Nash Equilibria in Competitive Societies, with Applications to Facility Location, Traffic Routing and Auctions

Adrian Vetta*
McGill University
avetta@math.mit.edu

Abstract

We consider the following class of problems. The value of an outcome to a society is measured via a submodular utility function (submodularity has a natural economic interpretation: decreasing marginal utility). Decisions, however are controlled by non-cooperative agents who seek to maximise their own private utility. We present, under some basic assumptions, guarantees on the social performance of Nash equilibria. For submodular utility functions, any Nash equilibrium gives an expected social utility within a factor 2 *of optimal, subject to a function-dependent additive term. For non-decreasing, submodular utility functions, any Nash equilibrium gives an expected social utility within a factor* $1 + \delta$ *of optimal, where* $0 \leq \delta \leq 1$ *is a number based upon the discrete curvature of the function. A condition under which all sets of social and private utility functions induce pure strategy Nash equilibria is presented. The case in which agents, themselves, make use of approximation algorithms in decision making is discussed and performance guarantees given. Finally we present some specific problems that fall into our framework. These include the competitive versions of the facility location problem and* k*-median problem, a maximisation version of the traffic routing problem studied by Roughgarden and Tardos [9], and multiple-item auctions.*

1 Introduction

Computer scientists have long studied the costs incurred by the lack of *complete information* or the lack of *unbounded computational resources*. For example, the fields of on-line algorithms and approximation algorithms were developed in response to these two problems. However, these fields presume a single authority (or agent) whose goal is to optimise some objective function. What happens when there is a clear social objective function but no single authority? In particular, what if there are many agents whose goals are to optimise their own private objective functions, rather than to collectively optimise the social objective function? Motivated by examples of this type concerning the internet, Koutsoupias and Papadimitriou [5] proposed applying game-theoretic techniques in order to analyse the costs resulting from a lack of *coordination*. Specifically, they proposed the study of non-cooperative games via the use of Nash equilibria (where the agents' strategies are mutual best responses to each other). Given the non-cooperative nature of these games and the fact that such games may have many Nash equilibria, they proposed studying such equilibria from a worst case perspective. That is, how bad can a Nash equilibrium be, with respect to the social objective, in comparison to the best cooperative solution (or solution produced in presence of a single authority). The study of Nash equilibria is especially fruitful for problems in which the actions of the agents may be changed quickly and at little cost. This is because it is in such circumstances that Nash equilibria are most likely to arise in practice. Such problems abound in the high-tech economy. From a theoretical viewpoint, notable amongst them is the traffic routing problem which has been studied with great success by Roughgarden and Tardos [9].

In this paper we consider a large class of problems with the following structure. Decisions are made by a set of non-cooperative agents whose action spaces are subsets of an underlying groundset. The actions of the agents induce some social utility, measured by a set function. The goal of the agents, though, is not to maximise the overall social utility; rather, they seek to maximise their own private utility functions. The only assumptions we make are

• The social utility and private utility functions are measured in the same standard unit: this standard utility unit may be money, gold, cake etc. Clearly, such a condition is necessary. For example, no guarantees can be obtained if the value to society is measured in terms of the number of oranges but the agents seek to maximise the number of

*This work was carried out with the support of a Li fellowship whilst the author was at the Massachusetts Institute of Technology.

Proceedings of the 43 rd Annual IEEE Symposium on Foundations of Computer Science (FOCS'02)

apples.

• The social utility function is submodular: submodularity corresponds to a property that arises frequently in economics: *decreasing marginal utility*. Here, the additional value accruing from an action decreases as the overall level of activity in the society rises. For example, the additional benefit to a town of an extra taxi company is greater if there are currently no taxi firms in the town rather than if there are already one hundred taxi firms.

• The private utility of an agent is at least the change in social utility that would occur if the agent declined to participate in the game: we remark that, equivalently, we require that the private utility of an agent is at least the *Vickery utility* with respect to that agent. This concept is often considered in the study of auction mechanisms (see Vickery-Clarke-Groves payment mechanisms). Moreover, this condition arises in other practical situations as we will see in our examples.

Problems for which these three assumptions hold are called *utility systems*. For a utility system, it is possible to provide some strong guarantees concerning the social utility provided by any Nash equilibrium (we will also show that good guarantees arise if we relax the third assumption). Specifically, for non-decreasing, submodular objective functions, any Nash equilibrium will give a solution with expected social utility within a factor $1 + \delta$ of the optimal solution, (where $0 \leq \delta \leq 1$ is a number based upon the "discrete curvature" of the function). Hence, any Nash equilibrium is always at least half as good as the optimal social solution. For submodular functions in general, the expected social utility of a Nash equilibrium is within a factor 2 of optimal, subject to a function-based additive term (which, as we will see in our examples, often has a clear economic interpretation). An alternative form of guarantee that has interesting interpretations in certain problems (for example, the traffic routing problem) is also given. These results are shown to be tight.

The other main result in the paper is to show that, given a simple condition, utility systems have the desirable property that they possess pure strategy Nash equilibria. We also discuss and provide performance guarantees for instances in which the agents apply approximation algorithms in determining their strategies.

An outline of the paper is as follows. In Section 2 we introduce the necessary background on game theory and submodular functions. In Section 3 we prove our results concerning the social performance of Nash equilibria. In Section 4 we discuss pure strategy Nash equilibria and mixed strategy Nash equilibria. We then present the simple condition under which a utility system will have pure strategy Nash equilibria. In Section 5 we relax our third assumption and present results for the situation in which the private utility of an agent is comparable to the Vickery utility with respect to that agent (loss in social utility that would result from the agent dropping out of the game). Since our three assumptions concerning the utility system are not very restrictive, the results are widely applicable. We illustrate this by presenting a range of problems that fit into our framework. Our first examples are competitive versions of the facility location problem and the k-median problem, which we introduce in Section 6. One implication of the results in this section is that competitive markets are less efficient in industries with high fixed costs and high marginal profits. Practical examples of such social inefficiencies include the duplication of work, as well as the over-supply of lucrative markets (and under-supply of less valuable markets) by firms. Our next example, given in Section 7, concerns traffic routing in networks. In Section 8, we consider the issues of polynomial time implementations. These issues include the time it takes to obtain Nash equilibria and also the consequences of agents using approximation algorithms for strategy determination. One example in which speed considerations are of great importance is auctions. Thus, our last example, given in Section 9, is that of multiple-item auctions. We present a simple polynomial time auction that fits into our overall framework. It follows that the allocation of items given by the auction (in the presence of competing agents who bid in a greedy manner) is at least half as efficient as the optimal allocation given by a single authority. This matches the performance guarantee Lehmann, Lehmann and Nisan [6] gave for the problem where a single authority chooses an allocation.

