On the Market Power of Network Coding in P2P Content Distribution Systems

Xinyu Zhang, Baochun Li
Department of Electrical and Computer Engineering
University of Toronto
Email: {xzhang, bli}@eecg.toronto.edu

***Abstract*—Network coding is emerging as a promising alternative to traditional content distribution approaches in P2P networks. By allowing information mixture in peers, it simplifies the block scheduling problem, resulting in more efficient data delivery. Existing protocols have validated such advantages assuming altruistic and obedient peers. In this paper, we develop an analytical framework that characterizes a coding based P2P content distribution market where peers selfishly seek for individual payoff maximization. Through virtual monetary exchanges, agents in the market buy the coded blocks from others and resell their possessions to those in need. We model such transactions as decentralized strategic bargaining games, and derive the equilibrium prices between arbitrary pairs of agents when the market enters the steady state. We identify the traditional P2P content distribution approach as a special case of network coding, and characterize the relations between coding complexity and market performance metrics, including agents' entry price and expected payoff, thus providing operation guidelines for a real P2P market. Our analysis reveals that the major power of network coding lies in its ability to maintain stability of the market with impatient and selfish agents, and to incentivize agents with lower price and higher payoff, at the cost of reasonable coding complexity.**

I. Introduction

P2P content distribution systems are built atop the basic premise of voluntary resource contribution by participating peers. Two critical problems are inherent in this presumption: the scheduling decision of individual peers (*i.e.*, choosing which data blocks to share) and the incentives for sharing.

Existing P2P content distribution systems tackled the scheduling problem using random or rarest-first strategies. Such heuristic local algorithms tend to result in suboptimal uploading or downloading decisions that waste network resources [1]. Network coding circumvents the scheduling problem by allowing each peer to encode and deliver a random linear combination of the data on hand. As long as one block is fresh, the entire encoded block is useful to the requester with high probability. Therefore, the risk of uploading duplicate information can be significantly reduced without sophisticated scheduling. Existing protocols (see, *e.g.*, [1]) have identified network coding based content distribution as a workable idea, but without rigorous theoretical quantification of its advantages. They have also assumed altruistic resource sharing among peers, which is inconsistent with the greedy and selfish behavior that dominates real-world P2P systems [2].

In this paper, we analyze the performance of network coding based P2P content distribution protocols from an economic and game theoretic perspective. We envision the P2P system as a decentralized content distribution market. Each peer acts as a market agent, namely a seller and buyer. Before entering the market, a peer must pay an initial service fee (referred to as *entry price*) that is used to obtain at least one block. Afterwards, he can resell the blocks he already possesses and purchase additional blocks with the money on hand. Whenever a seller and a buyer meet, they bargain over the blocks of interest for a consensus price. Both sides of the bargaining game take into account the availability of alternative sellers and buyers, and the potential resale value of the good once the transaction succeeds. Such a model resembles an exchange economy for digital information goods, and sheds lights on the deployment and evolution of practical P2P markets.

We classify peers in the market according to their possessions, *i.e.*, the availability of blocks on them. By modeling the transactions between peers as noncooperative games, we derive the equilibrium pricing strategies for different types of peers. We find that unlike traditional centrally managed market economy, no uniform price exists under strategic bargaining. Instead, the price depends on not only the availability of the goods, but also the valuation of each type of peer on each good. Furthermore, we extend the game to a market scale, and characterize a market equilibrium in which individual peers adopt stationary strategies, and no one can profit more by deviating over time. We then approximate the evolution of such a market using a system of differential equations, and derive the availability of goods when the market enters steady state.

The above theoretical framework results in closed-form equations that quantify the impact of various design parameters on the stable operations of the market. Through these equations, we observe that the fundamental advantage of network coding lies in maintaining the availability of data blocks even when peers are highly impatient and even in the absence of content servers. Translated into market terms, coding based protocols induce a higher level of competition among content sellers, thereby avoiding the monopoly or oligopoly scenarios in which a limited number of content holders force up the price. Furthermore, network coding incentivizes the peers by increasing their expected payoff, and reasonable coding complexity is sufficient to harvest such an advantage. Unfortunately, we also find that network coding is against the interests of content servers as their profit decreases with increasing coding complexity.

The remainder of the paper is organized as follows. In Sec. II we contrast our work with existing analyses of P2P systems, especially those with network coding and from an

This work was supported in part by LG Electronics, and by Bell Canada through its Bell University Laboratories R&D program.

economic perspective. We then give a brief introduction to practical network coding protocols and our system models in Sec. III. Sec. IV analyzes the decentralized bargaining game and characterizes its equilibrium and the corresponding market equilibrium. Sec. V models the evolution of goods availability, and then presents a comprehensive analysis on the market's properties at steady state. Finally, Sec. VI concludes the paper.

II. Related Work

Since the pioneering work by Ho *et al.* [3], randomized network coding has received substantial attention from P2P protocol designers. The Avalanche [1] system implemented a primitive form of random linear code that encodes all data blocks in a file. However, when the file size exceeds a few hundred blocks, such a full-coding scheme causes intolerable computational overhead even for modern processors [4]. More recent protocols have extended the idea of *segment-based network coding* [5], which splits the file into multiple segments, each allowing for efficient encoding and decoding. This approach has demonstrated its effectiveness in not only file sharing, but also elastic content distribution systems like P2P streaming [6].

Despite its wide applications, the fundamental benefits of network coding in such systems have not been fully explored with theoretical rigor. Chiu *et al.* [7] abstracted a P2P system as a *static* star topology, and claimed that coding does not increase the network capacity compared with routing. Through mean-field analysis of a dynamic P2P system, Niu *et al.* [4] claimed that network coding can alleviate the imbalance of block distributions in traditional content distribution protocols, thereby improving the resilience to network dynamics. Both analytical works, as well as the existing system implementations, have relied on the premise of cooperative peers, while measurement of real P2P systems exhibits a dominant portion of selfish free-riders [2], [8]. In this paper, we aim at quantifying the fundamental advantages of network coding in such non-cooperative environment.

Our work is partly inspired by Rubinstein [9], who analyzed the impact of strategic price settings on the equilibrium of a market economy. Traditional market economy has assumed agents leaving the market after a successful transaction, with the buyer owning the goods while the seller earning the payment. In contrast, P2P systems feature *copiable* and *resalable* products that propagate their values over time, thus requiring the support of a brand new model.

Game theoretic analysis of peer behaviors has been widely employed (see [2] for a survey). This line of research has mostly focused on designing incentives that encourage cooperation and lead the peers towards a socially optimal point. The results have inspired commercial P2P systems to adopt incentive mechanisms (*e.g.*, virtual payment) that motivate peers' willingness to share [10]. We are less concerned with designing such payment protocols, and instead, more focused on the equilibrium analysis assuming a virtual payment scheme is available. Our work differs from existing game theoretical framework not only in an emphasis on network coding, but also in its equilibrium analysis under a *decentralized market* setting. We consider not just the strategic behavior of individual peers, but also how their self-interested pricing strategies influence the content distribution market as a whole.

