Using Multithreading for the Automatic Load Balancing of Adaptive Finite Element Meshes

Gerd Heber[1], Rupak Biswas[2], Parimala Thulasiraman[1], and Guang R. Gao[1]

[1] CAPSL, University of Delaware, 140 Evans Hall, Newark, DE 19716, USA
[2] MRJ, NASA Ames Research Center, MS T27A-1, Moffett Field, CA 94035, USA

Abstract. In this paper, we present a multithreaded approach for the automatic load balancing of adaptive finite element (FE) meshes. The platform of our choice is the EARTH multithreaded system which offers sufficient capabilities to tackle this problem. We implement the adaption phase of FE applications on triangular meshes, and exploit the EARTH token mechanism to automatically balance the resulting irregular and highly nonuniform workload. We discuss the results of our experiments on EARTH-SP2, an implementation of EARTH on the IBM SP2, with different load balancing strategies that are built into the runtime system.

1 Introduction

In this paper, we present and examine a multithreaded approach for the automatic load balancing of adaptive finite element (FE) meshes. During a FE adaption phase, the unstructured mesh is refined/coarsened (according to some application-specific criteria), and the workload may become (seriously) unbalanced on a multiprocessor system. This significantly affects the overall efficiency of parallel adaptive FE calculations. Some of the difficulties encountered when using the traditional approach to resolve the load imbalance problem can be summarized as follows:

- It is necessary to assemble global information to get an accurate estimate of the load distribution – a step that often creates serious bottlenecks in practice.
- Significant effort is usually required to preserve data locality while making the load balancing decisions – an optimization which is computationally difficult.
- To make matters worse, the evolution of computational load and data locality requirements are dynamic, irregular, and unpredictable at compile time.

Multithreaded architectures, such as the EARTH (*Efficient Architecture for Running THreads*) system [8], offer new capabilities and opportunities to tackle this problem. They strive to hide long latency operations by overlapping computation and communication with the help of *threads*. A thread is a (small) sequence of instructions. EARTH provides mechanisms to enable automatic load balancing for applications that do not allow a good static (compile time) task distribution. The programmer can simply encapsulate a threaded function invocation as a *token*. The advantage of this token mechanism is that they can flexibly migrate over processors, by means of a load balancer.

Given our experience with EARTH, we decided to pursue a novel approach to implement a two-dimensional version of the dynamic mesh adaption procedure, called 3D_TAG [1]. Our decisions were based on the following observations:

- EARTH provides a runtime load balancing mechanism at a very fine-grain level, i.e., the token can be an ultra light-weight function invocation. Therefore, as long as the computation generates a large number of such tokens, it is likely that the EARTH runtime system can automatically distribute them to the processors to keep them usefully busy. It is our hypothesis that we can thus eliminate an explicit repartitioning/remapping phase (in the first order) and the gathering of global mesh information.
- Locality is always an important issue for efficiency. However, since EARTH has the ability to tolerate latency using multithreading, its performance is less sensitive to the amount of non-local accesses by each processor. In other words, the optimality of locality or data distribution is less critical. Instead of optimally trying to repartition/remap the data to minimize communication, the user should take advantage of the EARTH token mechanism and structure the algorithm/code such that the input arguments to a token can be migrated easily with the token itself. The token should thus have good *mobility*.
- The difficulty of handling irregular and unpredictable load evolution in traditional systems is partly due to the limitations of the data parallel programming model and the SPMD execution paradigm where the computation in each processor should be loosely synchronous at a coarse-grain level. The EARTH execution model is fundamentally asynchronous and does not rely on the SPMD model. Function-level parallelism can be naturally expressed at a fine grain in the EARTH programming paradigm. The possible association of a parallel function invocation with a token at such a fine grain enables the runtime load balancing mechanism to work smoothly with the evolving computational load.

We implement the adaption phase of a simulated FE calculation and exploit the EARTH token mechanism to automatically balance the highly nonuniform workload. To achieve good token mobility, we apply a new indexing methodology for FE meshes that does not assume a particular architecture or programming environment. During mesh adaption, tokens that migrate to another processor to be executed leave the processed data in that processor instead of transferring the data back to the processor that generated the token. This is a novel approach in solving this problem: we avoid doing work twice and do a load-driven remapping without explicit user intervention. A major part of this paper is devoted to a discussion of the *quality* of such an approach which can be achieved using different load balancing strategies (built into the EARTH runtime system).

