The Impact of Incomplete Information on Games in Parallel Relay Networks

Hongda Xiao and Edmund M. Yeh

Electrical Engineering, Yale University
{hongda.xiao,edmund.yeh}@yale.edu

Abstract. This work considers the impact of incomplete information on incentives for node cooperation in parallel relay networks with one source node and multiple relay nodes. All nodes are selfish and strategic, interested in maximizing their own profit instead of the social welfare. We consider the practical situation where a node cannot observe the state of links adjacent to other nodes. We examine a general game setting where the source has full bargaining power, and propose a framework for analyzing the efficiency loss induced by incomplete information.

1 Introduction

There is now widespread awareness of the importance of incentives in the management of communication networks. Network nodes often cannot be relied upon to cooperatively implement network algorithms in the service of the social good. Instead, selfish nodes will behave in a given manner only if it is profitable for them to do so. Of clear interest is the impact of such selfish actions on the social good. From the network point of view, it is important to design incentives such as pricing schemes, which induce selfish behavior aligned with the social good.

In single-hop and multi-hop networks, the incentive issue and its impact on social efficiency have been extensively studied[1,2,3,4,5,6,7,8,9]. All these papers, however, assume a complete information setting where players in the network game have complete knowledge about quantities such as the state of network links. In practice, this assumption is often too strong. Information regarding network quantities is typically incomplete and imperfect. In a multi-hop network such as the Internet, a source does not typically have perfect information on the congestion state of links a few hops away [10]. In wireless networks, the source usually cannot observe or test the channel state from a relay to the destination. Neither can a relay observe the channel state from other relays to the destination. Given the above, it is clear that in analyzing selfish behavior in network settings, the role of incomplete information must be emphasized.

One approach to network design problems with incomplete information is through dominant implementable mechanisms [11]. This idea has been used in the context of spectrum auctions [12] and communication networks [13]. These mechanisms, however, require a centralized authority and extra funding from an outsider. This makes the extension to general multi-hop networks difficult.

R. Jain and R. Kannan (Eds.): GameNets 2011, LNICST 75, pp. 475–489, 2012.

Another approach, based on the idea of Bayesian Nash Equilibrium, a generalization of the Nash Equilibrium concept, is advocated in [14]. Here, the authors consider selfish routing in a single-hop network, where every source node knows only its own traffic requirement, but has knowledge of the traffic distribution of other sources. While the results in [14] are appealing, it remains unclear how they might extend to the multi-hop network situation.

In this work, we investigate the impact of incomplete information on the problem of incentives in a two-hop parallel relay network. In our setting, the state of the links adjacent to a given relay is not observable by the source or the other relays, although the prior distribution of the link state is known. We consider a game with full source bargaining power, where the source offers a general contract mapping signals from the relays to traffic allocations and transfer payments. Given the proposed contract, the relays decide whether to accept or not. Those relays which choose to accept the contract then send link-state-dependent signal functions to the source, which then allocates traffic and transfer payments according to the proposed contract. This general setting includes as a specific case of pricing games where the signals from the relays consist of nonlinear traffic-dependent charging functions. We study the Bayesian Nash equilibria corresponding to the general game. To provide a benchmark, we first show that in the game with complete information, (Bayesian) Nash equilibria exist and are all efficient. Next, we investigate the game with incomplete information. To deal with the difficulty of characterizing the Bayesian Nash Equilibria in this case, we first show that if a resource allocation outcome can be realized by a Bayesian Nash equilibrium, then there exists a "truth telling" Bayesian Nash equilibrium that realizes the outcome. We then show that the outcome of the "truth telling" Bayesian Nash equilibrium coincides with that of the Nash equilibrium for a *complete information game*, in which the link cost functions are replaced by a specified "virtual cost functions." Using this approach, we obtain for a symmetric network scenario a bound on the amount of inefficiency which may result from incomplete information.

2 Network Model and Problem Formulation

Consider a parallel relay network modeled by a directed graph $G = (V, E)$, with a single source s, destination d, and a set of relays I, where $|I| = n$. We assume that there is no direct link between s and d. Instead, The relays in I are used to forward traffic in a two-hop fashion from s to d.

We shall consider two scenarios. In the first *inelastic* scenario, the source has a fixed rate r_s of transmission. This rate must be carried by the relays in I, where the traffic rate forwarded by relay i is r_i, and $\sum_{i=1}^{n} r_i = r_s$. In the second *elastic* scenario, the source may be willing to withhold some of its transmission rate, according to how the cost of sending traffic affects it overall utility. Let r_0 denote the amount of traffic withheld or rejected. Then $r_s - r_0$ is the total admitted traffic from the source. A traffic vector $\mathbf{r} \triangleq (r_0, r_1, \ldots, r_n) \in \mathbb{R}_+^{n+1}$ is a feasible routing of the source traffic if it satisfies $r_0 + \sum_{i=1}^{n} r_i = r_s$.

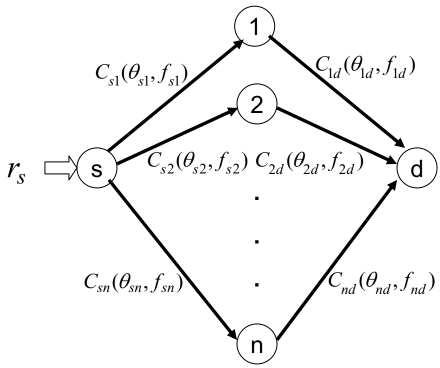

Fig. 1. Relay channel

2.1 Cost Function and Utility Function

In general, for any relay node i, there is a cost involved in forwarding traffic for source s. This cost typically depends both on the properties of the links adjacent on relay i and the amount of traffic flowing through those links. Denote the traffic flow on link $(i,j) \in E$ by f_{ij}. We assume that link (i,j) has a cost function $C_{ij}(\theta_{ij}, f_{ij})$ with $C_{ij}(\theta_{ij}, 0) = 0$, where θ_{ij} is a measure of the quality of link (i,j). This quality may have different physical meanings in different contexts. For example, if the cost function reflects the queuing delay on (i,j), then using the M/M/1 approximation, $C_{ij}(\theta_{ij}, f_{ij}) = \frac{f_{ij}}{k_{ij} - f_{ij}}$. Here, θ_{ij} denotes the link capacity k_{ij}. For another example, consider the cost of power assumption required for transmitting traffic of rate f_{ij} over a wireless link with channel gain g_{ij}, bandwidth W, and receiver noise power N. Using the Shannon capacity formula, we have $f_{ij} = W \log(1 + g_{ij} P_{ij}/N)$, where P_{ij} is transmission power required on link (i,j). Thus, the link cost is

$$C_{ij}(\theta_{ij}, f_{ij}) = \frac{N}{g_{ij}}(2^{f_{ij}/W} - 1).$$

Here, θ_{ij} denotes the channel gain g_{ij}.

