

Locally Adapted Hierarchical Basis Preconditioning

Richard Szeliski*
Microsoft Research



Figure 1: Piecewise smooth color interpolation in a colorization application: (a) input gray image with color strokes overlaid; (b) solution after 20 iterations of conjugate gradient; (c) using 1 iteration of hierarchical basis function preconditioning; (d) using 1 iteration of locally adapted hierarchical basis functions.

Abstract

This paper develops locally adapted hierarchical basis functions for effectively preconditioning large optimization problems that arise in computer graphics applications such as tone mapping, gradient-domain blending, colorization, and scattered data interpolation. By looking at the local structure of the coefficient matrix and performing a recursive set of variable eliminations, combined with a simplification of the resulting coarse level problems, we obtain bases better suited for problems with inhomogeneous (spatially varying) data, smoothness, and boundary constraints. Our approach removes the need to heuristically adjust the optimal number of preconditioning levels, significantly outperforms previously proposed approaches, and also maps cleanly onto data-parallel architectures such as modern GPUs.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Display algorithms; G.1.8 [Numerical Analysis]: Partial Differential Equations—Multigrid and multilevel methods

Keywords: Computational photography, Poisson blending, colorization, fast PDE solution, multilevel techniques, parallel algorithms, GPU acceleration

1 Introduction

A large number of computer graphics problems can be formulated as the solution of a set of spatially-varying linear or non-linear partial differential equations (PDEs). Such problems include computational photography algorithms such as high-dynamic range tone mapping [Fattal et al. 2002], Poisson and gradient-domain blending [Pérez et al. 2003; Levin et al. 2004b], and colorization [Levin et al. 2004a]. They also include various physically-based modeling problems such as elastic surface dynamics [Terzopoulos and

Witkin 1988], as well as the regularized solution of computer vision problems such as surface reconstruction, optic flow, and shape from shading [Terzopoulos 1986].

The discretization of these continuous variational problems using finite element or finite difference techniques yields large sparse systems of linear equations. Direct methods for sparse systems [Duff et al. 1986] become inefficient for large multi-dimensional problems because of excessive amounts of *fill-in* that occurs during the solution. In most cases, iterative solution are a better choice [Saad 2003]. Unfortunately, as the problems become larger, the performance of the iterative algorithms usually degrades due to the increased condition number of the associated systems.

Two different classes of techniques, both based on multi-level (hierarchical) representations, have traditionally been used to accelerate the convergence of iterative techniques [Saad 2003]. The first are multigrid techniques [Briggs et al. 2000], which interpolate between different levels of resolution in an effort to alternately reduce the low- and high-frequency components of the error. Unfortunately, these techniques work best for smoothly varying (homogeneous) problems, which is often not the case for computer graphics applications.

The second class of techniques are optimization algorithms such as conjugate gradient preconditioned using a variety of techniques [Saad 2003]. Among these techniques, multi-level preconditioners such as hierarchical basis functions and wavelets [Yserentant 1986; Szeliski 1990; Pentland 1994; Gortler and Cohen 1995] are well suited for graphics applications because they exploit the natural multi-scale nature of many visual problems.

While multi-level preconditioners have proven to be quite effective at accelerating the solution of such tasks, some problems remain. First, the choice of basis functions and the number of levels is problem-dependent [Szeliski 1990]. Second, these algorithms perform poorly on problems with large amounts of inhomogeneity, such as local discontinuities or irregularly spaced data.

In this paper, we go back to basics and ask the question: is there an *optimal* set of hierarchical basis functions that can perfectly precondition an arbitrary multi-dimensional variational problem? It turns out that for a one-dimensional first order system, we can construct such a multi-resolution basis. For higher-dimensional problems, we can use a similar methodology to derive locally adapted hierarchical

*e-mail: szeliski@microsoft.com

basis functions that *approximate* an optimal preconditioning of arbitrary problems and significantly outperform previously published techniques (Figure 1). The key to our approach is to carefully examine the local structure of the coefficient matrix to perform a recursive set of variable eliminations combined with a simplification of the resulting coarse level problems.

We begin this paper with a review of variational problems and their discretization (Section 2). In Section 3, we review previously developed techniques, both direct and iterative, for solving such linear systems, including the use of multigrid and multi-level preconditioners. In Section 4, we derive optimal hierarchical bases for the one-dimensional first order problem. In Section 5, we develop a methodology for constructing locally adaptive hierarchical bases for two-dimensional first order problems such as Poisson image blending, HDR tone mapping, and colorization. We also discuss how to map these algorithms onto GPUs. Section 6 presents some experimental results, including comparisons with previous approaches, while Section 7 discusses areas for further research.

2 Problem formulation

Many problems in computer graphics can be formulated using a *variational approach*, which involves defining a continuous two-dimensional optimization problem, adding appropriate *regularization terms*, and then discretizing the problem on a regular grid [Gortler and Cohen 1995; Zhang et al. 2002; Fattal et al. 2002; Pérez et al. 2003; Levin et al. 2004a; Levin et al. 2004b]. The discretization of these variational formulations leads to quadratic forms with large sparse symmetric positive definite (SPD) coefficient (a.k.a. Hessian or stiffness) matrices that have a natural embedding in a two-dimensional grid and are hence amenable to multi-level approaches.

