

# Fast Poisson Blending using Multi-Splines

Richard Szeliski, Matt Uyttendaele, and Drew Steedly

Microsoft Research

April 2008

Technical Report

MSR-TR-2008-58

We present a technique for fast Poisson blending and gradient domain compositing. Instead of using a single (piecewise-smooth) offset map to perform the blending, we associate a separate map with each input source image. Each individual offset map is itself smoothly varying and can therefore be represented using a low-dimensional spline. The resulting linear system is much smaller than either the original Poisson system or the quadtree spline approximation of a single (unified) offset map. We demonstrate the speed and memory improvements available with our system and apply it to large panoramas. We also show how robustly modeling the multiplicative gain rather than the offset between overlapping images leads to improved results.

Microsoft Research

Microsoft Corporation

One Microsoft Way

Redmond, WA 98052

<http://www.research.microsoft.com>

# 1 Introduction

While Poisson blending (Pérez *et al.* 2003) (also known as *gradient domain compositing*) was originally developed to support the seamless insertion of objects from one image into another (Pérez *et al.* 2003, Jia *et al.* 2006), it has found widespread use in hiding seams due to exposure differences in image stitching applications (Levin *et al.* 2004b, Agarwala *et al.* 2004, Agarwala 2007, Hays and Efros 2007). Poisson blending usually produces good results for these applications. Unfortunately, the resulting two-dimensional optimization problems require large amounts of memory and time to solve.

A number of approaches have been proposed in the past to speed up the solution of the resulting sparse system of equations. One approach is to use multigrid (Levin *et al.* 2004a) or multi-level preconditioners (Szeliski 2006) (which can be implemented on GPUs (Bolz *et al.* 2003, Buatois *et al.* 2007)) to reduce the number of iterations required to solve the system. More recently, Agarwala (2007) proposed using a quadtree representation of the offset field, which is the difference between the original unblended images and the final blended result. Another possibility is to solve the Poisson blending problem at a lower resolution, and to then upsample the resolution while taking the location of seams into account (Kopf *et al.* 2007a).

In this paper, we propose an alternative approach, which further reduces the number of variables involved in the system. Instead of using a single offset field as in (Agarwala 2007), we associate a separate low-resolution offset field with each source image. We then simultaneously optimize over all of the (coupled) offset field parameters. Because each of the offset fields is represented using a low-dimensional spline, we call the resulting representation a *multi-spline*.

The basis of our approach is the observation the offset field between the original unblended solution and the final blended result is piecewise smooth except at the seams between source regions (Pérez *et al.* 2003, Agarwala 2007). In his paper on efficient gradient-domain compositing, Agarwala (2007) exploits this property to represent the offset field using a quadtree. In this paper, we observe that if the offset field is partitioned into separate per-source correction fields, each of these will be truly smooth rather than just piecewise smooth. (See Appendix A for a simple proof under some commonly occurring conditions.)




Figure 1: *One dimensional examples of Poisson blending and offset maps: (a) the original Poisson blend of two source images  $u_i^1$  and  $u_i^2$  produces the blended function  $f_i$ ; (b) the offset image  $h_i$  is fitted to zero gradients everywhere except at the source image discontinuity, where it jumps by an amount equal to the average difference across the region boundary; (c) the multiple offset images  $h_i^1$  and  $h_i^2$ , each of which is smooth, along with the inter-image constraint at the boundary; the offsets are defined by the spline control vertices  $c_k^1$  and  $c_k^2$ .*

Our method is thus related to previous work in exposure and vignetting compensation (Eden *et al.* 2006, Goldman and Chen 2005, Kopf *et al.* 2007a), as it computes a per-image correction that reduces visible seams at region boundaries. Motivated by this observation, we investigate the use of multiplicative rather than additive corrections and show that these generally produce better results for image stitching applications.

The remainder of this paper is structured as follows. To begin, we formulate the Poisson blending problem and show how it can be reformulated as the computation of a piecewise-smooth offset field (Section 2). We then introduce the concept of multiple offset maps (Section 3) and show how these can be represented using tensor product splines (Section 4). In Section 5, we discuss efficient methods for solving the resulting sparse set of linear equations. In Section 6, we apply our technique to a variety of large-scale image stitching problems, demonstrating both the speed and memory improvements available with our technique, as well as the quality improvements available from using multiplicative (gain) compensation. We close with a discussion of the results and ideas for possible extensions.

