

Market Sharing Games Applied to Content Distribution in Ad Hoc Networks

Michel X. Goemans, Li (Erran) Li, Vahab S. Mirrokni, and Marina Thottan, *Member, IEEE*

Abstract—In third-generation (3G) wireless data networks, repeated requests for popular data items can exacerbate the already scarce wireless spectrum. In this paper, we propose an architectural and protocol framework that allows 3G service providers to host efficient content distribution services. We offload the spectrum intensive task of content distribution to an ad hoc network. Less mobile users (resident subscribers) are provided incentives to cache popular data items, while mobile users (transit subscribers) access this data from resident subscribers through the ad hoc network. Since the participants of this data distribution network act as selfish agents, they may collude to maximize their individual payoff. Our proposed protocol discourages potential collusion scenarios. In this architecture, the goal (social function) of the 3G service provider is to have the selfishly motivated resident subscribers service as many data requests as possible. However, the choice of which set of items to cache is left to the individual user. The caching activity among the different users can be modeled as a market sharing game.

In this work, we study the Nash equilibria of market sharing games and the performance of such equilibria in terms of a social function. These games are a special case of congestion games that have been studied in the economics literature. In particular, pure strategy Nash equilibria for this set of games exist. We give a polynomial-time algorithm to find a pure strategy Nash equilibrium for a special case, while it is NP-hard to do so in the general case. As for the performance of Nash equilibria, we show that the *price of anarchy*—the worst case ratio between the social function at any Nash equilibrium and at the social optimum—can be upper bounded by a factor of 2. When the popularity follows a Zipf distribution, the price of anarchy is bounded by 1.45 in the special case where caching any item has a positive reward for all players. We prove that the selfish behavior of computationally bounded agents converges to an approximate Nash equilibrium in a finite number of improvements. Furthermore, we prove that, after each agent computes its response function once using a constant factor approximation algorithm, the outcome of the game is within a factor of $O(\log n)$ of the optimal social value, where n is the number of agents. Our simulation scenarios show that the price of anarchy is 30% better than that of the worst case analysis and that the system quickly (1 or 2 steps) converges to a Nash equilibrium.

Index Terms—Mobile ad hoc networks, Nash equilibrium, price of anarchy, third-generation (3G) wireless networks, unified architecture.

Manuscript received February 15, 2005; revised January 15, 2006. This work was supported in part by the Office of Naval Research under Grant N00014-05-1-0148. The work of L. Li was supported in part by the National Science Foundation (NSF) under Grant ANI-0335244. This paper was presented at the ACM Mobile Ad Hoc Networking and Computing Symposium, 2004.

M. X. Goemans and V. S. Mirrokni are with the Department of Mathematics and the Computer Science and Artificial Intelligence Laboratories, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: goemans@math.mit.edu; mirrokni@theory.csail.mit.edu).

L. Li and M. Thottan are with the Center for Networking Research, Bell Laboratories, Lucent Technologies, Murray Hill, NJ 07974 USA (e-mail: erranli@research.bell-labs.com; marinat@research.bell-labs.com).

Digital Object Identifier 10.1109/JSAC.2006.872884

I. INTRODUCTION

THIRD-GENERATION (3G) wide-area wireless networks have recently experienced tremendous growth, with the number of subscribers reaching more than 70 million worldwide [1]. Many 3G service providers have started offering content-rich services such as sports replays, news headlines, music videos, and movie trailers [1].

The 3G subscriber market can be categorized into groups with shared interest in location-based services, e.g., the preview of movies in a theater or the scene of the beach nearby. Since the 3G radio resources are limited, it is expensive to repeatedly transmit large quantities of data over the air interface from the base station (BS). It is more economical for the service provider to offload such repeated requests on to the ad hoc network comprised of its subscribers, where some of them may have recently acquired a copy of the data. In this scenario, the goal for the service provider is to give incentives for peer subscribers in the system to cache and forward the data to the requesting subscribers. Since each data item is large in size and transit subscribers are mobile, we assume that the data transfer occurs in within a range of a few hops.

We envision a system consisting of two groups of subscribers: resident and transit subscribers. Resident subscribers are less mobile and mostly confined to a certain geographical area. Resident subscribers have incentives to cache data items that are specific to this geographical region since the service provider gives monetary rewards for satisfying the queries of transit subscribers. Transit subscribers request their favorite data items when they visit a particular region. Since the service provider does not have knowledge of the spatial and temporal distribution of requests, it is difficult if not impossible for the provider to stipulate which subscriber should cache which set of data items. Therefore, the decision of what to cache is left to each individual subscriber. The realization of this content distribution system depends on two main issues. First, since subscribers are selfish agents, they may collude to increase their individual payoffs. Collusion can result in other subscribers being cheated of their rewards. In this work, we address the problem of colluding subscribers by providing a protocol framework that discourages or prevents collusions. The second issue is that the payoff of each item for each agent must be a function of the number of cache requests it services. This in turn depends on the number of agents who cache a given item since each agent has limited storage space (due to form factor, current 3G devices only have a couple of MB flash memory [1]). Therefore, each selfish agent may change the set of items it caches in response to the set of items cached by others. This leads to a noncooperative caching scenario which we model as a market sharing game.

In the market sharing game, the primary questions are whether the system converges, i.e., results in a pure strategy Nash equilibrium and how long it will take to converge. The goal of the service provider is to offload as many cache requests as possible to the ad hoc network. We refer to this goal as the social optimum. However, given the selfish behavior of the agents, it is unlikely that the system will result in a social optimum. Therefore, we would like to bound the ratio between the optimal social value and the outcome of the selfish behavior of players. We refer to this ratio as the *price of anarchy*. Furthermore, when computing the selfish behavior of individual players, it is essential to consider the computational constraint on the individual subscribers. We model computationally bounded agents using approximation algorithms and evaluate how fast the selfish behavior of such agents will converge to an approximate Nash equilibrium or arrive at an approximate solution to the social function.

The main contributions of this paper are as follows. First, we study the applicability of noncooperative caching in wireless networks and propose a detailed protocol that provides incentives for selfish agents to service other agents, while discouraging collusion among participating agents. We model the caching game among the different mobile users as a market sharing game. We show that this is a special case of congestion games. It is known that pure strategy Nash equilibria exist for these games [21]. We give a polynomial-time algorithm to find a pure strategy Nash equilibrium for the special case of uniform market sharing games, where all items are of the same size. We know that it is NP-hard to find such an equilibrium in the general market sharing game. As for the performance of Nash equilibria, we obtain an upper bound on the price of anarchy for two different cases. In the general case where the popularity of the items is arbitrary, the price of anarchy is upper bounded by 2. When the popularity follows a Zipf distribution, we have an upper bound of 1.45 in the special case where caching any item has a positive reward for all players. We give an example that shows that the factor 2 is tight in the general case even with Zipf distribution. We also consider the case of computationally bounded agents by modeling them using approximation algorithms. We prove that, after each agent computes its response function once using a constant factor approximation algorithm, the outcome of the game is within a factor of $O(\log n)$ of the optimal social value, where n is the number of agents. We also show that this bound is tight up to a constant factor. We also evaluate our algorithms using simulation scenarios and show that the price of anarchy in our network setting is less than 70% of the worst case analysis. We also demonstrate that in most cases the system converges to a Nash equilibrium in one or two steps of each player. A step is defined as the greedy action of a player to change to a different set of items given the actions of other players.

The rest of this paper is organized as follows. In Section II, we propose a detailed protocol for offloading content distribution from 3G wireless networks to ad hoc networks. In Section III, we discuss our incentive mechanism and show that it prevents or discourages collusion. In Section IV, we formulate the noncooperative caching problem as a general market sharing game. In Section V, we analyze the price of anarchy of our market sharing

game and in Section VI, we show that pure strategy Nash equilibrium exists, but it is NP-hard to find such an equilibrium. We also provide a polynomial-time algorithm to find a pure strategy Nash equilibrium for the special case when all the items have the same size. In Section VII, we model computationally bounded agents by approximation algorithms and analyze the outcome of their selfish behavior. We observe that such agents converge to an approximate Nash equilibrium. We also show that, after one step of improvements the outcome of the game is within an $O(\log n)$ factor of the optimum. In Section VIII, we investigate the price of anarchy for a set of sample instances and the convergence rate to exact and approximate Nash equilibria. Related work is provided in Section IX. A discussion on our results and other relevant issues are described in Section X. We conclude in Section XI.

II. SYSTEM ARCHITECTURE, TRUST MODEL, AND PROTOCOLS

In this section, we briefly discuss the system architecture and the trust model for distributed noncooperative caching. We then outline the protocols required for offloading popular data items from 3G networks to the multihop ad hoc network. Our architecture is largely the same as UCAN [14]. In our system architecture, we assume the following.

- A service provider operates a 3G data network, e.g., 1xEV-DO, CDMA2000-1X, etc. The cost of retrieving an item from the service provider is assumed to be significantly higher than obtaining the same item from the ad hoc network.
- The mobile device of each subscriber has a 3G interface, as well as an IEEE 802.11 radio interface so that it can receive data from the 3G data network and can participate in ad hoc forwarding.
- Users are selfish and act rationally. A user's primary objective is to maximize its own utility. The utility is in terms of obtaining the item of interest or monetary gains for forwarding or serving the requested item to another subscriber in the network.
- Users trust the service provider network for correct accounting, authentication, and packet transmission. We do not deal with malicious users in this paper. However, users could collude arbitrarily if they can benefit by doing so.

Each node u has a shared key K_u with the service provider. When node u requests an item from the service provider, the provider first checks whether the item has been cached by other local resident subscribers (subscribers keep the providers updated about the items they cache). If there are no cached copies for the item, the provider will service u directly. However, if the item is available through a subscriber cache, the provider will instruct u to look for the item in the ad hoc network. Along with this response the provider also sends descriptive information regarding the item requested, for example, the item ID and the size. Alternatively, if a node knows the item ID, it can request the item from the ad hoc network directly. Our protocol messages are illustrated in Fig. 1. The first message is optional if the node knows the item ID.

