Research Statement

Elliot Anshelevich

My research interests center in the design and analysis of algorithms, especially for large decentralized networks. Recently I have focused on network design questions and on algorithms for networks involving strategic agents.

The past few years have seen a great increase in the amount of research concerning the behavior of strategic agents in networks. Contrary to the bulk of research in network algorithm design, many real networks currently contain a large number of self-interested independent agents which neither blindly follow an imposed protocol nor maliciously try to hurt the network. The Internet, for example, is developed, built, and maintained by a large number of agents, all of which act in their own interests and have relatively limited goals. This fact and the enormous growth in the complexity of networked systems have made game theoretic techniques and self-interested agents an emerging theme in computer science. The outcomes of such self-interested behavior often have very different properties from the centrally designed or managed networks. Therefore, dealing with and managing such systems requires quite different methods and considerations. For instance, a routing protocol for self-interested agents (such as BGP) has to take into account the preferences of the agents and not just try to form good routes. A major part of my work has been concerned with the question of how to get these agents to behave appropriately without exerting too much pressure on them.

Network Design and Management with Strategic Agents

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. Consider the case where a network of potential links with fixed link building costs is given, and there is a collection of k agents, each of whom desires to connect certain terminal nodes which may differ for every agent. Each agent ensures that these nodes are connected by offering to pay for some links in the network to be built, but it would like to pay as little as possible, so it would prefer for other agents to pay for the links it needs. This models the creation process of some local networks. If this were a centrally designed network, finding a set of links which has the least cost and yet connects all the terminal nodes (and thereby maximizes social welfare) would be a well-studied optimization problem known as the minimum-cost Steiner forest. The agents are not centrally controlled, however, and their behavior can result in outcomes which are very different from the network which maximizes social welfare. In fact, there are stable solutions formed by these agents which differ from the centralized optimum by a factor of k, and there are even networks where no stable solutions exist. By stable solutions in this context we mean outcomes where each agent is satisfied with the current network. In particular, such solutions are known as Nash equilibria.

Together with Dasgupta, Tardos, and Wexler, I explored the *price of stability* in this problem, which represents the tradeoff between the global properties of the network (in this case, the overall cost), and network stability. The price of stability is a very different notion from the price of anarchy, and yet similar to it in concept. My research was one of the first attempts to quantify the price that a network pays for the stability of such self-interested agents. I derived bounds on the price of stability for this general model, as well as much stronger bounds for the model where all agents wish to connect a personal node to a common source node, such as when the network being designed is a multicast network. We also showed how to find and implement good approximate Nash equilibria in the cases where finding an exact one is NP-Hard. These results let us find and suggest good stable outcomes to the agents, and thereby, ensure that the resulting network is good without exerting any force on the agents. Our methods make heavy use of techniques from both game theory and approximation algorithms. In essence, because of the tradeoff between how stable a solution is and how close it can be to the centralized optimum, optimizing both at once means finding a bicriteria approximation. Despite our results, for the general case it is still an open question whether there exists a $(1+\varepsilon)$-approximate Nash equilibrium that is close to optimum in cost.

This question illustrates well the relation between game theory and approximation algorithms in this network formation context. One of my major research goals is to explore other notions of approximation in this and similar game theoretic contexts. Suppose there exists a centralized authority that desires good global network properties, but has very limited power to affect the agents. A statement about low price of stability, such as the ones above, means that such an authority can encourage the agents to form a globally good outcome by simply suggesting a good stable solution to them, and the only power that this authority would need to accomplish this is the power of simultaneous suggestion (such as through an underlying network communication protocol). I would like to consider other possible powers that this authority might have, such as the ability to pay agents money from a limited budget, the ability to destroy or create a small number of links, the ability to take over a constant number of agents, or no abilities at all besides limited computational power, and explore how much this authority would be able to alter the behavior of the agents. In the above context, each of these abilities corresponds to a different notion of approximation, so this research would have to make heavy use of approximation algorithmic techniques. This research should provide a greater understanding of how an entity (whether altruistic, selfish, or malicious) could affect the above network design process to its benefit.

