11-22-95

SANDIA REPORT

SAND95-1559 • UC-405

Unlimited Release

Printed October 1995

Aztec User's Guide

Version 1.0

Scott A. Hutchinson, John N. Shadid, Ray S. Tuminaro

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550
for the United States Department of Energy
under Contract DE-AC04-94AL85000

Approved for public release; distribution is unlimited.

SF2900Q(8-81)

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

21

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
Office of Scientific and Technical Information
PO Box 62
Oak Ridge, TN 37831

Prices available from (615) 576-8401, FTS 626-8401

Available to the public from
National Technical Information Service
US Department of Commerce
5285 Port Royal Rd
Springfield, VA 22161

NTIS price codes
Printed copy: A04
Microfiche copy: A01

Aztec User's Guide* Version 1.0

Scott A. Hutchinson[†] John N. Shadid[§] Ray S. Tuminaro[‡]
Massively Parallel Computing Research Laboratory
Sandia National Laboratories
Albuquerque, NM 87185

Abstract

Aztec is an iterative library that greatly simplifies the parallelization process when solving the linear systems of equations $Ax = b$ where A is a user supplied $n \times n$ sparse matrix, b is a user supplied vector of length n and x is a vector of length n to be computed. **Aztec** is intended as a software tool for users who want to avoid cumbersome parallel programming details but who have large sparse linear systems which require an efficiently utilized parallel processing system. A collection of data transformation tools are provided that allow for easy creation of distributed sparse unstructured matrices for parallel solution. Once the distributed matrix is created, computation can be performed on any of the parallel machines running **Aztec**: nCUBE 2, IBM SP2 and Intel Paragon, MPI platforms as well as standard serial and vector platforms.

Aztec includes a number of Krylov iterative methods such as conjugate gradient (CG), generalized minimum residual (GMRES) and stabilized biconjugate gradient (BiCGSTAB) to solve systems of equations. These Krylov methods are used in conjunction with various preconditioners such as polynomial or domain decomposition methods using LU or incomplete LU factorizations within subdomains. Although the matrix A can be general, the package has been designed for matrices arising from the approximation of partial differential equations (PDEs). In particular, the **Aztec** package is oriented toward systems arising from PDE applications.

* This work was supported by the Applied Mathematical Sciences program, U.S. Department of Energy, Office of Energy Research, and was performed at Sandia National Laboratories, operated for the U.S. Department of Energy under contract No. DE-AC04-94AL85000. The **Aztec** software package was developed by the authors at Sandia National Laboratories and is under copyright protection

[†] Parallel Computational Sciences Department; sahutch@cs.sandia.gov; (505) 845-7996

[§] Parallel Computational Sciences Department; jnshadi@cs.sandia.gov; (505) 845-7876

[‡] Applied & Numerical Mathematics Department; tuminaro@cs.sandia.gov; (505) 845-7298

Contents.

1	Overview	1
2	Aztec: High Level View	1
2.1	Aztec Options	2
2.2	Aztec parameters	6
2.3	Return status	7
3	Data Formats	7
3.1	Distributed Modified Sparse Row (DMSR) Format	9
3.2	Distributed Variable Block Row (DVBR) Format	9
4	High Level Data Interface	10
5	Examples	12
6	Advanced Topics	15
6.1	Data Layout	15
6.2	Reusing factorizations	17
6.3	Important Constants	18
6.4	AZ_transform Subtasks	18
7	Aztec Functions	19
	AZ_broadcast	20
	AZ_check_input	21
	AZ_check_msr	21
	AZ_check_vbr	22
	AZ_defaults	22
	AZ_exchange_bdry	23
	AZ_find_index	23
	AZ_find_local_indices	24
	AZ_find_procs_for_externs	24
	AZ_free_memory	25
	AZ_gavg_double	25
	AZ_gdot	25
	AZ_gmax_double	26
	AZ_gmax_int	26
	AZ_gmax_matrix_norm	26
	AZ_gmax_vec	27
	AZ_gmin_double	27
	AZ_gmin_int	27
	AZ_gsum_double	28
	AZ_gsum_int	28
	AZ_gsum_vec_int	28
	AZ_gvector_norm	29
	AZ_init_quick_find	29
	AZ_matvec_mult	30
	AZ_msr2vbr	31

AZ_order_ele	31
AZ_print_error	32
AZ_processor_info	33
AZ_quick_find	33
AZ_read_msr_matrix	34
AZ_read_update	34
AZ_reorder_matrix	35
AZ_set_message_info	36
AZ_solve	37
AZ_sort	37
AZ_transform	38

Notation Conventions

Different fonts are used to indicate program fragments, keys words, variables, or parameters in order to clarify the presentation. The table below describes the meaning denoted by these different fonts.

Convention	Meaning
<code>typewriter</code>	File names, code examples and code fragments.
<code>sans serif</code>	C language elements such as function names and constants when they appear embedded in text or in function definition syntax lines.
<i>italics</i>	Parameter and variable names when they appear embedded in text or function definition syntax lines.
<code>AZ_</code>	C language elements such as function names and constants which are supplied by the Aztec library.

Code Distribution

Aztec is publicly available for research purposes and may be licensed for commercial application. The code is distributed along with technical documentation, example C and Fortran driver routines and sample input files via the internet. It may be obtained by contacting one of the authors listed on page i of this report.

1. Overview. *Aztec* is an iterative library that greatly simplifies the parallelization process when solving the linear system of equations

$$Ax = b$$

where A is a user supplied $n \times n$ sparse matrix, b is a user supplied vector of length n and x is a vector of length n to be computed. *Aztec* is intended as a software tool for users who want to avoid cumbersome parallel programming details but who have large sparse linear systems requiring efficient use of a parallel processing system. The most complicated parallelization task for an *Aztec* user is the distributed matrix specification for the particular application. Although this may seem difficult, a collection of data transformation tools are provided that allow creation of distributed sparse unstructured matrices for parallel solution with ease of effort that is similar to a serial implementation. Background information regarding the data transformation tools can be found in [5]. Once the distributed matrix is created, computation can occur on any of the parallel machines running *Aztec*: nCUBE 2, IBM SP2, Intel Paragon, and MPI platforms. In addition, *Aztec* can be used on standard serial and vector platforms such as SUN, SGI and CRAY computers.

Aztec includes a number of Krylov iterative methods such as conjugate gradient (CG), generalized minimum residual (GMRES) and stabilized biconjugate gradient (BiCGSTAB) to solve systems of equations. These Krylov methods are used in conjunction with various preconditioners such as polynomial preconditioners or domain decomposition using LU or incomplete LU factorizations within subdomains. Background information concerning the iterative methods and the preconditioners can be found in [4]. Although the matrix A can be general, the package has been designed for matrices arising from the approximation of partial differential equations (PDEs). In particular, the preconditioners, iterative methods and parallelization techniques are oriented toward systems arising from PDE applications. Lastly, *Aztec* can use one of two different sparse matrix notations – either a point-entry modified sparse row (MSR) format or a block-entry variable block row (VBR) format. These two formats have been generalized for parallel implementation and, as such, are referred to as “distributed” yielding DMSR and DVBR references.

The remainder of this guide describes how *Aztec* is invoked within an application. *Aztec* is written in ANSI-standard c and as such, all arrays in the descriptions which follow begin indexing with 0. Also, all function prototypes (loosely, descriptions) are presented in ANSI c format. Section 2 discusses iterative method, preconditioning and convergence options. Section 3 explains vectors and sparse matrix formats supported by *Aztec*. In Section 4 we discuss the data transformation tool for creating distributed vectors and matrices. A concrete detailed programming example using this tool is given in Section 5 and some advance topics are discussed in Section 6. Finally, Section 7 gives a glossary of *Aztec* functions available to users.

2. *Aztec*: High Level View. The following tasks must be performed to successfully invoke *Aztec*:

- describe the parallel machine (e.g. number of processors).
- initialize matrix and vector data structures.
- choose iterative methods, preconditioners and the convergence criteria.
- initialize the right hand side and initial guess.
- invoke the solver.

Example

```
#include "az_aztec.h"

void main(void) {

    AZ_processor_info(proc_config);

    init_matrix_vector_structures(bindx, val, update, external,
                                update_index, extern_index, data_org);
    init_options(options, params);

    init_guess_and_rhs(x, b, data_org, update, update_index);

    AZ_solve(x, b, options, params, bindx, val, data_org, status,
            proc_config);
}
```

FIG. 1. *High level code for Aztec application.*

A sample C program is shown in Figure 1 omitting declarations and some parameters¹. The functions `init_matrix_vector_structures`, `init_options`, and `init_guess_and_rhs` are supplied by the user. In this section, we give an overview of **Aztec**'s features by describing the user input arrays, *options* and *params*, that are set by the user in the function `init_options`. A discussion of the other subroutines is deferred to Sections 4 and 5.

2.1. Aztec Options. *options* is an integer array of length `AZ_OPTIONS_SIZE` set by the user. It is used (but not altered) by the function `AZ_solve` to choose between iterative solvers, preconditioners, etc. Below we discuss each of the possible options. In some of these descriptions, reference is made to a user-defined *options* or *params* value which is yet to be introduced. These descriptions will follow but the reader may wish to "jump ahead" and read the descriptions if the immediate context is not clear.

Specifications

<i>options</i> [<code>AZ_solver</code>]	Specifies solution algorithm. DEFAULT: <code>AZ_gmres</code> .
<code>AZ_cg</code>	Conjugate gradient (only applicable to symmetric positive definite matrices).
<code>AZ_gmres</code>	Restarted generalized minimal residual.
<code>AZ_cgs</code>	Conjugate gradient squared.
<code>AZ_tfqmr</code>	Transpose-free quasi-minimal residual.

¹ The entire main program with specific sample problems is distributed with the package in the file `az_main.c`

AZ_bicgstab	Bi-conjugate gradient with stabilization.
AZ_lu	Sparse direct solver (single processor only).
<i>options</i> [AZ_scaling]	Specifies scaling algorithm. The entire matrix is scaled (overwriting the old matrix). Additionally, the right hand side, the initial guess and the final computed solution are scaled if necessary. DEFAULT: AZ_none.
AZ_none	No scaling.
AZ_Jacobi	Point Jacobi scaling.
AZ_BJacobi	Block Jacobi scaling where the block size corresponds to the VBR blocks. Point Jacobi scaling is performed when using the MSR format.
AZ_row_sum	Scale each row so the magnitude of its elements sum to 1.
AZ_sym_diag	Symmetric scaling so diagonal elements are 1.
AZ_sym_row_sum	Symmetric scaling using the matrix row sums.
<i>options</i> [AZ_precond]	Specifies preconditioner. DEFAULT: AZ_none.
AZ_none	No preconditioning.
AZ_Jacobi	k step Jacobi (block Jacobi for DVBR matrices where each block corresponds to a VBR block). The number of Jacobi steps, k , is set via <i>options</i> [AZ_poly_ord].
AZ_Neumann	Neumann series polynomial where the polynomial order is set via <i>options</i> [AZ_poly_ord].
AZ_ls	Least-squares polynomial where the polynomial order is set via <i>options</i> [AZ_poly_ord].
AZ_lu	Domain decomposition preconditioner (additive Schwarz) using a sparse LU factorization in conjunction with a drop tolerance <i>params</i> [AZ_drop] on each processor's submatrix. The treatment of external variables in the submatrix is determined by <i>options</i> [AZ_overlap]. The current sparse lu factorization is provided by the package y12m [6].
AZ_ilu	Similar to AZ_lu using ilu(0) instead of LU.
AZ_bilu	Similar to AZ_lu using block ilu(0) instead of LU where each block corresponds to a VBR block.

