On the Dynamic Multicast Problem for Coded Networks

Desmond S. Lun and Muriel Médard
Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
E-mail: {dslun, medard}@mit.edu

David R. Karger
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
E-mail: karger@mit.edu

***Abstract*— We consider the problem of finding minimum-cost time-varying subgraphs that can deliver continuous service to dynamic multicast groups in coded networks (i.e. networks that use network coding). This problem is relevant for applications such as real-time media distribution. We formulate the problem within the framework of dynamic programming and apply dynamic programming theory to suggest how it may be solved.**

I. Introduction

The fact that network coding can not only improve the throughput of multicast connections but also the complexity involved in establishing optimal multicast connections according to some cost criterion has been commented upon by several authors [1], [2], [3], [4]. Their key observation is that, while finding minimum-cost multicast trees in traditional routed networks (i.e. networks that do not use network coding) requires solving the directed Steiner tree problem, which is known to be NP-complete, minimum-cost multicast subgraphs in coded networks (i.e. networks that do use network coding) can be found by solving a linear program, which is therefore tractable, and which is moreover amenable to decentralized computation [5], [2].

This body of work, however, looks only at the case of static multicast, where a connection is set up for the use of a multicast group and then discarded. The membership of the multicast group is assumed to be constant for the duration of the connection. In many applications, however, the membership of the multicast group changes in time, with nodes joining and leaving the group. And, under these dynamic conditions, we often cannot simply re-establish the connection with every membership change because doing so would cause an unacceptable disruption in the service being delivered to those nodes remaining in the group. A good example of an application where such issues arise is real-time media distribution. Thus, we desire to find minimum-cost time-varying trees or subgraphs that can deliver continuous service to dynamic multicast groups. This is the dynamic multicast problem.

This work was supported by the National Science Foundation under grant nos. CCR-0093349 and CCR-0325496, by Hewlett-Packard under grant no. 008542-008, and by the Air Force Office of Scientific Research under grant no. F49620-01-1-0365.

Though our objective is clear, our description of the problem is currently vague. Indeed, one of the principal hurdles to tackling the dynamic multicast problem is in formulating the problem in such a way that it is suitable for analysis and addresses our objective. For routed networks, the dynamic multicast problem is generally formulated as the dynamic Steiner tree problem, which was first proposed in [6]. Under this formulation, the focus is on worst-case behavior and modifications of the multicast tree are allowed only when nodes join or leave the multicast group. The formulation is adequate, but not compelling; indeed, there does not appear to be any compelling reason for the restriction on when the multicast tree can be modified.

In this paper, we propose a formulation of the dynamic multicast problem for coded networks. We draw some inspiration from [6], but we focus on expected behavior rather than worst-case behavior, and we do not restrict modifications of the multicast subgraph to when nodes join or leave the multicast tree. We cast the problem in the framework of dynamic programming and, while the dynamic programming formulation that we arrive at is by no means straightforward to solve, it does appear as though it is amenable to a number of different approximate dynamic programming methods, which could lead to an implementable solution to the problem of dynamic multicast for coded networks.

We begin, in the following section, with a description of the problem formulation. In Section III, we discuss how the problem may be solved.

II. Problem formulation

We model the network with a directed graph $G = (N, A)$, where N is the set of nodes and A is the set of arcs. There is a special node s called the source node; we denote the remainder the of nodes by $N' := N \setminus \{s\}$. Each arc (i, j) represents a lossless point-to-point link from node i to node j, and we associate with it a convex, monotonically increasing cost function f_{ij} taking values in the extended real numbers: When packets are injected into arc (i, j) at rate z_{ij} for one unit time, a cost of $f_{ij}(z_{ij})$ is incurred. The cost might represent, for example, delay, energy, monetary cost, or imaginary weight cost. We call the vector z, whose elements are the rates

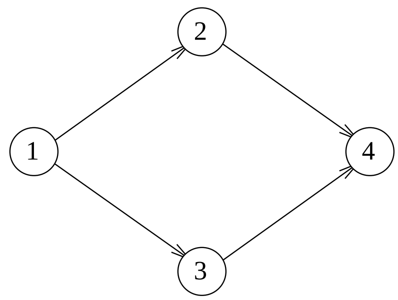

Fig. 1. A four node network.

at which packets are injected into each arc, the multicast subgraph.

