Some Observations on Local Least Squares

Richard H. Bartels	Gene H. Golub	Faramarz F. Samavati
Computer Science	Computer Science	Computer Science
University of Waterloo	Stanford University	University of Calgary
Waterloo, Ontario	Stanford, California	Calgary, Alberta
N2L 3G1	94305	T2N 1N4
Canada	U.S.A.	Canada
rhbartel@uwaterloo.ca	golub@stanford.edu	samavati@cpsc.ucalgary.ca

March 29, 2004

Abstract

In a previous work [2] we introduced a construction to produce biorthogonal multiresolutions from given subdivisions. The approach we employed could be interpreted as estimating the solution to a least squares problem by means of a number of smaller least squares approximations on local portions of the data. In this work we use a result by Dahlquist, et. al. [3], on the *method of averages* to explore how and to what extent this local least squares estimation approaches the full least squares approximation. Two example problem domains are used:

1. Data reduction processes – characterized by the removal of knots from splines, the first stage in the construction of a biorthogonal multiresolution, and the simplification of geometric objects according to subdivision schemes;
2. Approximation processes – characterized by data sampling and discrete least squares.

We observe in these examples that the least squares solution is often well estimated by local least squares, that the estimation rapidly improves with the size of the local least squares problems, and that the quality of the estimate is independent of the size of the full problem in uniform situations.

keywords: subdivision, interpolation, least squares, projection, subspace angles, left inverse.

Dedication

This work is dedicated to Germund Dahlquist on the occasion of his 80^{th} birthday.

1 Introduction

The material developed in [2, 9] proposed a local, least-squares construction for processing geometric data into a multiresolution form. A number of linear algebraic issues arise in that construction, some of which have a wider implication.

The construction, briefly, begins with an overdetermined $(m \times n)$ system of equations

$$\mathbf{P}\mathbf{z} \approx \mathbf{f} \tag{1.1}$$

and has the goal of producing associated matrices **A**, **B**, and **C** such that

$$\mathbf{A}\mathbf{f} \to \mathbf{g} \tag{1.2}$$

represents a solution comparable to the ordinary least squares solution **z** of (1.1), and such that

$$\mathbf{f} \equiv \mathbf{P}\mathbf{g} + \mathbf{C}\mathbf{d} \tag{1.3}$$

where

$$\mathbf{B}\mathbf{f} \to \mathbf{d} \tag{1.4}$$

In this regard $\begin{bmatrix} \mathbf{A} \\ \mathbf{B} \end{bmatrix}$ and $[\ \mathbf{P}\ \ \mathbf{C}\ ]$ are inverses.

The construction was designed for banded matrices **P** of full column rank. Experience has been gained up to this point with matrices defining certain forms of subdivision and with B-spline interpolation. In both these contexts the matrices are "σ-shifted," meaning that each column is displaced at least $(\sigma > 1)$ rows downward from its predecessor.

The construction consists of two parts; the first part constructs **A**, and the second part constructs **B** and **C** together. The approximate solution vector **g**, and the vector **d** by which the residuals are represented, are simple byproducts of matrix multiplication. It is the construction of **A** that will concern us here. Briefly, in the construction each row of **A** is produced by selecting the next column in sequence from **P**, forming an overdetermined submatrix of **P** containing the nonzeros of that column, and determining the nonzero row elements of **A** from a least squares problem based on that submatrix. The motivation is simple: the j^{th} column of **P** represents the j^{th} unknown's contribution to the right hand side,

$$\mathbf{f} = \cdots + z_j \mathbf{p}_j + \cdots$$

If **P** is banded, this contribution is restricted to a portion of **f**

$$\begin{bmatrix} \vdots \\ f_{i_a} \\ \vdots \\ f_{i_b} \\ \vdots \end{bmatrix} = \cdots + z_j \begin{bmatrix} \mathbf{0} \\ p_{i_a,j} \\ \vdots \\ p_{i_b,j} \\ \mathbf{0} \end{bmatrix} + \cdots$$

We would hope to get a good estimate of z_j from a least squares subproblem localized upon the nonzero portion of column $\mathbf{p}_j$. And, if the matrix **P** is σ-shifted, any selection of the columns having nonzeros in the rows containing the nonzeros of $\mathbf{p}_j$ will tend to yield overdetermined submatrices. As such, the construction of **A** is designed to construct an approximate solution of (1.1) by generating each component of the approximation through a local least squares estimate.

In Section 2 we give more details on the formation of **A**, and in Section 3 we use the material on subspace angles and the theory developed for the method of averages in [1] and [3] to establish how to compare the resulting vector **g** of (1.2) with the ordinary least squares solution vector **z** of (1.1). Section 4 will consider some examples of **A** for some choices of **P** and give estimates of what quality of solution **g** can be obtained. Section 5 will conclude with a mention of how the construction of **A** can be realized in the context of a finite element mesh.

2 Formation of A

We view the solution of (1.1) as requiring the estimation of each z_j in $\mathbf{z}$ individually. Consider rows indexed i_a through $i_b > i_a$ of column j of $\mathbf{P}$ (banded), chosen so that there is a nontrivial intersection with the nonzero elements of the j^{th} column, which are contained beginning with row $i_a + \lambda$ and running through row $i_b - \mu$ for integers λ and μ.

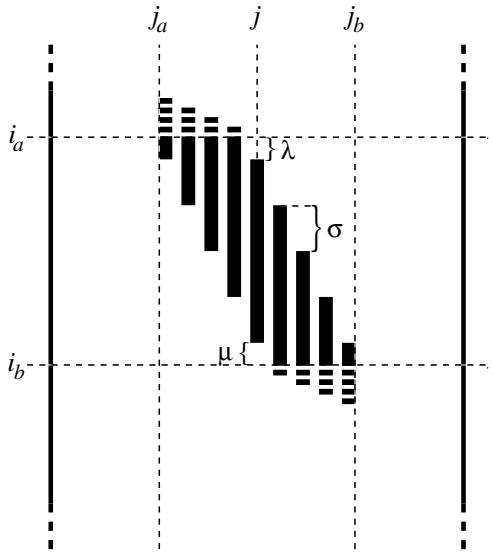

Figure 1: Local submatrix $\mathbf{P}_j$ for estimating z_j

Clearly, the values of λ and μ are fixed by the location of the nonzeros in the j^{th} column and the choice of i_a and i_b. If λ and μ are both nonnegative, then all nonzero elements of the j^{th} column have been included in the selection of rows i_a through i_b. If either or both are negative, then some number of nonzero entries from column j will have been ignored in the construction.

The matrices on which we carry out our construction are to be banded. We expect them to have only zero entries in rows i_a through i_b for all columns sufficiently far to the left and right of the j^{th}, say columns with indices ℓ satisfying $|\ell - j| > K$ for a positive integer K. And we expect that the nonzeros of each successive column are located further down the matrix. Our experience has been with matrices $\mathbf{P}$ where any column's nonzero content is shifted down σ rows from the content of its predecessor. Figure 1 gives an impression when λ and μ are positive.

Having chosen i_a and i_b, we have a starting submatrix. This submatrix may have to be be narrowed by casting out columns on the left or right that have only zeros, or it may have to be expanded if columns to the left and/or right intrude on the submatrix to impact its rank or its overdetermined structure. Assuming success with these adjustments, we are left with columns j_a through j_b. We denote the resulting submatrix by $\mathbf{P}_j$, and we use $\mathbf{P}_j$ to estimate z_j by least squares.

Assumption 1 *The choice of i_a and i_b causes $\mathbf{P}_j$ to have no fewer rows than columns and to be of full rank.*

An estimate of z_j can be extracted from

$$\left(\mathbf{P}_j{}^T\mathbf{P}_j\right)^{-1}\mathbf{P}_j{}^T\mathbf{f}_j \tag{2.1}$$

where $\mathbf{f}_j$ is the vector comprising $f_{i_a}, \ldots, f_{i_b}$. This estimates $z_{j_a}, \ldots, z_j, \ldots, z_{j_b}$. However, all except z_j are discarded, so that we retain only

$$z_j \approx \mathbf{e}_j{}^T \left(\mathbf{P}_j{}^T \mathbf{P}_j\right)^{-1} \mathbf{P}_j{}^T \mathbf{f}_j \tag{2.2}$$

where $\mathbf{e}_j$ is a vector of the appropriate length with 1 in position j and zeros elsewhere. This is repeated for each component of the vector $\mathbf{z}$.

The row vector $\mathbf{a}_j = \mathbf{e}_j{}^T \left(\mathbf{P}_j{}^T \mathbf{P}_j\right)^{-1} \mathbf{P}_j{}^T$ will have $i_b - i_a + 1$ elements, and it can be expanded into a vector $[\mathbf{0}\ \mathbf{a}_j\ \mathbf{0}]$ of length n by prefixing zeros in positions 1 through $i_a - 1$ and suffixing zeros in positions $i_b + 1$ through n. The matrix $\mathbf{A}$ is constructed by assembling all such expanded rows j, $j = 1, \ldots, n$. $\mathbf{g} = \mathbf{A}\mathbf{f}$ is to be taken as an estimate of $\mathbf{z}$.

Theorem 1 $\mathbf{A}$ *is a left inverse of* $\mathbf{P}$

Proof: The j^{th} row of $\mathbf{A}$ has zeros initially and finally, by construction, so that it could only form nonzero products by left multiplication with the elements in rows i_a through i_b of $\mathbf{P}$. Within these rows, the only nontrivial elements of $\mathbf{P}$ are to be found in columns j_a through j_b. Thus, the j^{th} row of $\mathbf{A}$ has the same product with $\mathbf{P}$ that the vector $\mathbf{a}_j$ has with the submatrix $\mathbf{P}_j$; namely, 1 for column j and zero for all other columns. The construction produces a row $\mathbf{a}_j$ for each column j of $\mathbf{P}$. The product $\mathbf{AP}$ is, as a consequence, the $n \times n$ identity. □

3 Analysis of A

Dahlquist [3] has compared the "method of averages" to the ordinary method of least squares for overdetermined linear systems. Briefly, the method of least squares solves (1.1) as

$$\left(\mathbf{P}^T \mathbf{P}\right)^{-1} \mathbf{P}^T \mathbf{f} \to \mathbf{z} \tag{3.1}$$

while the method of averages solves (1.1) as

$$\left(\mathbf{G}^T \mathbf{P}\right)^{-1} \mathbf{G}^T \mathbf{f} \to \mathbf{h} \tag{3.2}$$

The $n \times m$ matrix $\mathbf{G}$ is a *summation matrix*, meaning that any row of $\mathbf{G}^T$ has the form $[0, \ldots, 0, 1, \ldots, 1, 0, \ldots 0]$ for some number of consecutive unit entries located at some position in the row. Only those summation matrices for which $\mathbf{G}^T\mathbf{P}$ is nonsingular are to be considered, of course. Anderson [1] has extended the class of matrices $\mathbf{G}$ useful for problems where $\mathbf{P}$ arises from approximating data by a Chebyshev system of functions to *rectangular smoothing matrices*. These are defined by requiring $\mathbf{G}$ to have full rank and requiring every $n \times n$ minor of $\mathbf{G}$ formed by deleting rows to be nonnegative.

