

Panoramas (photo-stitching) – Part 2

T.O.M.S.

In Part 1 we related that, whatever software package we used to do the image blending or ‘stitching’ to produce a composite panorama, we found we could improve the results quite significantly by manual intervention; firstly by paying extra attention to the way in which we captured the source images and, secondly, by applying pre-processing to at least some of the images before the final rendering.

The gamut of possible ‘tweaks’ is quite extensive, some more effective than others, so please bear with us if we spend a while working through the repertoire.

AutoStitch (PanoramaPlus 2 demo)

But first an apology; sorreeee. Due to an oversight, we omitted to mention in Part 1 that a free, demo version of *PanoramaPlus 2* is available. It’s called *AutoStitch* (v2.185 demo), is distributed by the original authors (*not* Serif), and the 1MB file can be downloaded from:

www.cs.ubc.ca/~mbrown/autostitch/autostitch.html

There’s some debate as to whether this reportedly limited-feature demo is as effective in operation as the fully-fledged commercial release (licensed by Serif) but it will certainly give you a very good flavour of what is possible. Plus all the ‘tweaks’ which we’ve deduced for *PanoramaPlus 2* can be applied to it – and indeed to the other titles we considered in Part 1.

Note however that the *AutoStitch* Edit-Options window is seriously unfriendly, with numerous meaningless terms (e.g. RANSAC Parameters; *Que?*), although the *Read_Me* file does occasionally assist. Perhaps many of these options are irrelevant anyway as they don’t appear in *PanoramaPlus 2* which is *much* more polished and user-friendly. But don’t be put off by all this; do try before you buy.

The more the merrier

A point we cannot overstress is that you can never have too many source images. The beauty of a digital camera is that running costs are negligible, so *always* take as many piccies as you can; you can always delete the unused but, other than doing a complete re-run, you can’t add to them.

In any event, once you’ve bought a powerful PC out of the £500-£1,500 you’ve saved on not having to buy a wide-angle lens (!), there should be enough left over for an extra media card or two, plus a spare set of batteries.

Even then, you probably won’t need the extra batteries as you’re unlikely to be using the flash and, even with a camcorder, the zoom control won’t be needed (to be discussed).

So do take as many images as you can, perhaps in different conditions where one sweep might prove to be unusable but another (maybe if the sun goes behind a cloud) turns out to be just the ticket.

For reasons we’ll discuss later under ‘*Cropping*’, we also recommend adding one or two further images – at least – to both the left and right ends of a planned, horizontal panorama.

It doesn’t matter if these ‘extras’ have little or no aesthetic merit – a bunch of boring bushes will do – as they will later be cropped off intentionally anyway.

Use of a tripod & spirit level

In the *Tips* section of their *Companion* document, Serif say the use of a tripod is “useful” but we feel that’s quite an understatement.

In fact, we’d be inclined to say that, if you need to stitch four or more photos together, a tripod becomes increasingly important, and near-essential for genuine panoramas (say, anything wider than 90°).

By the same token, if a video camera is used, a tripod is also near-essential. Without one, with the best will in the world, the image outline will be uneven. And having to crop off big chunks, maybe spoiling the whole effect, becomes virtually unavoidable.

For similar reasons, a spirit level can be beneficial for extra-wide, horizontal panoramas, where the plane of the tripod head in both axes becomes crucial to avoid having to crop off large areas of the composite image and risk marring it (Fig 1).

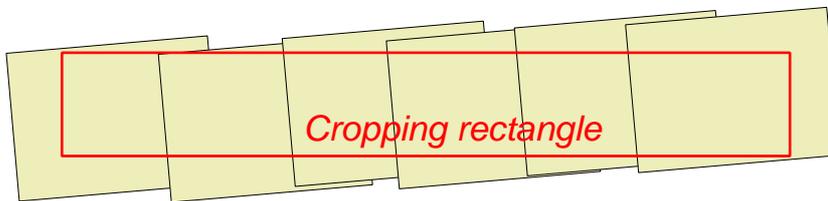


Fig 1 “I do wish I’d used a tripod ...” (NB: The cropping rectangle *is* horizontal!)

