

Simon Schofield

Paper: An approach to creating very large, high resolution artistic printed images**Topic: Visual Arts****Authors:**

Simon Schofield
 Department of
 Computing,
 Nottingham Trent
 University, Nottingham,
 UK

simonschofield.net

References:

[1] Cline D et al. (1997).
*Dictionary of Visual
 Science*. 4th ed.
 Butterworth-Heinemann,
 Boston

[6] Wikipedia authors
 (2011) *Largest
 photographs in the
 world*, Accessed 2012.
[http://en.wikipedia.org/w
 iki/Largest_photo](http://en.wikipedia.org/wiki/Largest_photo)

[8] Porter, T. and Duff, T
 (1984). *Compositing
 Digital Images*.
 Computer Graphics 18

(3): 253–259.
[doi:10.1145/800031.808
 606](https://doi.org/10.1145/800031.808606). ISBN 0897911385.

Abstract:

Contemporary printing processes, such as wall and floor graphics, offer artists the potential to create very large-scale pieces. However creating images that exploit both the potential size (say many metres square) AND resolution of the media (say 200 dpi), is very difficult using standard bitmap editing software as both the creative processes involved and file sizes become too cumbersome to manage. We present an approach to this problem that uses a combination of algorithmic techniques to control the generation of such an image as a set of non-repeating, seamlessly tiled sections, and facilitate a high degree of artistic authorship throughout the process. Our approach also necessitates the use of generative techniques, primarily in generating a very high degree of local detail over the entire surface of the image. We also hypothesise a further generative technique for building potentially limitless images at high resolutions. A working example is given.




Image of Generated Artwork

Contact:

[simon.schofield@ntu.a
 c.uk](mailto:simon.schofield@ntu.ac.uk)

Keywords:

Artworks, Print, resolution, generative texture, bitmap manipulation

An approach to creating very large, high resolution artistic printed images

Dr Simon Schofield MA, PhD

Department of Computing, Nottingham Trent University, Nottingham, UK

www.simonschofield.net

simon.schofield@ntu.ac.uk

Abstract

Contemporary printing processes, such as wall and floor graphics, offer artists the potential to create very large-scale pieces. However creating images that exploit both the potential size (say many metres square) AND resolution of the media (say 200 dpi), is very difficult using standard bitmap editing software as both the creative processes involved and file sizes become too cumbersome to manage. We present an approach to this problem that uses a combination of algorithmic techniques to control the generation of such an image as a set of non-repeating, seamlessly tiled sections, and facilitate a high degree of artistic authorship throughout the process. Our approach also necessitates the use of generative techniques, primarily in generating a very high degree of local detail over the entire surface of the image. We also hypothesise a further generative technique for building potentially limitless images at high resolutions. A working example is given.

Introduction

Increasingly accessible large-format printing processes, such as those producing billboard size wall and floor graphics, offer artists the potential for creating large scale artworks. Such systems use standard ink-jet technology and so enjoy the potential capability of printing at near-photographic resolutions (300 dots per inch or greater). However it is highly unusual that such high resolutions are used in the production of billboard size graphics. The traditionally held view is that such high resolutions would be squandered at a billboard's intended viewing distance of, say, 10 meters. Also, sourcing and manipulating the enormous digital image file necessary to provide such a high degree of resolution over such a large expanse creates a number of problems that we discuss.

In most extant cases, billboard sized prints use a resolution only adequate for distant viewing (say 4000 pixels by 2300 pixels, over 6 x 4 meters, viewed at 10 meters) with a resultant print density equivalent of around 17 dots per inch (dpi). This low-resolution printing strategy suffices for the purpose of exterior billboards and hoardings. However, it is not really

an option if we consider producing images for the floors or walls of a domestic setting, or an art gallery. Looking down on a floor graphic at a typical viewing distance of 1.5 meters, a normally sighted person can resolve detail of up to around 100 dpi [1] (see footnote 1). In a domestic setting, a wall graphic might be inspected as if it were a painting, perhaps from a distance of 0.75 meters, where a higher resolution of 200 dpi might suffice. At this higher resolution, if the floor or wall image were 6 meters by 3 meters, this would require an overall digital image size of 48,000 by 24,000 pixels (1,150 Mega Pixels); what we refer to, for convenience, as a "massive image".