Both papers compare $\mathbf{z}$, the solution of (1.1) via (3.1), and $\mathbf{h}$, the solution of (1.1) via (3.2), by looking at how much their corresponding residuals could differ. More precisely, they look at

$$\eta = \inf_{\mathbf{f}} \left(\frac{\|\mathbf{f} - \mathbf{P}\mathbf{z}\|^2}{\|\mathbf{f} - \mathbf{P}\mathbf{h}\|^2} \right) \tag{3.3}$$

A simple geometric argument, proceeding from the observation that $\mathbf{f} - \mathbf{P}\mathbf{z}$ is orthogonal to $\mathcal{P}$, the column space of $\mathbf{P}$, and that $\mathbf{f} - \mathbf{P}\mathbf{h}$ is orthogonal to $\mathcal{G}$, the column space of $\mathbf{G}$, establishes that

$$\frac{\|\mathbf{f} - \mathbf{P}\mathbf{z}\|}{\|\mathbf{f} - \mathbf{P}\mathbf{h}\|} = \cos\left[\mathcal{P}, \mathcal{G}\right] \tag{3.4}$$

where $[\mathcal{P},\mathcal{G}]$ is the final principal angle between the column spaces as ordered in the definition of [7] (Chapter 12, Section 12.4.3).

An orthonormal basis for $\mathcal{P}$ is given by the columns of the "skinny" QR factorization of $\mathbf{P}$: $\mathbf{P}=\mathbf{Q_P R_P}$ ([7], Chapter 5, Section 5.2.6), and likewise for $\mathcal{G}$. The singular values of the product matrix $\mathbf{Q_P^T Q_G}$ provide the cosines of the principal angles between $\mathcal{P}$ and $\mathcal{G}$ ([7] Chapter 12, Section 12.4.3). The cosine in (3.4) is given by the smallest singular value.

In Theorem 1 of [3] it is established that η is given by $\cos^2[\mathcal{P},\mathcal{G}]$.

Theorem 2

$$\gamma=\inf_{\mathbf{f}}\left(\frac{\|\mathbf{f}-\mathbf{Pz}\|^2}{\|\mathbf{f}-\mathbf{Pg}\|^2}\right)=\cos^2[\mathcal{P},\mathcal{A}] \tag{3.5}$$

where $\mathbf{g}$ *is the solution to (1.1) by* $\mathbf{Af}\rightarrow\mathbf{g}$ *and* $\mathcal{A}$ *is the column space of* $\mathbf{A}^T$.

Proof: By Theorem 1 and the definition of $\mathbf{g}$, $\mathbf{A}(\mathbf{f}-\mathbf{Pg})=\mathbf{g}-\mathbf{APg}=\mathbf{g}-\mathbf{g}$. Consequently, $\mathbf{f}-\mathbf{Pg}$ is orthogonal to the column space of $\mathbf{A}^T$, and the geometric argument of [3] applies to yield

$$\frac{\|\mathbf{f}-\mathbf{Pz}\|}{\|\mathbf{f}-\mathbf{Pg}\|}=\cos\left[\mathbf{P},\mathbf{A}^T\right] \tag{3.6}$$

Theorem 1 of [3] applies directly to our construction, since our $\mathbf{A}^T$ is simply a substitute (of the same size and rank) for the matrix $\mathbf{G}$, and we are solving $\mathbf{APg}=\mathbf{Af}$ where the method of averages solves $\mathbf{G}^T\mathbf{Ph}=\mathbf{G}^T\mathbf{f}$. Nevertheless, to confirm that our $\mathbf{A}^T$ is a proper stand-in for $\mathbf{G}$, we check that $\mathbf{A}^T$ satisfies all the properties satisfied by $\mathbf{G}$ during the proof of Theorem 1 in [3].

Let $\mathcal{P}'$ be the orthogonal complement of $\mathcal{P}$ and $\mathcal{A}'$ be the orthogonal complement of $\mathcal{A}$. Let $\mathbf{Q_A R_A}$ be the skinny QR factorization of $\mathbf{A}^T$, whence $\mathbf{Q_P Q_P^T}=\Pi_{\mathcal{P}}$, is the orthogonal projector onto $\mathcal{P}$, $\mathbf{I}-\mathbf{C_P Q_P^T}=\Pi_{\mathcal{P}'}$, is the orthogonal projector onto $\mathcal{P}'$, and similarly for projectors $\Pi_{\mathcal{A}}$ and $\Pi_{\mathcal{A}'}$ formed from $\mathbf{Q_A}$. The construction of $\mathbf{A}$ is designed to provide one linearly independent row of the matrix for each column in $\mathbf{P}$ so that $\dim\mathcal{P}=\dim\mathcal{A}$. $\Pi_{\mathcal{P}}\mathbf{f}=\mathbf{Pz}$ holds as a general result of least squares solutions (1.1). $\Pi_{\mathcal{P}'}\mathbf{Pg}=\mathbf{0}$ trivially, since $\mathbf{Pg}\in\mathcal{P}$. $\mathbf{f}-\mathbf{Pg}\in\mathcal{A}'$ since $\mathbf{A}(\mathbf{f}-\mathbf{Pg})$ is equivalent to $\mathbf{A}(\mathbf{f}-\mathbf{PAf})$, and $\mathbf{A}$ is a left inverse of $\mathbf{P}$. Conversely, any $\mathbf{v}\in\mathcal{A}'$ can be expressed in the form $\mathbf{f}-\mathbf{PAf}$ for some $\mathbf{f}$, since we have simply arranged, in effect through our construction of $\mathbf{A}$, to solve $\mathbf{APg}=\mathbf{Af}$ in place of the normal equations corresponding to $\mathbf{P}$, where $\mathbf{AP}$ is nonsingular. This completes the check on $\mathbf{A}$. □

4 Examples

We offer a few examples to compare the effectiveness of using local least squares; that is, using $\mathbf{A}$ in place of $\left(\mathbf{P}^T\mathbf{P}\right)^{-1}\mathbf{P}^T$, on problems whose structures conform to the assumptions underlying our construction. The problems are

1. Data approximation using B-splines,
2. Data reduction by knot removal using B-splines,
3. Geometric simplification by reversal of butterfly subdivision (a subdivision not designed from B-splines).

For each example problem we have explored two issues:

a. For fixed locality; that is, for fixed and consistent choice of i_a and i_b with respect to each column j, how do the subspace angles and residuals compare between the local and the ordinary least squares methods of solution as a function of the size of $\mathbf{P}$? (The fixed locality needed to be relaxed somewhat in the case of problems with random structure – matrices of type iii below – but the amount by which deviation was allowed was kept as small as possible.)

b. For fixed size of $\mathbf{P}$, how do the subspace angles and residuals compare between the local and the ordinary methods of solution as the size of locality is increased; that is, as the number of rows in the selection $i_a \to i_b$ is increased?

If the locality size is increased as far as possible, we would eventually get $\mathbf{A} = \left(\mathbf{P}^T\mathbf{P}\right)^{-1}\mathbf{P}^T$, yielding the ordinary least squares solution. Our sample progressions are provided to give an impression of how quickly the solution to the local least squares problem approaches the full solution to the ordinary least squares problem as a function of the size of locality and the size of the problem.

Our example matrices are of three types:

i. Those that are cyclic and have identical columns resulting from interpolation or subdivison problems of regular structure,

ii. Those that are not cyclic but result from interpolation or subdivision problems of regular structure,

iii. Those constructed from data approximation or reduction problems of random structure.

In all example problems, the rows i_a and i_b have been chosen so that $\lambda = \mu$; that is, the intent was to site the nonzeros of column j centrally with respect to the submatrix $\mathbf{P}_j$. (This, too, necessarily had to be relaxed in the case of the beginning and ending columns of $\mathbf{P}$ matrices of types ii and iii.)

4.1 B-Spline Based Data Approximation

Our first collection of examples is taken from data sampling on the model

$$\sum_{i=0}^{L} c_i B_{i,k,\mathcal{Z}}(t_j) \approx f(t_j) \tag{4.1.1}$$

where $B_{i,k,\mathcal{Z}}$ are the B-splines of order k with the integers $\mathcal{Z}$ as knots [4], t_j are the evaluation points, f is the function to be approximated, and c_i are the unknowns. The i, j entries of $\mathbf{P}$ are given by the values $B_{i,k,\mathcal{Z}}(t_j)$, which can be calculated using the de Boor recurrence [4].

4.1.1 The Regular, Cyclic Case

One possible view of (4.1.1) yields an example that is regular and cyclic; that is, the t_j may be taken regularly spaced through the domain, and the functions may be taken as periodic. To generate regular, cyclic examples of (4.1.1), we took the t_j to be points of the form ℓ/λ (with integer ℓ modulo the periodicity in the domain and values of $\lambda = 2, 3, 4, 5, 6$), several orders of splines (values of $k = 3, 4, 5$), and various lengths of periodicity (accomplished by varying the numbers of columns in $\mathbf{P}$, since in the cyclic case the

number of rows is fixed by the shift σ and the number of columns). For example, a 6-column version of **P** for cubic B-splines ($k = 4$) and t_j of the form $\ell/2$ has entries:

$$\begin{bmatrix}
\frac{2}{3} & \frac{1}{6} & 0 & 0 & 0 & \frac{1}{6} \\
\frac{23}{48} & \frac{23}{48} & \frac{1}{48} & 0 & 0 & \frac{1}{48} \\
\frac{1}{6} & \frac{2}{3} & \frac{1}{6} & 0 & 0 & 0 \\
\frac{1}{48} & \frac{23}{48} & \frac{23}{48} & \frac{1}{48} & 0 & 0 \\
0 & \frac{1}{6} & \frac{2}{3} & \frac{1}{6} & 0 & 0 \\
0 & \frac{1}{48} & \frac{23}{48} & \frac{23}{48} & \frac{1}{48} & 0 \\
0 & 0 & \frac{1}{6} & \frac{2}{3} & \frac{1}{6} & 0 \\
0 & 0 & \frac{1}{48} & \frac{23}{48} & \frac{23}{48} & \frac{1}{48} \\
0 & 0 & 0 & \frac{1}{6} & \frac{2}{3} & \frac{1}{6} \\
\frac{1}{48} & 0 & 0 & \frac{1}{48} & \frac{23}{48} & \frac{23}{48} \\
\frac{1}{6} & 0 & 0 & 0 & \frac{1}{6} & \frac{2}{3} \\
\frac{23}{48} & \frac{1}{48} & 0 & 0 & \frac{1}{48} & \frac{23}{48}
\end{bmatrix}$$

For this class of examples, obviously, $\sigma = 2$. (For t_j having the form $\ell/3$, $\sigma = 3$, and a 6-column **P** would have 18 rows; for $\ell/4$, $\sigma = 4$, and 6 columns, we would have 24 rows; etc.. Such correlations show up in the way results are presented in Figure 9.)

