Algorithmic Spectral Graph Theory (in Fall 2014)

Final Program Report, May 2016

James R. Lee and Prasad Raghavendra (Organizing Chairs)

Summary

The field of **spectral graph theory** is based on a simple idea: The eigenvalues and eigenvectors of a graph's adjacency matrix (and other related matrices) contain information that cannot easily be gleaned from directly examining the graph's combinatorial structure. The reasons for this are classical, deep, and grounded in physical principles (thermodynamics, electromagnetism, and quantum mechanics). This connection carries **algorithmic** import because the spectral data can be computed very efficiently (both on a personal computer and on large-scale distributed systems). Since the spectral representation has a global character, it finds remarkable application in a range of optimization problems.

From its initial conception, the Algorithmic Spectral Graph Theory (SGT) program at the Simons Institute had a lofty and diverse set of goals. The scientific goals were largely reflected in the content of our three workshops and the Boot Camp, where we sought to explore and expand on the following four themes.

From theory to practice: Spectral graph methods form both a fundamental building block of algorithmic theory, and a central tool employed across an array of application areas (machine learning and scientific computing being the most prominent). We felt the need to respect both aspects by fostering interaction between the core theoretical members of our program and practitioners and applied researchers.

SDPs and algebraic geometry: Our sister program in the Fall semester explored the role of algebraic geometry in the theory of computation. We shared with them a strong common interest in exploring the strength of semi-definite programming (SDP), an optimization method that can be seen as a powerful generalization of spectral methods. Understanding the geometry of such programs is connected to one of the most compelling open questions in computational complexity: the Unique Games Conjecture.

Iterative methods and scientific computing: A significant wave of recent progress in spectral graph theory has taken place at the interface with scientific computing. Extremely fast solvers for Laplacian systems have led to remarkable algorithms that break long-standing barriers in algorithms for classical optimization problems (such as computing maximum flows and minimum cuts). Moreover, this connection has fomented a revolution in the application and understanding of iterative methods (such as interior point methods for linear programming).

Kadison-Singer and interlacing polynomials: Last but not least, our semester began on the heels of the solution of the famous Kadison-Singer conjecture (by Marcus, Spielman and Srivastava) in operator theory and the foundations of quantum mechanics. Their proof employs methods and intuition cultivated in spectral graph theory, combined with the theory of real-stable polynomials.

Fundamental progress—often exciting and unexpected—was made in each of these directions. Moreover, the diverse set of participants sparked new interactions that led to unanticipated developments in a number of related computational areas. Indeed, the SGT program involved two weekly seminars, a reading group, a boot camp, and three workshops; as one can see below, these gatherings were often inspiration for new collaborations. And finally, as a general byproduct of the environment at the Institute and the fact that it played host to a collection of brilliant and ambitious scientists, the participants naturally produced some groundbreaking research that did not fall neatly within the major themes of the program.

All in all, the SGT program led to publications in every major theoretical computer science conference (FOCS, STOC, SODA, CCC, SoCG, ITCS, ...). This body of work has been recognized so far by STOC 2015 and FOCS 2015 best paper awards, and a FOCS 2015 best student paper award. And, as one demonstration of the interaction between theoretical and applied research at the institute, the SGT program has so far produced four publications in NIPS (considered, along with ICML, as one of the two top international conferences in machine learning).

Research Highlights

One would be hard-pressed not to label the semester a resounding success. For at least a few months, the Simons Institute was the world epicenter for research in the theory of algorithms. We now review some highlights from the extensive list of accomplishments arising from the SGT program.

Convex optimization

Contemporary research in algorithmic spectral graph theory relies heavily on intimate connections with convex optimization. Moreover, general purpose principles for optimizing convex functions have frequently been discovered first in the spectral setting (largely related to fast algorithms for Laplacian linear systems and the application of such systems to combinatorial problems).
Some of the most groundbreaking work in this direction was performed by two of our research fellows, Yin Tat Lee and Aaron Sidford. As the program began, they had just completed a project described as the "biggest breakthrough in interior point methods in two decades."

