

Research Statement

Sourav Chakraborty

My field of research is Theoretical Computer Science. My focus has been in the classical and quantum complexity of Boolean functions (including property testing, sensitivity and block sensitivity of Boolean functions and quantum database search), in electronic commerce, in graph algorithms and in coding theory. I have designed effective algorithms as well as proved lower bounds for the complexity of problems in this area.

1 Combinatorial Complexity Measures of Boolean Functions

In my work in the field of combinatorial complexity measures of Boolean functions I strive to obtain a better understanding of different measures of hardness for Boolean functions, and their relation to the amount of resources needed to compute them in several combinatorial models. Previous results in the area have been obtained by insightful identification of the right measures of complexity and the choice of appropriate mathematical tools. Similarly, in my study I apply mathematical ideas and develop tools for analyzing the complexity of Boolean functions.

1.1 Property Testing

Many data sets that arise in the fields of biology, geology, astronomy, climatology, artificial intelligence, etc are massive. In fact they are so huge that even reading the whole data requires impossibly large resources. At the same time we need to analyze and have a good estimate of whether the data has certain property. We can use statistical methods to analyze the data. But statistical methods has it limitations. One alternative is “property testing”: looking at a very small (ideally constant) part of the graph/data and use combinatorial techniques to decide whether the graph/data has a predetermined property or is “far” from satisfying the property.

This way of analyzing the data started in the late eighties and early nineties [BLR, BFL, RS93]. Subsequently it had important application in the fields of probabilistically checkable proofs (PCP) [ArS, ALMSS] and learning and approximation [GGR96]. In recent years the field of combinatorial property testing has enjoyed a rapid growth (see, e. g., [AF+00, AS+06, AS03, AS05], cf. [R01, F01])

Graphs are combinatorial objects that are used to represent relations over a data set. Often we need to determine whether the graph has a predetermined property. For a fixed graph property \mathcal{P} , the *query-complexity for \mathcal{P}* is the number of queries to the input graph that the best algorithms needs, to determine, with high probability, whether the graph has the property or is “far” from having the property. In my research I aim to categorize various important graph properties according to their query complexities.

The query complexity can depend on the representation of the input graph. The two different models of representation that I have worked on are the dense graph model (or adjacency matrix model) and the orientation model. While the dense graph model is good for testing properties only in dense graphs, the orientation model is more suited for testing properties in directed graphs of all sizes.

1.1.1 Testing of *st*-connectivity in the Orientation Model

Hayley *et al* [HLNT] introduced the orientation model for property testing: Given an multi-graph with each edge oriented in exactly one direction, we have to distinguish between the case where the directed graph has a particular property \mathcal{P} and the case where we have to change the orientation of at least a constant fraction

of the edges to make the directed graph have the property. In this model we have prior knowledge of the whole underlying undirected graph and we get the orientations of each edge by querying.

We study various graph properties in this model and try to categorize them according to their query complexity. In particular we want to characterize the graph properties that can be tested with constant number of queries in this model.

Fischer, Lachish, Matsliah, Newman and I [CFL+] show how to test the property of st -connectivity using constant number of queries in the orientation model. In other words, we test whether there is a directed path between two specified vertices s and t in the graph by looking at the orientations of a constant number of edges. Our result gives the first non-trivial graph-property that can be tested using constant number of queries in this orientation model.

Our algorithm uses new combinatorial techniques, a new number-theoretic lemma and a reduction to Newman's algorithm [New] for testing membership in languages that have constant width branching programs. The algorithm crucially uses the prior knowledge of the underlying graph.

We are now trying to understand the query complexity in the orientation model for simple properties like acyclicity and strong-connectedness.

1.1.2 Testing of Generalization of Graph-Isomorphism in the Dense Graph Model

In the dense-graph model the input is an adjacency matrix of a graph and the access to each entry of the matrix requires a separate query.

