

Parallel Dynamic Graph Partitioning for Adaptive Unstructured Meshes

C. Walshaw, M. Cross, and M. G. Everett

Centre for Numerical Modelling and Process Analysis, University of Greenwich, London, SE18 6PF, United Kingdom

A parallel method for the dynamic partitioning of unstructured meshes is described. The method introduces a new iterative optimization technique known as relative gain optimization which both balances the workload and attempts to minimize the inter-processor communications overhead. Experiments on a series of adaptively refined meshes indicate that the algorithm provides partitions of an equivalent or higher quality to static partitioners (which do not reuse the existing partition) and much more rapidly. Perhaps more importantly, the algorithm results in only a small fraction of the amount of data migration compared to the static partitioners. © 1997 Academic Press

Key Words: graph-partitioning; adaptive unstructured meshes; load-balancing; parallel computing.

1. INTRODUCTION

The use of unstructured mesh codes on parallel machines can be one of the most efficient ways to solve large computational fluid dynamics (CFD) and computational mechanics (CM) problems. Completely general geometries and complex behavior can be readily modeled and, in principle, the inherent sparsity of many such problems can be exploited to obtain excellent parallel efficiencies. An important consideration, however, is the problem of distributing the mesh across the memory of the machine at runtime so that the computational load is evenly balanced and the amount of interprocessor communication is minimized. It is well known that this problem is NP complete, so in recent years much attention has been focused on developing suitable heuristics, and some powerful methods, many based on a graph corresponding to the communication requirements of the mesh, have been devised, e.g., [2, 9, 15].

An increasingly important area for mesh partitioning arises from problems in which the computational load varies throughout the evolution of the solution. For example, heterogeneity in either the computing resources (e.g., processors which are unevenly matched or not dedicated to single users) or in the solver (e.g., solving for flow or stress in different parts of the domain in a multiphysics casting simulation) can result in load imbalance and poor performance. Alternatively, time-dependent unstructured mesh codes which use adaptive refinement can give rise to a series of meshes in which the position and density of the data points varies dramatically over the

course of an integration and which may need to be frequently repartitioned for maximum parallel efficiency. This dynamic partitioning problem has not been nearly as thoroughly studied as the static problem but related work can be found in [4, 5, 11, 12, 15, 18].

The dynamic evolution of load has three major influences on possible partitioning techniques; cost, reuse, and parallelism. First, frequent load balancing may be required and so must have a low cost relative to that of the solution algorithm in between. This could potentially restrict the use of high quality partitioning algorithms but fortunately, if the mesh has not changed too much, it is a simple matter to interpolate the existing partition from the old mesh to the new and use this as the starting point for repartitioning, [18]. In fact, not only is the load balancing likely to be unnecessarily computationally expensive if it fails to use this information, but also the mesh elements will be redistributed without any reference to their previous “home processor” and heavy data migration may result. Finally, the data is distributed and so should be repartitioned *in situ* rather than incurring the expense of transferring it back to some host processor for load balancing and some powerful arguments have been advanced in support of this proposition, [11]. Collectively these issues call for parallel load balancing and, if a high quality partition is desired, a parallel optimization algorithm.

In this paper we describe such a parallel optimization technique (Sect. 2) which incorporates a distributed load-balancing algorithm and which provides an extremely fast solution to the problem of dynamically load-balancing unstructured meshes. In addition, a parallel graph contraction technique (described in Sect. 3) can be employed to enhance the partition quality and the resulting strategy (which can also be applied to static partitioning problems) outperforms or matches results from existing state-of-the-art static mesh partitioning algorithms.

Here, in particular, we focus on the case arising from adaptively refined meshes where we assume that the mesh will be repartitioned after each refinement phase. However, the method is also applicable to the more general case where load may be constantly varying, and in [1] a method for determining how frequently to partition (for maximum efficiency) is described, together with examples using the same partitioning techniques.

1.1. Notation and Definitions

Let $G = G(V, E)$ be an undirected graph of V vertices with E edges which represent the data dependencies in the mesh and let P be a set of processors. We assume that both vertices and edges are weighted (with positive integer values) and that $|v|$ denotes the weight of a vertex v , $|S| := \sum_{v \in S} |v|$ the weight of a subset $S \subset V$ and similarly for edges. Once the vertices are partitioned into $|P|$ sets we denote the subdomains by S_p , for $p \in P$ and the optimal subdomain weight is given by $W := \lceil |V|/|P| \rceil$. We denote the set of cut (or intersubdomain) edges by E_c and the border of each subdomain, B_p , is defined as the set of vertices in S_p which have an edge in E_c . We shall use the notation \leftrightarrow to mean “is adjacent to,” for example, for $u, v \in V$, $u \leftrightarrow v$ if there exists $(u, v) \in E$.

