Auction-based P2P VoD Streaming: Incentives and Optimal Scheduling

Chuan Wu, The University of Hong Kong
Zongpeng Li, University of Calgary
Xuanjia Qiu, The University of Hong Kong
Francis C.M. Lau, The University of Hong Kong

Real-world large-scale Peer-to-Peer (P2P) Video-on-Demand (VoD) streaming applications face more design challenges as compared to P2P live streaming, due to higher peer dynamics and less buffer overlap. The situation is further complicated when we consider the selfish nature of peers, who in general wish to download more and upload less, unless otherwise motivated. Taking a new perspective of distributed dynamic auctions, we design efficient P2P VoD streaming algorithms with simultaneous consideration of peer incentives and streaming optimality. In our solution, media block exchanges among peers are carried out through local auctions, in which budget-constrained peers bid for desired blocks from their neighbors, which in turn deliver blocks to the winning bidders and collect revenue. With strategic design of a discriminative second price auction with seller reservation, a supplying peer has full incentive to maximally contribute its bandwidth to increase its budget; requesting peers are also motivated to bid in such a way that optimal media block scheduling is achieved effectively in a fully decentralized fashion. Applying techniques from convex optimization and mechanism design, we prove (a) the incentive compatibility at the selling and buying peers, and (b) the optimality of the induced media block scheduling in terms of social welfare maximization. Large-scale empirical studies are conducted to investigate the behavior of the proposed auction mechanisms in dynamic P2P VoD systems based on real-world settings.

Categories and Subject Descriptors: C.2.4 [**Computer-Communication Networks**]: Distributed Systems—*Distributed Applications*

General Terms: Algorithm, Design, Performance

Additional Key Words and Phrases: Peer-to-peer streaming, auction, incentive, optimal scheduling

ACM Reference Format:

Wu, C. and et al. 2011.Auction-based P2P VoD Streaming: Incentives and Optimal Scheduling. ACM TOMCCAP. 2, 3, Article 1 (May 2010), 20 pages.
DOI = 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

This work is supported in part by Hong Kong RGC GRF grant under the contract 718710E.
Author's address: C. Wu, X. Qiu, F. C.M. Lau, Department of Computer Science, The University of Hong Kong, Hong Kong; email: {cwu, xjqiu, fcmlau}@cs.hku.hk; Z. Li, Department of Computer Science, University of Calgary, 2500 University Drive NW, Calgary AB T2N 1N4, Canada; email: zongpeng@ucalgary.ca

DOI 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

1. INTRODUCTION

Following the success of P2P live media streaming, large-scale P2P VoD streaming applications have been successfully deployed over the Internet [PPL ecom][UUS ecom], providing thousands of videos to millions of users. The state-of-the-art P2P VoD applications are based on the design philosophy of having peers watching the same video exchange available media blocks, in order to alleviate the server load. As compared to live streaming, P2P VoD system design presents new fundamental challenges, due to its higher level of streaming instabilities caused by VCR operations, and the inherently lower levels of content overlap caused by asynchronous peer playback [Huang et al. 2008]. Consequently, effective streaming strategy design becomes particularly important for a P2P VoD system.

A related challenge that further complicates the P2P VoD design is to handle the selfish nature of real-world streaming peers. Unless otherwise regulated, a peer naturally wishes to maximize its download rate and minimize upload contribution, compromising the foundation of a P2P system [Habib and Chuang 2006][Sung et al. 2006]. The situation is exacerbated by the emerging P2P software war, in which the system designers tend to furnish their client software with increasing end-user control, in the hope of attracting a larger market share.

Given these challenges, an ideal P2P VoD solution should contain a set of judiciously designed algorithms, under which (a) peers are automatically willing to contribute upload bandwidth for their own benefit, and (b) the induced block scheduling optimizes the utilization of such upload bandwidth, leading to maximum aggregated playback “satisfaction”. For practical applicability, such algorithm design should also be simple and fully decentralized.

Auction-based algorithm design is known to provide elegant solutions to a collection of network optimization problems, which are simple, efficient, fully distributed, and allow intuitive interpretations [Nisan et al. 2007][Lazar and Semret 1999]. We recently applied variations of auction algorithms to bandwidth allocation in P2P live streaming [Wu et al. 2008], with a sole focus on solution optimality but not on peer incentivization. While previous studies are based on *media flow* auctions, the higher level of peer dynamics coupled with a lower level of playback synchrony in P2P VoD streaming require a more realistic model of *media block* scheduling. In this paper, we approach the block scheduling and peer incentive problems in P2P VoD streaming simultaneously, through a set of auction-based streaming strategies.

We model media block exchanges among neighboring peers into a decentralized collection of dynamic, iterated auctions, where block valuations depend on block rareness and playback deadlines. Each peer is furnished with a budget and assumes dual roles of both a *bidder* and a *seller*. As a bidder, it bids in neighboring auctions for desired blocks, and pays prices out of its available budget. As a seller, it auctions buffered blocks, honors winning bids by delivering the corresponding blocks, and charges strategically-set prices to re-build its budget. Towards the two separate goals of upload incentives and effective block scheduling, we carefully tailor a *discriminative second price auction with seller reservation* for the sellers, and design a *truthful start with iterative price discovery* strategy for the bidders.

We show that our auction mechanism and bidding strategy represent equilibrium strategies [Nisan et al. 2007] at the sellers and the bidders, respectively. Furthermore, the auction mechanism we tailor is better than other possible representative auction mechanisms for revenue maximizing at a seller, and the bidding strategy is utility maximizing for a bidder. Based on convex optimization theory, in particular by utilizing the KKT optimality conditions [Boyd 2004] and by analyzing the integrality gap of the integer programs modeling local and global welfare optimization, we prove that the locally administrated auctions may act in concert and achieve social welfare maximization in terms of block distribution and streaming quality, under peer budget constraints. Finally, extensive empirical studies

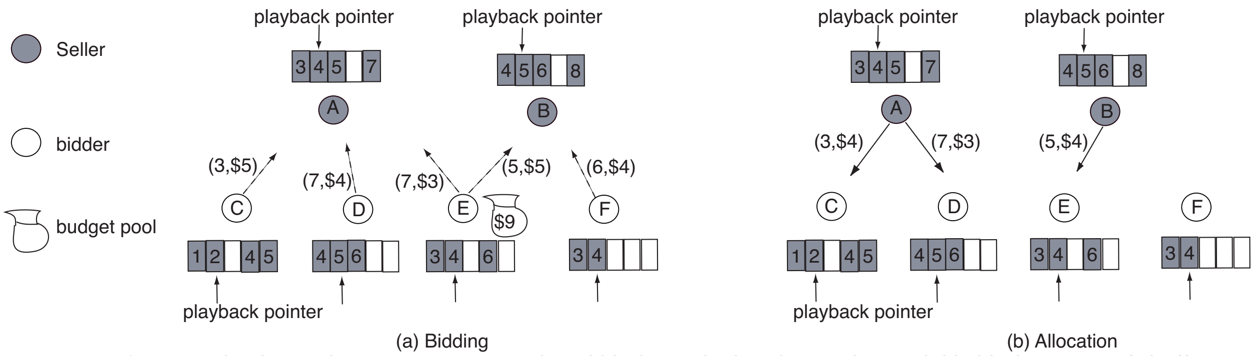

Fig. 1. An example of two adjacent auctions: numbered blocks in shadow denote the available blocks in a peer's buffer.

are conducted to investigate behaviors of the auction protocols in dynamic P2P VoD settings resembling real-world systems.

In the rest of the paper, we present our P2P VoD model and algorithm design in Sec. 2 and Sec. 3, perform theoretical and empirical studies of the proposed auction algorithms in Sec. 4 and Sec. 5, review related research in Sec. 6, and present concluding discussions in Sec. 7.

2. THE P2P VOD AUCTION MODEL

Consider a P2P VoD streaming system based on the typical *pull-based* design over mesh overlay topologies [Huang et al. 2008]. Each video is a sequence of *media blocks* streamed from the media server into a P2P overlay of interconnected peers playing the video. Upon joining the overlay, a peer obtains from a tracker a set of neighbors with similar playback progresses. A peer caches recently downloaded blocks in its buffer, periodically exchanges buffer availability bitmaps (*buffer maps*) with its neighbors, and requests desired blocks from them.

We organize block exchanges in a P2P VoD overlay into a collection of locally administrated auctions, each representing a market where a peer sells media blocks to neighbors. Each peer is furnished with a budget based on virtual currencies, for participating in adjacent auctions and bidding for desired blocks. A peer's initial budget is established when it first registers an account in the application, and can be used for block purchase in multiple video overlays. When the peer goes offline, any surplus remains in its account, and becomes its starting budget upon re-join.

We assume that the budget can be implemented by an existing virtual currency protocol [Vishnumurthy et al. 2003][Turner and Ross 2004] or as the simple credits in private BitTorrent systems [Liu et al. 2010]. These work study the implementation of a robust budget/credit system, which are complementary to our work: They strike to ensure the integrity and robustness of the payments, while leaving pricing strategies to the specific application; our work focuses on the design of incentive-compatible pricing schemes, while relying on these existing work for implementing secure payments and transactions between peers.

The auction at each peer is carried out periodically. In each round, it receives bids from requesting neighbors, transfers blocks to winning bidders, then computes and collects charges. Along with the media streaming process, each auction executes round by round with continuously changing blocks and possibly different bidders, constituting a dynamic *multi-unit multi-round auction*. Fig. 1 shows an example of two auctions, where C, D, and E bid at A, and E and F bid at B.