The unit of time that we employ is related to the time that it takes for changes in the multicast subgraph to settle. In particular, suppose that at a given time the multicast subgraph is z and that it is capable of supporting a multicast connection to sink nodes T. Then, in one unit time, we can change the multicast subgraph to z', which is capable of supporting a multicast connection to sink nodes T', without disrupting the service being delivered to $T \cap T'$ provided that (componentwise) $z \geq z'$ or $z \leq z'$. The interpretation of this assumption is that we allow, in one time unit, only for the subgraph to increase, meaning that any sink node receiving a particular stream will continue to receive it (albeit with possible changes in the code, depending on how the coding is implemented) and therefore facing no significant disruption to service; or for the subgraph to decrease, meaning that any sink node receiving a particular stream will be forced to reduce to a subset of that stream, but one that is sufficient to recover the source's transmission provided that the sink node is in T', and therefore again facing no significant disruption to service. We do not allow for both operations to take place in a single unit of time (which would allow for arbitrary changes) because, in that case, sink nodes may face temporary disruptions to service when decreases to the multicast subgraph follow too closely to increases.

As an example, consider the four node network shown in Figure 1. Suppose that $s = 1$ and that, at a given time, we have $T = \{2, 4\}$. We support a multicast of unit rate with the subgraph

$$(z_{12}, z_{13}, z_{24}, z_{34}) = (1, 0, 1, 0).$$

Now suppose that the group membership changes, and node 2 leaves while node 3 joins, so $T' = \{3, 4\}$. As a result, we decide that we wish to change to the subgraph

$$(z_{12}, z_{13}, z_{24}, z_{34}) = (0, 1, 0, 1).$$

If we simply make the change naïvely in a single time unit, then node 4 may face a temporary disruption to its service as packets on $(2, 4)$ stop arriving and before packets on $(3, 4)$ start arriving. The assumption that we have made on allowed operations ensures that we must first increase the subgraph to

$$(z_{12}, z_{13}, z_{24}, z_{34}) = (1, 1, 1, 1),$$

allow for the change to settle by waiting for one time unit, then decrease the subgraph to

$$(z_{12}, z_{13}, z_{24}, z_{34}) = (0, 1, 0, 1).$$

With this series of operations, node 4 maintains continuous service throughout the subgraph change.

We discretize the time axis into time intervals of a single time unit. We suppose that at the beginning of each time interval, we receive zero or more requests from sink nodes that are not currently part of the multicast group to join and zero or more requests from sink nodes that are currently part of the multicast group to leave. We model these join and leave requests as a discrete stochastic process and make the assumption that, once all the members of the multicast group leave, the connection is over and remains in that state forever. Let T_m denote the sink nodes in the multicast group at the end of time interval m. Then, we assume that

$$\lim_{m \to \infty} \Pr(T_m \neq \emptyset | T_0 = T) = 0 \tag{1}$$

for any initial multicast group T. A possible, simple model of join and leave requests is to model $|T_m|$ as a birth-death process with a single absorbing state at state 0, and to choose a node uniformly from $N' \setminus T_m$ at each birth and from T_m at each death.

Let $z^{(m)}$ be the multicast subgraph at the beginning of time interval m, which, by the assumptions made thus far, means that it supports a multicast connection to sink nodes T_{m-1}. Let V_{m-1} and W_{m-1} be the join and leave requests that arrive at the end of time interval $m-1$, respectively. Hence, $V_{m-1} \subset N' \setminus T_{m-1}$, $W_{m-1} \subset T_{m-1}$, and $T_m = (T_{m-1} \setminus W_{m-1}) \cup V_{m-1}$. We choose $z^{(m+1)}$ from $z^{(m)}$ and T_m using the function μ_m, so $z^{(m+1)} = \mu_m(z^{(m)}, T_m)$, where $z^{(m+1)}$ must lie in a particular constraint set $U(z^{(m)}, T_m)$.

To characterize the constraint set $U(z, T)$, we make use of Theorem 1 of [1], which we now quote.

Theorem 1: Consider the following optimization problem:

$$\begin{aligned}
&\text{minimize} \sum_{(i,j) \in A} f_{ij}(z_{ij}) \\
&\text{subject to } z_{ij} \geq x_{ij}^{(t)}, \qquad \forall\, (i,j) \in A,\ t \in T, \\
&\sum_{\{j | (i,j) \in A\}} x_{ij}^{(t)} - \sum_{\{j | (j,i) \in A\}} x_{ji}^{(t)} = \sigma_i^{(t)}, \\
&\qquad\qquad\qquad\qquad \forall\, i \in N,\ t \in T, \\
&x_{ij}^{(t)} \geq 0, \qquad \forall\, (i,j) \in A,\ t \in T,
\end{aligned} \tag{2}$$

where

$$\sigma_i^{(t)} = \begin{cases} R & \text{if } i = s, \\ -R & \text{if } i = t, \\ 0 & \text{otherwise.} \end{cases}$$

The vector z is part of a feasible solution if and only if there exists a network code that sets up a multicast connection in graph G at rate arbitrarily close to R from source s to sinks in the set T and that injects packets at rate arbitrarily close to z_{ij} on each arc (i, j).

