A Non-Monetary Protocol for Peer-to-Peer Content Distribution in Wireless Broadcast Networks with Network Coding

I-Hong Hou, Yao Liu, and Alex Sprintson
CESG and Department of ECE
Texas A&M University
College Station, TX 77843, USA
Email:{ihou,steveliu,spalex}@tamu.edu

***Abstract*—This paper studies the problem of content distribution in wireless peer-to-peer networks with selfish nodes. In this problem a group of wireless nodes exchange data over a lossless broadcast channel. Each node aims to increase its own download rate and minimize its upload rate. We propose a distributed protocol that provides incentives for the nodes to participate in the content distribution. Our protocol does not require any exchange of money, reputation, etc., and hence can be easily implemented without additional infrastructure. Moreover, our protocol can be easily modified to employ network coding.**

Focusing on the important case in which the system contains two files that need to be distributed, we derive a closed-form expression of Nash Equilibria. We also derive the prices of anarchy, both from each node's perspective and the whole system's perspective. Furthermore, we propose a distributed mechanism where the strategy of each node is only based on the local information and show that the mechanism converges to a Nash Equilibrium. We also introduce an approach for calculating Nash Equilibria for systems that incorporate network coding when more than two files need to be distributed.

I. Introduction

Recently, there has been a significant interest in using wireless peer-to-peer (P2P) networks to distribute information between mobile devices. The peer-to-peer content distribution can improve the system performance in many different ways. For example, in cellular networks it is often the case that mobile phones can retrieve the required information from their peers instead of downloading it from the remote servers. Since exchanging data between local devices requires much less power and results in less interference to other devices, such an approach may reduce power consumption and increase spatial reuse.

Many existing studies (e.g., [1]–[3]) have demonstrated the benefits of wireless P2P networks. However, these studies have assumed that all nodes are cooperative and do not require additional incentives to cooperate. In practice, nodes may be selfish and have little incentive to help other nodes to obtain the data. Therefore, a major challenge for wireless P2P networks is to provide incentives to nodes in the network so that they are willing to contribute to the network by sharing their data with other nodes. While there are many studies, such as [4]–[6], on this topic for wired P2P networks, these works cannot be applied to wireless P2P networks. Due to the broadcast nature of the wireless medium, when a node transmits a packet, all nodes within the proximity are able to receive the packet. Therefore, in wireless P2P networks, data exchange involves all nodes within the system, rather than only two nodes as in wired P2P networks.

In this paper, we study wireless P2P networks composed of selfish nodes. We first provide a model that considers the broadcast nature of wireless transmissions and the incentives of selfish nodes. Each node in the system aims to increase its download rate and decrease its upload rate, so as to reduce its own power consumption. We then propose a protocol for content distribution for this setting. Our protocol does not require the exchange of money, reputation, etc., and hence can be implemented without the need of additional infrastructure. This non-monetary feature further distinguishes our work from other studies that rely on additional infrastructure to set prices or payoffs [6]–[8], or to punish uncooperative nodes [9]. Moreover, our protocol can be easily modified to employ network coding.

We provide a detailed performance analysis for our protocol. For the practically important case with two files in the system, we derive closed-form expressions for each node's strategies under a series of Nash Equilibria. We also derive the prices of anarchy under these Nash Equilibria, both from a node's selfish perspective and the whole system's perspective.

To compute its strategy under a Nash Equilibrium, a node needs information of all other nodes, which is not always available to it. To address this challenge, we propose a distributed mechanism where each node updates its strategy only based on its private information and the history of the system. We show that this distributed mechanism converges to a Nash Equilibrium. Moreover, this mechanism is also consistent with each node's incentive, as the expected cost of each node reduces with each update.

We then consider systems that have more than two files

The work of Alex Sprintson was partially supported the NSF under grant CNS-0954153 and by the AFOSR under contract No. FA9550-13-1-0008.

and employ network coding. We propose a systematic approach to compute the Nash Equilibria. The performance of such systems is further investigated through numerical studies.

The rest of the paper is organized as follows. Section II proposes our system model and protocol for content distribution. Section III studies the Nash Equilibria for systems with only two files. Section IV studies the prices of anarchy under these Nash Equilibria. Section V discusses implementation issues and provides a distributed mechanism for nodes to update their strategies. Section VI studies the Nash Equilibria when there are more than two files in systems that employ network coding. Section VII provides some numerical results. Finally, Section VIII concludes the paper.

II. System Model and Protocol Overview

We consider the *direct data exchange problem* [10] in which a group of wireless nodes are collaborating to exchange the set of files $X = \{A, B, C, \dots\}$. Each node has a subset of files in X available to it and needs to obtain all other files in X. The nodes use a lossless broadcast channel to transmit files to other nodes. We assume that the files are very large, hence a large number of packets need to be broadcasted over the channel to deliver every file to other nodes. For clarity of presentation, we assume that each file contains infinitely many packets. For long files, this assumption incurs a very small penalty in terms of the accuracy of the obtained results.

In this paper we focus on the special case in which each node n requires a single file, denoted by X_n. This is an important case that captures many of the salient features of the problem at hand. This case is of a significant practical importance since in many settings wireless nodes need to recover only a small number of packets, (e.g., packets that are lost due to fading or interference).

The broadcasting nature of wireless transmissions may result in a "*free rider*" problem as illustrated in the following example. Due to the free-rider problem, the existing protocols designed for wired networks cannot be applied directly to the wireless setting.

Example 1: Consider a system with four nodes and two files where $X_1 = X_2 = A$ and $X_3 = X_4 = B$. Suppose that node 1 and node 3 exchange data, that is, node 1 transmits packets of file B, and node 3 transmits packets of file A in return. As all nodes can receive all transmissions, node 2 and node 4 can obtain packets of A and B without transmitting any packet. Therefore, we say that nodes 2 and 4 are *free riders*.

In addition to being unfair, the possibility of being a free rider may prevent selfish nodes from transmitting data and contributing to the network. In the above example, each node may refrain from transmitting data, in the hope that other nodes participate in exchanging data, making itself a free rider.

In this paper, we propose a non-monetary protocol for P2P content distribution and study its performance when all nodes are selfish. Before introducing the protocol, we first formally describe the goal of each node.

A. Definitions

We define the time needed for transmitting a data packet to be one unit time. Let $r_n(k)$ denote the time when node n downloads the k-th packet from the required file. We define the *download rate* of node n as $\liminf_{k\to\infty} \frac{k}{r_n(k)}$. The download rate captures the average number of packets that node n downloads per time unit.

The goal of each node is to maximize its download rate and, at the same time, to reduce the number of transmissions it makes. To be more specific, we assume that whenever node n transmits a data packet, it needs to pay a *transmission cost* g_n. The transmission cost can be chosen, for example, to reflect the amount of power needed for making a transmission. On the other hand, node n pays a *waiting cost* w_n per unit time through the course of the protocol.