Let $\mathcal{K}_{ij}$ be the set of blocks available at i but not j. A bid from j to i is a pair $b_{ij}^{(k)} = (I_{ij}^{(k)}, p_{ij}^{(k)})$. Here $k \in \mathcal{K}_{ij}$ is the block requested, with ID $I_{ij}^{(k)}$ and bidding price $p_{ij}^{(k)}$. We denote the set of all possible bids as $\mathcal{B}$. Let O_i be the maximum upload capacity at i, and $o_i \in [0, O_i]$ be the bandwidth i decides to

Table I. Notation

$a_{ij}^{(k)}$	1/0: i allocates k to j/not	$\mathcal{K}_{ij}$	blocks i can help j with
$b_{ij}^{(k)}$	a bid from j to i for k	$\mathcal{D}_i$	i's neighbor set
$c_{ij}^{(k)}$	i's charge to j for k	$\tilde{p_i}$	market price at i
$v_{ij}^{(k)}$	j's valuation of k at i	O_i	i's max. upload capacity
$x_{ij}^{(k)}$	1/0: j bids for k at i/not	o_i	i's upload contribution
$p_{ij}^{(k)}$	bidding price	e_j	j's budget
$\ddot{q}_{ij}$	market price estimate	$\hat{e}_j$	j's used budget
z_j^k	1/0: bidder j requests block k from one neighbor or not		

contribute in a round. We assume media blocks are of equal size taking exactly one round to transmit. Therefore O_i (o_i) equals the number of blocks i can (wishes to) transmit in a round. Let $a_{ij}^{(k)}$ be a binary variable indicating whether bid $b_{ij}^{(k)}$ is successful (1 = yes, 0 = no), and $c_{ij}^{(k)}$ be the price i decides to charge if successful. $\mathcal{D}_i$ is the set of neighbors of i. Given a vector of bids $\mathbf{b_i} = (b_{ij}^{(k)})$ (let $m = |\mathbf{b_i}|$) received by i, an auction mechanism decides the number of bids to accept, the winning bids, and the prices to charge:

Definition 1. A *P2P VoD auction mechanism* M is a triple $(\mathsf{A}, \mathsf{C}, \mathsf{U})$, where:

—$\mathsf{A} : \mathcal{B}^m \to [0,1]^m$ is the allocation rule that maps $\mathbf{b_i} = (b_{ij}^{(k)})$ to $\mathbf{a_i} = (a_{ij}^{(k)})$.

—$\mathsf{C} : \mathcal{B}^m \to \mathbb{R}_+^m$ is the charging scheme that maps $\mathbf{b_i} = (b_{ij}^{(k)})$ to $\mathbf{c_i} = (c_{ij}^{(k)})$.

—$\mathsf{U} : \mathcal{B}^m \times \mathbb{R}_+ \to [0, O_i]$ is the upload capacity contribution strategy that maps $\mathbf{b_i} = (b_{ij}^{(k)})$ and O_i to o_i.

A peer j may bid at multiple neighbors in $\mathcal{D}_j$. It has a valuation $v_{ij}^{(k)}$ for each block $k \in \mathcal{K}_{ij}$. Let e_j be the budget j owns. Let $h_{ij} \in \mathcal{H}$ represent historical information regarding j's bidding experience at i, including its own previous bids and whether they were successful. Given a vector of n valuations $\mathbf{v_j} = (v_{ij}^{(k)})$ for the n blocks j desires from its neighbors, a history vector $\mathbf{h}_j = (h_{ij})$ for the l auctions j participates in, and the budget e_j, a bidding strategy at j is defined as follows:

Definition 2. A *bidding strategy* $s_j : \mathbb{R}_+^n \times \mathcal{H}^l \times \mathbb{R}_+ \to \mathcal{B}^n$ for j is a function that maps $(\mathbf{v}_j, \mathbf{h}_j, e_j)$ to j's bid vector $\mathbf{b_j}$. It describes how j determines its bids to each neighbor i, based on its block valuations, historical information, and budget.

Finally, we require that a peer j bids for a block at no more than one neighbor in the same round to avoid duplicate transmissions, and that a peer's total bidding price can not exceed its budget e_j. We summarize important notation in Table 2.

3. THE P2P VOD AUCTION DESIGN

We now present the design of the auction mechanism and bidding strategy, for efficient block scheduling and incentive compatibility at selfish peers.

3.1 Auction Mechanism at a Seller

Since bidding for desired blocks requires a budget, a peer naturally wishes to maximize its budget through revenue making from its own auction.

Definition 3. The *revenue* a seller i obtains in one round of an auction M is $\sum_{j \in \mathcal{D}_i} \sum_{k \in \mathcal{K}_{ij}} a_{ij}^{(k)} c_{ij}^{(k)}$, *i.e.*, total charge of the winning bids.

At the end of a round, i's budget e_i is incremented by its revenue, which can be used for i's future block bidding. We now describe the VoD auction mechanism $\mathsf{M} = (\mathsf{A}, \mathsf{C}, \mathsf{U})$ for revenue maximization at each seller.

3.1.1 *The allocation rule* A. Since charges are upper-bounded by bidding prices, a seller prefers the highest bids. Assume $o_i \leq O_i$ is the amount of upload capacity i decides to contribute using strategy U. The allocation should be feasible:

$$\sum_{j \in \mathcal{D}_i} \sum_{k \in \mathcal{K}_{ij}} a_{ij}^{(k)} \leq o_i, a_{ij}^{(k)} \in \{0, 1\}.$$

We design the allocation rule at seller i as: order received bids by bidding prices, and maximally sell blocks in that order, within the available upload bandwidth o_i (break ties randomly).

In Fig. 1, A receives three bids, with bidding prices 5, 4, 3 for blocks 3, 7, 7, from C, D, and E, respectively. Assume $o_A = 2$, A will sell block 3 to C and 7 to D.

3.1.2 *The charging scheme* C. Towards the seller's goal of revenue maximization, we design a *discriminative second price auction* mechanism to decide the charge for each winning bid. A seller may charge different prices during the same round: it charges a winning bid (with $a_{ij}^{(k)} = 1$) the bidding price from the immediately lower bid, *i.e.*, $c_{ij}^{(k)} = p_{ij'}^{(k')}$. Here $p_{ij'}^{(k')}$ is the bidding price in bid $b_{ij'}^{(k')}$, which is immediately lower than $p_{ij}^{(k)}$ in bid $b_{ij}^{(k)}$, among all the bids seller i has received. In Fig. 1, the winning bids from C and D will be charged the bidding prices from D and E, *i.e.*, 4 and 3, respectively.

3.1.3 *The upload capacity contribution strategy* U. Another rational idea for revenue maximization is to sell as many blocks as possible. In a multi-round auction where selfish bidders may probe the *market price*, the seller needs to act more strategically though.

Definition 4. The *market price* in an auction round at a seller i, $\tilde{p}_i$, is the highest bidding price among the losing bids at i.

For example, in Fig. 1, the market price at seller A is 3, the bidding price in the losing bid from E. In order to maintain a competitive local market, a seller i decides its upload capacity contribution as follows: when the number of bids m is larger than O_i (demand higher than supply), i contributes all O_i; when O_i is larger than or equal to m (sufficient supply), the seller may keep bandwidth supply slightly lower than the demand, *i.e.*, delivering one fewer block than requested, which we refer to as *seller reservation*:

$$o_i = \begin{cases} O_i, & \text{if } m > O_i, \\ m - 1, & \text{if } m \leq O_i \end{cases} \quad (2)$$

In Fig. 1, even if O_A is 3 instead of 2, A still wishes to sell 2 blocks only. Otherwise, if all three bids are entertained, A's market price becomes 0 (assuming a virtual 4th bid with bidding price of 0).

Similar selling strategies have been discussed in the existing literature for multi-product auction [Armstrong 1996][Kirkegaard and Overgaard 2008]. Analysis in Sec. 4.1 will show that such seller reservation is the best strategy for a seller in a practical dynamic P2P VoD market with intense market competitions, outperforming no reservation or more aggregative reservation. We summarize the seller's auction protocol in Algorithm 1 at the end of the section.

3.2 Bidding Strategy at a Bidder

In each round, a utility-maximizing bidder j strategically determines the blocks to bid for and prices to offer, within its budget e_j. It first computes the price it is willing to pay for each block, and then selects blocks with the highest marginal utilities to actually bid for from different sellers.

Definition 5. The *utility* of a block $k \in \mathcal{K}_{ij}$ to a bidder j is $v_{ij}^{(k)} - c_{ij}^{(k)}$, *i.e.*, j's valuation of the block minus the charge it pays to seller i.

A bidder's valuation $v_{ij}^{(k)}$ of a block k captures the "value" associated with receiving k from i. In P2P VoD streaming, the valuation can reflect (a) the urgency level of downloading the block, *i.e.*, the closer the playback deadline, the higher a block is valued, and (b) the rareness level of the block, *i.e.*, the rarer a block k is among bidder j's neighbor set $\mathcal{D}_j$, the higher resale opportunity k represents. Concrete valuation methods are presented in simulations in Sec. 5. For now we only need to assume that $v_{ij}^{(k)}$ is a function over $[0, 1]$ that is differentiable, non-decreasing, and quasi-linear [Lazar and Semret 1999]. $v_{ij}^{(k)}(1)$ (also short as $v_{ij}^{(k)}$ hereafter) is the valuation for the block when it is completely downloaded, and $v_{ij}^{(k)}(0)$ equals 0.

