Anup Rao*

Research Statement

My research interests are primarily in theoretical computer science. It is an exciting time to be working in theory — there are many opportunities to discover new connections and many opportunities to make progress on foundational questions. I have had success with two such broad questions so far.

Question 1: To what extent is the use of randomness necessary in computer science?

By now it is clear that randomness plays a key role in algorithm design, cryptography and distributed computing. So it is worthwhile to investigate exactly what can be done with and without randomness, and to identify the minimal assumptions on the randomness under which we can still get the benefits of randomized methods. This is the aim of the research area called **pseudorandomness**.

Question 2: How can we make computationally hard problems even harder?

Finding explicit examples of computationally hard problems is a major project in complexity and is the central goal of the research area called **lower bounds**. Such examples would be useful for generating pseudorandomness, and for cryptography. One way to find such an explicit problem is to start with a mildly hard problem, and amplify its hardness by transforming it in some simple way. Similar techniques are also useful to prove that it is hard to approximate the answer to optimization problems.

My research has led to many discoveries, and many avenues for new progress remain. I start by giving some highlights of my past research.

Pseudorandomness: I have spent most of my efforts in this area on the design of efficient *randomness extractors*.

Randomness Extractors for Independent Sources Randomness is an essential resource for solving computer science problems, yet physical sources of randomness can at best give us samples from a distribution with some entropy. In [Rao06a], I gave the best known algorithms to extract true randomness from a few independent defective physical sources, showing how to extract randomness even when each source that supplies n bits has $n^{\Omega(1)}$ entropy. The best results prior to my work required $\Omega(n)$ entropy in each source.

Ramsey Graphs In 1947, Erdős introduced the *Probabilistic Method* by showing that most graphs on n vertices are Ramsey Graphs — they do not have any cliques or independent sets of size $2 \log n$. The quest to find explicit Ramsey graphs was born. In [BRSW06], my coauthors and I gave an algorithm to build the best known explicit Ramsey Graph today— a graph of size n that has no cliques or independent sets of size $2^{\log^{o(1)} n}$. The previous bound of $2^{\sqrt{\log n}}$ [FW81] had survived for 25 years, despite many attempts to surpass it.

*Institute for Advanced Study, arao@ias.edu

Lower bounds: Most of my results in this area relate to certain types of *amplification methods*.

Parallel Repetition Proving that it is NP-hard to approximate the value to optimization problems is a major research area. I gave an optimal analysis [Rao08b] for parallel repetition, a crucial component of such hardness of approximation results. My result gave the first positive evidence for the famous *Unique Games Conjecture* of Khot.

Direct Sums in Communication Complexity If two parties need to communicate c bits with each other to compute some joint function of their inputs, how many bits of communication does it take to compute the function on k different inputs? Proving that the communication must grow for the general model of randomized communication has been a longstanding open problem. In [BBCR08], my coauthors and I made the first progress on this question, showing that roughly $\Omega(c\sqrt{k})$ bits of communication are required.

Foams What is the least surface area of a body that tiles space like a cube? In [KORW08], my coauthors and I used techniques developed to understand parallel repetition to show the surprising fact that there is a body that tiles space like a cube, yet has surface area that is similar to that of a sphere (which is best possible). Similar ideas led us to a new probabilistic rounding scheme that can round points in Euclidean space to nearby integer points in such a way that the probability that two points get different roundings is proportional to the distance between them.

Next, I give a more detailed description of my research interests.

1 Pseudorandomness and Extractors

It is widely acknowledged that the physical universe is unpredictable. In computer science this unpredictability, modeled as randomness, turns out to be an enabling feature in algorithm design, cryptography and distributed computing. Can we weaken the assumption that we have access to true randomness? One way to do this is to show how to simulate any solution that requires true randomness with a solution that requires a weaker type of randomness, or no randomness at all. Some of the most celebrated recent results in theoretical computer science have come as a result of taking some randomized solution (usually a very natural one) and replacing the randomness with something that is pseudorandom. The deterministic primality test of Agrawal, Kayal and Saxena [AKS04] and the log-space algorithm for undirected connectivity of Reingold [Rei05] are both examples of this approach.