AZ_sym_GS	Non-overlapping domain decomposition (additive Schwarz) k step symmetric Gauss-Siedel. In particular, a symmetric Gauss-Siedel domain decomposition procedure is used where each processor independently performs one step of symmetric Gauss-Siedel on its local matrix, followed by communication to update boundary values before the next local symmetric Gauss-Siedel step. The number of steps, k , is set via <i>options</i> [AZ_poly_ord].
<i>options</i> [AZ_conv]	Determines the residual expression used in convergence checks and printing. DEFAULT: AZ_r0. The iterative solver terminates if the corresponding residual expression is less than <i>params</i> [AZ_tol]:
AZ_r0	$\ r\ _2 / \ r^{(0)}\ _2$
AZ_rhs	$\ r\ _2 / \ b\ _2$
AZ_Anorm	$\ r\ _2 / \ A\ _\infty$
AZ_sol	$\ r\ _\infty / (\ A\ _\infty * \ x\ _1 + \ b\ _\infty)$
AZ_weighted	$\ r\ _{WRMS}$ where $\ \cdot\ _{WRMS} = \sqrt{(1/n) \sum_{i=1}^n (r_i/w_i)^2}$, n is the total number of unknowns, w is a weight vector provided by the user via <i>params</i> [AZ_weights] and $r^{(0)}$ is the initial residual.
<i>options</i> [AZ_output]	Specifies information (residual expressions - see <i>options</i> [AZ_conv]) to be printed. DEFAULT: 1.
AZ_all	Print out the matrix and indexing vectors for each processor. Print out all intermediate residual expressions.
AZ_none	No intermediate results are printed.
AZ_last	Print out only the final residual expression.
> 0	Print residual expression every <i>options</i> [AZ_output] iterations.
<i>options</i> [AZ_pre_calc]	Indicates whether to use factorization information from previous calls to AZ_solve. DEFAULT: AZ_calc.
AZ_calc	Use no information from previous AZ_solve calls.
AZ_recalc	Use preprocessing information from a previous call but recalculate preconditioning factors. This is primarily intended for factorization software which performs a symbolic stage.

AZ_reuse	Use preconditioner from a previous AZ_solve call, do not recalculate preconditioning factors. Also, use scaling factors from previous call to scale the right hand side, initial guess and the final solution.
<i>options</i> [AZ_max_iter]	Maximum number of iterations. DEFAULT: 500.
<i>options</i> [AZ_poly_ord]	The polynomial order when using polynomial preconditioning. Also, the number of steps when using Jacobi or symmetric Gauss-Seidel preconditioning. DEFAULT: 3.
<i>options</i> [AZ_overlap]	Determines the submatrices factored with the domain decomposition algorithms: AZ_lu, AZ_ilu, AZ_bilu. DEFAULT: AZ_none.
AZ_none	Factor the local submatrix defined on this processor discarding column entries that correspond to external elements.
AZ_diag	Factor the local submatrix defined on this processor augmented by a diagonal (block diagonal for VBR format) matrix. This diagonal matrix corresponds to the diagonal entries of the matrix rows (found on other processors) associated with external elements. This can be viewed as taking one Jacobi step to update the external elements and then performing domain decomposition with AZ_none on the residual equations.
AZ_full	Factor the local submatrix defined on this processor augmented by the rows (found on other processors) associated with external variables (discarding column entries associated with variables not defined on this processor). The resulting procedure is an overlapped additive Schwarz procedure.
<i>options</i> [AZ_kspace]	Krylov subspace size for restarted GMRES. DEFAULT: 30.
<i>options</i> [AZ_orthog]	GMRES orthogonalization scheme. DEFAULT: AZ_classic.
AZ_classic	Classical Gramm-Schmidt orthogonalization.
AZ_modified	Modified Gramm-Schmidt orthogonalization.
<i>options</i> [AZ_aux_vec]	Determines $\tilde{r}$ (a required vector within some iterative methods). The convergence behavior varies slightly depending on how this is set. DEFAULT: AZ_resid.
AZ_resid	$\tilde{r}$ is set to the initial residual vector.

AZ_rand $\tilde{r}$ is set to random numbers between -1 and 1.
 NOTE: When using this option, the convergence depends on the number of processors (i.e. the iterates obtained with x processors differ from the iterates obtained with y processors if $x \neq y$).

2.2. Aztec parameters. *params* is a double precision array set by the user and normally of length AZ_PARAMS_SIZE. However, when a weight vector is needed for the convergence check (i.e. *options*[AZ_conv] = AZ_weighted), it is embedded in *params* whose length must now be AZ_PARAMS_SIZE + # of elements updated on this processor. In either case, the contents of *params* are used (but not altered) by the function AZ_solve to control the behavior of the iterative methods. The array elements are specified as follows:

Specifications

<i>params</i> [AZ_tol]	Specifies tolerance value used in conjunction with convergence tests. DEFAULT: 10^{-6} .
<i>params</i> [AZ_drop]	Specifies drop tolerance used in conjunction with LU preconditioner. DEFAULT: 0.0.
<i>params</i> [AZ_weights]	When <i>options</i> [AZ_conv] = AZ_weighted, the i 'th local component of the weight vector is stored in the location <i>params</i> [AZ_weights+i].

Figure 2 illustrates a sample function *init_options* where the Aztec function AZ_defaults sets the default options.

Example

```
void init_options(int options[AZ_OPTIONS_SIZE],
                 double params[AZ_PARAMS_SIZE])
{
    AZ_defaults(options, params);
    options[AZ_solver]      = AZ_cgs;
    options[AZ_scaling]     = AZ_none;
    options[AZ_precond]     = AZ_ls;
    options[AZ_output]      = 1;
    options[AZ_max_iter]    = 640;
    options[AZ_poly_ord]    = 7;
    params[AZ_tol]          = 0.0000001;
}
```

FIG. 2. Example option initialization routine (*init_options*).

2.3. Return status. *status* is a double precision array of length `AZ_STATUS_SIZE` returned from `AZ_solve`². The contents of *status* are described below.

Specifications

<i>status</i> [<code>AZ_its</code>]	Number of iterations taken by the iterative method.
<i>status</i> [<code>AZ_why</code>]	Reason why <code>AZ_solve</code> terminated.
<code>AZ_normal</code>	User requested convergence criteria is satisfied.
<code>AZ_param</code>	User requested option is not available.
<code>AZ_breakdown</code>	Numerical breakdown occurred.
<code>AZ_loss</code>	Numerical loss of precision occurred.
<code>AZ_maxits</code>	Maximum iterations taken without convergence.
<i>status</i> [<code>AZ_r</code>]	The true residual norm corresponding to the choice <i>options</i> [<code>AZ_conv</code>] (this norm is calculated using the computed solution).
<i>status</i> [<code>AZ_scaled_r</code>]	The true residual ratio expression as defined by <i>options</i> [<code>AZ_conv</code>].
<i>status</i> [<code>AZ_rec_r</code>]	Norm corresponding to <i>options</i> [<code>AZ_conv</code>] of final residual or estimated final residual (recursively computed by iterative method). Note: When using the 2-norm, <code>tfqmr</code> computes an estimate of the residual norm instead of computing the residual.

When `AZ_solve` returns abnormally, the user may elect to restart using the current computed solution as an initial guess.

3. Data Formats. In this section we describe the matrix and vector formats used internally by `Aztec`. In Section 4 we discuss a tool that transforms data from a simpler format to this format. Here, the terms “element” and “component” are used interchangeably to denote a particular entry of a vector.

The sparse matrix-vector product, $y \leftarrow Ax$, is the major kernel operation of `Aztec`. To perform this operation in parallel, the vectors x and y as well as the matrix A must be distributed across the processors. The elements of any vector of length n are assigned to a particular processor via some partitioning method (e.g. `Chaco` [2]). When calculating elements in a vector such as y , a processor computes only those elements in y which it has been assigned. These vector elements are explicitly stored on the processor and are defined by a set of indices referred to as the processor’s *update* set. The *update* set is further divided into two subsets: *internal* and *border*. A component corresponding to an index in the *internal* set is updated using only information on the

² All integer information returned from `AZ_solve` is cast into double precision and stored in *status*.

current processor. As an example, the index i is in *internal* if, in the matrix-vector product kernel, the element y_i is updated by this processor and if each j defining a nonzero A_{ij} in row i is in *update*. The *border* set defines elements which would require values from other processors in order to be updated during the matrix vector product. For example, the index i is in *border* if, in the matrix-vector product kernel, the element y_i is updated by this processor and if there exists at least one j associated with a nonzero A_{ij} found in row i that is not in *update*. In the matrix-vector product, the set of indices which identify the off-processor elements in x that are needed to update components corresponding to *border* indices is referred to as *external*. They are explicitly stored by and are obtained from other processors via communication whenever a matrix-vector product is performed. Figure 3 illustrates how a set of vertices in a partitioning of a grid would be used to define these sets. Since these sets of indices are used exclusively

FIG. 3. Example partitioning of a finite element grid.

to reference specific vector components, the same names (i.e., *update*, *internal*, *border* and *external*) are sometimes used below to describe the vector elements themselves. Having generalized these labels, the three types of vector elements are distinguished by locally storing the *internal* components first, followed by the *border* components and finally by the *external* components. In addition, all *external* components received from the same processor are stored consecutively. Below we summarize the nomenclature for a processor with N total elements where N_{internal} , N_{border} , and N_{external} elements are distributed over the sets *internal*, *border* and *external* respectively.

set	description	local numbering
<i>internal</i>	updated w/o communication	0 to $N_internal - 1$.
<i>border</i>	updated with communication	$N_internal$ to $N_internal + N_border - 1$.
<i>external</i>	not updated but used to update <i>border</i>	$N_internal + N_border$ to $N - 1$. elements received from the same processor are numbered consecutively.

