THESIS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

Topics in Hardness of Approximation and Social Choice Theory

MARCUS ISAKSSON

DEPARTMENT OF MATHEMATICAL SCIENCES
CHALMERS UNIVERSITY OF TECHNOLOGY AND UNIVERSITY OF GOTHENBURG
GOTHENBURG, SWEDEN 2010

Topics in Hardness of Approximation and Social Choice Theory
Marcus Isaksson

ISBN 978-91-7385-392-7
Doktorsavhandlings vid Chalmers tekniska högskola. Ny serie nr 3073.
ISSN 0346-718X

Department of Mathematical Sciences
Chalmers University of Technology and University of Gothenburg
412 96 Gothenburg, Sweden
Phone: +46 (0)31-772 10 00

Printed in Gothenburg, Sweden, 2010

Topics in Hardness of Approximation and Social Choice Theory

Marcus Isaksson
Department of Mathematical Sciences
Chalmers University of Technology and University of Gothenburg

ABSTRACT

Tools from Fourier analysis of Boolean functions have commonly been used to prove results both in hardness of approximation in computer science and in the study of voting schemes in social choice theory. In this thesis we consider various topics in both these contexts. In hardness of approximation we study the asymptotic approximation curve of MAX-CSP's for predicates given by linear threshold functions and prove upper and lower bounds for this curve for majority-like threshold functions. We also relate the hardness of MAX-q-CUT to a conjecture in Gaussian isoperimetry and the plurality is stablest conjecture in social choice. In particular the Frieze-Jerrum semidefinite programming based algorithm for MAX-q-CUT achieves the optimal approximation factor assuming the unique games conjecture if plurality is indeed stablest. In social choice theory we show a quantitative version of the Gibbard-Satterthwaite Theorem, showing that for election schemes in elections with more than 2 candidates, situations in which a voter has an incentive to manipulate by not voting according to his true preference are common enough that they cannot completely be masked behind computational hardness. We also prove a generalization of a Gaussian isoperimetric result by Borell and show that it implies that the majority function is optimal in Condorcet voting in the sense that it maximizes the probability that there is a single candidate which the society prefers over all other candidates.

Keywords: Hardness of approximation, social choice theory, Gibbard-Satterthwaite, max-q-cut, Condorcet voting, linear threshold functions.

List of papers

This thesis consists of the following papers:

I. **Marcus Isaksson and Elchanan Mossel,** ***Maximally Stable Gaussian Partitions with Discrete Applications*****. (Submitted).**

II. **Marcus Isaksson, Guy Kindler and Elchanan Mossel,** ***The Geometry of Manipulation - a Quantitative Proof of the Gibbard Satterthwaite Theorem.*** **(Submitted).**

III. **Mahdi Cheraghchi, Johan Håstad, Marcus Isaksson and Ola Svensson,** ***Approximating Linear Threshold Predicates*****. (Submitted).**

Acknowledgments

I would like to thank my adviser Jeffrey Steif for inspiration, support and for many valuable comments on the draft of this thesis. I also want to thank my co-adviser Devdatt Dubhashi, for your support and for bringing me to Chalmers in the first place, and of course the co-authors of the papers in this thesis; Elchanan Mossel, Johan Håstad, Guy Kindler, Ola Svensson and Mahdi Cheraghchi. I had a great time visiting you and working with you on the topics of this thesis. I also enjoyed thinking about the problems we did not solve.

Furthermore, I want to thank all my colleagues at the department; Oscar, Janeli, Emilio, Alexandra, Marcus, the floorball players and many more for making my life here so much more fun.

Finally, I want to thank my family, my parents and my sisters for all your support and for always believing in me, and most of all, Eugenia for brightening up my days and for sharing your life with me.

Marcus Isaksson
Gothenburg, April 19, 2010

Contents

List of Figures

1

Introduction

An approximation algorithm is an algorithm that computes approximate solutions to a given optimization problem. Consider for instance the problem of coloring regions or countries on a map using a fixed set of colors so that adjacent regions have different colors. It is well-known that four colors suffice if the map is planar and the regions are contiguous. For more general maps more colors may be needed. Thus, if we decide that we want to color the map using only three colors, we cannot in general hope to find a coloring. Instead we may consider the optimization problem of finding a coloring that maximizes the number of adjacent pairs of regions which are differently colored. The optimal such coloring can be found by trying all possibilities but this is a very inefficient algorithm - requiring time exponentially increasing with the number of regions. Frieze and Jerrum gave an efficient (polynomial time) approximation algorithm for (a more general version of) this problem which guarantees to always find a coloring in which the number of adjacent pairs of regions which are differently colored is at least 83.6% of the optimal number. One may ask if this number can be improved upon. This is the type of question studied in hardness of approximation, where the objective is to prove upper bounds on the approximation guarantees that can be achieved by *any* efficient algorithm for a given optimization problem.

In social choice theory one studies methods for collective decision making. An example is an election of say a president or a mayor from a fixed number of candidates. For this example one of the most obvious methods is *plurality voting*, where every individual votes on one candidate and then the winner is selected to be the candidate with the most number of votes. But there are also many other methods in use. For example electoral voting systems, where (in a simple case) individuals in subregions select a winning candidate by plurality voting, and then the winner is selected to be the candidate that won the most number of subregions. Another alternative is *instant runoff voting*, where each individual presents a preference ordering on the candidates and then, repeatedly, the candidate which is preferred by the least number of individuals is eliminated until one candidate is preferred by a majority of the individuals among the remaining candidates. Although plurality voting seems very reasonable it does have some possibly undesirable properties. For instance, if all individuals have an implicit preference ordering of the candidates, there could be situations in which the runner-up in the election is actually preferred over the winner by a majority of the voters. Furthermore, there can be situations in which a well-informed individual has an incentive to vote strategically. For instance, if his top preference is a candidate which is very unlikely to win, but he prefers the likely runner-up more than the likely winner, he would have a higher probability of changing the outcome to his advantage by voting on the likely runner-up. Both of these situations commonly occur when a third candidate which have no chance of winning "steals" votes from one of the two top candidates. Arrow and Gibbard/Satterthwaite, respectively, showed that both these situations cannot be avoided by changing the voting method without giving up some other desirable property. On the more positive side it is conjectured that, under certain requirements, plurality is the most stable voting method under noise, where a small fraction of the votes are assumed to be incorrectly registered.

These are examples of the topics considered in this thesis. In the next section we give some more formal background to the theory used in the papers. Then follows an overview of the included papers, and finally the three papers themselves.

2

Background

2.1 Boolean functions

Many of the results in this thesis involve Boolean functions. In this section we introduce Boolean functions and some common notation and tools used to deal with them.

A Boolean function of arity n is a function $f\colon \{0,1\}^n \to \{0,1\}$, mapping a set of n Boolean values to another Boolean value. 0 is usually thought of as *false* and 1 as *true*. In many settings it is more convenient to work with ± 1 instead, so for the remainder of the section a Boolean function is a function $f\colon \{-1,1\}^n \to \{-1,1\}$.