1D problems The simplest version of this problem is one-dimensional curve interpolation, where the problem is to reconstruct a function $f(x)$ that approximately interpolates a set of data points $\{d_k\}$ at locations $\{x_k\}$ with associated weights $\{w_k\}$. To make the problem well-posed, we add a first order smoothness penalty (functional),

$$s = \int s(x) f_x^2(x) dx, \quad (1)$$

where $s(x)$ is the spatially varying *smoothness* weight, which can be used to locally control discontinuities. Second order smoothness functions can also be used, but we leave their discussion and solution to future work [Szeliski 2006].

To discretize the above problem, we choose a regular grid spacing of size h and place all of our data constraints at discrete grid locations, which results in the discrete energy

$$E_d = \sum_i w_i (f_i - d_i)^2, \quad (2)$$

with $w_i = 0$ where there are no data points. Figure 2 shows a sample set of input data values d_i and weights w_i , which we can think of points pulling on the final surface f_i with strong springs. (Weaker springs connect the remaining points to the 0 baseline.)

To discretize the smoothness functional, we use finite element analysis [Bathe and Wilson 1976], which models the function $f(x)$ as a piecewise linear spline [Terzopoulos 1983]. After analytically integrating the energy functional s , we obtain

$$E_s = \sum_i s_i (f_{i+1} - f_i)^2, \quad (3)$$

with the value for s_i given by $s_i = h^{-2} \int_{x_i}^{x_{i+1}} s(x) dx = s/h$, for a constant $s(x) = s$.




Figure 2: *Input data and output solution.* The input data values are $d_i = \{150, 50, 200\}$ at $i = \{3, 8, 15\}$ and 0 elsewhere. The data weights are $w_i = 100$ at $i = \{3, 8, 15\}$ and 1 elsewhere, which causes the interpolated curve f to sag towards 0. The smoothness function is $s_i = 10$ except for $s_{11} = 0$, which causes a “tear” in the reconstructed f curve between $i = \{11, 12\}$. (The tear is not readily apparent because of the linear interpolation used by the plotting software.)

2D problems Two-dimensional variational problems can be formulated in a similar manner. Again, we wish to reconstruct a function $f(x, y)$ through a set of data points $\{d_k\}$ at locations $\{(x_k, y_k)\}$ with associated weights $\{w_k\}$. The two-dimensional first order smoothness functional (*membrane* model [Terzopoulos 1983]) can be written as

$$s = \int s^x(x, y) f_x^2(x, y) + s^y(x, y) f_y^2(x, y) dx dy. \quad (4)$$

Discretizing this functional leads to a first order *smoothness* term,

$$E_s = \sum_{i,j} s_i^x (f_{i+1,j} - f_{i,j} - g_{i,j}^x)^2 + s_i^y (f_{i,j+1} - f_{i,j} - g_{i,j}^y)^2. \quad (5)$$

The data values $g_{i,j}^x$ and $g_{i,j}^y$ are gradient data terms (constraints) used by algorithms such as HDR tone mapping, Poisson blending, and gradient-domain blending [Fattal et al. 2002; Pérez et al. 2003; Levin et al. 2004b]. (They are 0 when just discretizing the conventional first order smoothness functional (4).) The smoothness weights s_i^x and s_i^y control the location of horizontal and vertical tears (or weaknesses) in the surface. The two dimensional discrete data energy is written as

$$E_d = \sum_{i,j} w_{i,j} (f_{i,j} - d_{i,j})^2. \quad (6)$$

To obtain the discrete energies using finite element analysis, Terzopoulos uses piecewise linear right-angled *triangular* spline elements for the membrane model. Because the squared derivatives get integrated over local two-dimensional patches of area h^2 , the discrete first order squared smoothness terms become independent of h .

For both the one and two-dimensional problems, the total energy of the discretized problem can be written as a *quadratic form*

$$E = E_d + E_s = x^T A x - 2x^T b + c, \quad (7)$$

where $x = [f_0 \dots f_{n-1}]$ is called the *state vector*.

The sparse symmetric positive-definite matrix A is called the *Hessian* since it encodes the second derivative of the energy function. (In numerical analysis, A is called the *coefficient* matrix [Saad 2003], while in finite element analysis [Bathe and Wilson 1976],

it is called the *stiffness* matrix.) For the one-dimensional first order problem, A is tridiagonal, while for the two-dimensional problem, it is multi-banded with 5 non-zero entries per row. We call b the *weighted data vector*. Minimizing the above quadratic form is equivalent to solving the sparse linear system

$$Ax = b. \quad (8)$$

3 Direct and iterative techniques

The solution of sparse positive definite (SPD) linear systems of equations breaks down into two major classes of algorithms: direct and iterative.

Direct methods use sparse matrix representations together with stable factorization algorithms such as LU (lower-upper) or Cholesky decomposition to obtain sparse symmetric factor matrices for which forward and backward substitution can be efficiently performed [Duff et al. 1986]. Unfortunately, for two-dimensional problems, the factored matrices suffer from *fill-in*, i.e., the zero entries separating the main diagonal band and the off-diagonal bands (which result when the 2D variables are placed in raster order) get replaced with non-zero values. Furthermore, direct solvers, being sequential in nature, are not directly amenable to data-parallel execution on GPUs.

Iterative relaxation algorithms such as gradient descent, successive over-relaxation (SOR), and conjugate gradient descent minimize the quadratic energy function (7) using a series of steps that successively refine the solution by reducing its energy [Saad 2003]. Basic iterative algorithms can be accelerated using a variety of techniques, including multigrid, hierarchical basis preconditioning, and tree-based preconditioning.