Alon and Shapira [AS05] gave a complete characterization of graph properties that can be tested with constant number of queries. One important graph property that cannot be tested using constant number of queries is "graph isomorphism." Given two adjacency matrix we want to distinguish between the case where the graphs are isomorphic and the case where we need to change at least a constant fraction of the matrix to make the graphs isomorphic. Eldar and Matsliah [FM06] obtained the exact bounds on the query complexity for testing graph isomorphism.

A study a natural and important generalization of the graph isomorphism is testing isomorphisms under group actions. Given two n bit strings x and y and a permutation group G we want to test whether there is a permutation π in G such that by permuting the indices of x by the π we obtain the string y .

This abstract generalization has applications in many fields. For example, in image recognition if we want to check whether one images is a shifts of another we consider the group of all shifts of the pixels of the images and test for isomorphism under this group action. Also graph and hyper-graph isomorphism are special cases.

Babai and I [BaC] study this problem when G is a primitive group. Primitive groups acts as building blocks for other permutation groups. Hence understanding query complexity under primitive group action is the first but significant step toward the main problem. In fact various finite geometric transformations (like projective, orthogonal, symplectic, etc) corresponds to primitive groups. Graph and hyper-graph isomorphism also corresponds to primitive groups actions.

In the primitive group case we prove tight bounds for query complexity when the algorithm is allowed to make only one sided errors. The techniques are similar to that in [FM06], but we need to use the structure theorem of primitive groups [Ca] to obtain tight lower bounds. We also prove some other upper and lower bounds for the more general cases.

We want to obtain tight bounds in the case when the group is transitive and not just primitive. It seems that the problem is significantly more difficult than the primitive group case and would need a better understanding of a particular group structure.

1.1.3 Testing of Distributions

One of the central problem in statistical hypothesis testing is to understand how many samples are needed to determine whether two distributions on a given set M are close or far in the ℓ_1 distance. The probability distributions are given in the form of an oracle $f : [n] \rightarrow [m]$ that one can query. The probability of an outcome $j \in [m]$ is the fraction of the domain that is mapped to j by f (hence sampling the distribution can be done by querying a uniformly random element of the domain).

Batu, Fortnow *et al* [BFR+] and Batu, Fischer *et al* [BFF+] gave tight bounds on the number of samples necessary, for the case when both the distributions have to be sampled and also for the case when only one of the distributions has to be sampled while the other distribution is known explicitly.

Fischer, Matsliah, de Wolf and I [CFMW] gave quantum algorithms for testing whether two such distributions are identical or at least ϵ -far in variation distance. We considered the case where one of the distribution is known, and show that testing can be done with roughly $m^{1/3}$ quantum queries, which is essentially optimal. In contrast, it is known that classical testing algorithms need about \sqrt{m} queries. This also helps us obtain tight bound on the quantum query complexity for testing graph isomorphism testing in the dense graph model when one graph is known in advance.

We are currently trying to obtain similar tight bounds on number of quantum queries to the distribution needed if both the distributions have to be sampled. It is known that classical testing algorithms need about $m^{2/3}$ queries in this case.

1.2 Sensitivity of Cyclically Invariant Boolean Function

Sensitivity of a Boolean function is the number of bits of the input on which the function is “sensitive”, i.e, it is the number of bits of an inputs such that if we change any one of those bits the function value changes. Block sensitivity is a similarly defined simple combinatorial complexity measures for Boolean functions. They were introduced to obtain non-trivial lower bounds in different models of computation. For example sensitivity helps in providing lower bounds on the time needed on a CREW PRAM model [CDR] while block sensitivity and CREW PRAM complexity are polynomially related [Nis]. Sun, Yao and Zhang [SYZ] used block sensitivity to obtain good lower bounds on the quantum query complexity for certain class of functions. (Quantum query model is another model of computation where a quantum machine computes the function by repeatedly asking quantum queries to the input. The quantum query complexity is the number of queries made.)

Unfortunately our knowledge about most of these measures is very limited, even for simple classes of functions, like cyclically invariant functions. Cyclically invariant functions are those functions whose value remains unchanged by cyclically shifts of the bits of the input. For a long time the sensitivity of cyclically invariant functions was thought to be greater than square-root of the size of the input [Tur, KK].

