Internet Video Multicast via Constrained Space Information Flow

Yaochen Hu[1], Di Niu[1] and Zongpeng Li[2]
1. University of Alberta, Canada
2. University of Calgary, Canada
{*yaochen,dniu*}*@ualberta.ca, zongpeng@ucalgary.ca*

1. Introduction

Many multimedia streaming and video webcasting applications can be modeled as a multicast session, where a single source node multicasts the same multimedia content to all the participating terminals, possibly with the help of relay servers. Since video content consumes a lot of network resources, it is desirable to minimize the congestion that the video session imposes on the Internet. The degree of congestion imposed on each link can be modeled by the bandwidth delay product, i.e., the network volume on it. A min-cost video multicast network should be constructed to minimize the sum of bandwidth delay products over all the links in the multicast network.

The min-cost multicast problem is traditionally solved on a given graph formed by the terminals, the source and all the available candidate relay servers in the Internet to find the optimal topology and flow assignments. Such an approach appears to be efficient when the number of candidate relay servers is relatively small. However, recent years have witnessed a rapid growth of the utilizable server pool including CDN nodes and small to medium datacenters, making traditional approaches incapable of handling the large scale of the graph at hand.

An alternative approach is to map the nodes onto a delay space using a network coordinate system, in which the distance between two nodes estimates their pairwise delay [2]. As such, the min-cost multicast network can first be constructed via geometric optimization in a delay space, where we can insert relay servers at arbitrary positions. Such optimal relay positions found in the delay space can then be mapped back to the closest real Internet servers. As long as the servers are densely distributed, this geometric optimization approach can yield a good approximation to the original min-cost multicast problem on a graph.

The remaining geometric problem in the delay space is similar to the *Space Information Flow* (SIF) problem [1], which aims to minimize the sum of bandwidth-distance products in a (geographic) space, allowing network coding and free insertion of relay nodes. The work in [1] presents a heuristic solution to the space information flow problem. However, it has two main drawbacks in practice. *First*, in real applications, introducing more relay server nodes will clearly lead to a higher cost. Neither does [1] consider such cost, nor is it able to solve the problem under a relay number constraint. *Second*, the solution in [1] approximates the geometric problem with a graph version of the problem by dividing the space with grids. This introduces a large number of intermediate variables, as such grids need to be fine-grained to increase accuracy. Although the overall mean complexity of the grid-based approach appears to be polynomial (depending on the optimization accuracy), it still has a large complexity especially when the dimension of the space is high.

In this paper, we propose the *Constrained Space Information Flow* problem, which aims to find the min-cost multicast network under a constraint on the number of relay servers, allowing operators to adjust the cost of using relay servers through such a constraint. We propose an effective EM algorithm to solve the problem *directly* in the geometric space, which can converge to the local optimal solutions with high efficiency.

2. Problem Formulation and Algorithms

In this section, we formulate the constrained space information flow problem, show some important properties of the optimal solutions to it, based on which we present our EM algorithm.

Constrained space information flow problem.
Although our idea can be extend to a general space, in this letter we focus our work on the min-cost multicast problem in a Euclidean space. Given N terminal nodes $T_1, T_2, \ldots, T_N$ with coordinates in a space (e.g., in a delay space where the distance between two nodes models their pairwise delay on the Internet) and a multicast session from one source node S to the N terminals as sinks, the objective is to construct a min-cost network in the space, allowing the introduction of at most M extra relay nodes, and allowing any form of coding including network coding to be performed. We define the total cost of the network as

$$\sum_e w(e) f(e),$$

where $f(e)$ is the information flow rate on link e, and $w(e)$ is the weight of the link. In a delay space, we set $w(e)$ as the link length $\|e\|$, i.e., the delay on link e.

