

# Hypergraph-Partitioning-Based Decomposition for Parallel Sparse-Matrix Vector Multiplication

Ümit V. Çatalyürek and Cevdet Aykanat, *Member, IEEE*

**Abstract**—In this work, we show that the standard graph-partitioning-based decomposition of sparse matrices does not reflect the actual communication volume requirement for parallel matrix-vector multiplication. We propose two computational hypergraph models which avoid this crucial deficiency of the graph model. The proposed models reduce the decomposition problem to the well-known hypergraph partitioning problem. The recently proposed successful multilevel framework is exploited to develop a multilevel hypergraph partitioning tool PaToH for the experimental verification of our proposed hypergraph models. Experimental results on a wide range of realistic sparse test matrices confirm the validity of the proposed hypergraph models. In the decomposition of the test matrices, the hypergraph models using PaToH and hMeTiS result in up to 63 percent less communication volume (30 to 38 percent less on the average) than the graph model using MeTiS, while PaToH is only 1.3–2.3 times slower than MeTiS on the average.

**Index Terms**—Sparse matrices, matrix multiplication, parallel processing, matrix decomposition, computational graph model, graph partitioning, computational hypergraph model, hypergraph partitioning.

## 1 INTRODUCTION

ITERATIVE solvers are widely used for the solution of large, sparse, linear systems of equations on multicomputers. Two basic types of operations are repeatedly performed at each iteration. These are linear operations on dense vectors and sparse-matrix vector product (SpMxV) of the form  $\mathbf{y} = \mathbf{A}\mathbf{x}$ , where  $\mathbf{A}$  is an  $m \times m$  square matrix with the same sparsity structure as the coefficient matrix [3], [5], [8], [35], and  $\mathbf{y}$  and  $\mathbf{x}$  are dense vectors. Our goal is the parallelization of the computations in the iterative solvers through *rowwise* or *columnwise* decomposition of the  $\mathbf{A}$  matrix as

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_1^r \\ \vdots \\ \mathbf{A}_k^r \\ \vdots \\ \mathbf{A}_K^r \end{bmatrix} \quad \text{and} \quad \mathbf{A} = [\mathbf{A}_1^c \cdots \mathbf{A}_k^c \cdots \mathbf{A}_K^c],$$

where processor  $P_k$  owns row stripe  $\mathbf{A}_k^r$  or column stripe  $\mathbf{A}_k^c$ , respectively, for a parallel system with  $K$  processors. In order to avoid the communication of vector components during the linear vector operations, a symmetric partitioning scheme is adopted. That is, all vectors used in the solver are divided conformally with the row partitioning or the column partitioning in rowwise or columnwise decomposition schemes, respectively. In particular, the  $\mathbf{x}$  and  $\mathbf{y}$  vectors are divided as  $[\mathbf{x}_1, \dots, \mathbf{x}_K]^t$  and  $[\mathbf{y}_1, \dots, \mathbf{y}_K]^t$ , respectively. In rowwise decomposition, processor  $P_k$  is responsible for computing  $\mathbf{y}_k = \mathbf{A}_k^r \mathbf{x}$  and the linear operations on the  $k$ th blocks of the vectors. In columnwise decomposition, processor  $P_k$  is responsible for computing  $\mathbf{y}^k = \mathbf{A}_k^c \mathbf{x}_k$

(where  $\mathbf{y} = \sum_{k=1}^K \mathbf{y}^k$ ) and the linear operations on the  $k$ th blocks of the vectors. With these decomposition schemes, the linear vector operations can be easily and efficiently parallelized [3], [35] such that only the inner-product computations introduce global communication overhead of which its volume does not scale up with increasing problem size. In parallel SpMxV, the rowwise and columnwise decomposition schemes require communication before or after the local SpMxV computations, thus they can also be considered as *pre-* and *post-*communication schemes, respectively. Depending on the way in which the rows or columns of  $\mathbf{A}$  are partitioned among the processors, entries in  $\mathbf{x}$  or entries in  $\mathbf{y}^k$  may need to be communicated among the processors. Unfortunately, the communication volume scales up with increasing problem size. Our goal is to find a rowwise or columnwise partition of  $\mathbf{A}$  that minimizes the total volume of communication while maintaining the computational load balance.

The decomposition heuristics [32], [33], [37] proposed for computational load balancing may result in extensive communication volume because they do not consider the minimization of the communication volume during the decomposition. In one-dimensional (1D) decomposition, the worst-case communication requirement is  $K(K-1)$  messages and  $(K-1)m$  words, and it occurs when each submatrix  $\mathbf{A}_k^r$  ( $\mathbf{A}_k^c$ ) has at least one nonzero in each column (row) in rowwise (columnwise) decomposition. The approach based on 2D checkerboard partitioning [15], [30] reduces the worst-case communication to  $2K(\sqrt{K}-1)$  messages and  $2(\sqrt{K}-1)m$  words. In this approach, the worst-case occurs when each row and column of each submatrix has at least one nonzero.

The *computational graph* model is widely used in the representation of computational structures of various scientific applications, including repeated SpMxV computations, to decompose the computational domains for parallelization [5], [6], [20], [21], [27], [28], [31], [36]. In this

• The authors are with the Computer Engineering Department, Bilkent University, 06533 Bilkent, Ankara, Turkey.  
E-mail: {cumit,aykanat}@cs.bilkent.edu.tr.

Manuscript received 14 May 1998.

For information on obtaining reprints of this article, please send e-mail to: [tpds@computer.org](mailto:tpds@computer.org), and reference IEEECS Log Number 106847.

model, the problem of sparse matrix decomposition for minimizing the communication volume while maintaining the load balance is formulated as the well-known  $K$ -way graph partitioning problem. In this work, we show the deficiencies of the graph model for decomposing sparse matrices for parallel SpMxV. The first deficiency is that it can only be used for structurally symmetric square matrices. In order to avoid this deficiency, we propose a generalized graph model in Section 2.3 which enables the decomposition of structurally nonsymmetric square matrices as well as symmetric matrices. The second deficiency is the fact that the graph models (both standard and proposed ones) do not reflect the actual communication requirement as will be described in Section 2.4. These flaws are also mentioned in a concurrent work [16]. In this work, we propose two *computational hypergraph* models which avoid all deficiencies of the graph model. The proposed models enable the representation and, hence, the decomposition of rectangular matrices [34], as well as symmetric and nonsymmetric square matrices. Furthermore, they introduce an exact representation for the communication volume requirement as described in Section 3.2. The proposed hypergraph models reduce the decomposition problem to the well-known  $K$ -way hypergraph partitioning problem widely encountered in circuit partitioning in VLSI layout design. Hence, the proposed models will be amenable to the advances in the circuit partitioning heuristics in VLSI community.

Decomposition is a preprocessing introduced for the sake of efficient parallelization of a given problem. Hence, heuristics used for decomposition should run in low order polynomial time. Recently, multilevel graph partitioning heuristics [4], [13], [21] are proposed leading to fast and successful graph partitioning tools Chaco [14] and MeTiS [22]. We have exploited the multilevel partitioning methods for the experimental verification of the proposed hypergraph models in two approaches. In the first approach, MeTiS graph partitioning tool is used as a black box by transforming hypergraphs to graphs using the randomized clique-net model as presented in Section 4.1. In the second approach, the lack of a multilevel hypergraph partitioning tool at the time that this work was carried out led us to develop a multilevel hypergraph partitioning tool PaToH for a fair experimental comparison of the hypergraph models with the graph models. Another objective in our PaToH implementation was to investigate the performance of multilevel approach in hypergraph partitioning as described in Section 4.2. A recently released multilevel hypergraph partitioning tool hMeTiS [24] is also used in the second approach. Experimental results presented in Section 5 confirm both the validity of the proposed hypergraph models and the appropriateness of the multilevel approach to hypergraph partitioning. The hypergraph models using PaToH and hMeTiS produce 30 percent–38 percent better decompositions than the graph models using MeTiS, while the hypergraph models using PaToH are only 34 percent–130 percent slower than the graph models using the most recent version (Version 3.0) of MeTiS, on the average.

## 2 GRAPH MODELS AND THEIR DEFICIENCIES

### 2.1 Graph Partitioning Problem

An undirected graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  is defined as a set of vertices  $\mathcal{V}$  and a set of edges  $\mathcal{E}$ . Every edge  $e_{ij} \in \mathcal{E}$  connects a pair of distinct vertices  $v_i$  and  $v_j$ . The degree  $d_i$  of a vertex  $v_i$  is equal to the number of edges incident to  $v_i$ . Weights and costs can be assigned to the vertices and edges of the graph, respectively. Let  $w_i$  and  $c_{ij}$  denote the weight of vertex  $v_i \in \mathcal{V}$  and the cost of edge  $e_{ij} \in \mathcal{E}$ , respectively.

$\Pi = \{\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_K\}$  is a  $K$ -way partition of  $\mathcal{G}$  if the following conditions hold: Each part  $\mathcal{P}_k$ ,  $1 \leq k \leq K$ , is a nonempty subset of  $\mathcal{V}$ , parts are pairwise disjoint ( $\mathcal{P}_k \cap \mathcal{P}_\ell = \emptyset$  for all  $1 \leq k < \ell \leq K$ ) and union of  $K$  parts is equal to  $\mathcal{V}$  (i.e.,  $\bigcup_{k=1}^K \mathcal{P}_k = \mathcal{V}$ ). A  $K$ -way partition is also called a *multiway* partition if  $K > 2$  and a *bipartition* if  $K = 2$ . A partition is said to be balanced if each part  $\mathcal{P}_k$  satisfies the *balance criterion*

$$W_k \leq W_{avg}(1 + \varepsilon), \quad \text{for } k = 1, 2, \dots, K. \quad (1)$$

In (1), weight  $W_k$  of a part  $\mathcal{P}_k$  is defined as the sum of the weights of the vertices in that part (i.e.  $W_k = \sum_{v_i \in \mathcal{P}_k} w_i$ ),  $W_{avg} = (\sum_{v_i \in \mathcal{V}} w_i) / K$  denotes the weight of each part under the perfect load balance condition, and  $\varepsilon$  represents the predetermined maximum imbalance ratio allowed.

In a partition  $\Pi$  of  $\mathcal{G}$ , an edge is said to be *cut* if its pair of vertices belong to two different parts and *uncut*, otherwise. The cut and uncut edges are also referred to here as *external* and *internal* edges, respectively. The set of external edges of a partition  $\Pi$  is denoted as  $\mathcal{E}_E$ . The *cutsizes* definition for representing the cost  $\chi(\Pi)$  of a partition  $\Pi$  is

$$\chi(\Pi) = \sum_{e_{ij} \in \mathcal{E}_E} c_{ij}. \quad (2)$$

In (2), each cut edge  $e_{ij}$  contributes its cost  $c_{ij}$  to the cutsize. Hence, the graph partitioning problem can be defined as the task of dividing a graph into two or more parts such that the cutsize is minimized while the balance criterion (1) on part weights is maintained. The graph partitioning problem is known to be NP-hard even for bipartitioning unweighted graphs [11].

### 2.2 Standard Graph Model for Structurally Symmetric Matrices

A structurally symmetric sparse matrix  $\mathbf{A}$  can be represented as an undirected graph  $\mathcal{G}_A = (\mathcal{V}, \mathcal{E})$ , where the sparsity pattern of  $\mathbf{A}$  corresponds to the adjacency matrix representation of graph  $\mathcal{G}_A$ . That is, the vertices of  $\mathcal{G}_A$  correspond to the rows/columns of matrix  $\mathbf{A}$  and there exists an edge  $e_{ij} \in \mathcal{E}$  for  $i \neq j$  if and only if off-diagonal entries  $a_{ij}$  and  $a_{ji}$  of matrix  $\mathbf{A}$  are nonzeros. In rowwise decomposition, each vertex  $v_i \in \mathcal{V}$  corresponds to atomic task  $i$  of computing the inner product of row  $i$  with column vector  $\mathbf{x}$ . In columnwise decomposition, each vertex  $v_i \in \mathcal{V}$  corresponds to atomic task  $i$  of computing the sparse SAXPY/DAXPY operation  $\mathbf{y} = \mathbf{y} + x_i \mathbf{a}_{*i}$ , where  $\mathbf{a}_{*i}$  denotes column  $i$  of matrix  $\mathbf{A}$ . Hence, each nonzero entry in a row and column of  $\mathbf{A}$  incurs a multiply-and-add operation during the local SpMxV computations in the pre- and post-communication schemes, respectively. Thus, computational load  $w_i$  of row/column  $i$  is the number of nonzero entries in

row/column  $i$ . In graph theoretical notation,  $w_i = d_i$  when  $a_{ii} = 0$  and  $w_i = d_i + 1$  when  $a_{ii} \neq 0$ . Note that the number of nonzeros in row  $i$  and column  $i$  are equal in a symmetric matrix.

