Price of Anarchy*

Mohammad Ali Safari

Department of Computer Science
University of British Columbia
safari@cs.ubc.ca

1 Introduction

Internet users, auction bidders, and stock buyers are examples of agents who are selfishly trying to maximize their own benefit without necessarily caring the global objective of the underlying game such as the overall load on the network or the overall satisfaction of agents. How does these two points of view (agents vs. global) correlate? How much does the society suffer by the lack of coordination between players?

The optimal social utility function happens when we have a single authority who dictates every agent what to do. In contrast, when agents choose their own action, we should study their behavior and compare the obtained social utility with the optimal one.

There are currently two main approaches for studying behavior of agents. Koutsoupias and Papadimitriou [KP99] assume that agents play according to a Nash equilibrium; thus they, pessimistically, consider the Nash equilibrium that gives the worst social utility and compare it with the optimal social utility. In contrast, Goemans et al. [GMV05] consider games in which agents are reluctant to choose mixed strategies and repeatedly choose their action by playing their best pure responses even if there is no pure strategy Nash equilibrium in the game.

In this note we review these two different approaches: *price of anarchy* and *price of sinking*. In section 2 we survey some results related to the price of anarchy. In particular, we discuss previous researches on unsplittable flow problem, congestion games, and a class of games, called valid utility systems, whose price of anarchy is constant. Then, in Section 3, we introduce the price of sinking and its value on unsplittable routing games, congestion games, and valid utility systems. The two approaches, price of sinking and price of anarchy, are compared at the end.

2 Price of Anarchy

It's fair to assume that in an n-agent game agents play according to Nash equilibria of the game. So, one can study the effect of selfish agents by considering

* CPSC532A Course Project, Fall 2005.

the value of social utility on all possible Nash equilibria of the game. In particular, Koutsoupias and Papadimitriou [KP99] initially suggested the ratio of the optimal social utility and the worst Nash equilibrium as a measure, in fact a lower bound, for the amount of suffering to a society due to lack of coordination in a game. They support their definition by computing it on simple unsplittable flow games in which we have parallel links between two vertices. Later on, other researches extended their result to more complex games. In the next subsection we review the price of anarchy on unsplittable flow problem. As unsplittable flow games are special versions of congestion games, we then mention some result regarding the price of anarchy on symmetric and asymmetric congestion games. Finally, we present a broad class of games, called valid utility systems, whose price of anarchy is bounded.

2.1 Unsplittable Flow Problem

Koutsoupias and Papadimitriou [KP99] studied the price of anarchy on unsplittable flow games. In the general version of this game we have n agents; agent i wants to transfer an amount w_i of flow between a source s_i and a destination t_i and chooses, as his strategy, a path from s_i to t_i to transfer his flow. Unsplittable means agents are not allowed to break their flow and transfer it partially across more than one paths. As a result, each edge e carries some amount of traffic, say l_e, going through it, and there would be a latency caused by this edge which is a function of l_e and is denoted by $f_e(l_e)$. Assume agent i is using path P_i to transfer his flow. The latency corresponding to him would be the sum of latencies over all edges in P_i, i.e. $w_i \sum_{e \in P_i} f_e(l_e)$. The aim of each agent is to minimize his associated latency whereas the social objective is to minimize either the total latency or the maximum latency over all edges in the network.

Lets works with a toy example to illustrate the problem. Assume, as in Fig. 1, they are two parallel links between two vertices s and t and two agents with equal amounts $w_1 = w_2 = 1$ of flows and $f_e(x) = x$. Moreover, assume the social goal is to minimize the maximum latency. The optimal routing is to transfer agent i's flow through edge i, so the total latency would be $\max_e f_e(l_e) = 1$. A strategy in which each agent is choosing one of the links uniformly at random is an equilibrium. With probability $1/2$, one of the edges will carry two units of flow and with probability $1/2$ both edges carry one unit. So, the expected maximum latency equals $2 \times 1/2 + 1 \times 1/2 = 3/2$. Consequently, the price of anarchy in this network is at least $3/2 = 1.5$. Koutsoupias and Papadimitriou [KP99] show that the price of anarchy is exactly 1.5 in this case.

