Getting started with Quicktime VR

By Andrew Nemeth - Andrew Nemeth must be the only mac developer who does not own a cat. He lives in Warrimoo Australia and can be reached at <azn@nemeng.com>. His WWW page is located at <www.nemeng.com>.

Introduction

Apple's Quicktime VR technology is possibly one of the coolest things to happen to digital imaging in the last five years. With it you can produce and view 360 degree panoramas of real or computer generated scenes, without the need for expensive panoramic cameras or any other of the fancy equipment usually associated with VR.

You create a panorama by photographing a scene with a standard 35mm camera atop a tripod, taking 12-18 photographs to capture a full 360 degree view. Then using Apple's software, you 'stitch' these images into a single PICT file, which is then processed further to create the final, user-navigable movie which is playable on both Macintosh and WINTEL personal computers.

As there are already a number of general introductory articles on Quicktime VR (see the 'See Also' list further down), the following will concentrate instead on the tools and techniques used to create a photographic panorama. Multi-node panoramas (where the user can jump from image to image by clicking on pre-defined hot-spots) will be covered at a later date.

Assembling the equipment to photograph a panorama

To keep costs down and for the sake of portability, use of the following is suggested:

- 35mm SLR camera
- 24mm f2.8 lens
- (Lots of) 400 ISO low-contrast C41 negative film
- Sturdy tripod with a built-in spirit levels
- Kaidan QP-1A QuickPan Panoramic Base
- Vertical camera mounting bracket

(Note that some of the above is different to that recommended by Apple in their Quicktime VR 1.0 Authoring Tools. The reasons for these differences will be noted below.)

35mm Camera
I use two cameras for my work: a 25 year old Nikon F (pictured in Figure 4) and a more recent Nikon F90x. These are used depending on the locality, with the F90x for general use whilst the older - indestructible - F reserved for more rugged activity.

Although not absolutely necessary, a 35mm SLR camera gives you the option of being able to use different focal length lenses on the same camera body, meaning you can use wider lenses indoors and longer ones for outside work. They also tend to be better

built and feature more accurate focusing and metering. A 35mm camera also gives you a much higher quality image than the current crop of digital cameras.

With regards light meters, I prefer to use a hand-held model. Again this is not essential but I find it to be more convenient than the built-in meters found in cameras (especially the Nikon F, which does not have *any* kind of meter!).

24mm lens
As the first of many departures from Apple's suggestions, a relatively inexpensive Sigma f2.8 24mm lens is favoured for outdoors work rather than the Apple-sanctioned Nikkor f3.5 15mm. Although the wider 'Apple' lens does have greater vertical and horizontal coverage (allowing for the reduction in the number of individual photographs required capture a scene), there are two major problems with it; firstly, at list price of $2000 (US) the 15mm is *outrageously* expensive. Secondly, its field of view is *too* wide for outdoors work, with everything more than 3 meters away appearing to look as though it was on the edge of the observable horizon.

Although other wide-angle lenses could be used (18, 20, 28, 35mm), the 24mm is a good compromise on a cost/image-width basis. The down-side is that because Apple has assumed everyone will be doing Quicktime VR with a 15mm lens, adjustments have to be made to the MPW scripts which come with the Authoring tools in order to render the images properly (more on this further down).

400 ISO Low contrast film
A faster ISO film gives you more scope to capture scenes whatever the prevailing light. Although a slower film (50 or 100 ISO) is sharper and has less grain, for Quicktime VR such advantages are lost during the software processing. Furthermore, slower film is *too* slow to use indoors.

A low contrast C41 film is recommended because it is the best kind of colour film for capturing all the levels of brightness in a scene. A film like Fuji's NPH 400 is sharp, fine-grained and - as it is designed for professional flash-lit photography - more than capable of dealing with harsh outdoor light.

By the way, be warned that Quicktime VR literally *chews* through film - at 16 images per panorama you can fit only two scenes onto a roll.

Tripod and head
The sturdier the better. The one pictured (a Manfrotto, see Figure 1) is a 5kg behemoth which features two built-in spirit levels to allow you to quickly level the head prior to attatching the Quicktime VR related brackets. In case you think you can get away with using a lighter tripod, think again. You really do need a heavy duty tripod for this kind of work because you don't want it flapping around in the breeze or moving suddenly in the middle of a sequence of shots. (Okay I relent: for people with bad backs and deep pockets, the Gitzo model 1228 tripod is a good carbon fibre substitute.)