We will usually think of the input as being uniformly distributed over $\{-1,1\}^n$ and denote it by X. By the Fourier-Walsh transform, any *real-valued function* on the *Boolean hypercube*, $f\colon \{-1,1\}^n \to \mathbb{R}$, can be written uniquely as a multilinear polynomial in the input variables

$$f(x) = \sum_{S \subseteq [n]} \hat{f}(S) \prod_{i \in S} x_i. \tag{2.1}$$

To see this it is enough to note that $\left(\prod_{i\in S} x_i\right)_{S\subseteq[n]}$ is an orthonormal basis for the vector space of functions $f\colon \{-1,1\}^n \to \mathbb{R}$ equipped with the inner product

$$\langle f, g\rangle = \mathbf{E}[f(X)g(X)].$$

By (2.1), f can be viewed as a multilinear polynomial over the variables $x_1, \ldots, x_n$, and we define the degree of f as

$$\deg f = \max_{S \mid \hat{f}(S) \neq 0} |S|.$$

For any coordinate $i \in [n]$ we define its influence on a Boolean function $f\colon \{-1,1\}^n \to \{-1,1\}$ as the probability that changing the value of that coordinate will change the value of f, i.e.

$$\mathrm{Inf}_i(f) = \mathbf{P}(f(X) \neq f(X^{(i)})),$$

where $X^{(i)}$ is obtained from X by flipping the i:th coordinate. Note that for a *dictator* function $\mathrm{DICT}_{n,i}(x) := x_i$ exactly one coordinate has influence 1 while the others have influence 0. For the *majority* function $\mathrm{MAJ}_n(x) := \mathrm{sgn}(\sum_{i=1}^n x_i)$,which is defined for odd n, one can show that each coordinate has influence $\Theta\left(\frac{1}{\sqrt{n}}\right)$. Thinking of the functions as social choice functions, that given n voters' preferences between two candidates determines the winning candidate it is natural to ask which function minimizes the most influential voter. One candidate is the *tribes* function TRIBES_n which, rougly, divides the n variables into $\Theta\left(\frac{n}{\log n}\right)$ groups of size $\Theta(\log n)$ and returns $+1$ if and only if all variables in one group is $+1$. For this function the influence of every variable is $\Theta\left(\frac{\log n}{n}\right)$. The KKL theorem [9] showed that this is asymptotically optimal:

THEOREM 2.1 (KKL). *For any $f\colon \{-1,1\}^n \to \{-1,1\}$ there exists an $i \in [n]$ such that*

$$\mathrm{Inf}_i(f) \geq \Omega\left(\mathbf{Var}(f)\frac{\log n}{n}\right).$$

The definition of influences can also be extended to functions taking real values, $f\colon \{-1,1\}^n \to \mathbb{R}$, in the following way:

$$\mathrm{Inf}_i(f) = \mathbf{E}[\mathbf{Var}[f(X) \mid X_1, \ldots, X_{i-1}, X_{i+1}, \ldots X_n]].$$

2.1.1 The Invariance Principle

The invariance principle of [13], which can be seen as a generalization of central limit theorems to multilinear polynomials, is a very useful tool for mapping discrete problems to continuous problems. In its simplest form, it states that if $f\colon \{-1,1\}^n \to \mathbb{R}$ is of low degree and each coordinate has small influence on f, then the distribution of $f(X)$ will not change by much if we replace the X_i's in the multilinear polynomial expansion given by (2.1) by i.i.d. standard Gaussians $Z_i \sim \mathrm{N}(0,1)$. The change of the distribution is measured by an arbitrary $\mathcal{C}^3$ function Ψ having bounded third order derivatives.

THEOREM 2.2. *([13], special case of Theorem 3.18)*
Suppose $X_1, \dots, X_n$ *are i.i.d. uniform on* $\{-1, 1\}$*,* $f\colon \{-1,1\}^n \to \mathbb{R}$ *has* $\deg f \le d$ *and* $\mathrm{Inf}_i f \le \tau, \forall i$*. Let* $\Psi : \mathbb{R} \to \mathbb{R}$ *be a* $\mathcal{C}^3$ *function with* $|\Psi^{(\mathbf{r})}| \le B$ *for* $|\mathbf{r}| = 3$*. Then,*

$$\left| \mathbf{E}\,\Psi(f(X)) - \mathbf{E}\,\Psi\left(\sum_{S\subseteq[n]} \hat{f}(S) \prod_{i\in S} Z_i\right)\right| \le B10^d\tau,$$

where $Z_1, \dots Z_n$ *are i.i.d* $\mathrm{N}(0,1)$*.*

The theorems in [13] and [11] are much more general. For example

- The underlying probability space is generalized to an arbitrary finite product space $(\Omega, \mu) = (\prod_{i=1}^n \Omega_i, \prod_{i=1}^n \mu_i)$ where $|\Omega_i| < \infty, \forall i$. Functions $f\colon \Omega \to \mathbb{R}$ can still be written as a multilinear polynomial by constructing an orthonormal basis $\mathcal{X}_i = (\mathcal{X}_{i,0} = 1, \mathcal{X}_{i,1}, \dots, \mathcal{X}_{i,|\Omega_i|-1})$ for the space of functions $\Omega_i \to \mathbb{R}$ equipped with the inner product $\langle f, g\rangle = \mathbf{E}_{X\sim\mu}[f(X_i)g(X_i)]$ and expressing f as
$$f(x) = \sum_{\sigma} \hat{f}(\sigma) \prod_{i=1}^{n} \mathcal{X}_{i,\sigma_i}(x),$$
where the sum is over all tuples $\sigma = (\sigma_1, \dots, \sigma_n)$ such that $0 \le \sigma_i < |\Omega_i|$.

- Multidimensional functions $f\colon \Omega \to \mathbb{R}^k$ can be handled similarly using a test function $\Psi : \mathbb{R}^k \to \mathbb{R}$.

In Paper I a few more generalizations that are useful in applications are introduced:

- the $\mathcal{C}^3$ restriction on Ψ in Theorem 2.2 is removed and replaced with a Lipschitz continuity requirement.

- non-orthonormal bases for the functions spaces $\Omega_i \to \mathbb{R}$ are handled (this was also discussed in [11]).

2.2 Social Choice Theory

Social choice theory is the study of methods of collective decision making. An important example, which is studied in this thesis, are elections where n voters choose between q candidates. In this case the method of decision making can be described by a social choice function which maps the preferences of all voters to a preference of the whole society. Depending on the setting, a *preference* could be either a single preferred candidate, a linear ordering of all candidates, or some more general binary relation on the set of candidates.

Assuming that the candidates are numbered $1, \dots, q$, we let $[q]$ denote the set $\{1, \dots, q\}$, L_q denote the set of all $q!$ linear orderings of the set $[q]$ and G_q denote the set of all $2^{\binom{n}{2}}$ tournaments on the set $[q]$, i.e. binary relations such that for every pair of candidates $a, b \in [q]$, either $a\, G_q\, b$ or $b\, G_q\, a$, indicating whether a is preferred over b or vice versa.