There is another interesting power that a central authority may possess, which I have already begun studying. The model above allows the agents to decide how much they contribute to any link, and in particular an agent can use an existing link to connect its terminals even if it pays nothing for it. It makes sense to ask the question here: what if instead there were a fixed rule for sharing the cost of a link among the agents using it? With Dasgupta, Kleinberg, Tardos, Wexler, and Roughgarden, I explored the case where this cost sharing is done on an link-by-link basis via a fair-division scheme that can be derived from the Shapley value. This new sharing mechanism induces fairness, as well as ensuring that a stable solution always exists. In fact, we prove that one always exists with the cost at most $(\log k)$ away from the centralized optimum in the general case. This result is tight and holds even if the costs of the links are not fixed, but are instead arbitrary concave functions of the number of agents using that link.

If we consider convex link costs, however, then we are faced with a problem where the costs can now be thought of as latencies, and the agents are attempting to route their traffic on a path of smallest latency. Therefore, the questions described above do not apply only to network design, but to routing on an existing network as well. For this routing model, we proved results akin to the network routing results of Roughgarden and Tardos, and showed various bounds on the price of stability for different classes of latency functions.

Besides this cost-sharing mechanism, I would like to explore the existence of other such mechanisms, which might guarantee even better global properties. These mechanisms must be fully distributed, use only local information, and use only limited power. There is a tradeoff here between the power and the knowledge available to the mechanism, on the one hand, and the quality of the solutions it can produce, on the other hand. One of my goals is to examine this tradeoff and find mechanisms with little power and yet significant effect.

Agents With Local Interaction So far in this section we considered only network design questions where an agent has global information and effect, since an agent can choose the route to connect its terminals through the entire network, and can pay for links anywhere in the network.

In many real networks such as the network of Autonomous Systems (AS) which comprises the Internet, however, each agent only operates locally. With Shepherd and Wilfong, I recently began exploring a model of contracts and incentives which define the local routing relationships between AS's. As before, my main goal is to determine when limited power can be used to manipulate the AS's into a favorable configuration. We have produced some preliminary results in the form of sufficient conditions for an entity with no power to manipulate the AS's into an optimal (for the entity) and stable solution, and given an algorithm for finding the best stable solution that ends in all the traffic getting to its destination. My next goal here is to consider this model in the presence of entities with different objectives and powers, and determine how much they are able to manipulate the situation. This should tell us if the current arrangements in the Internet are stable with respect to manipulation, or if perhaps a few AS's could do much better if they join forces, etc.

Optical Network Design

Networks with strategic agents introduce new concerns into network design and management questions. The same is true for optical networks. The profoundly different technology used in optical networks requires very different approaches from traditional networks, since, for example, instead of capacities we now have the wavelength assignment problem. An optical network also often has extra structure, such as the line system network which must consist of edge-disjoint paths called line systems. With Zhang, I considered the question of designing such an optical network with small building cost. Note that the cost in this case is not simply a sum of edge costs, but is a rather interesting and complicated function. We developed both approximation and exact algorithms for various cases of this problem. Ours was one of the first theoretical works to address network design of such optical networks. There are many more questions left to answer in this context, and in general, I am interested in addressing important network design and routing questions for emerging technologies, such as optical networks, wireless networks, or sensor networks.

Decentralized Processes in Networks

Load Balancing and Packet Routing One of my major general interests is proving static properties of dynamic decentralized processes. In load balancing, jobs are inserted into a network of processors, and a distributed algorithm moves them along the capacitated links to balance the overall load. In an abstract sense this is similar to packet routing, where the jobs correspond to packets and the goal of the algorithm is to get them to their destinations. With Kempe and Kleinberg, I showed that a simple decentralized algorithm keeps the system stable, i.e., it avoids having too many jobs/packets at one node of the network, in both load balancing and packet routing contexts. I am very interested in exploring other good properties of such algorithms, perhaps combining stability with small packet delivery time, or showing stability for the case of multiple packet destinations. Designing decentralized algorithms, and proving such global properties about them, is a long-standing area that only becomes more important as more systems and protocols become decentralized. Note that decentralized interaction of strategic agents also fits into this framework.

Influence Propagation The 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 propagation in networks. The usual models of influence propagation apply not only to social networks, but also to cascading failures and gossip protocols in various computer networks, as well as information propagation in sensor networks. I am interested in analyzing this process of spreading information and influence, and finding how it can be manipulated or enhanced.