3.2.1 *Deciding bidding prices.* In the first round a peer j bids at a neighbor i (*e.g.*, upon joining a VoD overlay), a reasonable strategy for j is to bid its true valuation for a block, since j has essentially zero information regarding other buyers' bids and the market price:

$$p_{ij}^{(k)} = v_{ij}^{(k)}, \forall i \in \mathcal{D}_j, k \in \mathcal{K}_{ij}. \quad (3)$$

In the subsequent rounds, the bidder may probe the lowest prices to win desired blocks based on its past bids, in order to make the most of its budget. We design the following *truthful start with iterative price discovery* strategy at the bidders: The first time j bids at i, j decides the bidding prices for blocks in $\mathcal{K}_{ij}$ as the true block valuations computed by (3). In each subsequent round, j maintains a market price estimate for i, $\ddot{q}_{ij}$. j may bid for multiple blocks at i in round t. If there are successful bids, $\ddot{q}_{ij}$ is set to slightly lower than the lowest charge, *i.e.*, $\ddot{q}_{ij} = \min_k c_{ij}^{(k)}[t] - \delta, \delta > 0$; if all the bids fail, $\ddot{q}_{ij}$ is larger than the highest bidding price, *i.e.*, $\ddot{q}_{ij} = \max_k p_{ij}^{(k)}[t] + \delta$. In round t+1, j computes the new bidding price for a block $k \in \mathcal{K}_{ij}$ as the minimum between its true valuation and the market price estimate, *i.e.*,

$$p_{ij}^{(k)}[t+1] = \min(v_{ij}^{(k)}[t+1], \ddot{q}_{ij}), \forall i \in \mathcal{D}_j, k \in \mathcal{K}_{ij}. \quad (4)$$

3.2.2 *Block request strategy.* Given bidding prices computed by (4), bidder j selects blocks from $\mathcal{K}_{ij}$ to actually bid for at each seller $i \in \mathcal{D}_j$. Let $x_{ij}^{(k)}$ be a binary variable indicating whether k is selected (1 = yes, 0 = no). The optimal block request strategy can be modeled as an integer program:

Maximize $\quad \sum_{i \in \mathcal{D}_j} \sum_{k \in \mathcal{K}_{ij}} (v_{ij}^{(k)}(x_{ij}^{(k)}) - p_{ij}^{(k)} x_{ij}^{(k)}) \quad (5)$

Subject to:

$$\begin{cases} \sum_{i \in \mathcal{D}_j} \sum_{k \in \mathcal{K}_{ij}} p_{ij}^{(k)} x_{ij}^{(k)} \leq e_j & (6) \\ \sum_{i \in \mathcal{D}_j} x_{ij}^{(k)} = z_j^k \quad \forall k \in \mathcal{K}_{ij} & (7) \\ x_{ij}^{(k)}, z_j^k \in \{0, 1\} \quad \forall i \in \mathcal{D}_j, \forall k \in \mathcal{K}_{ij} & (8) \end{cases}$$

The objective function in (5) represents the maximization of j's overall utility from all requested blocks. Comparing it to Definition 5, we note that bidding prices, $p_{ij}^{(k)}$'s computed using (4), are used as block charge predictions. Constraint (6) models the budget limitation. To enforce that each block is requested from at most one neighbor, we include constraint (7) with auxiliary binary variable z_j^k, where $z_j^k = 1$ indicates that bidder j bids for block k from one (and only one) of its neighbors in $\mathcal{D}_j$ (*i.e.*, at most one $x_{ij}^{(k)}$ can be 1, $\forall i \in \mathcal{D}_j$), and $z_j^k = 0$ means that block k is not selected by peer j to bid for in this round.

The integer program in (5) is essentially a *0-1 knapsack* problem. Given a set of blocks with respective prices and utilities, a bidder is to determine the blocks to include in its bidding set, so that

ALGORITHM 1: Protocol at Seller i (in every interval T)

(a) Allocation:
receive bids $\mathbf{b_i}$ from neighbors in $\mathcal{D}_i$;
order $\mathbf{b_i}$ in non-increasing order of bidding prices into list $\mathbf{l_s}$;
set $o_i = O_i$ if $m = |\mathbf{b_i}| > O_i$; otherwise set $o_i = m - 1$;
while $o_i > 0$ **do**
    select next bid $b_{ij}^{(k)} = (I_{ij}^{(k)}, p_{ij}^{(k)})$ in list $\mathbf{l_s}$;
    let charge $c_{ij}^{(k)}$ be $p_{ij'}^{(k')}$, price in the subsequent bid in $\mathbf{l_s}$;
    send charge $c_{ij}^{(k)}$ to bidder j;
    start transfer of block $I_{ij}^{(k)}$ to bidder j;
    $o_i \leftarrow o_i - 1$;
end

(b) Upon receiving payment from bidder j for block $I_{ij}^{(k)}$:
update budget, $e_i \leftarrow e_i + c_{ij}^{(k)}$;

its budget is respected and the total utility is maximized. A pseudo-polynomial time solution can be designed using dynamic programming [Papadimitriou and Steiglitz 1998]. To derive a more intuitive algorithm at the bidder, we consider a greedy algorithm: A bidder j sorts all blocks it can potentially bid for in $\cup_{i\in\mathcal{D}_j}\mathcal{K}_{ij}$ in non-increasing order by marginal utility $v_{ij}^{(k)}/p_{ij}^{(k)}$. For a block available at multiple neighbors, only one instance with the largest marginal utility is included in the list. j then selects blocks from the list in order, until the remaining budget is not sufficient to bid further.

Let $\hat{e}_j \leq e_j$ represent the amount of budget actually used in this round. The greedy algorithm can derive the optimal solution to the optimization problem in (5), with constraint (6) replaced by:

$$\sum_{i\in\mathcal{D}_j}\sum_{k\in\mathcal{K}_{ij}} p_{ij}^{(k)} x_{ij}^{(k)} \leq \hat{e}_j. \qquad (9)$$

The bidding strategy is summarized in Algorithm 2, which is carried out periodically together with Algorithm 1.

4. ANALYSIS OF THE AUCTION ALGORITHMS

We now analyze the auction-based algorithms according to the two design goals. In Sec. 4.1, we examine the auction mechanism and bidding strategy within the framework of mechanism design and incentive engineering. In Sec. 4.2, we show effectiveness of block scheduling from the perspective of global welfare maximization.

4.1 Incentive Engineering at VoD Peers

We assume peers in the P2P VoD system are rational, *i.e.*, they aim to use the best strategies to maximize their own benefits, *e.g.*, revenues and utilities for sellers and bidders, respectively.

Seller Incentive Compatibility

A rational seller wishes to maximize its revenue in the auctions, in order to gain more budget for its own block purchase. We show the discriminative second price auction employed in Algorithm 1 is no worse (if not better) than other possible representative auction mechanisms, including sealed-bid first price auctions or the Vickrey auctions, in terms of a number of merits including seller revenue.

Discriminative second price auction is known to be an appropriate choice for auctioning virtual objects (*e.g.*, units of upload bandwidth, advertisement banners for online keyword search) in general [Edelman et al. 2007]. While first price auctions usually lead to longer-term contracts, such second price auctions allow sales to happen quickly, dynamically, and at small scales [Edelman et al. 2007].

ALGORITHM 2: Protocol at Bidder j

Initialization
set $\ddot{q}_{ij}$ to a MAX value, $\forall i \in \mathcal{D}_j$;

Every Interval T
(a) Bidding
for *each neighbor* $i \in \mathcal{D}_j$ **do**
  exchange buffer map with i and derive $\mathcal{K}_{ij}$;
  set $p_{ij}^{(k)} = \min(v_{ij}^{(k)}, \ddot{q}_{ij}), \forall k \in \mathcal{K}_{ij}$;
end
order blocks in $\cup_{i \in \mathcal{D}_j} \mathcal{K}_{ij}$ in non-increasing order of marginal utility $v_{ij}^{(k)}/p_{ij}^{(k)}$ into list $\mathbf{l_b}$ (excluding duplicates);
$p_{ij}^{(k)} \leftarrow$ price of the first block in list $\mathbf{l_b}$;
$\hat{e}_j \leftarrow p_{ij}^{(k)}$;
while $e_j \geq \hat{e}_j$ **do**
  send bid $(I_{ij}^{(k)}, p_{ij}^{(k)})$ to the corresponding seller i;
  $p_{ij}^{(k)} \leftarrow$ price of the next block in list $\mathbf{l_b}$;
  $\hat{e}_j \leftarrow \hat{e}_j + p_{ij}^{(k)}$;
end
(b) After Bidding
$p_i^{max} \leftarrow$ highest bidding price sent to neighbor $i, \forall i \in \mathcal{D}_j$;
set the lowest charge at i, $c_i^{min} = p_i^{max}, \forall i \in \mathcal{D}_j$;
for *each charge* $c_{ij}^{(k)}$ *received from* $i, \forall i \in \mathcal{D}_j$ **do**
  deduct e_j by $c_{ij}^{(k)}$ received;
  pay $c_{ij}^{(k)}$ to i;
  $c_i^{min} \leftarrow \min(c_i^{min}, c_{ij}^{(k)})$;
end
for *each neighbor* $i \in \mathcal{D}_j$ **do**
  if no bid was successful (no charge received from i), $\ddot{q}_{ij} = p_i^{max} + \delta$; otherwise, $\ddot{q}_{ij} = c_i^{min} - \delta$;
end

Therefore, despite the fact that no dominant strategy equilibrium or truthful equilibrium generally exists, they are the practical auction mechanism of choice for Internet advertising by Google and Yahoo!.

Below we show that our discriminative second price auction is a better choice for a seller than its natural alternative: uniform second price auction, which is equivalent to the Vickrey-Clarke-Groves (VCG) mechanism. A VCG mechanism is known to induce truthful bids, *i.e.*, each bidder bids its true valuation for the object it wants. In a VCG mechanism, the seller charges a winner the externality the winner imposes on others, *i.e.*, the decrease of the overall valuation of other bidders due to this winner's presence. When one object is on sale (single-object auctions), the charge to the winner is the second highest bid. In our context, o_i units of upload capacity are on sale at each seller i. If we apply a VCG mechanism, all the first o_i bids with highest bidding prices (the bidders' true valuations) win and pay the $(o_i + 1)$st highest bid, which constitutes a uniform second price auction.

Theorem 1 gives that Algorithm 1 always weakly dominates the VCG mechanism. Detailed proof of the theorem can be found in Appendix 8.1.

Theorem 1. *The discriminative second price auction in Algorithm 1 weakly dominates the VCG mechanism, by always generating equal or higher revenue for the seller than VCG does in Nash equilibrium states,* i.e.*, let* $R_{alg1} = \sum_{j \in \mathcal{D}_i} \sum_{k \in \mathcal{K}_{ij}} a_{ij}^{(k)} c_{ij}^{(k)}$ *be the revenue seller* i *obtains with Algorithm 1 in Nash* *equilibrium, and R_{vcg} be the revenue obtained with the VCG mechanism in Nash equilibrium, we have $R_{alg1} \geq R_{vcg}$.*

In Algorithm 1, we also consider the upload capacity contribution strategy at the sellers: when m is no larger than O_i, seller i refuses to sell the mth unit upload capacity. However, the total revenue is unaffected since the price charged for the mth unit capacity is zero even if it is sold. On the other hand, we may prevent an equal-to-zero market price in the auction, considering price probing at the bidders.