Spirit levels are available from camera retailers for around £10 and fit conveniently onto the camera’s hot-shoe.

However, if you don’t have one, you can achieve acceptable results by panning the camera (laterally) across a horizon, before capturing your images, and ensuring there’s no “up/down” error.

If there is, adjust the tripod leg(s) and repeat until it’s spot on. Accuracy at this stage means significantly less cropping of the blended image (Fig 1 again).

Camera exposure setting

Serif recommend “lock[ing] the exposure”. After *lots* of trial and error, we would tend to disagree with that and advise the opposite.

As we see it, the problem with a locked exposure is that, if you pan the camera within (say) $\pm 60^\circ$ of the sun’s azimuth, even if it’s behind cloud, there are such *massive* changes in light levels that the resulting panorama will inevitably be far too overexposed in that segment, and heavily underexposed elsewhere.

We would therefore recommend the alternative of letting the camera (digital or video) sort out an *average* exposure by using an automatic exposure setting (AE or P).

The only problem then is that, as the ambient light levels change so rapidly near the sun’s azimuth, the ‘joins’ between adjacent exposures become more and more evident. However, this can be allowed for (see next).

Even if you don’t point the camera directly towards the sun’s azimuth, there may still be

significant changes in light levels. This is because the area directly opposite the sun’s azimuth is, er, sunlit, so does tend to be much brighter than the two segments at $\pm 90^\circ$ to the sun.

Overlapping images

Serif recommend an overlap of 30% to 50% between adjacent images. Provided you have time on your hands (both for taking the extra pics and the subsequent blending), we would recommend *not less than 50%* overlap, so call it 55%.

Furthermore, if you’re pointing within, say, $\pm 60^\circ$ of the sun’s azimuth, we’d then recommend increasing the overlap, in that segment, to at least 60% – or even 70% (Fig 2) – to help counter the “you can see the joins” effect mentioned in the previous topic.

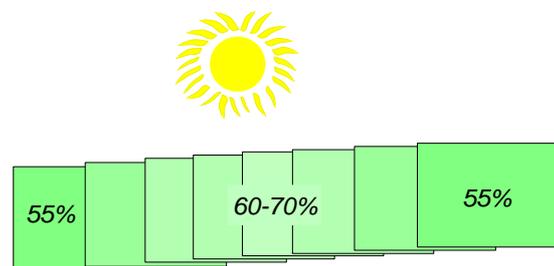


Fig 2 Image overlap (offset for clarity)

To combine this and the previous point, we feel that setting auto-exposure helps to even out the average light level throughout the panorama; while increasing the image overlap – especially approaching the sun’s azimuth – reduces differences in exposure between any one image and its immediate neighbour.

Taken together – maybe with some pre-processing as well – in our experience this really does help to avoid the “you can still see the join” phenomenon.

Extra images

An incidental advantage of increasing the general overlap to around 55% or more is that, should a particular image turn out to be unsuitable (for whatever reason), it could in principle be dumped and the two adjacent images stitched together; albeit with minimum overlap.

It’s far safer to gather more images than you think you may need, even repeating a ‘sweep’, just in case you need to fill in with extras.

In *PanoramaPlus 2*, if using a video camera, this is resolved during the blending process, simply by increasing the setting for the number of snapshots grabbed per second before processing.

Focal length setting

Serif say “set the camera’s zoom” – but they don’t say what to...! By trial and error, we’ve concluded that by far the best results are obtained if the camera zoom (digital or video) is set closest to what the human eye will normally see.

The focal length of the Mk 1 eyeball (and incidentally of 35mm film) is 50mm *but*, because the digital camera sensor is a CCD, there’s an inherent magnification of $\times 1.6$.

So the lens’s focal length needs to be set to *31mm* instead. Some digital camera zoom lenses are marked to show where the 31mm setting is.

Setting the focal length of a camcorder with a motorised zoom lens isn’t so straightforward as they rarely give a readout in millimetres. However, if you set the zoom to the point that it ‘looks right’ (i.e. what you will see with the naked eye), then zoom in a little more (towards the telephoto setting), that should suffice.