During the program, they simplified, improved, and clarified this line of work, leading to their remarkable new paper, "Efficient inverse maintenance and faster algorithms for linear programming" [29], which stands as the state of the art in linear programming. The paper employs many insights from spectral graph theory and computational linear algebra.

In further monumental work carried out in part at the Simons Institute, Lee, Sidford, and UC Berkeley student Wong presented a faster algorithm for convex programming with a separation oracle, the first such improvement in over 25 years [55]. This leads to the fastest-known algorithms for fundamental optimization problems like minimizing a submodular function and semi-definite programming. For this, the authors received the Best Student Paper award at FOCS 2015.

As a prime example of the cross-area pollination the program encouraged, during the workshop on "Fast algorithms via spectral methods," long-term participant Lorenzo Orecchia presented joint work with his student Zeyuan Allen-Zhu on using first-order optimization and online learning techniques in the design of fast spectral algorithms. Inspired by this presentation, Berkeley professor Peter Bartlett (also a long-term participant in the program) and his Ph.D. student Walid Krichene began discussions during the workshop that led to the development of accelerated mirror descent methods for solving large-scale optimization problems in machine learning [47].

Semi-definite programming and sums-of-squares optimization

Semi-definite programs (SDPs) can be seen as combining the rich expressiveness of linear programs with the global geometric power of spectral methods. Given the interest from the concurrent program (on Algorithms and Complexity in Algebraic Geometry), understanding the computational efficacy of SDPs was a major focus of the SGT program.

Two program organizers, James R. Lee and Prasad Raghavendra, together with long-term participant David Steurer, discovered a proof that polynomial-size SDPs cannot solve NP-complete problems (like the Traveling Salesman Problem). Such unconditional lower bounds on a powerful computational model are rare in complexity theory. The main technical obstacle was to prove a lower bound on the *positive semi-definite rank* of an explicit matrix, a problem that featured prominently in the sister program as well. The resulting paper "Lower bounds on the size of semi-definite programming relaxations" [53] won the Best Paper award at the 2015 *ACM STOC* Symposium.

Their method of proof relates general SDPs to those arising from the sum-of-squares (SoS) SDP hierarchy. Pablo Parrillo (one of the discoverers of the SoS hierarchy) was a long-term participant in both the SGT and AG programs. After the semester, he remarked that "*In my view, one of the great (non-technical) outcomes of the program was the new genuine interactions between what were previously fairly disjoint research communities (namely, TCS and mathematical optimization). I expect many of these to deepen in the coming years.*"

While SDPs are not able to efficiently solve NP-complete problems, there was evidence that they might be well suited to tackling some basic computational tasks arising in machine learning. Indeed, this was one key component of the first workshop on "Semidefinite optimization, approximation, and applications." Workshop organizers Boaz Barak, Jon Kelner, and David Steurer demonstrated this in [16], at least in the theoretical sense (polynomial-time algorithms). In work begun during the SGT program, Steurer, along with visiting graduate students Sam Hopkins, Tselil Schramm and Jonathan Shi showed that SoS methods can be used as inspiration to obtain much faster *spectral* algorithms for the underlying problems [41].

The Kadison-Singer problem and TSP

The resolution of the Kadison-Singer conjecture referred to earlier did not take long to find a staggering algorithmic application: UC Berkeley Ph.D. student Nima Ahmadipouranari and long-term program participant Shayan Oveis Gharan first proved an elegant generalization of the Kadison-Singer conjecture to strongly Rayleigh measures [5]. Combined with a mastery of spectral graph techniques, they employed this generalization to give a $(\log\log n)^{O(1)}$-approximation to the value of the optimal tour for instances of the *asymmetric* traveling salesman problem (TSP) [4]. Asymmetric TSP is one of the few classical optimization problems whose approximability has remained largely a mystery. This work was selected for plenary presentation at the 2015 *IEEE FOCS* Symposium.