A submatrix based upon 5 rows of **P** is

$$\begin{bmatrix}
\frac{1}{6} & \frac{2}{3} & \frac{1}{6} & 0 & 0 \\
\frac{1}{48} & \frac{23}{48} & \frac{23}{48} & \frac{1}{48} & 0 \\
0 & \frac{1}{6} & \frac{2}{3} & \frac{1}{6} & 0 \\
0 & \frac{1}{48} & \frac{23}{48} & \frac{23}{48} & \frac{1}{48} \\
0 & 0 & \frac{1}{6} & \frac{2}{3} & \frac{1}{6}
\end{bmatrix}$$

No widening or shrinking of the number of columns is necessary with this choice of i_a, i_b, and the submatrix is square and nonsingular rather than overdetermined, which is a rare but acceptable. The corresponding **A** is:

$$\begin{bmatrix}
\frac{10}{3} & -\frac{4}{3} & \frac{1}{6} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{6} & -\frac{4}{3} \\
\frac{1}{6} & -\frac{4}{3} & \frac{10}{3} & -\frac{4}{3} & \frac{1}{6} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & \frac{1}{6} & -\frac{4}{3} & \frac{10}{3} & -\frac{4}{3} & \frac{1}{6} & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & \frac{1}{6} & -\frac{4}{3} & \frac{10}{3} & -\frac{4}{3} & \frac{1}{6} & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{6} & -\frac{4}{3} & \frac{10}{3} & -\frac{4}{3} & \frac{1}{6} & 0 \\
\frac{1}{6} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{6} & -\frac{4}{3} & \frac{10}{3} & -\frac{4}{3}
\end{bmatrix}$$

that is, $\mathbf{a}_j$ has 5 entries:

$$\left[\begin{array}{ccccc} \frac{1}{6} & -\frac{4}{3} & \frac{10}{3} & -\frac{4}{3} & \frac{1}{6} \end{array}\right]$$

and the rows of $\mathbf{A}$ reflect the columns of $\mathbf{P}$ by shifting two positions right for each successive i.

If the submatrix $\mathbf{P}_j$ is expanded to 7 rows by setting i_a smaller and i_b larger by one, the next larger size for which $\lambda = \mu$, the resulting $\mathbf{a}_j$ is the same as above. We have observed this at times in regular problems, where one setting of locality and the next larger one yield the same results. This shows up, for example, in Figure 7, where for two pairs of localities the same minimum ratios are produced.

If the submatrix $\mathbf{P}_j$ is expanded to 9 rows with an additional row at the top and bottom, a new $\mathbf{a}_j$ results that has entries:

$$\left[\begin{array}{ccccccccc} -\frac{134}{3299} & \frac{1072}{3299} & -\frac{15997}{19794} & \frac{2884}{9897} & \frac{14498}{9897} & \frac{2884}{9897} & -\frac{15997}{19794} & \frac{1072}{3299} & -\frac{134}{3299} \end{array}\right]$$

Expansion to 11 rows yields the same result as this, but expansion to 13 rows provides a new $\mathbf{a}_j$ results that has entries:

$$\left[\begin{array}{ccccccc} \frac{609395}{35207268} & -\frac{1218790}{8801817} & \frac{3091228}{8801817} & -\frac{1572814}{8801817} & -\frac{19744189}{35207268} & \frac{1185728}{2933939} & \frac{3551866}{2933939} \\ \frac{1185728}{2933939} & -\frac{19744189}{35207268} & -\frac{1572814}{8801817} & \frac{3091228}{8801817} & -\frac{1218790}{8801817} & \frac{609395}{35207268} & \end{array}\right]$$

When t_j is of the form $\ell/4$, the nonzero entries of a column of $\mathbf{P}$ are

$$\left[\begin{array}{cccccccccccccccc} \frac{1}{384} & \frac{1}{48} & \frac{9}{128} & \frac{1}{6} & \frac{121}{384} & \frac{23}{48} & \frac{235}{384} & \frac{2}{3} & \frac{235}{384} & \frac{23}{48} & \frac{121}{384} & \frac{1}{6} & \frac{9}{128} & \frac{1}{48} & \frac{1}{384} \end{array}\right]^T$$

and the entries shift down four rows each successive column ($\sigma = 4$). Various choices of i_a and i_b yield various $\mathbf{a}_j$, of course, and the rows of $\mathbf{A}$ reflect the columns of $\mathbf{P}$ by shifting four positions right for each successive i.

4.1.2 The Regular Non-Cyclic Case

If (4.1.1) is not structured to have a periodic domain, yet the sample points t_j are still taken regularly spaced, then examples are generated that have the character of Section 4.1.1 for the interior columns of $\mathbf{P}$, but which begin and end with special columns. The number of rows in $\mathbf{P}$ is fixed, as in the cyclic case, once σ and the number columns is fixed. In turn, σ is fixed by ℓ/λ.

An 11×8 version of $\mathbf{P}$ generated from t_j having the form $\ell/2$ is

$$\begin{bmatrix}
\frac{1}{6} & \frac{2}{3} & \frac{1}{6} & 0 & 0 & 0 & 0 & 0 \\
\frac{1}{48} & \frac{23}{48} & \frac{23}{48} & \frac{1}{48} & 0 & 0 & 0 & 0 \\
0 & \frac{1}{6} & \frac{2}{3} & \frac{1}{6} & 0 & 0 & 0 & 0 \\
0 & \frac{1}{48} & \frac{23}{48} & \frac{23}{48} & \frac{1}{48} & 0 & 0 & 0 \\
0 & 0 & \frac{1}{6} & \frac{2}{3} & \frac{1}{6} & 0 & 0 & 0 \\
0 & 0 & \frac{1}{48} & \frac{23}{48} & \frac{23}{48} & \frac{1}{48} & 0 & 0 \\
0 & 0 & 0 & \frac{1}{6} & \frac{2}{3} & \frac{1}{6} & 0 & 0 \\
0 & 0 & 0 & \frac{1}{48} & \frac{23}{48} & \frac{23}{48} & \frac{1}{48} & 0 \\
0 & 0 & 0 & 0 & \frac{1}{6} & \frac{2}{3} & \frac{1}{6} & 0 \\
0 & 0 & 0 & 0 & \frac{1}{48} & \frac{23}{48} & \frac{23}{48} & \frac{1}{48} \\
0 & 0 & 0 & 0 & 0 & \frac{1}{6} & \frac{2}{3} & \frac{1}{6}
\end{bmatrix}$$

The version of $\mathbf{A}$ with 5 entries in the rows would be

$$\begin{bmatrix}
\frac{47}{6} & -\frac{44}{3} & \frac{34}{3} & -4 & \frac{1}{2} & 0 & 0 & 0 & 0 & 0 & 0 \\
-\frac{1}{2} & 4 & -\frac{11}{3} & \frac{4}{3} & -\frac{1}{6} & 0 & 0 & 0 & 0 & 0 & 0 \\
\frac{1}{6} & -\frac{4}{3} & \frac{10}{3} & -\frac{4}{3} & \frac{1}{6} & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & \frac{1}{6} & -\frac{4}{3} & \frac{10}{3} & -\frac{4}{3} & \frac{1}{6} & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & \frac{1}{6} & -\frac{4}{3} & \frac{10}{3} & -\frac{4}{3} & \frac{1}{6} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{6} & -\frac{4}{3} & \frac{10}{3} & -\frac{4}{3} & \frac{1}{6} \\
0 & 0 & 0 & 0 & 0 & 0 & -\frac{1}{6} & \frac{4}{3} & -\frac{11}{3} & 4 & -\frac{1}{2} \\
0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{2} & -4 & \frac{34}{3} & -\frac{44}{3} & \frac{47}{6}
\end{bmatrix}$$

where the first two and last two local least squares problems need to be distinct, of course, to achieve this.

4.1.3 The Irregular Non-Cyclic Case

For our irregular examples of this class of problems, we continued using the B-splines $B_{i,k,\mathbf{z}}(t)$, but we generated the sample points t_j uniformly throughout the domain, using enough samples so that there were, on average, λ samples per unit interval. This imitated the examples of Section 4.1.2 with sample points placed randomly rather than at the locations ℓ/λ. For such problems, the shift σ could vary slightly from column to column, as could the number of entries per column. We required that each local least squares subproblem of "width w" have at least enough rows in the submatrix $\mathbf{P}_j$ to include all the nonzero elements of the chosen column, or at least w rows containing these nonzero elements, whichever was larger. This allowed the local least squares problem to vary in size, but only if necessary, from column to column while always including at least w rows.

4.2 B-Spline Data Reduction

Our next collection of examples is taken from a form of approximation on a function subspace known variously as "data reduction", "reverse subdivision", and (in the case of splines) "knot removal". The context is one in which a function $\phi(t)$ is approximated by a linear combination of the basis functions for some space $\mathcal{V}^{k+1}$

$$\phi(t) \approx \sum_{\ell} f_{\ell} \beta_{\ell}^{k+1}(t)$$

and there is a proper subspace $\mathcal{V}^k$ on which one wishes to find an approximation to $\phi(t)$. Since any basis for $\mathcal{V}^k$ can be represented in terms of the basis for $\mathcal{V}^{k+1}$,

$$\beta_j^k = \sum_i p_{i,j} \beta_i^{k+1}$$

we know that

$$\phi(t) \approx \sum_j z_j \beta_j^k(t) = \sum_j z_j \sum_i p_{i,j} \beta_i^{k+1}(t) = \sum_i \left(\sum_j p_{i,j} z_j \right) \beta_i^{k+1}(t)$$

and from we take

$$\sum_j p_{i,j} z_j \approx f_i$$

as defining the form of the approximation we wish. This is, of course, exactly (1.1).

For any matrix **P** to be generated in this class of problems, the rows of **P** represent the *discrete B-splines* generated by the subspace's knot sequence and the refined knot sequence of the containing space. The discrete B-splines may be computed using the *Oslo algorithm*, which is a variation on the de Boor recurrence. [6].