Should a seller be more aggressive and withhold more than one units of its upload capacity, for perhaps a higher market price and hence higher revenues? The answer is *no*. In a P2P VoD market, intense market competitions exist, *i.e.*, sellers may well offer the same blocks through their own auctions, and collusion among all sellers (who are also the buyers) with ubiquitous aggressive capacity reservation is unlikely to happen. A buyer can choose from multiple competing sellers to bid for a given block; if it discovers a high market price at one seller, it may well purchase the block (and later blocks) from other cheaper sellers. Therefore, a seller that reserves a high fraction of its upload bandwidth runs a high risk of failing to boost up the market price, but may loose sales to other competing sellers in the interconnected markets. In a practical dynamic system where neither sellers or bidders have complete information of real-time market status, the best strategy for a seller is to maximally contribute its bandwidth in each round, as in Algorithm 1, in order to glean more revenue over time.

In conclusion, Algorithm 1 is therefore *incentive compatible* with the VoD sellers, *i.e.*, sellers have the incentive to apply it as their best auction mechanism, as compared to other representative auction mechanisms, to achieve maximized revenue.

Bidder Incentive Compatibility

The bidding strategy in Algorithm 2 naturally leads to a *truthful start* followed by *market price probing* at each bidder. In the first round upon joining an auction, a bidder j bids its true valuation of a block. Even if the bid is temporarily higher than necessary, the overpayment will be corrected soon with progressive price adjustments. Below we further discuss Algorithm 2's bidding strategy for the subsequent rounds.

Theorem 2. *In the auction at seller i described in Algorithm 1, for each block $k \in \mathcal{K}_{ij}$, bidding a price equal to $\min(v_{ij}^{(k)}, \tilde{p}_i)$,* i.e.*, the minimum between the block valuation and the market price at i, is an equilibrium strategy for every bidder j.*

Please refer to Appendix 8.2 for detailed proof of the theorem.

In each subsequent round, though the ideal action for j is to bid at $\min(v_{ij}^{(k)}, \tilde{p}_i)$, $\tilde{p}_i$ is unknown at this stage. Hence j can only probe it using its market price estimate $\ddot{q}_{ij}$, which is captured in the bid $p_{ij}^{(k)} = min(v_{ij}^{(k)}, \ddot{q}_{ij})$ in (4) and the adjustment of $\ddot{q}_{ij}$ in Algorithm 2. A tradeoff of the probing is that when $\ddot{q}_{ij}$ is eventually adjusted to slightly below $\tilde{p}_i$, j may temporarily lose the auction before it re-adjusts its bidding price upward.

4.2 Social Welfare Maximization

We now analyze the effectiveness of the resulting block scheduling based on Algorithms 1 and 2. We first show that a Nash Equilibrium exists at the stable state of VoD streaming. We note that without peer joins/departures and VCR operations, the equilibrium defined as E^* in the proof of Theorem 1 constitutes a long-term stable operating state for the auction at a seller i, in which the market price $\tilde{p}_i$ is the true block valuation in the highest losing bid, each participating bidder j with $v_{ij}^{(k)} \geq \tilde{p}_i$ bids and pays $\tilde{p}_i$, and all other participating bidders bid their true valuations and lose. In the rest of the analysis, we also refer to such a Nash Equilibrium as the *stable state* of the VoD system. All bidding prices in the winning bids received by a seller i are equal in the stable state.

Given the budget constraints imposed at the VoD peers, the problem of optimizing the global social welfare can be formulated into the following integer program. Here $\mathcal{N}$ is the set of all peers in the overlay.

$$\text{Maximize} \quad \sum_{j\in\mathcal{N}}\sum_{i\in\mathcal{D}_j}\sum_{k\in\mathcal{K}_{ij}} v_{ij}^{(k)}(a_{ij}^{(k)}) \tag{10}$$

$$\text{Subject to:} \quad \mathcal{P}_{global}\begin{cases}\sum_{i\in\mathcal{D}_j}\sum_{k\in\mathcal{K}_{ij}} p_{ij}^{(k)}a_{ij}^{(k)} \le \hat{e}_j & \forall j\in\mathcal{N}\\ \sum_{j\in\mathcal{D}_i}\sum_{k\in\mathcal{K}_{ij}} a_{ij}^{(k)} \le o_i & \forall i\in\mathcal{N}\\ \sum_{i\in\mathcal{D}_j} a_{ij}^{(k)} = z_j^k & \forall j\in\mathcal{N}, \forall k\in\mathcal{K}_{ij}\end{cases}$$
$$a_{ij}^{(k)}, z_j^k \in \{0,1\}, \forall j\in\mathcal{N}, \forall i\in\mathcal{D}_j, \forall k\in\mathcal{K}_{ij}$$

Theorem 3 shows that the distributed local optimization carried out through block auctions in Algorithms 1 and 2 leads to an efficient system design, in that aggregated "satisfaction" $v_{ij}^{(k)}(a_{ij}^{(k)})$ among all peers are implicitly maximized.

Theorem 3. *Algorithms 1 and 2 solve (10),* i.e.*, achieve social welfare maximization, in a stable P2P VoD overlay.*

Please refer to Appendix 8.3 for the detailed proof.

In summary, the analysis in this section shows that our auction-based algorithm design achieves the following desirable objectives: (1) A seller has full incentive to apply our auction mechanism in Algorithm 1, in order to achieve maximized revenue for itself. (2) A bidder is also incentivized to apply our bidding strategies in Algorithm 2, which obtain the best blocks with the largest aggregate utility for itself. (3) The auction mechanism and bidding strategies together enable efficient block scheduling among peers, that leads to social welfare maximization in the entire P2P system at stable state. Social welfare reflects balanced block distribution or timely block download to meet playback deadlines at all peers, depending on the valuation functions used. On the other hand, in a dynamic system where VCR operations and peer joins/departures persist, Algorithms 1 and 2 still strive to pursue optimal block scheduling, chasing the objective of social welfare maximization, which is a moving target.

5. PERFORMANCE EVALUATION

For evaluating the auction-based P2P VoD algorithms, we have implemented an efficient multi-threaded P2P network simulator in Java, based partly on the source codes of *PlanetSim* [Pla tsim]. Our simulator supports all peer dynamics, including VCR operations, peer joins and departures, by events scheduled at their respective times. With careful optimization, it can simulate P2P systems with more than $10,000$ simultaneous peers in highly dynamic scenarios, distinguishing itself from representative existing P2P simulators [Naicken et al. 2007] which may support 3000 peers at most.

Our simulation settings are intended to model realistic VoD systems. An 80-minute video is streamed, the playback bitrate is 450 Kbps, and each block equals $1/3$ second of playback [Wu et al. 2008]. Peer upload capacities follow a heavy-tailed Pareto distribution in the major range of [256 Kbps, 10 Mbps] with shape parameter $k = 2$ or 3 (our default), corresponding to a mean upload bandwidth of 512 Kbps or 384 Kbps, respectively. Peers join the overlay following a Poisson arrival model with 1.8-second mean inter-arrival time. Peer lifetime follows an exponential distribution with a mean of 30 minutes [Hei et al. 2007]. A peer issues a random-seek VCR command periodically, with inter-command time following an exponential distribution with a 5-minute mean. Each peer maintains around 30 neighbors and may buffer blocks for up to 20-minute playback [Wu et al. 2008]. There is a streaming server with 10 Mbps upload capacity. Buffer maps are exchanged and bidding and allocation are carried out every 5 seconds. Though our simulator can support $10,000$ and more peers, 3000 peers are used in our

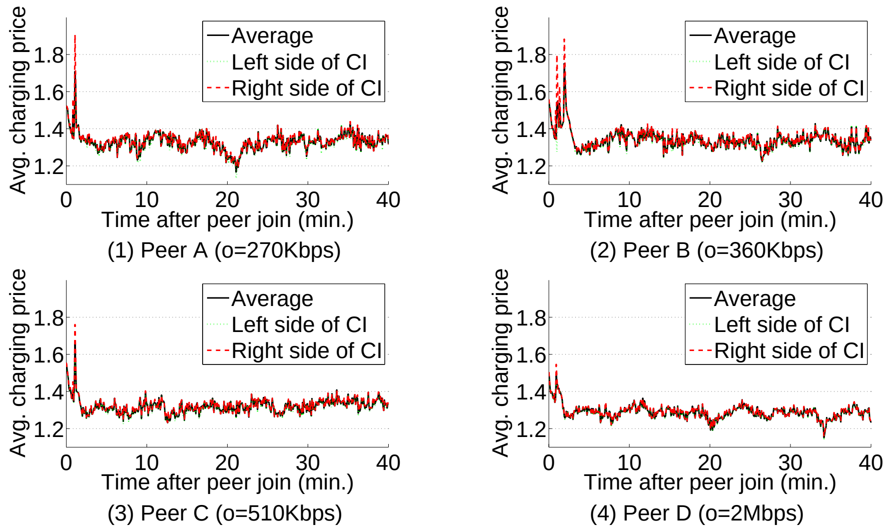

Fig. 2. Average charges at sellers with different upload capacities.

experiments for timely execution without loss of intriguing observations. The default price adjustment δ (used in the estimate of market prices) is set to 0.05.

We study different block valuation methods for $v_{ij}^{(k)}$: (i) *Deadline-based valuation* $\frac{\alpha_d}{\log(\beta_d+d)}$, emphasizing how urgent a block is. Here d is the time to playback deadline of the block, normalized by dividing 20 minutes (buffer capacity), and α_d and β_d are tunable constants. (ii) *Rareness-based valuation* $\frac{\alpha_r}{\log(\beta_r+r)}$, emphasizing how rare the block is among peer j's neighbor set $\mathcal{D}_j$. Here r is the number of neighbors possessing the block divided by $|\mathcal{D}_j|$, and α_r and β_r are tunable constants. (iii) A hybrid valuation (our default) based on linear combination of (i) and (ii). The default constant values are: $\alpha_d = 1$, $\beta_d = 1.2$, $\alpha_r = 1$ and $\beta_r = 1.2$. Initial peer budget is set to 2000, sufficient to buy blocks for 5-min playback.