The definition of the graph-partitioning problem is to partition the vertices, V , into $|P|$ disjoint sets, one per processor, such that the load or vertex weight in each subdomain is evenly balanced while the communications cost is minimized. More precisely we seek a partition such that $S_p \leq W$ for $p \in P$ (although this is not always possible for graphs with non-unitary vertex weights) and such that $|E_c|$ is minimized (though see Sect. 2.2).

1.2. Parallelization

The algorithms described run equally well in parallel or in serial and for the parallel version we use the single program, multiple data (SPMD) paradigm with message passing in the (reasonable) expectation that the underlying unstructured mesh application will do the same. (If the mesh application does not use partitioned data, why would it need to call a graph partitioner?) To this end, each processor is assigned to a subdomain and stores a double linked-list of the vertices within that subdomain. However, each processor also maintains a read only “halo” of neighboring vertices in other subdomains. For the serial version the migration of vertices simply involves transferring them from one linked-list to another. In parallel, however, this task is far more complicated as migrating vertices, together with newly created halo vertices, must be packed into messages, sent off to the destination processor, unpacked, and the pointer based data structure recreated there. In addition, “halo updates” must be regularly carried out to inform neighboring processors of new values attached to their halo vertices. Since the existence of $|V|$ length arrays is simply not scalable in memory terms, location of vertices locally on a processor is carried out using their global index and some sophisticated hash table and binary tree searches.

2. OPTIMIZATION

In this section we present a new parallel iterative algorithm for load-balancing and optimizing unstructured mesh partitions. The method is based on the concept of *relative gain*, also described in [16] and, in common with similar techniques, e.g., [9], we localize the vertex migration with respect to the current partition by only allowing vertices on subdomain bor-

ders to migrate to neighboring subdomains (i.e., to subdomains in which a neighboring vertex lies).

2.1. Load-balancing

The method includes an iterative load-balancing scheme which runs alongside the optimization to evenly distribute the workload over the processors. The generalized load-balancing problem is a very important area for research in its own right with a vast range of applications and here we use an elegant technique developed by Hu and Blake, [8]. It is related to the commonly used diffusive methods, e.g., [3], but has faster convergence and minimizes the Euclidean norm of the transferred weight. The algorithm simply involves solving the system $L\mathbf{x} = \mathbf{b}$, where L is the Laplacian of the subdomain graph, ($L_{pp} = \text{degree}(S_p)$; $L_{pq} = -1$ if $S_p \leftrightarrow S_q$, $L_{pq} = 0$ otherwise), $b_p = |S_p| - W$ and the weight to be transferred across edge (S_p, S_q) is then given by $x_p - x_q$. Typically this system is solved iteratively with a conjugate gradient algorithm and is similar to diffusive methods in that each step is like a diffusive step, but with the diffusion coefficients being determined iteratively using a conjugate gradient search rather than being fixed throughout the procedure. Since the subdomain graph is connected (which it must be since we have assumed that the original graph is connected) Hu and Blake demonstrate that the conjugate gradient iterations are guaranteed to converge in less than $|P| - 1$ iterations and often much faster than that if the subdomain graph has special structure, [8]. In practice, we found that convergence was usually reached within 10 iterations for the examples presented in Section 4; in the configurations using graph reduction (Sect. 3) this is reduced still further in the final levels of optimization, as balance is already achieved on the coarser graphs.

This algorithm (or, in principle, any other distributed load-balancing algorithm) then determines how much weight to transfer across the edges of the subdomain graph and we use f_{pq} to denote the required flow from S_p to S_q . Note that the flow is positive ($f_{pq} \geq 0$) and unidirectional (i.e., if $f_{pq} > 0$ then $f_{qp} = 0$ and vice versa) and even if this is not the case we can force it to be so by setting $f_{pq} = f_{pq} - \max(f_{pq}, f_{qp})$ and $f_{qp} = f_{qp} - \max(f_{pq}, f_{qp})$. Note also that if there is not a sufficient weight of vertices in a particular border to satisfy the required flow, then outstanding flow is recorded and added in at the next iteration.

