Network Coding - How much do I win the non-cooperative game?

A. Vijayaraghavan arvindv@cse.iitm.ernet.in[1]
Advisor: C. Pandu Rangan rangan@shiva.iitm.ernet.in[1]

Indian Institute of Technology Madras

Abstract. We consider the problem of network coding in a single-source multicast network where each of the vertices act greedily in order to maximize their own payoff. We introduce a model for network coding from a game-theoretic perspective with strategies for each of the nodes(intermediate and receivers). The information flow in the network is defined by the topology of the network and the strategies that the intermediate nodes and multicast receivers choose to pursue. We then study the existence and characterization the pure Nash Equilibrium for these games. To the best of my knowledge, this is the first time that such a game-theoretic treatment of network coding has been done.

1 Introduction

Sending data through a network is a task that pervades much of our modern lives. The question of how to send data efficiently remains a central question in several fields of research, from information theory to computer networking. Store-and-forward had been the predominant technique for transmitting information through a network, where independent data streams maintain their identity and are unaltered, though they may share network resources. However, this has been proven to be non-optimal especially in the case of multiple sources and multicasting. Network coding, which was first introduced in [1] offers a new paradigm for network communications and has generated abundant research interest in information theory, coding theory, networking and many other fields. Network coding promises to transfer information optimally (at the rate given by the max-flow min-cut theorem to each receiver) by allowing intermediate nodes to mix and perform coding on incoming data streams such that outgoing data streams are functions(say linear) of these incoming streams. Hence the incoming data streams lose their identity and the receivers decode the orginal messages after receiving the modified data streams.

Network Coding has primarily been studied from a more algorithmic perspective where the objective has been to come up with protocols which ensure the transfer of data from the source to the multicast receivers optimally with respect to information flow rate [4, 5]. The nodes in most of the earlier works are honest, while some of them analyse and devise rate-optimal protocols in the midst of Byzantine adversaries [8] and Passive adversaries. [11, 2, 6] look at a game-theoretic treatment to networks. However, network coding has not been studied much from a game-theoretic standpoint.

In real life, each of the nodes in such a network would correspond to routers or computers owned by people who would prefer to act greedily in a way that is beneficial to them. This encourages a game-theoretic treatment where each of the nodes chooses a strategy from his strategy set to maximize his individual payoff. [3] studied the scenario where each of the multicast receivers greedily routes his flows to maximize his individual throughput. However, the decisions were made only by the receivers and it also required knowledge of the entire network topology. But in most practical scenarios, each packet(data) sent from a node would entail some cost which needs to be borne by the receivers and other nodes. Thus it is more realistic to model each of the non-sink nodes also as players in the game. Further, the networks that we encounter in real life are too big for the receivers to have a complete idea of the network topology. With these in mind, we introduce a model in which each of the receivers and intermediate nodes choose a strategy which would decide which of the edges incident on that node would be used to send data. Further, the model also studies the problem of network coding in a decentralized setting, where it is necessary for a node to know only its immediate surroundings(neighbours) by using the notion of Randomized Network coding. To the best of our knowledge, this is the first time that such a game-theoretic treatment of network coding has been done.

2 Model

We adopt the model of [10] where the network is represented as a directed graph $G(V, E)$. For sake of simplicity, we assume that there is only one sender and there are $d \geq 1$ receiver nodes and the flow rate from the source is ω. We denote by $U = \{1, 2, \ldots, n\}$ the set of all the nodes apart from the source(having at least one incoming edge). With each node $i \in U$, we associate the strategy set $S_i = 2^{E_i}$, where $E_i = \{e \in E | e = (r, i) \text{ is incident on node } i\}$. We assume that each of the edges have unit capacity and edges with multiple capacity are replaced by unit-capacity edges. The strategy set S_i corresponds to choosing a subset of edges along which node i wants to receive data streams. Once, we describe a strategy profile μ of an instance of the game, by specifying the strategies of all players, the underlying directed graph $H(\mu)$ representing the information flow is determined(using algorithms in [9, 7]). We also assume that the intermediate nodes usually use random network coding[7] to encode information from its incoming nodes for its outgoing nodes or we get the flows from deterministic network coding algorithms. However, these may not give the optimal flow[3]. Given the graph H, the sinks which receive the data is known ($\sigma(i) = 1$ if $maxflow(i) \geq \omega$) and these nodes get a utility u_i. Let us denote by R the set of all sinks that receive the ω data streams. Further the cost of communicating along these edges are borne by the nodes in the network. The payoff for each node i is represented by $p_i(\mu) = p_i(H(\mu))$ where $\mu = (s_1, s_2, \ldots, s_{|U|})$ represents the strategy profile of that game instance. In most instances, this payoff p_i depends on the cost allocation mechanism for each node x_i and the utility u_i only.