This graph model displays a bidirectional computational interdependency view for SpMxV. Each edge  $e_{ij} \in \mathcal{E}$  can be considered as incurring the computations  $y_i \leftarrow y_i + a_{ij}x_j$  and  $y_j \leftarrow y_j + a_{ji}x_i$ . Hence, each edge represents the bidirectional interaction between the respective pair of vertices in both inner and outer product computation schemes for SpMxV. If rows (columns)  $i$  and  $j$  are assigned to the same processor in a rowwise (columnwise) decomposition, then edge  $e_{ij}$  does not incur any communication. However, in the precommunication scheme, if rows  $i$  and  $j$  are assigned to different processors then cut edge  $e_{ij}$  necessitates the communication of two floating-point words because of the need of the exchange of updated  $x_i$  and  $x_j$  values between atomic tasks  $i$  and  $j$  just before the local SpMxV computations. In the post-communication scheme, if columns  $i$  and  $j$  are assigned to different processors then cut edge  $e_{ij}$  necessitates the communication of two floating-point words because of the need of the exchange of partial  $y_i$  and  $y_j$  values between atomic tasks  $i$  and  $j$  just after the local SpMxV computations. Hence, by setting  $c_{ij} = 2$  for each edge  $e_{ij} \in \mathcal{E}$ , both rowwise and columnwise decompositions of matrix  $\mathbf{A}$  reduce to the  $K$ -way partitioning of its associated graph  $\mathcal{G}_A$  according to the cutsizes definition given in (2). Thus, minimizing the cutsizes is an effort towards minimizing the total volume of interprocessor communication. Maintaining the balance criterion (1) corresponds to maintaining the computational load balance during local SpMxV computations.

Each vertex  $v_i \in \mathcal{V}$  effectively represents both row  $i$  and column  $i$  in  $\mathcal{G}_A$  although its atomic task definition differs in rowwise and columnwise decompositions. Hence, a partition  $\Pi$  of  $\mathcal{G}_A$  automatically achieves a symmetric partitioning by inducing the same partition on the  $\mathbf{y}$ -vector and  $\mathbf{x}$ -vector components since a vertex  $v_i \in \mathcal{P}_k$  corresponds to assigning row  $i$  (column  $i$ ),  $y_i$ , and  $x_i$  to the same part in rowwise (columnwise) decomposition.

In matrix theoretical view, the symmetric partitioning induced by a partition  $\Pi$  of  $\mathcal{G}_A$  can also be considered as inducing a partial symmetric permutation on the rows and columns of  $\mathbf{A}$ . Here, the partial permutation corresponds to ordering the rows/columns assigned to part  $\mathcal{P}_k$  before the rows/columns assigned to part  $\mathcal{P}_{k+1}$ , for  $k = 1, \dots, K - 1$ , where the rows/columns within a part are ordered arbitrarily. Let  $\mathbf{A}^\Pi$  denote the permuted version of  $\mathbf{A}$  according to a partial symmetric permutation induced by  $\Pi$ . An internal edge  $e_{ij}$  of a part  $\mathcal{P}_k$  corresponds to locating both  $a_{ij}$  and  $a_{ji}$  in diagonal block  $\mathbf{A}_{kk}^\Pi$ . An external edge  $e_{ij}$  of cost 2 between parts  $\mathcal{P}_k$  and  $\mathcal{P}_\ell$  corresponds to locating nonzero entry  $a_{ij}$  of  $\mathbf{A}$  in off-diagonal block  $\mathbf{A}_{k\ell}^\Pi$  and  $a_{ji}$  of  $\mathbf{A}$  in off-diagonal block  $\mathbf{A}_{\ell k}^\Pi$ , or vice versa. Hence, minimizing the cutsizes in the graph model can also be considered as permuting the rows and columns of the matrix to minimize the total number of nonzeros in the off-diagonal blocks.

Fig. 1 illustrates a sample  $10 \times 10$  symmetric sparse matrix  $\mathbf{A}$  and its associated graph  $\mathcal{G}_A$ . The numbers inside the circles indicate the computational weights of the

respective vertices (rows/columns). This figure also illustrates a rowwise decomposition of the symmetric  $\mathbf{A}$  matrix and the corresponding bipartitioning of  $\mathcal{G}_A$  for a two-processor system. As seen in Fig. 1, the cutsizes in the given graph bipartitioning is 8, which is also equal to the total number of nonzero entries in the off-diagonal blocks. The bipartition illustrated in Fig. 1 achieves perfect load balance by assigning 21 nonzero entries to each row stripe. This number can also be obtained by adding the weights of the vertices in each part.

### 2.3 Generalized Graph Model for Structurally Symmetric/Nonsymmetric Square Matrices

The standard graph model is not suitable for the partitioning of nonsymmetric matrices. A recently proposed *bipartite* graph model [17], [26] enables the partitioning of rectangular as well as structurally symmetric/nonsymmetric square matrices. In this model, each row and column is represented by a vertex and the sets of vertices representing the rows and columns form the bipartition, i.e.,  $\mathcal{V} = \mathcal{V}_R \cup \mathcal{V}_C$ . There exists an edge between a row vertex  $i \in \mathcal{V}_R$  and a column vertex  $j \in \mathcal{V}_C$  if and only if the respective entry  $a_{ij}$  of matrix  $\mathbf{A}$  is nonzero. Partitions  $\Pi_R$  and  $\Pi_C$  on  $\mathcal{V}_R$  and  $\mathcal{V}_C$ , respectively, determine the overall partition  $\Pi = \{\mathcal{P}_1, \dots, \mathcal{P}_K\}$ , where  $\mathcal{P}_k = \mathcal{V}_{R_k} \cup \mathcal{V}_{C_k}$  for  $k = 1, \dots, K$ . For rowwise (columnwise) decomposition, vertices in  $\mathcal{V}_R$  ( $\mathcal{V}_C$ ) are weighted with the number of nonzeros in the respective row (column) so that the balance criterion (1) is imposed only on the partitioning of  $\mathcal{V}_R$  ( $\mathcal{V}_C$ ). As in the standard graph model, minimizing the number of cut edges corresponds to minimizing the total number of nonzeros in the off-diagonal blocks. This approach has the flexibility of achieving nonsymmetric partitioning. In the context of parallel SpMxV, the need for symmetric partitioning on square matrices is achieved by enforcing  $\Pi_R \equiv \Pi_C$ . Hendrickson and Kolda [17] propose several bipartite-graph partitioning algorithms that are adopted from the techniques for the standard graph model and one partitioning algorithm that is specific to bipartite graphs.

In this work, we propose a simple yet effective graph model for symmetric partitioning of structurally nonsymmetric square matrices. The proposed model enables the use of the standard graph partitioning tools without any modification. In the proposed model, a nonsymmetric square matrix  $\mathbf{A}$  is represented as an undirected graph  $\mathcal{G}_R = (\mathcal{V}_R, \mathcal{E})$  and  $\mathcal{G}_C = (\mathcal{V}_C, \mathcal{E})$  for the rowwise and columnwise decomposition schemes, respectively. Graphs  $\mathcal{G}_R$  and  $\mathcal{G}_C$  differ only in their vertex weight definitions. The vertex set and the corresponding atomic task definitions are identical to those of the symmetric matrices. That is, weight  $w_i$  of a vertex  $v_i \in \mathcal{V}_R$  ( $v_i \in \mathcal{V}_C$ ) is equal to the total number of nonzeros in row  $i$  (column  $i$ ) in  $\mathcal{G}_R$  ( $\mathcal{G}_C$ ). In the edge set  $\mathcal{E}$ ,  $e_{ij} \in \mathcal{E}$  if and only if off-diagonal entries  $a_{ij} \neq 0$  or  $a_{ji} \neq 0$ . That is, the vertices in the adjacency list of a vertex  $v_i$  denote the union of the column indices of the off-diagonal nonzeros at row  $i$  and the row indices of the off-diagonal nonzeros at column  $i$ . The cost  $c_{ij}$  of an edge  $e_{ij}$  is set to 1 if either  $a_{ij} \neq 0$  or  $a_{ji} \neq 0$ , and it is set to 2 if both  $a_{ij} \neq 0$  and  $a_{ji} \neq 0$ . The proposed scheme is referred to here as a generalized model since it automatically produces the standard graph




Fig. 1. Two-way rowwise decomposition of a sample structurally symmetric matrix  $\mathbf{A}$  and the corresponding bipartitioning of its associated graph  $\mathcal{G}_A$ .






Fig. 2. Two-way rowwise decomposition of a sample structurally nonsymmetric matrix  $\mathbf{A}$  and the corresponding bipartitioning of its associated graph  $\mathcal{G}_R$ .

representation for structurally symmetric matrices by computing the same cost of 2 for every edge.

Fig. 2 illustrates a sample  $10 \times 10$  nonsymmetric sparse matrix  $\mathbf{A}$  and its associated graph  $\mathcal{G}_R$  for rowwise decomposition. The numbers inside the circles indicate the computational weights of the respective vertices (rows). This figure also illustrates a rowwise decomposition of the matrix and the corresponding bipartitioning of its associated graph for a two-processor system. As seen in Fig. 2, the cutsize of the given graph bipartitioning is 7, which is also equal to the total number of nonzero entries in the off-diagonal blocks. Hence, similar to the standard and bipartite graph models, minimizing cutsize in the proposed graph model corresponds to minimizing the total number of nonzeros in the off-diagonal blocks. As seen in Fig. 2, the bipartitioning achieves perfect load balance by assigning 16 nonzero entries to each row stripe. As mentioned earlier, the  $\mathcal{G}_C$  model of a matrix for columnwise decomposition differs from the  $\mathcal{G}_R$  model only in vertex weights. Hence, the graph bipartitioning illustrated in Fig. 2 can also be considered as incurring a slightly imbalanced (15 versus 17 nonzeros) columnwise decomposition of sample matrix  $\mathbf{A}$  (shown by vertical dash line) with identical communication requirement.



### 2.4 Deficiencies of the Graph Models

Consider the symmetric matrix decomposition given in Fig. 1. Assume that parts  $\mathcal{P}_1$  and  $\mathcal{P}_2$  are mapped to processors  $P_1$  and  $P_2$ , respectively. The cutsize of the bipartition shown in this figure is equal to  $2 \times 4 = 8$ , thus estimating the communication volume requirement as eight words. In the pre-communication scheme, off-block-diagonal entries  $a_{4,7}$  and  $a_{5,7}$  assigned to processor  $P_1$  display the same need for the nonlocal  $\mathbf{x}$ -vector component  $x_7$  twice. However, it is clear that processor  $P_2$  will send  $x_7$  only once to processor  $P_1$ . Similarly, processor  $P_1$  will send  $x_4$  only once to processor  $P_2$  because of the off-block-diagonal entries  $a_{7,4}$  and  $a_{8,4}$  assigned to processor  $P_2$ . In the post-communication scheme, the graph model treats the off-block-diagonal nonzeros  $a_{7,4}$  and  $a_{7,5}$  in  $\mathcal{P}_1$  as if processor  $P_1$  will send two multiplication results  $a_{7,4}x_4$  and  $a_{7,5}x_5$  to processor  $P_2$ . However, it is obvious that processor  $P_1$  will compute the partial result for the nonlocal  $\mathbf{y}$ -vector component  $y'_7 = a_{7,4}x_4 + a_{7,5}x_5$  during the local SpMxV phase and send this single value to processor  $P_2$  during the post-communication phase. Similarly, processor  $P_2$  will only compute and send the single value  $y'_4 = a_{4,7}x_7 + a_{4,8}x_8$  to processor  $P_1$ . Hence, the actual communication volume is in fact six words instead of eight in both pre- and post-communication schemes. A similar analysis of the rowwise decomposition of the nonsymmetric matrix given in Fig. 2 reveals the fact that the actual communication requirement

is five words ( $x_4, x_5, x_6, x_7,$  and  $x_8$ ) instead of seven, determined by the cutsize of the given bipartition of  $\mathcal{G}_R$ .

In matrix theoretical view, the nonzero entries in the same column of an off-diagonal block incur the communication of a single  $x$  value in the rowwise decomposition (pre-communication) scheme. Similarly, the nonzero entries in the same row of an off-diagonal block incur the communication of a single  $y$  value in the columnwise decomposition (post-communication) scheme. However, as mentioned earlier, the graph models try to minimize the total number of off-block-diagonal nonzeros without considering the relative spatial locations of such nonzeros. In other words, the graph models treat all off-block-diagonal nonzeros in an identical manner by assuming that each off-block-diagonal nonzero will incur a distinct communication of a single word.