While a 1150 MP image can be theoretically declared in memory (Photoshop currently potentially allows users to declare an image 300k pixels x 300k pixels, a 9000 MP image [2]), an image of such size would be extremely cumbersome to manipulate within current interactive environments such as Photoshop or Gimp [3]. File loading, manipulations, redraws and save times become frustratingly slow under such conditions. Worse still if one starts to build up a number of layers to manipulate the content of the image (as is standard practice in most image editing software), as each layer adds further demands to the memory. Under such duress, both software and hardware may also become unreliable, and this coupled with the overall slowness of response makes for an untenable situation.

The next problem posed by massive images is in sourcing the desired image data. Clearly there are extant processes that can *mechanically* provide the image-matter to fill the space of such massive images. A single continuous photographic scene can be assembled from a set of abutting photographs such as Google Maps [4]. There are many other large composite images, such as the current record holder, a 272 Giga Pixel image of Shanghai [5] (See "The Largest Photographs in the World" page of Wikipedia [6]) that could be used to populate a highly detailed continuous printed images of this sort of scale. Many visual artists, however, seek to create quite arbitrary images, such as digital murals, designs, paintings or complex montages where high degrees of visual creativity and arbitrary intervention come in to play over all parts of the image. They wish to work freely gathering imagery from many disparate sources, and so have only limited need of this sort of mechanically produced content.

Another methodology to create high levels of detail over a large physical expanse is to use the step-repeat method of wall paper; while this allows for a high degree of creativity over the single repeat pattern, there is no scope for variation over the image as a whole.

Without recourse to repeating sections, or using mechanically harvested imagery, the artist is left to apply the detail over every part of the image as if “by hand”, so that any section can be viewed close up and provide visual engagement. Such an undertaking via standard image editing techniques would be oppressively laborious over such a large surface.

Forbearing of the above, once the piece has been fully composed, it must be printed out as a series of partially overlapping or abutting *print-tiles* (or “pages” as they are sometimes referred to in the print trade), each generated from a separate print file. These print-tiles are then physically re-assembled into the whole image on-site.

We present an approach for creating extremely large non-repeating murals that addresses many of these technical and creative problems. The output image from our approach is arbitrarily large; both logically and physically, highly detailed over every part of the surface, and contains no discernible repeat patterns. Our approach also affords a good degree of artistic control and visual feedback during the production process. Our technique makes necessary and good use of generative techniques, primarily in redressing the image-detail problem, and speculatively in generating images of potentially limitless size.

We present an actual example of our system at work, generating an image of 54,000 x 12,000 pixels for a printed piece 18 meters by 4 meters at 100 dpi. We also hypothesize that our approach enables the production of almost limitlessly large printed surfaces, where every part of the printed surface can be unique.

Method

In our approach, the memory problems of massive images are mitigated through generating the full-scale imagery only during the production of the final printed output, one print-tile at a time, and therefore, one *print-file* at a time. Each print-file is of a known and manageable size; the larger the overall image, the more print-files are generated. As the system only needs to cope with one print file at a time it is the in-memory size of an individual print-file, rather than the overall image, which is the limiting factor. Consequently, in our system, the full size image never actually exists in memory in its entirety.

The overall design, composition and user-interaction for the whole image all takes place in a drastically reduced, small-scale “key” image

assembled and manipulated by the artist in a standard layer-based image editing system, such as Photoshop. The completed design, held in the key-image, is composed from a number of layers, each layer containing a single layer-patch of "image matter". Each layer-patch of the key-image has a smallish bounding box in relation to the overall key-image size (see Fig. 2), and each layer-patch corresponds, by its layer name, to a much larger version of the layer stored as a distinct "full scale patch" file in a local directory. The layer-patches, both small and full-scale, contain alpha to determine transparency, and give them visual irregularity. The layer-patch positions, relative drawing order and layer names stored in the key image are made available to a specially developed scriptable image-compositing engine through the use of Python/COM [7]. Using this data, the image-compositing engine is able to produce a full-scale render of a specified print tile. For each print tile, the engine inspects the layers occupying that section of the key-image (see Fig. 2). The system then creates a print file, using the full-scale patches, based on their relative positions within the key-image to create a "flattened" image using Porter-Duff [8] alpha compositing. Each print file is saved out in turn.