2 The EARTH System

EARTH (*Efficient Architecture for Running THreads*) [8] supports a multithreaded program execution model in which user code is divided into threads that are scheduled atomically using dataflow-like synchronization operations. These "EARTH operations" comprise a rich set of primitives, including remote loads and stores, synchronization operations, block data transfers, remote function

calls, and dynamic load balancing. EARTH operations are initiated by the threads themselves. Once a thread is started, it runs to completion, and instructions within it are executed sequentially.[1] Therefore, a conventional processor can execute a thread efficiently even when the thread is purely sequential. It is thus possible to obtain single-node performance close to that of a purely sequential implementation, as was shown in our earlier work [10].

Conceptually, each EARTH node consists of an *Execution Unit* (EU), which executes the threads and a *Synchronization Unit* (SU), which performs the EARTH operations requested by the threads (cf. Fig. 1). The EARTH runtime system is presently available on the MANNA architecture [7], the IBM SP2 multiprocessor [3], and the Beowulf workstation clusters.

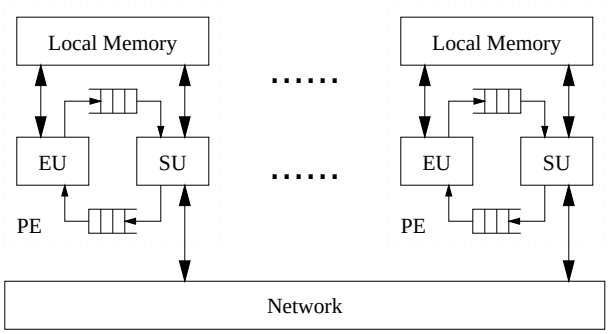

Fig. 1. The EARTH architecture

Currently, programs running on EARTH are written in Threaded-C [13], a C extension containing multithreading instructions. It is clean and powerful enough to be used as a user-level, explicitly parallel programming language. The keyword THREADED specifies that the function's frame should be allocated from the heap and that the function may contain multiple threads.[2] Each thread, except for the first, begins with a label THREAD_*n*; the first thread begins at the start of the function, and is automatically executed when the function is invoked using an INVOKE or TOKEN instruction. Each thread ends with either END_THREAD(), which simply executes a **fetch_next** instruction, or END_FUNCTION(), which executes an **end_function** instruction to deallocate the frame before executing a **fetch_next**.

In addition, for those applications for which a good task distribution cannot be determined statically by the programmer and communicated to the compiler, EARTH provides an automatic load balancing mechanism. The programmer can simply encapsulate a function invocation as a token. The token is sent to the SU, which puts it on top of the local token queue. When there are no more threads in the ready queue, the SU removes a token from the top of the token queue and

[1] Instructions may be executed out of order, as on a superscalar machine, as long as the semantics of the sequential ordering are obeyed.

[2] Non-threaded functions are specified as regular C functions.

invokes the function as specified locally. This load balancing technique is derived from the token management strategy used in the ADAM architecture [9]. Note that putting locally-generated tokens on top of the queue and then removing tokens from the top leads to a *depth-first* traversal of the call-graph. This generally leads to better control of functional parallelism, i.e., it diminishes the likelihood of parallelism explosion that can exhaust the memory resources of a node.

When both the ready queue and the token queue are empty, the SU sends a message to a neighboring processor requesting a token, in effect performing *work stealing*. The neighboring processor, if it has tokens in its queue, extracts one from the bottom of its queue and sends it back to the requester. In this manner, a *breadth-first* traversal of the call-graph is implemented across processors, hopefully resulting in a better distribution of tasks. If the neighboring processor does not have any tokens to satisfy a request, the neighbor's neighbor is queried, and so on. Either a token will be found, or the request cannot be fulfilled. Idle processors periodically query their neighbors for work.