I [Cha] construct a cyclically invariant function which has sensitivity cube-root of the size of the input. Thus answering a couple of old questions in the negative. On the other hand for minterm-transitive functions (a natural class of Boolean functions including our example) this is shown to be tight. Similar techniques have helped to obtain tight bounds on block sensitivity on weakly symmetric functions [Sun].

Our understanding of combinatorial measures are so limited that almost everything in this field is open [C-MT]. I plan to use some advanced tools like Fourier analysis to study these combinatorial measures.

1.3 Quantum Query Complexity - Database Search

Radhakrishnan, Raghunathan and I [CRRS] consider the quantum database search problem, where we are given a Boolean function $f : [N] \rightarrow \{0, 1\}$, and are required to return an $x \in [N]$ (a target address) such that $f(x) = 1$. We consider the amount of error that we will make if we are just allowed to make just a few queries.

Grover [Gro] showed that there is an algorithm that after making one quantum query to the database, returns an $X \in [N]$ (a random variable) such that

$$\Pr[f(X) = 0] = \epsilon^3,$$

where $\epsilon = |f^{-1}(0)|/N$. Using the same idea, Grover (and subsequently Grover et al [GTP]) derived a t -query quantum algorithm that errs with probability only ϵ^{2t+1} . This method can be placed in a more general framework, where given any algorithm that produces a target state for some database f with probability of error ϵ , one can obtain another that makes t queries to f , and errs with probability ϵ^{2t+1} . For this method to work, we do not require prior knowledge of ϵ . Note that no classical randomized algorithm can reduce the error probability to significantly below ϵ^{t+1} , even if ϵ is known. So this problem is very important in understanding how powerful is a quantum machine compared to a classical randomized machine.

In our paper, we obtain *lower bounds* that show that the amplification achieved by these quantum algorithms is essentially optimal. We also present simple alternative algorithms that achieve the same bound as those in Grover [Gro], and have some other desirable properties. We then study the best reduction in error that can be achieved by a t -query quantum algorithm, when the initial error ϵ is known to lie in an interval of the form $[\ell, u]$. We generalize our basic algorithms and lower bounds, and obtain nearly tight bounds in this setting.

Our lower bound result uses the polynomial method of Beals et al [BBC+] combined with an elementary analysis based on the roots of low degree polynomials, but unlike previous proofs using this method, we do not rely on any special tools for bounding the rate of growth of low degree polynomials.

2 Electronic Commerce

Electronic commerce, commonly known as e-commerce consists of the buying and selling of products or services over electronic systems such as the Internet and other computer networks. The amount of trade conducted electronically has grown dramatically since the spread of the Internet. A wide variety of commerce is conducted in this way, spurring and drawing on innovations in electronic funds transfer, supply chain management, internet marketing, online transaction processing, electronic data interchange, automated inventory management systems, and automated data collection systems. These have in fact given birth to important algorithmic and complexity problems many of which had immediate effects to the industry.

The problems in e-commerce I have worked are related to multi-unit auctioning, market clearance pricing on a metric and allocation of objects according to the preference list of the people.

2.1 Improved Algorithms for Multi-Unit Auction with Unknown Supplies.

Mahdian and Saberi [MS] studied the problem of multi unit auction for perishable good where supplies arrives in an online fashion. This problem is motivated by application in advertisement auctions on the internet.

Let C be a perishable commodity (for example, a advertisement space on a search page) and a number of the bidder want one copy of the commodity each. So the bidders gives their bids. But we have no knowledge

how many copies of the commodity will be produced during the day. As soon as a copy gets produced it should either be given to a bidder or we discard it. To be fair to the bidders, at the end of the day all the bidders who have received a copy of the commodity are charged the same amount. On the other hand no bidder can be charged more than his bid. At the end of the day we compare the income with the best possible return that could have been achieved in hindsight.