Let $\alpha_n(k)$ be the number of transmissions that node n makes until time $r_n(k)$. Consider the time period $(r_n(k-1), r_n(k)]$. During this time period, node n receives one packet and transmits $\alpha_n(k) - \alpha_n(k-1)$ packets. Thus, the *total cost* for node n to download packet k during this time period is $(\alpha_n(k) - \alpha_n(k-1))g_n + (r_n(k) - r_n(k-1))w_n$. The *average total cost* of node n is then defined as the long-term average total cost per downloaded packet.

Definition 1: The average total cost of node n is defined as

$$\limsup_{k\to\infty} \frac{\alpha_n(k)g_n + r_n(k)w_n}{k}.$$

The goal of each node is to minimize its average total cost. More specifically, node n wishes to reduce $\limsup_{k\to\infty} \frac{\alpha_n(k)}{k}$ in order to achieve small transmission cost. On the other hand, $\limsup_{k\to\infty} \frac{r_n(k)}{k}$ is the average delay between two successive packet downloads, which is also the inverse of the download rate. Therefore, to have small delay and high download rate, node n would like to minimize $\limsup_{k\to\infty} \frac{r_n(k)}{k}$.

B. Protocol description

We now describe our protocol for P2P content distribution. The process consists of *rounds* such that during each round one packet from A and one packet from B is broadcasted over the channel. At the beginning of a round, each node n secretly picks a back-off time, τ_n. Node n then waits and listens to the channel for τ_n time. If no transmissions take place in τ_n time, node n transmits a control packet that contains the name X_n of the file required by n. The control packet can be interpreted as an obligation of node n to transmit the next packet from a file held by n if any other node transmits the next packet from file X_n (the *next* packet refers to the packet which has not been broadcasted over the channel). Since the control packets are very small, we assume that their transmission does not incur any cost and that the time required for transmission of such packets is negligible. We refer to the time between the beginning of the round and transmission of the control packet as the *initiation phase*.

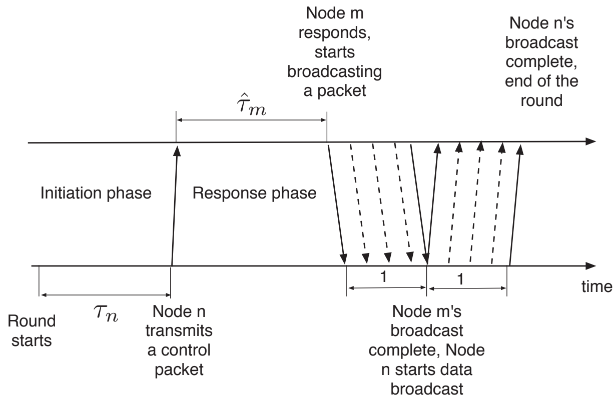

Fig. 1: An example of a round.

After node n transmits the control packet, every node m that has file X_n and whose required packet X_m is available at node n secretly and randomly picks a back-off timer $\hat{\tau}_m$. Every node m then waits and listens to the channel for $\hat{\tau}_m$ time. If no other nodes transmit in $\hat{\tau}_m$ time, node m transmits a data packet of X_n, and piggybacks its value of X_m. Upon receiving the data packet from node m, node n responds with a packet of X_m, as promised in its control packet. The round ends after node n completes broadcast of packet X_m and a new round begins. We refer to the time interval between the end of the initiation phase and the transmission of a packet by node m as the *response phase*. Note that τ_n and $\hat{\tau}_m$ are lengths of the initialization and response phase, respectively. The protocol execution is demonstrated in Fig. 1.

Intuitively, when a node n chooses large values of τ_n and $\hat{\tau}_n$, it is likely that node n does not transmit, which increases its chance of being a free rider and reduces its transmission cost. However, large values of τ_n and $\hat{\tau}_n$ might result in large waiting times for n, which, in turn, might result in a large waiting cost. By taking waiting costs into account, our protocol provides incentive for the nodes to choose reasonably small values of τ_n and $\hat{\tau}_n$, and hence enables the nodes to exchange data in an efficient manner. We also note that our protocol is non-monetary and can be easily implemented for modern wireless networks without the need of additional infrastructure.

Finally, we show that our protocol can be modified for the settings in which the network coding technique is used. The modified protocol is very similar to the one described above. Still, at each round, exactly two data packets will be transmitted. With the network coding technique, the packet broadcasted at the response phase will be a linear combination of the packets available at the transmitting node. For example, suppose that X contains three files A, B, C and node n requires file A, i.e., $X_n = A$. Then, in the response phase, node n can send a combination of the next packet from B and the next packet from C (the combination can be a linear operation underlying finite field). The use of the network coding technique will increase the number of nodes that benefit from each transmission, but can also lead to more free riders.

C. Price of Anarchy

Under our protocol, each node's strategy consists of two parts: choosing a distribution to generate its back-off timer in the initiation phase, τ_n, and choosing a distribution to generate its back-off timer in the response phase, $\hat{\tau}_n$. We say that the strategies of all nodes in the system form a *Nash Equilibrium* if, for each node n, its strategy minimizes its average cost, given the strategies of all other nodes.

In this paper, we analyze the performance of our protocol under Nash Equilibrium. We consider the performance from both nodes' perspectives and the system's perspective. A node's performance is based on its average cost. On the other hand, the system's performance is characterized by the average download rate of all nodes. We also define the *price of anarchy on node cost* and *price of anarchy on per-node download rate* by comparing the average cost and download rate under the Nash Equilibrium with the average code and download rate that can be achieved under a cooperative scenario. In the cooperative scenario, we assume that there are two nodes, n and m, that choose $\tau_n = 0$ and $\hat{\tau}_m = 0$, respectively, in each round. Therefore, all other nodes do not transmit, and the length of each round is two time units. The prices of anarchy are defined as follows:

Definition 2: The *price of anarchy on node cost* for node n under a Nash Equilibrium is the ratio of the average total cost of n under the Nash Equilibrium and that under a cooperative scenario where n never transmits.

Definition 3: The *price of anarchy on per-node download rate* under a Nash Equilibrium is defined as the ratio of per-node download rate under a cooperative scenario and per-node download rate under the Nash Equilibrium.

III. Nash Equilibria for Bilateral File Exchanges

In this section, we analyze the performance of our protocol. We will focus on the special case in which there are only two files, A and B in the system. We can divide nodes into two groups, nodes in one group, indexed by $a_1, a_2, \dots, a_I$, have the file B and need the file A, i.e., $X_{a_i} = A$, while nodes in the other group, indexed by $b_1, b_2, \dots, b_J$ have $X_{b_j} = B$. In this setting, network coding is not employed.

We will show that there is a series of Nash Equilibria where each node n chooses each of its back-off timers in the initiation phase and response phase as an exponential random variable. We focus on exponential random variables due to its memoryless property, which makes it easy to implement.