Node u then broadcasts a CacheReq message for the specific item to all k neighbors in the ad hoc network. Node u includes

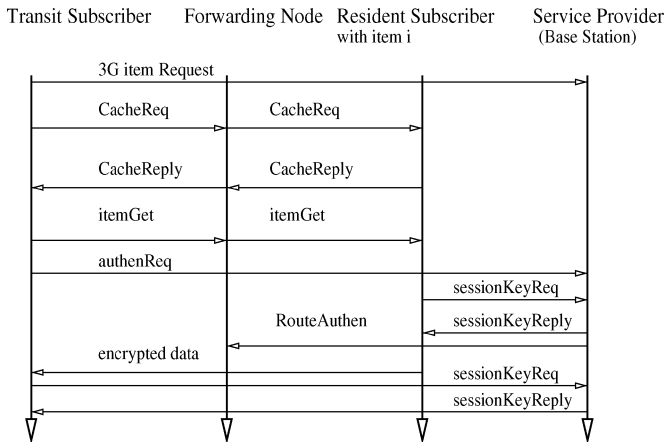


Fig. 1. Protocol messages.

a Message authentication code (MAC) in the CacheReq message. See [14] for details on how to construct MAC codes. The CacheReq message includes the address of u and the requested item ID. Upon receiving the message any resident node that has the requested item will send a CacheReply message to u and attach its MAC computed using its shared key with the provider. The ad hoc network performs source routing, and therefore the final message received by the node u will contain a string of MACs or a layered MAC [24] for all the intermediate hops.

Node u may receive multiple CacheReply messages. It may choose to be serviced by the resident subscriber that has the minimum number of hops. Node u asks the service provider to authenticate the chosen route by sending an authenReq message. The service provider in turn registers the transaction along with the route used. When the BS receives the authenReq and sessionKeyReq message, the BS constructs a RouteAuthen message. The message contains a list of [Session ID (SID), source u , destination v , previous hop, next hop, item size] encrypted using the shared key of each node along the path. An intermediate node w forward traffic for session SID only if it can decode the message(SID, u , v , previous hop, next hop, item size) using K_w shared with the service provider. Once the route authentication has been completed, node u sends an itemGet message to the source node v . When node v receives the itemGet message, it requests a session key K_s from the BS by sending a sessionKeyReq message. This message contains the session ID and the source ID v and destination ID u . The BS responds with a session key encrypted using K_v .

Upon receiving the session key the data item is encrypted using this key. Node v appends a MAC to each packet of the data item using key K_v , and it then sends the packet on the ad hoc network. Each intermediate node w computes a MAC over the received packet using key K_w it shares with the service provider. It replaces the MAC in the received packet with this new MAC (for details of MAC layering, see [24]) and sends the packet to the next hop. During data transfer in the ad hoc network each intermediate node monitors the packets that are being forwarded by the next node in order to prevent misuse of the authenticated ad hoc path. When the entire file transfer has been completed, node u requests the session key from the BS to decode the data received. When an intermediate relay node

forward data by more than a threshold of the itemSize, it will inform the BS. Forwarding nodes also randomly sample packets with a low probability and report the packet along with its contents to the BS.

Our protocol is lightweight. For one item transfer, we have at most six short messages between the service provider and the nodes involved plus a small percentage of random sampled packets reported to the BS.

III. CREDIT ACCOUNTING AND AUDITING

In this section, we present schemes for charging and rewarding. We also discuss mechanisms that aid in tracking the credits for resident subscribers and mechanisms that prevent or discourage cheating and collusion among the participants to gain false credits. Since security is not the main focus, we only provide an informal discussion. For a more rigorous treatment, please refer to [10].

There are two types of monetary rewards: 1) R_i for successfully servicing the query for item i and 2) f_i for forwarding the query. $R_i > f_i$ and for a given item the total forwarding reward is fixed and shared equally among the forwarding nodes. Both types of rewards are a function of the file size of the item. The receiving node is charged $C_0(i)$ if it obtains the item from subscribers in the ad hoc network. While if the data item is obtained from the service provider the receiving node is charged $C_S(i)$. $C_S(i)$ is much larger than $C_0(i)$. Participating nodes are charged or rewarded only if BS determines that the item transfer is successfully completed. The service provider can tune R_i and f_i so that a critical number of resident subscribers are willing to participate in caching the popular items.

Our data transfer protocol proposed in the previous section provides no incentive for the following types of behavior.

- 1) *Stealing reward from forwarding nodes*: The repeated collusion of the sender and receiver to leak the session key over a covert channel, thus cheating the forwarding nodes of their rewards is prevented by the record keeping done by the BS. This record keeping is accomplished by the random sampling of packets in the intermediate nodes and the information obtained is used to black mark source-destination pairs for attempting to cheat the forwarding nodes.
- 2) *Refusal to pay by the receiver*: The receiver u cannot refuse to pay if it requires the session key from the BS in order to decrypt the content. If it resorts to a covert channel between the sender and receiver to obtain the session key, then BS can determine with high probability that the session is complete through the sampled packets from forwarding nodes.
- 3) *Impersonate the sender*: There is no incentive to impersonate the sender since a sender only gets paid if the receiver acknowledges the service provider about a successful completion of the data transfer. The provider will only pay the original sender that was registered in its database.
- 4) *Packet dropping*: There is no incentive for intermediate nodes to drop packets since they are credited only if the

item is transferred completely and the credit is based on the size of the item, not the amount of data relayed.

- 5) *Free riding*: (a) Since intermediate nodes only forward traffic for a session if it receives authentication from the BS, the source cannot collude with the receiver such that the intermediate nodes do not get paid. (b) The source node or any intermediate node cannot collude with another intermediate node to piggyback data. The “man-in-the-middle” victim knows the size of the item, and it can report this to the provider. Although an individual node is not to be trusted, if many nodes report the same problem with the source node, the provider can punish the source. An intermediate node has no incentive to implicate the source falsely since reporting to the provider incurs a cost. However, if all the intermediate relay nodes are colluding nodes, then they merely form a cooperative ad hoc group. In our protocol, we require that intermediate nodes sample packets and report them to the BS. If the BS cannot decrypt the packet correctly due to piggybacked extra data, the BS knows that someone has changed the packet. (c) If a destination node colludes with an intermediate node to get a free ride for their own data and drops legitimate traffic, the destination will have a high session failure rate. The BS can blacklist all such nodes with a high session failure rate. In addition, due to the fact that packets are sampled to the BS, the BS knows that the packet has been changed. (d) If two relay nodes collude and get a free ride on their data and drop legitimate traffic, monitoring by the previous hop node who sent legitimate traffic will reveal that packets other than the one it sent got transmitted. The monitoring node can then report this to the provider. Based on such reports from many sessions, the provider can determine a cheating node with high probability. In all the above cases, if the BS knows that data has been tampered by intermediate nodes from MAC decoding failure, the BS can take immediate actions, e.g., asking the destination to locate another node with the item or find an alternative route.
- 6) *Suboptimal routes*: There is less incentive for a node to add other nodes in the route since a fixed reward is divided among all forwarding nodes. In addition, since the receiver is likely to select the shortest route, there is no advantage to announcing a longer route.

IV. PROBLEM FORMULATION

In this section, we formalize our caching problem abstractly in terms of a competitive market sharing game.

In a *market sharing game*, there is a set U of n agents or players and a set H of m markets. Each market $i \in H$ has a query rate q_i , i.e., the rate at which market i is requested per unit time. Each market i also has a cost C_i corresponding to the cost for servicing this market. Each agent j has a total budget B_j . We are also given a bipartite graph $G = (H \cup U, E)$ in which an edge between agent j and market i in G ($(j, i) \in E(G)$) means that i is of interest to agent j (we write j is interested in market i). Each agent should decide which subset of markets to service. Agent j can service a subset S_j of markets, if the sum

of the costs C_i of the markets in S_j is less than or equal to B_j , i.e., $\sum_{i \in S_j} C_i \leq B_j$. Agent j gets a reward $R_i = q_i/n_i$ for providing service to market i , where n_i is the number of agents that service market i . Observe that the total reward received by all players equal the total query rate of the markets being serviced (by at least one player).

We now define the necessary game theoretic notations to formally describe the problem. A strategic game is defined as the tuple $(U, \{A_j\}, \{P_j(\cdot)\})$, where U is the set of players or agents. A_j is the set of actions or strategies for player j and $P_j : \prod_j A_j \rightarrow R$ is the payoff or utility function for agent j given the set of strategies of all the players. In our model, feasible actions are the set of markets that can be serviced under the given budget constraint. We denote player j 's action by S_j . Thus, S_j is a feasible action if $\sum_{i \in S_j} C_i \leq B_j$. Given the set, \mathcal{S} , of actions for all agents, we can find the number of agents n_i that serve market i , and hence find the reward of each market. The payoff or utility function of player j is the sum of rewards it gets from the markets it serves, i.e., $P_j(\mathcal{S}) = \sum_{i \in S_j} (q_i/n_i)$, where S_j is the set of markets j serves. We call the above game a *market sharing game* and the special case of uniform ($C_i = C$) cost for all markets a *uniform market sharing game*. In this game, each agent wants to maximize its own payoff. The *social function* is the total amount of queries satisfied in the market, i.e., $\gamma(\mathcal{S}) = \sum_{i \in \cup_j S_j} q_i$ for $\mathcal{S} = (S_1, S_2, \dots, S_n)$. Notice that this is also the sum of the utility functions of all players, i.e., $\gamma(\mathcal{S}) = \sum_{j \in U} P_j(\mathcal{S}) = \sum_{i \in \cup_j S_j} q_i$. We refer to a strategy profile that maximizes the social function as *social optimum* or *optimum* and the value of this profile as *optimal social value* or *the optimum*. We denote the social optimum by **OPT**.

It is obvious that, in a market sharing game, given the set of actions of other players, the best action of an agent can be obtained by solving a knapsack problem where the value of market (item) i is equal to q_i/n_i or $q_i/(n_i + 1)$ depending on whether market i is currently serviced by this player or not. The size of i in the knapsack instance is C_i , and the knapsack capacity is equal to B_j .

Example: The subscriber caching game described in the introduction is one example of a market sharing game. Agents or players correspond to subscribers and markets correspond to data items. Query rate of market i is the query rate of item i . The cost C_i corresponds to the size of item i and the total budget B_j of agent j corresponds to the available storage space on the subscriber's device.