Mahdian and Saberi [MS] gave a $1/4$ -competitive algorithm for the problem and also proved that the competitive ratio cannot be better than $e/(e + 1)$. Nikhil and I [CDe] have designed an algorithm that achieves $1/2$ -competitive ration. Our algorithm is much simple than the earlier algorithm. I worked on this problem as an intern at Microsoft Research India in summer 2007. Currently we are in the process of implementing this algorithm and applying for patent on this. There is still scope for improvement as our algorithm does not match the current lower bound. We are working to close the gap.

There are other variations of this problem which seem more natural and we would like to consider. For example if the number of copies produced each day is distributed according to a probability distribution.

2.2 Market Clearance Pricing on a Metric

At the annual budget the government of each country has allocated certain amount of money for buying petroleum for that year. During that year a certain amount of petroleum has been produced in various countries but to transport petroleum from a country to another country will cost additional money depending on various factors. Of course each country will like to pay the least amount possible for every unit of petroleum it buys. The question is does there exists a pricing scheme such that at the end all the amount of petroleum is bought and all the money sanctioned in the budgets are spend and everybody is happy. It is called a envy-free market clearance pricing scheme on a metric.

It was not even clear whether such a pricing exists; computing one quickly is a harder problem. Envy-free pricing has been studied a lot in the field of e-commerce and market clearance pricing is a very important component of modern day economics. But only in very restricted setting was the market clearance pricing studied in e-commerce earlier.

Devanur, Karande and I [CDK] proved that on any metric a envy-free market clearance pricing scheme exists, although our algorithm takes exponential time to compute it. We have also given a polynomial time algorithm to compute an approximation to the pricing scheme. Both the algorithms as well as the analysis are quite involved and the proofs uses various combinatorial tools.

2.3 Allocation problem in rental markets.

In an online DVD rental shop every customer gives their preference list for the DVD's and the rental shop is supposed to allocate the DVDs among its customers in an "optimal" way. But there is no universally accepted definition of "optimality". Since the DVDs has no preference on whom it is assigned to, the well known optimality conditions like stable matching does not make much sense in this case.

Different definition of optimality has been considered. Irving et al [IKM+] introduced a notion called popular matching based on voting schemes. The advantage of this optimality notion is that it is very easy to find the optimal if it exists. Unfortunately a given set of preference lists may not have a popular matching. Mahdian [Mah] proved that if the preference lists are drawn from a uniform distribution then with high probability there is a popular matching.

Dani and I [CDa] considered this problem when the DVD rental shop has more than one copy of each DVD. Mahdian's result cannot be applied in the case of multiple copies. In our paper we have given a characterization to popular matchings for multiple copies case just as in [IKM+]. Using it we prove a result

similar to Mahdian's: if we have k copies for each DVD and the preference lists are drawn uniformly at random then with high probability a popular matching exists.

We want to generalize our results to the case when the preference lists are drawn from an arbitrary distribution (not necessarily uniformly). In that case how many copies of each DVD do we need, to ensure that a popular matching exist with high probability. This will help the rental companies to understand how many copies of each DVD they need to have in their inventory.

3 Graph Algorithms

The subject of graph algorithms has been one of the most well studied subjects in the theoretical computer science and combinatorics and yet so many unexplored areas and unsolved problems exist. I have worked on a couple of problems in graph algorithms. The first is on rainbow coloring of graphs and the second on parametric weighted graphs.

3.1 Rainbow Coloring of Graphs

An edge-colored graph G is rainbow connected if any two vertices are connected by a path whose edges have distinct colors. The rainbow connectivity of a connected graph G , denoted $rc(G)$, is the smallest number of colors that are needed in order to make G rainbow connected. In addition to being a natural combinatorial problem, the rainbow connectivity problem is motivated by applications in cellular networks.

