

Strategies of Conflict in Coexisting Streaming Overlays

Chuan Wu, Baochun Li

Department of Electrical and Computer Engineering

University of Toronto

{chuanwu, bli}@eecg.toronto.edu

Abstract—In multimedia applications such as IPTV, it is natural to accommodate multiple coexisting peer-to-peer streaming overlays, corresponding to channels of programming. With coexisting streaming overlays, one wonders how these overlays may efficiently share the available upload bandwidth on peers, in order to satisfy the required streaming rate in each overlay, as well as to minimize streaming costs. In this paper, we seek to design simple, effective and decentralized strategies to resolve conflicts among coexisting streaming overlays. Since such strategies of conflict are game theoretic in nature, we characterize them as a decentralized collection of dynamic auction games, in which downstream peers submit bids for bandwidth at the upstream peers. With extensive theoretical analysis and performance evaluation, we show that the outcome of these local games is an optimal topology for each overlay that minimizes streaming costs. These overlay topologies evolve and adapt to peer dynamics, fairly share peer upload bandwidth, and can be prioritized.

I. INTRODUCTION

Peer-to-peer streaming applications have recently become a reality in the Internet, in which large numbers of peers self-organize into streaming overlays. It is natural to consider multiple coexisting streaming overlays (sessions) in such applications, each of which corresponds to a channel of television programming or live events. Generated with a modern codec such as H.264, each overlay distributes a live media stream with a specific streaming rate, such as 800 Kbps for a Standard-Definition stream and 1700 Kbps for a 480p (848 × 480 pixels) High-Definition stream. To meet such exacting *demands* of bandwidth that have to be satisfied at all participating peers, a streaming overlay relies on available upload bandwidth *supplies* of both dedicated streaming servers and regular participating peers. Smooth streaming playback is not possible unless such *supplies* meet the demand for streaming bandwidth.

It only becomes more challenging when coexisting streaming overlays are considered, sharing the available upload bandwidth in the peer-to-peer network. Consider a typical scenario where multiple peers from different overlays are in conflict with one another, competing for limited upload bandwidth at the same streaming server or upstream peer in the network. Apparently, the allocation of such upload bandwidth needs to be meticulously mediated with appropriate strategies, such that the streaming rate requirement of each overlay is satisfied at all participating peers. It would be best if, at the same time, fairness or prioritization can be achieved across different overlays, and certain costs of streaming (*e.g.*, latencies) can be minimized. It goes without saying that if such tactical

strategies are not implemented, the conflict among streaming overlays may not be resolved satisfactorily.

In this paper, we seek to design simple, decentralized, but nonetheless effective tactical strategies to resolve inherent conflicts among coexisting streaming overlays. Much inspired by the seminal work of the Nobel Prize winner Thomas Schelling “*The Strategy of Conflict*,” we believe that it is best to characterize such conflicts in a game theoretic setting, and with *dynamic auction games*. Such games evolve over time, and involve repeated *auctions* in which *bids* are submitted by competing downstream peers from different overlays to the same upstream peer. In these dynamic auction games, an upstream peer allocates its upload bandwidth based on bids from downstream peers, and a downstream peer may optimize and place its bids to multiple upstream peers, and subsequently compete in multiple auctions. Each of these auctions is locally administered, and leads to cleanly decentralized strategies.

With extensive theoretical analysis and performance evaluation using simulations, we show that these decentralized game-theoretic strategies not only converge to a Nash equilibrium, but also lead to favorable outcomes: we are able to obtain an optimal topology for each coexisting streaming overlay, in the sense that streaming rates are satisfied, and streaming costs are minimized. These topologies of coexisting overlays evolve and adapt to peer dynamics, fairly share peer upload bandwidth, and can be prioritized. In contrast to existing game theoretic approaches that are largely theoretical in nature, we show that our proposed strategies can be practically implemented in realistic streaming overlays. Indeed, our focus in this paper is not on reasoning about the rationality and selfishness of peers, nor on incentive engineering to encourage contribution. We seek to devise practical strategies that may be realistically implemented, and use game theoretic tools only to facilitate the design of such conflict-resolving strategies.

The remainder of this paper is organized as follows. In Sec. II, we present our network model and motivate the design of distributed auction games. In Sec. III, we discuss in details the bidding and allocation strategies, prove their convergence, and then discuss their practical implementation issues. Sec. IV is dedicated to an in-depth study of the proposed strategies in realistic settings, with respect to interactions of multiple dynamic streaming overlays. We then discuss related work and conclude the paper in Sec. V and Sec. VI, respectively.

II. MODEL

In this paper, we consider multiple coexisting streaming overlays, each consisting of streaming servers and participating




Fig. 1. Two concurrent peer-to-peer streaming overlays: an example.

peers. Each server may serve more than one overlay, while each peer may also participate in multiple overlays. In each streaming overlay, participating servers and peers form a mesh topology, in which any peer is served by its *upstream* peers (servers can be deemed as special upstream peers), and may serve one or more *downstream* peers at the same time. Fig. 1 shows an example of two coexisting streaming overlays, each with two streaming servers and four participating peers.

Let \mathcal{S} denote the set of all coexisting streaming overlays. The topology of each overlay $s \in \mathcal{S}$ can be modeled as a directed graph $\mathcal{G}_s = (\mathcal{V}_s, \mathcal{N}_s, \mathcal{A}_s)$, where \mathcal{V}_s is the set of servers serving overlay s , \mathcal{N}_s represents the set of participating peers, and \mathcal{A}_s denotes the set of application-layer links in overlay s . Let R_s be the required streaming rate of the media stream distributed in overlay s . Let \mathcal{V} be the set of all streaming servers in the network, *i.e.*, $\mathcal{V} = \cup_{s \in \mathcal{S}} \mathcal{V}_s$, and \mathcal{N} be the set of all existing peers, *i.e.*, $\mathcal{N} = \cup_{s \in \mathcal{S}} \mathcal{N}_s$. U_i denotes the upload bandwidth at peer i , $\forall i \in \mathcal{V} \cup \mathcal{N}$. Realistically, we assume that the last-mile upload bandwidth on each peer constitutes the “supply” of bandwidth in the overlays. We are not concerned with insufficient peer download bandwidth, as it is not possible to achieve required streaming rates in case of such lack of bandwidth, with any solution.