If necessary, the focal length can be increased without seriously upsetting things. But we would strongly recommend not *reducing* it to significantly less than the standard setting as the focal length would then tend towards the ‘fish-eye’ phenomenon, where the apparent curvature of straight-line objects will play havoc with the subsequent blending process.

If you have a need to reduce the focal length to less than 50mm (31mm for a digital camera CCD) to capture the full object being photographed, then – instead – consider capturing a ‘matrix’ of source images (to be discussed).

Vertical panoramas

Serif discuss the possibility of taking vertical montages as an alternative to the traditional horizontal panoramas. This is perfectly feasible, using exactly the same procedures and camera settings, except that the camera is panned in the vertical axis instead of the horizontal.

The significant difference compared with the more conventional horizontal viewpoint is that, as the camera is panned upwards from the near-horizontal, the distance from the lens to the object being photographed will increase, so it will appear to get smaller.

In practice this results in an apparent ‘curving’ of all vertical lines (e.g. the sides of a building) towards a vanishing point in the overhead, making it much more difficult for the blending software to make sense of it.

Again, experience shows that increasing the degree of image overlap – even to as much as 60-70% – will significantly improve matters.

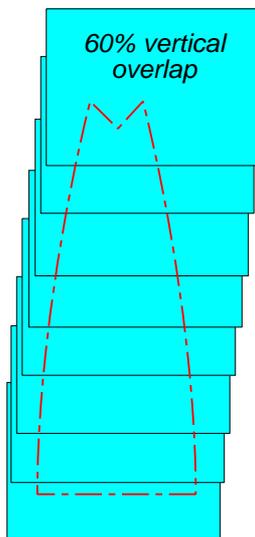


Fig 3 Vertical panorama
(offset for clarity)

The increased overlap between images is illustrated in Fig 3.

In this context, especially if you can't stand back from the object being photographed, it is even more important to resist the natural temptation to go for a reduced, 'fish-eye' focal length; it simply screws up the blending process.

Note that, because of the *apparent* curvature of the underlying object, the boundaries of the resulting panorama image to be markedly uneven and this will need to be removed by relatively extensive cropping, so allow for this by extending the image coverage well beyond the *width* of the object being photographed.

If in any doubt, don't hesitate to play safe and opt to capture a 'matrix' of images instead, as discussed in the next topic.

There's an example vertical panorama in the Articles-Panoramas sub-directory on the CD.

Matrix panoramas

If you simply cannot capture both the full width *and* the full height of an object to be photographed, without resorting to reducing the focal length significantly, then instead we would recommend capturing a matrix of source images by panning the camera in both the horizontal and vertical axes.

So, for example, if the source object is a tall building which is too wide to be captured fully by a single, vertical column of images (Fig 3 again), then the answer is to capture two overlapping columns of images (Fig 4).

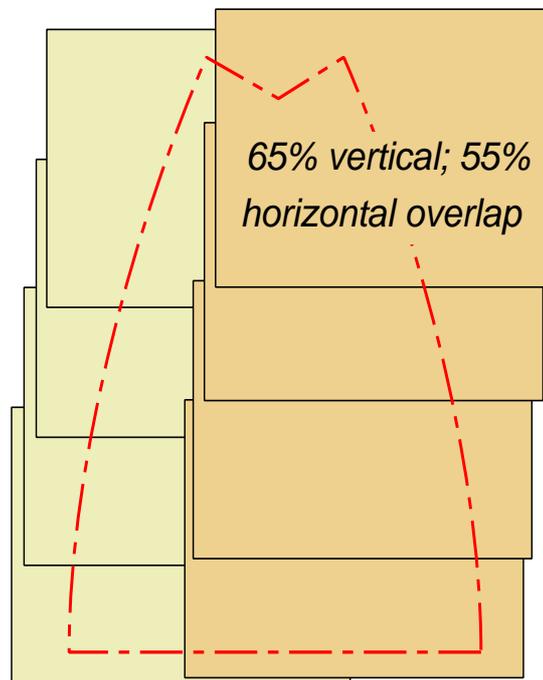


Fig 4 Matrix panorama (offset for clarity)

Each vertical column could be based on an overlap of 60-70%, as discussed earlier, while the overlap between two adjacent columns should be not less than the 50%; as if you were panning in the horizontal.