Now consider the overall cost $C_i(\theta_i, r_i)$ for relay node i to forward traffic of rate r_i from source s to destination d, where θ_i measures the quality or *type* of the *path* from s to d through i. We assume that $C_i(\theta_i, r_i) = C_{si}(\theta_{si}, r_i) + C_{id}(\theta_{id}, r_i)$. The costs $C_i(\theta_i, r_i)$ are particularly amenable to analysis if θ_i can be expressed as a simple scalar function of θ_{si} and θ_{id}: $\theta_i = h(\theta_{si}, \theta_{id})$. This is true in the example of the power consumption cost function given above, where $\theta_{ij} = g_{ij}$ is the channel gain on link (i,j). Normalizing the bandwidth and receiver noise power to 1, we have

$$C_i(\theta_i, r_i) = P_{si} + P_{id} = (2^{r_i} - 1)(g_{si}^{-1} + g_{id}^{-1}) = (2^{r_i} - 1)\theta_i^{-1}, \tag{1}$$

where $\theta_i \triangleq (g_{si}^{-1} + g_{id}^{-1})^{-1} = (\theta_{si}^{-1} + \theta_{id}^{-1})^{-1}$. In this paper, we focus on situations where the path quality θ_i can be expressed as a scalar function of θ_{si} and θ_{id}. We further assume that θ_i belongs to a compact interval $[\underline{\theta}_i, \overline{\theta}_i]$.

Motivated by the power consumption example, we assume that $C_i(\theta_i, r_i)$ is twice continuously differentiable on $[\underline{\theta}_i, \overline{\theta}_i] \times [0, r_s]$, and strictly increasing and strictly convex in r_i: $\partial C_i(\theta_i, r_i)/\partial r_i > 0$ and $\partial^2 C_i(\theta_i, r_i)/\partial r_i^2 > 0$. Also, assume that $C_i(\theta_i, r_i)$ is strictly decreasing in θ_i: $\partial C_i(\theta_i, r_i)/\partial \theta_i < 0$. Furthermore, assume $\partial^2 C_i(\theta_i, r_i)/\partial \theta_i \partial r_i \leq 0$.

Now consider the source s. Let the utility function of the source be given by $W_s(\theta_s, r)$, where $\theta_s \in [\underline{\theta}_s, \overline{\theta}_s]$ parameterizes the utility for the source, and r is the source rate admitted into the network. For example, the source utility may be $W_s(\theta_s, r) = \theta_s \log(1 + r)$. Assume that $W_s(r) = W_s(r_s)$ for all $r \geq r_s$, i.e. r_s is the maximum desired source rate. $W_s(\theta_s, r)$ is assumed to be continuously differentiable, strictly increasing and strictly concave in r on $[0, r_s]$. Let $C_s(\theta_s, r_0) \triangleq W_s(r_s) - W_s(r_s - r_0)$ denote the source's *utility loss* from having traffic of rate r_0 withheld from the network. Since $W_s(r_s)$ is a constant, it can be seen that $C_s(\theta_s, r_0)$ is continuously differentiable on $[\underline{\theta}_s, \overline{\theta}_s] \times [0, r_s]$, strictly increasing and strictly convex in r_0: $\partial C_s(\theta_s, r_0)/\partial r_0 > 0$ and $\partial^2 C_s(\theta_s, r_0)/\partial r_0^2 > 0$. Furthermore, we assume that $C_s(\theta_s, r_0)$ is strictly decreasing in θ_s: $\partial C_s(\theta_s, r_0)/\partial \theta_s < 0$. Finally, it can be seen that $C_s(\theta_s, 0) = 0$ for all θ_s.

2.2 Socially Optimal Allocation

A socially optimal traffic allocation in a parallel relay network is an allocation which minimizes the total network cost, assumed to be the sum of the link costs. Let $R \triangleq \{(r_0, r_1, ..., r_n) : r_j \geq 0 \ \forall j = 0, \ldots, n, \sum_{j=0}^n r_j = r_s\}$ be the set of feasible traffic allocations, and let $\mathbf{r} = (r_0, r_1, ..., r_n) \in R$ denote the vector of traffic rates in the network.

Definition 1. *A traffic allocation vector* $\mathbf{r}^*$ *is called socially optimal if*

$$\mathbf{r}^* \in \arg\min_{\mathbf{r} \in R} C_s(\theta_s, r_0) + \sum_{i=1}^{n} C_i(\theta_i, r_i). \tag{2}$$

Since the link cost functions $C_i(\theta_i, r_i)$ as well as $C_s(\theta_s, r_0)$ are all strictly increasing and strictly convex, the socially optimal allocation $\mathbf{r}^*$ exists and is unique. The conditions for specifying $\mathbf{r}^*$ can be obtained using the Kuhn-Tucker conditions. Let $c_i(\theta_i, r_i) \triangleq \partial C_i(\theta_i, r_i)/\partial r_i$ and $c_s(\theta_s, r_0) \triangleq \partial C_s(\theta_s, r_0)/\partial r_0$ denote the marginal cost function of link i and the marginal cost function of the overflow link for source s, respectively.

