D I S S E R T A T I O N   P R O P O S A L

Thomas Lavastida

"On Scalable Algorithms and Algorithms Augmented with Predictions"

Tuesday, December 8, 2020
9:00am EST
Zoom: https://cmu.zoom.us/s/96055450877

The first three chapters of this dissertation are concerned with finding scalable algorithms for common problems arising in data analysis and machine learning.

In the first chapter, we consider the weighted longest common subsequence problem and its generalization, all-substrings weighted longest common subsequence. This problem arises in computational biology applications, where there is a need to find methods for scaling to massive inputs. We give efficient parallel algorithms which yield nearly optimal solutions for both problems. Additionally, we demonstrate the scalability of this method via an empirical evaluation.

The second chapter is focused on finding parallel and distributed methods for hierarchical clustering, one of the most popular methods for clustering data in practice. In order to get around natural barriers for scaling these methods, approximations are considered. We give efficient parallel and distributed implementations for divisive k-clustering methods as well as the agglomerative centroid-linkage method.

In the third chapter, we propose studying how to scale the popular average-linkage method for hierarchical clustering. To accomplish this, we introduce a general technique we call cluster embeddings.

The final three chapters are concerned with incorporating machine learned predictions into the design of combinatorial algorithms. We give formal models to analyze algorithms augmented with predictions and apply this model to several problems. Our model identifies the following key properties that the predictions must satisfy: (1) for each instance of a problem, there must exist some prediction which is *useful* for solving that instance, (2) the prediction must be *robust* to errors or small changes in the input, and (3) the prediction must be *efficiently learnable* from past instances of the problem.

In the fourth chapter, we consider online load balancing with restricted assignments. We show the existence of predictions which can be used to guide the online algorithm in constructing a fractional solution to this problem. These predictions are robust to small errors and they can be efficiently learned from past instances of the problem. Finally, we show how to round fractional solutions for this problem to integer solutions in the online setting.

In the fifth chapter, we consider the online flow allocation problem in DAG's. Again, we identify predictions for this problem that are useful, robust, and efficiently learnable.

In the sixth chapter, we propose studying augmenting algorithms with predictions to improve their runtime. This can also be viewed as giving theoretical foundations for analyzing warm-start algorithms. In particular, we plan to study primal-dual algorithms.