Approximation algorithms and the hardness of approximation

Chandra Chekuri
(Dept. of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL),
Joseph Cheriyan
(Dept. of Combinatorics and Optimization, University of Waterloo, Waterloo, ON),
Ryan O'Donnell
(School of Computer Science, Carnegie Mellon University, Pittsburgh, PA),
Mohammad R. Salavatipour
(Dept. of Computing Science, University of Alberta, Edmonton, AL),
David Williamson
(School of Operations Research and Information Engineering, Cornell University, Ithaca, NY)

27 November 2011 – 2 December 2011

1 Overview of the Field

Most of the many discrete optimization problems arising in the sciences, engineering, and mathematics are NP-hard, that is, there exist no efficient algorithms to solve them to optimality, assuming the P≠NP conjecture. The area of approximation algorithms focuses on the design and analysis of efficient algorithms that find solutions that are within a guaranteed factor of the optimal one. Loosely speaking, in the context of studying algorithmic problems, an approximation guarantee captures the "goodness" of an algorithm – for every possible set of input data for the problem, the algorithm finds a solution whose cost is within this factor of the optimal cost. A hardness threshold indicates the "badness" of the algorithmic problem – no efficient algorithm can achieve an approximation guarantee better than the hardness threshold assuming that P≠NP (or a similar complexity assumption). Over the last two decades, there have been major advances on the design and analysis of approximation algorithms, and on the complementary topic of the hardness of approximation, see [33], [34].

The long-term agenda of our area (approximation algorithms and hardness results) is to classify all of the fundamental NP-hard problems according to their approximability and hardness thresholds. This agenda may seem far-fetched, but remarkable progress has been made over the last two decades. Approximation guarantees and hardness thresholds that "match" each other have been established for key problems in topics such as:

- covering and partitioning (the set covering problem [11]),
- algebra (overdetermined system of equations [14])
- graphs (clique, colouring [35]),
- combinatorial optimization (maximum cut [13],[17]),

- constraint satisfaction (maximum sat problems [13],[14]), etc.

Even more significant than these specific successes is the impact of the results and techniques from this area on related areas of mathematics. We list a few instances.

• Combinatorial Optimization: Combinatorial optimization is a mature body of research that has been developed by some of the leading researchers in discrete math over more than five decades; it has many deep mathematical results as well as many "real world" applications. The technique of iterative rounding has been developed in our area (starting with Jain [15]) to give remarkably good results for problems beyond the reach of classical combinatorial optimization. Recently, iterative rounding combined with the uncrossing technique has been used to give new proofs for several of the classic results in combinatorial optimization, including Edmonds' matching polyhedron theorem, which is one of the keystone results of the area; see the manuscript by Lau, Ravi and Singh [20].

• Metric Embeddings: Structure-preserving embeddings between various geometric spaces have been studied intensively for decades, in fields like differential geometry and functional analysis. Starting with the work of Linial, London, and Rabinovich [22], many applications of metric embeddings have been found in computer science, especially in the area of approximation algorithms. Moreover, the interaction between these fields has increased in recent years, with techniques developed in our area leading to the solution of open problems in non-linear functional analysis, e.g., the work of Brinkman and Charikar [9] solving the problem of dimension reduction in L_1, and the work of Arora, Rao, Vazirani [5] and Arora, Lee, Naor [4] leading to the near-solution of the Euclidean embedding problem for finite subsets of L_1.

• Analysis of Boolean Functions: Recent progress on hardness of approximation has come with the development of new tools in the area of "analysis of Boolean functions". This area combines techniques from harmonic analysis, probability theory, and functional analysis to study basic properties of Boolean functions. One recently developed tool, the Invariance Principle (Mossel-O'Donnell-Oleszkiewicz [25]), allows transfer of results from Gaussian probability spaces to results on Boolean probability spaces. This has led to fruitful connections between hardness of approximation and the geometry of Gaussian space. Although the original motivation for the Invariance Principle was proving matching hardness thresholds, it has also led to the solution of problems in other areas of mathematics and computer science. For example, in the mathematical theory of Voting and Social Choice, the Invariance Principle was used to prove: the "Majority Is Stablest Conjecture" and the "It Ain't Over Till It's Over Conjectures" regarding the optimality of Majority voting (Mossel-O'Donnell-Oleszkiewicz [25]); new results on the predictability of voting in many-party elections (Mossel [26]); and new quantitative bounds for Arrow's Theorem (Mossel [28]). Recently developed methods in the analysis of Boolean functions have also led to intractability results in the areas of property testing and learning theory, as well as positive results in the area of derandomization.