What if there are m parallel links between s and t and m agents with equal amounts of flows? One equilibrium happens when each agent chooses one of the m links uniformly at random. It can be proven that, expectedly, $\theta(\log m / \log \log m)$ of agents choose a common link causing a total latency of at least $\theta(\log m / \log \log m)$. But, in an optimal algorithm, agent i chooses link i to transfer his flow; hence, the total latency would be 1. Consequently, the price of anarchy is at least $\Omega(\log m / \log \log m)$. Koutsoupias and Papadimitriou [KP99] propose this lower bound as well as an upper bound $3 + 4\sqrt{4m \ln m}$ for the price

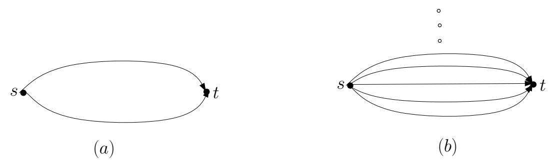

Fig. 1. A network with (a)2 (b)m parallel links.

<table>
<tr><th></th><th>Measure</th><th>Lower Bound</th><th>Upper Bound</th></tr>
<tr><td>2 parallel links</td><td>Maximum Latency</td><td colspan="2">1.5[KP99]</td></tr>
<tr><td>2 parallel links(different speeds)</td><td>Maximum Latency</td><td colspan="2">1.618[KP99]</td></tr>
<tr><td>m parallel links</td><td>Maximum Latency</td><td colspan="2">$\theta(\log m/\log\log m)$[CV02]</td></tr>
<tr><td>Linear latency function</td><td>Total Latency</td><td colspan="2">2.618[AAE05]</td></tr>
<tr><td>Linear latency function (Unweighted)</td><td>Total Latency</td><td>2.5[AAE05]</td><td>2.618[AAE05]</td></tr>
<tr><td>Polynomial latency function of degree d</td><td>Total Latency</td><td>$\Omega(d^{d/2})$[AAE05]</td><td>$O(2^d d^{d+1})$[AAE05]</td></tr>
</table>

Table 1. Price of Anarchy for various cases.

of anarchy in this case. Later on, Czumaj et al. [CV02] proved that the price of anarchy is tightly $\theta(\log m/\log\log m)$ for m parallel links and m agents.

A more realistic model is when links have different speeds, i.e. the latency corresponding to an edge e is computed as $f_e(l_e)/s_e$, where s_e is the speed of link e. In this case Koutsoupias and Papadimitriou [KP99] proved that the price of anarchy increases in the case of 2 parallel links to the golden ratio $\varphi = \frac{1+\sqrt{5}}{2} \simeq 1.6111$. Notice that in all previously mentioned results the social utility is the **maximum** latency over all edges, i.e. $max_e(l_e)$, where l_e is the load on edge e.

Recently Awerbuch et al. [AAE05] considered general networks with polynomial latency functions and considered the **total** latency as the social utility function, i.e.

$$\sum_e f_e(l_e) l_e \tag{1}$$

For linear latency functions, i.e. $f_e(x) = a_e x + b_e$, for non-negative a_e and b_e, they prove a tight bound $(3+\sqrt{5})/2 \simeq 2.618$. A lower bound example is depicted in Fig. 2. Assume there are four agents with demands (U, V, ϕ), (U, W, ϕ), $(V, W, 1)$, and $(W, V, 1)$, i.e. agent 1 wants to transfer ϕ units of flow from U to V and so on. The best way is that agents 1, 2, 3, and 4 use paths UV, UW, VW, and WV, respectively, which gives a total latency of $2\phi^2 + 2$. On the other hand, one Nash equilibrium happens when agents use paths UWV, UVW, VUW, and WUV, respectively, which gives a total latency $4\phi^2 + 4\phi + 2$. The ratio $(4\phi^2 + 4\phi + 2)/(2\phi^2 + 2)$ is maximized when $\phi = (1 + \sqrt{5})/2$ and equals $(3 + \sqrt{5})/2 \simeq 2.618$ for that value. Notice that agents have different amounts of flows. If we restrict agents to request equal flow amounts, we get a lower bound 2.5 on the network in Fig. 2 by setting $\phi = 1$.

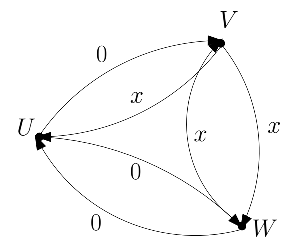

Fig. 2. A network with linear latency function and price of anarchy 2.618.

Awerbuch et al. [AAE05] also consider the more general case when f_e is a polynomial of degree d and prove an upper bound $O(2^d d^{d+1})$ and a lower bound $\Omega(d^{d/2})$ for the price of anarchy in this case. All results are shown in Table 1.