Multigrid Multigrid methods operate on a pyramid and alternate between solving the problem on the original (fine) level and projecting the error to a coarser level where the lower-frequency components can be reduced [Briggs et al. 2000; Saad 2003]. This alternation between fine and coarse levels can be extended recursively by attacking the coarse level problem with its own set of inter-level transfers, which results in the commonly used *V-cycle* and *W-cycle* algorithms [Briggs et al. 2000].

While multigrid techniques are provably optimal for simple homogeneous problems, their performance degrades as the problem becomes more irregular. Algebraic multigrid solvers (AMG), which locally adapt the interpolation functions to the local structure of the coefficient matrix, can be more effective in these cases [Briggs et al. 2000; Saad 2003]. However, to our knowledge, AMG solvers have not previously been applied to computer graphics problems. The techniques we develop are related to AMG (as we discuss later), although we use preconditioned conjugate gradient as our basic algorithm rather than multigrid, since it can compensate for the imperfect coarsening that is endemic in most real-world problems [Saad 2003].

Hierarchical basis preconditioning An alternative approach to solving sparse multi-dimensional relaxation problems also draws its inspiration from multi-level techniques. However, instead of solving a series of multi-resolution sub-problems interspersed with inter-level transfers, preconditioned conjugate gradient uses the multi-level hierarchy to *precondition* the original system, i.e., to make the search directions more independent and better scaled. These techniques were first developed in the numerical analysis community by Yserentant [1986] and applied to computer vision problems by Szeliski [1990]. More recent versions using different




Figure 3: *Multiresolution pyramid with half-octave (quincunx sampling (odd levels are colored gray for easier visibility). Hierarchical basis function control variables are shown as black dots.*

basis functions (such as wavelets) and some amount of local adaptation have been proposed by Pentland [1994], Yaou and Chang [1994], Gortler and Cohen [1995], and Lai and Vemuri [1997].

The basic idea is to replace the original *nodal* variables $x = [f_0 \dots f_{n-1}]$ with a set of *hierarchical* variables y that lie on a multi-resolution pyramid, which has been sub-sampled so that it has the same number of samples as the original problem (Figure 3). Because some of these new variables have larger bases (influence regions), the solver can take larger low-frequency steps as it iteratively searches for the optimal solution.

To convert between the hierarchical (multi-resolution) and nodal bases, we use a set of local interpolation steps,

$$x = Sy = S_1 S_2 \dots S_{L-1} y. \quad (9)$$

The columns of S can be interpreted as the *hierarchical basis functions* corresponding to the nodal variables [Yserentant 1986; Szeliski 1990]. In the wavelet community, this operation is called the *lifting scheme* for *second generation wavelets* [Schröder and Sweldens 1995; Sweldens 1997].

The form of the individual interlevel interpolation functions S_l is usually similar to that used in multigrid, e.g., linear interpolation for one-dimensional problems and bilinear interpolation for two-dimensional problems. Figure 4 shows a set of hierarchical bases that correspond to using uniform linear interpolation at each stage.

Hierarchical basis preconditioning often outperforms multigrid relaxation on the kind of non-uniform scattered data interpolation problems that arise in low-level vision [Szeliski 1990]. However, as more levels of preconditioning are used, the performance starts to degrade once the hierarchical bases become larger than the “natural” local spacing of the data. This is due to the fact that the bases are not properly adapted to particular interpolation problem at hand. While attempts have been made to take into account both the strength of the data constraints [Lai and Vemuri 1997] and local discontinuities [Szeliski 1990; Zhang et al. 2002], the question still remains whether one can derive an *optimal* set of hierarchical basis functions, or at least a more principled way to adapt them to the local structure of the underlying problem.

4 One-dimensional problems

The basic idea in using a hierarchical basis is to interpolate the solution obtained with a coarser level approximation and to then add a local *correction* term. However, what if we were given the exact solution at the coarser level? Could we then *perfectly* predict the solution at the finer level?




Figure 4: *Regular (non-adaptive) hierarchical basis functions.* Note how some of the functions (at the ends and in the middle) are very broad, corresponding to the coarse-level bases, while others become progressively narrower.

For the first order one-dimensional problem, it turns out that we can because each variable only interacts with its two nearest neighbors. Thus, if we place all the even variables in the coarse level and keep the odd variables at the fine level, we can solve for each fine level variable independent of the others. Consider the i th row from the tridiagonal system of linear equations (8),

$$a_{i,i-1}x_{i-1} + a_{i,i}x_i + a_{i,i+1}x_{i+1} = b_i. \quad (10)$$

We can write the solution for x_i as

$$x_i = a_{i,i}^{-1} [b_i - a_{i,i-1}x_{i-1} - a_{i,i+1}x_{i+1}]. \quad (11)$$

Note that this “interpolant” does *not*, in general, interpolate the two parents x_{i-1} and x_{i+1} since $a_{i,i-1}/a_{i,i} + a_{i,i+1}/a_{i,i} < 1$ whenever $w_i > 0$, i.e., whenever there is any data present at a node. In fact, as the data weight at a point $w_i \rightarrow \infty$, we get $x_i \rightarrow b_i/a_{i,i} = d_i$ and the target data value is interpolated exactly without using *any* smoothing. This is an example of what Sweldens calls *weighted wavelets* [Sweldens 1997].