Fischer, Matsliah, Yuster and I [CFMY] gave the first proof that computing $rc(G)$ is NP-Hard. In fact, we prove that it is already NP-Complete to decide if $rc(G) = 2$, and also that it is NP-Complete to decide whether a given edge-colored (with an unbounded number of colors) graph is rainbow connected. On the positive side, we prove that for every $\epsilon > 0$, a connected graph with minimum degree at least ϵn has bounded rainbow connectivity, where the bound depends only on ϵ , and the corresponding coloring can be constructed in polynomial time. For the proof we needed a modified version of the Regularity Lemma which we also proved.

3.2 Parametric Weighted Graph

Now a days our graphs occurring in real life computations are becoming bigger and bigger and we want our computation done faster. One way out is doing a preprocessing that will help making the actual computation faster. Many times during the preprocessing time we don't know the exact input but we have some partial information about the input. To model this we considered parametric weighted graph.

A parametric weighted graph is a graph whose edges are labeled with continuous real functions of a single common variable. For any instantiation of the variable, one obtains a standard edge-weighted graph. Parametric weighted graph problems are generalizations of weighted graph problems, and arise in various natural scenarios. Parametric weighted graph algorithms consist of two phases. A preprocessing phase whose input is a parametric weighted graph, and whose output is a data structure, the advice, that is later used by the instantiation phase, where a specific value for the variable is given. The instantiation phase outputs the solution to the (standard) weighted graph problem that arises from the instantiation. The goal is to have the running time of the instantiation phase supersede the running time of any algorithm that solves the weighted graph problem from scratch, by taking advantage of the advice. This model was earlier studied by Carstensen [Car].

Fischer, Lachish, Yuster and I [CFLY] constructed several parametric algorithms for the shortest path problem. For the case of linear function weights we present an algorithm for the single source shortest

path problem. Its preprocessing phase runs in $\tilde{O}(V^4)$ time, while its instantiation phase runs in only $O(E + V \log V)$ time. The fastest standard algorithm for single source shortest path runs in $O(VE)$ time. For the case of weight functions defined by degree d polynomials, we present an algorithm with sub-exponential preprocessing time and instantiation time only $\tilde{O}(V)$. We also gave a randomized algorithm whose preprocessing time is $\tilde{O}(V^{3.5})$ and outputs an all-pairs shortest path solution, up to an additive constant error in $O(V^2)$ time.

Our algorithms shows that parametric algorithms can in fact help us in doing our computation faster even for classical problems like all pair shortest path. We plan to continue working on this area and understand the potential of parametric algorithms better.

4 Coding Theory

The subject of coding theory started with the seminal papers of Shannon and Hamming. They studied the problem of transmitting messages through a noisy channel. The goal is to find the best encoding assuming that the amount of noise is bounded. Various models on how the errors are introduced have been studied. In the past few decades unbelievable advancement has been done in the subject of coding theory.

I have studied two such models: one model is called the memoryless noisy channel introduced by Shannon and in the other model where the errors are introduced by a computationally bounded adversary.

4.1 Coding for memoryless noisy channel

In 1953 Shannon [Sha] introduced a different model where the message is sent through a finite memoryless noisy channel and the noise is only of a particular kind. Such a channel can be modeled as a bipartite graph (V, W, E) , where (v, w) is in E iff the letter w can be received when the letter v is transmitted. The goal then, is to encode messages as strings of letters from the input alphabet V and recover it from the received message. We want to use as few input letters as possible and still recover the intended message perfectly. Shannon [Sha] and Lovasz [L79] determined the best rate of transmission achievable under this model for several specific channels.

Radhakrishnan, Raghunathan, Sasatte and I [CRRS] consider the list-decoding version of this problem, studied by Elias[Eli]. In list decoding we output a small list of possible message words with the guarantee that the list contains the actual message word. Not much work has been done on list decoding in this model. Where as, in recent times extensive amount of work has been done in the area of list decoding in the Hamming model.

The $q/(q-1)$ channel is the complete bipartite graph $K(q, q)$ minus a perfect matching. We study the zero-error list-decoding capacity of this channel. For the $3/2$ channel, Elias gave lower and upper bounds for the zero error capacity when the list size is 2.