In graph theoretical view, the graph models treat all cut edges of equal cost in an identical manner while computing the cutsize. However,  $r$  cut edges, each of cost 2, stemming from a vertex  $v_{i_1}$  in part  $\mathcal{P}_k$  to  $r$  vertices  $v_{i_2}, v_{i_3}, \dots, v_{i_{r+1}}$  in part  $\mathcal{P}_\ell$  incur only  $r+1$  communications instead of  $2r$  in both pre- and post-communication schemes. In the pre-communication scheme, processor  $P_k$  sends  $x_{i_1}$  to processor  $P_\ell$  while  $P_\ell$  sends  $x_{i_2}, x_{i_3}, \dots, x_{i_{r+1}}$  to  $P_k$ . In the post-communication scheme, processor  $P_\ell$  sends  $y'_{i_2}, y'_{i_3}, \dots, y'_{i_{r+1}}$  to processor  $P_k$  while  $P_k$  sends  $y'_{i_1}$  to  $P_\ell$ . Similarly, the amount of communication required by  $r$  cut edges, each of cost 1, stemming from a vertex  $v_{i_1}$  in part  $\mathcal{P}_k$  to  $r$  vertices  $v_{i_2}, v_{i_3}, \dots, v_{i_{r+1}}$  in part  $\mathcal{P}_\ell$  may vary between 1 and  $r$  words instead of exactly  $r$  words, determined by the cutsize of the given graph partitioning.

### 3 HYPERGRAPH MODELS FOR DECOMPOSITION

#### 3.1 Hypergraph Partitioning Problem

A hypergraph  $\mathcal{H}=(\mathcal{V}, \mathcal{N})$  is defined as a set of vertices  $\mathcal{V}$  and a set of nets (hyperedges)  $\mathcal{N}$  among those vertices. Every net  $n_j \in \mathcal{N}$  is a subset of vertices, i.e.,  $n_j \subseteq \mathcal{V}$ . The vertices in a net  $n_j$  are called its *pins* and denoted as  $\text{pins}[n_j]$ . The size of a net is equal to the number of its pins, i.e.,  $s_j = |\text{pins}[n_j]|$ . The set of nets connected to a vertex  $v_i$  is denoted as  $\text{nets}[v_i]$ . The degree of a vertex is equal to the number of nets it is connected to, i.e.,  $d_i = |\text{nets}[v_i]|$ . Graph is a special instance of hypergraph such that each net has exactly two pins. Similar to graphs, let  $w_i$  and  $c_j$  denote the weight of vertex  $v_i \in \mathcal{V}$  and the cost of net  $n_j \in \mathcal{N}$ , respectively.

Definition of  $K$ -way partition of hypergraphs is identical to that of graphs. In a partition  $\Pi$  of  $\mathcal{H}$ , a net that has at least one pin (vertex) in a part is said to *connect* that part. *Connectivity set*  $\Lambda_j$  of a net  $n_j$  is defined as the set of parts connected by  $n_j$ . *Connectivity*  $\lambda_j = |\Lambda_j|$  of a net  $n_j$  denotes the number of parts connected by  $n_j$ . A net  $n_j$  is said to be *cut* if it connects more than one part (i.e.,  $\lambda_j > 1$ ) and *uncut*, otherwise (i.e.,  $\lambda_j = 1$ ). The cut and uncut nets are also referred to here as *external* and *internal* nets, respectively. The set of external nets of a partition  $\Pi$  is denoted as  $\mathcal{N}_E$ . There are various *cutsizes* definitions for representing the cost  $\chi(\Pi)$  of a partition  $\Pi$ . Two relevant definitions are:

$$(a) \chi(\Pi) = \sum_{n_j \in \mathcal{N}_E} c_j \text{ and } (b) \chi(\Pi) = \sum_{n_j \in \mathcal{N}_E} c_j (\lambda_j - 1). \quad (3)$$

In (3.a), the cutsize is equal to the sum of the costs of the cut nets. In (3.b), each cut net  $n_j$  contributes  $c_j(\lambda_j - 1)$  to the cutsize. Hence, the hypergraph partitioning problem [29] can be defined as the task of dividing a hypergraph into two or more parts such that the cutsize is minimized while a given balance criterion (1) among the part weights is maintained. Here, part weight definition is identical to that of the graph model. The hypergraph partitioning problem is known to be NP-hard [29].

#### 3.2 Two Hypergraph Models for Decomposition

We propose two computational hypergraph models for the decomposition of sparse matrices. These models are referred to here as the *column-net* and *row-net* models proposed for the rowwise decomposition (pre-communication) and columnwise decomposition (post-communication) schemes, respectively.

In the column-net model, matrix  $\mathbf{A}$  is represented as a hypergraph  $\mathcal{H}_R=(\mathcal{V}_R, \mathcal{N}_C)$  for rowwise decomposition. Vertex and net sets  $\mathcal{V}_R$  and  $\mathcal{N}_C$  correspond to the rows and columns of matrix  $\mathbf{A}$ , respectively. There exist one vertex  $v_i$  and one net  $n_j$  for each row  $i$  and column  $j$ , respectively. Net  $n_j \subseteq \mathcal{V}_R$  contains the vertices corresponding to the rows that have a nonzero entry in column  $j$ . That is,  $v_i \in n_j$  if and only if  $a_{ij} \neq 0$ . Each vertex  $v_i \in \mathcal{V}_R$  corresponds to atomic task  $i$  of computing the inner product of row  $i$  with column vector  $\mathbf{x}$ . Hence, computational weight  $w_i$  of a vertex  $v_i \in \mathcal{V}_R$  is equal to the total number of nonzeros in row  $i$ . The nets of  $\mathcal{H}_R$  represent the *dependency* relations of the atomic tasks on the  $\mathbf{x}$ -vector components in rowwise decomposition. Each net  $n_j$  can be considered as incurring the computation  $y_i \leftarrow y_i + a_{ij}x_j$  for each vertex (row)  $v_i \in n_j$ . Hence, each net  $n_j$  denotes the set of atomic tasks (vertices) that need  $x_j$ . Note that each pin  $v_i$  of a net  $n_j$  corresponds to a unique nonzero  $a_{ij}$ , thus enabling the representation and decomposition of structurally nonsymmetric matrices, as well as symmetric matrices, without any extra effort. Fig. 3a illustrates the dependency relation view of the column-net model. As seen in this figure, net  $n_j = \{v_h, v_i, v_k\}$  represents the dependency of atomic tasks  $h, i, k$  to  $x_j$  because of the computations  $y_h \leftarrow y_h + a_{hj}x_j$ ,  $y_i \leftarrow y_i + a_{ij}x_j$ , and  $y_k \leftarrow y_k + a_{kj}x_j$ . Fig. 4b illustrates the column-net representation of the sample  $16 \times 16$  nonsymmetric matrix given in Fig. 4a. In Fig. 4b, the pins of net  $n_7 = \{v_7, v_{10}, v_{13}\}$  represent nonzeros  $a_{7,7}$ ,  $a_{10,7}$ , and  $a_{13,7}$ . Net  $n_7$  also represents the dependency of atomic tasks 7, 10, and 13 to  $x_7$  because of the computations  $y_7 \leftarrow y_7 + a_{7,7}x_7$ ,  $y_{10} \leftarrow y_{10} + a_{10,7}x_7$ , and  $y_{13} \leftarrow y_{13} + a_{13,7}x_7$ .

The row-net model can be considered as the dual of the column-net model. In this model, matrix  $\mathbf{A}$  is represented as a hypergraph  $\mathcal{H}_C=(\mathcal{V}_C, \mathcal{N}_R)$  for columnwise decomposition. Vertex and net sets  $\mathcal{V}_C$  and  $\mathcal{N}_R$  correspond to the columns and rows of matrix  $\mathbf{A}$ , respectively. There exists one vertex  $v_i$  and one net  $n_j$  for each column  $i$  and row  $j$ , respectively. Net  $n_j \subseteq \mathcal{V}_C$  contains the vertices corresponding to the columns that have a nonzero entry in row  $j$ . That is,  $v_i \in n_j$  if and only if  $a_{ji} \neq 0$ . Each vertex  $v_i \in \mathcal{V}_C$  corresponds to atomic task  $i$  of computing the sparse




Fig. 3. Dependency relation views of (a) column-net and (b) row-net models.

SAXPY/DAXPY operation  $\mathbf{y} = \mathbf{y} + x_i \mathbf{a}_{*i}$ . Hence, computational weight  $w_i$  of a vertex  $v_i \in \mathcal{V}_C$  is equal to the total number of nonzeros in column  $i$ . The nets of  $\mathcal{H}_C$  represent the *dependency* relations of the computations of the  $\mathbf{y}$ -vector components on the atomic tasks represented by the vertices of  $\mathcal{H}_C$  in columnwise decomposition. Each net  $n_j$  can be considered as incurring the computation  $y_j \leftarrow y_j + a_{ji} x_i$  for each vertex (column)  $v_i \in n_j$ . Hence, each net  $n_j$  denotes the set of atomic task results needed to accumulate  $y_j$ . Note that each pin  $v_i$  of a net  $n_j$  corresponds to a unique nonzero  $a_{ji}$ , thus enabling the representation and decomposition of structurally nonsymmetric matrices as well as symmetric matrices without any extra effort. Fig. 3b illustrates the dependency relation view of the row-net model. As seen in this figure, net  $n_j = \{v_h, v_i, v_k\}$  represents the dependency of accumulating  $y_j = y_j^h + y_j^i + y_j^k$  on the partial  $y_j$  results  $y_j^h = a_{jh} x_h$ ,  $y_j^i = a_{ji} x_i$ , and  $y_j^k = a_{jk} x_k$ . Note that the row-net and column-net models become identical in structurally symmetric matrices.

By assigning unit costs to the nets (i.e.,  $c_j = 1$  for each net  $n_j$ ), the proposed column-net and row-net models reduce the decomposition problem to the  $K$ -way hypergraph partitioning problem according to the cutsizes definition given in (3.b) for the pre- and post-communication schemes, respectively. Consistency of the proposed hypergraph models for accurate representation of communication volume requirement while maintaining the symmetric partitioning restriction depends on the condition that “ $v_j \in n_j$  for each net  $n_j$ .” We first assume that this condition holds in the discussion throughout the following four paragraphs and then discuss the appropriateness of the assumption in the last paragraph of this section.

The validity of the proposed hypergraph models is discussed only for the column-net model. A dual discussion holds for the row-net model. Consider a partition  $\Pi$  of  $\mathcal{H}_R$  in the column-net model for rowwise decomposition of a matrix  $\mathbf{A}$ . Without loss of generality, we assume that part  $\mathcal{P}_k$  is assigned to processor  $P_k$  for  $k = 1, 2, \dots, K$ . As  $\Pi$  is defined as a partition on the vertex set of  $\mathcal{H}_R$ , it induces a complete part (hence, processor) assignment for the rows of matrix  $\mathbf{A}$  and, hence, for the components of the  $\mathbf{y}$  vector. That is, a vertex  $v_i$  assigned to part  $\mathcal{P}_k$  in  $\Pi$  corresponds to assigning row  $i$  and  $y_i$  to part  $\mathcal{P}_k$ . However, partition  $\Pi$  does not induce any part assignment for the nets of  $\mathcal{H}_R$ . Here, we consider partition  $\Pi$  as inducing an assignment for the internal nets of  $\mathcal{H}_R$ , hence, for the respective  $\mathbf{x}$ -vector

components. Consider an internal net  $n_j$  of part  $\mathcal{P}_k$  (i.e.,  $\Lambda_j = \{\mathcal{P}_k\}$ ) which corresponds to column  $j$  of  $\mathbf{A}$ . As all pins of net  $n_j$  lie in  $\mathcal{P}_k$ , all rows (including row  $j$  by the consistency condition) which need  $x_j$  for inner-product computations are already assigned to processor  $P_k$ . Hence, internal net  $n_j$  of  $\mathcal{P}_k$ , which does not contribute to the cutsizes (3.b) of partition  $\Pi$ , does not necessitate any communication if  $x_j$  is assigned to processor  $P_k$ . The assignment of  $x_j$  to processor  $P_k$  can be considered as permuting column  $j$  to part  $\mathcal{P}_k$ , thus respecting the symmetric partitioning of  $\mathbf{A}$  since row  $j$  is already assigned to  $\mathcal{P}_k$ . In the 4-way decomposition given in Fig. 4b, internal nets  $n_1, n_{10}, n_{13}$  of part  $\mathcal{P}_1$  induce the assignment of  $x_1, x_{10}, x_{13}$  and columns 1, 10, 13 to part  $\mathcal{P}_1$ . Note that part  $\mathcal{P}_1$  already contains rows 1, 10, 13, thus respecting the symmetric partitioning of  $\mathbf{A}$ .