Each upstream peer i , $\forall i \in \mathcal{V} \cup \mathcal{N}$, organizes a *dynamic auction game*, referred to as auction i , in order to mediate competition for its *upload bandwidth* — the “goods” for sale, with a total quantity U_i . The players in auction i are all the downstream peers of peer i in each overlay it participates in. Let j^s represent peer j in overlay s . The set of players in auction i can be expressed as $\{j^s, \forall j : (i, j) \in \mathcal{A}_s, \forall s \in \mathcal{S}\}$. A player j^s may submit its bids to multiple upstream peers in their respective auction games. In case a downstream peer j participates in multiple overlays, it is viewed as multiple players, each for one overlay.

The dynamic auction games at the peers are repeatedly carried out over time. In each *bidding round* of auction game i , each player submits its bid to peer i , declaring its requested share of upload bandwidth, as well as the unit price it is willing to pay. The upstream peer i then allocates shares of its upload capacity U_i to the players based on their bids. Let x_{ij}^s denote the upload bandwidth that player j^s requests from peer i , and p_{ij}^s denote the unit price it is willing to pay to peer i . The bid for player j^s in auction i can be represented as a 2-tuple $b_{ij}^s = (p_{ij}^s, x_{ij}^s)$.

Such a distribute game model can be illustrated with the example in Fig. 2. In the example, there are 7 auction games,




Fig. 2. Decentralized auction games in the example streaming overlays.

two of which are marked: auction 1 at v_1 with 5 players 3^1 (peer n_3 in overlay 1), 5^1 , 5^2 , 6^1 and 6^2 , auction 2 at v_2 , with 4 players 4^2 , 5^1 , 5^2 and 7^1 , respectively.

III. THE AUCTION GAME

We are now ready to propose the allocation strategy taken by an upstream peer, and the bidding strategy by downstream peers in distributed auction games. We show these game-theoretic strategies can be readily implemented in practical streaming overlays, and analyze their convergence to a Nash equilibrium.

A. Allocation strategy

In auction i , the seller, upstream peer i , aims to maximize its revenue by selling its upload bandwidth U_i at the best prices. Given bids $b_{ij}^s = (p_{ij}^s, x_{ij}^s)$'s from all the players j^s ($\forall j : (i, j) \in \mathcal{A}_s, \forall s \in \mathcal{S}$), upstream peer i 's allocation strategy can be represented by the following revenue maximization problem. Here, a_{ij}^s ($\forall j : (i, j) \in \mathcal{A}_s, \forall s \in \mathcal{S}$) is the bandwidth share to be allocated to each downstream peer j in each competing overlay s .

Allocation i:

$$\max \sum_{s \in \mathcal{S}} \sum_{j: (i, j) \in \mathcal{A}_s} p_{ij}^s a_{ij}^s \quad (1)$$

subject to

$$\begin{aligned} \sum_{s \in \mathcal{S}} \sum_{j: (i, j) \in \mathcal{A}_s} a_{ij}^s &\leq U_i, \\ 0 &\leq a_{ij}^s \leq x_{ij}^s, \quad \forall j : (i, j) \in \mathcal{A}_s, \forall s \in \mathcal{S}. \end{aligned}$$

Such an allocation strategy can be achieved in the following fashion:

Upstream peer i selects the highest bid price, e.g., p_{ij}^s from player j^s , and allocates bandwidth $a_{ij}^s = \min(U_i, x_{ij}^s)$ to it. Then if it still has remaining bandwidth, it selects the second highest bid price and assigns the requested bandwidth to the corresponding player. This process repeats until peer i has allocated all its upload capacity, or bandwidth requests from all the players have been satisfied. \square

The above allocation strategy can be formally stated in the following formula:

$$\begin{aligned} a_{ij}^s &= \min(x_{ij}^s, U_i - \sum_{p_{ik}^s \geq p_{ij}^s, k^s \neq j^s} a_{ik}^s), \\ &\forall j : (i, j) \in \mathcal{A}_s, \forall s \in \mathcal{S}. \end{aligned} \quad (2)$$

B. Bidding strategy

In each overlay $s \in \mathcal{S}$, a peer j may place its bids to multiple upstream peers. As a common objective, it wishes to acquire the required streaming rate for the overlay, and experience minimum costs. We consider two parts of costs when peer j streams from peer i in overlay s : streaming cost — denoted by streaming cost function $D_{ij}^s(x_{ij}^s)$ — represents the streaming latency actually experienced by j ; bidding cost — calculated by $p_{ij}^s x_{ij}^s$ — represents the bid peer j submits to peer i in overlay s . The bidding cost reflects the degree of *competition* and *demand* for bandwidth in the auction games at upstream peers. The overall cost at player j^s is the sum of the two parts from all its upstream peers, $\forall i : (i, j) \in \mathcal{A}_s$.

In this way, the preference for player j^s in deciding its bids in the auctions can be expressed by the following cost minimization problem. Practically, we assume cost functions D_{ij}^s are non-decreasing, twice differentiable and strictly convex.

Bidding j^s :

$$\min \sum_{i:(i,j) \in \mathcal{A}_s} (D_{ij}^s(x_{ij}^s) + p_{ij}^s x_{ij}^s) \quad (3)$$

subject to

$$\sum_{i:(i,j) \in \mathcal{A}_s} x_{ij}^s \geq R_s, \quad (4)$$

$$x_{ij}^s \geq 0, \quad \forall i : (i, j) \in \mathcal{A}_s. \quad (5)$$

The bidding strategy of player j^s consists of two main components: bandwidth requests and price adjustments.