You also need the spirit levels because the Quicktime VR software insists on the camera being absolutely level when photographing a scene.

Fig 1 Tripod head with double spirit levels

Camera mounting brackets
For the Authoring Tools to work, each 360 degree scene must be photographed by taking a series of overlapping shots using a vertically oriented camera. To avoid parallax errors when doing this, you must position the optical centre of the lens directly over the axis of rotation (see Figure 4).

For an example of parallax error due to off-centre rotation, hold one finger 10 cm from your face, close one eye, focus onto the background and then turn your head slowly from side to side - notice how the background appears to shift from side to side behind your hand? Now keeping your head and hand still, turn your eyes from left to right - this time the background *doesn't* move relative to your finger! In the first case your eyes are not centred on the axis of rotation and hence the parallax error when you turn your head. In the second example you are turning your eyes, the centres of which *do* coincide with the turning axis!

The set-up I use is shown in Figures 2, 3 and 4. A quick-mount plate has been added onto the Kaidan QP-1A Base (Figure 2 RHS) to make it easier to attatch to the tripod head. Black gaff-tape has also been stuck over the 'deck' of the plate as its highly reflective silver finish caused lens-flare whenever the camera was rotated over it.

I prefer to use the QP-1A rather than Apple's scheme of mounting a second tripod head because it is smaller, lighter and you can set it to 'click-stop' at the number of shots you wish to take in a pan. (For the 24mm lens it has been set for 16 detents. It can also be set for 8, 12, 14, or 18).

The bracket used to mount the camera onto the QP-1A is something I made myself from quarter-inch steel brackets, bolts and Araldite (Figure 3). This is a more durable and *much* cheaper solution than Kaidan's QPU-1 ($180 US) bracket. However, because it is not adjustable (unlike the QPU-1), separate brackets have had to be made for each camera used.

The entire rig is compact, strong and very easy to assemble (Figure 4). Because the vertical 'home-brew' bracket has already been pre-aligned (when it was made), the only calibration

required when setting up on-site is the leveling of the tripod head.

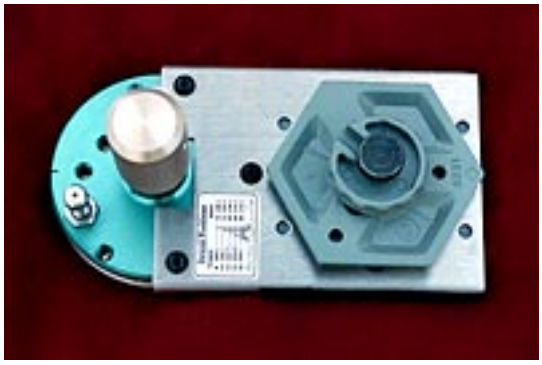

Fig 2 Kaidan QP-1A Base with Manfrotto quickmount plate

Fig 3 Kaidan QP-1A base and (custom) 35mm bracket

Fig 4 Full Quicktime VR rig with camera and 24mm lens

Photographing the scene

With the rig set-up and the camera level, you will need to take a series of overlapping, vertically-oriented photographs to cover a scene. The amount of overlap is in the order of 30-50% and is needed by the stitching tool to help it align adjacent images when it pieces together the final panorama. The more overlap the better, but too much overlap equals wasted film and long stitching times.

Figuring out the amount of overlap (and thus how many photographs to take per panorama) is very much an educated guess, with 16 shots for a 24mm lens appearing to be adequate. Some people shoot 18, some live (very) dangerously and try to scrape by with only 12.

There are a number of things to keep in mind when photographing a scene:

- Your camera rig must be (and remain) asolutely dead-level. The more out of level you are, the harder it will be for the stitcher to piece the images together.
- Photograph in a left-to-right direction. The stitching tool assumes this and you will get very wierd results otherwise.
- Keep the exposure identical for each shot. This means you will have to meter the entire scene carefully prior to shooting and then take all your exposures using an averaged value. If you alter the exposure for each individual shot then you will get a banding effect when you come to peice all the images together.
- Use a small lens aperture to make sure everything is in focus. A value of f16-22 gives you coverage from 1.5m to infinity.
- Take your shots as fast as you can without knocking your camera out of alignment. You have to be fast because the light and clouds and cars and people have a habit of not staying still.

