The Price of Selfishness in Network Coding

Jason R. Marden and Michelle Effros

***Abstract*—We introduce a game theoretic framework for studying a restricted form of network coding in a general wireless network. The network is fixed and known, and the system performance is measured as the number of wireless transmissions required to meet n unicast demands. Game theory is here employed as a tool for improving distributed network coding solutions. The approach involves designing cost functions for individual "players" (unicast sessions) so that players independently pursuing their own selfish interests (i.e., minimizing their individual costs) achieve a desirable system performance in a shared network environment. We propose a family of cost functions and compare the performance of the resulting distributed algorithms to the best performance that could be found and implemented using a centralized controller. We focus on the performance of stable solutions – where stability here refers to a form of Nash equilibrium defined below. Results include bounds on the best- and worst-case stable solutions as compared to the optimal centralized solution. We show that our bound on the worst-case stable performance cannot be improved using cost functions that are independent of the network structure. Results in learning in games prove that the best-case stable solution can be learned by self-interested players with probability approaching 1.**

I. Introduction

The network coding literature treats the design and performance of network codes aimed at goals such as maximizing capacity, minimizing power consumption, or improving the robustness of communication in a network environment. While early results primarily treat the multicast problem – where a single source transmits the same information to all sinks in the network – more recent work changes the focus to problems where multiple independent communication sessions share the network environment. A multiple unicast problem, characterized by a list of source-sink pairs with a distinct information flow to be established from each source to its corresponding sink, is one example of such a coding scenario. Multi-session network coding problems like the multiple unicast problem differ from the single-session network coding problems in that they establish competition between independent information flows for shared network resources.

A spectrum of approaches ranging from centralized control systems to totally distributed design and operation is possible for tackling multi-session coding problems. In centralized approaches like [1], the codes at all nodes of the network are designed by a code designer with access to complete information about the network and all sessions competing for network resources. At the other extreme end of the spectrum are distributed approaches like [2], which tackle large optimization problems by treating nodes or possibly sessions as independent decision makers in a shared network environment. One advantage of distributed algorithms is the savings in computation and coordination achieved by taking a large central optimization problem and dividing it into smaller problems. One disadvantage of this approach is that dividing the central optimization problem in this fashion may induce undesirable equilibrium points.

Recently, non-cooperative game theory has been proposed as a tool for cooperative control of distributed systems [3]. The interactions of a distributed/multi-agent control system are modeled as a non-cooperative game among agents where agents are *self-interested*. The challenge of modeling a multi-agent system as a non-cooperative game is designing *local cost functions*, which may very well be in conflict with one another, and *implementable distributed learning dynamics* such that the resulting global behavior is desirable with respect to the global objective.

In this paper, we explore the applicability of non-cooperative game theory for designing distributed algorithms for network coding in a wireless network. We formulate the network coding problem as a non-cooperative game where the self-interested players are the individual unicast sessions. We propose a family of local cost functions aimed at minimizing total network power consumption using a simple form of network coding. These cost functions are designed without knowledge of the specific network or the demands traversing the network but with the aim of encouraging individual behavior to optimize a centralized network objective. We evaluate the desirability of the cost functions by examining the performance at the stable solutions, i.e, the equilibria. We compare the global performance of both the optimal equilibrium and the worst-case equilibrium to the performance of of the optimal centralized solution.

In this paper, we primarily focus on analyzing equilibrium behavior. Of equal importance is understanding how the players can reach an equilibrium in a distributed fashion. The theory of learning in games provides several distributed learning algorithms that provide such guarantees. When modeling the network coding problem as a non-cooperative game, we can appeal to these distributed algorithms to guarantee convergence to an equilibrium. We direct the readers to [4]–[6] for a comprehensive review.

Several papers have used game theoretic methods for analyzing network coding problems by viewing either the individual unicasts or individual nodes in the network as selfish decision makers [7]–[10]. However, most of these results are only applicable in restricted settings. For example, the authors of [8] derive a cost mechanism for single-source multicast with network coding and analyze the efficiency of the Nash equilibrium. The authors of [7] propose a decentralized coding scheme for a class of two-user networks and analyze the equilibrium behavior.

The remainder of the paper is organized as follows. In Section II, we give a brief overview of the game theoretic concepts used in this paper. In Section III, we describe a simple wireless network coding problem called the reverse carpooling problem. In Section IV, we formulate the reverse carpooling problem as a state based game. In Section V, we derive bounds on the efficiency of the equilibrium points. In Section VI, we provide some concluding remarks.

J. R. Marden is a junior fellow with the Social and Informational Sciences Laboratory, California Institute of Technology, M/C 136–93, Pasadena, CA 91125, marden@caltech.edu.

M. Effros is a professor with the Department of Electrical Engineering, California Institute of Technology, M/C 136–93, Pasadena, CA 91125, effros@caltech.edu.