2 Background

2.1 Some game theory.

Suppose we have k agents and disjoint groundsets $V_1, \ldots, V_k$. Each element in V_i represents an act that agent i may make, $1 \leq i \leq k$; let $a_i \subseteq V_i$ be an *action* (set of acts) available to agent i. We may wish to restrict the set of actions an agent may make; thus we may not allow every subset of V_i to be a feasible action. Towards this end, we let $\mathcal{A}_i = \{a_i \subseteq V_i : a_i \text{ is a feasible action}\} = \{a_i^1, a_i^2, \ldots, a_i^{n_i}\}$ be the set of all actions available to agent i. We call $\mathcal{A}_i$ the *action space* for agent i. A *pure strategy* is one in which the agent decides to carry out a specific action. A *mixed strategy* is one in which the agent decides upon an action according to some probability distribution. The *strategy space* $\mathcal{S}_i$ of agent i is the set of mixed strategies. So we represent $\mathcal{S}_i$ as $\mathcal{S}_i = \{s_i \in \mathbb{R}^{n_i} : \sum_{j=1}^{n_i} s_i^j = 1, s_i^j \geq 0\}$. Thus $s_i \in \mathcal{S}_i$ corresponds to the mixed strategy in which action a_i^1 is implemented with probability s_i^1, action a_i^2 is implemented with probability s_i^2, etc. Hence, a pure strategy corresponds to $(0, 1)$-vector in $\mathcal{S}_i$. Now let $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2 \times \cdots \times \mathcal{A}_k$ and let $\mathcal{S} = \mathcal{S}_1 \times \mathcal{S}_2 \times \cdots \times \mathcal{S}_k$.

Proceedings of the 43 rd Annual IEEE Symposium on Foundations of Computer Science (FOCS'02)
0272-5428/02 $17.00 © 2002 IEEE

In addition, we let $V = V_1 \cup V_2 \cup \cdots \cup V_k$. Then for a function $f : 2^V \to \mathbb{R}$, we define $\bar{f} : \mathcal{S} \to \mathbb{R}$ as follows $\bar{f}(S) = \sum_{A \in \mathcal{A}} f(A) \Pr(A|S)$. where $\Pr(A|S)$ is the probability that action set $A = \{a_1, a_2, \ldots, a_k\}$ is implemented given that the agents are using the strategy set $S = \{s_1, s_2, \ldots, s_k\}$. Thus $\bar{f}(S)$ is just the expected value of f on the strategy set S.

Given an action set $A = \{a_1, a_2, \ldots, a_k\} \in \mathcal{A}$, let $A \oplus a'_i$ denote the action set obtained if agent i changes its action from a_i to a'_i. Formally, $A \oplus a'_i = \{a_1, \ldots, a_{i-1}, a'_i, a_{i+1}, \ldots, a_k\}$. Similarly, given a strategy set $S = \{s_1, s_2, \ldots, s_k\} \in \mathcal{S}$, let $S \oplus s'_i = \{s_1, \ldots, s_{i-1}, s'_i, s_{i+1}, \ldots, s_k\}$, i.e. the strategy set obtained if agent i changes its strategy from s_i to s'_i. In this paper we will denote by $\gamma : 2^V \to \mathbb{R}$ the social utility function. In addition, for each agent $1 \le i \le k$, there is a private utility function $\alpha_i : 2^V \to \mathbb{R}$. The goal of each agent is, therefore, to select a strategy in order to maximise its private utility. Clearly, though, such strategies may not produce a good solution with respect to social utility γ. We say that set of strategies $S \in \mathcal{S}$ is a *Nash equilibrium* if no agent has an incentive to change strategy. That is, for any agent i,

$$\bar{\alpha}_i(S) \ge \bar{\alpha}_i(S \oplus s'_i) \qquad \forall s'_i \in \mathcal{S}_i$$

Equivalently, given the other agents strategies, s_i is the best response of agent i. We say that a Nash equilibrium $\{s_1, s_2, \ldots, s_k\}$ is a *pure strategy Nash equilibrium* if, for each agent i, s_i is a pure strategy. Otherwise we say that the Nash equilibrium is a *mixed strategy Nash equilibrium*. The following result, due to Nash [7], shows that the task of comparing the performance of Nash equilibria against a socially optimal solution is feasible.

Theorem 1. *Any finite, k-person, non-cooperative game has at least one Nash equilibrium.* □

2.2 Submodular functions.

A function with the form $f : 2^V \to \mathbb{R}$ is called a *set function*. A set function f is *submodular* if $f(X) + f(Y) \ge f(X \cap Y) + f(X \cup Y)$, $\forall X, Y \subseteq V$. It is *supermodular* if this inequality is reversed. A set function f is *non-decreasing* if $f(X) \le f(Y)$, $\forall X \subseteq Y \subseteq V$. For a set function f, the *discrete derivative* at $X \subseteq V$ in the direction $D \subseteq V - X$ is defined as $f'_D(X) = f(X \cup D) - f(X)$. The following result is standard. Condition (III) shows that, in economic terms, submodularity corresponds to the property of decreasing marginal utility, that is, the additional value accruing from an action decreases as the overall level of activity in the society rises.

Lemma 1. *The following are equivalent: (I) f is submodular. (II) $A \subseteq B$ implies $f'_D(A) \ge f'_D(B)$, $\forall D \subseteq V - B$. (III) $A \subseteq B$ implies $f'_i(A) \ge f'_i(B)$, $\forall i \in V - B$.* □

2.3 Utility systems.

Given our competitive game, let the optimal social solution be $\Omega = \{\sigma_1, \sigma_2, \ldots, \sigma_k\}$, with optimal value $\textsc{opt} = \gamma(\Omega)$. Here we consider the private utilities of the agents in a solution $S = \{s_1, s_2, \ldots, s_k\}$. First, we introduce some more notation. We denote by $\emptyset_i$, the *null strategy (action)* for agent i; such a strategy corresponds to agent i declining to take part in the game. We denote by $\emptyset = \{\emptyset_1, \emptyset_2, \ldots, \emptyset_k\}$ the strategy set in which each player has a null strategy. Now take an arbitrary ordering of the agents. Without loss of generality, we may assume that the ordering is $\{1, 2, \ldots, k\}$. Now given $A \in \mathcal{A}$ we set $A^i = \{a_1, a_2, \ldots, a_i, \emptyset_{i+1}, \ldots, \emptyset_k\}$. Similarly given $S \in \mathcal{S}$ we set $S^i = \{s_1, s_2, \ldots, s_i, \emptyset_{i+1}, \ldots, \emptyset_k\}$. Then, by construction, it follows that

Lemma 2. *For an action set $A \in \mathcal{A}$ and set function γ, we have $\gamma(A) = \sum_{i=1}^k \gamma'_{a_i}(A^{i-1})$.* □

Corollary 1. *For any strategy set $S \in \mathcal{S}$ and set function γ, we have $\bar{\gamma}(S) = \sum_{i=1}^k \bar{\gamma}'_{s_i}(S^{i-1})$.* □

Now take our submodular, social utility function $\gamma : 2^V \to \mathbb{R}$ (we remark that $\bar{\gamma}$ is also submodular) and our collection of private utility functions $\alpha_i : 2^V \to \mathbb{R}$, $1 \le i \le k$. Recall that our third assumption regarding the utility functions states that the private utility to an agent is at least as great as the loss in social utility resulting from the agent dropping out of the game. That is, the system $(\gamma, \cup_i \alpha_i)$ has the property

$$\bar{\alpha}_i(S) \quad \ge \quad \bar{\gamma}'_{s_i}(S \oplus \emptyset_i) \tag{1}$$

Given condition (1), we say that the system $(\gamma, \cup_i \alpha_i)$ is a *utility system*. The utility system $(\gamma, \cup_i \alpha_i)$ is said to be *basic* if we have equality in condition (1), that is $\bar{\alpha}_i(S) = \bar{\gamma}'_{s_i}(S \oplus \emptyset_i)$. Observe that, since we are assuming that utilities are measured in the same units, we may view the game in the following manner. The function γ represents the value of the game (or size of the cake), and α_i represents the return to the agent i (i.e. the size of agent i's piece of the cake). Therefore we also require that the sum of the sizes of the pieces must be smaller than the total size of the cake. That is we require that the sum of the private utilities of the agents is at most the social utility

$$\sum_i \bar{\alpha}_i(S) \quad \le \quad \bar{\gamma}(S) \tag{2}$$

In such a circumstance we say that the utility system $(\gamma, \cup_i \alpha_i)$ is valid. Note, we do not require that $\sum_i \bar{\alpha}_i(S) = \bar{\gamma}(S)$. In fact, as we shall see the value $\bar{\gamma}(S) - \sum_i \bar{\alpha}_i(S)$ often has a clear meaning. For the moment we may view $\bar{\gamma}(S) - \sum_i \bar{\alpha}_i(S)$ as the utility of some non-agent, say the utility of the general public. Observe that conditions (1) and (2) must hold if the following two conditions hold

$$\alpha_i(A) \geq \gamma'_{a_i}(A \oplus \emptyset_i) \quad (3)$$

$$\sum_i \alpha_i(A) \leq \gamma(A) \quad (4)$$

We now show that valid utility systems do exist.