We can now substitute the above formula for x_i for all odd values of i simultaneously and obtain a new energy function that does not involve the odd variables. This process is commonly known as *cyclic reduction* [Golub and Van Loan 1996].

Another way to view this operation is to say that (11) defines a locally adapted basis (interpolation) function [Kobbelt and Schröder 1998]. Figure 5 shows the locally adapted hierarchical basis functions for the data weighting and smoothness function introduced in Figure 2. Notice how none of the functions span the discontinuity in the smoothness function and how the local shape of the functions is controlled by the data spacing and strength.

If we permute the entries in A and S so that the fine level variables come first and the coarse level variables come second, we can rewrite these matrices as

$$A = \begin{bmatrix} D & E \\ E^T & F \end{bmatrix} \quad \text{and} \quad S_1 = \begin{bmatrix} I & -D^{-1}E \\ 0^T & I \end{bmatrix}. \quad (12)$$

The resulting *preconditioned* coefficient matrix can therefore be written as

$$\hat{A}_1 = S_1^T A S_1 = \begin{bmatrix} D & 0 \\ 0^T & F - E^T D^{-1} E \end{bmatrix}. \quad (13)$$

For the one-dimensional first order problem, D is a diagonal matrix, E has a bandwidth of two, and F is also diagonal. The resulting (smaller) coarse-level coefficient matrix $\hat{A}_1^c = F - E^T D^{-1} E$ is therefore tri-diagonal, just like the original coefficient matrix A . We can therefore recurse on this matrix to construct a locally adapted set of hierarchical basis functions that that convert the original coefficient into a diagonal matrix \hat{A} . In the iterative methods community, the above algorithm is called a *complete factorization* of




Figure 5: *Locally adapted hierarchical basis functions.* Note how none of the functions span the break at $i = \{11, 12\}$ and how the shape of individual functions varies according to the local data weights and smoothness values. This particular set of bases is an exact preconditioner for the 1-D problem.

\hat{A} , since no approximations are used in this derivation [Ciarlet Jr. 1994]. It is a special instance of the more general *incomplete LU factorization with multi-elimination* (ILUM) algorithm [Saad 2003, §12.5], which we describe in the next section.

The construction described above leads to the following re-interpretation of the well known cyclic reduction algorithm [Kobbelt and Schröder 1998]:


For a one-dimensional first-order problem, using the preconditioner $S\hat{A}^{-1}S^T$, where S is defined using the set of locally adapted hierarchical basis functions $S_1 = [I \quad -D_1^{-1}E_1]$ as described above and \hat{A} is the resulting diagonal (preconditioned) multi-level coefficient matrix, results in perfect preconditioning and hence a one-step solution.

Unfortunately, this result only holds for one-dimensional first order problems. However, the general methodology developed above can be modified so that it can be applied to two-dimensional problems, as we describe next.

5 Two-dimensional problems

In two dimensions, most hierarchical basis function approaches (as well as pyramid-based algorithms) eliminate odd rows and columns simultaneously, so that each coarser level has one quarter as many variables as the finer one [Szeliski 1990]. However, the fine-level variables being eliminated are not independent, which makes it difficult to apply the previous methodology.

Instead, we eliminate the red elements in a red-black checkerboard first, and then eliminate the other half of the variables to move a full level up in the pyramid. In the numerical analysis community, this is known as *repeated red-black (RRB) ordering* (or elimination) [Brand 1992; Ciarlet Jr. 1994; Notay and Amar 1997], and is a particular instance of more general *multicoloring* strategies that can be applied in conjunction with *incomplete LU with multi-elimination* (ILUM) preconditioning [Saad 2003]. (Half-octave pyramids are also used in image processing [Crowley and Stern 1984], where they are sometimes called *quincunx* sampling [Feilner et al. 2005].) Pictorially, we can view this as



where the \times 's are eliminated in the first pass and the diamonds are eliminated in the second. The gray lines indicate the connections (non-zero entries in A) between variables, while the blue lines indicate undesirable connections, which must be eliminated, as described below.




Figure 6: Two-dimensional cliques of neighbors and how they are treated during the simplification process. (a–c): The left hand side of each arrow diagram shows the original (non-zero) connections between the crosses and the circle along with their original weights. The right hand side shows the connections after variable elimination. The non-diagonal connections shown in blue are eliminated (redistributed to the black ones) in the final step. (d): General case of diagonal edge weight redistribution.

Using this scheme, the D matrix is diagonal, the E matrix is 4-banded (4 nearest neighbors), and the F matrix is diagonal (crosses and diamonds are independent of each other). The resulting \hat{A}_1^c matrix is 9-banded, i.e., each circle (or diamond) depends on not only its immediate 4-neighbors (gray lines in the above diagram and Figure 6), but also on its diagonal 2-D neighbors (light blue lines).

In order to keep the bandwidth from growing, we would like to eliminate these diagonal connections. But how? The usual answer [Brand 1992; Saad 2003] is to perform *incomplete factorization*, which means simply dropping the undesired element (ILU0), or perhaps adding these dropped values to the diagonal to preserve affine/interpolatory properties of the resulting system (known as *modified ILU* or MILU). (A blend of these two approaches, known *relaxed ILU* or RILU, is to add a fraction ω of the dropped elements to the diagonal [Saad 2003].)