II. Game Theory Background

We consider a finite non-cooperative game [11]. The n players are represented by the set $N := \{1, ..., n\}$. Each player $i \in N$ has an action set $\mathcal{A}_i$ and a cost function $J_i : \mathcal{A} \to \mathbb{R}$ where $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$ denotes the joint action set. For an action profile $a = (a_1, a_2, ..., a_n) \in \mathcal{A}$, let a_{-i} denote the profile of player actions *other than* player i, i.e., $a_{-i} = (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$. With this notation, we sometimes write a profile a of actions as (a_i, a_{-i}). Similarly, we may write $J_i(a)$ as $J_i(a_i, a_{-i})$. We also use $\mathcal{A}_{-i} = \prod_{j \neq i} \mathcal{A}_j$ to denote the set of possible collective actions of all players other than player i.

The most well known form of an equilibrium is the Nash equilibrium.

Definition 1 (Pure Nash Equilibrium)**.** *An action profile $a^* \in \mathcal{A}$ is called a pure Nash equilibrium if for each player $i \in N$,*

$$J_i(a_i^*, a_{-i}^*) = \min_{a_i \in \mathcal{A}_i} J_i(a_i, a_{-i}^*). \tag{1}$$

A Nash equilibrium represents a scenario for which no player has an incentive to unilaterally deviate.

In many applications, players' cost functions are directly influenced by an exogenous state variable. In this paper, we consider the framework of *state based games* introduced in [12] which generalizes the non-cooperative game setting to such an environment. State based games are a simplification of the class of stochastic games [13]. In a state based game, there exists a finite state space X. Each player $i \in N$ has an action set $\mathcal{A}_i$ and a state dependent cost function $J_i : \mathcal{A} \times X \to \mathbb{R}$. We assume that the state evolves according to a state-transition function $P : \mathcal{A} \times X \to \Delta(X)$ where $\Delta(X)$ denotes the set of probability distributions over the finite state space X.

A state based game proceeds as follows. Let the state at time $t \in \{0, 1, ...\}$ be denoted by $x(t) \in X$. At any time t, each player i selects an action $a_i(t) \in \mathcal{A}_i$ randomly based on available information. The state $x(t)$ and the action profile $a(t) := (a_1(t), ..., a_n(t))$ together determine each player's cost $J_i(a(t), x(t))$ at time t. Each player selects an action $a_i(t)$ simultaneously seeking to minimize his one-stage expected cost $\mathbf{E}[J_i(a(t), x(t))]$, where the expectation is over player i's belief regarding the action choice of the other players, i.e., $a_{-i}(t)$. In this case, a player's strategy is unaffected by how his current action impacts the state dynamics and potential future rewards. After each player selects his respective action, the ensuing state $x(t+1)$ is chosen randomly according to the probability distribution $P(a(t), x(t)) \in \Delta(X)$. In this paper, we restrict our attention to state dynamics that satisfy

$$a(t) = a(t-1) \;\Rightarrow x(t+1) = x(t). \tag{2}$$

This paper focuses on analyzing equilibrium behavior in such games. We consider state based Nash equilibria, which generalize pure Nash equilibria to the state based setting [12].

Definition 2 (State Based Nash Equilibrium)**.** *The action state pair $[a^*, x^*]$ is a state based Nash equilibrium if for every player $i \in N$ and every state x' in the support of $P(a^*, x^*)$*

$$J_i(a_i^*, a_{-i}^*, x') = \min_{a_i \in \mathcal{A}_i} J_i(a_i, a_{-i}^*, x').$$

If $[a^*, x^*]$ is a state based Nash equilibrium, then no player $i \in N$ will have a unilateral incentive to deviate from a_i^* provided that all other players play a_{-i}^* regardless of the state that emerges according to the transition function $P(a^*, x^*)$. We use the term equilibrium to mean state based Nash equilibrium in the discussion that follows.

Given a state based game, an equilibrium may or may not exist. We consider the framework of *state based potential games*, introduced in [12], for which an equilibrium is guaranteed to exist. State based potential games generalize potential games [14] to the state based setting.

Definition 3 (State Based Potential Games)**.** *A state based game with state transition function P is a state based potential game if there exists a potential function $\phi : \mathcal{A} \to \mathbb{R}$ such that for any action state pair $[a, x] \in \mathcal{A} \times X$, player $i \in N$, and action $a_i' \in \mathcal{A}_i$*

$$J_i(a_i', a_{-i}, x) - J_i(a, x) < 0 \Rightarrow \phi(a_i', a_{-i}) - \phi(a) < 0.$$

This condition states that players' cost functions are aligned with the potential function. To see that an equilibrium exists in any state based potential game, let $[a^*, x^*]$ be any action state pair such that $a^* \in \arg\min_{a \in \mathcal{A}} \phi(a)$. The action state pair $[a^*, x^*]$ is an equilibrium.

