Experiments on Hypergraph Models for Parallelizing Preconditioned Iterative Methods*

Bora Uçar and Cevdet Aykanat
Department of Computer Engineering,
Bilkent University, 06800, Ankara, Turkey
{ubora,aykanat}@cs.bilkent.edu.tr

Abstract

We have developed hypergraph models to efficiently parallelize the preconditioned iterative methods that use explicit preconditioners. The models are discussed elsewhere. Here, we report our experiments.

Keywords. matrix partitioning, preconditioning, iterative method, parallel computing

Technical Report:	No: BU-CE-0413
	Department of Computer Engineering,
	Bilkent University, 06800 Ankara, Turkey.
	Available at http://www.cs.bilkent.edu.tr/publications/
	Date: October 12, 2004.
	Updated: September 27, 2005.

*This work is partially supported by the Scientific and Technical Research Council of Turkey under grant 103E028

1 Experiments

We chose the right preconditioned BiCGStab method given in Fig. 1 to evaluate the proposed simultaneous partitioning method. We used a set of unsymmetric sparse matrices which were obtained from University of Florida Sparse Matrix Collection [6]. Approximate inverse preconditioners were obtained using SPAI version 3.0 [7]. Factored approximate inverses were obtained using AINV [1]. These two programs have parameters that affect the quality of the preconditioner matrices. However, we set the parameters in such a way that the number of nonzeros of the approximate inverse or the total number of nonzeros of the factors of the approximate inverse is at most twice and at least half of the number of nonzeros of the coefficient matrix. We adjusted the tolerance parameter eps, number of nonzero entries allowed per step mn, and the number of steps ns in SPAI. In AINV, we adjusted the drop tolerance parameter τ. The properties of the matrices, approximate inverses, and factors of the approximate inverses are given in Table 2. In the table, the coefficient matrices are listed with a suffix of A; the approximate inverse matrices are listed with a suffix of M; the factors of the approximate inverse matrices are listed with suffixes of Z and W, where approximate inverse is equivalent to ZW. The composite hypergraphs were partitioned using PaToH [3] with default parameters. The imbalance among processors' loads is kept below 10% in all partitioning instances. Throughout this section, we use "SPAI-matrices" to refer to a pair of a coefficient matrix and its approximate inverse preconditioner. Similarly, we use "AINV-matrices" to refer to a triplet of coefficient matrix and the factors of its approximate inverse preconditioner.

For SPAI-matrices, we choose the partitioning dimensions as columnwise-rowwise (CR) and rowwise-columnwise (RC) for the A and M matrices in the given order. Since the partitioning tool PaToH [3] incorporates randomized algorithms it was run 20 times starting from different random seeds for 8-, 16-, 32-, and 64-way partitionings of every hypergraph. Averages of the resulting communication patterns of these runs are displayed in the following tables. Although the main objective in the simultaneous partitioning method is the minimization of the total communication volume, the results for other communication cost metrics such as the total number of messages, the maximum volume and the maximum number of messages handled by a single processor are also given.

2 Composite versus simple hypergraph partitioning

Here we evaluate two alternative approaches to partitioning composite hypergraphs. These alternative approaches are based on partitioning simple hypergraphs, i.e., partitioning hypergraphs of a single matrix. The first alternative is to obtain independent partitionings on the matrices by partitioning the simple hypergraph models of the coefficient and the preconditioner matrices independently. This approach requires vector reordering in between the two matrix-vector multiplies. We discuss this alternative in §2.1. The second alternative is to obtain the same partition for the coefficient and preconditioner matrices. For this purpose, we partition the coefficient matrices symmetrically by rows or columns using hypergraph models and apply the resulting partitions to the preconditioner matrices as well. Note that since we partition a single matrix, graph model can also be used. Observe that this alternative avoids vector reordering operation by partitioning all vector conformably. We discuss this alternative in §2.2.

2.1 Simple hypergraph partitioning: Independent partitions on the matrices

Tables 3, 4, 13, and 14 display the average communication patterns of the simultaneous and the individual partitionings for SPAI-matrices. The tables also show the volume of communication required to reorder the vector elements—in an iteration of the BiCGStab method—when the matrices are partitioned individually. Suppose that symmetric partitionings PAP^T and QMQ^T were obtained on the A and M matrices. Then, for each iteration we have to reorder $\hat{p}$ and $\hat{s}$ from Q to P after the matrix-vector multiplies at lines 12 and 17 of the BiCGStab method (see Fig. 1), respectively. We also have to reorder v and t from P to Q before the vector update operation at line 15 and the inner product at line 19, respectively. The volume of communication in the reordering operation is obtained according to the permutation matrices that give the best volume for the individual matrices. The actual total volume of communication in the individual partitioning method can be obtained by adding the volume of reordering operations to the total volumes of the individual partitionings. In all of the partitioning instances, the volume of communication in the reordering operation itself is higher than the volume of communication in the simultaneous partitioning. These high volumes of communication and the associated message start-up overheads prohibit the application of the individual partitioning method. For example, the individual partitioning method incurs higher total communication volume than the proposed simultaneous partitioning method by factors that vary between 2.8 (`cage12`) and 24.6 (`epb3`) with an overall average factor of 8.6 for 32-way CR partitioning. The average factor in 64-way CR partitioning is 6.7. For RC partitioning, the average factors are 5.8 and 4.2 for 32- and 64-way partitionings, respectively.

Tables 5, 6, 16, and 15 display the averages of the communication patterns of the 8-, 16-, 32- and 64-way simultaneous and individual partitionings for AINV-matrices. We give the experiments in which the partitioning dimensions are chosen as columnwise-rowwise-columnwise (CRC) and rowwise-columnwise-rowwise (RCR) for the A, Z, and W matrices in the given order. For AINV-matrices, the individual partitioning method requires two additional reordering operations which are necessary for the chains of matrix-vector multiplies at lines 12 and 17 of the BiCGStab method. For AINV-matrices the minimum ratio of the communication volumes in the individual partitionings to that of the simultaneous partitionings is 2.7 which is obtained for the 64-way partitioning of the `cage11` matrix. The maximum ratio is 14.9 which is obtained for the 32-way partitioning of the `big` matrix. The average of the ratios for 32- and 64-way partitionings are 10.1 and 7.2, respectively. The extremely high communication volumes definitely prohibit the application of the individual partitioning method. The minimum and maximum ratios in the 32- and 64-way RCR partitionings are sightly better than those in CRC case (2.21 and 13.0) implying a slightly better average (8.6 and 6.5 for 32- and 64-way partitionings, respectively) but still not tolerable.

Consider the difference between the total communication volumes of the simultaneous and individual partitionings (without the reordering cost). The increases in the total communication volume values for the simultaneous partitionings remain below 26% of those of the individual partitionings, on the average, for the 32-way CR partitioning instances. The minimum and the maximum of these increases are 13% (`Zhao1`) and 61% (`epb3`). The 64-way CR partitionings give better ratios. The average increase is 20% with the minimum and the maximum being 12% and 43% which are obtained for the same matrices. In fact, for each matrix the 64-way CR partitioning gives smaller percent increase than the 32-way CR partitioning. We investigated the 8- and 16-way CR partitionings as well and observed that for each matrix in our data set the larger the number of parts, the smaller the percent increases. The same relation holds for $K = 8, 16, 32$, and 64-way RC partitioning case, except for the 32- and 64-way partitionings of the `epb2` and `mark3jac060` matrices. It also holds for most of the CRC and RCR partitioning of AINV-matrices. See Tables 13, 14, 15, and 16. The reason behind this may be the following. The cutsize function monotonically increases with the increasing number of parts. In other words, the flexibility of finding better partitions reduces by the increasing number of parts. At the limit, where $K = \mathcal{V}$ and all the nets are in cut, the cutsize of a composite hypergraph will be equivalent to the sum of the cutsizes of the individual hypergraphs (e.g., $nnz(A) + nnz(M) - 2m$) that forms it. Therefore the difference between the total communication volumes has to converge to zero.

2.2 Simple hypergraph partitioning: The same partition on the matrices

Obtaining a symmetric partition on A and then applying the resulting partition to M results in columnwise-columnwise (CC) or rowwise-rowwise (RR) partitioning on the A and M matrices. Recall that for CC and RR partitioning schemes, there is a communication phase in between the two matrix-vector multiplies. Since the A and M matrices have comparable number of nonzeros (see Table 2), processors' loads for the two matrix-vector multiplies should be balanced separately, i.e., a two-constraint formulation is necessary.

Observe that the approach evaluated here disregards the sparsity pattern of the preconditioner matrices. However, the sparsity patterns of the approximate inverse preconditioners are usually related to the sparsity patterns of the

Table 1: Relation between the sparsity patterns of the coefficient matrices and the approximate inverses. We use A and M to represent the set of the positions of the nonzeros in the corresponding matrices.

