

Evdokia Nikolova

Ph.D. Candidate

Massachusetts Institute of Technology

RESEARCH STATEMENT

Summary

My goal is to extend the frontier of theoretic research and unify methods from *combinatorial and continuous optimization, algorithms* and *economics* in what I call an emerging theory of *strategic algorithms for complex computing domains*. My graduate training in mathematics and theoretical computer science and my undergraduate training in applied mathematics and economics have provided me with a unique combination of research techniques and perspectives to achieve my research goals.

Approach and Impact. My approach emphasizes interdisciplinary collaboration in order to integrate problems and solution techniques from a variety of disciplines. In my work so far I have actively sought continual communication and collaboration with experts from computer science and economics, mathematics, engineering, artificial intelligence, networking and operations research. I have also formed collaborations with industry, which has focused my research on some directly applicable problems such as the design of sponsored search auctions at Google, prediction markets at Yahoo! Research and route planning at Mitsubishi Electric Research Lab. My contributions to stochastic combinatorial optimization have also been applied in other research projects of direct practical impact such as the CarTel project at MIT for traffic prediction and navigation.

Contributions. My work in computational economics proposes several new designs for mechanisms, as well as algorithms for computing the optimal strategies and profits for *sponsored search, prediction markets* and *network auctions* [5, 4, 3, 13, 1, 2]. These are all complex computational settings whose importance has exploded recently with the boom of computer networks and the Internet. In each, there are a number of selfish agents, each wanting to maximize its own profit. The challenges are to design mechanisms under which selfish behavior will lead to cooperation or efficiency of the system, and to compute optimal strategies in these mechanisms. This calls for an inherently interdisciplinary approach that combines economics, algorithms and optimization.

One of my main contributions to algorithms and optimization is discovering a novel connection between stochastic, combinatorial and nonconvex optimization. Traditional algorithm design has not taken into account the risk inherent in stochastic combinatorial optimization problems. My work considers realistic objectives that incorporate risk, and shows that these are equivalent to certain hard nonconvex problems from the realm of continuous optimization. Hence it requires a mix of discrete and continuous techniques and provides the first approximation algorithms for a wide class of stochastic and nonconvex optimization problems [11, 12, 8, 7, 10].

1 Computational Economics

My work in computational economics proposes several new designs for mechanisms, as well as algorithms for computing the optimal strategies and profits for *sponsored search*, *prediction markets* and *network auctions*. These are all complex computational settings whose importance has exploded recently with the boom of computer networks and the Internet. In each there are a number of selfish agents, each wanting to maximize its own profit. Thus the three main challenges are to:

1. *Model* realistically the economic interactions in these complex systems;
2. *Design* mechanisms with good properties (*e.g.*, truthfulness, revenue-maximization, etc.);
3. *Compute* the optimal strategies of agents.

Thus, an inherently interdisciplinary approach is necessary, that combines economics, algorithms and optimization.

Sponsored Search. Search engines like Google, Yahoo, MSN and others have revolutionized the way we access information. The business model for sustaining this incredible improvement in search efficiency is sponsored search: advertisers pay for search keywords and queries via an elaborate repeated auction. As sponsored search is a very recent and complex phenomenon, both economically and computationally, it has spurred urgent and exciting research in both mechanism design and algorithmic optimization. However we are still far from a good understanding of its strategic properties and implications. Google colleagues Jon Feldman, S. Muthukrishnan, Martin Pál and I [2] offer the first analysis of sponsored search that incorporates both multiple ad slots and advertiser budgets, for click-maximizing advertisers. We propose a truthful mechanism for allocating bidders to slots during a period of time such as a day (as opposed to a single round of the auction). We further prove that a natural revenue-maximizing mechanism is not truthful, however it has a Nash equilibrium identical to that of our truthful mechanism. In contrast to traditional mechanism design in Economics, our mechanism and its analysis are heavily based on *combinatorial optimization* and properties of *scheduling* on parallel machines.