Theorem 2. *For any submodular function γ, there exist functions α_i, $1 \leq i \leq k$ such that $(\gamma, \cup_i \alpha_i)$ is a valid utility system. In particular, the basic utility system is valid.*

Proof. So we need to show that, for the basic utility system, Condition (2) holds. Now (by by Lemma 2, by Lemma 1 and the basic-ness of $(\gamma, \cup_i \alpha_i)$) we have

$$\begin{aligned}\bar{\gamma}(S) &= \sum_{i=1}^{k} \bar{\gamma}'_{s_i}(S^{i-1}) = \sum_{i=1}^{k} \bar{\gamma}'_{s_i}(S^i \oplus \emptyset_i) \\ &\geq \sum_{i=1}^{k} \bar{\gamma}'_{s_i}(S \oplus \emptyset_i) = \sum_{i=1}^{k} \bar{\alpha}_i(S) \quad \square\end{aligned}$$

3 Main Results

In this section we present our guarantees concerning the social value of a Nash equilibrium. In particular, for a valid utility system with a non-decreasing, submodular, social utility function we will show that any Nash equilibrium has an expected social value of at least half that of an optimal social solution. In fact, following an approach of Conforti and Cornuéjols [1], we obtain a tighter bound (although it provides the same guarantee in the worst case) with respect to a parameter based upon the discrete curvature of the non-decreasing, submodular function. For a valid utility system with a submodular, social utility function it is not possible to obtain a simple multiplicative guarantee. However, the expected social value of the Nash equilibrium is at least half the social optimal subject to an additive term. This additive term is function-dependent and often has a clean social/economic interpretation; for example, we will see in the Section 6 that, for the competitive facility location problem, it is bounded by the fixed investment costs. We begin with the following result, concerning any strategy set $S \in \mathcal{S}$.

Lemma 3. *Let γ be a submodular set function. Then, for any $S \in \mathcal{S}$,*

$$\bar{\gamma}(\Omega) \leq \bar{\gamma}(S) + \sum_{i:\sigma_i \in \Omega - S} \bar{\gamma}'_{\sigma_i}(S \oplus \emptyset_i) - \sum_{i:s_i \in S-\Omega} \bar{\gamma}'_{s_i}(\Omega \cup S^{i-1})$$

Proof. Observe that, by Lemma 1, $\bar{\gamma}(\Omega \cup S) \leq \bar{\gamma}(S) + \sum_{i:\sigma_i \in \Omega - S} \bar{\gamma}'_{\sigma_i}(S) \leq \bar{\gamma}(S) + \sum_{i:\sigma_i \in \Omega - S} \bar{\gamma}'_{\sigma_i}(S \oplus \emptyset_i)$. The lemma then follows from the following observation that $\bar{\gamma}(\Omega \cup S) = \bar{\gamma}(\Omega) + \sum_{i:s_i \in S-\Omega} \bar{\gamma}'_{s_i}(\Omega \cup S^{i-1})$. $\square$ Now let us focus specifically on the case of Nash equilibria. We then obtain the following guarantee concerning the social value of a Nash equilibrium.

Theorem 3. *Let γ be a submodular set function. If $(\gamma, \cup_i \alpha_i)$ is a valid utility system then for any Nash equilibrium $S \in \mathcal{S}$ we have* OPT $\leq 2\,\bar{\gamma}(S) - \sum_i \bar{\gamma}'_{s_i}(S \cup \Omega - s_i)$.

Proof. First note that Ω is a strategy set consisting of pure strategies. Therefore OPT $= \gamma(\Omega) = \bar{\gamma}(\Omega)$. So we have

$$\begin{aligned}\text{OPT} &\leq \bar{\gamma}(S) + \sum_{i:\sigma_i \in \Omega - S} \bar{\gamma}'_{\sigma_i}(S \oplus \emptyset_i) - \sum_{i:s_i \in S-\Omega} \bar{\gamma}'_{s_i}(\Omega \cup S^{i-1}) \\ &\leq \bar{\gamma}(S) + \sum_{i:\sigma_i \in \Omega - S} \bar{\alpha}_i(S) - \sum_{i:s_i \in S-\Omega} \bar{\gamma}'_{s_i}(\Omega \cup S^{i-1}) \\ &= \bar{\gamma}(S) + \sum_{i:s_i \in S-\Omega} \bar{\alpha}_i(S) - \sum_{i:s_i \in S-\Omega} \bar{\gamma}'_{s_i}(\Omega \cup S^{i-1}) \\ &\leq \bar{\gamma}(S) + \left(\bar{\gamma}(S) - \sum_{i:s_i \in S\cap\Omega} \bar{\gamma}'_{s_i}(S \oplus \emptyset_i)\right) \\ &\quad - \sum_{i:s_i \in S-\Omega} \bar{\gamma}'_{s_i}(\Omega \cup S^{i-1}) \\ &\leq 2\bar{\gamma}(S) - \sum_{i:s_i \in S\cap\Omega} \bar{\gamma}'_{s_i}(S \cup \Omega - s_i) \\ &\quad - \sum_{i:s_i \in S-\Omega} \bar{\gamma}'_{s_i}(S \cup \Omega - s_i) \\ &= 2\bar{\gamma}(S) - \sum_i \bar{\gamma}'_{s_i}(S \cup \Omega - s_i) \quad \square\end{aligned}$$

Observe that, for a general submodular function γ, the term $\sum_i \bar{\gamma}'_{s_i}(S \cup \Omega - s_i)$ may be negative. Thus, the social value of the Nash equilibrium is at least half the social optimal subject to a function-dependent additive term. As mentioned, this additive term often has a economic/social meaning. An alternative type of guarantee is also available. This result has clean implications in certain problems, for example, the traffic routing problem of Section 7 is also available.