With 16 exposures per panorama, you can fit two scenes comfortably onto a 36-exposure roll of film, with a few shots left over as spares.

Developing and digitising the photographs

Apple advises you to develop/scan your images with Kodak's PhotoCD process - which may be great for people who live in the U.S. but is Less Than Optimal for everybody else.

In Australia for example, the turn-around time for PhotoCD is in the order of 2++ weeks, with the cost averaging $80 (AUS) per roll of 36-exposure film. Thus if you are intending to do lots of Quicktime VR, then you may be better off investing in a 35mm film scanner and digitising the images yourself!

This is what I and a few others have done. After development of the negatives, they are scanned with a Polaroid SprintScan 35 via Photoshop (16 images require 25 minutes). When the scans are no longer needed, they are archived with Stuffit and 100 MB Zip discs.

To justify the purchase of a scanner you really must be looking at doing more than just a handful of QuicktimeVR panoramas, otherwise PhotoCD is the way to go. A tip: try to use the 'Portfolio CD' service as this will allow you to cram 500 768 x 512 resolution images onto a single CD.

Assembling the required computer tools

Quicktime VR requires the following:

- Fast PowerPC/ 48MB physical RAM/ 100MB free disc space
- Colour monitor/ millions of colours
- MPW (native 3.4 and beyond)
- Adobe Photoshop 3.0.x
- HyperCard 2.3

As you can see, these are pretty stiff requirements!

Fast PowerPC (etc)
Although Apple suggests you can do Quicktime VR development using a 68k based computer (68030, 68040), in reality the tools run so slowly that you cannot. For example it takes 4 minutes to stitch 16 images on a PowerPC 7500, on a Quadra 800 it requires slightly over an hour. Dicing is similarly glacial on a 68k machine: 25 minutes as opposed to 3 on a PowerPC.

You need a minimum 48MB of physical RAM because you will have to allocate at least 40MB of it MPW. Yes, MPW...

Native MPW
At a time when Macintosh developers are abandoning it, for Apple to release a multi-media development suite dependent on MPW is perverse to say the least - yet this is exactly what they have done with the Quicktime VR 1.0 Authoring Suite.

In order to piece together the images you use the 'stitch' MPW tool. To make a Quicktime VR movie, you first dice (compress) the stitched image with the 'p2mv' tool and then make the movie with the 'msnm' tool. You drive these and other tools in the time-discredited MPW fashion of executing lines of script in a worksheet. Joy.

APDA offers specials on bundles of Quicktime VR kits with MPW Pro - ignore them. Since you only need to run a small sub-set of MPW (still a bloated 8MB), save your money and use the MPW which comes with the CodeWarrior (or Symantec?) tools.

Adobe Photoshop 3.0
This is *essential* for Quicktime VR development. After stitching you will have to retouch the image, balance the colour, adjust the gama and a host of other things. It is surely no coincidence that the stitch tool outputs its resultant panorama with a Photoshop creator and file type!

HyperCard 2.3
(You will only need this if you intend to produce mult-node movies. This will not be dealt with in this article).

An overview of the Quicktime VR software process

Figure 5 outlines the procedure required to convert your sequence of scanned images into a single Quicktime VR movie. As you can see, extensive use is made of MPW.

