

Graph Partitioning Models for Parallel Computing

Bruce Hendrickson

*Parallel Computing Sciences Dept., Sandia National Labs,
Albuquerque, NM 87185-1110. Email: bah@cs.sandia.gov.*

Tamara G. Kolda

*Computer Science and Mathematics Division, Oak Ridge National Laboratory,
Oak Ridge, TN 37831-6367. Email: kolda@msr.epm.ornl.gov.*

Abstract

Calculations can naturally be described as graphs in which vertices represent computation and edges reflect data dependencies. By partitioning the vertices of a graph, the calculation can be divided among processors of a parallel computer. However, the standard methodology for graph partitioning minimizes the wrong metric and lacks expressibility. We survey several recently proposed alternatives and discuss their relative merits.

1 Introduction

Graphs are widely used to describe the data dependencies within a computation. Recall that a *graph*, $G = (V, E)$, consists of a set of *vertices*, $V = \{v_1, v_2, \dots, v_n\}$, and a set of pairwise relationships, $E \subset V \times V$, called *edges*. If $(v_i, v_j) \in E$, then we say that vertices v_i and v_j are *neighbors*. For our purposes, the vertices of the graph will represent units of computation, and the edges will encode data dependencies. Sometimes it is appropriate to associate *weights* with the nodes and/or edges of the graph to indicate the amount of work and/or data, respectively.

¹ This work was supported by the Applied Mathematical Sciences Research Program, Office of Energy Research, U.S. Department of Energy, under contracts DE-AC05-96OR22464 and DE-AC04-94AL85000 with Lockheed Martin Energy Research Corp.

For example, differential equations are usually solved numerically on a grid. During each iteration in the process towards a solution, all the grid points are updated using neighboring values in the mesh. In Fig. 1 we show the mesh and an associated data dependency graph for a symmetric 7-point stencil. Here each vertex in the graph at right represents the computation to update the associated point on the grid. Each vertex has edges connecting it to the vertices from which it needs information. Outputs from one iteration serve as inputs for the next.

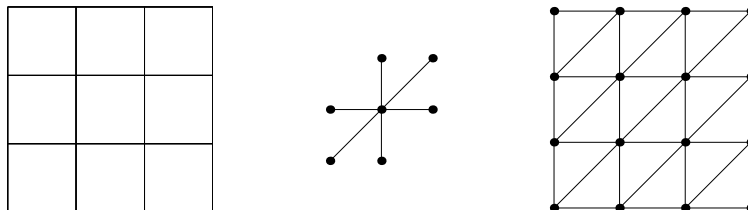


Fig. 1. Grid, Stencil, and Graph.

Once we have a graph model of a computation, *graph partitioning* can be used to determine how to divide up the work and data for an efficient parallel computation. Our objectives, stated loosely, are to evenly distribute the computations over p processors by partitioning the vertices into p equally weighted sets while minimizing interprocessor communication which is represented by edges crossing between partitions.

It is this simple relationship between graphs and computations which explains the ubiquity of graph partitioning in parallel computing. Graph partitioning is universally employed in the parallelization of calculations on unstructured grids including finite element, finite difference and finite volume techniques using both explicit and implicit methods. It is used in the parallelization of matrix-vector multiplication for all types of iterative solvers. It is also used to parallelize neural net simulations, particle calculations, circuit simulations, and a variety of other computations.

Until recently only the *standard graph partitioning approach* has been employed. The standard approach is to model the problem using a graph as described above and partition the vertices of the graph into equally weighted sets in such a way that the weight of the edges crossing between sets is minimized. Well-known software packages such as Chaco [13] and METIS [19] can be used for this purpose. Note that the graph partitioning problem is NP-hard [9], so these tools merely apply heuristics to generate approximate solutions.

Unfortunately, the standard graph partitioning approach has several significant shortcomings which are discussed in detail in §2. The edge cut metric that it tries to minimize is, at best, an imperfect model of communication in

a parallel computation. The model also suffers from a lack of expressibility which limits the applications it can address.

This paper is an extension of and an elaboration upon Hendrickson’s critique of the standard partitioning model [10]. Whereas Hendrickson restricted his concerns to matrix-vector products, in the current paper we show that the same issues plague virtually all applications of graph partitioning to parallel computation. In §3 we survey some recent work on alternative models which address some of the limitations of the standard approach. We follow with a brief discussion of algorithms in §4, and suggest some fertile areas for further research in §5.

2 Shortcomings of the Standard Graph Partitioning Approach

We discuss several shortcomings of the standard graph partitioning approach. We begin with flaws associated with using the edge cuts metric (§2.1) and continue with limitations of the standard graph model (§2.2).