4.2.1 The Regular, Cyclic Case

A convenient example is provided by taking β_i^{k+1} to be the periodic, uniform cubic B-splines with knots at the integers and half integers, while taking β_j^k to be the uniform cubic B-splines with knots at the integers

alone. A 12×6 version of the matrix $\mathbf{P}$ for this situation has entries:

$$\begin{bmatrix}
\frac{3}{4} & \frac{1}{8} & 0 & 0 & 0 & \frac{1}{8} \\
\frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 & 0 \\
\frac{1}{8} & \frac{3}{4} & \frac{1}{8} & 0 & 0 & 0 \\
0 & \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 \\
0 & \frac{1}{8} & \frac{3}{4} & \frac{1}{8} & 0 & 0 \\
0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 & 0 \\
0 & 0 & \frac{1}{8} & \frac{3}{4} & \frac{1}{8} & 0 \\
0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 \\
0 & 0 & 0 & \frac{1}{8} & \frac{3}{4} & \frac{1}{8} \\
0 & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} \\
\frac{1}{8} & 0 & 0 & 0 & \frac{1}{8} & \frac{3}{4} \\
\frac{1}{2} & 0 & 0 & 0 & 0 & \frac{1}{2}
\end{bmatrix}$$

The corresponding $\mathbf{A}$ based upon 5 rows of $\mathbf{P}$ is:

$$\begin{bmatrix}
2 & -\frac{1}{2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -\frac{1}{2} \\
0 & -\frac{1}{2} & 2 & -\frac{1}{2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & -\frac{1}{2} & 2 & -\frac{1}{2} & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & -\frac{1}{2} & 2 & -\frac{1}{2} & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & -\frac{1}{2} & 2 & -\frac{1}{2} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -\frac{1}{2} & 2 & -\frac{1}{2}
\end{bmatrix}$$

being composed from vectors $\mathbf{a}_j$ with entries: $\mathbf{a}_j = \begin{bmatrix} 0 & -\frac{1}{2} & 2 & -\frac{1}{2} & 0 \end{bmatrix}$. Using 7 rows provides the following $\mathbf{a}_j$:

$$\begin{bmatrix} \frac{23}{196} & -\frac{23}{49} & \frac{9}{28} & \frac{52}{49} & \frac{9}{28} & -\frac{23}{49} & \frac{23}{196} \end{bmatrix}$$

For 11 rows the $\mathbf{a}_j$ is:

$$\begin{bmatrix} -\frac{569}{12038} & \frac{1138}{6019} & -\frac{141}{926} & -\frac{2024}{6019} & \frac{4479}{12038} & \frac{5714}{6019} & \frac{4479}{12038} & -\frac{2024}{6019} & -\frac{141}{926} & \frac{1138}{6019} & -\frac{569}{12038} \end{bmatrix}$$

4.2.2 The Regular Non-Cyclic Case

If the domain is not wrapped around to produce periodic functions, then the first and last few columns of $\mathbf{P}$ become special. For example, if we continue to use B-splines with half-integral knots to define the containing space, and B-splines with integral knots to define the subspace (ℓ/λ with $\lambda = 2$), this produces

matrices **P** of the following character (cubics):

$$\begin{bmatrix}
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
\frac{1}{6} & \frac{5}{6} & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 & 0 & 0 \\
0 & \frac{1}{8} & \frac{3}{4} & \frac{1}{8} & 0 & 0 & 0 & 0 \\
0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 & 0 \\
0 & 0 & \frac{1}{8} & \frac{3}{4} & \frac{1}{8} & 0 & 0 & 0 \\
0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 \\
0 & 0 & 0 & \frac{1}{8} & \frac{3}{4} & \frac{1}{8} & 0 & 0 \\
0 & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 & 0 \\
0 & 0 & 0 & 0 & \frac{1}{8} & \frac{3}{4} & \frac{1}{8} & 0 \\
0 & 0 & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & \frac{5}{6} & \frac{1}{6} \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1
\end{bmatrix}$$

and a corresponding **A** designed to have at least 5 elements in its rows would be

$$\begin{bmatrix}
\frac{597}{599} & \frac{12}{599} & -\frac{25}{599} & \frac{20}{599} & -\frac{5}{599} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
-\frac{105}{599} & \frac{630}{599} & \frac{185}{599} & -\frac{148}{599} & \frac{37}{599} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & -\frac{47}{140} & \frac{47}{35} & \frac{17}{35} & -\frac{23}{35} & \frac{23}{140} & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & -\frac{1}{2} & 2 & -\frac{1}{2} & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & -\frac{1}{2} & 2 & -\frac{1}{2} & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & \frac{23}{140} & -\frac{23}{35} & \frac{17}{35} & \frac{47}{35} & -\frac{47}{140} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \frac{37}{599} & -\frac{148}{599} & \frac{185}{599} & \frac{630}{599} & -\frac{105}{599} \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -\frac{5}{599} & \frac{20}{599} & -\frac{25}{599} & \frac{12}{599} & \frac{597}{599}
\end{bmatrix}$$

4.2.3 The Irregular Non-Cyclic Case

For our random, irregular examples of this class of problems, we used the B-splines with integral knots, $B_{i,k,\mathcal{Z}}(t)$, to define the subspace; that is, with the unit intervals as the domains of spline segments. The containing space was generated by inserting new knots, taken randomly, uniformly-distributed throughout the domain so that, on average, λ new knots would be inserted in every unit interval for various $\lambda = 2, 3, 4 \ldots$. For such problems, the shift σ and the entries per column in **P** could vary slightly from column to column. Our concept of the "width *w*" of the local least squares problem followed that of Section 4.1.3 for examples of this type.

4.3 Butterfly Subdivision Reversal

The examples in this section follow exactly the same model as the one used in Section 4.2, except here basis functions are those generated by a **P**-defined recurrence. The recurrence is *Butterfly subdivison* [5], and our example comprises *reverse Butterfly subdivision*. A sample 12×6 cyclic version of **P** is:

$$\begin{bmatrix}
1 & 0 & 0 & 0 & 0 & 0 \\
\frac{9}{16} & \frac{9}{16} & -\frac{1}{16} & 0 & 0 & -\frac{1}{16} \\
0 & 1 & 0 & 0 & 0 & 0 \\
-\frac{1}{16} & \frac{9}{16} & \frac{9}{16} & -\frac{1}{16} & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 \\
0 & -\frac{1}{16} & \frac{9}{16} & \frac{9}{16} & -\frac{1}{16} & 0 \\
0 & 0 & 0 & 1 & 0 & 0 \\
0 & 0 & -\frac{1}{16} & \frac{9}{16} & \frac{9}{16} & -\frac{1}{16} \\
0 & 0 & 0 & 0 & 1 & 0 \\
-\frac{1}{16} & 0 & 0 & -\frac{1}{16} & \frac{9}{16} & \frac{9}{16} \\
0 & 0 & 0 & 0 & 0 & 1 \\
\frac{9}{16} & -\frac{1}{16} & 0 & 0 & -\frac{1}{16} & \frac{9}{16}
\end{bmatrix}$$

A trivial $\mathbf{a}_j$ can be based upon a single row of **P** centered on the unit entries: $\mathbf{a}_j = [1]$. Interestingly, widening the local problem to 3, 5, and 7 rows all give the same result: $\mathbf{a}_j = [1]$. Increasing the locality to include 9 rows of **P** produces the following $\mathbf{a}_j$:

$$\begin{bmatrix} \frac{3}{161} & 0 & -\frac{24}{161} & \frac{48}{161} & \frac{107}{161} & \frac{48}{161} & -\frac{24}{161} & 0 & \frac{3}{161} \end{bmatrix}$$

This same solution results with a locality of 11 rows. Increasing the locality to include 13 rows of **P** produces the following $\mathbf{a}_j$:

$$\begin{bmatrix} -\frac{148}{33553} & 0 & \frac{1971}{33553} & -\frac{2368}{33553} & -\frac{3780}{33553} & \frac{10224}{33553} & \frac{21755}{33553} & \frac{10224}{33553} & -\frac{3780}{33553} & -\frac{2368}{33553} & \frac{1971}{33553} & 0 & -\frac{148}{33553} \end{bmatrix}$$

Only regular, cyclic Butterfly examples were generated.

4.4 Results

We confine ourselves to showing only a sampling of results that are representative of those obtained from all our various examples. Recall that γ is the minimum possible ratio of the least squares ratio to the local least squares ratio, and the closer γ comes to 1, the closer is the local approximation to the standard least squares approximation (our examples are all well conditioned, and so residuals really do reflect the quality of the approximation).

4.4.1 Random Non-Cyclic Examples

The quality of the local approximations varied most strongly in the random examples. Both classes of examples, the B-spline data approximation and B-spline data reduction problems, showed the same behavior for all orders of spline; so for brevity, we present only a sample of this behavior for quadratic B-spline data reduction. Figure 2 shows an overview: a smooth surface fit to the data of roughly 4000 randomly generated test examples. Each test contributing to this figure had at least 40 rows and 10 columns, with 11 rows being taken as the (lower bound) width of locality. For each row-column pair, 30 test examples were generated and their minimum ratios (γ) were averaged. The width setting of $i_b - i_a + 1 = 11$, for the random problems of course, constituted only a minimum for the locality actually used on any randomly generated example. The nonzeros in any column varied in number, and each submatrix may have been expanded or contracted as necessary to yield a viable local least squares problem. The averaging process was used to smooth out such variations.

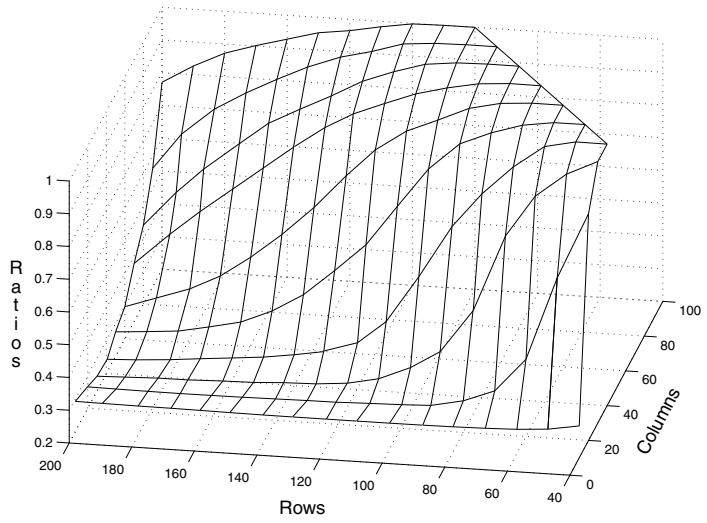

Figure 2: γ vs. problem size for 11-row locality

To provide a closer look, Figure 3 shows the variation of minimum ratios, for 11-row (lower bound) locality with respect to number of rows in the problem taken for those problems generated with 70 columns, and with respect to number of columns taken for problems with 90 rows. In other words, Figure 3 displays two slices through the data that yielded the surface of Figure 2.

As the minimum setting of that width is increased, the ratios approach 1, since the subproblems approach the full least squares problem. Figure 4 shows this variation on problems having 250 rows and 50 columns with widths varying from 5 to 50. The points, as in the earlier figures, are averages of 30 randomly generated examples.

