Preface and Acknowledgment

This thesis is about randomized optimization, and even though it is doubtful whether the work itself approaches any reasonable optimization criteria, I can definitely say that a lot of randomization was involved in the process of writing it. In fact, people have been the most valuable source of randomness. To begin with, the subject of this thesis is almost completely random. When I started back in October 1991, I had just finished my Diplomarbeit with Emo Welzl. This was about VC-dimension [19], and the only thing clear at that time was that I had to find something else to work on. As it happened, Micha Sharir and Emo had just developed a simple and fast randomized algorithm to solve linear programming and a whole class of other optimization problems (including e.g. the minimum spanning ball problem) in a unified abstract framework [45]. Independently (another accident!), Gil Kalai – working on the polytope diameter problem where he had obtained a striking new result [28] – came up with a randomized variant of the simplex algorithm, capable of solving linear programs in an expected subexponential [1] number of steps [27]. Since all prior bounds were exponential, this was a major step forward on the longstanding open question about the worst-case performance of the simplex algorithm. Kalai's simplex variant turned out to be closely related to the Sharir-Welzl algorithm, and together with Jirka Matoušek, they were able to prove that this algorithm was actually subexponential as well, at least for linear programming [33].

This little 'at least' led to my involvement in this research. It started with Emo asking me – quite innocently – whether I had any idea how to make the subexponential analysis applicable to the minimum spanning ball problem as well. The difficulty was that although the *algorithm* worked, the *analysis* didn't go through because it was using a special feature of linear programming.

Within short time, I was convinced that the problem was solved, and I proposed a solution, which quickly turned out to be wrong (like many ideas and 'solutions' still to come). This failure made me dig even deeper into the problem, and during the winter, the minimum spanning ball problem became 'my' problem. By March 1992 I had collected all kinds of half-baked ideas but none of them seemed to be any good. I was stuck. During that time, Franz Aurenhammer was my office-mate (remarkably, he left Berlin only a month later), and I kept telling him the troubles I had with the 'miniball' problem. On one of these occasions, from the depths of his memory, Franz recalled a paper by Rajan about d-dimensional Delaunay triangulations (this surely was no accident) where he remembered

[1] A function is called subexponential if its logarithm is sublinear.

having seen something mentioned about the miniball problem [42]. When I looked up the paper, I found that Rajan had noted a close connection between the miniball problem and the problem of finding the point in a polytope closest to the origin. I immediately had an idea about how to solve the latter problem![2] Not only that: it turned out that even before, I had, without noticing it, already collected all the ideas to make the subexponential analysis applicable to the whole class of abstract problems originally considered by Micha and Emo – except for the concrete miniball problem on which I had focused too hard to see beyond it. Only in the process of preparing a talk about all this for our notorious 'noon seminar' did I notice that even more general problems could be handled. This led to the concept of *abstract optimization problems* (AOPs) and a subexponential algorithm for solving them. Without Franz innocently leading me on the right track, I would not have achieved this.

Still, I was not quite happy. After having spent so many months on the problem, I found the solution too obvious and the algorithm – containing a recursive 'restart scheme' based on an idea in Kalai's paper – still too complicated. I spent another couple of weeks trying to simplify it, without any success. Meanwhile, Emo was trying to push forward publication of the new subexponential bounds and suggested that I should submit the results to the forthcoming *33rd Symposium on Foundations of Computer Science* (FOCS). I didn't want to submit it until I could simplify the algorithm, but I kept on trying in vain. Three days before the submission deadline, I happened to talk to Johannes Blömer. He was about to submit a paper to FOCS for his second time, and he told me that the first time he had been very unsure about the quality of his paper – just like me – but had (successfully) submitted it anyway [7]. He urged me to do the same, his main argument being: it can't hurt. This was quite convincing; within a day and a night I wrote down the current status of my research, and we got our submissions delivered together and on time, by express mail. In the end, both our papers got accepted [8, 18]; for Johannes, this was the beginning of the final phase of his thesis. For me it was a first and very encouraging start that I owe to Johannes.

After I returned from FOCS – which was an exciting event for me, partially because it was my first visit ever to the US – I tried to build on the results I had for AOPs. In particular, I was looking for lower bounds. After all, the AOP framework was so weak that it seemed possible to establish a superpolynomial lower bound on the runtime of any algorithm for solving them. Although I spent weeks and weeks, I must admit that my insights into this problem are still practically nonexistent. Another issue I tackled was trying to improve the subexponential upper bound. The exponent was of the order $\sqrt{d}$. Maybe one could get $\sqrt[3]{d}$, or even $\log^2 d$? There was reason to believe that such bounds were not impossible, at least for linear programming and other concrete problems. Well, I couldn't prove any.

On the one hand, I wanted to give up on this line of research, and on the other, I had no idea how to continue with my thesis if I did. For several months, I kept muddling

[2] Later I noticed that this was basically the idea already pursued by Wolfe [48].

along and made no considerable progress. Unfortunately, Emo was on sabbatical leave and therefore did not notice what was going on (and I didn't tell him, either).