2.1 Flaws of the Edge Cut Metric

Minimizing edge cuts has several major flaws. First of all, although it is not widely acknowledged, edge cuts are *not* proportional to the total communication volume. The scenario is illustrated in Fig. 2. The ovals correspond to different processors among which the vertices of the graph are partitioned. Assume that each edge has a weight of two corresponding to one unit of data being communicated in each direction. So the weight of the cut edges is ten. However, observe that the data from node v_2 on processor P_1 need only be communicated *once* to processor P_2 ; similarly with nodes v_4 and v_7 . Thus, the actual communication volume is only seven. In general, the edge cut metric does not recognize that two or more edges may be representing the same information flow, so it over counts the true volume of communication.

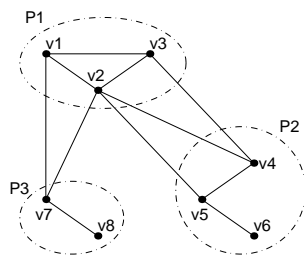


Fig. 2. Edge cuts versus communication volume

Secondly, the time to send a message on a parallel computer is a function of the *latency* (or *start-up time*) as well as the size of the message. Graph partitioning approaches try to (approximately) minimize the total volume, but not the total number of messages. Depending on the machine architecture and problem size, message latencies can be more important than message volume.

Thirdly, the performance of a parallel application is generally limited by the slowest processor. Even if the computational work is well balanced, the communication effort might not be. So, rather than minimizing the total communication volume or even the total number of message, we may instead wish to minimize the maximum volume and/or number of messages handled by any *single* processor. The standard edge cuts measure does not encapsulate this type of objective.

Lastly, on many architectures the time to send a message depends upon the distance between the sending and receiving processors. Geographic distance is not the issue here, but rather the number of switches the message is routed through. Although most modern machines have some form of cut-through or wormhole routing which enables a single message to travel quickly between distant processors, the communication network is usually handling many messages simultaneously. A message between distant processors ties up many wires which cannot be used by other messages. So to avoid message contention and improve the overall throughput of the message traffic, it is preferable to have communication restricted to processors which are near each other. So, for the problem illustrated in Fig. 2, on a one-dimensional row of processors, the layout $P_3 - P_1 - P_2$ would be preferable to $P_1 - P_2 - P_3$.

In actuality, we are interested in all of these metrics to varying degrees, depending on how they affect the overall speed of the application. So we will likely want to minimize an objective function with several components (e.g., total volume *and* total number of messages), weighted to reflect the importance of each measure. In even more complicated settings, we may wish to balance the sum of the computational and communications work on each processor while minimizing these combined objectives.

Despite these problems with the edge cut measure, the standard partitioning approach has proved successful for the parallel solution of differential equations and grid-based problems in general. There are several reasons for this success. First, grid points generally have only a small number of neighbors, so the number of edge cuts is within a small multiple of the actual communication volume. This is not true of more general problems with more complex data dependencies. Second, computational grids generally exhibit a high degree of geometric locality which ensures that good partitions exist [26]. If the grid size, n , is increased while the number of processors is held fixed, the ratio of

communication volume to computational work grows as $n^{-1/3}$ in three dimensions and $n^{-1/2}$ in two dimensions. Similarly, geometric locality ensures that the number of messages each processor sends is bounded. Lastly, the communication volume per processor is fairly evenly distributed since there usually is not an enormous difference in the size of the boundary of each piece of the grid. For all these reasons large grid computations tend to be limited by computational performance, so the details of the communication (and hence the partition) are not critical. For other applications with more complex dependency patterns the quality of the partition can have a much more dramatic impact on overall performance.

2.2 Limitations of the Standard Graph Model

Besides minimizing the wrong objective function, the standard graph partitioning approach suffers from limitations due to the lack of expressibility in the model.

One limitation of the undirected graph model is that it can only express symmetric data dependencies. For example, the graph associated with a symmetric matrix is shown in Fig. 3. For the computation $y = Ax$, vertex v_i is associated with the computation of the inner product between row i of the matrix A with the vector x . Observe that the edge between node v_1 and v_2 symbolizes a symmetric dependency: v_1 needs x_2 , and v_2 needs x_1 .

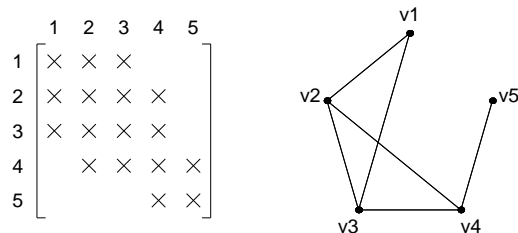


Fig. 3. Graph of a Symmetric Matrix

However, if the matrix is square but unsymmetric, then the dependencies are unsymmetric as well: v_1 might need x_2 , while v_2 does not need x_1 . This situation can be easily represented in a *directed* graph, but not in the standard model. In a directed graph, edges are directed *from* the data producing vertex *to* the data consuming vertex. There are two work-arounds to make the standard model ‘fit’ unsymmetric dependencies. The first is to convert the directed edges to undirected edges. The second is a slight extension of the first; an edge that represents only a one-way communication gets a weight of one, and an edge that represents two-way communication gets a weight of two. Unsymmetric dependencies show up in other settings as well. For example, flow calculations often involve unsymmetric stencils as depicted in Fig. 4.