A key step in many such problems is the efficient deterministic construction of *pseudorandom* objects, examples of which include *randomness extractors* and *Ramsey graphs*.

1.1 Explicit Extractors

Randomness extractors are algorithms that can help close the gap between the physical assumption of unpredictability and the computer science model of pure randomness. Physical sources don't yield independent truly random bits, instead they produce samples with some entropy. A randomness extractor is an algorithm that extracts truly random bits from such defective sources, which can then be used in computer science applications.

Unfortunately, it is impossible to extract randomness from a single defective source. So most work has focussed on finding the most general assumptions under which extractor algorithms can be designed.

A very natural case under which it is possible to extract is when we have access to 2 or more independent defective sources. In [Rao06a], I designed an algorithm that has the best known tradeoff between the number of sources and the entropy available in each source. It only requires a constant number of sources of length n, each with entropy n^ϵ for arbitrary $\epsilon > 0$. Results prior to my work [CG88, BIW04, BKS+05, Raz05, Bou05] all required entropy at least $\Omega(n)$, and the best of them relied on major results from arithmetic combinatorics. The technique that I developed to solve this problem was elementary, and turned out to be useful for several other randomness extractor constructions. In subsequent works [RZ08, Rao08a] we managed to show how to extract even when the number of sources is as small as 3 or 2, though these results have some additional restrictions.

A good model for randomness obtained from physical devices is to model the randomness as coming from a process with small memory. In work with Jesse Kamp, Salil Vadhan and David Zuckerman [KRVZ06] we showed how to extract from such a source when the entropy is as small as n^α, for some fixed constant $\alpha < 1$.

Another context in which such extractors are useful is when we have some truly random string that is corrupted by an adversary. For example, it might be the case that the adversary has learnt some unknown part of our secret key in a cryptographic protocol [Dod00, CDH+00]. In [Rao06b], I showed how to extract many private bits even when the adversary knows all but $\log^{\Omega(1)} n$ of the n bits. Prior to my work, the best results [KZ07, GRS04] required that the adversary does not know at least $\sqrt{n}$ of the bits.

1.2 Ramsey Graphs

A *Ramsey Graph* is a prototypical example of a pseudorandom object — it is a graph that has no large cliques or independent sets. In 1928 Ramsey [Ram28] proved that every graph on n vertices must have a clique or independent set of size at least $(1/2)\log n$. In 1947 Erdős published his paper inaugurating the *Probabilistic Method* with a few examples, including a proof that complemented Ramsey's discovery: *most* graphs on n vertices avoid cliques and independent sets of size $2\log n$. This prompted many attempts to find an explicit family of graphs matching Erdős' bound. For a long time, the best record was a construction of Frankl and Wilson [FW81] from 1981. They used intersection theorems for set systems to construct n-vertex graphs that avoid cliques/independent sets of size $2^{\Omega(\sqrt{\log n})}$. This bound was matched by several other works [Alo98, Gro00, Bar06, Gop06]. Remarkably all of these attempts got stuck at essentially the same bound.

Ramsey graphs are intimately connected with extractors for two sources — if the adjacency matrix for the graph is viewed as the truth table of an extractor, the extractor property guarantees that every large set of vertices in the graph induces a subgraph with edge density close to 50%. In particular, this ensures that there can be no large cliques or independent sets. Barak, Shaltiel, Wigderson and myself [BRSW06] used this connection to give a new explicit Ramsey Graph that avoids cliques and independent sets of size $2^{\log^{o(1)} n}$, finally surpassing the bound of Frankl and Wilson.

1.3 Network Extractors

Distributed computing provides many examples of problems that are impossible to solve without the use of randomness. Can we relax the assumption that we have access to truly random bits in these settings? This kind of question was first considered by Goldwasser et al. [GSV05] who gave some results for a restricted model of impure randomness.