Prediction Markets. In a separate direction of mechanism design, Sami and I [13] began investigating the design of prediction markets, which resemble financial markets but are specifically designed to aggregate traders' information and predict future events (such as, who will win the upcoming presidential election?). We defined an abstract model of prediction markets that had an intuitive geometric visualization and thus became a useful new tool for analysis of strategic behavior and properties of existing prediction market models. In particular, it allowed us to prove the absence of arbitrage (*i.e.*, one cannot profit from a net zero trade) in the dynamic parimutuel market, an earlier version of which had unintentionally admitted arbitrage opportunities. Y. Chen, L. Fortnow, D. Pennock and I further investigated computational issues in prediction markets [1]. We considered the auctioneer problem of risklessly matching trader orders in the context of permutation combinatorics (when the event being predicted is the final ranking of candidates such as the outcome of an election, or a tournament). We obtained a counterintuitive result that the auctioneer problem may be hard for a seemingly simple bidding language (betting on ordered pairs of candidates), and easy for an exponentially large language (betting on which subsets of candidates may finish in a given position).

Network Auctions. David Karger and I [5, 6] considered how much it would cost to purchase a path in a graph in which edges are selfish agents who need to be paid for routing packets through them. We investigated the payments under the celebrated Vickrey-Clarke-Groves (VCG) mechanism in several random graph models and random edge costs. We resolved a conjecture by Mihail, Papadimitriou and Saberi [9] in the negative and sharpened their result, showing that the average VCG overpayment in the Erdos-Renyi $G(n, p)$ random graph model is lower-bounded by constant, and is in fact exactly equal to $\frac{p}{2-p}$ for graphs with sufficiently large average degree. N. Immorlica, D. Karger, R. Sami and I [4] further investigated resulting payments for the same problem of path auctions, but instead under a first-price auction mechanism. We showed that the equilibrium payments are lower than the corresponding payments in the VCG mechanism.

In contrast to the VCG mechanism however, where it is optimal for edges to reveal their true costs and thus it is straightforward to reach an equilibrium, it is open to design a mechanism for converging to equilibrium under first-price auctions.

2 Algorithms and Optimization

One of my main contributions is discovering a novel connection between stochastic, combinatorial and non-convex optimization. Traditional algorithm design has not taken into account the risk inherent in stochastic combinatorial optimization problems. My work considers realistic objectives that incorporate risk, and shows that these are equivalent to certain hard nonconvex problems from the realm of continuous optimization. Hence it requires a mix of discrete and continuous techniques and provides the first approximation algorithms for a wide class of stochastic and nonconvex optimization problems, that evolved in a series of papers with colleagues Matt Brand, David Karger, Jonathan Kelner and Michael Mitzenmacher [11, 12, 8, 7, 10]. Some of our models and algorithms on stochastic shortest paths have been integrated as a key component in a traffic prediction and navigation system developed by a team of computer networking and artificial intelligence experts from *MIT's CarTel project*.

Stochastic Shortest Paths. In one of the earlier developments, we focused on finding shortest paths in the presence of uncertainty. In particular, the question was to find the shortest path from a given source node to a given destination node in a graph, in which the edge weights are random and coming from independent specified distributions. One possible analogy in practice would be finding a route in a transportation network with uncertain travel times that depend on traffic conditions. We classified the complexity and offered different algorithms for a number of possible user utility functions and edge-weight distributions [11]. We then designed and analyzed algorithms for a utility function that is related to a commonly used value-at-risk objective in finance [12], for a number of possible distributions. In particular, our goal was to maximize the probability that the path length does not exceed a given threshold t (in the traffic analogy, the objective is to “arrive on-time”). We showed that an equivalent mathematical formulation for this stochastic shortest paths problem with normally-distributed edge weights, was to find the path maximizing $\frac{t - \text{path mean}}{\sqrt{\text{path variance}}}$, that is maximizing a non-convex function over the set of feasible paths. We gave an exact algorithm with an unusual subexponential running time $n^{O(\log n)}$. The algorithm was based on enumerating certain extreme points from the projection of the feasible set onto a subspace.

Concave minimization. No efficient approximations had been known for the kind of nonconvex problem that arose in the stochastic shortest paths setting above, and we proceeded to give the first polynomial-time approximation algorithms for a much broader class of *low-rank concave minimization* [8] problems. These problems minimize a low-rank concave or quasi-concave function over a polytope, where the low-rank means that the function can be essentially thought as being defined over a constant dimensional subspace (where the dimension equals the rank). A very special instance of concave minimization is Linear Programming (which has rank 1), other instances include the stochastic shortest paths problem above (which has rank 2) as well as the very different 2-stage Stochastic Minimum Spanning Tree problem.