Matrix	number of nonzeros			
	$A \cup M$	$A \setminus M$	$M \setminus A$	$\frac{A \cap M}{M}$
Zhao1	234205	67752	53217	0.63
big	147632	56167	38544	0.49
cage11	780776	221054	356068	0.48
cage12	2784199	751663	1339549	0.48
epb2	333794	158767	89341	0.35
epb3	773107	309482	240256	0.42
mark3jac060	397706	227011	121120	0.18
olafu	1357370	342214	637497	0.52
stomach	5182305	2160657	2272022	0.26
xenon1	1520936	339816	642793	0.61

coefficient matrices [4, 8]. Therefore, the partitions on the coefficient matrices are expected to be effective for the preconditioners. To justify this reasoning, we show the relation between the sparsity patterns of the coefficient matrices and the approximate inverses in Table 1. As seen in the table, the relation between the sparsity patterns of the coefficient and preconditioner matrices varies; 63% of the nonzeros of `Zhao1-M` are covered by the nonzeros of `Zhao1-A`, and only 18% of the nonzeros of `mark3jac060-M` are covered by the nonzeros of `mark3jac060-A`. Another reason for using the same partition for the coefficient and preconditioner matrices is the following. Parallel construction of the approximate inverse preconditioners produces preconditioners in such a way that the initial partitions on the coefficient matrices become partitions on the preconditioner matrices. For example, the left approximate inverse preconditioners can be efficiently constructed rowwise when the coefficient matrix A is partitioned rowwise [5]. The construction yields the same rowwise partition on the approximate inverse M. Equivalently, a right approximate inverse preconditioner can be efficiently constructed columnwise when the coefficient matrix A is partitioned columnwise.

The row-net hypergraph model of A can be used to obtain a CC partitioning on A and M. In order to obtain load balance for the two multiplies, the vertices of A are assigned two weights which correspond to the number of nonzeros in the respective columns of A and M matrices. That is, v_i has weights $\langle |c_i(A)|, |c_i(M)| \rangle$. Similarly, the column-net hypergraph model of A, with two weights on the vertices, can be used to obtain an RR partitioning on A and M.

The composite hypergraph model for the CC partitioning scheme is built first, by creating enhanced row-net hypergraph models of $y \leftarrow Ax$ and $w \leftarrow Mz$; second, by applying vertex amalgamation operation to the vertices y_i and $c_i(M)/z_i$, and also to the vertices $c_i(A)/x_i$ and w_i for each i; third, by applying vertex weighting operation in such a way that the vertex $y_i/c_i(M)/z_i$ has weights $\langle 0, |c_i(M)| \rangle$, and the vertex $c_i(A)/x_i$ has weights $\langle |c_i(A), 0| \rangle$. Theoretically, a third constraint formulation is necessary to balance the vector operations, however, we use a two constraint formulation in order to ease the job of the hypergraph partitioning tool. The composite hypergraph model for the RR partitioning scheme is built similarly.

Tables 7 and 8 display the averages of the communication volumes of the 32- and 64-way partitioning of the SPAI-matrices with the composite hypergraphs and simple hypergraph models of the coefficient matrices. The "% gain" columns in these tables show the improvements achieved by the composite hypergraph partitioning as the percentage of the total communication volumes found by partitioning the simple hypergraph of A. The minimum percent improvements are obtained for the `Zhao1` matrices in all cases. The maximum percent improvements are obtained for the `mark3jac060` matrices in all cases. As seen in Table 1, the `Zhao1` matrices have the highest number of common nonzeros, and the `mark3jac060` matrices have the least number of common nonzeros. Although `xenon1-A` covers 61% of `xenon1-M` (second maximum), the percent improvements achieved for these matrices are quite satisfactory. The average of the improvements is 20% in Tables 7 and 8 both for the CC and RR partitioning choices.

We have also experimented with the 32- and 64-way, CC and RR partitionings using single constraint formulation. The communication patterns for these experiments are given in Tables 9 and 10. In the single constraint formulation, the weight of a vertex v_i is set to the sum of the number of nonzeros in the ith columns (rows) of A and M for the CC (RR) partitioning. Both the composite hypergraph and the simple hypergraph formulations were able to obtain balance on the total loads of the processors. Since these formulations ignore the fact that there is a local synchronization, both formulations could not obtain balance on the loads of the processors for the individual matrix-vector multiplies. The composite hypergraph partitioning approach obtained 17% better solutions than the simple hypergraph partitioning on the average. The best and worst improvements are again obtained for the `mark3jac060` and `Zhao1` matrices.

3 Impact of the partitioning dimension

Comparing the lower and upper halves of the Tables 3 and 4 we see that CR partitioning yields better total communication volume than the RC one. The ratio of the average total communication volume in CR partitioning to the average total communication volume in RC partitioning is around 0.74 for the data given in the Tables 3 and 4. This ratio remains the same for the 8- and 16-way partitionings given in the Tables 13 and 14. The standard deviation of these ratios is around 0.13 for each $K = 8, 16, 32, 64$. Note that the matrices do not have dense rows or dense columns. Therefore, it is expected that the rowwise and columnwise partitioning of the matrices in our data set will obtain comparable results. This theoretical expectation is verified by the data given in the total communication volume column under the individual partitioning header in Tables 3, 4, 13, and 14. In the light of this observation we can deduce that the performance difference between the CR and RC partitioning schemes is mainly due to the two-constraint formulation in the RC scheme. This degradation in the multi-constraint formulation is in concordance with the previously reported results [9, 11]. The degradation in our case stems from two facts. First, the additional balance constraints shrink the search space. Second, the heuristics are not very well tailored toward handling the multiple weights.

4 Parallelization results

It is important to see whether the theoretical improvements obtained by the proposed simultaneous partitioning method hold in practice. For this purpose, we have implemented a parallel program for the BiCGStab method. The program uses LAM/MPI 6.5.6 [2] message passing library. The tests were carried out on a Beowulf class [10] PC cluster with 24 nodes. Each node has a 400MHz Pentium-II processor and 128 MB memory. The interconnection network is comprised of a 3COM Superstack II 3900 managed switch connected to Intel Ethernet Pro 100 Fast Ethernet network interface cards at each node. The system runs Linux kernel 2.4.20 and the Debian GNU/Linux 3.0 distribution.

We are not concerned with the numerics of the preconditioners and the BiCGStab method. Therefore, for each matrix we let the BiCGStab run for at most 100 iterations and measure the average running time of a single iteration. In some of the matrices, BiCGStab failed due to a zero ρ observed at line 5 of the method (see Fig. 1). For those matrices, we measure the average running time of the iterations before failure. The running times of the parallel program with the SPAI preconditioners are given in Table 11 in milliseconds. The given running times are the averages of 20 runs corresponding to different partitionings. In order to show how the improvements obtained by the proposed method relate to parallel running times, we give the average communication patterns of the partitionings in Tables 13 and 14.

As seen from Table 11, the CR partitioning gives better speedup values than the RC partitioning for all matrices. On the average, CR obtains speedup values of 6.3 and 9.9 for 8- and 16-way partitionings, respectively, where the highest speedups are 7.3 (`epb3`) and 14.1 (`stomach`). Meanwhile, RC obtains speedups of 5.8 and 8.4, on the average, for 8-way and 16-way partitionings, respectively, where the the highest speedups are 7.2 (`epb3`) and 12.7 (`stomach`). The worst speedups for 8-way partitioning are obtained for the `cage11` matrix by both of the partitioning schemes. The worst speedups for 16-way partitioning are obtained for the `big` and `cage11` matrices by the CR and RC schemes, respectively. As seen from Tables 13 and 14, the `cage11` matrix-pair has inferior communication pattern than all but the `cage12` matrix-pair in terms of the total and maximum number of messages metrics. Therefore, we were already expecting to have the worst speedups with the `cage` matrices. The `big` matrix has the smallest number of nonzeros. This low granularity of computations may be the reason behind having the worst speedup with 16-way CR partitioning of the `big` matrix. The same reasoning may also explain why we had better speedups in `cage12` than those in `cage11`.

We have also experimented with the CRC and RCR partitioning schemes for AINV-matrices and give the speedup values in Table 12. The speedup values are not as good as those for SPAI-matrices as expected because of the three load balance constraints. The best speedups for 8- and 16-way CRC partitionings are 6.7 and 8.9, respectively. The best speedups for 8- and 16-way RCR partitionings are 6.4 and 8.9, respectively.

Lastly, we comment on the execution time of the simultaneous partitioning method. Let the sum of the times elapsed in individual partitionings (of the SPAI- and AINV-matrices) be 1.0. Then, the average running times of the simultaneous partitioning of the SPAI-matrices with the CR and RC schemes are 1.4 and 1.2, respectively, for all $K = 8, 16, 32$, and 64. The average running times of the simultaneous partitioning of the AINV-matrices with both the CRC and RCR schemes are close to 1.4. These increases are acceptable because the simultaneous partitioning method obtains much smaller total communication volume than the individual partitioning method combined with the reordering cost. Timings for the CR and RC partitionings for SPAI-matrices are given in Tables 17 and 18. Timings for the CRC and RCR partitioning schemes of AINV-matrices are given in Tables 19 and 20.