Economic models, in particular the market models for P2P systems have been explored by the MMAPPS project [11], which proposed market management techniques to encourage cooperation. Within MMAPPS, Antoniadis *et al.* [12] developed a theoretical framework that abstracted the shared content as public goods. Assuming that peers' valuation of the goods follows uniform distribution, a social planner determines incentive-compatible prices and services for them. However, the mechanism lacks a support for network dynamics and a concrete modeling of the peers' valuations.

Aperjis *et al.* [13] proposed a comprehensive exchange economy model that captures the optimal equilibrium price in a content distribution market. They analyzed a one-shot exchange market without resale, and without discriminating goods according to their availability. The economic implication of network coding has been discussed in recent work [14], focusing on centralized cellular networks with price-taking agents. To our knowledge, there does not exist any previous work on the power of network coding in a decentralized P2P content distribution market with strategical participants.

III. Coding Based P2P Content Distribution Market

In this section, we first introduce the widely used segment-based network coding protocol for P2P content distribution. When running such a protocol, peers purchase and resell the coded data blocks, thus forming a content distribution market. We specify the various elements of such a P2P market economy, including the classification of peers and the price formation.

A. P2P Content Distribution with Network Coding

Existing coding based P2P content distribution protocols have mostly adopted the following segment based scheme. Before transmission, the original data file is grouped into segments, each containing K blocks of size E bytes. K and E are termed *segment size* and *block size*, respectively. *The coding operations are performed within each segment*. We represent each segment as a matrix B, a $K \times E$ matrix, with rows being the K blocks, and columns the bytes (integers from 0 to 255) of each block. The encoding operation produces a linear combination of the original blocks in this segment by $X = R \cdot B$, where R is a $K \times K$ matrix composed of random coefficients in the Galois field $GF(2^8)$. The *coded blocks* (rows in X), together with the *coding coefficients* (rows in R), are packetized and delivered to other peers.

The decoding operation at each peer is the matrix multiplication $B = R^{-1} \cdot X$, where each row of X represents a coded block and each row of R represents the coding coefficients accomplished with it. The successful recovery of the original segment B requires that the matrix R be of full rank, *i.e.*, each peer must collect K independent coded blocks for this segment. However, a peer can upload coded blocks even if the segment is not ready to decode yet. It produces a new block by *re-encoding* existing blocks it has collected in this segment. The re-encoding operation replaces the coding coefficients accomplished with the original coded packets with another set of random coefficients. For instance, consider the existing coded packets as rows in the matrix Y, which from the viewpoint of the source was obtained using $Y = R_y \cdot B$ (B is the original uncoded packets and R_y is the random coefficients). Then the current holder may produce a new code block by re-encoding existing packets as $Y' = R' \cdot R_y \cdot B = R'_y \cdot B$. As a result, the original coefficients R_y are replaced by R'_y. The re-encoding operation circumvents the block-level scheduling problem in traditional content distribution protocols, because by randomly mixing information from all existing blocks, a newly generated coded block is innovative to the downstream peer with high probability [4].

Although randomized network coding solves the block selection problem within a segment, a scheduling algorithm must be in position to decide which segment to upload or download. We assume each peer adopts a push-based random scheduling protocol, which randomly selects a segment, generates a coded block, and then upload it to his partner. This assumption does not limit the generality of our major analysis, as our game theoretical models conclude with pricing strategies that adapt to general scheduling policies.

Note that *traditional non-coding scheme can be considered as a special case of segment-based network coding where $K = 1$*, *i.e.*, each segment has a single data block. By contrast, the Avalanche [1] protocol corresponds to the other polar, *i.e.*, the *full-coding* case, where the entire file is encoded into a single segment.

B. The Decentralized Market

We envision the existence of an online market place where peers act as the *agents* who purchase and sell coded data blocks originated from a single file. The file consists of F blocks and is grouped into M segments, *i.e.*, the *segment size* $K = \frac{F}{M}$. Since all data blocks within each segment are equally useful to the buyers, *each segment corresponds to one type of good*, *i.e.*, the total types of goods circulating in the market equals M. A virtual currency, such as the lightweight currency in [15], serves as the medium of transaction. Whenever a buyer and a seller meet, they initiate a pairwise bargaining process over the segment of mutual interest. If both peers agree upon a certain price, then the seller uploads a coded block from the segment of interest, and the buyer will pay the money in return. An agent may act as a seller and buyer simultaneously, resulting in an exchange transaction, as illustrated in Fig. 1. If on the other hand the negotiation ends with a disagreement, then both peers have to switch to alternative partners.

Without loss of generality, we mainly discuss the pricing of a single good. We classify the market agents into $(K+1)$ types. A type-i agent ($0 \leq i \leq K$) possesses a total number of i coded blocks of the good. Obviously, A type-0 agent can only purchase goods, while a type-K agent who has fulfilled the segment only sells goods to others.

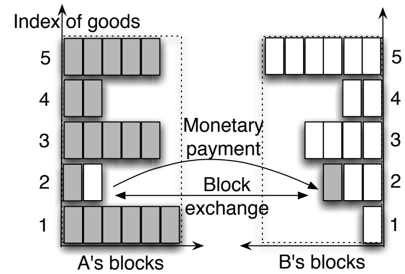

Fig. 1. The pairwise transaction in a coding based P2P market with $M = 5$ segments and segment size $K = 6$. Agents A and B randomly select a good (segment 2) for bargaining and transaction.

As in a real-world market, the outcome of any pairwise bargaining depends on the current market condition, *i.e.*, the availability of the goods. If a good is abundant in the market, the buyer can easily find an alternative seller, and the buyer may be better-off searching for alternative providers if the price proposed by the current seller is too high. Conversely, scarce goods will be charged higher prices than abundant ones. As the market evolves, we can expect that an equilibrium exists that specifies a stationary per-block price for each good. Once the market evolves to a steady-state, all peers agree upon a common set of prices and no actual negotiation takes place (Sec. IV). We will formalize the equilibrium point in the following section.

IV. The Decentralized Bargaining Game

In this section, we describe the elemental transaction procedure on the P2P market, *i.e.*, the pairwise bargaining game. We rigorously characterize the equilibrium pricing strategies of each type of agents, and then extend the pairwise bargaining to a market setting.