2.2. The Gain and Preference Functions

A key concept in many graph partition optimization algorithms is the idea of gain and preference functions. Loosely, the gain $g(v, q)$ of a vertex v in subdomain S_p can be calculated for every other subdomain, S_q , $q \neq p$, and expresses some “estimate” of how much the partition would be “improved” were v to migrate to S_q . The preference $f(v)$ is then just the value of q which maximizes the gain—i.e., $f(v) = q$ where $g(v, q)$ attains $\max_{r \in P} g(v, r)$.

The gain is usually directly related to some cost function which measures the quality of the partition and which we aim

to minimize. Typically the cost function used is simply the total weight of cut edges, $|E_c|$, and then the gain expresses the change in $|E_c|$. More recently, however, there has been some debate about the most important quantity to minimize and in [14], Vanderstraeten *et al.* demonstrated that it can be extremely effective to vary the cost function based on a knowledge of the solver. Meanwhile, in [17] we showed that the architecture of the parallel machine and how the partition is mapped down onto its communications network can also play an important role. Whichever cost function is chosen, however, the idea of gains is generic. For the purposes of this paper we shall assume that the gain $g(v, q)$ just expresses the reduction in the cut-edge weight, $|E_c|$.

2.3. Relative Gain Optimization

Having determined the required flow across the edges of the subdomain graph (Sect. 2.1) we need to migrate vertices from adjacent subdomains in order to satisfy that flow. Choosing appropriate vertices to migrate is not an easy task because we also wish to optimize the partition quality with respect to the cost function. Indeed, in order to obtain partitions of the highest quality, it is likely that vertices will need to be exchanged even if there is no flow required. Simply moving vertices with the highest gain is not a satisfactory solution, however, as it means that adjacent vertices may be swapped simultaneously (an event often known as a collision) and this may lead to an *increase* in the cost. We have previously addressed this problem by using a Kernighan–Lin (KL) type algorithm run in the boundary regions alone, [15], but this causes a loss of efficiency in parallel because, in order to retain the hill-climbing abilities of the algorithm, processors must maintain edges between halo vertices and this turns out to be a costly task. In addition, we have tried a red–black coloring strategy where, at each iteration, vertices from only one of each pair of neighboring subdomains are migrated, [16]. The algorithm is fast and efficient but we have not succeeded in generating the highest quality of partitions using this technique which has been independently investigated by Karypis and Kumar, [10]. Here, however, we introduce a new strategy which uses the concept of *relative gain*.

Bulk Migration. The optimization takes a series of iterative steps in which vertices may be migrated from every subdomain, S_p , to each of its neighbors $\{S_q : S_q \leftrightarrow S_p\}$. The first part of each iterative step is to use a simple formula based on both the flow and the total weight of vertices with positive gain to determine how much total load to migrate. First, for the interface between subdomains S_p and S_q , define border regions, B_{pq} , as the set of vertices in B_p (the border of S_p) whose preference is q , i.e., $B_{pq} = \{v \in B_p : f(v) = q\}$ and let g_{pq} be the total weight of vertices in B_{pq} with gain > 0 (and similarly for B_{qp} and g_{qp}). Then if $d = \max(g_{pq} - f_{pq} + g_{qp} - f_{qp}, 0)$, the load to be migrated from S_p to each neighbor S_q is set to $a_{pq} = f_{pq} + d/2$.

To motivate this formula a little consider the following. First of all, the amount of load to be migrated, a_{pq} , is decided by satisfying any required flow, f_{pq} , and we assume that this takes place by migrating vertices with the highest positive gain. Thus, after the flow has been satisfied the amount of vertices with positive gain is approximately given by $G_{pq} = g_{pq} - f_{pq} + g_{qp} - f_{qp}$. It could be argued that this will be an underestimate if $f_{pq} > g_{pq}$, but in this case the scheme is cautious rather than reckless. At this point we wish, in a similar manner to the KL algorithm, to swap vertices so that none with a positive gain remain. After some experimentation we have found that simply moving $G_{pq}/2$ from S_p to S_q and vice versa, ensures fast and effective optimization provided the vertices are chosen carefully.