Alternative versions of selfish routing problem are also widely studied. Roughgarden et al. [CDR03a,CDR03b,Rou01] consider the splittable selfish routing with various latency functions and different social utility functions.

2.2 Finite Congestion Games

One natural generalization of unsplittable flow problem is finite congestion games. In this game agents are allowed to pick single edges instead of a whole path. For an introduction to congestion games the reader is referred to [MS96,Ros73]. A *congestion game* is a tuple $(N, E, (\Sigma_i)_{i \in N}, (f_e)_{e \in M})$ where N is the set of players, E is a set of facilities, $\Sigma_i \subseteq 2^E$ is a collection of actions for player i, i.e. an action is choosing a set of facilities, and f_e is a cost function associated with facility j. Assume $A = (A_1, A_2, \cdots, A_n)$ is a strategy profile. The cost of A to agent i is $c_i(A) = \sum_{e \in A_i} f_e(n_e(A))$ where $n_e(A)$ means the number of agents who use facility e. The social utility is either the sum of player's costs, i.e. $SUM(A) = \Sigma_{i \in N} c_i(A)$, or the maximum of players' costs, i.e. $MAX(A) = max_{i \in N} c_i(A)$. Finally, a game is *symmetric* if all Σ_i's are equal and is *asymmetric*, otherwise.

Christodoulou and Koutsoupias [CK05] recently studied the price of anarchy for finite congestion games for symmetric and asymmetric cases as well as linear and polynomial latency functions. Their results are listed in Table 2

2.3 Valid Utility Systems

One obvious requirement in order to have bounded price of anarchy is that the social utility and agent's individual utilities must be of the same type. For example, if agents try to maximize the number of apples but society wants the maximum number of oranges then there is no hope of getting bounds on the

	Linear Latencies		Polynomial Latencies	
	SUM	MAX	SUM	MAX
Symmetric	5/2	5/2	$d^{\theta(d)}$	$d^{\theta(d)}$
Asymmetric	5/2	$\theta(\sqrt{n})$	$d^{\theta(d)}$	$\Omega(n^{d/(d+1)})$, $O(n)$

Table 2. The price of anarchy on finite ingestion games.

price of anarchy. Assuming the above trivial requirement, Adrian Vetta [Vet02] in a seminal work defines a broad class of games called *valid utility systems* and proves that their price of anarchy is at most 2.0, i.e. any Nash equilibrium yields a social utility at least half as much as the optimal social utility. Lets define utility systems and then state his results.

In valid utility games, for each agent i, there exist a ground set V_i. An action of agent i corresponds to choosing a subset of V_i. In particular, we restrict agents to choose any subset by defining a feasible subset $F_i \subseteq 2^{V_i}$ from which agent i chooses his actions. Let $V = \cup_i V_i$. Both agent's utilities function, α_j's, and social utility function, γ, are defined as function from 2^V to R^+.

A game is called *valid utility system* if it satisfies the following three basic conditions.

1. Cake Condition: The social utility and private utility functions must be of the same type; moreover, the sum of private utilities over all agents by playing any strategy S does not exceed the social utility over that strategy, i.e.

$$\sum_i \alpha_i(S) \leq \gamma(S) \tag{2}$$

   where $\alpha_i(S)$ and $\gamma(S)$ are the expected utility of agent i and the social utility when players play according to the strategy S, respectively.
2. Submodularity: The social utility function is sub-modular, i.e. $f(X \cup D) - F(X) \geq f(Y \cup D) - F(Y)$ when $X \subseteq Y$.
3. Vickrey Condition: The private utility of an agent is not less than the change in the social utility if he refuses to participate.

The intuition behind the second condition is simple: If a town does not have any grocery store then establishing one has a lot more benefit than when the town has hundreds of grocery stores.

The third condition is similar to the *no single-agent effect* condition that we had in studying VCG mechanism. It says for any strategy S if agent i refuses to participate then the difference in the social utility is not bigger than the utility of agent i, i.e. $\alpha_i(S) \geq \gamma(S) - \gamma(S_{-i})$.

The three major results in [Vet02] are listed below. Let OPT be the value of the optimal social utility.