Theorem 4. *Let γ be a submodular set function. If $(\gamma, \cup_i \alpha_i)$ is a valid utility system then for any Nash equilibrium $S \in \mathcal{S}$ we have $2\bar{\gamma}(S) \geq \bar{\gamma}(\Omega \cup S) + \sum_{i:\sigma_i \in \Omega\cap S} \bar{\gamma}'_{\sigma_i}(S \oplus \emptyset_i)$.*

Proof.

$$\begin{aligned}2\bar{\gamma}(S) &\geq \gamma(\Omega) + \sum_{i:s_i \in S-\Omega} \bar{\gamma}'_{s_i}(\Omega \cup S^{i-1}) + \sum_{i:s_i \in \Omega\cap S} \bar{\gamma}'_{s_i}(S \oplus \emptyset_i) \\ &= \bar{\gamma}(\Omega \cup S) + \sum_{i:s_i \in \Omega\cap S} \bar{\gamma}'_{s_i}(S \oplus \emptyset_i) \\ &= \bar{\gamma}(\Omega \cup S) + \sum_{i:\sigma_i \in \Omega\cap S} \bar{\gamma}'_{\sigma_i}(S \oplus \emptyset_i) \quad \square\end{aligned}$$

For non-decreasing, submodular functions the additive term in Theorem 3 is positive and, hence, we obtain a factor 2 guarantee. We can, however, do better in this specific case. We define the *discrete curvature* of a non-decreasing, submodular set function f to be

$$\delta(f) = \max_{i:T\subseteq V, f'_T(\emptyset)>0} \left(\frac{f'_T(\emptyset) - f'_T(V-T)}{f'_T(\emptyset)} \right) = \max_{i:v\in V, f'_v(\emptyset)>0} \left(\frac{f'_v(\emptyset) - f'_v(V-v)}{f'_v(\emptyset)} \right)$$

Theorem 5. *Let γ be a non-decreasing, submodular set function. If $(\gamma, \cup_i \alpha_i)$ is a valid utility system then for any Nash equilibrium $S \in \mathcal{S}$ we have* $\text{OPT} \le (1 + \delta(\gamma))\,\bar{\gamma}(S) \le 2\,\bar{\gamma}(S)$.

Proof. Now $\text{OPT} \le 2\bar{\gamma}(S) - \sum_i \bar{\gamma}'_{s_i}(S \cup \Omega - s_i)) \le 2\bar{\gamma}(S) - (1 - \delta_0(\bar{\gamma}))\sum_i \bar{\gamma}'_{s_i}(S^{i-1})$, where $\delta_0(\bar{\gamma}) = \max_{i:\bar{\gamma}'_{s_i}(S^{i-1})>0} \left(\frac{\bar{\gamma}'_{s_i}(S^{i-1}) - \bar{\gamma}'_{s_i}(\Omega\cup S^{i-1})}{\bar{\gamma}'_{s_i}(S^{i-1})} \right)$. Observe that, since γ is submodular, we have $0 \le \delta_0(\bar{\gamma}) \le \delta(\gamma) \le 1$. Thus, we obtain $\text{OPT} \le (1 + \delta_0(\bar{\gamma}))\,\bar{\gamma}(S) \le (1 + \delta(\gamma))\,\bar{\gamma}(S)$. □

Thus, for a non-decreasing, submodular set function γ any Nash equilibrium provides a social utility within a $1 + \delta(\gamma)$ factor of the optimal social utility. Theorems 3 and 5 are both tight. We give examples to show this in the full paper.

4 Pure Strategy Nash Equilibria

Recall Theorem 1 which states that finite, non-cooperative, k-agent games have a Nash equilibrium. Unfortunately this is just an existence result and offers no help in actually finding Nash equilibria. In addition, the result just guarantees the existence of a mixed strategy Nash equilibrium. It is not the case that there need be pure strategy Nash equilibria; in fact, generally complex games will not have a pure strategy Nash equilibria. The existence of pure strategy Nash equilibria is of interest for several reasons. In many practical situations, e.g. decisions concerning the location of facilities, agents are likely to adopt pure strategies. They are unlikely to chose one action amongst many on the basis of a coin toss. Furthermore, the strategy space of pure strategies is much smaller than the strategy space of mixed strategies. Thus, the discovery of pure strategy Nash equilibria may become a feasible. Moreover, given this smaller space, it is more reasonable to imagine that the agents can and will act in such a way as to generate a pure strategy Nash equilibria. In this section, we will show that any basic utility system has pure strategy Nash equilibria. We will also discuss how such equilibria may be realised in practice.

Theorem 6. *Take a valid utility system $(\gamma, \cup_i \alpha_i)$. If the utility system is basic then there are pure strategy Nash equilibria.*

Proof. Consider a directed graph D, each node of which corresponds to one of the possible pure strategy sets (i.e. action sets). There is an arc from node $\{a_1, a_2, \ldots, a_i, \ldots, a_k\}$ to node $\{a_1, a_2, \ldots, a'_i, \ldots, a_k\}$ if $\alpha_i(\{a_1, a_2, \ldots, a_i, \ldots, a_k\}) < \alpha_i(\{a_1, a_2, \ldots, a'_i, \ldots, a_k\})$, for some agent i. It follows that a node $\{a_1, a_2, \ldots, a_k\}$ in D corresponds to a pure strategy Nash equilibrium if and only if the node has out-degree zero. In particular, the system has a pure strategy Nash equilibrium if D is acyclic. We will show that for basic utility systems this is indeed the case. Suppose D is not acyclic. Then take a directed cycle C in D. Suppose the cycle contains nodes corresponding to the action sets $A_0 = \{a_1^0, \ldots, a_k^0\}, A_1 = \{a_1^1 \ldots, a_k^1\}, \ldots, A_t = \{a_1^t, \ldots, a_k^t\}$ where $A_0 = A_t$. It follows that the action sets A_r and A_{r+1} differ in only the action of one agent, say agent i_r. Thus $a_i^r = a_i^{r+1}$ if $i \ne i_r$, and $\alpha_{i_r}(A_r) < \alpha_{i_r}(A_{r+1})$, that is $\alpha_{i_r}(\{a_1^r, a_2^r, \ldots, a_k^r\}) < \alpha_{i_r}(\{a_1^{r+1} = a_1^r, \ldots, a_{i_{r-1}}^{r+1} = a_{i_{r-1}}^r, a_{i_r}^{r+1}, a_{i_{r+1}}^{r+1} = a_{i_{r+1}}^r, \ldots, a_k^{r+1} = a_k^r\})$. In particular, it must be the case that $\sum_{r=0}^{t-1} \alpha_{i_r}(A_{r+1}) - \alpha_{i_r}(A_r) > 0$. We will obtain a contradiction by showing that, in fact, $\sum_{r=0}^{t-1} \alpha_{i_r}(A_{r+1}) - \alpha_{i_r}(A_r) = 0$. Now $\alpha_{i_r}(A_{r+1}) = \gamma'_{a_{i_r}^{r+1}}(A_{r+1} \oplus \emptyset_{i_r})$ and $\alpha_{i_r}(A_r) = \gamma'_{a_{i_r}^r}(A_r \oplus \emptyset_{i_r})$. Thus $\alpha_{i_r}(A_{r+1}) - \alpha_{i_r}(A_r)$

$$\begin{aligned}
&= \gamma'_{a_{i_r}^{r+1}}(A_{r+1} \oplus \emptyset_{i_r}) - \gamma'_{a_{i_r}^r}(A_r \oplus \emptyset_{i_r}) \\
&= (\gamma(A_{r+1}) - \gamma(A_{r+1} \oplus \emptyset_{i_r})) - (\gamma(A_r) - \gamma(A_r \oplus \emptyset_{i_r})) \\
&= (\gamma(A_{r+1}) - \gamma(A_r)) + (\gamma(A_r \oplus \emptyset_{i_r}) - \gamma(A_{r+1} \oplus \emptyset_{i_r})) \\
&= \gamma(A_{r+1}) - \gamma(A_r)
\end{aligned}$$