Similar to vectors, a subset of matrix non-zeros is stored on each processor. In particular, each processor stores only those rows which correspond to its *update* set. For example, if vector element i is updated on processor p , then processor p also stores all the non-zeros of row i in the matrix. Further, the local numbering of vector elements on a specific processor induces a local numbering of matrix rows and columns. For example, if vector element k is locally numbered as k_l , then all references to row k or column k in the matrix would be locally numbered as k_l . Thus, each processor contains a submatrix whose row and column entries correspond to variables defined on this processor.

The remainder of this section describes the two sparse matrix formats that are used to store the local renumbered submatrix. These two sparse matrix formats correspond to common formats used in serial computations.

3.1. Distributed Modified Sparse Row (DMSR) Format. The DMSR format is a generalization of the MSR format [3]. The data structure consists of an integer vector *bindx* and a double precision vector *val* each of length $N_nonzeros + 1$ where $N_nonzeros$ is the number of nonzeros in the local submatrix. For a submatrix with m rows the DMSR arrays are as follows:

bindx :

$$\begin{aligned}
bindx[0] &= m + 1 \\
bindx[k+1] - bindx[k] &= \text{number of nonzero off-diagonal elements in } k\text{'th row, } k < m \\
bindx[k_s \dots k_e] &= \text{column indices of the off-diagonal nonzeros in row } k_l \text{ where } k_s = bindx[k] \text{ and } k_e = bindx[k+1]-1.
\end{aligned}$$

val :

$$\begin{aligned}
val[k] &= A_{kk}, k < m \\
val[k_i] &= \text{the } (k, bindx[k_i])\text{'th matrix element where } k_s \leq k_i \leq k_e \text{ with } k_s \text{ and } k_e \text{ as defined above.}
\end{aligned}$$

Note: $val[m]$ is not used. See [1] for a detailed discussion of the MSR format.

3.2. Distributed Variable Block Row (DVBR) Format. The Distributed Variable Block Row (DVBR) format is a generalization of the VBR format [1]. The data structure consists of a double precision vector *val* and five integer vectors: *indx*, *bindx*, *rpnt*, *cpnt* and *bpnt*. The format is best suited for sparse block matrices of

the form

$$A = \begin{pmatrix} A_{00} & A_{01} & \cdots & A_{0k} \\ A_{10} & A_{11} & \cdots & A_{1k} \\ \vdots & & \ddots & \vdots \\ A_{m0} & \cdots & \cdots & A_{mk} \end{pmatrix}$$

where A_{ij} denotes a block (or submatrix). In a sparse block matrix, some of these blocks would be entirely zero while others may be dense. The DVBR vectors are described below for a matrix with $M \times K$ blocks.

rpntnr[0 ... *M*] :

$$\textit{rpntnr}[0] = 0$$

$$\textit{rpntnr}[k+1] - \textit{rpntnr}[k] = \text{number of rows in } k\text{'th block row}$$

cpntnr[0 ... *K*] :

$$\textit{cpntnr}[0] = 0$$

$$\textit{cpntnr}[k+1] - \textit{cpntnr}[k] = \text{number of columns in } k\text{'th block column}$$

bpntnr[0 ... *M*] :

$$\textit{bpntnr}[0] = 0$$

$$\textit{bpntnr}[k+1] - \textit{bpntnr}[k] = \text{number of nonzero blocks in the } k\text{'th block row}$$

bindx[0 ... *bpntnr*[*M*]] :

$$\textit{bindx}[k_s \dots k_e] = \text{block column indices of nonzero blocks in block row } k \\ \text{where } k_s = \textit{bpntnr}[k] \text{ and } k_e = \textit{bpntnr}[k+1]-1$$

indx[0 ... *bpntnr*[*M*]] :

$$\textit{indx}[0] = 0$$

$$\textit{indx}[k_i+1] - \textit{indx}[k_i] = \text{number of nonzeros in the } (k, \textit{bindx}[k_i])\text{'th block} \\ \text{where } k_s \leq k_i \leq k_e \text{ with } k_s \text{ and } k_e \text{ as defined above.}$$

val[0 ... *indx*[*bpntnr*[*M*]]] :

$$\textit{val}[i_s \dots i_e] = \text{nonzeros in the } (k, \textit{bindx}[k_i])\text{'th block stored in} \\ \text{column major order where } k_i \text{ is as defined above,} \\ i_s = \textit{indx}[k_i] \text{ and } i_e = \textit{indx}[k_i+1]-1$$

See [1] for a detailed discussion of the VBR format.

4. High Level Data Interface. Setting up the distributed format described in Section 3 for the local submatrix on each processor can be quite cumbersome. In particular, the user must determine a mapping between the global numbering scheme and a local scheme which facilitates proper communication. Further, a number of additional variables must be set for communication and synchronization (see Section 6). In this section we describe a simpler data format that is used in conjunction with a transformation function to generate data structures suitable for *Aztec*. The new format allows the user to specify the rows in a natural order as well as to use global column numbers in the *bindx* array. To use the transformation function the user supplies the

update set and the submatrix for each processor. Unlike the previous section, however, the submatrix is specified using the global coordinate numbering instead of the local numbering required by **Aztec**. This procedure greatly facilitates matrix specification and is the main advantage of the transformation software.

On a given processor, the *update* set (i.e. vector element assignment to processors) is defined by initializing the array *update* on each processor so that it contains the global index of each element assigned to the processor. The *update* array must be sorted in ascending order (i.e. $i < j \Rightarrow \text{update}[i] < \text{update}[j]$). This sorting can be performed using the **Aztec** function **AZ_sort**. Matrix specification occurs using the arrays defined in the previous section. However, now the local rows are defined in the same order as the *update* array and column indices (e.g. *bindx*) are given as global column indices. To illustrate this in more detail, consider the following example matrix:

$$A = \begin{pmatrix} a_{00} & a_{01} & & a_{03} & a_{04} & & \\ a_{10} & a_{11} & & a_{13} & & & \\ & & a_{22} & a_{23} & a_{24} & a_{25} & \\ a_{30} & a_{31} & a_{32} & a_{33} & a_{34} & a_{35} & \\ a_{40} & & a_{42} & a_{43} & a_{44} & & \\ & & a_{52} & a_{53} & & a_{55} & \end{pmatrix}.$$

Figure 4 illustrates the information corresponding to a particular matrix partitioning that is specified by the user as input to the data transformation tool. Using this

Examples

proc 0:

N_update:	3													
update:	0	1	3											
bindx:	4	7	9	14	1	3	4	0	3	1	0	4	2	5
val:	a_{00}	a_{11}	a_{33}	-	a_{01}	a_{03}	a_{04}	a_{10}	a_{13}	a_{30}	a_{31}	a_{32}	a_{34}	a_{35}

proc 1:

N_update:	1				
update:	4				
bindx:	2	5	0	3	2
val:	a_{44}	-	a_{40}	a_{42}	a_{43}

proc 2:

N_update:	2							
update:	2	5						
bindx:	3	6	8	4	3	5	3	2
val:	a_{22}	a_{55}	-	a_{23}	a_{24}	a_{25}	a_{52}	a_{53}

FIG. 4. User input (MSR format) to initialize the sample matrix problem.

information, **AZ_transform**

- determines the sets *internal*, *border* and *external*.
- determines the local numbering: *update_index[i]* is the local numbering for *update[i]* while *extern_index[i]* is the local numbering for *external[i]*.

- permutes and renumbers the local submatrix rows and columns so that they now correspond to the new ordering.
- computes additional information (e.g. the number of internal, border and external components on this processor) and stores this in *data_org* (see Section 6).

A sample transformation is given in Figure 5 and is found in the file *az_app_utils.c*.

Example

```
init_matrix_vector_structures(bindx, val, update, external,
                             update_index, extern_index, data_org);
{
    AZ_read_update(update, N_update);
    create_matrix(bindx, val, update, N_update);
    AZ_transform(bindx, val, update, external, update_index,
                extern_index, data_org, N_update);
}
```

FIG. 5. *init_matrix_vector_structures*.

AZ_read_update is an Aztec utility which reads a file and assigns elements to *update*. The user supplied routine *create_matrix* creates an MSR or VBR matrix using the global numbering. Once transformed the matrix can now be used within Aztec.

5. Examples. A sample program is described by completing the program fragments given earlier (Figures 1, 2 and 5). In Figure 1, *AZ_processor_info* is an Aztec utility which initializes the array *proc_config* to reflect the number of processors being used and the node number of this processor. The function *AZ_solve* is also supplied by Aztec to solve the user supplied linear system. Thus, the only functions that the user must supply which have not already been discussed include: *init_guess_and_rhs* in Figure 1 and *create_matrix* in Figure 5.

The function *init_guess_and_rhs* initializes the initial guess and the right hand side.

Example

```
void init_guess_and_rhs(x, rhs, data_org, update, update_index)
{
    N_update = data_org[AZ_N_internal] + data_org[AZ_N_border];
    for (i = 0; i < N_update ; i = i + 1) {
        rhs[update_index[i]] = (double) update[i];
        x[i] = 0.0;
    }
}
```

FIG. 6. *init_guess_and_rhs*.

In Figure 6, a sample routine is given which sets the initial guess vector to zero and sets the right hand side vector equal to the global indices (where the local element *update_index[i]* corresponds to global element *update[i]*, see Section 4).

A `create_matrix` function to initialize an MSR matrix is illustrated in Figure 7. Different matrix problems can be implemented by changing the function `add_row` which computes the MSR entries corresponding to a new row of the matrix. The specific

Example

```
void create_matrix(bindx, val, update, N_update)

{
    N_nonzeros = N_update + 1;
    bindx[0] = N_nonzeros;

    for (i = 0; i < N_update; i = i + 1)
        add_row(update[i], i, val, bindx);
}
```

FIG. 7. `create_matrix`.

`add_row` function for implementing a 5-point 2D Poisson operator on an $n \times n$ grid is shown in Figure 8 (n is a global variable set by the user). With these few lines of code

Example

```
void add_row(row, location, val, bindx)
{
    k = bindx[location];

    /* check neighboring points in each direction and add nonzero */
    /* entry if neighbor exists.                                     */

    bindx[k] = row + 1;    if (row%n != n-1) val[k++] = -1.;
    bindx[k] = row - 1;    if (row%n != 0) val[k++] = -1.;
    bindx[k] = row + n;    if ((row/n)%n != n-1) val[k++] = -1.;
    bindx[k] = row - n;    if ((row/n)%n != 0) val[k++] = -1.;

    bindx[location+1] = k;
    val[location] = 4.;    /* matrix diagonal */
}
```

FIG. 8. `add_row` for a 2D Poisson problem

and the functions described earlier, the user initializes and solves a 2D Poisson problem. While for simplicity of presentation this specific example is structured the **Aztec** library does not assume any structure in the sparse matrix. All the communication and variable renumbering is done automatically without the assumption of structured communication.

