

2D Parallel Constrained Delaunay Mesh Generation: A Multigrain Approach on Deep Multiprocessors.

Christos D. Antonopoulos, Nikos Chrisochoides, Dimitrios S. Nikolopoulos

Parallel Constrained Delaunay Mesh (PCDM) is a 2D adaptive and irregular meshing algorithm. In PCDM one can explore concurrency using three different levels of granularity: (i) coarse-grain at the sub-mesh level, (ii) medium-grain at the cavity level and (iii) fine-grain at the element level. The medium- and fine-grain approaches can be used to improve the single-processor performance of coarse-grain PCDM.

The coarsest grain of PCDM has been parallelized using a pure message passing approach, and is best suited for distributed memory clusters. It is characterized by good scalability with respect to the number of nodes. The application has minimal communication requirements between nodes processing neighboring subdomains. The communication is, however, unstructured and thus can not be directly mapped to any physical processor topology. The exchanged messages are typically small (28 bytes), however message aggregation can be applied. Computation is predictably unbalanced for graded-quality meshes. A middleware runtime layer (PREMA) has been successfully used to transparently address load balancing issues. The coarse-grain PCDM implementation (on a single processor) is two to three times slower than the best publicly available 2D sequential Delaunay mesh generation software, Triangle from Berkeley.

The fine-grain PCDM approach targets hardware level multithreading for SMT or CMP processors. In-processor execution contexts can (ideally) be exploited to close the single-processor performance gap between coarse-grain PCDM and Triangle. Each cavity expansion lasts between 4 and 6 usec if performed by a single thread on a modern Intel P4 processor. Fine-grain PCDM is a pointer-chasing code with a very irregular memory access pattern and moderate memory bandwidth requirements (around 96 MBps/thread). Fine-grained synchronization is required between threads simultaneously expanding the same cavity. We find that current commercially available SMTs are not capable of exploiting parallelism at a fine granularity similar to that available in PCDM. A simulated system study indicates that minor, realistic hardware support for efficient in-processor synchronization and hardware thread creation / management allows fine-grain PCDM to outperform both sequential PCDM and the MPI-only coarse grain parallelization on a single SMT processor. Moreover, we find efficient, low-overhead execution context suspend/resume functionality to be necessary for the effective exploitation of execution contexts in alternative ways, such as speculative precomputation.