Long-term participant Nikhil Srivastava, in joint work with collaborators Adam Marcus and Dan Spielman, continued the development of their theory surrounding the Kadison-Singer problem and interlacing polynomials. In particular, they employed the framework of free probability theory [62] to demonstrate the existence of bipartite Ramanujan graphs of every degree and every number of vertices [63], finally resolving a long-standing open problem whose history and import are too broad to be adequately covered here.

Core topics in Laplacian systems and spectral graph theory

One of the major advances at the interface of spectral graph theory and numerical linear algebra is the development, starting with the seminal work of Spielman and Teng, of near-linear time solvers for diagonally dominant linear systems. A primary challenge is to extend these solvers to more general families of matrices (and, optimistically, to general positive semi-definite systems). A significant step was made by program fellow Yin Tat Lee in collaboration with his coauthors [49]; they developed new algorithms that yield near-linear time solvers for *connection* Laplacians, a generalization of Laplacians that arises in image and signal processing.

A key tool in the design of fast spectral algorithms involves *sparsification* of a graph while preserving its spectral structure. Batson, Spielman, and Srivastava (BSS) had already shown how to find linear-sized sparsifiers (indeed, this was an early precursor to later work on the Kadison-Singer problem), but it was unknown how efficiently they could be constructed. Allen-Zhu, Liao, and Orecchia [3] showed how general purpose online convex optimization methods could find BSS sparsifiers in quadratic time.

Visiting student Allen-Zhu presented this result in one of the weekly seminars. Around midnight that evening, office mates and institute fellows Yin Tat Lee and He Sun were walking back to their nearby residences when their discussion turned to the day's talk, and how one might construct BSS sparsifiers even faster. This marked the beginning of a joint project that led to a dramatic conclusion: a near linear-time algorithm for spectral graph sparsification [56].

Connections to applications

There are a number of prominent examples of applied work that was performed wholly or initiated at the Institute. During the workshop on "Spectral algorithms: From theory to practice," workshop participant Kevin Chen and long-term visitor Kamalika Chaudhuri began discussions about whether they could use spectral algorithms for learning parameters of more complex processes that arise in bioinformatics. In joint work with Zhang and Song, they design a spectral algorithm to learn properties of human genomes [80].

Research fellow Mihai Cucuringu describes his joint work with long-term participant Yiannis Koutis thus: "*I greatly enjoyed my stay at the Simons Institute. It gave me an opportunity to meet new and old collaborators, and start on several different projects. Perhaps most importantly, the constrained clustering project gave me the opportunity to learn about Laplacian linear system solvers from one of the experts in the field*." Cucuringu and Koutis showed how multi-way spectral partitioning algorithms developed by program participants can be used to obtain practical algorithms for clustering data sets with millions of points, hundreds of times faster than state of the art software (and, in addition, these new algorithms find better clusterings).

Long-term participant Marina Meila remarks, "*This program revived my interest in spectral graph theory, a topic on which I had worked a decade ago. On my return to UW, I started a research program in clustering graphs, from which resulted [78] and [79]*." These publications revisit and extend well-known spectral clustering algorithms for the Stochastic Block Model to a larger class of models with more expressive community structures.

Finally, we mention joint work of research fellow He Sun, together with Richard Peng and Luca Zanetti [68], who present the first rigorous analysis of the widely-employed k-means heuristic in the setting of spectral clustering. In practice, it is often observed that a gap between the k th and $(k+1)$ st eigenvalues is an indication that the data should admit a clustering into k clusters. Previous theoretical results (established earlier by other program participants) showed that such a gap guarantees the existence of a good clustering. In this work, Peng et. al. show that not only does such a clustering exist, but it can be found by the popular k-means heuristic run on the k-dimensional spectral embedding of the graph. This offers an elegant mathematical explanation for the effectiveness of such algorithms in practice.