Consider an external net  $n_j$  with connectivity set  $\Lambda_j$ , where  $\lambda_j = |\Lambda_j|$  and  $\lambda_j > 1$ . As all pins of net  $n_j$  lie in the parts in its connectivity set  $\Lambda_j$ , all rows (including row  $j$  by the consistency condition) which need  $x_j$  for inner-product computations are assigned to the parts (processors) in  $\Lambda_j$ . Hence, contribution  $\lambda_j - 1$  of external net  $n_j$  to the cutsizes according to (3.b) accurately models the amount of communication volume to incur during the parallel SpMxV computations because of  $x_j$  if  $x_j$  is assigned to any processor in  $\Lambda_j$ . Let  $map[j] \in \Lambda_j$  denote the part and, hence, processor assignment for  $x_j$  corresponding to cut net  $n_j$ . In the column-net model together with the pre-communication scheme, cut net  $n_j$  indicates that processor  $map[j]$  should send its local  $x_j$  to those processors in connectivity set  $\Lambda_j$  of net  $n_j$  except itself (i.e., to processors in the set  $\Lambda_j - \{map[j]\}$ ). Hence, processor  $map[j]$  should send its local  $x_j$  to  $|\Lambda_j| - 1 = \lambda_j - 1$  distinct processors. As the consistency condition “ $v_j \in n_j$ ” ensures that row  $j$  is already assigned to a part in  $\Lambda_j$ , symmetric partitioning of  $\mathbf{A}$  can easily be maintained by assigning  $x_j$ , hence permuting column  $j$  to the part which contains row  $j$ . In the 4-way decomposition shown in Fig. 4b, external net  $n_5$  (with  $\Lambda_5 = \{\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3\}$ ) incurs the assignment of  $x_5$  (hence, permuting column 5) to part  $\mathcal{P}_1$  since row 5 ( $v_5 \in n_5$ ) is already assigned to part  $\mathcal{P}_1$ . The contribution  $\lambda_5 - 1 = 2$  of net  $n_5$  to the cutsizes accurately models the communication volume to incur due to  $x_5$  because processor  $P_1$  should send  $x_5$  to both processors  $P_2$  and  $P_3$  only once since  $\Lambda_5 - \{map[5]\} = \Lambda_5 - \{\mathcal{P}_1\} = \{\mathcal{P}_2, \mathcal{P}_3\}$ .




Fig. 4. (a) A  $16 \times 16$  structurally nonsymmetric matrix  $\mathbf{A}$ . (b) Column-net representation  $\mathcal{H}_R$  of matrix  $\mathbf{A}$  and 4-way partitioning  $\Pi$  of  $\mathcal{H}_R$ . (c) 4-way rowwise decomposition of matrix  $\mathbf{A}^{\Pi}$  obtained by permuting  $\mathbf{A}$  according to the symmetric partitioning induced by  $\Pi$ .

In essence, in the column-net model, any partition  $\Pi$  of  $\mathcal{H}_R$  with  $v_i \in \mathcal{P}_k$  can be safely decoded as assigning row  $i$ ,  $y_i$  and  $x_i$  to processor  $P_k$  for rowwise decomposition. Similarly, in the row-net model, any partition  $\Pi$  of  $\mathcal{H}_C$  with  $v_i \in \mathcal{P}_k$  can be safely decoded as assigning column  $i$ ,  $x_i$ , and  $y_i$  to processor  $P_k$  for columnwise decomposition. Thus, in the column-net and row-net models, minimizing the cutsize according to (3.b) corresponds to minimizing the actual volume of interprocessor communication during the pre- and post-communication phases, respectively. Maintaining the balance criterion (1) corresponds to maintaining the computational load balance during the local SpMxV computations. Fig. 4c displays a permutation of the sample matrix given in Fig. 4a according to the symmetric partitioning induced by the 4-way decomposition shown in Fig. 4b. As seen in Fig. 4c, the actual communication volume for the given rowwise decomposition is six words since processor  $P_1$  should send  $x_5$  to both  $P_2$  and  $P_3$ ,  $P_2$  should send  $x_{11}$  to  $P_4$ ,  $P_3$  should send  $x_7$  to  $P_1$ , and  $P_4$  should send  $x_{12}$  to both  $P_2$  and  $P_3$ . As seen in Fig. 4b, external nets  $n_5$ ,  $n_7$ ,  $n_{11}$ , and  $n_{12}$  contribute 2, 1, 1, and 2 to the cutsize since  $\lambda_5 = 3$ ,  $\lambda_7 = 2$ ,  $\lambda_{11} = 2$ , and  $\lambda_{12} = 3$ , respectively. Hence, the cutsize of the 4-way decomposition given in Fig. 4b is 6, thus leading to the accurate modeling of the communication requirement. Note that the graph model will estimate the total communication volume as 13 words for the 4-way decomposition given in Fig. 4c since the total number of nonzeros in the off-diagonal blocks is 13. As seen in Fig. 4c, each processor is assigned 12 nonzeros thus achieving perfect computational load balance.

In matrix theoretical view, let  $\mathbf{A}^{\Pi}$  denote a permuted version of matrix  $\mathbf{A}$  according to the symmetric partitioning induced by a partition  $\Pi$  of  $\mathcal{H}_R$  in the column-net model. Each cut-net  $n_j$  with connectivity set  $\Lambda_j$  and  $map[j] = \mathcal{P}_\ell$  corresponds to column  $j$  of  $\mathbf{A}$  containing nonzeros in  $\lambda_j$  distinct blocks ( $\mathbf{A}_{k\ell}^{\Pi}$ , for  $\mathcal{P}_k \in \Lambda_j$ ) of matrix  $\mathbf{A}^{\Pi}$ . Since connectivity set  $\Lambda_j$  of net  $n_j$  is guaranteed to contain part  $map[j]$ , column  $j$  contains nonzeros in  $\lambda_j - 1$  distinct off-diagonal blocks of  $\mathbf{A}^{\Pi}$ . Note that multiple nonzeros of column  $j$  in a particular off-diagonal block contributes only

one to connectivity  $\lambda_j$  of net  $n_j$  by definition of  $\lambda_j$ . So, the cutsize of a partition  $\Pi$  of  $\mathcal{H}_R$  is equal to the number of nonzero column segments in the off-diagonal blocks of matrix  $\mathbf{A}^{\Pi}$ . For example, external net  $n_5$  with  $\Lambda_5 = \{\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3\}$  and  $map[5] = \mathcal{P}_1$  in Fig. 4b indicates that column 5 has nonzeros in two off-diagonal blocks  $\mathbf{A}_{3,1}^{\Pi}$  and  $\mathbf{A}_{3,1}^{\Pi}$ , as seen in Fig. 4c. As also seen in Fig. 4c, the number of nonzero column segments in the off-diagonal blocks of matrix  $\mathbf{A}^{\Pi}$  is 6, which is equal to the cutsize of partition  $\Pi$  shown in Fig. 4b. Hence, the column-net model tries to achieve a symmetric permutation which minimizes the total number of nonzero column segments in the off-diagonal blocks for the pre-communication scheme. Similarly, the row-net model tries to achieve a symmetric permutation which minimizes the total number of nonzero row segments in the off-diagonal blocks for the post-communication scheme.

Nonzero diagonal entries automatically satisfy the condition “ $v_j \in n_j$  for each net  $n_j$ ,” thus enabling both accurate representation of communication requirement and symmetric partitioning of  $\mathbf{A}$ . A nonzero diagonal entry  $a_{jj}$  already implies that net  $n_j$  contains vertex  $v_j$  as its pin. If, however, some diagonal entries of the given matrix are zeros, then the consistency of the proposed column-net model is easily maintained by simply adding rows, which do not contain diagonal entries, to the pin lists of the respective column nets. That is, if  $a_{jj} = 0$  then vertex  $v_j$  (row  $j$ ) is added to the pin list  $pins[n_j]$  of net  $n_j$  and net  $n_j$  is added to the net list  $nets[v_j]$  of vertex  $v_j$ . These pin additions do not affect the computational weight assignments of the vertices. That is, weight  $w_j$  of vertex  $v_j$  in  $\mathcal{H}_R$  becomes equal to either  $d_j$  or  $d_j - 1$  depending on whether  $a_{jj} \neq 0$  or  $a_{jj} = 0$ , respectively. The consistency of the row-net model is preserved in a dual manner.

## 4 DECOMPOSITION HEURISTICS

Kernighan-Lin (KL)-based heuristics are widely used for graph/hypergraph partitioning because of their short runtimes and good quality results. The KL algorithm is an iterative improvement heuristic originally proposed for

graph bipartitioning [25]. The KL algorithm, starting from an initial bipartition, performs a number of passes until it finds a locally minimum partition. Each pass consists of a sequence of vertex swaps. The same swap strategy was applied to the hypergraph bipartitioning problem by Schweikert-Kernighan [38]. Fiduccia-Mattheyses (FM) [10] introduced a faster implementation of the KL algorithm for hypergraph partitioning. They proposed vertex move concept instead of vertex swap. This modification, as well as proper data structures, e.g., bucket lists, reduced the time complexity of a single pass of the KL algorithm to linear in the size of the graph and the hypergraph. Here, *size* refers to the number of edges and pins in a graph and hypergraph, respectively.

The performance of the FM algorithm deteriorates for large and very sparse graphs/hypergraphs. Here, sparsity of graphs and hypergraphs refer to their average vertex degrees. Furthermore, the solution quality of FM is not *stable (predictable)*, i.e., average FM solution is significantly worse than the best FM solution, which is a common weakness of the move-based iterative improvement approaches. Random multistart approach is used in VLSI layout design to alleviate this problem by running the FM algorithm many times starting from random initial partitions to return the best solution found [1]. However, this approach is not viable in parallel computing since decomposition is a preprocessing overhead introduced to increase the efficiency of the underlying parallel algorithm/program. Most users will rely on one run of the decomposition heuristic, so the quality of the decomposition tool depends equally on the worst and average decompositions than on just the best decomposition.

These considerations have motivated the *two-phase* application of the move-based algorithms in hypergraph partitioning [12]. In this approach, a clustering is performed on the original hypergraph  $\mathcal{H}_0$  to induce a coarser hypergraph  $\mathcal{H}_1$ . Clustering corresponds to coalescing highly interacting vertices to supernodes as a preprocessing to FM. Then, FM is run on  $\mathcal{H}_1$  to find a bipartition  $\Pi_1$ , and this bipartition is projected back to a bipartition  $\Pi_0$  of  $\mathcal{H}_0$ . Finally, FM is rerun on  $\mathcal{H}_0$  using  $\Pi_0$  as an initial solution. Recently, the two-phase approach has been extended to *multilevel* approaches [4], [13], [21], leading to successful graph partitioning tools Chaco [14] and MeTiS [22]. These multilevel heuristics consist of three phases: *coarsening*, *initial partitioning*, and *uncoarsening*. In the first phase, a multilevel clustering is applied starting from the original graph by adopting various matching heuristics until the number of vertices in the coarsened graph reduces below a predetermined threshold value. In the second phase, the coarsest graph is partitioned using various heuristics, including FM. In the third phase, the partition found in the second phase is successively projected back towards the original graph by refining the projected partitions on the intermediate level uncoarser graphs using various heuristics, including FM.

In this work, we exploit the multilevel partitioning schemes for the experimental verification of the proposed hypergraph models in two approaches. In the first approach, multilevel graph partitioning tool MeTiS is used

as a black box by transforming hypergraphs to graphs using the randomized clique-net model proposed in [2]. In the second approach, we have implemented a multilevel hypergraph partitioning tool PaToH, and tested both PaToH and multilevel hypergraph partitioning tool hMeTiS [23], [24] which was released very recently.

#### 4.1 Randomized Clique-Net Model for Graph Representation of Hypergraphs

In the clique-net transformation model, the vertex set of the target graph is equal to the vertex set of the given hypergraph with the same vertex weights. Each net of the given hypergraph is represented by a clique of vertices corresponding to its pins. That is, each net induces an edge between every pair of its pins. The multiple edges connecting each pair of vertices of the graph are contracted into a single edge of which cost is equal to the sum of the costs of the edges it represents. In the *standard* clique-net model [29], a uniform cost of  $1/(s_i - 1)$  is assigned to every clique edge of net  $n_i$  with size  $s_i$ . Various other edge weighting functions are also proposed in the literature [1]. If an edge is in the cut set of a graph partitioning then all nets represented by this edge are in the cut set of hypergraph partitioning, and vice versa. Ideally, no matter how vertices of a net are partitioned, the contribution of a cut net to the cutsize should always be one in a bipartition. However, the deficiency of the clique-net model is that it is impossible to achieve such a *perfect* clique-net model [18]. Furthermore, the transformation may result in very large graphs since the number of clique edges induced by the nets increase quadratically with their sizes.