Other `add_row` functions corresponding to a 3D Poisson equation and a high order 2D Poisson equation are distributed with *Aztec* (file `az_examples.c`). We recommend that potential users review at these examples. In many cases, new applications can be written by simply editing these programs. The interested reader should note that only a few lines of code are different between the functions for the 5-pt Poisson, the high order Poisson and the 3D Poisson codes. Further, the `add_row` routines are essentially identical to those that would be used to set up sparse matrices in serial applications and that there are no references to processors, communications or anything specific to parallel programming.

While *Aztec* simplifies the parallel coding associated with structured problems, it is for unstructured problems that *Aztec* makes a significant programming difference. To illustrate this, a 2D finite element example is given where the underlying grid is a triangulation of a complex geometry. Unlike the previous example `create_matrix` defines a sparsity pattern (i.e. *bindx*) but not the actual nonzero entries (i.e. *val*) as interprocessor communication is required before they can be computed. Thus, in this example **AZ_transform** takes the sparsity pattern and initializes the communication data structures. Using these structures, communication can be performed at a later stage in computing the matrix nonzeros.

Figure 9 depicts `create_matrix` while Figure 10 depicts an additional function ma-

Example

```
void create_matrix(bindx, val, update, N_update);
{
    read_triangles(T, N_triangles);
    init_msr(val, bindx, N_update);

    for (triangle = 0; triangle < N_triangles; triangle = triangle + 1)
        for (i = 0; i < 3; i = i + 1) {
            row = AZ_find_index(T[triangle][i], update, N_update);
            for (j = 0; j < 3; j = j + 1) {
                if (row != NOT_FOUND)
                    add_to_element(row, T[triangle][j], 0.0, val, bindx, i==j);
            }
        }
    compress_matrix(val, bindx, N_update);
}
```

FIG. 9. `create_matrix` for the Poisson finite element problem.

`trix_fill` that must be included before `AZ_solve` is invoked in Figure 1. We have not made any effort to optimize these routines. In both figures the new lines that have been added specifically for a parallel implementation are underlined. That is, `create_matrix` and `matrix_fill` have been created by taking a serial program that creates the finite element discretization, splitting this program over the two functions and adding a few new lines necessary for the parallel implementation. The only additional change is to replace the single data file containing the triangle connectivity read using `read_triangles`

by a set of data files containing the triangle connectivity for each processor. We do not discuss the details of this program but only wish to draw the readers attention to the small number of lines that need changing to convert the serial unstructured application to parallel. Most of the main routines such as `setup_Ke` which computes the element contributions and `add_to_element` which stores the element contributions in the MSR data structures remain the same. In fact, almost all the new lines of code correspond to adding the communication (`AZ_exchange_bdry`) (which was the main reason that the calculation of the matrix nonzeros was deferred) and the conversion of global index values by local index values with the help of `AZ_find_index`. As in the Poisson example, all of the details with respect to communication are hidden from the user.

6. Advanced Topics.

6.1. Data Layout. The `Aztec` function `AZ_transform` initializes the integer array `data_org`. This array specifies how the matrix is set up on the parallel machine. In many cases, the user need not be concerned with the contents of this array. However, in some situations it is useful to initialize these elements without the use of `AZ_transform`, to access these array elements (e.g. determine how many *internal* components are used), or to change these array elements (e.g. when reusing factorization information, see Section 6.2). When using the transformation software, the user can ignore the size of `data_org` as it is allocated in `AZ_transform`. However, when this is not used, `data_org` must be allocated of size `AZ_COMM_SIZE` + number of vector elements sent to other processors during matrix-vector multiplies. The contents of `data_org` are as follows:

Specifications

<code>data_org[AZ_matrix_type]</code>	Specifies matrix format.
<code>AZ_VBR_MATRIX</code>	Matrix corresponds to VBR format.
<code>AZ_MSR_MATRIX</code>	Matrix corresponds to MSR format.
<code>data_org[AZ_N_internal]</code>	Number of elements updated by this processor that can be computed without information from neighboring processors (<i>N_internal</i>). This also corresponds to the number of internal rows assigned to this processor.
<code>data_org[AZ_N_border]</code>	Number of elements updated by this processor that use information from neighboring processors (<i>N_border</i>).
<code>data_org[AZ_N_external]</code>	Number of <i>external</i> components needed by this processor (<i>N_external</i>).
<code>data_org[AZ_N_int_blk]</code>	Number of internal VBR block rows owned by this processor. Set to <code>data_org[AZ_N_internal]</code> for MSR matrices.
<code>data_org[AZ_N_bord_blk]</code>	Number of border VBR block rows owned by this processor. Set to <code>data_org[AZ_N_border]</code> for MSR matrices.

Example

```
void matrix_fill(bindx, val, N_update, update, update_index,
                N_external, external, extern_index)

/* read the x and y coordinates from an input file */

for (i = 0; i < N_update; i = i + 1){
    read_from_file(x[update_index[i]], y[update_index[i]]);
}
AZ_exchange_bdry(x);
AZ_exchange_bdry(y);

/* Locally renumber the rows and columns of the new sparse matrix */

for (triangle = 0; triangle < N_triangles; triangle = triangle + 1)
    for (i = 0; i < 3; i = i + 1) {
        row = AZ_find_index(T[triangle][i], update, N_update);
        if (row == NOT_FOUND) {
            row = AZ_find_index(T[triangle][i], external, N_external);
            T[triangle][i] = extern_index[row];
        }
        else T[triangle][i] = update_index[row];
    }
}

/* Fill the element stiffness matrix Ke */

for (triangle = 0; triangle < N_triangles; triangle = triangle + 1){
    setup_Ke(Ke, x[T[triangle][0]], y[T[triangle][0]],
            x[T[triangle][1]], y[T[triangle][1]],
            x[T[triangle][2]], y[T[triangle][2]]);

/* Fill the sparse matrix by scattering Ke to appropriate locations */

    for (i = 0; i < 3; i = i + 1) {
        for (j = 0; j < 3; j = j + 1){
            if (T[triangle][i] < N_update){
                add_to_element(T[triangle][i], T[triangle][j], Ke[i][j],
                    val, bindx, i==j);
            }
        }
    }
}
```

FIG. 10. matrix_fill for the Poisson finite element problem.

<i>data_org</i> [AZ_N_ext_blk/	Number of external VBR block rows on this processor. Set to <i>data_org</i> [AZ_N_external] for MSR matrices.
<i>data_org</i> [AZ_N_neigh/	Number of processors with which we exchange information (send or receive) in performing matrix-vector products.
<i>data_org</i> [AZ_total_send/	Total number of vector elements sent to other processors during matrix-vector products.
<i>data_org</i> [AZ_name/	Name of the matrix. This name is utilized when deciding which previous factorization to use as a preconditioner (see Section 6.2). (positive integer value).
<i>data_org</i> [AZ_neighbors/	Start of vector containing node i.d.'s of neighboring processors. That is, <i>data_org</i> [AZ_neighbors+i] gives the node i.d. of the (i+1)'th neighbor.
<i>data_org</i> [AZ_rec.length/	Start of vector containing the number of elements to receive from each neighbor. We receive from the (i+1)'th neighbor <i>data_org</i> [AZ_rec.length+i] elements.
<i>data_org</i> [AZ_send.length/	Start of vector containing the number of elements to send to each neighbor. We send to the (i+1)'th neighbor <i>data_org</i> [AZ_rec.length+i] elements.
<i>data_org</i> [AZ_send.list/	Start of vector indicating the elements that we will send to other processors during communication. The first <i>data_org</i> [AZ_send.length] components correspond to the elements for the first neighbor and the next <i>data_org</i> [AZ_send.length+1] components correspond to element indices for the second neighbor, and so on.

6.2. Reusing factorizations. When solving a problem, Aztec may create certain information that can be reused later. In most cases, this information corresponds to either matrix scaling factors or preconditioning factorization information for LU or ILU. This information is saved internally and referenced by the matrix name given by *data_org*[AZ_name]. By changing *options*[AZ_pre_calc] and *data_org*[AZ_name] a number of different Aztec possibilities can be realized. As an example, consider the following situation. A user needs to solve the linear systems in the order shown below:

$$A_1x = b, A_2y = x, \text{ and } A_1z = y.$$

The first and second systems are solved with *options*[AZ_pre_calc] set to AZ_calc. However, the name (i.e. *data_org*[AZ_name]) is changed between these two solves. In this way, scaling and preconditioning information computed from the first solve is not overwritten during the second solve. By then setting *options*[AZ_pre_calc] to AZ_reuse and *data_org*[AZ_name] to the name used during the first solve, the third system is solved reusing the scaling information (to scale the right hand side, initial guess, and rescale

the final solution³) and the preconditioning factorizations (e.g. ILU) used during the first solve. While in this example the same matrix system is solved for the first and third solve, this is not necessary. In particular, preconditioners can be reused from previous nonlinear iterates even though the linear system being solved are changing. Of course, many times information from previous linear solves is not reused. In this case the user must explicitly free the space associated with the matrix or this information will remain allocated for the duration of the program. Space is cleared by invoking `AZ_clear(data_org/AZ_name/)`.

6.3. Important Constants. *Aztec* uses a number of constants which are defined in the file `az_aztec_defs.h`. Most users can ignore these constants. However, there may be situations where they should be changed. Below is a list of these constants with a brief description:

<code>AZ_MAX_NEIGHBORS</code>	Maximum number of processors with which information can be exchanged during matrix-vector products.
<code>AZ_MSG_TYPE</code> <code>AZ_NUM_MSGS</code>	All message types used inside <i>Aztec</i> lie between <code>AZ_MSG_TYPE</code> and <code>AZ_MSG_TYPE + AZ_NUM_MSGS - 1</code> .
<code>AZ_MAX_BUFFER_SIZE</code>	Maximum message information that can be sent by any processor at any given time before receiving. This is used to subdivide large messages to avoid buffer overflows.
<code>AZ_MAX_MEMORY_SIZE</code>	Maximum available memory. Used primarily for the LU-factorizations where a large amount of memory is first allocated and then unused portions are freed after factorization.
<code>AZ_TEST_ELE</code>	Internal algorithm parameter that can effect the speed of the <code>AZ_find_procs_for_externs</code> calculation. Reduce <code>AZ_TEST_ELE</code> if communication buffers are exceeded during this calculation.