In work with Yael Tauman Kalai, Xin Li and David Zuckerman [KLRZ08], we showed that under certain cryptographic assumptions, *any* cryptographic task involving the use of true randomness by

many parties can be simulated if each participant has just one defective source of randomness with entropy n^ϵ. In particular, this implies that we can perform reliable cryptography even if all but n^ϵ of the information in the private randomness of each player is leaked to the adversary. We proved this by designing a *network extractor protocol* — a protocol that the players can run to extract randomness from their sources, even in the presence of adversaries in the network. The protocol succeeds even if a large fraction of the players collude in order to disrupt the protocol or learn information about the private randomness of the honest players. An open problem, on which we have made significant progress, is to get a similar result based on standard cryptographic assumptions, like the existence of one way functions.

2 Amplification by Repetition

The central (and often elusive) goal of complexity theory is to understand exactly how hard or easy it is to compute functions under various models of computation. Since we understand very little about how the complexity of functions behaves, it makes sense to look for some structure that may help us build a coherent picture. My own research in this area has involved understanding a very natural way to take a function and use it to get a function that is even harder to compute, a process called *amplification*.

2.1 Parallel Repetition

A significant body of work in theoretical computer science has gone in to show that even mildly approximating the correct answer to many optimization problems is NP-hard [FGL+91, ALM+98, FL92, LY93]. In order to prove such a *hardness of approximation* result for a specific problem, one typically needs to show how any approximation algorithm with a bad approximation guarantee can be used to get an approximation algorithm with a much stronger approximation guarantee. A key step in carrying out this kind of amplification involves proving a *parallel repetition* theorem.

For example, suppose we are given as input a system of equations over variables $x_1, \ldots, x_m$, where each equation has the form $x_i + x_j = a$. Is there an algorithm that can distinguish when at most a $(1-\epsilon)$ fraction of the equations can be satisfied from when at least $(1-\gamma)$ fraction of the equations can be satisfied? A natural way to make this problem easier is to *amplify the gap* by parallel repetition: we consider a new system of vector equations, where each equation $\mathbf{x}_i + \mathbf{x}_j = \mathbf{a}$ corresponds to satisfying k of the original equations simultaneously. We might expect that if the fraction of satisfiable equations was less than $(1-\epsilon)$ to start with, the fraction of satisfiable equations in the new system is less than $(1-\epsilon)^k$. Thus, it would be enough to distinguish the case that this fraction is less than $(1-\epsilon)^k$ or more than $(1-\gamma)^k$, which can be quite a large gap for a carefully chosen k.

However, actually proving that this type of gap amplification works took considerable effort. In a celebrated paper, Raz [Raz98] showed that if the optimization problem is ϵ-close to having a perfect solution, the satisfiable fraction after repeating the problem k times in parallel is at most $(1-\epsilon^{32})^{\Omega(k/c)}$, where $c \geq 1$ is a constant that depends on the structure of the problem. Raz's result was enough to prove several hardness of approximation results and was subsequently simplified by Holenstein [Hol07], to get the bound $(1-\epsilon^3)^{\Omega(k/c)}$.

In recent years, many optimal hardness of approximation results have come to be based on the *Unique Games Conjecture* of Khot [Kho02], for example for Max 2-Lin [Kho02], Vertex Cover [KR03], Max-Cut [Kho02, KKMO04, MOO05], Approximate Coloring [DMR06], Sparsest Cut [CKK+06, KV05] and Max 2-Sat [Aus07]. The conjecture renewed interest in the exact bound for parallel

repetition, since a bound that was independent of the structure of the problem would weaken the conjecture.

I helped to settle the question of how much parallel repetition can help in this situation. In [Rao08b], I showed that the truth $\left((1-\epsilon^2)^{\Omega(k)}\right)$ *is* independent of the structure of the particular optimization problems being considered. Raz [Raz08a] then discovered that my bounds were optimal. In subsequent work with Barak, Haviv, Hardt, Regev and Steurer [BHH+08], we showed how to use semi-definite programs to classify exactly which problems do not behave well under parallel repetition. We discovered that the situations in which parallel repetition fails to work well are exactly the situations that we would like it to work in (namely when known algorithms give bad approximation guarantees).

