GRAPH PRODUCTS, HARDNESS OF APPROXIMATION, AND BEYOND

PARINYA CHALERMSOOK

We revisit an old-fashioned reduction for proving hardness of approximation. Suppose we have a reduction from maximum independent set (or graph coloring) that "almost" works but does not. We show how to use graph products to make the reduction work without blowing up the size of the reduction.

This idea has been used in proving hardness of approximation for problems in various areas of computer science. Sometimes it also implies other types of lower bounds (e.g. impossibility of PAC learning). The power of this technique is illustrated by proving the following results:

- Deterministic Finite Automata (DFA) is not PAC-learnable unless NP=RP. This resolves an open question raised 25 years ago by Pitt and Warmuth.
- Edge-Disjoint Paths on DAGs are hard to approximate to within $n^{1/2-\epsilon}$. This matches the upper bound by [Cherkuri, Khanna and Sherpherd 2005].
- Tight hardness of k-Cycle Packing for large k.
- Tight hardnesses of Induced Matching, Poset Dimension
- Tight hardness of approximating Bipartite dimension of a graph (a.k.a. the Boolean rank of a matrix).
- An improved hardness of Strong Edge Coloring and colorings of power graphs.
- Alternate (and arguably simpler) proofs of many other results, such as hardness of learning DNF, CNF and intersection of halfspaces; hardness of some pricing problems.

Our technique reduces the task of proving hardnesses to merely analyzing graph product inequalities, which are often as simple as textbook exercises.

These results are mostly based on joint works with B. Laekhanukit and D. Nanongkai that appeared in SODA 2013, LATIN 2014, and FOCS 2014 (to appear). The bipartite dimension result is based on joint work with S. Heydrich, E. Holm, and A. Karrenbauer [ESA 2014]