A. The Rules of Bargaining

We model the P2P market as a discrete time system. The duration of a period equals the time needed to transmit a single block. To capture the market dynamics, we assume an agent is impatient, remaining in the market in each period with probability θ, *i.e.*, the churn rate (peer join and departure rate) $\mu = 1 - \theta$. Agents have homogeneous upload and download bandwidth, which equals 1 block per period. We abstract the peer selection as a matching process in which an agent is randomly matched to another agent in each period. Upon matching, the pair of agents select one good of mutual interest and propose to *exchange one block* of the good (agents are only allowed to exchange blocks of the same good because different goods may experience different availability and distinct prices). The outcome of the matching depends on two factors: the usefulness of the block, and the bargaining result.

First, before transaction, both agents need to make sure they can provide at least one useful block to each other. This can be trivially satisfied if they are of type i and j, respectively, where $0 < i, j < K$. If one of them is of type 0 or K, then the transaction becomes a unilateral sale, instead of bilateral exchange. Note that even in a unilateral sale, the bandwidth is not wasted because two goods may be simultaneously on sale.

The second and most critical factor in the transaction is whether the bargaining between the pair of agents results in an agreement. Since the agents may have different valuations of the good, the one who gain more has to pay for the other. To avoid unfair advantages of the initiator, one agent (referred to as *proposer*) is randomly selected to propose a price. The opposite agent (referred to as *responder*) responds by either accepting or rejecting the proposal. In case of rejection, both agents continue to the next period, looking for new partners. The ability to switch to alternative partners enhances the agents' bargaining power, since they can threat to abandon the current partner, thus making it a "take-it-or-leave-it" offer. Therefore, whether the bargaining results in agreement or disagreement depends on the availability of alternative partners on the market.

As the market evolves to a steady-state, each type of agents adopt stationary strategies. To be specific, a *stationary strategy* implies that each type of proposer or responder maintains the same reservation prices when facing the same type of partners. The *reservation prices* of a type-i agent include a proposer price p^*_{ij}, the optimal price he can bid that is acceptable to a type-j agent; and a responder price q^*_{ij}, the optimal price that is proposed by his partner j and is acceptable to him. For consistency, the subscript ij always indicates the price that i should pay to j, hence $p_{ij} = -q_{ji}$. With stationary strategies, whenever an agent i is matched to agent j, he proposes p^*_{ij} to agent j if he is selected as the proposer; and he accepts a proposal q_{ij} from j if and only if $q_{ij} \le q^*_{ij}$. Hence, the negotiation procedure is no longer needed in a steady-state market. In what follows, we characterize the reservation prices p^*_{ij} and q^*_{ij} corresponding to the unique stationary strategy that satisfies subgame-perfectness. We further justify that it is not profitable for an agent to use non-stationary strategies at equilibrium.

B. The Subgame-Perfect Nash Equilibrium in Pairwise Transactions

The classic concept of *Nash equilibrium* in game theory characterizes the strategy profiles in which no players can profit more by unilaterally deviating from his current strategy. However, Nash equilibrium strategies may include *incredible threats*, which the threatener himself does not prefer to issue, but which may still deter the actions of the one under threat. In the above bargaining game, the stationary strategy that rejects all transactions constitutes a Nash equilibrium, since every agent receives payoff 0 and no one can profit more than 0 by changing his own strategy. However, threatening to resort to such strategies are incredible since the agents are aware that any alternatives that encourage transactions can be more benefitial. The concept of Subgame-Perfect Nash Equilibrium (SPNE) [9] refines Nash equilibrium by ruling out such incredible threats.

Specifically, a *subgame* in the above bargaining is a game starting from an arbitrary proposer and lasts one time slot, ending up with either a disagreement or a successful transaction. The *strategy* for an agent i in the subgame are the pricing proposal (p_{ij} or q_{ij}) and a response (accept or reject). The *payoff* in each single transaction equals the utility minus cost. More precisely, for a transaction between proposer i and responder j, the payoff equals $S_{ij} - p_{ij}$ for agent i and $T_{ij} + p_{ij}$ for agent j. Here S_{ij} is the *utility* of the payer i, which equals to the number of blocks (either 0 or 1) i downloads from j in the transaction. Similarly, T_{ij} is the number of blocks the payee j downloads from i, hence $S_{ij} = T_{ji}$. Note that the *expected payoff* depends not only on the payoff in a single transaction, but also the potential payoff he can gain by reselling the blocks he gets, and the possibility of switching to alternative partners.

Given the above elements of the game, a strategy profile constitutes a *Subgame-Perfect Nash Equilibrium* (SPNE) if it induces a Nash equilibrium in every subgame, each corresponding to a possible round of negotiation. In what follows, we establish the necessary and sufficient condition for a stationary SPNE strategy in Lemma 1 and Lemma 2, respectively. We summarize the results concerning the existence and uniqueness of the SPNE in Theorem 1.

Lemma 1. *Any stationary SPNE strategy must satisfy:*

$$p^*_{ij} = \begin{cases} -\eta^{j-K} - 1, & \text{if } i > 0 \text{ and } 0 < j \le K, \\ 0, & \text{if } i = 0, \\ -V_x - 1, & \text{if } j = 0. \end{cases} \tag{1}$$

and: $$q^*_{ij} = -p^*_{ji} \tag{2}$$

where $\eta = 1 + \frac{\frac{2}{\theta} - 2}{\rho - \alpha_0}$*;* α_i *is the probability of meeting a type-*i *agent.* $\rho = \sum_{i=0}^{K} \alpha_i = \frac{1}{M}$*, and*

$$V_x = \left(\frac{2}{\theta} + \theta + \rho - 4\right)^{-1}\left(\rho - \alpha_K - \sum_{i=1}^{K-1} \frac{\alpha_i}{\eta^{K-i}} - \frac{\rho - \alpha_0}{\eta^{K-1}}\right)$$

Proof: Let U_i be the expected equilibrium payoff of a type-i agent that is unmatched, *i.e.*, he has no partner that can provide the good of interest in the current time slot. Denote M_{ij} as the equilibrium payoff of a type-i agent when he is matched with a type-j agent.

For a unmatched agent i, all payoff begins only from the next time slot, where with probability $\mu = 1 - \theta$, he leaves the market and gets zero payoff. Conditioned on the event that he remains in the market, he may either be matched to an agent of type-j with probability α_j ($0 \le j \le K$), or remains unmatched with probability $(1 - \rho)$. Therefore, the expected equilibrium payoff for a unmatched type-i agent is:

$$U_i = (1 - \theta) \cdot 0 + \theta \cdot \left[\sum_{j=0}^{K} \alpha_j M_{ij} + (1 - \rho) U_i\right] \tag{3}$$

For a matched agent i, the equilibrium payoff consists of the payoff in the current transaction, plus the expected payoff in the forthcoming time slots. The current payoff equals his utility minus the expected cost: $S_{ij} - \frac{1}{2}(p_{ij} + q_{ij})$. If he obtains one block in the current transaction, then he becomes type-$(i + 1)$ beginning from the next period and the expected payoff equals to U_{i+1}. Otherwise if $j = 0$, *i.e.*, he is matched to an agent with zero blocks, then his future payoff remains to be U_i. Therefore, the expected equilibrium payoff for a type-i agent with a type-j partner is:

$$M_{ij} = \begin{cases} S_{ij} - \frac{1}{2}(p_{ij} + q_{ij}) + U_i, & \text{if } j = 0. \\ S_{ij} - \frac{1}{2}(p_{ij} + q_{ij}) + U_{i+1}, & \text{otherwise }. \end{cases} \tag{4}$$

Note that when $i = K$, *i.e.*, the agent collects a full set of blocks for the good, then he remains in type-K until leaving the market. In equation (4) and what follows, we equate a type-$(K+1)$ agent with a type-K agent.