As mentioned above, in this market sharing game, a feasible action S_j is a feasible solution to the knapsack problem of player j . An action profile (or strategy profile) is a vector of actions of all players: $\mathcal{S} = (S_1, S_2, \dots, S_n)$. A action profile \mathcal{S} is a pure strategy Nash equilibrium, if for any player j , given the strategies of all other players, j has no incentive to change its strategy S_j to any other subset to improve its payoff. A strategy S'_j is an *improvement move* for player j in strategy profile $\mathcal{S} = (S_1, S_2, \dots, S_n)$ if the payoff of player j strictly increases in \mathcal{S} by changing his strategy from S_j to S'_j , i.e., $P_j(\mathcal{S}) < P_j(S_1, \dots, S_{j-1}, S'_j, S_{j+1}, \dots, S_n)$.

A mixed strategy Nash equilibrium is a randomization over different pure strategies such that fixing other players' strategy, each player maximizes its expected payoff using the current

strategy. Nash equilibria characterize the set of candidate action profiles that selfish players will eventually converge upon if at all.

As mentioned earlier, we focus on market sharing games that have a central authority with a social objective function to be accomplished by means of competitive market sharing. Therefore given a social function, we want to bound the *price of anarchy*, i.e., the worst case ratio between the social function value at the optimum and at any Nash equilibrium.

The general market sharing game described here is a special case of congestion games [21]. These games have been studied in the economics literature and the existence of a pure Nash equilibrium has been established. The markets of the market sharing game correspond to the factors (or elements) in the definition of congestion games [21]. We will use this fact to prove that the best response dynamics in this game converge to pure Nash equilibria. Furthermore, we prove upper bounds on the efficiency of Nash equilibria in this game and provide some results on the speed of convergence to approximately efficient solutions.

V. PRICE OF ANARCHY

A. General Distribution

Studying the outcome of games in terms of their price of anarchy is a natural way to evaluate the outcome of a game. It is not hard to verify that the pure Nash equilibrium of the market sharing game is not necessarily unique. As a result, the price of anarchy of this game is concerned with the worst case equilibria. In this section, we provide some bounds on the performance of the worst Nash equilibria of the market sharing game by bounding the price of anarchy. We do this by showing that the market sharing game is a *valid-utility game*.

Valid-utility games, introduced by Vetta [26], is a set of non-cooperative games for which Vetta proved upper bounds for their price of anarchy. We show that the market sharing game is in this class of games which implies that its price of anarchy is at most 2. In fact, this bound for the price of anarchy can be obtained without referring to valid-utility games. The main advantage of proving this bound using valid-utility games is that it can be generalized to more general variants of market sharing games. In order to state this result, we need the following definitions.

Definition 5.1: A function of the form $f : 2^{\mathcal{H}} \rightarrow \mathbb{R}$ is called a set function. A set function $f : 2^{\mathcal{H}} \rightarrow \mathbb{R}$ is submodular if $f(\emptyset) = 0$ and for any two sets $A, B \subseteq \mathcal{H}$, $f(A) + f(B) \geq f(A \cap B) + f(A \cup B)$. This function is a nondecreasing function, if for any $X \subseteq Y \subseteq \mathcal{H}$, $f(X) \leq f(Y)$.

Given an action profile $\mathcal{S} = (S_1, \dots, S_n)$, the set $\mathcal{H}_{\mathcal{S}} = \{(k, i) : 1 \leq k \leq n, i \in S_k\}$ is called the *pair set for \mathcal{S}* . Given a function $f : \prod_k S_k \rightarrow \mathbb{R}$, the corresponding set function, f^s , of f is a set function of the form $2^{\mathcal{H}} \rightarrow \mathbb{R}$, where $\mathcal{H} = \{(k, i) : 1 \leq k \leq n, 1 \leq i \leq m\}$ and $f^s(\mathcal{H}_{\mathcal{S}}) = f(\mathcal{S})$. In other words, $f^s(K) = f((A_1, A_2, \dots, A_n))$ if $A_j = \{i : (j, i) \in K\}$.

Definition 5.2: Let $\mathcal{G}(U, \{S_j\}, \{P_j(\prod_k S_k)\})$ be a noncooperative game with social function $\gamma : \prod_k S_k \rightarrow \mathbb{R}$. \mathcal{G} is called a *valid-utility game* if it satisfies these properties.

- γ^s , the corresponding set function of γ , is submodular and nondecreasing.

- The payoff of a player is at least equal to the difference in the social function when the player participates versus when it does not participate.
- The sum of the utility or payoff functions for any set of strategies should be less than or equal to the social function.

Theorem 5.3: [26] Let \mathcal{G} be a valid-utility game, then for any mixed strategy Nash equilibrium, the social function at this equilibrium is at least half the optimal social function, i.e., in a Nash equilibrium A , $\gamma(\mathbf{OPT}) \leq 2E[\gamma(A)]$.

Theorem 5.4: The market sharing game is a valid-utility game.

Proof: We need to show that our social function given by $\gamma(\mathcal{S}) = \sum_{i \in \cup_{j \in U} S_j} q_i$ satisfies the three properties itemized above.

- First, it is clear that γ^s is nondecreasing. To show its submodularity, we use an equivalent definition of submodular functions: A set function f is submodular if for any two subsets A and B such that $A \subset B$ and for any element $i \notin B$, $f(A \cup \{i\}) - f(A) \geq f(B \cup \{i\}) - f(B)$ (see, for example, Lemma 2.2 of [15]). Thus, in order to prove that γ^s is submodular, it is enough to prove that for two (possibly infeasible) action profiles $\mathcal{S} = (S_1, \dots, S_n)$ and $\mathcal{S}' = (S'_1, \dots, S'_n)$ such that $S_j \subseteq S'_j$ for all $j \in U$, by adding a new market i to the action set of any player, j , the increase in γ^s for \mathcal{S} is no less than the increase for \mathcal{S}' . If market i is not in $\cup_{j \in U} S'_j$, its addition to one of the S'_j increases γ^s by q_i . In this case, market $i \notin \cup_{j \in U} S_j$, because $\cup_{j \in U} S_j \subseteq \cup_{j \in U} S'_j$. Thus, adding market i to \mathcal{S} also increases γ^s by q_i . If i is in $\cup_{j \in U} S'_j$, adding i will not increase γ^s . Thus, in any case, the increase for \mathcal{S} is no less than the increase for \mathcal{S}' .
- The difference in the social function when j plays S_j or empty (does not play at all) is equal to $\sum_{i \in S_j, n_i=1} q_i$ and this is indeed less than or equal to $\sum_{i \in S_j} (q_i/n_i)$.
- By the definition of our social function, we have $\sum_{j \in U} P_j(\mathcal{S}) = \gamma(\mathcal{S})$, and therefore the third property is satisfied as well. ■

To deal with the situation in which agents only compute an approximate solution to the knapsack problem, we consider approximate Nash equilibria. A β -approximate Nash equilibrium is an action profile in which no player can unilaterally change his action and get a payoff that is more than β times his current payoff. As an extension of Theorem 5.3, Vetta proved that the optimal social value is at most $\beta + 1$ times the social value of a β -approximate Nash equilibrium [26]. Using this result and Theorem 5.4, we have the following corollary.

Corollary 5.5: The optimal social value is at most $\beta + 1$ times the social value of a β -approximate Nash equilibrium. This corollary can also be proved directly without referring to valid-utility games, see below. The proof based on valid-utility games can be generalized to more general variants of market sharing game and is discussed in Section X.

Proof: Let Y_j be the payoff of player j in the approximate Nash equilibrium. Let X_j be the sum of the q_i 's over the set of markets that player j serves in \mathbf{OPT} but are not served by any agent in the equilibrium being considered. By the definition of approximate Nash equilibrium, we have $\beta Y_j \geq X_j$. Otherwise, player j would switch action. Furthermore, $\sum_{j \in U} Y_j \geq$

$\gamma(\mathbf{OPT}) - \sum_{j \in U} X_j$, as $\sum_{j \in U} X_j$ is at least the total amount of query of the markets serviced by \mathbf{OPT} but not by the approximate Nash equilibrium. Thus, $\gamma(\mathbf{OPT}) \leq (\beta + 1) \sum_{j \in U} Y_j$, as desired. ■

B. Zipf Distribution

In the previous section, we proved that the price of anarchy in the general market sharing game is at most 2. This provides an upper bound on the ratio and is a worst case analysis. However, we can prove better bounds for special cases. One important special case is when query rates q_i follow the power law (Zipf) distribution, namely, $q_i = 1/i^\alpha$ for a parameter $0 < \alpha \leq 1$. These distributions are important as it has been observed that in many practical situations, demand curves follow these distributions [4].

In an instance of the uniform market sharing game, let r be the number of markets that are served in an optimal solution. In this section, we prove that in a uniform market sharing game, if all players are interested in all markets, the price of anarchy is less than $1.45 + o(1)$ in the worst case, where $o(1)$ depends on r , i.e., $o(1)$ tends to 0 as $r \rightarrow \infty$. Furthermore, for cases in which the given bipartite graph is not complete or markets have different costs, we prove that the factor 2 is tight.

Theorem 5.6: In the uniform market sharing game, if the given bipartite graph is complete and query rates are from a Zipf distribution, then the following.

- The price of anarchy is less than or equal to $(1/(1 - \alpha)^{1-\alpha}) + o(1)$ for any $\alpha < 1$. In particular, it is less than $e^{1/e} + o(1) < 1.45 + o(1)$ for any $\alpha < 1$ and it tends to $1 + o(1)$ when $\alpha \rightarrow 1$.
- For $\alpha = 1$, the price of anarchy is $(1 + (\ln \ln(r)/\ln(r))(1 + o(1)))$.