## 2 Problem formulation

The Poisson blending problem can be written in discrete form as

$$E_1 = \sum_{i,j} s_{i,j}^x [f_{i+1,j} - f_{i,j} - g_{i,j}^x]^2 + s_{i,j}^y [f_{i,j+1} - f_{i,j} - g_{i,j}^y]^2, \quad (1)$$

where  $f_{i,j}$  is the desired Poisson blended (result) image,  $g_{i,j}^x$  and  $g_{i,j}^y$  are the target gradient values, and  $s_{i,j}^x$  and  $s_{i,j}^y$  are the (potentially per-pixel) gradient constraint (smoothness) weights. (This notation is adapted from (Szeliski 2006).)

In the original formulation (Pérez *et al.* 2003), the weights are all set uniformly, and the gradients are computed from the gradients of the source image being blended in, with additional hard constraints along the boundary of the cut-out region to match the enclosing image. In the general multi-image formulation of Poisson blending (Agarwala *et al.* 2004), the gradients are obtained from the gradients of whichever image is being composited inside a given region,

$$g_{i,j}^x = u_{i+1,j}^{l_{i,j}} - u_{i,j}^{l_{i,j}}, \quad (2)$$

$$g_{i,j}^y = u_{i,j+1}^{l_{i,j}} - u_{i,j}^{l_{i,j}}, \quad (3)$$

where  $\{u^1 \dots u^L\}$  are the original *unblended* (source) images and  $l_{i,j}$  is the *label* (indicator variable) for each pixel, which indicates which image is being composited. At the boundaries between regions, the average of the gradients from the two adjacent images is used,

$$g_{i,j}^x = (u_{i+1,j}^{l_{i,j}} - u_{i,j}^{l_{i,j}} + u_{i+1,j}^{l_{i+1,j}} - u_{i,j}^{l_{i+1,j}})/2, \quad (4)$$

$$g_{i,j}^y = (u_{i,j+1}^{l_{i,j}} - u_{i,j}^{l_{i,j}} + u_{i,j+1}^{l_{i,j+1}} - u_{i,j}^{l_{i,j+1}})/2. \quad (5)$$

Note how these equations reduce to the previous case (2) and (3) on the interior, since the indicator variables are the same. We then substitute (2–5) into (1) and minimize the resulting cost function. The resulting function  $f$  reproduces the high-frequency variations in the input images while feathering away low-frequency intensity offsets at the seam boundaries (Figure 1a).

The per-pixel weights can be tweaked to allow the final image to match the original image with less fidelity around strong edges (Agarwala *et al.* 2004), where the eye is less sensitive to variations, resulting in what is sometimes called the *weak membrane* (Szeliski 2006). (This can

also be used to perform dynamic range compression (Fattal *et al.* 2002), which can be advantageous when compositing images with widely different exposures.) In this paper, as in (Agarwala *et al.* 2004, Agarwala 2007), we set the weights to be constant inside each source region, but optionally allow them to be weaker along high gradient seam boundaries,

$$s_{i,j}^x = \frac{1}{1 + (g_{i,j}^x/a)^2}. \quad (6)$$

In our current implementation, we use  $a = 5$  (when using 8-bit pixel values).

If only gradient constraints are used (1), the gradient-domain reconstruction problem is under-constrained. Pérez *et al.* (2003) use hard constraints along the region boundary, while Agarwala *et al.* (2004) have the user specify a single pixel value to match. In our work, we add a weak constraint towards the colors in the original image  $u_{i,j}^{l_{i,j}}$ ,

$$E_0 = \sum_{i,j} w_{i,j} [f_{i,j} - u_{i,j}^{l_{i,j}}]^2, \quad (7)$$

typically with  $w_{i,j} = 10^{-4}$ , which reduces unwanted low-frequency variations in the result and helps ensure that the final composite does not get too light or dark.

## 2.1 Offset formulation

As noted in (Pérez *et al.* 2003, Eqn.(5)) and (Agarwala *et al.* 2004, Eqn.(3)), we can replace the solution  $\{f_{i,j}\}$  with an *offset* from the original (unblended) image,

$$f_{i,j} = u_{i,j}^{l_{i,j}} + h_{i,j} \quad (8)$$

and solve for the offset image  $\{h_{i,j}\}$  instead. The new criterion being minimized becomes

$$E_2 = \sum_{i,j} s_{i,j}^x [h_{i+1,j} - h_{i,j} - \tilde{g}_{i,j}^x]^2 + s_{i,j}^y [h_{i,j+1} - h_{i,j} - \tilde{g}_{i,j}^y]^2 + w_{i,j} [h_{i,j}]^2, \quad (9)$$

where the modified gradients  $\tilde{g}_{i,j}^x$  and  $\tilde{g}_{i,j}^y$  are zero away from the boundaries and

$$\tilde{g}_{i,j}^x = (u_{i,j}^{l_{i,j}} - u_{i,j}^{l_{i+1,j}} + u_{i+1,j}^{l_{i,j}} - u_{i+1,j}^{l_{i+1,j}})/2, \quad (10)$$

$$\tilde{g}_{i,j}^y = (u_{i,j}^{l_{i,j}} - u_{i,j}^{l_{i,j+1}} + u_{i,j+1}^{l_{i,j}} - u_{i,j+1}^{l_{i,j+1}})/2, \quad (11)$$

at the boundaries between regions. (This can be verified by substituting (8) into (1) and (4–5).)