Relative Gain. To determine which vertices to migrate we use the concept of relative gain which we define as follows; for a vertex $v \in B_{pq}$ let $\Gamma_q(v)$ be the set of vertices in B_{qp} adjacent to v , i.e., $\Gamma_p(v) = \{u \in B_{qp} : u \leftrightarrow v\}$. The relative gain of a vertex v is then defined as $g(v, q) - \sum_{u \in \Gamma_q(v)} g(u, p) / O[\Gamma_q(v)]$ (where $O[\Gamma_q(v)]$ represents the number of vertices in $\Gamma_q(v)$). Put more simply, the relative gain of a vertex v is just the gain of v less the average gain of opposing vertices, and gives an indication of which are the best vertices to move in order to avoid collisions. Thus, to prioritize the migration, for each subdomain S_p , vertices in each border B_{pq} are sorted by relative gain, largest first, and a weight of a_{pq} is migrated to S_q according to this ordering. The sorting carried out need not be a full sort since it is only necessary to determine the level of relative gain below which no vertices will be moved and we have implemented a simple set-based sort on this basis. For the parallel version, at each iteration the gains of border vertices are calculated locally (in parallel) and propagated to neighboring processors via a halo update operation (Sect. 1.2) and then the relative gains may be calculated locally.

Convergence. The algorithm is not as predictable as either Kernighan–Lin optimization, or the red–black scheme mentioned above, which can both predict exactly the improvement in cost for a bipartition and fairly accurately for a multiway partition. However, although the relative gain gives no more than an indication of which vertices to move, in practice it works very effectively and collisions are rare since the formulation takes into account the likelihood of opposite vertices migrating. The method usually converges robustly with the global cost, $|E_c|$, decreasing monotonically, but because of this impreciseness, it is necessary to prevent cyclic “thrashing” by terminating the optimization after a couple of iterations if the cost has not decreased.

3. GRAPH REDUCTION

The algorithm described above provides what is essentially very localized optimization and it has been recognized for some time that an effective way of both speeding up optimization and, perhaps more importantly, giving it a more global

perception is to use graph reduction. The idea is to group vertices together to form *clusters*, use the clusters to define a new graph, recursively iterate this procedure until the graph size falls below some threshold, and then successively optimize these reduced size graphs. It is a common technique and has been used by several authors in various ways—for example, in a multilevel way analogous to multigrid techniques, [2, 7], and in an adaptive way analogous to dynamic refinement techniques, [18]. Several algorithms for carrying out the reduction can be found in [9].

Reduction. To create a coarser graph $G'(V', E')$ from $G(V, E)$ we use a variant of the edge contraction algorithm proposed by Hendrickson and Leland [7], and improved by Karypis and Kumar in [9]. The idea is to find a maximal independent subset of graph edges and then collapse them. The set is independent because no two edges in the set are incident on the same vertex (so no two edges in the set are adjacent) and maximal because no more edges can be added to the set without breaking the independence criterion. Having found such a set, each selected edge is collapsed and the vertices, $u_1, u_2 \in V$ say, at either end of it are merged to form a new vertex $v \in V'$ with weight $|v| = |u_1| + |u_2|$. Edges which have not been collapsed are inherited by the reduced graph and, where they become duplicated, are merged with their weight summed. This occurs if, for example, the edges (u_1, u_3) and (u_2, u_3) exist when edge (u_1, u_2) is collapsed. Because of the inheritance properties of this algorithm, it is easy to see that the total graph weight remains the same, $|V| = |V'|$, and the total edge weight is reduced by an amount equal to the weight of the collapsed edges.

Parallel Matching. A simple way to construct a maximal independent subset of edges is to visit the vertices of the graph in a random order and pair up or match unmatched vertices with an unmatched neighbor. It has been shown, [9], that it can be beneficial to the optimization to collapse the most heavily weighted edges and our matching algorithm uses this heuristic. For the parallel version we use more or less the same procedure: each processor visiting in parallel the vertices that it owns. We modify the matching algorithm, however, by always matching with a local vertex in preference to a vertex owned by another processor. The local matching can then take place entirely in parallel but usually leaves a few boundary vertices who have no unmatched local neighbors but possibly some unmatched nonlocal neighbors.

The simplest solution would be to terminate the matching at this point. However, in the worst-case scenario if the initial partition is particularly bad and most vertices have no local neighbors (for example a random partition), little or no matching may have taken place. We therefore continue the matching with a parallel iterative procedure which finishes only when there are no vertices unmatched. Vertices which are matched across interprocessor boundaries are migrated to one of the two owning processors and then the construction

of the reduced graph can take place entirely in parallel. The algorithm is fully described in [16].

