's going to notice a 0.1% speed change? So all of this information is just for geeks, yes?

Unfortunately, no. That tiny difference adds up to a video field every 16.7 seconds, and an entire frame every 33.3 seconds. That means if you ignore the differences between 29.97 and 30, you may need to either repeat or skip an image field or frame that often! Note that you have no control over when that skip or repeat happens - it might be the very first frame of the resulting animation. Also, even the most tone-deaf person can tell when video falls two frames out of sync with its soundtrack (many can tell before then), and that magic number is reached in just over a minute.

Alas, there are many ways that incorrect frame rate video can still creep into our workflows. For example:

- Many 3D animation programs (not to mention their users)

employ 30 fps, not 29.97 fps, as the NTSC frame rate.

• Some applications derive the frame rate of a movie by looking at how many frames long it is, and dividing it by its duration in seconds. However, some video systems can introduce a small error in the duration of just the first or last frame during a video capture, resulting in the frame rate of the entire movie now appearing to be slightly off. This is also a problem with screen capture software: Many of our SnapZ captures - which are supposed to be 12 fps - come into After Effects labeled as 12.06 fps, causing audio sync problems down the line.

• If you receive a sequence of film frames from a film scanner, you might assume you should assign it a frame rate of 24 fps. However, if you have been working from a video dub of the film to make your initial edit, the actual frame rate is 23.976 fps. This difference can wreak havoc when you have to synchronize an edit to, say, a soundtrack arriving on CD or via a digital file - did the music editor assume 24 fps, or 23.976?

• Thanks to the insanity that is the high-definition video spec, 60, 59.94, 30, and 29.97 fps are all legal video rates; 24 and 23.976 are common production rates. And many cameras have switches to change the capture frame rate between these numbers. Just to make life interesting, video playback decks aren't always aware of the frame rate video was captured at. So it is possible that footage was captured at the "wrong" rate without you knowing.

Protect Yourself At All Times

In light of these potential frame rate problems, here are a series of defensive practices you can employ to ensure you don't fall victim to skipped or repeated frames, or audio falling out of sync. Many involve the important step of "conforming" the frame rate of visual assets you import into your NLE or compositing program. This in essence means overriding the file's frame rate and assigning it a "new" rate, causing every frame to be labeled as precisely having your specified duration.

• If you receive a 3D render or stock footage clip, conform its frame rate to match that of your project - for example, conform 30 fps 3D renders to 29.97 fps for an NTSC video job.

• If you receive a video clip that appears to be labeled with the correct rate (such as "29.97" fps), conform it anyway to get around the strange math errors mentioned earlier. Rogue frame rates are one of the biggest causes of mysterious skipped or repeated frames in After Effects, for example.

• Label all of your tapes with the frame rate used when they were captured. If someone hands you a tape, have them track it back upstream to the camera to find out what frame rate was used.

• If you alter the frame rate of a video asset, make sure any audio that is supposed to be attached gets changed by the same amount. For example, if the 3D animator in the first example was keyframing their movements at 30 fps against a soundtrack, conforming their render will now cause some of the audio hit points to drift out of time - so slow down the audio by 0.1% as well to re-synchronize them.

• If you receive a soundtrack from someone working on a film rate project, ask them if they were assuming a frame rate of 24 or 23.976 fps, and alter the soundtrack's rate accordingly if you are working at a different rate.

• Use 23.976 fps for "film-like" rates, not 23.98 fps. Some software round 23.976 to two digits after the decimal point for the sake of display, even though they use the longer number internally; don't let this display error lull you into using the wrong frame rate in practice.

Zooming Out

There is one more potential problem that we may have to deal with in the future. To this day, most software uses the number 29.97 - as in 2997/100 - as the video frame rate. Unfortunately, the actual frame rate of NTSC color video is 30,000/1001.

The day will come when software starts to move to the correct 30,000/1001 number. Indeed, Apple has been recommending it since 1999 (see their document at developer.apple.com/quicktime/icefloe/dispatch019.html), although few have actually followed it as of the time this article was written. Beware when you start receiving content tagged in this fashion! Just to test the potential problems, years ago we hacked a test movie to have this frame rate and fed it into programs such as earlier versions of After Effects. The result was the very first frame being repeated. In the short term, the conforming trick mentioned above will cure it; in the long term, our new-fangled computers will eventually work just like the electronics of the 50s...

###

Chris and Trish Meyer are the founders and owners of CyberMotion (www.cybmotion.com), an award-winning motion graphics studio in Los Angeles that has created a wide variety of work for film, broadcast, corporate events, and special venues. They were one of the original development sites for After Effects, wrote the highly-acclaimed books Creating Motion Graphics and After Effects Apprentice (2007), and are long-time Artbeats users.