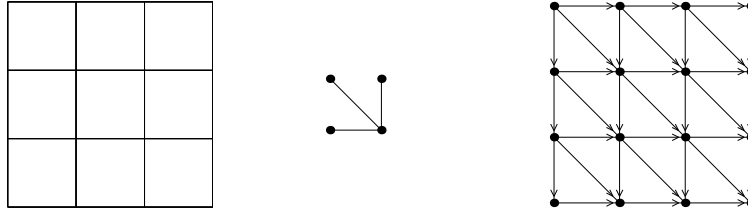


Fig. 4. Grid, Stencil, and Directed Graph.

Secondly, the symmetric model forces the partition of the input and output data to be identical. This is often desirable, particularly when the output from the previous computation is the input to the next computation. But in many situations it is an unnecessary restriction. For instance, the standard model generates the identical partitions of x and y when computing $y = Ax$ for a square matrix. For unsymmetric matrices, communication may be reduced by allowing the two partitions to differ. For example, The input x may be the result of a previous two-part operation which first computes $y = Ax$ and then $z = A^T y$; this effectively maps from x -space to y -space and back to x -space. (The data layout and communication for application of A and A^T is identical; see [12]). This brings us to the last two important issues.

The third limitation of the standard model is that it assumes that the input and output of the calculation are the size. For example, when A is rectangular in the calculation of $y = Ax$, the x - and y -spaces are of different dimensions. Recall that the standard model handles symmetric matrix-vector multiplication ($y = Ax$) by having a single vertex v_i represent *both* x_i and y_i . When the matrix is not square, x and y are of different lengths, and the standard model is inapplicable.

Lastly, even within the general framework of calculations which are repeated over and over again, it is common for the calculation to consist of several distinct phases. Examples include the application of a matrix *and* a preconditioner in an iterative method, solving a differential equation *and* applying boundary conditions, simulating different phenomena in a multi-physics calculation, and advancing a grid *and* detecting contacts in a transient dynamics computation. The union of multiple phases cannot generally be described via an undirected graph. As we will see in the next section, some alternatives to the standard model retain its basic simplicity while enabling some of these more complex situations to be handled.

3 Alternative Graph Partitioning Models

Some of the shortcomings of the standard graph partitioning model can be addressed by using recently developed alternatives. We describe four such non-standard models below.

3.1 A Bipartite Graph Model

As we noted in §2.2, the standard model using an undirected graph can only encode symmetric data dependencies and symmetric partitions. These limitations are particularly problematic for iterative solvers on unsymmetric or non-square matrices. When using preconditioners, the inability of the standard model to capture multiple phase calculations are also problematic. In [11,12,23] Kolda and Hendrickson propose a bipartite graph model for describing matrix-vector multiplication which addresses some of these shortcomings. The bipartite model can also be applied to other applications involving unsymmetric dependencies and multiple phases.

A *bipartite* graph, $G = (V_1, V_2, E)$, is a special type of graph in which the vertices are divided into two disjoint subsets, V_1 and V_2 , and $E \subset V_1 \times V_2$. So, no edges connect two vertices in the same subset; instead, all the edges cross between V_1 and V_2 .

This bipartite graph representation is most useful when the initial tasks are logically distinct from the final tasks. This occurs in the transfer between phases of the multi-phase calculations described in §2.2. An important example is matrix-vector multiplication with non-square matrices. Fig. 5 shows the bipartite graph representation of a rectangular matrix. Here, the sets V_1 and V_2 correspond to the row and column vertices respectively. Each row vertex in V_1 is weighted with the number of nonzeros in its row; e.g., row vertex r_4 has a weight of one. This weighting reflects the computational work required in the matrix-vector product. Whichever processor owns vertex r_i will own the piece y_i of the resulting solution vector $y = Ax$. The partitioning of the column vertices (V_2) affects the layout of the input vector, x . The column vertices may be left unweighted so that x may be partitioned in the optimal way to minimize edge cuts. Better yet, the column vertices may be weighted to distribute the computation of another operation on the input data such as level-1 BLAS operations or multiplication by another matrix such as a preconditioner in an iterative method.