Our solution was based on a smoothed analysis of the problem, namely showing that when the input parameters defining the objective function (such as the mean and variance vectors above) are slightly perturbed, an exact algorithm for the problem runs in expected polynomial time. In particular, we gave a smoothed bound on the number of extreme points appearing on the projection of the feasible polytope onto a lower-dimensional subspace. The bound is polynomial in the original polytope dimension and the inverse of the perturbation size, and exponential in the dimension of the projection subspace. In addition to the novel proof techniques, our result had two important implications: that the exact algorithms (which are exponential in the worst case, like the simplex method for linear programming) are likely to have good practical performance. Secondly, it allowed us to design an efficient approximation scheme.

Stochastic and Nonconvex optimization. The randomized approximation scheme obtained via smoothed analysis above was restricted to functions with bounded gradient. It was thus complemented by my deterministic algorithms for stochastic combinatorial optimization incorporating risk [10], and its corresponding

nonconvex optimization equivalents with unbounded gradients. As C. Swamy and D. Shmoys conclude in their recent survey on stochastic optimization [14], “It would be interesting to explore stochastic models that incorporate risk,” in other words to move away from simply taking expectations of the random variables encountered. In addition, one of the key advantages of these algorithms is that they use the solution to the underlying deterministic algorithm as a black-box and thus immediately apply to a wide range of combinatorial problems.

3 Future Plan

By nature, Computational Economics—and the mechanisms and strategy analysis that my work on it provides—is interdisciplinary and requires tools from economics, algorithms and optimization. Similarly, my discovery of the novel connection between stochastic and nonconvex optimization, and the resulting algorithms for both, relied on a synthesis of continuous and discrete mathematics techniques. In the following two sections I will first outline questions and research directions that arise from my previous work, and in the last section I will explain how they fit together in what I see as a unifying theory of strategic algorithms.

Algorithms and Optimization. For a theoretical computer scientist, an $n^{O(\log n)}$ problem complexity is rare and intriguing. It remains an open question whether an exact polynomial-time algorithm exists for the stochastic shortest path problem above, or if not—to prove a super-polynomial lower bound on its solution.

Another challenging goal in the general stochastic optimization models would be to incorporate correlations between the participating probability distributions, as well as different distributions and objective functions. Further, all models I have described so far are offline in that the solution is computed with the input information available at the start. In many real-world settings new input information keeps arriving over time and we would be interested in an online solution that makes use of subsequent information. David Karger and I [7] are already exploring one such model of stochastic shortest paths, in which upon arrival at a node the agent sees the actual lengths of all adjacent edges and needs to make an optimal (adaptive) decision of where to go next; such extension to online optimization would also be natural in the general models of stochastic and nonlinear optimization I have considered.

A key feature in my work is the interplay between continuous and discrete methods. One example that triggered our hardness results on concave minimization [8] was the connection between convexity and submodularity, which I plan to explore further. The theory of submodularity has been dubbed the discrete analog of convexity theory, although the direct correspondence seems to stop at the very special class of submodular set functions, which are defined only over the vertices of the unit hypercube. The significance of submodularity seems to reach much deeper. In particular an intriguing connection is its use in both combinatorial optimization and economics & game theory, and a better understanding has the potential of bridging together and advancing the applications in these areas.

Computational Economics. In the work on VCG path auctions I described above, Karger and I provided experimental results that the overpayment in power-law graphs generated by the preferential attachment model is bounded by 1. Preferential attachment has become one of the most common mathematical models for complex real-world networks such as the webgraph, the Internet AS-level topology, social networks and others. This model has proven to be very challenging mathematically, requiring new techniques for analysis to answer even some of the most basic questions such as what is the diameter and degree sequence of the graph. I remain interested in proving a constant bound on the VCG overpayment in such graphs—this would have an exciting impact for justifying the use of VCG mechanisms in economics settings with real-world networks. A number of related open questions remain as well, such as considering more complex models of the utilities and economic relationships between autonomous agents in the network, in the spirit of my work with Alex Hall and Christos Papadimitriou [3].

The above research areas illustrate the richness of design and optimization problems in the field of computational economics, as well as the importance of rigorous theoretical analysis for a multitude of practical applications. I plan to continue investigating the computational and incentive issues both in the models above

as well as any new problems that arise from the intersection of economics and computer science. For example, sponsored search auctions have become the defining characteristic of the Internet search industry, which in turn is becoming increasingly influential in our access and perception of information. Understanding the incentives guiding search and sponsored search, and informing new sponsored search auction designs would thus have a huge impact not only academically, but in a much broader socio-economic and policy-making aspect.