This paper takes an alternative approach. Rather than thinking of the undesired matrix elements as abstract quantities, we view them as springs whose impact on the system must be approximated by stiffening adjacent (non-eliminated) springs. The goal is to make the coarsened *energy function* (quadratic form with enforced zero entries) be as good as possible an approximation to the original energy.

Let us start with some special cases. If the data constraint is so strong that the connections in $D^{-1}E$ are essentially zero, we do not have a problem. The same is true when a circle only depends on an adjacent pair of neighbors (Figure 6a), since this induces a single new link (dependence).

When three out of the four connections are non-zero (Figure 6b), we can view this configuration as a triangular finite element. If the original weights between neighbors are 1, it is easy to demonstrate (by substituting the average value $(x_j + x_k + x_l)/3$ for x_i)¹ that the resulting energy has connections of strengths $1/3$ between the neighbors shown in Figure 6b (right size of the first arrow). For a triangular element of area h^2 with unit smoothness $s = 1$, the diagonal connections should be $1/2$, so that they add up to 1 after the other side of each link is also eliminated. (Recall that the finite element analysis in Section 2 indicates that the membrane energy in 2D does not scale with h .) Therefore, a sensible heuristic is to “distribute” half of the strength along the vertical blue line to each of its two neighbors, which results in the desired pair of values $(1/2, 1/2)$.

For the fully four-connected case (Figure 6c), a similar analysis shows that the vertical and horizontal strength should be evenly distributed between the adjacent neighbors *after being scaled up by a factor of 2*, which results in a final strength $1/2$ on each of the diagonal (black) edges.

What about the case when the strengths are not all the same? A general heuristic rule can be formulated that reduces to the above

¹Note how we now use single subscripts to index the variables

special cases:

Heuristic 1: For the two-dimensional membrane, after multi-elimination of the red (cross) variables, distribute any “diagonal” connection strengths in the rotated grid to each connection’s nearest neighbors, using the relative weighting of the original 4-connections, with appropriate scaling.

More formally, if a variable x_i originally has connections to its 4 neighbors of weight $a_{ij} \dots a_{im}$ (Figure 6d), and the interpolation (adapted hierarchical basis) weights are $s_{ij} \dots s_{im}$, the entries in the new coarser-level coefficient matrix \hat{A}_1 (13) corresponding to that finite element will be

$$\hat{a}_{jk} = a_{ii}s_{ij}s_{ik}, \quad \text{etc.} \quad (14)$$

(The entry \hat{a}_{jk} will also receive a contribution from the finite element bordering it on the other side.)

This formula works fine, except that it introduces undesired entries between nodes j and l (and also k and m , not shown in Figure 6d), which must be re-distributed to the neighboring pairs $jk \dots jm$. To do this, we first compute the relative weights for the four allowable links,

$$w_{jk} = \hat{a}_{jk} / \sum_{lm} \hat{a}_{lm}, \quad (15)$$

and then update each of the allowable link weights by

$$\hat{a}_{jk} \leftarrow \hat{a}_{jk} + s_N w_{jk} \hat{a}_{jl}. \quad (16)$$

(The diagonal entries are also adjusted to maintain the balance between off-diagonal and diagonal entries in \hat{A}_1 .) The scaling factor $s_N = 2$ is chosen so that the coarse level coefficient matrix in homogenous (constant smoothness) regions is the same as would be obtained by applying finite element analysis at the coarser level (see [Szeliski 2006] for more details and a discussion of alternatives).

5.1 GPU implementation

Implementing the locally adaptive hierarchical basis preconditioned conjugate gradient algorithm is straightforward, as it involves only a small extension to the conjugate gradient and multigrid algorithms already developed by Bolz *et al.* [2003]. The basic operations in conjugate gradient, namely the evaluation of the residual vector r , which involves the sparse matrix multiplication $Ax - b$, and the computation of the new conjugated direction d and step size, which involve the computation of two inner products, map naturally onto the GPU. (Note that the A matrix, which has only 4 non-zero entries per row, is stored naturally alongside all of the other vector (image) quantities as a 4-banded float image.) Similarly, the transformation of the current direction vector by the S^T and S matrices [Szeliski 1990] involves nothing more than the inter-level transfer operations present in multigrid, where the spatially varying entries in each S_l matrix can be again stored in a 4-banded image aligned with the quantities they are being applied to.

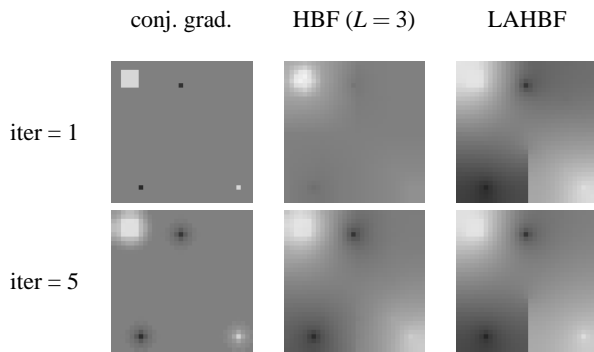


Figure 7: Simple two-dimensional discontinuous interpolation problem. The two rows show the solution after the first and fifth iterations, and the three columns show regular conjugate gradient, hierarchical basis preconditioning (3 levels), and LAHBF preconditioning. Note how conjugate gradient hardly makes any progress, while regular HBFs over-smooth the solution.