Proof: Suppose the given bipartite graph is complete. Consider a Nash equilibrium and let p be the least index such that the players do not select p but select all markets 1 to $p - 1$. No market beyond p can be selected by any player, since otherwise such a player would have an incentive to switch to market p , thus $q_i/n_i \geq q_p$ or $1/n_i i^\alpha \geq 1/p^\alpha$ for $1 \leq i \leq p - 1$. Summing over all markets $1 \leq i \leq p - 1$, we get $\sum_{i=1}^{p-1} (1/i^\alpha) \geq (\sum_{i=1}^{p-1} n_i)(1/p^\alpha)$. Letting $Q_\alpha(k) = \sum_{i=1}^k (1/i^\alpha)$, we get $p^\alpha Q_\alpha(p-1) \geq \sum_{i=1}^{p-1} n_i \geq r$ since \mathbf{OPT} can at most serve $\sum_{i=1}^{p-1} n_i$ markets as our market sharing game is uniform. In the optimal solution, markets 1, 2, ..., and r are served. Thus, we need to bound the price of anarchy which is equal to $Q_\alpha(r)/Q_\alpha(p-1)$.

We consider the two cases $\alpha = 1$ and $\alpha < 1$ separately. We start with $\alpha = 1$. We need to compute $Q_1(r)/Q_1(p-1)$. Using the inequalities $\ln(q) \leq Q_1(q) \leq \ln(q) + 1$, we get $p \geq r/(\ln(r) + 1)$. Therefore

$$\begin{aligned} \frac{Q_1(r)}{Q_1(p-1)} &\leq \frac{\ln(r) + 1}{\ln\left(\frac{r - \ln(r) - 1}{\ln(r) + 1}\right)} \\ &\leq \frac{\ln(r) + 1}{\ln(r - \ln(r) - 1) - \ln(\ln(r) + 1)} \\ &= 1 + \frac{\ln \ln(r)}{\ln r} + o\left(\frac{\ln \ln(r)}{\ln r}\right) \end{aligned}$$

where the last step can be proved using L'Hôpital's rule when $r \rightarrow \infty$.

We now consider the case $\alpha < 1$. Let $L_\alpha(k) = (1/(1 - \alpha))(k^{1-\alpha} - 1)$, then it is easy to see $L_\alpha(k) \leq Q_\alpha(k) \leq L_\alpha(k) + 1$. Using this fact, we can bound the ratio $\lim_{r \rightarrow \infty} (Q_\alpha(r)/Q_\alpha(p-1))$. Observing the facts that $p \rightarrow \infty$ and $L_\alpha(p-1) \rightarrow \infty$ and α is a positive constant less than 1, we can compute the bound as follows:

$$\begin{aligned} \lim_{r \rightarrow \infty} \frac{Q_\alpha(r)}{Q_\alpha(p-1)} &\leq \lim_{r \rightarrow \infty} \frac{L_\alpha(r) + 1}{L_\alpha(p-1)} \\ &= \lim_{r \rightarrow \infty} \frac{r^{1-\alpha} - 1}{(p-1)^{1-\alpha} - 1} \\ &\leq \lim_{r \rightarrow \infty} \frac{(p^\alpha Q_\alpha(p-1))^{1-\alpha}}{(p-1)^{1-\alpha} - 1} \\ &\leq \lim_{p \rightarrow \infty} \frac{(p^\alpha (L_\alpha(p-1) + 1))^{1-\alpha}}{(p-1)^{1-\alpha} - 1} \\ &\leq \lim_{p \rightarrow \infty} \frac{p^{\alpha(1-\alpha)} \frac{1}{(1-\alpha)^{1-\alpha}} (p^{1-\alpha} - 1)^{(1-\alpha)}}{(p-1)^{1-\alpha} - 1} \\ &= \lim_{p \rightarrow \infty} \frac{p^{(1-\alpha)\alpha} \frac{1}{(1-\alpha)^{1-\alpha}}}{((p-1)^{1-\alpha} - 1)^\alpha} \\ &= \lim_{p \rightarrow \infty} \frac{1}{(1-\alpha)^{1-\alpha}} \left(\frac{p^{1-\alpha}}{(p-1)^{1-\alpha} - 1} \right)^\alpha \\ &= \frac{1}{(1-\alpha)^{1-\alpha}}. \end{aligned}$$

Now, one can observe that this bound is less than $e^{1/e}$ for any $\alpha < 1$ and is equal to 1 as α tends to 1. ■

Theorem 5.7: The factor 2 for the price of anarchy of uniform market sharing game is tight even for Zipf distributions if the given bipartite graph is not complete. Moreover, this bound is tight for a general market sharing game even if the given bipartite graph is complete.

Proof: We give an example with $q_i = 1/i$. There are $n^2 + n$ markets and n^2 players. Players are partitioned in n groups of size n . Players in group k are interested in markets $k, kn + 1, kn + 2, \dots, kn + n$. All budgets and costs are equal, i.e., for all $1 \leq i \leq n^2 + n$ and $1 \leq j \leq n^2$, $B_j = C_i = c$. Now, it is not hard to see that there exists a Nash equilibrium in which all players of group k provide market k . Thus, there exists a Nash equilibrium with the social value $H_n = 1 + (1/2) + \dots + (1/n)$. However, if we assign a new market to each player, all markets are provided except n of them. The value of this assignment is $H_{n^2+n} - \sum_{i=1}^n (1/(in + n)) = H_{n^2+n} - ((H_{n+1} - 1)/n)$. Thus, the ratio is $(H_{n^2+n}/H_n) - ((H_{n+1} - 1)/nH_n)$ which is equal to $\ln(n^2 + n)/\ln(n) = 2$ as $n \rightarrow \infty$.

The proof that this bound is tight for general market sharing games is based on an example similar to the above one. The given bipartite graph is a complete bipartite graph. The only difference is that the budgets of players in group k are $n - k$. The cost of markets $k, kn + 1, kn + 2, \dots, kn + n - 1$ is also $n - k$ for $1 \leq k \leq n$. The cost of the market $kn + n$ is large. The query rate of market i is $q_i = 1/i$. Similar to the previous example, if all players of group k provide market k , no player has incentive to change his strategy. It follows that the price of anarchy for this pure Nash equilibrium is arbitrarily close to 2. ■

VI. FINDING A NASH EQUILIBRIUM

Nash proved that any strategic game has a mixed Nash equilibrium [17]. However, there are strategic games that have no pure strategy Nash equilibrium. Furthermore, the complexity of finding a (mixed) Nash equilibrium for strategic games is still an open problem.

In the market sharing game, if we have only one player, finding a Nash equilibrium corresponds to solving optimally a knapsack problem. Thus, the problem of finding a Nash equilibrium in this market sharing game is NP-hard.

However, a Nash equilibrium always exists. An existence proof comes from a general framework defined by Rosenthal [21] for congestion games, as our market sharing game is a special case of such games. For completeness, we sketch the proof here.

Theorem 6.1: [21] A pure strategy Nash equilibrium always exists for the market sharing game.

Proof: The proof is based on defining a potential function and proving that it decreases as agents improve their payoffs by changing their strategies. The potential function for a strategy profile \mathcal{S} is $F(\mathcal{S}) = \sum_{i=1}^m \sum_{t=1}^{n_i} (q_i/t)$, where n_i is the number of player playing market i in \mathcal{S} . Consider player i who changes her strategy from a set S to S' . Let $A = S - S'$ and $B = S' - S$. Let F be the value of the potential function when i plays S and F' be the value when i plays S' . It is easy to see that $F' - F = \sum_{i \in B} (q_i/(n_i + 1)) - \sum_{i \in A} (q_i/n_i)$. This value is exactly the increase in the payoff of i . Assuming the fact that i improves her payoff, we know $F' > F$. This shows that the pure strategy profile that maximizes this potential function is a Nash equilibrium, since if some player was able to increase her own payoff, she would improve the potential function by the same amount. ■

In the rest of this section, we give a polynomial-time algorithm for finding a pure strategy Nash equilibrium for the *uniform* market sharing game (and this will at the same time prove the existence of such a Nash equilibrium). One main feature of the uniform variant is that it is easy for player j to determine its optimal strategy, given the set of strategies for other players. Indeed, player j only needs to solve an easy maximization problem corresponding to selecting the k_j most rewarding markets, where $k_j = \lfloor B_j/C \rfloor$. We could, therefore, let players repeatedly and optimally improve their strategy, but the main issue is to show that such a process converges to a Nash equilibrium in polynomial time. In fact, we will not analyze this algorithm. Instead, we analyze an iterative algorithm in which each agent is restricted to a set of changes at each step. This proves that if players change according to these restrictions, they will converge to a Nash equilibrium in polynomially many steps. It would imply that a Nash equilibrium can be found in polynomial time. An iterative algorithm to find a Nash equilibrium navigates the *state graph* defined below.

Definition 6.2: The *state graph* $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ is a directed graph. Each vertex in \mathcal{V} corresponds to a strategy profile. There is an arc from vertex \mathcal{S} to vertex \mathcal{S}' with label j if the only difference between \mathcal{S} and \mathcal{S}' is the action of agent j and the payoff of player j in \mathcal{S} is strictly less than her payoff in \mathcal{S}' , i.e., player j plays an improvement move from \mathcal{S} to \mathcal{S}' .

Any vertex in the state graph without any outgoing arc (a sink) corresponds to a Nash equilibrium. Thus, in order to prove that a Nash equilibrium exists, we can show the existence of a sink in the state graph. In the following theorem, we give a polynomial-time algorithm to find such a vertex in the state graph.

Theorem 6.3: For the uniform market sharing game, a pure strategy Nash equilibrium always exists and can be found in polynomial time. Furthermore, it can be obtained by traversing a path of length at most m^2n in the state graph.

Proof: Our algorithm for traversing the state graph and finding a sink proceeds in rounds. The first round starts at the vertex \emptyset corresponding to the set of empty actions. In each round, the first arc traversed corresponds to a player, say j , switching from S_j to S'_j , where $S'_j = S_j \cup \{i\}$. In other words, player j only adds precisely one market to its strategy. We refer to this first arc as an *add* arc. After this first arc, subsequent arcs in a round are *change* arcs. These correspond to a player, say j , replacing S_j by $S_j \cup \{i\} \setminus \{k\}$, where $i \notin S_j$ and $k \in S_j$; player j exchanges market k for market i . Furthermore, given j and k , i is selected among all possible markets i of interest to j and not currently in S_j in order to maximize j 's payoff. A round finishes when there are no *change* arcs out of the current vertex. Subsequent rounds start at the vertex where the previous round finishes, unless this vertex has no *add* arc outgoing from it in which case this is the last round.