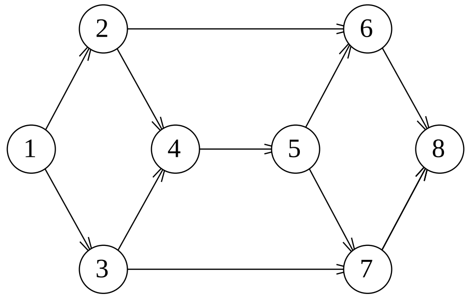

Fig. 2. A network used for dynamic multicast.

From Theorem 1, it follows that we can write $U(z,T) = U_+(z,T) \cup U_-(z,T)$, where

$$U_+(z,T) = \{z' \in Z(T) | z' \geq z\},$$
$$U_-(z,T) = \{z' \in Z(T) | z' \leq z\},$$

and $Z(T)$ is the feasible set of problem (2) for a given T.

It is worth noting that, as a consequence of the characterization of the constraint set $U(z,T)$, we can, from any $z^{(m)}$ such that $\sum_{(i,j)\in A} f_{ij}(z_{ij}^{(m)}) < \infty$, choose $z^{(m+1)}$ such that $\sum_{(i,j)\in A} f_{ij}(z_{ij}^{(m+1)}) < \infty$ provided that there exists $z \in Z(T_m)$ such that $\sum_{(i,j)\in A} f_{ij}(z_{ij}) < \infty$; that is, from any subgraph incurring finite cost at stage m, it is possible to go to a subgraph incurring finite cost at stage $m+1$ provided that one exists for the multicast group T_m. But while this is the case for coded networks, it is not always the case for routed networks. Indeed, if multiple multicast trees are being used (as discussed in [7], for example), then it is definitely possible to find ourselves in a state where we cannot incur finite cost at stage $m+1$ even though static multicast to T_m at finite cost is possible using multiple multicast trees.

As an example of this phenomenon, consider the network depicted in Figure 2. Suppose that a finite cost is incurred if the rate on each arc does not exceed 1 and that the cost is infinite otherwise. Suppose further that $s = 1$ and that, at a given time, we have $T = \{6, 8\}$. We support a multicast of rate 2 with the two trees $\{(1,3),(3,4),(4,5),(5,6),(5,7),(7,8)\}$ and $\{(1,2),(2,6),(6,8)\}$, each carrying unit rate. Now suppose that the group membership changes, and node 6 leaves while node 7 joins, so $T' = \{7, 8\}$. It is clear that static multicast to T' at finite cost is possible using multiple multicast trees (we simply reflect the solution for T), but we cannot achieve multicast to T' at finite cost by only adding edges to the two existing trees. Our only recourse at this stage is to abandon the existing trees and establish new ones, which causes a disruption to the service of node 8, or to gradually reconfigure the existing trees, which causes a delay before node 7 is actually joined to the group.

Returning to the problem at hand, we see that our objective is to find a policy $\pi = \{\mu_0, \mu_1, \ldots,\}$ that minimizes the cost function

$$J_\pi(z^{(0)}, T_0) = \lim_{M\to\infty} \mathbb{E}\left[\sum_{m=0}^{M-1} \sum_{(i,j)\in A} f_{ij}(z_{ij}^{(m+1)}) \chi_{2^{N'}\setminus\{\emptyset\}}(T_m)\right],$$

where $\chi_{2^{N'}\setminus\{\emptyset\}}$ is the characteristic function for $2^{N'}\setminus\{\emptyset\}$ (i.e. $\chi_{2^{N'}\setminus\{\emptyset\}}(T) = 1$ if $T \neq \emptyset$, and $\chi_{2^{N'}\setminus\{\emptyset\}}(T) = 0$ if $T = \emptyset$).

We impose the assumption that there exists $z \in Z(N')$ such that $\sum_{(i,j)\in A} f_{ij}(z_{ij}) < \infty$; that is, we assume that there exists a subgraph that supports broadcast at finite cost. This assumption, coupled with (1) ensures that there exists at least one policy π (namely, the one that uses z until the multicast group is empty) such that $J_\pi(z^{(0)}, T_0) < \infty$ and places an upper bound on the cost of the optimal policy.

It is now not difficult to see that we are dealing with an undiscounted, infinite-horizon dynamic programming problem (see, for example, [8, Chapter 3]), and we apply the theory developed for such problems to our problem in the following section.