5.1 Dynamics of Prices and Budgets

We first study the evolution of prices and budgets in the microscopic level. During an 80-min run of the VoD system, we pick four representative peers with upload capacity of 270 Kbps, 360 Kbps, 510 Kbps and 2 Mbps, respectively, who join the overlay at around 18 minutes. We temporarily disable VCR functions and allow peers to stay till the end, to investigate prices/budgets in relatively stable states.

5.1.1 *Evolution of charges.* Fig. 2 shows the average and 95% confidence intervals of all charges to their bidders at the four sellers in each round. All average charges and confidence intervals at the sellers are calculated at the same time in each round. We observe that the confidence intervals are very small (largely overlapping with the averages) and charges at each seller stabilize quickly to the vicinity of the respective market prices, due to the dynamic price probing strategy taken by each bidder.

5.1.2 *Evolution of estimate market prices.* We further explore the impact of price adjustment granularity, δ, on the convergence of estimated market prices at the bidders. In Fig. 3, we plot the average market price estimate among all bidders at seller C in each round, when different δ values are used in the estimation. We observe that the smaller δ is, the slower the average market price estimate converges. Though it always stabilizes to the vicinity of the "true" market price when different δ values are used, how close it is to the market price is decided by δ: the smaller δ is, the closer the market price estimate can converge to the "true" market price. Therefore, the choice of δ results in a tradeoff between

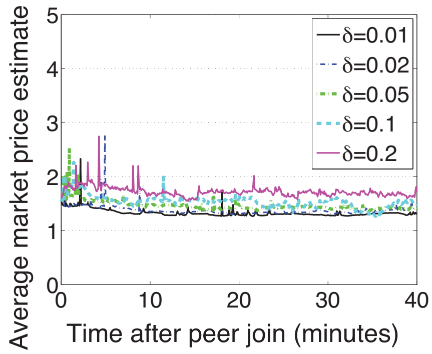

Fig. 3. Average market price estimate at peer C (o=510 Kbps) with different δ.

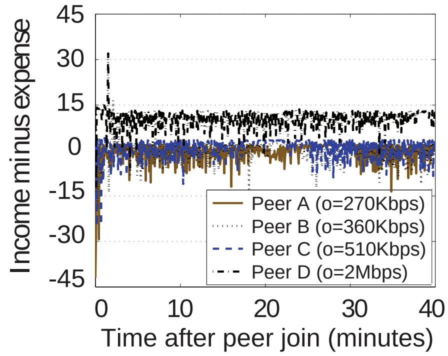

Fig. 4. Evolution of net profit at peers with different upload capacities: k=3, hybrid valuation.

convergence speed and accuracy. We choose a default $\delta = 0.05$ in our experiments. Nevertheless, an intelligent price adjustment algorithm may be devised to achieve fast convergence with high accuracy, *e.g.*, larger δ can be used at the first few bidding rounds and smaller δ is applied in the following. This investigation is orthogonal to the focus of the current paper, which we leave as future work.

5.1.3 *Evolution of budgets.* We next investigate how the budgets at the four peers evolve under various bandwidth abundance levels and block valuation methods. The budget is sampled at the end of each round, after prices paid and charges collected. Fig. 5 shows a rising budget at the high-capacity peers and a decreasing budget at the low-capacity ones. Such an interesting phenomenon can be further illustrated by Fig. 4 which plots peers' net profit in each round. Peers with larger capacities are making profit over time (income$\geq$expense), while low-capacity ones have deficits (income$\leq$expense). This is due to the higher upload bandwidth contribution at the high-capacity peers, than their download bandwidth consumption.

When the network bandwidth is insufficient (Fig. 5(b)), peer D with the largest upload capacity continuously aggregates the budget in the system, while the budgets at low-capacity peers (*e.g.*, A, B) decreases. Upload capacity is saturated at all peers in the case that network bandwidth is insufficient; thus high-capacity peers earn more than with sufficient bandwidth supply, and can prefetch more blocks. Their larger block diversity further attract more bidders. All these lead to an interesting "the rich get richer" phenomenon.

In Fig. 5(c) and (d), we employ deadline-only and rareness-only block valuations, respectively. Comparing Fig. 5(c) with Fig. 5(a), we notice a lower level of budget aggregation by the high-capacity peers, as their block diversity is lower when only blocks needed for playback are prefetched, leading to less revenue at those peers. Comparing Fig. 5(d) with Fig. 5(a), budget aggregation is more evident. Peer block diversity is highest with rareness-only evaluation. After joining the system, high-capacity peers with larger revenue can increase their block diversity faster, as they can afford to buy more rare blocks; this edge is lost (their budgets stabilize) when the block diversity at low-capacity peers catches up.

5.2 Incentives

To incentivize peer contribution, it is important to guarantee that peers who contribute more upload bandwidth can benefit more in streaming. We have observed that peers with larger upload capacity (thus large upload bandwidth contribution) aggregate more budget over time in Fig. 5. Fig. 6 and Fig. 7 show that the larger a peer's lifetime average upload bandwidth is, the higher download rate it enjoys on average (calculated as the average rate during its video download period, which could be shorter than its lifetime), helping more blocks meet playback deadlines. Each sample in the figures represents one of the 3000 peers that join the system, with arrival interval of 1.8 seconds.

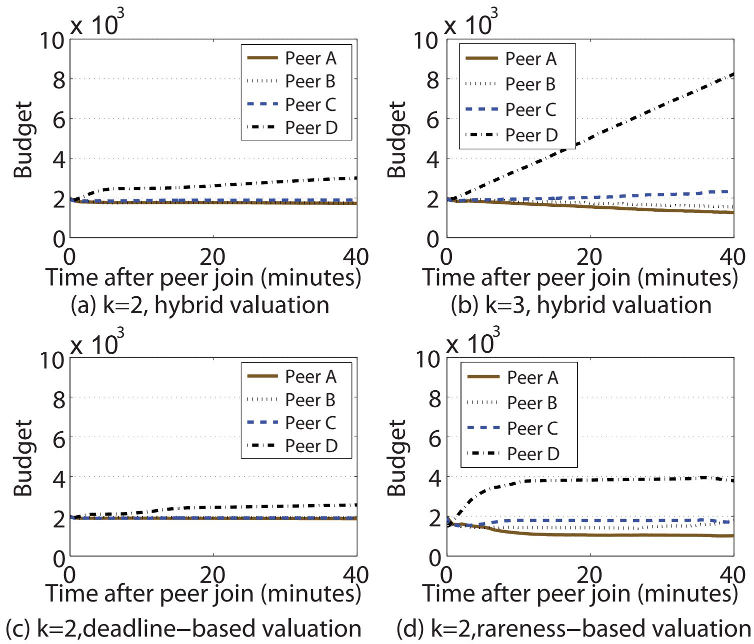

Fig. 5. Evolution of budget at peers with different upload capacities.

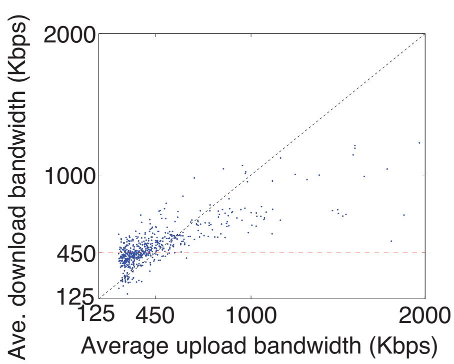

Fig. 6. Ave. download bandwidth vs. ave. upload bandwidth at each peer.

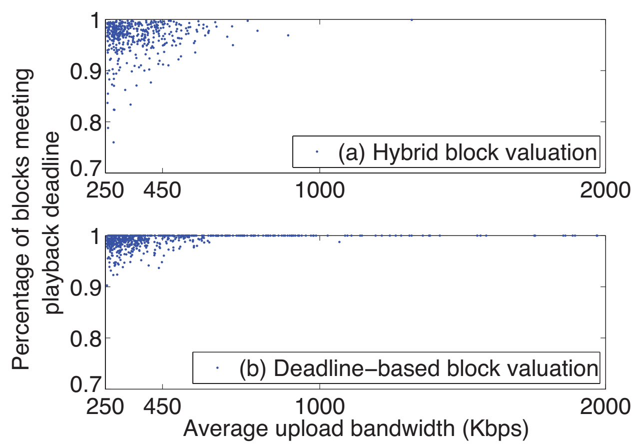

Fig. 7. Average percentage of blocks meeting playback deadline vs. average upload bandwidth at each peer.

Fig. 7 depicts the percentage of blocks meeting playback deadlines, during a peer's continuous viewing periods (excluding re-buffering periods upon a VCR operation). Comparing Fig. 7(b) to Fig. 7(a), we see that the percentage is higher with deadline-based valuation. In both cases, better playback smoothness is achieved at peers with larger upload bandwidth contributions.

5.3 Social Welfare

We now investigate the social welfare achieved by our auction-based algorithm, and compare it with two scheduling algorithms representing commonly used protocols: (1) *Algorithm S*, identical to our algorithm except for block scheduling: each downstream peer exchanges buffer maps with neighbors

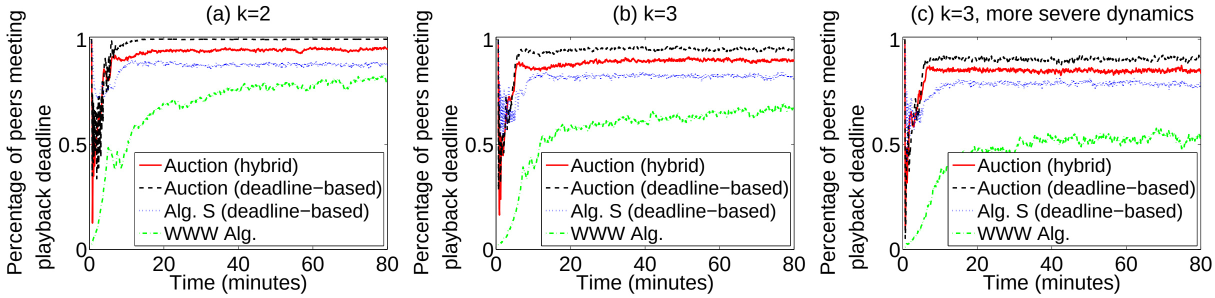

Fig. 8. Evolution of the percentage of peers meeting playback deadline.