Here the last equality follows from the observation that $a_i^r = a_i^{r+1}$ if $i \ne i_r$. Then, since $A_0 = A_t$, we obtain

$$\sum_{r=0}^{t-1} \alpha_{i_r}(A_{r+1}) - \alpha_{i_r}(A_r) = \sum_{r=0}^{t} \gamma(A_{r+1}) - \gamma(A_r) = \gamma(A_t) - \gamma(A_0) = 0 \quad \square$$

Observe that Theorem 6 states not only that a pure strategy Nash equilibrium exists, but the proof also shows how one may be obtained. Specifically, if we start with any pure strategy set S (for example, $S = \{\emptyset_1, \emptyset_2, \ldots \emptyset_k\}$) and the agents sequentially alter their actions in order to maximise their own profits then we will automatically converge to a pure strategy Nash equilibrium. In addition, this is true even if the agents do not chose an optimal response at each step, but rather just chose any action that leads to an improvement in their private utility. So suppose that agents can quickly adapt their actions. Then pure strategy Nash equilibria can be generated just by the agents acting in any greedy fashion. We note that for Theorem 6 we do require that the utility system be basic. For example, suppose we have a utility system $(\gamma, \cup_i \alpha_i)$ in which $\gamma(A) = M$, for some large constant M. Hence g is a constant function and is, therefore, submodular. Consequently, we have $\gamma'_{a_i}(A \oplus \emptyset_i) = 0$. It follows that, if the system is not basic, the only constraints on the private utility functions are that $\sum_i \alpha_i(A) \leq M$, $\forall A \in \mathcal{A}$ and that $\alpha_i(A) \geq 0$, $\forall i$. However, this presents no real restriction on the game, other than that the private payoffs must be non-negative. It is, therefore, easy to give examples with no pure strategy Nash equilibria.

5 A Broader Framework

In this section we relax our third assumption, that is $\bar{\alpha}_i(S) \geq \bar{\gamma}'_{s_i}(S \oplus \emptyset_i)$. Instead we will consider the situation in which the private utility of an agent is comparable to the Vickery utility with respect to that agent (loss in social utility that would result from the agent withdrawing from the game). We say that $(\gamma, \cup_i \alpha_i)$ is a *(P,Q)-utility system* if, for some constants $P, Q > 0$,

$$\bar{\alpha}_i(S) \quad \geq \quad \frac{1}{P}\,\bar{\gamma}'_{s_i}(S \oplus \emptyset_i) - Q \qquad (5)$$

A (P,Q)-utility system is *(P,Q)-basic* if we have equality in condition (5): $\bar{\alpha}_i(S) = \frac{1}{P}\bar{\gamma}'_{s_i}(S \oplus \emptyset_i) - Q$. The system is *valid* if $\sum_i \bar{\alpha}_i(S) \leq \bar{\gamma}(S)$. Then we easily obtain the following results.

Theorem 7. *Let γ be a submodular set function. If $(\gamma, \cup_i \alpha_i)$ is a valid (P,Q)-utility system then for any Nash equilibrium $S \in \mathcal{S}$ we have* $\text{OPT} \leq (1 + P)\bar{\gamma}(S) + (kQ - \sum_i \bar{\gamma}'_{s_i}(S \cup \Omega - s_i))$. □

Theorem 8. *Let γ be a non-decreasing, submodular set function. If $(\gamma, \cup_i \alpha_i)$ is a valid (P,Q)-utility system then for any Nash equilibrium $S \in \mathcal{S}$ we have* $\text{OPT} \leq (P + \delta(\gamma))\bar{\gamma}(S) + kQ$. □

Theorem 9. *Take a valid (P,Q)-utility system. If it is (P,Q)-basic then there are pure strategy Nash equilibria.* □

6 The Competitive Facility Location and k-Median Problems

In this section we consider the facility location and k-median problems. First we will describe the problems and then introduce competitive versions of the problems. We will then show that these competitive problems fit into the framework given in the previous sections.

6.1 The base problems.

Both these facility location problems have the following form. We are given a bipartite graph $G = (W \cup U, E)$ with vertex partition W and U. The set W consists of locations at which facilities may be built. The set U consists of locations at which consumers are found. For clarity, we will refer to vertices in W as locations and the vertices in U as markets. In the base problems we have a single agent or monopolistic firm. The monopolist wishes to construct facilities at various locations in W in order to maximise its profits. Each market u in U has an associated *value* π_u. A facility may be built at a location v for a *fixed cost* c_v. A facility at location v is able to service a market located at u for the *marginal cost* λ_{vu}. The *marginal profit* of the firm is its revenue minus its marginal costs. The *profit* of the firm is its marginal profit minus its fixed costs (i.e. revenue minus total costs). The *consumer surplus* is defined to be the total value minus total price. The *social surplus* is defined to be profits plus consumer surplus or, equivalently, total value minus total costs.

Let us examine these terms in more detail. Consider the revenue of the firm. This is just the sum of the prices it charges each market for servicing it. What will this price be, though, in the monopolistic case? Observe that consumers in market u have no choice but to be serviced by the monopolist. Their only constraint is that they will not pay more that π_u; thus, the firm will charge u a price $p_u = \pi_u$. It follows that consumer surplus is zero in the monopolist case. Thus a firm maximising profits is also, inadvertently, maximising the social surplus. Observe that the firm will refuse to service a market u from a facility v if $\lambda_{vu} > \pi_u$. Thus a firm can always obtain a marginal revenue of zero with respect to each market. Thus our objective function will not be affected if we assume that our bipartite graph is complete and we have $\lambda_{vu} \leq \pi_u$ for each edge vu (that is setting $\lambda_{vu} = \pi_u$ where $\lambda_{vu} > \pi_u$ will not affect the outcome). For the facility location problem, the firm may open whichever facilities it desires. So, formally, the *facility location problem* is

$$\max_{A \subseteq W} \mu(A) = \max_{A \subseteq W} \left(\sum_u \max_{v \in A}(\pi_u - \lambda_{vu}) - \sum_{v \in A} c_v \right)$$

In the k-median problem the firm faces an additional constraint in that it can open at most k facilities. Formally, the *k-median problem* is

$$\max_{A \subseteq W\, |A| \leq k} \mu(A) = \max_{A \subseteq W\, |A| \leq k} \left(\sum_u \max_{v \in A}(\pi_u - \lambda_{vu}) - \sum_{v \in A} c_v \right)$$

The performance of algorithms for these problems has been widely studied, (see, for example, [2] and [1]). Note, it is often assumed that for the k-median problem there are no fixed costs i.e. $c_v = 0, \forall v$. We also remark that, recently, the minimisation versions of both these problems have also received widespread attention (see, for example, [3]). The minimisation problems correspond to minimising the total costs of servicing all the markets. The broader economic viewpoint implied by the traditional maximisation problem, though, allows for very clean competitive formulations. It is these formulations that we will now introduce.

6.2 The competitive problems.

The base problems correspond to the monopolistic situation. The corresponding competitive problem is as follows. Instead of a single monopoly, suppose we have k competing firms (or agents). In the *competitive facility location problem* the number of facilities each firm may open is unrestricted; whereas in the *competitive k-median problem* each firm may build at most one facility (in fact, our results hold for a more general problem in which firm i can open at most m_i facilities). We allow firms to build at the same location, but assume, however, that the costs differ for each firm. Thus firm i, $1 \leq i \leq k$, may build a facility at location v for a fixed cost c_v^i. In addition, the marginal cost of firm i servicing a market u from a facility at location v is λ_{vu}^i. Again, the value of market u is π_u.