We proceed to characterize the SPNE prices which are closely related with the above payoff functions. Consider any subgame with agent i being the proposer, who bids price p_{ij} for the transaction. If $i>0$, then the total expected payoff of agent j from current and future payoff is $(T_{ij}+p_{ij}+U_{j+1})$. Subgame perfection requires agent i to propose a price which gives agent j no less payoff than if he rejects the proposal and remains in type-j, *i.e.*, $(T_{ij}+p_{ij}+U_{j+1}) \geq U_j$. However, if $(T_{ij}+p_{ij}+U_{j+1}) > U_j$, agent i can gain more by proposing a price that is less than p_{ij} but still acceptable by agent j. Therefore, we must have $(T_{ij}+p_{ij}+U_{j+1}) = U_j$. For the case $i=0$, agent j remains to be type j after the transaction, and thus $(T_{ij}+p_{ij}+U_j) = U_j$. In consequence,

$$U_j = \begin{cases} T_{ij}+p_{ij}+U_j, & \text{if } i=0. \\ T_{ij}+p_{ij}+U_{j+1}, & \text{otherwise }. \end{cases} \tag{5}$$

Using a symmetric argument (with roles of i and j reversed), we can obtain the SPNE price when i is the responder:

$$U_i = \begin{cases} S_{ij}-q_{ij}+U_i, & \text{if } j=0. \\ S_{ij}-q_{ij}+U_{i+1}, & \text{otherwise }. \end{cases} \tag{6}$$

In summary, *any stationary SPNE strategy must necessarily satisfy* (3), (4), (5) *and* (6). This necessary condition involves $K+1+2(K+1)^2$ linear equations and the same number of variables, including U_i, M_{ij}, p_{ij} and q_{ij}, $(0 \leq i \leq K, 0 \leq j \leq K)$. By solving this system of equations, we obtain the equilibrium prices p^*_{ij}, q^*_{ij} and the corresponding equilibrium payoff. Due to space constraint, interested readers are referred to [16] for details of the solution procedure. □

Lemma 2. *The stationary strategy with reservation prices defined in* (1) *and* (2) *is a SPNE strategy for every pairwise bargaining game.*

Proof: To prove that the threshold based stationary strategy is SPNE, it is sufficient to show that in an arbitrary subgame, either proposer in the matched pair is willing to adopt the prices p^*_{ij} and q^*_{ij}, and cannot profit more by unilaterally deviating from such equilibrium prices. The latter condition is straightforward following the equilibrium argument in proving Lemma 1, thus we only prove the former condition.

To verify that the proposer j indeed has the incentive to propose price q^*_{ij}, we need to ensure that the profit from this proposal is no less than if he remains inactive and wait for the next transaction, *i.e.*,

$$T_{ij}+q^*_{ij} \geq 0, \text{ if } j=0. \tag{7}$$

$$T_{ij}+q^*_{ij}+U_{j+1} \geq U_j, \text{ if } j>0 \tag{8}$$

For brevity, we only present the general cases where $j>0$ and $0<i<K$ (see [16] for a complete proof). Equation (8) is equivalent to:

$$T_{ij}+(U_{i+1}-U_i)+S_{ij}+(U_{j+1}-U_j) \geq 0 \tag{9}$$

From Lemma 1, we have $U_{i+1}-U_i = -\eta^{i-K}, (0<i<K-1)$, where $\eta = \frac{2a-\alpha_0-\rho}{\rho-\alpha_0} = 1+\frac{2(\theta^{-1}-1)}{\rho-\alpha_0}$. Recall that $\rho = \sum_{i=0}^{K}\alpha_i \geq \alpha_0$. Therefore, we have $\eta \geq 1$ and subsequently $-1 \leq U_{i+1}-U_i \leq 0$. Similarly $-1 \leq U_{j+1}-U_j \leq 0$. Since in this case $T_{ij}=1$, equation (9) follows directly. By a symmetric argument, we can also prove that the proposer i has the incentive to propose price p^*_{ij}, thus completing the proof for Lemma 2. □

Lemma 1 establishes that (1) and (2) are the necessary condition for an SPNE strategy, while Lemma 2 justifies the sufficiency of the condition. Since the systems of equations corresponding to the condition has a unique solution, we have the following result.

Theorem 1. *The unique stationary subgame perfect Nash equilibrium strategy is the threshold based strategy with reservation prices defined in* (1) *and* (2).

From (1) and (2), we conclude that the SPNE price depends on the coding complexity (reflected in segment size K), availability of the good (reflected in α_i), as well as the degree of market dynamics (reflected in μ). The intricate relations implied by the above theorem will be further clarified in Sec. V.

C. The Market Equilibrium and Its Stability

The analyses above have revolved around the strategically stable configurations, *i.e.*, the SPNE of each pairwise "take-it-or-leave-it" bargaining game. In this subsection, we extend the equilibrium to a temporally stable configuration, claiming that the equilibrium is insensitive to strategical manipulations of any individual agent over time.

Towards this end, we define the *expected payoff* of an agent as $R(h) = \sum_{t=0}^{\infty} R(h(t))$, where $R(h(t))$ is the payoff within time slot t when the agent adopts strategy h. Assume agents are expected payoff maximizers, then following the microeconomics literature [9], we define *market equilibrium* as a stationary strategy profile h^* that is adopted by all agents and that maximizes the expected payoff of each agent. More precisely, for each agent Υ, $R(h^*_\Upsilon, h^*_{-\Upsilon}) \geq R(h_\Upsilon, h^*_{-\Upsilon})$ for all possible strategies h, where $h^*_{-\Upsilon}$ indicates that all agents other than Υ adopt the same stationary strategy. Essentially, in a market equilibrium all agents adopt the same stationary strategy, and no single agent can gain more by strategically vary his proposals and responses during his lifetime. With this concept, we have:

Theorem 2. *In the P2P content distribution market, the threshold based strategies with reservation prices defined in* (1) *and* (2) *constitute a market equilibrium.*

Proof: Consider an agent Υ entering the market with zero blocks of the good. In searching for a payoff-maximizing policy, Υ essentially faces a Markov decision process (Fig. 2). The state space includes P_i, A_i, Z_i, and *leave*. Z_i denotes that Υ has evolved to type i and has no partner yet. P_i denotes that the agent has evolved to type i and has been selected as the proposer in a bargaining game. Each state P_i includes a subset of states $P_{ij}(0 \leq j \leq K)$, indicating Υ is matched to a partner of type j. Similar definition applies for A_i, where the agent has been selected as the responder.

