Acknowledgments

In the next several paragraphs I will thank the people who supported me during the time of pursuing my PhD. Then, I will share some of my personal experience from the last five years as a PhD student, so this will be a bit longer than a typical acknowledgments section.

First, I would like to deeply thank my advisor Micha Sharir, who took me under his wings and guided me on how to approach theoretical computer science. The most valuable thing for me was the numerous research discussions with Micha, having a peek at how he approaches research problems, the way he simplifies them and getting fast to their bottleneck, the core one should be focused on when trying to discover new results. I perceive these discussions as "private lessons" for a deep analytical thinking, and in particular, for research in theoretical computer science, regardless of the problems in hand. I feel it totally changed the way I approach analytical problems, and not only related to research, but also in other complex matters I encounter, such as in programming, and solving engineering problems. The value of these "lessons" is priceless.

I would like to thank Liam Roditty, who discussed with me about graph diameter algorithms, and connected me with Keerti Choudhary. This has led to the discovery of non-trivial "diameter spanners" in directed graphs, and to the joint paper with Keerti [65], which Chapter 6 in this thesis is based on. This is also the right place to thank my co-author Keerti Choudhary, for her dedication to our project and the fruitful joint work.

I would like to thank Reuven Cohen, my master's advisor from Bar-Ilan University, who encouraged me to pursue a PhD, and provided all the assistance needed to make it happen. I thank also Moshe Lewenstein who gave me a valuable feedback for my PhD research proposal, and invited me to a full week workshop on "Structure and Hardness in P" at the Dagstuhl castle in Germany.

I would like to thank the senior researchers who reviewed this thesis. I felt honored to receive your reviews. A special thanks goes to the reviewer who noticed that we can shave-off the $\log\log\log n$ factor from our Dynamic Time Warping and Geometric Edit Distance algorithms that appear in Chapter 4, using the SMAWK algorithm for totally monotone matrices [9].

I would also like to thank fellow PhD students with whom I had many interesting research conversation with: Sarel Cohen, who introduced to me some cool research questions in graph algorithms, and sparked my interest in the field. Dor Minzer, who I shared with some research problems I have been working on, and gave me valuable feedbacks. Orr Fischer, with whom I had many interesting conversations about theoretical computer science and academia in general.

Last but by no means the least, I would like to thank my parents, Michael and Neomi, who supported me throughout this long academic track, from starting as an undergraduate student in Ben-Gurion University more than 10 years ago, throughout pursuing my master's in Bar-Ilan University, and finally, throughout the extensive journey of doing a PhD in Tel Aviv University. Thank you for your endless support, encouragement, and being there for me whenever needed.

My Personal Experience. Computer science is a relatively new scientific field, and it is emerging with many important discoveries every year. The amount of new significant discoveries in theoretical computer science that were discovered only during the time I was a PhD student really amazed me. I feel lucky I had the opportunity to do research in theoretical computer science during this time in history, as I could witness plenty of new interesting results and breakthroughs being published by the community in "real time", and sometimes even to contribute a little bit. I do not know for how long this rate of new significant discoveries in computer science will continue, but thinking again about how young this scientific field is (which only 100 years ago nobody knew about), I guess that this rate will not decline anytime soon.

I remember that when I began my PhD studies in late 2014, my advisor Micha told me about a recent breakthrough on the 3SUM problem (determining whether there are three numbers that sum to zero in a given set of n real numbers). Allan Grønlund and Seth Pettie [104] showed that 3SUM can be solved in subquadratic time. Although only small polylogarithmic factors were improved over the well-known $\Theta(n^2)$ time bound, this result made huge strides, since the 3SUM problem is well-known for basing conditional lower bounds for many other problems, and therefore, it raised doubts on the optimality of many other algorithms, such as for determining whether n given points in the plane are located in a general position (i.e., no three points lie on a common line).