3 An Indexing Technique for FE Meshes

Dynamic mesh adaption is a common and powerful technique for efficiently solving FE problems, and is achieved by *coarsening* and/or *refining* some elements of the computational mesh. In 2D FE meshes, triangular elements are the most popular. Once a triangle is targeted for refinement, it is generally subdivided by bisecting its three edges to form four congruent triangles as shown in Fig. 2(a). This type of subdivision is called *isotropic* (or 1:4) and the resulting triangles are referred to as being *red*. A problem is that the refined mesh will be *nonconforming* unless all the triangles are isotropically subdivided. To get a consistent triangulation without global refinement, a second type of subdivision is allowed. For example, a triangle may be subdivided into two smaller triangles by bisecting only one edge as shown in Fig. 2(b). This type of subdivision is called *anisotropic* (or 1:2 in this case) and the resulting triangles are referred to as being *green*. The process of creating a consistent triangulation is defined as a *closure* operation. Note that several iterations may be necessary to achieve closure [1].

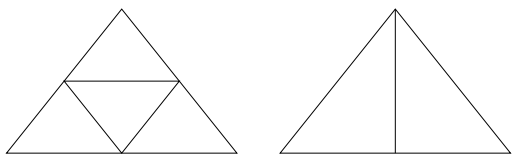

Fig. 2. (a) Isotropic and (b) anisotropic subdivision of a triangle

A couple of additional rules are applied, primarily to assure that the quality of the adapted mesh does not deteriorate drastically with repeated refinement:

1. All triangles with exactly two bisected edges have their third edge also bisected. Thus, such triangles are isotropically refined.

2. A green triangle cannot be further subdivided. Instead, the previous subdivision is discarded and isotropic subdivision is first applied to the (red) ancestor triangle [2].

It is the task of an *index scheme* to properly name or label the various objects (vertices, edges, triangles) of a mesh. We prefer the term index scheme instead of *numbering* to stress that the use of natural numbers as indices is not sufficient to meet the naming requirements of the FE objects on parallel architectures.

We give a brief description of our indexing scheme for the sake of completeness; refer to [6] for a detailed discussion. Our technique is a combination of *coarse* and *local* schemes, and is intended for hierarchical meshes. The coarse scheme labels the initial mesh objects in such a way that the incidence relations can be easily derived from the labels. For the example in Fig. 3, the set of vertices for the coarse triangulation consists of the following numbers:

$$\text{vertices} = \{1,\ 2,\ 3,\ 4,\ 5,\ 6,\ 7,\ 8\}.$$

The edges of the coarse triangulation are indexed by ordered pairs of integers corresponding to the endpoints of the edges. The ordering is chosen so that the first index is less than the second:

$$\begin{aligned}\text{edges} = \{&(1,2),\ (1,3),\ (1,5),\ (2,5),\ (3,4),\ (3,5),\ (3,6),\\ &(4,6),\ (5,6),\ (5,7),\ (5,8),\ (6,7),\ (7,8)\}.\end{aligned}$$

The same principles are applied to index the coarse triangles. They are denoted by the triple consisting of their vertex numbers in ascending order:

$$\text{triangles} = \{(1,2,5),\ (1,3,5),\ (3,4,6),\ (3,5,6),\ (5,6,7),\ (5,7,8)\}.$$

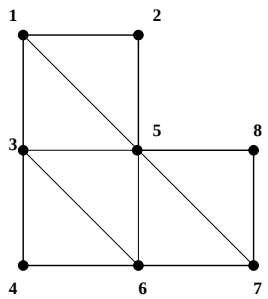

Fig. 3. An L-shaped domain and its coarse triangulation

The local scheme exploits the regularity (and the finiteness) of the refinement rules to produce names for the objects at subsequent refinement levels [6]. We use (scaled) natural coordinates as indices in the local scheme. Again, this is done in a way such that the incidence relations are *encoded* in the indices of the objects. For example, the set of vertices at level k in the local model is given by:

$$V_k = \{(a,b,c) \in \mathrm{N}^3 \mid a+b+c = 2^k\}.$$

We do not know which of these will actually be present; however, we already have names for them. Figure 4 shows the local indices for the vertices and the triangles on the first two refinement levels (for isotropic subdivision). The coarse and local schemes are combined by taking the union of the Cartesian products of the coarse mesh objects with their corresponding local schemes. Ambiguities are resolved by using a *normal* form of the index.