It is not surprising that a "fixed" locality (i.e., within slight variation) yields poorer results as the number of rows in a least squares approximation is increased, as the left-hand plot of Figure 3 shows, since the subproblems will depart further and further from the entire problem. It is also intuitive that the results will become better as the width of locality is increased, as is seen in Figure 4, since this implies that fewer subproblems of larger sizes are employed, tending to one subproblem comprising the entire

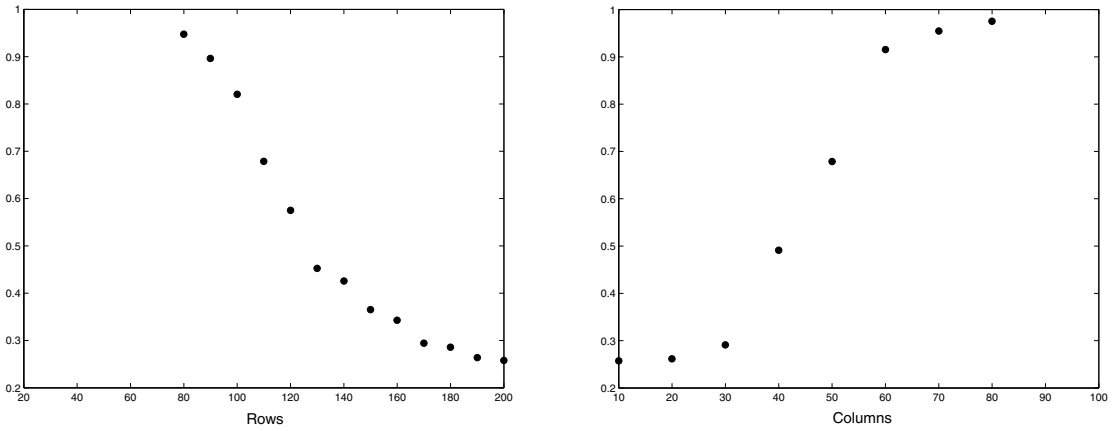

γ vs. number of rows γ vs. number of columns

Figure 3: γ vs. problem size for 11-row (minimum) locality

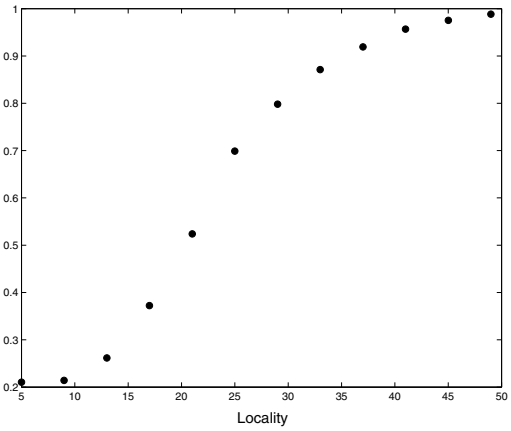

Figure 4: γ vs. width of locality on a 250×50 problem

least squares approximation. What is surprising is that, as the number of columns is increased for a fixed number of rows, the locally-based approximation for a "fixed" locality becomes better and better, as is seen in the right-hand plot of Figure 3. We need to study further whether this is an artifact of the way we are generating random problems, or whether it says something fundamental about local least squares.

4.4.2 Regular Non-Cyclic Examples

The situation for regular problems, cyclic as well as non-cyclic, becomes much simpler. In both the cyclic and the non-cyclic cases we would, of course, expect no independence of results between numbers of columns and numbers of rows in **P**, since the number of rows is strictly determined by the number of columns for each instance of σ, depending in turn on parameters such as k or λ. The astonishing

observation, however, was that there was almost no dependence whatsoever on the size of **P**. For non-cyclic problems, as the number of columns grew (and the population of the few special columns at the beginning and the end decreased relative to the generic interior columns), the minimum ratio approached the minimum ratio of the corresponding cyclic problem (consisting entirely of interior columns). Even for small numbers of columns, the difference between minimum ratios in corresponding cyclic and non-cyclic examples was quite small. Figure 5 shows an example of this for quadratic B-spline data reduction (arising from knot doubling in the uniform case).

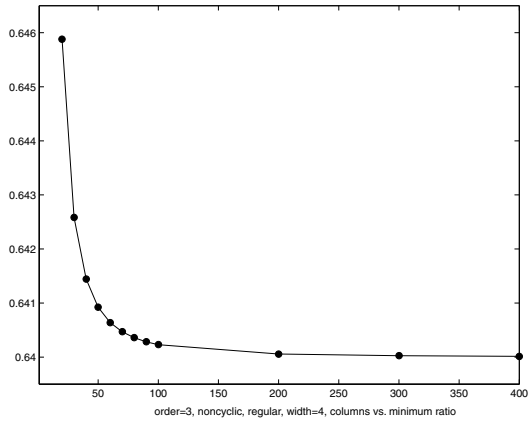

Figure 5: γ vs. number of columns for uniform, quadratic, B-spline refinement

This being the case, we show results for quadratic B-spline data reduction in Figure 6, which are typical of all the results we saw. The axis relating to the number of rows in **P** has the label "`Dimension Increase`".

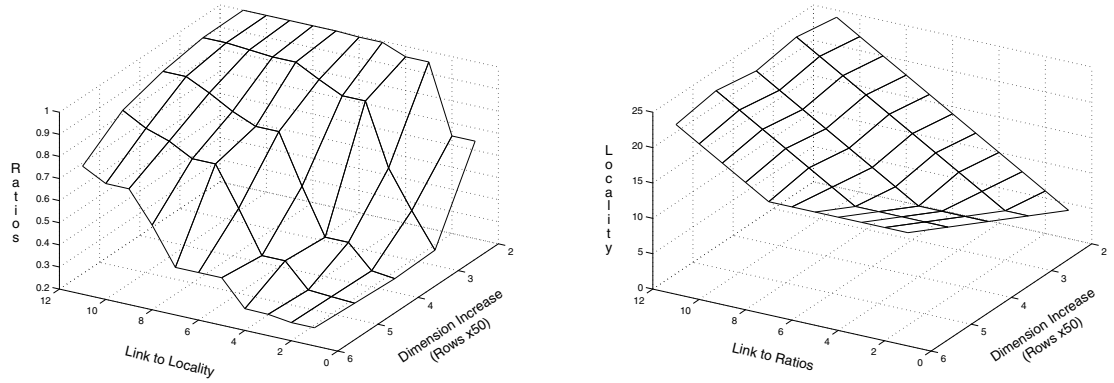

Figure 6: γ vs. matrix and locality

The matrix **P** is generated by inserting regularly-spaced knots into the parameter range of a basis of uniform B-splines. Inserting a single new knot at the midpoints of the existing knot intervals will effectively double the number of B-splines and raise the dimension of the spline space by two. In matrix

terms, this will give **P** effectively twice as many columns as rows. Inserting two new knots per interval will effectively triple the dimensionality of the spline space, tripling the ratio of rows to columns, and so on. The number of insertions will, of course, also influence the number of nonzeros in any column. One additional knot per interval corresponds to four nonzeros per column. Two additional knots corresponds to seven nonzeros per column, and so on. The actual problem sizes represented by Figure 6 ranged from 18×10 through 982×100, with nonzeros per column ranging from 4 to 19.

For any of these problem sizes, we wished to explore a range of sizes of locality. However, a desired locality for any size of problem may have been adjusted as necessary to obtain full-rank subproblems, in the interior of the matrix, due to the shift and the numbers of nonzeros in a column, and at the first and last columns of the matrix in the non-cyclic problems, due to the special character of these boundary columns. But due to the regularity of these problems, the actual width of locality that resulted from a width that was specified and subsequently adjusted was entirely deterministic and could be recorded along with the results. The widths of locality are presented in the right-hand plot of Figure 6 in a way that is coordinated with the minimum ratios shown in the left-hand plot. The coordination is through the axis with the label "`Link to Locality`" in the left-hand plot and with the label "`Link to Ratios`" in the right-hand plot. All sizes of problems generated by inserting 4 additional knots, for example, correspond to the central line on the base planes bisecting the "`Dimension Increase`" axis. Those producing a minimum ratio of around 0.8, for example (corresponding approximately to position 6 on the "`Link to Locality`" axis) were achieved with a locality width shown at the 4-6 intersection in the right-hand plot (between 10 and 15 rows in the local subproblems).

As in the case of fully random problems, the approximations become better as the width of locality is increased. Figure 7 shows an example of this. To stress what Figure 7 implies, for the regular, data

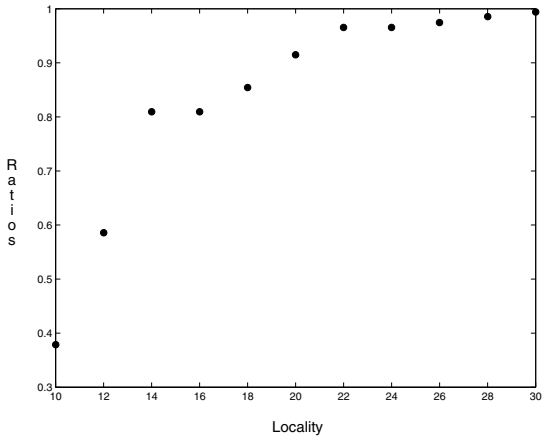

Figure 7: γ vs. width of locality for 4-knot insertion on quadratic B-splines for any size of **P**

reduction problem shown, a least squares problem of any size can be solved to within $\gamma \approx 0.99$ by solving *a few subproblems* with only 28 rows (one subproblem for each special column at the beginning to find special first and last rows of **A**, and one subproblem to obtain all the interior rows of **A**).

4.4.3 Regular, Cyclic Examples

The characteristic feature of all regular, cyclic examples is that there are no special columns; and because the number of rows is determined by the number of columns and the shift σ from one column to the next, it is not possible to vary the number of rows and columns independently, as was also the case in the problems of previous section.

There is typically a slight variation in minimum ratios for the cyclic problems. For large problems, the ratios asymptote to a characteristic value for the problem, but they do this in a way that is determined by the width of locality (the number of rows in the local least squares problem). Figure 8 shows examples of this. The left plot in Figure 8 is taken from the data reduction problem derived by knot doubling in uniform quadratic B-splines with a 4-row local least squares subproblem. The right plot in Figure 8 is the corresponding situation for 8-row local least squares subproblems. The variations in minimum ratios are not great. The 4-row locality shows a 4-cycle variation in minimum ratios, while the 8-row locality shows an 8-cycle variation.

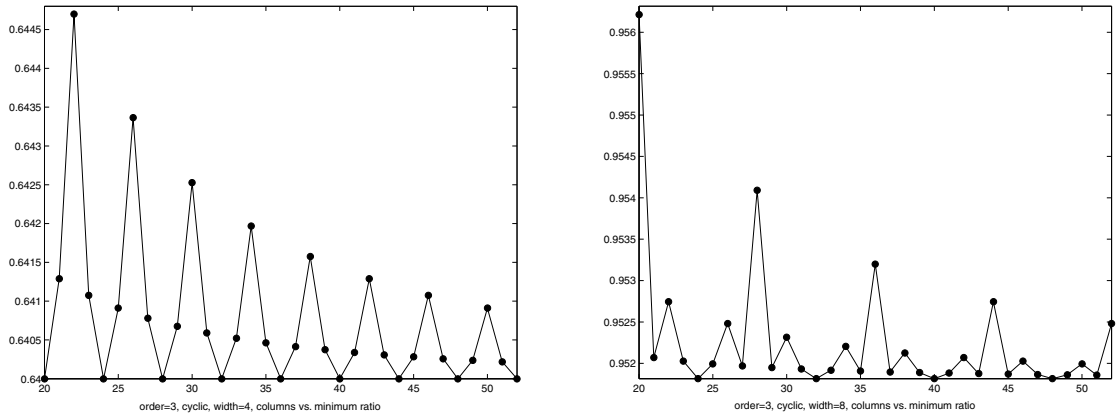

4 width locality

8 width locality

Figure 8: γ vs. number of columns in uniform, quadratic B-spline knot doubling

Results, corresponding to Figure 6 in Section 4.4.2, for a representative cyclic case covering a span of problem sizes is shown in Figure 9 using the same linking conventions as for Figure 6.