It turns out that a variant of the parallel repetition question is relevant to quantum physicists trying to design experiments that verify quantum entanglement [Bel64, CHSH69, BCH+02, Gil03]. My results also provided the first useful bound for that context.

2.2 Direct Sums in Communication Complexity

Communication is at the heart of most computational processes, and so the study of *communication complexity* has had many applications in computer science. Communication complexity measures how many bits two parties need to send each other in order to compute some joint function of their inputs. If the communication complexity of a function is c, what can we say about the communication complexity of computing the function on k different inputs simultaneously? No answer better than the trivial lower bound of c was known for the general model of randomized communication complexity, despite much effort [KKN95, FKNN95, CSWY01, BYJKS04, SS02, Abl93, HJMR07].

In work with Barak, Braverman, and Chen [BBCR08], we showed that computing the function on k inputs requires roughly $\Omega(c\sqrt{k})$ bits of communication. If we are interested in computing the function on average under the uniform distribution, our bound is improved to roughly $\Omega(ck)$, essentially best possible. A key part of our proof involved using the insights and intuitions developed in the work on parallel repetition.

2.3 Foams

My work on parallel repetition led me to a beautiful question in geometry — what is the least possible surface area of a body in d-dimensional Euclidean space that tiles space like the cube? Here we are looking for a body that can cover all of space by integer translations, without overlaps. The cube itself has surface area $2d$, and any such body must have surface area larger than the unit volume sphere $O(\sqrt{d})$. It was known [Cho89] that the best such shape in the plane has perimeter $\sqrt{2}+\sqrt{6} \approx 3.86$, which is less than 4, the perimeter of the unit square. However, the best candidate for the case of large d had surface area $\Omega(d)$.

In work with Kindler, O'Donnell and Wigderson [KORW08], we used ideas from Raz's proof of optimality for my parallel repetition theorem to show the surprising fact that there is a such a body with surface area that is about 3 times larger than that of the d dimensional sphere ($O(\sqrt{d})$). Our work was particularly satisfying because it illustrated how work in computer science could have an impact on mathematics.

Our methods also allowed us to show that there is a way to round points in Euclidean space to nearby points with integer coordinates in such a way that the probability that two points get different roundings is proportional to the distance between them. This kind of rounding scheme would be useful

in algorithms for geometric problems like nearest neighbor search, though we need to find an efficient implementation of our scheme before it becomes useful for this application.

3 The Future

There are a several broad questions that I plan to explore in the next few years. The first is the central question in the area of pseudorandomness:

Question 1: Can every randomized algorithm be efficiently simulated by a deterministic algorithm?

I intend to use this as the guiding question for my work in this area over the next few years. One way to approach this issue is to find an efficient *pseudorandomness generator*, namely a deterministic algorithm that produces bits that looks random to efficient algorithms. Recently, there has been some progress on constructing such generators for circuits of depth 2 [Baz07, Raz08b], showing that this area might be open to attack.

A target that I feel is closer at hand is to relax the assumption that we have true randomness in distributed environments:

Question 2: Can every distributed computing protocol be efficiently simulated by players with defective sources of randomness?

While it is known that a defective source of randomness is enough to simulate every randomized algorithm, it is still unclear that every protocol in a distributed environment can be simulated by players with defective sources of randomness. I believe that the answer is positive, but many ideas are needed to come up with a formal proof of this fact. We have already made a lot of progress towards proving the analogous result for cryptographic protocols, where the adversary is assumed to be computationally bounded.

The main goal in the area of lower bounds is to find an explicit function that is hard to compute by efficient algorithms. This is a major open problem and it is not likely to be easily resolved. A likely easier question, that would still be a major step, is the following:

Question 3: Are boolean circuits of large depth strictly more powerful than boolean circuits of small depth?

