INFO4990 Literature Review for the project "Game Theoretic Analysis of Internet Network Problems"

Enoch Lau
SID 200415765

30 April 2007

Abstract

In a network, greedy, independent agents aim to minimise their own personal cost (such as travel time between source and destination) without regard to wider, societal impacts of their behaviour. The inefficiency due to this behaviour can be studied through such measures as the price of anarchy, which is the ratio of the cost of the worst-case Nash equilibrium to that of the optimal flow. One application of this game theoretic analysis is in allowing network operators to charge users equitably and profitably for multicast traffic sent through their networks, because multicast traffic along a link cannot be simply attributed to one particular user. The aim of the project is to extend and modify existing models of multicast pricing, to improve the resultant societal cost even with greedy, independent agents and their applicability to real-world multicast uses. Theoretical properties, such as time and network overhead complexity, will be examined to determine the tractability of the models.

Contents

1 Introduction

The Internet is fertile ground for the intersection of theoretical computer science and the economic notions of game theory, as introduced by von Neumann and Morgenstern. As it is a network of wildly diverse users, some in competition and others in collaboration, game theoretical analysis of the Internet sheds light on the inefficiencies caused by the primacy of self-interest, and economic measures that can be deployed to encourage selfish individuals to behave in a socially optimum way without the need for centralised control [14]. Other aspects of game theory that impact on concepts in computer science include auctions and mechanism design, where the aim is to create a game with desired outcomes [12].

We begin by describing the price of anarchy, which is a measure that captures the inefficiency of selfish routing over the socially optimum behaviour. Various results, depending on the underlying properties of the network, in particular, the steepness of functions used to attribute cost, have been proved in the literature. Secondly, we examine proposals for charging users for the use of a multicast stream that alleviate the inefficiencies posed by selfish routing, as a particular subproblem in the field. Thirdly, we consider the types of games where a Nash equilibrium exists, and the computational complexity of computing them, as their study would be fruitless if they did not exist, or at least could not be computed. Finally, we briefly summarise some of the notable techniques used in some of the proofs in the literature.

2 The Price of Anarchy

2.1 The Price of Anarchy on a Traffic-Based Model

2.1.1 The Model

We first describe a network model that has been used as the basis for several results about the price of anarchy, as formulated in [17]. In a network G, there are source-destination pairs $\{s_i, t_i\}$, called *commodities*, between which the network must route r_i units of traffic. For any commodity i, there may be several *paths* to choose from, and traffic may be split between these paths. *Flow* is the aggregation of the routes chosen by a large number of agents, and each agent is assumed to control a negligible amount of flow. A flow is said to be *feasible* if for all i, all required units of traffic are routed, that is:

$$\sum_{P\in\mathscr{P}_i} f_P = r_i \tag{1}$$

Edges have an associated *cost function* $c_e(\cdot)$ that may depend on the amount of flow on the edge, and the cost of a path P is denoted by c_P. If f_P is the amount of agents using path P, and f_e is the amount of flow using paths that include the edge e, then we can say that the total cost of a flow f is:

$$C(f) = \sum_{P\in\mathscr{P}} c_P(f) f_P = \sum_{e\in E} c_e(f_e) f_e \tag{2}$$

An *instance* is a triple of the form (G, r, c).

2.1.2 Nash Equilibrium

In general, a game consists of n players, and player i may choose from a set of strategies S_i [14]. Each player has a payoff function $u_i : S_1 \times \cdots \times S_n \to \mathbb{R}$ that measures the amount that player

i obtains given the strategies chosen by the other players as well. Rational behaviour is deemed to be one in which no player has incentive to deviate, and this is the Nash equilibrium. Formally, it is a combination of strategies $x_1 \in S_1, \ldots, x_n \in S_n$ for which $u_i(x_1, \ldots, x_i, \ldots, x_n) \geq u_i(x_1, \ldots, x_i', \ldots, x_n)$ for all i and $x_i' \in S_i$. Specifically, flows at Nash equilibrium are flows such that if an agent were to switch paths, its cost would only increase.