It turns out that this problem is related to the perfect hashing problem studied by Fredman-Komlos[Fre]. Their result proves that the zero-error capacity for the channel is exponentially small in q when the list size is q . We prove a generalization of their theorem using a completely different technique (using probabilistic methods) and give a similar upper bound for list size as large as $1.58q$. We conjecture that for any constant c , if the list size is less than cq then the zero error list-decoding capacity is exponentially small in q .

4.2 Coding against computationally bounded adversary

In the Hamming model the channel is allowed to tamper at most a certain number of bits of the message, but the channel can decide what bits to tamper arbitrarily. On the other hand in the Shannon model each bit

is tampered with a certain probability. But to model real life setting many different models have been introduced that are somewhat between Hamming model and the Shannon model. One such model is when the adversary is allowed to tamper a certain number of bits but the adversary is also computationally bounded.

Micali *et al* [MPSW] showed that in this model the rate that is, the amount of information that can be sent per bit of the encoded message is strictly greater than in the case of Hamming model.

In locally decodable codes the goal to decode a single bit of the message by looking at only a small number of bits of the encoded message. In the Hamming model it is a big open problem whether there exists locally decodable codes of polynomial rate when the number of queries is constant. Even for 3-query locally decodable codes the best lower bound on the rate we have is quadratic. Recently Yekhanin [Yek], Kedlaya and Yekhanin [KY] and Efremenko [Efr] proved a sub-exponential upper bound on the rate for 3-query locally decodable codes.

Hemenway and Ostrovsky [HO] consider the problem of public-key locally decodable codes when the adversary is computationally bounded. They proved that when the number of queries is logarithmic in the message size linear rate codes exists.

Bhattacharyya and I [BhC] tried to answer the question “Assuming certain cryptographic hardness is there a constant query locally decodable code against a computationally bounded adversary with polynomial rate?” This is an ongoing project and we have proved a number of lower bounds on the rate. Our proofs use techniques from cryptography, private information retrieval and LDCs. Our hope is that better understanding of this problem for computationally bounded adversary will help us understand the locally decodable codes in the Hamming model.

References

- [AF+00] N. Alon, E. Fischer, M. Krivelevich, and M. Szegedy. *Efficient testing of large graphs*. *Combinatorica* 20 (2000), 451-476.
- [AS+06] N. Alon, E. Fischer, I. Newman and A. Shapira. *A Combinatorial Characterization of the Testable Graph Properties: It’s All About Regularity* Proc. of STOC (2006), 251-260.
- [AS03] N. Alon and A. Shapira. *Testing subgraphs in directed graphs* STOC (2003), 700-709.
- [AS05] N. Alon and A. Shapira. *A Characterization of the (natural) Graph Properties Testable with One-sided Error*. Proceedings of the 46th IEEE FOCS (2005).
- [ArS] S. Arora and S. Safra: Probabilistic checking of proofs: a new characterization of NP. *J. ACM* 45 (1998) 70-122. (Prelim. FOCS 1992)
- [ALMSS] S. Arora, C. Lund, R. Motwani, M. Sudan and M. Szegedy: Proof verification and hardness of approximation problems. *J. ACM* 45 (1998) 501–555. (Prelim. FOCS 1992)
- [BaC] L. Babai and S. Chakraborty. *Property Testing of Equivalence under a Permutation Group Action*, To appear in *The ACM Transactions on Computation Theory (ToCT)*.
- [BhC] R. Bhattacharyya and S. Chakraborty. *Constant Query Locally Decodable Codes against a Computationally Bounded Adversary*, under preparation 2009.
- [BFL] L. Babai, L. Fortnow and C. Lund. *Nondeterministic exponential time has two-prover interactive protocols*. *Computational Complexity* 1 (1991), 3–40. (Prelim. FOCS 1990)