1) *Bandwidth requests*: If the bid prices p_{ij}^s 's are given, the requested bandwidths at player j^s towards each of its upstream peers in overlay s , i.e., $x_{ij}^s, \forall i : (i, j) \in \mathcal{A}_s$, can be optimally decided by solving the problem **Bidding j^s** . This can be done efficiently with a water-filling approach, in which player j^s acquires the required streaming rate R_s by requesting from upstream peers that incur minimum marginal costs:

Let $f_j^s(x)$ denote the overall cost at player j^s , i.e., $f_j^s(x) = \sum_{i:(i,j) \in \mathcal{A}_s} (D_{ij}^s(x_{ij}^s) + p_{ij}^s x_{ij}^s)$. The marginal cost with respect to x_{ij}^s is $\frac{df_j^s(x)}{dx_{ij}^s} = D_{ij}^s(x_{ij}^s) + p_{ij}^s$. Beginning with $x_{ij}^s = 0$ ($\forall i : (i, j) \in \mathcal{A}_s$), the player identifies one x_{ij}^s that achieves the smallest marginal cost and increases its value. As $D_{ij}^s(x_{ij}^s)$ is strictly convex, $D_{ij}^s(x_{ij}^s)$ increases with the increase of x_{ij}^s . The player increases the x_{ij}^s until its marginal cost is no longer the smallest. Then it finds a new x_{ij}^s with current smallest marginal cost and increases its value. This process repeats until the sum of all x_{ij}^s 's ($\forall i : (i, j) \in \mathcal{A}_s$) reaches R_s . \square




Fig. 3. Bandwidth requesting strategy at player j^s : an illustration of the water-filling approach.

The water-filling approach can be illustrated in Fig. 3, in

which the height of each bin represents the marginal cost for player j^s to stream from each upstream peer i . To fill water at a total quantity of R_s into these bins, the bins with the lowest heights are flooded first, until all bins reach the same water level. Then the same water level keeps increasing until all the water has been filled in.

Theorem 1. Given bid prices $p_{ij}^s, \forall i : (i, j) \in \mathcal{A}_s$, the water-filling approach obtains a unique optimal requested bandwidth assignment at player j^s , i.e., $(x_{ij}^s, \forall i : (i, j) \in \mathcal{A}_s)$, which is the unique optimal solution to the problem **Bidding j^s** .

Due to space constraints, interested readers are referred to our technical report [1] for a complete proof of Theorem 1.

2) *Price adjustments*: The next critical question to address is how each player is to determine the bid price to each of its upstream peers. A price adjustment scheme is designed for this purpose, by which each player tactically adjusts its prices in participating auctions based on bids placed by its opponents in the previous bidding round.

When player j^s first joins the auction at an upstream peer i , it sets its bid price p_{ij}^s to 0. Together with its prices towards other upstream peers in overlay s , it calculates the current optimal requested bandwidth assignment with the *water-filling approach*, and then sends its bids to upstream peers. After upstream peer i allocates its upload capacity with the *allocation strategy*, it sends allocated bandwidth values to corresponding players. Upon receiving an allocated bandwidth, player j^s increases the corresponding bid price if its “demand” is higher than the “supply” from the upstream peer, and otherwise decreases the price. Meanwhile, it recomputes its requested bandwidth assignment for all its upstream peers with the water-filling approach. Such price adjustment is carried out in an iterative fashion, until the player’s bandwidth requests may all be granted if the new prices are bid.

Using the water-filling approach as a building block, the price adjustment scheme is summarized in the *bidding strategy* to be carried out by player j^s in each round of its participating auctions, as presented in Table I.

TABLE I
BIDDING STRATEGY AT PLAYER j^s

1. Receive allocated bandwidths a_{ij}^s from all upstream peers $i, \forall i : (i, j) \in \mathcal{A}_s$.
2. Adjust prices and bandwidth requests by:
 - Repeat
 - (a) For each upstream peer i
 - If $x_{ij}^s > a_{ij}^s$, increase price p_{ij}^s by a small amount δ ;
 - If $x_{ij}^s \leq a_{ij}^s$ and $p_{ij}^s > 0$, decrease price p_{ij}^s by δ .
 - (b) Adjust requested bandwidth assignment $(x_{ij}^s, \forall i : (i, j) \in \mathcal{A}_s)$ with the water-filling approach.
 - (c) For each upstream peer i
 - Calculate new allocation a_{ij}^s that can be acquired from i if the current price p_{ij}^s is bid, based on Eqn. (2), with queried bids of some other players in the previous round of auction i .
- Until: all requested bandwidths x_{ij}^s 's, are to be achieved with current prices p_{ij}^s 's, i.e., $x_{ij}^s \leq a_{ij}^s, \forall i : (i, j) \in \mathcal{A}_s$, and prices p_{ij}^s 's are the lowest possible to achieve it.
3. Submit new bids $b_{ij}^s = (p_{ij}^s, x_{ij}^s), \forall i : (i, j) \in \mathcal{A}_s$, to respective upstream peers.
4. Go to 1.

We next briefly remark on the computation and messaging complexity of implementing the bidding strategy.

First, although the requested bandwidths x_{ij}^s 's are to be recomputed each time the prices change, the water-filling approach does not need to be executed from the very beginning. To illustrate this based on Fig. 3, when the price p_{ij}^s to upstream peer i is raised, the water level of its corresponding bin (*i.e.*, marginal cost for x_{ij}^s) rises; when the price is decreased, the water level is lowered. Therefore, x_{ij}^s corresponding to a raised bin is reduced while that of a lowered bin is increased, in order to achieve a same water level again. Such local adjustments of x_{ij}^s 's are small in magnitude, and involve low computation overhead.

Second, to calculate the new achievable allocation a_{ij}^s , player j^s needs to know bids placed by some of its opponents in the previous bidding round in auction i . Instead of asking upstream peer i to send all received bids, player j^s can query such information gradually only when necessary. If p_{ij}^s is to be increased, it asks for the bid of opponent $m^{s'}$ whose price $p_{im}^{s'}$ is immediately higher than p_{ij}^s in auction i . While p_{ij}^s is still below $p_{im}^{s'}$, player j^s 's achievable bandwidth is unchanged; only when p_{ij}^s exceeds $p_{im}^{s'}$, its achievable bandwidth is increased by $a_{im}^{s'}$, and player j^s queries upstream peer i again for the bid containing a price immediately higher than the current value of p_{ij}^s . Similar bid inquiries can be implemented for the case that p_{ij}^s is to be reduced. In this way, the price adjustments can be achieved practically with little messaging overhead.

