

Uncovering Hidden Structure through Parallel Problem Decomposition for the Set Basis Problem*

Yexiang Xue
Cornell University
yexiang@cs.cornell.edu

Stefano Ermon
Stanford University
ermon@cs.stanford.edu

Carla P. Gomes
Cornell University
gomes@cs.cornell.edu

Bart Selman
Cornell University
selman@cs.cornell.edu

Abstract

Exploiting parallelism is a key strategy for speeding up computation. However, on hard combinatorial problems, such a strategy has been surprisingly challenging due to the intricate variable interactions.

In this paper we introduce a novel way in which parallelism can be used to exploit hidden structure of hard combinatorial problems, orthogonal to divide-and-conquer and portfolio approaches. We demonstrate the success of this approach on the minimal set basis problem, which has a wide range of applications e.g., in optimization, machine learning, and system security. We also show the effectiveness of our approach on a related application problem from materials discovery.

In our approach, a large number of smaller sub-problems are identified and solved concurrently. We then aggregate the information from those solutions, and use this information to initialize the search of a global, complete solver. We show that this strategy leads to a significant speed-up over a sequential approach since the aggregated sub-problem solution information often provides key structural insights to the complete solver. Our approach also greatly outperforms state-of-the-art incomplete solvers in terms of solution quality. Our work opens up a novel angle for using parallelism to solve hard combinatorial problems.

*This work was supported by National Science Foundation CCF-0832782, CNS-1059284.
Copyright © 2015, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.