4. EXPERIMENTAL RESULTS

The software tool which we have used to test the optimization and graph reduction algorithms is known as JOSTLE and is available from <http://www.gre.ac.uk/~c.walshaw/jostle>. For the purposes of this paper it is run in three configurations, dynamic (JOSTLE-D), multilevel-dynamic (JOSTLE-MD) and multilevel-static (JOSTLE-MS). The dynamic configuration, JOSTLE-D, reads in an existing partition and uses the algorithm described in Section 2 to balance and optimize the partition. The multilevel-dynamic, JOSTLE-MD, uses the same procedure but additionally uses graph reduction (Sect. 3) down to a threshold of 20 vertices per processor to improve the partition quality. The static version, JOSTLE-MS, carries out graph reduction on the unpartitioned graph to reduce the graph to $|P|$ vertices which are randomly assigned to partitions and then uses the algorithm described in Section 2 to optimize each of the multilevel graphs.

The test meshes have been taken from an example contained in the DIME (distributed irregular mesh environment) software package freely available by anonymous ftp. The particular application solves Laplace's equation with Dirichlet boundary conditions on a square domain with an S-shaped hole and using a triangular finite element discretization. The problem is repeatedly solved by Jacobi iteration, refined based on this solution, and then load balanced. A very similar set of meshes has previously been used for testing mesh partitioning algorithms and details about the solver, the domain, and DIME can be found in [19]. The particular series of ten meshes and the resulting graphs that we used range in size from the first one which contains 23,787 vertices and 35,281 edges to the final one which contains 224,843 vertices and 336,024 edges. We were unable to run the adaptivity and solver in parallel, but the serial runtime a Sun SPARC Ultra with a 140 MHz CPU was about 12,000 s. The example includes only mesh refinement and not mesh coarsening, but the same algorithms have been successfully tested on time-dependent problems with both coarsening and refinement although the results are commercially sensitive and cannot be used here. In addition, dynamic load-balancing results on a fixed mesh with variable processing resources can be found in [1].

4.1. Comparison Results

In order to demonstrate the quality of the partitions we have compared the method with three of the most popular partitioning schemes, METIS, GREEDY, and multilevel recursive spectral bisection (MRSB). Of the three METIS is the most similar to JOSTLE, employing a graph reduction technique and iterative optimization. The version used here is `kmetis` from the most recent public distribution available by anonymous ftp, [9]. The GREEDY algorithm, [6], is actually performed as

TABLE I
Average Results over the 9 Meshes

Method	$ P = 16$			$ P = 32$			$ P = 64$		
	$ E_c $	$t(s)$	$M\%$	$ E_c $	$t(s)$	$M\%$	$ E_c $	$t(s)$	$M\%$
JOSTLE-D	942	0.51	0.54	1551	0.64	1.80	2598	0.85	3.76
JOSTLE-MD	846	2.39	4.92	1447	2.60	6.26	2410	3.02	8.82
JOSTLE-MS	879	3.96	93.96	1488	4.19	92.77	2417	4.95	99.00
METIS	913	4.83	94.36	1543	4.91	95.94	2427	5.15	97.95
MRSB	939	55.85	83.54	1577	71.42	90.01	2520	87.34	95.07
GREEDY	1816	0.77	81.62	2897	0.83	90.64	4300	1.00	94.42

part of the JOSTLE-MS configuration and is fast but not particularly good at minimizing $|E_c|$. MRSB, on the other hand, is a highly sophisticated method, good at minimizing $|E_c|$ but suffering from relatively high runtimes, [2]. The MRSB code was made available to us by one of its authors, Horst Simon, and run unchanged with a contraction thresholds of 100.

The following experiments were carried out in serial on a Sun SPARC Ultra with a 140 MHz CPU and 64 Mbytes of memory. We use three metrics to measure the performance of the algorithms—the total weight of cut edges, $|E_c|$, the execution time in seconds of each algorithm, $t(s)$, and the percentage of vertices which need to be migrated, M .

For the two dynamic configurations, the initial mesh is partitioned with the static version—JOSTLE-MS. Subsequently at each refinement, the existing partition is interpolated onto the new mesh using the techniques described in [18] and the new partition is then optimized and balanced. For the experiments reported here, since there is only refinement and no coarsening, the mapping of the new partition to the processors is the canonical one and new elements are assigned to the processor to which their parent was assigned while existing elements are simply assigned to the same partition as before; for more complicated situations the use of a “smart” mapping such as that employed by JOVE might be beneficial, [13].