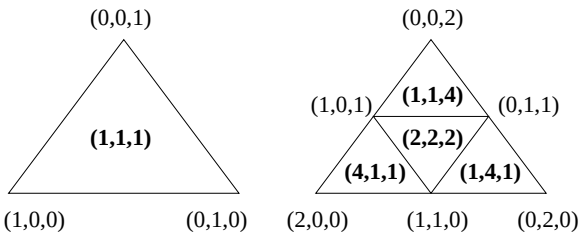

Fig. 4. Examples of the local index scheme for triangular elements (the vertices and triangles are denoted by integer triples, triangles in bold)

The key features of such a scheme are:

- Each object is assigned a *global name* that is independent of any architectural considerations or implementation choices.
- Combinatorial information is translated into simple arithmetic.
- It is *well-behaved* under (adaptive) refinement. No artificial synchronization/serialization is introduced.
- It can be extended (with appropriate modifications) to three dimensions [6].

4 Description of the Test Problems

4.1 Mesh, Partitioning, and Mapping

As the initial mesh for our experiments, we choose a rectangular mesh similar to, but much larger than, the one shown in Fig. 5. Mesh generation simplicity was an overriding concern. However, we never, neither in the algorithms nor in the implementation, exploit this simple connectivity. Recall that geometry information is irrelevant for our index scheme. Mesh adaption on even a simple case like this can fully exercise several aspects of the EARTH multithreaded system and its dynamic load balancing capabilities.

Fine-grain multithreading would allow us to have one thread per element. However, to control granularity and to prevent *fragmentation*[3] of the triangulation, we choose partitions as our smallest migratable unit. A good overview of common partitioning techniques and mapping strategies is given in [4]. Based on observations made there, we decided to use a $P \times Q$ block partitioning.

Initially, the entire mesh resides on node 0. A marking procedure (an *error estimator/indicator* in real applications) marks edges/elements for refinement or

[3] The EARTH runtime system does not support a restricted migration of tokens.

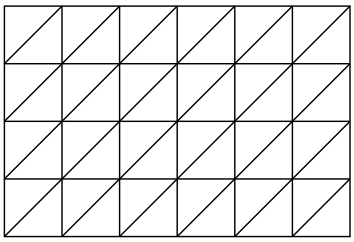

Fig. 5. A small sample of our test mesh

coarsening (on higher levels). Based on this marking, tokens (one for each partition) are generated that perform the actual refinement. This situation immediately serves as a *stress* test for the EARTH token mechanism to automatically balance the highly nonuniform workload. If a token migrates to and is executed on another processor, it transfers the partition and refinement information to the executing node. The actual refinement/coarsening is done there. It is thus absolutely crucial to *compress* the *metacontext* as much as possible. The property of our index scheme to provide global indices becomes extremely useful here: it is not necessary to transform names if processor/memory boundaries are crossed. Furthermore, only the indices of the elements have to be transferred since the indices of all other objects (edges, vertices) can be easily recomputed [6]. (This can be further optimized since a parent element can, in conjunction with the refinement information, rapidly calculate the indices of its children.)

For mesh refinement, one can distinguish between two strategies. In the *first* case, parent elements are deleted from the data structure when they are refined. In the *second* case, the refinement history is maintained. The advantage of the second strategy is that parent objects do not have to be reconstructed from scratch during mesh coarsening. We have implemented this second strategy [1].

When all elements in a FE mesh are subdivided isotropically, it is referred to as *uniform* refinement and automatically leads to a conforming triangulation. In the case of *nonuniform* refinement, a (iterative) consistency procedure is required since adaptive coarsening/refinement generally leads to inconsistent triangulations. Moreover, to guarantee mesh quality using our adaption algorithm, the parents of green triangles are reinstated before further refinement [2].

In real applications, mesh adaption is driven by an error indicator that is obtained either geometrically or from the numerical solution. The computational mesh is refined in regions where the indicator is high, and coarsened in regions where it is low. In our experiments, elements are randomly selected for adaption. The behavior of an error indicator can generally be mimicked by suitably adjusting the distribution pattern of the random numbers.

