Research Statement

Elliot Anshelevich
September 2011

My research interests center in the design and analysis of algorithms, especially in large decentralized networks. My main focus has largely been on algorithmic game theory, and more specifically on networks involving strategic agents. In order to understand such networks, I also study diffusion processes in social networks, and more generally approximation algorithms.

An increasing number of networks at the forefront of scientific research, and of great importance to our world, consist of interactions of a large number of independent agents *that we cannot control.* Contrary to this, traditional study of network design assumes that there is a single network designer, and that its design is accepted and implemented by all parts of the network. Similarly, study involving network management assumes that some protocol is implemented in the network, and all parts of the network follow this protocol (or almost all parts, while a small part of the network is malicious). The lack of control of the agents can occur because of the self-interest of these independent agents (such as in peer-to-peer file sharing schemes, business contracts between companies, transportation networks consisting of individual cars, or social networks representing relationships between groups of people), or because the process involved inherently lacks control by outside forces (such as in epidemics spreading through a population). The outcomes and dynamics of such behavior often have very different properties from those of the centrally designed or managed networks. Therefore, dealing with such systems requires quite different methods and considerations, leading to new and exciting algorithmic questions. This fact and the enormous growth in the complexity of networked systems have made game theoretic techniques and networks of self-interested agents become a common theme in computer science.

The major goal of my research is to design algorithms and develop mechanisms for dealing with and understanding such networks. While these networks cannot be fully controlled, they can often be *influenced* in a limited way, sometimes resulting in dramatic improvement of the global network behavior. I have focused on algorithms determining how such influence can be applied, as well as studying agent dynamics, and understanding the properties of equilibrium outcomes formed by such agents. The emphasis of my work is usually on algorithms and network properties with provable guarantees. My current major research topics are described below.

Network Formation by Strategic Agents. (See presentation slides at `http://www.cs.rpi.edu/~eanshel/abstracts.html#GNFG`) Network design and management is one of the major areas where the appearance of strategic agents requires the development of novel theoretical and practical methods to ensure good network properties. During the last decade, I studied many different contexts in which independent self-interested agents create a network together. For these contexts, some of my major goals were to confirm the existence of equilibrium solutions, to understand the quality of these solutions compared to the optimum centralized solution (as measured by the classic notion of "price of anarchy"[1]), as well as to study convergence properties of various game dynamics. Understanding the price of anarchy is often crucial in game theoretic contexts, since it gives a bound on how much the stable solutions can be improved by influencing the agent behavior. In a paper together with Dasgupta, Kleinberg, Tardos, Wexler, and Roughgarden, we

---

[1] The price of anarchy is the ratio of the social welfare in the optimum centralized solution with that of the worst Nash equilibrium.

also introduced the notion of "price of stability"[2], and I was also able to provide algorithms that generate *approximate* Nash equilibria, especially desirable in instances where exact Nash equilibrium does not exist. The price of stability and quality of approximate equilibria are important measures, since they correspond to the quality of a stable solution that could be formed if some external authority were able to influence the players, e.g., by giving them incentives not to deviate (this corresponds to approximate equilibrium) or by suggesting a good stable solution to the players (thus realizing the price of stability).

More specifically, I began by formulating and analyzing the *Connection Game* [1], a general model of network formation that captures the essence of agent interactions when building a network together. In this game, a network of potential links with link building costs is given, and each agents desires to connect certain terminal nodes, which may differ for every agent [1,2]. To do this, and to lower their costs for building the network, agents are willing to cooperate and share the cost of certain links [1, 2]. Together with my students at RPI, I continued to extend the techniques for analyzing these games to more general connectivity requirements, such as forming networks resilient to link failure [3], forming groups [4], and disconnecting terminal nodes instead of connecting them [5]. The algorithms I designed make it possible to find good stable (or approximately stable) solutions that can be implemented in the system. My work on the Connection Game and its extensions spawned a plethora of followup work on network formation games (see for example [6, Chapters 17,19] and references therein; my papers [1–5] on this topic have well over 500 citations), and the price of stability has become a standard notion in algorithmic game theory [6].

