Weighted aggregation for multi-level graph partitioning

Cédric Chevalier* Ilya Safro†

Nov 2008

Graph partitioning is a technique which is widely used in many fields of computer science and engineering. The goal is to partition the vertices of a graph into a certain number of disjoint sets of approximately the same size, so that a cut metric is minimized. Due to the NP-hardness of the problem and its practical importance, many different heuristics (spectral, combinatorial, evolutionist, etc.) have been developed to provide an approximate result in a decent (near linear) computational time. However, only the introduction of the multilevel methods during the nineties has really provided a break-through in efficiency and quality. We present a coarsening strategy inspired by methods used in algebraic multigrid solvers that is radically different from current methods used in partitioners, and show that this can improve both quality and robustness in graph partitioners.

A multilevel algorithm may be viewed as a process of graph topology learning at different scales in order to generate a better approximation for any approximation method incorporated in the framework. It usually consists in three steps, the *coarsening* which aims at decreasing the number of vertices then the *initial partitioning* on the coarsest graph and then the *uncoarsening* which projects back the solution to the finer levels.

Typically, almost all previously developed multilevel schemes for simple graphs possess exactly the same *"strict"* coarsening. It is carried out by matching groups (usually pairs) of vertices together and representing each group with a single vertex in the coarsened space. Another class of multilevel schemes used for several combinatorial optimization problems is based on the *Algebraic multigrid* (AMG). This class allows to do not only strict aggregation (SAG) but also weighted aggregation (WAG). In SAG process (also called edge contraction or matching of vertices) the nodes are blocked in small disjoint subsets, called aggregates. In WAG, each node can be divided into *fractions*, and different fractions belong to different aggregates, i.e., the set of vertices V will be covered by (presumably) small intersecting subsets of V. The nodes that belongs to more than one subset will be divided among corresponding coarse aggregates.

The main goal of our work was a systematic comparison of the AMG based scheme (WAG) versus strict scheme based on heavy edge matching (HEM, that was adopted since 1995 and implemented in many multilevel packages) for the partitioning problem while having the uncoarsening parts (based on the most popular sequential algorithm called Fiduccia-Mattheyses) exactly the same in both cases. This issue still was not empirically studied in contrast to many other works in which many good uncoarsening and postprocessing procedures were suggested.

As a main result of this work, we can recommend the adoption of WAG instead of a classical HEM. On a large set of test problems, WAG improves the (average) quality of the partitions and provides better chance of finding a near-optimal partition.

*Scalable Algorithms Dept, Sandia National Labs, NM, USA, ccheval@sandia.gov
†Mathematics and Computer Science Division, Argonne National Lab, IL, USA, safro@mcs.anl.gov

Dagstuhl Seminar Proceedings 09061
Combinatorial Scientific Computing
http://drops.dagstuhl.de/opus/volltexte/2009/2082