The competitive situation differs markedly from the monopolistic case. Consider, for example, the pricing strategies of firms in non-competitive and competitive markets. We have seen that in the monopolistic case there is no consumer surplus; the monopoly gets all of the social surplus for itself. In a competitive market, though, firms have to compete for the market u. Let $\lambda_u^1, \lambda_u^2, \ldots, \lambda_u^k$ be the lowest marginal costs with which the firms can supply market u, i.e. $\lambda_u^i = \text{in}(\ \lambda_{v,u}^i : \text{firm } i \text{ has an open facility at } v)$, and let $\lambda_u = \min_i \lambda_u^i$. What will happen in such a situation? Well, not surprisingly firm $i_u^* = \text{argmin}_i \lambda_u^i$ will compete most efficiently and will, thus, service market u. However, the firm will not be able to charge π_u; instead, it will only be able to charge the marginal cost of the second most efficient firm. Thus u will pay a price of $p_u = \min_{i \neq i_u^*} \lambda_u^i$ in order to be serviced. If the firm i_u^* tries to charge more than this it will be under-cut by another firm. Since the price p_u may be less than π_u, positive consumer surpluses may now arise. Hence, the social surplus is indeed shared between the individual firms and the consumers; market u contributes $\pi_u - p_u$ to the consumer surplus and $p_u - \lambda_u$ to the marginal profits of the firm that services it. (It may be the case that multiple firms all have the lowest marginal costs with respect to a market u. In such circumstances we will assume that customers in u randomly allocate their custom between these firms. The marginal profits for these firms will, though, be zero with respect to a market u, since they will compete away each others profits.)

Let $\Gamma_i = \{u : i = i_u^*\}$, $N_u = \{i : i = i_u^*\}$ and $n_u = |N_u|$. Then, given a set of actions $A = A_1 \times A_2 \times \cdots \times A_k$ we have that the profit of firm i is $\omega_i(A) = \sum_{u \in \Gamma_i}(p_u - \lambda_u^i) - \sum_{v \in a_i} c_v^i$. The consumer surplus is $\zeta(A) = \sum_i \sum_{u \in \Gamma_i}(\pi_u - p_u)$, and the social surplus is $\mu(A) = \sum_i \sum_{u \in \Gamma_i}(\pi_u - \lambda_u^i) - \sum_i \sum_{v \in a_i} c_v^i = \sum_u \sum_{i \in N_u} \frac{(\pi_u - \lambda_u^i)}{n_u} - \sum_i \sum_{v \in a_i} c_v^i$. So from a social viewpoint it would be best for a single authority to direct where each firm should locate in order to maximise the social surplus (utility). However, the firms themselves will choose strategies according to their own private profit (utility) functions. We next show, however, that these competitive formulations fit into the framework we have developed and, thus, we are able to obtain guarantees concerning the social performance on Nash equilibria in these facility location problems.

6.3 Social value of NE in facility location.

First we show that we can formulate both the competitive facility location problem and the competitive k-median problem appropriately for our purposes. Then we will show that our social utility (surplus) function is submodular.

Lemma 4. *The competitive facility location and k-median problems can be formulated in the action set framework.*

Proof. Consider first the competitive facility location problem. Recall that for our base problems we have a bipartite graph $G = (W \cup U, E)$. It follows that each agent i has a groundset $V_i = W$. Now, since a firm may open facilities at any set of locations, we have $\mathcal{A}_i = \{X : X \subseteq V_i\}$. Next consider the competitive k-median problem. Again, each agent i has a groundset $V_i = W$. Since each firm may open at most one facility we have $\mathcal{A}_i = \emptyset \cup \{v : v \in V_i\}$. □

Lemma 5. *The social surplus function μ is submodular.*

Proof. So $\mu(A) = \sum_i \sum_{u \in \Gamma_i}(\pi_u - \lambda_u^i) - \sum_i \sum_{v \in a_i} c_v^i = h(A) - g(A)$. Now, clearly, $g(A) + g(B) = g(A \cap B) + g(A \cup B)$, for $A, B \subseteq V = \cup_i V_i$. So it suffices to show that h is submodular i.e. $\sum_i \sum_{u \in \Gamma_i} \lambda_u^i$ is supermodular. In what follows, we add an action set descriptor to distinguish between the four types of action set (A, B, $A \cap B$ and $A \cup B$). Let $i \in N_u(A \cup B)$. Without loss of generality, assume that $\lambda_u(A \cup B) = \lambda_{v_i,u}^i$ where $v_i \in A$. Then $\lambda_u(A \cup B) = \lambda_u(A)$. Clearly, however, $\lambda_u(A \cap B) \geq \lambda_u(B)$. It follows that $h(A) + h(B) \geq h(A \cap B) + h(A \cup B)$. □

As mentioned, traditionally, the k-median problem is usually presented in the absence of fixed costs i.e. $c_v^i = 0$, $\forall i \, \forall v$. Such a formulation gives the following property.

Corollary 2. *In the absence of fixed costs, the social surplus function μ is non-decreasing.*

Proof. In the absence of fixed costs we have $g(A) = 0$, $\forall A \subseteq V$. Clearly h is a non-decreasing function, and hence μ is also non-decreasing. $\square$

Lemma 6. *The system $(\mu, \cup_i \omega_i)$ is a valid utility system. In particular, the utility system is basic.*

Proof. Recall that our private utility (profit) functions are $\omega_i(A) = \sum_{u \in \Gamma_i}(p_u - \lambda_u^i) - \sum_{v \in a_i} c_v^i$. We now show that $(\mu, \cup_i \omega_i)$ is a basic utility system, that is $\mu_i'(A) = \mu(A) - \mu(A \oplus a_i) = \sum_{u \in \Gamma_i}(p_u - \lambda_u^i) - \sum_{v \in a_i} c_v^i$. The change in the social utility is the increase in the total marginal profits minus the increase in the total fixed cost, when agent i changes its action from the null action to action a_i. The increase in total marginal profits, though, is just the sum over all markets of the extra efficiency gained by i having action a_i. This, in turn, is the difference between the marginal costs of i, in those markets where it is the most efficient firm, and the marginal costs of the next most efficient firm. This is just $\sum_{u \in \Gamma_i}(p_u - \lambda_u^i)$. Clearly the total change in the fixed costs is $\sum_{v \in a_i} c_v^i$, as required. Hence, the system $(\mu, \cup_i \omega_i)$ is a basic utility system. By Theorem 2, the utility system is also valid. $\square$

We are now in the position to apply Theorems 3 and 5. If we denote by $FC(S)$ and $MP(S)$ the expected fixed costs and expected marginal profits, respectively, associated with a solution $S \in \mathcal{S}$, then

Theorem 10. *For the competitive k-median and facility locations problems, any Nash equilibrium $S \in \mathcal{S}$ satisfies*

$$\text{OPT} \;\leq\; 2\,\bar{\mu}(S) + FC(S) \;=\; \bar{\mu}(S) + MP(S) + \bar{\zeta}(S)$$