This new problem has a natural interpretation: the offset value should be everywhere smooth, except at the region boundaries, where it should jump by an amount equal to the (negative) *average difference* in intensity between the overlapping source images. The resulting offset function is piecewise smooth (Figure 1b), which makes it amenable to being represented by a quadtree spline, with smaller grid cells closer to the region boundaries (Agarwala 2007).

### 3 Multiple offset maps

In this paper, instead of using a single offset map, as suggested in (Pérez *et al.* 2003, Agarwala 2007), we use a different offset map for each source image, i.e.,

$$f_{i,j} = u_{i,j}^{l_{i,j}} + h_{i,j}^{l_{i,j}}, \quad (12)$$

where the  $\{h^1 \dots h^l\}$  are now the *per-source image offset maps* (see Figure 1c).

The optimization problem (9) now becomes

$$E_3 = \sum_{i,j} s_{i,j}^x [h_{i+1,j}^{l_{i+1,j}} - h_{i,j}^{l_{i,j}} - \tilde{g}_{i,j}^x]^2 + \quad (13)$$

$$s_{i,j}^y [h_{i,j+1}^{l_{i,j+1}} - h_{i,j}^{l_{i,j}} - \tilde{g}_{i,j}^y]^2 + w_{i,j} [h_{i,j}^{l_{i,j}}]^2,$$

Notice that in this problem, whenever two adjacent pixels, say  $(i, j)$  and  $(i + 1, j)$  come from the same source and hence share the same offset map, the gradient  $\tilde{g}_{i,j}^x$  is 0, and so the function is encouraged to be smooth. When two adjacent pixels come from different regions, the *difference* between their offset values is constrained to be the average difference in source values at the two pixels (10). This is illustrated schematically in Figure 1c.

What is the advantage of re-formulating the problem using a larger number of unknowns? There is none if we keep all of the  $h_{i,j}^l$  as independent variables.

However, under normal circumstances, e.g., when working with log intensities and multiplicative exposure differences, each of the individual per-source offset maps will be *smooth*, and not just *piecewise smooth* as in the case of a single offset map (see Appendix A for a proof). Therefore, each offset map can be represented at a much lower resolution, as we describe next.

## 4 Spline offset maps

To take advantage of the smoothness of each offset image, we represent each map with a tensor-product spline that covers the visible extent of each region, as shown in Figure 2. The choice of pixel spacing (subsampling)  $S$  is problem dependent, i.e., it depends on the amount of unmodeled variations in the scene and acquisition process, e.g., the severity of lens vignetting or the amount of inconsistent texture along the seam, but is largely independent of the actual pixel (sensor) resolution. We can either align each grid with each region’s bounding box (Figure 2a) or use a globally consistent alignment (Figure 2b). We use the latter, since it makes the nested dissection algorithm discussed in Section 5 easier to implement.

Once we have chosen  $S$  and the control grid locations, we can re-write each pixel in an individual offset map as the linear combination of the per-level spline control vertices  $c_{k,m}^l$  (Figure 1c),

$$h_{i,j}^l = \sum_{km} c_{k,m}^l B(i - kS, j - mS) \quad (14)$$

where

$$B(i, j) = b(i)b(j) \quad (15)$$

is a 2D tensor product spline basis and  $b(i)$  is a 1-D B-spline basis function (Farin 1996). (In image processing,  $B(i, j)$  is known as a separable interpolating impulse response function.) For example, when bilinear (first order) interpolation is used, as is the case in our experiments, the first order 1-D B-spline is the usual tent function, and each pixel is a linear blend of its 4 adjacent control vertices (Figure 2c).

The values of  $h_{i,j}^l$  in (14) can be substituted into (13) to obtain a new energy function (omitted for brevity) that only depends on the spline control variables  $c_{k,m}^l$ . This new energy function can be minimized as a sparse least squares system to compute a smooth spline-based approximation to the offset fields. Once the sparse least squares system has been solved, as described in Section 5, the per-pixel offset values can be computed using regular spline interpolation (14).