First, observe that when the last round finishes, the current vertex has no *add* or *change* arcs, and therefore must be a pure Nash equilibrium. This implicitly uses the fact that we are dealing with a *uniform* market sharing game, and therefore any maximal strategy for player j can be obtained from any other maximal strategy by exchanging in and out two markets at a time. Furthermore, at the end of each round, the current state has no *change* arcs outgoing from it, which implies that it corresponds to a pure Nash equilibrium if we suitably modify the budgets of each player (so that they cannot add markets).

As one player adds a market at the beginning of each round, the number of rounds cannot be greater than mn . We now show that each round ends, and in fact ends after traversing at most $m - 1$ *change* arcs (and one *add* arc).

Let us focus on one round and let n_i be the number of players servicing market i at the *beginning* of a round. For simplicity, we assume that the markets are sorted in such a way that $q_1/(n_1 + 1) \geq q_2/(n_2 + 1) \geq \dots \geq q_m/(n_m + 1)$. Consider any vertex in the round after the first *add* arc. Let S_j be the markets currently served by j and let T_j be the markets of interest to j (i.e., i with $(j, i) \in E$) not in S_j . For player j , let $s(j)$ denote $\min\{i \in T_j\}$. We show by induction that the following properties hold throughout the round.

- 1) For any player j , $s(j)$ does not decrease during the round.
- 2) Every market i is covered by n_i players, except one, denoted by p which is covered by $n_p + 1$.
- 3) For any player j , any market $i \in S_j$ and any market $k \in T_j$, we have $q_i/n_i \geq q_k/(n_k + 1)$.

Properties 1 and 2 are obviously true after the first *add* arc corresponding to player, say l , adding market p . Property 3 is also true after the first *add* arc. Indeed, the condition reduces to the fact that the last round ended in a pure Nash equilibrium except

for the case where $j = l$ and $i = p$, where it follows from the choice of $p : q_p/n_p > q_p/(n_p + 1) \geq q_k/(n_{k+1})$.

We see now what happens when we traverse a *change* arc corresponding to player l exchanging two markets. Condition 3 implies that l leaves market p since all other markets do not increase l 's payoff. Thus, Property 2 is maintained after the change (with a different value for p). Second, l will now serve market $s(l)$ by definition of $s(l)$. This implies that $q_p/(n_p+1) < q_{s(l)}/(n_{s(l)} + 1)$, i.e., $s(l) < p$. This means that Property 1 is also maintained. To verify that Property 3 is still maintained, we only need to consider the cases in which $j = l$ and either $k = p$ or $i = s(l)$. If $i = s(l)$, Property 3 follows from the definition of $s(l) : q_{s(l)}/n_{s(l)} > q_{s(l)}/(n_{s(l)} + 1) \geq q_k/(n_{k+1})$. If $k = p$, it follows from $q_i/n_i \geq q_{s(l)}/(n_{s(l)} + 1) > q_p/(n_p + 1)$.

All three properties are maintained during the round. Furthermore, since player l replaces market p by market $s(l)$ and $s(l) < p$, we have that p decreases as we traverse *change* arcs. This implies that we have at most $m - 1$ *change* arcs in a round. This proves our bound of nm^2 on the length of the traversal we construct before reaching a pure Nash equilibrium. ■

In order to run the algorithm, we actually do not need to construct the entire state graph. We only need to be able to find the next arc to traverse, and this can be done in $O(m + n)$ time, resulting in a total running time of $O((m + n)m^2n)$.

VII. GREEDY BEHAVIOR ANALYSIS

Although we have proved that the price of anarchy for a Nash equilibrium is upper bounded by 2, there are still two main issues that are left unaddressed. First, converging to an exact Nash equilibrium in a way described in Theorem 6.1 requires agents to solve optimally an NP-hard knapsack problem in order to find their best response. Considering the computational constraints of rational agents, we consider agents that choose their action based on an approximation algorithm for knapsack. The second issue is the speed with which the agents can converge to a Nash equilibrium or an approximate Nash equilibrium. We consider both issues in this section.

In the proof of Theorem 6.1, we see that if agents improve their payoff, they will converge to a Nash equilibrium after finitely many steps (since the potential function can take only finitely many values and strictly increases with every improvement). This result can be generalized to the setting in which each agent chooses its action using a β -approximation algorithm for knapsack (keeping the actions of all the other players unchanged). A β -approximation algorithm is a polynomial-time algorithm guaranteed to return a feasible solution of value (i.e., payoff) at least $1/\beta$ times the optimum.

Theorem 7.1: If all players rely on a β -approximation algorithm to improve their actions, they will converge to a β -approximate Nash equilibrium in finitely many steps.

From Corollary 5.5, we know that such a β -approximate Nash equilibrium gives a social function value within a factor $1/(\beta + 1)$ of the optimal social value.

The proof of the above theorem is similar to that of Theorem 6.1. Using the same potential function, one observes that, after

finitely many steps, no agent will be able to improve its action, and therefore, we will have reached a β -approximate Nash equilibrium.

The above theorem, as well as Theorem 6.1, do not bound the number of steps needed for players to converge to an equilibrium. In the remainder of this section, we analyze a greedy behavior of the players, *greedy upon arrival*, in which players enter the game one by one and select their first action by using a β -approximation algorithm for knapsack for some constant β . We show that after the players have selected their first action, i.e., after each player makes just one decision, the resulting strategy profile has a social function value within a logarithmic factor of the optimal social value.

Theorem 7.2: If each player chooses its first action using a β -approximation algorithm then the social function of the resulting strategy profile is at least $1/(\beta + 1)H_n$ of the optimal social value, where $H_n = 1 + (1/2) + \dots + (1/n) \leq \ln(n) + 1$. Moreover, this logarithmic bound is tight, up to a constant factor.

Proof: For the purpose of the analysis, we consider the following restricted multiple knapsack problem (RMKP) corresponding to an instance of our market sharing game. In an instance of the RMKP problem, there are m groups of n items each for a total of mn items. The n items in group i all have size C_i , for $1 \leq i \leq m$. The values of items in group i are $(q_i, q_i/2, q_i/3, \dots, q_i/n)$. There are n bins with capacities (B_1, B_2, \dots, B_n) . Bin j can get an item in group i iff agent j is interested in market i ($(j, i) \in E(G)$). We are only allowed to place at most one item of each group in a bin, and the goal is to place (some of the) items in bins so that the size assigned to every bin is at most its capacity and the total value of items assigned is maximized. Given the objective function, we can focus our attention on algorithms for which an item of group i is not assigned to any bin if a higher valued item of the same group is not assigned to any bin. We prove our claims using the following lemma.

Lemma 7.3: Consider the following greedy algorithm for RMKP problem: arbitrarily order bins, pack bins one at a time using a β -approximation algorithm for the single knapsack problem instance for the current bin (considering only the highest valued item remaining for each group) and discard the packed items. Then, the resulting assignment is a $\beta + 1$ -approximation for RMKP.

Proof: Let \mathbf{OPT}' be an optimal solution of RMKP. Let A_j be the set of items in bin j in \mathbf{OPT}' which do not appear in the greedy solution in any bin, and X_j their total value. Let D_j be the set of items in bin j in the greedy solution, and Y_j their total value. Observe that if A_j and D_j each contain a different item of group i then the value of the item, say p , in A_j is smaller than the value of the item, say q in D_j (since p does not appear in the greedy solution in any bin). Since we use a β -approximation algorithm for filling each bin, $\beta Y_j \geq X_j$. Thus, $\beta \sum_{1 \leq j \leq n} Y_j \geq \sum_{1 \leq j \leq n} X_j \geq \gamma(\mathbf{OPT}') - \sum_{1 \leq j \leq n} Y_j$. Thus, $(\beta + 1) \sum_{1 \leq j \leq n} Y_j \geq \gamma(\mathbf{OPT}')$ as desired. ■

We are ready to continue the proof of Theorem 7.2. If n_i is the current number of agents servicing market i , the current value of this market is $q_i/(n_i + 1)$ for a new agent. Equivalently, we can consider a new market of value $q_i/(n_i + 1)$. In this way, we add different copies of the same market with

different query rates and the same size. The only restriction is that no player can provide two markets of the same group. Now, one can observe that after one step of greedy behavior by each player, the resulting strategy profile is equivalent to the assignment of the greedy algorithm from Lemma 7.3 for RMKP. Let the resulting assignment be \mathcal{G} . Now, we prove that the sum of Y_j 's over all players is at most H_n times the sum of payoffs of players in \mathcal{G} (where Y_j is defined in Lemma 7.3). In $\sum_{j \in U} Y_j$, the sum of terms corresponding to markets of group i is at most $q_i + (q_i/2) + \dots + (q_i/n) = q_i * H_n$, while market i contributes to the sum of payoffs in \mathcal{G} as q_i . Thus, $\sum_{j \in U} Y_j \leq H_n \sum_{j \in U} P_j(\mathcal{G})$.

In an instance of RMKP, there are extra markets. Therefore, the optimal solution of RMKP, $\gamma(\mathbf{OPT}')$, is greater than or equal to the optimal social value, $\gamma(\mathbf{OPT})$. Thus

$$(1+\beta)H_n \sum_{j \in U} P_j(\mathcal{G}) \geq (1+\beta) \sum_{j \in U} Y_j \geq \gamma(\mathbf{OPT}') \geq \gamma(\mathbf{OPT})$$

as desired.

Furthermore, we show that this bound is tight up to a constant factor. Consider an instance of the market sharing game with n markets and n players: Let $q_i = (n/i) - \epsilon$ for all $1 \leq i \leq n$ where ϵ is sufficiently small. Assume that player j is interested in market 1 and markets $j, j+1, \dots, n$ and all costs and budgets are the same and are equal to 1. If we start from all empty strategies and players choose their best strategy in the order $1, 2, \dots, n$, all players provide market number 1 and the social value of this assignment is $n - \epsilon$. However, in the optimal assignment, agent j services market j . As a result, the value of the optimal solution is $nH_n - n\epsilon$. Thus, the ratio between the optimum and the value of the resulting assignment after one step for each player is arbitrarily close to $H_n = \Omega(\log n)$. ■

The following theorem shows that in market sharing games, if players improve their own payoffs, the social value does not decrease by a factor more than H_n .