6.4. AZ_transform Subtasks. The function `AZ_transform` described in Section 4 is actually made up of 5 subtasks. In most cases the user need not be concerned with the individual tasks. However, there might arise situations where additional information is available such that some of the subtasks can be omitted. In this case, it is possible for the user to edit the code for `AZ_transform` located in the file `az_tools.c` to suit the application. In this section we briefly describe the five subroutines which make up the transformation function. More detailed descriptions are given in [5]. Prototypes for these subroutines as well as for `AZ_transform` are given in Section 7.

`AZ_transform` begins by identifying the *external* set needed by each processor. Here, each column entry must correspond to either an element updated by this processor or

³ The matrix does not need to be rescaled as the scaling during the first solve overwrites the original matrix.

an *external* component. The function `AZ_find_local_indices` checks each column entry. If a column is in *update*, its number is replaced by the appropriate index into *update* (i.e. $update[new\ column\ index] = old\ column\ index$). If a column number is not found in *update*, it is stored in the *external* list and the column number is replaced by an index into *external* (i.e. $external[new\ column\ index - N_update] = old\ column\ index$).

`AZ_find_procs_for_extrns` queries the other processors to determine which processors update each of its *external* components. The array *extern_proc* is set such that *extern_proc[i]* indicates which processor updates *external[i]*.

`AZ_order_ele` reorders the *external* components such that elements updated by the same processor are contiguous. This new ordering is given by *extern_index* where *extern_index[i]* indicates the local numbering of *external[i]*. Additionally, *update* components are reordered so the *internal* components precede the *border* components. This new ordering is given by *update_index* where *update_index[i]* indicates the local numbering of *update[i]*.

`AZ_set_message_info` initializes *data_org* (see Section 6.1) This is done by computing the number of neighbors, making a list of the neighbors, computing the number of values to be sent and received with each neighbor and computing the list of elements which will be sent to other processors during communication steps.

Finally, `AZ_reorder_matrix` permutes and reorders the matrix nonzeros so that its entries correspond to the newly reordered vector elements.

7. Aztec Functions . In this section we describe the **Aztec** functions available to the user. Certain variables appear many times in the parameter lists of these frequently used functions. In the interest of brevity we describe these variables at the beginning of this section and then proceed with the individual function descriptions.

Frequently Used Aztec Parameters

<i>data_org</i>	Array describing the matrix format (Section 6.1). Allocated and set <code>AZ_set_message_info</code> and <code>AZ_transform</code> .
<i>extern_index</i>	<i>extern_index[i]</i> gives the local numbering of global element <i>external[i]</i> . Allocated and set by <code>AZ_order_ele</code> and <code>AZ_transform</code> .
<i>extern_proc</i>	<i>extern_proc[i]</i> is updating processor of <i>external[i]</i> . Allocated and set by <code>AZ_find_procs_for_extrns</code> .
<i>external</i>	Sorted list (global indices) of external elements on this node. Allocated and set by <code>AZ_find_local_indices</code> and <code>AZ_transform</code> .
<i>N_external</i>	Number of <i>external</i> components. Set by <code>AZ_find_procs_for_extrns</code> and <code>AZ_transform</code> .
<i>N_update</i>	Number of <i>update</i> components assigned to this processor. Set by <code>AZ_read_update</code> .
<i>options, params</i>	Arrays describing <code>AZ_solve</code> options (Section 2).

<i>proc_config[AZ_node]</i>	Node i.d. of this processor.
<i>proc_config[AZ_N_proc]</i>	Total number of processors used in current simulation. Allocated and set by <i>AZ_processor_info</i> .
<i>update_index</i>	<i>update_index[i]</i> gives the local numbering of global element <i>update[i]</i> . Allocated and set by <i>AZ_order_ele</i> and <i>AZ_transform</i> .
<i>update</i>	Sorted list of elements (global indices) to be updated on this processor. Allocated and set by <i>AZ_read_update</i> .
<i>val, bindx, bpntr, cpntr, indx, rpntr</i>	Arrays used to store matrix. For MSR matrices <i>bpntr, cpntr, indx, rpntr</i> are ignored (Section 3).

Prototype

```
void AZ_broadcast(char *ptr, int length, int *proc_config, int action)
```

Description

Used to concatenate a buffer of information and to broadcast this information from processor 0 to the other processors. The four possible actions are

- *action* == *AZ_PACK*
 - *proc_config[AZ_node]* == 0: store *ptr* in the internal buffer.
 - *proc_config[AZ_node]* ≠ 0: read from the internal buffer to *ptr*. If the internal buffer is empty, first receive the broadcast information.
- *action* == *AZ_SEND*
 - *proc_config[AZ_node]* == 0: broadcast the internal buffer (filled by *AZ_broadcast*) and then clear it.
 - *proc_config[AZ_node]* ≠ 0: clear internal buffer.

Sample Usage:

The following code fragment broadcasts the information in 'a' and 'b'.

```
if (proc_config[AZ_node] == 0) {
    a = 1;
    b = 2;
}
AZ_broadcast(&a, sizeof(int), proc_config, AZ_PACK);
AZ_broadcast(&b, sizeof(int), proc_config, AZ_PACK);
AZ_broadcast(NULL, 0, proc_config, AZ_SEND);
```

NOTE: There can be no other communication calls between the `AZ_PACK` and `AZ_SEND` calls to `AZ_broadcast`.

Parameters _____

<i>ptr</i>	On input, data string of size <i>length</i> . Information is either stored to or retrieved from <i>ptr</i> as described above.
<i>length</i>	On input, length of <i>ptr</i> to be broadcast/received.
<i>action</i>	On input, determines <code>AZ_broadcast</code> behavior.

Prototype _____

```
int AZ_check_input(int *data_org, int *options, double *params, int *proc_config)
```

Description _____

Perform checks for iterative solver library. This is to be called by the user of the solver library to check the values in *data_org*, *options*, *params*, and *proc_config*. If all the values are valid `AZ_check_input` returns 0, otherwise it returns an error code which can be deciphered using `AZ_print_error`.

Prototype _____

```
void AZ_check_msr(int *bindx, int N_update, int N_external, int option,  
                 int *proc_config)
```

Description _____

Check that the number of nonzero off-diagonals in each row and that the column indices are nonnegative and not too large (see *option*).

Parameters _____

option

AZ_LOCAL	On input, indicates matrix uses local indices. The number of nonzeros in a row and the largest column index must not exceed the total number of elements on this processor.
AZ_GLOBAL	On input, indicates matrix uses global indices. The number of nonzeros in a row and the largest column index must not exceed the total number of elements in the simulation.

Prototype

```
void AZ_check_vbr(int N_update, int N_external, int option, int *bindx,
                 int *bpntr, int *cpntr, int *rpntr, int *proc_config )
```

Description

Check VBR matrix for the following:

- number of columns within each block column is nonnegative.
- $rpntr[i] == cpntr[i]$ for $i \leq N_update$.
- number of nonzero blocks in each block row is nonnegative and not too large.
- block column indices are nonnegative and not too large.

Parameters

option

AZ_LOCAL	On input, indicates matrix uses local indices. The number of block nonzeros in a row and the largest block column index must not exceed the total number of blocks columns on this processor.
AZ_GLOBAL	On input, indicates matrix uses global indices. The number of block nonzeros in a row and the largest block column index must not exceed the total number of blocks rows in the simulation.

Prototype

```
int AZ_defaults(double *options, int *params )
```

Description

Set *options* and *params* so that the default options are chosen.

Parameters _____

<i>options</i>	On output, set to the default options.
<i>params</i>	On output, set to the default parameters.

Prototype _____

```
void AZ_exchange_bdry(double *x, int *data_org)
```

Description _____

Locally exchange the components of the vector x so that the *external* components of x are updated.

Parameters _____

x	On input, vector defined on this processor. On output, <i>external</i> components of x are updated via communication.
-----	---

Prototype _____

```
int AZ_find_index(int key, int *list, int length )
```

Description _____

Returns the index, i , in *list* (assumed to be sorted) which matches the key (i.e. $list[i] == key$). If *key* is not found AZ_find_index returns -1. See also AZ_quick_find.

Parameters _____

<i>key</i>	On input, element to be search for in list.
<i>list</i>	On input, sorted list to be searched.
<i>length</i>	On input, length of list.

Prototype

```
void AZ_find_local_indices(int N_update, int *bindx, int *update,
                          int **external, int *N_external, int mat_type,
                          int *bpntr)
```

Description

Given the global column indices for a matrix and a list of elements updated on this processor, compute the *external* set and change the global column indices to local column indices. Specifically,

- allocate *external*, compute and store the external components in *external*.
- renumber column indices so that column entry *k* is renumbered as *j* where either *update[j] == k* or *external[j-N_update] == k*.

Called by AZ_transform.

Parameters

<i>mat_type</i>	On input, indicates whether matrix format is MSR (= AZ_MSR_MATRIX) or VBR (= AZ_VBR_MATRIX).
<i>external</i>	On output, allocated and set to sorted list of the external elements.
<i>bindx</i>	On input, contains global column numbers of MSR or VBR matrix (Section 3). On output, contains local column numbers as described above.

Prototype

```
void AZ_find_procs_for_extens(int N_update, int *update, int *external,
                              int N_external, int *proc_config, int **extern_proc)
```

Description

Determine which processors are responsible for updating each external element.

Called by AZ_transform.

Parameters

<i>extern_proc</i>	On output, <i>extern_proc[i]</i> contains the node number of the processor which updates <i>external[i]</i> .
--------------------	---

Prototype

```
void AZ_free_memory(int name)
```

Description

Free Aztec memory associated with matrices with *data_org*[AZ_name] = *name*. This is primarily scaling and preconditioning information that has been computed on earlier calls to AZ_solve.

Parameters

<i>name</i>	On output, all preconditioning and scaling information is freed for matrices which have <i>data_org</i> [AZ_name] = <i>name</i> .
-------------	---

Prototype

```
double AZ_gavg_double(double value, int *proc_config )
```

Description

Return the average of the numbers in *value* on all processors.

Parameters

<i>value</i>	On input, <i>value</i> contains a double precision number.
--------------	--

Prototype

```
double AZ_gdot(int N, double *r, double *z, int *proc_config )
```

Description

Return the dot product of *r* and *z* with unit stride. This routine calls the BLAS routine *ddot* to do the local vector dot product and then uses the global summation routine *AZ_gsum_double* to obtain the required global result.

Parameters _____

N	On input, length of r and z on this processor.
r, z	On input, vectors distributed over all the processors.

Prototype _____

```
double AZ_gmax_double(double value, int *proc_config )
```

Description _____

Return the maximum of the numbers in *value* on all processors.

Parameters _____

<i>value</i>	On input, <i>value</i> contains a double precision number.
--------------	--

Prototype _____

```
int AZ_gmax_int(int value, int *proc_config )
```

Description _____

Return the maximum of the numbers in *value* on all processors.