The intuition behind the bidding strategy is that, each player places different bid prices to different upstream peers, considering both the streaming cost and the overall demand at each upstream peer. If the streaming cost is low from an upstream peer, the player is willing to pay a higher price and strives to acquire more upload bandwidth from this peer. On the other hand, if the bandwidth competition at an upstream peer is intense such that the bidding cost becomes excessive, the player will forgo its price increases and request more bandwidths from other peers. At all times, the marginal cost of streaming from each upstream peer is kept the same, as achieved by the water-filling process.

C. Convergence analysis

The distributed auction games in the coexisting streaming overlays are carried out in a repeated fashion, as these are *dynamic games*. They are correlated with each other as each player optimally places its bids in multiple auctions. A critical question to investigate is: *Does there exist a stable "operating point" of the decentralized games, that achieves efficient partition of network upload bandwidths?* In what follows, we seek to investigate the convergence of the dynamic resource allocation from the global point of view.

We consider upload bandwidth competition in the entire network as one *extended* dynamic non-cooperative strategic game (referred to as G_{ext}), containing all the distributed correlated auctions. The set of players in the extended game can be represented as

$$\mathcal{I} = \{j^s, \forall j \in \mathcal{N}_s, \forall s \in \mathcal{S}\}. \quad (6)$$

The action profile taken by player j^s is a vector of bids, in which each component is the bid to place to one upstream peer. Formally, the set of action profiles for player j^s is defined as

$$\Gamma_j^s = \{B_j^s | B_j^s = (b_{ij}^s, \forall i : (i, j) \in \mathcal{A}_s), \\ b_{ij}^s = (p_{ij}^s, x_{ij}^s) \in [0, +\infty) \times [0, R_s], \sum_{i:(i,j) \in \mathcal{A}_s} x_{ij}^s \geq R_s\}. \quad (7)$$

Then, let B denote the bid profile in the entire network, *i.e.*, $B = (B_j^s, \forall j \in \mathcal{N}_s, \forall s \in \mathcal{S}) \in \times_{j,s} \Gamma_j^s$. The preference relation \succsim_j^s for player j^s can be defined by the following overall cost function, which is the objective function in the problem **Bidding j^s** in (3)

$$\text{Cost}_j^s(B) = \sum_{i:(i,j) \in \mathcal{A}_s} (D_{ij}^s(x_{ij}^s) + p_{ij}^s x_{ij}^s). \quad (8)$$

Therefore, we say two bid profiles $B \succsim_j^s B'$ if $\text{Cost}_j^s(B) \leq \text{Cost}_j^s(B')$.

Definition 1. A bid profile B in the network, $B = (B_j^s, \forall j \in \mathcal{N}_s, \forall s \in \mathcal{S}) \in \times_{j,s} \Gamma_j^s$, is feasible if its bandwidth requests further satisfy upload capacity constraints at all the upstream peers, *i.e.*,

$$\sum_{s \in \mathcal{S}} \sum_{j:(i,j) \in \mathcal{A}_s} x_{ij}^s \leq U_i, \quad \forall i \in \mathcal{V} \cup \mathcal{N}.$$

When a bid profile is feasible, from the allocation strategy discussed in Sec. III-A, we can see the upload bandwidth allocations will be equal to the requested bandwidths.

Using \widetilde{B}_j^s to represent action profiles of all players other than player j^s in \mathcal{I} , *i.e.*, $\widetilde{B}_j^s = (B_m^k, \forall m^k \in \mathcal{I} \setminus \{j^s\})$, we have the following definition of Nash equilibrium.

Definition 2. A feasible bid profile $B^* = (B_j^{s*}, \forall j \in \mathcal{N}_s, \forall s \in \mathcal{S})$ is a Nash equilibrium of the extended game $G_{\text{ext}}(\mathcal{I}, (\Gamma_j^s), (\succsim_j^s))$ if for every player $j^s \in \mathcal{I}$, we have $\text{Cost}_j^s(B_j^{s*}, \widetilde{B}_j^{s*}) \leq \text{Cost}_j^s(B_j'^s, \widetilde{B}_j^{s*})$ for any other feasible bid profile $B' = (B_j'^s, \widetilde{B}_j^{s*})$.

We next show the convergence of the extended game to such an equilibrium. We focus on feasible streaming scenarios as stated in the following assumption:

Assumption 1. The total upload bandwidth in the peer-to-peer network is sufficient to support all the peers in all overlays to stream at required rates, *i.e.*, there exists a feasible bid profile in the peer-to-peer network.

Theorem 2. In the extended game $G_{\text{ext}}(\mathcal{I}, (\Gamma_j^s), (\succsim_j^s))$ in which distributed auctions are dynamically carried out with the allocation strategy in (2) and the bidding strategy in Table I, there exists a Nash equilibrium under Assumption 1.

For a detailed proof of Theorem 2, interested readers are referred to our technical report [1] due to space constraints.

The next theorem shows that at equilibrium, the upload bandwidth allocation in the network achieves the minimization of the global streaming cost.

Theorem 3. At Nash equilibrium of the extended game $G_{\text{ext}}(\mathcal{I}, (\Gamma_j^s), (\succsim_j^s))$, upload bandwidth allocation in the entire network achieves streaming cost minimization, as achieved by the following global streaming cost minimization problem:

$$\min \sum_{s \in \mathcal{S}} \sum_{j \in \mathcal{N}_s} \sum_{i: (i,j) \in \mathcal{A}_s} D_{ij}^s(y_{ij}^s) \quad (9)$$

subject to

$$\sum_{s \in \mathcal{S}} \sum_{j: (i,j) \in \mathcal{A}_s} y_{ij}^s \leq U_i, \quad \forall i \in \mathcal{V} \cup \mathcal{N}, \quad (10)$$

$$\sum_{i: (i,j) \in \mathcal{A}_s} y_{ij}^s \geq R_s, \quad \forall j \in \mathcal{N}_s, \forall s \in \mathcal{S}, \quad (11)$$

$$y_{ij}^s \geq 0, \quad \forall (i,j) \in \mathcal{A}_s, \forall s \in \mathcal{S} \quad (12)$$