BiCGStab(A,M,x,b)#Solve Ax=b using the right preconditioner M
begin
$r^{(0)} = b - AMx^{(0)}$ for some initial $x^{(0)} = x$
$\tilde{r} = r^{(0)}$
for $i = 1, 2, \cdots$ **do**
  $\rho_{i-1} = \tilde{r}^T r^{(i-1)}$
  if $\rho_{i-1} = 0$ method fails
  if $i = 1$
    $p^{(i)} = r^{(i-1)}$
  else
    $\beta_{i-1} = (\rho_{i-1}/\rho_{i-2})(\alpha_{i-1}/\omega_{i-1})$
    $p^{(i)} = r^{(i-1)} + \beta_{i-1}\left(p^{(i-1)} - \omega_{i-1} v^{(i-1)}\right)$
  endif
  $\hat{p} = Mp^{(i)}$
  $v^{(i)} = A\hat{p}$
  $\alpha_i = \rho_{i-1}/\tilde{r}^T v^{(i)}$
  $s = r^{(i-1)} - \alpha_i v^{(i)}$
  check norm of s; if small enough; set $x^{(i)} = x^{(i-1)} + \alpha_i p^{(i)}$ and stop
  $\hat{s} = Ms$
  $t = A\hat{s}$
  $\omega_i = t^T s / t^T t$
  $x^{(i)} = x^{(i-1)} + \alpha_i p^{(i)} + \omega_i s$
  $r^{(i)} = s - \omega_i t$
  check convergence; continue if necessary
  for continuation it is necessary that $\omega_i \neq 0$
endfor
end

Figure 1: Preconditioned BiCGStab using the approximate inverse M as a right preconditioner.

Table 2: Properties of test matrices.

Matrix	number of	number of nonzeros					
	rows/cols	total	average	row		col	
			row/col	min	max	min	max
Zhao1-A	33861	166453	4.9	3	6	2	7
big-A	13209	91465	6.9	3	12	3	12
cage11-A	39082	559722	14.3	3	31	3	31
cage12-A	130228	2032536	15.6	5	33	5	33
epb2-A	25228	175027	6.9	3	87	3	87
epb3-A	84617	463625	5.5	3	6	3	7
mark3jac060-A	27449	170695	6.2	2	44	2	47
olafu-A	16146	1015156	62.9	24	89	24	89
stomach-A	213360	3021648	14.2	7	19	6	22
xenon1-A	48600	1181120	24.3	1	27	1	27
SPAI							
Zhao1-M	33861	180988	5.3	1	11	1	16
big-M	13209	109088	8.3	2	22	1	21
cage11-M	39082	424708	10.9	2	51	2	21
cage12-M	130228	1444650	11.1	1	62	2	21
epb2-M	25228	244453	9.7	2	177	2	21
epb3-M	84617	532851	6.3	2	20	2	20
mark3jac060-M	27449	276586	10.1	1	37	1	21
olafu-M	16146	719873	44.6	5	114	4	46
stomach-M	213360	2910283	13.6	2	120	2	46
xenon1-M	48600	878143	18.1	1	35	1	21
AINV; $M = ZW$							
Zhao1-Z	33861	179803	5.3	1	13	1	28
Zhao1-W	33861	57832	1.7	1	5	1	6
big-Z	13209	56302	4.3	1	11	1	13
big-W	13209	56314	4.3	1	13	1	11
cage11-Z	39082	302775	7.7	1	26	1	110
cage11-W	39082	299939	7.7	1	26	1	32
epb2-Z	25228	116161	4.6	1	13	1	22
epb2-W	25228	107620	4.3	1	36	1	19

Table 3: Communication patterns for 32-way simultaneous and individual partitionings for SPAI-matrices.

	Simultaneous partitioning				Individual partitioning				
	Volume		Message		Volume		Message		Reorder
Matrix	tot	max	tot	max	tot	max	tot	max	Volume
	CR				C/R				
Zhao1-A	9419	453	248.1	13.2	8131	340	217.6	11.4	
Zhao1-M	7927	415	251.8	13.4	7174	334	250.8	13.8	135412
big-A	2455	128	162.1	8.9	2071	91	150.9	7.5	
big-M	2357	133	169.4	9.8	1946	91	155.8	7.8	52828
cage11-A	47979	2473	601.5	26.4	42424	2068	444.1	21.3	
cage11-M	29021	1515	640.8	27.2	24460	1189	495.0	23.3	148924
cage12-A	162030	8313	783.8	29.9	142434	6771	614.1	27.4	
cage12-M	93273	4783	815.0	30.4	76776	3519	660.2	28.8	488032
epb2-A	4846	317	233.8	15.8	4162	243	212.7	15.0	
epb2-M	4943	287	186.8	10.7	3918	188	161.8	9.4	100912
epb3-A	5938	315	168.1	9.2	3705	166	126.5	6.1	
epb3-M	7262	380	169.2	9.1	4478	203	141.9	7.3	317008
mark3jac060-A	13519	631	347.7	17.5	9735	377	266.7	12.1	
mark3jac060-M	14578	697	324.0	16.9	11648	460	298.3	14.2	109508
olafu-A	10390	672	155.2	9.2	8394	444	127.9	7.2	
olafu-M	18197	1180	197.6	11.2	15023	890	152.4	8.4	62312
stomach-A	34872	1864	187.8	10.4	26075	976	178.9	7.6	
stomach-M	41181	2022	193.7	10.7	30306	1221	152.6	7.1	853440
xenon1-A	21833	1085	291.4	14.5	19090	824	242.6	11.9	
xenon1-M	29525	1431	314.1	15.8	23634	1003	262.6	13.3	180032
	RC				R/C				
Zhao1-A	9815	564	245.8	12.8	7801	347	218.1	11.7	
Zhao1-M	10386	568	244.2	12.7	8326	386	234.7	13.2	135444
big-A	3536	214	185.4	10.1	2083	92	152.9	7.7	
big-M	5022	292	189.3	9.9	3521	173	161.9	8.2	52824
cage11-A	59783	3988	787.0	30.6	42539	1993	446.6	21.6	
cage11-M	46953	2263	776.5	29.4	31419	1483	557.5	25.4	150560
cage12-A	192970	10335	923.9	31.0	142734	6222	613.9	26.9	
cage12-M	141307	7303	915.5	31.0	95652	4397	722.0	29.8	511804
epb2-A	6919	516	331.0	19.8	3944	221	207.7	11.4	
epb2-M	7889	515	240.9	14.1	4823	232	173.6	9.9	100912
epb3-A	12771	1248	244.7	16.2	4840	204	143.3	7.5	
epb3-M	13461	1383	242.3	16.1	4720	218	141.4	7.5	337660
mark3jac060-A	14277	760	390.1	19.7	9693	394	327.8	15.8	
mark3jac060-M	15896	756	355.2	18.3	12589	524	311.7	14.9	109796
olafu-A	15169	999	209.7	13.1	8469	451	126.8	7.2	
olafu-M	23002	1325	259.0	15.6	14601	817	153.4	8.8	62648
stomach-A	53375	3853	225.2	13.3	26217	978	177.9	7.7	
stomach-M	62102	3981	230.0	13.7	32674	1329	161.1	7.5	853440
xenon1-A	26536	1398	349.9	19.1	19018	813	242.8	12.1	
xenon1-M	34745	1734	376.1	20.1	23484	983	264.2	12.8	191872

Table 4: Communication patterns for 64-way simultaneous and individual partitionings for SPAI-matrices.