Parameters _____

<i>value</i>	On input, <i>value</i> contains an integer.
--------------	---

Prototype _____

```
double AZ_gmax_matrix_norm(double *val, int *indx, int *bindx, int *rpntr, int *cpntr,  
int *bpntr, int *proc_config, int *data_org)
```


Description _____

Returns the maximum matrix norm $\|A\|_{\infty}$ for the distributed matrix encoded in *val*, *indx*, *bindx*, *rpnt*, *cpnt*, *bpnt* (Section 3).

Prototype _____

`double AZ_gmax_vec(int N, double *vec, int *proc_config )`

Description _____

Return the maximum of all the numbers located in *vec*[*i*] (*i* < *N*) on all processors.

Parameters _____

<i>vec</i>	On input, <i>vec</i> contains a list of numbers.
<i>N</i>	On input, length of <i>vec</i> .

Prototype _____

`double AZ_gmin_double(double value, int *proc_config )`

Description _____

Return the minimum of the numbers in *value* on all processors.

Parameters _____

<i>value</i>	On input, <i>value</i> contains a double precision number.
--------------	--

Prototype _____

`int AZ_gmin_int(int value, int *proc_config )`

Description _____

Return the minimum of the numbers in *value* on all processors.

Parameters _____

value On input, *value* contains an integer.

Prototype _____

```
double AZ_gsum_double(double value, int *proc_config )
```

Description _____

Return the sum of the numbers in *value* on all processors.

Parameters _____

value On input, *value* contains a double precision number.

Prototype _____

```
int AZ_gsum_int(int value, int *proc_config )
```

Description _____

Return the sum of the integers in *value* on all processors.

Parameters _____

value On input, *value* contains an integer.

Prototype _____

```
void AZ_gsum_vec_int(int *values, int *workspace, int length, int *proc_config )
```

Description _____

$values[i]$ is set to the sum of the input numbers in $values[i]$ on all processors ($i < length$).

Parameters _____

<i>values</i>	On input, <i>values</i> contains a list of integers. On output, <i>values[i]</i> contains the sum of the input <i>values[i]</i> on all the processors.
<i>workspace</i>	On input, workspace array of size <i>length</i> .
<i>length</i>	On input, length of <i>values</i> and <i>workspace</i> .

Prototype _____

```
double AZ_gvector_norm(int n, int p, double *x, int *proc_config)
```

Description _____

Returns the p norm of the vector x distributed over the processors:

$$\|x\|_p = (x[0]^p + x[1]^p + \dots + x[N-1]^p)^{1/p}$$

where N is the total number of elements in x over all processors.

NOTE: For the $\|\cdot\|_\infty$ norm, set $p = -1$.

Parameters _____

<i>n</i>	On input, number of <i>update</i> components of x on this processor.
<i>p</i>	On input, order of the norm to perform, i.e., $\ x\ _p$.
<i>x</i>	On input, vector whose norm will be computed.

Prototype _____

```
void AZ_init_quick_find(int *list, int length, int *shift, int *bins )
```

Description

shift and *bins* are set so that they can be used with `AZ_quick_find`. On output, *shift* satisfies

$$\frac{range}{2^{shift-1}} > \left\lfloor \frac{length}{4} \right\rfloor \quad \text{and} \quad \frac{range}{2^{shift}} \leq \left\lfloor \frac{length}{4} \right\rfloor$$

where $range = list[length - 1] - list[0]$. The array *bins* must be of size $2 + length/4$ and is set so that

$$bins[k] \leq list[j] < bins[k + 1]$$

where $k = (list[j] - list[0]) / 2^{shift}$.

This routine is used in conjunction with `AZ_quick_find`. The idea is to use *bins* to get a good initial guess as to the location of *value* in *list*.

Parameters

<i>list</i>	On input, sorted <i>list</i> .
<i>length</i>	On input, length of <i>list</i> .
<i>shift</i>	On output, <i>shift</i> is set as described in above.
<i>bins</i>	On input, array of size $2 + length/4$. On output, <i>bins</i> is set as described above.

Prototype

```
void AZ_matvec_mult(double *val, int *indx, int *bindx, int *rpnt, int *cpnt,
                   int *bpnt, double *b, double *c, int exchange_flag,
                   int *data_org )
```

Description

Perform the matrix-vector multiply

$$c \leftarrow Ab$$

where the matrix *A* is encoded in *val*, *indx*, *bindx*, *rpnt*, *cpnt*, *bpnt* (Section 3).

Parameters

<i>b</i>	On input, distributed vector to use in multiplication.
----------	--

<i>c</i>	On output, result of matrix-vector multiplication.
<i>exchange_flag</i>	On input, dictates whether communication needs to occur. If <i>exchange_flag</i> == 1, communication occurs. If <i>exchange_flag</i> == 0, no communication occurs.

Prototype

```
void AZ_msr2vbr(double *val, int *indx, int *rpnt, int *cpnt, int *bpnt, int *bindx,
               int *bindx2, double *val2, int total_blk_rows, int total_blk_cols,
               int blk_space, int nz_space, int blk_type)
```

Description

Convert the DMSR matrix defined in (*val2*, *bindx2*) to a DVBR matrix defined in (*val*, *indx*, *rpnt*, *cpnt*, *bpnt*, *bindx*).

Parameters

<i>val2</i> , <i>bindx2</i>	On input, DMSR arrays holding the matrix to be converted.
<i>cpnt</i>	On input, <i>cpnt</i> [<i>i</i>] is the block size of the <i>i</i> th block in the resulting DVBR matrix. Columns 0 to <i>cpnt</i> [0] – 1 form the first block column, columns <i>cpnt</i> [0] to <i>cpnt</i> [0] + <i>cpnt</i> [1] – 1 form the second block column, etc. On output, <i>cpnt</i> corresponds to the resulting DVBR matrix.
<i>val</i> , <i>indx</i> , <i>rpnt</i> , <i>bpnt</i> , <i>bindx</i>	On output, DVBR arrays of converted DMSR matrix.
<i>total_blk_rows</i>	On input, number of block rows in resulting local VBR matrix.
<i>total_blk_cols</i>	On input, number of block columns in resulting local VBR matrix.
<i>blk_space</i>	On input, length allocated for <i>bindx</i> and <i>indx</i> .
<i>nz_space</i>	On input, length allocated for <i>val</i> .
<i>blk_type</i>	On input, if <i>blk_type</i> > 0, indicates that all block rows (and columns) have the same size given by <i>blk_type</i> . If <i>blk_type</i> < 0, the block rows have different sizes.

Prototype

```
void AZ_order_ele(int *update_index, int *extern_index, int *N_internal,
                  int *N_border, int N_update, int *bpntr, int *bindx,
                  int *extern_proc, int N_external, int option, int mat_type)
```

Description

Find orderings for *update* and *external*. *external* are ordered so that elements updated by the same processor are contiguous. If *option* == AZ_ALL, *update* are ordered so that the *internal* components have the lowest numbers followed by the *border* components. Otherwise, the order of *update* is unchanged. The ordering information is placed in *update_index* and *extern_index* (Section 4). Called by AZ_transform.

Parameters

<i>N_internal</i>	On output, number of <i>internal</i> components on processor.
<i>N_border</i>	On output, number of <i>border</i> components on processor.
<i>update_index</i>	On output, <i>update_index[i]</i> indicates the local index (or order) of <i>update[i]</i> .
<i>extern_index</i>	On output, <i>extern_index[i]</i> indicates the new local index (or order) of <i>external[i]</i> .
<i>option</i>	On input, indicates whether to reorder <i>update</i> .
AZ_ALL	Order <i>update</i> and <i>external</i> .
AZ_EXTERNS	Order only external elements.
<i>mat_type</i>	On input, indicates whether matrix format is MSR (= AZ_MSR_MATRIX) or VBR (= AZ_VBR_MATRIX).

Prototype

```
void AZ_print_error(int error_code)
```

Description

Prints out an error message corresponding to *error_code*. Typically, *error_code* is generated by *AZ_check_input*.

Parameters _____

<i>error_code</i>	On input, error code generated by <i>AZ_check_input</i> .
-------------------	---

Prototype _____

```
void AZ_processor_info(int *proc_config)
```

Description _____

proc_config[AZ_node] is set to the node name of this processor. *proc_config[AZ_N_proc]* is set to the number of processors used in simulation.

Prototype _____

```
int AZ_quick_find(int key, int *list, int length, int shift, int *bins )
```

Description _____

Return the index, *i*, in *list* (assumed to be sorted) which matches the key (i.e. *list[i] = key*). If *key* is not found *AZ_quick_find* returns -1.

NOTE: This version is faster than *AZ_find* but requires *bins* to be set and stored using *AZ_init_quick_find*.

Parameters _____

<i>key</i>	On input, element to search for in <i>list</i> .
<i>list</i>	On input, sorted list to be searched.
<i>length</i>	On input, length of list.
<i>shift</i>	On input, used for initial guess (computed by previous <i>AZ_init_quick_find</i> call).

bins

On input, computed by `AZ_init_quick_find` for initial guess. *bins* is set so that $list[bins[k]] \leq key < list[bins[k+1]]$ where $k = (key - list[0])/2^{shift}$.

Prototype

```
void AZ_read_msr_matrix(int *update, double **val, int **bindx, int N_update,
                        int *proc_config )
```

Description

Read the file `.data` and create a matrix in the MSR format. Processor 0 reads the input file. If the new row to be added resides in processor 0's *update*, it is added to processor 0's matrix. Otherwise, processor 0 determines which processor has requested this row and sends it to this processor for its local matrix.

The form of the input file is as follows:

```
num_rows
col_num1  entry1  col_num2  entry2
col_num3  entry3  -1
col_num4  entry4  col_num5  entry5
col_num6  entry6  -1
```

This input corresponds to two rows: 0 and 1. Row 0 contains `entry1` in column `col_num1`, `entry2` in column `col_num2` and `entry3` in column `col_num3`. Row 1 contains `entry4` in column `col_num4`, `entry5` in column `col_num5` and `entry6` in column `col_num6`.

NOTE: row and column numbers must start from 0.

NOTE: `AZ_read_msr_matrix()` is inefficient for large matrices.

Parameters

val, *bindx*

On output, these two arrays are allocated and filled with the MSR representation corresponding to the file `.data`.

Prototype

```
void AZ_read_update(int *N_update, int **update, int *proc_config,
                    int N, int chunk, int input_option )
```


Description

This routine initializes *update* to the global indices updated by this processor and initializes *N_update* to the total number of elements to be updated.