4.5 Projector Alignment and Subspace Angles

We have observed in the foregoing examples that the minimum ratio depends only on the size of locality in regular problems. As a consequence, a problem with 50 columns, for example, will have almost exactly the same smallest singular value associated with **A** and **P** as a problem with 5000 columns. This implies that the other 4999 singular values must be sandwiched into the same span of values as were the 49 singular values of the smaller problem, suggesting that the projectors **PA** and $\mathbf{P}(\mathbf{P}^T\mathbf{P})^{-1}\mathbf{P}^T$ come, in a relative sense, more into alignment as **P** grows larger. The following theorem shows how the cosines of the angles between the column subspaces of **P** and **A**, which measure this alignment, provide a comparison between the least squares solution for (1.1), **z**, and the local least squares solution, **g**. These subspace angles are provided by singular values associated with **P** and **A** together:

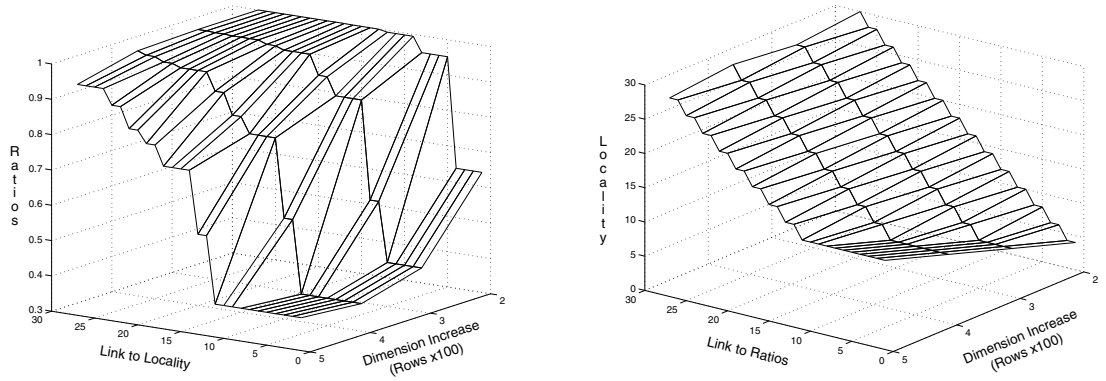

Ratios Localities

Figure 9: γ vs. matrix and locality

Theorem 3 *Let the "skinny" QR factorizations of* $\mathbf{P}$ *and* $\mathbf{A}^T$ *be* $\mathbf{Q_P R_P}$ *and* $\mathbf{Q_A R_A}$ *respectively. Let* $\mathbf{UDV}$ *be the singular value decomposition of* $\mathbf{Q_P^T Q_A}$. *Then*

$$\begin{aligned} \mathbf{U}^T\mathbf{R_P z} &= \mathbf{DVQ_A^T f} \\ \mathbf{U}^T\mathbf{R_P g} &= \mathbf{D}^{-1}\mathbf{VQ_A^T f} \end{aligned}$$

Proof: Note that $\mathbf{U}$, $\mathbf{D}$, and $\mathbf{V}$ are $n \times n$ square matrices. Note also that

$$\mathbf{UDVQ_A^T} = \mathbf{Q_P^T}$$

From the least squares problem we have

$$\begin{aligned} \mathbf{Pz} &\approx \mathbf{f} \\ \mathbf{P}^T\mathbf{Pz} &= \mathbf{P}^T\mathbf{f} \\ \mathbf{R_P^T Q_P^T Q_P R_P z} &= \mathbf{R_P^T Q_P^T f} \\ \mathbf{R_P z} &= \mathbf{Q_P^T f} \\ \mathbf{R_P z} &= \mathbf{UDVQ_A^T f} \\ \mathbf{U}^T\mathbf{R_P z} &= \mathbf{DVQ_A^T f} \end{aligned}$$

From the local least squares problem we have

$$\begin{aligned} \mathbf{Pg} &\approx \mathbf{f} \\ \mathbf{APg} &= \mathbf{Af} \\ \mathbf{R_A^T Q_A^T Q_P R_P g} &= \mathbf{R_A^T Q_A^T f} \\ \mathbf{V}^T\mathbf{DU}^T\mathbf{R_P g} &= \mathbf{Q_A^T f} \\ \mathbf{U}^T\mathbf{R_P g} &= \mathbf{D}^{-1}\mathbf{VQ_A^T f} \end{aligned}$$

□

This theorem gives another view of the fact that $\mathbf{g} \rightarrow \mathbf{z}$ only as the singular values of $\mathbf{Q}_\mathbf{P}^T\mathbf{Q}_\mathbf{A}$ approach 1; that is, only as the angles between the subspaces $\mathcal{P}$ and $\mathcal{A}$ all go to zero; that is, as the projectors $\mathbf{PA}$ and $\mathbf{P}(\mathbf{P}^T\mathbf{P})^{-1}\mathbf{P}^T$ become aligned. This alignment, in turn, progresses as the width of the locality is increased. We show this in Figures 10 and 11.

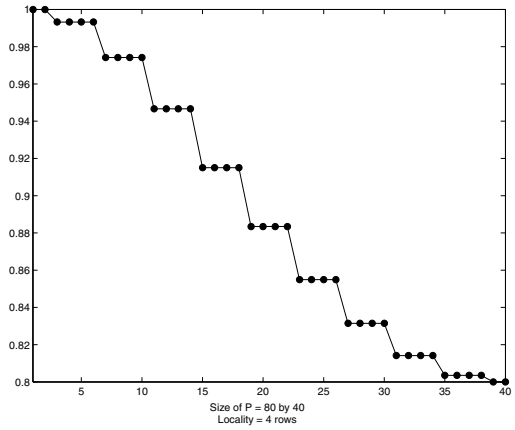

Width 4 for 40 columns

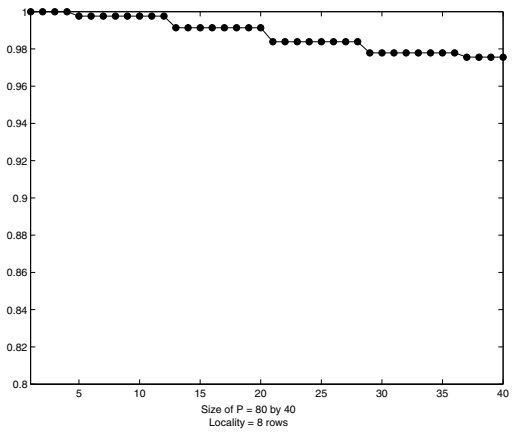

Width 8 for 40 columns

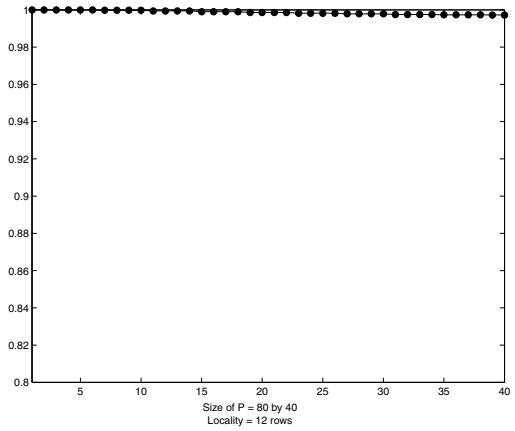

Width 12 for 40 columns

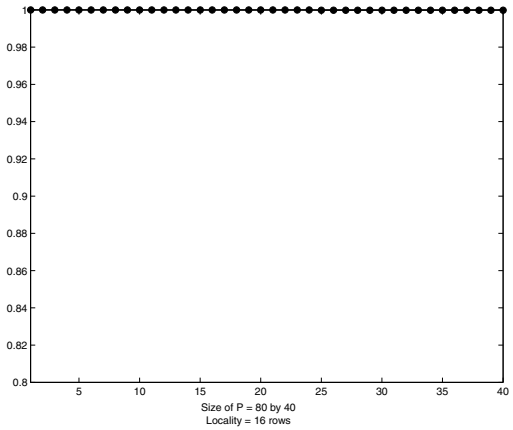

Width 16 for 40 columns

Figure 10: Cosine angles in quadratic B-spline data reduction as the width of locality is increased

The way in which the subspaces for $\mathbf{A}$ and $\mathbf{P}(\mathbf{P}^T\mathbf{P})^{-1}\mathbf{P}^T$ align, shown by the subspace angles whose example cosines are being shown, suggests that the minimum ratio may be pessimistic. If a right-hand vector $\mathbf{f}$ is rich in components of the columns of $\mathbf{Q}_\mathbf{A}$ associated with the higher singular values, we might expect that $\mathbf{A}$ will more closely approximate the behavior of $\mathbf{P}(\mathbf{P}^T\mathbf{P})^{-1}\mathbf{P}^T$ on that $\mathbf{f}$. From the representative graphs of the subspace cosines we have given, we see that relatively many of the singular values are higher in value, and the smaller values are represented in only a few subspace components. This might suggest that random right-hand vectors would provide better ratios than the minimum ratio. This sugges-

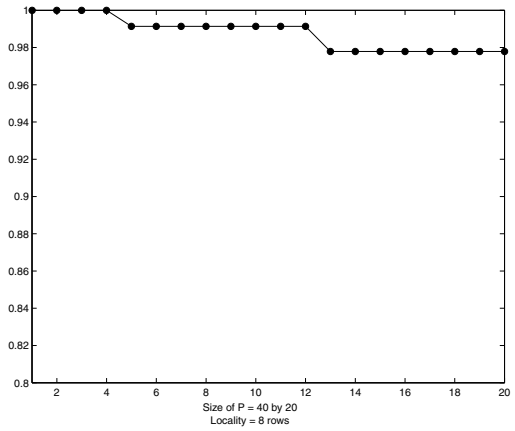

Width 8 for 20 columns

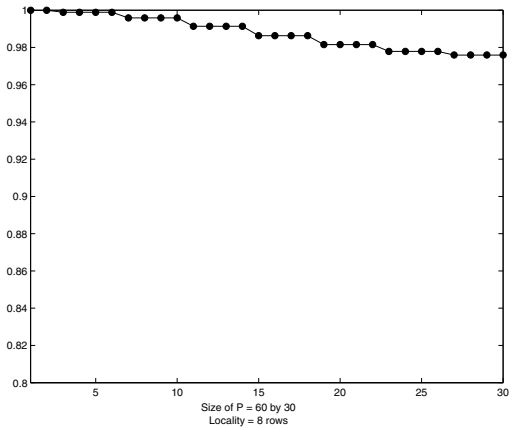

Width 8 for 30 columns

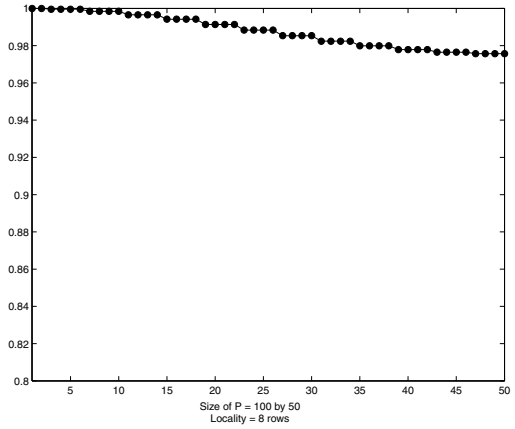

Width 8 for 50 columns

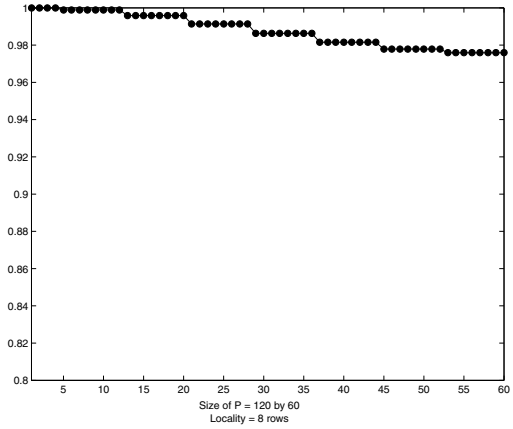

Width 8 for 60 columns

Figure 11: Cosine angles in quadratic B-spline data reduction as the size of problem is increased