The actual inner loop of the least squares system setup simply involves iterating over all the pixels, pulling out the  $(K + 1)^d$  non-zero B-spline basis function values (where  $K$  is the order of




Figure 2: A spline grid is overlaid on top of each source region: (a) the grids are aligned with each region’s bounding box; (b) the grids are aligned with the final composite bounding box (shown in gray). The grid cell and control variables (nodes at cell corners) that are inactive are shown in pale colors. (c) Detail inside two of the spline patches: inside the left square (blue), each pixel depends on four neighboring control vertices; along a seam boundary (mixed colors), each pixel depends on eight.

the interpolant, and  $d$  is the dimensionality of the field), forming each of the linear equations in the control variables inside each squared term, and then updating the appropriate entries in the normal equations (stiffness matrix and right-hand side) (Szeliski and Coughlan 1997, Szeliski 2006).

Figure 2c illustrates this process for bilinear splines, where each offset value  $h_{i,j}^l$  (blue pixel inside the large blue square) depends on its four neighboring control vertices. Thus, each gradient constraint in (13) depends on either four or eight control vertices, the latter occurring only at seam boundaries where offsets from two different spline maps are being compared (i.e., the pair of black pixels inside the mixed color region). Note that spline vertices that do not influence any visible pixels can be eliminated from the variable set and are shown as pale colors in Figures 2a-b.

#### 4.1 Simplifying the constraints

Inside spline patches where all the pixels come from the same source (the blue patch in Figure 2c) and the smoothness and data weights are homogeneous,  $s_{i,j}^x = s_{i,j}^y = s$  and  $w_{i,j} = w$ , we can pre-compute the effect of all the individual per-pixel gradient and smoothness constraints ahead of



time. This is similar to the process of analytic integration performed during finite element analysis to determine the effects of continuous deformations on the discrete control variables (Terzopoulos 1983). Performing this analysis for a bilinear interpolant yields a *nine-point stencil*, which couples together vertices that are diagonally adjacent.

In our current implementation, we instead adopt a simpler approach and assume, for the purpose of computing the internal smoothness constraints, that the interpolating spline uses conforming triangular linear elements, as in (Terzopoulos 1983, Szeliski 2006); for the data constraint, we use a piecewise constant interpolant. This results in a simple per-layer energy function that is a coarsened version of the original fine-level Poisson blending energy,

$$E^l = \sum_{k,m} s[c_{k+1,m}^l - c_{k,m}^l]^2 + s[c_{k,m+1}^l - c_{k,m}^l]^2 + S^2 w[c_{k,m}^l]^2. \quad (16)$$

To further speed up the formulation of the least squares system, we apply this same discrete energy to *all* spline patches within each offset layer. We then add in the original gradient constraints (13) only along seam boundary pixels, i.e., pairs of pixels that have two different labels, as shown by the pair of black pixels in Figure 2c.

## 4.2 Multiplicative (gain) offsets

The observant reader may have noticed that multi-spline offset fields are just a generalization of the bias-gain correction fields commonly used to adjust the exposures of images before blending (Eden *et al.* 2006, Goldman and Chen 2005, Kopf *et al.* 2007a). A single linear spline patch per source image is equivalent to the commonly used affine intensity variation model. If full spline fields are used, one can view our new technique as simply constraining the spline correction fields to map overlapping pixels near boundary regions to the same value.

The important difference, however, between our derivation from the original Poisson blending equations and a more heuristic implementation is that our approach directly tells us how to set the tradeoff between the smoothness in the final correction field and the strength of the seam matching term. Without this balance, the resulting correction fields can result in composites with either visible tears or flat spots.



(a)

(b)

(c)

Figure 3: *Comparison of multiplicative gain vs. additive offset blending: (a) unblended image; (b) multiplicative gain blending; (c) additive offset blending. Note how the additive result has a visible contrast change across the seam.*

Viewing Poisson blending as the computation of per-image correction fields suggests that computing multiplicative gain fields may be preferable to computing additive offset fields. In fact, most visible seams in panoramas are due to camera exposure variation, vignetting, or illumination changes (the sun going behind a cloud), all of which are better modeled as multiplicative gains rather than additive offsets. This is true even if the images are gamma-corrected, since the resulting images are still related by a multiplicative factor, i.e.,  $I_2 = kI_1 \Rightarrow I_2^\gamma = k^\gamma I_1^\gamma$ .

To estimate a multiplicative gain, we simply take the logarithm of each input image before computing the seam difference, or, equivalently, take the logarithm of the ratio of overlapping seam pixels. (Dark pixels can be clamped to a minimum value such as 1.) The resulting computed offset field is then exponentiated and used as a multiplicative gain field.

Figure 3 compares the result using the multiplicative gain approach vs. the traditional additive Poisson blending approach. Because the additive offset does not model the differing amounts of contrast in the two source images (which are related by a multiplicative exposure difference), the blended result in Figure 3c has a visible contrast change in the vicinity of the seam (more muddy looking colors to the right of the seam, which are better seen in the supplementary materials).

