Approximate Matchings in Massive Graphs via Local Structure

Clifford Stein
Columbia University, New York City, USA

Abstract

Finding a maximum matching is a fundamental algorithmic problem and is fairly well understood in traditional sequential computing models. Some modern applications require that we handle massive graphs and hence we need to consider algorithms in models that do not allow the entire input graph to be held in the memory of one computer, or models in which the graph is evolving over time.

We introduce a new concept called an "Edge Degree Constrained Subgraph (EDCS)", which is a subgraph that is guaranteed to contain a large matching, and which can be identified via local conditions. We then show how to use an EDCS to find 1.5-approximate matchings in several different models including Map Reduce, streaming and distributed computing. We can also use an EDCS to maintain a 1.5-optimal matching in a dynamic graph.

This work is joint with Sepehr Asadi, Aaron Bernstein, Mohammad Hossein Bateni and Vahab Marrokni.