	Simultaneous partitioning				Individual partitioning				
	Volume		Message		Volume		Message		Reorder
Matrix	tot	max	tot	max	tot	max	tot	max	Volume
	CR				C/R				
Zhao1-A	13026	327	592.2	17.0	11421	237	529.0	13.1	
Zhao1-M	10809	305	598.0	17.4	9857	234	583.8	16.4	135080
big-A	3666	98	336.5	9.2	3210	69	326.3	8.6	
big-M	3562	102	352.8	9.8	3054	70	339.4	8.8	52824
cage11-A	63779	1732	1585.7	39.5	58177	1463	1173.2	30.9	
cage11-M	38714	1249	1709.0	43.5	32937	835	1262.2	33.5	153784
cage12-A	207077	6111	2229.6	51.1	185531	4711	1626.4	40.4	
cage12-M	119287	3678	2366.2	52.5	98493	2552	1731.2	44.3	504628
epb2-A	6731	259	463.5	22.2	5984	172	439.5	18.6	
epb2-M	7215	211	381.2	12.4	5967	138	343.6	10.8	99904
epb3-A	8167	227	365.7	10.9	5713	130	298.1	7.5	
epb3-M	9759	282	370.2	11.1	6846	164	305.9	8.0	338468
mark3jac060-A	17447	498	945.7	24.3	13331	319	724.1	16.9	
mark3jac060-M	18970	533	923.2	25.4	15567	358	887.5	20.1	109788
olafu-A	16743	569	363.4	11.3	14012	356	294.3	8.2	
olafu-M	29348	1102	518.2	16.1	25137	696	399.6	11.2	63492
stomach-A	47689	1343	424.9	12.5	36800	706	371.3	8.0	
stomach-M	57755	1588	447.1	12.9	44232	966	391.1	8.7	853440
xenon1-A	29644	769	663.5	18.1	26710	593	542.4	15.3	
xenon1-M	40270	1066	744.2	20.8	33597	754	614.0	16.5	194380
	RC				R/C				
Zhao1-A	13734	399	619.0	17.0	10811	238	517.4	13.3	
Zhao1-M	14551	420	612.5	16.4	11756	280	568.7	15.4	135348
big-A	5453	197	410.0	12.8	3215	68	327.9	8.8	
big-M	7610	223	422.6	12.2	5447	140	347.1	9.1	52804
cage11-A	79052	3482	2265.8	56.0	58272	1452	1164.6	31.7	
cage11-M	61304	1627	2203.7	48.6	42512	1057	1470.5	38.8	151840
cage12-A	248253	8193	2959.2	60.6	185191	4217	1617.2	40.9	
cage12-M	180128	4976	2879.7	58.6	122513	3246	1991.8	48.8	510824
epb2-A	10610	515	766.9	29.9	5527	166	430.2	13.7	
epb2-M	11694	437	492.7	17.1	7411	181	353.9	11.5	100912
epb3-A	17911	985	525.7	19.5	7250	155	312.0	7.5	
epb3-M	18512	938	518.4	19.3	7057	172	295.6	7.9	333320
mark3jac060-A	19503	643	1094.3	29.6	12676	285	953.9	24.6	
mark3jac060-M	21096	525	996.0	24.7	16563	408	891.5	20.1	109780
olafu-A	23870	1041	532.2	18.8	13912	370	292.6	8.2	
olafu-M	36427	1071	696.5	21.8	24735	641	399.1	11.3	57372
stomach-A	77080	3239	552.2	18.7	37219	717	370.3	8.0	
stomach-M	89479	3418	574.5	18.8	47965	1028	391.2	8.9	853440
xenon1-A	36044	1049	808.0	24.5	26669	594	545.3	14.8	
xenon1-M	46970	1183	904.5	25.9	33241	729	604.0	15.1	178388

Table 5: Communication patterns for 32- and 64-way CRC simultaneous and individual partitionings for AINV-matrices.

	Simultaneous partitioning				Individual partitioning				
	CRC				C/R/C				
	Volume		Message		Volume		Message		Reorder
Matrix	tot	max	tot	max	tot	max	tot	max	Volume
	$K = 32$								
Zhao1-A	11191	515	261.2	13.4	8131	340	217.6	11.4	
Zhao1-Z	9132	476	278.6	14.8	7877	1134	293.9	23.4	
Zhao1-W	1711	108	211.4	11.2	76	11	17.8	2.1	198860
big-A	2443	113	157.5	8.4	2071	91	150.9	7.5	
big-Z	1486	85	150.7	8.4	1217	63	146.7	7.3	
big-W	1496	80	149.1	8.2	1218	60	147.8	7.2	76250
cage11-A	49562	2381	508.1	23.9	42424	2068	444.1	21.3	
cage11-Z	21612	1200	545.5	25.9	16277	1119	354.2	19.6	
cage11-W	20323	1050	537.9	25.1	16127	808	336.9	16.9	220128
epb2-A	7028	393	470.1	27.1	4162	243	212.7	15.0	
epb2-Z	2637	174	174.2	9.8	1395	131	105.8	8.6	
epb2-W	2454	162	168.3	9.1	923	75	116.7	9.6	147538
	$K = 64$								
Zhao1-A	15258	365	635.0	17.9	11421	237	529.0	13.1	
Zhao1-Z	12811	332	698.4	20.7	10808	1630	676.9	38.1	
Zhao1-W	2494	84	464.9	13.7	170	13	43.5	2.6	198614
big-A	3730	91	343.9	10.4	3210	69	326.3	8.6	
big-Z	2262	71	323.2	9.4	1861	50	309.8	8.8	
big-W	2305	68	319.4	9.4	1859	49	309.7	7.8	77630
cage11-A	65430	1828	1276.5	32.5	58177	1463	1173.2	30.9	
cage11-Z	28640	925	1367.2	40.9	22023	956	799.8	26.9	
cage11-W	27397	766	1351.0	36.4	21676	572	743.6	21.6	227770
epb2-A	10313	328	1050.8	39.2	5984	172	439.5	18.6	
epb2-Z	3967	152	372.7	11.8	2058	121	233.9	9.3	
epb2-W	3786	143	349.3	10.3	1431	71	209.4	12.8	150522

Table 6: Communication patterns for 32- and 64-way RCR simultaneous and individual partitionings for AINV-matrices.

	Simultaneous partitioning				Individual partitioning				
	RCR				R/C/R				
	Volume		Message		Volume		Message		Reorder
Matrix	tot	max	tot	max	tot	max	tot	max	Volume
	$K = 32$								
Zhao1-A	11146	507	252.8	13.1	7801	347	218.1	11.7	
Zhao1-Z	12153	620	257.1	13.3	9078	441	200.7	11.0	
Zhao1-W	2932	201	199.0	10.2	257	47	60.1	5.0	196614
big-A	2476	123	161.9	8.3	2083	92	152.9	7.7	
big-Z	1561	92	149.7	7.8	1275	63	147.7	7.6	
big-W	1568	72	153.3	7.9	1273	64	148.5	7.5	68492
cage11-A	48379	2432	630.5	27.6	42539	1993	446.6	21.6	
cage11-Z	36660	1890	612.5	26.9	30658	1838	462.7	23.6	
cage11-W	36204	1643	497.9	23.8	30173	1568	454.3	22.6	210112
epb2-A	7288	398	526.5	28.1	3944	221	207.7	11.4	
epb2-Z	3764	248	163.4	8.8	1414	99	83.2	6.6	
epb2-W	3227	211	169.6	8.7	683	36	79.2	4.8	146564
	$K = 64$								
Zhao1-A	14977	372	652.0	17.9	10811	238	517.4	13.3	
Zhao1-Z	17376	447	660.0	18.4	13580	341	510.8	14.2	
Zhao1-W	4380	140	459.6	13.2	377	38	107.5	6.7	199538.0
big-A	3633	99	342.5	9.8	3215	68	327.9	8.8	
big-Z	2314	68	308.8	8.8	1947	49	314.6	8.2	
big-W	2325	58	314.9	8.9	1948	51	316.9	8.1	78060
cage11-A	65052	1965	1758.6	46.8	58272	1452	1164.6	31.7	
cage11-Z	48783	1370	1697.7	43.6	40428	1279	1125.5	33.4	
cage11-W	48453	1180	1273.3	34.7	39731	1193	1096.5	30.3	219988
epb2-A	10490	364	1160.0	41.5	5527	166	430.2	13.7	
epb2-Z	5672	207	347.8	10.4	2478	98	204.3	7.2	
epb2-W	4935	169	359.1	10.4	1174	37	176.2	5.3	148912

Table 7: Communication patterns for 32-way CC and RR composite and individual hypergraph partitionings for SPAI-matrices.

	Individual partitioning				Simultaneous partitioning				
	Volume		Message		Volume		Message		Percent
Matrix	tot	max	tot	max	tot	max	tot	max	Gain
	CC								
Zhao1-A	9228	385	205.0	11.0	9256	392	217.6	11.2	7
Zhao1-M	10847	489	207.8	11.0	9331	420	217.4	11.2	
big-A	2691	128	177.1	9.4	2519	136	161.8	8.4	19
big-M	5358	256	188.3	10.0	4016	189	169.7	8.9	
cage11-A	49634	2138	554.0	23.8	48096	2147	545.5	24.6	19
cage11-M	55307	2482	663.1	27.8	36739	1875	534.7	24.4	
cage12-A	164456	7199	805.3	29.9	161042	7224	736.4	28.6	18
cage12-M	172098	7725	885.2	30.9	114446	6103	718.6	29.8	
epb2-A	6410	317	382.4	19.1	6259	313	341.4	18.1	15
epb2-M	8620	410	220.7	11.3	6587	368	215.4	12.4	
epb3-A	8169	555	199.3	11.8	7851	519	190.5	11.1	25
epb3-M	14290	850	202.5	11.8	8992	577	193.8	11.4	
mark3jac060-A	11155	471	310.8	16.6	13146	518	318.1	15.1	31
mark3jac060-M	29590	1241	303.2	16.8	15158	644	290.8	14.6	
olafu-A	12455	670	174.3	10.2	10275	593	130.7	7.2	28
olafu-M	24120	1307	240.2	13.8	15905	930	162.2	9.8	
stomach-A	34714	1596	221.8	11.7	40865	2618	213.5	11.4	16
stomach-M	72997	3376	233.7	12.1	49987	3099	216.2	11.5	
xenon1-A	21809	933	265.8	14.1	19571	829	252.8	12.4	22
xenon1-M	35264	1513	306.7	15.8	25116	1144	269.6	13.8	
	RR								
Zhao1-A	8829	372	201.1	10.0	8697	381	224.1	11.9	6
Zhao1-M	8787	383	203.4	10.1	7871	356	224.6	12.2	
big-A	2682	126	178.4	9.7	2484	127	165.4	8.4	20
big-M	3641	170	191.1	10.1	2543	131	169.1	8.8	
cage11-A	50521	2179	571.8	25.1	46893	2621	506.9	23.6	19
cage11-M	45057	1935	684.1	28.0	30468	1390	531.1	24.2	
cage12-A	166189	7185	790.0	29.8	156384	8709	696.9	30.1	18
cage12-M	142088	6160	875.0	30.9	96906	4172	704.6	28.6	
epb2-A	6186	371	365.1	20.2	5817	470	297.9	17.4	17
epb2-M	8009	409	223.9	12.2	5939	341	209.6	11.4	
epb3-A	8967	678	220.0	12.8	8376	593	187.7	11.2	21
epb3-M	12786	836	224.5	13.2	8788	592	193.9	11.2	
mark3jac060-A	12851	526	483.6	24.3	11793	501	344.6	16.4	36
mark3jac060-M	26559	1158	507.2	26.1	13569	566	323.0	15.8	
olafu-A	12762	719	172.5	10.6	9846	598	132.2	8.2	29
olafu-M	24535	1239	237.9	13.9	16657	888	160.4	8.7	
stomach-A	36049	1718	217.2	11.2	40160	2652	215.1	11.2	13
stomach-M	62622	2857	229.5	11.9	45415	3101	218.8	11.6	
xenon1-A	20993	904	249.5	12.4	19683	907	249.3	12.8	17
xenon1-M	33878	1427	288.9	14.8	25768	1108	265.9	13.8	