Theorem 7.4: Let Q be a state of social value T . If Q' is a state resulting from the state Q after a sequence of improvement moves by players, then the social value of Q' is at least T/H_n . In particular, if we start from the state with the optimal social value, the social value of any state resulting from some improvement moves by players is at least \mathbf{OPT}/H_n .

Proof: We use the potential function used in the proof of Theorem 6.1. Let n_i and n'_i be the number of players that have market i in their strategy set in states Q and Q' , respectively. Let the potential function for states Q and Q' be $F(Q) = \sum_{i=1}^m \sum_{t=1}^{n_i} (q_i/n_i)$ and $F(Q') = \sum_{i=1}^m \sum_{t=1}^{n'_i} (q_i/n'_i)$, respectively. Since Q' is the result of a sequence of improvement moves by players starting from the state Q , similar to the proof of Theorem 6.1, one can show that $F(Q') > F(Q)$. Let the social value of states Q and Q' be $\gamma(Q)$ and $\gamma(Q')$, respectively. It is easy to verify that $F(Q) \geq \gamma(Q)$ and $F(Q') \leq H_n \gamma(Q')$. As a result, we have $\gamma(Q') \geq F(Q')/H_n > F(Q)/H_n \geq \gamma(Q)/H_n = T/H_n$. ■

VIII. EVALUATION

In this section, we investigate the efficiency of our decentralized caching scheme in a simulated network scenario. We are interested in obtaining the price of anarchy and determining

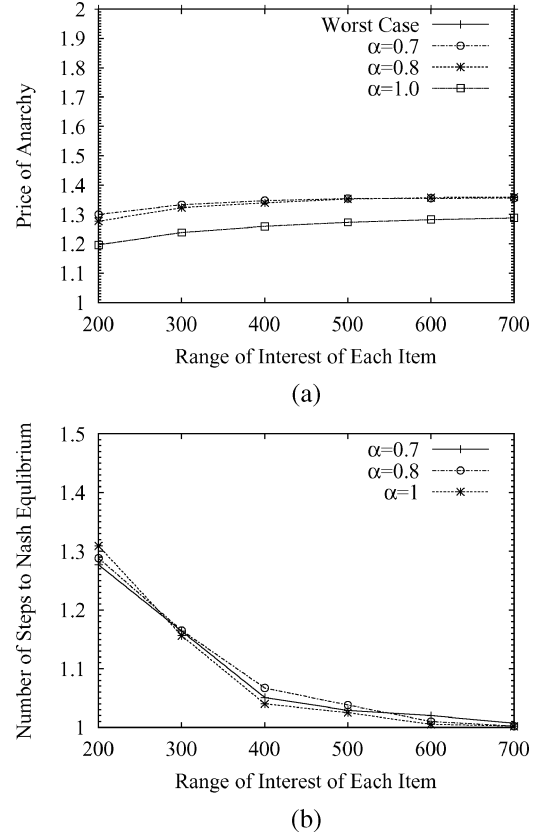


Fig. 2. Price of anarchy and number of steps to convergence to a Nash equilibrium for uniform game. (a) Price of anarchy. (b) Number of steps to convergence to a Nash equilibrium.

how quickly the game converges to an exact or approximate pure strategy Nash equilibrium. We investigate the advantages of the greedy behavior in terms of the payoff function, as well as evaluate how far it is from the optimum.

We generated 10 random networks, each with 100 nodes (resident subscribers). These nodes are randomly placed in a 800×800 rectangular region. Each node has an 802.11b interface and has a maximum transmission radius of 115 units. Each item has a radius of interest R_I , i.e., only nodes within that radius of interest benefit from caching the item. We vary this radius from 200 to 700. In our simulation, the effect of the transit subscribers is captured by the query rate. We assume a query can originate from anywhere in the network within the item's radius of interest. We assume that we have 1000 items whose popularity follows a Zipf distribution. We vary the α parameter of the Zipf distribution. The items' size is 1 unit for the uniform market sharing game and for the nonuniform market sharing game it follows a lognormal distribution with mean 7.5 and standard deviation 1.5. Each resident subscriber can cache at most 5 items in the uniform case, 20 units for the nonuniform case. We focus our study on issues related to the impact of selfish behavior on network performance, and therefore, do not perform packet level simulations. All our results are averaged over ten random networks.

A. Uniform Market Sharing Game

We would like to investigate the price of anarchy and how quickly the players converge to a Nash equilibrium. Fig. 2(a)

TABLE I
UNIFORM GAME: PERCENTAGE OF NODES WITH THE SAME NUMBER OF STEPS

	Number of Steps			
	1	2	3	4
percentage of Nodes	78%	20%	1%	1%

TABLE II
UNIFORM GAME: PERCENTAGE OF NODES WITH THE SAME NUMBER OF CACHE REPLACEMENTS. RANGE OF INTEREST $R_I = 200$


	Number of Cache Replacements						
	0	1	2	3	4	5	6
percentage of Nodes	78%	0%	19%	0%	2%	0%	1%

shows the price of anarchy with respect to the item’s range of interest R_I . The larger R_I is, the more edges in the bipartite graph between items and resident subscribers. For Zipf distribution with $\alpha = 0.7$, the worst case price of anarchy from our analysis is 1.45 and 2 for complete and incomplete bipartite graphs, respectively. The price of anarchy ρ_a from our simulation is between 1.30 and 1.36. Therefore, ρ_a is within 90% to 95% of 1.45 and 65% to 68% of 2. Therefore, the price of anarchy is far from the worst case for the incomplete bipartite graph case and close to the worst case for the complete bipartite graph. Note that the price of anarchy function $\rho_w = 1/(1 - \alpha)^{(1-\alpha)}$ increases in interval $[0, 0.63]$ and decreases in $[0.63, 1]$. However, according to [4], in the case of popular items, α is between 0.65 and 0.8. In this particular range of α , we observed that the price of anarchy ρ_a increases only slightly as R_I increases. This is because popular items are within the range of more players and as more players cache common items, ρ_a increases. To see the impact of α on the price of anarchy, we ran our simulation with $\alpha = 0.8$ and $\alpha = 1$. The observations are similar to the case of $\alpha = 0.7$.


Fig. 2(b) shows the average number of steps required of each player to converge to the pure strategy Nash equilibrium. As we can see, each player takes an average of 1.3 steps or less. Recall that, by a step, we mean that the player changes to a different set of items to cache in response to the action taken by other players. Table I shows for a given network setting the percentage of nodes requiring the same number of steps after an equilibrium is attained. We see that, 78% of the players never change the set of items to cache. 20% of them change once. Only 1% change twice and three times. Table II shows the percentage of nodes with the same number of cache replacements. We see that 78% of the players stay with the original set of items. Only 19% replaced two items. 2% changed 4 items and 1% changed 6 items in the cache.

B. General Market Sharing Game

For the general market sharing game, as finding the best response for a player is NP-hard, we choose to investigate a 2-approximate Nash equilibrium based on a simple 2-approximate greedy algorithm. The algorithm [25] works as follows. Order the items by payoff over size. Let the sorted order of objects be a_1, a_2, \dots, a_m . Find the lowest k such that the size of the first k objects exceeds the cache space. Pick the more profitable of $\{a_1, a_2, \dots, a_{k-1}\}$ or $\{a_k\}$. We refer to the sum of the payoffs of the items cached when the players reach this



(a)



(b)

Fig. 3. The approximate social versus Nash payoff and number of steps to converge to a 2-approximate Nash equilibrium for general market sharing game. (a) Approximate social versus Nash payoff. (b) Number of steps to converge to a 2-approximate Nash equilibrium.

TABLE III
GENERAL MARKET SHARING GAME: PERCENTAGE OF NODES WITH THE SAME NUMBER OF STEPS

	Number of Steps	
	1	2
percentage of Nodes	83%	17%

2-approximate Nash equilibrium as 2-approximate Nash payoff and denote it as γ_{N_2} .

Computing the optimal social value involves solving a multiple knapsack problem. As this problem is also NP-hard, we choose a simple 3-approximation algorithm [6]. The algorithm greedily applies the 2-approximate knapsack algorithm for each knapsack. We refer to sum of the payoffs of the cached items computed by this algorithm as the 3-approximate social payoff and denote it as γ_{S_3} .

The performance of the approximate Nash payoff with respect to the approximate social payoff is given by the ratio $\gamma_{S_3}/\gamma_{N_2}$. Fig. 3(a) shows that this ratio is about 1.3 in our simulation. Fig. 3 also shows the number of steps it takes to converge to the 2-approximate Nash equilibrium as the range of interest R_I increases. Consider the specific networking setting of $R_I = 200$. Table III shows that 83% of the players only take one step, while the remaining 17% take 2 steps to reach the approximate Nash equilibrium. Furthermore from Table IV, we see that, 83% of players stayed with the initial set of items,

TABLE IV
GENERAL MARKET SHARING GAME: PERCENTAGE OF NODES
WITH THE SAME NUMBER OF CACHE REPLACEMENTS

	Number of Cache Replacements			
	0	1	2	3
percentage of Nodes	83%	7%	9%	1%

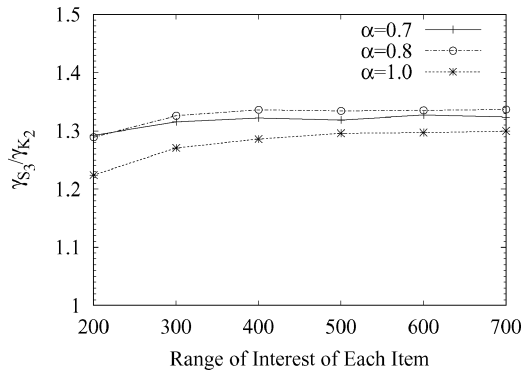


Fig. 4. Compare greedy payoff with approximate social payoff.

7% replaced just one item, 9% replaced two, and 1% replaced three items.

C. One Step of Greedy Behavior

We would also like to see the payoff for one step of greedy behavior. Each player, when it decides to participate, caches items according to the 2-approximate knapsack algorithm. Fig. 4 compares the corresponding one step greedy payoff γ_{K_2} with the 3-approximate social payoff γ_{S_3} . We see that the behavior of this payoff ratio is similar to the ratio $\gamma_{S_3}/\gamma_{N_2}$ obtained in the previous section. In fact, the one step social payoff γ_{K_2} is only at most 2% worse than the 2-approximate Nash payoff.