### 4.3 Blending Gigapixel images

In order to apply our multi-spline based Poisson blending technique to Gigapixel images (Kopf *et al.* 2007b), we make one further approximation. Instead of computing the seam costs at the final Gigapixel resolution, we compute these costs at the same resolution as the graph cut optimization performed in (Kopf *et al.* 2007b), which is  $1/8$ th the horizontal and vertical resolution of the final panorama.

The decision to accumulate the seam costs on a lower resolution image is actually well-justified. Since the relative contribution of each seam constraint to the spline vertices is a slowly varying function, summing these contributions over a slightly coarser grid than the pixels (but still finer than the spline) does not affect the results very much. Because we are computing a least squares correction, summing up the least squares contributions over regions does not affect the final cost (except for the replacement of the spline weights with a slightly more discretized version).

Another way of looking at this is that if we are estimating the offset or gain adjustment between overlapping images, a similar results will be obtained if we look at lower-resolution versions of these images (so long as the lowered resolution is still significantly higher than the spline grid).

Once we have computed our multi-spline correction fields, these are then upsampled to the final Gigapixel resolution during the tile-based final compositing process (Kopf *et al.* 2007b). An alternative could be to use joint bilateral upsampling (Kopf *et al.* 2007a) applied to a single offset field at the lower resolution (e.g., using (Agarwala 2007)), but this would be more complicated and error-prone.

## 5 Solving the system

A variety of techniques can be used to solve the small sparse positive definite system of equations arising from the multi-spline correction fields. For large sparse system, iterative techniques (potentially enhanced with multi-grid or multi-level preconditioning) can be used (Saad 2003). When the systems are smaller, direct factorization techniques such as Cholesky decomposition are more efficient (Davis 2006).

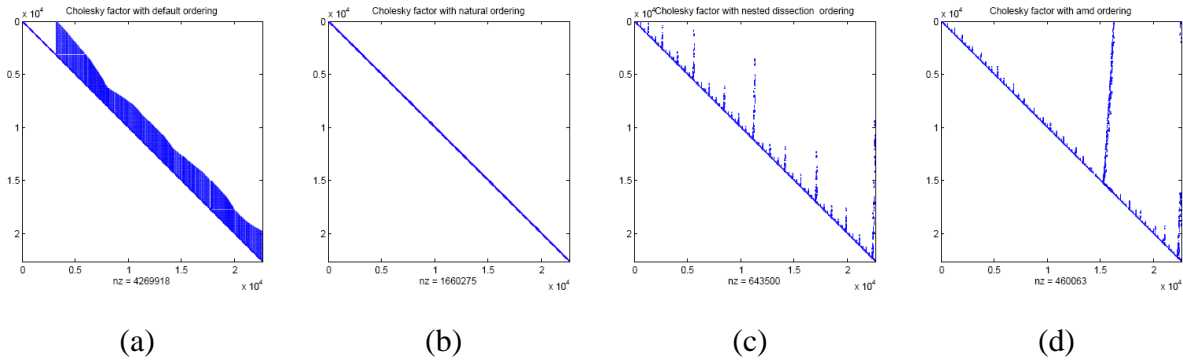


Figure 4: *Fill in patterns (Cholesky factor matrix) for four different variable orderings: (a) default, (b) raster, (c) nested dissection, and (d) AMD. Please see the supplementary materials for full images.*

Because his sparse systems are larger, Agarwala (2007) uses the conjugate gradient sparse iterative solver, which partially accounts for his longer run times. Because our multi-spline systems are much smaller, we use direct techniques. The efficiency of these techniques depends on the amount of *fill-in* during the factorization process, which can be reduced by an appropriate reordering of the variables (Davis 2006).

For two-dimensional grid problems, *nested dissection*, which recursively splits the problem along small length rows or columns, results in good performance, i.e.,  $O(n \log n)$  space and  $O(n^{3/2})$  time (or better for asymmetrically shaped domains), where  $n$  is the number of variables. In order for this technique to work, we need to ensure that all the spline variables line up in the same rows and columns, which is why we use the aligned spline grid shown in Figure 2b.

In our experimental results section, we compare nested dissection with three other techniques, namely the default ordering our system uses, which is to raster order all variables within each layer, the joint raster ordering of all variables on a grid (with zero or more variables at each grid location), and the more complex *approximate minimum degree* (AMD) algorithm described in (Davis 2006). Figure 4 shows the Cholesky factor patterns for the four different strategies.