Table 8: Communication patterns for 64-way CC and RR composite and individual hypergraph partitionings for SPAI-matrices.

	Individual partitioning				Simultaneous partitioning				
	Volume		Message		Volume		Message		Percent
Matrix	tot	max	tot	max	tot	max	tot	max	Gain
	CC								
Zhao1-A	12982	271	543.7	13.8	12892	283	558.2	14.6	8
Zhao1-M	15406	359	550.6	13.8	13116	311	560.6	14.5	
big-A	4089	101	378.0	11.3	3849	109	332.8	9.3	18
big-M	8082	204	412.9	12.5	6114	154	358.9	10.3	
cage11-A	67198	1610	1511.1	36.9	64928	1641	1456.4	36.9	18
cage11-M	72710	1803	1886.3	43.7	49361	1362	1410.2	37.2	
cage12-A	213360	5079	2249.8	49.0	208209	5054	2004.5	44.9	18
cage12-M	218227	5396	2659.7	54.9	146726	4357	1904.1	48.6	
epb2-A	9820	268	907.0	27.4	9357	263	782.1	27.9	16
epb2-M	12913	335	491.7	15.8	9766	314	457.6	15.0	
epb3-A	11764	483	508.9	17.0	11293	414	425.0	13.1	26
epb3-M	21010	761	521.2	18.1	13063	443	427.1	13.8	
mark3jac060-A	15676	367	900.8	24.3	18090	380	911.1	22.8	34
mark3jac060-M	42608	954	1009.5	28.2	20183	506	894.6	22.9	
olafu-A	19802	562	402.3	13.0	16733	476	321.9	8.8	26
olafu-M	38318	1102	653.0	20.2	26470	731	431.6	13.4	
stomach-A	50980	1239	488.4	13.6	57777	1916	488.6	13.4	19
stomach-M	107817	2664	536.4	15.2	71060	2310	501.3	13.8	
xenon1-A	30510	671	608.3	17.1	27683	601	567.9	14.9	21
xenon1-M	49568	1111	753.9	20.7	35520	836	624.6	16.8	
	RR								
Zhao1-A	12502	269	541.0	13.9	12131	284	557.6	14.8	8
Zhao1-M	12523	277	550.4	14.2	10896	251	561.4	14.7	
big-A	4068	98	377.6	11.6	3783	104	347.1	10.3	20
big-M	5548	138	414.2	12.6	3904	112	359.5	10.8	
cage11-A	68374	1607	1536.0	37.6	63612	1877	1324.9	35.8	19
cage11-M	60385	1384	1932.8	44.7	41030	974	1376.5	34.2	
cage12-A	216501	4888	2248.7	48.3	203197	6290	1857.8	49.2	18
cage12-M	183548	4205	2670.7	54.1	125836	2979	1887.8	43.1	
epb2-A	9611	367	873.5	30.4	8545	402	640.8	25.4	19
epb2-M	11649	308	484.2	16.1	8590	257	431.9	14.3	
epb3-A	12441	519	468.8	15.2	11607	434	396.2	12.7	22
epb3-M	18036	637	480.6	15.5	12201	420	410.1	12.8	
mark3jac060-A	18107	426	1396.2	40.0	15891	399	1016.0	26.2	36
mark3jac060-M	34929	899	1483.6	46.8	18005	420	960.5	25.2	
olafu-A	20273	585	410.3	12.8	16173	478	325.9	10.2	26
olafu-M	39026	982	642.5	20.0	27911	724	435.4	12.1	
stomach-A	52582	1301	484.9	12.8	58230	2007	486.5	13.1	13
stomach-M	90605	2103	532.9	14.0	66561	2253	502.3	13.2	
xenon1-A	29597	660	590.9	16.4	27615	648	559.5	15.1	17
xenon1-M	47909	1047	729.2	20.1	36437	797	620.5	16.1	

Table 9: Communication patterns for 32-way CC and RR composite and individual hypergraph partitionings for SPAI-matrices with single constraint.

	Individual partitioning				Simultaneous partitioning				
	Volume		Message		Volume		Message		Percent
Matrix	tot	max	tot	max	tot	max	tot	max	Gain
	CC								
Zhao1-A	9228	385	205.0	11.0	9256	392	217.6	11.2	7
Zhao1-M	10847	489	207.8	11.0	9331	420	217.4	11.2	
big-A	2691	128	177.1	9.4	2519	136	161.8	8.4	19
big-M	5358	256	188.3	10.0	4016	189	169.7	8.9	
cage11-A	49634	2138	554.0	23.8	48096	2147	545.5	24.6	19
cage11-M	55307	2482	663.1	27.8	36739	1875	534.7	24.4	
cage12-A	164456	7199	805.3	29.9	161042	7224	736.4	28.6	18
cage12-M	172098	7725	885.2	30.9	114446	6103	718.6	29.8	
epb2-A	6410	317	382.4	19.1	6259	313	341.4	18.1	15
epb2-M	8620	410	220.7	11.3	6587	368	215.4	12.4	
epb3-A	8169	555	199.3	11.8	7851	519	190.5	11.1	25
epb3-M	14290	850	202.5	11.8	8992	577	193.8	11.4	
mark3jac060-A	11155	471	310.8	16.6	13146	518	318.1	15.1	31
mark3jac060-M	29590	1241	303.2	16.8	15158	644	290.8	14.6	
olafu-A	12455	670	174.3	10.2	10275	593	130.7	7.2	28
olafu-M	24120	1307	240.2	13.8	15905	930	162.2	9.8	
stomach-A	34714	1596	221.8	11.7	40865	2618	213.5	11.4	16
stomach-M	72997	3376	233.7	12.1	49987	3099	216.2	11.5	
xenon1-A	21809	933	265.8	14.1	19571	829	252.8	12.4	22
xenon1-M	35264	1513	306.7	15.8	25116	1144	269.6	13.8	
	RR								
Zhao1-A	7796	338	10.9	219.8	8253	393	12.3	228.0	-1
Zhao1-M	7746	349	11.1	223.4	7414	342	12.3	231.3	
big-A	2072	90	7.7	150.9	2151	101	7.8	152.5	16
big-M	2905	126	7.8	158.4	2034	95	7.8	154.3	
cage11-A	42865	1974	22.4	453.6	42933	2328	22.4	475.6	14
cage11-M	37557	1674	25.6	573.2	26514	1328	24.1	510.9	
cage12-A	143478	6286	27.4	625.2	144026	6963	28.9	652.0	13
cage12-M	117062	5078	29.1	719.2	83090	4288	29.3	668.1	
epb2-A	3974	236	11.2	207.8	3921	239	10.8	193.2	24
epb2-M	6681	328	9.9	187.3	4130	196	9.2	167.1	
epb3-A	4827	208	7.3	143.2	5043	246	7.5	144.2	20
epb3-M	7484	358	7.5	143.7	4779	233	7.5	144.8	
mark3jac060-A	9632	382	15.5	329.0	10279	394	13.1	295.6	33
mark3jac060-M	24689	1043	14.7	330.9	12648	542	11.9	275.4	
olafu-A	8462	451	7.3	129.2	8480	628	7.8	126.8	12
olafu-M	18226	980	9.8	171.0	15089	939	8.8	151.2	
stomach-A	26188	986	8.0	177.9	27689	1098	7.0	157.4	20
stomach-M	48162	2189	8.2	184.4	31734	1375	7.2	159.8	
xenon1-A	19117	832	11.9	240.7	18970	862	12.4	245.3	14
xenon1-M	31167	1353	13.6	277.6	24038	1106	13.1	262.3	

Table 10: Communication patterns for 64-way CC and RR composite and individual hypergraph partitionings for SPAI-matrices witn single constraint.