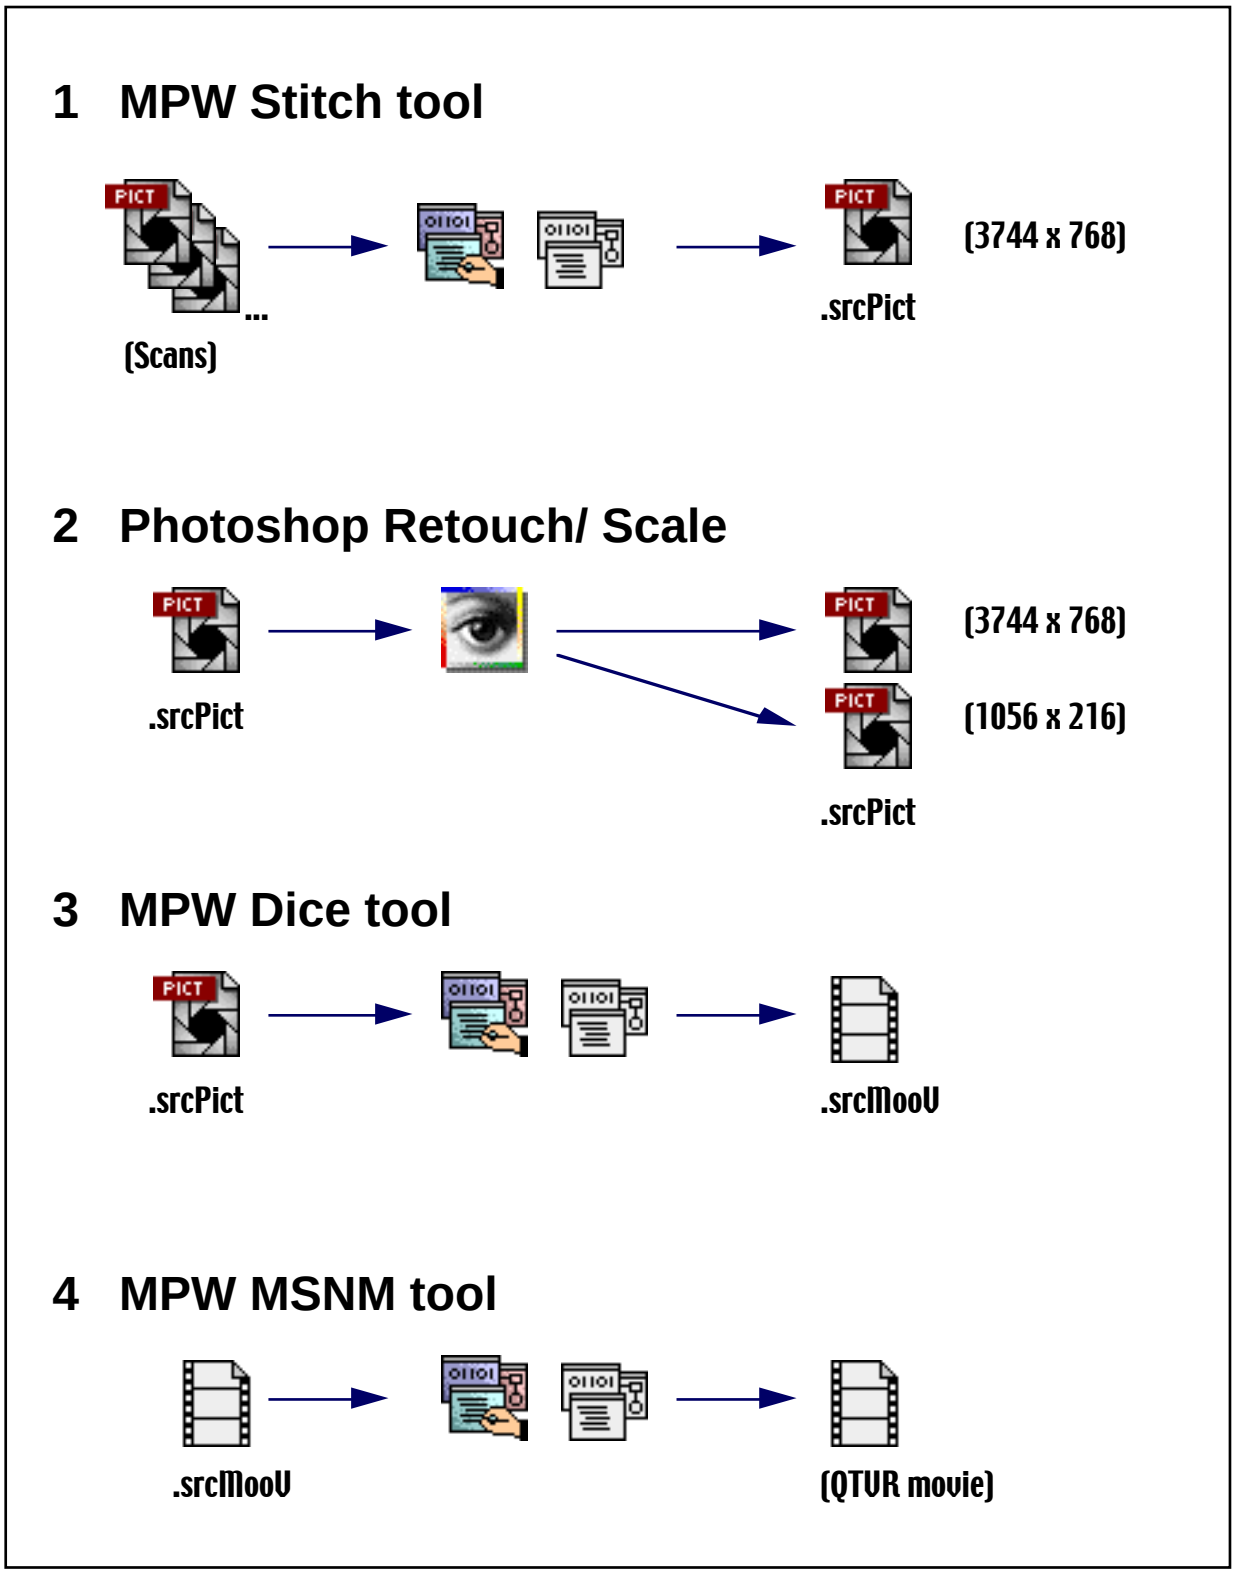

Fig 5 The Quicktime VR software process (24mm capture lens)

1 Stitching the images

This is where you take the separate images (see the Sydney Opera House sequence in Figure 6) and - using software - meld them into a seamless whole (Figures 7, 8). In essence what you are doing is using software to produce the same kind of cylindrically distorted image a specialist panoramic camera would create, for approximately one quarter of the cost.

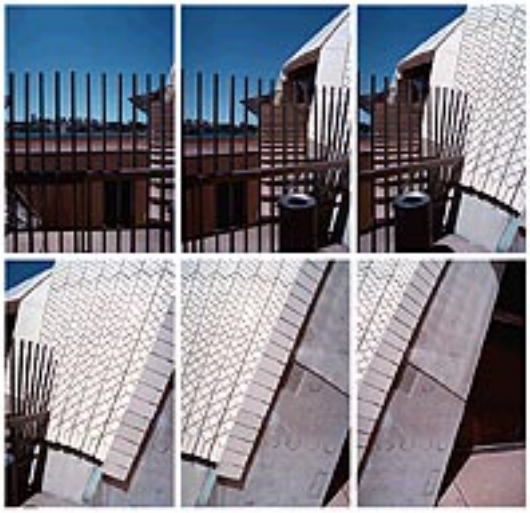

Fig 6 Six images prior to stitching

Fig 7 The six images after stitching

The advantage of doing it this way is that you can use almost any camera to photograph a scene. The disadvantage is that you have to employ a lot of computing grunt to make it happen. Hence the fast PowerPC.

I use the following in my MPW worksheet to do a stitch. Although Apple supplies MPW script files to make the stitching process 'easier', I prefer to call the stitch tool directly so I can see exactly what parameters are fed in.

```
# STITCH
#
# Because of 24mm lens, FOV = 60 & outHeight = 3744
#
set scansPath        "Schnieder HD:myQTVR:Raw Scans ƒ"
set panOutFolder           "Schnieder HD:myQTVR:WIP"
set outName                "aznLatest.xx"
set scanFilenames          "01-16"

set myFOV                  60
set myOutV                 768
set myOutH                 3744

stitch -fovy {myFOV} ∂
   -offset 250 0 -range 60 30 -rotate -90 ∂
   -dir "{scansPath}" -files "{scanFilenames}" ∂
   -wrap -cropOut -outHeight {myOutH} -outWidth {myOutV} ∂
   -sharpen -blend -fill -show ∂
   -out "{panOutFolder}:{outName}.srcPict" ∂
   #-i               #Uncomment for INTERACTIVE stitching
```

As Apple assumes a 15mm lens for all its MPW scripts and tools, adjustments have been made to accommodate the 24mm lens. Specifically, the Field Of View (fovy) value has been reduced from 97 degree to 60 degree and the offset and matching range have been widened to allow for the 24mm lens smaller vertical and horizontal coverage.