6 Experimental results

To evaluate the performance of our new algorithm, we have compared it to regular conjugate gradient descent, regular (non-adaptive) hierarchical basis preconditioning, and three different variants of incomplete factorization (ILU0, MILU, and RILU with $\omega = 0.5$). Figure 7 shows a simple 2D interpolation problem that consists of interpolating a set of 4 sparse constraints on a square grid with a tear halfway across the bottom edge. As you can see, conjugate gradient hardly makes any progress, even after 20 iterations. Regular hierarchical basis preconditioning starts off strong, but after 5 iterations, it has not yet accurately modeled the discontinuity. Locally adaptive hierarchical basis preconditioning, on the other hand, achieves a visually acceptable solution in a single iteration.

Figure 8 shows a plot of the RMS error (relative to the minimum energy solution, expressed in gray levels) for all of these algorithms. As you can see, LAHBF significantly outperforms all previously published algorithms, which can also be verified by looking at the semi-log plots of the error curves [Szeliski 2006].

We next evaluate the performance of our algorithm on a number of computer graphics applications.

Colorization Colorization is the process of propagating a small number of color strokes to a complete gray-scale image, while attempting to preserve discontinuities in the image [Levin et al. 2004a]. As such, is it a perfect match to the controlled-continuity interpolators developed in this paper. The actual smoothness term being minimized in [Levin et al. 2004a] involves differences between a pixel’s value and the average of its eight neighbors. As such, it is not a first-order smoothness term and is not directly amenable to our acceleration technique. Instead, we replace the term used in [Levin et al. 2004a] with a simpler term, i.e., we set the horizontal and vertical smoothness strengths inversely proportional to the horizontal and vertical grayscale gradients.

Figure 1 visually shows the result of running our algorithm, as well as conventional preconditioned conjugate gradient on the sparse set of color strokes shown in Figure 1a. As you can see, conjugate gradient hardly makes any progress, even after 20 iterations, while LAHBF preconditioned conjugate gradient converges in just a few iterations. Figure 9 shows the convergence of these algorithms (RMS error in the reconstructed chrominance signals) as a function of the number of iterations.




Figure 8: Error plot for the simple two-dimensional problem. The LAHBF preconditioned solver converges in just a few iterations, while regular HBFs converge more slowly and require manual selection of the number of levels for best performance. The MILU algorithm performs better than non-adapted bases, but still significantly underperforms the new algorithm.

Poisson (gradient domain) blending Poisson blending is a technique originally developed to remove visual artifacts due to strong intensity or color differences when a cutout is pasted into a new background [Pérez et al. 2003]. It has since been extended to non-linear variants [Levin et al. 2004b], as well as applied to image stitching, where it is used to reduce visual discontinuities across image seams [Agarwala et al. 2004]. Poisson blending relies on reconstructing an image that matches, in a least-squares sense, a gradient field that has been computed from the various images that contribute to the final mosaic.

Figure 10 shows an original two-image stitch where blending has been applied. The left image shows the unblended result. The middle image shows the label field, indicating where each pixel is drawn from. The right image shows the blended result, obtained using just a few iteration of LAHBF preconditioned conjugate gradient. (Notice that a slight seam is still visible, because the *contrast* is not properly matched across images. This suggest that blending in the log-luminance domain might be more appropriate.) Figure 10 also shows the corresponding error plots. As before, we see that LAHBF converges in just a handful of steps. In these examples, rather than selecting a single pixel as a hard constraint to remove the overall shift ambiguity, we used a weak constraint towards the unblended original image.

Tone mapping Our final graphics application is gradient domain high dynamic range compression [Fattal et al. 2002]. In this technique, the gradients of a log-luminance function are first computed and then compressed through a non-linearity (Figure 11a–b), A compressed log-luminance function is then reconstructed from these gradients. Figure 11c shows the sparse set of constraints used to clamp the dark and light values, Figure 11d shows the reconstructed log-luminance function, and Figure 11e shows the con-




Figure 9: Error plot for the simple colorization problem. The LAHBF preconditioned solver converges in just a few iterations, while regular HBFs and MILU converge more slowly. (For a semi-log plot, please see [Szeliski 2006].)

vergence plots. As before, we see that our new approach offers significant performance improvements. We have also applied our algorithm to a newly developed interactive tone mapping algorithm [Lischinski et al. 2006] with good results [Szeliski 2006].

Performance The performance of our algorithm on the three real-world examples discussed above is summarized in Table 1. These numbers were obtained on a Xeon 2.8GHz PC running on unoptimized C++ code. The time taken and memory size are roughly proportional to the problem size. Computing the locally adaptive basis functions takes a little longer than actually using them as a preconditioner, which is about twice as slow as the basic conjugate gradient update step. The convergence plots shown previously should therefore have their top lines (regular conjugate gradient $L = 0$) squished left by a factor of three to make the comparisons sensible in terms of run times. However, the advantages of LAHBF over conjugate gradient still remain quite dramatic

7 Discussion

Our experimental results demonstrate that our locally adapted hierarchical basis functions significantly outperform unadapted hierarchical bases as preconditioners for the kinds of first-order optimization problems presented in the previous section.

