

High Quality Graph and Hypergraph Partitioning

Wednesday, October 11, 2017 2:00 PM (30 minutes)

In computer science, engineering, and related fields graph partitioning (GP) is a common technique for processing very large graphs, e.g. networks stemming from finite element methods, route planning, social networks or web graphs. For example, in parallel computing good partitionings of unstructured graphs are very valuable. In this area, graph partitioning is mostly used to partition the underlying graph model of computation and communication. Generally speaking, nodes in this graph represent computation units and edges denote communication. This graph needs to be partitioned such that there are few edges between the blocks (pieces). In particular, if we want to use k processors we want to partition the graph into k blocks of about equal size.

Hypergraphs are a generalization of graphs, where each (hyper)edge (also called *net*) can connect more than two vertices.

The hypergraph partitioning (HGP) problem is the generalization of the graph partitioning problem. However, allowing nets of arbitrary size makes it more difficult in practice.

HGP has a wide range of applications. Two prominent areas are

VLSI design and scientific computing (e.g. accelerating sparse matrix-vector multiplications).

While the former is an example of a field where small optimizations can lead to significant savings, the latter exemplifies problems where hypergraph-based modeling is more flexible than graph-based approaches.

Since (hyper)graph partitioning is NP-hard and since it is even NP-hard to find good approximate solutions for graphs, heuristic *multilevel* algorithms are used in practice to partition large instances.

These algorithms consist of three phases: In the *coarsening phase*, the (hyper)graph is coarsened to obtain a hierarchy of

smaller graphs that reflect the basic structure of the input. After applying an *initial partitioning* algorithm to the smallest graph in the second phase, coarsening is undone and, at each level, a *local search* method is used to improve the partition induced by the coarser level.

In this presentation intended for the BDAHM track, we briefly introduce the (hyper)graph partitioning problem

along with the multilevel framework and present the KaHIP (Karlsruhe High Quality Partitioning) family of graph partitioning programs

as well as the hypergraph partitioning framework KaHyPar (Karlsruhe Hypergraph Partitioning).

Both systems provide world class solution quality. For example, KaHIP has been able to improve or reproduce most of the

entries reported in the broadly accepted Walshaw benchmark as well as the best marks both with respect to quality and running

time versus quality among all participants in the 10th DIMACS Implementation Challenge, while KaHyPar is the method of choice for a wide range of hypergraph

partitioning tasks, computing better solutions than the widely used general purpose tools hMetis and PaToH.

Track

BDAHM

Authors: Dr SCHULZ, Christian (KIT); Mr SCHLAG, Sebastian (ITI, KIT); Prof. SANDERS, peter (KIT)

Presenter: Mr SCHLAG, Sebastian (ITI, KIT)

Session Classification: Analytics