The bipartite graph model is useful principally where the standard model fails to be a good representation, and it has three main advantages. First, it can encode a class of problems that the standard graph model cannot.

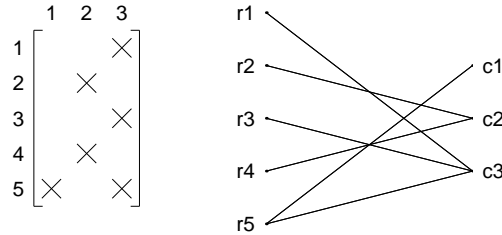


Fig. 5. Rectangular Matrix and Bipartite Graph.

Specifically, the initial (or input) vertices can be different from the final (or output) vertices. Second, even if the initial vertices are identical to the final vertices, the bipartite model allows for the initial partition to differ from the final partition. It achieves this by representing each vertex twice, once as an initial vertex and once as a final vertex. This freedom can allow for a reduction in communication. However, in many applications a symmetric partition is preferable, and this model cannot provide that. Third, by partitioning both the initial and the final vertices, it can ensure load balance in two separate operations, as mentioned above.

Although the bipartite model has expressibility that the standard model lacks, the algorithms in [12] still optimize the flawed metric of edge cuts (as well as sharing the other problems of the standard model described in §2.1). As we will see in the next section, this problem can be addressed by optimizing a graph quantity other than cut edges.

Although the bipartite model is good for describing two computational operations, it is not able to accurately encode more. One possible generalization is to use a k -partite graph in which the first set of vertices is connected to a second set, which is connected to a third set, and so on. An alternative is the multi-constraint methodology described below in §3.3.

3.2 A Hypergraph Model

Recall that edge cuts are not equal to communication volume, as illustrated in Fig. 2. In the figure, vertex v_2 on processor P_1 , for example, has two edges connecting to vertices on processor P_2 , but v_2 need only be communicated once. The true communication volume is not a function the number of edges being cut, but rather the sum of the number of processors to which each vertex has connections. More formally, the total communication volume is $\sum_i b_i$ where b_i is the number of external partitions in which vertex v_i has neighbors. We will call this quantity the *boundary cut* of a partition. The observation that boundary cuts are the more appropriate metric was made in [10] and motivated a modification in METIS to minimize this more accurate metric [17]. Boundary cuts can also be employed in the bipartite graph model from §3.1.

A more elegant expression of this metric is in the hypergraph model proposed by Çatalyürek, Aykanat, Pinar, and Pinar [3,4,25]. A *hypergraph* is a generalization of a graph in which edges can include more than two vertices. A hypergraph, $G = (V, H)$, consists of a set of vertices, V , and a set of *hyperedges*, H . Each hyperedge comprises a subset of vertices. Note that graphs are special cases of hypergraphs in which each hyperedge only contains two vertices. For our purposes, a hyperedges allow an alternative representation of the data dependencies. The partitioning problem is now to divide the vertices into equally weighted sets so that few hyperedges cross between partitions.

As we will discuss below, the hypergraph model has broader applicability than the standard approach. But even for problems that can be described with the standard model, the hypergraph model is preferable since it correctly minimizes the communication volume. To see this, consider a computation like the one in Fig. 2 which can be describe by a standard undirected graph $G = (V, E)$. Now construct an equivalent hypergraph (V, H) with $|V|$ hyperedges. Each vertex v_i G , corresponds to a hyperedge h_i consisting of v_i and all its neighbors in G . A hyperedge reflects all of the entities that either produce or consume a piece of data. When the vertices are partitioned among processors, that piece of data must be communicated from the processor which produced it to to all those consume it. Thus, the communication associated with a hyperedge is one less than the number of processors its constituent vertices are partitioned among. (This corresponds to the boundary cut value from the discussion above.) So by partitioning the hypergraph in such a way that hyperedges are split among as few processors as possible, the model correctly minimizes communication volume.

In [4], Çatalyürek and Aykanat apply this model to symmetric matrix-vector multiplication. For a set of highly unstructured matrices from linear programming problems they report that the hypergraph model reduces communication by over 30% on average over the standard partitioning approach. However, for reasons discussed in §2.1, the gains were more modest for matrices from grid calculations, generally less than 10% [1].

In addition to resolving the principle problem of the edge cut metric, the hypergraph approach is more expressive than the standard model. It can encode problems in with unsymmetric dependencies and even problems in which the initial vertices differ from the final vertices. In Figure 6, we show two different sketches of a hypergraph relating the data dependencies for the rectangular matrix-vector multiply in Fig. 5. For example, hyperedge h_2 contains all the vertices that need x_2 , i.e., $\{v_2, v_4\}$. In the left figure the hyperedges are illustrated by the ovals. In the right figure the vertices are on one side and hyperedges on the other, and each hyperedge is connected to the vertices which comprise it. We include this second hypergraph representation to underline the one-to-one relationship between hypergraphs and bipartite graphs.