Table I compares the six different partitioning methods for $|P| = 16, 32,$ and 64 with the results averaged over the last 9 meshes (i.e., not including the static partitioning results for the first mesh). The high quality partitioners—both JOSTLE multilevel configurations, METIS, and MRSB—all give similar values for $|E_c|$ with MRSB giving marginally the worst results and JOSTLE-MD giving the best. It is not clear why JOSTLE-MS gives slightly worse results than JOSTLE-MD, possibly this is as a result of using parallel graph reduction rather than serial. In general, JOSTLE-D, without the benefit of graph reduction, provides slightly lower quality partitions but approximately equivalent to those of MRSB. In terms of execution time, JOSTLE-D is slightly faster than GREEDY with both of them being much faster than any of the multilevel algorithms. Of these multilevel algorithms, however, JOSTLE-MD is considerably faster than JOSTLE-

MS and METIS, and MRSB is by far the slowest. It is the final column which is perhaps the most telling though. Because the static partitioners take no account of the existing distribution they result in a vast amount of data migration. The dynamic configurations, JOSTLE-D and JOSTLE-MD, on the other hand, migrate very few of the vertices. As could be expected JOSTLE-MD migrates somewhat more than JOSTLE-D since it does a more thorough optimization.

Taking the results as a whole, the multilevel-dynamic configuration, JOSTLE-MD, provides the best partitions very rapidly and with very little vertex migration. If a slight degradation in partition quality can be tolerated, the JOSTLE-D configuration load balances and optimizes even more rapidly, faster than the GREEDY algorithm, with even less vertex migration. However, JOSTLE-D is essentially only a local optimization method and for robustness JOSTLE-MD is to be preferred. Possibly the ideal solution (which we have not tested) would be a combination of the two—using JOSTLE-D most of the time and JOSTLE-MD occasionally both on a regular basis or if the mesh is known to have changed a great deal.

With regard to the load balancing, all the algorithms resulted in partitions with less than 1% load imbalance (i.e., $\max|S_p|/W < 1.01$), with the direct partitioners GREEDY and MRSB giving exact balance. The dynamic JOSTLE configurations, JOSTLE-D and JOSTLE-MD, started their partitioning on unbalanced partitions with varying degrees of imbalance, which averaged out to about 7% for $|P| = 16$, 12% for $|P| = 32$, and 17% for $|P| = 64$. This is not an extreme example of imbalance but is still enough to badly slow down the computation, especially in the light of the computational runtime relative to the costs of repartitioning.

4.2. Parallel Timings

Achieving high parallel performance for parallel partitioning codes such as JOSTLE is not as easy as, say, a typical CFD or CM code. For a start the algorithms use only integer operations and so there are no MFlops to “hide behind.” In addition, most of the work is carried out on the subdomain boundaries and so very little of the actual graph is used. Also the partitioner itself may not necessarily be well load balanced and the communi-

TABLE II
Serial and Parallel Timings for the JOSTLE-D and JOSTLE-MD Configurations

V	E	P = 16			P = 32			P = 64		
		$t_s(s)$	$t_p(s)$	speed up	$t_s(s)$	$t_p(s)$	speed up	$t_s(s)$	$t_p(s)$	speed up
JOSTLE-D										
31172	46309	0.26	0.06	4.33	0.32	0.06	5.33	0.42	0.07	6.00
40851	60753	0.34	0.07	4.86	0.44	0.07	6.29	0.59	0.08	7.37
53338	79415	0.38	0.06	6.33	0.55	0.08	6.88	0.78	0.11	7.09
69813	104034	0.56	0.10	5.60	0.71	0.09	7.89	0.85	0.13	6.54
88743	132329	0.71	0.13	5.46	0.82	0.10	8.20	1.00	0.09	11.11
115110	171782	0.82	0.11	7.45	1.03	0.11	9.36	1.30	0.11	11.82
146014	218014	1.14	0.16	7.12	1.29	0.13	9.92	1.60	0.13	12.31
185761	277510	1.47	0.21	7.00	1.58	0.15	10.53	2.06	0.16	12.88
224843	336024	1.63	0.19	8.58	1.97	0.18	10.94	2.30	0.14	16.43
JOSTLE-MD										
31172	46309	0.93	0.35	2.66	1.12	0.26	4.31	1.42	0.25	5.68
40851	60753	1.25	0.40	3.12	1.50	0.32	4.69	1.99	0.32	6.22
53338	79415	1.53	0.97	1.58	1.73	0.30	5.77	2.25	0.32	7.03
69813	104034	1.99	0.48	4.15	2.22	0.32	6.94	2.73	0.33	8.27
88743	132329	2.44	0.49	4.98	2.83	0.40	7.08	3.34	0.38	8.79
115110	171782	3.15	0.61	5.16	3.51	0.44	7.98	4.13	0.39	10.59
146014	218014	3.98	0.75	5.31	4.58	0.56	8.18	5.39	0.55	9.80
185761	277510	5.03	0.87	5.78	5.50	0.63	8.73	6.45	0.55	11.73
224843	336024	6.04	0.95	6.36	6.66	0.67	9.94	7.66	0.59	12.98