Recently, a randomized clique-net model implementation was proposed [2] which yields very promising results when used together with graph partitioning tool MeTiS. In this model, all nets of size larger than  $T$  are removed during the transformation. Furthermore, for each net  $n_i$  of size  $s_i$ ,  $F \times s_i$  random pairs of its pins (vertices) are selected and an edge with cost one is added to the graph for each selected pair of vertices. The multiple edges between each pair of vertices of the resulting graph are contracted into a single edge as mentioned earlier. In this scheme, the nets with size smaller than  $2F + 1$  (small nets) induce a larger number of edges than the standard clique-net model, whereas the nets with size larger than  $2F + 1$  (large nets) induce a smaller number of edges than the standard clique-net model. Considering the fact that MeTiS accepts integer edge costs for the input graph, this scheme has two nice features.<sup>1</sup> First, it simulates the uniform edge-weighting scheme of the standard clique-net model for small nets in a random manner since each clique edge (if induced) of a net  $n_i$  with size  $s_i < 2F + 1$  will be assigned an integer cost close to  $2F/(s_i - 1)$  on the average. Second, it prevents the quadratic increase in the number of clique edges induced by large nets in the standard model since the number of clique edges induced by a net in this scheme is linear in the size of the net. In our implementation, we use the parameters  $T = 50$  and  $F = 5$  in accordance with the recommendations given in [2].

1. Private communication with C.J. Alpert.



## 4.2 PaToH: A Multilevel Hypergraph Partitioning Tool

In this work, we exploit the successful multilevel methodology [4], [13], [21] proposed and implemented for graph partitioning [14], [22] to develop a new multilevel hypergraph partitioning tool, called PaToH (PaToH: **P**artitioning **T**ools for **H**ypergraphs).

The data structures used to store hypergraphs in PaToH mainly consist of the following arrays. The *NETLIST* array stores the net lists of the vertices. The *PINLIST* array stores the pin lists of the nets. The size of both arrays is equal to the total number of pins in the hypergraph. Two auxiliary index arrays *VTXS* and *NETS* of sizes  $|\mathcal{V}|+1$  and  $|\mathcal{N}|+1$  hold the starting indices of the net lists and pin lists of the vertices and nets in the *NETLIST* and *PINLIST* arrays, respectively. In sparse matrix storage terminology, this scheme corresponds to storing the given matrix both in *Compressed Sparse Row (CSR)* and *Compressed Sparse Column (CSC)* formats [27] without storing the numerical data. In the column-net model proposed for rowwise decomposition, the *VTXS* and *NETLIST* arrays correspond to the CSR storage scheme, and the *NETS* and *PINLIST* arrays correspond to the CSC storage scheme. This correspondence is dual in the row-net model proposed for columnwise decomposition.

The  $K$ -way graph/hypergraph partitioning problem is usually solved by recursive bisection. In this scheme, first a 2-way partition of  $\mathcal{G}/\mathcal{H}$  is obtained and, then, this bipartition is further partitioned in a recursive manner. After  $\lg_2 K$  phases, graph  $\mathcal{G}/\mathcal{H}$  is partitioned into  $K$  parts. PaToH achieves  $K$ -way hypergraph partitioning by recursive bisection for any  $K$  value (i.e.,  $K$  is not restricted to be a power of 2).

The connectivity cutsizes metric given in (3.b) needs special attention in  $K$ -way hypergraph partitioning by recursive bisection. Note that the cutsizes metrics given in (3.a) and (3.b) become equivalent in hypergraph bisection. Consider a bipartition  $\mathcal{V}_A$  and  $\mathcal{V}_B$  of  $\mathcal{V}$  obtained after a bisection step. It is clear that  $\mathcal{V}_A$  and  $\mathcal{V}_B$  and the internal nets of parts  $\mathcal{A}$  and  $\mathcal{B}$  will become the vertex and net sets of  $\mathcal{H}_A$  and  $\mathcal{H}_B$ , respectively, for the following recursive bisection steps. Note that each cut net of this bipartition already contributes 1 to the total cutsize of the final  $K$ -way partition to be obtained by further recursive bisections. However, the further recursive bisections of  $\mathcal{V}_A$  and  $\mathcal{V}_B$  may increase the connectivity of these cut nets. In parallel SpMxV view, while each cut net already incurs the communication of a single word, these nets may induce additional communication because of the following recursive bisection steps. Hence, after every hypergraph bisection step, each cut net  $n_i$  is split into two pin-wise disjoint nets  $n'_i = pins[n_i] \cap \mathcal{V}_A$  and  $n''_i = pins[n_i] \cap \mathcal{V}_B$  and, then, these two nets are added to the net lists of  $\mathcal{H}_A$  and  $\mathcal{H}_B$  if  $|n'_i| > 1$  and  $|n''_i| > 1$ , respectively. Note that the single-pin nets are discarded during the split operation since such nets cannot contribute to the cutsize in the following recursive bisection steps. Thus, the total cutsize according to (3.b) will become equal to the sum of the number of cut split nets at every bisection step by using the above cut-net split method. Fig. 5 illustrates two cut nets  $n_i$  and  $n_k$  in a bipartition and their splits into




Fig. 5. Cut-net splitting during recursive bisection.

nets  $n'_i, n''_i$  and  $n'_k, n''_k$ , respectively. Note that net  $n''_k$  becomes a single-pin net and it is discarded.

Similar to multilevel graph and hypergraph partitioning tools Chaco [14], MeTiS [22], and hMeTiS [24], the multilevel hypergraph bisection algorithm used in PaToH consists of three phases: coarsening, initial partitioning, and uncoarsening. The following sections briefly summarize our multilevel bisection algorithm. Although PaToH works on weighted nets, we will assume unit cost nets both for the sake of simplicity of presentation and for the fact that all nets are assigned unit cost in the hypergraph representation of sparse matrices.

### 4.2.1 Coarsening Phase

In this phase, the given hypergraph  $\mathcal{H} = \mathcal{H}_0 = (\mathcal{V}_0, \mathcal{N}_0)$  is coarsened into a sequence of smaller hypergraphs  $\mathcal{H}_1 = (\mathcal{V}_1, \mathcal{N}_1), \mathcal{H}_2 = (\mathcal{V}_2, \mathcal{N}_2), \dots, \mathcal{H}_m = (\mathcal{V}_m, \mathcal{N}_m)$  satisfying  $|\mathcal{V}_0| > |\mathcal{V}_1| > |\mathcal{V}_2| > \dots > |\mathcal{V}_m|$ . This coarsening is achieved by coalescing disjoint subsets of vertices of hypergraph  $\mathcal{H}_i$  into *multinodes* such that each multinode in  $\mathcal{H}_i$  forms a single vertex of  $\mathcal{H}_{i+1}$ . The weight of each vertex of  $\mathcal{H}_{i+1}$  becomes equal to the sum of its constituent vertices of the respective multinode in  $\mathcal{H}_i$ . The net set of each vertex of  $\mathcal{H}_{i+1}$  becomes equal to the union of the net sets of the constituent vertices of the respective multinode in  $\mathcal{H}_i$ . Here, multiple pins of a net  $n \in \mathcal{N}_i$  in a multinode cluster of  $\mathcal{H}_i$  are contracted to a single pin of the respective net  $n' \in \mathcal{N}_{i+1}$  of  $\mathcal{H}_{i+1}$ . Furthermore, the single-pin nets obtained during this contraction are discarded. Note that such single-pin nets correspond to the internal nets of the clustering performed on  $\mathcal{H}_i$ . The coarsening phase terminates when the number of vertices in the coarsened hypergraph reduces below 100 (i.e.,  $|\mathcal{V}_m| \leq 100$ ).

Clustering approaches can be classified as *agglomerative* and *hierarchical*. In the agglomerative clustering, new clusters are formed one at a time, whereas in the hierarchical clustering, several new clusters may be formed

simultaneously. In PaToH, we have implemented both randomized matching-based hierarchical clustering and randomized hierarchic-agglomerative clustering. The former and latter approaches will be abbreviated as matching-based clustering and agglomerative clustering, respectively.

The matching-based clustering works as follows: Vertices of  $\mathcal{H}_i$  are visited in a random order. If a vertex  $u \in \mathcal{V}_i$  has not been matched yet, one of its unmatched adjacent vertices is selected according to a criterion. If such a vertex  $v$  exists, we merge the matched pair  $u$  and  $v$  into a cluster. If there is no unmatched adjacent vertex of  $u$ , then vertex  $u$  remains unmatched, i.e.,  $u$  remains as a singleton cluster. Here, two vertices  $u$  and  $v$  are said to be adjacent if they share at least one net, i.e.,  $\text{nets}[u] \cap \text{nets}[v] \neq \emptyset$ . The selection criterion used in PaToH for matching chooses a vertex  $v$  with the highest connectivity value  $N_{uv}$ . Here, connectivity  $N_{uv} = |\text{nets}[u] \cap \text{nets}[v]|$  refers to the number of shared nets between  $u$  and  $v$ . This matching-based scheme is referred to here as *Heavy Connectivity Matching (HCM)*.

The matching-based clustering allows the clustering of only pairs of vertices in a level. In order to enable the clustering of more than two vertices at each level, we have implemented a randomized agglomerative clustering approach. In this scheme, each vertex  $u$  is assumed to constitute a singleton cluster  $C_u = \{u\}$  at the beginning of each coarsening level. Then, vertices are visited in a random order. If a vertex  $u$  has already been clustered (i.e.,  $|C_u| > 1$ ), it is not considered for being the source of a new clustering. However, an unclustered vertex  $u$  can choose to join a multinode cluster as well as a singleton cluster. That is, all adjacent vertices of an unclustered vertex  $u$  are considered for selection according to a criterion. The selection of a vertex  $v$  adjacent to  $u$  corresponds to including vertex  $u$  to cluster  $C_v$  to grow a new multinode cluster  $C_u = C_v = C_v \cup \{u\}$ . Note that no singleton cluster remains at the end of this process as far as there exists no isolated vertex. The selection criterion used in PaToH for agglomerative clustering chooses a singleton or multinode cluster  $C_v$  with the highest  $N_{u,C_v}/W_{u,C_v}$  value, where  $N_{u,C_v} = |\text{nets}[u] \cap \bigcup_{x \in C_v} \text{nets}[x]|$  and  $W_{u,C_v}$  is the weight of the multinode cluster candidate  $\{u\} \cup C_v$ . The division of  $N_{u,C_v}$  by  $W_{u,C_v}$  is an effort to avoiding the polarization towards very large clusters. This agglomerative clustering scheme is referred to here as *Heavy Connectivity Clustering (HCC)*.

The objective in both HCM and HCC is to find highly connected vertex clusters. Connectivity values  $N_{uv}$  and  $N_{u,C_v}$  used for selection serve this objective. Note that  $N_{uv}$  ( $N_{u,C_v}$ ) also denotes the lower bound in the amount of decrease in the number of pins because of the pin contractions to be performed when  $u$  joins  $v$  ( $C_v$ ). Recall that there might be additional decrease in the number of pins because of single-pin nets that may occur after clustering. Hence, the connectivity metric is also an effort towards minimizing the complexity of the following coarsening levels, partitioning phase, and refinement phase since the size of a hypergraph is equal to the number of its pins.

In rowwise matrix decomposition context (i.e., column-net model), the connectivity metric corresponds to the number of common column indices between two rows or

row groups. Hence, both HCM and HCC try to combine rows or row groups with similar sparsity patterns. This in turn corresponds to combining rows or row groups which need similar sets of x-vector components in the pre-communication scheme. A dual discussion holds for the row-net model. Fig. 6 illustrates a single level coarsening of an  $8 \times 8$  sample matrix  $\mathbf{A}_0$  in the column-net model using HCM and HCC. The original decimal ordering of the rows is assumed to be the random vertex visit order. As seen in Fig. 6, HCM matches row pairs  $\{1, 3\}$ ,  $\{2, 6\}$ , and  $\{4, 5\}$  with the connectivity values of 3, 2, and 2, respectively. Note that the total number of nonzeros of  $\mathbf{A}_0$  reduces from 28 to 21 in  $\mathbf{A}_1^{HCM}$  after clustering. This difference is equal to the sum  $3+2+2=7$  of the connectivity values of the matched row-vertex pairs since pin contractions do not lead to any single-pin nets. As seen in Fig. 6, HCC constructs three clusters  $\{1, 2, 3\}$ ,  $\{4, 5\}$ , and  $\{6, 7, 8\}$  through the clustering sequence of  $\{1, 3\}$ ,  $\{1, 2, 3\}$ ,  $\{4, 5\}$ ,  $\{6, 7\}$ , and  $\{6, 7, 8\}$  with the connectivity values of 3, 4, 2, 3, and 2, respectively. Note that pin contractions lead to three single-pin nets  $n_2$ ,  $n_3$ , and  $n_7$ , thus columns 2, 3, and 7 are removed. As also seen in Fig. 6, although rows 7 and 8 remain unmatched in HCM, every row is involved in at least one clustering in HCC.