The hypergraph partitioning model is closely related to the bipartite model from §3.1, but the partitioning objectives are different.

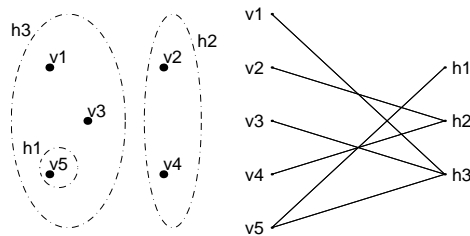


Fig. 6. Two Hypergraph Representations.

The guiding principle in the construction of a hypergraph is that each hyperedge contains the set of vertices which generate or need some data. This principle applies equally well to the case when dependencies are uni-directional, and it continues to correctly model the communication volume. However, there is a subtle requirement that the data is produced by one of the vertices that depends on it. For example, in Fig. 6 we assume that the data associated with hyperedge h_1 will live on the processor that owns vertex v_5 . If that is not the case for some reason, e.g., h_1 is the output of v_1 , then h_1 should also include its producing vertex, e.g., v_1 .

The hypergraph model can also be used even in cases where the input and output data partitions are not identical, although it is perhaps not as natural as the bipartite model in this case. We simply find the best partition for the computation nodes using a hypergraph partitioner, and this yields a partition of the output data. Then, rather than assuming the input data has the same partition as the output data, we can calculate the optimal input data partitioning as an assignment problem. So, the hypergraph model can be an alternative to the bipartite model when we are only encoding one operation; however, the bipartite (or k -partite) models are still best when encoding multi-step operations.

In summary, we find the hypergraph model to be uniformly superior to the standard model. It is also an attractive alternative to the bipartite model for unsymmetric problems when only one operation is being encoded. However, the bipartite (or k -partite) models are still more powerful when encoding multi-phase operations. This is particularly true when the bipartite model minimizes the boundary cut value as discussed above.

3.3 Multi-Constraint Partitioning

The bipartite model from §3.1 is able to describe some types of multi-phase calculations. An alternative approach is the multi-constraint partitioning model

of Karypis and Kumar [21]. Strictly speaking, the multi-constraint approach is not an alternative to other models but rather an augmentation. In the multi-constraint model, each vertex is assigned a vector of k weights which represent the work associated with that vertex in each of k computational phases. The goal is now to partition the vertices of that graph in such a way that communication is minimized and that each of the k weights is balanced. In this way, each phase of the computation will be load balanced. The edges in the graph represent data dependencies in all the computational phases.

This is a very general and powerful model. For instance, when solving a differential equation *and* also applying boundary conditions, each vertex can have two weights. The first weight will reflect the work required by a grid point in the solver, and the second can encode the work required for the boundary condition. For vertices not on the boundary, the value of the second weight will be zero. So partitioning this problem will ensure that the equation solver is balanced in such a way that each processor has an equal fraction of the surface vertices.

The multi-constraint model includes the bipartite (and k -partite) approaches as a special case. Given a bipartite graph $G = (V_1, V_2, E)$, an equivalent multi-constraint model would have a set of vertices $V = V_1 \cup V_2$, and edges identical to those in the bipartite graph. Each vertex would be assigned two weights, one for the phase modeled by V_1 and the second for the phase modeled by V_2 . Hence, each vertex would have one of its weights set to zero. More generally, the multi-constraint model can encode multiple phases with distinct vertices via a model in which it includes the union of all vertices in all phases.

As originally proposed by Karypis and Kumar, the multi-constraint model attempts to minimize edge cuts, but this is an unnecessary restriction. Hyper-edges could be used or, equivalently, the boundary cut value from §3.2.

Although the power of the model is attractive, partitioning general multi-criteria problems is difficult. When other, simpler models can be applied, they may be easier to work with.

3.4 *Skewed partitioning*

Yet another alternative to the standard partitioning model is the *skewed* partitioning approach developed by Pellegrini [24] and Hendrickson, Leland and Van Driessche [15]. As with the multi-constraint model, skewed partitioning is really an augmentation of any of the other graph partitioning models rather than a true alternative. In the skewed model, each vertex is allowed to have a set of k preference values expressing its respective desire to be in each of the k sets. When determining how to partition the vertices, these

preference values are considered along with the metrics of communication cost.

Preference values can be used in several different ways to achieve different objectives. In dynamic load balancing it is desirable that the new partition be similar to the existing one to limit the amount of data that needs to be moved. This can be encoded in the preference values by giving each datum a preference to remain in its current partition [27]. The magnitude of the preference values can be adjusted to trade off between partition quality and reduction in data movement.