2.2.1 Transitivity

Here we will assume that the voters present linear orderings on the candidates which the social choice functions maps to a tournament describing the society's preference between every pair of candidates. Such a social choice function $f\colon L_q^n \to G_q$ is said to be

- Independent of Irrelevant Alternatives (IIA), if the society's preference between any two candidates only depends on the voters' relative preferences between these two candidates.
- unanimous, if for any two candidates $a, b \in [q]$, the society prefers a to b whenever all voters prefer a to b.
- transitive, if the society's preference is always a linear order, i.e. $f(x) \in L_q, \forall x$.
- neutral, if it is invariant under renumberings of the candidates.
- a dictator, if there is a voter $i \in [n]$ such that the society's preference only depends on that voter's preference.

One of the first results in Social Choice Theory is Condorcet's paradox which applies to the function PAIRWISE-MAJ_n for which the society's preference between every two candidates is determined by the majority among the voters' preferences between those two candidates. Even for this simple and natural function, there can be situations in which the society's preference is cyclic, e.g. the society prefers candidate a over b, b over c and c over a. Thus, this particular social choice function is not transitive.

Arrow's Theorem states that this situation cannot be avoided unless we sacrifice one of the arguably desirable properties IIA, unanimity or non-dictatorship.

THEOREM 2.3 (Arrow's Theorem, [3, 4]). *Any social choice function $f\colon L_q^n \to G_q$ which is IIA, unanimous and transitive must be a dictator.*

A related, but weaker requirement than transitivity is that of having a Condorcet winner. We say that a social choice function $f\colon L_q^n \to G_q$

- has a Condorcet winner on input x, if the society's preference $f(x)$ has one candidate which is preferred over all the other candidates.

In paper I we show that the function $\mathrm{PAIRWISE\text{-}MAJ}_n$ maximizes the probability of having a Condorcet winner among social choice functions $f\colon L_q^n \to G_q$ that are IIA, neutral and low-influential (in that every voter has a small probability of being able to change the outcome), when all voters' preferences are assumed to be selected uniformly at random.

2.2.2 Manipulability

Another desirable property of a social choice function is that of non-manipulability. Here we consider social choice functions $f\colon L_q^n \to [q]$ where the society selects a single winning candidate. We say that such a function is

- manipulable, if there are situations in which a voter i, who knows the votes of all other voters, has an incentive to change his vote, in that by doing so the society's preference will change to a candidate which voter i originally preferred more than the candidate that would win if voter i voted according to his true preference.

An example of a non-manipulable function is a dictator which always selects the top preference of a specific voter.

The *Gibbard-Satterthwaite Theorem* shows that manipulability cannot be avoided when there are at least three candidates, unless the social choice function is a dictator or never selects any but two of the candidates:

THEOREM 2.4 (Gibbard-Satterthwaite, [8,16]). *Any social choice function $f\colon L_q^n \to [q]$ which takes on at least three values and is not a dictator is manipulable.*

In paper II we show a quantitative version of this theorem.

2.2.3 Noise Stability

Consider a simple election with n voters choosing between q candidates of which only one can be selected by the society. The method of decision making can then be described by a social choice function $f\colon [q]^n \to [q]$.

The noise stability of such functions measures the stability of the output when the votes are chosen independently and uniformly at random, and then re-randomized with probability $1-\rho$. In social choice this can for instance be used to model the situation where a certain fraction of the votes are incorrectly counted.

DEFINITION 2.5. *For $\rho \in [0,1]$, the noise stability of $f\colon [q]^n \to [q]$ is*

$$\mathbb{S}_\rho(f) = \mathbf{P}(f(\omega) = f(\lambda)),$$

where ω is uniformly selected from $[q]^n$ and each λ_i is independently selected using the conditional distribution

$$\mu(\lambda_i|\omega_i) = \rho 1_{\{\lambda_i=\omega_i\}} + (1-\rho)\frac{1}{q}.$$

We say that a social choice function $f\colon [q]^n \to [q]$ is balanced if for any j, $\mathbf{P}(f(\omega) = j) = \frac{1}{q}$ when $\omega \in [q]^n$ is chosen uniformly at random.

It is natural to require that a social choice function has low influence in each coordinate, so that a single voter has a small chance of changing the outcome of the election. This avoids dictatorships and other functions that mainly depend on a few voters. Another natural requirement is for the function to be as noise stable as possible, so that even if a small fraction of the votes are miscounted the result is unlikely to change.

It is conjectured that for balanced low-influential functions $f\colon [q]^n \to [q]$, this noise stability is maximized by the plurality function $\mathrm{PLUR}_{n,q}$, which selects the most popular candidate (ties broken arbitrarily):

CONJECTURE 2.6 (Plurality is Stablest).
For any $q \geq 2$, $\rho \in [0,1]$ and $\epsilon > 0$ there exists a $\tau > 0$ such that if $f\colon [q]^n \to [q]$ is a balanced function with $\mathrm{Inf}_i(1_{\{f(\cdot)=j\}}) \leq \tau$, $\forall i \in [n], j \in [q]$, then

$$\mathbb{S}_\rho(f) \leq \lim_{m\to\infty} \mathbb{S}_\rho(\mathrm{PLUR}_{m,q}) + \epsilon.$$

The special case for $q = 2$, the *Majority is stablest* theorem, was proved in [13].

In paper I the invariance principle is used to show that the Plurality is Stablest conjecture is equivalent to a conjecture on *Gaussian noise stability*.

2.3 Hardness of Approximation

2.3.1 Introduction to computational complexity theory

In computational complexity theory, one is interested in the asymptotics of the amount of time (or space) required to compute discrete functions. For simplicity we will assume that all combinatorial objects used (numbers, sets, graphs, formal mathematical proofs etc.) are represented as binary strings, i.e. elements in $\Sigma^* = \bigcup_{n\in\mathbb{N}}\{0,1\}^n$. The exact encoding used for different objects does not matter for our purposes (as long as it is a reasonable one). The length of a string $x \in \Sigma^*$ is denoted by $|x|$.

In general a *computational problem* is defined by a function $f\colon \Sigma^* \to \Sigma^*$. A *decision problem* is a problem which can be answered by *yes* or *no*. For instance,

- 3-COLOR: given a graph, can the vertices be colored using 3 colors such that no neighboring vertices have the same color?

- TRUE$_\Gamma$: given a proposition T in a formal mathematical theory Γ and an empty proof consisting of n zeroes [1] , does there exist a formal proof of T of length at most n?

By identifying a given decision problem with the subset of all strings (usually called *language*) for which the answer is *yes* we may define decision problems as follows:

DEFINITION 2.7. *A* decision problem L *is a subset of* Σ^*.

The complexity class P consists of all decision problems that can be computed in polynomial time (on any (and thus all) universal Turing machines, which the reader may think of as a regular computer equipped with unlimited amount of memory). If an algorithm's running time is bounded above by a polynomial in the length of the input (for some fixed universal Turing machine) we say that it is a polynomial time algorithm.