Strategic Algorithms for Complex Computing Domains. I view my work in Computational Economics and in Algorithms & Optimization as a single long-term research direction: In a more realistic setting there are typically multiple users affecting each others strategic choices, while each individually is trying to solve a complex optimization problem. My long-term research goal is to develop a unifying theory for the strategies and algorithms arising in complex computing domains.

As an example and first steps in this direction, Nicolas Stier-Moses and I have begun exploring stochastic shortest paths equilibria: my previous work concerned the optimal path choice of a single user in a stochastic network. Clearly though, different users would increase congestion and affect each other's optimal path choice. We are starting to examine the existence of equilibria and algorithms for computing the optimal paths comprising an equilibrium. Of course, this model is not restricted to shortest paths per se, and would extend to any setting with multiple users who use or compete for common resources and thereby affect each other's choices and the parameters of the environment.

Another example comes from prediction markets: some of my work considers a deterministic offline scenario in which bidders submit bids and the auctioneer needs to make a decision (e.g., maximize revenue) with the full knowledge of all bids. While this may already be a hard optimization problem, it would be more realistic to consider a stochastic model of bid arrival; bids together with a nonlinear (risk-averse) utility for the auctioneer and perform the optimization and strategic analysis in that setting.

Sponsored search in itself offers a rich body of unexplored problems, in which queries, bids, customer traffic and advertiser revenues are all stochastic, and call for a strong unified approach of analysis and solutions using high-performance Algorithms, Economics and Optimization.

Thus, each area I have described in the previous sections is a component of a broader theory of *strategic algorithms for complex computing domains* which I am keenly interested in and plan to pursue in the long-term.

References

- [1] Y. Chen, L. Fortnow, E. Nikolova, and D. M. Pennock. Betting on permutations. In *EC '07: Proceedings of the 8th ACM conference on Electronic commerce*, pages 326–335, New York, NY, USA, 2007. ACM Press.
- [2] J. Feldman, S. Muthukrishnan, E. Nikolova, and M. Pál. A truthful mechanism for ne ad slot scheduling. *Manuscript*, 2007.
- [3] A. Hall, E. Nikolova, and C. Papadimitriou. Incentive-compatible interdomain routing with linear utilities. In *Lecture Notes in Computer Science (3rd International Workshop On Internet And Network Economics)*, volume 4858, pages 232–244, 2007.
- [4] N. Immorlica, D. Karger, E. Nikolova, and R. Sami. First-price path auctions. In *EC '05: Proceedings of the 6th ACM conference on Electronic commerce*, pages 203–212, New York, NY, USA, 2005. ACM Press.
- [5] D. Karger and E. Nikolova. Brief announcement: on the expected overpayment of VCG mechanisms in large networks. In *PODC '05: Proceedings of the twenty-fourth annual ACM symposium on Principles of distributed computing*, pages 126–126, New York, NY, USA, 2005. ACM Press.
- [6] D. Karger and E. Nikolova. On the expected VCG overpayment in large networks. In *45th IEEE Conference on Decision and Control*, pages 2831–2836, 2006.
- [7] D. Karger and E. Nikolova. Exact algorithms for the Canadian Traveller problem on special graphs. *Manuscript*, 2007.
- [8] J. A. Kelner and E. Nikolova. On the hardness and smoothed complexity of quasi-concave minimization. In *Proceedings of the 48st Annual Symposium on Foundations of Computer Science*, Providence, RI, USA, 2007.
- [9] M. Mihail, C. Papadimitriou, and A. Saberi. On certain connectivity properties of the internet topology. In *FOCS '03: Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, page 28, Washington, DC, USA, 2003. IEEE Computer Society.
- [10] E. Nikolova. Black-box approximation algorithms for stochastic and nonconvex combinatorial optimization. *Manuscript*, September 2007.
- [11] E. Nikolova, M. Brand, and D. R. Karger. Optimal route planning under uncertainty. In *Proceedings of the International Conference on Automated Planning and Scheduling*, 2006.
- [12] E. Nikolova, J. A. Kelner, M. Brand, and M. Mitzenmacher. Stochastic shortest paths via quasi-convex maximization. In *Lecture Notes in Computer Science 4168 (ESA 2006)*, pages 552–563, Springer-Verlag, 2006.
- [13] E. Nikolova and R. Sami. A strategic model for information markets. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 316–325, New York, NY, USA, 2007. ACM Press.
- [14] C. Swamy and D. B. Shmoys. Approximation algorithms for 2-stage stochastic optimization problems. *ACM SIGACT News*, 37(1):33–46, 2006.