- For any valid utility system and any Nash equilibrium S, pure or mixed, the social utility obtained by playing according to S is at least half as much as

OPT minus some additive functional values. The exact inequality is

$$OPT \leq 2\gamma(S) - \sum_{t:s_i=\sigma_i} (\gamma(S) - \gamma(S_{-i})) - \sum_{t:s_i \neq \sigma_i} (\gamma(\Omega \cup S^i) - \gamma(\Omega \cup S^{i-1})) \quad (3)$$

where

- $\Omega = (\sigma_1, \sigma_2, \cdots, \sigma_n)$ is the strategy that gives the optimal utility.
- S^i is the strategy in which players $1, 2, \cdots, i$ play according to S and other players take the empty set $\emptyset$ as their action.
- $\Omega \cup S$ is defined in the following way. Let $S = (s_1, s_2, \cdots, s_n)$ and s_j be a mixed strategy in which agent i plays actions $a_j^1, a_j^2, \cdots, a_j^t$ with probabilities $p_j^1, p_j^2, \cdots, p_j^t$, respectively. Then $\Omega \cup S$ is a strategy in which agent j plays $a_j^k \cup \sigma_j$ with probaility p_j^k for $k = 1, 2, \cdots, t$. Notice that a_j^k is a subset of V_j and σ_j is the action of agent j in Ω.

– In case that the utility function γ is increasing, i.e. $\gamma(X) \leq \gamma(Y)$ whenever $X \subseteq Y$, then the additive terms in the above equations can be omitted; thus, for any valid utility system and increasing utility function

$$OPT \leq 2\gamma(S) \quad (4)$$

– He also considers the existence of pure strategy Nash equilibria in the special case that a valid utility system is *basic*, i.e. the equality holds in equation 2. In this case pure strategy Nash equilibria always exist.

– He also proposes competitive versions of facility location and k-median problems and show that they are valid utility systems; hence their price of anarchy is at most 2.0.

A wide range of games fall in the category of valid utility systems. Examples are market sharing games[GLMT04], distributed caching games[FGMS05], and, as we saw earlier, facility location games, traffic routing, and auctions [Vet02].

3 Price of Sinking

There is one major assumption behind considering price of sinking for measuring the lack of coordination in a game. We assume that agents tend to paly according to Nash equilibria of the game. There are, however, several drawbacks in this assumption. It often happens in practice, for example in auctions, that agents are reluctant to adopt mixed strategies; instead they are willing to repeatedly play according to pure strategies even if there is no pure strategy Nash equilibrium in the game. By the way, Nash equilibria are stable points in a game rather than optimal points, so agents might not necessarily be looking for those points.

By assuming that agents play repeatedly according to some pure strategies, Goemans et al. [GMV05] introduce a new measure for the lack of coordination, *price of sinking*. Lets denote each pure strategy profile $S = (a_1, a_2, \cdots, a_n)$ by a node in a directed graph and connect node A to node

$$A \oplus a_i' = (a_1, a_2, \cdots, a_{i-1}, a_i', a_{i+1}, \cdots, a_n)$$

only if a'_i is agent i's best response to S_{-i}. We denote this digraph by $\mathcal{D}$. Agents' repeated moves correspond to walks across the above described digraph. What happens ultimately? Agents may end up in a node u without any outgoing edge. Such a node u obviously corresponds to a pure strategy Nash equilibrium in the game. This happens only when $\mathcal{D}$ is acyclic. What if it is not?

A strongly connected component in a digraph is a maximal set of vertices that are mutually connected, i.e. there is a path between any two of them. It is well known that every digraph can be partitioned into strongly connected components. If we replace each strongly connected component x by a single vertex v_x and connect v_x to v_y if there is an edge from a vertex in x to a vertex in y then we obtain an acyclic digraph, say $B(\mathcal{D})$. This is often referred as the *block digraph* in Graph theory literature. A digraph with three strongly connected components is depicted in Fig. 3. As you see it has no pure strategy Nash equilibrium.

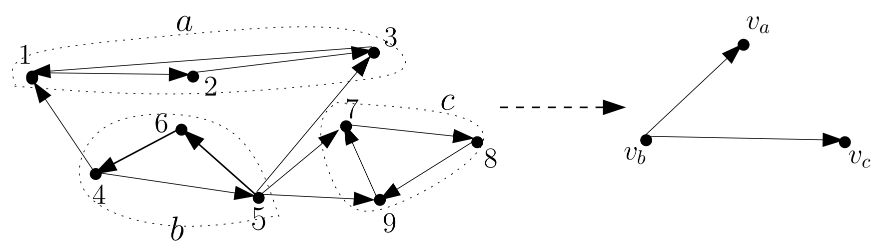

Fig. 3. A digraph with its block digraph.