The compositing engine operates by giving access to a set of low level image handling functions (such as load, save, image transforms, filters, and compositing functions) through Lua script [9], and is used both the full scale rendering and, as described later, in the rendering of synthetic textures. Lua sessions can be saved and loaded, making for a convenient and flexible image processing system.

This technique of working small, and rendering big, through the use of small scale "image-proxies" has been utilised as a core methodology in previous systems such as HSC Software's *Live Picture* [10] although the technique has somewhat fallen from common usage today, presumably because systems such as Photoshop suffice in most practical cases.

In our actual example given below, the completed key-image was composed of around 100 layers (Figure 1). However, we found that we could use duplicate layers several times within the overall image without the duplication becoming visually obvious, as the layers visually intermingled in different ways. This reduced the number of different full-scale patches required to around only 25.

Once this process pipeline has been established, it is easy to view and control the overall composition of the image, and edit and make new patches as required. So, while the final rendering process is script-driven and mechanistic, the process whereby the overall image is constructed and manipulated enjoys most of the interactivity of a Photoshop session. Indeed, the process provides a very fluid method for working from sketch

to final version, so long as the artist maintains the relationship between the small layer-patches in the key-image, and their full-scale counterparts.



Fig 1. The overall key-image, composed from around 100 layers



Fig 2. The surface broken into print tiles (defined by the white lines), with two "layer-patches" of duckweed highlighted in yellow. The full-scale render of the region within the red box is shown in Fig 3.




Fig 3. Indicating the level of detail at all points of the final image

Generating Detail

Apart from managing the large size of the overall image and providing a good degree of interactive authorship dealt with so far, our process also hinges on being able to generate a large variety of full-scale image patches containing extremely high levels of detail (from which are derived their small scale layers in the key-image). In our working example we use three techniques for generating the full-scale patches; natural photography, standard compositing, and a generative technique for texture synthesis.

In the example image, the pebbles and Elodea pondweed (the long flowing weed) were obtained by straightforward photography and given alpha masks by standard matting techniques. The patches of Fish were created originally from natural photographs, but then cut out and arranged into fluid shoal-like compositions using standard image editing software.

The patches of Duckweed and Junk (Figures 4 and 5) were synthetic textures generated using the scriptable image processing engine. The basic process was to generate a detailed final texture (output image) from the repeated compositing of smaller elements (input images) according to a set of stochastic functions. The input images, in the case of the duckweed patches, is a set around 20 alpha-masked Duckweed leaves stored as alpha-masked PNG files. The system builds a list of these input

images, selects one at random, scales and rotates the image within bounds, and attempts to place it onto the output image at a random point. A larger mask-image imposes an overall distribution of the duckweed, and is responsible for the overall shape of the patch. The final output rafts of duckweed were constructed using something in the region of 5000 iterations. A "junk carpet" was generated using around 100 individually photographed, and alpha-masked differing pieces of junk, and can be seen laying at the bottom of the pond in the final image.




Fig 4. A generative texture: Duckweed patch and detail



Fig 5. A generative texture: The "junk carpet" patch and detail

This method of texture synthesis, where textures are generated through the repeated addition of "patches" of image, is generally known as "Texture Bombing" and has several precedents, notably the work of Paul

Haeberli [11] in creating painterly surfaces, and owes something to Xu *et al* [12] patch-based texture synthesis and Efros and Freeman's [13] Image quilting. It is most similar to Dischler's [14] Texture Particles, but in our context is not intended for 3-D texture mapping, and used in a purely 2-D sense.

Results – A working Example

The techniques described were used to create a floor-based artwork (*The Kipple Pond*) originally commissioned by a UK Museum and shown in October 2009 and re-commissioned by a New York Museum in April 2010 (details withheld to maintain anonymity). The first version was 18 meters by 4 meters, and produced at a print resolution of 100 dpi, using Scotchprint 2000 Floor Graphic, in 12 sections. The NY version was 12 meters by 3 meters, in six sections using Scotchprint Pavement Graphics.

Upon encountering the piece, visitors were able to walk across the surface of a large and highly detailed millpond. At first glance the image seemed to be a huge natural photograph of a real pond, full of weeds, insects and fish. On closer inspection, the undulating bed of the millpond was seen to be composed of an accretion of thousands of items of man-made detritus or "kipple" as P.K. Dick called it [15]; jewels, coins, toys, tin cans, electrical components, bottles, screws, cutlery, pieces of machines and so on; the natural and the unnatural existing in close visual harmony, suggesting that nature might have a restorative effect over man's promiscuous outpourings. The resultant effect was poised between a single *Trompe-l'oeil* image, expansive decorative surface and a visual puzzle.