tion is given some support in Figure 12, which shows a histogram for the residual ratios found for 200 randomly-generated (components uniformly distributed in $[0,1]$) vectors $\mathbf{f}$ in a regular, cyclic problem of data-reduction type for midpoint-knot-insertion in quadratic B-splines.

5 Constructing A on Meshes

We close with a brief mention of how these local approaches may be extended to meshes; e.g., for possible use in finite element settings. Further details may be found in [9].

The local least squares construction outlined in Section 2 can be described in an equivalent way. Referring to Figure 1: the row $\mathbf{a}_j$ of $\mathbf{A}$, which we extract from the subproblem $\mathbf{P}_j$ formed from rows i_a

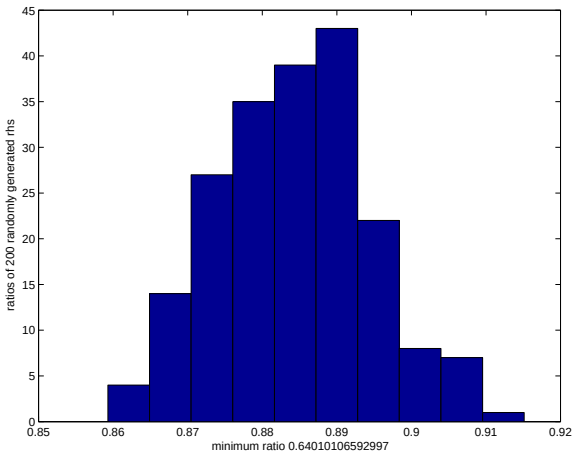

Figure 12: Histogram of actual residual ratios for 200 random right-hand sides in quadratic, B-spline data reduction

through i_b, can be generated equivalently by solving

$$\mathbf{a}_j \mathbf{P}_j = [0, \ldots, 0, 1, 0, \ldots, 0] \tag{5.1}$$

where the entry 1 in the right hand vector corresponds in position to the row position to be extracted from $(\mathbf{P}^T\mathbf{P})^{-1}\mathbf{P}^T$. This equivalence is established in [2].

Starting from this point of view, it is only necessary to generate a few equations, amounting to the interaction (inner product) of the unknown vector $\mathbf{a}_j$ with the nonzero elements of a few columns of $\mathbf{P}$. In the setting of a mesh whose nodes represent the unknowns $\mathbf{f}$ in (1.1), any column j of $\mathbf{P}$ can be represented in terms of a stencil on that mesh that shows how the j^{th} element of $\mathbf{z}$ contributes to the components of $\mathbf{f}$. The stencil of Figure 13 gives an example for $\mathbf{P}$ defining a subdivision due to Litke *et. al* [8] on a regular, hexagonal mesh. In this figure, the component of $\mathbf{f}$ represented by the middle node is defined by $\frac{3}{2}$ times z_j plus contributions from other values of $\mathbf{z}$, the component of $\mathbf{f}$ represented by the node directly to the "north" is defined by $\frac{3}{8}$ times z_j plus contributions from other values of $\mathbf{z}$, and so forth.

Corresponding to such "column stencils" for $\mathbf{P}$ we may define "row stencils" for the rows of $\mathbf{A}$. Such a stencil describes how the inner product $\mathbf{a}_j \cdot \mathbf{f}$ is formed by summing the values of components of $\mathbf{f}$ with the components of the stencil as coefficients. For example, Figure 15 defines the inner product

$$\frac{5}{9} f_{i_0} + \frac{2}{27} f_{i_1} + \frac{2}{27} f_{i_2} + \frac{2}{27} f_{i_3} + \frac{2}{27} f_{i_4} + \frac{2}{27} f_{i_5} + \frac{2}{27} f_{i_6}$$

for some indexing i_μ of the mesh. With these preliminaries, the equations of (5.1) are found through the interactions of stencils such as is shown in Figure 14, which represents the equation

$$a_5 p_{24} + a_6 p_{23} = 0$$

(The right hand side would be 1 only when the stencil for $\mathbf{a}_j$ interacts with the column stencil from $\mathbf{P}_j$ corresponding to the particular row of $\mathbf{A}$ to be extracted.)

Two examples of $\mathbf{A}$ are embodied in Figures 15 and 16 representing two different "widths" of locality. For the "shorter" of these two versions of $\mathbf{A}$, that shown in Figure 15, we have carried out the minimum

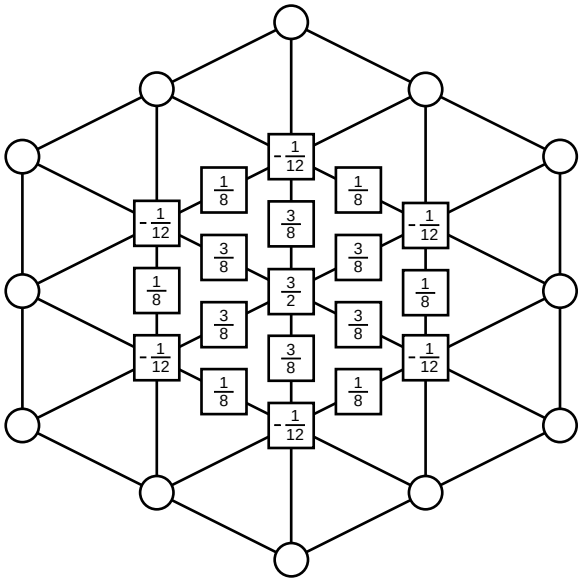

Figure 13: Litke *et. al* subdivision stencil (a column of **P**)

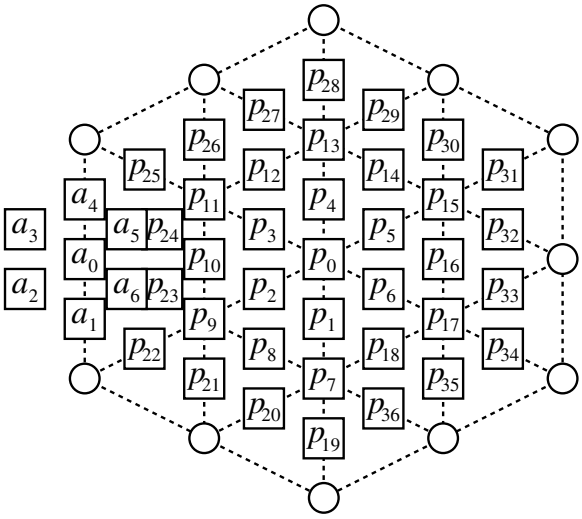

Figure 14: Interaction of an **A** and a **P** stencil

ratio analysis. There is no way of creating a cyclic topology on the hexagonal mesh, so we had to resort to a regular mesh with boundaries. Since the effects of the boundary will diminish as the mesh grows larger, providing an approach to the minimum ratio of an infinite mesh, we chose to take the simplest boundary conditions we could find; namely, reflection across the boundary. If any exterior node is required by **P** or **A**, the corresponding interior node found by reflecting across the boundary is used. This yields the stencils shown in Figure 17 through Figure 23. In addition to the interior stencils of **P** and **A** shown in Figure 13 and Figure 15 respectively, there are five special boundary and near boundary **P** stencils (Figures 17 through 21) and two special boundary **A** stencils (Figures 22 and 23).

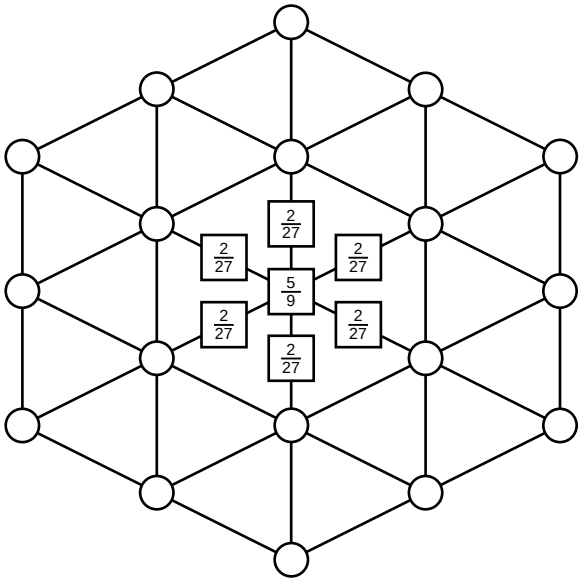

Figure 15: A minimal-width local inverse (a row of **A**)

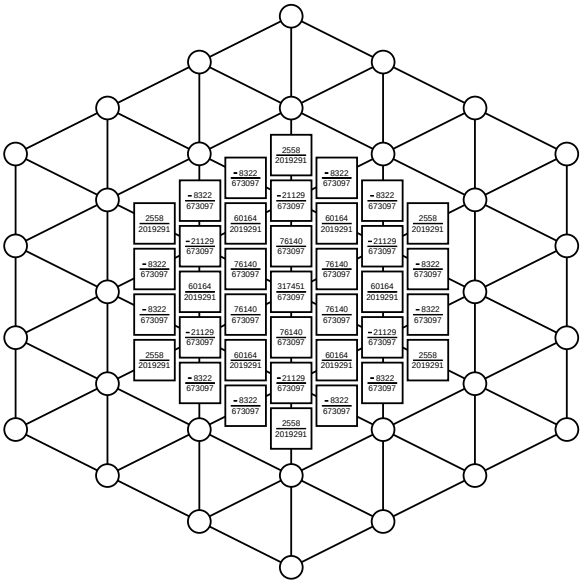

Figure 16: A larger local inverse

Figure 24 shows the minimum ratio for a number of mesh sizes, measured in “Cycles”; that is, in numbers of rings of nodes around the central node of the mesh. Finally, Figure 25 shows the subspace cosines for a mesh of 7 cycles.