Grid size $S$	RMS error	max error
8	0.0886	11.20
16	0.1039	12.80
32	0.1841	13.70
64	0.2990	14.40
128	0.4118	13.90

Table 1: *RMS and maximum error comparisons to the ground truth Poisson blend for different grid sizes  $S$ . For these experiments, we used the 9.7 MPixel St. Emilion dataset shown in Figure 5.*




Figure 5: Fast Poisson blending using multi-splines: (a) unblended composite; (b) piecewise smooth multi-spline offset image; (c) final blended composite.

## 6 Experiments

In order to validate our approach and to see how much of a speedup could be expected, we first obtained the four large panoramas shown in Figure 6 from the author of (Agarwala 2007). For these images, we used an additive offset field to match the results presented in (Agarwala 2007) as closely as possible. We also used a spline spacing of  $S = 64$  and bilinear splines.

The results of running our algorithm on these four data sets are shown in Table 2. As you can see, our multi-spline technique is about 5–10 $\times$  faster and requires about 10 $\times$  less memory than the quadtree-based approach developed in (Agarwala 2007). The two techniques produce results of comparable visual quality, as can be seen by inspecting the large full-size images provided in the supplementary material.

The timings also reveal that because of fill-in, variable reordering strategies (Davis 2006) can significantly affect both the memory requirements and the run time of the direct solver. Taking the time to implement a simple reordering strategy like nested dissection (which involves a known



Figure 6: *Thumbnails of the four large panorama test images: (a) Sedona, (b) Edinburgh, (c) Crag, (d) RedRock.*

fixed pattern) is therefore worthwhile. While with our current large decimation factor ( $S = 64$ ), system solving takes less than 10% of the time (even with the worst reordering), changing  $S$  to a smaller value of say  $S = 16$  (which may be preferable for smaller images) can increase this time by a factor of  $16\times$ .

Table 1 shows how the RMS (root mean square) and maximum error (in gray levels) in the solution depend on the grid size  $S$ . For these experiments, we used the solution to the *St. Emilion* data set provided by Agarwala (2007) (Figure 5) as our ground truth. We then ran our fast multi-spline-based solver using a variety of grid sizes,  $S = \{8, 16, \dots, 128\}$  and computed both the RMS and maximum error between the offset field we computed and the full solution. As you can see, the RMS error grows steadily with the grid size, while the maximum error does not vary that much. Visual inspection of the full-resolution results (which are available as part of the supplementary materials) shows that the maximum error is concentrated at isolated pixels along seam boundaries where highly textured regions are mis-aligned. Fortunately, these “errors” are masked by the textures themselves, so that the final blended images appear of identical quality to the eye.

Next, we applied our technique to the Seattle Skyline image shown in Figure 4 of (Kopf *et al.* 2007b), using the multiplicative gain (log intensity) formulation because of the large exposure differences. In this case, because the seam costs were computed on a  $1/8$ th (on side) resolution image, the seam cost evaluation (shown as Setup in Table 2) and system solving times as well as the

Dataset	#	Mpix	Quadtree			Multi-spline			Setup		Default		Raster		ND		AMD		Render
			V (%)	T (s)	M	V (%)	T (s)	M	T (s)	M	T (s)	M	T (s)	M	T (s)	M	T (s)	M	T (s)
Sedona	6	34.6	0.47	29	52	0.0271	6.12	4	3.33	3	0.29	22	0.13	8	0.28	5	0.09	4	2.70
Edinburgh	25	39.7	1.15	122	123	0.0315	7.99	10	3.99	10	0.39	26	0.23	11	0.41	7	0.17	5	3.84
Crag	7	62.7	0.47	78	96	0.0271	11.18	7	6.16	6	0.83	53	0.34	22	0.54	9	0.20	7	4.82
RedRock	9	83.7	0.46	118	112	0.0270	14.85	10	8.11	8	1.15	71	0.50	29	0.75	13	0.32	10	6.42
Seattle	650	3186.9				0.0009	7.33+	57	6.35	57	2.04	114	1.50	73	1.42	34	0.98	22	56m

Table 2: *Performance of the quadtree vs. multi-spline based solver. The first three columns list the dataset name, the number of source images, and the number of pixels. The next three columns show the results for the quadtree-based acceleration, including the ratio of variables to original pixels (as a percentage), the total run time (in seconds), and the memory usage (in MBytes). The next three columns show the corresponding results for our multi-spline based approach (for all of our experiments,  $S = 64$ ). As you can see, our results are roughly  $5 - 10\times$  faster and smaller. The final sets of columns break down the time and memory requirements of the three multi-spline blending stages, namely the setup (computation of seam boundary constraints), the direct solution using the default, raster, nested dissection, and AMD reordering strategies, and the final rendering (compensation) stage. Note that the Gigapixel Seattle total time does not include the i/o bound rendering stage, which took 56 minutes to produce the final image tile set.*

memory requirement are comparable to that of the 84 Mpixel RedRock panorama. The rendering time required to read back and warp the source images, apply the spline-based correction, and write out the resulting tiles is significantly longer, and is a component we plan to speed up.