Let $d(k)$ be the duration of the k-th round. Let $u_n(k)$ be the number of transmissions that n makes in the k-th round. Since there are only two files in the system, every node downloads one packet in each round. Therefore, in this setting the average total cost of node n (as defined by Definition 1) is equal to

$$\limsup_{K\to\infty} \frac{\sum_{k=1}^{K}[u_n(k)g_n + d(k)w_n]}{K}.$$

Moreover, since each node uses the same strategy in each round, we also have

$$\limsup_{K\to\infty}\frac{\sum_{i=1}^{K}[u_n(i)g_n+d(i)w_n]}{K}=E[u_n(k)g_n+d(k)w_n]$$

which is the expected total cost of n in round k, for each k. Therefore, in this section, we derive the average total cost of a node by considering its expected total cost in a round.

Our analysis consists of two parts: We first derive the nodes' strategies and the system behavior in the response phase. We then derive those in the initiation phase by taking into account of the system behavior in the response phase. We apply the following theorem to obtain Nash Equilibrium:

Theorem 1: [11, Lemma 3.1] The strategies of nodes in a game form a Nash Equilibrium if, for each node n, given the strategies of other nodes, the expected cost of node n is the same regardless of its chosen values of back-off timers.

A. Analysis for the Response Phase

We first consider the nodes' strategies on choosing the back-off timers in the response phase, $\hat{t}_n$. Without loss of generality, we assume that node b_1 has sent the control packet in the initiation phase. Each node a_i secretly and randomly chooses a back-off timer $\hat{t}_{a_i}$. Assume that the timer chosen by node a_{i^*} is the smallest among all back-off timers, that is, $\hat{t}_{a_{i^*}}=\min\{\hat{t}_{a_1},\hat{t}_{a_2},\dots\}$. Node a_{i^*} will transmit after time $\hat{t}_{a_{i^*}}$, and the *additional waiting time*, which is the length of the response phase, for each node a_i to download a packet is $\hat{t}_{a_{i^*}}$. In the expression of the additional waiting time, we exclude the time needed waiting for b_1 to transmit its control packet and the time needed for transmitting data packets, as these times are not influenced by the values of $\hat{t}_{a_i}$. In the following, we also exclude the waiting cost incurred by the time waiting for b_1 to transmit its control packet and the time needed for transmitting data packets when discussing *additional waiting cost* and *additional total cost*.

We show that there is a Nash Equilibrium where each node a_i chooses its back-off timer as an exponential random variable. In particular, we assume that $\hat{t}_{a_i}$ is exponentially distributed with expectation λ_{a_i}, i.e., $\hat{t}_{a_i}\sim EXP(\lambda_{a_i})$, where the value of λ_{a_i} will be determined in the sequel.

We now apply Theorem 1 to determine the values of λ_{a_i}. Suppose that node a_{i^*} chooses $\hat{t}_{a_{i^*}}=t$. If t is smaller than all $\hat{t}_{a_i}$, $i\neq i^*$, that is, $t<\min_{i\neq i^*}\hat{t}_{a_i}$, node a_{i^*} needs to transmit a packet after waiting for t units of time, and thus the additional total cost for node a_{i^*} on this packet is $g_{a_{i^*}}+w_{a_{i^*}}t$. On the other hand, if $t>\min_{i\neq i^*}\hat{t}_{a_i}$, a node other than a_{i^*} transmits the packet, and node a_{i^*} becomes the free rider. In this case, the additional total cost of node a_{i^*} on this packet is $w_{a_{i^*}}(\min_{i\neq i^*}\{\hat{t}_{a_i}\})$. Note that we have $\min_{i\neq i^*}\{\hat{t}_{a_i}\}\sim EXP(\sum_{i\neq i^*}\lambda_{a_i})$. Let $\lambda_{-i^*}:=\sum_{i\neq i^*}\lambda_{a_i}$. The expected additional total cost of node i^* can be written as:

$$\begin{aligned}&\int_{s=0}^{t}w_{a_{i^*}}s\lambda_{-i^*}e^{-\lambda_{-i^*}s}ds\\&+\int_{s=t}^{\infty}(g_{a_{i^*}}+w_{a_{i^*}}t)\lambda_{-i^*}e^{-\lambda_{-i^*}s}ds\\=&\frac{w_{a_{i^*}}}{\lambda_{-i^*}}+e^{-\lambda_{-i^*}t}(g_{a_{i^*}}-\frac{w_{a_{i^*}}}{\lambda_{-i^*}}).\end{aligned}\tag{1}$$

Thus, if $g_{a_i}=\frac{w_{a_i}}{\lambda_{-i}}$, for all i, the strategies form a Nash Equilibrium. This can be done by choosing

$$\lambda_{a_i}=\frac{1}{I-1}\sum_l\frac{w_{a_l}}{g_{a_l}}-\frac{w_{a_i}}{g_{a_i}},\tag{2}$$

where I is the number of nodes in group $\{a_1,a_2,\dots\}$. We can also conclude that the expected value of $\min_i\{\hat{t}_{a_i}\}$, which is the expected duration of the response phase, given that a node in the group $\{b_1,b_2,\dots\}$ transmits the control packet, is

$$\hat{T}_A:=\frac{1}{\sum_i\lambda_{a_i}}=\frac{I-1}{\sum_i w_{a_i}/g_{a_i}},\tag{3}$$

and that the expected additional total cost of node a_i is g_{a_i}.

Similarly, if a node a_i transmits a control packet, each node b_j selects $\hat{t}_{b_j}\sim EXP(\lambda_{b_j})$, where

$$\lambda_{b_j}=\frac{1}{J-1}\sum_l\frac{w_{b_l}}{g_{b_l}}-\frac{w_{b_j}}{g_{b_j}},\tag{4}$$

where J is the number of nodes in group $\{b_1,b_2,\dots\}$. The expected duration of the response phase is

$$\hat{T}_B:=\frac{1}{\sum_j\lambda_{b_j}}=\frac{J-1}{\sum_j w_{b_j}/g_{b_j}},\tag{5}$$

and the expected additional total cost of node b_j is g_{b_j}.

B. Analysis for the Initiation Phase

Next, we consider the choice of back-off timer in the initiation phase, that is, the choice of t_n for a node n. We will show that there is a Nash Equilibrium where each node a_i selects $t_{a_i}\sim EXP(\gamma_{a_i})$ and each node b_j selects $t_{b_j}\sim EXP(\gamma_{b_j})$.

Assume that a node, say, node a_{i^*}, selects $t_{a_{i^*}}=t$. If t is the smallest timer among all timers, that is, $t<\min_{i\neq i^*}t_{a_i}$ and $t<\min_j t_{b_j}$, node a_{i^*} transmits the control packet after time t. After which time, it needs to wait one of the nodes in $\{b_1,b_2,\dots\}$ to respond with a data packet, and then a_{i^*} needs to transmit a data packet. By the analysis above, we know that the expected time that a_{i^*} waits for one of the nodes in $\{b_1,b_2,\dots\}$ to respond is $\hat{T}_B$. Thus, the expected total cost for node a_{i^*} is $g_{a_{i^*}}+w_{a_{i^*}}(t+\hat{T}_B)+2w_{a_{i^*}}$, where the last term, $2w_{a_{i^*}}$, accounts for the waiting cost caused by transmission delays, as it takes two units time to transmit two data packets.

