# Review of Doctoral Thesis of Pavel Dvořák

Amit Chakrabarti

## 1 Overview

This thesis—titled "Limits of Data Structures, Communication, and Cards"—is in the broad area of computational complexity theory and, within it, focuses on lower bounds in concrete models of computation. The thesis provides results in three very popular computational models, namely static data structures, dynamic data structures, and two-party communication, plus a fourth less popular model, namely secure two-party computation using cards. After a short chapter titled *Preliminaries* that establishes some basic definitions and notation, the thesis spends one chapter per computational model to develop these results.

## 2 Dynamic Data Structures

The chapter *Lower Bounds for Semi-adaptive Data Structures via Corruption* studies a family of problems motivated by a specific approach of Pătraşcu that, if successful, would imply strong (at the level of a major breakthrough) lower bounds for various dynamic data structure problems. Specifically, Pătraşcu gave a way of building a "multiphase" data-structure problem out of any two-party communication problem $f$ and conjectured that good lower bounds on the communication complexity of $f$ could translate into such strong lower bounds for the multiphase problem. To date, however, this conjecture remains unproven, and indeed, there is some work showing that a strong version of Pătraşcu's conjecture is false.

The result of this chapter is that Pătraşcu's plan does lead to a restricted kind of data structure lower bound, namely for "semi-adaptive" data structures, a somewhat technical definition tailored to the multiphase problem. These were previously defined and considered by Ko and Weinstein, who proved a good deterministic lower bound for the data-structure problem arising from SET-DISJOINTNESS. Here, the authors extend the Ko–Weinstein result to a broader class of communication problems that have high complexity as given by the smooth corruption bound, a well-studied quantity in communication complexity theory. They also extend the argument to handle randomized query algorithms (for the data structures).

To prove these results, the authors delve deep into the Ko–Weinstein argument where a particular combinatorial property of the DISJOINTNESS function is used. They show how to get by with a less specific combinatorial property that roughly says that large near-monochromatic rectangles do not exist.

Even though no groundbreaking data structure lower bounds are prove here, this chapter makes a good conceptual contribution by somewhat de-mystifying the Ko–Weinstein result and making it clear that there is nothing magical about the DISJOINTNESS function that makes it work.

## 3 Two-Party Communication Complexity

The chapter *Lower Bound for Elimination of Greater-Than via Weak Regularity* considers a direct-sum-like problem in communication complexity and gives a new lower bound for it. Specifically, suppose $f$ is a two-player communication problem with Boolean output and Alice and Bob are given their respective portions of $k$ independent inputs to $f$. A direct-sum or direct-product question would ask about the complexity of evaluating $f$ on all $k$ inputs. Here, instead, the authors consider the *elimination* problem, where Alice and Bob are to eliminate *one* of the $2^k$ possible combinations of outputs, using a low-error randomized protocol. Trivially, this problem is at best easier than computing a single copy of $f$ to low error; the goal then is to show a comparable lower bound.

Earlier works had given lower bounds for the elimination version of some specific problems, which were good for small $k$. The result of this chapter applies to the GREATER-THAN problem $\mathrm{GT}_n$ and fairly large $k$, namely $o(n^\alpha)$ for some fixed $\alpha$. In fact, the candidate has a published work with a more general result, applying to elimination of any function with a large *discrepancy* bound, but for this thesis chapter he has chosen to focus on $\mathrm{GT}_n$ so as to give a fully self-contained proof.

The proof itself is based on showing that under a carefully chosen input distribution $\mu$ (tailored to $\textsc{gt}_n$), all large enough rectangles inside the communication matrix of $\textsc{gt}_n^k$ have the following property. Suppose the rectangle is partitioned into $2^k$ subsets, based on the $k$-bit output of $\textsc{gt}_n^k$ at each input. Then each part has at least roughly $2^{-k}$ fraction of the $\mu^k$-mass of the rectangle, up to some additive tolerance. This property is termed "weak regularity." With this fact in hand, the lower bound for elimination is straightforward.

While the weak regularity proof is still somewhat technical, it is quite a bit more elementary than the unrolled version of the argument that would give a comparable lower bound for elimination of $\textsc{gt}_n$ using discrepancy, because the latter uses a powerful direct product theorem for discrepancy, which in turn goes through SDP duality and non-elementary matrix analysis.

# 4 Static Data Structures

The chapter *Network Coding Conjecture Implies Data Structure Lower Bounds* considers an approach to proving strong lower bounds for static data structures, trading off storage space against query time. Such lower bounds are conditional on a popular information-theoretic conjecture called the *network coding conjecture* (NCC), which asserts that a certain natural lower bound (achievability result) on the rate at which multi-source multi-target data transmission can be accomplished in an undirected network is in fact tight, i.e., it is the largest possible rate. This lower bound is simply the optimal (i.e., maximum) multicommodity flow rate in the network, so the conjecture asserts that no benefit is to be gained from doing any coding at the nodes, in the undirected case.

It is intuitive that such a conjecture might have something to say about the hardness of propagating information in some data structure from a collection of nodes in a graph representing memory cells to other nodes representing queries. At a very high level, the results of this chapter use this intuition to obtain new lower bounds for three basic data structure problems. In greater detail, the argument uses an interesting twist to enable the use of NCC: the query algorithm of the data structure is invoked twice in succession in a careful way. The resulting information propagation is modeled as transmission along a three-layered graph. The double invocation of the query algorithm results in information encoded in the data structure at memory-cell nodes (layer 1) propagating to a middle layer of nodes and then to an output layer (layer 3) of nodes. Invoking NCC, it follows that the graph modeling the data struture's query strategy must be large, which then translates into certain data structure lower bounds.