Parameters

<i>N_update</i>	On output, number of elements updated by processor.
<i>update</i>	On output, <i>update</i> is allocated and contains a list of elements updated by this processor in ascending order.
<i>chunk</i>	Number of indices within a group. For example, $chunk == 2 \Rightarrow chunk_0 = \{0, 1\}$, and $chunk_1 = \{2, 3\}$.
<i>N</i>	Total number of chunks in the vector.
<i>input_option</i>	
AZ_LINEAR	Processor 0 is assigned the first $\left\lfloor \frac{N+P-1}{P} \right\rfloor$ chunks, processor 1 is assigned the next $\left\lfloor \frac{N+P-2}{P} \right\rfloor$ chunks, etc. where $P = \text{proc_config}[\text{AZ_N_proc}]$.
AZ_BOX	The processor system is viewed as a $p_2 \times p_1 \times p_0$ where $p_i = 2^{\lfloor (k+i)/3 \rfloor}$ (so $\text{proc_config}[\text{AZ_N_proc}]$ must equal 2^k). The chunks are viewed as an $n \times n \times n$ cube where n is divisible by each p_i . Chunks are distributed into uniform boxes such that each processor has the same number of chunks.
AZ_FILE	Read the $\text{proc_config}[\text{AZ_N_proc}]$ lists contained in the file <i>update</i> . Each list contains a set of global indices preceded by the of number of indices in this set. List 0 is sent to processor $\text{proc_config}[\text{AZ_N_proc}] - 1$, list 1 is sent to processor $\text{proc_config}[\text{AZ_N_proc}] - 2$, etc. Note: A graph partitioning package named Chaco [2] produces files in this format.

Prototype

```
void AZ_reorder_matrix(int N_update, int *bindx, double *val, int *update_index,
                      int *extern_index, int *indx, int *rpnt, int *bpnt,
                      int N_external, int *cpnt, int option, int mat_type)
```

Description

Reorder the matrix so that it corresponds to the new ordering given by *update_index* and *extern_index*. Specifically, global matrix entry (*update*[*i*], *update*[*j*]) which was stored as local matrix entry (*i*, *j*) is stored as (*update_index*[*i*], *update_index*[*j*]) on output. Likewise, global matrix entry (*update*[*i*], *external*[*k*]) which was stored as local matrix entry (*i*, *k* + *N_update*) is stored locally as (*update_index*[*i*], *extern_index*[*k*]) on output. Called by AZ_transform.

IMPORTANT: This routine assumes that *update_index* contains two sequences of numbers that are ordered but intertwined. For example,

<i>update_index</i> :	4	5	0	6	1	2	3	7
sequence 1:			0		1	2	3	
sequence 2:	4	5		6				7

Parameters

<i>option</i>	On input, indicates whether to reorder update elements.
AZ_ALL	All the rows and columns are renumbered.
AZ_EXTERNS	Only columns corresponding to external elements are renumbered.
<i>mat_type</i>	On input, indicates matrix format.
AZ_MSR_MATRIX	DMSR matrix format.
AZ_VBR_MATRIX	DVBR matrix format.
<i>bindx</i> , <i>val</i> , <i>indx</i> , <i>rpnt</i> , <i>bpnt</i> , <i>cpnt</i>	On input, matrix ordered as described above. On output, matrix reordered using <i>update_index</i> and <i>extern_index</i> as described above.

Prototype

```
void AZ_set_message_info(int N_external, int *extern_index, int N_update,
                        int *external, int *extern_proc, int *update,
                        int *update_index, int *proc_config, int *cpnt,
                        int **data_org, int mat_type)
```

Description

Initialize *data_org* so that local communications can occur to support matrix vector products. This includes:

- determine neighbors with which we send or receive.
- determine the total number of elements that we send and allocate *data_org*.
- initialize *data_org* as described in Section 6.1.

Note: *data_org*[AZ_name] is set to a number (starting from 1) that is incremented each time AZ_set_message_info is called.

Called by AZ_transform.

NOTE: Implicitly the neighbors are numbered using the ordering of the external elements (which have been previously ordered such that elements updated by the same processor are contiguous).

Parameters

<i>data_org</i>	On output, <i>data_org</i> is allocated and completely initialized as described in Section 6.1.
<i>mat_type</i>	On input, indicates matrix format.
AZ_MSR_MATRIX	DMSR matrix.
AZ_VBR_MATRIX	DVBR matrix.

Prototype

```
void AZ_solve(double *x, double *b, int *options, double *params, int *indx,
              int *bindx, int *rpntr, int *cpntr, int *bpntr, double *val,
              int *data_org, double *status, int *proc_config)
```

Description

Solve the system of equations $Ax = b$ via an iterative method where the matrix A is encoded in *indx*, *bindx*, *rpntr*, *cpntr*, *bpntr* and *val* (see Section 3 and Section 2).

Parameters

<i>x</i>	On input <i>x</i> contains the initial guess. On output <i>x</i> contains the solution to linear system.
<i>b</i>	Right hand side of linear system.
<i>options</i> , <i>params</i>	Options and parameters used during the solution process (Section 2).
<i>status</i>	On output, status of iterative solver (Section 2).

Prototype

```
void AZ_sort(int *list1, int N, int *list2, double *list3 )
```

Description

Sort the elements in *list1*. Additionally, move the elements in *list2* and *list3* so that they correspond with the moves done to *list1*. NOTE: If *list2* == NULL, *list2* is not manipulated. If *list3* == NULL, *list3* is not manipulated.

Parameters

<i>list1</i>	On input, values to be sorted. On output, sorted values (i.e. $list1[i] \leq list1[i+1]$)
<i>N</i>	On input, length of lists to be sorted.
<i>list2</i>	On input, a list associated with <i>list1</i> . On output, if <i>list1[k]</i> on input is now stored in <i>list1[j]</i> on output, <i>list2[k]</i> on input is also stored as <i>list2[j]</i> on output.
<i>list3</i>	On input, a list associated with <i>list1</i> . On output, if <i>list1[k]</i> on input is now stored in <i>list1[j]</i> on output, <i>list3[k]</i> on input is also stored as <i>list3[j]</i> on output. Note: if <i>list3</i> == NULL on input, it is unchanged on output.

Prototype

```
void AZ_transform(int *proc_config, int **external, int *bindx,  
                 double *val, int *update, int **update_index,  
                 int **extern_index, int **data_org, int N_update,  
                 int *indx, int *bpntr, int *rpntr, int **cpntr, int mat_type)
```

Description

Convert the global matrix description to a distributed local matrix format (see Section 2 and Section 6.4).

Parameters

<i>external</i>	On output, allocated and set to components that must be communicated during the matrix vector multiply.
<i>bindx, val, index, bpntr, rpntr</i>	On input, matrix arrays (MSR or VBR) corresponding to global format. On output, matrix arrays (DMSR or DVBR) corresponding to local format. See Section 2.
<i>update_index</i>	On output, allocated and set such that <i>update_index</i> [<i>i</i>] is the local numbering corresponding to <i>update</i> [<i>i</i>].
<i>extern_index</i>	On output, allocated and set such that <i>extern_index</i> [<i>i</i>] is the local numbering corresponding to <i>external</i> [<i>i</i>].
<i>data_org</i>	On output, allocated and set to data layout information, see Section 6.1.
<i>cpntr</i>	On output, allocated and set for VBR matrices to the column pointer array.
<i>mat_type</i>	On input, matrix format: either AZ_VBR_MATRIX or AZ_MSR_MATRIX.

REFERENCES

- [1] S. Carney, M. Heroux, and G. Li. A proposal for a sparse BLAS toolkit. Technical report, Cray Research Inc., Eagan, MN, 1993.
- [2] B. Hendrickson and R. Leland. The Chaco user's guide - version 1.0. Technical Report Sand93-2339, Sandia National Laboratories, Albuquerque NM, 87185, August 1993.
- [3] J. N. Shadid and R. S. Tuminaro. Sparse iterative algorithm software for large-scale MIMD machines: An initial discussion and implementation. *Concurrency: Practice and Experience*, 4(6):481-497, September 1992.
- [4] J. N. Shadid and R. S. Tuminaro. Aztec - a parallel preconditioned Krylov solver library: Algorithm description version 1.0. Technical Report Sand95:in preparation, Sandia National Laboratories, Albuquerque NM, 87185, August 1995.
- [5] R. S. Tuminaro, J. N. Shadid, and S. A. Hutchinson. Parallel sparse matrix vector multiply software for matrices with data locality. *Submitted to BIT*, August 1995.
- [6] Z. Zlatev, V.A. Barker, and P.G. Thomsen. SSLEST - a FORTRAN IV subroutine for solving sparse systems of linear equations (user's guide). Technical report, Institute for Numerical Analysis, Technical University of Denmark, Lyngby, Denmark, 1978.

EXTERNAL DISTRIBUTION:

Steve Ashby
Lawrence Livermore Nat. Lab.
M/S L-316
PO Box 808
Livermore, CA 94551-0808

D. M. Austin
Army High Per. Comp. Res. Cntr.
University of Minnesota
1100 S. Second St.
Minneapolis, MN 55415

Rob Bisseling
Department of Mathematics
Budapestlaan 6, De Uithof, Utrecht
PO Box 80.010, 3508 TA Utrecht
The Netherlands

Petter Bjorstad
University of Bergen
Institutt for Informatikk
Thomohlengst 55
N-5008 Bergen, Norway

Randall Bramley
Dept of CSci
Indiana University
Bloomington IN 47405

G. F. Carey
ASE/EM Dept., WRW 305
University of Texas
Austin, TX 78712

Steven P. Castillo
Klipsch School of Electrical & Computer Eng.
New Mexico State University
Box 30001
Las Cruces, NM 88003-0001

J. M. Cavallini
US Department of Energy
OSC, ER-30, GTN
Washington, DC 20585

T. Chan
UCLA
405 Hilgard Ave.
Los Angeles, CA 90024-7009

Warren Chernock
Scientific Advisor DP-1
US Department of Energy
Forestal Bldg. 4A-045
Washington, DC 20585

Doug Cline
The University of Texas System
Center for High Performance Computing
10100 Burnett Road, CMS 1.154
Austin, Texas 78758

Tom Coleman
Dept. of Computer Science
Upton Hall
Cornell University
Ithaca, NY 14853

Pedro Diniz
Computer Science Department
Engineering I Bldg, Room 2106
University of California at Santa Barbara
Santa Barbara, CA 93106