Theorem 3 can be proven by showing that the set of KKT conditions for the global streaming cost minimization problem is the same as that satisfied by the equilibrium bid profile $B^* = ((p_{ij}^{s*}, x_{ij}^{s*}), \forall (i,j) \in \mathcal{A}_s, \forall s \in \mathcal{S})$, and the equilibrium bid prices at each upstream peer i (i.e., $p_{ij}^{s*}, \forall j : (i,j) \in \mathcal{A}_s, \forall s \in \mathcal{S}$) have the same value as the Lagrangian multiplier associated with the upload capacity constraint (10) at peer i . Again, interested readers are referred to our technical report [1] for the detailed proof of Theorem 3. From Theorem 3, we can derive the following corollary:

Corollary. *At Nash equilibrium, the bid prices to each upstream peer i from all competing players that are allocated non-zero bandwidths are the same, i.e., $\exists t_i^*, p_{ij}^{s*} = t_i^*$ if $x_{ij}^{s*} > 0, \forall j : (i,j) \in \mathcal{A}_s, \forall s \in \mathcal{S}$.*

This corollary can also be intuitively illustrated: If a player in auction i is paying a price higher than some other player who is also allocated non-zero bandwidth, the former can always acquire more bandwidth from the latter with a price lower than its current price. Thus at equilibrium, when no one can unilaterally alter its price, all players must be paying the same price.

D. Implementation concerns

Towards a practical implementation of the distributed auctions in realistic peer-to-peer networks, two scenarios need to be considered:

1) *Peer asynchrony*: With different processing speeds and message passing latencies, peers in real world are inherently asynchronous. As bids and allocated bandwidth updates may arrive at each upstream or downstream peer at different times, each auction is carried out in a completely asynchronous fashion.

In our design, bids and allocation updates are passed by messages sent over TCP, such that their arrival is guaranteed. In each auction, the upstream peer allocates bandwidth after it has received new bids from all its existing downstream peers in all the overlays it participates in, or a timeout value, T , has passed since the previous allocation. If a bid does not arrive before the timeout, the upstream peer assumes the corresponding downstream peer is not interested in requesting bandwidth from itself in this round. Similarly, at each downstream peer in each streaming overlay, it starts its price adjustment and requested bandwidth reallocation after all allocated bandwidth updates have arrived from all its requested upstream peers, or time T has passed since the last time it placed all the bids.

In this way, each upstream peer only needs to synchronize its allocation across its own upload links, and each downstream

peer synchronizes its bid updates across its own download links in each overlay. No synchronization is required among different upstream or downstream peers.

2) *Peer dynamics*: In a practical network, the players in each auction may change dynamically, due to new peers joining the network, existing peers switching upstream peers, or peer failures and departures. Our asynchronous implementation design can readily adapt to all such dynamics.

When a new peer joins a streaming overlay, it initiates bid prices towards all known upstream peers to 0 and calculates its initial bandwidth requests. Then it sends bid to each upstream peer from which it requests a non-zero bandwidth. Similarly, in the case that an existing peer in an overlay decides to bid at a new upstream peer, it initializes the bid price to 0, computes requested bandwidth together with those to other upstream peers, and then forwards its bid to the new upstream peer. In this way, new peers can immediately participate in the respective auctions.

When a peer fails or departs from an overlay, its upstream peer(s) can detect this based on the broken connections; when a peer quits the auction at an upstream peer, the upstream peer will not receive its new bids within the time bound T . In either case, a corresponding upstream peer allocates upload bandwidth in a new round to the bidding peers only, naturally excluding the departed peer from the auction game. At the downstream side, after detecting the failure or departure of an upstream peer, or discovering that an upstream peer can no longer provide new media content of an overlay, a downstream peer simply excludes it from its bandwidth request calculation.

IV. PRACTICALITY OF PROPOSED STRATEGIES: AN EMPIRICAL STUDY

We dedicate this section to in-depth investigations of how the proposed auction strategies perform in practical scenarios. Using simulations under real-world asynchronous settings, the focus of this study is to show that, as an outcome of our proposed strategies, coexisting overlay topologies can fairly share network bandwidth, evolve under various network dynamics, and be prioritized.

A. Limited visibility of upstream peers

In Assumption 1 of our convergence analysis in Sec. III-C, we assume that upload capacities in the network can be fully utilized to support all the peers to stream at required rates. This is generally achievable when each peer knows a lot of other peers in each overlay it participates in. However, in practical scenarios, a peer only has knowledge of a limited number of upstream peers in each overlay. We now study the convergence and optimality of the proposed strategies in such practical cases. The *set* of all known upstream peers to a peer is henceforth referred to as the *upstream vicinity* of the peer.

In our investigations, peers in each upstream vicinity are randomly selected by a bootstrapping server from the set of all possible upstream peers. We seek to answer the following questions with empirical studies. First, what is the appropriate size of the upstream vicinity, such that the required streaming rate can be achieved at all peers in an overlay? Second, if the




Fig. 4. Outcomes of distributed auctions in networks of different sizes, and with various numbers of upstream peer candidates in the upstream vicinities.

upstream vicinity is smaller, do peers need to bid for more rounds before the auction games converge? Finally, if auction game strategies are used with upstream vicinities of limited sizes, how different is the resulting topologies from the ones achieved when upstream vicinities contain all the other peers in the overlay, with respect to total streaming cost?

Evaluation. We investigate by experimenting in networks with 100 to 10,000 peers with various sizes of upstream vicinities. We set up the following realistic environment for our forthcoming experiments: Each network includes two classes of peers, 30% Ethernet peers with 10 Mbps upload capacities and 70% ADSL/Cable modem peers with heterogeneous upload capacities in the range of 0.4–0.8 Mbps. There exists one server — which is an Ethernet peer — serving a 1 Mbps media stream to all the peers (we now consider a single overlay). We use delay-bandwidth products to represent streaming costs (M/M/1 delays), with streaming cost functions in the form of $D_{ij}^s = x_{ij}^s / (C_{ij} - x_{ij}^s)$. Here, C_{ij} is the available overlay link bandwidth, chosen from the distribution of measured capacities between PlanetLab nodes [2].