The network cost is determined by two types of variables, one being the positions of the relay nodes, and the other being the flow assignments on the links. We call these two factors *positions* and *flow assignments*. Note that the flow assignments will also determine the connection topology of all nodes, since a link with a zero rate indicates that the link does not exist. Our problem is to tune these two sets of variables with no more than M relay servers to achieve the minimum cost. Denote V_R as the set of M candidate relay nodes to be found and V as the set of all nodes. Our optimization problem can be stated as

$$\textbf{Minimize} \quad \sum_{u,v\in V} \| x_u - x_v \| f(\overrightarrow{uv}) \quad (1)$$

Subject to :

$$\begin{cases} \sum_{v\in V} f_i(\overrightarrow{vu}) = \sum_{v\in V} f_i(\overrightarrow{uv}), & \forall i, \forall u \in V, \quad (2) \\ f_i(\overrightarrow{T_iS}) = r, & \forall i, \quad (3) \\ f_i(\overrightarrow{uv}) \le f(\overrightarrow{uv}), & \forall i, \forall u,v \in V, \quad (4) \\ f(\overrightarrow{uv}) \le c(\overrightarrow{uv}), & \forall u,v \in V, \quad (5) \\ f(\overrightarrow{uv}) \ge 0, f_i(\overrightarrow{uv}) \ge 0, & \forall i, \forall u,v \in V, \quad (6) \\ |V_R| \le M. & \quad (7) \end{cases}$$

In (1), x_u is the position of node u. The positions of the source node and the terminal nodes are fixed input vectors, while the positions of the relay server nodes are variables to be optimized. For every network information flow $S \rightarrow T_i$, there is a *conceptional* flow $f_i(uv)$. We call it *conceptional* because different conceptional flows share the bandwidth on the same link. As stated in (4), the final flow rate $f(uv)$ of a link uv should be no less than the maximum conceptional rate, which will directly affect the total cost. The constraint (2) guarantees the conceptional flow equilibrium property for every node and every conceptional flow i. The assigned "feedback" flow in (3) characterizes the desired receiving rate at each terminal. (5) is the trivial link capacity constraint. For every pair of nodes, we have both $f_i(uv)$ and $f_i(vu)$ to indicate the flows in two directions, so that (6) gives another trivial bound. Finally, (7) indicates the constraint on the maximum number of relay server nodes.

We need to solve this problem over both the variables x_u (*relay positions*) for u in V_R and all the conceptional flow assignments on all the links (*flow assignment*). Note that any feasible flow assignment satisfying (2)-(7) can be achieved with linear network coding in a single multicast session [3].

Properties of the Solutions.

It is hard to simultaneously obtain the optimal values for both relay positions and flow assignments, since it is not hard to verify that the problem (1)-(7) is non-convex. However, we have some good properties for the problem once we fix one set of variables.

When the *positions* of the relay server nodes are fixed, the proposed optimization problem (1)-(7) is reduced to a simple *linear program* (LP). The number of variables is $N+1$ times the number of links, i.e., $O(N(M+N)^2)$. The number of linear constraints is also $O(N(M+N)^2)$. Therefore, we can solve it efficiently with common LP solvers.

When the *flow assignment* of the network is fixed, the cost function in (1) is the sum of norms and all the constraints in (2) to (6) are irrelevant to the position variables. The optimization problem now reduces to a convex optimization problem. There are many efficient algorithms to solve such kind of problems. More specifically, for the sum of norms in this problem, an *Equilibrium* method has been proposed in [1], which can efficiently converge to the optimum.

Based on these observations, we propose our EM heuristic algorithm. In the EM algorithm, the above two local optimizations for relay positions and flow assignments are alternately performed.

An EM Algorithm.

Our proposed EM algorithm is shown in **Algorithm 1**. Initially, Step 1 randomly assigns the positions of the relay server nodes in the smallest box region containing all the terminals and the source. The following steps are iterative operations. In each iteration, there are three major steps. We first solve the LP in (1)-(7) with the relay positions fixed to obtain the flow rate assignments. Then with these flow rates fixed, we solve a convex optimization problem for the relay server node positions. Finally, for each relay node that has no throughput on it, we randomly reassign a new position to it and repeat the iterations. The ε in Step 5 is a small positive threshold to exclude the fake non-zeros, since in our LP solver there is always a small non-zero value on an actually zero-valued variable.