Related areas

The program led to some remarkable works that do not fall squarely under the "spectral graph theory" umbrella. In an impressive collaboration between the two Fall programs, research fellows Ben Rossman (Algebraic Geometry) and Li-Yang Tan (SGT), in joint work with Rocco Servedio, proved that the polynomial hierarchy is infinite relative to a random oracle, solving a famous open problem in computational complexity theory posed separately by Håstad, Cai and Babai in 1986. Their work [70] received the Best Paper Award at FOCS 2015.

During the "Spectral algorithms: From theory to practice" workshop, Hui Han Chin presented an algorithm for image denoising that he developed based on fast solvers for Laplacian linear systems. Visiting graduate student Jakub Pachocki had a spark of inspiration: maybe some of Chin's ideas could be applied to finding the geometric median of a set of points in Euclidean space, a classical problem in computational geometry. Even for the case of three points in the plane, this problem goes back to Fermat, who famously sent a letter to Italian mathematician Torricelli asking for a solution. (The solution is now known as the Fermat-Torricelli point of a triangle, i.e., the one that minimizes the sum of the Euclidean distances to its three vertices.) In joint work with program organizer Gary Miller, research fellows Lee and Sidford, and MIT graduate student Michael Cohen, they eventually produced a spectacular advance: a nearly-linear time algorithm in any fixed dimension [27].

A picture of science at the Institute

In addition to the Boot Camp and three very well attended workshops, the program was overflowing with scientific activity. The second and third floor common areas in Calvin Lab were frequently utilized from early in the morning to very late at night. Certainly intense collaborations provided the intellectual meat of the semester. More than a few projects were initiated during the daily afternoon tea and cookies!

There were two concurrent weekly seminars: one on theoretical aspects of spectral graph theory, and the other specially designed to highlight the application of spectral graph methods in industry (with speakers including researchers from Google, Yahoo and Netflix, and the VP of HeartFlow—a company using spectral methods for medical imaging).

We also ran a weekly "as long as it takes" reading group. After the first instantiation ran for 3.5 hours, it was agreed that, by popular vote, the group could be paused and resumed the next afternoon. There were a number of smaller, specialized reading groups as well.

Finally, the two Open Lectures provided the larger UC Berkeley community with an opportunity to learn about the high-level scientific content of the program. The first lecture (by program organizer Gary Miller) covered the state of the art in using spectral graph methods to solve linear systems. The second lecture (by program organizer James R. Lee) came near the end of the semester and was able to report on much of the remarkable scientific progress that the program achieved.

Algorithmic Spectral Graph Theory, Fall 2014

[1] N. AGARWAL, K. CHANDRASEKARAN, A. KOLLA, and V. MADAN. On the expansion of group based lifts. 2016. In submission.

[2] N. AGARWAL, A. S. BANDEIRA, K. KOILIARIS, and A. KOLLA. Multisection in the stochastic block model using semidefinite programming. *arXiv preprint arXiv:1507.02323,* 2015.

[3] Z. ALLEN-ZHU, Z. LIAO, and L. ORECCHIA. Linear-Sized Spectral Sparsification in Almost Quadratic Time and Regret Minimization Beyond Matrix Multiplicative Weight Updates. In *Proceedings of the 47th Annual Symposium on the Theory of Computing (STOC)*, 2015.

[4] Z. ALLEN-ZHU and L. ORECCHIA. Nearly-Linear Time Packing and Covering LP Solver with Faster Convergence Rate Than $O(1/\varepsilon^2)$. In *Proceedings of the 47th Annual Symposium on the Theory of Computing (STOC)*, 2015.

[5] N. ANARI and S. OVEIS GHARAN. The Kadison-Singer Problem for Strongly Rayleigh Measures and Applications to Asymmetric TSP. *arXiv preprint arXiv:1412.1143v1,* 2014.

[6] N. ANARI and S. OVEIS GHARAN. Effective-resistance-reducing flows, spectrally thin trees, and asymmetric TSP. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 20–39, IEEE, 2015.