Autonomous Systems and the Internet. The modern Internet is composed of tens of thousands of sub-networks called *Autonomous Systems* (AS), each under a single administrative authority with its own distinct goals in controlling the traffic entering and leaving its network. The complex system of business relationships between various Internet entities (e.g., Autonomous Systems, enterprize networks, residential customers) is at the heart of Internet connectivity. In a sense, this system of contracts *is* the Internet, since a link between two routers is useless without a contract between the companies that own these routers. Because of this, it is crucial that we understand not only the behavior of these self-interested AS's, but also how we could subtly influence them to improve the quality of the Internet. In addition to looking at abstract network formation games, I also considered specifically the games played by Autonomous Systems and ISP's when forming contracts in the Internet [7,8], as well as when pricing those contracts [9]. In all of the above contexts, I was able to show either that the stable solutions of the game have good quality, or ways to influence the agents in the network in order to form good global solutions. My recent work on pricing and routing [9] stands in contrast to more traditional "selfish routing", since in [9] the agents only determine the next-hop of their route, not the entire route from start to destination. This represents current Internet routing much more closely, and may give us insight into the behavior and routing policies of ISP's.

Contribution Games in Social Networks. (See presentation slides at `http://www.cs.rpi.edu/~eanshel/abstracts.html#Effort`) Very recently, I have begun to analyze several kinds of network contribution games, where agents not only form local links to their neighbors, but also determine the strength of these links. More specifically, each node/agent has a budget of effort that it can allocate to different incident edges representing its friendships, relationships, collaborations, etc. The amount of effort allocated to a link by its endpoints determines the strength of this relationship, as well as the happiness of the participants with this relationship.

---

[2] The price of stability is the ratio of the social welfare in the optimal centralized solution with that of the *best* Nash equilibrium.

This simple framework and its many variants illustrate many of the complexities that arise when agents in a social network choose the intensity of the links they form. Together with Martin Hoefer, I was able to show the existence of various types of coalitional equilibria in such games, as well as to quantify the price of anarchy, and prove various convergence properties of agent interactions [10]. These games have a close relationship to the classic problem of stable matching with cardinal preferences, which arises in such applications as matching medical residents to hospitals, online dating, and kidney exchange. I was able to use similar techniques to provide price of anarchy bounds for both exact and approximate stable matching [11,12]. The "freedom of choice" that the agents enjoy in contribution games by determining the strength of their ties with others, and not simply forming links, results in greater insight into agent behavior, and into the properties of social networks. I believe that extending these games to hypergraphs, determining effective ways to influence the agents of these games, and considering a combination of agent incentives (including average distance to other nodes, betweenness, etc.), has great potential to improve our understanding of the structure of some social networks. Together with one of my Ph.D. students, we are currently preparing a manuscript for submission that quantifies techniques for dealing with several of these complexities [13].

Selfish Flow over Time and Applications in Transportation. Routing games are used to understand the impact of individual users' decisions on network efficiency, when each user decides the route it wants to take in a network. Most prior work on routing games uses a simplified model of network flow where all flow exists simultaneously, and users care about either their maximum delay or their total delay. Both of these measures are surrogates for measuring how long it takes to get all of a user's traffic through the network. We attempt a more direct study of how competition affects network efficiency by examining routing games in a flow over time model (also known as dynamic flow) [14, 15]. While much work exists on flow over time, and on selfish routing, very little prior work exists on *selfish flow over time.* Such flows differ greatly from more traditional "static selfish flows," since we assume that the network flow can change over time, and that it takes time for traffic to traverse a link. Selfish flow over time is much better at modeling traffic congestion, and is of great interest to transportation scientists [15].

Diffusion Processes and Approximation Algorithms. (See presentation slides at `http://www.cs.rpi.edu/~eanshel/abstracts.html#firefighter`) In addition to algorithmic game theory, I am also extremely interested in various diffusion processes in networks. The high-level questions I am interested in can often be summed up as: How can we take advantage of the network structure (and in the case of strategic agents, the interests of the agents) to ensure good global behavior? This question is especially explicit in the context of influence and information propagation in networks. Most recently, I studied the Firefighter problem, and provided the first approximation algorithms for several versions of this problem [17] (some of my results were subsequently improved by [18]). In the Firefighter problem, an epidemic (which also can be thought of as a rumor, information, computer virus, etc.) is spreading through a population, and because of a limited supply of vaccine (or simply a lack of time to administer it), it is necessary to decide whom to vaccinate. The goal of the vaccinations is to stop the spread of the epidemic through the underlying social network of people's interactions, and save as many lives as possible. Together with one of my students, I am now using the insight from this problem to analyze the ways in which evacuation alerts spread through a population, and to design efficient techniques to stop false alarms. This research builds on my previous expertise both in graph theory [19] and in diffusion processes [20].

References

[1] Elliot Anshelevich, Anirban Dasgupta, Éva Tardos, and Tom Wexler. Near-Optimal Network Design with Selfish Agents. In *Theory of Computing,* Volume 4 (2008), pp. 77-109. Conference version appeared in *Proc. 35th ACM Symposium on Theory of Computing* (STOC 2003).