Both HCM and HCC necessitate scanning the pin lists of all nets in the net list of the source vertex to find its adjacent vertices for matching and clustering. In the column-net (row-net) model, the total cost of these scan operations can be as expensive as the total number of multiply and add operations which lead to nonzero entries in the computation of  $\mathbf{A}\mathbf{A}^T$  ( $\mathbf{A}^T\mathbf{A}$ ). In HCM, the key point to efficient implementation is to move the matched vertices encountered during the scan of the pin list of a net to the end of its pin list through a simple swap operation. This scheme avoids the revisits of the matched vertices during the following matching operations at that level. Although this scheme requires an additional index array to maintain the temporary tail indices of the pin lists, it achieves substantial decrease in the run-time of the coarsening phase. Unfortunately, this simple yet effective scheme cannot be fully used in HCC. Since a singleton vertex can select a multinode cluster, the revisits of the clustered vertices are partially avoided by maintaining only a single vertex to represent the multinode cluster in the pin-list of each net connected to the cluster, through simple swap operations. Through the use of these efficient implementation schemes the total cost of the scan operations in the column-net (row-net) model can be as low as the total number of nonzeros in  $\mathbf{A}\mathbf{A}^T$  ( $\mathbf{A}^T\mathbf{A}$ ). In order to maintain this cost within reasonable limits, all nets of size greater than  $4s_{avg}$  are not considered in a bipartitioning step, where  $s_{avg}$  denotes the average net size of the hypergraph to be partitioned in that step. Note that such nets can be reconsidered during the further levels of recursion because of net splitting.

The cluster growing operation in HCC requires disjoint-set operations for maintaining the representatives of the clusters, where the union operations are restricted to the union of a singleton source cluster with a singleton or a multinode target cluster. This restriction is exploited by always choosing the representative of the target cluster as the representative of the new cluster. Hence, it is sufficient

$$\mathbf{A}_0 = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \\ 6 \\ 7 \\ 8 \end{matrix} & \begin{bmatrix} x & & x & & x & & & x \\ & x & x & & & x & & x \\ x & x & x & & x & x & & \\ & & & x & x & & & \\ x & & & x & x & & & \\ x & & & x & & x & x & x \\ & & & x & & x & x & \\ x & & & & & & & x \end{bmatrix} \end{matrix}$$
  

$$\mathbf{A}_1^{HCM} = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 \end{matrix} \\ \begin{matrix} 1, 3 \\ 2, 6 \\ 4, 5 \\ 7 \\ 8 \end{matrix} & \begin{bmatrix} x & x & x & & x & x & & x \\ x & x & x & x & & x & x & x \\ x & & & x & x & & & \\ & & & x & & x & x & \\ x & & & & & & & x \end{bmatrix} \end{matrix}$$
  

$$\mathbf{A}_1^{HCC} = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 \end{matrix} \\ \begin{matrix} 1, 2, 3 \\ 4, 5 \\ 6, 7, 8 \end{matrix} & \begin{bmatrix} x & x & x & & x & x & & x \\ x & & & x & x & & & \\ x & & & x & & x & x & x \end{bmatrix} = \begin{matrix} & \begin{matrix} 1 & 4 & 5 & 6 & 8 \end{matrix} \\ \begin{matrix} 1, 2, 3 \\ 4, 5 \\ 6, 7, 8 \end{matrix} & \begin{bmatrix} x & & x & x & x \\ x & x & x & & \\ x & x & & x & x \end{bmatrix} \end{matrix}$$

Fig. 6. Matching-based clustering  $\mathbf{A}_1^{HCM}$  and agglomerative clustering  $\mathbf{A}_1^{HCC}$  of the rows of matrix  $\mathbf{A}_0$ .

to update the representative pointer of only the singleton source cluster joining to a multinode target cluster. Therefore, each disjoint-set operation required in this scheme is performed in  $O(1)$  time.

### 4.2.2 Initial Partitioning Phase

The goal in this phase is to find a bipartition on the coarsest hypergraph  $\mathcal{H}_m$ . In PaToH, we use the *Greedy Hypergraph Growing (GHG)* algorithm for bisecting  $\mathcal{H}_m$ . This algorithm can be considered as an extension of the GGGP algorithm used in MeTiS to hypergraphs. In GHG, we grow a cluster around a randomly selected vertex. During the course of the algorithm, the selected and unselected vertices induce a bipartition on  $\mathcal{H}_m$ . The unselected vertices connected to the growing cluster are inserted into a priority queue according to their FM gains. Here, the gain of an unselected vertex corresponds to the decrease in the cutsize of the current bipartition if the vertex moves to the growing cluster. Then, a vertex with the highest gain is selected from the priority queue. After a vertex moves to the growing cluster, the gains of its unselected adjacent vertices that are currently in the priority queue are updated and those not in the priority queue are inserted. This cluster growing operation continues until a predetermined bipartition balance criterion is reached. As also mentioned in MeTiS, the quality of this algorithm is sensitive to the choice of the initial random vertex. Since the coarsest hypergraph  $\mathcal{H}_m$  is small, we run GHG four times, starting from different random vertices and select the best bipartition for refinement during the uncoarsening phase.

### 4.2.3 Uncoarsening Phase

At each level  $i$  (for  $i = m, m-1, \dots, 1$ ), bipartition  $\Pi_i$  found on  $\mathcal{H}_i$  is projected back to a bipartition  $\Pi_{i-1}$  on  $\mathcal{H}_{i-1}$ . The constituent vertices of each multinode in  $\mathcal{H}_{i-1}$  are assigned to the part of the respective vertex in  $\mathcal{H}_i$ . Obviously,  $\Pi_{i-1}$  of  $\mathcal{H}_{i-1}$  has the same cutsize with  $\Pi_i$  of  $\mathcal{H}_i$ . Then, we refine this bipartition by running a *Boundary FM (BFM)* hypergraph bipartitioning algorithm on  $\mathcal{H}_{i-1}$  starting from initial

bipartition  $\Pi_{i-1}$ . BFM moves only the boundary vertices from the overloaded part to the under-loaded part, where a vertex is said to be a boundary vertex if it is connected to at least one cut net.

BFM requires maintaining the *pin-connectivity* of each net for both initial gain computations and gain updates. The pin-connectivity  $\sigma_k[n] = |n \cap \mathcal{P}_k|$  of a net  $n$  to a part  $\mathcal{P}_k$  denotes the number of pins of net  $n$  that lie in part  $\mathcal{P}_k$ , for  $k = 1, 2$ . In order to avoid the scan of the pin lists of all nets, we adopt an efficient scheme to initialize the  $\sigma$  values for the first BFM pass in a level. It is clear that initial bipartition  $\Pi_{i-1}$  of  $\mathcal{H}_{i-1}$  has the same cut-net set with  $\Pi_i$  of  $\mathcal{H}_i$ . Hence, we scan only the pin lists of the cut nets of  $\Pi_{i-1}$  to initialize their  $\sigma$  values. For each other net  $n$ ,  $\sigma_1[n]$  and  $\sigma_2[n]$  values are easily initialized as  $\sigma_1[n] = s_n$  and  $\sigma_2[n] = 0$  if net  $n$  is internal to part  $\mathcal{P}_1$ , and  $\sigma_1[n] = 0$  and  $\sigma_2[n] = s_n$ , otherwise. After initializing the gain value of each vertex  $v$  as  $g[v] = -d_v$ , we exploit  $\sigma$  values as follows. We rescan the pin list of each external net  $n$  and update the gain value of each vertex  $v \in pins[n]$  as  $g[v] = g[v] + 2$  or  $g[v] = g[v] + 1$  depending on whether net  $n$  is *critical* to the part containing  $v$  or not, respectively. An external net  $n$  is said to be critical to a part  $k$  if  $\sigma_k[n] = 1$  so that moving the single vertex of net  $n$  that lies in that part to the other part removes net  $n$  from the cut. Note that two-pin cut nets are critical to both parts. The vertices visited while scanning the pin-lists of the external nets are identified as boundary vertices and only these vertices are inserted into the priority queue according to their computed gains.

In each pass of the BFM algorithm, a sequence of unmoved vertices with the highest gains are selected to move to the other part. As in the original FM algorithm, a vertex move necessitates gain updates of its adjacent vertices. However, in the BFM algorithm, some of the adjacent vertices of the moved vertex may not be in the priority queue because they may not be boundary vertices before the move. Hence, such vertices which become boundary vertices after the move are inserted into the priority queue according to their updated gain values. The

TABLE 1  
Properties of Test Matrices

matrix name	description	number of rows/cols	number of nonzeros									
			total	avg. per row/col	per column				per row			
					min	max	std	cov	min	max	std	cov
Structurally Symmetric Matrices												
SHERMAN3	[9] 3D finite difference grid	5005	20033	4.00	1	7	2.66	0.67	1	7	2.66	0.67
KEN-11	[7] linear programming	14694	82454	5.61	2	243	14.54	2.59	2	243	14.54	2.59
NL	[19] linear programming	7039	105089	14.93	1	361	28.48	1.91	1	361	28.48	1.91
KEN-13	[7] linear programming	28632	161804	5.65	2	339	16.84	2.98	2	339	16.84	2.98
CQ9	[19] linear programming	9278	221590	23.88	1	702	54.46	2.28	1	702	54.46	2.28
CO9	[19] linear programming	10789	249205	23.10	1	707	52.17	2.26	1	707	52.17	2.26
CRE-D	[7] linear programming	8926	372266	41.71	1	845	76.46	1.83	1	845	76.46	1.83
CRE-B	[7] linear programming	9648	398806	41.34	1	904	74.69	1.81	1	904	74.69	1.81
FINAN512	[38] stochastic programming	74752	615774	8.24	3	1449	20.00	2.43	3	1449	20.00	2.43
Structurally Nonsymmetric Matrices												
GEMAT11	[9] optimal power flow	4929	38101	7.73	1	28	2.96	0.38	1	29	3.38	0.44
LHR07	[38] light hydrocarbon recovery	7337	163716	22.31	1	64	26.19	1.17	2	37	16.00	0.72
ONETONE2	[38] nonlinear analog circuit	36057	254595	7.06	2	34	5.13	0.73	2	66	6.67	0.94
LHR14	[38] light hydrocarbon recovery	14270	321988	22.56	1	64	26.26	1.16	2	37	15.98	0.71
ONETONE1	[38] nonlinear analog circuit	36057	368055	10.21	2	82	14.32	1.40	2	162	17.85	1.75
LHR17	[38] light hydrocarbon recovery	17576	399500	22.73	1	64	26.32	1.16	2	37	15.96	0.70
LHR34	[38] light hydrocarbon recovery	35152	799064	22.73	1	64	26.32	1.16	2	37	15.96	0.70
BCSSTK32	[9] 3D stiffness matrix	44609	1029655	23.08	1	141	10.10	0.44	1	192	10.45	0.45
BCSSTK30	[9] 3D stiffness matrix	28924	1036208	35.83	1	159	21.99	0.61	1	104	15.27	0.43

refinement process within a pass terminates when no *feasible* move remains or the sequence of last  $\max\{50, 0.001|V_i|\}$  moves does not yield a decrease in the total cutsize. A move is said to be feasible if it does not disturb the load balance criterion (1) with  $K=2$ . At the end of a BFM pass, we have a sequence of tentative vertex moves and their respective gains. We then construct from this sequence the maximum prefix subsequence of moves with the maximum prefix sum which incurs the maximum decrease in the cutsize. The permanent realization of the moves in this maximum prefix subsequence is efficiently achieved by rolling back the remaining moves at the end of the overall sequence. The initial gain computations for the following pass in a level is achieved through this rollback. The overall refinement process in a level terminates if the maximum prefix sum of a pass is not positive. In the current implementation of PaToH, at most two BFM passes are allowed at each level of the uncoarsening phase.

## 5 EXPERIMENTAL RESULTS

We have tested the validity of the proposed hypergraph models by running MeTiS on the graphs obtained by randomized clique-net transformation and running PaToH and hMeTiS directly on the hypergraphs for the decompositions of various realistic sparse test matrices arising in different application domains. These decomposition results are compared with the decompositions obtained by running MeTiS using the standard and proposed graph models for the symmetric and nonsymmetric test matrices, respectively. The most recent version (version 3.0) of MeTiS [22] was used in the experiments. As both hMeTiS and PaToH achieve  $K$ -way partitioning through recursive bisection, recursive MeTiS (pMeTiS) was used for the sake of a fair

comparison. Another reason for using pMeTiS is that direct  $K$ -way partitioning version of MeTiS (kMeTiS) produces 9 percent worse partitions than pMeTiS in the decomposition of the nonsymmetric test matrices, although it is 2.5 times faster, on the average. pMeTiS was run with the default parameters: sorted heavy-edge matching, region growing, and early-exit boundary FM refinement for coarsening, initial partitioning, and uncoarsening phases, respectively. The current version (version 1.0.2) of hMeTiS [24] was run with the parameters: greedy first-choice scheme (GFC) and early-exit FM refinement (EE-FM) for coarsening and uncoarsening phases, respectively. The V-cycle refinement scheme was not used because, in our experiments, it achieved at most 1 percent (much less on the average) better decompositions at the expense of approximately three times slower execution time (on the average) in the decomposition of the test matrices. The GFC scheme was found to be 28 percent faster than the other clustering schemes while producing slightly (1 percent–2 percent) better decompositions on the average. The EE-FM scheme was observed to be 30 percent faster than the other refinement schemes without any difference in the decomposition quality on the average.