DEFINITION 2.8. *The complexity class P consists of all decision problems L for which there exists a polynomial time algorithm A such that*

$$\begin{cases} x \in L \Rightarrow A(x) = yes \\ x \notin L \Rightarrow A(x) = no \end{cases}$$

The complexity class NP consists of all decision problems for which *yes*-instances have proofs that can be verified in polynomial time.

DEFINITION 2.9. *The complexity class NP consists of all decision problems L for which there exists a polynomial q and a polynomial time algorithm (verifier) V such that*

$$\begin{cases} x \in L \Rightarrow \exists \Pi \in \Sigma^* : |\Pi| \leq q(|x|) \text{ and } V(\Pi) = \textit{yes} \\ x \notin L \Rightarrow \forall \Pi \in \Sigma^* : V(\Pi) = \textit{no} \end{cases}$$

Note that both 3-COLOR and TRUE$_\Gamma$ are in NP. For instance, for 3-COLOR the verifier V can be taken to be an algorithm that simply checks that Π is a string that describes a coloring of all vertices in the graph in a way such that no neighboring vertices have the same color. Clearly, such a Π exists if and only if $x \in$ 3-COLOR.

Further, P $\subseteq$ NP, since for $L \in$ P we can simply ignore the proof Π and use the algorithm A as verifier. It remains an open problem whether P $=$ NP, although equality would be very surprising (implying e.g. that proofs of mathematical theorems can be found in time polynomial in the length of the statement and the length of the proof).

In hardness of approximation one is interested in showing non-existence of polynomial time algorithms for approximating combinatorial optimization problems (assuming P $\neq$ NP). Let us first define combinatorial optimizations problems.

[1] The reason that we include an empty proof of length n in the instance and not just the number n is that the number n is encoded by a string of length $\Theta(\log(n))$ but we later want a polynomial in the length of the instance to be polynomial in n.

DEFINITION 2.10. *A combinatorial maximization problem is defined by a function $f\colon \Sigma^* \times \Sigma^* \to \mathbb{R} \cup \{-\infty\}$ assigning a value $f(x,l)$ to any solution l of an instance x such that for each x, there are only a finite number of solutions l (called* **feasible** *for x) for which $f(x,l) \neq -\infty$.*
An instance x is said to be valid if it has a feasible solution l.
The optimal value for an instance $x \in \Sigma$ is

$$\mathrm{OPT}(x) = \max_l f(x,l).$$

A minimization problem is defined similarly by replacing the $\max$ by $\min$ and $-\infty$ by $+\infty$.

We can now define the corresponding complexity classes PO and NPO.

DEFINITION 2.11. *The complexity class NPO consists of all combinatorial optimization problems f for which there exist*

i) a polynomial time algorithm that determines whether an instance x is valid,

ii) a polynomial q such that for any instance x, all feasible solutions l satisfy $|l| \leq q(|x|)$, and

iii) a polynomial time algorithm that computes f.

PO is the subset of NPO for which $\mathrm{OPT}(x)$ is computable by a polynomial time algorithm.

There is a natural pre-ordering of computational problems given by polynomial time reducibility.

DEFINITION 2.12. *Given two computational problems X and Y, we say that X is polynomial time reducible to Y, denoted $X \leq_P Y$, if there exists a polynomial time algorithm A which computes the value of instances $x \in X$ in polynomial time, given access to an oracle for Y (i.e. a hypothetical algorithm that computes Y in constant time).*

From this we may define the complexity classes NP-complete consisting of the hardest problems in NP and NP-hard consisting of all problems that are at least as hard as NP. More generally,

DEFINITION 2.13. *Let $\mathcal{C}$ be a complexity class. Then $\mathcal{C}-\mathrm{hard}$ consists of all computational problems Y such that $X \leq_P Y, \forall X \in \mathcal{C}$. Further, $\mathcal{C}-\mathrm{complete} = \mathcal{C}-\mathrm{hard} \cap \mathcal{C}$*

2.3.2 Approximation algorithms

Many NP-hard optimization problems (for which no polynomial time algorithm exists unless $P = NP$) are possible to approximate within a constant factor in polynomial time. For instance, for the Euclidean Traveling Salesman Problem where one is given a set of points in Euclidean space, computing the shortest round-trip route visiting all points is NP-hard. However, for any $\epsilon > 0$ there exist a polynomial time approximation algorithm that computes a route no more than $1+\epsilon$ times longer than the optimal route.

DEFINITION 2.14. *If $f\colon \Sigma^* \times \Sigma^* \to \mathbb{R}^+ \cup \{-\infty\}$ is a maximization problem in NPO, A is an algorithm and $r \in [0, 1)$, we say that A is an r-approximation algorithm for f if for all valid instances x,*

$$f(x, A(x)) \geq r\,\mathrm{OPT}(x).$$

Similarly, if $f\colon \Sigma^* \times \Sigma^* \to \mathbb{R}^+ \cup \{\infty\}$ is a minimization problem and $r > 1$ we say that A is an r-approximation algorithm for f if for all valid instances x,

$$f(x, A(x)) \leq r\,\mathrm{OPT}(x).$$

Thus, for any $\epsilon > 0$, the Euclidean Traveling Salesman Problem has a polynomial time $1+\epsilon$-approximation algorithm.

Other problems can only be efficiently approximated up to a certain approximation constant. For instance, consider MAX-3-SAT defined as

DEFINITION 2.15. *An instance of the MAX-3-SAT problem consists of m clauses, each being a disjunction (logical or) of at most three literals, where each literal is either a variable or the negation of a variable from a set of n Boolean variables $b_1, \ldots, b_n$. A feasible solution is an assignment $l : [n] \to \{0, 1\}$ to these variables. The value $f(x, l)$ of an assignment is the fraction of clauses that are satisfied by the assignment.*

For MAX-3-SAT there exist a $\frac{7}{8}$ approximation algorithm based on semidefinite programming [18]. For the restricted problem MAX-E3-SAT, where we require that each clause contains exactly three (different) variables, this can be achieved by picking a random assignment which will satisfy a $\frac{7}{8}$ fraction of the clauses in expectation. Note that this algorithm can easily be derandomized to get a deterministic polynomial time algorithm that is guaranteed to find an assignment satisfying at least a $\frac{7}{8}$ of the clauses. This is done in the following way, using the *method of conditional expectation*: On after another, set each variable to the value which maximizes the conditional expectation over the remaining variables.