For the case of an inelastic source, $\mathbf{r}^* = (0, r_1^*, \ldots, r_n^*)$ is the socially optimal allocation if and only if for each $i = 1, \ldots, n$,

$$c_i(\theta_i, r_i^*) = c^* \text{ if } r_i^* > 0, \quad c_i(\theta_i, r_i^*) > c^* \text{ if } r_i^* = 0. \tag{3}$$

For the case of an elastic source, $\mathbf{r}^* = (r_0^*, r_1^*, \ldots, r_n^*)$ is the socially optimal allocation if and only if (3) holds and furthermore,

$$c_s(\theta_s, r_0^*) = c^* \text{ if } r_0^* > 0, \quad c_s(\theta_s, r_0^*) > c^* \text{ if } r_0^* = 0.$$

2.3 Game Structure

We assume that the (maximum) source input rate r_s and the parameter θ_s are known to all nodes. We also assume that θ_i is randomly distributed according to distribution function $F_i(\theta_i)$. In practical network scenarios, the exact realization of θ_i is typically known only to relay i, and not to the source or to the relays other than i. Thus, θ_i is *private information* to relay i. Nevertheless, the source and other relays may still have knowledge of the distribution $F_i(\theta_i)$. For instance, a wireless source or a relay $j \neq i$ may know the distribution of the channel gains for relay i, but typically does not know the realization of those channel gains.

In order for the source node to allocate its traffic intelligently in the presence of incomplete information regarding the θ_i's, it needs to observe some "signal" from the relay nodes. This can be realized by having the relay node send a signal according to the realization of its type to the source.[1] Let M_i be the set of signals for relay i, where M_i is a subset of the set of differentiable functions on $[0, r_s]$. The signal map for relay i is

$$s_i : \Theta_i \to M_i,$$

where $\Theta_i \triangleq [\underline{\theta}_i, \overline{\theta}_i]$ and $s_i(\theta_i) = m_i(\cdot)$.

Given the signals $m_i(\cdot), i = 1, \ldots, n$, the source decides on an allocation of its traffic as well as a vector of transfer payments to the relays. This allocation is called a *contract*. Let $\mathbf{r} = (r_0, r_1, ..., r_n) \in R$ denote the vector of traffic rates in the network. Note that for the inelastic case, $r_0 = 0$. Now let $\mathbf{t} = (t_1, t_2, ..., t_n) \in \mathbb{R}^n_+$ be the vector of transfer payments, where t_i is the transfer payment to relay i. Let $M \triangleq M_1 \times \cdots \times M_n$ and $T \triangleq \mathbb{R}^n_+$. Then the allocation map of the source node is

$$g : M \to R \times T,$$

where $g(m_1(\cdot), \ldots, m_n(\cdot)) = (\mathbf{r}, \mathbf{t})$.

The above framework encompasses many forms of pricing games explored in previous literature. For instance, in [9], the relay signals are simply charging functions $P_i(\cdot)$, and the transfer payments are required to equal the charges demanded by the relays, i.e. $t_i = P_i(r_i)$.

The signal maps of the relays along with the allocation map of the source realize a corresponding network allocation map

$$f : \Theta \to R \times T,$$

where $f(\theta_1, \ldots, \theta_n) = g(s_1(\theta_1), \ldots, s_n(\theta_n)) = (\mathbf{r}, \mathbf{t})$.

In the game with incomplete information corresponding to the above setting, the utility of the source is given by

$$U_s(\theta_s, g(s_1(\theta_1), \ldots, s_n(\theta_n)) = W_s(r_s) - C_s(\theta_s, r_0) - \sum_{i=1}^{n} t_i.$$

[1] One can also consider the possibility of the source sending a signal according to its type θ_s. However, since we assume θ_s is known to all network nodes, we do not consider this possibility here.

The utility of relay i is given by

$$U_i(\theta_i, g(s_1(\theta_1), \ldots, s_n(\theta_n))) = t_i - C_i(\theta_i, r_i).$$

The game with incomplete information proceeds as follows. First, each relay i observes its own private information θ_i. Second, the source provides a contract for the relay nodes. The contract announces the source allocation rule $g : M \to R \times T$. Third, the relays simultaneously decide to either accept or reject the contract. If a given relay accepts the contract, then it will participate in the game which follows. Otherwise, the relay quits and receives zero utility.[2] Fourth and finally, the relay nodes simultaneously send their signals to the source, and the source allocates rates and transfer payments according to the announced g.

In the following, we give the formal definition of the Bayesian Nash equilibrium corresponding to the game with incomplete information described above. Let $\theta \triangleq (\theta_1, \ldots, \theta_n)$, $\theta_{-i} \triangleq (\theta_j)_{j \neq i}$, and $s_{-i}(\theta_{-i}) \triangleq (s_j(\theta_j))_{j \neq i}$.

Definition 2. *A Bayesian Nash Equilibrium of the above game is a set of strategies* $\{s_1, \ldots, s_n, g\}$ *satisfying*

1. *for each relay node i and for every feasible $\widetilde{s}_i : \Theta_i \to M_i$,*

$$E_{\theta_{-i}} \{U_i(\theta_i, g(s_i(\theta_i), s_{-i}(\theta_{-i})))\} \geq E_{\theta_{-i}} \{U_i(\theta_i, g(\widetilde{s}_i(\theta_i), s_{-i}(\theta_{-i})))\}; \quad (4)$$

2. *for every feasible $\widetilde{g} : M \to R \times T$,*

$$E_\theta \{U_s(\theta_s, g(s(\theta)))\} \geq E_\theta \{U_s(\theta_s, \widetilde{g}(s(\theta)))\}. \quad (5)$$

3 Games with Complete Information

We examine games where the source has full bargaining power, in the sense that source can offer any contract $g \in R \times T$.[3] We first investigate the (Bayesian) Nash equilibria which can result from games with source bargaining power in the case of complete information. Here, we show that all (Bayesian) Nash equilibria are efficient. Then, we proceed to the case of incomplete information, and characterize the potential inefficiencies associated with that case.