J. J. Dongarra
Computer Science Dept.
104 Ayres Hall
University of Tennessee
Knoxville, TN 37996-1301

I. S. Duff
CSS Division
Harwell Laboratory
Oxfordshire, OX11 0RA
United Kingdom

Alan Edelman
Dept. of Mathematics
MIT
Cambridge, MA 02139

Steve Elbert
US Department of Energy
OSC, ER-30, GTN
Washington, DC 20585

H. Elman
Computer Science Dept.
University of Maryland
College Park, MD 20842

R. E. Ewing
Mathematics Dept.
University of Wyoming
PO Box 3036 University Station
Laramie, WY 82071

Charbel Farhat
Dept. Aerospace Engineering
UC Boulder
Boulder, CO 80309-0429

J. E. Flaherty
Computer Science Dept.
Rensselaer Polytech Inst.
Troy, NY 12181

G. C. Fox
Northeast Parallel Archit. Cntr.
111 College Place
Syracuse, NY 13244

R. F. Freund
NRaD- Code 423
San Diego, CA 99152-5000

D. B. Gannon
Computer Science Dept.
Indiana University
Bloomington, IN 47401

Horst Gietl
nCUBE Deutschland
Hanauer Str. 85
8000 Munchen 50
Germany

Paul Giguere
Group TSA-8
MS K575
Los Alamos National Laboratory
Los Alamos, NM 87545

John Gilbert
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304

G. H. Golub
Computer Science Dept.
Stanford University
Stanford, CA 94305

Anne Greenbaum
New York University
Courant Institute
251 Mercer Street
New York, NY 10012-1185

Satya Gupta
Intel SSD
Bldg. CO6-09, Zone 8
14924 NW Greenbrier Parkway
Beaverton, OR 97006

J. Gustafson
Computer Science Dept.
236 Wilhelm Hall
Iowa State University
Ames, IA 50011

Doug Harless
NCUBE
2221 East Lamar Blvd., Suite 360
Arlington, TX 76006

Michael Heath
Univ. of Ill., Nat. CSA
4157 Bechman Institute
405 North Matthews Ave.
Urbana, IL 61801-2300

Mike Heroux
Cray Research Park
655F Lone Oak Drive
Eagan, MN 55121

Dan Hitchcock
US Department of Energy
SCS, ER-30 GTN
Washington, DC 20585

Fred Howes
US Department of Energy
OSC, ER-30, GTN
Washington, DC 20585

Christopher R. Johnson
Department of Computer Science
3484 MEB
University of Utah
Salt Lake City, UT 84112

David Keyes
Dept. of Mechanical Engineering
Yale University
PO Box 2159, Yale Station
New Haven, CT 06520-2159

David Kincaid
Center for Numerical Analysis
RLM 13.150
University of Texas
Austin, TX 78713-8510

T. A. Kitchens
US Department of Energy
OSC, ER-30, GTN
Washington, DC 20585

Vipin Kumar
Computer Science Department
Institute of Technology
200 Union Street S.E.
Minneapolis, MN 55455

Joanna Lees
Intel Corp.
Scalable Systems Division
CO1-15
15201 NW Greenbrier Parkway
Beaverton, OR 97006

John Lewis
Boeing Corp.
M/S 7L-21
P.O. box 24346
Seattle, WA 98124-0346

T. A. Manteuffel
Department of Mathematics
University of Co. at Denver
Denver, CO 80202

S. F. McCormick
Computer Mathematics Group
University of CO at Denver
1200 Larimer St.
Denver, CO 80204

Robert McLay
University of Texas at Austin
Dept. ASE-EM
Austin, TX 78712

P. C. Messina
158-79
Mathematics & Comp Sci. Dept.
Caltech
Pasadena, CA 91125

C. Moler
The Mathworks
24 Prime Park Way
Natick, MA 01760

Gary Montry
Southwest Software
11812 Persimmon, NE
Albuquerque, NM 87111

D. B. Nelson
US Department of Energy
OSC, ER-30, GTN
Washington, DC 20585

Kwong T. Ng
Klipsch School of Electrical & Computer Eng.
New Mexico State University
Box 30001
Las Cruces, NM 88003-0001

J. M. Ortega
Applied Mathematics Dept.
University of Virginia
Charlottesville, VA 22903

Linda Petzold
L-316
Lawrence Livermore Natl . Lab.
Livermore, CA 94550

Barry Peyton
Mathematical Sciences Section
Oak Ridge National Laboratory
P.O. Box 2008, Bldg. 6012
Oak Ridge, TN 37831-6367

Paul Plassman
Math and Computer Science Division
Argonne National Lab
Argonne, IL 60439

Claude Pommerell
AT&T Bell Labs
600 Mountain Ave, Room 2C-548A
Murray Hill, NJ 07974-0636

Alex Pothen
Department of Computer Science
Old Dominion University
Norfolk, VA 23529-0162
J. Rattner
Intel Scientific Computers
15201 NW Greenbriar Pkwy.
Beaverton, OR 97006

Patrick Riley
Intel-SSD
600 S. Cherry St., Suite 700
Denver, CO 80222

Ed Rothberg
Silicon Graphics, Inc.
MS 7L-580
2011 N. Shoreline Blvd.
Mountain View, CA 94043

Y. Saad
University of Minnesota
4-192 EE/CSci Bldg.
200 Union St.
Minneapolis, MN 55455-0159

P. Sadayappan
Ohio State University
Comp. & Inf. Sci., 228 Boltz Hall
2036 Neil Avenue
Columbus, OH 43210-1277

Joel Saltz
Computer Science Department
A.V. Williams Building
University of Maryland
College Park, MD 20742

A. H. Sameh
CSR/D, University of Illinois
305 Talbot Laboratory
104 S. Wright St.
Urbana, IL 61801

P. E. Saylor
Dept. of Comp. Science
222 Digital Computation Lab
University of Illinois
Urbana, IL 61801

Carl Scarnick
San Diego Supercomputer Center
P.O. Box 85608
San Diego, CA 92186-9784

Rob Schreiber
RIACS
NASA Ames Research Center
Mail Stop T045-1
Moffett Field, CA 94035-1000

Elliott Schulman
nCUBE Corp.
3575 9th St.
Boulder, Co. 80304

M. H. Schultz
Department of Computer Science
Yale University
PO Box 2158
New Haven, CT 06520

Mark Seager
LLNL, L-80
PO box 803
Livermore, CA 94550

Horst Simon
Silicon Graphics
Mail Stop 7L-580
2011 N. Shoreline Blvd.
Mountain View, CA 94043

Richard Sincovec
Mathematical Sciences Section
Oak Ridge Nat. Lab.
P.O. Box 2008, Bldg. 6012
Oak Ridge, TN 37831-6367

Vineet Singh
HP Labs, Bldg. 1U, MS 14
1501 Page Mill Road
Palo Alto, CA 94304

Anthony Skjellum
Mississippi State University
Computer Science
PO Drawer CS
Mississippi State, MS 39762

L. Smarr
Director, Supercomputer Apps.
152 Supercomputer Applications
Bldg. 605 E. Springfield
Champaign, IL 61801

Burton Smith
Tera Computer Co
400 N. 34th St., Suite 300
Seattle, WA 98103

Barry Smith
Department of Mathematics
UCLA
Los Angeles, CA 90024-1555

Harold Trease
Los Alamos National Lab
PO Box 1666, MS F663
Los Alamos, NM 87545

C. VanLoan
Department of Computer Science
Cornell University, Rm. 5146
Ithaca, NY 14853

John VanRosendale
ICASE, NASA Langley Research Center
MS 132C
Hampton, VA 23665

Steve Vavasis
Department of Computer Science / ACRI
722 Engineering and Theory Center
Cornell University
Ithaca, NY 14853

R. G. Voigt
MS 132-C
NASA Langley Resch Cntr, ICASE
Hampton, VA 36665

Phuong Vu
Cray Research, Inc.
19607 Franz Road
Houston, TX 77084

Steven J. Wallach
Convex Computer Corp.
3000 Waterview Parkway
PO Box 833851
Richardson, TX 75083-3851

G. W. Weigand
DARPA/CSTO
3701 N. Fairfax Ave.
Arlington, VA 22203-1714

Olof B. Widlund
Dept. Computer Science
Courant Inst., NYU
251 Mercer St.
New York, NY 10012

Roy Williams
California Institute of Technology
206-49
Pasadena, CA 91104

INTERNAL DISTRIBUTION:

1	MS 0360	A.R.C. Westwood, 1000
1	MS 0151	Gerold Yonas, 9000
1	MS 0321	Ed Barsis, 1400
1	MS 0321	William Camp, 1400
1	MS 0601	Harry K. Moffat, 1126
1	MS 1111	Sudip Dosanjh, 1421
10	MS 1111	Scott Hutchinson, 1421
10	MS 1111	John N. Shadid, 1421
1	MS 1111	Andrew G. Salinger, 1421
1	MS 1111	Gary L. Hennigan, 1421
1	MS 1111	Martin Lewitt, 1421
1	MS 1111	Mark P. Sears, 1421
1	MS 1111	Daniel Barnette, 1421
1	MS 1111	Steven J. Plimpton, 1421
1	MS 1111	David R. Gardner, 1421
1	MS 1110	Richard C. Allen, 1422
1	MS 1110	Bruce A. Hendrickson, 1422
1	MS 1110	David E. Womble, 1422
10	MS 1110	Ray S. Tuminaro, 1422
1	MS 1110	Lydie Prevost, 1422
1	MS 1109	Art Hale, 1424
1	MS 1109	Ted Barragy, 1424
1	MS 1109	Robert W. Leland, 1424
1	MS 1109	Karen Devine, 1424
1	MS 1109	Courtenay Vaughn, 1424
1	MS 1109	James Tomkins, 1424
1	MS 0819	J. Michael McGlaun, 1431
1	MS 0819	James S. Perry, 1431
1	MS 0819	Allen C. Robinson, 1431
1	MS 00439	David R. Martinez, 1434
1	MS 0833	Johnny H. Biffle, 1503
1	MS 0827	Dave K. Gartling, 1511
1	MS 0827	Randy Schunk, 1511
1	MS 0827	Phil Sackinger, 1511
1	MS 0827	Mario Martinez, 1511
1	MS 0827	Mike Glass, 1511
1	MS 0827	Bob McGrath, 1511
1	MS 0827	Poly Hopkins, 1511
1	MS 0827	Jim Schutt, 1511
1	MS 0827	Melinda Sirmar, 1511
1	MS 0834	Robert B. Campbell, 1512
1	MS 0835	Roy E. Hogan Jr., 1513
1	MS 0835	Mark A. Christon, 1513
1	MS 0826	Robert J. Cochran, 1514
1	MS 0750	Greg A. Newman, 6116
1	MS 0750	David L. Alumbaugh, 6116
1	MS 9214	Juan Meza, 8117
1	MS 9042	Joseph F. Grcar, 8745
1	MS 9042	Greg Evans, 8745
1	MS 1166	Joseph Kotulski, 9352

1	MS 9018	Central Technical Files, 8523-2
5	MS 0899	Technical Library, 13414
1	MS 0619	Print Media, 12615
2	MS 0100	Document Processing, 7613-2
		For DOE/OSTI