On the other hand it is known [17] that no $\frac{7}{8} + \epsilon$ polynomial time approximation can be achieved (unless $P = NP$), for any $\epsilon > 0$.

$$(b_1 \vee \neg b_2 \vee b_4) \wedge (\neg b_1 \vee \neg b_3 \vee b_2) \wedge (\neg b_2 \vee b_3 \vee b_5)$$

Figure 2.1: ***A MAX-E3-SAT instance. All 3 clauses can be satisfied simultaneously so the optimal value is 1.***

MAX-3-SAT is an example of class of optimization problems called Constraint Satisfaction Problems (CSP's). The maximization version of CSP's can be defined as follows:

DEFINITION 2.16. ***A MAX-CSP*** $\Lambda = (\mathcal{P}, q)$ ***is specified by a set of predicates*** $\mathcal{P}$ ***over the finite domain*** $[q]$***. The arity of*** Λ ***is the maximal arity of the predicates in*** $\mathcal{P}$***.***
An instance $\mathcal{I}$ ***of*** Λ ***consists of a set of variables*** $b_1, \dots, b_n$ ***and a set of constraints, each formed by a predicate from*** $\mathcal{P}$ ***applied to a subset of the variables and their negations. The optimal value*** $\mathrm{OPT}(\mathcal{I})$ ***for*** $\mathcal{I}$ ***is the maximum number of constraints satisfied by any assignment*** $b \in [q]^n$***.***

Thus, MAX-3-SAT, which can be described as a MAX-CSP $(\mathcal{P}, q)$ **with** $q = 2$ **and** $\mathcal{P} = \{x_1 \vee x_2 \vee x_3\}$**, is a ternary MAX-CSP over a Boolean domain.**

2.3.3 The PCP Theorem and the Unique Games Conjecture

The $\frac{7}{8}+\epsilon$ **inapproximability result for MAX-3-SAT (and similar results for other MAX-CSP's) is obtained by a reduction from a standard problem called the Label Cover problem for which arbitrarily good inapproximability results exist.**

DEFINITION 2.17.
An instance of the **Label Cover** ***problem,*** $\mathcal{L}(V, W, E, M, N, \{\sigma_{v,w}\}_{(v,w)\in E})$***, consists of a bipartite graph*** $(V \cup W, E)$ ***with a function*** $\sigma_{v,w} : [M] \to [N]$ ***associated with every edge*** $(v, w) \in E \subseteq V \times W$***. A labeling*** $l = (l_V, l_W)$***, where*** $l_V : V \to [M]$ ***and*** $l_W : W \to [N]$***, is said to satisfy an edge*** (v, w) ***if***

$$\sigma_{(v,w)}(l_W(w)) = l_V(v).$$

The value of a labeling l***,*** $\mathrm{VAL}_l(\mathcal{L})$***, is the fraction of edges satisfied by*** l ***and the optimal value for*** $\mathcal{L}$ ***is the maximal fraction of edges satisfied by any labeling,***

$$\mathrm{OPT}(\mathcal{L}) = \max_l \mathrm{VAL}_l(\mathcal{L}).$$

The PCP (Probabilistically Checkable Proofs) theorem [1,2] asserts that the Label Cover problem is NP-hard to approximate within any constant $\epsilon > 0$**, for suitable choices of** M **and** N**.**

THEOREM 2.18 (Label Cover version of the PCP Theorem). ***For any*** $\epsilon > 0$ ***there exist*** M ***and*** N ***such that it is NP-hard to distinguish between instances*** $\mathcal{L}$ ***of the Label***

Cover problem with label set sizes M and N having $\mathrm{OPT}(\mathcal{L}) = 1$ from those having $\mathrm{OPT}(\mathcal{L}) \leq \epsilon$.

This implies that any problem in NP (for instance TRUE_Γ) has a probabilistically checkable proof, which can be verified by looking only at a constant (depending on ϵ, but not on the length of the instance $|x|$) number of bits in such a way that a false proof is accepted with probability ϵ while a correct proof is always accepted. The proof structure is given by the polynomial time reduction from the NP problem to a Label Cover problem for which a correct proof (assignment) satisfies all edges while any other (incorrect) proof satisfies at most an ϵ fraction of the edges.

However, the PCP theorem is not strong enough to give sharp inapproximability results for binary MAX-CSP's (2-CSP's). One promising direction forward is the Unique Games Conjecture (UGC), a strengthened form of the PCP Theorem introduced by Khot [10].

DEFINITION 2.19. *A* Label Cover*problem $\mathcal{L}(V, W, E, M, N, \{\sigma_{v,w}\}_{(v,w)\in E})$ is called* unique *if $M = N$ and each $\sigma_{v,w} : M \to M$ is a permutation.*

CONJECTURE 2.20 (Unique Games Conjecture). *For any $\eta, \gamma > 0$ there exists $M = M(\eta, \gamma)$ such that it is NP-hard to distinguish instances $\mathcal{L}$ of the Unique Label Cover problem with label set size M having $\mathrm{OPT}(\mathcal{L}) \geq 1-\eta$ from those having $\mathrm{OPT}(\mathcal{L}) \leq \gamma$.*

It was recently shown [15] how to obtain optimal approximation algorithms for any MAX-CSP including 2-CSP's assuming the Unique Games Conjecture. However, the optimal approximation constants in [15] are generally not very explicit but given as the optima of certain optimization algorithms whose running time is doubly exponential in $1/\epsilon$. where ϵ is the desired precision.

It should be noted that, although many hardness of approximation results have been based on the Unique Games Conjecture, it is still not known whether this conjecture holds.

3

Summary of Papers

3.1 Paper I

Maximally Stable Gaussian Partitions with Discrete Applications.
Marcus Isaksson and Elchanan Mossel.

Gaussian noise stability measures the stability of partitions of Gaussian space under noise. In the simplest form we have two jointly standard Gaussian vectors X and Y in $\mathbb{R}^n$, with a covariance matrix $\mathbf{Cov}(X,Y) = \mathbf{E}[XY^T] = \rho I_n$, i.e. the coordinate pairs (X_i, Y_i) are i.i.d. $\mathrm{N}\left(0, \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}\right)$. The stability of a subset A of $\mathbb{R}^n$ is defined to be the probability that both X and Y fall into A. Borell [5] proved that for sets of fixed Gaussian measure, half-spaces maximize this stability . For simplicity, in this overview, we restrict attention to balanced partitions, i.e. sets of Gaussian measure $\frac{1}{2}$.

THEOREM 3.1. *[5] Fix $\rho \in [0,1]$. Suppose $X, Y \sim \mathrm{N}(0, I_n)$ are jointly normal and $\mathbf{Cov}(X,Y) = \rho I_n$. Let $A \subseteq \mathbb{R}^n$ with $\mathbf{P}(X \in A) = \frac{1}{2}$. Then*

$$\mathbf{P}(X \in A, Y \in A) \leq \mathbf{P}(X \in H, Y \in H),$$

where $H = \{x \in \mathbb{R}^n | x_1 \geq 0\}$.