At this point – when my depression reached a local maximum in the middle of 1993 – Günter Ziegler came to the rescue. He only happened to intervene because Jirka Matoušek was visiting us in Berlin at that time, and Günter took the opportunity to tell us both about one of his favorite open problems. It was about the behavior of the Random-Edge-Simplex algorithm on the so-called Klee-Minty cube. Fortunately, he mentioned this application only in a side remark but otherwise presented the problem as a randomized flipping game on binary numbers that he would like to have analyzed. In this formulation, its pure combinatorial flavor immediately attracted me. Moreover, Günter cleverly adorned his presentation with stories about people having either claimed wrong solutions or tried in vain. Although I saw no immediate connection to my previous research, I instantly started working on the problem, partly because I was just lucky to have found something new and interesting to think about. Günter had appeared just at the right time, with the right problem.

History was beginning to repeat itself. Just like in the miniball case almost two years ago, the problem looked doable at first sight, but very soon turned out to be much harder than I had originally thought. This time the fall found me at my desk, filling one notepad after another. It was late December when the only visible result of my investigations was a vague idea how to analyze another flipping game that I thought might have something to do with the original one.

Emo had returned from his sabbatical and had gone back to the old habit of letting his graduate students report on their work in regular intervals. When I was due, I told him about my idea and what I thought it might give in the end. Almost instantly, Emo came up with a way of formalizing my vague thoughts, and within half an hour we set up a recurrence relation. By the next day, Emo had found a beautiful way of solving it, incidentally rediscovering Chebyshev's summation inequalities, Appendix (7.9), (7.10). Still, it was not clear what the bound would mean for the original game, but after this preliminary success, my mind was free enough to realize that the correspondence was in my notes already. I just had to look at things another way. Talking for half an hour with Emo had made all the difference.

Günter, who had noticed since December that things were gaining momentum, meanwhile had a couple of other results, and we put everything together in a technical report. Just to inform Jirka that we now had a solution to this problem that Günter had told us about months ago, I sent him a copy. Jirka not only liked the solution, he also found that the correspondence between the two flipping games was actually much more explicit (and much nicer) than I had realized. We gratefully incorporated his remarks into the paper, and in this version it finally ended up at FOCS '94 [21]. Like Johannes two years before, I had the end of my thesis in sight.

However, it took another year until it was finally written, during which two more accidents crucially determined the way it looks now. First, I discovered Chvátal's beautiful book on linear programming [10] and second, I visited Jirka in Prague. From Chvátal I learned about the simplex method, and this had more impact on the style of this thesis than any other text I read. Of course, I always had in mind the geometric interpretation of the simplex algorithm as a path-following method on the boundary of a polyhedron, but never got around learning the actual algebra, mostly because I thought it was complicated and unnecessary. Just following Chvátal's initial example explaining how the method really works, was a revelation. Not only did I find it amazingly simple, I also discovered that arguing on the algebraic level is nothing to be afraid of! Geometry helps, but restricting oneself to geometric arguments can impede the view. Without this insight I would not have been able to write the chapter about convex optimization. Moreover, the whole approach of this thesis ('simplex is not everything, but everything is simplex') would have been different without Chvátal's book.

Finally, I still had the problem that my results on the AOPs did not naturally fit together with the material on the Klee-Minty cube (except probably under the very general 'everything is simplex' philosophy). Then I visited Jirka in Prague for a week. We talked about lower bounds for the Random-Edge-Simplex algorithm in an abstract setting and what kinds of input might be useful in order to obtain such bounds. At first, this didn't seem to lead anywhere, but back home I decided in this context to reread Jirka's paper about lower bounds for an abstraction of the Random-Facet-Simplex algorithm [34]. Reading it then, I suddenly realized that the Klee-Minty cube was just a special case of the general examples Jirka had considered and that the techniques I had applied to it did actually work in this general setting as well. Although this insight has not so far led to new results, it provided exactly the missing link for embedding my result on the Klee-Minty cube very naturally into the chapter about lower bounds on abstract optimization.

All the people mentioned above have crucially influenced this thesis. Of course, there are other persons that contributed to it in one or the other way.[3] I would explicitly like to thank Günter Rote who showed an ongoing interest in any developments concerning the AOPs. In particular, he advertised them for me by giving a talk about the subject in China. Günter always used to look at AOPs in a different (and more general) way than I did; the conversations with him and the careful handwritten notes he made during preparation of his talk have been very inspiring. The only reason 'his' AOPs did not make it into this thesis is that I felt they would take me just a little to far out into the 'sea of abstraction' to be able to get back home safely.

Finally, the most decisive thanks go to my advisor Emo Welzl, for never loosing faith in me. I knew whatever I did would in one way or the other make sense to him. Even in the depressing times of zero progress, I didn't need to explain myself – we had a quiet agreement that things would go on sooner or later. I was free to think about whatever I wanted to, and I was free to do it in my style and at my own pace. Only at the very end did Emo urge me to get it over with. As I know now, this was necessary – and hopefully sufficient.

[3] Nina Amenta etwa hat das Vorwort sprachlich überarbeitet, mit Ausnahme dieser Fußnote.