4.2 Load Balancing Strategies

For a detailed description of the load balancing strategies that have been implemented in the EARTH-SP[4] runtime system, we refer the reader to [3]. This is

[4] This refers to the version of Threaded-C available for the IBM SP2 and SP3.

a comprehensive study that provides an excellent overview of the implemented strategies and their qualitative behavior under several stress tests.

Some questions we would like to answer for our particular simulation are:

- What portion of the total number of elements are finally located on nodes other than node 0? How do the numbers of nodes and tokens generated per node affect this behavior? Is the load acceptably balanced outside of node 0?
- How does the token generation process affect the runtime? Is there a measurable system overhead?
- Does an increased number of tokens decrease the *variance* of the load balancer (as one would expect)?
- How does the system behave if there are less tokens than processors?

5 Results

The data structure used to keep the indices of our FE mesh objects is a red-black tree [5], as it is used in the implementation of set container in C++ STL. The code for the uniform refinement case (including that for the red-black tree) consists of about 2000 lines of Threaded-C. In the nonuniform case, closure of the triangulation process is added, and the code is about 2700 lines.

The program contains certain optimizations for the case when a token is executed on the node it originates. Executing the program on one node implies that all tokens are both generated and executed on the same node. In that case, no remote operations or tree compression/expansion are necessary.

The initial mesh is the same for all the experiments and consists of 2048 vertices, 5953 edges, and 3906 triangular elements. At the beginning of each experiment, the entire mesh resides on node 0. It is the responsibility of the EARTH runtime system to scatter partitions across the available nodes. Given the space limitations, we analyze only the uniform refinement case in this paper. A total of four mesh refinement steps were performed. The final mesh contained 999,936 triangles since each triangle was refined isotropically. All measurements were done upon completion of these four steps.

Currently, there are eight different load balancers available with EARTH-SP, and we tested all of them. As a representative for the discussions in this section, we chose the DUAL load balancer. It is one of the simplest load balancers (a *virtual ring* [3]) and does not provide the best quantitative results; however, its general behavior is comparable with the other load balancers in EARTH-SP.

Discussion of Table 1: Runtime versus the number of tokens generated per node.

- Even on one node, the program runs faster as the number of tokens increases! Recall that the number of tokens corresponds to the number of partitions. Hence, for a fixed mesh size, more partitions imply less elements per partition. Since the elements in a partition are organized in a tree, it follows that the trees will be shallower. This, in turn, accelerates the tree operations. The entries for one node in the table confirm that the underlying data structure must be of logarithmic complexity.

Table 1. Execution time (in seconds) versus the number of tokens generated per node

	Tokens									
Nodes	4	8	16	32	64	128	256	512	1024	2048
1	20.78	19.77	18.73	17.74	16.74	15.68	14.91	14.39	13.71	12.96
2	18.13	14.85	13.79	12.37	11.53	11.92	10.91	10.89	10.38	9.94
3	11.50	11.97	9.72	9.44	9.14	7.90	8.54	8.03	7.98	8.06
4	19.48	10.25	7.36	7.99	7.08	6.82	7.25	6.06	6.35	6.87
5	11.30	10.12	8.06	6.69	5.81	5.47	6.24	5.40	5.92	5.35
6	11.96	10.56	10.13	6.94	5.90	5.40	5.84	5.78	5.36	4.89
7	26.83	9.97	5.25	6.89	5.20	5.26	5.49	4.94	4.69	5.07
8	22.62	10.76	7.91	5.91	5.55	5.44	5.05	5.15	5.00	5.04

- Reasonable speedup is obtained for up to five nodes. However, we must consider the additional quality information (discussed below) for our evaluation, since quality has rather strong implications for other parts (solver/preconditioner) of a complete FE application. It might therefore be useful to invest some more resources that do not necessarily increase the speedup but improve quality.
- A certain instability is observed in the region where the number of tokens is not much greater than the number of processors.