When all other agents adopt the same stationary strategies defined in Lemma 1, Υ only has two policies in each state A_{ij} and P_{ij}. He either chooses *agreement* by proposing p^*_{ij} and accepting q^*_{ij}, or chooses *disagreement* by proposing $p_{ij} > p^*_{ij}$

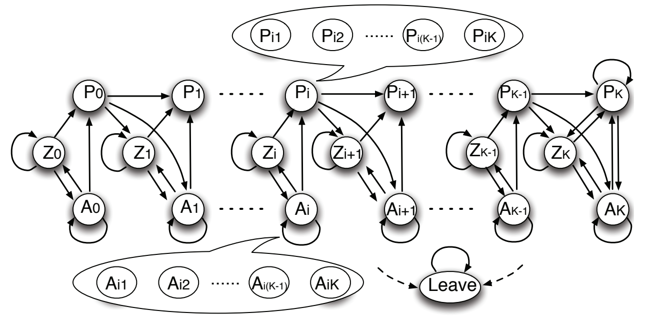

Fig. 2. The state transition diagram of an agent, assuming he adopts the *disagreement* policy in every state A_i and *agreement* in every state P_i. Each state A_i includes K substates. Substate A_{ij} is reached with probability α_j. Similar definition applies for P_i. *Leave* is an absorbing state that can be reached from any other state with probability μ.

and rejecting $q_{ij} \leq q^*_{ij}$. His policies have no impact on the states Z_i and *leave*.

Denote X^a and X^d as the expected payoff in state X when choosing *agreement* and *disagreement*, respectively. The expected payoff in each state and for each policy equals the payoff gained within the state plus the expected payoff after the policy is taken. More specifically, when the agreement policy is taken in state P_{ij}, the expected payoff is:

$$P^a_{ij} = S_{ij} - p^*_{ij} + \theta[\sum_{k=0}^{K} \frac{\alpha_k}{2}(P^*_{(i+1)k} + A^*_{(i+1)k}) + (1-\rho)U_{i+1}]$$

In this equation, $(S_{ij} - p^*_{ij})$ represents the average payoff within state P_{ij} and when the agreement policy is enforced. $P^*_{(i+1)k}$ and $A^*_{(i+1)k}$ are the corresponding optimal payoff in the next period when the agent Υ is selected as the proposer and responder, respectively. The agent Υ obtains payoff U_{i+1} if this good is not selected for transaction in the next period, which happens with probability $(1-\rho)$.

In a similar vein, we can derive the expected payoff when enforcing the disagreement policy, which is independent of j.

$$P^d_{ij} = \theta[\sum_{k=0}^{K} \frac{\alpha_k}{2}(P^*_{ik} + A^*_{ik}) + (1-\rho)U_i]$$

Consequently, the payoff difference for the two policies is:

$$\begin{aligned} P^a_{ij} - P^d_{ij} = S_{ij} - p^*_{ij} + \theta[\sum_{k=0}^{K} \alpha_k M_{(i+1)k} + (1-\rho)U_{i+1}] \\ - \theta[\sum_{k=0}^{K} \alpha_k M_{ik} + (1-\rho)U_i] \\ = S_{ij} - p^*_{ij} + U_{i+1} - U_i = S_{ij} - p^*_{ij} + q^*_{ij} - S_{ij} \geq 0 \end{aligned}$$

The last inequality follows from the definition of the equilibrium prices in Lemma 1. An intuitive explanation can be derived by contradiction. Suppose $q^*_{ij} < p^*_{ij}$, then agent j can propose q'_{ij} such that $p^*_{ij} > q'_{ij} > q^*_{ij}$, which contradicts the optimality of q^*_{ij}. Similarly, we can justify the optimality of the agreement policy in each state A_{ij}. Given that the agreement action is optimal for every state, it constitutes a stationary policy that solves the following revenue-maximizing Bellman equations in a dynamic control problem [17]:

$$J^*(P_{ij}) = \max\{P^d_{ij}, P^a_{ij}\}, J^*(A_{ij}) = \max\{A^d_{ij}, A^a_{ij}\} \quad (10)$$

Following Proposition 7.2.1 in [17], it can be easily verified that the stationary policy of agreement is the optimal policy for the payoff-maximizing problem corresponding to the market equilibrium. □

V. The Equilibrium Price and Payoff

In this section, we analyze the steady state distribution of goods availability in the coding based P2P market, and then integrate it with the previous game theoretic analysis. This leads us to a comprehensive understanding of the relation between the scarcity of goods and the equilibrium price, and the market power of network coding in this context.

A. Availability of Goods at Steady State

Our basic approach is a continuous time approximation to the evolution of the market using differential equations. We focus on a steady state of the peer population, in which the total number of agents N remains roughly constant. Assume the peers join and depart the market following a Poisson process, then the arrival rate equals the departure rate, and corresponds to the departing probability μ in the game model. Suppose the goods (segments) are *randomly selected* for downloading upon the encounter of two agents. Then one could expect that each good experiences a similar level of availability. In the following analysis, we only focus on the block distribution of a single good. We will justify the above modeling assumptions using simulations.

Denote s_i as the number of agents having at most i blocks of the good, and n_i as the number of agents holding exactly i blocks. Consider the evolution of the market during a short period Δt. The increase of s_i within Δt equals the number of departing peers each holding at least $(i+1)$ blocks, and each subsequently replaced by a new peer with zero block, which amounts to: $\mu N \Delta t \frac{\sum_{j=i+1}^{K} n_j}{N}$.

The decrease of s_i equals to the total number of peers that have i blocks and download one more block. The probability that such a peer is chosen equals $\frac{n_i}{N}$, while the probability that a segment i is chosen is $\rho = \frac{1}{M}$. The total decrease of s_i thus equals: $\Delta t \sum_{j=1}^{K} n_j \frac{n_i}{NM}$.