In paper I, two generalizations of this theorem are considered from which various applications in social choice and hardness of approximation are derived. The first generalization, also stated in a simplified form here, considers the probability of $k > 2$ correlated vectors falling into A. We prove that half-spaces are optimal even in this case:

THEOREM 3.2. *Fix* $\rho \in [0, 1]$. *Suppose* $X_1, \dots, X_k \sim \mathrm{N}(0, I_n)$ *are jointly normal and* $\mathrm{Cov}(X_i, X_j) = \rho I_n$ *for* $i \neq j$. *Let* $A \subseteq \mathbb{R}^n$ *with* $\mathbf{P}(X_i \in A) = \frac{1}{2}$. *Then*

$$\mathbf{P}(\forall i : X_i \in A) \leq \mathbf{P}(\forall i : X_i \in H),$$

where $H = \{x \in \mathbb{R}^n | x_1 \geq 0\}$.

Note that such a distribution on $X_1, \dots, X_k$ exists for any $\rho \in [0, 1]$ since it can be constructed by letting $X_i = \sqrt{\rho} Z_0 + \sqrt{1-\rho} Z_i$ where $Z_0, Z_1, \dots, Z_k$ are i.i.d. standard Gaussians.

We prove two applications of this theorem.

- First, in the context of social choice, we show that the function $\mathrm{PAIRWISE\text{-}MAJ}_n$ defined in Section 2.2.1 maximizes the probability of having a Condorcet winner among social choice functions $f\colon L_q^n \to G_q$ that are IIA, neutral and low-influential, if all voters' preferences are selected uniformly at random.

- The second application is in the context of *cosmic coin flipping [12, 14]* where k players receive noisy copies of the same n bits and want to agree on a single uniformly random bit without communicating. We show that each player applying the majority function on his received bits maximizes the probability of all players agreeing among low influence functions.

The second generalization considers a partition of $\mathbb{R}^n$ into $q > 2$ subsets (instead of just the two A and A^C), and asks for the probability that all k vectors fall into the same subset. We will still restrict attention to balanced partitions, i.e. into disjoint sets $A_1, \dots A_q \subseteq \mathbb{R}^n$ with equal Gaussian measure $\frac{1}{q}$.

It is conjectured that for $n \geq q - 1$ and $\rho \geq 0$ the most stable partition is a *standard simplex partition*, which divides $\mathbb{R}^n$ into q partitions depending on which of q maximally separated unit vectors are closest (ties may be broken arbitrarily):

DEFINITION 3.3. *For* $n+1 \geq q \geq 2$, $A_1, \dots, A_q$ *is a standard simplex partition of* $\mathbb{R}^n$ *if for all* i

$$A_i = \{x \in \mathbb{R}^n | x \cdot a_i > x \cdot a_j, \forall j \neq i\},$$

where $a_1, \dots a_q \in \mathbb{R}^n$ *are* q *vectors satisfying* $a_i \cdot a_j = \begin{cases} 1 & \textit{if } i = j \\ -\frac{1}{q-1} & \textit{if } i \neq j \end{cases}$

The standard simplex partition is also conjectured to be the least stable partition when $\rho < 0$.

When $n \geq q$ a standard simplex partition can be formed by picking q orthonormal vectors $e_1, \ldots, e_q$, subtracting their mean and scaling appropriately, i.e.

$$a_i = \sqrt{\frac{q}{q-1}} \left(e_i - \frac{1}{q} \sum_{i=1}^{q} e_i \right)$$

and for $n = q - 1$ it is enough to project these vectors onto the $q - 1$ -dimensional space which they span.

When $q = 3$ the standard simplex partition, also known as the *standard Y partition* or the *peace sign partition*, is described in $\mathbb{R}^2$ by three half-lines meeting at an 120 degree angle at the origin (Figure 3.1) and in $\mathbb{R}^n$, where $n > 2$, it can be exemplified by taking the Cartesian product of the peace sign partition and $\mathbb{R}^{n-2}$.

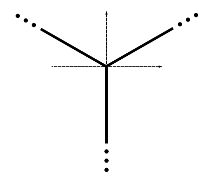

Figure 3.1: *The peace sign partition*

The conjecture, which is still open, can slightly simplified be stated as:

CONJECTURE 3.4. *Fix $\rho \in [0,1]$ and $3 \leq q \leq n+1$. Suppose $X, Y \sim \mathrm{N}(0, I_n)$ are jointly normal and $\mathrm{Cov}(X,Y) = \rho I_n$. Let $A_1, \ldots, A_q \subseteq \mathbb{R}^n$ be a balanced partition of $\mathbb{R}^n$. Then,*

$$\mathbf{P}((X,Y) \in A_1^2 \cup \cdots \cup A_q^2) \leq \mathbf{P}\left((X,Y) \in S_1^2 \cup \cdots \cup S_q^2\right), \qquad (3.1)$$

where $S_1, \ldots, S_q$ is a standard simplex partition of $\mathbb{R}^n$. Further, for $\rho \in [-1,0]$, (3.1) *holds in reverse:*

$$\mathbf{P}((X,Y) \in A_1^2 \cup \cdots \cup A_q^2) \geq \mathbf{P}\left((X,Y) \in S_1^2 \cup \cdots \cup S_q^2\right).$$

In the paper we derive two applications of this conjecture:

- First we show that it implies that *Plurality is Stablest* (Conjecture 2.6).
- We also show that it implies that an approximation algorithm by Frieze and Jerrum [7] for the optimization problem MAX-q-CUT (described in the next section) achieves the optimal approximation ratio assuming the unique games conjecture.

Since it is not known whether this conjecture holds and the standard simplex partition is optimal, it should be pointed out that one of the main contributions of paper I is to show that the optimality of certain discrete problems can be reduced to the question of finding optimal partitions with respect to Gaussian noise stability.

The main tool for the proofs of the applications is the invariance principle described in Section 2.1.1 and the paper also includes proofs for the generalizations described in that section.

3.1.1 MAX-q-CUT

In the MAX-q-CUT problem or Approximate q-Coloring, one is given a (possible edge weighted) graph and seeks a coloring of the vertices using q colors that minimizes the number of edges between nodes of the same color (i.e. maximizes the number of edges between different colors).

DEFINITION 3.5.
An instance of the weighted MAX-q-CUT problem, $\mathcal{M}_q(V,E,w)$***, consists of a graph*** (V,E) ***with a weight function*** $w : E \to [0,1]$ ***assigning a weight to each edge. A q-cut*** $l : V \to [q]$ ***is a partition of the vertices into q parts. The value of a q-cut*** l ***is***

$$\mathrm{VAL}_l(\mathcal{M}_q) = \sum_{(u,v)\in E: l(u)\neq l(v)} w_{(u,v)}.$$

The optimal value for $\mathcal{M}_q$ ***is***

$$\mathrm{OPT}(\mathcal{M}_q) = \max_l \mathrm{VAL}_l(\mathcal{M}_q).$$

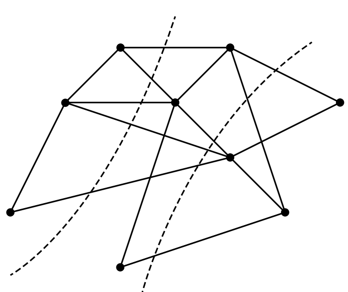

Figure 3.2: ***In MAX-3-CUT we want to find a partition of the vertices into 3 sets so as to maximize the weight of edges between different sets.***

Note that MAX-q-CUT is a (weighted) binary MAX-CSP over the alphabet $[q]$.