III. A Simple Wireless Network Coding Problem

We consider the distributed design of network codes for multiple unicasts in a shared wireless network. We restrict our attention to the simplest form of network codes, where any node relaying one message in each direction between a pair of neighboring nodes can reduce the power required for transmission by broadcasting the bit-wise binary sum of the received messages in a single transmission rather than transmitting the two messages sequentially. Each neighbor can then determine its intended message by adding the information that it sent to the received sum, as illustrated in Figure 1. This type of coding is sometimes called "reverse carpooling" since it allows two flows share a single transmission provided that the two flows traverse the node in opposite directions. The goal of our network code design for this wireless network is to minimize the power required to simultaneously satisfy a given collection of unicast flow demands. For simplicity, we measure the cost of a network coding solution by evaluating the number of transmissions per packet required under steady state flow conditions. The following notation helps make these ideas concrete.

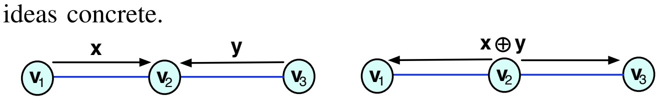

Fig. 1. Illustration of Reverse Carpooling

We describe a network by a set $V = \{v_1, ..., v_m\}$ of vertices, or nodes, and for each $v_i \in V$, the neighbors of node v_i, denoted as $\mathcal{N}(v_i) \subseteq V$; each transmission by node v_i is heard by all the nodes in $\mathcal{N}(v_i)$ and only those nodes.

Suppose the network needs to be shared by a finite set of players $N := \{1, ..., n\}$. Each player i represents a single unicast from sender s_i to receiver t_i, where $(s_i, t_i) \in V^2$. A path a_i from source s_i to terminal t_i equals a set of nodes capable of transmitting the information, i.e., $a_i = \{v_1, v_2, \ldots, v_{|a_i|}\}$ where $|a_i|$ denotes the number of nodes in a_i, $v_1 = s_i$, $v_{|a_i|} = t_i$, and $v_{k+1} \in \mathcal{N}(v_k)$ for all $k \in \{1, 2, ..., |a_i| - 1\}$. We use $\mathcal{A}_i$ to denote all paths from s_i to t_i available to player i. Let $\mathcal{A} := \prod_i \mathcal{A}_i$ denote the paths available to all players.

For analysis, it is convenient to label each transmission from v by the node from which the information was obtained and the node for which it is next intended so that we can recognize coding opportunities. To that end, let the detailed path of $a_i = \{v_1, v_2, \ldots, v_{|a_i|}\}$ be defined as

$$I(a_i) := \{v_1[\emptyset, v_2], v_2[v_1, v_3], \ldots, v_{|a_i|-1}[v_{|a_i|-2}, t_i]\}$$

where $v_k[v_{k-1}, v_{k+1}]$ represents a transmission by node v_k to node v_{k+1} of information that was received from node v_{k-1}. The transmission $v_1[\emptyset, v_2]$ represents a transmission by node v_1 to node v_2 of information that originated at node v_1. Notice that for any path a_i there is a unique detailed path $I(a_i)$ consisting of $|a_i| - 1$ elements.

A node that participates in more than one path may have the opportunity to combine messages and save on transmission costs if the paths traverse the node in opposite directions. Suppose players i and j traverse node v in opposite directions, i.e., $v \in a_i \cap a_j$, where player i sends message z_i using the transmission $v[v'_i, v''_i]$ and player j sends message z_j using the transmission $v[v'_j, v''_j]$. If $v'_i = v''_j$ and $v'_j = v''_i$, then node v can transmit $z_{ij} = z_i \oplus z_j$ and both v''_i and v''_j will be able to decode their intended message. This allows node v to serve players i and j with one transmission instead of two. This is true since node v'_i knows z_i and receives $z_i \oplus z_j$, allowing it to decode its intended message $z_j = z_i \oplus (z_i \oplus z_j)$.

We assume that each vertex $v \in V$ has a power cost $C_v : \mathcal{A} \to \mathbb{R}$ that measures the number of transmissions by that vertex necessary for any routing profile. For our problem, this cost takes on the simplified form $C_v : (V \times V)^n \to \mathbb{R}$, which depends only on each player's transmission through the vertex v. Before defining the structure of the cost function for the reverse carpooling setting, we first introduce some notation. For a given routing profile $a \in \mathcal{A}$, let $\sigma(a, v[v_x, v_y])$ be defined as the number of players sending information from v_x to v_y through v, i.e.,

$$\sigma(a, v[v_x, v_y]) := |\{i \in N : v[v_x, v_y] \in I(a_i)\}|.$$

Using only reverse carpooling codes, the transmission cost at node $v \in V$ for profile $a \in \mathcal{A}$ is defined as

$$C_v(a) := \sum_{(v_x, v_y) \in \mathcal{N}(v)^2 : x \geq y} \max\left(\sigma\left(a, v\left[v_x, v_y\right]\right),\ \sigma\left(a, v\left[v_y, v_x\right]\right)\right). \quad (3)$$

Hence, the power consumption at a vertex depends only on the number of players using each transmission through the vertex. The system cost for a routing profile $a \in \mathcal{A}$ is

$$C(a) := \sum_{v \in V} C_v(a) \quad (4)$$

In general, a global planner would like to use a profile $a \in \mathcal{A}$ that minimizes the system cost.