The goals of the workshop are as follows:

1. To bring together researchers in the fields of approximation algorithms (who work on finding algorithms with good approximation guarantees) and complexity theory (who work on finding hardness thresholds), and to stimulate the exchange of ideas and techniques between the two groups.

2. To highlight some of the new technical/mathematical directions in approximation guarantees (hierarchies of linear programming and semidefinite programming relaxations, uses of convex programming) and hardness thresholds (boolean functions, noise stability). These directions will be the subject of either survey or focus talks.

3. One specific topic on which we will focus is the status of the "Unique Games Conjecture", which states that unless P is equal to NP, there is no efficient algorithm to distinguish whether one can satisfy almost all constraints or almost no constraints of a particular type of constraint satisfaction problem (Khot [16]). It is not currently known whether Unique Games Conjecture is true. However, several quite exciting results have been shown by assuming the truth of the Unique Games Conjecture; in particular, matching hardness thresholds have been shown for several fundamental problems in combinatorial optimization, such as the maximum cut problem and the minimum vertex cover problem (Khot-Kindler-Mossel-O'Donnell [17], Khot-Regev [19]). There has been significant work on the Unique Games Conjecture both on the part of algorithms researchers and complexity theorists. The algorithms researchers are finding approximation algorithms for the Unique Games Problem; approximation algorithms with particular approximation guarantees will disprove the conjecture. The complexity theorists

have been using the Unique Games Conjecture to find additional hardness thresholds for fundamental problems. As part of the workshop, we will survey the current status of the conjecture, and devote time to approaches to resolve the conjecture, as well as further applications of it.

4. Another focus topic for the workshop is the hierarchy of LP (linear programming) and SDP (semidefinite programming) relaxations. The key idea here is to start from an LP relaxation of an NP-hard problem, and then obtain a series of tighter and tighter LP or SDP relaxations, by adding auxiliary variables and linear or semidefinite constraints. These methods were developed by Balas [7], Sherali and Adams [32], Lovasz and Schrijver [23], Lasserre [21], etc., and capture most of the known LP and SDP relaxations that have been exploited in the design of approximation algorithms. A stream of exciting research starting from the work of Arora, Bollobas and Lovasz [2], [3] has developed techniques to prove lower bounds on the approximation guarantees achievable by these methods. An important direction is to capture the relationship between hardness thresholds and the lower bound results on hierarchies of convex relaxations.

5. To include many younger researchers, and foster a relaxed interaction with established researchers. Our goal is to have a third of the workshop participants from this group.

6. To allow groups of Canadian researchers working in this area to meet, and either initiate or renew collaborations.

The workshop will present an opportunity to bring together some of the experts in related fields with the hope of initiating collaboration on some of the major open problems, and to explore the wider ramifications of the evolving body of powerful results and techniques from our area.

2 Recent Developments and Open Problems

The study of approximation algorithms and the hardness of approximation is one of the most exciting areas among researchers in theoretical computer science; every major conference in the field has several papers on these topics. Significant progress is being made. We give two examples of recent, dramatic innovations:

(1) Raghavendra [30] has recently shown that there is a fixed efficient semidefinite programming algorithm which achieves the best approximation guarantee for all constraint satisfaction problems – assuming the Unique Games Conjecture.

(2) Asadpour, Goemans, Madry, Oveis Gharan, and Saberi [6] recently achieved the first significant improvement on the approximation guarantee of the asymmetric traveling salesman problem in over twenty five years. Shortly after, Mömke and Svensson [24] obtained a 1.461-approximation algorithm, for the graphic TSP, based on a novel use of matchings. The analysis of this algorithm has been further improved by Mucha [29] who shows an approximation ratio of 13/9.

3 Presentation Highlights

3.1 Lift and Project

Konstantinos (Costis) Georgiou gave a tutorial presentation on the Lift and Project method, on the first day of the workshop. He introduced all popular lift-and-project systems deriving LP and SDP hierarchies, including the Lovasz-Schrijver LP and SDP system, the Sherali-Adams system, and the Lasserre system. The goal of his talk was to describe in a crystal-clear way the definitions of the systems, with many examples, as well as the proof of convergence to the integral hull – the proof illustrates some of the essential features of lift-and-project systems.