IX. RELATED WORK

Caching has been widely studied in the context of wired networks [27]. Caching in wireless networks has recently become a popular research area. Sailhan and Issarny [23] propose protocols to enable cooperative caching in ad hoc networks. They also study cache management strategies that aim to reduce energy consumption and network load. Xiang *et al.* [29] propose a cooperative cache management scheme in cellular networks. They present a simulation study of cache replication strategies in BSs for streaming services. Nuggehalli *et al.* [19] present an energy-efficient cache placement scheme. Our work studies caching in a noncooperative environment. The problem of noncooperative caching is fundamentally different from its cooperative counterpart.

The problem of how to stimulate cooperation among selfish nodes in ad hoc networks and multihop cellular networks has received significant attention recently. Zhong *et al.* [30] propose a credit-based system which relies on a central authority to collect receipts from forwarding nodes. Charges and rewards are based on the receipts. They assume the availability of public and private keys to compute message signatures as receipts. We cannot directly apply their scheme since we use symmetric keys rather than asymmetric keys in our integrated 3G and ad hoc networks. Salem *et al.* [24] propose a charging and rewarding

scheme to make collaboration rational for selfish nodes. Their solution is based on symmetric cryptography. However, they require all data packets to go through the BS. Since our goal is to offload the data distribution from the BS as much as possible, their scheme cannot be applied directly in our context. Jakobsson *et al.* [12] propose a micropayment scheme for multihop cellular networks that encourages collaboration in packet forwarding. The sender attaches a payment for each packet. In our context, the sender gets paid by the receiver through the BS. In addition, it is natural to charge per session in our context. Therefore, their scheme does not apply directly.

Reputation-based schemes such as [5] are not applicable in our context since it cannot deal with certain collusion scenarios, e.g., collusion between the sender and receiver where forwarding nodes do not get remunerated.

Recently, game theory has found applications in wireless networks [2], [7], [13]. Andereg and Eidenbenz [2] applied the Vickrey–Clark–Grove (VCG) mechanism from mechanism design theory to design a routing protocol. However, they have overlooked the issue of “mutually dependent types” which resulted in their scheme not being truthful. Zhong *et al.* [31] model routing and forwarding in ad hoc networks as one “ad hoc” game. They propose the first routing and forwarding protocols in ad hoc networks that is cooperation-optimal. That is, following the protocol is a dominant-subaction in the routing stage, and is a subgame-perfect Nash equilibrium in the forwarding stage. Wang *et al.* [28] propose the first multicast routing protocol without using VCG mechanism such that each agent maximizes its profit when it truthfully reports its cost. Eidenbenz *et al.* [7] proposed a topology control game that models user’s selfish behavior in forming ad hoc network topology. Lin *et al.* [13] used game theory to propose an admission and rate control framework for code-division multiple-access (CDMA) data networks. To the best of our knowledge, we are the first to propose a game-theoretic study of caching in wireless networks.

The issues of the efficiency of computing Nash equilibria and considering computational constraints of the central authority and selfish agents, and the outcome of the selfish behavior of these agents have been investigated recently in computer science [18], [20]. The market sharing game is a special case of congestion games introduced by Rosenthal [21]. In these games, each player chooses a particular combination of factors out of a common set of factors. The payoff from each factor is a function of the number of players who have chosen this factor in their set. Milchataich [15] studied more general settings and the length of best response paths in this set of games. Congestion games are in a general class of potential games [16]. None of these works considered the efficiency of the outcome of the selfish behavior of agents or the speed of convergence to an approximate solution or the computational constraints of selfish agents. We address all of the above issues for the special case of market sharing games. We consider a social function and the outcome of selfish behavior of players in terms of this social function. Comparing the value of the social function obtained from the outcome of selfish behavior with the optimal social value is done via the notion of the price of anarchy. For example, for a variant of nonatomic congestion games, Roughgarden and Tardos [22] give some upper bounds on the price of anarchy.

After this work, there have been some new results on congestion games. Fabrikant *et al.* [8] have studied the convergence to equilibria in (network) congestion games with delay functions. Goemans *et al.* [11] have studied the convergence to approximate solutions in (network) congestion games with polynomial delay functions. Awerbuch *et al.* [3] and Koutsoupias and Christodoulou have studied the price of anarchy in (network) congestion games with polynomial delay functions. All of these papers consider social and payoff functions that are different from market sharing games. In particular, in network congestion games, each player is minimizing his delay instead of maximizing his profit. Finally, Fleischer *et al.* [9] have developed approximation algorithms and mechanisms for the distributed caching problem. The distributed caching problem is a generalization of the (geometric) market sharing games.

X. DISCUSSION AND FUTURE RESEARCH

There are two main issues discussed in this paper. One is the architecture for enabling mobile content distribution in ad hoc networks, and the other one involves the game theoretic formulation of the noncooperative caching problem. In this section, we discuss issues in modeling the problem as a strategic game and point out some future research directions.

A. Implementation Aspects of the Caching Game

The architectural framework for mobile content distribution described in this paper assumes that each player can determine the amount of requests it will handle if it caches an item. In reality, the player may not know this information. However, this information can be estimated by observing the cache request and reply messages of our protocol. We also assumed that only one player changes the set of cached items at a given time. Once the change is made, all other relevant players are assumed to become aware of the changed payoff. In reality, multiple players may make decisions at the same time. This issue can be addressed by using a random timer to decide when to change thus preventing synchronization effects. Using a timer would allow the effect of the change to stabilize before players respond to the new change. If the estimated payoff changes continuously, the player will have no incentive to change since replacing a cached item will incur a cost. The cost is incurred since the player has to either request the changed item from the 3G service provider or from some other players who have cached the item. We assume the cost of changing cached items are amortized. Therefore, players have an incentive to change if the payoff of an item is stable and it increases his utility.

B. Refined Models of the Caching Game

In market sharing games, we assume that the query of markets are uniformly distributed. This roughly models certain routing protocols where the first returned route is selected. In this case, a client wants to optimize its response time. In some networking setting, a client may prefer the closest player. There are many incentives to choose the closest one, e.g., routes can be more stable if the client is very mobile. In terms of implementation, the routing protocol can wait longer and select the closest player

from the set of returned routes. Here, we define a variant of this game called *geometric market sharing game*.

Definition 10.1: In the geometric market sharing game, there is a set U of n agents or players and a set H of m markets and a set of k market types. Market $i \in H$ is of type t_i , where $1 \leq t_i \leq k$ and has query rate q_i . Each market type $1 \leq t \leq k$ has a cost C_t corresponding to the cost for servicing this type of market. For each agent j and market i , we are given a distance d_{ij} . Each agent j has a total budget B_j . Each agent should decide which subset of market types to service. Agent j can service a subset S_j of market types, if the sum of the costs C_t of the market types in S_j is less than or equal to B_j , i.e., $\sum_{t \in S_j} C_t \leq B_j$. Each agent chooses a feasible set of market types to provide. Given the strategies of all agents, each market i will get the service from an agent who services the market type t_i with the minimum distance to market i ; and as a result, the reward of the market i (which is equal to q_i) goes to this agent. Each agent would like to provide a set of market types that maximizes its own payoff. Similar to the market sharing game, the total payoff (reward) received by all players is equal to the total query rate of the markets being serviced (by at least one player).

Consider the geometric market sharing game and the following social function: the sum of payoffs of players. This social function is the same as the total market that is satisfied by players. It is not hard to see that this game satisfies all the properties of the valid-utility games:

Theorem 10.2: The geometric market sharing game is a valid-utility game and thus, the price of anarchy of this game (for mixed Nash equilibria) is at most 2.

Proof: The proof is similar to that of Theorem 5.4. Here, we give a proof sketch. Clearly, the social function is nondecreasing. For submodularity, we observe that adding a new market type to the strategy of a player increases the social function if and only if this market type is not provided by any other player. This shows that if $A \subset B$, the marginal increase of adding a market type to A is greater than or equal to the marginal increase for adding a market type to B . Using this observation, one can prove that the social function is submodular. Moreover, the effect of playing a player j in the social function corresponds to the set of market types that j provides and no other player provides. Player j will get the whole payoff of markets of these market types (that only he provides). Therefore, the payoff a player is at least the difference that he makes in the social function by playing his strategy (compared to playing the empty strategy). Finally, by definition, the social function is equal to the sum of payoffs of players. Hence, the geometric market sharing game satisfies all the properties of valid-utility game. ■

Unlike the market sharing game, we cannot prove existence of pure Nash equilibria and convergence to it for all instances of the geometric market sharing game. In fact, we know examples in which there is no pure Nash equilibrium for this game.

Theorem 10.3: There are instances of the geometric market sharing game in which there is no pure Nash equilibria.

Proof: Consider the geometric market sharing game with two players. The budget of players 1 and 2 are 11 and 8, respectively. There are four market types and one market of each type. The cost of market types are 6, 8, 3, and 4, respectively. The

query rate of markets of these four types are 4, 6, 2, and 3, respectively. Market types 2 and 4 are closer to player 1. Market types 1 and 3 are closer to player 2. It is not hard to check that there is no pure Nash equilibrium in the geometric market sharing game corresponding to this example. In fact, the only best-response strategies for player 1 are (2,3) and (1,4) and the only best-response strategies for player 2 are (3,4) and (2). Any combination of these best responses is not a pure Nash equilibrium, and therefore, there exists no Nash equilibrium in this game. ■

Goemans *et al.* [11] have introduced an equilibrium concept, namely, the sink equilibria, for games without pure Nash equilibria. Bounding the price of anarchy for sink equilibria and studying convergence to approximate solutions in the geometric market sharing games are interesting problems for future research.

XI. CONCLUSION

In this paper, we propose a mobile content distribution architecture to offload popular data items from a 3G network to its subscriber-based ad hoc network. We present a protocol to enable decentralized caching and propose novel incentive mechanisms for proper accounting and crediting. Our incentive mechanisms prevent or discourage players from colluding for selfish gains at the expense of other users.