Given the limited coding structure considered in this paper, it is possible to solve for the optimal routing profile in a centralized fashion using linear programming. However, in large networks, centrally controlling all entities in the network may be infeasible. Therefore, we seek to address the degree to which localizing decisions impacts performance.

IV. The Reverse Carpooling Problem: A State Based Game Formulation

In this section, we model the reverse carpooling problem as a state based game by assigning each player a local cost function. We restrict our attention to cost functions where each player is charged a fixed cost for each node in his path where he carpools and a different fixed cost for each node in his path where he does not carpool. This cost function does not fall into the traditional non-cooperative game setting because a player's cost function is not only influenced by other players, but also by whether the player has an opportunity to carpool. We model this scenario as a state based game where the state captures the notion of carpooling assignments. We show that this cost structure leads to a state based potential game.

We define a set of states, X, that identifies reverse carpooling assignments over the network. For each $v, v', v'' \in V$ with $v \in \mathcal{N}(v')$ and $v'' \in \mathcal{N}(v)$, state x describes the players using transition $v[v', v'']$ and their order of priority for any carpooling opportunities that may arise. For a given allocation $a \in \mathcal{A}$, we define the set of admissible states as $X(a) \subset X$ where $X(a)$ is nonempty and a state $x \in X(a)$ defines for each $v[v', v'']$ an order of priority for the players using transition $v[v', v'']$ in the allocation a. The oder of priority for each $v[v', v'']$ is described by a queue.

If one player changes his action, the state evolves deterministically as follows: (i) if a player stops using transition $v[v', v'']$, then each player in the queue behind him moves forward one spot in the queue and (ii) if a player joins a queue, then the player becomes the last individual in that queue. If multiple players seek to enter a queue simultaneously, the order of the entering players is randomly chosen. The state dynamics satisfy (2). We refer to these state dynamics as *first in first out (FIFO)*.

For a given profile $a \in \mathcal{A}$ and an admissible state $x \in X(a)$, let $N_i^C(a, x)$ represent the number of transmissions for which player i carpools and let $N_i^{NC}(a, x)$ represent the number of transmissions for which player i does not carpool. Carpooling assignment are made first-come, first-serve using the given state of the queue. Priority is given by the ordering in the queue. For example, suppose player 1 uses the transmission $v[v', v'']$ and players 2 and 3 use the transmission $v[v'', v']$. If player 2 has a higher priority than player 3, then players 1 and 2 would carpool while player 3 would not. This definition implies that $N_i^C(a, x) + N_i^{NC}(a, x) = |I(a_i)| = |a_i| - 1$ for any admissible action state pair $[a, x] \in \mathcal{A} \times X(a)$. We assign each player $i \in N$ the cost function

$$J_i(a, x) = N_i^C(a, x) + \alpha\ N_i^{NC}(a, x), \quad (5)$$

for some $\alpha > 0$. Therefore, each player pays a cost of 1 for each transmission where he carpools and a cost of α for each transmission where he does not carpool. Note that for the special case of $\alpha = 2$

$$C(a) = \frac{1}{2} \sum_{i \in N} J_i(a, x),$$

where $C(a)$ is the system cost in (4).

In any state based game, each player is assigned a state based cost function over all admissible action state pairs according to the action profile. We extend (5) as follows: for any player $i \in N$, admissible action state pair $[a, x] \in \mathcal{A} \times X(a)$, and action profile $a' \in \mathcal{A}$,

$$J_i(a', x) = \mathbf{E}_{x'} J(a', x'),$$

where the state x' is chosen randomly according to $P(a', x)$. We refer to the state based game formulation of the reverse carpooling problem as the *reverse carpooling game*.

Theorem 1. *The reverse carpooling game with state dependent cost functions as defined in (5) and FIFO state dynamics is a state based potential game for any α.*

Proof: Consider the following potential function $\phi : \mathcal{A} \to \mathbb{R}$

$$\phi(a) = (\alpha - 1)\, C(a) + \sum_{i \in N} |I(a_i)|. \tag{6}$$

We now show that for any admissible action state pair $[a, x] \in \mathcal{A} \times X(a)$, player $i \in N$, and action $a_i' \in \mathcal{A}_i$,

$$J_i(a_i', a_{-i}, x) - J_i(a, x) < 0 \;\Rightarrow\; \phi(a_i', a_{-i}) - \phi(a) < 0.$$

For notational ease, let $a' = (a_i', a_i)$ and x' be the state selected according to the deterministic transition $P(a', x)$. Note that $J_i(a', x') = J_i(a', x)$.