	Individual partitioning				Simultaneous partitioning				
	Volume		Message		Volume		Message		Percent
Matrix	tot	max	tot	max	tot	max	tot	max	Gain
	CC								
Zhao1-A	11460	244	14.2	528.2	12091	281	14.8	542.1	3
Zhao1-M	13763	321	14.2	545.5	12297	294	14.7	551.0	
big-A	3217	70	8.8	327.2	3325	80	8.7	324.9	12
big-M	6751	153	9.5	349.8	5477	139	9.2	337.9	
cage11-A	58329	1526	31.4	1153.4	58828	1915	36.6	1354.8	15
cage11-M	63072	1486	39.0	1526.5	44674	1324	34.8	1318.4	
cage12-A	186001	4750	41.5	1652.1	188732	6878	47.5	1875.0	15
cage12-M	189229	4349	46.8	2020.1	131323	3960	44.8	1781.8	
epb2-A	6019	176	18.7	436.2	6129	267	21.0	407.1	17
epb2-M	10484	274	12.1	388.8	7586	200	10.5	355.1	
epb3-A	5930	134	7.9	298.8	6772	174	8.0	310.6	24
epb3-M	13436	334	8.0	300.4	8036	201	7.8	305.3	
mark3jac060-A	13412	293	17.1	720.6	15276	431	16.8	767.8	36
mark3jac060-M	38534	786	18.4	808.4	18131	411	16.6	768.6	
olafu-A	14021	360	8.3	290.5	13969	539	8.8	296.3	11
olafu-M	29629	754	13.1	461.8	24738	858	12.8	398.8	
stomach-A	36679	697	7.8	371.6	38852	815	8.6	385.0	22
stomach-M	77211	1809	8.2	393.9	49977	1115	8.6	391.2	
xenon1-A	26864	587	15.2	545.6	26721	631	15.4	552.0	15
xenon1-M	44262	974	18.2	674.1	33481	797	17.2	608.4	
	RR								
Zhao1-A	10828	239	13.6	523.6	11400	284	15.4	538.5	0
Zhao1-M	10794	251	13.7	537.5	10227	244	15.1	546.6	
big-A	3212	68	8.7	326.2	3310	80	8.7	327.1	16
big-M	4494	102	9.4	348.5	3161	81	8.8	331.2	
cage11-A	58380	1442	31.4	1170.0	58187	1798	33.1	1214.3	14
cage11-M	50886	1196	39.5	1547.1	35663	1069	35.0	1296.8	
cage12-A	186304	4183	41.8	1651.8	186668	5304	43.1	1678.8	13
cage12-M	152480	3528	47.5	2026.0	107676	3532	43.8	1755.8	
epb2-A	5533	170	13.7	433.7	5540	184	12.7	391.9	24
epb2-M	9888	259	12.0	403.9	6204	147	11.4	356.0	
epb3-A	7254	156	7.6	312.2	7534	179	8.0	311.7	20
epb3-M	11348	267	7.6	313.0	7308	181	8.0	312.4	
mark3jac060-A	12635	290	24.2	938.2	13435	307	18.5	848.9	36
mark3jac060-M	34831	814	25.1	1053.0	16994	440	18.8	826.2	
olafu-A	13924	364	8.2	290.9	14000	591	9.3	294.1	12
olafu-M	30491	769	12.9	456.6	25169	839	11.8	398.9	
stomach-A	37149	715	7.8	370.8	38407	794	8.3	383.8	18
stomach-M	66012	1612	8.3	391.2	46361	1052	8.4	396.3	
xenon1-A	26815	602	15.1	543.7	26687	643	15.8	549.9	14
xenon1-M	44110	992	17.9	666.2	34128	815	17.1	608.3	

Table 11: Speedups for the BiCGStab method with SPAI-matrices.

		CR		RC	
Matrix	K	Time	Speedup	Time	Speedup
Zhao1	1	113	1.0	113	1.0
	8	18	6.2	19	5.9
	16	13	8.7	14	8.1
big	1	51	1.0	51	1.0
	8	9	5.7	10	5.2
	16	7	7.3	8	6.4
cage11	1	210	1.0	210	1.0
	8	38	5.5	50	4.2
	16	26	8.1	38	5.5
cage12	1	744	1.0	745	1.0
	8	127	5.9	170	4.4
	16	79	9.4	121	6.2
epb2	1	103	1.0	103	1.0
	8	16	6.4	17	6.1
	16	12	8.6	12	8.6
epb3	1	286	1.0	300	1.0
	8	39	7.3	42	7.2
	16	23	12.4	26	11.7
mark3jac060	1	122	1.0	122	1.0
	8	21	5.8	21	5.7
	16	14	8.7	16	7.4
olafu	1	255	1.0	255	1.0
	8	38	6.7	43	6.0
	16	24	10.6	30	8.6
stomach	1	1258	1.0	1258	1.0
	8	177	7.1	177	7.1
	16	89	14.1	99	12.7
xenon1	1	368	1.0	368	1.0
	8	55	6.7	61	6.1
	16	33	11.2	40	9.2

Table 12: Speedups for the BiCGStab method with AINV-matrices.

		CRC		RCR	
Matrix	K	Time	Speedup	Time	Speedup
Zhao1	1	133	1.0	134	1.0
	8	20	6.7	21	6.4
	16	15	8.9	15	8.9
big	1	50	1.0	50	1.0
	8	10	5.0	10	5.0
	16	8	6.3	8	6.3
cage11	1	227	1.0	227	1.0
	8	43	5.3	50	4.5
	16	30	7.6	38	6.0
epb2	1	104	1.0	104	1.0
	8	17	6.1	18	5.8
	16	12	8.7	13	8.0

Table 13: Communication patterns for 8- and 16-way CR simultaneous and individual partitionings for SPAI-matrices.

	Simultaneous partitioning				Individual partitioning				
	CR				C/R				
	Volume		Message		Volume		Message		Reorder
Matrix	tot	max	tot	max	tot	max	tot	max	Volume
	$K = 8$								
Zhao1-A	4098	746	32.2	5.7	3028	483	28.2	4.8	
Zhao1-M	3514	694	32.2	5.6	2890	480	30.9	5.2	101836
big-A	1032	201	31.4	5.7	763	131	28.0	5.4	
big-M	989	191	31.9	5.6	780	133	28.6	5.5	52832
cage11-A	24424	4144	54.6	7.0	20989	3435	51.9	7.0	
cage11-M	14663	2439	55.1	7.0	12238	2104	52.4	7.0	138076
cage12-A	87542	14306	56.0	7.0	76240	13057	55.5	7.0	
cage12-M	50962	7839	56.0	7.0	41788	6797	55.5	7.0	434264
epb2-A	2326	429	39.0	6.4	1897	312	39.1	6.2	
epb2-M	2242	438	35.0	6.5	1561	274	31.8	5.5	79936
epb3-A	2354	442	23.9	4.3	1239	195	14.2	2.1	
epb3-M	3003	536	23.9	4.3	1814	339	26.1	5.2	289592
mark3jac060-A	5249	960	35.2	6.3	2648	409	18.4	3.1	
mark3jac060-M	6323	1182	32.2	6.0	4230	657	24.6	4.5	109448
olafu-A	3908	960	25.8	5.0	2933	631	22.3	4.5	
olafu-M	6749	1449	28.0	5.4	4849	1151	19.4	3.8	62616
stomach-A	14614	2815	21.1	4.0	8938	1321	14.0	2.0	
stomach-M	16193	3206	21.4	4.0	9263	1688	14.0	2.0	651820
xenon1-A	10848	2037	36.2	6.5	8765	1360	32.5	5.6	
xenon1-M	14437	2523	37.7	6.7	10471	1758	30.4	5.4	88592
	$K = 16$								
Zhao1-A	6444	586	96.7	9.3	5491	438	91.7	8.5	
Zhao1-M	5478	551	97.0	9.3	4957	421	95.2	9.6	134356
big-A	1581	156	73.2	7.5	1315	113	68.9	6.8	
big-M	1527	150	75.3	7.5	1255	110	70.8	7.2	51912
cage11-A	34835	3314	201.2	14.8	30101	2646	165.9	13.4	
cage11-M	21010	1917	208.9	15.0	17687	1628	176.6	14.0	120612
cage12-A	122878	11925	230.3	15.0	107373	9804	206.7	15.0	
cage12-M	71066	6136	233.1	15.0	58355	4912	213.4	15.0	453388
epb2-A	3357	371	102.8	9.6	2864	273	93.2	8.9	
epb2-M	3335	335	84.7	8.4	2495	238	75.5	7.5	100908
epb3-A	3971	393	66.0	6.5	2302	186	56.9	4.8	
epb3-M	5023	496	66.3	6.5	2877	260	64.7	6.8	338468
mark3jac060-A	9370	786	115.0	11.3	5958	441	81.8	7.1	
mark3jac060-M	10287	964	105.3	11.0	7981	596	86.0	8.0	109784
olafu-A	6489	781	66.2	6.8	5065	523	57.0	6.0	
olafu-M	11258	1285	77.8	7.8	8870	1025	58.9	6.3	53124
stomach-A	24436	2351	67.2	7.0	17415	1264	69.6	5.0	
stomach-M	28014	2652	67.8	7.1	18400	1611	42.8	4.5	841300
xenon1-A	15998	1496	113.2	11.3	13288	1088	100.9	9.6	
xenon1-M	21459	2032	117.8	11.8	16465	1412	103.2	9.9	189004

Table 14: Communication patterns for 8- and 16-way RC simultaneous and individual partitionings for SPAI-matrices.