The 'outHeight' parameter has also been increased to reflect the greater number of shots required to cover a 360 degree view. The value of 3744 is the nearest multiple of 96 (4x24) to the output height produced whenever a stitch is done 'raw' (ie. without height/length concatenation). You need to pay attention to the height and width dimensions as the dicing tool will fail if you get them wrong (see below). RTFM.

I recommend sticking with Apple's 'outWidth' of 768 for the sake of RAM and movie run-time efficiency. Although you could work with larger image sizes, in practice the final movie quality does not appear to benefit from it. Likewise, the 'sharpen', 'fill' and 'cropOut' activities could be performed later in Photoshop, but again little practical benefit has been found in doing so.

You will notice that the script has provision for doing stitches interactively. This allows you - rather than the software - to determine precisely where images should be joined. Some developers always stitch with this option on, I prefer instead to stitch automatically at first and then re-stitch interactively only if there are problems.

After highlighting the script lines in MPW, you hit the enter key and then watch as the stitcher does its magic and pieces together the separate scans into one panorama (the 'show' switch above puts up a window during stitching to display the panorama being built). As mentioned earlier, it takes roughly four minutes for this to happen on a PowerPC 7500, at the end of which the resulting 8 MB file (see Figure 8) is placed into the 'panOutFolder'. Notice how the image is turned onto its side? This is deliberate and is required by the dicing tool.

2 Retouching the stitched image in Photoshop

As no scan is perfect, you will always have to adjust the image gama, balance the colours and remove dust spots. Similarly, no stitch is perfect and sometimes you will also have to touch up any blend 'failures' (see Figure 9).

Another thing you can (and should!) do in Photoshop is produce scaled down versions of the panorama for low-res versions of the Quicktime VR movie. An image size of 1056 x 216 will compact well into a small enough movie to be useable on the internet (see below).

3 Dicing the Panorama

As an intermediary step, you must run the stitched panorama through the dicing tool to convert it from a PICT file into a Quicktime movie made up of 24 compressed frames.

The compression applied is of the Cinepak 'cvid' kind. It is particularly aggressive (8MB -> 800KB) and results in a moderate degree of image deterioration - so there is little point in being overly precious when retouching the panorama.

For the dicing step to work properly - as well enable your movies to be playable across platforms - you must make certain that the height of the panorma is exactly divisible by 96 and its width exactly by 4. In addition, if the height is not exactly divisible by 24 then you will get a dicing error and then you will have to re-stitch (or go back to Photoshop) to re-scale the image.

Fig 8 The final 360 degree stitched panorama

Because this is a straight-forward step, I call the 'SrcPictToMovie' MPW script supplied by Apple:

```
#DICING
#
set myBasename      "aznLatest.xx"
set myRoot          "Schnieder HD:myQTVR:WIP"
```

```
SrcPictToMovie    "{myRoot}:{myBasename}.srcPict"
"{myRoot}:{myBasename}.srcMooV"
```

On a PowerPC this step requires 20 MB of RAM and takes 3 minutes. On 68k machines it takes eight times longer(!).

Fig 9 The white box encloses a stitch 'failure'

4 Making the Quicktime VR movie

The final step is the conversion of the standard Quicktime movie produced by the dicing tool into the specially formatted Quicktime VR version:

```
#SINGLE-NODE MOVIE (24mm lens, WIDE size)
#
set myRoot           "Schnieder HD:myQTVR:WIP"
set myBasename       "aznLatest.xx"
set myPan            32
set myWind_H   400
set myWind_V   240

msnm "{myRoot}:{myBasename} (wide)" -source ∂
"{myRoot}:{myBasename}.srcMooV" ∂
   1 24 -vPanRange {myPan} -{myPan} ∂
   -windowSize {myWind_H} {myWind_V} ∂
   -defaultView 0 0 50
```

Again because of the non-15mm lens, some of the parameters have been adjusted in order to get the perspective right. The 'vPanRange' has been reduced from the more usual ±42.5 to ±32 degrees to allow for the shallower field of view for the 24mm lens. In addition, the default size of the completed movie has been enlarged slightly from the standard 320x200 to 400x240. The '0 0 50' parameters tell the tool where to set the default pan angle(s) and zoom ratio.