3.2 Routing in Undirected Graphs with Constant Congestion

The first plenary talk was given by Julia Chuzhoy, on a result that was acclaimed as a breakthrough by experts on the subject, [10]. The Maximum Edge Disjoint Paths Problem (MEDP) is a fundamental routing problem in Combinatorial Optimization. The input consists of an undirected graph $G = (V, E)$ and k node pairs $(s_1, t_1), (s_2, t_2), \ldots, (s_k, t_k)$. The goal is to maximize the number of node pairs that can be connected by edge-disjoint paths. Let n denote the number of nodes in the graph. The approximability of MEDP in undirected graphs is still not well understood: the best known approximation ratio is $O(\sqrt{n})$ while the best known hardness of approximation threshold is $\Omega(\log^{1/2-\epsilon} n)$, assuming that NP is not contained in randomized quasi-polynomial time. An impediment to improving the approximation ratio is that almost all algorithms are based on a multicommodity flow relaxation for the problem. A grid-like example shows that the maximum fractional (multicommodity) flow can be $\Omega(\sqrt{n})$ times the maximum integral flow even in planar graphs. This topological obstruction goes away if we consider a relaxation of MEDP, namely, the MEDP with congestion problem (MEDPwC): here we allow the paths for the pairs to use an edge up to c times for an integer c.

Chuzhoy presented an efficient randomized algorithm to route $\Omega(OPT/\text{polylog}(k))$ source-sink pairs with congestion at most 14, where OPT is the maximum number of pairs that can be simultaneously routed on edge-disjoint paths. The best previous algorithm that routed $\Omega(OPT/\text{polylog}(n))$ pairs required congestion $\text{poly}(\log \log n)$, and for the setting where the maximum allowed congestion is bounded by a constant c, the best previous algorithms could only guarantee the routing of $\Omega(OPT/n^{O(1/c)})$ pairs. The key technical result of the paper shows that a graph has a large "routing structure" if it has a large well-linked set; this is proved by embedding an expander graph in any graph that has a large well-linked set.

Chuzhoy's result resolves an important open problem, though much work still remains to be done to get a full understanding of MEDPwC and related problems.

3.3 The Traveling Salesman Problem

The Traveling Salesman Problem (TSP) is one of the most well-studied problems in combinatorial optimization. Given a set of cities $\{1, 2, \ldots, n\}$, and distances $c(i, j)$ for traveling from city i to j, the goal is to find a tour of minimum length that visits each city exactly once. An important special case of the TSP is the case when the distance forms a metric, i.e., $c(i, j) \le c(i, k) + c(k, j)$ for all i, j, k, and all distances are symmetric, i.e., $c(i, j) = c(j, i)$ for all i, j. If the distances are not symmetric, they are said to be asymmetric.

For thirty years, a $\frac{3}{2}$-approximation algorithm due to Christofides has been the best known approximation algorithm for the TSP when edge costs are symmetric and obey the triangle inequality. Only in the past year has there been any significant progress in improving the state-of-the-art in approximating this important problem, and we had an entire day of talks dedicating to discussing the progress that has been made.

A sub-case of the TSP is the *graphic* TSP. In this case, the distances $c(i, j)$ are derived from an undirected graph G given as input, with $\{1, 2, \ldots, n\}$ as its vertex set. The cost $c(i, j)$ is the number of edges on the shortest path between i and j in G.

In a plenary talk, Amin Saberi described his December 2010 result with his student Shayan Oveis Gharan and with Mohit Singh which gave a $(\frac{3}{2} - \epsilon)$-approximation algorithm for the graphic TSP, in which $\epsilon \approx 10^{-12}$, [27]. The analysis of the algorithms builds on a variety of ideas such as properties of strongly Rayleigh measures from probability theory, graph theoretical results on the structure of near minimum cuts, and the integrality of the T-join polytope from polyhedral theory. It is worth noting that this result also bounds the integrality gap of the Held-Karp linear programming relaxation (worst case ratio between integer and linear programming relaxation solution) of the graphic TSP by the same ratio.

In a second plenary talk, Ola Svensson described his 1.461-approximation algorithm (April 2011), jointly with Tobias Mömke, for the graphic TSP, based on a novel use of matchings, [24]. Traditionally, matchings have been used to add edges in order to make a given graph Eulerian, whereas their approach also allows for the removal of certain edges. For the TSP on graphic metrics (graph-TSP), the approach yields a 1.461-approximation algorithm with respect to the Held-Karp LP lower bound. For graph-TSP restricted to a class of graphs that contains degree three bounded and claw-free graphs, they show that the integrality gap of the Held-Karp relaxation matches the conjectured ratio 4/3. The framework allows for generalizations in a natural way and also leads to a 1.586-approximation algorithm for the traveling salesman path problem on graphic metrics where the start and end vertices are prespecified. Later, Mucha [29] improved on one part of the analysis (but not the algorithm) to achieve an approximation guarantee of $\frac{13}{9}$.