Proof. By Theorem 3, $\bar{\mu}(\Omega) \leq 2\bar{\mu}(S) - \sum_i \bar{\mu}_{s_i}'(S \cup \Omega - s_i)$. Now $\sum_i \bar{\mu}_{s_i}'(S \cup \Omega - s_i)) = \sum_i \bar{\mu}(S \cup \Omega) - \bar{\mu}(S \cup \Omega - s_i) \geq \sum_i FC(S \cup \Omega - s_i) - FC(S \cup \Omega) = -FC(S)$. The result follows as $MP(S) - FC(S) + \bar{\zeta}(S) = \bar{\mu}(S)$. $\square$

This theorem tell us that our guarantee is good when either the fixed costs or the marginal profits plus consumer surplus induced by the solution S are small compared to OPT. Conversely, if the fixed costs and marginal profits plus consumer surplus are both large then the overall social performance may be very poor. Such a situation may arise in industries in which there are high start-up costs combined with markets that contain a collection of highly valuable customers. As a result, firms may over-supply the valuable customers (at the expense of less valuable customers) leading to a wasteful duplication of services. Such examples are common in the high-tech industry where the occurrence of high initial costs often allows a firm access to lucrative markets. In the absence of fixed costs we have

Theorem 11. *For the competitive k-median problem in the absence of fixed costs, any Nash equilibrium $S \in \mathcal{S}$ satisfies* OPT $\leq (1 + \delta(\mu))\bar{\mu}(S) \leq 2\bar{\mu}(S)$. $\square$

6.4 Pure strategy NE and facility location.

Observe that both facility location problems have the desirable property that they possess pure strategy Nash equilibria. This follows from Theorem 6 and Lemma 6.

Theorem 12. *The competitive facility location and k-median problems have pure strategy Nash equilibria.* $\square$

7 The Selfish Routing Problem

In this section we consider the problem of routing traffic in a network. Congestion in the network causes delays and is costly for individual agents and society as a whole. It would help, therefore, if the traffic could be directed by a single authority. However, it is individual agents who make their own routing decisions. Thus the problem appears suitable for analysis via our techniques. In particular, here we sketch how a maximisation version of the selfish routing problem of Roughgarden and Tardos [9] fits into our framework. They considered the following network routing problem. There is a directed network $G = (V, A)$ and k source-destination vertex pairs, $\{s_1, t_1\}, \ldots, \{s_k, t_k\}$ (note that we do not require k to be large). The collection of paths from s_i to t_i is denoted by $\mathcal{P}_i$ with $\mathcal{P} = \cup_i \mathcal{P}_i$. A flow is a function $f : \mathcal{P} \rightarrow \mathbb{R}^+$; for a fixed flow f, we have $f_a = \sum_{P \in \mathcal{P}: a \in P} f_P$. Now $f = \cup_i f_i$ where f_i is a flow from s_i to t_i. We will abuse our notation slightly and also denote by f_i the value of the flow f_i; given the context this should not cause any confusion.

Each arc $a \in A$ has a load-dependent *latency function*, denoted by $l_a(f)$. The latency of a path P with respect to a flow f is defined as the sum of the latencies of the edges in the path, denoted by $l_P(f) = \sum_{a \in A} l_a(f_a)$. The latency with respect to an agent i is $l_i(f) = \sum_{P_i \in \mathcal{P}_i} l_{P_i}(f) f_{P_i}$. The latency $l(f)$ of a flow f is the total latency incurred by f i.e. $l(f) = \sum_{a \in A} l_a(f_a) f_a = \sum_{P \in \mathcal{P}} l_P(f) f_P = \sum_i l_i(f)$. In [9] the social objective is to minimise the total latency, given that a flow of value r_i must be routed from s_i to t_i. The private objective of an agent i is to minimise its own latency i.e. $l_i(f)$. We consider a maximisation version of this problem. Each agent may route a flow of weight at most r_i from s_i to t_i. Associated with each source-destination $\{s_i, t_i\}$ pairing is a value π_i that signifies the revenue (utility) from routing one unit of flow from s_i to t_i. However, we still associate with a routing the latency-based cost. Thus, a flow f that successfully routes f_i units of flow from s_i to t_i will induce a profit to agent i of $\zeta_i(f) = \pi_i\, f_i - l_i(f)$. Hence, the social objective is to maximise the function

$$\kappa(f) \;=\; \sum_i \zeta_i(f) \;=\; \sum_i \pi_i\, f_i - l_i(f)$$

and agent i seeks to maximise the private objective function ζ_i. We will now show that this problem also fits into our

Proceedings of the 43 rd Annual IEEE Symposium on Foundations of Computer Science (FOCS'02)

framework. To do this we will discretise the problem by assuming that flow may be sent only in whole unit increments; for this problem it is not difficult to generalise the results to continuous space.

Lemma 7. *The routing problem can be formulated in the action set framework.*

Proof. The action space $\mathcal{A}_i$ of agent i consists of any flow f_i of value at most r_i from s_i to t_i. We now show how this fits into our framework. For each agent i we have a collection of paths $\mathcal{P}_i$ from s_i to t_i. The agent assigns a weight to each path $p_i \in \mathcal{P}_i$. Let the groundset V_i consist of r_i copies of each path p_i i.e. $p_i^1, \ldots, p_i^{r_i}$. Here the choice of p_r^t correspond to the routing of t units of flow on path p_i.

We may allow an agent to select multiple copies of a path. In such a circumstance only the action corresponding to the copy with the greatest amount of flow is implemented. (Alternatively, we may restrict the action space of agent i to allow for the choice of at most one copy of each path p_i). Note that if no copy of p_i is chosen then no flow is sent along that path. □

Now consider that latency functions $l_a(f)$. We will assume that these functions are non-negative, non-decreasing and convex. Note that these assumptions correspond to some natural properties of traffic systems. The non-decreasing property implies that the costs incurred increase as the volume of the traffic increases; the convexity property implies that the additional costs incurred (by adding an additional unit of traffic) increase as the volume of the traffic increases. Observe that convexity implies that the latency functions are supermodular when restricted to our discretised space. It follows easily that

Lemma 8. *For the selfish routing problem, the social objective function κ is submodular.* □

Lemma 9. *For the selfish routing problem, the system (κ, ζ_i) is a valid utility system.*

Proof. We show, for each agent i, that $\zeta_i(f) \geq \kappa'_{f_i}(f \oplus \emptyset_i)$. Now $\kappa'_{f_i}(f \oplus \emptyset_i) = \kappa(f) - \kappa(f - f_i) = \sum_j (f_j \pi_j - l_j(f)) - \sum_{j:j \neq i} (f_j \pi_j - l_j(f - f_i)) = f_i \pi_i - l_i(f) + \sum_{j:j \neq i} (l_j(f - f_i) - l_j(f)) \leq f_i \pi_i - l_i(f) = \zeta_i(f)$. Thus (κ, ζ_i) is a utility system. We have already seen that $\kappa(f) = \sum_i \zeta_i(f)$ and, thus, the utility system is valid. □

So we then obtain the following guarantees.

Theorem 13. *For the selfish routing problem, any Nash equilibrium $S \in \mathcal{S}$ satisfies* $\text{OPT} \leq 2\bar{\kappa}(S) - \sum_i \bar{\kappa}'_{s_i}(S \cup \Omega - s_i)$. □

Thus we obtain a factor 2 guarantee if, for example, $\bar{\kappa}'_{s_i}(S \cup \Omega - s_i) \geq 0, \forall i$. An alternative guarantee follows from Theorem 4. This compares the value of a Nash equilibrium S against the social value of a particular solution, $S + \Omega$, that routes twice as much traffic.