cations cost may dominate on the coarsest reduced graphs since at this stage there are very few vertices per processor. On the other hand, as was explained in Section 1, partitioning on the host may be impossible or at least much more expensive and if the cost of load balancing is regarded (as it should be) as a parallel overhead, it is usually extremely inexpensive relative to the overall solution time of the problem. Indeed if the reverse is true, it may not be worth load balancing at all, e.g., [1].

Table II gives serial and parallel timings for the JOSTLE-D and JOSTLE-MD configurations on the 512 node Cray T3E at HLRS, the High Performance Computer Centre at the University of Stuttgart. The parallel version uses the MPI communications library although we are working on a shmem version which could be expected to show even faster timings. The parallel timings generally decrease as $|P|$ increases although this is not so true on the smaller meshes and especially for JOSTLE-D. We believe that this is because there is so little computational work that these figures just show parallel communication overhead; dividing the serial time by $|P|$ to estimate the parallel computational work suggests that this is the case. However, these figures show good speedups for this sort of code and more importantly, very low overheads

(always less than a second) for the parallel partitioning. Finally note that the partitions obtained for the parallel version of JOSTLE are exactly the same as those of the serial version.

5. CONCLUSION

We have described a new method for optimizing and load balancing graph partitions with a specific focus on its application to the dynamic mapping of unstructured meshes onto parallel computers. In this context the graph-partitioning task can be very efficiently addressed by reoptimizing the existing partition, rather than starting the partitioning from afresh. For the experiments reported in this paper, which are somewhat limited in that they only involve adaptive mesh refinement and not coarsening, the dynamic procedures are much faster than static techniques, provide partitions of similar or higher quality and, in comparison, involve the migration of a fraction of the data.

ACKNOWLEDGMENTS

We thank HLRS, the High Performance Computer Centre at the University of Stuttgart, for access to the Cray T3E.