The *PanoramaPlus 2* blending software is well capable of sorting out the complex matrix of source images and, in our experience, the end results will be far better than reducing the focal length and risk throwing the processing due to the fish-eye effects.

If you're using a video camera, once again there's no reason why you shouldn't make two or even more 'sweeps' with the camera, although you may need to tinker with the frames per second setting, prior to processing the final image.

With matrix panoramas, the boundaries of the resultant image tend to be yet more uneven and, again, this is removed by cropping. Once more, allow for that when capturing the source images.

There's an example matrix panorama in the Articles-Panoramas sub-directory on the CD, showing this marked irregularity of the composite image boundary.

Image movement

It can be important to avoid moving objects if at all possible. For example, if you photograph a seascape, the movement of wave crests between adjacent exposures can confuse and throw the stitching process.

On land, moving objects can also badly throw things if they appear at different places on two adjacent source images. If possible, wait for the object(s) to depart but, if that's not possible (e.g. a stream of traffic) take more overlapping exposures and pick-and-choose the best before starting the blending process.

Cropping the image

The composite image will never be perfectly rectangular, for various reasons, a common factor being that *any* camera movement in the axis opposite to the panning direction will result in higgledy-piggledy edges along the longer sides. Accidentally kicking the tripod legs while panning is a favourite!

The final image can of course be cropped to suit using a suitable utility, for example David Pilling's scanning software (Windows or RISC OS). Both *PhotoStitch* for Windows and *PanoramaPlus 2* call up a cropping rectangle which can be adjusted, as required, before final processing (each has drag boxes along the four sides).

For some reason we've not been able to suss out, with wide-angle/horizontal panoramas, errors seem to build up noticeably at both ends.

So to cut back to a perfect rectangle, you need to either crop the image height, perhaps more than you would ideally wish, or alternatively crop a large-ish chunk off each end. Neither option is very satisfactory.

For that reason, we suggest fooling the software by adding one or two extra 'dummy' images at the left and right ends of the panorama (as mentioned earlier).

Although an error build-up at each end will be virtually unavoidable, in practice it will now become most marked in the extended area, containing the dummy image[s] (depicted in darker green in Fig 5).

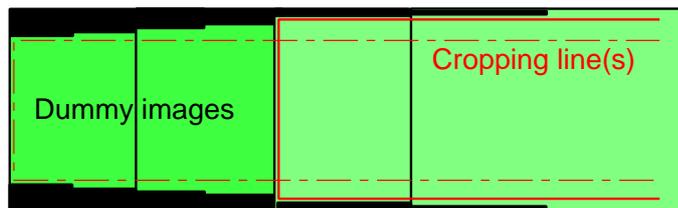


Fig 5 Use of dummy images (repeated at right hand end)

But that means we can retain significantly more of the usable portion of the image (inside the *continuous* red cropping area), without being forced to cut off much more in height (the dashed red area). Cunning or what?

Choosing the image resolution

The two *PhotoStitch* applications output a high-resolution image, based on the underlying source images from the camera (still or video). The composite image can of course be scaled down if desired, using any number of utilities.

As described in Part 1, p4, *PanoramaPlus 2* gives total freedom on what resolution the rendered image(s) will have, ranging from the relatively low-res 'semi-thumbnail' (which in itself is fine for use in a document, for example), to the high-res version based on the source images (in size terms, identical to *PhotoStitch*). Any interim size can be specified and entered manually (by width or height).

To cover all likely eventualities, we tend to save both the low-res and the high-res images, in the default TIF format, although they can then easily be converted to JPEGs in, say, *DPScan*.

Viewing the panorama

Even if you have a widescreen TV, you're not going to get sensible results from viewing panoramas as the image height, as displayed on the screen, will of course be very small, and you simply won't see much detail.

The only exception we can think of is if, for example, you blend just two or three standard 4:3 images into a 16:9 composite.

But that would apply only if you can't decrease the focal length of the camera to widen the view (and then crop the result), or if you can't stand back from the object to be photographed (to increase the field of view).

Alternatively, viewing a panorama on a computer monitor can be a pleasure, especially if that also is a widescreen model.