	Simultaneous partitioning				Individual partitioning				
	RC				R/C				
	Volume		Message		Volume		Message		Reorder
Matrix	tot	max	tot	max	tot	max	tot	max	Volume
	$K = 8$								
Zhao1-A	4146	905	33.2	5.8	2946	480	27.9	4.8	
Zhao1-M	4362	771	33.1	5.6	3101	541	30.4	5.0	72820
big-A	1467	320	33.9	5.8	766	133	28.7	5.4	
big-M	2084	433	34.0	5.8	1402	242	30.2	5.3	36372
cage11-A	30373	6054	55.9	7.0	20894	3316	51.8	7.0	
cage11-M	24447	4302	55.9	7.0	15994	2588	53.9	7.0	125040
cage12-A	107187	17621	56.0	7.0	76755	13076	55.2	7.0	
cage12-M	80949	15047	56.0	7.0	53204	8362	56.0	7.0	308444
epb2-A	3074	646	46.7	6.8	1859	300	38.3	6.3	
epb2-M	3789	814	43.8	6.5	1989	365	33.2	5.8	59428
epb3-A	6347	1907	39.8	6.8	1858	296	25.9	4.4	
epb3-M	6766	1937	39.5	6.8	1937	362	26.6	5.2	315808
mark3jac060-A	5568	981	40.7	6.7	3430	540	26.1	4.6	
mark3jac060-M	6417	1213	38.6	6.8	4481	701	28.4	5.2	109668
olafu-A	5605	1088	30.5	5.8	2932	613	22.6	4.5	
olafu-M	9075	1947	33.5	6.2	4814	1038	21.1	4.6	39624
stomach-A	21139	4354	27.4	5.3	9108	1372	14.2	2.0	
stomach-M	24538	5221	27.6	5.7	10387	1732	14.3	2.1	829212
xenon1-A	12654	2322	39.6	6.9	8720	1387	31.2	5.5	
xenon1-M	16486	2974	40.8	7.0	10432	1754	28.4	5.2	193772
	$K = 16$								
Zhao1-A	6728	734	95.5	9.3	5196	436	94.7	8.9	
Zhao1-M	7141	716	95.5	9.3	5710	500	86.8	8.8	135444
big-A	2408	280	84.7	8.8	1318	114	69.7	7.1	
big-M	3326	366	85.9	8.4	2266	214	71.4	6.7	52520
cage11-A	43335	4599	227.0	15.0	30147	2692	166.0	13.3	
cage11-M	34339	3151	227.4	15.0	22943	2074	194.8	14.7	146952
cage12-A	147504	12516	239.8	15.0	107450	9356	208.0	15.0	
cage12-M	109472	10885	239.7	15.0	73429	6118	222.8	15.0	491724
epb2-A	4555	560	132.1	12.9	2766	269	99.1	8.8	
epb2-M	5388	633	109.5	10.6	3129	303	81.8	7.8	96272
epb3-A	8490	1244	108.3	11.7	3116	282	60.6	5.8	
epb3-M	8991	1480	108.0	11.8	3029	275	63.3	6.8	281644
mark3jac060-A	9792	906	137.4	13.0	6697	515	98.3	9.2	
mark3jac060-M	11099	1027	126.3	12.2	8759	655	97.9	9.1	109724
olafu-A	9325	1012	86.0	9.1	5062	560	55.5	5.7	
olafu-M	14636	1636	99.8	10.2	8661	979	59.1	6.8	59208
stomach-A	35856	4255	81.5	8.8	17565	1266	69.7	5.1	
stomach-M	41254	4660	82.8	8.9	20613	1699	46.4	4.6	595580
xenon1-A	18800	1797	126.2	12.2	13301	1094	102.8	9.5	
xenon1-M	24477	2286	131.6	12.9	16530	1395	93.7	9.3	154180

Table 15: Communication patterns for 8- and 16-way RCR simultaneous and individual partitionings for AINV-matrices.

	Simultaneous partitioning				Individual partitioning				
	RCR				R/C/R				
	Volume		Message		Volume		Message		Reorder
Matrix	tot	max	tot	max	tot	max	tot	max	Volume
	$K = 8$								
Zhao1-A	4325	771	29.9	5.0	2946	480	27.9	4.8	
Zhao1-Z	4344	749	30.4	5.0	3703	651	29.0	5.0	
Zhao1-W	855	205	27.4	4.8	159	67	17.0	3.5	182620
big-A	968	176	29.1	5.5	766	133	28.7	5.4	
big-Z	613	126	28.4	5.4	490	89	28.1	5.2	
big-W	613	104	28.4	5.3	503	94	27.9	5.4	64062
cage11-A	24318	4149	54.8	7.0	20894	3316	51.8	7.0	
cage11-Z	19393	3448	54.4	7.0	15619	3486	49.8	7.0	
cage11-W	19136	3019	52.6	7.0	15341	2812	50.9	7.0	197472
epb2-A	3056	510	54.2	7.0	1859	300	38.3	6.3	
epb2-Z	1298	293	25.6	4.7	366	92	18.4	4.0	
epb2-W	1144	236	26.2	4.6	156	37	14.5	3.2	127480
	$K = 16$								
Zhao1-A	7466	647	86.2	8.0	5196	436	94.7	8.9	
Zhao1-Z	7412	742	87.5	8.1	5742	536	81.3	8.0	
Zhao1-W	1576	231	75.9	7.5	180	48	31.2	4.5	193166
big-A	1575	150	73.3	7.5	1318	114	69.7	7.1	
big-Z	992	108	69.7	7.2	826	78	67.4	7.0	
big-W	997	85	70.7	7.2	818	75	68.3	6.8	73962
cage11-A	34610	3226	203.8	14.8	30147	2692	166.0	13.3	
cage11-Z	26798	2638	204.2	14.9	22276	2534	165.1	14.1	
cage11-W	26361	2299	176.9	14.2	21958	2080	161.7	14.3	182386
epb2-A	4821	455	188.7	14.9	2766	269	99.1	8.8	
epb2-Z	2395	286	70.0	7.0	693	95	41.6	5.3	
epb2-W	2009	243	72.3	6.8	331	39	31.4	3.5	136834

Table 16: Communication patterns for 8- and 16-way CRC simultaneous and individual partitionings for AINV-matrices.

	Simultaneous partitioning				Individual partitioning				
	RCR				R/C/R				
	Volume		Message		Volume		Message		Reorder
Matrix	tot	max	tot	max	tot	max	tot	max	Volume
	$K = 8$								
Zhao1-A	5395	910	36.0	6.0	3028	483	28.2	4.8	
Zhao1-Z	4319	805	37.5	6.0	3936	1379	44.2	7.0	
Zhao1-W	695	125	33.1	5.7	2	2	0.5	0.3	185612
big-A	1030	177	30.1	5.5	763	131	28.0	5.4	
big-Z	620	128	29.4	5.5	469	88	27.5	5.0	
big-W	632	129	29.8	5.5	466	82	27.5	5.3	76000
cage11-A	25588	4336	53.6	7.0	20989	3435	51.9	7.0	
cage11-Z	11063	1965	54.1	7.0	8264	1780	49.5	7.0	
cage11-W	10412	1740	54.0	7.0	8308	1403	47.0	7.0	184414
epb2-A	2876	537	52.1	7.0	1019	178	26.4	4.8	
epb2-Z	1136	267	30.8	5.6	495	133	25.1	5.0	
epb2-W	943	208	30.1	5.7	360	89	26.4	6.2	128132
	$K = 16$								
Zhao1-A	7864	673	101.5	9.1	5491	438	91.7	8.5	
Zhao1-Z	6374	626	106.5	9.8	5497	1228	119.0	13.5	
Zhao1-W	1097	116	87.2	8.2	21	6	6.0	1.1	194176
big-A	1636	148	74.2	7.8	1315	113	68.9	6.8	
big-Z	995	110	71.9	7.4	783	73	68.5	6.8	
big-W	994	112	70.5	7.5	780	76	67.5	6.8	79102
cage11-A	36008	3282	182.4	14.7	30101	2646	165.9	13.4	
cage11-Z	15505	1495	191.7	14.8	11912	1395	139.8	13.2	
cage11-W	14604	1372	187.8	14.6	11738	1041	135.8	12.1	207118
epb2-A	4694	451	183.7	14.9	2864	273	93.2	8.9	
epb2-Z	1742	198	78.2	7.3	796	127	53.1	6.5	
epb2-W	1577	196	75.7	7.4	615	93	54.5	8.3	136990

Table 17: Average CR partitioning times for the SPAI-matrices in seconds.