Table 1 illustrates the properties of the test matrices listed in the order of increasing number of nonzeros. In this table, the “description” column displays both the nature and the source of each test matrix. The sparsity patterns of the Linear Programming matrices used as symmetric test matrices are obtained by multiplying the respective rectangular constraint matrices with their transposes. In Table 1, the total number of nonzeros of a matrix also denotes the total number of pins in both column-net and row-net models. The minimum and maximum number of nonzeros per row (column) of a matrix correspond to the










Fig. 7. Relative run-time performance of the proposed column-net/row-net hypergraph model (Clique-net, hMeTiS, PaToH-HCM, and PaToH-HCC) to the graph model (pMeTiS) in rowwise/columnwise decomposition of symmetric test matrices. Bars above 1.0 indicate that the hypergraph model leads to slower decomposition time.

incurring 37 percent less concurrent communication volume, on the overall average.

Although the hypergraph models perform better than the graph models in terms of number of messages, the performance gap is not as large as in the communication volume metrics. However, the performance gap increases with increasing  $K$ . As seen in Table 2, in the 64-way decomposition of the symmetric test matrices, the hypergraph model using PaToH-HCC incurs 32 percent and 10 percent less total and concurrent number of messages than the graph model, respectively. As seen in Table 3, in the rowwise decomposition of the nonsymmetric test matrices, PaToH-HCC incurs 32 percent and 26 percent less total and concurrent number of messages than the graph model, respectively.

The performance comparison of the graph/hypergraph partitioning-based 1D decomposition schemes with the conventional algorithms based on 1D and 2D [15], [30] decomposition schemes is as follows: As mentioned earlier, in  $K$ -way decompositions of  $m \times m$  matrices, the conventional 1D and 2D schemes incur the total communication volume of  $(K - 1)m$  and  $2(\sqrt{K} - 1)m$  words, respectively. For example, in 64-way decompositions, the conventional 1D and 2D schemes incur the total communication volumes of  $63m$  and  $14m$  words, respectively. As seen at the bottom of Tables 2 and 3, PaToH-HCC reduces the total communication volume to  $1.91m$  and  $0.90m$  words in the 1D 64-way

decomposition of the symmetric and nonsymmetric test matrices, respectively, on the overall average. In 64-way decompositions, the conventional 1D and 2D schemes incur the concurrent communication volumes of approximately  $m$  and  $0.22m$  words, respectively. As seen in Tables 2 and 3, PaToH-HCC reduces the concurrent communication volume to  $0.052m$  and  $0.025m$  words in the 1D 64-way decomposition of the symmetric and nonsymmetric test matrices, respectively, on the overall average.

Fig. 7 illustrates the relative run-time performance of the proposed hypergraph model compared to the standard graph model in the rowwise/columnwise decomposition of the symmetric test matrices. Figs. 8 and 9 display the relative run-time performance of the column-net and row-net hypergraph models compared to the proposed graph models in the rowwise and columnwise decompositions of the nonsymmetric test matrices, respectively. In Figs. 7, 8, and 9, for each decomposition instance, we plot the ratios of the average execution times of the tools using the respective hypergraph model to that of pMeTiS using the respective graph model. The results displayed in Figs. 7, 8, and 9 are obtained by assuming that the test matrix is given either in CSR or in CSC form, which are commonly used for SpMxV computations. The standard graph model does not necessitate any preprocessing since CSR and CSC forms are equivalent in symmetric matrices and both of them correspond to the adjacency list representation of the






Fig. 8. Relative run-time performance of the proposed column-net hypergraph model (Clique-net, hMeTiS, PaToH-HCM, and PaToH-HCC) to the graph model (pMeTiS) in rowwise decomposition of symmetric test matrices. Bars above 1.0 indicate that the hypergraph model leads to slower decomposition time than the graph model.

standard graph model. However, in nonsymmetric matrices, construction of the proposed graph model requires some amount of preprocessing time, although we have implemented a very efficient construction code which totally avoids index search. Thus, the execution time averages of the graph models for the nonsymmetric test matrices include this preprocessing time. The preprocessing time constitutes approximately 3 percent of the total execution time on the overall average. In the clique-net model, transforming the hypergraph representation of the given matrices to graphs using the randomized clique-net model introduces considerable amount of preprocessing time, despite the efficient implementation scheme we have adopted. Hence, the execution time averages of the clique-net model include this transformation time. The transformation time constitutes approximately 23 percent of the total execution time on the overall average. As mentioned earlier, the PaToH and hMeTiS tools use both CSR and CSC forms such that the construction of the other form from the given one is performed within the respective tool.

As seen in Figs. 7, 8, and 9, the tools using the hypergraph models run slower than pMeTiS using the the graph models in most of the instances. The comparison of Fig. 7 with Figs. 8 and 9 shows that the gap between the run-time performances of the graph and hypergraph models is much less in the decomposition of the nonsymmetric test matrices than that of the symmetric test matrices.

These experimental findings were expected, because the execution times of graph partitioning tool pMeTiS, and hypergraph partitioning tools hMeTiS and PaToH are proportional to the sizes of the graph and hypergraph, respectively. In the representation of an  $m \times m$  square matrix with  $Z$  off-diagonal nonzeros, the graph models contain  $|\mathcal{E}| = Z/2$  and  $Z/2 < |\mathcal{E}| \leq Z$  edges for symmetric and nonsymmetric matrices, respectively. However, the hypergraph models contain  $p = m + Z$  pins for both symmetric and nonsymmetric matrices. Hence, the size of the hypergraph representation of a matrix is always greater than the size of its graph representation, and this gap in the sizes decreases in favor of the hypergraph models in nonsymmetric matrices. Fig. 9 displays an interesting behavior that pMeTiS using the clique-net model runs faster than pMeTiS using the graph model in the column-wise decomposition of four out of nine nonsymmetric test matrices. In these four test matrices, the edge contractions during the hypergraph-to-graph transformation through randomized clique-net approach lead to less number of edges than the graph model.

As seen in Figs. 7, 8, and 9, both PaToH-HCM and PaToH-HCC run considerably faster than hMeTiS in each decomposition instance. This situation can be most probably due to the design considerations of hMeTiS. hMeTiS mainly aims at partitioning VLSI circuits of which hypergraph representations are much more sparse than the




Fig. 9. Relative run-time performance of the proposed row-net hypergraph model (Clique-net, hMeTiS, PaToH-HCM, and PaToH-HCC) to the graph model (pMeTiS) in columnwise decomposition of symmetric test matrices. Bars above 1.0 indicate that the hypergraph model leads to slower decomposition time than the graph model.

hypergraph representations of the test matrices. In the comparison of the HCM and HCC clustering schemes of PaToH, PaToH-HCM runs slightly faster than PaToH-HCC in the decomposition of almost all test matrices except in the decomposition of symmetric matrices KEN-11 and KEN-13, and nonsymmetric matrices ONETONE1 and ONETONE2. As seen in Fig. 7, PaToH-HCM using the hypergraph model runs 1.47–2.93 times slower than pMeTiS using the graph model in the decomposition of the symmetric test matrices. As seen in Figs. 8 and 9, PaToH-HCM runs 1.04–1.63 times and 0.83–1.79 times slower than pMeTiS using the graph model in the rowwise and columnwise decomposition of the nonsymmetric test matrices, respectively. Note that PaToH-HCM runs 17 percent, 8 percent, and 6 percent faster than pMeTiS using the graph model in the 8-way, 16-way, and 32-way columnwise decompositions of nonsymmetric matrix LHR34, respectively. PaToH-HCM achieves 64-way rowwise decomposition of the largest test matrix BCSSTK32 containing 44.6K rows/columns and 1030K nonzeros in only 25.6 seconds, which is equal to the sequential execution time of multiplying matrix BCSSTK32 with a dense vector 73.5 times.

The relative performance results of the hypergraph models with respect to the graph models are summarized in Table 5 in terms of total communication volume and execution time by averaging over different  $K$  values. This table also displays the averages of the best and worst

performance results of the tools using the hypergraph models. In Table 5, the performance results for the hypergraph models are normalized with respect to those of pMeTiS using the graph models. In the symmetric test matrices, direct approaches PaToH and hMeTiS produce 30 to 32 percent better partitions than pMeTiS using the graph model, whereas the clique-net approach produces 16 percent better partitions, on the overall average. In the nonsymmetric test matrices, the direct approaches achieve 34 to 38 percent better decomposition quality than pMeTiS using the graph model, whereas the clique-net approach achieves 21 to 24 percent better decomposition quality. As seen in Table 5, the clique-net approach is faster than the direct approaches in the decomposition of the symmetric test matrices. However, PaToH-HCM achieves nearly equal run-time performance as pMeTiS using the clique-net approach in the decomposition of the nonsymmetric test matrices. It is interesting to note that the execution time of the clique-net approach relative to the graph model decreases with increasing number of processors  $K$ . This is because of the fact that the percent preprocessing overhead due to the hypergraph-to-graph transformation in the total execution time of pMeTiS using the clique-net approach decreases with increasing  $K$ .

As seen in Table 5, hMeTiS produces slightly (2 percent) better partitions at the expense of considerably

TABLE 5  
Overall Performance Averages of the Proposed Hypergraph Models Normalized with Respect to Those of the Graph Models Using pMeTiS

K	pMeTiS (clique-net model)				hMeTiS				PaToH-HCM				PaToH-HCC			
	Tot. Comm. Volume			Time	Tot. Comm. Volume			Time	Tot. Comm. Volume			Time	Tot. Comm. Volume			Time
	best	worst	avg		best	worst	avg		best	worst	avg		best	worst	avg	
Symmetric Matrices: Column-net Model $\equiv$ Row-net Model																
8	0.86	0.84	0.85	2.08	0.73	0.70	0.71	8.13	0.73	0.73	0.73	2.19	0.73	0.73	0.73	2.42
16	0.86	0.84	0.83	1.90	0.70	0.66	0.66	8.95	0.70	0.69	0.68	2.25	0.71	0.69	0.69	2.43
32	0.85	0.84	0.84	1.79	0.68	0.65	0.66	9.72	0.69	0.68	0.68	2.33	0.69	0.68	0.68	2.44
64	0.85	0.84	0.84	1.78	0.71	0.68	0.69	10.64	0.72	0.69	0.70	2.41	0.72	0.69	0.70	2.56
avg	0.86	0.84	0.84	1.89	0.70	0.67	0.68	9.36	0.71	0.70	0.70	2.30	0.71	0.70	0.70	2.46
Nonsymmetric Matrices: Column-net Model																
8	0.78	0.78	0.78	1.48	0.68	0.63	0.64	5.31	0.67	0.64	0.64	1.32	0.66	0.62	0.63	1.50
16	0.80	0.78	0.78	1.44	0.66	0.63	0.64	5.53	0.67	0.64	0.65	1.37	0.65	0.62	0.63	1.56
32	0.79	0.78	0.78	1.34	0.66	0.64	0.66	5.88	0.67	0.65	0.66	1.44	0.65	0.63	0.64	1.61
64	0.80	0.79	0.79	1.34	0.69	0.68	0.68	6.17	0.69	0.68	0.68	1.45	0.67	0.66	0.66	1.62
avg	0.79	0.78	0.79	1.40	0.67	0.64	0.66	5.72	0.67	0.65	0.66	1.39	0.66	0.63	0.64	1.57
Nonsymmetric Matrices: Row-net Model																
8	0.75	0.74	0.76	1.25	0.64	0.62	0.63	5.22	0.64	0.63	0.63	1.29	0.62	0.60	0.61	1.50
16	0.75	0.74	0.75	1.15	0.65	0.63	0.64	5.34	0.65	0.63	0.65	1.33	0.62	0.61	0.62	1.54
32	0.75	0.75	0.75	1.12	0.67	0.65	0.66	5.55	0.66	0.64	0.66	1.38	0.63	0.62	0.63	1.58
64	0.76	0.77	0.76	1.09	0.67	0.67	0.67	5.84	0.66	0.65	0.66	1.36	0.64	0.63	0.63	1.50
avg	0.75	0.75	0.76	1.15	0.66	0.64	0.65	5.49	0.65	0.64	0.65	1.34	0.63	0.61	0.62	1.53

In total communication volume, a ratio smaller than 1.00 indicates that the hypergraph model produces better decompositions than the graph model. In execution time, a ratio greater than 1.00 indicates that the hypergraph model leads to slower decomposition time than the graph model.

larger execution time in the decomposition of the symmetric test matrices. However, PaToH-HCM achieves the same decomposition quality as hMeTiS for the nonsymmetric test matrices, whereas PaToH-HCC achieves slightly (2 to 3 percent) better decomposition quality. In the decomposition of the nonsymmetric test matrices, although PaToH-HCC performs slightly better than PaToH-HCM in terms of decomposition quality, it is 13 to 14 percent slower.