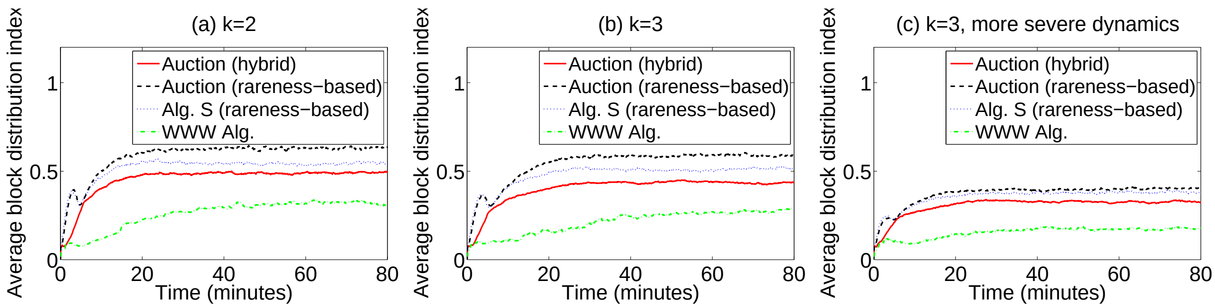

Fig. 9. Evolution of the average block distribution index.

and requests blocks according to the order of deadlines or rareness; each upstream peer decides the requests to serve randomly. (2) The scheduling algorithm in [Annapureddy et al. 2007], referred to as *the WWW algorithm*, by which each video is divided into s segments and each segment into h blocks. Blocks in a segment are encoded with network coding such that only segment level scheduling is needed. Its segment scheduling is based on a combined consideration of playback deadline and rareness: rarest segment of immediate interest is uploaded first. We simulate the WWW algorithm by setting s to 80 and h to 180, such that the total number of blocks in the video is the same as in our other experiments. In the following set of experiments, 3000 peers join the overlay simultaneously at the beginning, simulating the flash crowd of users at new video release. All other settings remain the same.

5.3.1 *Playback deadline.* We first compare overall playback smoothness at the peers. Fig. 8(a) and (b) show that the percentage of peers meeting deadlines is higher in our algorithm when deadline-based block valuation is applied, as compared to algorithm S (with deadline-based block valuation) and the WWW algorithm, regardless of the bandwidth abundance level. We also studied more severe system dynamics by setting the arrival interval of peers, peer life time, and the interval between VCR operations to half of their default values. Fig. 8(c) shows our algorithm still outperforms the others. This exhibits the effectiveness of our algorithm in utilizing prices to regulate the supply and demand at each peer, such that their upload bandwidth are most efficiently used to serve highest valued blocks.

5.3.2 *Block distribution.* Next we evaluate how well blocks are distributed in the system to serve peers' needs, with the three scheduling algorithms. We evaluate the block distribution around peer i by *the block distribution index*, which represents the average probability that a neighbor may serve peer i a block it needs:

$$\frac{\textit{total \# of copies of } i\textit{'s needed blocks in } \mathcal{D}_i}{\textit{\# of blocks } i \textit{ currently needs} \times |\mathcal{D}_i|}.$$

Fig. 9 shows that the average block distribution index across all peers is higher in our system based on rareness-based block valuation, as compared to algorithm S (with rareness-based block valuation)

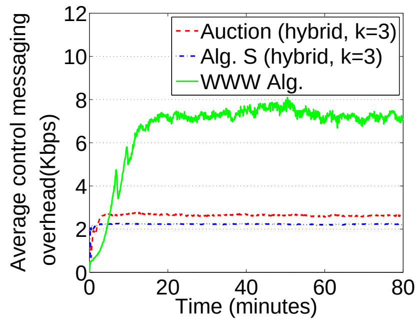

Fig. 10. Evolution of control messaging overhead.

and the WWW algorithm under all settings. This again validates the effectiveness of our auction-based scheduling, in making the best use of bandwidth in distributing blocks effectively.

5.3.3 *Control messaging overhead.* We now compare the control messaging overhead incurred in the three algorithms. For control messaging, buffer map exchange messages are common among all three algorithms. Our auction algorithm further involves delivery of bids and charges among sellers and buyers; Algorithm S includes block requests sent to the upstream nodes and responses of block allocation to the downstream nodes; for the WWW algorithm, the overhead includes segment requests sent to upstream nodes and coding coefficients sent together with coded blocks.

Fig. 10 shows that control messaging in the WWW algorithm consumes the most network bandwidth, as random coefficients need to be delivered together with coded blocks in this network coding scheme. In the WWW algorithm, there are 180 blocks in each segment and each coefficient is 2 bytes long; the size of random coefficients to be delivered with each coded block could be as large as 360 bytes. On the other hand, control messaging overhead is much lower with our auction algorithm and algorithm S. As compared to algorithm S, our algorithm involves more control messages to send bids/charges among peers. However, the overhead with both algorithms is only as low as $0.4 - 0.6\%$ of the streaming rate.

5.3.4 *Impact of block valuation methods.* We now explore how different block valuation methods, $\frac{\alpha_d}{log(\beta_d+d)} + \frac{\alpha_r}{log(\beta_r+r)}$ with different values of α_d, β_d, α_r and β_r, influence the overall playback smoothness and block distribution in the system. For this purpose, we repeat experiments with our auction scheme, when different valuation methods are adopted. Fig. 11 shows that when the radio of α_d over α_r is larger, *i.e.*, when the deadline-based term dominates the rareness-based term, the playback smoothness in the system is better.

On the other hand, Fig. 12 shows that in general, the larger the ratio of α_r over α_d is (the rareness-based term dominates the deadline-based term), the larger the block distribution index is (more efficiently blocks are distributed). Nevertheless, for the purely deadline-based evaluation (when $\alpha_r = 0$), the block distribution index is not as expected the smallest. Our close investigation reveals the reason: neighbors of a peer are mostly peers with similar playback progresses, and therefore valuable blocks at a peer (based on deadline-based valuations) may well be needed by its neighbors as well.

From the above observations, we can derive the following rules of setting parameters in the block valuation: if playback smoothness is more desirable in the system, more weight should be placed on the deadline-based valuation part (*i.e.*, larger α_d over α_r), and the optimal setting to achieve the best playback smoothness is to set $\alpha_r = 0$ (*i.e.*, only considering deadline-based valuation); if more efficient block distribution in the network is in demand, more weight should be put on the rareness-based evaluation part (*i.e.*, larger α_r over α_d).

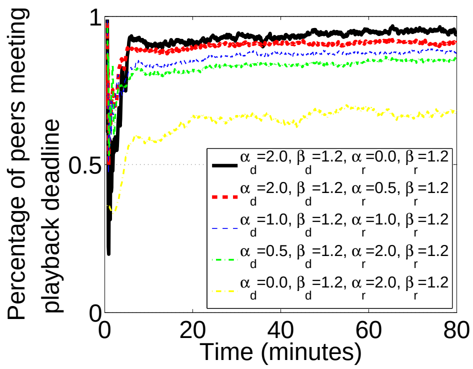

Fig. 11. Evolution of the percentage of peers meeting playback deadline with different block valuation methods.

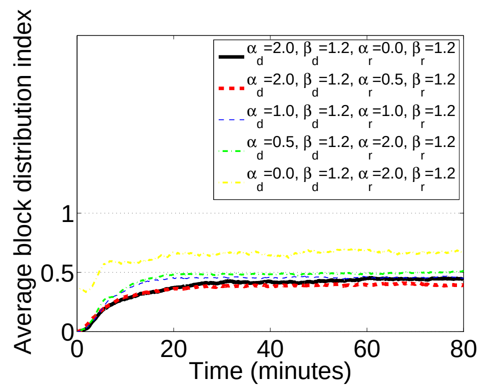

Fig. 12. Evolution of average block distribution index with different block valuation methods.

6. RELATED WORK

Incentive engineering in P2P streaming has attracted increasing research attention in recent years [Golle et al. 2001; Lai et al. 2003; Habib and Chuang 2006]. In the context of P2P live streaming, it was shown that simple *tit-for-tat* from P2P file sharing does not work well for streaming applications, which have additional delay and bandwidth requirements [Pai and Mohr 2006]. New incentive mechanisms were then proposed, connecting a peer's bandwidth contribution to its streaming quality, by differentiating the supplying peers it can download from [Habib and Chuang 2006], or the amount of bandwidth/substreams it can receive [Sung et al. 2006]. For P2P VoD streaming, Give-to-Get [Mol et al. 2008] advocates that peers favor uploading to proven good uploaders; the recent proposal of iPass [Liang et al. 2009] differentiates pre-fetching ability at the peers, allowing peers with high contribution to prefetch at high speed. In comparison, our budget-based incentive provides a quantitative approach for evaluating a peer's cumulative contribution over time (the budget a peer aggregates in the past can be used to purchase desired services in later streaming). We also show that maximal contribution is the best strategy at the peers with our algorithms, which has not been shown in existing proposals.

Budget-based mechanisms have been proposed for incentivizing peer contribution in P2P content distribution systems. PACE [Aperjis et al. 2008] proposes a market-based mechanism where currencies are exchanged for content based on a simple pricing mechanism, *i.e.*, a single charging price per seller. Such simple pricing may not be able to achieve optimal block scheduling if applied to P2P VoD streaming (with fast changing demand and supply relationship), while our algorithm has been proven to achieve so. Tan *et al.* [Tan and Jarvis 2008] design a payment-based incentive mechanism for P2P live streaming, based on first-price auctions. Targeting at substream auction in live streaming scenarios, their auctions cannot address the challenges of VoD streaming well. With respect to BitTorrent-like file sharing systems, a few auction-based models have been discussed for peer incentive engineering [Levin et al. 2006; Hausheer and Stiller 2005]. However, rigorous analysis of either the incentive compatibility or the optimality of chunk scheduling, is missing from these work.