As for the *termination condition* in Step 9, we introduce a counter to help us monitor the termination condition. In each iteration, we first calculate the ratio between the cost in the previous iteration and the cost in the current iteration, and increase the counter if the ratio is less than some threshold. Once the counter reaches some number, the whole algorithm terminates. Finally, we delete the relay nodes that have no throughput and output solution..

Algorithm 1 EM Heuristic Algorithm

Input: N terminals, the source node, the constraint M on the number of relay servers

Output: a solution to the Constrained SIF problem Relay positions, flow assignments

1: Randomly generate M relays in the smallest box containing the source and all the terminals;
2: (**Flow Assignment**) Fix the relay positions, and solve the LP in (1)-(7) to get the flow rates;
3: (**Relay Position Tuning**) Fix the flow rates assigned in Step 2, and solve the convex optimization problem for the relay positions, e.g., using the Equilibrium method in [1].

(**Random Seed Generation**)
4: **for** i = 1 to M, **do**
5: **if** the total flow on Relay $i < \varepsilon$ **then**
6: Randomly generate a new relay i in the smallest box containing the source and all the terminals;
7: **end if**
8: **end for**
9: **if** the *termination condition* does not hold, **go to** Step 2;

10: Delete the relay server nodes that have no flow on them.

3. Simulation and Performance

We have simulated the EM algorithm in a 2-D Euclidean space. We choose 11 nodes uniformly at random in the unit box. 10 of them are set as the terminals, while the remaining one is set as the source. We set 4 as the constraint on the number of relays. Since the multicast cost will be proportional to the receiving rate, we assume the receiving rate is 1 at each terminal and assign a large capacity 10 to each link. For the termination condition, we will increment the counter if the former cost is no greater than 1.05 times the current cost. We stop the algorithm when the counter reaches 1000. Fig. 1 shows the resulted relay server positions and topology. The final total cost in terms of the sum of bandwidth-delay products is 6.72.

4. Conclusions

In this letter, we present a solution for the min-cost video multicast problem by considering the problem in a delay space based on network coordinates. We solve the resulted Constrained Space Information Flow problem using an EM algorithm, that can take into account the constraint on the number of relay servers and can scale to a large pool of candidate servers. Preliminary simulations show that our algorithm can yield fast convergence to the local optimal solutions to the non-convex combinatorial optimization problem.

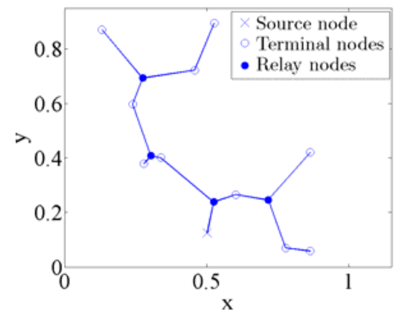

Fig. 1 The simulation result for randomly chosen 10 terminals and 1 source. The delay node number constraint is 4.

References

[1] Jiaqing Huang, Xunrui Yin, Xiaoxi Zhang, Xu Du and Zongpeng Li, "On Space Information Flow: Single Multicast," in the Proc. of Network Coding (NetCod), 2013 International Symposium on 7-9 June 2013.

[2] F. Dabek, R. Cox, F. Kaashoek, and R. Morris, "Vivaldi: A decentralized network coordinate system," in Proc. of ACM SIGCOMM, 2004.

[3] R. Ahlswede, N. Cai, S.T.R Li, and R. W. Yeung, Network Information Flow, IEEE Trans. On Information Theory, 46(4):1204-1216, 2000.

[4] S. Boyd and L. Vandenberghe, Convex Optimization (7e), Cambridge University Press, 2007.