- [BFF+] T. Batu, E. Fischer, L. Fortnow, R. Kumar, R. Rubinfeld and P. White. *Testing random variables for independence and identity*, Proceedings of the 42nd FOCS (2001), 442-451.
- [BFR+] T. Batu, L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White. *Testing that distributions are close*, FOCS 2000, p. 259.
- [BBC+] R. Beals, H. Buhrman, R. Cleve, M. Mosca and R. de Wolf. *Quantum lower bounds by polynomials*, FOCS (1998): 352–361. quant-ph/9802049.
- [BLR] M. Blum, M. Luby and R. Rubinfeld, *Self testing/correcting with applications to numerical problems*, Journal of Computer and System Sciences 47 (1993), 549–595.
- [Ca] Peter J. Cameron. *Finite Permutation Groups and Finite Simple Groups*. Bull. London Math. Soc., 13 (1981), 1–22.
- [Car] P. Carstensen, *The complexity of some problems in parametric linear and combinatorial programming*, Ph.D. Thesis, Mathematics Dept., U. of Michigan, Ann Arbor, Mich., 1983.
- [Cha] S. Chakraborty. *On the Sensitivity of Cyclically Invariant Boolean Functions*, Computational Complexity Conference (CCC), 2005.
- [C-MT] S. Chakraborty. *Sensitivity, Block Sensitivity and Certificate Complexity of Boolean Functions*, Masters Thesis, 2005.
- [CRR] S. Chakraborty, J. Radhakrishnan and N. Raghunathan. *Bounds for Error Reduction with Few Quantum Queries*, 9th International Workshop on Randomization and Computation, 2005.
- [CRRS] S. Chakraborty, J. Radhakrishnan, N. Raghunathan and P. Sasatte. *Zero error list-decoding capacity of the $q/(q - 1)$ channel*, Foundations of Software Technology and Theoretical Computer Science (FSTTCS), 2006.
- [CDa] S. Chakraborty and V. Dani. *On popular matchings*, under preparation, 2009.
- [CDe] S. Chakraborty and N. Devanur. *Improved Algorithms for Multi-unit Auction with unknown supplies*, Online version available at <http://arxiv.org/abs/0901.1427>. Preliminary version appeared at the Forth Workshop on Ad Auctions 2008.
- [CDK] S. Chakraborty, N. Devanur and C. Karande. *Market Clearance Pricing in a Metric*, under preparation, 2009.
- [CFL+] S. Chakraborty, E. Fischer, O. Lachish, A. Matsliah and I. Newman. *Testing st -connectivity*, 11th International Workshop on Randomization and Computation, 2007.
- [CFLY] S. Chakraborty, E. Fischer, O. Lachish, and R. Yuster. *Two-phase algorithms for the parametric shortest path problem*, submitted 2009.
- [CFMW] S. Chakraborty, E. Fischer, A. Matsliah and R. de Wolf. *Quantum Query Complexity for Testing Distribution*, submitted 2009.
- [CFMY] S. Chakraborty, E. Fischer, A. Matsliah and R. Yuster. *Hardness and Algorithms for Rainbow Connectivity*, 26th International Symposium on Theoretical Aspects of Computer Science (STACS), 2009.