Theorem 14. *For any Nash equilibrium $S \in \mathcal{S}$, we have* $2\bar{\kappa}(S) \geq \bar{\kappa}(\Omega \cup S) + \sum_{i:\sigma_i \in \Omega \cap S} \bar{\kappa}'_{\sigma_i}(S \oplus \emptyset_i) \geq \bar{\kappa}(S + \Omega)$. □

A result of this flavour also follows from the work of [9]; the social value of a Nash equilibrium is at least the social value of the optimal solution that routes twice as much traffic when the all the rewards π_i are halved. If κ is non-decreasing (hence, it is always in the interest of agent i to route all r_i units of flow), then from Theorem 5 we obtain

Theorem 15. *If κ is non-decreasing then, for the selfish routing problem, any Nash equilibrium $S \in \mathcal{S}$ satisfies* $\text{OPT} \leq (1 + \delta(\kappa))\bar{\kappa}(S) \leq 2\bar{\kappa}(S)$. □

8 Polynomial Time Considerations

Our discussion regarding pure strategy Nash equilibria touched upon the importance of speed considerations in the strategy determination. We discuss this in more detail in this section. Let us measure the size of the problem input in terms of the size of the groundsets V_i, $1 \leq i \leq k$. It would be useful if we obtained a Nash equilibria in polynomial time in the problem size. Two factors are important here:
(i) Bounding the number of times an agent changes strategy before a Nash equilibria is obtained.
(ii) Bounding the time an agent takes to decide upon a strategy.
How to bound the number of iterations required before convergence to a Nash equilibria is an important open question. In the presence of pure Nash equilibria, as we have seen, the overall size of the state space gives one upper bound. We note, however, that good guarantees may be obtained within a constant number of iterations (we only need each agent to change strategies a constant number of times). That is, solutions that arise long before we reach a Nash equilibria also provide good guarantees. Thus, although these solutions may not be stable, they do give good performance. We omit the details here.

Regarding the second factor, if the size of the action space $\mathcal{A}_i$ of agent i is polynomial in $|V_i|$, then the agent can easily find its best strategy in polynomial time. However, the action space $\mathcal{A}_i$ may be as large as $2^{|V_i|}$. Thus in some circumstances it may not be possible to find an optimal strategy quickly. It may, though, be possible to obtain approximately optimal strategies in polynomial time. We will show that the use of approximation algorithms by the agents in their strategy determination does lead to guarantees on the social performance of Nash equilibria. We

Proceedings of the 43 rd Annual IEEE Symposium on Foundations of Computer Science (FOCS'02)
0272-5428/02 $17.00 © 2002 IEEE

have one difficulty to overcome though. The use of approximately optimal strategies is not consistent with the concept of a Nash equilibria. That is, approximately optimal strategies are *not* the optimal best response strategies required by Nash equilibria. Thus, we are really using *approximate Nash equilibria*. They are equilibria in the sense that no agent can find (by whatever methods they are using) a better alternative strategy in polynomial time. So suppose that each agent has access to an approximation algorithm at each stage. Let these algorithms have an approximation guarantee of ξ, say. Then, Theorem 3, Theorem 4 and Theorem 5 apply (with slightly weaker guarantees) to approximate Nash equilibria. For example, if our social utility function is non-decreasing, we have the following theorem.

Theorem 16. *Let γ be a non-decreasing, submodular set function, and $(\gamma, \cup_i \alpha_i)$ be a valid utility system. If the agents can generate ξ-approximate solutions, then for any approximate Nash equilibrium $S \in \mathcal{S}$ we have* $\text{OPT} \leq (\xi + \delta(\gamma))\bar{\gamma}(S) \leq (\xi + 1)\bar{\gamma}(S)$. □

For an example consider the case of matroids. A *matroid* $\mathcal{T}$ is a family of subsets of V such that (i) $\emptyset \in \mathcal{T}$, (ii) If $Y \in \mathcal{T}$ and $X \subseteq Y$, then $X \in \mathcal{T}$ and (iii) If $X, Y \in \mathcal{T}$ and $|X| < |Y|$, then $\exists y \in Y - X$ such that $X \cup \{y\} \in \mathcal{T}$. Fisher, Nemhauser and Wolsey [4] gave a simple 2-approximation algorithm for the problem of maximising a non-decreasing, submodular function over a matroid. Thus, if each action set $\mathcal{A}_i$ is a matroid then we have

Corollary 3. *Let γ be a non-decreasing, submodular set function, and $(\gamma, \cup_i \alpha_i)$ be a valid utility system. If each $\mathcal{A}_i$ is a matroid, then we obtain an approximate Nash equilibrium $S \in \mathcal{S}$ with* $\text{OPT} \leq (2 + \delta(\gamma))\bar{\gamma}(S) \leq 3\bar{\gamma}(S)$. □

9 Multiple-Item Auctions

Consider the following class of auction: there is one seller (auctioneer) with a set J of n different items, and a set of k potential buyers (agents) who have a private valuation for each subset of items. Recently, Lehmann, Lehmann and Nisan [6] considered the *allocation problem* induced by this framework. There, a single authority wishes to find an allocation of optimal efficiency (social value). They present a polynomial time algorithm that produces an allocation with social value at least one half that of the optimal solution, provided that the agents valuations are submodular. Again, our interest is in the competitive situation in which the seller and buyers all seek to maximise their own utility. We present a simple class of *multi-round auction* that is guaranteed to produce an allocation within a factor 2 of optimal, despite the valuation functions being private knowledge and with the sellers and buyers acting in a selfish manner. Moreover, the allocation procedure of [6] can easily be implemented within this class of auction. Complete details can be found in the full paper.

Acknowledgments. I am especially grateful to Tim Roughgarden for comments and suggestions that greatly improved this paper. I am also very grateful to Bruce Shepherd, Andreas Schulz and Santosh Vempala for interesting discussions on this work. Finally, I would like to thank Eva Tardos for suggesting auctions as a possible application of the method.

References

[1] M. Conforti and G. Cornuéjols, "Submodular set functions, matroids and the greedy algorithm: tight worst-case bounds and some generalizations of the Rado-Edmonds theorem", *Discrete Applied Mathematics*, **7**, pp251-274, 1984.

[2] G. Cornuéjols, M. Fischer and G. Nemhauser, "Location of bank accounts to optimize float: an analytic study of exact and approximation algorithms", *Manag. Sci.*, **23**, pp789-810, 1977.

[3] M. Charikar, S. Guha, E. Tardos and D. Shmoys, "A constant-factor approximation algorithm for the k-median problem", *Proc.* 31^{st} *STOC*, pp1-10, 1999.

[4] M. Fisher, G. Nemhauser and L. Wolsey, "An analysis of approximations for maximising submodular set functions II", *Math. Prog. Study*, **8**, pp73-87, 1978.

[5] E. Koutsoupias and C. Papadimitriou, "Worst-case equilibria", *Proc.* 16^{th} *STACS*, pp404-413, 1999.

[6] B. Lehmann, D. Lehmann and N. Nisan, "Combinatorial auctions with decreasing marginal utilities", *ACM Conf. on Electronic Commerce*, 2001.

[7] J. Nash, "Non-cooperative games", *Annals of Mathematics*, **54**, pp268-95, 1951.

[8] C. Papadimitriou, "Algorithms, games, and the internet", *Proc.* 33^{rd} *STOC*, 2001.

[9] T. Roughgarden and E. Tardos, "How bad is selfish routing?", *Proc.* 41^{st} *FOCS*, pp93-102, 2000.