We study the selfish behavior of the players involved in caching and forwarding by using game theoretic approaches. We formulate the caching problem as a general market sharing game. We observe that pure strategy Nash equilibrium exists in this game and present a polynomial-time algorithm to find it in the case when item sizes are the same. Finding a pure strategy Nash equilibrium for the general game is NP-hard. We model computationally bounded agents by using approximation algorithms and observe that they also converge to an approximate Nash equilibrium. We show that starting from the empty assignment, after one step of the greedy behavior by the players, the social value of the resulting assignment is within an $O(\log n)$ of the optimum. Mirrokni and Vetta [11] have extended this result by proving the same bound starting from any assignment. We show that the price of anarchy is bounded by 2 for any arbitrary popularity distribution of the cached items. For the special case where the popularity of items is according to a Zipf distribution and caching any item has a positive reward for all players, we obtain a bound of 1.45. Our simulations show that for most real and practical instances of the game, the price of anarchy is better than the worst case. These simulations also show that typically most players converge to the equilibrium in one or two steps most of the time.

ACKNOWLEDGMENT

The authors would like to thank Prof. J. Y. Halpern at Cornell University for initial discussions on the problem. They are also very grateful to Dr. J. Gary, Dr. E. Grosse, and Dr. S. Paul at Bell Laboratories, and S. Zhong at Yale University for helpful discussions and comments on the security aspects of the paper.

REFERENCES


- [1] 3G Today. [Online]. Available: <http://www.3gtoday.com/>.
- [2] L. Anderegg and S. Eidenbenz, "Ad hoc-VCG: A truthful and cost-efficient routing protocol for mobile ad hoc networks with selfish agents," in *Proc. MobiCom*, 2003, pp. 245–259.
- [3] B. Awerbuch, Y. Azar, and A. Epstein, "The price of routing unsplittable flow," in *Proc. STOC*, 2005, pp. 57–66.
- [4] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web caching and Zipf-like distributions: Evidence and implications," in *Proc. INFOCOM*, 1999, pp. 126–134.
- [5] S. Buchegger and J.-Y. Le Boudec, "Performance analysis of the CON-FIDANT protocol: Cooperation of nodes-fairness in dynamic ad hoc networks," in *Proc. MobiHoc*, 2002, pp. 80–91.
- [6] C. Chekuri and S. Khanna, "A PTAs for the multiple knapsack problem," in *Proc. ACM-SIAM Sympo. Discrete Algorithms*, 2000, pp. 213–222.
- [7] S. Eidenbenz, V. S. Anil Kumar, and S. Züst, "Equilibria in topology control games for ad hoc networks," in *Proc. Joint Workshop on Foundations of Mobile Computing*, 2003, pp. 2–11.
- [8] A. Fabrikant, C. Papadimitriou, and K. Talwar, "The complexity of pure equilibria," in *Proc. ACM STOC*, 2004, pp. 604–612.
- [9] L. Fleischer, M. Goemans, V. S. Mirrokni, and M. Sviridenko, "Tight approximation algorithms for maximizing general assignment problems," in *SODA*, 2006.
- [10] H. Garay and L. Li, "Simple and efficient security mechanisms for hybrid wireless networks," Bell Labs Tech. Memorandum, 2005.
- [11] M. Goemans, V. Mirrokni, and A. Vetta, "Sink equilibria and convergence," in *Proc. IEEE FOCS*, Washington, DC, 2005, pp. 142–154.
- [12] M. Jakobsson, J.-P. Hubaux, and L. Buttyan, "A micro-payment scheme encouraging collaboration in multihop cellular networks," in *Proc. Financial Cryptography*, 2003, pp. 15–33.
- [13] H. Lin, M. Chatterjee, S. K. Das, and K. Basu, "ARC: an integrated admission and rate control framework for cdma data networks based on noncooperative games," in *Proc. MOBICOM*, 2003, pp. 326–338.
- [14] H. Luo, R. Ramjee, P. Sinha, L. Li, and S. Lu, "UCAN: A unified cellular and ad-hoc network architecture," in *Proc. MOBICOM*, 2003, pp. 353–367.
- [15] I. Milchtaich, "Congestion games with player-specific payoff functions," *Games and Economic Behavior*, vol. 13, pp. 111–124, 1996.
- [16] D. Monderer and L. Shapley, "Potential games," *Games and Economic Behavior*, vol. 14, pp. 124–143, 1996.
- [17] J. F. Nash, "Equilibrium points in N-person games," in *Proc. NAS*, vol. 36, 1950, pp. 48–49.
- [18] N. Nisan and A. Ronen, "Algorithmic mechanism design," in *Proc. ACM STOC*, 1999, pp. 129–140.
- [19] P. Nuggehalli, V. Srinivasan, and C. Chiasserini, "Energy-efficient caching strategies in ad hoc wireless networks," in *Proc. MobiHoc*, 2003, pp. 25–34.
- [20] C. Papadimitriou, "Algorithms, games, and the Internet," in *Proceedings of ACM STOC*, 2001, pp. 749–753.
- [21] R. W. Rosenthal, "A class of games possessing pure-strategy Nash equilibria," *Int. J. Game Theory*, vol. 2, pp. 65–67, 1973.
- [22] T. Roughgarden and E. Tardos, "How bad is selfish routing?," *J. ACM*, vol. 49, no. 2, pp. 236–259, 2002.
- [23] F. Saitlan and V. Issarny, "Cooperative caching in ad hoc networks," in *Proc. 4th Int. Conf. Mobile Data Manage.*, 2003, pp. 13–28.
- [24] N. B. Salem, L. Buttyan, J. P. Hubaux, and M. Jakobsson, "A charging and rewarding scheme for packet forwarding," in *Proc. fMobiHoc*, Annapolis, MD, June 2003, pp. 13–24.
- [25] V. Vazirani, *Approximation Algorithms*. New York: Springer-Verlag, Jun. 1999. ISBN: 3540653678.
- [26] A. Vetta, "Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions," in *Proc. FOCS*, 2002, pp. 416–425.
- [27] J. Wang, "A survey of web caching schemes for the Internet," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 29, no. 5, pp. 36–46, 1999.
- [28] W. Wang, X.-Y. Li, and Y. Wang, "Truthful multicast in selfish wireless networks," in *Proc. MOBICOM*, 2004, pp. 245–259.
- [29] Z. Xiang, Z. Zhong, and Y. Zhong, "A cache cooperation management for wireless multimedia streaming," in *Proc. IEEE Int. Conf. Info-Tech and Info-Net ICI*, 2001, pp. 328–333.
- [30] S. Zhong, J. Chen, and Y. R. Yan, "Sprite: a simple, cheat-proof, credit-based system for mobile ad-hoc networks," in *Proc. INFOCOM*, vol. 3, San Francisco, CA, Apr. 2003, pp. 1987–1997.

- [31] S. Zhong, L. Li, Y. Liu, and Y. Yang, "On designing incentive-compatible routing and forwarding protocols in wireless ad-hoc networks: an integrated approach using game theoretical and cryptographic techniques," in *Proc. MOBICOM*, 2005, pp. 117–131.



Michel X. Goemans is a Professor of Applied Mathematics at the Massachusetts Institute of Technology (MIT), Cambridge, and a faculty member of the MIT Computer Science and Artificial Intelligence Laboratory (CSAIL) and the MIT Operations Research Center (ORC). He has held an Adjunct Professorship at the University of Waterloo and a Professorship at the University of Louvain. He has been an invited speaker at many conferences, including the International Congress of Mathematicians in 1998. He has been on the program committee of several major theoretical computer science conferences, including as chair of the 2003 Symposium on Theory of Computing. His research—in the areas of discrete algorithms and combinatorial optimization—has been rewarded by several prizes.

Dr. Goemans received several prizes including the 2000 AMS-MPS Fulkerson Prize and twice the SIAM Optimization Prize.



Li (Erran) Li received the B.E. degree in automatic control from Beijing Polytechnic University, Beijing, China, in 1993, the M.E. degree in pattern recognition from the Institute of Automation, Chinese Academy of Sciences, Beijing, in 1996, and the Ph.D. degree in computer science from Cornell University, Ithaca, NY, in 2001, where J. Y. Halpern was his advisor.

During his graduate study at Cornell University, he worked at Microsoft Research and Bell Laboratories, Lucent Technologies, as an Intern, and at AT&T Research Center at ICSI Bekerley as a Visiting Student.

He is presently a member of the Networking Research Center at Bell Laboratories. He has published over 30 papers. His research interests are in networking with a focus on wireless networking and mobile computing.

Dr. Li has served as a program committee member for several conferences including ACM MobiCom, ACM MobiHoc, IEEE INFOCOM, and IEEE ICNP. He is a Guest Editor for the *JOURNAL ON SELECTED AREAS IN COMMUNICATIONS* (Special Issue on Noncooperative Behavior in Networking).



Vahab S. Mirrokni received the B.S. degree from Sharif University, Tehran, Iran, in 2001 and the Ph.D. degree from the Massachusetts Institute of Technology (MIT), Cambridge, in 2005, under the supervision of Michel X. Goemans.

He is a Postdoctoral Associate at the Computer Science and Artificial Intelligence Laboratory (CSAIL) and the Mathematics Department at the MIT. He has worked at Bell Laboratories, Microsoft Research, IBM T. J. Watson, and Amazon.com.

He has published more than 25 papers in refereed conferences and journals. His research interests include algorithmic aspects of game theory and economics, approximation algorithms for combinatorial optimization, network optimization, mobile computing, and graph theory.

Dr. Mirrokni is the co-winner of the Best Student Paper Award at the Symposium on Discrete Algorithms (SODA'05).



Marina Thottan (M'99) received the M.S. degree in biomedical engineering and the Ph.D. degree in electrical and computer systems engineering from Rensselaer Polytechnic Institute, Troy, NY.

She is a Member of Technical Staff in the Center for Networking Research, Bell Laboratories, Lucent Technologies, Murray Hill, NJ. She is active in the fields of integrated data over optical networks and network management and has served as a program committee member for several conferences in these areas. Her research publications have appeared in a

number of ACM, IEEE, and other professional conferences and journals. Her research interest is in the application of statistical, combinatorial and game theoretic algorithms to problems in data networking, specifically in network measurements, network design, and QoS.