Next, consider the case that $t>\min\{\min_{i\neq i^*}t_{a_i},\min_j t_{b_j}\}$. We have that $\min_{i\neq i^*}t_{a_i}\sim EXP(\sum_{i\neq i^*}\gamma_{a_i})$ and $\min_j t_{b_j}\sim EXP(\sum_j\gamma_{b_j})$. By the memoryless property of exponential functions, we have that

$$\begin{aligned}&P_{A\backslash\{i^*\}<B}:=\\&Prob\{\min_{i\neq i^*}t_{a_i}<\min_j t_{b_j}|t>\min\{\min_{i\neq i^*}t_{a_i},\min_j t_{b_j}\}\}\\=&\frac{\sum_{i\neq i^*}\gamma_{a_i}}{\sum_{i\neq i^*}\gamma_{a_i}+\sum_j\gamma_{b_j}}.\end{aligned}$$

That is, with probability $P_{A\setminus\{i^*\}<B}$, one of the nodes in $\{a_1, a_2, \dots\}$ other than a_{i^*} transmits the control packet, and, with probability $1 - P_{A\setminus\{i^*\}<B}$, one of the nodes in $\{b_1, b_2, \dots\}$ transmits the control packet. If it is the former case, node a_{i^*} does not need to transmit any packets, and its expected cost is $w_{a_{i^*}}(\min\{\min_{i\neq i^*} t_{a_i}, \min_j t_{b_j}\} + \hat{T}_B + 2)$. If it is the later case, the expected cost is $w_{a_{i^*}}(\min\{\min_{i\neq i^*} t_{a_i}, \min_j t_{b_j}\} + 2) + g_{a_{i^*}}$, since we have shown that the expected additional total cost of a_{i^*} is $g_{a_{i^*}}$. Hence, given that $t > \min\{\min_{i\neq i^*} t_{a_i}, \min_j t_{b_j}\}$, the expected cost is $w_{a_{i^*}}(\min\{\min_{i\neq i^*} t_{a_i}, \min_j t_{b_j}\} + 2) + P_{A\setminus\{i^*\}<B} w_{a_{i^*}} \hat{T}_B + (1 - P_{A\setminus\{i^*\}<B}) g_{a_{i^*}}$.

As we have $\min\{\min_{i\neq i^*} t_{a_i}, \min_j t_{b_j}\} \sim EXP(\sum_{i\neq i^*} \gamma_{a_i} + \sum_j \gamma_{b_j})$, by letting $\gamma_{-i^*} := \sum_{i\neq i^*} \gamma_{a_i} + \sum_j \gamma_{b_j}$, the expected cost of a_{i^*} can be computed as:

$$\begin{aligned}
&\int_{s=0}^{t} [w_{a_{i^*}}(s+2) + P_{A\setminus\{i^*\}<B} w_{a_{i^*}} \hat{T}_B] \gamma_{-i^*} e^{-\gamma_{-i^*} s} ds \\
&+ \int_{s=0}^{t} [(1 - P_{A\setminus\{i^*\}<B}) g_{a_{i^*}}] \gamma_{-i^*} e^{-\gamma_{-i^*} s} ds \\
&+ \int_{s=t}^{\infty} (g_{a_{i^*}} + w_{a_{i^*}}(t + \hat{T}_B) + 2 w_{a_{i^*}}) \gamma_{-i^*} e^{-\gamma_{-i^*} s} ds \\
=& \frac{w_{a_{i^*}}}{\gamma_{-i^*}} + \frac{\sum_{i\neq i^*} \gamma_{a_i}}{\gamma_{-i^*}} w_{a_{i^*}} \hat{T}_B + \frac{\sum_j \gamma_{b_j}}{\gamma_{-i^*}} g_{a_{i^*}} + 2 w_{a_{i^*}} \\
&+ e^{-\gamma_{-i^*} t} \left(\frac{\sum_j \gamma_{b_j}}{\gamma_{-i^*}} w_{a_{i^*}} \hat{T}_B + \frac{\sum_{i\neq i^*} \gamma_{a_i}}{\gamma_{-i^*}} g_{a_{i^*}} - \frac{w_{a_{i^*}}}{\gamma_{-i^*}} \right). \quad (6)
\end{aligned}$$

We wish to find $\{\gamma_{a_1}, \gamma_{a_2}, \dots, \gamma_{b_1}, \gamma_{b_2}, \dots\}$ so that the expected cost of a_{i^*} is the same for all t. Hence, we require that

$$(1 - P_{A\setminus\{i^*\}<B}) w_{a_{i^*}} \hat{T}_B + P_{A\setminus\{i^*\}<B} g_{a_{i^*}} = \frac{w_{a_{i^*}}}{\gamma_{-i^*}} \quad (7)$$

$$\Leftrightarrow (\sum_j \gamma_{b_j}) \hat{T}_B w_{a_{i^*}} + (\sum_{i\neq i^*} \gamma_{a_i}) g_{a_{i^*}} = w_{a_{i^*}} \quad (8)$$

$$\Leftrightarrow \frac{w_{a_{i^*}}}{g_{a_{i^*}}} = \frac{\sum_{i\neq i^*} \gamma_{a_i}}{1 - (\sum_j \gamma_{b_j}) \hat{T}_B}, \quad (9)$$

for all a_{i^*}. Similarly, by studying the expected cost of a node b_{j^*}, we also require that

$$\frac{w_{b_{j^*}}}{g_{b_{j^*}}} = \frac{\sum_{j\neq j^*} \gamma_{b_j}}{1 - (\sum_i \gamma_{a_i}) \hat{T}_A}, \quad (10)$$

for all b_{j^*}.

Summing the (9) over all a_{i^*} and we have

$$\sum_i \frac{w_{a_i}}{g_{a_i}} = \frac{(I-1)\sum_i \gamma_{a_i}}{1 - (\sum_j \gamma_{b_j}) \hat{T}_B} \quad (11)$$

$$\Leftrightarrow (\sum_i \gamma_{a_j}) \hat{T}_A + (\sum_j \gamma_{b_j}) \hat{T}_B = 1. \quad (12)$$

Assume that $(\sum_i \gamma_{a_j}) \hat{T}_A = \alpha$ and $(\sum_j \gamma_{b_j}) \hat{T}_B = 1 - \alpha$, for some $\alpha \in (0, 1)$. Using (9) and (11), we obtain

$$\gamma_{a_{i^*}} = \sum_i \gamma_{a_i} - \sum_{i\neq i^*} \gamma_{a_i} = \alpha \left(\frac{\sum_i w_{a_i}/g_{a_i}}{I-1} - \frac{w_{a_{i^*}}}{g_{a_{i^*}}} \right). \quad (13)$$

Similarly, we also obtain

$$\gamma_{b_{j^*}} = (1-\alpha) \left(\frac{\sum_j w_{b_j}/g_{b_j}}{J-1} - \frac{w_{b_{j^*}}}{g_{b_{j^*}}} \right). \quad (14)$$

It is easy to check that, for every $\alpha \in (0, 1)$, setting γ_{a_i} and γ_{b_j} according to (13) and (14) satisfies (9) and (10) for all nodes, and the expected cost of each node is the same regardless the actual back-off timer it chooses. Hence, (13) and (14) form a Nash Equilibrium. Further, as α can be any number in $(0, 1)$, this game has infinitely many Nash Equilibria.