As we mentioned before, our technique is closely related to both algebraic multigrid and ILUM. Incomplete LU (lower-upper) factorization with multi-elimination (ILUM) on a red-black (multicolor) grid performs the same first step of problem reduction/coarsening as our algorithm. The major difference comes in how the undesired off-diagonal entries are treated. In ILUM, it is traditional to simply drop coefficient matrix entries that are of small magnitude relative to other entries in a given row. (This is known as ILU with thresholding, or ILUT [Saad 2003, p. 306].) In our approach, rather than dropping such terms we *re-distribute* these entries to other off-diagonal and on-diagonal entries in a way that *preserves* a good finite-element approximation to the original variational problem.

Algebraic multigrid techniques attempt to remedy some of the shortcomings of traditional multigrid techniques, which assume a fixed uniform set of coarsening and interpolation kernels, by adapting these kernel to the underlying structure (continuity and stiffness) of the finite element problem being solved. As discussed in [Saad 2003, p. 444], more recent AMG techniques use the same basic multi-elimination equations to define the interlevel transfer

Applic.	Prob. size	Setup	LAHBF	CG	Prec.	Mem (MB)
Color.	$320 \times 262 \times 2$	31	157	30	94	15 / 35
Blend.	$686 \times 686 \times 3$	262	797	280	563	106 / 220
Compr.	$423 \times 636 \times 1$	359	1984	150	1200	174 / 373

Table 1: Performance of LAHBF on three sample problems. The problem size is given in pixels (width \times height) by the number functions (r.h.s. columns) being reconstructed. The setup (coefficient matrix computation), locally adaptive HBF pre-computation, conjugate gradient iteration, and hierarchical preconditioning steps are given in milliseconds (msec). The memory footprint is in MB, and shows the memory without and with hierarchical preconditioning.

operators as ILUM. As with ILUM, small terms in the resulting coarser level coefficient matrix are again neglected in order to keep the system bandwidth from growing excessively. Again, the approach presented in this paper suggests a more principled and effective way for re-distributing these neglected terms.

Another class of recently developed pre-conditioning techniques relies on *trees* defined over the original connectivity graph [Gremban 1996]. These techniques, also known as *combinatorial preconditioners* because of their close connection to graph theoretic algorithms, are known to perform well on sparse irregular linear systems. (<http://www.preconditioners.com> contains a good tutorial introduction as well as pointers to the most recent literature.) While we have not yet had a chance to compare the performance of our approach to these new techniques, we believe that for the 2D grid-based first order problems that we address in the paper, our technique will perform better because it does not neglect a subset of connections in the original problem, as the tree-based techniques do.

Future work The basic approach developed in this paper can be extended in a number of ways. We are currently working on extending our technique to second order problems in both one and two dimensions.

We are also interested in extending our approach to non-quadratic energy minimization problems such as dynamic programming (which can be used for stereo matching or speech recognition). While 1-D versions of these problems are known to have effective sequential solutions, our approach can be used to parallelize these algorithms onto architectures such as GPUs. The extension to 2-D MRFs is likely to be quite challenging, but could have large payoffs for problems such as stereo matching and graph cut segmentation.

Finally, we would like to apply our techniques to a wider range of computer graphics applications such as fluid simulations.

8 Conclusions

In this paper, we have shown how to locally adapt hierarchical basis functions to inhomogeneous problems so that they are more effective at preconditioning large first-order optimization problems that arise in computer graphics. Our approach is based on an extension of repeated red-black (RRB) ILUM preconditioners that uses finite element analysis to redistribute off-diagonal terms in a manner that preserves good approximations to the underlying variational problem. The resulting algorithm performs significantly better than previously proposed approaches, is simple to implement, and also maps naturally onto a GPU. We have demonstrated that our current algorithm has a wide range of applications in computer graphics, and we plan to extend its range of applicability in the future.



Figure 10: Poisson blending example. The left image shows the unblended result, the middle image shows the pixel labels, and the right image shows the final blended image.

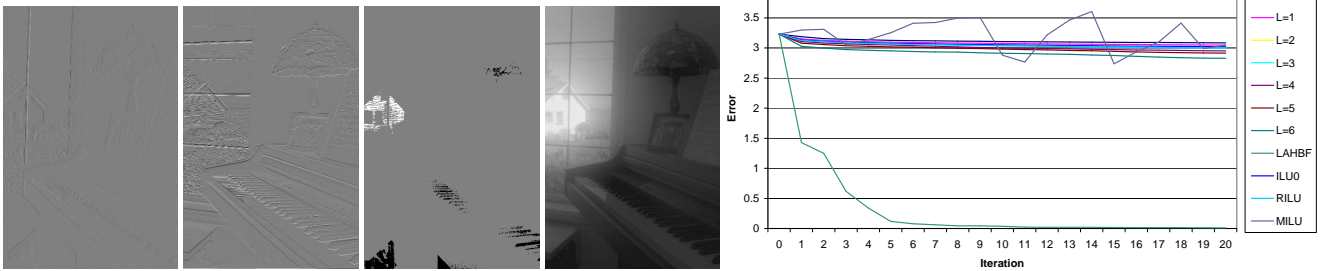


Figure 11: Gradient domain high dynamic range compression: Leftmost two images: horizontal and vertical compressed gradient fields; middle image: light and dark value constraints; right image: reconstructed compressed log-luminance function.