This step only takes a few seconds and results in a 800KB movie.

The Completed Quicktime VR movie

I lied eariler because there is still one more step before the movie is completed! You must open it in the QTVRPlayer application which comes with the tools and re-save it as 'Self contained' and 'Playable on non-Apple computers'. This final step will allow our QTW equipped WINTEL friends to see what they are missing out on.

Quicktime VR on the Internet

At 500-800KB, a 'hi-res' Quicktime VR movie is simply too large to put on the internet. Although there are many ways of going about it, the following works well in producing scaled-down low-res versions (assuming a 24mm capture lens):

- In Photoshop, save a copy of the stitched panorama
- Re-size it to be 1056x216 (note that 1056 is exactly divisible by 96 and 216 by 4)
- Run the dicing and msnm tools on this smaller image, setting the msnm default window size to 240x140
- Re-save the movie as self-contained & playable on non-Apple computers, giving the resulting file either a '.mov' or '.qvr' extension

The image quality isn't incredible, but at a 98KB size, the resulting movie is small enough for anyone to download.

On my home page I have also provided a small drag 'n drop application to allow users to restore the creator/file type of the downloaded movie to Quicktime VR's 'vrod'/'MooV' - which makes sure users will be able to play the movies no matter what happens!

Other than supply movies for others to view, the current state of the technology does not allow you to do much else. You cannot launch URLs from hotspots embedded in the movies, neither can you play sounds or link to graphic or text files - for this you need to construct MacroMedia Director, Apple Media Tool 2.0 or Hypercard extravaganzas.

Quicktime VR will aquire a more formidable net presence though once a 'C' API is released. When that happens, Netscape plub-ins can be written and then VRML will be 0xDEADBEEF.

The Future of Quicktime VR

According to postings to the Quicktime Mailing List (see the 'See Also' heading below), a phallanx of Apple DTS Engineers is working furiously to complete a 'C' API similar to that for Quicktime. Currently (Febuary 1996) only Hypercard and Director XCMDs are supplied with the authoring tools - meaning if you want to incorporate your movies with anything else, then you can only do via environments which support these externals (N.B. version 2.0 of the Apple Medial Tool features built-in support).

When (if?) the 'C' API becomes available, you will be able to integrate panoramas with anything you like. You could launch URLs, play sounds, link to movies or still pictures or generally write your own player applications which do *exactly* what you want. Which is why we became developers in the first place right?

To be fair, Apple has recently done the right thing in lowering the price of the Authoring tools from a stratospheric $2000 (US) to a more reasonable $500 (US). As of 12 Dec 1995 they have also waived most of the run-time license requirements. The MPW requirement should go and apparently will - with the next(ish) version of the authoring suite.

Meanwhile, we wait…

See Also

http://qtvr.quicktime.apple.com
(Apple's Quicktime VR site should be your first port of call for up-to-date information for developers, the latest versions of the Quicktime VR players and links to various organisations working in this area.)

listproc@solutions.apple.com
(No subject, the following in the message body)
subscribe quicktime-dev (Your Organic Name)
(The Quicktime Development mailing list features extensive discussion of Quicktime VR related issues - in fact there is often so much that it swamps the more general Quicktime related postings.)

Tom R. Halfhill "See you Around" *Byte Magazine* May 1995 pp.85-90.
http://www.byte.com
(A good overview of of the what/how of Quicktime VR. Includes a comparison with Microsoft's "Surround Video".)

Kaidan@aol.com
http://www.kaidan.com
(Manufacturers of Quicktime VR camera mounting accessories.)

http://www.usit.net/omniview
(Omniview's 'PhotoShpere' competition to Quicktime VR. Instead of a wide angle lens you use a super-expensive 'fisheye' to capture image hemi-spheres.)

http://www.zeta.org.au/~aznemeng
(The Man From Warrimoo's home page featuring ten or so low resolution Quicktime VR images of Sydney and environs.)

END