2.1.3 Price of Stability and Price of Anarchy

The *price of stability* is defined as the ratio between the best flow at Nash equilibrium and that of an optimal flow [8]. This has a natural interpretation, as game theory, in particular, mechanism design, was used early on to design strategies in games to create certain desired outcomes.

With structures like the Internet, however, where the network was not designed, but instead evolved, a different measure of the anarchy that would result from the lack of centralised control was needed. The *price of anarchy* is defined as the worst-case ratio between the cost of a flow at Nash equilibrium and that of an optimal flow [18, 7]:

$$\rho(G, r, c) = \frac{C(f)}{C(f^*)} \tag{3}$$

2.1.4 Bounds on the Price of Anarchy

[18] proves that for edge cost functions of the form $c_e(x) = a_e(x) + b_e$ for some $a_e, b_e \geq 0$, the price of anarchy is:

$$\rho(G, r, c) \leq \frac{4}{3} \tag{4}$$

This cannot be improved, as demonstrated by Pigou's Example, which has a price of anarchy of 4/3.

Further from linear edge cost functions, the paper then demonstrates that if $\mathscr{I}_p$ is the set of instances with cost functions that are polynomials up to a degree of p, then the price of anarchy is:

$$\sup_{(G,r,c)\in\mathscr{I}_p} \rho(G, r, c) = \Theta\left(\frac{p}{\ln p}\right) \tag{5}$$

However, for more general sets of functions, we introduce the concept of the anarchy value, which, for a set, captures how ill-behaved a set of allowable functions is. If c is a cost function, the *anarchy value* of c is defined, in [17], to be:

$$\alpha(c) = \sup_{x,r\geq 0} \frac{r \cdot c(r)}{x \cdot c(x) + (r - x)c(r)} \tag{6}$$

If $\mathscr{C}$ is a set of cost functions, then the *anarchy value* of $\mathscr{C}$ is:

$$\alpha(\mathscr{C}) = \sup_{0\neq c\in\mathscr{C}} \alpha(c) \tag{7}$$

This is useful for characterising the price of anarchy for general sets of allowable functions. If $\mathscr{C}$ is a set of cost functions with anarchy value $\alpha(\mathscr{C})$, and (G, r, c) is an instance with cost functions in $\mathscr{C}$, then:

$$\rho(G, r, c) \leq \alpha(\mathscr{C}) \tag{8}$$

In general, this can be unbounded; that is, the cost of the Nash equilibrium can be arbitrarily worse than that of the optimal flow.

2.1.5 Bicriteria Bound in General Networks

[18] proves that if f is a flow at Nash equilibrium for (G, r, c), and f^* is feasible for $(G, 2r, c)$, then $C(f) \leq C(f^*)$; that is, a flow at Nash equilibrium is no worse than an optimal flow forced to route twice as much traffic.

2.1.6 Marginal Cost Function

[17] defines a *marginal cost function* (which captures the cost for new flow, as well as the additional congestion costs for existing traffic) for a differentiable function c:

$$c^* = \frac{d}{dx}(x \cdot c(x)) = c(x) + x \cdot c'(x) \tag{9}$$

The work shows that a flow f feasible for an instance (G, r, c) is optimal iff it is at Nash equilibrium for the instance (G, r, c^*), a result which relates the optimal and Nash flows.

2.2 Variations on the Traffic Model

2.2.1 Non-atomic Congestion Games

Non-atomic congestion games (NCGs) are a generalisation of the above traffic model onto games that do not involve networks. [19] demonstrates that many of the above results carry through to NCGs. This result is interesting, because the non-atomic congestion game is the limit of the general congestion game as n, the number of players, goes to infinity.

