

Executive Summary of the Dagstuhl Seminar 06111 “Complexity of Boolean Functions”

Matthias Krause, Mannheim
Dieter van Melkebeek, Madison, Wisconsin
Pavel Pudlák, Prague, Czech Republic
Rüdiger Reischuk, Lübeck

Keywords:

Boolean and quantum circuits, discrete problems, computational complexity, lower bounds, communication complexity, proof and query complexity, randomization, pseudo-randomness, derandomization, approximation, cryptography, computational learning

1 Introduction and Goals

Estimating the computational complexity of discrete functions is one of the central and classical topics in the theory of computation. Mathematicians and computer scientists have long tried to classify natural families of Boolean functions according to fundamental complexity measures like Boolean circuit size and depth. A variety of other nonuniform computational models with individual bit operations have been considered: bounded fan-in circuits, formulae, branching programs, binary decision diagrams (BDDs), span programs, etc.

The analysis and relative power of these models remains a major challenge. For models of low expressive power, non-trivial efficient realizations of certain hardware-relevant functions have been found, but this question is still open in many cases. Several lower bound techniques for explicitly defined Boolean functions have been developed – most of them are of combinatorial nature. Such negative results are not only of theoretical value, but would have constructive implications, for example in cryptography and derandomization.

Methods that were originally designed to analyze the expressive power of restricted circuit models have also yielded interesting applications in other areas, such as hardware design and verification, algorithmic learning, neural computing, and quantum computing. This leads to the problem as to what type of proof method might be developed and applied at all in this setting. For higher complexity classes, we now know that the existence of natural lower bound arguments would disprove widely believed hardness assumptions. Thus, novel approaches are needed to establish lower bounds for more expressive models in discrete computational complexity.

Nowadays, investigations on the computational complexity of discrete functions have diverged and specialized into many different branches such that it becomes hard to

keep a close look at all approaches. Thus, it is important to bring together researchers from different subareas in a more relaxed atmosphere Dagstuhl provides (as compared to the situation at the major international conferences in this field like STACS, STOC, CC or FOCS) to foster interaction and exchange of new ideas that might be applied in other settings as well. On the one hand, we wanted to present some of the most recent results in the different subareas to a broader audience, in particular in currently fast developing areas like, for example, approximation, communication and proof complexity or quantum computing. Secondly, we wanted to give the opportunity to discuss extensions of different proof methods as well as their applications to other fields.

2 Organization of the Meeting

About 60 researchers accepted our invitation to come together in Dagstuhl for this meeting. Half of them had the chance to present their results in a plenary talk. The length varied between 25 and 60 minutes, the spectrum of their focus ranged from an overview on the state of the art in a larger subarea up to the recent solution of a specific problem. In addition, there were many discussions in smaller groups inbetween talks or after dinner.

The plenary events were structured into five morning and three afternoon sessions. Each session focussed on a special topic. We had two sessions on the basic subject circuit complexity including related models like BDDs and another one on machine-based complexity. Further topics were communication complexity, randomness, algorithms in general, and algorithmic learning. A special session was devoted to cryptography, quantum computing and quantum protocols.

3 Topics Discussed and Achievements

In the following we list some of the major topics that have been considered during the meeting. More details can be found in the Abstracts Collection, which are ordered alphabetically by authors’ names. It also contains additional material of the participants that has been presented in smaller groups or has been evolved from discussions in Dagstuhl.

Proving lower bounds for unrestricted Boolean circuits seems unlikely to be resolved within the next years. Despite the simplicity of the computational model there are only few cases for which it is known what optimal circuits look like. Quadratic functions can be computed by a single-level of AND-gates and there has been a long standing conjecture that this circuit design is close to optimal. Stasys Jukna showed that this conjecture is far from being true by establishing an almost linear gap between optimal circuits and single-level circuits.

A problem of similar flavour is the task to realize a set of monomials by AND-gate circuits of minimal size. Now, we have many outputs since the different monomials are not connected by OR-gates – they have to be evaluated separately. For this family of Boolean functions Jan Arpe showed upper and lower bounds on the best possible circuit design that can be computed efficiently. As usual, for the lower bound one has to assume some intractability property – in this case $\mathcal{P} \neq \mathcal{NP}$ suffices. Lisa Hellerstein discussed another approximation problem, to find a smallest DNF formula that is consistent with a given truth table. Whereas it was known for a long time that finding the optimum is \mathcal{NP} -hard, she now showed strong lower bounds on the best possible approximation ratio under a slightly stronger intractability assumption. In addition, one can find examples for which the obvious greedy strategy performs extremely bad.

Emanuele Viola proved lower bounds for approximating the majority function by depth-3 circuits and explained how this circuit result translates into a lower bound for the classical result that \mathcal{BPP} is contained in $\Sigma_2^{\mathcal{P}}$. Approximation techniques have also been considered by Stephan Waack for the case of Parity-BDDs. Ilan Newman presented a nontrivial method to approximate the maximal number of clauses in a CNF formula that can be satisfied simultaneously.