We first focus on the system cost when player i switches from a_i to a_i'. Let a_i^0 be defined as the common vertices (with consistent transmissions) between a_i and a_i', i.e., $I(a_i^0) = I(a_i) \cap I(a_i')$. We bound $C(a') - C(a)$ by bounding $C(a) - C(a_i^0, a_{-i})$ and $C(a') - C(a_i^0, a_{-i})$. When player i switches from a_i to a_i^0, the system cost decreases by at least

$$C(a) - C(a_i^0, a_{-i}) \geq N_i^{NC}(a, x) - N_i^{NC}(a_i^0, a_{-i}, x), \tag{7}$$

since removing a player from a vertex at which the player was not carpooling decreases the system cost by exactly 1. Removing a player from a vertex at which the player was carpooling may reduce the system cost by 1 or leave it unchanged depending on the state. Hence, (7) is an upper bound on the net change in the system cost when player i switches from a_i to a_i^0. When player i switches from a_i^0 to a_i', the system cost increases by exactly

$$C(a') - C(a_i^0, a_{-i}) = N_i^{NC}(a', x') - N_i^{NC}(a_i^0, a_{-i}, x) \tag{8}$$

because player i enters the end of the queue for each transmission in his new path segment $I(a') \setminus I(a_i^0)$. Therefore, combining (7) and (8) results in

$$C(a') - C(a) \quad \leq \quad N_i^{NC}(a', x') - N_i^{NC}(a, x). \tag{9}$$

Focusing on the second half of (6), we have that

$$\sum_{i \in N} |I(a_i')| - \sum_{i \in N} |I(a_i)| = \tag{10}$$
$$N_i^C(a', x') + N_i^{NC}(a', x') - N_i^C(a, x) - N_i^{NC}(a, x).$$

Plugging (9) and (10) into (6) and simplifying we obtain

$$\begin{aligned} &\phi(a') - \phi(a) \\ &\leq \alpha \left(N_i^{NC}(a', x') - N_i^{NC}(a, x) \right) + N_i^C(a', x') - N_i^C(a, x), \\ &= J_i(a', x') - J_i(a, x). \end{aligned}$$

■

V. Efficiency – Price of Anarchy and Price of Stability

The results of Theorem 1 guarantee the existence of an equilibrium in any reverse carpooling game with cost functions as defined in (5). We gauge the efficiency of equilibria using the well known worst case measures called *price of anarchy (PoA)* and *price of stability (PoS)* [15]. In terms of the reverse carpooling game, the PoA gives an upper bound on the system cost achieved by any equilibrium while the PoS gives an upper bound on the system cost of the best equilibrium for any reverse carpooling game. Specifically, let $\mathcal{G}$ denote the set of reverse carpooling games. For any particular game $G \in \mathcal{G}$ let $\mathcal{E}(G)$ denote the set of equilibria, $PoA(G)$ denote the price of anarchy, and $PoS(G)$ denote the price of stability for the game G, where

$$\begin{aligned} PoA(G) &:= \max_{a \in \mathcal{E}(G)} \frac{C(a)}{C(a^{\text{opt}})}, \\ PoS(G) &:= \min_{a \in \mathcal{E}(G)} \frac{C(a)}{C(a^{\text{opt}})}, \end{aligned}$$

and $a^{\text{opt}} \in \arg\min_{a^* \in \mathcal{A}} C(a^*)$. We define the PoA and PoS for the family of reverse carpooling games as $PoA := \inf_{G \in \mathcal{G}} PoA(G)$ and $PoS := \inf_{G \in \mathcal{G}} PoS(G)$.

Theorem 2. *Consider all reverse carpooling games with cost functions as in (5) and FIFO state dynamics. The price of anarchy is*

$$PoA(\alpha) = \begin{cases} \frac{2}{\alpha}, & \alpha \in (0, 1), \\ 2, & \alpha \in [1, 2], \\ \alpha, & \alpha \in (2, \infty). \end{cases}$$

Proof: Before beginning, we define player i's contribution to the system cost as

$$V_i(a, x) := \frac{1}{2} N_i^C(a, x) + N_i^{NC}(a, x).$$

Using this definition, for any admissible action state pair $[a, x]$, we have $C(a) = \sum_{i \in N} V_i(a, x)$.

Let $a^{\text{opt}} \in \arg\min_{a \in \mathcal{A}} C(a)$ be any optimal assignment that minimizes the system cost. Let $[a^{\text{ne}}, x^{\text{ne}}]$ be any state based Nash equilibrium. Trivially, the total system cost of a state based Nash equilibrium is higher than the system cost of an optimal assignment, i.e., $C(a^{\text{ne}}) \geq C(a^{\text{opt}})$. The system cost at the optimal allocation satisfies

$$C(a^{\text{opt}}) \geq \sum_{i \in N} \frac{1}{2} \min_{a_i \in \mathcal{A}_i} |I(a_i)|. \tag{11}$$

This bounds represents the system cost resulting from an assignment where each player selects his shortest path, i.e., path consisting of the minimal number of nodes, and reverse carpools at each transmission in his path. Note that achieving this cost in general is not possible since a player can never carpool at his originating vertex. However, in the ensuing analysis, we ignore these “edge” effects, thereby deriving loose bounds. Note that (11) is independent of α. However, an upper bound on the system cost associated with a state base Nash equilibrium is highly dependent on α. Because of the lack of space, we will only analyze the case when $\alpha \in [1, 2]$. The other cases can be analyzed in a similar fashion.