We summarize our results for systems with only two files in the following theorem.

Theorem 2: For any $\alpha \in (0, 1)$, if each node a_{i^*} chooses $t_{a_{i^*}} \sim EXP(\gamma_{a_{i^*}})$ and $\hat{t}_{a_{i^*}} \sim EXP(\lambda_{a_{i^*}})$, and each node b_{j^*} chooses $t_{b_{j^*}} \sim EXP(\gamma_{b_{j^*}})$ and $\hat{t}_{b_{j^*}} \sim EXP(\lambda_{b_{j^*}})$, where $\gamma_{a_{i^*}}, \lambda_{a_{i^*}}, \gamma_{b_{j^*}}$, and $\lambda_{b_{j^*}}$ are chosen by (13), (2), (14), and (4), then these strategies form a Nash Equilibrium.

IV. Price of Anarchy for Bilateral File Exchanges

We have found Nash Equilibria for our protocol when there are only two files in the system. We now discuss the performances of these Nash Equilibria.

Suppose all nodes choose their back-off timers according to Theorem 2, for some $\alpha \in (0, 1)$. The duration of the initiation phase is $\min\{t_{a_1}, t_{a_2}, \dots, t_{b_1}, t_{b_2}, \dots\}$, which is an exponential random variable with mean

$$T_{A,B} := \frac{1}{\sum_i \gamma_{a_i} + \sum_j \gamma_{b_j}}. \quad (15)$$

Also, the probability that one of the nodes in group $\{a_1, a_2, \dots\}$ transmits the control packet is

$$Prob\{\min_i t_{a_i} < \min_j t_{b_j}\} = \frac{\sum_i \gamma_{a_i}}{\sum_i \gamma_{a_i} + \sum_j \gamma_{b_j}}. \quad (16)$$

Using (3) and (5), we can express the expected amount of time for the two groups of nodes to exchange two packets as

$$\begin{aligned}
&T_{A,B} + \frac{\sum_i \gamma_{a_i}}{\sum_i \gamma_{a_i} + \sum_j \gamma_{b_j}} \hat{T}_B + \frac{\sum_j \gamma_{b_j}}{\sum_i \gamma_{a_i} + \sum_j \gamma_{b_j}} \hat{T}_A + 2 \\
&= \hat{T}_A + \hat{T}_B + 2, \quad (17)
\end{aligned}$$

where the last term in the equation accounts for the time needed for transmitting two packets. Since every node downloads a packet in each round, the per-node download rate of our protocol is then $1/(\hat{T}_A + \hat{T}_B + 2)$ packets per unit time under the described Nash Equilibria. On the other hand, the download rate of a node is 0.5 packet per unit time under a cooperative scenario. Therefore, the price of anarchy on per-node download rate is $(\hat{T}_A + \hat{T}_B + 2)/2$. To better understand the price of anarchy on per-node download rate, we consider the special case where all nodes have the same parameters for waiting cost and for transmission cost, that is, $w_n \equiv w$ and $g_n \equiv g$ for all n. In this case, we have $\hat{T}_A = \frac{(I-1)(g/w)}{I}$ and $\hat{T}_B = \frac{(J-1)(g/w)}{J}$, and the price of anarchy on system throughput is $(\frac{(I-1)(g/w)}{I} + \frac{(J-1)(g/w)}{J} + 2)/2 \leq 1 + \frac{g}{w}$.

We now compute the average total costs of nodes. The probability that a node a_{i^*} transmits a packet in a round can be expressed as

$$
\begin{aligned}
&Prob\{\min_i t_{a_i} < \min_j t_{b_j}\}Prob\{t_{a_{i^*}} < t_{a_i}, \forall i \neq i^*\}\\
&+ Prob\{\min_i t_{a_i} > \min_j t_{b_j}\}Prob\{\hat{t}_{a_{i^*}} < \hat{t}_{a_i}, \forall i \neq i^*\}\\
=&(\frac{\sum_i w_{a_i}/g_{a_i}}{I-1} - \frac{w_{a_{i^*}}}{g_{a_{i^*}}})/(\frac{\sum_i w_{a_i}/g_{a_i}}{I-1}).
\end{aligned}
$$

The average total cost of node a_{i^*} is then

$$
\begin{aligned}
&w_{a_{i^*}}(\hat{T}_A + \hat{T}_B + 2)\\
&+ g_{a_{i^*}}(\frac{\sum_i w_{a_i}/g_{a_i}}{I-1} - \frac{w_{a_{i^*}}}{g_{a_{i^*}}})/(\frac{\sum_i w_{a_i}/g_{a_i}}{I-1})\\
=&g_{a_{i^*}} + \frac{(J-1)w_{a_{i^*}}}{\sum_j w_{b_j}/g_{b_j}} + 2w_{a_{i^*}}. \qquad (18)
\end{aligned}
$$

On the other hand, under our protocol, the download rate of node a_{i^*} is at most one packet per 2 unit times. Hence, the average total cost of node a_{i^*} is at least $2w_{a_{i^*}}$. We then have that the price of anarchy on node cost of a_{i^*} is at most $(\frac{g_{a_{i^*}}}{w_{a_{i^*}}} + \frac{(J-1)}{\sum_j w_{b_j}/g_{b_j}} + 2)/2$. In the special case where $w_n \equiv w$ and $g_n \equiv g$, for all nodes, the price of anarchy on node cost of a_{i^*} is at most

$$
(\frac{g}{w} + \frac{(J-1)(g/w)}{J} + 2)/2 \leq 1 + \frac{g}{w}.
$$

We note that both the price of anarchy on system throughput and the price of anarchy on node cost increase with $\frac{g}{w}$. Intuitively, when g is small compared to w, nodes focus more on improving their download rates than on reducing transmission costs. Hence, nodes tend to choose small back-off timers, which makes the prices of anarchy small. On the other hand, when g is much larger than w, transmission costs become an important factor of nodes' costs. Hence, each node tends to choose large back-off timers to increase its chance of becoming a free rider, which results in large prices of anarchy.

As a final remark, we note that both the system throughput and total average costs of nodes remain the same for all $\alpha \in (0, 1)$. Therefore, the performance of the system is the same for all Nash Equilibria described by Theorem 2.

V. Implementation Issues and Convergence

Section III has described a Nash Equilibrium for a system with two files. However, for a node n to derive its strategies, that is, to compute the values of γ_n and λ_n, node n needs to know information of the whole network, including the private values of w_m and g_m for all other nodes m. In this section, we propose a distributed mechanism for each node to update its values of γ_n and λ_n only based on its values of w_n and g_n and the history of the system. We show that this mechanism is compatible to the node's incentive, in the sense that the updated γ_n and λ_n achieve smaller average total cost for the node. Moreover, we also show that the system converges to a Nash Equilibrium when all nodes apply this mechanism.