Another use for preference values is to encourage communicating objects to be assigned to architecturally close processors to reduce message congestion [24,15]. Assume you partitioning for p processors by recursive application of a k -way partitioner. After the first partition, the graph is divided into k parts which are assigned to k portions of the parallel machine. When doing subsequent partitions, each vertex can be assigned a preference to be assigned to a portion of the machine which is near its neighbors. In this way, the partitioning step is coupled with the problem of assigning partitions to processors. The result is a partition which exhibits better message locality. As before, the magnitude of the preferences can be altered to trade off between partition quality and message locality.

This same idea was developed independently in the circuit placement community to place circuit elements on a chip with short overall wire lengths [7]. Several algorithms for this problem have been devised including multilevel and spectral approaches [15].

4 Partitioning Algorithms

The different graph partitioning models reviewed in §3 are only viable if efficient and effective algorithms can be developed to partition them. Fortunately, the *multilevel* paradigm for partitioning has proven to be quite robust and general. The multilevel approach was devised independently by several researchers in the early 90s [2,6,14] and popularized by the the Chaco [13] and METIS [19] partitioning tools. The basic idea is quite simple. A large graph is approximated by a sequence of smaller and smaller graphs. The smallest graph is partitioned using any suitable algorithm. This partition is then propagated back through the sequence of larger and larger graphs, being refined along the way.

Adapting the multilevel approach to a particular partitioning problem requires the following tools.

- (1) A method for generating a sequence of smaller graphs which preserve the essential properties of the original.
- (2) An algorithm for partitioning the smallest graph.
- (3) A refinement technique for improving the partition as it is propagated back up to the original graph.

These tools are generally straightforward to devise; however, the precise details of these tools require some attention to the nature of the partitioning problem being addressed. The generation of smaller graphs is typically done with some kind of edge contraction scheme. Any existing algorithm which handles weights on edges and vertices can be used to partition the smallest graph. The refinement often involves a greedy algorithm in the spirit of Kernighan-Lin [22].

Following the multilevel paradigm, efficient and effective partitioners have been developed for partitioning graphs to minimize edge cuts [2,14], minimize vertex cuts [16], and perform multi-constraint partitioning [21]. The same approach has been successfully used to partition hypergraphs to minimize cut hyperedges [6,18,5] and to partition bipartite graphs [12]. The flexibility of the technique makes it well suited to address a range of different partitioning models and metrics.

5 Conclusions and Directions for Further Research

In many respects, those of us working in the partitioning field have been fortunate. The dominant application for our algorithms and tools has been differential equation solvers. Whether solved implicitly or explicitly, these applications produce dependency graphs which are fairly easy to partition, and large problems are computation rather than communication bound. The applications achieved good parallel performance despite the limitations of our approaches.

But different applications are becoming common which are much more sensitive to partition quality. Challenging partitioning problems that arise from interior point methods for linear programming, least squares problems, circuit simulation, truncated singular value computations for latent semantic indexing in information retrieval, and other applications are revealing the limitations of our traditional approaches. The standard graph partitioning methodology optimizes an inappropriate quantity, and its expressibility is too limited to address some important classes of applications.

We surveyed several alternative models which address some of the problems with the standard methodology. The bipartite model and the hypergraph

model can both handle unsymmetric dependencies. The hypergraph approach correctly encodes communication volume, while the bipartite model and its k -partite generalization have the advantage of being able to represent some multi-phase calculations. The multi-constraint approach offers an alternative way to represent multiple phases, while the skewed partitioning model provides a mechanism for including extra information in a partitioning problem to, for example, reduce message congestion. However, these new models only start to address the problems detailed in §2. A number of important open problems remain, including the following.

- (1) **Partitioning for alternate objectives or multi-objectives.** Models which are well suited to minimizing the number of messages or the maximum communication per processor instead of the total communication are still needed. Of further value would be hybrid models which encapsulate several metrics. New partitioning metrics may lead to new algorithmic challenges.
- (2) **Partitioning for alternative architectures.** Most of the work in partitioning techniques have been motivated by distributed memory architectures and has tried to minimize interprocessor communication. Similar, but not identical, issues occur in shared memory machines (SMPs). It is advantageous to partition the shared memory between the processors to minimize cache coherence overhead. However, the precise objectives in the shared memory setting may differ from those for distributed memory machines. There is little published work on this problem.

Other architectural trends pose different challenges for partitioners. One important development is the growing importance of heterogeneous machines. Many current parallel machines consist of a collection of shared memory nodes networked together. These machines exhibit significant network heterogeneity. Accesses within an SMP are fast, but between SMPs are slow. It is unclear how best to partition for these architectures.