For any $\alpha \in [1, 2]$ and any equilibrium, each player's cost satisfies

$$J_i(a^{\text{ne}}, x^{\text{ne}}) \leq \alpha \min_{a_i \in \mathcal{A}_i} |I(a_i)|, \tag{12}$$

where $\alpha \min_{a_i \in \mathcal{A}_i} |I(a_i)|$ represents the cost player i would incur if the player selected his shortest path and was unable to carpool at any vertex in the path a_i. Since $J_i(a, x) = N_i^C(a, x) + \alpha N_i^{NC}(a, x)$, this bound implies $N_i^C(a^{\text{ne}}, x^{\text{ne}}) + \alpha N_i^{NC}(a^{\text{ne}}, x^{\text{ne}}) \leq \alpha \min_{a_i \in \mathcal{A}_i} |I(a_i)|$, which in turn implies the following bound on a player's contribution to the system cost

$$V_i(a^{\rm ne}, x^{\rm ne}) \leq \min_{a_i \in \mathcal{A}_i} |I(a_i)|.$$

Therefore, the system cost of an equilibrium is bounded above by

$$C(a^{\rm ne}) \quad \leq \quad \sum_{i \in N} \min_{a_i \in \mathcal{A}_i} |I(a_i)|. \qquad (13)$$

Combining (11) with (13), we obtain

$$\sum_{i\in N} \min_{a_i \in \mathcal{A}_i} |I(a_i)| \;\geq\; C(a^{\rm ne}) \;\geq\; C(a^{\rm opt}) \;\geq\; \sum_{i \in N} \frac{1}{2} \min_{a_i \in \mathcal{A}_i} |I(a_i)|.$$

This implies that for any $\alpha \in [1,2]$ the price of anarchy satisfies $PoA(\alpha) \leq 2$.

Before proving the tightness of this bound, we introduce the following approximation. Figure 2 illustrates 2 players traversing in opposite directions over a series of nodes. If there are m interior nodes, the system cost associated with this transmission is $m+2$ and each player's cost is $\alpha + m$, as each player carpools on the interior m nodes but does not carpool on the boundary nodes. For large m, the constants of 2 and α are insignificant. Normalizing by m, we say that this scenario approximately yields a system cost of 1 and player cost of 1.

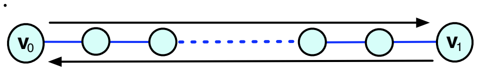

Fig. 2. Illustration of Cost Approximations

To see that this bound is tight, consider the example in Figure 3 where each edge segment consists of several internal nodes as illustrated in Figure 2. This example consists of 8 players located at the 8 exterior nodes. Each player's destination is the exterior node two hops clockwise from the player's source. A player's viable paths to his destination are (i) traverse two segments on the exterior or (ii) traverse two segments through the interior. If each player traverses on the exterior, utilizing our approximation the system cost is 16 and each player's cost is 2α as there are no carpooling opportunities for any player. If a player unilaterally switched to his interior path, his cost is still 2α since he still has no carpooling opportunities. Therefore, this is an equilibrium. If each player traverses through the interior, this is also an equilibrium. The system cost of this scenario is 8 and each player's cost is 2 since he carpools on both segments. Therefore, this example exhibits a price of anarchy of 2 for any $\alpha \in [1,2]$. ■

Theorem 2 is a negative result in the sense that it demonstrates that the worst case equilibrium can achieve a cost two times the cost of the optimal solution when $\alpha \in [1,2]$ (and even higher when $\alpha \notin [1,2]$). Since the cost associated with players choosing their shortest paths and applying no network coding is at most twice the cost of the optimal solution, the price of anarchy achieved here is very high. It turns out that this negative result is not a result of the particular cost function chosen in (5). Let $x_{v[v',v'']}$ be the component of the state for the transmission $v[v',v'']$. Consider any player cost function of the form

$$J_i(a,x) = \sum_{v[v' \to v''] \in I(a_i)} f\left(x_{v[v' \to v'']}, x_{v[v'' \to v']}\right). \qquad (14)$$

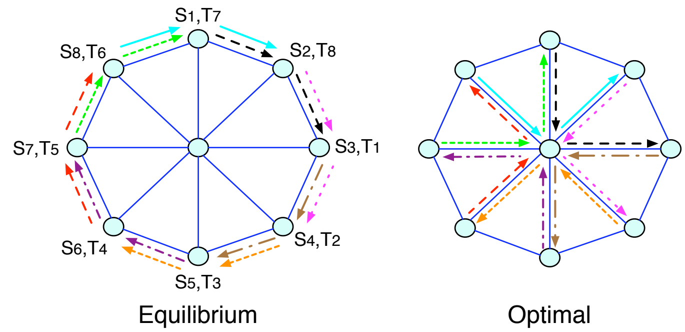

Fig. 3. Tightness of Price of Anarchy for $\alpha \in [1,2]$

The structure of the cost function in (14) is such that the cost of utilizing an edge is solely based on the players using that edge. The cost is not based on the edge itself. We will call such cost functions *anonymous cost functions*. Note that (5) is a special case of (14).