[7] A. ANDONI, J. CHEN, R. KRAUTHGAMER, B. QIN, D. P. WOODRUFF, and Q. ZHANG. On Sketching Quadratic Forms. In *Proceedings of the 7th ACM Conference on Innovations in Theoretical Computer Science (ITCS)*, pp. 311–319, ACM, 2016.

[8] A. ANDONI, R. KRAUTHGAMER, and I. RAZENSHTEYN. Sketching and Embedding are Equivalent for Norms. In *Proceedings of the 47th Annual Symposium on the Theory of Computing (STOC)*, 2015.

[9] A. ANDONI and I. RAZENSHTEYN. Tight Lower Bounds for Data-Dependent Locality-Sensitive Hashing. *arXiv preprint arXiv:1507.04299v1,* 2015.

[10] A. ANDONI and I. RAZENSHTEYN. Optimal Data-Dependent Hashing for Approximate Near Neighbors. In *Proceedings of the 47th Annual Symposium on the Theory of Computing (STOC)*, 2015.

[11] P. AWASTHI, M. CHARIKAR, R. KRISHNASWAMY, and A. K. SINOP. Spectral Embedding of k-Cliques, Graph Partitioning and k-Means. In *Proceedings of the 7th ACM Conference on Innovations in Theoretical Computer Science (ITCS)*, pp. 301-310, 2016.

[12] P. AWASTHI, M. CHARIKAR, R. KRISHNASWAMY, and A. K. SINOP. The Hardness of Approximation of Euclidean k-Means. In *Proceedings of the 31st International Symposium on Computational Geometry (SoCG)*, pp. 754-767, 2015.

[13] Y. BABICHENKO, C. PAPADIMITRIOU, and A. RUBINSTEIN. Can Almost Everybody be Almost Happy? PCP for PPAD and the Inapproximability of Nash. In *Proceedings of the 7th ACM Conference on Innovations in Theoretical Computer Science (ITCS)*, pp. 1-9, 2016.

[14] D. BACON, S. T. FLAMMIA, A. W. HARROW, and J. SHI. Sparse Quantum Codes from Quantum Circuits. In *Proceedings of the 47th Annual Symposium on the Theory of Computing (STOC)*, 2015.

[15] T. BANSAL, C. BHATTACHARYYA, and R. KANNAN. A provable SVD-based algorithm for learning topics in dominant admixture corpus. In *Advances in Neural Information Processing Systems 27 (NIPS)*, pp. 1997–2005, December 2014.

[16] B. BARAK, J. A. KELNER, and D. STEURER. Dictionary Learning and Tensor Decomposition via the Sum-of-Squares Method. In *Proceedings of the 47th Annual Symposium on the Theory of Computing (STOC)*, 2015.

[17] M. BRAVERMAN, Y. K. KO, A. RUBINSTEIN, and O. WEINSTEIN. ETH Hardness for Densest-k-Subgraph with Perfect Completeness. *arXiv preprint arXiv:1504.08352v1,* 2015.

[18] S. BROOKS, E. LE MASSON, and E. LINDENSTRAUSS. Quantum ergodicity and averaging operators on the sphere. *International Mathematics Research Notices,* RNV337, 2015.

[19] D. CALANDRIELLO, A. LAZARIC, M. VALKO, and I. KOUTIS. Incremental Spectral Sparsification for Large-Scale Graph-Based Semi-Supervised Learning. *arXiv preprint arXiv:1601.05675v1,* 2016. In submission.

[20] P. CHALERMSOOK, F. GRANDONI, and B. LAEKHANUKIT. On Survivable Set Connectivity. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 25–36, 2015.

[21] S. O. CHAN, T. C. KWOK, and L. C. LAU. Random Walks and Evolving Sets: Faster Convergence and Limitations. *arXiv preprint arXiv:1507.02069v1,* 2015.