In the symmetric test matrices, the use of the proposed hypergraph model instead of the graph model achieves 30 percent decrease in the communication volume requirement of a single parallel SpMxV computation at the expense of 130 percent increase in the decomposition time by using PaToH-HCM for hypergraph partitioning. In the nonsymmetric test matrices, the use of the proposed hypergraph models instead of the graph model achieves 34 to 35 percent decrease in the communication volume requirement of a single parallel SpMxV computation at the expense of only 34 to 39 percent increase in the decomposition time by using PaToH-HCM.

## 6 CONCLUSION

Two computational hypergraph models were proposed to decompose sparse matrices for minimizing communication volume while maintaining load balance during repeated parallel matrix-vector product computations. The proposed models enable the representation and, hence, the decomposition of structurally nonsymmetric matrices as well as structurally symmetric matrices. Furthermore, they introduce a much more accurate

representation for the communication requirement than the standard computational graph model widely used in the literature for the parallelization of various scientific applications. The proposed models reduce the decomposition problem to the well-known hypergraph partitioning problem, thus enabling the use of circuit partitioning heuristics widely used in VLSI design. The successful multilevel graph partitioning tool MeTiS was used for the experimental evaluation of the validity of the proposed hypergraph models through hypergraph-to-graph transformation using the randomized clique-net model. A successful multilevel hypergraph partitioning tool PaToH was also implemented and both PaToH and recently released multilevel hypergraph partitioning tool hMeTiS were used for testing the validity of the proposed hypergraph models. Experimental results carried out on a wide range of sparse test matrices arising in different application domains confirmed the validity of the proposed hypergraph models. In the decomposition of the test matrices, the use of the proposed hypergraph models instead of the graph models achieved 30 to 38 percent decrease in the communication volume requirement of a single parallel matrix-vector multiplication at the expense of only 34 to 130 percent increase in the decomposition time by using PaToH, on the average. This work was also an effort towards showing that the computational hypergraph model is more powerful than the standard computational graph model as it provides a more versatile representation for the interactions among the atomic tasks of the computational domains.

## ACKNOWLEDGMENTS

This work is partially supported by the Commission of the European Communities, Directorate General for Industry under contract ITDC 204-82166, and Turkish Science and Research Council under grant EEEAG-160.

We would like to thank Bruce Hendrickson, Cleve Ashcraft, and the anonymous referees for their constructive remarks that helped improve the quality of the paper.

## REFERENCES

- [1] C.J. Alpert and A.B. Kahng, "Recent Directions in Netlist Partitioning: A Survey," *VLSI J.*, vol. 19, nos. 1-2, pp. 1-81, 1995.
- [2] C.J. Alpert, L.W. Hagen, and A.B. Kahng, "A Hybrid Multilevel/Genetic Approach for Circuit Partitioning," technical report, UCLA Computer Science Dept., 1996.
- [3] C. Aykanat, F. Ozguner, F. Ercal, and P. Sadayappan, "Iterative Algorithms for Solution of Large Sparse Systems of Linear Equations on Hypercubes," *IEEE Trans. Computers*, vol. 37, no. 12, pp. 1,554-1,567, Dec. 1988.
- [4] T. Bui and C. Jones, "A Heuristic for Reducing Fill in Sparse Matrix Factorization," *Proc. Sixth SIAM Conf. Parallel Processing for Scientific Computing*, pp. 445-452, 1993.
- [5] T. Bultan and C. Aykanat, "A New Mapping Heuristic Based on Mean Field Annealing," *J. Parallel and Distributed Computing*, vol. 16, pp. 292-305, 1992.
- [6] W. Camp, S.J. Plimpton, B. Hendrickson, and R.W. Leland, "Massively Parallel Methods for Engineering and Science Problems," *Comm. ACM*, vol. 37, pp. 31-41, Apr. 1994.
- [7] W.J. Carolan, J.E. Hill, J.L. Kennington, S. Niemi, and S.J. Wichmann, "An Empirical Evaluation of the Korb Algorithms for Military Airlift Applications," *Operations Research*, vol. 38, no. 2, pp. 240-248, 1990.
- [8] Ü.V. Çatalyürek and C. Aykanat, "Decomposing Irregularly Sparse Matrices for Parallel Matrix-Vector Multiplications," *Proc. Third Int'l Workshop Parallel Algorithms for Irregularly Structured Problems (IRREGULAR '96)*, pp. 175-181, 1996.
- [9] I.S. Duff, R. Grimes, and J. Lewis, "Sparse Matrix Test Problems," *ACM Trans. Mathematical Software*, vol. 15, pp. 1-14, Mar. 1989.
- [10] C.M. Fiduccia and R.M. Mattheyses, "A Linear-Time Heuristic for Improving Network Partitions," *Proc. 19th ACM/IEEE Design Automation Conf.*, pp. 175-181, 1982.
- [11] M. Garey, D. Johnson, and L. Stockmeyer, "Some Simplified NP-Complete Graph Problems," *Theoretical Computer Science*, vol. 1, pp. 237-267, 1976.
- [12] M.K. Goldberg and M. Burstein, "Heuristic Improvement Techniques for Bisection of VLSI Networks," *Proc. IEEE Int'l Conf. Computer Design*, pp. 122-125, 1983.
- [13] B. Hendrickson and R. Leland, "A Multilevel Algorithm for Partitioning Graphs," technical report, Sandia Nat'l Laboratories, 1993.
- [14] B. Hendrickson and R. Leland, "The Chaco User's Guide, version 2.0," Technical Report SAND95-2344, Sandia Nat'l Laboratories, 1995.
- [15] B. Hendrickson, R. Leland, and S. Plimpton, "An Efficient Parallel Algorithm for Matrix-Vector Multiplication," *Int'l J. High Speed Computing*, vol. 7, no. 1, pp. 73-88, 1995.
- [16] B. Hendrickson, "Graph Partitioning and Parallel Solvers: Has the Emperor No Clothes?," *Lecture Notes in Computer Science*, vol. 1,457, pp. 218-225, 1998.
- [17] B. Hendrickson and T.G. Kolda, "Partitioning Rectangular and Structurally Nonsymmetric Sparse Matrices for Parallel Processing," submitted to *SIAM J. Scientific Computing*.
- [18] E. Ihler, D. Wagner, and F. Wagner, "Modeling Hypergraphs by Graphs with the Same Mincut Properties," *Information Processing Letters*, vol. 45, pp. 171-175, Mar. 1993.
- [19] IOWA Optimization Center, "Linear Programming problems," <ftp://col.biz.uiowa.edu/pub/testprob/lp/gondzio>.
- [20] M. Kaddoura, C.W. Qu, and S. Ranka, "Partitioning Unstructured Computational Graphs for Nonuniform and Adaptive Environments," *IEEE Parallel and Distributed Technology*, pp. 63-69, 1995.
- [21] G. Karypis and V. Kumar, "A Fast and High Quality Multilevel Scheme for Partitioning Irregular Graphs," *SIAM J. Scientific Computing*, to appear.
- [22] G. Karypis and V. Kumar, *MeTiS: A Software Package for Partitioning Unstructured Graphs, Partitioning Meshes, and Computing Fill-Reducing Orderings of Sparse Matrices, Version 3.0*, Univ. of Minnesota, Dept. of Computer Science and Engineering, Army HPC Research Center, Minneapolis, Minn., 1998.
- [23] G. Karypis, V. Kumar, R. Aggarwal, and S. Shekhar, "Hypergraph Partitioning Using Multilevel Approach: Applications in VLSI Domain," *IEEE Trans. VLSI Systems*, to appear.
- [24] G. Karypis, V. Kumar, R. Aggarwal, and S. Shekhar, *MeTiS: A Hypergraph Partitioning Package, Version 1.0.1*, Univ. of Minnesota, Dept. of Computer Science and Engineering, Army HPC Research Center, Minneapolis, Minn., 1998.
- [25] B.W. Kernighan and S. Lin, "An Efficient Heuristic Procedure for Partitioning Graphs," *The Bell System Technical J.*, vol. 49, pp. 291-307, Feb. 1970.
- [26] T.G. Kolda, "Partitioning Sparse Rectangular Matrices for Parallel Processing," *Lecture Notes in Computer Science*, vol. 1,457, pp. 68-79, 1998.
- [27] V. Kumar, A. Grama, A. Gupta, and G. Karypis, *Introduction to Parallel Computing: Design and Analysis of Algorithms*. Redwood City, Calif.: Benjamin/Cummings, 1994.
- [28] V. Lakamsani, L.N. Bhuyan, and D.S. Linthicum, "Mapping Molecular Dynamics Computations on to Hypercubes," *Parallel Computing*, vol. 21, pp. 993-1,013, 1995.
- [29] T. Lengauer, *Combinatorial Algorithms for Integrated Circuit Layout*. Chichester, U.K.: Wiley, 1990.
- [30] J.G. Lewis and R.A. van de Geijn, "Distributed Memory Matrix-Vector Multiplication and Conjugate Gradient Algorithms," *Proc. Supercomputing '93*, pp. 15-19, 1993.
- [31] O.C. Martin and S.W. Otto, "Partitioning of Unstructured Meshes for Load Balancing," *Concurrency: Practice and Experience*, vol. 7, no. 4, pp. 303-314, 1995.
- [32] S.G. Nastea, O. Frieder, and T. El-Ghazawi, "Load-Balanced Sparse Matrix-Vector Multiplication on Parallel Computers," *J. Parallel and Distributed Computing*, vol. 46, pp. 439-458, 1997.
- [33] A.T. Ogielski and W. Aiello, "Sparse Matrix Computations on Parallel Processor Arrays," *SIAM J. Scientific Computing*, 1993.
- [34] A. Pinar, Ü.V. Çatalyürek, C. Aykanat, and M. Pinar, "Decomposing Linear Programs for Parallel Solution," *Lecture Notes in Computer Science*, vol. 1,041, pp. 473-482, 1996.
- [35] C. Pommerell, M. Annaratone, and W. Fichtner, "A Set of New Mapping and Coloring Heuristics for Distributed-Memory Parallel Processors," *SIAM J. Scientific and Statistical Computing*, vol. 13, pp. 194-226, Jan. 1992.
- [36] C.-W. Qu and S. Ranka, "Parallel Incremental Graph Partitioning," *IEEE Trans. Parallel and Distributed Systems*, vol. 8, no. 8, pp. 884-896, 1997.
- [37] Y. Saad, K. Wu, and S. Petiton, "Sparse Matrix Computations on the CM-5," *Proc. Sixth SIAM Conf. Parallel Processing for Scientific Computing*, 1993.
- [38] D.G. Schweikert and B.W. Kernighan, "A Proper Model for the Partitioning of Electrical Circuits," *Proc. Ninth ACM/IEEE Design Automation Conf.*, pp. 57-62, 1972.
- [39] T. Davis, Univ. of Florida Sparse Matrix Collection, <http://www.cise.ufl.edu/davis/sparse/>, *NA Digest*, vols. 92/96/97, nos. 42/28/23, 1994/1996/1997.



**Ümit V. Çatalyürek** received the BS and MS degrees in computer engineering and information science from Bilkent University, Ankara, Turkey, in 1992 and 1994, respectively. He is currently working towards the PhD degree in the Department of Computer Engineering and Information Science, Bilkent University, Ankara, Turkey. His current research interests are parallel computing and graph/hypergraph partitioning.



**Cevdet Aykanat** received the BS and MS degrees from Middle East Technical University, Ankara, Turkey, in 1977 and 1980, respectively, and the PhD degree from Ohio State University, Columbus, in 1988, all in electrical engineering. He was a Fulbright scholar during his PhD studies. He worked at the Intel Supercomputer Systems Division, Beaverton, Oregon, as a research associate. Since October 1988, he has been with the Department of Computer Engineering and Information Science, Bilkent University, Ankara, Turkey, where he is currently an associate professor. His research interests include parallel computer architectures, parallel algorithms, applied parallel computing, neural network algorithms, and graph/hypergraph partitioning. He is a member of the ACM, the IEEE, and the IEEE Computer Society.