K	Matrix	Individual partitioning		Simultaneous partitioning	Ratio
		A	M		
8	Zhao	2.28	2.00	6.76	1.6
	big	0.69	0.80	2.56	1.7
	cage11	7.23	5.19	15.64	1.3
	cage12	34.35	22.78	65.08	1.1
	epb2	1.24	1.73	4.41	1.5
	epb3	3.58	4.54	12.67	1.6
	mark3jac060	1.88	2.21	5.64	1.4
	olafu	4.81	5.11	12.19	1.2
	stomach	22.98	25.36	63.52	1.3
	xenon1	6.49	8.54	17.06	1.1
16	Zhao	2.86	2.57	8.66	1.6
	big	0.92	1.03	3.20	1.6
	cage11	9.23	6.50	19.86	1.3
	cage12	43.88	28.90	83.61	1.1
	epb2	1.63	2.16	5.72	1.5
	epb3	4.81	5.99	16.70	1.5
	mark3jac060	2.40	2.85	7.38	1.4
	olafu	6.32	6.73	15.87	1.2
	stomach	30.66	33.38	84.00	1.3
	xenon1	8.75	9.76	23.33	1.3
32	Zhao	3.70	3.08	10.17	1.5
	big	1.07	1.23	3.77	1.6
	cage11	10.93	7.57	23.66	1.3
	cage12	52.27	33.82	100.56	1.2
	epb2	2.04	2.71	6.91	1.5
	epb3	5.93	7.23	20.54	1.6
	mark3jac060	2.89	3.58	8.85	1.4
	olafu	7.71	8.54	19.35	1.2
	stomach	37.86	40.91	104.06	1.3
	xenon1	10.58	11.72	28.50	1.3
64	Zhao	4.02	3.56	11.58	1.5
	big	1.33	1.52	4.36	1.5
	cage11	12.61	8.74	27.02	1.3
	cage12	58.87	43.86	116.01	1.1
	epb2	2.43	3.18	8.05	1.4
	epb3	6.85	8.64	23.95	1.5
	mark3jac060	3.29	4.17	10.22	1.4
	olafu	9.33	10.28	22.70	1.2
	stomach	46.33	48.32	123.63	1.3
	xenon1	12.45	13.80	33.21	1.3

Table 18: Average RC partitioning times for the SPAI-matrices in seconds.

K	Matrix	Individual partitioning		Simultaneous partitioning	Ratio
		A	M		
8	Zhao	2.24	2.34	5.90	1.3
	big-A	0.71	0.95	2.30	1.4
	cage11-A	7.27	5.99	14.89	1.1
	cage12-A	34.70	26.32	62.06	1.0
	epb2-A	1.25	1.73	4.04	1.4
	epb3-A	3.98	4.59	11.54	1.3
	mark3jac060-A	1.79	2.39	4.90	1.2
	olafu-A	4.85	5.30	11.54	1.1
	stomach-A	22.99	25.21	57.59	1.2
	xenon1-A	6.69	7.34	16.07	1.1
16	Zhao	2.86	3.04	7.51	1.3
	big-A	0.88	1.14	2.93	1.4
	cage11-A	9.08	7.46	19.05	1.2
	cage12-A	43.84	32.48	79.21	1.0
	epb2-A	1.65	2.32	5.23	1.3
	epb3-A	5.24	6.06	15.17	1.3
	mark3jac060-A	2.28	3.04	6.40	1.2
	olafu-A	6.33	7.01	14.93	1.1
	stomach-A	30.55	32.97	75.71	1.2
	xenon1-A	8.64	9.75	21.38	1.2
32	Zhao	3.30	3.62	8.88	1.3
	big-A	1.09	1.46	3.50	1.4
	cage11-A	10.83	8.92	22.76	1.2
	cage12-A	51.80	38.05	95.57	1.1
	epb2-A	1.99	2.81	6.25	1.3
	epb3-A	6.45	7.31	18.59	1.4
	mark3jac060-A	2.74	3.86	7.71	1.2
	olafu-A	7.77	8.68	18.19	1.1
	stomach-A	38.10	40.81	93.57	1.2
	xenon1-A	10.69	11.68	26.44	1.2
64	Zhao	3.77	4.16	10.27	1.3
	big-A	1.34	1.69	4.13	1.4
	cage11-A	12.59	10.20	26.17	1.1
	cage12-A	59.33	43.34	110.08	1.1
	epb2-A	2.39	3.31	7.33	1.3
	epb3-A	7.55	8.69	21.73	1.3
	mark3jac060-A	3.25	4.45	9.03	1.2
	olafu-A	9.32	10.46	21.43	1.1
	stomach-A	44.92	48.42	117.00	1.3
	xenon1-A	12.48	13.80	30.87	1.2

Table 19: Average CRC partitioning times for the AINV-matrices in seconds.

K	Matrix	Individual partitioning			Simultaneous partitioning	Ratio
		A	Z	W		
8	Zhao	2.28	1.71	0.60	6.44	1.4
	big	0.69	0.50	0.47	2.37	1.4
	cage11	7.23	3.52	3.51	16.61	1.2
	epb2	1.24	1.01	0.79	4.41	1.5
16	Zhao	2.86	2.28	0.77	8.09	1.4
	big	0.92	0.61	0.60	2.99	1.4
	cage11	9.23	4.48	4.40	21.22	1.2
	epb2	1.63	1.20	1.01	5.63	1.5
32	Zhao	3.70	2.75	0.90	9.57	1.3
	big	1.07	0.74	0.79	3.64	1.4
	cage11	10.93	5.31	5.35	25.16	1.2
	epb2	2.04	1.51	1.21	6.79	1.4
64	Zhao	4.02	3.28	1.01	11.08	1.3
	big	1.33	0.84	0.89	4.27	1.4
	cage11	12.61	6.29	6.21	29.25	1.2
	epb2	2.43	1.68	1.46	7.98	1.4

Table 20: Average RCR partitioning times for the AINV-matrices in seconds.

K	Matrix	Individual partitioning			Simultaneous partitioning	Ratio
		A	Z	W		
8	Zhao	2.24	2.30	0.59	7.40	1.4
	big	0.71	0.54	0.46	2.308	1.3
	cage11	7.27	4.78	4.68	19.17	1.1
	epb2	1.25	1.01	0.84	4.59	1.5
16	Zhao	2.86	2.85	0.73	9.34	1.5
	big	0.88	0.62	0.62	2.94	1.4
	cage11	9.08	5.80	5.86	24.32	1.2
	epb2	1.65	1.31	1.06	5.91	1.5
32	Zhao	3.30	3.53	0.83	11.19	1.5
	big	1.09	0.75	0.76	3.53	1.4
	cage11	10.83	6.99	6.90	29.16	1.2
	epb2	1.99	1.58	1.30	7.16	1.5
64	Zhao	3.77	4.10	1.01	12.93	1.5
	big	1.34	0.89	0.91	4.17	1.3
	cage11	12.59	8.01	7.89	33.36	1.2
	epb2	2.39	1.82	1.54	8.44	1.5

References

[1] Michele Benzi and Miroslav Tuma. Private communication, 2003.

[2] Greg Burns, Raja Daoud, and James Vaigl. LAM: an open cluster environment for MPI. In John W. Ross, editor, *Proceedings of Supercomputing Symposium '94*, pages 379–386. University of Toronto, 1994.

[3] Ü. V. Çatalyürek and C. Aykanat. PaToH: A multilevel hypergraph partitioning tool, version 3.0. Technical Report BU-CE-9915, Computer Engineering Department, Bilkent University, 1999.

[4] Edmond Chow. A priori sparsity patterns for parallel sparse approximate inverse preconditioners. *SIAM J. Sci. Comput.*, 21(5):1804–1822, 2000.

[5] Edmond Chow. Parallel implementation and practical use of sparse approximate inverse preconditioners with a priori sparsity patterns. *Int. J. High Perf. Comput. Apps.*, 15:56–74, 2001.

[6] Tim Davis. University of Florida Sparse Matrix Collection. http://www.cise.ufl.edu/research/sparse/matrices/.

[7] Marcus Grote. SPAI: SParse Approximate Inverse Preconditioner. http://www.sam.math.ethz.ch/˜grote/spai/.

[8] Thomas Huckle. Approximate sparsity patterns for the inverse of a matrix and preconditioning. Technical Report 342/12/98 A, Technische Universität München, 1998.

[9] George Karypis and Vipin Kumar. Multilevel algorithms for multi-constraint graph partitioning. Technical Report 98-019, University of Minnesota, Department of Computer Science/Army HPC Research Center, May 1998.

[10] T. Sterling, D. Savarese, D. J. Becker, J. E. Dorband, U. A. Ranaweke, and C. V. Packer. BEOWULF: A parallel workstation for scientific computation. In *Proceedings of the 24th International Conference on Parallel Processing*, 1995.

[11] C. Walshaw, M. Cross, and K. McManus. Multiphase mesh partitioning. *Applied Mathematical Modeling*, 25:123–140, 2000.