Table 2. Migration (as a percentage of the total number of triangles) from node 0 versus the number of tokens generated per node

	Tokens									
Nodes	4	8	16	32	64	128	256	512	1024	2048
2	24	37	37	31	32	25	27	25	25	24
3	50	49	61	50	53	51	44	45	45	41
4	52	75	61	58	60	58	52	58	58	49
5	100	88	68	75	70	66	59	63	58	60
6	75	100	81	75	71	68	62	61	66	65
7	75	88	75	80	75	68	68	68	68	73
8	100	100	75	84	68	67	71	66	66	69

Discussion of Table 2: Migration from node 0 versus the number of tokens generated per node.

- Migration is measured as the percentage of all triangles that are, at the end, owned by nodes other than node 0. This measure is somewhat imprecise since a partition may migrate back and forth.[5]
- Migration increases with the number of nodes, but might decrease with the number of tokens if not enough processors are available. This is probably because the DUAL load balancer is unable to handle the load within these resource limitations.

[5] Such fluctuations occur especially if there are only a few processors and/or tokens.

- It appears that there is a region of stabilization at about 70 percent of migration. Asymptotically, the migration percentage from node 0 when using P nodes is $100 \times (P-1)/P$. Migration can be increased toward this asymptotic by using more processors or more tokens. Generating more tokens implies downsizing partitions, which would consist of only one triangle in the extreme case. However, we would then inevitably be faced with a fragmentation problem (and a total loss of locality).
- Once again, there is a region of instability when the number of tokens is not much greater than the number of processors.

Table 3. Total variance (as a percentage of the total number of triangles) versus the number of tokens generated per node

	Tokens									
Nodes	4	8	16	32	64	128	256	512	1024	2048
2	26	13	13	19	18	25	23	25	25	26
3	12	24	7	14	16	12	15	15	16	19
4	17	9	10	11	9	10	14	10	10	15
5	10	6	7	4	5	7	10	9	11	10
6	12	9	5	5	7	7	10	10	8	9
7	18	5	7	4	4	8	8	7	8	5
8	13	9	7	4	8	8	7	9	9	7

Table 4. Unbiased variance (as a percentage of the total number of triangles) outside of node 0 versus the number of tokens generated per node. Note that results are not shown for two nodes since node 0 is excluded when computing this variance

	Tokens									
Nodes	4	8	16	32	64	128	256	512	1024	2048
3	1	24	7	8	7	7	5	6	6	5
4	12	10	8	5	2	3	3	4	4	3
5	1	5	5	4	3	2	1	3	3	3
6	12	6	5	4	3	3	3	3	3	4
7	19	5	6	4	3	2	3	2	3	3
8	12	8	6	4	2	3	4	4	4	3

Discussion of Tables 3 and 4: Total variance and unbiased variance outside of node 0 versus the number of tokens generated per node.

- Total variance is defined as the expected deviation from the mean value of the number of triangles per node measured as a percentage of the total number of triangles. Since this measure is biased towards node 0, the variance outside of node 0 is computed to evaluate the uniformity of the migration (cf. Table 4).
- The variances oscillate rapidly if not enough tokens are available. If not enough

processors are available, increasing the number of tokens tends to overload the load balancer.
- A large number of tokens seems to guarantee the stability of, and possibly a slight improvement in, the variances.
- Perhaps the most important result is the quality of load balancing. Experimental values of the unbiased variance presented in Table 4 show that the load is extremely well balanced (between 3 and 6 percent). This is generally acceptable for actual adaptive FE calculations.

5.1 Summary of Results

We summarize by addressing each of the three observations made in Sec. 1:
- The experimental results confirm our hypothesis outlined in the introduction: although we did not have an explicit repartitioning/remapping phase, the processors appear to have a good and balanced utilization as long as there are enough tokens being processed in the system (cf. Tables 2 and 3).
- The experiments also show that the natural indexing method used to enhance token mobility does appear to work; i.e., the remote communication due to token migration does not seem to have a major impact on the overall performance (see the region of good speedup and balanced workload implied in Table 1). Of course, the FE solution phase is usually more sensitive to data locality which is not included in the scope of this paper.
- Finally, the experimental results have shown that the load balancing works smoothly as long as there are enough computational resources available, even when the workload evolution is irregular. This can be observed in Table 4, where we see a stable and small variance in the expected workload (measured by the number of triangles).