We order the two files by lexicographical order. If file A has higher order than B, we impose that $\gamma_{a_i} = 0$, for all i, and $\lambda_{b_j} = 0$, for all j. This corresponds to the case where $\alpha = 0$ in Section III. Therefore, in every round, a node in $\{b_1, b_2, \dots\}$ transmits a control packet in the initiation phase and a node in $\{a_1, a_2, \dots\}$ transmits a data packet in the response phase. On the other hand, if file B has higher order than A, we impose that $\lambda_{a_i} = 0$, for all i, and $\gamma_{b_j} = 0$, for all j, which corresponds to the case where $\alpha = 1$. Without loss of generality, we assume that file A has higher order than B.

Using (2), (3), (14) , we have that $\hat{T}_A = \frac{1}{\sum_i \lambda_{a_i}} = \frac{I-1}{\sum_i \frac{w_{a_i}}{g_{a_i}}}$, $\lambda_{a_{i^*}} = \frac{1}{\hat{T}_A} - \frac{w_{a_{i^*}}}{g_{a_{i^*}}}$, $\hat{T}_B = \frac{J-1}{\sum_j \frac{w_{b_j}}{g_{b_j}}}$, and $\gamma_{b_{j^*}} = \frac{1}{\hat{T}_B} - \frac{w_{b_{j^*}}}{g_{b_{j^*}}}$ at the Nash Equilibrium, where I and J are the number of nodes in groups $\{a_1, a_2, \dots\}$ and $\{b_1, b_2, \dots\}$, respectively. As we set $\alpha = 0$, $\hat{T}_A$ is the average back-off time in the response phase, and $\hat{T}_B$ is that in the initiation phase.

We now introduce our mechanism for a node a_{i^*}. Node a_{i^*} first guesses that the average amount of back-off time in the response phase is $\hat{T}_{a_{i^*},0}$, and sets $\lambda_{a_{i^*}} = \frac{1}{\hat{T}_{a_{i^*},0}} - \frac{w_{a_{i^*}}}{g_{a_{i^*}}}$. Node a_{i^*} then observes the system behaviors and updates its value of $\lambda_{a_{i^*}}$ every M rounds. When a_{i^*} updates $\lambda_{a_{i^*}}$ the k-th time, it computes the average back-off time in the response phase since it last updates its values, denoted by $\hat{T}_{A,k-1}$. Node a_{i^*} then sets $\hat{T}_{a_{i^*},k}$ so that

$$
\frac{1}{\hat{T}_{a_{i^*},k}} = \frac{1}{\hat{T}_{a_{i^*},k-1}} - \delta_k(\frac{1}{\hat{T}_{A,k-1}} - \frac{1}{\hat{T}_{a_{i^*},k-1}}),
$$

and $\lambda_{a_{i^*}} = \frac{1}{\hat{T}_{a_{i^*},k}} - \frac{w_{a_{i^*}}}{g_{a_{i^*}}}$, where the values of δ_k are chosen so that $\sum_{k=1}^{\infty} \delta_k = \infty$ and $\sum_{k=1}^{\infty} \delta_k^2 < \infty$. For example, one can choose $\delta_k = \frac{\epsilon}{k}$, where ϵ is a small constant. The mechanism for a node b_{j^*} can be derived similarly. The only difference is that node b_{j^*} updates its value of $\gamma_{b_{j^*}}$ based on the average back-off time in the initiation phase.

As described in the following theorems, this mechanism has two important features. First, this mechanism is compatible to the node's incentive. Second, this mechanism converges to a Nash Equilibrium.

Theorem 3: Fix the values of γ_n and λ_n, for all $n \neq a_{i^*}$, such that $\gamma_{a_i} = 0$, for all i, and $\lambda_{b_j} = 0$, for all j. The average total cost of node a_{i^*} is smaller when it sets $\lambda_{a_{i^*}} = \frac{1}{\hat{T}_{a_{i^*},k+1}} - \frac{w_{a_{i^*}}}{g_{a_{i^*}}}$, than when it sets $\lambda_{a_{i^*}} = \frac{1}{\hat{T}_{a_{i^*},k}} - \frac{w_{a_{i^*}}}{g_{a_{i^*}}}$, for all k.

Proof: Since we impose $\gamma_{a_{i^*}} = 0$, node a_{i^*} has no control on the amount of back-off time in the initiation phase. Hence, it suffices to show that setting $\lambda_{a_{i^*}} = \frac{1}{\hat{T}_{a_{i^*},k+1}} - \frac{w_{a_{i^*}}}{g_{a_{i^*}}}$ achieves smaller average additional total cost in the response phase than setting $\lambda_{a_{i^*}} = \frac{1}{\hat{T}_{a_{i^*},k}} - \frac{w_{a_{i^*}}}{g_{a_{i^*}}}$.

We have

$$
\hat{T}_{A,k} = \frac{1}{\sum_{i \neq i^*} \lambda_{a_{i^*}} + (\frac{1}{\hat{T}_{a_{i^*},k}} - \frac{w_{a_{i^*}}}{g_{a_{i^*}}})},
$$

and

$$
g_{a_{i^*}} - \frac{w_{a_{i^*}}}{\sum_{i \neq i^*} \lambda_{a_{i^*}}} = \frac{g_{a_{i^*}}}{\sum_{i \neq i^*} \lambda_{a_{i^*}}}(\frac{1}{\hat{T}_{A,k}} - \frac{1}{\hat{T}_{a_{i^*},k}}).
$$

By (1), if $g_{a_{i^*}} - \frac{w_{a_{i^*}}}{\sum_{i \neq i^*} \lambda_{a_{i^*}}} > 0$, the expected additional cost of a_{i^*} strictly decreases with the back-off timer of a_{i^*},

t. Moreover, if $g_{a_{i^*}} - \frac{w_{a_{i^*}}}{\sum_{i\neq i^*}\lambda_{a_{i^*}}} > 0$, we have $\hat{T}_{a_{i^*},k} > \hat{T}_{A,k}$ and $\hat{T}_{a_{i^*},k} < \hat{T}_{a_{i^*},k+1}$. For every positive constant C, $Prob(t < C)$ is smaller when $\lambda_{a_{i^*}} = \frac{1}{\hat{T}_{a_{i^*},k+1}} - \frac{w_{a_{i^*}}}{g_{a_{i^*}}}$ than when $\lambda_{a_{i^*}} = \frac{1}{\hat{T}_{a_{i^*},k}} - \frac{w_{a_{i^*}}}{g_{a_{i^*}}}$. Therefore, the average additional total cost is smaller when $\lambda_{a_{i^*}} = \frac{1}{\hat{T}_{a_{i^*},k+1}} - \frac{w_{a_{i^*}}}{g_{a_{i^*}}}$ than when $\lambda_{a_{i^*}} = \frac{1}{\hat{T}_{a_{i^*},k}} - \frac{w_{a_{i^*}}}{g_{a_{i^*}}}$.