In the afternoon, we then had a sequence of three 30 minute talks about TSP-related problems. David Shmoys gave a talk about the traveling salesman path problem. In this problem, in addition to the usual input for the TSP, one is also given two cities s and t. The goal is to find the minimum-cost that starts at s, ends at t, and visits all the other cities in-between. For nearly 20 years, the best known approximation algorithm for this problem has been a $\frac{5}{3}$-approximation algorithm due to Hoogeveen, which is simply an adaptation of Christofides' algorithm. Shmoys described his work with his student Hyung-Chan An and with Robert Kleinberg to obtain a $\frac{1+\sqrt{5}}{2}$-approximation algorithm for the problem, [1].

Anke van Zuylen gave a talk on "A proof of the Boyd-Carr conjecture." A long-standing conjecture related to the TSP concerns the ratio of the cost of an optimal tour to the solution of a well-known linear programming relaxation of the TSP. It is known that this ratio is always at most $\frac{3}{2}$, and is for some instances at least $\frac{4}{3}$, but the exact worst-case ratio is unknown. Boyd and Carr considered the minimum-cost cycle cover problem; a tour is a cycle cover, but there can be cheaper cycle covers than a tour. Boyd and Carr conjectured that the ratio of the minimum-cost cycle cover to the linear programming relaxation for the TSP is always at most $\frac{10}{9}$. Van Zuylen described her proof of this conjecture in work together with Frans Schalekamp and David Williamson, [31].

Finally, Zachary Friggstad described his work on the asymmetric traveling salesman path problem with multiple salesmen, [12]. In addition to the input given for the traveling salesman path problem described above, there is a number k of salesman who can be sent to visit the cities. All k salesman start at s and end at t, and the union of their paths must visit all the cities. Friggstad described his bicriteria algorithm for this problem that in one extreme case uses k salesmen and has cost at most $O(k \log n)$ times the optimal, and in the other extreme uses $2k$ salesman and has a cost only $O(\log n)$ times the optimal.

3.4 Semidefinite Programming Hierarchies and the Unique Games Conjecture

David Steurer, in a plenary talk, presented his new results, jointly with Boaz Barak and Prasad Raghavendra, on the Unique Games Conjecture, [8].

Semidefinite programming (SDP) relaxation are a form of convex relaxation that found many uses in algorithms for combinatorial optimization. In the early 1990's several researchers proposed stronger forms of SDP relaxation known as SDP hierarchies. Steurer presented a new way of taking algorithmic advantage of these hierarchies to solve constraint satisfaction problems for 2-variable constraints such as Label-Cover, Max-Cut, and Unique-Games. Specifically, he described an algorithm based on an SDP-hierarchy that provides arbitrarily good approximation to all these problems in time $poly(n) * \exp(r)$, where r is the number of eigenvalues in the constraint graph larger than some constant threshold (depending on the accuracy parameter and type of constraint used). In particular, quasi-polynomial-time algorithms are obtained for instances whose constraint graph is hyper-contractive, as is the case for all the canonical "hard instances" for MAX-CUT and UNIQUE-GAMES. This result gives more reason to consider relatively low levels of an SDP hierarchy as candidate algorithms for refuting Khot's Unique Games Conjecture.

3.5 The Sliding Scale Conjecture From Intersecting Curves

In another plenary talk, Dana Moshkovitz discussed the Sliding Scale Conjecture; this conjecture was posed by Bellare, Goldwasser, Lund and Russell in 1993 and has been open since. It says that there are PCPs with constant number of queries, polynomial alphabet and polynomially small error. She showed that the conjecture can be proved assuming a certain geometric conjecture about curves over finite fields. The geometric conjecture states that there are small families of low degree curves that behave, both in their distribution over points and in the intersections between pairs of curves from the family, similarly to the family of all low degree curves.

3.6 Hardness of Approximating the Closest Vector Problem (with Pre-Processing)

The final plenary talk was given by Nisheeth Vishnoi, on results obtained jointly with Subhash Khot and Preyas Popat, [18].

In the Closest Vector Problem (CVP) one is given as input a basis B for a lattice and a target vector t, and the goal is to find the lattice point closest to t, say in the l_2 norm. This problem is NP-hard; the best approximation algorithm achieves roughly $2^{O(n)}$ while the best hardness is $2^{\log n / \log \log n}$. It is an outstanding open problem to close this gap.