2.2.2 Approximate Nash Flows

A flow f feasible for an instance (G, r, c) is at ϵ-approximate Nash equilibrium if the cost of a path is smaller than a factor of $(1+\epsilon)$ of any other path for that commodity. [18] extends the previous result for networks with linear cost functions to show that the price of anarchy is (tightly) bounded from above by:

$$C(f) \leq \frac{4+4\epsilon}{3-\epsilon} \cdot C(f^*) \tag{10}$$

2.2.3 Atomic Selfish Routing

In comparison to the traffic routing model, we now have a finite number of agents, each of which controls a non-negligible amount of flow. There are two cases:

- Splittable flow:
  - The difference between this case and the traffic routing model is that in the atomic splittable case, the agents will not take into account the congestion it causes for its own traffic.

    - In [18], it was shown that the bicriteria bound and the upper bound on the price of anarchy both extend to the atomic splittable case.
- Unsplittable flow:
    - In [18], the bicriteria bound and the upper bound on the price of anarchy both do not extend to the atomic unsplittable case.
    - However, the *best* Nash flow obeys both, subject to certain constraints regarding the cost functions and all commodities share the same source vertex).
    - In [2], the authors were able to further this work and find expressions for the price of anarchy for the general network case for unsplittable flow.
        * For linear latency functions, for weighted demand, the price of anarchy is $\frac{3+\sqrt{5}}{2}$ for pure and mixed strategies, and for unweighted demand, the price of anarchy is 2.5 for pure strategies only.
        * For polynomials of degree d latency, the price of anarchy is $d^{\Theta(d)}$ for pure and mixed is bounded from above by $O(2^d d^{d+1})$ and from below by $\Omega(d^{d/2})$.

3 Charging for Multicast

3.1 Multicast

Traditional Internet routing is based on unicast, where for each receiver, one stream of data is sent from the source. This can result in duplication of data, especially when the same information is sent down the same link twice. Multicast alleviates this problem by building a tree, typically to minimise the multicast rate or to minimise the multicast cost. When confronted with self-interested agents, the selfish multicast problem is one of designing economic incentives so that self-interested participants are guided to behave in a way that leads to the social optimum. Crucially, the difference between selfish multicast routing and the selfish routing model presented above is that in this case, per-user costs for an edge are typically non-increasing in the number of users of the edge.

3.2 Selfish Users

[3] models a selfish multicast game in which the cost of each edge is split evenly between those who use it; this is based on the Shapley value [20]. Players are added one-by-one into the network, and they use best response dynamics to pick the path to the source. There are two cases top consider, the integral (unsplittable) multicast case, and the fractional (splittable) multicast case.

3.2.1 Integral Multicast Game

The paper presents a proof that the price of anarchy of a Nash equilibrium obtained from best response dynamics has an upper bound of:

$$O(\sqrt{n}\log^2 n) \tag{11}$$

and a lower bound of:

$$\Omega\left(\frac{\log n}{\log\log n}\right) \tag{12}$$

3.2.2 Fractional Multicast Game

The price of anarchy of a Nash equilibrium in this case has an upper bound of:

$$O(\log n) \tag{13}$$

3.3 Incorporating Receivers' Utilities

In [5], we have the added consideration of factoring the users' utilities, u, of participating in the multicast transmission. A cost-sharing mechanism is a pair of formulas that calculate $x(u)$, the cost that an agent or user has to pay for participation in the network, and $\sigma(u)$, a binary vector that indicates whether a user is included or not. There are several desirable properties for the choice of mechanism:

- Strategyproofness
- No Positive Transfers
- Voluntary Participation
- Consumer Sovereignty
- Budget-Balance
- Efficiency

However, there are complications caused by combinations of these properties; for example, there is no strategyproof mechanism that is both budget-balanced and efficient [10]. These considerations suggest only the Marginal Cost mechanism, which is based on the marginal contribution to the overall welfare provided by a user having non-zero utility for the transmission, and the Shapley Value mechanism [20], which is based on sharing the cost of a link by those users downstream from the link. The Marginal Cost mechanism needs only a linear number of messages, while the Shapley Value mechanism requires a quadratic number of messages for communication.