REFERENCES

1. A. Arulananthan, S. Johnson, K. McManus, C. Walshaw, and M. Cross, A generic strategy for dynamic load balancing of distributed memory parallel computational mechanics using unstructured meshes. *Proc. Parallel CFD '97*, in press.
2. S. T. Barnard and H. D. Simon, A fast multilevel implementation of recursive spectral bisection for partitioning unstructured problems. *Concurrency: Practice and Experience* **6**, 2 (1994), 101–117.
3. G. Cybenko, Dynamic load balancing for distributed memory multiprocessors. *J. Parallel Distrib. Comput.* **7**, 2 (1989), 279–301.
4. R. Diekmann, D. Meyer, and B. Monien, Parallel decomposition of unstructured FEM-meshes. In A. Ferreira and J. Rolim (Eds.), *Proc. Irregular '95: Parallel Algorithms for Irregularly Structured Problems*. Springer-Verlag, Berlin/New York, 1995, pp. 199–215.
5. P. Diniz, S. Plimpton, B. Hendrickson, and R. Leland, Parallel algorithms for dynamically partitioning unstructured grids. In D. Bailey *et al.* (Eds.), *Parallel Processing for Scientific Computing*, SIAM, 1995, pp. 615–620.
6. C. Farhat, A simple and efficient automatic FEM domain decomposer. *Comp. Struct.* **28**, 5 (1988), 579–602.
7. B. Hendrickson and R. Leland, A multilevel algorithm for partitioning graphs. In *Proc. Supercomputing '95*, 1995.
8. Y. F. Hu and R. J. Blake, An optimal dynamic load balancing algorithm. Preprint DL-P-95-011, Daresbury Laboratory, Warrington, WA4 4AD, UK, 1995. (*Concurrency: Practice & Experience*, in press.)
9. G. Karypis and V. Kumar, A fast and high quality multilevel scheme for partitioning irregular graphs. TR 95-035, Computer Science Department, University of Minnesota, Minneapolis, MN 55455, 1995.
10. G. Karypis and V. Kumar, A coarse-grain parallel formulation of multilevel k -way graph partitioning algorithm. In M. Heath *et al.* (Eds.), *Parallel Processing for Scientific Computing*. SIAM, Philadelphia, 1997.
11. R. Lohner, R. Ramamurti, and D. Martin, A parallelizable load balancing algorithm. AIAA-93-0061, American Institute of Aeronautics and Astronautics, Washington, DC, 1993.
12. H. Simon, A. Sohn, and R. Biswas, HARP: A fast spectral partitioner. In *Proc. 9th ACM Symposium on Parallel Algorithms and Architectures*, 1997, pp. 43–52.
13. A. Sohn, R. Biswas, and H. Simon, Impact of load balancing on unstructured adaptive grid computations for distributed-memory multiprocessors. In *Proc. 8th IEEE Symposium on Parallel and Distributed Processing*, 1996, pp. 26–33.
14. D. Vanderstraeten and R. Keunings, Optimized partitioning of unstructured computational grids. *Internat. J. Numer. Methods Engrg.* **38** (1995), 433–450.
15. C. Walshaw, M. Cross, and M. Everett, Dynamic mesh partitioning: A unified optimisation and load-balancing algorithm. Tech. Rep. 95/IM/06, University of Greenwich, London SE18 6PF, UK, 1995.
16. C. Walshaw, M. Cross, and M. Everett, Parallel dynamic graph-partitioning for unstructured meshes. Tech. Rep. 97/IM/20, University of Greenwich, London SE18 6PF, UK, March 1997.
17. C. Walshaw, M. Cross, M. Everett, S. Johnson, and K. McManus, Partitioning and mapping of unstructured meshes to parallel machine topologies. In A. Ferreira and J. Rolim (Eds.), *Proc. Irregular '95: Parallel Algorithms for Irregularly Structured Problems*, Lecture Notes in Computer Science, Vol. 980, Springer-Verlag, Berlin/New York, 1995, pp. 121–126.
18. C. H. Walshaw and M. Berzins, Dynamic load-balancing for PDE solvers on adaptive unstructured meshes. *Concurrency: Practice and Experience* **7**, 1 (1995), 17–28.
19. R. D. Williams, Performance of dynamic load balancing algorithms for unstructured mesh calculations. *Concurrency: Practice and Experience* **3** (1991), 457–481.

CHRIS WALSHAW is a Research Fellow in the School of Computing and Mathematical Sciences at the University of Greenwich. He graduated from Bath University with a B.Sc. in mathematics and then moved to Edinburgh where he gained an M.Sc. from Edinburgh University and a Ph.D. from Heriot-Watt University where his doctoral thesis concerned parallel algorithms for systems of differential equations. His postdoctoral work has extended this theme into parallel methods for adaptive unstructured meshes and, in particular, mesh partitioning. Since joining the University of Greenwich in 1993 he has developed the publicly-available JOSTLE mesh partitioning software. He is the author of some 25 research papers.

MARK CROSS is Professor of Numerical Modelling and Director of the Centre for Numerical Modelling and Process Analysis in the School of Computing and Mathematical Sciences at the University of Greenwich. The Centre has about 100 staff and graduate students of which about 10 are associated with the Parallel Processing Group whose work is focused on the development of software tools to support the exploitation of such systems by computational modeling software. Professor Cross was educated at the University of Wales, Cardiff, and received a Ph.D. in 1972 for work on the modeling of semiconductor lasers. Since then he has worked in industry and academia in both the UK and USA and has been at Greenwich since 1982. His research interests cover computational modeling of metals/materials processes, computational mechanics algorithms and software tools, and the exploitation of HPC systems. The editor of the archival journal *Applied Mathematical Modelling*, published by Elsevier, he is the author of some 200 research publications.

MARTIN EVERETT is Professor of Applicable Mathematics and Head of the School of Computing and Mathematical Sciences at the University of Greenwich. He graduated from Loughborough University with a B.Sc. in Mathematics with first class honors and then gained an M.Sc. in mathematics and a D.Phil. degree both from Oxford University in 1981. His graduate work involved numerical analysis and applications of graph theory. His doctoral thesis focused on the use of graph theory in the analysis of social networks. Since coming to the University of Greenwich in 1980, Professor Everett has both maintained a research program in social network analysis and extended his interests to other areas of graph theory applications; one of these is mesh partitioning for parallel computing. A visiting professor at U.C. Irvine, he is the author of over 60 research papers.