- [CDR] S. Cook, C. Dwork and R. Reischuk. *Upper and lower time bounds for parallel random access machines without simultaneous writes*, SIAM J.Comput. 15 (1986), no. 1, 87-97.
- [Efr] K. Efremenko. *3-Query Locally Decodable Codes of Subexponential Length* Electronic Colloquium on Computational Complexity Report TR08-069.
- [Eli] P. Elias. *Zero Error Capacity Under List Decoding*, IEEE Transactions on Information Theory, vol. 34, No. 5, (1988): 1070-1074.
- [FO1] E. Fischer. *The art of uninformed decisions: A primer to property testing*. The bulletin of the European Association for Theoretical Computer Science 75 (2001), 97–126.
- [FM06] E. Fischer and A. Matsliah. *Testing Graph Isomorphism*. SODA 2006.
- [Fre] M. Fredman and J. Komlós. *On the Size of Separating Systems and Families of Perfect Hash Functions*, SIAM J. Alg. and Disc. Meth., Vol. 5, No. 1 (1984): 61-68.
- [GGR96] O. Goldreich, S. Goldwasser and D. Ron. *Property testing and its connection to learning and approximation*. J. ACM 45, 1998. Preliminary version appeared in Proc. 37th FOCS, 1996.
- [Gro] L.K. Grover. *A different kind of quantum search*, March 2005. quant-ph/0503205.
- [GTP] L.K. Grover, T. Tulsi and A. Patel. *A new algorithm for directed quantum search*, May 2005. quant-ph/0505007.
- [HLNT] S. Halevy, O. Lachish, I. Newman and D. Tsur, *Testing Properties of Constraint-Graphs*, Proceedings of the 22nd IEEE Annual Conference on Computational Complexity (CCC 2007).
- [HO] B. Hemenway and R. Ostrovsky. *Public-Key Locally-Decodable Codes*, Crypto 2008, 126-143.
- [IKM+] R. W. Irving, T. Kavitha, K. Mehlhorn, D. Michail and K. Paluch. *Rank-maximal matchings*, ACM-SIAM Symposium on Discrete Algorithms, 2004.
- [KY] K. S. Kedlaya and S. Yekhanin. *Locally Decodable Codes From Nice Subsets of Finite Fields and Prime Factors of Mersenne Numbers*. IEEE Conference on Computational Complexity 2008: 175-186
- [KK] C. Kenyon and S. Kutin. *Sensitivity, block sensitivity, and ℓ -block sensitivity of Boolean functions*, Information and Computation, vol 189 (2004), no. 1, 43-53.
- [L79] L. Lovász. *On the Shannon Capacity of a Graph*, IEEE Trans. Inform. Theory, Vol. IT-25 (1979): 1-7.
- [Mah] M. Mahdian. *Random Popular Matchings* Electronic Commerce, 2006.
- [MS] M. Mahdian and A. Saberi. *Multi-unit Auction with Unknown Supply*, ACM Conference on Electronic Commerce, p.243-249, (2006).
- [MPSW] S. Micali, C. Peikert, M. Sudan and D. A. Wilson. *Optimal Error Correction Against Computationally Bounded Noise*, Theory of Cryptography Conference (TCC), 2005.

- [New] I. Newman. *Testing Membership in Languages that Have Small Width Branching Programs*, SIAM Journal on Computing 31(5):1557–1570, 2002.
- [Nis] N. Nisan. *CREW PRAMs and decision trees*, SIAM J. Comput. 20 (1991), no. 6, 999-1070.
- [R01] D. Ron. *Property Testing (a tutorial)*, In: *Handbook of Randomised Computing* (S. Rajasekaran, P.M. Pardalos, J.H. Reif and J.D.P. Rolim eds), Kluwer Press (2001), Vol II Chapter 15.
- [RS93] R. Rubinfeld and M. Sudan. *Robust characterization of polynomials with applications to program testing* SIAM Journal of Computing 25 (1996), 253-271. (Prelim: tech rep., Cornell University, 1993).
- [Sha] C.E. Shannon. *The zero error capacity of a noisy channel*, IEEE Trans. Inform. Theory, Vol. IT-2, no. 3, (1956): 8-19. (Reprinted in D. Slepian, Ed., *Key Papers in the Development of Information Theory*. New York: IEEE Press (1974): 112-123)
- [SYZ] X. Sun, A.C. Yao and S. Zhang. *Graph Properties and Circular Functions: How Low Can Quantum Query Complexity Go?*, IEEE CCC 2004.
- [Sun] X. Sun. *Block sensitivity of weakly symmetric functions* Theoretical Computer Science, vol 384. P 87-91.
- [Tur] G. Turán. *The critical complexity of graph properties*, Inform. Process. Lett. 18 (1984), 151-153.
- [Yek] S. Yekhanin. *Towards 3-query locally decodable codes of subexponential length*. STOC 2007: 266-274