A cropped portion of our result is shown in Figure 5, and the unblended, offset, and blended images at the  $1/8$ th working resolution, along with some cropped portions of the final Gigapixel image are shown in the supplementary materials.

The most visible artifacts in these results, besides the saturated regions and gross misalignments caused by the moving crane, are the occasional seams visible in the sky regions near dark buildings, which are due to some of the original source images having saturated pixels (usually in the blue channel) in these regions. Unfortunately, since the values at these pixels do not reflect the true irradiance, the multiplicative gain computed in areas that border unsaturated pixels is inconsistent, and cannot simultaneously hide both kinds of seams.

## 7 Discussion and extensions

Unfortunately, many cameras do not have a pure linear or gamma response function. Because of this, multiplicative gain cannot always account for all of the exposure differences between overlapping regions. This can be seen by closer inspection of the Seattle skyline panorama (see the supplementary material), where some visible seams remain in high-contrast areas such as sunny buildings. First linearizing the camera response, as was done in (Kopf *et al.* 2007b), is the principled way to solve this problem, and we plan to investigate this. Another possibility would be to model exposure differences as an affine (multiplicative plus additive) difference.

As we can see from our experiments, the biggest difference between our multi-spline approach and full Poisson blending (and its quadtree approximation) is that we enforce piecewise smoothness in both the  $x$  and  $y$  dimensions, whereas Poisson blending can tolerate irregular offsets *along* the seam. While our approach can occasionally lead to artifacts, e.g., in images that are not log-linear, Poisson blending can introduce different artifacts, such as “ruffles” that sometime propagate away from seam boundaries when they disagree. Ultimately, improving the alignment between



images using an optic-flow deghosting technique (Shum and Szeliski 2000) followed by robustly estimating the mapping between overlapping images will probably produce the best results.

In terms of computational complexity, as the resolution of photographs continues to increase, our multi-spline based approach has better scaling properties than the quadtree based approach. Because the number of spline control vertices depends on the smoothness of the unmodeled inter-exposure variation and not the pixel density, we expect it to remain fixed. In the quadtree-based approach, the number of variables increases linearly with the on-side (as opposed to pixel count) resolution. Hybrid techniques, which coarsen the spline grid away from the boundaries, or conversely, work first on lower-resolution images, might therefore be worth investigating.

## 8 Conclusions

In this paper, we have developed a new approach to gradient domain compositing that allocates a separate smoothly varying spline correction field for each source image. We also investigated the benefits of using a multiplicative gain formulation over the more traditional additive offset formulation. Using our approach, we obtain linear systems an order of magnitude smaller than those obtained with a unified quadtree representation of the offset map, while producing results of comparable visual quality. We also suggest areas for further investigations into better quality algorithms for seam blending.

## References

- Agarwala, A. (2007). Efficient gradient-domain compositing using quadtrees. *ACM Transactions on Graphics*, 28(3).
- Agarwala, A., Dontcheva, M., Agrawala, M., Drucker, S., Colburn, A., Curless, B., Salesin, D. H., and Cohen, M. F. (2004). Interactive digital photomontage. *ACM Transactions on Graphics*, 23(3), 292–300.

- Bolz, J., Farmer, I., Grinspun, E., and Schröder, P. (2003). Sparse matrix solvers on the GPU: Conjugate gradients and multigrid. *ACM Transactions on Graphics*, 22(3), 917–924.
- Buatois, L., Caumon, G., and Lévy, B. (2007). Concurrent number cruncher: An efficient sparse linear solver on the gpu. In *High Performance Computing Conference (HPCC 2007)*, pages 358–371, Springer-Verlag.
- Davis, T. A. (2006). *Direct Methods for Sparse Linear Systems*. SIAM.
- Eden, A., Uyttendaele, M., and Szeliski, R. (2006). Seamless image stitching of scenes with large motions and exposure differences. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'2006)*, pages 2498–2505, New York, NY.
- Farin, G. E. (1996). *Curves and Surfaces for Computer Aided Geometric Design: A Practical Guide*. Academic Press, Boston, Massachusetts, 4th edition.
- Fattal, R., Lischinski, D., and Werman, M. (2002). Gradient domain high dynamic range compression. *ACM Transactions on Graphics (TOG)*, 21(3), 249–256.
- Goldman, D. B. and Chen, J.-H. (2005). Vignette and exposure calibration and compensation. In *Tenth International Conference on Computer Vision (ICCV 2005)*, pages 899–906, Beijing, China.
- Hays, J. and Efros, A. A. (2007). Scene completion using millions of photographs. *ACM Transactions on Graphics*, 28(3).
- Jia, J., Sun, J., Tang, C.-K., and Shum, H.-Y. (2006). Drag-and-drop pasting. *ACM Transactions on Graphics*, 27(3), 631–636.
- Kopf, J., Cohen, M. F., Lischinski, D., and Uyttendaele, M. (2007a). Joint bilateral upsampling. *ACM Transactions on Graphics*, 28(3).
- Kopf, J., Uyttendaele, M., Deussen, O., and Cohen, M. F. (2007b). Capturing and viewing gigapixel images. *ACM Transactions on Graphics*, 28(3).