As $\Delta t \to 0$, the following system of differential equations captures the evolution of the market:

$$\frac{\mathrm{d}s_i(t)}{\mathrm{d}t} = \mu N \frac{\sum_{j=i+1}^{K} n_j}{N} - \sum_{j=1}^{K} n_j \cdot \frac{n_i}{NM} \quad (11)$$

$$\frac{\mathrm{d}s_0(t)}{\mathrm{d}t} = \mu N \frac{N - n_0}{N} - (N - n_0)\frac{n_0}{NM} \quad (12)$$

Solving for the steady state, and let $\phi = \frac{n_0}{N-n_0}$, we have:

$$n_0 = \mu NM, n_K = \frac{n_0}{\phi(1+\phi)^{K-1}}, \quad (13)$$

$$n_i = \frac{n_0}{(1+\phi)^i}, (0 < i < K). \quad (14)$$

To evaluate the accuracy of the above model, we have developed a C-based P2P network simulator. We simulate a dynamic P2P market following the random peer selection and segment selection policy with population $N = 10^4$ and file size $F = 1000$. The simulation lasts for 6000 periods. A server is online in the beginning and leaves after 1000 periods. The results are sampled after the market evolves to a steady state,

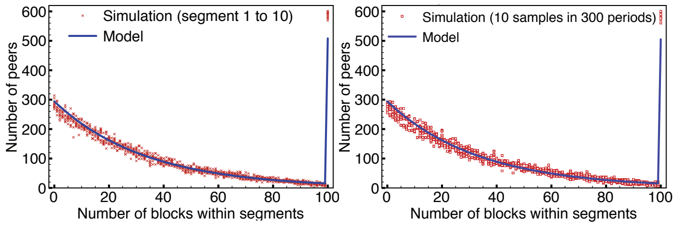

Fig. 3. The availability among different goods (left) and availability of a single good over time(right), reflected by the number of type-i ($0 \le i \le K$) agents on the market. Segment size $K = 100$ and churn rate $\mu = 0.003$.

which usually takes around several hundred periods. Fig. 3 plots the steady-state availability of goods on the market.

The availability demonstrates little variation over time and across different goods, and the model is able to capture the average number of each type of agents. More detailed experiments in [16] reveal that the variation of availability across goods generally increases with smaller segment size K. However, even in the extreme case where $K = 1$, the variation is still negligible, especially when considering $\alpha_i = \frac{n_i}{NM}$. Therefore, in the forthcoming analysis, we equate the availability value predicted by the differential equation model with the availability of each good. Note that in a streaming system such as [6], the segments are prioritized in sequence, and demonstrate considerable variation. The pricing analysis in previous section still applies to such systems, where the α_i value can be determined online by localized probing or steady-state analysis. Such extensions are beyond the scope of our current paper.

For the case with *content servers* (*i.e.*, agents who hold the entire file and never leave the market), the decrease of s_i in Δt is $(N + N_s - n_0) \cdot \frac{n_i \Delta t}{NM}$, where N_s is the number of servers online. The corresponding steady-state solutions to n_i can be easily derived similar to the case without servers, and are omitted in this paper. Interested readers are referred to [16] for detailed analysis and simulation results.

B. Equilibrium Properties of the Steady-State Market

We proceed to integrate the SPNE and market equilibrium analysis in Sec. IV with the steady state model. Our emphasis is on how network coding affects the equilibrium properties of the market. We use asymptotic approximations to derive theoretical insights, and use exact numerical simulations to crystalize such effects. We focus on three metrics: *entry price*, *lifetime payoff* and *seeder's payoff*, which will be defined below. The former two metrics are closely related with agents' incentive to join in a market economy, while the latter is closely related with a seeder's incentive to serve others after he obtains all the goods, and with the server's incentive to keep the market online.

1) Entry Price: When entering the market, an agent has no blocks to exchange with others, and thus must bring an initial capital that allows him to buy one block of a certain good. The amount of initial capital needed to start transacting a good is referred to as the *entry price* of that good. For the steady state market with SPNE strategies, the entry price equals $\max\{p_{01}^*, p_{02}^* \cdots p_{0K}^*, q_{01}^*, q_{02}^*, \cdots, q_{0K}^*\}$. Since $p_{0j}^* = 0$ for all $0 \le j \le K$, we only need to focus on q_{0j}^*.

From Lemma 1, we know $q_{0j}^* = 1 + V_x$, and q_{0j}^* is independent of j. By integrating with the steady state analysis in Sec. V-A, and noting that $\alpha_i = \frac{n_i}{NM}$, we have:

$$\begin{aligned}V_x &= \frac{1}{\frac{2}{\theta} + \theta + \rho - 4}[\rho - \frac{1}{\phi(1+\phi)^{K-1}} \\ &- \frac{\mu(\eta^{1-K} - (1+\phi)^{1-K})}{1+\phi-\eta} - \eta^{1-K}(\rho-\mu)] \\ &= \frac{1}{\frac{2}{1-\mu} - 2 + \rho - 2\mu}[\rho - \rho(1-\frac{\mu}{\rho})^K - \\ &\frac{(2\mu + 2 - \frac{2}{1-\mu})(\rho - \mu + 2 + \frac{2}{1-\mu}) - \mu\rho^{1-K}}{2 + \mu - \frac{2}{1-\mu}(\rho-\mu)^{-K}}]\end{aligned} \quad (15)$$

To avoid more complex exposition, we mainly focus on the closed form solutions to two extreme cases, namely the non-coding and full-coding case. We evaluate the general partial coding cases through numerical simulation. For the non-coding case (*i.e.*, $K = 1, M = F$), the above can be reduced to:

$$V_x = \frac{2\mu - \rho}{\frac{2}{1-\mu} - 2 - (2\mu - \rho)} \quad (16)$$

Considering that the file size F is usually very large, the entry price equation can be further reduced by ignoring the second order terms of μ:

$$\begin{aligned}q_{0j}^* = 1 + V_x &= \frac{2 - 2(1-\mu)}{2 - 2(1-\mu) + \frac{1-\mu}{M} - 2\mu(1-\mu)} \\ &\approx \frac{2\mu M}{1-\mu} \approx 2\mu M\end{aligned} \quad (17)$$

Therefore, for the non-coding case, when file size is fixed, entry price increases approximately linearly with the churn rate, namely the impatience of agents.

For the full-coding case, the entire file is a single segment (*i.e.*, $M = 1, K = F$), and V_x can be reduced to:

$$\begin{aligned}V_x &\approx \frac{1}{\frac{2}{1-\mu} - 1 - 2\mu}(1 - (1-\mu)^K - \frac{(1-\mu)^{K+1}}{1+\mu}) \\ &= \frac{1}{\frac{2}{1-\mu} - 1 - 2\mu}(1 - \frac{2(1-\mu)^K}{1+\mu})\end{aligned}$$

When μ is close to 0, the above can be simplified to $V_x \approx 1 - 2(1-\mu)^K$. When μ is close to 1, we can ignore the K-th order terms, and obtain the Taylor series of V_x at $\mu = 1$:

$$V_x = -\frac{1}{2}(\mu - 1) + \frac{3}{4}(\mu-1)^2 + O((\mu-1)^2) \quad (18)$$

which is a decreasing function when μ approaches 1. Therefore, for full-coding, the entry price has distinct properties in two regions roughly defined with respect to churn rate. In the *low churn rate region* (μ close to 0), entry price increases with churn rate and decreases with file size. However, in the *high churn rate region* (μ close to 1), file size is irrelevant, and entry price decreases with churn rate. We proceed to numerically justify these intuitions with more accuracy and for the partial coding case.