Fig 6. The artwork in situ in the UK 2009



Fig 7. The artwork under the harshest of scrutiny

Conclusions and Future work

Our creative methodology works well as an expressive medium for the artist by allowing the creation of an infinite variety of fully detailed, non-repeating, large-scale artworks. The process enjoys a good degree of fluid interactivity and visual feedback, and is reasonably amenable to continual edits and updates to content. There are, of course, repeats of elements at certain levels within the image; the duckweed leaves repeat at the small scale, and the large patches may be used several times each within the overall image. But at every level the repeat is concealed. The large image patches all intermingle in different ways, making each part of the surface unique, and the viewer is unconcerned about repeated elements (such as the Duckweed leaves) at the lowest level of detail.

Because the actual “final” image is never realised as a whole within the computer’s memory, but rendered and saved out as uniformly sized print-ready files (each one a manageable size for most computers/printers) the potential size of the complete output image is unlimited. Indeed, one could imagine a system where the key-image is replaced, functionally, by an algorithm generating the necessary descriptions of patches to render a print-tile (and its abutting neighbours) out of an arbitrarily large surface. Such a system would potentially be able to churn out limitlessly big images. One would however have to ensure that the algorithm produced something suitably varied and interesting over the actual extents of the final image, and this may not be a trivial undertaking. Indeed such an algorithm might in turn rely on a worked up key-image to generate a statistical model of the distribution of elements. Such a system would go some way towards the notion of generating endless, non-repeating surface designs that might be used instead of traditional wallpaper or floor covering.

A “deep zoom” version of *I Kipple Pond* can be found online at www.simonschofield.net, along with other generative artworks.

References

- [1] Cline D et al. (1997). *Dictionary of Visual Science*. 4th ed. Butterworth-Heinemann, Boston
- [2] Knoll, T. et al. (2010). *Adobe Photoshop CS5* [computer software]. CA: Adobe Inc.

- [3] GIMP.org (2001) *GNU Image Manipulation Program 2.7.4* [computer software]. Open Source Gimp.org
- [4] Google Inc. (2009). *Google Earth* (Version 5.1.3533.1731) [Software]. CA: Google Inc.
- [5] Gigapan Org (2011) *zoomable image of Shanghai* (website) available at <http://www.gigapan.org/gigapans/66626/> accessed 2012
- [6] Wikipedia authors (2011) *Largest photographs in the world*, Accessed 2012. http://en.wikipedia.org/wiki/Largest_photo
- [7] Python.org (2011) *Python programming environment v 3.2.2* [Computer software]. Open source. Python.com
- [8] Porter, T. and Duff, T (1984). *Compositing Digital Images*. Computer Graphics 18 (3): 253–259. [doi:10.1145/800031.808606](https://doi.org/10.1145/800031.808606). [ISBN 0897911385](https://www.amazon.com/dp/0897911385).
- [9] Lua.org (2011) *Lua embedded programming language 5.2* [computer software]. Brazil: Lablua, www.lua.org/about.html accessed Jan 2012
- [10] HSC Software (1998) *Live Picture* [computer software]. CA: HSC Software
- [11] Haeberli, P. (1990, August) *Painting by Numbers: Abstract Image Representation*. Computer Graphics Vol. 24, ACM Press (pp207-214).
- [12] [Xu, Guo and Shum, \(2000\)](#) *Patch-based texture synthesis*. Microsoft technical paper, Microsoft, SA
- [13] Alexei A. Efros and William T. Freeman (2001) *Image Quilting for Texture Synthesis and Transfer*, Proceedings of SIGGRAPH '01, Los Angeles, California, August, 2001.
- [14] Jean-Michel Dichler, et al (2002) *Texture Particles* , Eurographics conference proceedings
- [15] Dick, Philip K, (1968) *Do Androids Dream of Electric Sheep?* Novel

All images are fully owned by the author.

Footnote

To calculate the size of the smallest object that can be perceived at a distance of 1.5m: The smallest angle subtended by a 'normal' eye is 1 minute of arc (or 0.01667 degrees). This calculates an object of around 0.25 per mm, the visual equivalent of around 100 dpi.