In the complete information game, the source can observe the type vector $\theta = (\theta_1, \ldots, \theta_n)$ of the relays, and then design the allocation map g according to θ. Since the type θ_i is not private to relay i, relay i cannot manipulate this information in designing its signalling strategy s_i. Since the source can observe θ, it can effectively ignore the strategies of the relays in designing g. Nevertheless, the source needs to ensure that the relays will accept its proposed contract and

[2] Note that the relays which quit can simply be left out of the game formulation. Thus, without loss of generality, we assume for the rest of the paper that the source plays the game in a manner which gives non-negative expected utility to all relays, so that all relays stay in the game.

[3] In games with partial bargaining power, such as auctions where the transfers are determined by bids, the source can propose only a subset of $R \times T$.

stay in the game. The latter will hold as long as $E_{\theta_{-i}}\{U_i(\theta_i, g(\theta))\} = E_{\theta_{-i}}\{t_i - C_i(\theta_i, r_i)\} \geq 0$ for all i. That is, all relays receive non-negative expected utility by accepting the contract proposed by the source, and therefore are willing to participate in the game.

Lemma 1. *In any (Bayesian) Nash Equilibrium of the complete information game with source bargaining power, all relays receive zero utility.*

Proof. Suppose that there exists a (Bayesian) Nash Equilibrium where the source allocation rule

$$g(m_1(\cdot), ..., m_n(\cdot)) = (\mathbf{r}, \mathbf{t})$$

is such that $U_i(\theta_i, r_i, t_i) = t_i - C_i(\theta_i, r_i) > 0$ for some i. Since the source can observe θ, it could select another allocation rule $g'(m_1(\cdot), ..., m_n(\cdot)) = (\mathbf{r}', \mathbf{t}')$ such that

$$r_i' = r_i, i = 1, \ldots, n; \quad t_i' = t_i - \epsilon, \;\; t_j' = t_j \text{ for all } j \neq i$$

where ϵ is small enough so that $t_i' - C_i(\theta_i, r_i') > 0$. Note that the set of relays which would opt to accept contract g and stay in the game is the same as the set for contract g'. On the other hand, by shifting its allocation rule from g to g', the source has strictly decreased its total transfer payment, while keeping the same traffic allocation. Thus, the source's utility is strictly increased. This contradicts our assumption of being at a Nash equilibrium.

Theorem 1. *In the complete information game with source bargaining power, all (Bayesian) Nash equilibria are efficient.*

Proof. At any Nash equilibrium, the source maximizes its utility

$$U_s(\theta_s, g(s_1(\theta_1), \ldots, s_n(\theta_n)) = W_s(\theta_s, r_s) - C_s(\theta_s, r_0) - \sum_{i=1}^{n} t_i.$$

By Lemma 1, at the equilibrium, we have $t_i = C_i(\theta_i, r_i)$ for all i. Thus, the traffic allocation by the source minimizes $C_s(\theta_s, r_0) + \sum_{i=1}^{n} C_i(\theta_i, r_i)$, and therefore the equilibrium is efficient.

Using Lemma 1 and Theorem 1, we can easily solve for the Nash equilibrium of the complete information game with source bargaining power. By Theorem 1, the source allocation rule at the equilibrium may be obtained by solving for the socially optimal traffic allocation $\mathbf{r}^*$, where $\mathbf{r}^* = \arg\min_{\mathbf{r} \in R} C_s(\theta_s, r_0) + \sum_{i=1}^{n} C_i(\theta_i, r_i)$. As noted in Section 2.2, due to the strict convexity of the optimization problem, $\mathbf{r}^*$ exists and is unique. By Lemma 1, at the equilibrium, the transfer payment $t_i = C_i(\theta_i, r_i^*)$ for every $i = 1, \ldots, n$.[4] For the relays, any feasible signal map s_i may be chosen for the equilibrium.

[4] Recall that $C_i(\theta_i, 0) = 0$.

4 Games with Incomplete Information

We now turn to the case where the source cannot observe the types of the relay. In this case, the relay nodes can manipulate their types in order to obtain more utility, and the source can no longer design the allocation map according to θ. Instead, the source must maximize its expected utility according to the signals sent by the relays. The characterization of Bayesian Nash Equilibria for this case is very difficult due to the complexity of the strategy set and the possible behaviors of source and relays. Nevertheless, we devise a method for characterizing outcomes corresponding to the Bayesian Nash Equilibria which avoids the difficulty of calculating the the equilibria explicitly. We shall do this in two steps. First, we show that if a resource allocation outcome can be realized by a Bayesian Nash equilibrium for a game with source bargaining in which every relay receives non-negative expected utility, then there exists a "truth telling" Bayesian Nash equilibrium that realizes the outcome. Second, we show that the outcome of the "truth telling" Bayesian Nash equilibrium coincides with that of the Nash equilibrium for a *complete information game*, in which the link cost functions are replaced by a specified "virtual cost functions."

Definition 3. *A Bayesian Nash Equilibrium of the game with source bargaining power is* truth telling *if $M = \Theta$ and every relay node is willing to report its true type to the source node.*

Theorem 2. *If a resource allocation outcome f can be realized by a Bayesian Nash Equilibrium of the game with source bargaining power, in which every relay receives non-negative expected utility, then there exists a truth telling Bayesian Nash Equilibrium which realizes f.*

Proof. Suppose there is a Bayesian Nash Equilibrium which realizes the allocation outcome $f(\theta)$. By the definition of the Bayesian Nash Equilibrium, we have (4) and (5). Now observe that by (4), we must have

$$\theta_i \in \arg\max_{\widetilde{\theta}_i} E_{\theta_{-i}} \left\{ U_i(\theta_i, g(s_i(\widetilde{\theta}_i), s_{-i}(\theta_{-i}))) \right\} \quad \text{for all } i.$$