Fig. 4 plots the curves derived directly from (15). As can be induced from the figure and the steady-state analysis, a content distribution protocol is stable only if the churn rate μ is less than $\frac{1}{M}$, otherwise the agents holding zero blocks will eventually dominate the market and the good will vanish. Therefore, under

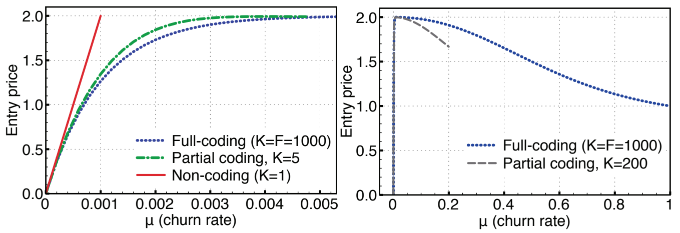

Fig. 4. The entry price under different coding complexity and churn rate. K and F are segment size and file size, respectively. Number of segments $M = \frac{F}{K}$. A protocol is stable only if $\mu \leq \frac{1}{M}$.

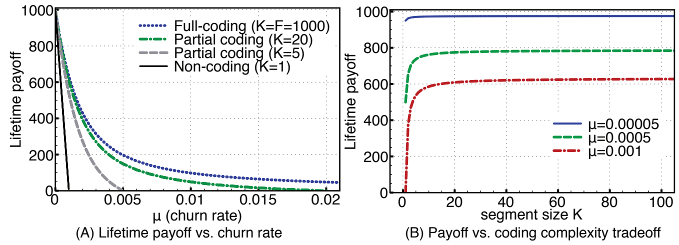

Fig. 5. The lifetime payoff as a function of churn rate and coding complexity.

a fixed file size F ($F = 1000$ in all our numerical simulation), higher coding complexity (larger K) corresponds to smaller M, allowing for larger churn rate μ. This means that *a P2P content distribution market is more tolerant to agents' impatience when using network coding, especially the full-coding protocol*.

Since $F = M$ for the non-coding protocol, it is only stable for $\mu \in (0, \frac{1}{F})$. In this region, the full-coding protocol has the lowest entry price as it results in the highest availability level. Entry price increases as the coding complexity decreases, meaning that *lower entry price is obtained at the cost of coding complexity*. When μ is sufficiently large, however, using smaller segment size may result in lower entry price. This is because the resale value of goods is degraded in the high churn rate region. With smaller K, the resale value is shared by a larger number of goods, hence the per-good value decreases, resulting in lower entry price. In the extreme case $\mu = 1$, a good has no resale value, and its entry price equals the utility value 1.

We remark that a real-world P2P market tends to survive in the low churn rate region. The following back-of-the-envelop calculation, based on the measurements in [8], may be convincing. Consider a file with $F = 1000$ blocks, each of size 1 MB. Suppose peer upload/download bandwidth is 0.5 MB/s, then the duration of a period in our discrete model is 2 seconds. According to [8], an agent's average lifetime equals 60 minutes, corresponding to 1800 periods, thus $\mu = \frac{1}{1800} \approx 5 \times 10^{-4}$, which is obviously in the low churn rate region.

2) Lifetime Payoff: We define *lifetime payoff* as an agent's expected payoff when he enters the steady-state market. Initially, an agent holds zero block, hence his expected payoff for each good equals U_0, and the lifetime payoff equals to MU_0 as M represents the total number of goods on sale.

From the equilibrium analysis established when proving Lemma 1, we have $U_1 - U_0 = V_x$ and

$$(\frac{1}{1-\mu} + \rho - 1)U_0 = \alpha_0 U_0 + \frac{1}{2}(U_0 + U_1 + 1)(\rho - \alpha_0)$$

By solving these two equations, we obtain:

$$U_0 = \frac{1}{2}((1-\mu)^{-1} - 1)^{-1}(\rho - \mu)(1 + V_x) \quad (19)$$

For the full-coding case, by expanding V_x, we have the following approximation:

$$U_0 \approx \frac{1}{2}\mu^{-1}(1-\mu)^2, 0 < \mu < 1 \quad (20)$$

For non-coding, we have:

$$U_0 \approx (1-\mu)(1-\mu M) \approx 1 - \mu M, 0 < \mu < \frac{1}{M} \quad (21)$$

We conclude from (20) and (21) that the lifetime payoff monotonically decreases as churn rate increases from 0 to $\frac{1}{M}$. The rate of decreasing is approximately linear for non-coding and approximately sublinear (for $0 < \mu < 1$) for full-coding. Therefore, *network coding can alleviate the market's instability facing churns, and can expand the region in which the agents have positive payoff and are motivated to join*.

From the general cases plotted in Fig. 5(A), we can see that higher coding complexity always induces higher level of payoff. For any configuration, payoff approaches 0 as churn rate approaches $\frac{1}{M}$. As churn rate approaches 0, all configurations approach the highest possible payoff, which equals to the file size F. In summary, *the advantages of network coding are best demonstrated in a dynamic market with impatient agents, and such advantages diminish as the agents become more patient*.

Fig. 5(B) characterizes the tradeoff between lifetime payoff and coding complexity. In general, payoff increases with coding complexity, namely the segment size K. However, the increase is negligible when K is beyond a small threshold that decreases with churn rate. This implies that *encoding a small number of blocks is sufficient to harvest the major benefit of network coding*.

3) The Seeder's Payoff: We refer to an agent who has collected all blocks of all goods as a *seeder*. At the moment an agent has fulfilled a single good, his expected payoff during the residual lifetime is U_K. Therefore, after he becomes a seeder, the expected payoff equals MU_K.

From the proof of Lemma 1, we have:

$$U_K = (U_K - U_1) + (U_1 - U_0) + U_0$$
$$= \sum_{k=1}^{K-1}(U_{k+1} - U_k) + V_x + U_0 = \frac{(1-\eta^{1-K})}{1-\eta} + V_x + U_0$$

For non-coding, we have: $MU_K = M(V_x + U_0) \approx \mu M^2$. When file size is fixed, the seeder's payoff is approximately linearly increasing with churn rate. For the full-coding case, we can easily verify, following the approximations in the above subsections, that the seeder's payoff demonstrates different characteristics depending on the churn rate. However, we only present the numerical results due to space limitation.