[2] Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The Price of Stability for Network Design with Fair Cost Allocation. In *SIAM Journal on Computing,* Volume 38, Issue 4 (November 2008), pp. 1602-1623. Conference version appeared in *Proc. 45th Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2004).

[3] Elliot Anshelevich and Bugra Caskurlu. Price of Stability in Survivable Network Design. *Theory of Computing Systems,* Volume 49, Number 1 (July 2011), pp. 98-138. Conference version appeared in *Proc. 2nd International Symposium on Algorithmic Game Theory* (SAGT 2009).

[4] Elliot Anshelevich and Bugra Caskurlu. Exact and Approximate Equilibria for Optimal Group Network Formation. *Theoretical Computer Science,* to appear. Conference version appeared in *Proc. 17th Annual European Symposium on Algorithms* (ESA 2009).

[5] Elliot Anshelevich, Bugra Caskurlu, and Ameya Hate. Strategic Multiway Cut and Multicut Games. In *Proc. 8th Workshop on Approximation and Online Algorithms* (WAOA 2010).

[6] N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani (eds.), *Algorithmic Game Theory,* Cambridge University Press.

[7] Elliot Anshelevich, Bruce Shepherd, and Gordon Wilfong. Strategic Network Formation through Peering and Service Agreements. *Games and Economic Behavior,* Volume 73, Issue 1, September 2011, Pages 17-38. Conference version appeared in *Proc. 47th Annual IEEE Symposium on Foundations of Computer Science* (FOCS 2006).

[8] Elliot Anshelevich and Gordon Wilfong. Network Formation and Routing by Strategic Agents using Local Contracts. In *Proc. 4th International Workshop On Internet And Network Economics* (WINE 2008).

[9] Elliot Anshelevich, Ameya Hate, and Koushik Kar. Strategic Pricing in Next-hop Routing with Elastic Demands. In *Proc. 4th International Symposium on Algorithmic Game Theory* (SAGT 2011).

[10] Elliot Anshelevich and Martin Hoefer. Contribution Games in Social Networks. *Algorithmica,* to appear. Conference version appeared in *Proc. 18th Annual European Symposium on Algorithms* (ESA 2010).

[11] Elliot Anshelevich, Sanmay Das and Yonatan Naamad. Anarchy, Stability, and Utopia: Creating Better Matchings. In *Proc. 2nd International Symposium on Algorithmic Game Theory* (SAGT 2009).

[12] Elliot Anshelevich and Sanmay Das. Matching, Cardinal Utility, and Social Welfare. *ACM SIGecom Exchanges,* Volume 9.1, June 2010.

[13] Elliot Anshelevich, Onkar Bhardwaj, and Martin Hoefer. Contribution Games in the Presence of Complex Utility Structures and Unequal Reward Sharing. In preparation.

[14] Umang Bhaskar, Lisa Fleischer, and Elliot Anshelevich. A Stackelberg Strategy for Routing Flow over Time. In *Proc. ACM-SIAM Symposium on Discrete Algorithms* (SODA 2011).

[15] Elliot Anshelevich and Satish Ukkusuri. Equilibria in Dynamic Selfish Routing. In *Proc. 2nd International Symposium on Algorithmic Game Theory* (SAGT 2009).

[16] Elliot Anshelevich, Bugra Caskurlu, and Ameya Hate. Partition Equilibrium Always Exists in Resource Selection Games. In *Proc. 3rd International Symposium on Algorithmic Game Theory* (SAGT 2010).

[17] Elliot Anshelevich, Deeparnab Chakrabarty, Ameya Hate, and Chaitanya Swamy. Approximability of the Firefighter Problem: Computing Cuts over Time. *Algorithmica,* to appear.

[18] P. Chalermsook and J. Chuzhoy. Resource Minimization for Fire Containment. In *Proc. 21st Annual ACM-SIAM Symposium on Discrete Algorithms* (SODA 2010).

[19] Elliot Anshelevich, Adriana Karagiozova. Terminal Backup, 3D Matching, and Covering Cubic Graphs. *SIAM Journal on Computing,* Volume 40, Issue 3 (2011), pp. 678-708.

[20] Elliot Anshelevich, David Kempe, and Jon Kleinberg. Stability of Load Balancing Algorithms in Dynamic Adversarial Systems. *SIAM Journal on Computing*, Volume 37, Issue 5 (January 2008), pp. 1656-1673.