[22] S. CHAWLA, K. MAKARYCHEV, T. SCHRAMM, and G. YAROSLAVTSEV. Near Optimal LP Rounding

Algorithm for Correlation Clustering on Complete and Complete k-partite Graphs. In *Proceedings of the 47th Annual Symposium on the Theory of Computing (STOC)*, 2015.
[23] W. CHEN, F. LI, T. LIN, and A. RUBINSTEIN. Combining Traditional Marketing and Viral Marketing with Amphibious Influence Maximization. In *Proceedings of the 16th ACM Conference on Economics and Computation (EC)*, pp. 779-796, 2015.
[24] X. CHEN, A. DE, R. A. SERVEDIO, and L.-Y. TAN. Boolean function monotonicity testing requires (almost) n^½ non-adaptive queries. In *Proceedings of the 47th Annual Symposium on the Theory of Computing (STOC)*, 2015.
[25] F. CHUNG. A generalized Alon-Boppana bound and weak Ramanujan graphs. Preprint, 2016.
[26] F. CHUNG and R. GRAHAM. The matrix cover polynomial. 2016. J. of Combinatorics, to appear.
[27] M. B. COHEN, Y. T. LEE, G. L. MILLER, J. W. PACHOCKI, and A. SIDFORD. Geometric Median in Nearly Linear Time. To appear in STOC'16.
[28] M. CUCURINGU, I. KOUTIS, S. CHAWLA, G. MILLER, and R. PENG. Simple and Scalable Constrained Clustering: A Generalized Spectral Method. *arXiv preprint arXiv:1504.00653,* 2016. To appear in AISTATS'16.
[29] M. CUCURINGU and J. WOODWORTH. Ordinal embedding of unweighted kNN graphs via synchronization. *Machine Learning for Signal Processing (MLSP), 2015 IEEE 25th International Workshop,* 2015.
[30] R. ELDAN, J. R. LEE, and J. LEHEC. A transportation-entropy inequality for Markov chains with a local contraction condition. Preprint, 2016.
[31] R. ELDAN and J. R. LEE. Regularization under diffusion and anti-concentration of temperature. *arXiv preprint arXiv:1410.3887v2,* 2014.
[32] R. ELDAN, M. RÁCZ, and T. SCHRAMM. Braess's paradox for the spectral gap in random graphs and delocalization of eigenvectors. *arXiv preprint arXiv:1504.07669v2,* 2015.
[33] A. R. ENE, G. L. MILLER, J. W. PACHOCKI, and A. SIDFORD. Routing under Balance. 2015. To appear in STOC'16.
[34] H. FAWZI, J. SAUNDERSON, and P. A. PARRILO. Equivariant semidefinite lifts of regular polygons. *arXiv preprint arXiv:1409.4379v1,* 2014.
[35] H. FAWZI, J. SAUNDERSON, and P. A. PARRILO. Sparse sum-of-squares certificates on finite abelian groups. *arXiv preprint arXiv:1503.01207v1,* 2015.
[36] R. FROSTIG, R. GE, S. M. KAKADE, and A. SIDFORD. Competing with the empirical risk minimizer in a single pass. In *Proceedings of the 28th Conference on Learning Theory (COLT)*, 2015.
[37] N. HARVEY and J. VONDRÁK. An Algorithmic Proof of the Lovasz Local Lemma via Resampling Oracles. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2015.
[38] D. S. HOCHBAUM. Faster weighted isotonic regression with fused Lasso under the L1 norm. Preprint. 2016.
[39] D. S. HOCHBAUM and C. LYU. Improved Algorithms Solving Generalizations of Isotonic Median Regression. Preprint. 2015.
[40] S. B. HOPKINS, J. SHI, and D. STEURER. Tensor principal component analysis via sum-of-squares proofs. In *Proceedings of the 28th Conference on Learning Theory (COLT)*, 2015.
[41] S. B. HOPKINS, T. SCHRAMM, J. SHI, and D. STEURER. Fast spectral algorithms from sum-of-squares proofs: tensor decomposition and planted sparse vectors. *arXiv preprint arXiv:1512.02337v2,* 2016. To appear in STOC'16.
[42] G. KINDLER, A. KOLLA, and L. TREVISAN. Approximation of non-boolean 2CSP. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1705-1714, 2016.
[43] B. KLARTAG, G. KOZMA, P. RALLI, and P. TETALI. Discrete curvature and abelian groups. *arXiv preprint arXiv:1501.00516,* 2015. Canadian J. Math. to appear (2016).
[44] R. KONDOR, N. TENEVA, and P. K. MUDRAKARTA. Parallel MMF: a Multiresolution Approach to Matrix Computation. *arXiv preprint arXiv:1507.04396,* 2015.
[45] I. KOUTIS. Simple parallel and distributed algorithms for spectral graph sparsification. In *Proceedings of the 26th ACM symposium on Parallelism in algorithms and architectures*, pp. 61–66, ACM, 2014. Revised journal version submitted.
[46] I. KOUTIS. Algebraic Algorithms for Out-Branching Problems. 2016. In submission.
[47] W. KRICHENE, A. BAYEN, and P. L. BARTLETT. Accelerated Mirror Descent in Continuous and Discrete Time. In *Advances in Neural Information Processing Systems 28 (NIPS)*, pp. 2827–2835, 2015.
[48] T. C. KWOK, L. C. LAU, and Y. T. LEE. Improved Cheeger's Inequality and Analysis of Local Graph Partitioning using Vertex Expansion and Expansion Profile. In *Proceedings of the ACM-SIAM Symposium on*