Theorem 3. *For any reverse carpooling game with anonymous cost functions of the form (14), the price of anarchy is at least 2.*

Proof: Consider the example in Figure 4 which is a slight variant of the example in Figure 3. There are 8 players with the same source and terminals as indicated in Figure 3. Each player only has the option of taking an exterior path or interior path, each consisting of 2 segments. The interior and exterior paths for player 1 are illustrated. The difference from the example in Figure 3 is that the exterior paths of different players are now disjoint. If the cost function is anonymous and if all players take their exterior path as highlighted, then we have an equilibrium. At this equilibrium, each player is alone on a path consisting of 2 segments. If a player unilaterally switched to his interior path from this equilibrium, the player would also be alone on a path consisting of 2 segments; hence it would result in the same cost since cost functions are anonymous. All players taking their exterior path results in a system cost of 16. All players taking their interior path results in a system cost of 8. Therefore, the price of anarchy is 2. ■

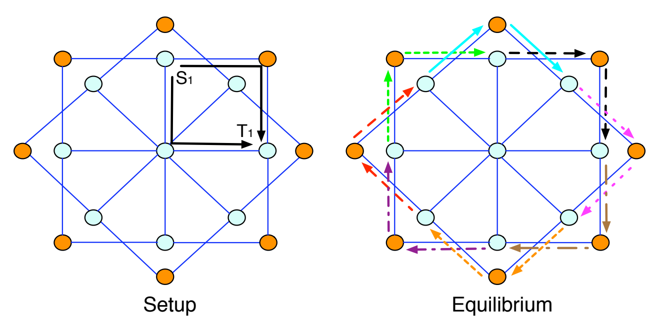

Fig. 4. Price of Anarchy of 2 for Anonymous Cost Functions

Our next result is more positive, showing in particular when $\alpha = 2$ that the best equilibrium has a cost at most 50% higher than the cost of the optimal solution. The existence of such an equilibrium point is desirable because there exists distributed learning algorithm that guarantee convergence to the best-case equilibrium, i.e., the minimizer of the potential function, with probability approaching 1 [12], [16], [17].

Theorem 4. *Consider all reverse carpooling games with cost functions as in (5) and FIFO state dynamics. The price of stability is*

$$PoS(\alpha) = \frac{\alpha+1}{\alpha}.$$

Proof: Any reverse carpooling with cost functions as in (5) and FIFO state dynamics is a state based potential game with potential function $\phi(a) = (\alpha-1)\,C(a) + \sum_{i\in N}|a_i|$. The existence of this potential function implies that any action profile that minimizes the potential function is also an equilibrium, i.e., any action profile $a \in \arg\min_{a\in\mathcal{A}} \phi(a)$ and admissible state $x \in X(a)$ constitutes an equilibrium. We next use this fact to bound the price of stability.

Expanding out the second component of $\phi(a)$ we have

$$\begin{aligned}\sum_{i\in N}|a_i| &= \sum_{i\in N}\left(N_i^C(a,x) + N_i^{NC}(a,x)\right), \\ &= 2C(a) - \sum_{i\in N} N_i^{NC}(a,x).\end{aligned}$$

Plugging into $\phi(a)$, we have

$$\begin{aligned}\phi(a) &= (\alpha-1)\,C(a) + 2C(a) - \sum_{i\in N} N_i^{NC}(a,x), \\ &= (\alpha+1)\left(C(a) - \frac{1}{\alpha+1}\sum_{i\in N} N_i^{NC}(a,x)\right). \quad (15)\end{aligned}$$

Let $a^{\text{ne}} \in \{a \in \mathcal{A} : \phi(a) = \min_{\tilde{a}\in\mathcal{A}} \phi(\tilde{a})\}$ be a Nash equilibrium that minimizes the potential function and $a^{\text{opt}} \in \{a \in \mathcal{A} : C(a) = \min_{\tilde{a}\in\mathcal{A}} C(\tilde{a})\}$ be an optimal action profile. Then $C(a^{\text{ne}}) \geq C(a^{\text{opt}})$. Also, since a^{ne} minimizes ϕ,

$$\phi(a^{\text{opt}}) \geq \phi(a^{\text{ne}}). \quad (16)$$

Rewriting (16) in terms of (15) we obtain

$$\begin{aligned}&C(a^{\text{opt}}) - \sum_{i\in N}\left(\frac{1}{\alpha+1}N_i^{NC}(a^{\text{opt}})\right) \\ &\quad \geq C(a^{\text{ne}}) - \sum_{i\in N}\left(\frac{1}{\alpha+1}N_i^{NC}(a^{\text{ne}})\right),\end{aligned}$$