Fig. 4 illustrates the outcome of our distributed auction strategies, either when they converge, or when a maximum number of bidding rounds per peer, 100, has been reached. In the latter case, we can assume the games have failed to converge, as there exist peers that cannot achieve the required streaming rate with their current size of upstream vicinities. Decreasing the size of upstream vicinities from $n - 1$ where n is the total number of peers in each network, we discover that with 15–20 peers in the upstream vicinity, the games can still converge and the streaming rate can still be satisfied at all peers in most networks, as shown in Fig. 4(A) and (B). Fig. 4(B) further reveals that convergence is always carried out rapidly in all networks with different sizes of upstream vicinities, as long as these games converge at all with a particular upstream vicinity size. Fig. 4(C) compares the optimality of resulting topologies in terms of streaming costs. Compared to the ultimate minimum streaming cost achieved when upstream vicinities contain all other peers in the overlay, costs experienced by using upstream vicinities of a much smaller size (20) are only 10% higher. Not shown in Fig. 4, another observation we made during the experiments is that the number of upstream peers selected from the upstream vicinities to serve a particular peer (with non-zero upload bandwidth allocation) is at most 4, with an average of 1–2.

Summary. From these empirical observations, it appears that the appropriate size of upstream vicinities is relatively

independent of network sizes, and only a very small number of upstream peers are actually selected, which leads to sparse overlay topologies. Both are good news when our game strategies are to be applied in realistic large-scale networks.

B. The case of multiple coexisting overlays

We now proceed to study how our game strategies resolve the bandwidth competition among multiple coexisting streaming overlays. In particular, how does the topology of each overlay evolve, if coexisting overlays are started in the network? Do multiple coexisting overlays fairly share network bandwidth, and experience similar streaming costs?

Evaluation 1. We introduce more and more streaming overlays onto a 1000-peer network, constructed under the same settings as used in previous experiments, and with 20 peers in the upstream vicinities. At the beginning, all peers participate in one overlay and start to bid for their streaming bandwidths. Then every 50 seconds (each bidding round takes approximately one second), the peers join one more new streaming overlay. To clearly show the effects of an increasing number of coexisting overlays on the achieved streaming rate of each overlay, the required streaming rates for all overlays are set to the same 1 Mbps.

Fig. 5 illustrates the evolution of the average achieved peer streaming rate in each overlay, when 5 overlays are sequentially formed in the network. We see that upload capacities in the network can support up to 3 overlays to stream at their required rates, and become insufficient when the 4th and 5th overlay join.

In the former case with 1–3 overlays, every time a new overlay is formed, the games converge again to new equilibria very quickly. Fig. 6 further shows the costs experienced by coexisting overlays when their topologies stabilize. We observe both streaming and bidding costs are very similar across the multiple coexisting overlays.

In the latter case with 4–5 overlays in the network, Fig. 5 shows that the games fail to converge. We observed during the experiment that peers in each overlay bid higher and higher prices at their upstream peers, but were nevertheless unable to achieve the required streaming bandwidths. Similar rate deficits can be observed in all coexisting overlays from Fig. 5.

In practical peer-to-peer applications, some streaming overlays might expect to receive better service quality than others. For example, live streaming of premium television channels should enjoy a higher priority and better quality than regular




Fig. 5. The evolution of streaming rates in multiple coexisting overlays, with an increasing number of overlays over time.




Fig. 7. The evolution of streaming rates for multiple coexisting overlays with different budgets, and with an increasing number of overlays over time.




Fig. 6. A comparison of costs among multiple coexisting overlays.




Fig. 8. A comparison of streaming costs among multiple coexisting overlays with different budgets.

ones. Since our game strategies can achieve fairness among various overlays (as observed from Fig. 5 and Fig. 6), we wonder if it is further possible to introduce a practical prioritization strategy in our games, such that differentiated service qualities can be provided to different overlays.

In our previous experiment, we have observed that overlays fairly share bandwidth for a simple reason: peers in different overlays are not constrained by a bidding *budget*, and they can all raise bid prices at will to acquire more bandwidth from their desired upstream peers, which leads to relative fair bandwidth allocation at the upstream peers.

Motivated by such insights, we introduce a budget-based strategy to achieve service differentiation, by offering higher budgets to peers in higher priority overlays. To introduce such budgets, we only need to make the following minor modification to the bidding strategy proposed in Sec. III-B:

When a peer j joins a streaming overlay s , it obtains a bidding budget W_s from its bootstrapping server. Such a budget represents the “funds” peer j can use to acquire bandwidth in overlay s , and its total bidding cost to all upstream peers cannot exceed this budget, i.e., $\sum_{i:(i,j) \in A_s} p_{ij}^s x_{ij}^s \leq W_s$. All peers in the same overlay receive the same budget, and the bootstrapping server assigns different levels of budgets to different overlays based on their priorities. During its price

adjustments in overlay s , peer j may only increase its bid price if the incurred total bidding cost does not exceed W_s .

Evaluation 2. Applying the budget-based bidding strategy, we perform the previous experiment again and show our new results in Fig. 7 and Fig. 8. The levels of budgets assigned to peers in overlay 1 to 5 range from low to high.

Comparing Fig. 7 with Fig. 5 in the cases when 1 to 3 overlays coexist, we see that overlays can still achieve their required streaming rates within their budgets. However, when comparing Fig. 8 to Fig. 6(A), we observe that the streaming costs are differentiated across overlays, i.e., overlays with larger budgets achieve lower streaming cost than those with smaller budgets. This is because the former can afford to pay higher prices and thus eclipse the latter in auctions at their commonly desired upstream peers.

A further comparison between Fig. 7 and Fig. 5 (when 4 or 5 overlays coexist) shows that, when upload capacities become insufficient, the overlay with the highest budget, overlay 4 or overlay 5 in respective phases, always achieves the highest and most stable streaming rates, while those for overlays with smaller budgets become less sufficient and less stable.

Summary. With respect to fairness without budgets, we have observed that, no matter if upload capacities are sufficient

or not, our game strategies achieve fair bandwidth sharing among multiple coexisting overlays. When overlays are able to achieve their required streaming rates, they also experience similarly costs, which further reveal their fair share of lower latency paths. Further, we show that by introducing budgets to our bidding strategy, we are able to differentiate service qualities among coexisting overlays.

C. Overlay interaction under peer dynamics

Finally, we study how coexisting streaming overlays evolve with peer arrivals and departures, with or without differentiated budgets.