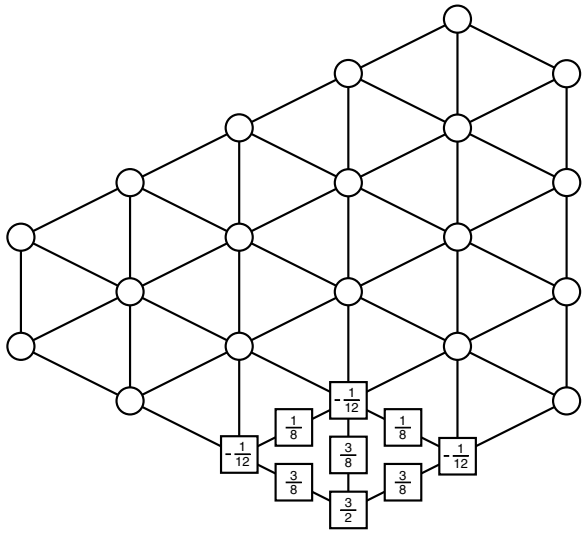

Figure 17: Boundary corner **P** stencil, resulting from reflection

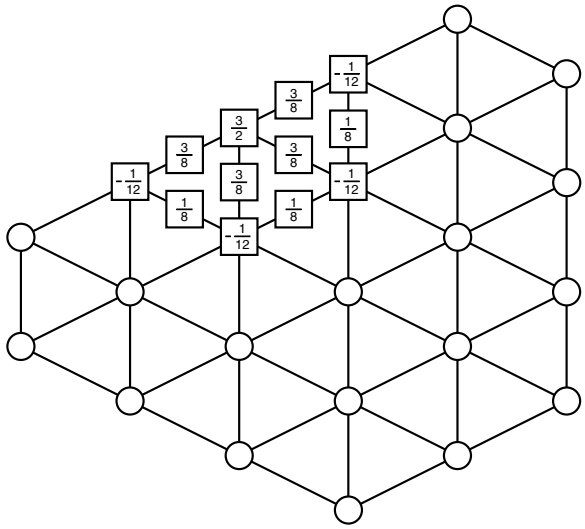

Figure 18: Boundary edge interior **P** stencil, resulting from reflection

Acknowledgments

This work was carried out under funding by the Natural Sciences and Engineering Research Council of Canada and by the National Science Foundation of the U. S.

References

[1] N. Anderson. A generalization of the method of averages for overdetermined linear systems. *Mathematics of Computation*, 29(130):607–614, April 1975.

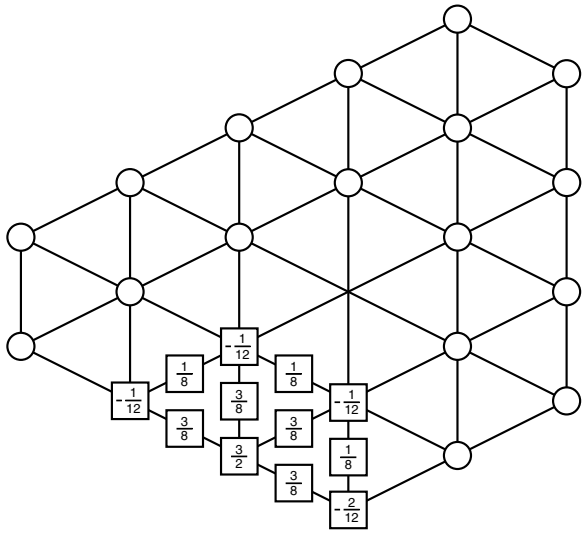

Figure 19: Boundary edge **P** stencil adjacent to a corner, resulting from reflection

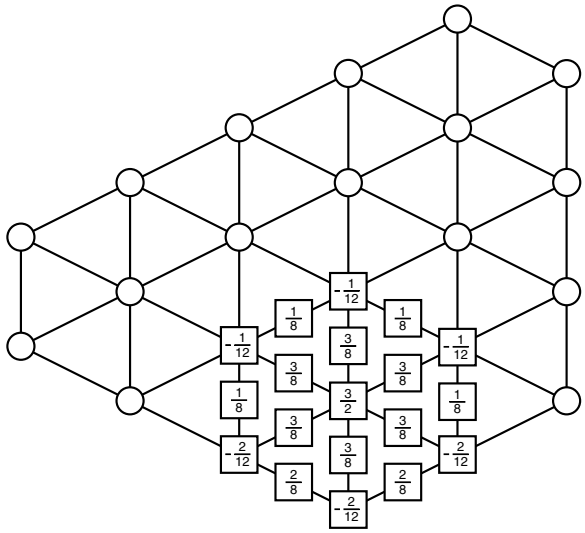

Figure 20: Near boundary **P** stencil adjacent to a corner, resulting from reflection

[2] R. H. Bartels and F. F. Samavati. Reversing subdivision rules: Local linear conditions and observations on inner products. *Journal of Computational and Applied Mathematics*, 119(1–2):29–67, July 2000.

[3] G. Dahlquist, B. Sjöberg, and P. Svensson. Comparison of the method of averages with the method of least squares. *Mathematics of Computation*, 22(104):833–845, October 1968.

[4] C. de Boor. *A Practical Guide to Splines*. Applied Mathematical Sciences 27. Springer-Verlag, 1978.

[5] N. Dyn, D. Levin, and J. Gregory. A 4-point interpolatory subdivision scheme for curve design. *Computer Aided Geometric Design*, 4(4):257–268, 1988.

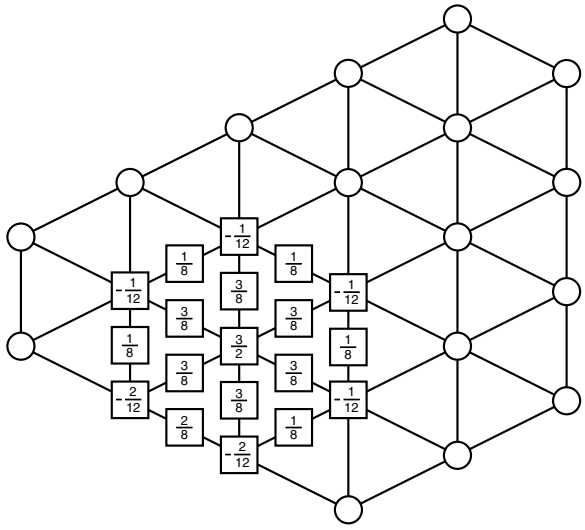

Figure 21: Near boundary **P** stencil adjacent to an edge, resulting from reflection

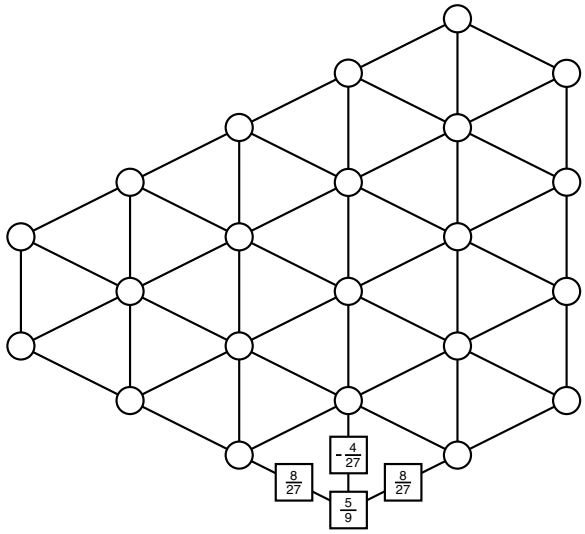

Figure 22: Boundary corner **A** stencil, resulting from reflection

[6] R. N. Goldman and T. Lyche, editors. *Knot Insertion and Deletion Algorithms for B-Spline Curves and Surfaces*. SIAM, 1993.

[7] G. H. Golub and C. F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, second edition, 1989.

[8] N. Litke, J. Levin, and P. Schröder. Trimming for subdivision surfaces. *Computer Aided Geometric Design*, 18(5):463–481, June 2001.

[9] F. F. Samavati and R. H. Bartels. Reversing subdivision using local linear conditions: Generating multiresolutions on regular triangular meshes. Technical report, University of Calgary, Computer

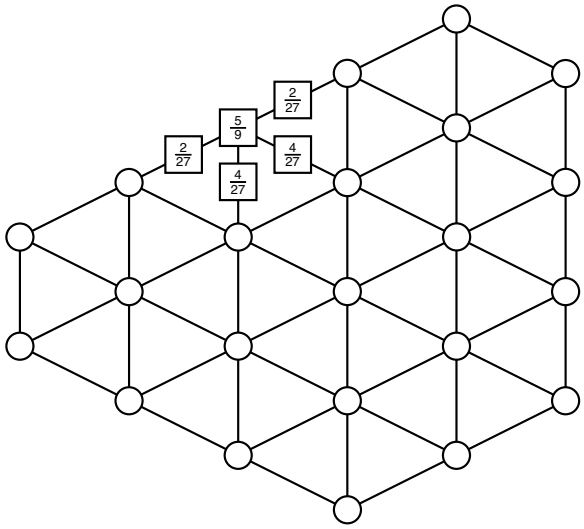

Figure 23: Boundary edge **A** stencil, resulting from reflection

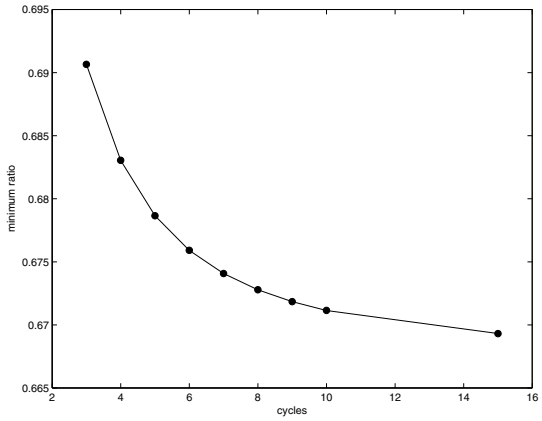

Figure 24: Minimum ratios for several sizes of regular, hexagonal mesh

Science, TR 2002-711-14, December 2002.

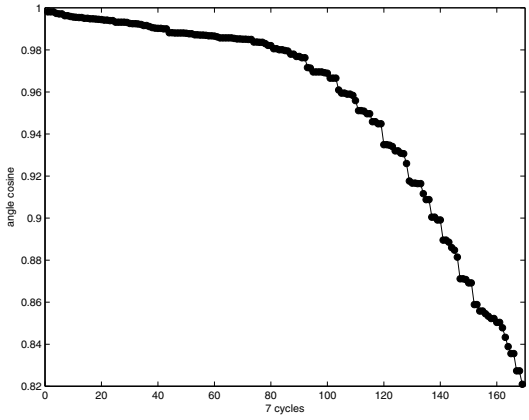

Figure 25: Subspace cosines for a representative mesh