Otherwise, if there exists some θ' such that $E_{\theta_{-i}} \{U_i(\theta_i, g(s_i(\theta'_i), s_{-i}(\theta_{-i})))\} > E_{\theta_{-i}} \{U_i(\theta_i, g(s_i(\theta_i), s_{-i}(\theta_{-i})))\}$, then there is another strategy $s'_i(\theta)$ satisfying $s'_i(\theta_i) = s_i(\theta'_i)$ and $s'_i(\theta) = s_i(\theta)$ for all $\theta \neq \theta_i$, such that $E_{\theta_{-i}}$ $\{U_i(\theta_i, g(s'_i(\theta_i), s_{-i}(\theta_{-i})))\} > E_{\theta_{-i}} \{U_i(\theta_i, g(s_i(\theta_i), s_{-i}(\theta_{-i})))\}$, violating (4). Therefore, since $g(s_1(\theta_1), \ldots, s_n(\theta_n)) = f(\theta)$, we have

$$\theta_i \in \arg\max_{\widetilde{\theta}_i \in \Theta_i} E_{\theta_{-i}} \left\{ U_i(\theta_i, f(\widetilde{\theta}_i, \theta_{-i})) \right\} \quad \text{for all } i, \tag{6}$$

$$f \in \arg\max_{\widetilde{f}} E_\theta \left\{ U_s(\theta_s, \widetilde{f}(\theta)) \right\}. \tag{7}$$

Thus, there exists a direct truth telling Bayesian Nash Equilibrium with the outcome $f(\theta)$.

Theorem 2 says that the set of outcomes corresponding to Bayesian Nash Equilibria for the game with source bargaining power and incomplete information is a subset of the outcomes corresponding to truth telling Bayesian Nash Equilibria, in which each relay proposes its type truthfully to the source, and the source optimally allocates rates according to the relays' types. This finding simplifies our analysis considerably, since we can now focus on the truth telling Bayesian Nash Equilibria in order to bound the efficiency loss introduced by incomplete information in games with source bargaining power.

We now investigate the outcomes which can be realized by truth telling Bayesian Nash Equilibria. Notice that these equilibria correspond to the solutions of the optimization problem given by (6) and (7), in addition to the non-negative expected utility constraint

$$E_{\theta_{-i}}\left\{U_i(\theta_i, r_i)\right\} = E_{\theta_{-i}}\left\{t_i(\theta, \mathbf{r}) - C_i(\theta_i, r_i(\theta))\right\} \geq 0 \quad \text{for all } i. \tag{8}$$

and feasibility constraint $\mathbf{r} \in R$.

Theorem 3. *The set of solutions for the optimization problem defined by* (6)-(8) *is the same as the set of outcomes corresponding to the Nash equilibria for the complete information game in which the link cost functions* $C_i(\theta_i, r_i)$ *are replaced by*

$$J_i(\theta_i, r_i) = C_i(\theta_i, r_i) - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} \frac{\partial C_i(\theta_i, r_i)}{\partial \theta_i}. \tag{9}$$

Proof. Please see the Appendix.

We refer to the functions J_i as *virtual cost functions*. The proof involves examining the first and second order optimality conditions, using the envelope theorem, and expressing the expected utility of the source in an equivalent form. We skip the proof here due to space constraints. Note that by Theorem 1, all Nash equilibria corresponding to games with complete information are efficient. Thus, the set of outcomes corresponding to the Nash equilibria for the complete information game with virtual link cost functions $J_i(\theta_i, r_i)$ is given by

$$\mathbf{r}' = \arg\min_{\mathbf{r} \in R} C_s(\theta_s, r_0) + \sum_{i=1}^{n} J_i(\theta_i, r_i).$$

5 Efficiency Analysis

In this section, we bound the amount of inefficiency in the outcomes for games with incomplete information. We focus on the inelastic scenario where $r_0 = 0$. Following [1], define the price of anarchy for type θ as:

$$\rho(\theta) = \frac{\max_{\mathbf{r} \in R^E} \sum_i C_i(\theta_i, r_i)}{\min_{\mathbf{r} \in R} \sum_i C_i(\theta_i, r_i)} \tag{10}$$

where R^E is the set of all Bayesian Nash Equilibria for the game with incomplete information, and $R^J \equiv \arg\min_{\mathbf{r}\in R} \sum_i J_i(\theta_i, r_i)$. By Theorems 2 and 3, we have $R^E \subseteq R^J$. Therefore,

$$\rho(\theta) \le \frac{\max_{\mathbf{r}\in R^J} \sum_i C_i(\theta_i, r_i)}{\min_{\mathbf{r}\in R} \sum_i C_i(\theta_i, r_i)} \tag{11}$$

Since the link cost functions are strictly convex, the socially optimal allocation $\mathbf{r}^*$ are given by the necessary and sufficient conditions in (3). An allocation $\mathbf{r}'$ in R^J must satisfy the following necessary conditions: for all $i \in \{1, \ldots, n\}$ such that $r'_i > 0$,

$$\frac{\partial C(\theta_i, r'_i)}{\partial r_i} - \frac{1-F_i(\theta_i)}{f_i(\theta_i)} \frac{\partial^2 C(\theta_i, r'_i)}{\partial \theta_i \partial r_i} \le \frac{\partial C(\theta_j, r'_j)}{\partial r_j} - \frac{1-F_j(\theta_j)}{f_j(\theta_j)} \frac{\partial^2 C(\theta_i, r'_j)}{\partial \theta_i \partial r_j} \quad \text{for all } j \tag{12}$$

We now bound the price of anarchy in the symmetric case.

Theorem 4. *Consider the symmetric case where the link cost functions $C_i(\theta_i, r_i)$ and the type distributions $F_i(\theta_i)$ are the same for all relays. If (i) $J(\theta_i, r_i)$ is convex in r_i and decreasing in θ_i, (ii) $X(\theta_i, r_i) \equiv J(\theta_i, r_i) - C(\theta_i, r_i)$ is concave in θ_i, (iii) $\frac{\partial X(\theta_i, r_i)}{\partial \theta_i \partial r_i} \le 0$, then the price of anarchy $\rho(\theta)$ can be upper bounded as follows. If the marginal cost function $c(\theta_i, r_i) = \frac{\partial C(\theta_i, r_i)}{\partial r_i}$ is concave, then $\rho(\theta) \le n$, where n is the number of relays. If $\frac{c(\underline{\theta}, r_s)}{c(\overline{\theta}, 0)} \le k$ for some constant k, then $\rho(\theta) \le k$.*