3.4 Selfish Information Flows

Selfish information flows, as described by [1], assumes that each part of the flow (which would correspond to a single packet) attempts to minimise its own cost as it travels through edges of varying weights, instead of minimising the cost from the users' perspective. In [8], two cases were studied for selfish information flows: the uncapacitied and capacitied networks.

3.4.1 Uncapacitied Networks

The Shapley value-based cost sharing does not enforce min-cost multicast flows, because the Shapley value is inherently a local cost sharing method. Instead, a pair of path-based LPs were formulated, assuming the unique replicable and encodable properties of information flows. The solution of the LP gives the shadow prices, which intuitively measure the amount that the objective function will grow if a constraint is relaxed by one unit. The shadow price method guarantees a price of stability of 1, but the price of anarchy is only 1 in specific examples, for example, where we adopt a two-tier charging scheme.

3.4.2 Capacitied Networks

The LP is modified to include constraints relating to finite edge capacities, and a vector t in the dual, that represents the taxes on edges. The addition of taxes, which can then be refunded, guides the formation of an equilibrium, and results in a strategyproof system, where edges have no incentive to lie about their cost. After the taxes are removed, the multicast flow is still at Nash equilibrium, and budget-balance is still maintained.

4 Existence and Computational Complexity of Nash Equilibria

4.1 Classical Results

Nash proved that all games have mixed Nash equilibria [11]. His proof was based on Kakutani's fixed point theorem [6], which is a generalisation of Brouwer's fixed point theorem. The proofs were all non-constructive proofs, and did not give any insight into how one might construct such equilibria [13].

Closer to the selfish routing problem, Rosenthal, in a seminal paper [15], proved that non-cooperative games, that is, congestion games, always have pure Nash equilibria. Monderer and Shapley [9] proved that any exact potential game (games are associated with a potential function that guides the players towards equilibrium) is isomorphic to a congestion game. However, [4] shows that this requirement that the differences be equal when considering potential functions is too strict, and argues that using *general potential functions*, where all that matters is the sign, is valid.

4.2 Polynomial-Time Algorithms in Restricted Classes

[4] shows that there was a polynomial time algorithm for calculating a pure Nash equilibrium in the symmetric network case, that is, the case where all players have the same source and destination. The paper also shows that it is possible to approximate the Nash equilibria of non-atomic congestion games in strongly polynomial time.

4.3 Complexity of Calculating a Nash Equilibrium in the General Case

The difficulty with finding Nash equilibria is that their proof is non-constructive. [4] proves that computing a Nash equilibrium in the general network case is PLS-complete (which, intuitively, corresponds to "as hard to compute as any object whose existence is guaranteed by a potential function").

4.4 Existence of Nash Equilibria for the Selfish Multicast Problem

In [3], it was noted that the integral mutlicast game is, by definition, a congestion game, and thus by Rosenthal, has a pure Nash equilibrium. For the fractional multicast game, the potential function is:

$$\Phi = \sum_{e \in s} \left(\sum_{j=1}^{n_e(s)} \sum_{i=1}^{n_e+1-j} c_e \frac{f_{e,j} - f_{e,j-1}}{i} \right) \tag{14}$$

and because this is an exact potential function, it is isomorphic to a congestion game. Finally, the weighted fractional multicast problem has a Nash equilibrium, with the proof via Kakutani's fixed point theorem [6].

4.5 Complexity of Calculating an Optimal Flow

[17] states that it is possible to compute, up to an arbitrarily small error term, a Nash flow for a given instance in time polynomial in the size of the instance and the number of bits of precision required.

5 Techniques

In general, much of the work done at the intersection of theoretical computer science and game theory utilises linear programming, convex optimisation and approximation algorithms to achieve the proofs required. However, we also examined a few specific techniques that were particularly useful in the works examined.