Discreet Algorithms (SODA), pp. 1848-1861, 2016.
[49] R. KYNG, Y. T. LEE, R. PENG, S. SACHDEVA, and D. A. SPIELMAN. Sparsified Cholesky and Multigrid Solvers for Connection Laplacians. *arXiv preprint arXiv:1512.01892,* 2015.
[50] B. LAEKHANUKIT. Tree Embedding and Directed Steiner Problems. *arXiv preprint arXiv:1511.06559,* 2015.
[51] K. G. LARSEN, J. NELSON, and H. L. NGUYEN. Time lower bounds for nonadaptive turnstile streaming algorithms. In *Proceedings of the 47th Annual Symposium on the Theory of Computing (STOC)*, 2015.
[52] J. R. LEE. Covering the large spectrum and generalized Riesz products. *arXiv preprint arXiv:1508.07109,* 2015.
[53] J. R. LEE, P. RAGHAVENDRA, and D. STEURER. Lower bounds on the size of semidefinite programming relaxations. In *Proceedings of the 47th Annual Symposium on the Theory of Computing (STOC)*, 2015. Best paper award.
[54] Y. T. LEE and A. SIDFORD. Efficient Inverse Maintenance and Faster Algorithms for Linear Programming. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 230-249, 2015.
[55] Y. T. LEE, A. SIDFORD, and S. C.-W. WONG. A faster cutting plane method and its implications for combinatorial and convex optimization. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 1049–1065, 2015. Best student paper award.
[56] Y. T. LEE and H. SUN. Constructing Linear-Sized Spectral Sparsification in Almost-Linear Time. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 250-269, 2015.
[57] A. LOUIS and S. S. VEMPALA. Accelerated Newton Iteration: Roots of Black Box Polynomials and Matrix Eigenvalues. *arXiv preprint arXiv:1511.03186,* 2015.
[58] K. MAKARYCHEV, Y. MAKARYCHEV, M. SVIRIDENKO, and J. WARD. A bi-criteria approximation algorithm for \$k\$ Means. *arXiv preprint arXiv:1507.04227v2,* 2015.
[59] K. MAKARYCHEV, Y. MAKARYCHEV, and A. VIJAYARAGHAVAN. Learning Communities in the Presence of Errors. *arXiv preprint arXiv:1511.03229v2,* 2016.
[60] K. MAKARYCHEV, Y. MAKARYCHEV, and A. VIJAYARAGHAVAN. Correlation clustering with noisy partial information. *Journal of Machine Learning Research,* 40, pp. 1–22, 2015. Also appeared in COLT'15.
[61] K. MAKARYCHEV, Y. MAKARYCHEV, and Y. ZHOU. Satisfiability of Ordering CSPs Above Average. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2015.
[62] A. MARCUS, D. A. SPIELMAN, and N. SRIVASTAVA. Finite free convolutions of polynomials. *arXiv preprint arXiv:1504.00350v1,* 2015.
[63] A. W. MARCUS, D. A. SPIELMAN, and N. SRIVASTAVA. Interlacing Families IV: Bipartite Ramanujan Graphs of All Sizes. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 1358–1377, IEEE, 2015.
[64] R. MEKA, A. POTECHIN, and A. WIGDERSON. Sum-of-squares lower bounds for planted clique. In *Proceedings of the 47th Annual Symposium on the Theory of Computing (STOC)*, pp. 87–96, ACM, 2015.
[65] W. MULZER, H. L. NGUYEN, P. SEIFERTH, and Y. STEIN. Approximate k-flat Nearest Neighbor Search. In *Proceedings of the 47th Annual Symposium on the Theory of Computing (STOC)*, 2015.
[66] C. PAPADIMITRIOU, G. PIERRAKOS, C.-A. PSOMAS, and A. RUBINSTEIN. On the Complexity of Dynamic Mechanism Design. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1458-1475, 2016.
[67] C. H. PAPADIMITRIOU and S. S. VEMPALA. Unsupervised Learning through Prediction in a Model of Cortex. In *Proceedings of the 28th Conference on Learning Theory (COLT)*, 2015.
[68] R. PENG, H. SUN, and L. ZANETTI. Partitioning Well-Clustered Graphs: Spectral Clustering Works! In *Proceedings of the 28th Conference on Learning Theory (COLT)*, pp. 1423-1455, 2015.
[69] Y. RABANI and L. J. SCHULMAN. The Invisible Hand of Laplace: the Role of Market Structure in Price Convergence and Oscillation. *arXiv preprint arXiv:1602.07628,* 2016.
[70] B. ROSSMAN, R. A. SERVEDIO, and L.-Y. TAN. An average-case depth hierarchy theorem for Boolean circuits. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2015. Best paper award.
[71] A. RUBINSTEIN. Inapproximability of Nash Equilibrium. In *Proceedings of the 47th Annual Symposium on the Theory of Computing (STOC)*, pp. 409-418, 2015.
[72] A. RUBINSTEIN and S. M. WEINBERG. Simple Mechanisms for a Subadditive Buyer and Applications to Revenue Monotonicity. In *Proceedings of the 16th ACM Conference on Economics and Computation (EC)*, pp.