Proof. Let $(r'_i) \in \arg\min_i \sum_i J(\theta_i, r_i)$, and r^*_i be the efficient allocation. Let $x(\theta_i, r_i) = \frac{\partial X(\theta_i, r_i)}{\partial r_i}$. Suppose $\theta_i > \theta_j$, and both carry traffic in optimal point, thus for the optimal allocation, $r^*_i > r^*_j > 0$. As $X(\theta_i, r_i)$ is concave in r_i, thus $x(\theta_i, r^*_i) < x(\theta_i, r^*_j)$. As the cost is decreasing in θ_i, we have $x(\theta_i, r^*_j) < x(\theta_j, r^*_j)$ Thus $x(\theta_i, r^*_i) < x(\theta_j, r^*_j)$. As $c(\theta_i, r^*_i) = c(\theta_j, r^*_j)$, $c(\theta_i, r^*_i) + x(\theta_i, r^*_i) < c(\theta_j, r^*_j) + x(\theta_j, r^*_j)$. But to maximize the virtual cost, we have $c(\theta_i, r'_i) + x(\theta_i, r'_i) \ge c(\theta_j, r'_j) + x(\theta_j, r'_j)$. As the virtual cost function is convex in r_i, we get $r'_i > r^*_i$. Thus $c(\theta_i, r'_i) > c(\theta_i, r^*_i) = c(\theta_j, r^*_j) > c(\theta_j, r'_j)$. As the cost function is convex in r_i and decreasing in θ_i, $c(\max_i \theta_i, r_s) \ge c(\theta_i, r'_i)$. Thus $C(\max_i \theta_i, r_s) \ge \sum_i C(\theta_i, r'_i)$. Thus the price of anarchy has the same bound with game 2.

As an example of a link cost function/type distribution pair which satisfies the assumptions in Theorem 4, consider the uniform type distribution over $[0, 1]$ and the cost function $C_i(\theta_i, r_i) = \frac{1}{\theta_i}(2^{r_i} - 1)$.

If either the virtual cost functions $J_i(\theta_i, r_i)$ are not convex, or the link cost functions and type distributions are not the same across relays, then higher prices of anarchy may result. Consider the situation in Figure 2. Here, there are only two relays. $(r^*_1, r_s - r^*_1)$ is the efficient allocation. Since the type distributions are not the same, the marginal virtual costs are as indicated in the figure. To minimize the sum of the virtual costs, the source allocates all traffic to relay 2, while this allocation is clearly the worst outcome for minimizing the sum of the link costs.

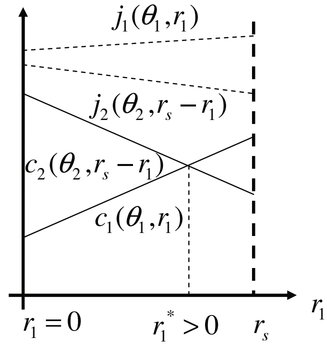

Fig. 2. Efficiency Loss in Asymmetric Case

6 Conclusion

This work investigated the impact of incomplete information on incentives for node cooperation in parallel relay networks. We considered a general game setting where the source has full bargaining power. We first showed that in the game with complete information, (Bayesian) Nash equilibria exist and are all efficient. Next, we investigated the game with incomplete information. To deal with the difficulty of characterizing the Bayesian Nash Equilibria in this case, we first showed that if a resource allocation outcome can be realized by a Bayesian Nash equilibrium, then there exists a "truth telling" Bayesian Nash equilibrium that realizes the outcome. We then showed that the outcome of the "truth telling" Bayesian Nash equilibrium coincides with that of the Nash equilibrium for a complete information game, in which the link cost functions are replaced by a specified "virtual cost functions." Using this approach, we obtained for a symmetric network scenario a bound on the amount of inefficiency which may result from incomplete information.

References

1. Roughgarden, T., Tardos, E.: How bad is selfish routing? In: Proceedings of the 41st Annual Symposium on Foundations of Computer Science (2000)
2. He, L., Walrand, J.: Pricing and revenue sharing strategies for internet service providers. In: Proceedings of the 24th Annual Joint Conference of the IEEE Computer and Communications Societies, vol. 1, pp. 205–216 (2005)
3. Shakkottai, S., Srikant, R.: Economics of network pricing with multiple isps. IEEE/ACM Transactions on Networking 14(6), 1233–1245 (2006)
4. Acemoglu, D., Ozdaglar, A.: Competition and efficiency in congested markets. Math. Oper. Res. 32(1), 1–31 (2007)
5. Neely, M.: Optimal pricing in a free market wireless network. In: IEEE 26th IEEE International Conference on Computer Communications, INFOCOM 2007, pp. 213–221 (May 2007)

6. Zhong, S., Chen, J., Yang, Y.: Sprite: a simple, cheat-proof, credit-based system for mobile ad-hoc networks. In: Proceedings of IEEE INFOCOM, vol. 3, pp. 1987–1997 (March 2003)
7. Ileri, O., Mau, S.-C., Mandayam, N.: Pricing for enabling forwarding in self-configuring ad hoc networks. IEEE Journal on Selected Areas in Communications 23(1), 151–162 (2005)
8. Wang, B., Han, Z., Liu, K.R.: Distributed relay selection and power control for multiuser cooperative communication networks using stackelberg game. IEEE Transactions on Mobile Computing 8(7), 975–990 (2009)
9. Xi, Y., Yeh, E.: Pricing, competition, and routing for selfish and strategic nodes in multi-hop relay networks. In: Proceedings of the 27th Annual Joint Conference of the IEEE Computer and Communications Societies, pp. 1463–1471 (2008)
10. Bertsekas, D.P., Gallager, R.: Data Networks, 2nd edn. Prentice Hall (1992)
11. Nisan, N., Ronen, A.: Algorithmic mechanism design (extended abstract). In: Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing, STOC 1999, pp. 129–140. ACM, New York (1999)
12. Archer, A., Papadimitriou, C., Talwar, K.: An approximate truthful mechanism for combinatorial auctions with single parameter agents. Internet Mathematics 1(2), 129–150 (2006)
13. Dobson, I., Carreras, B., Lynch, V., Newman, D.: Communication requirements of vcg-like mechanisms in convex environments. In: Proceedings of the Allerton Conference on Communication, Control, and Computing (2005)
14. Gairing, M., Monien, B., Tiemann, K.: Selfish routing with incomplete information. Theory of Computing Systems 42, 91–130 (2008), doi:10.1007/s00224-007-9015-8