Jehoshua Bruck devoted his contribution to the parity function and discussed in detail what has been known about it. He explained how parity functions can be used to design special codes to overcome faults in storage systems. The coding problem for binary sequences where only a subset of positions have to be specified was considered by Alexander Andreev. He presented codes of nearly optimal rates that possess efficient coding circuits. The complexity of multiplying two n -bit numbers is another classical problem. Ingo Wegener showed that Nechiporuk’s method can be applied to this function in the branching program model. This way one gets lower bounds of order $n^{3/2}$. A new variation of branching programs called *incremental branching programs* was advocated by Piere McKenzie. He showed tight connections of this model to classical machine models. For branching programs with a read-once restriction Martin Sauerhoff presented a construction that gives an exponential separation between the classical model and the quantum version.

Representing Boolean functions as polynomials and investigating properties of these polynomials like degree or Fourier coefficients has turned out to yield quite strong proof techniques, in particular for Boolean circuits of unbounded fanin. Frederic Green gave a longer survey talk on these methods. As new results he presented exponential lower bounds for special depth-3 $\text{mod } m$ circuits computing $\text{mod } q$ functions for m, q relatively prime. Analysing exponential sums he showed that the correlation between $\text{mod } m$ and $\text{mod } q$ functions is quite small.

Algorithmic learning of Boolean functions given a sample of function values has also been discussed. Traditionally, one has focussed on the behaviour of the learner, assuming that the teacher provides examples at random. Frank Balbach and Thomas

Zeugmann have considered the opposite scenario in which the teacher should select the examples carefully such that every learner – either on average or even a stupid one – learns the Boolean function as fast as possible. It turns out that in this setting the VC-dimension of Boolean concept classes has to be replaced by other properties of Boolean functions.

The distributed bit complexity of computing Boolean functions – also called communication complexity – has been investigated by Martin Dietzfelbinger in an average case setting and for probabilistic protocols by Ronald de Wolf. For example, we have learned that a common source for classical random bits (shared randomness) cannot be compensated by quantum bit communication. Anna Gal showed how multiparty communication complexity relates to combinatorial properties of matrices and Hadamard tensors. Privacy is an important issue when designing communication protocols. Andreas Jakob gave examples that previous characterizations of privately computable functions were incorrect and proved that a new definition is able to correctly measure the leakage of information when running an arbitrary protocol. Eike Kiltz considered privacy in an algebraic setting.

For efficiency and security, random bits are often essential. Michal Koucky discussed the problem to efficiently generate an almost unbiased random string by two agents that do not trust each other. Similarly, Ronen Shaltiel gave a new construction for disperser graphs from two weakly random sources. Several participants discussed randomness in a broader setting. Konstantin Pervyshev asked whether a single bit of advice can be helpful in randomized and quantum computations and gave a positive answer by establishing strong separations of corresponding complexity classes. Scott Diehl showed lower bounds for quantified Boolean formulae with a bounded number of quantifier alternations when given to a probabilistic TM with small space, even if a small error probability is allowed. Lance Fortnow, using the time-bounded variant of Kolmogorov complexity, established a nontrivial relation between worst-case and average-case complexity making an intractability assumption for a family of circuits with nonstandard gates.

For cryptographic applications Boolean functions should have the property that the input-output dependencies are highly masked. Claude Carlet explained what is known in this respect. He discussed in detail the notion of nonlinearity and algebraic degree. Mirosław Kutylowski showed how bit faults in Boolean circuits can be used to efficiently attack pseudorandom generators. Bit commitment is an important basic primitive for various more complex cryptographic applications. The realization of this primitive under different system settings is still an open problem. Maciej Liskiewicz presented a quantum protocol for bit commitment that is quite robust against cheating of either party. Coding issues were addressed by Carsten Damm improving on a classical cryptosystem proposed by McEliece, and by Kazuo Iwama who investigated the capacity of a network with respect to quantum bits. A tight relationship between list-decoding of error correcting codes and amplification of the hardness of Boolean functions was presented by Valentine Kabanets.

The complexity of resolution proofs was discussed by Jacob Nordström. He considered the complexity measures width and space and showed a separation between both measures. Finally, we had several contributions addressing more general complexity theoretic questions. For example, Chris Umans gave an overview on the state of the art in matrix multiplication and presented a group theoretic approach how the upper bound on the exponent might be reduced to the value 2. David Barrington considered the classical reachability problem for graphs when restricted to grid graphs and obtained better space upper bounds. Combinatorial techniques that might be useful in the analysis of Boolean circuits were discussed by Eldar Fischer (Szemerédi's famous Regularity Lemma) and Thomas Thierauf (matchings). Peter Miltersen addressed relations towards numerics and Thomas Hofmeister efficient algorithms for string problems.

4 Conclusion

Understanding the complexity of Boolean functions is still one of the fundamental tasks in the theory of computation. At present, besides classical methods like substitution or degree arguments a bunch of combinatorial and algebraic techniques have been introduced to tackle this extremely difficult problem. These techniques have also found applications in other areas of computational complexity – in some cases it worked also the other way around. There has been significant progress analysing the power of randomness and quantum bits or multiparty communication protocols that help to capture the complexity of Boolean functions. For tight estimations concerning the basic, most simple model – Boolean circuits – there still seems a long way to go.