377-394, 2015.

[73] S. SACHDEVA and N. K. VISHNOI. A Simple Analysis of the Dikin Walk. *arXiv preprint arXiv:1508.01977v1,* 2015.

[74] R. A. SERVEDIO, L.-Y. TAN, and J. WRIGHT. Adaptivity helps for testing juntas. In *30th Conference on Computational Complexity (CCC)*, 2015.

[75] A. K. SINOP. How to Round Subspaces: A New Spectral Clustering Algorithm. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1832-1847, 2016.

[76] D. STRASZAK and N. K. VISHNOI. Natural algorithms for flow problems. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1868–1883, 2016.

[77] N. TENEVA, P. K. MUDRAKARTA, and R. KONDOR. Multiresolution Matrix Compression. 2016. Accepted for oral presentation at AISTATS.

[78] Y. WAN and M. MEILA. Benchmarking recovery theorems for the DC-SBM. In *Proceedings of the International Symposium on Artificial Intelligence and Mathematics (ISAIM)*, 2016.

[79] Y. WAN and M. MEILA. A class of network models recoverable by spectral clustering. In *Advances in Neural Information Processing Systems 28 (NIPS)*, pp. 3267–3275, 2015.

[80] C. ZHANG, J. SONG, K. C. CHEN, and K. CHAUDHURI. Spectral Learning of Large Structured HMMs for Comparative Epigenomics. In *Advances in Neural Information Processing Systems 28 (NIPS)*, 2015.