7 Appendix

Proof of Theorem 3: The first and second-order conditions for (6) are:

$$\left.\frac{dE_{\theta_{-i}}\left\{U_i(\theta_i, f(\widetilde{\theta}_i, \theta_{-i}))\right\}}{d\widetilde{\theta}_i}\right|_{\widetilde{\theta}_i=\theta_i} = 0 \tag{13}$$

and

$$\left.\frac{d^2E_{\theta_{-i}}\left\{U_i(\theta_i, f(\widetilde{\theta}_i, \theta_{-i}))\right\}}{d\widetilde{\theta}_i^{\,2}}\right|_{\widetilde{\theta}_i=\theta_i} \leq 0. \tag{14}$$

The first-order condition is equivalent to

$$E_{\theta_{-i}} \left.\frac{dt_i(\widetilde{\theta}_i, \theta_{-i})}{d\widetilde{\theta}_i}\right|_{\widetilde{\theta}_i=\theta_i} \tag{15}$$

$$= E_{\theta_{-i}} \left.\left\{\frac{\partial C_i(\theta_i, r_i(\widetilde{\theta}_i, \theta_{-i}))}{\partial r_i} \frac{dr_i(\widetilde{\theta}_i, \theta_{-i})}{d\widetilde{\theta}_i}\right\}\right|_{\widetilde{\theta}_i=\theta_i}. \tag{16}$$

The second-order condition is equivalent to

$$E_{\theta_{-i}} \frac{d^2 t_i(\widetilde{\theta}_i, \theta_{-i})}{d\widetilde{\theta}_i^{\,2}} \bigg|_{\widetilde{\theta}_i = \theta_i}$$

$$\leq E_{\theta_{-i}} \left\{ \frac{\partial^2 C_i(\theta_i, r_i(\widetilde{\theta}_i, \theta_{-i}))}{\partial r_i^2} \left[\frac{dr_i(\widetilde{\theta}_i, \theta_{-i})}{d\widetilde{\theta}_i} \right]^2 \right. \tag{17}$$

$$\left. + \frac{\partial C_i(\theta_i, r_i(\widetilde{\theta}_i, \theta_{-i}))}{\partial r_i} \frac{d^2 r_i(\widetilde{\theta}_i, \theta_{-i})}{d\widetilde{\theta}_i^{\,2}} \right\} \bigg|_{\widetilde{\theta}_i = \theta_i}.$$

By evaluating the first-order condition at θ_i differentiating with respect to θ_i, we get:

$$E_{\theta_{-i}} \left\{ \frac{d^2 \left\{ t_i(\theta_i, \theta_{-i}) \right\}}{d\theta_i^2} \right\}$$

$$= E_{\theta_{-i}} \left\{ \frac{\partial^2 C_i(\theta_i, r_i(\theta_i, \theta_{-i}))}{\partial r_i^2} \left[\frac{dr_i(\theta_i, \theta_{-i})}{d\theta_i} \right]^2 \right. \tag{18}$$

$$+ \frac{\partial C_i(\theta_i, r_i(\theta_i, \theta_{-i}))}{\partial r_i} \frac{d^2 r_i(\theta_i, \theta_{-i})}{d\theta_i^2}$$

$$\left. + \frac{\partial^2 C_i(\theta_i, r_i(\theta_i, \theta_{-i}))}{\partial r_i \partial \theta_i} \frac{dr_i(\theta_i, \theta_{-i})}{d\theta_i} \right\}.$$

Comparing with the second-order condition, we get

$$E_{\theta_{-i}} \frac{\partial^2 C_i(\theta_i, r_i(\theta_i, \theta_{-i}))}{\partial r_i \partial \theta_i} \frac{dr_i(\theta_i, \theta_{-i})}{d\theta_i} \leq 0. \tag{19}$$

We have already assumed that

$$\frac{\partial^2 C_i(\theta_i, r_i(\theta_i, \theta_{-i}))}{\partial r_i \partial \theta_i} \leq 0 \text{ for each } \theta_{-i}. \tag{20}$$

Notice that when an outcome can be realized by a Bayesian Nash Equilibrium, the following condition must hold:

$$\frac{\partial r_i(\theta_i, \theta_{-i})}{\partial \theta_i} \geq 0 \text{ given any } \theta_{-i} \tag{21}$$

Otherwise, the source would allocate a higher rate to a lower type relay, which is not optimal. Notice that by (20) and (21), (19) automatically holds.

Thus, the following conditions are necessary for the first and second-order conditions to hold.

$$\frac{dE_{\theta_{-i}} \left\{ t_i(\widetilde{\theta}_i, \theta_{-i}) \right\}}{d\widetilde{\theta}_i}$$

$$= E_{\theta_{-i}} \frac{\partial C_i(\theta_i, r_i(\widetilde{\theta}_i, \theta_{-i}))}{\partial r_i} \frac{dr_i(\widetilde{\theta}_i, \theta_{-i})}{d\widetilde{\theta}_i} \bigg|_{\widetilde{\theta}_i = \theta_i}$$

$$\frac{\partial r_i(\theta_i, \theta_{-i})}{\partial \theta_i} \geq 0 \text{ given any } \theta_{-i}$$