6 Conclusions

The ability to dynamically adapt a finite element mesh is a powerful tool for solving computational problems with evolving physical features; however, an efficient parallel implementation is rather difficult. To address this problem, we examined a multithreaded approach for performing the load balancing automatically.

Preliminary results from experiments with 2D triangular meshes using the EARTH multithreaded system are extremely promising. The EARTH token mechanism was used to balance the highly irregular and nonuniform workload. To achieve good token mobility, a new indexing methodology was used. During mesh adaption, tokens that migrate to another processor to be executed leave the processed data in the destination processor. This is a novel approach that allows simultaneous load balancing and load-driven data remapping. Results showed that the load is extremely well balanced (3–6% of ideal) when a sufficient number of tokens is generated (at least 512 tokens per node). Token generation (up to 2048 tokens per node) does not significantly increase system overhead or degrade performance, independent of the number of nodes. Having a large number of tokens decreases the total variance of the load balancer and stabilizes the

uniformity of the migration. Finally, an underloaded system (one that has few tokens per node) causes instability and leads to unpredictable behavior by the load balancer.

These conclusions do not seem very surprising, but qualitatively confirm what was expected. Given the fact that very minimal programming effort (on the user level) was necessary to obtain these results and that this is one of the first multithreaded approaches to tackle unstructured mesh adaption, our findings and observations become extremely valuable. More extensive experiments will be done in the future, and the results compared critically with an explicit message-passing implementation [12] and a global load balancer [11]. One must remember however that dynamic mesh adaption comprises a small though significant part of a complete application. Further investigations are needed to determine whether the functionality of such an approach is viable for real applications.

References

1. Biswas, R., Strawn, R.: A new procedure for dynamic adaption of three-dimensional unstructured grids. Appl. Numer. Math. **13** (1994) 437–452
2. Biswas, R., Strawn, R.: Mesh quality control for multiply-refined tetrahedral grids. Appl. Numer. Math. **20** (1996) 337–348
3. Cai, H., Maquelin, O., Gao, G.: Design and evaluation of dynamic load balancing schemes under a multithreaded execution model. CAPSL Technical Memo TM-05 (1997)
4. Chrisochoides, N., Houstis, E., Rice, J.: Mapping algorithms and software environment for data parallel PDE iterative solvers. J. Par. Dist. Comput. **21** (1994) 75–95
5. Cormen, C., Leiserson, C., Rivest, R.: Introduction to Algorithms. MIT Press (1990)
6. Gerlach, J., Heber, G.: Fundamentals of natural indexing for simplex finite elements in two and three dimensions. RWCP Technical Report 97-008 (1997)
7. Hum, H., Maquelin, O., Theobald, K., Tian, X., Gao, G., Hendren, L.: A study of the EARTH-MANNA multithreaded system. Intl. J. of Par. Prog. **24** (1996) 319–347
8. Hum, H., Theobald, K., Gao, G.: Building multithreaded architectures with off-the-shelf microprocessors. Intl. Par. Proc. Symp. (1994) 288–297
9. Maquelin, O.: Load balancing and resource management in the ADAM machine. Advanced Topics in Dataflow Computing and Multithreading, IEEE Press (1995) 307–323
10. Maquelin, O., Hum, H., Gao, G.: Costs and benefits of multithreading with off-the-shelf RISC processors. Euro-Par'95 Parallel Processing, Springer-Verlag LNCS **966** (1995) 117–128
11. Oliker, L., Biswas, R.: PLUM: Parallel load balancing for adaptive unstructured meshes. NAS Technical Report NAS-97-020 (1997)
12. Oliker, L., Biswas, R., Strawn, R.: Parallel implementation of an adaptive scheme for 3D unstructured grids on the SP2. Parallel Algorithms for Irregularly Structured Problems, Springer-Verlag LNCS **1117** (1996) 35–47
13. Theobald, K., Amaral, J., Heber, G., Maquelin, O., Tang, X., Gao, G.: Overview of the Threaded-C language. CAPSL Technical Memo TM-19 (1998)