Frieze and Jerrum [7] devised a polynomial time approximation algorithm for MAX-q-CUT based on semi-definite programming, which for any fixed $\epsilon > 0$ achieves an approximation ratio of $\alpha_q - \epsilon$, where

$$\alpha_q = \inf_{-\frac{1}{q-1}\leq\rho\leq 1} \frac{q}{q-1}\frac{1-qI(\rho)}{1-\rho}.$$

Here, $qI(\rho)$ is the noise stability the standard simplex partition, i.e.

$$qI(\rho) = \mathbf{P}((X,Y) \in S_1^2 \cup \cdots \cup S_q^2),$$

where $X, Y \sim \mathrm{N}(0, I_{q-1})$ are jointly normal with $\mathbf{Cov}(X,Y) = \rho I_{q-1}$ and $S_1, \dots S_q$ is a standard simplex partition of $\mathbb{R}^{q-1}$. For instance, for $q = 3$ the approximation ratio can be made arbitrarily close to

$$\alpha_3 = \inf_{-\frac{1}{2} \le \rho \le 1} \frac{1 - \frac{9}{8\pi^2}(\arccos(-\rho)^2 - \arccos(\rho/2)^2)}{1-\rho} \approx 0.83601.$$

We show that Conjecture 3.4 implies that this is the optimal inapproximability constant for MAX-q-CUT assuming the unique games conjecture (and P $\neq$ NP):

THEOREM 3.6. ***Assume Conjecture 3.4, the UGC and P $\neq$ NP. Then, for any $\epsilon > 0$ no polynomial time algorithm exists that approximates MAX-q-CUT within $\alpha_q + \epsilon$.***

3.2 Paper II

The Geometry of Manipulation - a Quantitative Proof of the Gibbard Satterthwaite Theorem.
Marcus Isaksson, Guy Kindler and Elchanan Mossel.

In the second paper we give a quantitative version of Theorem 2.4, the Gibbard-Satterthwaite theorem. Recall that this theorem stated that any non-dictatorial social choice function $f\colon L_q^n \to [q]$ which takes on at least three values is manipulable. A consequence of this is that we shouldn't expect voters to always vote according to their true preferences in an election with three or more candidates. However, the theorem only guarantees that there is one set of orderings $x \in L_q^n$ for which some voter has an incentive to manipulate (such an x is called a manipulation point). Hence, this situation could be very unlikely.

In quantitative social choice some distribution on the voters' orderings is assumed and the probability of some undesirable situation is then estimated. In this paper we assume that the set of orderings X are selected uniformly at random from L_q^n and ask for the probability of X being a manipulation point. Letting

$$\mathrm{DICT} = \{f\colon L_q^n \to [q] \mid \exists i : f \textbf{ only depends on the } i\textbf{:th coordinate} \}$$

be the set of dictatorial functions, $\mathbf{D}(f,g) = \mathbf{P}(f(X) \neq g(X))$ be the fraction of $x \in L_q^n$ on which two functions $f, g : L_q^n \to [q]$ differ, and $\mathbf{D}(f, \mathrm{DICT}) = \min_{g \in \mathrm{DICT}} \mathbf{D}(f,g)$ the distance from f to the closest dictator, we show the following quantitative version of the Gibbard-Satterthwaite theorem under neutrality (recall that a function is neutral if it is invariant under renumbering of the candidates):

THEOREM 3.7. *Fix $q \geq 4$ and let $f\colon L_q^n \to [q]$ be a neutral social choice function with $\mathbf{D}(f, \mathrm{DICT}) \geq \epsilon$. Then,*

$$\mathbf{P}(f \textit{ is manipulable at } X) \geq \frac{\epsilon^2}{10^4 n^3 q^{30}},$$

where $X \in L_q^n$ is uniformly selected.

Note that although this probability bound is small, it is asymptotically (in q) much larger than the bound $\frac{1}{(q!)^n}$ that a direct application of the original Gibbard-Satterthwaite theorem gives.

In the paper we also prove a strengthened form of Theorem 3.7, which shows that the same bound holds if we only allow a small (linear in q) set of possible manipulations for a fixed voter and a given profile. This has implications for the approach of masking manipulability behind computational hardness; an algorithm that tries all such manipulations (linearly many in q) for a given voter will succeed finding a manipulation with non-negligible probability (at least $\epsilon^2/\mathrm{poly}(n, q)$).

3.3 Paper III

Approximating Linear Threshold Predicates.
Mahdi Cheraghchi, Johan Håstad, Marcus Isaksson and Ola Svensson.

Given a set of integer weights $w_1, \ldots, w_n$ with odd sum $W = w_1 + \ldots + w_n$ we can define a predicate $P : \{-1,1\}^n \to \{-1,1\}$ as

$$P(x) = \mathrm{sgn}\left(\sum_{i=1}^{n} w_i x_i\right)$$

where -1 is viewed as *false* and $+1$ as *true* so that the predicate is satisfied by x if and only if $P(x) = 1$. We call such a predicate a *homogenous linear threshold predicate*. Note that W being odd implies that the sum $\sum_{i=1}^{n} w_i x_i$ is never zero.

In the third paper we consider the hardness of approximating MAX-CSP's over homogenous linear threshold predicates. For such a predicate P the optimization problem MAX-CSP(P) can be defined as follows:

DEFINITION 3.8. *Fix a homogenous linear threshold predicate $P : \{-1,1\}^n \to \{-1,1\}$. An instance $\mathcal{I} = (m, N, l, s)$ of the optimization problem MAX-CSP(P) consists of N Boolean variables $x_1, \ldots, x_N$, and matrices $l \in N^{m\times n}$ and $s \in \{-1,1\}^{m\times n}$ describing m constraints. The i'th constraint is*

$$P(s_{i,1}x_{l_{i,1}}, \ldots, s_{i,n}x_{l_{i,n}}) = 1$$

The optimal value $\mathrm{OPT}(\mathcal{I})$ ***of*** $\mathcal{I}$ ***is the maximal fraction of constraints satisfied by any assignment*** $x \in \{-1,1\}^N$***.***

The simplest such predicate is the majority predicate $\mathrm{MAJ}_n = \mathrm{sgn}\left(\sum_{i=1}^n x_i\right)$, where n is odd, and for simplicity we will here describe the results for majority, letting MAX-MAJ-n denote the optimization problem MAX-CSP(MAJ_n). But first we need to define the *optimal approximation curve* which is a more refined characterisation of hardness of approximation than the optimal approximation constant defined in Definition 2.14:

DEFINITION 3.9. *If* $f\colon \Sigma^* \times \Sigma^* \to \mathbb{R}^+ \cup \{-\infty\}$ *is a maximization problem in NPO,* A *is an algorithm and* $s : \mathbb{R}^+ \to \mathbb{R}^+$*, we say that* A *is an* s*-approximation algorithm for* f *if for all valid instances* $\mathcal{I}$*,*

$$f(\mathcal{I}, A(\mathcal{I})) \geq s(\mathrm{OPT}(\mathcal{I})).$$

In the third paper we determine asymptotic upper and lower bounds on the optimal approximation curve $s^* : [0,1] \to [0,1]$ achievable by polynomial time algorithms for MAX-MAJ-n as $n \to \infty$, assuming the unique games conjecture and P $\neq$ NP. The bounds are almost tight in an asymptotic sense.