An easier sounding version of this problem, motivated from cryptography, is the Closest Vector Problem with Pre-processing (CVPP): Here B can be pre-processed arbitrarily and the input consists just of t. Could CVPP be much easier than CVP? For instance, one can compute the shortest vector in the lattice generated by B, a NP-hard problem, for free. Indeed, it was shown by Ahoronov and Regev how to approximate CVPP to within about $\sqrt{n}$ factor!

Vishnoi first surveyed the approximability of lattice problems and then showed that, as far as the best known hardness of approximation results are concerned, they can almost close the gap between CVP and CVPP: CVPP is hard to approximate to within a factor of $2^{(\log n)^(1-\epsilon)}$ unless NP is contained in quasi-poly time.

4 Scientific Progress Made and Outcome of the Meeting

The schedule of the workshop provided ample free time for participants to work on joint research projects. A number of new research projects were initiated during the workshop, while some other researchers used the opportunity to continue to work on projects started earlier. The research talks, the plenary talks, and the tutorial (on lift-and-project methods) were very well received.

There is a growing number of new results that are obtained using the lift and project method. It appears that several interesting new results could be obtained using LP and SDP hierarchies, perhaps with running times that are slightly super-polynomial (such as quasi-polynomial). Some very recent results (such as results on directed Steiner trees) are evidences of this.

Some interesting questions in this direction were also raised at the end of the tutorial. In particular, a number of researchers have started looking into approximation algorithms for some of the classical problems (e.g., set cover) that beat the hardness thresholds obtained so far if one allows the algorithm to run in super-polynomial time. More specifically, they have devised a $(1-\epsilon)\ln n + O(1)$-approximation for set cover with running time $2^{\tilde{O}(n^{\epsilon})}$ that is based on exploiting the Lovasz-Schrijver hierarchy of relaxations.

One of the participants reports that lift-and-project methods gave an SDP with an integrality gap of $O(\log n)$ for the unsplittable flow problem on a path. This does not match the recent constant factor approximation, but it is another demonstration that hierarchies can be used to improve integrality gaps.

Another group of researchers have started working on a promising approach toward subexponential constant-factor approximations for Sparsest Cut. More specifically, they have a conjecture on the existence of certain random walks on vertex expander graphs whose stationary distribution is approximately uniform. One of their results – discovered at the BIRS Workshop – uses $n^{O(\delta)}$-rounds of the Lasserre SDP hierarchy to achieve a relaxation that appears to be significantly stronger than the Goemans-Linial SDP.

Following the plenary talk on the "sliding scale conjecture from intersecting curves" there were several discussions among workshop participants regarding what would be needed from the family of intersecting curves in order to prove the sliding scale conjecture. Discussions continued also shortly after the workshop, with additional participants. The outcome is a manuscript clarifying some of the obstacles that one would need to overcome. In particular, a certain property that may have appeared to be desirable for the family of intersecting curves, termed in the manuscript as "well mixed", was proved to conflict with the sliding scale conjecture. As a possible remedy, the manuscript suggests a notion of "well separated" intersecting curves, that (if it exists) may potentially resolve the sliding scale conjecture.

At the end, we mention that the above are only a few examples of the research progress made during or after the workshop, and there could be several other ongoing projects that started at the workshop.

Acknowledgment: It is a pleasure to thank the BIRS staff for their support; this contributed to the success of the workshop. In particular, we thank Wynne Fong, Brenda Williams, and Brent Kearney.

References

[1] H-C.An, R.Kleinberg, D.B.Shmoys: Improving Christofides' Algorithm for the s-t Path TSP. ArXiv:1110.4604: (2011)

[2] S.Arora, B.Bollobas, L.Lovasz: Proving Integrality Gaps without Knowing the Linear Program. FOCS 2002: 313-322

[3] S.Arora, B.Bollobas, L.Lovasz, I.Tourlakis: Proving Integrality Gaps without Knowing the Linear Program. Theory of Computing 2(1): 19-51 (2006)

[4] S.Arora, J.R.Lee, A.Naor: Fre'chet Embeddings of Negative Type Metrics. Discrete & Computational Geometry 38(4): 726-739 (2007)

[5] S.Arora, S.Rao, U.V.Vazirani: Expander flows, geometric embeddings and graph partitioning. J. ACM 56(2): (2009)

[6] A.Asadpour, M.Goemans, A.Madry, S.Oveis Gharan, A.Saberi: An $O(\log n/\log\log n)$-Approximation Algorithm for the Asymmetric Traveling Salesman Problem. Manuscript (2009)