5.1 Lower Bounding the Price of Anarchy with Simple Networks

In order to find the price of anarchy, the literature, such as [18], proves upper bounds for the price of anarchy, and then shows that this bound cannot be improved upon by way of example, thus proving that the upper bound is tight. [18] proves that such examples can be realised by way of the simplest of networks:

- For a set $\mathscr{C}$ that contains the constant functions, the worst possible value of ρ is obtained by a single-commodity instance on a two-node, two-link network (up to an arbitrarily small additive factor).
- If we relax the constraint that $\mathscr{C}$ must contain the constant functions, the worst possible value of ρ is obtained by a single commodity instance on a network of parallel links.

In fact, the price of anarchy is independent of the complexity of the network topology [16].

5.2 Reducing the Size of a Linear Program

In [Li], the resulting linear program can be exponential to the graph size, thus destroying the possibility of solving it in polynomial time. The paper utilised two techniques, the second of which was found to be more effective:

- Reformulate an equivalent LP that uses (a polynomial number of) links instead of paths.
- Use Lagrangian relaxation, where, generally, constraints are removed and penalties are placed into the objective function, together with an iterative, subgradient algorithm.

5.3 Computing the Price of Anarchy

In [2], the authors noted the general techniques used to find the price of anarchy, in both pure and mixed strategy cases. For the pure case, the authors compared the delay encountered by each agent to the delay it would encounter if it changes to the optimal route, and these delays were combined in some weighted fashion. For mixed strategies, the authors considered the latency of the expected load of each edge, and this was augmented to the expectation of the total frequency.

6 Conclusion

We have described the scene that is at the intersection of computer science and game theory. Unregulated flows produce calculably worse results than regulated, optimal flows, and this is captured by the price of anarchy. The price of anarchy can be calculated for networks with different assumptions. On specific form of flow considered is the multicast routing problem, where several economic techniques were introduced to influence the selfish agents into globally optimum behaviour. Important when working with Nash equilibria, their existence under certain assumptions was discussed, as well as the occasions when an efficient algorithm for calculating them exist. Finally, some interesting techniques used in the papers were pointed out as being of potential reuse in future work.

References

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

[2] B. Awerbuch, Y. Azar, and A. Epstein. The price of routing unsplittable flow. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 57–66, New York, NY, USA, 2005. ACM Press.

[3] C. Chekuri, J. Chuzhoy, L. Lewin-Eytan, J. S. Naor, and A. Orda. Non-cooperative multicast and facility location games. In *Proceedings of the 7th ACM conference on Electronic commerce*, pages 72–81, New York, NY, USA, 2006. ACM Press.

[4] A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *Proceedings of the thirty-sixth annual ACM Symposium on Theory of Computing*, pages 604–612, New York, NY, USA, 2004. ACM Press.

[5] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmissions. *Journal of Computer and System Sciences*, 63:21–41, 2001.

[6] S. Kakutani. A generalization of brouwer's fixed point theorem. *Duke Mathematical Journal*, 8:457–459, 1941.

[7] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 404–413, 1999.

[8] Z. Li. Min-cost multicast of selfish information flows. In *Proceedings of IEEE INFOCOM 2007*, 2007.

[9] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.

[10] H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory*, 18(3):511–533, 2001.

[11] J. E. Nash Jr. Equilibrium points in n-person games. *Proc. Nat. Acad. Sci. US*, 36:48–49, 1950.

[12] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1/2):166–196, 2001.

[13] M. J. Osborne. *A Course in Game Theory*. MIT Press, Cambridge, Massachusetts, 1994.

[14] C. H. Papadimitriou. Algorithms, games, and the internet. In *Proceedings of the thirty-third annual ACM Symposium on Theory of Computing*, pages 749–753, New York, NY, USA, 2001. ACM Press.

[15] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

[16] T. Roughgarden. The price of anarchy is independent of the network topology. In *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, pages 428–437, New York, NY, USA, 2002. ACM Press.

[17] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, Cambridge, Massachusetts, 2005.

[18] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.

[19] T. Roughgarden and E. Tardos. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior*, 47(2):389–403, 2004.

[20] L. S. Shapley. A value for n-person games. *Contributions to the Theory of Games*, 2(28):307–317, 1953.