The same day, I started reading the paper of Grønlund and Pettie with enthusiasm, hoping that maybe a further improvement is possible. Since it was the first serious theory paper I read, it took a while until I controlled the details. It took months of thinking and many discussions with Micha until finally finding a way to improve their algorithm and decision tree bounds. Although the improvements were small (shaving polylogartihmic factors from both bounds), the exciting thing was that my first theory result was about a well-known problem.

Later, Timothy M. Chan [55] improved a bit further the algorithmic time bound for 3SUM (by another logarithmic factor). In 2017, a breakthrough on this problem came from Daniel Kane, Sachar Lovett, and Shay Moran [114], who showed that the decision tree complexity of 3SUM is near-linear, improving significantly our $O(n^{3/2})$ decision tree bound (and the $O(n^{3/2}\sqrt{\log n})$ bound of Grønlund and Pettie). I was very surprised that such a significant improvement is even possible. Their technique also gave near-linear decision tree complexity bounds for other core problems, such as "Sorting $X + Y$" and "All-Pairs-Shortest-Paths". What I described in this paragraph truly relates to what I mentioned in the first paragraph, about being lucky to witness breakthroughs during this time, especially when they are related to topics I have been working on.

Then, I have been working on extending the technique we used for the 3SUM problem, and looking for other fundamental problems to apply it, but did not find one. Until, one day Pankaj K. Agarwal gave a talk in our weekly computational geometry seminar about approximation algorithms for the Dynamic Time Warping and Geometric Edit Distance problems [8]. That was the first time I heard of these problems, and I discovered then that the best known algorithms to solve them use a standard dynamic programming approach that takes quadratic time [150]. During his talk I started to think about whether we can break this quadratic-time barrier. It took a dramatically more sophisticated use of the techniques used in our 3SUM paper, in conjunction with other techniques, until we finally managed to break the 50 years old quadratic time barrier for both Dynamic Time Warping and Geometric Edit Distance by a $\log \log n$ factor (actually it was a $\log \log n / \log \log \log n$ factor, but thanks again to one of the reviewers of this thesis, who noticed that we can in fact shave-off the $\log \log \log n$ factor). Now, when I look at this improvement factor it seems funny, as its growth rate (in proportion to the input size n) is very slow, but this is the nice thing about theoretical computer science, our goal is to find the optimal algorithm, the one that its runtime cannot be improved by *any* asymptotic factor. Practically, our algorithm can perhaps improve the runtime for very large inputs (depends also on the constant of proportionality in our time bound) over the standard quadratic-time algorithm.

The next result was on the high-dimensional L_∞ Closest Pair problem. That is, finding the closest pair of points under the L_∞ metric in a given set of n points in $\mathbb{R}^d$, where $d = \text{poly}(n)$ (for example $d = n$). We gave a new algorithm for this problem, improving a previous algorithm of Piotr Indyk, Moshe Lewenstein, Ohad Lipsky, and Ely Porat [112]. The thing I remember the most from this paper is that it appeared in the ISAAC 2017 conference that was held in Phuket, Thailand in a very nice suites hotel on the beach. It was definitely my most unforgettable academic trip to date. This trip has led me to travel more in Thailand, learn more about Southeast Asia, and to visit the Philippines for a whole month a year later. I had a blast in both Thailand and the Philippines. I met in both countries super friendly people and liked the general relaxed vibe.

After I finished working on these three papers, I felt eager to diversify my research and looked for areas I have not worked on before. I started exploring more seriously about graph algorithms. This has led to some interesting discussions with Liam Roditty, who also connected me with Keerti Choudhary. The work with Keerti has led to our joint SODA paper [65], in which we proved the existence of various non-trivial "diameter spanners" for directed graphs. That is, that any sufficiently dense directed graph has a significantly sparser subgraph that preserves the diameter of the original graph up to a factor that is strictly less than 2 (called also "stretch factor"). We showed how to efficiently compute such subgraphs with various non-trivial size-stretch trade-offs. This opens a large room for future work, and it will be interesting to see what new results on this topic will be further discovered.

Omer Gold
December 28, 2019