If the answer is positive (I believe that it is), this says that there is no generic way to increase the number of processors in a parallel computing environment to *always* give a speed up. In fact, it is known [KRW95] that this question can be cast in the framework of proving a direct sum type result for communication complexity. Given the progress on such direct sums that we have already made in [BBCR08], and the tools we discovered along the way to this result, I am optimistic that many doors will soon open.

Another area where I feel I can effectively employ my time is pushing the state of knowledge in hardness of approximation.

Question 4: Is the Unique Games Conjecture true?

As I discussed earlier, giving a positive answer to this question would give many optimal hardness of approximation results, so it is of considerable interest. My parallel repetition bounds have already given the first evidence that the answer is positive, and I'm hopeful that further progress is in reach.

Advances in theory often come from discovering new connections with other areas. In recent years, theory has begun a healthy exchange of ideas with arithmetic combinatorics and geometry. My own work on foams is an example of this type of connection. Several recent works in pseudorandomness have also benefited from this interaction [BV07, Vio08, RTTV08]. I believe that there are many more connections between these areas that have not yet been discovered, and I intend to play an active role in finding them.

My Work

[BBCR08] Boaz Barak, Mark Braverman, Xi Chen, and Anup Rao. A direct sum theorem for randomized communication complexity. *Manuscript*, 2008.

[BHH+08] Boaz Barak, Ishay Haviv, Moritz Hardt, Anup Rao, Oded Regev, and David Steurer. Rounding parallel repetitions of unique games. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*. IEEE Computer Society, 2008.

[BRSW06] Boaz Barak, Anup Rao, Ronen Shaltiel, and Avi Wigderson. 2 source dispersers for $n^{o(1)}$ entropy and Ramsey graphs beating the Frankl-Wilson construction. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, 2006.

[KLRZ08] Yael Tauman Kalai, Xin Li, Anup Rao, and David Zuckerman. Network extractor protocols. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*. IEEE Computer Society, 2008.

[KRVZ06] Jesse Kamp, Anup Rao, Salil Vadhan, and David Zuckerman. Deterministic extractors for small space sources. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, 2006.

[KORW08] Guy Kindler, Ryan O'Donnell, Anup Rao, and Avi Wigderson. Spherical cubes and rounding in high dimensions. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*. IEEE Computer Society, 2008.

[Rao06a] Anup Rao. Extractors for a constant number of polynomially small min-entropy independent sources. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, 2006.

[Rao06b] Anup Rao. Extractors for low-weight affine sources. *Manuscript*, 2006.

[Rao08a] Anup Rao. A 2-source almost-extractor for linear entropy. In *RANDOM 2008, 12th International Workshop on Randomization and Approximation Techniques in Computer Science*, 2008.

[Rao08b] Anup Rao. Parallel repetition in projection games and a concentration bound. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, 2008.

[RZ08] Anup Rao and David Zuckerman. Extractors for 3 uneven length sources. In *RANDOM 2008, 12th International Workshop on Randomization and Approximation Techniques in Computer Science*, 2008.

References

[Abl93] Farid Ablayev. Lower bounds for one-way probabilistic communication complexity. In Andrzej Lingas, Rolf Karlsson, and Svante Carlsson, editors, *Proceedings of the 20th International Colloquium on Automata, Languages, and Programming*, volume 700 of *LNCS*, pages 241–252. Springer-Verlag, 1993.

[AKS04] Manindra Agrawal, Neeraj Kayal, and Nitin Saxena. PRIMES is in P. *Annals of Mathematics*, 160, 2004.

[Alo98] Noga Alon. The Shannon capacity of a union. *Combinatorica*, 18, 1998.

[ALM+98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45, 1998.

[Aus07] Per Austrin. Balanced max 2-sat might not be the hardest. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pages 189–197. ACM, 2007.

[BYJKS04] Ziv Bar-Yossef, T. S. Jayram, Ravi Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. *Journal of Computer and System Sciences*, 68(4):702–732, 2004.