References

- AGARWALA, A., ET AL. 2004. Interactive digital photomontage. *ACM Transactions on Graphics* 23, 3 (August), 292–300.
- BATHE, K.-J., AND WILSON, E. L. 1976. *Numerical Methods in Finite Element Analysis*. Prentice-Hall, Inc., Englewood Cliffs, New Jersey.
- BOLZ, J., ET AL. 2003. Sparse matrix solvers on the GPU: Conjugate gradients and multigrid. *ACM Transactions on Graphics* 22, 3 (July), 917–924.
- BRAND, C. W. 1992. An incomplete factorization preconditioning using repeated red-black ordering. *Numer. Math.* 61, 433–454.
- BRIGGS, W. L., HENSON, V. E., AND MCCORMICK, S. F. 2000. *A Multigrid Tutorial*, second ed. Society for Industrial and Applied Mathematics, Philadelphia.
- CIARLET JR., P. 1994. Repeated red-black ordering: a new approach. *Numerical Algorithms* 7, 295–324.
- CROWLEY, J. L., AND STERN, R. M. 1984. Fast computation of the difference of low-pass transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6, 2 (March), 212–222.
- DUFF, I. S., ERISMAN, A. M., AND REID, J. K. 1986. *Direct Methods for Sparse Matrices*. Clarendon Press, Oxford.
- FATTAL, R., LISCHINSKI, D., AND WERMAN, M. 2002. Gradient domain high dynamic range compression. *ACM Transactions on Graphics (TOG)* 21, 3, 249–256.
- FEILNER, M., ET AL. 2005. An orthogonal family of quincunx wavelets with continuously adjustable order. *IEEE Transactions on Image Processing* 14, 4 (April), 499–520.
- GOLUB, G., AND VAN LOAN, C. F. 1996. *Matrix Computation, third edition*. The John Hopkins University Press, Baltimore and London.
- GORTLER, S. J., AND COHEN, M. F. 1995. Hierarchical and variational geometric modeling with wavelets. In *Symposium on Interactive 3D Graphics*, 35–43.
- GREMBAN, K. D. 1996. *Combinatorial Preconditioners for Sparse, Symmetric, Diagonally Dominant Linear Systems*. PhD thesis, Carnegie Mellon University. CMU-CS-96-123.
- KOBBELT, L., AND SCHRÖDER, P. 1998. A multiresolution framework for variational subdivision. *ACM Transactions on Graphics* 17, 4 (October), 209–237.
- LAI, S.-H., AND VEMURI, B. C. 1997. Physically based adaptive preconditioning for early vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 19, 6 (June), 594–607.
- LEVIN, A., LISCHINSKI, D., AND WEISS, Y. 2004. Colorization using optimization. *ACM Transactions on Graphics* 23, 3 (August), 689–694.
- LEVIN, A., ZOMET, A., PELEG, S., AND WEISS, Y. 2004. Seamless image stitching in the gradient domain. In *Eighth European Conference on Computer Vision (ECCV 2004)*, Springer-Verlag, Prague, vol. IV, 377–389.
- LISCHINSKI, D., FARBMAN, Z., UYTENDAELLE, M., AND SZELISKI, R. 2006. Locally adapted hierarchical basis preconditioning. *ACM Transactions on Graphics* 25, 3 (August).
- NOTAY, Y., AND AMAR, Z. O. 1997. A nearly optimal preconditioning based on recursive red-black ordering. *Numerical Linear Alg. Appl.* 4, 369–391.

- PENTLAND, A. P. 1994. Interpolation using wavelet bases. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 16, 4 (April), 410–414.
- PÉREZ, P., GANGNET, M., AND BLAKE, A. 2003. Poisson image editing. *ACM Transactions on Graphics* 22, 3 (July), 313–318.
- SAAD, Y. 2003. *Iterative Methods for Sparse Linear Systems*, second ed. SIAM.
- SCHRÖDER, P., AND SWELDENS, W. 1995. Spherical wavelets: Efficiently representing functions on the sphere. In *Proceedings of SIGGRAPH 95*, Computer Graphics Proceedings, Annual Conference Series, 161–172.
- SWELDENS, W. 1997. The lifting scheme: A construction of second generation wavelets. *SIAM J. Math. Anal.* 29, 2, 511–546.
- SZELISKI, R. 1990. Fast surface interpolation using hierarchical basis functions. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12, 6 (June), 513–528.
- SZELISKI, R. 2006. Locally adapted hierarchical basis preconditioning. Tech. Rep. MSR-TR-2006-38, Microsoft Research, May.
- TERZOPOULOS, D., AND WITKIN, A. 1988. Physically-based models with rigid and deformable components. *IEEE Computer Graphics and Applications* 8, 6, 41–51.
- TERZOPOULOS, D. 1983. Multilevel computational processes for visual surface reconstruction. *Computer Vision, Graphics, and Image Processing* 24, 52–96.
- TERZOPOULOS, D. 1986. Regularization of inverse visual problems involving discontinuities. *IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI-8*, 4 (July), 413–424.
- YAOU, M.-H., AND CHANG, W.-T. 1994. Fast surface interpolation using multiresolution wavelets. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 16, 7 (July), 673–689.
- YSERENTANT, H. 1986. On the multi-level splitting of finite element spaces. *Numerische Mathematik* 49, 379–412.
- ZHANG, L., ET AL. 2002. Single view modeling of free-form scenes. *Journal of Visualization and Computer Animation* 13, 4, 225–235.