First note that the trivial algorithm which selects a random assignment will satisfy a fraction $\frac{1}{2}$ of the constraints on average. Further, this algorithm can be derandomized using the method of conditional expectation (see Section 2.3.2) to deterministically find an assignment satsfying at least a fraction $\frac{1}{2}$ of the constraints. Hence $s^*(c) \geq \frac{1}{2}$ for all c. Our first result is that this is optimal for $c < 1 - \frac{1}{n+1}$.

THEOREM 3.10. *Assume the UGC and P* $\neq$ *NP. Then,*

$$s^*(c) = \frac{1}{2}\,, \textit{ for } c < 1 - \frac{1}{n+1}.$$

Then we give an algorithm that does slightly better when $c > 1 - \frac{1}{n+1}$, i.e. on instances where all but a small fraction (less than $\frac{1}{n+1}$) of the clauses can be simultaneously satisfied. Note that this establishes a critical threshold at $c = 1 - \frac{1}{n+1}$ above which non-trivial approximation is possible.

THEOREM 3.11. *For* $c = 1 - \frac{\delta}{n+1}$*, where* $\delta < 1$*, we have*

$$s^*(c) \geq \frac{1}{2} + \Omega\left(\frac{1}{\sqrt{n}}\right)$$

where the hidden constant depends on δ*.*

Although this only gives a small advantage over the random assignment algorithm, we also show that it is asymptotically optimal when δ is bounded away from 1.

THEOREM 3.12. ***Assume the UGC and P $\neq$ NP. Then, for*** $c = 1 - \frac{\delta}{n+1}$***, where*** $\delta < 1$***, we have***

$$s^*(c) \leq \frac{1}{2} + (1-\delta)\mathcal{O}\left(\frac{1}{\sqrt{n}}\right).$$

In the paper we also generalize these results to other majority-like homogenous linear threshold predicates. It is well-known that any linear threshold predicate is determined by its so called Chow parameters [6] which are the Fourier coefficients of size at most 1:

$$\hat{P}(\emptyset), \hat{P}(\{1\}), \ldots, \hat{P}(\{n\}).$$

For a homogenous linear threshold predicate $P(x) = \operatorname{sgn}\left(\sum_{i=1}^n w_i x_i\right)$, we have $\hat{P}(\emptyset) = 0$, and we say that P is Chow-robust if the function does not change when the weights are replaced by the corresponding Chow parameters, i.e. if for all $x \in \{-1,1\}^n$:

$$\operatorname{sgn}\left(\sum_{i=1}^n w_i x_i\right) = \operatorname{sgn}\left(\sum_{i=1}^n \hat{P}(\{i\}) x_i\right)$$

Note that majority is Chow robust since all its Fourier coefficients of size 1 are equal. In the paper we derive a sufficient condition for Chow robustness and show that our results generalize to Chow robust predicates satisfying some additional technical conditions.

References

[1] Arora and Safra, *Probabilistic checking of proofs; A new characterization of NP (draft)*, FOCS: IEEE Symposium on Foundations of Computer Science (FOCS), 1992.

[2] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy, *Proof verification and intractability of approximation problems*, Proc. 33rd IEEE Symp. on Foundations of Computer Science, 1992.

[3] K. Arrow, *A difficulty in the theory of social welfare*, J. of Political Economy 58 (1950), 328–346.

[4] ______, *Social choice and individual values*, John Wiley and Sons, 1963.

[5] C. Borell, *Geometric bounds on the ornstein-uhlenbeck velocity process*, Probability Theory and Related Fields 70 (1985), no. 1, 1–13.

[6] C. K. Chow, *On the characterization of threshold functions*, Proceedings of the 2nd Annual Symposium on Switching Circuit Theory and Logical Design, 1961, pp. 34–38.

[7] A. Frieze and M. Jerrum, *Improved approximation algorithms for MAX-k-CUT and MAX-BISECTION*, Integer Programming and Combinatorial Optimization (Egon Balas and Jens Clausen, eds.), vol. 920, Springer, 1995, pp. 1–13.

[8] A. Gibbard, *Manipulation of voting schemes: a general result*, Econometrica 41 (1973), no. 4, 587ñ–601.

[9] J. Kahn, G. Kalai, and N. Linial, *The influence of variables on boolean functions*, Proceedings of the 29th Annual Symposium on Foundations of Computer Science, 1988, pp. 68–80.

[10] S. Khot, *On the power of unique 2-prover 1-round games*, STOC '02: Proceedings of the thirty-fourth annual ACM symposium on Theory of computing (New York, NY, USA), ACM, 2002, pp. 767–775.

[11] E. Mossel, *Gaussian bounds for noise correlation of functions*, Submitted, 2008.

[12] E. Mossel and R. O'Donnell, *Coin flipping from a cosmic source: On error correction of truly random bits*, Random Structures Algorithms 26 (2005), no. 4, 418–436. MR MR2139875

[13] E. Mossel, R. O'Donnell, and K. Oleszkiewicz, *Noise stability of functions with low influences: invariance and optimality*, To appear in Ann. Math., 2008.

[14] E. Mossel, R. O'Donnell, O. Regev, J. E. Steif, and B. Sudakov, *Non-interactive correlation distillation, inhomogeneous Markov chains, and the reverse Bonami-Beckner inequality*, Israel J. Math. 154 (2006), 299–336. MR MR2254545

[15] Prasad Raghavendra, *Optimal algorithms and inapproximability results for every CSP?*, Proceedings of the 40th Annual ACM Symposium on Theory of Computing, ACM, 2008, pp. 245–254.

[16] M. A. Satterthwaite, *Strategy-proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions*, J. of Economic Theory 10 (1975), 187–ñ217.

[17] J. Håstad, *Some optimal inapproximability results*, Journal of the ACM, 1997, pp. 1–10.

[18] U. Zwick, *Computer assisted proof of optimal approximability results*, Proc. of 13th SODA, 2002, pp. 496–505.

PAPER I

Maximally Stable Gaussian Partitions with Discrete Applications

Marcus Isaksson and Elchanan Mossel

(Submitted)

PAPER II

The Geometry of Manipulation - a Quantitative Proof of the Gibbard Satterthwaite Theorem

Marcus Isaksson, Guy Kindler and Elchanan Mossel

(Submitted)

PAPER III

Approximating Linear Threshold Predicates

Mahdi Cheraghchi, Johan Håstad, Marcus Isaksson and Ola Svensson

(Submitted)