Various heuristic algorithms on scheduling of block transmissions have been proposed. In [Huang et al. 2008], a peer pulls chunks from others with a mixed strategy of first sequential and then rarest-first downloading. The scheduling algorithm in [Annapureddy et al. 2007] encodes blocks in a segment with network coding and prioritizes the rarest segment of immediate interest. In iPass [Liang et al. 2009], either rarest-first, oldest-first, or random selection of pieces can be adopted. In comparison, our algorithm implements optimal scheduling by quantitatively optimizing social welfare.

In our previous work on P2P live streaming [Wu et al. 2008], we employ auctions to model the bandwidth allocation across multiple P2P live streaming overlays. Chu et al. [Chu et al. 2009] apply auction algorithms to design a distributed solution to a min-cost media flow multicast problem. The two work above are based on auctions of abstract 'flows', while this work conducts auctions on realistic media blocks. In addition, they have both applied auctions to achieve optimal resource allocation, but not peer incentivization. We are not aware of any other existing work that achieves both incentive compatibility and optimal scheduling in one integral auction design.

7. CONCLUSIONS

We presented P2P VoD algorithms that target full upload incentives at peers and optimal block scheduling across the streaming overlay. Employing strategically designed dynamic auctions, we achieve the goals in a simple, efficient, and fully decentralized fashion. Leveraging theories of mechanism design and convex optimization, we prove the effectiveness of the auction algorithms, in that each peer is fully incentivized to upload blocks strategically and the block scheduling in the entire system achieves social welfare maximization, in terms of meeting playback deadlines or balancing block distribution. Our large-scale empirical studies support our conclusions in highly dynamic real-world settings.

Beyond what were discussed in the paper, we believe that our auction-based design motivates peers to contribute in a variety of ways, facilitating optimal streaming in practical P2P VoD systems. In a real-world P2P VoD application where a user has a large collection of videos to choose from, a wealthy user (who owns a large budget) will have the advantage of enjoying better streaming quality (*e.g.*, better playback smoothness and shorter buffering delays), among the flash crowd to view a newly published video. Towards getting wealthier, a peer is not only fully motivated to contribute its upload capacity during its own playback, but also may choose to stay longer in the system, just to make the best use of its idling capacity to continue building its budget through block sales. Similarly, a peer may also temporarily join another video channel where bandwidth demand exceeds supply (*e.g.*, a less popular channel), to assist in block distribution, again for the purpose of gleaning block resale revenue. In this way, the puzzle of how to design an effective *cross-overlay* help algorithm for maximizing global streaming performance in a multi-overlay system, is naturally solved in a fully decentralized fashion.

We have observed some interesting phenomena in our empirical studies of our algorithms, such as the possible wealth condensation phenomenon (Sec. 5.1.3), which warrant further exploration: How would the budget distribution in the system evolve over a very long term of months or even years? Given the observation that high-capacity peers may aggregate the "wealth" in the system, should we apply a taxation mechanism to counter-act this, while still guaranteeing full incentives for peers to contribute? We are studying these interesting questions in ongoing work.

8. APPENDIX

8.1 Proof of Theorem 1

Proof: We prove by showing that in the worst Nash Equilibrium for the seller under Algorithm 1, the revenue equals that of VCG. The following is a Nash Equilibrium of the auction under Algorithm 1 at seller i: all the o_i bidders with the highest block valuations ($v_{ij}^{(k)}$'s) bid and is charged a price equal to the (o_i+1)st highest block valuation among all bids (which is the market price), and win the auction; all other bidders bid their true block valuations and lose. No bidder can improve its utility by unilaterally changing to a different bidding price: peer j with a winning bid has no incentive to bid higher (still wins, pays the same price, utility unchanged) or lower (lose the auction, utility decreases to zero); peer j with a losing bid has no incentive to bid higher (either wins at a charge higher than true valuation that results in negative utility, or still loses with unchanged utility) or lower (still loses, no change in utility). We denote this equilibrium as E^*.

Next, one can translate the auction at a seller i into an *assignment* problem, in which a bid is to be assigned to each unit of i's upload capacity. Any equilibrium of the discriminative second price auction corresponds to a *stable* assignment [Edelman et al. 2007], in which there does not exist a bid (from j) and a unit capacity at seller i, such that j is willing to purchase that unit capacity at a price slightly higher than what another bidder currently pays. Consequently, E^* corresponds to a stable assignment.

In any stable assignment, the charge paid for any winning bid of j is at least the $(o_i + 1)$st highest block valuation among all bids, otherwise the bidder of the (o_i+1)st bid would be willing to purchase the unit capacity won by j, by paying a price slightly below its block valuation. Since the charge in E^* for any winning bid is equal to the (o_i+1)st highest block valuation, we derive that E^* is the worst-case equilibrium for the seller. Furthermore, the total seller's revenue in E^* equals that of VCG. We conclude that the discriminative second price auction in Algorithm 1 always generates equal or higher revenue for the seller than VCG does. □

8.2 Proof of Theorem 2

Proof: In each round, j can win a desired block k at i with a bidding price no lower than the market price $\tilde{p}_i$. To maximize its utility $(v_{ij}^{(k)} - c_{ij}^{(k)})$ under its budget constraint, j is incentivized to bid a price no higher than $\tilde{p}_i$. In addition, a rational bidder won't bid for a block with a price higher than its valuation $v_{ij}^{(k)}$, to avoid negative utility. There are two cases: (a) if $v_{ij}^{(k)} \geq \tilde{p}_i$, bidding $\tilde{p}_i$ wins the block with the minimum budget expenditure and the largest utility, as compared to bidding $p_{ij}^{(k)} > \min(v_{ij}^{(k)}, \tilde{p}_i) = \tilde{p}_i$; (b) if $v_{ij}^{(k)} < \tilde{p}_i$, bidding $v_{ij}^{(k)}$ leads to zero utility (the bid will be unsuccessful), which is better than bidding a price higher than the true valuation $(p_{ij}^{(k)} > v_{ij}^{(k)})$ and winning the block with a negative utility $(v_{ij}^{(k)} - c_{ij}^{(k)} \leq v_{ij}^{(k)} - \tilde{p}_i < 0)$.[1] In both cases, the bidding price is equal to $\min(v_{ij}^{(k)}, \tilde{p}_i)$, which leads to the largest utility for the block at bidder j. Therefore, bidding $\min(v_{ij}^{(k)}, \tilde{p}_i)$ constitutes a utility-maximizing equilibrium strategy for every bidder j. □

8.3 Proof of Theorem 3

Proof: We prove the theorem by showing that optimal solutions obtained at individual bidders through solving (5) can be combined to construct an optimal solution to the global social welfare maximization in (10) at the stable state. We first show that (i) the relaxation of (5) always has an integral optimal solution. We next show that (ii) the KKT conditions of the relaxation of (5), aggregated across all VoD peers, are equivalent to the KKT conditions of the relaxation of (10), *i.e.*, optimal solutions to (the relaxation of) (5) can be combined to construct an optimal solution to (the relaxation of) (10).

(i) Given a fixed set of $z_j^k, \forall k \in \mathcal{K}_{ij}$, the relaxed bidding preference program at each peer j (to (5)) is

$$\text{Maximize} \quad \sum_{i\in\mathcal{D}_j} \sum_{k\in\mathcal{K}_{ij}} (v_{ij}^{(k)}(a_{ij}^{(k)}) - p_{ij}^{(k)} a_{ij}^{(k)}) \qquad (11)$$

$$\text{Subject to:} \quad \begin{cases} \sum_{i\in\mathcal{D}_j} \sum_{k\in\mathcal{K}_{ij}} p_{ij}^{(k)} a_{ij}^{(k)} \leq \hat{e}_j & (12) \\ \sum_{i\in\mathcal{D}_j} a_{ij}^{(k)} = z_j^k \quad \forall k \in \mathcal{K}_{ij} & (13) \\ a_{ij}^{(k)} \geq 0 \quad \forall i \in \mathcal{D}_j, \forall k \in \mathcal{K}_{ij} & (14) \\ a_{ij}^{(k)} \leq 1 \quad \forall i \in \mathcal{D}_j, \forall k \in \mathcal{K}_{ij} & (15) \end{cases}$$

A bidder's original optimization in (5) involves vector x instead of a. However, for each optimal solution x*, the corresponding vector a* also satisfies (12)-(15); furthermore, replacing x* with a* does not change the objective function value in (11). The reason for the latter is that for an entry where x and

[1]We note that a peer would not go for negative utility with a higher bidding price, and would rather spend the budget on other blocks it needs. When the playback deadline or rareness of a block changes, a peer's valuation of the block $v_{ij}^{(k)}$ changes, and it may bid for it again in a later round using the same strategy $p_{ij}^{(k)} = \min(v_{ij}^{(k)}, \tilde{p}_i)$.

a differ, the corresponding bid is unsuccessful; at Nash equilibrium, a bid is unsuccessful if the bidder bids $p_{ij}^{(k)} = v_{ij}^{(k)}$ lower than market price, and therefore $v_{ij}^{(k)}(x_{ij}^{(k)}) - p_{ij}^{(k)} x_{ij}^{(k)} = v_{ij}^{(k)}(a_{ij}^{(k)}) - p_{ij}^{(k)} a_{ij}^{(k)} = 0$ for the corresponding bid.

At equilibrium, the following constraint is also satisfied:

$$\sum_{j \in \mathcal{D}_i} \sum_{k \in \mathcal{K}_{ij}} a_{ij}^{(k)} \le o_i, \forall i \in \mathcal{N}. \tag{16}$$

The constraint group (13)-(16) is *totally unimodular* [Papadimitriou and Steiglitz 1998], defining a polyhedron with integral vertices (integral binary vectors for a). Further, if based on the relaxed optimization, a bidder j receives fractions of a block k from different sellers (*e.g.*, both $a_{ij}^{(k)}$ and $a_{hj}^{(k)}$ are non-zero), bidding prices for block k to the sellers must be the same ($p_{ij}^{(k)} = p_{hj}^{(k)}$). Consequently, if a fractional solution a satisfies (12), there must exist an integral solution satisfying it too. We then conclude that (11), the relaxation of (5) for any given binary vector $\mathbf{z}$, always has integral optimal solutions.