The payoff function defined in the model is generic and different scenarios can be captured effectively by an appropriate choice of the payoff function. The decentralized setup described above can be captured by using local payoff functions. Consider the simple payoff function defined by

$$p_i(\mu) = k\Big(\big|\{l \in E(H(\mu))|head(l) = i\}\big| - \big|\{l' \in E(H(\mu))|tail(l') = i\}\big|\Big) + \sigma_i(\mu)u_i(\mu)$$

where $head(l)$ and $tail(l)$ represent the node that the link l starts from and the node that l is incident on respectively. This captures the simple scenario where each communication along an edge costs a constant k coins for the receiving node and the sinks have utility u_i for receiving a ω-flow. On the other hand, non-local payoffs can capture the centralized network coding setting effectively.

2.1 Results on Nash Equilibria

We present some of the results that analyse and characterize the Nash Equilibrium of the game describe above. We have omitted proofs to save space in this extended abstract.

Theorem 2.1. *Any game instance, over a single-sink network always has a Pure Nash equilibrium.*

Theorem 2.2. *For any single-source network, there exist local payoff functions such that the throughput-maximizing configuration is a Pure Nash Equilibrium.*

The significance of this result is that for any network, it is always possible to come up with local payoff functions such that every player is personally satisfied with a throughput-maximizing configuration (pure nash equilibrium).

3 Progress and Research directions

As indicated by the panel, we have looked at the problem of cost sharing and mechanism design for network coding. We have obtained distributed algorithms for implementing the Shapley value mechanism and the Marginal cost mechanism and obtained bounds on the message complexity and hardness of approximation results for finding the welfare-maximizing flows.

Theorem 3.1. *Given the set of ω-edge disjoint paths $\mathcal{P}$ in a directed acyclic graph G (n nodes and m links), the costs x_i and the receiver set R^* can be calculated in $O(n^2)$ messages($O(1)$ length) per link for the Marginal cost mechanism and $O(n)$ messages per link for the Shapley value mechanism.*

Theorem 3.2. *In the network represented by the directed acyclic graph $G(U, E)$, the welfare-maximizing network flow(with coding) is hard to approximate within a ratio of ϵ.*

The details of these results will be communicated shortly. For further research, it would be interesting to come up with bounds for the price of anarchy especially for local payoff functions and also analyse these games in the repeated games model. It would also be useful to study the computational complexity of algorithms for finding a pure nash equilibrium in the network coding games defined above and analyzing the Nash equilibrium of games for certain classes of payoff functions.

References

[1] R. Ahlswede, N. Cai, S.-Y. Li, and R. Yeung. Network information flow. *IEEE Transactions on Information Theory*, 46(4):1204–1216, 2000.

[2] A. Archer, J. Feigenbaum, A. Krishnamurthy, R. Sami, and S. Shenker. Approximation and collusion in multicast cost sharing. *Archer, A., Feigenbaum, J., Krishnamurthy, A., Sami, R., and Shenker, S. (2002). Approximation and Collusion in Multicast Cost Sharing, submitted. http://www.cs.yale.edu/homes/jf/AFKSS.ps*, 2002.

[3] S. Bhadra, S. Shakkottai, and P. Gupta. Min-cost selfish multicast with network coding. *Information Theory, IEEE Transactions on*, 52(11):5077–5087, 2006.

[4] N. Cai and R. Yeung. Network error correction, part i: Basic concepts and upper bounds. *Communications in Information and Systems*, 6(1):19–36, 2006.

[5] N. Cai and R. Yeung. Network error correction, part ii: Lower bounds. *Communications in Information and Systems*, 6(1):37–54, 2006.

[6] J. Feigenbaum, C. H. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmissions. *J. Comput. Syst. Sci.*, 63(1):21–41, 2001.

[7] T. Ho, R. Koetter, M. Medard, D. R. Karger, and M. Effros. Benefits of coding over routing in a randomized setting. *International Symposium on Information Theory*, pages 442–452, 2003.

[8] S. Jaggi, M. Langberg, S. Katti, T. Ho, D. Katabi, and M. Medard. Resilient network coding in the presence of byzantine adversaries. *To be presented at INFOCOM-07*, 2007.

[9] S. Jaggi, P. Sanders, P. Chou, M. Effros, S. Egner, K. Jain, and L. Tolhuizen. Polynomial time algorithms for multicast network code construction. *IEEE Transactions in Information Theory, July 2003.*, page 422, 2003.

[10] R. Koetter and M. Medard. Beyond routing: An algebraic approach to network coding. *21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOMM)*, 1:122–130, 2002.

[11] T. Roughgarden and E. Tardos. How bad is selfish routing? In *IEEE Symposium on Foundations of Computer Science*, pages 93–102, 2000.