III. Problem Solution

First, we note that the optimal cost function $J^* := \min_\pi J_\pi$ satisfies Bellman's equation; namely, we have

$$J^*(z,T) = \min_{u\in U(z,T)} \left\{ \sum_{(i,j)\in A} f_{ij}(u_{ij}) + \mathbb{E}[J^*(u, (T\setminus V)\cup W)] \right\}$$

if $T \neq \emptyset$, and $J^*(z,T) = 0$ if $T = \emptyset$. Moreover, the optimal cost is achieved by the stationary policy $\pi = \{\mu, \mu, \ldots\}$, where μ is given by

$$\mu(z,T) = \arg\min_{u\in U(z,T)} \left\{ \sum_{(i,j)\in A} f_{ij}(u_{ij}) + \mathbb{E}[J^*(u, (T\setminus V)\cup W)] \right\} \tag{3}$$

if $T \neq \emptyset$, and $\mu(z,T) = 0$ if $T = \emptyset$.

The fact that the optimal cost can be achieved by a stationary policy limits the space in which we need to search for optimal policies significantly, but we are still left with the difficulty that the state space is uncountably large; it is the space of all possible pairs (z,T), which is $\mathbb{R}^{|A|} \times 2^{N'}$. The size of the state space more or less eliminates the possibility of using techniques such as value iteration to obtain J^*.

On the other hand, given J^*, it does not seem at all implausible that we can compute the optimal decision at the beginning of each time interval using (3). Indeed, the constraint set is the union of two polyhedra, which can be handled by optimizing over each separately, and, although the objective function may not necessarily be convex owing to the term $\mathbb{E}[J^*(u, (T\setminus V)\cup W)]$, we are, at any rate, unable to obtain J^* precisely on account of the large state space, and can restrict our attention to approximations that make problem (3) tractable.

For dynamic programming problems, there are many approximations that have been developed to cope with large state spaces (see, for example, [8, Section 2.3.3]). In particular, we can approximate $J^*(z,T)$ by $\tilde{J}(z,T,r)$, where $\tilde{J}(z,T,r)$ is of some fixed form, and r is a parameter vector that is determined by some form of optimization, which can be performed offline if the graph G is static. The specific approximations $\tilde{J}(z,T,r)$ that we can use and their performance are beyond the scope of this paper.

IV. Conclusion

We have given a formulation of the dynamic multicast problem for coded networks that lies within the framework of dynamic programming. Our formulation addresses the desired objective of finding minimum-cost time-varying subgraphs that can deliver continuous service to dynamic multicast groups in coded networks and, because it lies within the framework of dynamic programming, can be approached using methods developed for general dynamic programming problems.

The solution that we propose uses such methods to approximate the optimal cost function, which is used to modify the objective function of an optimization that determines the multicast subgraph to use during each time interval. Depending upon the approximation that is used for the optimal cost function, this optimization conducted every time interval may be tractable and may even be amenable to decentralized computation using the techniques developed in [5], [2] (or simple modifications thereof).

We do not explore specific approximations for the optimal cost function, and leave it as a clear avenue for further work.

References

[1] D. S. Lun, M. Médard, T. Ho, and R. Koetter, "Network coding with a cost criterion," in *Proc. 2004 International Symposium on Information Theory and its Applications (ISITA 2004)*, Oct. 2004.

[2] Y. Cui, Y. Xue, and K. Nahrstedt, "Optimal distributed multicast routing using network coding: Theory and applications," *SIGMETRICS Perform. Eval. Rev.*, vol. 32, no. 2, pp. 47–49, 2004.

[3] Y. Wu, P. A. Chou, and S.-Y. Kung, "Minimum-energy multicast in mobile ad hoc networks using network coding," submitted to *IEEE Trans. Commun.*

[4] Z. Li, B. Li, D. Jiang, and L. C. Lau, "On achieving optimal throughput with network coding," accepted to *IEEE Infocom 2005*.

[5] D. S. Lun, N. Ratnakar, R. Koetter, M. Médard, E. Ahmed, and H. Lee, "Achieving minimum-cost multicast: A decentralized approach based on network coding," accepted to *IEEE Infocom 2005*.

[6] M. Imase and B. M. Waxman, "Dynamic Steiner tree problem," *SIAM J. Disc. Math.*, vol. 4, no. 3, pp. 369–384, Aug. 1991.

[7] Y. Wu, P. A. Chou, and K. Jain, "A comparison of network coding and tree packing," in *Proc. 2004 IEEE International Symposium on Information Theory (ISIT 2004)*, Chicago, USA, June–July 2004.

[8] D. P. Berksekas, *Dynamic Programming and Optimal Control*, 2nd ed. Belmont, MA: Athena Scientific, 2001, vol. 2.