Another important architectural development is the growing popularity of build-it-yourself parallel computers, epitomized by Beowulf-class machines. Machines built in this way can exhibit both network and processor heterogeneity. The partitioner will need to worry about differing processor speeds and memory sizes, as well as varying access times. Appropriate machine models and partitioning approaches for heterogeneous architectures is largely an untouched area.

- (3) **Parallel partitioning.** Most of the work on parallel partitioning has been done in the context of dynamic load balancing. Algorithmically, dynamic load balancing is more challenging than the problems we have been discussing since there is a pre-existing partition. If the new partition deviates significantly from the current one, then a large remapping cost is incurred. This consideration does not occur in static settings and complicates the evaluation of dynamic partitioning algorithms. Independent of dynamic problems, several trends are increasing the need for parallel par-

tioners. First is the interest in very large meshes, which will not easily fit on a sequential machine and so must be partitioned in parallel. Second, for a more subtle reason, is the growing interest in heterogeneous parallel architectures. Generally, partitioning is performed as a preprocessing step in which the user specifies the number of processors the problem will run on. With heterogeneous parallel machines, the number of processors is insufficient — the partitioner should also know their relative speeds and memory sizes. A user will want to run on whatever processors happen to be idle when the job is ready, so it is impossible to provide this information to a partitioner in advance. A better solution is to partition on the parallel machine when the job is initiated. A number of parallel partitioners have been implemented including Jostle [28] and ParMETIS [20]. This is an active area of research.

- (4) **Partitioning for domain decomposition.** Domain decomposition is a numerical technique in which a large grid is broken into smaller pieces. The solver works on individual subdomains first, and then couples them together. The properties of a good decomposition are not entirely clear, and they depend upon the details of the solution technique. But they are almost certainly not identical to the criteria used to minimize parallel communication. For instance, Farhat, et al. [8] argue that the domains must have good aspect ratios (e.g., not be long and skinny). It can also be important that subdomains are connected, even though the best partitions for parallel communication need not be. For the most part, practitioners of domain decomposition have made due with partitioning algorithms developed for other purposes, with perhaps some minor perturbations at the end. But a concerted effort to devise schemes which meet the need of this community could lead to significant advances. Progress in this area will probably require a combination of ideas from numerical analysis and graph algorithms.

Despite the general feeling that partitioning is a mature area, there are a number of open problems and many opportunities for significant advances in the state of the art. We expect to see a continuing stream of new insights and approaches which more closely fit the different classes of applications. As the hegemony of the standard approach crumbles, we foresee a balkanization of the partitioning field as different researchers choose to work on different models and applications. This is a positive development to the extent that this more focused work leads to better tools for specific applications.

References

- [1] C. Aykanat. Private Communication, 1998.
- [2] T. Bui and C. Jones. A heuristic for reducing fill in sparse matrix factorization.

- In *Proc. 6th SIAM Conf. Parallel Processing for Scientific Computing*, pages 445–452. SIAM, 1993.
- [3] U. V. Çatalyürek and C. Aykanat. Decomposing linear programs for parallel solution. In A. Ferreira et al., editors, *Parallel Algorithms for Irregularly Structured Problems, Third International Workshop, Irregular '96, Santa Barbara, CA, USA, August 19–21, 1996*, number 1117 in Lecture Notes in Computer Science, pages 75–86. Springer-Verlag, 1996.
 - [4] Ü. Çatalyürek and C. Aykanat. Hypergraph-partitioning based decomposition for parallel sparse-matrix vector multiplication. *IEEE Trans. Parallel Distrib. Sys.*, 1999. To appear.
 - [5] Ü. Çatalyürek and C. Aykanat. PaToH a multilevel hypergraph partitioning tool for decomposing sparse matrices and partitioning VLSI circuits. Technical Report BU-CEIS-9902, Dept. Computer Engineering and Information Science, Bilkent University, Turkey, 1999.
 - [6] J. Cong and M. L. Smith. A parallel bottom-up clustering algorithm with applications to circuit partitioning in VLSI design. In *Proc. 30th ACM/IEEE Design Automation Conference*, pages 755–760. ACM/IEEE, 1993.
 - [7] A. E. Dunlop and B. W. Kernighan. A procedure for placement of standard-cell VLSI circuits. *IEEE Trans. CAD*, 4(1):92–98, Jan. 1985.
 - [8] C. Farhat, N. Maman, and G. Brown. Mesh partitioning for implicit computation via domain decomposition: Impact and optimization of the subdomain aspect ratio. *Int. J. Num. Meth. Engrg.*, 38:989–1000, 1995.
 - [9] M. Garey, D. Johnson, and L. Stockmeyer. Some simplified NP-complete graph problems. *Theoretical Computer Science*, 1:237–267, 1976.
 - [10] B. Hendrickson. Graph partitioning and parallel solvers: Has the emperor no clothes? In A. Ferreira et al., editors, *Solving Irregularly Structured Problems in Parallel: 5th International Symposium, Irregular '98, Berkeley, California, USA, August 9–11, 1998*, number 1457 in Lecture Notes in Computer Science, pages 218–225. Springer-Verlag, 1998.
 - [11] B. Hendrickson and T. G. Kolda. Partitioning sparse rectangular matrices for parallel computations of Ax and $A^T v$. In B. Kågström et al., editors, *Applied Parallel Computing in Large Scale Scientific and Industrial Problems, 4th International Workshop, PARA'98, Umeå, Sweden, June 14–17, 1998*, number 1541 in Lecture Notes in Computer Science, pages 239–247. Springer-Verlag, 1998.
 - [12] B. Hendrickson and T. G. Kolda. Partitioning nonsquare and nonsymmetric matrices for parallel processing. *SIAM J. Sci. Comput.*, to appear.
 - [13] B. Hendrickson and R. Leland. The Chaco user's guide, version 2.0. Technical Report SAND95-2344, Sandia National Laboratories, Albuquerque, New Mexico, 1995.