In a little more detail, the lower bounds are against data structures that are both non-adaptive and "systematic," where the latter condition means that the data structure consists of an as-is copy of the input database plus a small amount of auxiliary info that can be accessed without paying any query cost: thus, the cost is only for accessing portions of the original input. The specific problems considered in this chapter are permutation inversion, polynomial evaluation, and polynomial interpolation. I will not go into a restatement of each of these problems, since the thesis does a very good job already. Each of these problems has good pedigree, with an established body of data structure lower bound results. This makes the contributions of this chapter quite compelling. If there is a criticism to be made here, it is that NCC abstracts out what might be the "hard part" of an unconditional proof, because after all any such lower bound is morally about identifying an information-theoretic bottleneck that prevents a too-efficient data structure. It would be good for the thesis to explicitly address this criticism.

# 5 Secure Computation Using Cards

Finally, the chapter *Barrington plays cards* considers computations in an unusual model where Alice and Bob jointly compute some function on an input split between them (as in communication complexity), with the computation occurring by moving and flipping "cards" according to a specified protocol. The cards provide a way for Alice and Bob to hide their respective inputs from the other player and it is conceivable that a function $f$ could be computed on the joint input $f(x,y)$ without Alice learning anything about $y$, nor Bob anything about $x$, beyond what is revealed by the value $f(x,y)$.

This chapter provides a few results about the computational power of this model, given natural constraints on the resources, i.e., the length of the protocol, the number of "work space" cards used in the protocol and the nature of the manipulations allowed on the input cards, e.g., that the protocol be read-only. The marquee result is that using polynomial length and $O(1)$ work space, a read-only protocol can securely compute any function in $\mathsf{NC}^1$ and that this characterization is tight. The proof uses Barrington's famous theorem characterizing $\mathsf{NC}^1$ using width-5 permutation branching programs. As a result, the main construction in this chapter can focus on implementing transpositions (which are especially simple permutations) using this card model. The latter portion of the chapter considers variants of the card model.

# 6 Assessment

Overall, this work makes several contributions to the theory of lower bounds. These contributions are at a fundamental enough level that they can be explained easily to a broad TCS audience, without the need for hyper-specialization. Several of the proof in the thesis are ingenious. I especially liked the direct proof the "elimination of GT" lower bound in Chapter 3 and the construction of the layered network from a data structure for function inversion in Chapter 4.

All of the work presented has been peer-reviewed by reputable international conferences. As such, the work has already been vetted as making meaningful contributions to theoretical computer science. The breadth and depth of work in this thesis is clearly deserving of a Ph.D. degree.

The writing is of very high quality and it made the thesis a pleasure to read. Every significant proof is preceded by high-level intuition and this intuition is given at a level that both aids in the understanding of the technical details and also serves as a useful overview for a first-pass reading.

While no major changes are required, I will close with some suggested improvements, most of which are about fixing small typos and grammar issues.

# 7 Suggested Changes and Edits

Below are comments to the candidate suggesting small fixes to the thesis.

- Page 12, "simpler version of the multiphase problem": Add a forward reference to where you describe this version.
- Theorem 2.3: "and each of them" → "such that each of them".
- Theorem 2.3: It would be good to explicity quantify $m$, e.g., "there is an integer $m$ and a random variable $\mathbf{Z} \in \{0,1\}^m$".
- Page 16, "two distribution" → "two distributions".
- Before Lemma 2.10: At this point, I suggest giving a longer explanation and more intuition for the bucketing argument to follow. You *did* give a short explanation; just make it more detailed, so a reader can see how the calculation to follow will finish the proof.
- Page 23, table headers: "Ballancedness" → "Balancedness"; "Lower bound of" → "Lower bound on".
- Page 29: "blocks of **V** such that" → "blocks of **V** of the form".
- Chapter 3 ends rather abruptly. Add a bit of text to recap what you did. This could also be a good place to discuss potential future work related to the topic of this chapter.
- Page 32: "involutions that are" → "involutions, i.e.,"
- Page 36: "than it is needed" → "than is needed".
- Page 37: "A coding rate" → "The coding rate".
- Page 37: The final sentence makes it sound like you are skipping something important, since NCC involves undirected communication networks. Consider rewording.
- Page 38: "A flow rate" → "The flow rate".
- Page 38: In the equation before Conjecture 4.1, put $\bar{c}$ on the left side, since you're defining $\bar{c}$.
- Page 38: Typo, "circuites" → "circuits".
- Page 40: "understand a graph" → "mean a graph".
- Page 41: Right after the first para, add a paragraph or subsection heading (basically, some sort of visual cue) to separate off the formal proof which begins here from the discussion thus far, so that the reader is alerted that the material so far is supposed to be mostly intuition.
- Page 41, before Fact 4.5: "well-known Stirling's formula" → "Stirling's well-known formula".

- Page 42, last line: “the same advice” → “this particular advice”.
- Page 45: The argument in the first paragraph feels a bit sketchy. You are fixing some coefficients and also fixing some evaluations of your polynomial. You need to add more detail to show that $|\mathscr{F}|$ is indeed lower-bounded as you claim.
- Page 47, item 3: “the non-deterministic” → “the class of functions computable in non-deterministic”.
- Page 50: “other player input” → “other player’s input”.
- Page 51: “the famous Barrington’s theorem” → “Barrington’s famous theorem”.
- Please add some text to Chapter 4 to address my comments made above.
- Please add some text to Chapter 5 to explain why the average theorist should care about this card model of computation. Is it at least tangentially related to some practical computational scenario?