The block digraph is always acyclic and, hence, has some vertices with out-degree zero. (v_a and v_c in Fig. 3) Such components are called *sink equilibria*; the reason for this terminology is simple: once we reach a sink equilibrium we never leave it. The price of sinking is defined as the worst case ratio of the social utility and any sink equilibrium. How do we compute the social utility over a sink equilibrium? Any infinitely long random walk over a strongly connected component Q reaches each vertex u with some fixed probability $\pi_Q(u)$. π_Q is is known as the *stationary distribution*. The social utility over a sink equilibrium is its expected value, with respect to the stationary distribution, of the social value on its states. In Fig. 3 all stationary distributions are uniform. Formally speaking, The social utility over a sink equilibrium Q is defined as

$$\gamma(Q) = \sum_{S \in Q} \pi_Q(S)\gamma(S) \tag{5}$$

where γ is the social utility function.

Gomeans et al. study the price of sinking over two category of games: unsplittable selfish routing and valid utility systems.

- For unsplittable selfish routing problem (as well as congestion games) they prove that the price of sinking is at most $O(2^{2d}d^{2d+3})$, for any latency func-

tion of degree d. Compare this with the $O(2^d d^{d+3})$ upper bound [AAE05] for the price of anarchy.
- As for valid utility games, they prove that price of sinking always lies between n and $n+1$.

They also prove some hardness results regarding the problem of computing price of sinking in general.

4 Price of Anarchy vs. Price of Sinking

How does price of sinking compare with the price of anarchy? As we said earlier, price of anarchy is based on the assumption that agents adopt (pure or mixed) strategy Nash equilibrium. In contrast, price of sinking is based on assuming that agents only adopt pure strategies but do repeated moves. Based on the latter assumption, Goemans et al. propose a valid utility game in which every possible outcome of the game, by repeatedly playing best-responses, is less than the optimal social utility by a factor of n. However the price of anarchy in such games is at most 2.0 according to [Vet02]. Consequently, the price of anarchy is unrealistic in this case.

References

[AAE05] Baruch Awerbuch, Yossi Azar, and Amir Epstein. The price of routing unsplittable flow. In *STOC '05: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 57–66, New York, NY, USA, 2005. ACM Press.

[CDR03a] Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. How much can taxes help selfish routing? In *ACM Conference on Electronic Commerce*, pages 98–107, 2003.

[CDR03b] Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. Pricing network edges for heterogeneous selfish users. In *STOC*, pages 521–530, 2003.

[CK05] George Christodoulou and Elias Koutsoupias. The price of anarchy of finite congestion games. In *STOC '05: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 67–73, New York, NY, USA, 2005. ACM Press.

[CV02] A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria, 2002.

[FGMS05] L. Fliescher, Michel X. Goemans, Vahab S. Mirrokni, and M. Sviridenko. Almost tight approximation algorithms for maximising general assignment problems. In *submitted*, 2005.

[GLMT04] Michel X. Goemans, Erran L. Li, Vahab S. Mirrokni, and Marina Thottan. Market sharing games applied to content distribution in ad-hoc networks. In *MobiHoc*, pages 55–66, 2004.

[GMV05] Michel Goemans, Vahab Mirrokni, and Adrian Vetta. Sink equilibria and convergence. In *FOCS '05: Proceedings of the 46th Symposium on Foundations of Computer Science*, Washington, DC, USA, 2005. IEEE Computer Society.

[KP99] E. Koutsoupias and CH. Papadimitriou. Worst-case equilibria. In *STACS '99: 16th International Symposium on Theoretical Aspects of Computer Science*, pages 404–413. LNCS, 1999.

[MS96] D. Monderer and L. S. Shapley. Potential games. In *Games and Economic Behavior*, pages 124–143, 1996.

[Ros73] Robert W. Rosenthal. A class of games possessing pure-strategy nash equilibria. In *International Journal of Game Theory*, pages 65–67, 1973.

[Rou01] Tim Roughgarden. Designing networks for selfish users is hard. In *FOCS*, pages 472–481, 2001.

[Vet02] Adrian Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *FOCS '02: Proceedings of the 43rd Symposium on Foundations of Computer Science*, page 416, Washington, DC, USA, 2002. IEEE Computer Society.