[7] E.Balas: Disjunctive programming and a hierarchy of relaxations for discrete optimization problems, SIAM Journal on Algebraic and Discrete Methods 6:466–485 (1985)

[8] B.Barak, P.Raghavendra, D.Steurer: Rounding Semidefinite Programming Hierarchies via Global Correlation. ArXiv:1104.4680: (2011)

[9] B.Brinkman, M.Charikar: On the impossibility of dimension reduction in ℓ_1. J. ACM 52(5): 766-788 (2005)

[10] J.Chuzhoy: Routing in Undirected Graphs with Constant Congestion. ArXiv:1107.2554v1 (2011)

[11] U.Feige: A Threshold of $\ln n$ for Approximating Set Cover. J. ACM 45(4): 634-652 (1998)

[12] Z.Friggstad: Multiple Traveling Salesmen in Asymmetric Metrics. ArXiv:1112.2930v2 (2011)

[13] M.X.Goemans, D.P.Williamson: Improved Approximation Algorithms for Maximum Cut and Satisfiability Problems Using Semidefinite Programming. J. ACM 42(6): 1115-1145 (1995)

[14] J.Hastad: Some optimal inapproximability results. J. ACM 48(4): 798-859 (2001)

[15] K.Jain: A Factor 2 Approximation Algorithm for the Generalized Steiner Network Problem. Combinatorica 21(1): 39-60 (2001)

[16] S.Khot: On the power of unique 2-prover 1-round games. STOC 2002: 767-775

[17] S.Khot, G.Kindler, E.Mossel, R.O'Donnell: Optimal Inapproximability Results for MAX-CUT and Other 2-Variable CSPs? SIAM J. Comput. 37(1): 319-357 (2007)

[18] S.Khot, P.Popat, N.K.Vishnoi: $2^{\log^{1-\epsilon} n}$ Hardness for Closest Vector Problem with Preprocessing. ArXiv:1109.2176: (2011)

[19] S.Khot, O.Regev: Vertex cover might be hard to approximate to within 2-epsilon. J. Comput. Syst. Sci. 74(3): 335-349 (2008)

[20] L.Lau, R.Ravi, M.Singh, Iterative methods in Combinatorial Optimization. Manuscript of monograph, 2009.

[21] J.B.Lasserre: An Explicit Exact SDP Relaxation for Nonlinear 0-1 Programs. IPCO 2001: 293-303

[22] N.Linial, E.London, Y.Rabinovich: The Geometry of Graphs and Some of its Algorithmic Applications. Combinatorica 15(2): 215-245 (1995)

[23] L.Lovasz and A.Schrijver, Cones of matrices and set functions, and 0-1 optimization. SIAM Journal on Optimization 1:166–190 (1991)

[24] T.Mömke, O.Svensson: Approximating Graphic TSP by Matchings. FOCS 2011, ArXiv:1104.3090 (2011)

[25] E.Mossel, R.O'Donnell, K.Oleszkiewicz: Noise stability of functions with low influences: invariance and optimality. FOCS 2005: 21-30

[26] E.Mossel: Gaussian Bounds for Noise Correlation of Functions and Tight Analysis of Long Codes. FOCS 2008: 156-165

[27] S.Oveis Gharan, A.Saberi and M.Singh: A Randomized Rounding Approach to the Traveling Salesman Problem. FOCS 2011.

[28] E.Mossel: A Quantitative Arrow Theorem. arXiv:0903.2574v3 [math.PR] (2009)

[29] M.Mucha: (13/9)-approximation for graphic TSP. STACS 2012.

[30] P.Raghavendra: Optimal algorithms and inapproximability results for every CSP? STOC 2008: 245-254

[31] F.Schalekamp, D.P.Williamson, A.van Zuylen: A Proof of the Boyd-Carr Conjecture. ArXiv:1107.1628 (2011)

[32] H.D.Sherali and W.P.Adams, A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. SIAM Journal on Discrete Mathematics 3 (3):411–430 (1990)

[33] V.V.Vazirani, Approximation Algorithms. Springer-Verlag, Berlin, (2001)

[34] D.P. Williamson, D.B. Shmoys. The Design of Approximation Algorithms. Cambridge University Press, to be published in 2010.

[35] D.Zuckerman, Linear Degree Extractors and the Inapproximability of Max Clique and Chromatic Number, Theory of Computing Volume 3 Article 6: 103-128 (2007)