Evaluation. We simulate a dynamic peer-to-peer streaming network, in which 2 servers concurrently broadcast 4 different 60-minute live streaming sessions, at the streaming rate of 300 Kbps, 500 Kbps, 800 Kbps and 1 Mbps, respectively. Starting from the beginning of the live broadcasts, 1000 peers join the network following a Poisson process. The inter-arrival times follow an exponential distribution with an expected length of *INTARRIV* seconds. Upon arrival, each peer randomly selects 2 broadcast sessions and joins the respective overlays; then the peer stays in the network for a certain period of time, following an exponential lifetime distribution with an expected length of *LIFETIME* seconds. In this way, we simulate 4 dynamically evolving streaming overlays with approximately the same number of participating peers at any time. All other settings of the experiment are identical to those in previous experiments. Applying our game strategies, we monitor achieved streaming rates at existing peers in each dynamic overlay during the 60-minute broadcasts.

The experiment was repeated with different values of *INTARRIV* and *LIFETIME*. First, when the bidding strategies are applied without budgets, Fig. 9 shows the results under two representative settings. For the scenario in Fig. 9(A), we observed during the experiment that with expected inter-arrival time of 1 second, 1000 peers have all joined the network in the first 10 minutes; with an expected lifetime of 30 minutes, approximately half of all the peers remain till the end of the broadcasts. For the scenario in Fig. 9(B), we observed peer arrivals last for 45 minutes with an expected inter-arrival time of 3 seconds, and most peers have left the network before the end with the expected lifetime of 10 minutes.

Comparing the two scenarios, the second represents more severe peer dynamics, as peers keep joining and leaving all the time during the broadcast; the overlays are more stable in the first scenario, as peer joins only occur at the beginning and each peer stays longer in the network. Therefore, Fig. 9(B) represents higher level of rate fluctuations than Fig. 9(A). In each scenario, a careful comparison of the rate fluctuation across different overlays reveals slightly larger fluctuations for overlays with larger streaming rate requirements. This is because different overlays fairly share upload capacities at common upstream peers, and the larger the required rate is, the harder it is to achieve.

However, when overlays with higher rate requirement are prioritized with higher budgets, Fig. 10 shows a different outcome. Compared to Fig. 9 under both settings, the prioritized

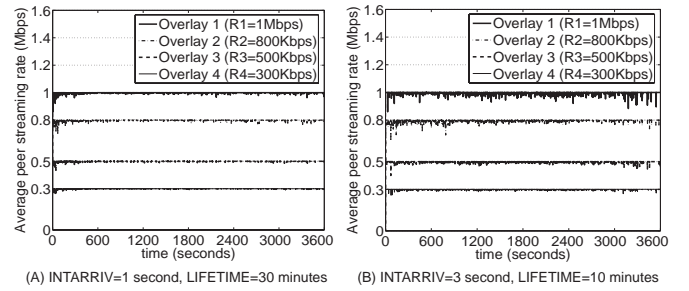


Fig. 9. Achieved streaming rates for 4 coexisting overlays: under peer dynamics without budget.

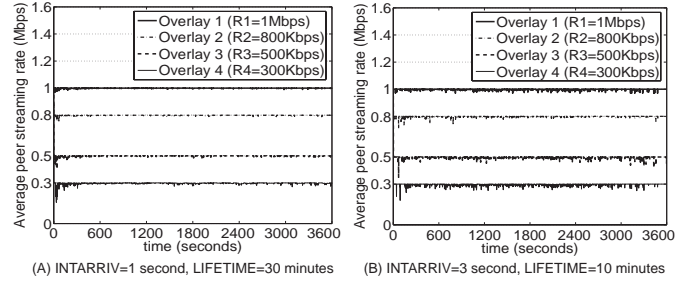


Fig. 10. Achieved streaming rates for 4 coexisting overlays: under peer dynamics with different budgets.

high-rate overlays always enjoy more stable rates, while low-rate overlays experience more severe rate fluctuations.

Summary. We have clearly demonstrated the effectiveness of our game strategies under high degrees of peer dynamics, which achieve stable streaming rates for all overlays at all times during such dynamics. Together with the budget-based strategy, they can further guarantee better streaming quality for prioritized overlays.

V. RELATED WORK

There exists very little literature that studies interactions and competitions among multiple coexisting overlays in a same peer-to-peer network. Recent work from Jiang *et al.* [3] and Keralapura *et al.* [4] is the most related, focusing on multiple overlay routing. In Jiang *et al.* [3], interactions among multiple selfish routing overlays are studied with a game theoretic model, where each overlay splits its traffic onto multiple paths and seeks to minimize its weighted average delay. In Keralapura *et al.* [4], route oscillations are investigated when multiple routing overlays inadvertently schedule their own traffic without knowledge of one another. Comparably, our work is significantly different, as we consider multiple peer-to-peer streaming overlays featuring many-to-many traffic, instead of point-to-point traffic in routing overlays.

Touching upon the topic of coexisting live streaming overlays, two recent pieces of work [5], [6] propose to encourage peers in different overlays to help each other by relaying media belonging to other overlays. While it is beneficial to improve network resource utilization at a specific time, there are questions remaining to be answered: How should each peer carefully allocate its upload capacity among concurrently requesting peers from different overlays? If new requests from peers in the same overlay come later, should the bandwidth

allocated to other overlays be deprived? From a more practical perspective, our work considers the case that each overlay consists of only receiving peers but each peer may participate in multiple overlays, and investigates bandwidth competition among the overlays at their common upstream peers.

Auction-based approaches have been proposed to allocate network bandwidth based on the demand and willingness to pay from competing users [7], [8], [9], [10]. A majority of such work are based on Progressive Second Price auctions, in which competitors decide their bids based on their true valuation. Aiming to solve the congestion problem on a single link or path, such existing work deals with elastic traffic, and competitors bid for their bandwidth share to maximize their utilities. In comparison, we design bandwidth auctions in a more complicated and practical scenario of constructing multiple streaming overlay topologies. Demanding an inelastic streaming rate at a lowest possible cost, each peer bids in multiple auctions, and adjusts its bid prices and requested bandwidths judiciously based on the current marginal cost of streaming from different upstream peers.