Clearly you can't view it all at once on screen, in any detail. However, most image viewers will allow you to select *Fit to Height* or similar, and then scroll left and right to view the image in segments.

Subjectively, the sheer amount of detail in the composite, high-res image (*c. 40 megapixels!*) makes this a fascinating hobby.

What image resolution is best for viewing? The low-res 'semi-thumbnail' image from *PanoramaPlus 2* is only *c. 200* pixels high, so that's too coarse. but the hi-res version may simply be too big to handle conveniently.

A half-way house setting for *PanoramaPlus 2*, before rendering, is to set the *Height* to whatever screen mode you use (e.g. 1024 pixels for a 1280×1024 mode), and save that for viewing on the monitor. This is a very effective compromise.

Printing panoramas

There's little doubt that by far the most impressive way of viewing a wide-angle panorama – in its entirety – is if printed at large 'banner' scale, on posh photo paper, either commercially (*very expensive for a single print*), or by resorting to DIY.

We don't know of a *current* solution for RISC OS, although we're still successfully using a six-year-old printer and *PhotoReal* driver under *VirtualRPC*, the only problem being getting consumables.

Under Windows, DIY gives us various options, only some of which may be feasible or suit your pocket, so are listed in order of decreasing costs:

– Get yourself an A3+ photo printer (around £300-£500 from Canon, Epson and HP) and – *if you can get it* – a 10m roll of 329mm wide photo paper. Under Windows (with all the paper rolls listed here), you will probably need to set up the desired paper size under the printer's *User Settings* or similar, but that's reasonably straightforward. Obtaining the paper roll (£47-£53!) may well prove a stopper; we think Epson were the only source but they no longer list it on their website.

– Get yourself an A4 photo printer (various brands, *c. £60-£90*) and a 10m roll of 210mm wide photo paper. For a wide-angle panorama, this is a good alternative, especially if the printer will support borderless printing. Again, obtaining the paper roll may be a problem; Epson's own brand (£37!) is currently "out of stock".

– If you can source it, obtain a pack of Epson '*Panoramic Photo Paper*' (10 sheets for *c. £12*) and use it with an A4 photo printer. The size is 594mm×210mm so an approx. 100° wide, semi-panorama will fill the sheet.

– Obtain a 10m roll of Jessops's '*Photo Gloss*' paper at *c. £10*. The medium is only 100mm wide but that will display a full, 360° panorama (around 1,150mm long).

If anyone knows of any other source(s) of 'banner' photo paper, we'd be very pleased to hear from you as we know people other than ourselves who will soon be needing replacement stocks – even more so if they start playing with these photo-stitching programmes...

The end results

As proof of the pudding, we'll try to combine all these points into a couple of comparative illustrations (Figs 6 and 7), which depict a roughly 90° segment of the demo, near-360° composite panorama contained in the *Articles-Panoramas* sub-directory.

Common factors are that exactly the same *source* images were used for both figures, taken using a tripod, automatic exposure, 55% overlap and 31mm focal length, all as we've recommended earlier.

Although at the outset we said we wouldn't compare individual packages, the evident hiccups in Fig 6 are typical of the sort of thing we see with freebie, 'bundled' titles. We've already looked at one 'ghost' image (Part 1, p3) but, if you increase the scale of Fig 6 to full-screen, you'll see many more.

Even more evident is that "you can see [some of] the joins", especially at the T-junction, and the clouds at top-right are looking 'mucky' (best viewed in a 16M colour screen mode).

On the other hand, the identical images dropped into *PanoramaPlus 2* (and the demo *AutoStitch*) show no ghosting whatsoever, nice 'clean' clouds and, on the initial blending, only the join at the T-junction was still there (Fig 7), although not as evident as in Fig 6.

This was 'ironed out' by pulling out the three adjoining source images, applying a bit of gamma correction by trial-and-error (0.7-0.8 in *DPScan*), saving them and re-blending.

The result is much-improved, as evidenced in Fig 7. With patience, it can be done.

T.O.M.S., toms@ndirect.co.uk



Fig 6 Blended by Canon PhotoStitch, without image pre-processing



Fig 7 Blended by Serif PanoramaPlus 2, with some image pre-processing