- Levin, A., Lischinski, D., and Weiss, Y. (2004a). Colorization using optimization. *ACM Transactions on Graphics*, 23(3), 689–694.
- Levin, A., Zomet, A., Peleg, S., and Weiss, Y. (2004b). Seamless image stitching in the gradient domain. In *Eighth European Conference on Computer Vision (ECCV 2004)*, pages 377–389, Springer-Verlag, Prague.
- Pérez, P., Gangnet, M., and Blake, A. (2003). Poisson image editing. *ACM Transactions on Graphics*, 22(3), 313–318.
- Rother, C., Bordeaux, L., Hamadi, Y., and Blake, A. (2006). Autocollage. *ACM Transactions on Graphics*, 27(3), 847–852.
- Saad, Y. (2003). *Iterative Methods for Sparse Linear Systems*. SIAM, second edition.
- Shum, H.-Y. and Szeliski, R. (2000). Construction of panoramic mosaics with global and local alignment. *International Journal of Computer Vision*, 36(2), 101–130. Erratum published July 2002, 48(2):151-152.
- Szeliski, R. (2006). Locally adapted hierarchical basis preconditioning. *ACM Transactions on Graphics*, 25(3), 1135–1143.
- Szeliski, R. and Coughlan, J. (1997). Spline-based image registration. *International Journal of Computer Vision*, 22(3), 199–218.
- Terzopoulos, D. (1983). Multilevel computational processes for visual surface reconstruction. *Computer Vision, Graphics, and Image Processing*, 24, 52–96.

## A Smooth per-region offset variation

The observation that each individual offset image  $h_{i,j}^l$  is typically smooth (rather than piecewise smooth) can be stated more formally in the following lemma.

**Lemma 1:** Assume that the source images are the result of taking (displaced) photographs of the same scene under smoothly varying illumination or exposure changes. In that case, the per-source offset maps are themselves smoothly varying, regardless of the shape of the boundaries or the variations in the original scene irradiance (complexity of the input images).

**Proof:** Let us denote the original source (unblended) images as

$$u_{i,j}^l = r_{i,j} + v_{i,j}^l, \quad (17)$$

where  $r_{i,j}$  is the scene irradiance and  $v_{i,j}^l$  are the per-image smooth intensity variations. (These conditions hold when we are operating in log-linear space and the smooth intensity variations are multiplicative, e.g., exposure, low-frequency lighting, and vignetting.)

We need to define a measure of smoothness, so we use Equation (13) evaluated at all the pixels *except* boundary pixels, i.e., we measure the deviation from zero gradient and offsets at all pixels interior to the source regions.

As an initial condition, set  $h_{i,j}^l = v_{i,j}^l$ . Under this condition, the terms in the full energy (13) corresponding to boundary pixels are exactly 0, since the offset functions  $h_{i,j}^l$  exactly compensate for the smooth intensity variations  $v_{i,j}^l$  added to the original scene irradiance  $r_{i,j}$ . Therefore, the particular solution  $h_{i,j}^l = v_{i,j}^l$  has the exact same smoothness energy as the intensity variations  $v_{i,j}^l$ .

Relaxing the system further by minimizing (13) over all possible  $h_{i,j}^l$  values will result in an even smoother solution (but may not match the actual smooth intensity variation  $v_{i,j}^l$ ). Therefore, the original Lemma that “the per-source offset maps are themselves smoothly varying” holds. **end-of-proof**

Note that the Lemma does not say that the true smooth offsets will be recovered by our algorithm, just that the recovered per-source offsets will themselves be at least as smooth as the true offsets. For example, if the exposure values are strongly varying across a large panorama, our technique will estimate a sawtooth offset function (for a sufficiently large  $w_{i,j}$ ), rather than a staircase function. Similarly, our technique may not correctly compensate for vignetting, so a more specialized technique may be preferable (Goldman and Chen 2005). Note also that it is always possible to construct images for which the offset maps will not actually be smooth, e.g., by placing seams between regions that actually do not match (Rother *et al.* 2006).