which simplifies as

$$\begin{aligned}C(a^{\text{ne}}) - C(a^{\text{opt}}) &\leq \frac{1}{\alpha+1}\sum_{i\in N}\left(N_i^{NC}(a^{\text{ne}}) - N_i^{NC}(a^{\text{opt}})\right), \\ &\leq \frac{1}{\alpha+1}\sum_{i\in N} N_i^{NC}(a^{\text{ne}}), \\ &\leq \frac{1}{\alpha+1}C(a^{\text{ne}}).\end{aligned}$$

Therefore

$$\frac{C(a^{\text{ne}})}{C(a^{\text{opt}})} \leq \frac{\alpha+1}{\alpha}.$$

It is straightforward to show that this bound is tight. ■

Figure 5 compares how the price of anarchy and price of stability vary with α. Notice that as $\alpha \to \infty$, we get a price of stability of 1 which is desirable; however, the price of anarchy is unbounded.

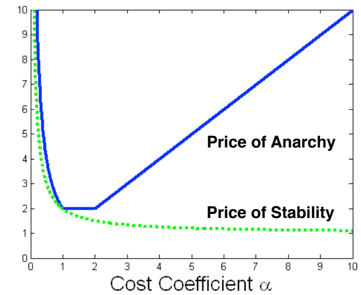

Fig. 5. Summary Price of Anarchy and Price of Stability

VI. Conclusion

We investigate the role of game theory as a distributed algorithm for dealing with multi-session coding problem. We formulate the reverse carpooling problem as a state based game and prove the existence and efficiency of equilibrium for a family of cost functions. For a particular cost function ($\alpha = 2$), our efficiency results proved that the best equilibrium has a system cost of at most 50% higher than the optimal system cost irrespective of the network structure or demands. While this paper focuses on the simple network coding structure of reverse carpooling, the theoretical foundations of state based games could easily be extended to alternative coding schemes and system costs.

References

[1] D. Traskov, N. Ratnakar, D. S. Lun, R. Koetter, and M. Medard, "Network coding for multiple unicasts: An approach based on linear optimization," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, July 2006.

[2] S. Katti, D. Katabi, W. Hu, H. Rahul, and M. Médard, "The importance of being opportunistic: practical network coding for wireless enviornments," in *43rd Allerton Annual Conference on Communications, Control, and Computing*. Monticello, IL: IEEE, Sept. 2005, invited paper.

[3] J. R. Marden, G. Arslan, and J. S. Shamma, "Connections between cooperative control and potential games," *IEEE Transactions on Systems, Man and Cybernetics. Part B: Cybernetics*, 2008, submitted.

[4] D. Fudenberg and D. Levine, *The Theory of Learning in Games*. Cambridge, MA: MIT Press, 1998.

[5] H. P. Young, *Individual Strategy and Social Structure*. Princeton, NJ: Princeton University Press, 1998.

[6] ——, *Strategic Learning and its Limits*. Oxford University Press, 2005.

[7] J. Price and T. Javidi, "Network coding games with unicast flows," *IEEE Journal on Selected Areas in Communications*, 2008, to appear.

[8] S. Bhadra, S. Shakkottai, and P. Gupta, *IEEE Transactions on Information Theory*, vol. 52, no. 11, pp. 5077–5087, 2006.

[9] D. S. Lun, N. Ratnakar, R. Koetter, M. Medard, E. Ahmed, and H. Lee, "Achieving minimum-cost multicast: A decentralized approach based on network coding," in *Proceedings of IEEE INFOCOM*, 2005, pp. 1607–1617.

[10] Y. E. Sagduyu and A. Ephremides, "Joint optimization of mac and network coding for cooperative and competitive wireless multicasting," in *Proceedings of Workshop on Modelling and Design of Multi-hop Wireless Networks*, 2006.

[11] D. Fudenberg and J. Tirole, *Game Theory*. Cambridge, MA: MIT Press, 1991.

[12] J. R. Marden and M. Effros, "State based potential games with applications to network coding," 2008, in preparation.

[13] L. S. Shapley, "Stochastic games," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 39, no. 10, pp. 1095–1100, 1953.

[14] D. Monderer and L. Shapley, "Potential games," *Games and Economic Behavior*, vol. 14, pp. 124–143, 1996.

[15] T. Roughgarden, *Selfish Routing and the Price of Anarchy*. Cambridge, MA, USA: MIT Press, 2005.

[16] J. R. Marden and J. S. Shamma, "Revisiting log-linear learning: Asynchrony, completeness and a payoff-based implementation,," *Game and Economic Behavior*, 2008, submitted.

[17] L. Blume, "The statistical mechanics of strategic interaction," *Games and Economic Behavior*, vol. 5, pp. 387–424, 1993.