From Fig. 6, we observe that the seeder's payoff increases monotonically with churn rate in the low churn rate region. Lower coding complexity results in higher revenue for the seeders, but at the cost of a lower level of tolerance to churns. The intuition behind is that with low coding complexity, the agent's impatience problem becomes more threatening, thus a seeder who holds all the goods has higher bargaining power on the market, and harvests more profit through the decentralized bargaining. In the high churn rate region, similar to entry price,

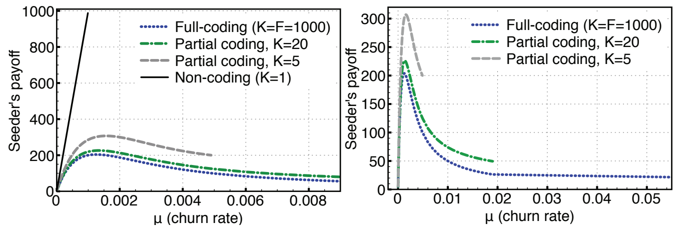

Fig. 6. The seeder's expected payoff on the steady-state market.

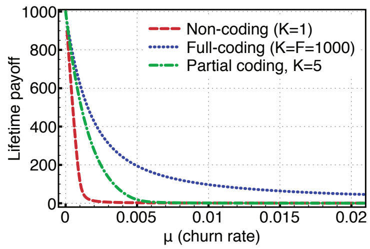

Fig. 7. The lifetime payoff in the presence of servers.

seeder's payoff decreases due to the dominant decrease of resale value.

4) Equilibrium with Servers: Due to space constraint, we only present numerical results for the case with servers. Fig. 7 plots the lifetime payoff when 100 servers facilitate 10^4 downloaders. Here all possible churn rate within $(0, 1)$ is supported by the market because the servers ensure that each good is constantly online. However, the lifetime payoff for low-complexity coding protocols and the non-coding protocol suffer from a steep decrease with churn rate, implying that agents are less motivated to join the market.

If we deem each server as a special seeder, who refreshes his life with probability μ every period, then the seeder's payoff is equivalent to the time-average payoff of the server, which is termed *per-server payoff*. The per-server payoff decreases as more servers join the market (Fig. 8). This is because the competition among servers reduces the individual bargaining power, thus reducing the revenue from each pairwise bargaining game. In addition, in the low churn rate region, non-coding has a much higher level of payoff than high-complexity coding protocols. This implies in a real-world P2P market, it is more beneficial for the servers to not use network coding, though the expected payoff of downloaders decreases with low coding complexity. Therefore, the two forces — content servers and downloaders — may need an additional bargaining game over the coding complexity to be employed.

VI. Conclusion

In this paper, we develop a theoretical framework that quantifies the market power of network coding in a non-cooperative P2P content distribution system. We model the network participants as market agents who purchase and resell goods (data segments), and strategically set prices according to availability of the goods. We then rigorously characterize the pricing strategies that constitute a subgame perfect Nash equilibrium, as well as a market equilibrium which is proof against individual temporal deviations. Combined with a steady-state modeling of the goods availability, this analysis allows us to derive closed-form solutions that capture the effects of network

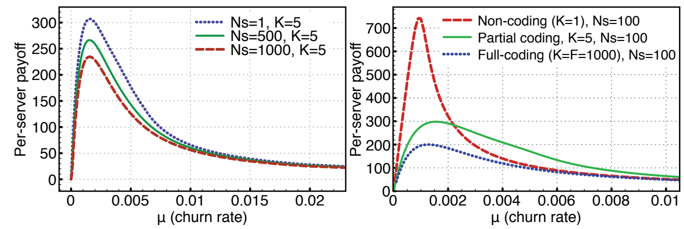

Fig. 8. The per-server payoff for varying churn rate (μ) and number of servers (N_s).

coding in a dynamic market. In particular, network coding improves the market's resilience to impatient agents, at the cost of high coding complexity. More importantly, it enhances the agents' incentive to join by lowering the entry price, and by increasing their expected payoff. Notably, such coding advantages diminish as the agents become more patient, *i.e.*, when the market experiences lesser dynamics. We have focused on a steady-state market in which agents adopt stationary strategies. An interesting future avenue is to understand the transient properties of the market and implement distributed pricing algorithms that lead the market to the stationary regime.

References

[1] C. Gkantsidis and P. R. Rodriguez, "Network Coding for Large Scale Content Distribution," in *Proc. of IEEE INFOCOM*, 2005.
[2] M. Feldman and J. Chuang, "Overcoming Free-riding Behavior in Peer-to-Peer Systems," *ACM SIGecom Exch.*, vol. 5, no. 4, 2005.
[3] T. Ho, M. Medard, J. Shi, M. Effros, and D. Karger, "On Randomized Network Coding," in *Proc. of Allerton Conference on Communication, Control and Computing*, 2003.
[4] D. Niu and B. Li, "On the Resilience-Complexity Tradeoff of Network Coding in Dynamic P2P Networks," in *Proc. of IEEE IWQoS*, 2007.
[5] P. A. Chou, Y. Wu, and K. Jain, "Practical Network Coding," in *Proc. of Allerton*, 2003.
[6] M. Wang and B. Li, "R2: Random Push with Random Network Coding in Live Peer-to-Peer Streaming," *IEEE JSAC*, vol. 25, no. 9, 2007.
[7] D. M. Chiu, R. W. Yeung, J. Huang, and B. Fan, "Can Network Coding Help in P2P Networks?" in *Proc. of NetCod*, 2006.
[8] S. Saroiu, P. K. Gummadi, and S. Gribble, "A Measurement Study of Peer-to-Peer File Sharing Systems," in *Proc. of MMCN*, 2002.
[9] M. Osborne and A. Rubinstein, *Bargaining and Markets*. Academic Press, 1990.
[10] K. Hosanagar, P. Han, and Y. Tan, "Diffusion Models for P2P Content Distribution," *to appear in Information Systems Research*.
[11] "MMAPPS - Market Management of Peer to Peer Services," 2004, http://www.mmapps.info.
[12] C. Courcoubetis and R. Weber, "Incentives for Large Peer-to-Peer Systems," *IEEE JSAC*, vol. 24, no. 5, 2006.
[13] C. Aperjis and R. Johari, "A Peer-to-Peer System As an Exchange Economy," in *Proc. of the Workshop on Game Theory for Communications and Networks (GameNets)*, 2006.
[14] E. Ahmed, A. Eryilmaz, A. Ozdaglar, and M. Medard, "Economic Aspects of Network Coding," in *Proc. of Allerton*, 2006.
[15] D. Turner and K. Ross, "A Light Weight Currency Paradigm for the P2P Resource Market," in *Proc. of International Conference on Electronic Commerce*, 2004.
[16] X. Zhang and B. Li, "On the Market Power of Network Coding in P2P Content Distribution Systems, http://www.eecg.toronto.edu/~bli/techreports/CodMart.pdf," Dept. of ECE, University of Toronto, Tech. Rep., Aug. 2008.
[17] D. P. Bertsekas, *Dynamic Programming and Optimal Control (Vol. I)*. Athena Scientific, 2005.