(ii) Introducing Lagrangian multiplier vectors [Boyd 2004] λ, β, μ and τ for constraints (12)-(15) respectively, we obtain the KKT conditions for (11) at every VoD peer $j \in \mathcal{N}$, which can be aggregated into:

$$KKT_{local} \begin{cases} \sum_{i \in \mathcal{D}_j} \sum_{k \in \mathcal{K}_{ij}} p_{ij}^{(k)} a_{ij}^{(k)} \le \hat{e}_j, \forall j \in \mathcal{N} \\ \sum_{i \in \mathcal{D}_j} a_{ij}^{(k)} = z_j^k, \forall j \in \mathcal{N}, \forall k \in \mathcal{K}_{ij} \\ 0 \le a_{ij}^{(k)} \le 1, \forall j \in \mathcal{N}, \forall i \in \mathcal{D}_j, \forall k \in \mathcal{K}_{ij} \\ \lambda, \mu, \tau \ge 0 \\ \lambda_j (\sum_{i \in \mathcal{D}_j} \sum_{k \in \mathcal{K}_{ij}} p_{ij}^{(k)} a_{ij}^{(k)} - \hat{e}_j) = 0, \forall j \in \mathcal{N} \\ \mu_{ij}^k a_{ij}^{(k)} = 0, \forall j \in \mathcal{N}, \forall i \in \mathcal{D}_j, \forall k \in \mathcal{K}_{ij} \\ \tau_{ij}^k (a_{ij}^{(k)} - 1) = 0, \forall j \in \mathcal{N}, \forall i \in \mathcal{D}_j, \forall k \in \mathcal{K}_{ij} \\ \sum_{j \in \mathcal{D}_i} \sum_{k \in \mathcal{K}_{ij}} a_{ij}^{(k)} \le o_i, \forall i \in \mathcal{N} \end{cases}$$

$$-v_{ij}^{(k)}(a_{ij}^{(k)}) + p_{ij}^{(k)} + \lambda_j p_{ij}^{(k)} - \mu_{ij}^k + \tau_{ij}^k + \beta_j^k = 0, \quad \forall j \in \mathcal{N}, \forall i \in \mathcal{D}_j, \forall k \in \mathcal{K}_{ij}. \tag{17}$$

Next, the relaxed global maximization problem to (10) is (given fixed z_j^k's):

$$\text{Maximize} \quad \sum_{j \in \mathcal{N}} \sum_{i \in \mathcal{D}_j} \sum_{k \in \mathcal{K}_{ij}} v_{ij}^{(k)}(a_{ij}^{(k)}) \tag{18}$$

$$\text{Subject to:} \quad \begin{cases} \mathcal{P}_{global} \\ 0 \le a_{ij}^{(k)} \le 1 \;\; \forall j \in \mathcal{N}, i \in \mathcal{D}_j, k \in \mathcal{K}_{ij} \end{cases}$$

Introducing Lagrangian multiplier vectors λ, ν, β, μ and τ for constraints of (18), we obtain the following KKT conditions for (18):

$$KKT_{global} \begin{cases} KKT_{local} \\ \nu \ge 0 \\ -v_{ij}^{(k)}(a_{ij}^{(k)}) + \nu_i + \lambda_j p_{ij}^{(k)} - \mu_{ij}^k + \tau_{ij}^k + \beta_j^k = 0, \\ \forall j \in \mathcal{N}, \forall i \in \mathcal{D}_j, \forall k \in \mathcal{K}_{ij}. \end{cases}$$

At each seller i, the winning bidding prices from different bidders are all equal to the market price at equilibrium, *i.e.*, $p_{ij}^{(k)} = \tilde{p}_i$ if $a_{ij}^{(k)} > 0, \forall j \in \mathcal{D}_i, k \in \mathcal{K}_{ij}$. Comparing KKT_{local} and (17) with KKT_{global}, we derive $\nu_i = \tilde{p}_i = p_{ij}^{(k)}, \forall i \in \mathcal{N}, j \in \mathcal{D}_i, k \in \mathcal{K}_{ij}$, and the two sets of KKT conditions are equivalent. □

REFERENCES

http://sourceforge.net/projects/planetsim/. *PlanetSim*.

http://www.pplive.com/. *PPLive*.

http://www.uusee.com/. *UUSee*.

ANNAPUREDDY, S., GUHA, S., GKANTSIDIS, C., GUNAWARDENA, D., AND RODRIGUEZ, P. 2007. Is High-Quality VoD Feasible using P2P Swarming? In *Proc. of the 16th International World Wide Web Conference*.

APERJIS, C., FREEDMAN, M. J., AND JOHARI, R. 2008. Peer-assisted Content Distribution with Prices. In *Proc. of the 4th ACM International Conference on emerging Networking EXperiments and Technologies (ACM CoNEXT)*.

ARMSTRONG, M. 1996. Multi-Product Nonlinear Pricing. *Econometrica 64,* 1, 51–75.

BOYD, S. 2004. *Convex Optimization*. Cambridge University Press.

CHU, X., ZHAO, K., LI, Z., AND MAHANTI, A. 2009. Auction Based On-Demand P2P Min-Cost Media Streaming with Network Coding. *IEEE Trans. on Parallel and Distributed Systems 20,* 12, 1816–1829.

EDELMAN, B., OSTROVSKY, M., AND SCHWARZ, M. 2007. Internet Advertising and the Generalized Second Price Auction. *American Economics Review 97,* 1, 242–259.

GOLLE, P., BROWN, K. L., MIRONOV, I., AND LILLIBRIDGE, M. 2001. Incentives for Sharing in Peer-to-Peer Networks. In *Proc. of the 2nd International Workshop on Electronic Commerce*.

HABIB, A. AND CHUANG, J. 2006. Service Differentiated Peer Selection: an Incentive Mechanism for Peer-to-Peer Media Streaming. *IEEE Transactions on Multimedia 8,* 3, 610–621.

HAUSHEER, D. AND STILLER, B. 2005. PeerMart: The Technology for a Distributed Auction-based Market for Peer-to-Peer Services. In *Proc. of IEEE ICC 2005)*.

HEI, X., LIANG, C., LIANG, J., LIU, Y., AND ROSS, K. W. 2007. A Measurement Study of a Large-Scale P2P IPTV System. *IEEE Trans. on Multimedia 9,* 8, 1672–1687.

HUANG, Y., FU, T. Z. J., CHIU, D.-M., LUI, J. C. S., AND HUANG, C. 2008. Challenges, Design and Analysis of a Large-Scale P2P-VoD System. In *Proc. of ACM SIGCOMM*.

KIRKEGAARD, R. AND OVERGAARD, P. B. 2008. Buy-out Prices in Auctions: Seller Competition and Multi-unit Demands. *The RAND Journal of Economics 3,* 3, 770–789.

LAI, K., FELDMAN, M., STOICA, I., AND CHUANG, J. 2003. Incentives for Cooperation in Peer-to-Peer Networks. In *Workshop on Economics of Peer-to-Peer Systems*.

LAZAR, A. AND SEMRET, N. 1999. Design and Analysis of the Progressive Second Price Auction for Network Bandwidth Sharing. *Telecommunications Systems: Special Issue on Network Economics*.

LEVIN, D., LACURTS, K., SPRING, N., AND BHATTACHARJEE, B. 2006. BitTorrent is an Auction: Analyzing and Improving BiTorrent's Incentives. In *Proc. of ACM SIGCOMM*.

LIANG, C., FU, Z., LIU, Y., AND WU, C. W. 2009. iPASS: Incentivized Peer-assisted System for Asynchronous Streaming. In *Proc. of IEEE INFOCOM 2009 Mini-Conference*.

LIU, Z., DHUNGEL, P., WU, D., ZHANG, C., AND ROSS, K. W. 2010. Understanding and Improving Incentives in Private P2P Communities. In *Proc. of the 30th International Conference on Distributed Computing Systems (ICDCS)*.

MOL, J., POUWELSE, J., MEULPOLDER, M., EPEMA, D., AND SIPS, H. 2008. Give-to-Get: Free-riding-resilient Video-on-Demand in P2P Systems. In *Proc. of MMCN 2008*.

NAICKEN, S., LIVINGSTON, B., BASU, A., RODHETBHAI, S., WAKEMAN, I., AND CHALMERS, D. 2007. The State of Peer-to-Peer Simulators and Simulations. *ACM SIGCOMM Computer Communication Review 37,* 2, 95 – 98.

NISAN, N., ROUGHGARDEN, T., TARDOS, E., AND VAZIRANI, V. V. 2007. *Algorithmic Game Theory*. Cambridge University Press.

PAI, V. AND MOHR, A. E. 2006. Improving Robustness of Peer-to-Peer Streaming with Incentives. In *Proc. of NetEcon'06*.

PAPADIMITRIOU, C. H. AND STEIGLITZ, K. 1998. *Combinatorial Optimization: Algorithms and Complexity*. Courier Dover.

SUNG, Y.-W., BISHOP, M., AND RAO, S. 2006. Enabling Contribution Awareness in an Overlay Broadcasting System. In *Proc. of ACM SIGCOMM*.

TAN, G. AND JARVIS, S. A. 2008. A Payment-Based Incentive and Service Differentiation Scheme for Peer-to-Peer Streaming Broadcast. *IEEE Transactions on Parallel and Distributed Systems 19,* 7, 940–953.

TURNER, D. AND ROSS, K. 2004. A Lightweight Currency Paradigm for the P2P Resource Market. In *Proc. of the 7th International Conference on Electronic Commerce Research*.

VISHNUMURTHY, V., CHANDRAKUMAR, S., AND SIRER, E. G. 2003. KARMA: A Secure Economic Framework for Peer-to-Peer Resource Sharing. In *Proc. of the 1st Workshop on the Economics of P2P Systems*.

WU, C., LI, B., AND LI, Z. 2008. Dynamic Bandwidth Auctions in Multi-overlay P2P Streaming with Network Coding. *IEEE Trans. on Parallel and Distributed Systems 19,* 6, 806–820.

WU, C., LI, B., AND ZHAO, S. 2008. Exploring Large-Scale P2P Live Streaming Topologies. *ACM Trans. on Multimedia Computing, Communications and Applications 4,* 19.