In peer-to-peer content distribution, game theory has been widely used to characterize peer selfishness and to provide incentives for peers to contribute their upload capacities (e.g., [11], [12], [13], [14]). As players in non-cooperative overlay construction games, each peer aims to maximize its own utility and tends to contribute if its utility is contingent on its contribution. Rather than modeling peer selfishness, our work utilizes the distributed and dynamical nature of auction games to design effective mechanisms for demand-driven dynamic bandwidth allocation, in which local games achieve globally optimal topology construction.

As the core of incentive mechanisms, service differentiation approaches are proposed to provide differentiated service quality to different peers, in terms of peer selection and bandwidth allocation [11], [12], [15]. Comparably, our work considers prioritization of an overlay as a whole and proposes an effective budget-based overlay service differentiation scheme, for which we are not aware of any existing studies.

Finally, pricing mechanisms [16], [17], [18] are proposed for a bandwidth provider to establish bandwidth prices to charge users, in order to regulate the behavior of selfish users and achieve social welfare maximization. Such pricing schemes are different from our auction games, in the sense that bandwidth prices are determined solely by the provider to maximize its revenue, rather than from bid prices placed by users.

VI. CONCLUDING REMARKS

This paper considers conflict-resolving strategies among coexisting overlays for streaming in peer-to-peer networks. Our objective is crystal clear: we wish to devise practical and completely decentralized strategies to allocate peer upload capacities, such that (1) the streaming rate can be satisfied in each overlay; (2) streaming costs can be globally minimized; and (3) overlays fairly share available upload bandwidths in the network. Most importantly, we wish to achieve global properties using localized algorithms. We use dynamic auction games to facilitate our design, and use game theory in our

analysis to characterize the conflict among coexisting overlays. Different from previous work, incentive engineering, selfishness and strategyproofness are not parts of our focus in this paper, whereas *practicality*, *simplicity* and *global optimality* are. We finally show that our proposed algorithm adapts well to peer dynamics, and can be augmented to provision service differentiation. Encouraged by our conclusions in this paper, we intend to work towards a real-world deployment of our proposal in coexisting streaming overlays in the near future.

REFERENCES

- [1] C. Wu and B. Li, "Strategies of Conflict in Coexisting Streaming Overlays," <http://iqua.ece.toronto.edu/papers/stratconflict.pdf>, ECE, University of Toronto, Tech. Rep., July 2006.
- [2] "PlanetLab IPerf," <http://jabber.services.planet-lab.org/php/iperf/>.
- [3] W. Jiang, D.-M. Chiu, and J. C. S. Lui, "On the interaction of multiple overlay routing," *Performance Evaluation*, vol. 62, no. 1-4, pp. 229–246, October 2005.
- [4] R. Keralapura, C.-N. Chuah, N. Taft, and G. Iannaccone, "Can Coexisting Overlays Inadvertently Step on Each Other," in *Proc. of 13th International Conference on Network Protocols (ICNP 2005)*, November 2005.
- [5] X. Liao, H. Jin, Y. Liu, L. M. Ni, and D. Deng, "AnySee: Peer-to-Peer Live Streaming," in *Proc. of IEEE INFOCOM 2006*, March 2006.
- [6] G. Tan and S. A. Jarvis, "Inter-Overlay Cooperation in High-Bandwidth Overlay Multicast," in *Proc. of the 35th International Conference on Parallel Processing (ICPP 06)*, August 2006.
- [7] A. Lazar and N. Semret, "Design and Analysis of the Progressive Second Price Auction for Network Bandwidth Sharing," *Telecommunications Systems: Special Issue on Network Economics*, 1999.
- [8] N. Semret, R. R.-F. Liao, A. T. Campbell, and A. A. Lazar, "Pricing, Provisioning and Peering: Dynamic Markets for Differentiated Internet Services and Implications for Network Interconnections," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 12, pp. 2499–2513, December 2000.
- [9] P. Maillé and B. Tuffin, "Multi-Bid Auctions for Bandwidth Allocation in Communication Networks," in *Proc. of IEEE INFOCOM 2004*, March 2004.
- [10] M. Bitsaki, G. D. Stamoulis, and C. Courcoubetis, "A New Strategy for Bidding in the Network-Wide Progressive Second Price Auction for Bandwidth," in *Proc. of the 1st Conference on Future Networking Technologies (CoNEXT 2005)*, October 2005.
- [11] R. T. B. Ma, S. C. M. Lee, J. C. S. Lui, and D. K. Y. Yau, "A Game Theoretic Approach to Provide Incentive and Service Differentiation in P2P Networks," in *Proc. of ACM SIGMETRICS/Performance'04*, June 2004.
- [12] G. Tan and S. A. Jarvis, "A Payment-based Incentive and Service Differentiation Mechanism for Peer-to-Peer Streaming Broadcast," in *Proc. of the 14th International Workshop on Quality of Service (IWQoS 2006)*, June 2006.
- [13] P. Golle, K. L. Brown, I. Mironov, and M. Lillibridge, "Incentives for Sharing in Peer-to-Peer Networks," in *Proc. of the 2nd International Workshop on Electronic Commerce*, 2001.
- [14] K. Lai, M. Feldman, I. Stoica, and J. Chuang, "Incentives for Cooperation in Peer-to-Peer Networks," in *Workshop on Economics of Peer-to-Peer Systems*, June 2003.
- [15] A. Habib and J. Chuang, "Incentive Mechanism for Peer-to-Peer Media Streaming," in *Proc. of the 12th International Workshop on Quality of Service (IWQoS 2004)*, 2004.
- [16] F. P. Kelly, A. Maulloo, and D. Tan, "Rate Control for Communication Networks: Shadow Prices, Proportional Fairness and Stability," *Journal of the Operational Research Society*, vol. 49, no. 3, pp. 237–252, March 1998.
- [17] T. Basar and R. Srikant, "Revenue-Maximizing Pricing and Capacity Expansion in a Multi-User Regime," in *Proc. of IEEE INFOCOM 2002*, June 2002.
- [18] Y. Qiu and P. Marbach, "Bandwidth Allocation in Ad Hoc Networks: A Price-Based Approach," in *Proc. of IEEE INFOCOM 2003*, April 2003.