[Bar06] Boaz Barak. A simple explicit construction of an $n^{\tilde{o}(\log n)}$-ramsey graph. Technical report, Arxiv, 2006. `http://arxiv.org/abs/math.CO/0601651`.

[BIW04] Boaz Barak, R. Impagliazzo, and Avi Wigderson. Extracting randomness using few independent sources. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pages 384–393, 2004.

[BKS+05] Boaz Barak, Guy Kindler, Ronen Shaltiel, Benny Sudakov, and Avi Wigderson. Simulating independence: New constructions of condensers, Ramsey graphs, dispersers, and extractors. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 1–10, 2005.

[BCH+02] Jonathan Barrett, Daniel Collins, Lucien Hardy, Adrian Kent, and Sandu Popescu. Quantum nonlocality, bell inequalities and the memory loophole. *Physical Review A*, 66:042111, 2002.

[Baz07] Louay Bazzi. Polylogarithmic independence can fool DNF formulas. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, pages 63–73. IEEE Computer Society, 2007.

[Bel64] John S. Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1(3):195–290, 1964.

[BV07] Andrej Bogdanov and Emanuele Viola. Pseudorandom bits for polynomials. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, pages 41–51. IEEE Computer Society, 2007.

[Bou05] Jean Bourgain. More on the sum-product phenomenon in prime fields and its applications. *International Journal of Number Theory*, 1:1–32, 2005.

[CDH+00] Ran Canetti, Yevgeniy Dodis, Shai Halevi, Eyal Kushilevitz, and Amit Sahai. Exposure-resilient functions and all-or-nothing transforms. In Bart Preneel, editor, *Advances in Cryptology — EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 453–469. Springer-Verlag, May 2000.

[CSWY01] Amit Chakrabarti, Yaoyun Shi, Anthony Wirth, and Andrew Yao. Informational complexity and the direct sum problem for simultaneous message complexity. In Bob Werner, editor, *Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science*, pages 270–278, Los Alamitos, CA, October 14–17 2001. IEEE Computer Society.

[CKK+06] Shuchi Chawla, Robert Krauthgamer, Ravi Kumar, Yuval Rabani, and D. Sivakumar. On the hardness of approximating multicut and sparsest-cut. *Proceedings of the 21th Annual IEEE Conference on Computational Complexity*, 15, 2006.

[Cho89] Jaigyoung Choe. On the existence and regularity of fundamental domains with least boundary area. *Journal of Differential Geometry*, 29, 1989.

[CG88] Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM Journal on Computing*, 17(2):230–261, 1988.

[CHSH69] John F. Clauser, Michael A. Horne, Abner Shimony, and Richard A. Holt. Proposed experiment to test local hidden-variable theories. *Phys. Rev. Lett.*, 23(15):880–884, Oct 1969.

[DMR06] Irit Dinur, Elchanan Mossel, and Oded Regev. Conditional hardness for approximate coloring. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*. ACM, 2006.

[Dod00] Yevgeniy Dodis. *Exposure-resilient cryptography*. PhD thesis, Massachusetts Institute of Technology, Dept. of Electrical Engineering and Computer Science, 2000.

[FKNN95] Tomàs Feder, Eyal Kushilevitz, Moni Naor, and Noam Nisan. Amortized communication complexity. *SIAM Journal on Computing*, 24(4):736–750, 1995.

[FGL+91] Uriel Feige, Shafi Goldwasser, Laszlo Lovasz, Shmuel Safra, and Mario Szegedy. Approximating clique is almost NP-complete (preliminary version). In *Proceedings of the 32nd Annual IEEE Symposium on Foundations of Computer Science*, 1991.

[FL92] Uriel Feige and Laszlo Lovasz. Two-prover one-round proof systems: their power and their problems. In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*, pages 733–744, 1992.

[FW81] Peter Frankl and R. M. Wilson. Intersection theorems with geometric consequences. *Combinatorica*, 1(4):357–368, 1981.

[GRS04] Ariel Gabizon, Ran Raz, and Ronen Shaltiel. Deterministic extractors for bit-fixing sources by obtaining an independent seed. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, 2004.