- [14] B. Hendrickson and R. Leland. A multilevel algorithm for partitioning graphs. In *Proc. 1995 ACM/IEEE Supercomputing Conference, San Diego, California, USA, December 3–8, 1995*. ACM, New York, 1995.
- [15] B. Hendrickson, R. Leland, and R. V. Driessche. Skewed graph partitioning. In M. Heath et al., editors, *Proc. 8th SIAM Conf. Parallel Processing for Scientific Computing*. SIAM, 1997.
- [16] B. Hendrickson and E. Rothberg. Improving the runtime and quality of nested dissection ordering. *SIAM J. Sci. Comput.*, 20(2):468–489, 1998.
- [17] G. Karypis. Private Communication, 1998.
- [18] G. Karypis, R. Aggarwal, V. Kumar, and S. Shekhar. Hypergraph partitioning using multilevel approach: Application in VLSI domain. In *Proc. 34th ACM/IEEE Design Automation Conference*. ACM/IEEE, 1997.
- [19] G. Karypis and V. Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. Technical Report 95-035, Dept. Computer Science, Univ. Minnesota, Minneapolis, Minnesota, 1995.
- [20] G. Karypis and V. Kumar. Parallel multilevel graph partitioning. Technical Report 95-036, Dept. Computer Science, Univ. Minnesota, Minneapolis, Minnesota, 1995.
- [21] G. Karypis and V. Kumar. Multilevel algorithms for multi-constraint graph partitioning. Technical Report 98-019, Dept. Computer Science, Univ. Minnesota, Minneapolis, Minnesota, May 1998.
- [22] B. W. Kernighan and S. Lin. An efficient heuristic procedure for partitioning graphs. *Bell System Technical Journal*, 1970.
- [23] T. G. Kolda. Partitioning sparse rectangular matrices for parallel processing. In A. Ferreira et al., editors, *Solving Irregularly Structured Problems in Parallel: 5th International Symposium, Irregular '98, Berkeley, California, USA, August 9–11, 1998*, number 1457 in Lecture Notes in Computer Science, pages 68–79. Springer-Verlag, 1998.
- [24] F. Pellegrini. Static mapping by dual recursive bipartitioning of process and architecture graphs. In *Proc. SHPCC'94*, pages 486–493. IEEE, 1994.
- [25] A. Pinar, U. V. Çatalyürek, C. Aykanat, and M. Pinar. Decomposing linear programs for parallel solution. In J. Dongarra et al., editors, *Applied Parallel Computing in Computations in Physics, Chemistry and Engineering Science: Second International Workshop, PARA '95, Lyngby, Denmark, August 21–24, 1995*, number 1041 in Lecture Notes in Computer Science, pages 473–482. Springer-Verlag, 1996.
- [26] S.-H. Teng. *Points, Spheres and Separators: A Unified Geometric Approach to Graph Partitioning*. PhD thesis, Dept. Computer Science, Carnegie Mellon Univ., Pittsburgh, PA, 1991.

- [27] R. Van Driessche and D. Roose. Dynamic load balancing with a spectral bisection algorithm for the constrained graph partitioning problem. In *High-Performance Computing and Networking*, number 919 in Lecture Notes in Computer Science, pages 392–397. Springer, 1995. Proc. International Conf. and Exhibition, Milan, Italy, May 1995.
- [28] C. Walshaw, M. Cross, and M. Everett. Mesh partitioning and load-balancing for distributed memory parallel systems. In B. Topping, editor, *Proc. Parallel & Distributed Computing for Computational Mechanics, Lochinver, Scotland, 1997*, 1998.