On the other hand, if $g_{a_{i^*}} - \frac{w_{a_{i^*}}}{\sum_{i\neq i^*}\lambda_{a_{i^*}}} < 0$, the expected additional cost of a_{i^*} strictly increases with the back-off timer of a_{i^*}, t. Moreover, if $g_{a_{i^*}} - \frac{w_{a_{i^*}}}{\sum_{i\neq i^*}\lambda_{a_{i^*}}} > 0$, we have $\hat{T}_{a_{i^*},k} < \hat{T}_{a_{i^*},k+1}$, and, for every positive constant C, $Prob(t < C)$ is larger when $\lambda_{a_{i^*}} = \frac{1}{\hat{T}_{a_{i^*},k+1}} - \frac{w_{a_{i^*}}}{g_{a_{i^*}}}$ than when $\lambda_{a_{i^*}} = \frac{1}{\hat{T}_{a_{i^*},k}} - \frac{w_{a_{i^*}}}{g_{a_{i^*}}}$. Therefore, the average additional total cost is also smaller when $\lambda_{a_{i^*}} = \frac{1}{\hat{T}_{a_{i^*},k+1}} - \frac{w_{a_{i^*}}}{g_{a_{i^*}}}$ than when $\lambda_{a_{i^*}} = \frac{1}{\hat{T}_{a_{i^*},k}} - \frac{w_{a_{i^*}}}{g_{a_{i^*}}}$. ■

Theorem 4: Fix the values of γ_n and λ_n, for all $n \neq b_{j^*}$, such that $\gamma_{a_i} = 0$, for all i, and $\lambda_{b_j} = 0$, for all j. The average total cost of node b_{j^*} becomes smaller when it updates its $\gamma_{b_{j^*}}$.

Proof: The proof is very similar to that of Theorem 3, and is hence omitted. ■

Next, we show that our mechanism converges to the Nash Equilibrium.

Theorem 5: If all nodes apply the proposed mechanism, then the value of $\hat{T}_{A,k}$ and $\hat{T}_{B,k}$ after each node updates k times converge to $\hat{T}_A = \frac{I-1}{\sum_i \frac{w_{a_i}}{g_{a_i}}}$ and $\hat{T}_B = \frac{J-1}{\sum_j \frac{w_{b_j}}{g_{b_j}}}$, respectively, as $k \to \infty$.

Proof: We only prove that $\hat{T}_{A,k}$ converges to $\hat{T}_A$. Under our mechanism, the value of λ_{a_i} after node a_i updates k times is

$$\begin{aligned}\lambda_{a_i} =& \frac{1}{\hat{T}_{a_i,k}} - \frac{w_{a_i}}{g_{a_i}}\\ =&(1+\delta_k)\frac{1}{\hat{T}_{a_i,k-1}} - \delta_k\frac{1}{\hat{T}_{A,k-1}} - \frac{w_{a_i}}{g_{a_i}}\\ =&(1+\delta_k)\lambda_{a_i,k-1} - \delta_k(\frac{1}{\hat{T}_{A,k-1}} - \frac{w_{a_i}}{g_{a_i}}),\end{aligned}$$

and hence

$$\begin{aligned}\frac{1}{\hat{T}_{A,k}} &= \textstyle\sum_i \lambda_{a_i}\\ &= (1+\delta_k)(\textstyle\sum_i \lambda_{a_i,k-1}) - \delta_k(\frac{I}{\hat{T}_{A,k-1}} - \sum_i \frac{w_{a_i}}{g_{a_i}})\\ &= \frac{1}{\hat{T}_{A,k-1}} + \delta_k(I-1)(\frac{1}{\hat{T}_A} - \frac{1}{\hat{T}_{A,k-1}}).\end{aligned}$$

As we have $\sum_k \delta_k = \infty$ and $\sum_k \delta_k^2 < \infty$, $\frac{1}{\hat{T}_{A,k}}$ converges to $\frac{1}{\hat{T}_A}$. ■

VI. Nash Equilibria for Multiple File Exchanges with Network Coding

In this section, we derive the Nash Equilibria for systems that incorporate network coding and have more than two files. We use $I_A, I_B, I_C, \ldots$ to denote the number of nodes that need $A, B, C, \ldots$, respectively. We will show that there exists a Nash Equilibrium where each node n chooses each of its back-off timers in both phases as exponential random variables. Many of derivations in this section are similar to those in Section III. Hence, due to the space constraints, we omit some details and only report the results.

We first consider the nodes' strategies in the response phase. Assume that a node n with $X_n = X$ has sent the control packet in the initiation phase. Each node m with $X_m \neq X$ secretly and randomly chooses a back-off timer $\hat{t}_m \sim EXP(\lambda_{m|X})$. Note that the value of $\lambda_{m|X}$ may depend on X. Let m_0 be the node that chooses the smallest value of $\hat{t}_m$. Then, m_0 will transmit a data packet of X after $\hat{t}_{m_0}$ time. After m_0 transmits the data packet, n will transmit a coded packet that contains one packet from each of the files except X, and every node m with $X_m \neq X$ can decode one packet from the transmission from n. Similar to the derivations of (2), we can show that, at a Nash Equilibrium, we have

$$\lambda_{m|X} = \frac{1}{\sum_{Y:Y\neq X} I_Y - 1}\sum_{l:X_l\neq X}\frac{w_l}{g_l} - \frac{w_m}{g_m}, \tag{19}$$

for all m such that $X_m \neq X$.

Next, we consider the nodes' strategies in the initiation phase. Assume that each node n chooses a back-off timer $t_n \sim EXP(\gamma_n)$. Let $\Gamma_X := \sum_{n:X_n=X}\gamma_n$. Similar to the derivations of (9), (10), and (12), we have that, at a Nash Equilibrium:

$$\frac{w_n}{g_n} = \frac{\Gamma_{X_n} - \gamma_n}{1 - (\sum_{X:X\neq X_n}\Gamma_X)\frac{\sum_{X:X\neq X_n} I_X - 1}{\sum_{m:X_m\neq X_n} w_m/g_m}}, \tag{20}$$

for all n, and

$$\frac{I_X - 1}{\sum_{n:X_n=X} w_n/g_n}\Gamma_X + \frac{\sum_{Y:Y\neq X} I_Y - 1}{\sum_{m:X_m\neq X} w_m/g_m}\sum_{Y:Y\neq X}\Gamma_Y = 1, \tag{21}$$

for all X. Equation (21) represents a series of linear equations where both the number of unknowns, $\{\Gamma_X\}$, and the number of equations equal to the number of files. We can use standard techniques for solving linear equations to obtain a solution of $\{\Gamma_X\}$ to (21). We can then use $\{\Gamma_X\}$ to obtain the values of $\{\gamma_n\}$ through (20). The derived $\{\gamma_n\}$ and $\{\lambda_{n|X}\}$ form a Nash Equilibrium.