[Gil03] Richard D. Gill. Accardi contra bell (cum mundi): The impossible coupling. *IMS LECTURE NOTES-MONOGRAPH SERIES*, 42, 2003.

[GSV05] Shafi Goldwasser, Madhu Sudan, and Vinod Vaikuntanathan. Distributed computing with imperfect randomness. In *Distributed Computing, 19th International Conference, DISC 2005, Cracow, Poland, September 26-29, 2005, Proceedings*, volume 3724 of *Lecture Notes in Computer Science*, pages 288–302. Springer, 2005.

[Gop06] Parkshit Gopalan. Constructing Ramsey graphs from boolean function representations. In *Proceedings of the 21th Annual IEEE Conference on Computational Complexity*, 2006.

[Gro00] Vince Grolmusz. Low rank co-diagonal matrices and ramsey graphs. *Electr. J. Comb*, 7, 2000.

[HJMR07] Prahladh Harsha, Rahul Jain, David A. McAllester, and Jaikumar Radhakrishnan. The communication complexity of correlation. In *Proceedings of the 22nd Annual IEEE Conference on Computational Complexity*, pages 10–23. IEEE Computer Society, 2007.

[Hol07] Thomas Holenstein. Parallel repetition: Simplifications and the no-signaling case. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, 2007.

[KZ07] Jesse Kamp and David Zuckerman. Deterministic extractors for bit-fixing sources and exposure-resilient cryptography. *SIAM Journal on Computing*, 36(5):1231–1247, 2007.

[KKN95] Mauricio Karchmer, Eyal Kushilevitz, and Noam Nisan. Fractional covers and communication complexity. *SIAM Journal on Discrete Mathematics*, 8(1):76–92, February 1995.

[KRW95] Mauricio Karchmer, Ran Raz, and Avi Wigderson. Super-logarithmic depth lower bounds via the direct sum in communication complexity. *Computational Complexity*, 5(3/4):191–204, 1995.

[Kho02] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 767–775, 2002.

[KKMO04] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal inapproximability results for max-cut and other 2-variable CSPs? In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, 2004.

[KR03] Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within 2-ϵ. In *IEEE Conference on Computational Complexity*, page 379. IEEE Computer Society, 2003.

[KV05] Subhash Khot and Nisheeth Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into ℓ_1. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, 2005.

[LY93] Carsten Lund and Mihalis Yannakakis. On the hardness of approximating minimization problems. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing*, pages 286–293, 1993.

[MOO05] Elchanan Mossel, Ryan O'Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low in.uences invariance and optimality. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 21–30. IEEE Computer Society, 2005.

[Ram28] Frank P. Ramsey. On a problem of formal logic. *Proceedings of the London Mathematical Society, Series 2*, 30(4):338–384, 1928.

[Raz98] Ran Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, June 1998.

[Raz05] Ran Raz. Extractors with weak random seeds. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 11–20, 2005.

[Raz08a] Ran Raz. A counterexample to strong parallel repetition. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*. IEEE Computer Society, 2008.

[Raz08b] Alexander Razborov. A simple proof of bazzi's theorem. Technical Report TR08-081, ECCC: Electronic Colloquium on Computational Complexity, 2008.

[Rei05] Omer Reingold. Undirected ST-connectivity in log-space. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 376–385. ACM, 2005.

[RTTV08] Omer Reingold, Luca Trevisan, Madhur Tulsiani, and Salil Vadhan. Dense subsets of pseudorandom sets. Technical Report TR08-045, ECCC: Electronic Colloquium on Computational Complexity, 2008.

[SS02] Michael Saks and Xiaodong Sun. Space lower bounds for distance approximation in the data stream model. In ACM, editor, *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 360–369. ACM Press, 2002.

[Vio08] Emanuele Viola. The sum of d small-bias generators fools polynomials of degree d. In *IEEE Conference on Computational Complexity*, pages 124–127. IEEE Computer Society, 2008.