Let $V_i(\theta_i, \theta_{-i}) = \max_{\widetilde{\theta}_i} U_i(\theta_i, r_i(\widetilde{\theta}_i, \theta_{-i}), t_i(\widetilde{\theta}_i, \theta_{-i}))$. We use the envelope theorem just as we did in the previous sections:

$$\begin{aligned} \frac{dE_{\theta_{-i}} V_i(\theta_i, \theta_{-i})}{d\theta_i} &= \left.\frac{\partial E_{\theta_{-i}} U_i(\theta_i, r_i(\widetilde{\theta}_i, \theta_{-i}), t_i(\widetilde{\theta}_i, \theta_{-i})}{\partial \theta_i}\right|_{\widetilde{\theta}_i = \theta_i} \\ &= -\left.\frac{\partial E_{\theta_{-i}} C_i(\theta_i, r_i(\widetilde{\theta}_i, \theta_{-i}))}{\partial \theta_i}\right|_{\widetilde{\theta}_i = \theta_i} \end{aligned} \tag{22}$$

Let $\overline{\theta_i}$ and $\underline{\theta_i}$ be the upper and lower bounds on relay node i's type, then

$$\begin{aligned} &E_{\theta_{-i}} V_i(\theta_i, \theta_{-i}) \\ &= E_{\theta_{-i}} V_i(\underline{\theta_i}, \theta_{-i}) - \int_{\underline{\theta_i}}^{\theta_i} \frac{\partial E_{\theta_{-i}} C_i(\theta_i, r_i(\theta_i, \theta_{-i}))}{\partial \theta_i} d\theta_i \end{aligned} \tag{23}$$

We see from the above equation that, as we already assumed $\frac{\partial C_i(\theta_i, r_i)}{\partial \theta_i} < 0$, the expected utility of relay i is non-decreasing with respect to θ_i. Thus, to guarantee that constraints (8) holds, the lowest type must receive non-negative profit. On the other hand, the relay with the lowest type can never receive a positive profit, otherwise the source will reduce its profit by some small amount and still guarantee that the contract is acceptable to all, which contradicts the definition of Bayesian Nash Equilibrium. Thus, the lowest type relay should receive zero profit.

$$E_{\theta_{-i}} V_i(\underline{\theta_i}, \theta_{-i}) = 0 \tag{24}$$

Plugging in, we get

$$E_{\theta_{-i}} V_i(\theta_i, \theta_{-i}) = -\int_{\underline{\theta_i}}^{\theta_i} \frac{\partial E_{\theta_{-i}} C_i(\theta_i, r_i(\theta_i, \theta_{-i}))}{\partial \theta_i} d\theta_i \tag{25}$$

Suppose the type distribution function of relay i is $F_i(\theta_i)$ and the density is $f_i(\theta_i)$. Let R be the expected revenue of the source node. Then,

$$\begin{aligned}
R &= E_\theta\left\{W_s(r_s) - C_s(\theta_s, r_0) - \sum_i t_i(\theta)\right\} \\
&= E_\theta\left\{W_s(r_s) - C_s(\theta_s, r_0) - \sum_i V_i(\theta) - \sum_i C_i(\theta)\right\} \\
&= E_\theta\left\{W_s(r_s) - C_s(\theta_s, r_0) - \sum_i C_i(\theta)\right\} + \sum_i \int_{\underline{\theta_i}}^{\overline{\theta_i}} f_i(\theta_i) E_{\theta_{-i}}\left[\int_{\underline{\theta_i}}^{\theta_i} \frac{\partial C_i(\theta_i', r_i(\theta_i', \theta_{-i}))}{\partial \theta_i'} d\theta_i'\right] d\theta_i \\
&= E_\theta\left\{W_s(r_s) - C_s(\theta_s, r_0) - \sum_i C_i(\theta)\right\} - \sum_i \int_{\underline{\theta_i}}^{\overline{\theta_i}} E_{\theta_{-i}}\left[\int_{\underline{\theta_i}}^{\theta_i} \frac{\partial C_i(\theta_i', r_i(\theta_i', \theta_{-i}))}{\partial \theta_i'} d\theta_i'\right] \times d(1 - F_i(\theta_i)) \\
&= E_\theta\left\{W_s(r_s) - C_s(\theta_s, r_s) - \sum_i C_i(\theta)\right\} - \sum_i E_{\theta_{-i}}\left[\int_{\underline{\theta_i}}^{\theta_i} \frac{\partial C_i(\theta_i', r_i(\theta_i', \theta_{-i}))}{\partial \theta_i'} d\theta_i'\right] \times (1 - F_i(\theta_i))|_{\underline{\theta_i}}^{\overline{\theta_i}} \\
&\quad + \sum_i E_{\theta_{-i}} \int_{\underline{\theta_i}}^{\overline{\theta_i}} (1 - F_i(\theta_i)) \times d\left[\int_{\underline{\theta_i}}^{\theta_i} \frac{\partial C_i(\theta_i', r_i(\theta_i', \theta_{-i}))}{\partial \theta_i'} d\theta_i'\right] \\
&= E_\theta\left\{W_s(r_s) - C_s(\theta_s, r_0) - \sum_i C_i(\theta)\right\} + \sum_i E_{\theta_{-i}} \int_{\underline{\theta_i}}^{\overline{\theta_i}} (1 - F_i(\theta_i)) \times d\left[\int_{\underline{\theta_i}}^{\theta_i} \frac{\partial C_i(\theta_i', r_i(\theta_i', \theta_{-i}))}{\partial \theta_i'} d\theta_i'\right] \\
&= E_\theta\{W_s(r_s) - C_s(\theta_s, r_0)\} - \sum_i E_{\theta_{-i}} \int_{\underline{\theta_i}}^{\overline{\theta_i}} C_i(\theta_i, r_i(\theta_i, \theta_{-i})) - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} \frac{\partial C_i(\theta_i, r_i(\theta))}{\partial \theta_i} f_i(\theta_i) d\theta_i \\
&= E_\theta[W_s(r_s) - C_s(\theta_s, r_0)] - E_\theta \sum_i \left(C_i(\theta_i, r_i(\theta)) - \frac{1 - F_i(\theta_i)}{f_i(\theta_i)} \frac{\partial C_i(\theta_i, r_i(\theta))}{\partial \theta_i}\right)
\end{aligned}$$

Thus, we obtain a game with complete information and full source bargaining power where the revenue function is changed to $J_i(\theta_i, r_i)$ rather than $C_i(\theta_i, r_i)$.