VII. Numerical Results

We now present our simulation results. We first consider a system with two files, A and B, and 20 nodes. We assume that there are 10 nodes that possess file A and need file B, and the other 10 nodes possess B and need A. We set $g_n = 1$ for all n, and set w_n to be uniformly distributed in $[1, 2]$. Each node applies the distributed mechanism introduced in Section V to determine its strategy. Each node n sets $\hat{T}_{n,0} = \frac{99}{100(w_n/g_n)}$, that is, it guesses that there are 100 nodes in its group, an overestimate by a factor of 10, and all nodes in its group have the same values of w_n and g_n as itself. Also, each node sets $\delta_k = 0.1/k$.

Fig. 2 shows the resulting per-node throughput after each update. It can be shown that, even though the initial strategies of nodes are far from the Nash Equilibrium, the per-node throughput under our mechanism converges

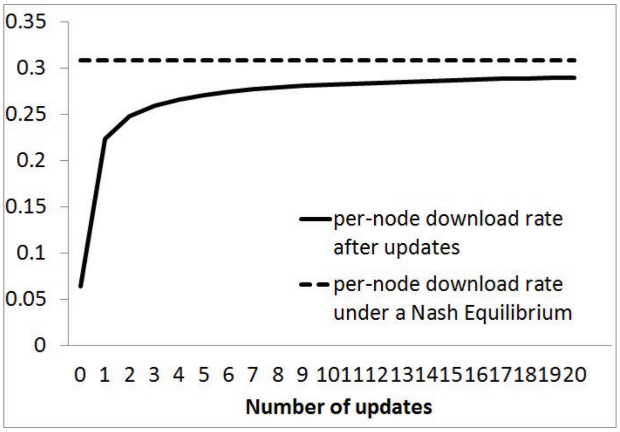

Fig. 2: Per-node download rate with two files.

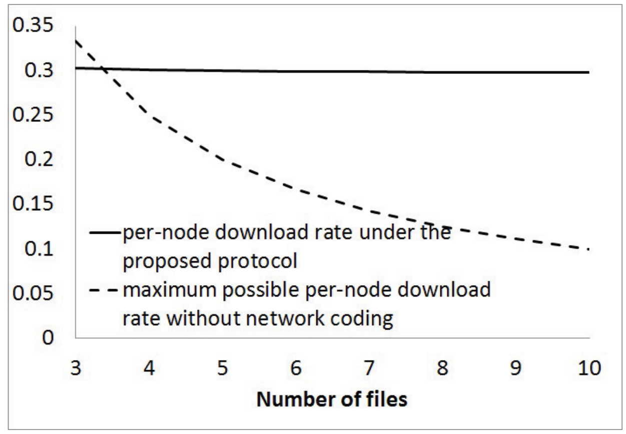

Fig. 3: Per-node dowload rate with multiple files and network coding.

to the Nash Equilibrium very quickly. With just three updates, the per-node throughput is about 85% of that under the Nash Equilibrium.

Next, we consider systems that have more than two files and employ network coding. We assume that, for each file, there are 10 nodes that need it. Therefore, there are a total number of $10 \times \{\text{number of files}\}$ nodes in the system. We also set $g_n = 1$ for all nodes, and w_n uniformly distributed in $[1, 2]$. We use the procedure described in Section VI to derive the values of γ_n and λ_n for all n, based on which we calculate the per-node throughput under a Nash Equilibrium.

Fig. 3 shows the per-node throughput under various numbers of files in the system. We compare the performance of our protocol against the maximum possible per-node throughput when network coding is not employed and when all nodes are cooperative. Without network coding, each transmission only contains one packet of one file. Hence, the maximum possible per-node throughput is $1/\{\text{number of files}\}$. Fig. 3 shows that, when there are only three files in the system, our protocol has slightly worse per-node throughput than the case when network coding is not employed and all nodes are cooperative. This is because our protocol is designed for selfish nodes which requires two phases of back-offs, while the compared scenario assumes all nodes are cooperative and hence there is no time spent on back-offs. However, as the number of files increases, the benefits of network coding outweigh the prices of anarchy. As a result, our protocol achieves better per-node throughput than the scenario where network coding is not employed and all nodes are cooperative.

VIII. Conclusions

The paper considers the problem of content distribution in wireless P2P networks and proposes a model that captures both the broadcast nature of wireless medium and the incentives of nodes. The paper presents a non-monetary protocol for content distribution in this model. The protocol provides incentives for selfish nodes to contribute to the network. We have studied the performance of our protocol when all nodes are selfish. For systems with only two files, we have derived closed-form expressions for Nash Equilibria and prices of anarchy. We have also proposed a distributed mechanism where each node updates its strategies based only on its private information and the history of the system. Our numerical results show that this mechanism converges to Nash Equilibria very quickly. For systems with more than two files, we propose a simple extension of our protocol to incorporate network coding and present a procedure to compute each node's strategy under a Nash Equilibrium. Numerical results show that our protocol may achieve better performance than scenarios where nodes are cooperative but do not employ network coding.

References

[1] H.-Y. Hsieh and R. Sivakumar, "On using peer-to-peer communication in cellular wireless data networks," *IEEE Transactions on Mobile Computing*, vol. 3, no. 1, pp. 57–72, 2004.

[2] M. J. Neely and L. Golubchik, "Utility optimization for dynamic peer-to-peer networks with tit-for-tat constraints," in *Proc. of IEEE INFOCOM*, 2011.

[3] J. Zhao, P. Zhang, and G. Cao, "On cooperative caching in wireless p2p networks," in *Proc. of ICDCS*, 2008.

[4] C. Aperjis, R. Johari, and M. J. Freedman, "Bilateral and multilateral exchanges for peer-assisted content distribution," *IEEE/ACM Transactions on Networking*, vol. 19, pp. 1290–1303, Oct. 2011.

[5] W. S. Lin, H. V. Zhao, and K. J. R. Liu, "Incentive cooperation strategies for peer-to-peer live multimedia streaming social networks," *IEEE Transactions on Multimedia*, vol. 11, pp. 396–412, Apr 2009.

[6] V. Misra, S. Ioannidis, A. Chaintreau, and L. Massoulié, "Incentivizing peer-assisted services: a fluid shapley value approach," in *Proc. of ACM SIGMETRICS*, 2010.

[7] J. Park and M. van der Schaar, "Pricing and incentives in peer-to-peer networks," in *Proc. of IEEE INFOCOM*, 2010.

[8] Z. Han and H. V. Poor, "Coalition games with cooperative transmission: a cure for the curse of boundary nodes in selfish packet-forwarding wireless networks," *IEEE Transactions on Communications*, vol. 57, pp. 203–213, Jan 2009.

[9] W. Li, J. Chen, and B. Zhou, "Game theory analysis for graded punishment mechanism restraining free-riding in p2p networks," in *Proc. of ISCCS*, 2011.

[10] A. Sprintson, P. Sadeghi, G. Booker, and S. El Rouayheb, "Deterministic algorithm for coded cooperative data exchange," in *Proceedings of the 7th International Conference on Heterogeneous Networking for Quality, Reliability, Security and Robustness (QShine 2010)*, 2010.

[11] Z. Han, D. Niyato, W. Saad, T. Basar, and A. Hjorungnes, *Game Theory in Wireless and Communication Networks*. Cambridge University Press, 2012.