

A Comparison of Network Coding and Tree Packing

Yunnan Wu
 Dept. of Electrical Engineering
 Princeton University
 Princeton, NJ 08544 USA
 e-mail: yunnanwu@princeton.edu

Philip A. Chou
 Microsoft Research
 One Microsoft Way
 Redmond, WA 98052-6399 USA
 e-mail: pachou@microsoft.com

Kamal Jain
 Microsoft Research
 One Microsoft Way
 Redmond, WA 98052-6399 USA
 e-mail: kamalj@microsoft.com

Abstract — We compare network coding solutions and routing solutions, namely packing distribution trees, for the problem of information multicast. To enable the comparison, we develop greedy tree packing algorithms that repeatedly pack the maximum-rate distribution tree and a greedy tree packing algorithm based on Lovasz' proof to Edmonds' Theorem. We then investigate the potential advantages of network coding over routing. In terms of throughput, tree packing performs comparably to network coding on the network graphs of six Internet service providers. However, network coding offers additional benefits, including fewer network resources consumed, ease of management, and robustness.

I. PACKING DISTRIBUTION TREES

Consider information multicast from a sender s to a set of receivers T in a communication network represented as $G = (V, E, c)$, where V and E are the set of vertices and edges respectively and associated with each directed edge $e \in E$ is a non-negative edge capacity $c(e)$.

Conventionally, this is done by routing information over one or more multicast distribution trees, each connecting the sender s to the receivers T . A distribution tree, represented as $G_k = (V_k, E_k, c_k)$, has an associated multicast rate equal to the minimum edge capacity of the tree, i.e., $\min_{e \in E_k} c_k(e)$. Multiple distribution trees, say $G_k, k = 1, \dots, K$, may be used to achieve a high rate by using different trees to distribute different information streams, as long as their sum fits in the network, i.e., $\sum c_k(e) \leq c(e)$.

Finding the best collection of multicast distribution trees, providing the maximum sum rate, has been shown to be NP-hard (e.g., [1]). In this paper, we show that the *maximum-rate distribution tree*, i.e., the single tree providing the maximum rate, can be found in polynomial time. We now describe a procedure based on Prim's algorithm for the minimum spanning tree problem. During the process, each node v is classified as either “reached” or “non-reached,” showing if v has already been reached or not. Initially, the set of reached nodes U includes only the sender s . In each subsequent step, we select the edge uv with maximum capacity, among those pointing from a reached node to a non-reached node. Then the node v is added to the set of reached nodes U , and the edge uv is recorded as an edge on the distribution tree. The process continues until all the receivers have been reached.

An alternate algorithm to find the maximum-rate distribution tree is as follows. Consider all the edges one by one, in increasing order of their capacities. If the deletion of an edge does not disconnect a receiver from the sender then delete the edge else keep it. The edges that survive the deletion process form the maximum-rate distribution tree.

We also develop a greedy tree-packing algorithm based on Lovasz's constructive proof [2] to Edmonds' Theorem [3] on

packing spanning trees. Assuming all edges have unit-capacity and allowing multiple edges for each ordered node pair, the algorithm packs unit-capacity trees one by one and each tree is constructed by greedily augmenting a tree edge by edge, similar to the greedy tree-packing algorithm based on Prim's algorithm. The distinction lies in the rule of selecting the edge among those pointing from the set of reached nodes U to the set of non-reached nodes $V - U$, which is to choose the unit-capacity edge whose removal leads to the least reduction in the multicast capacity (see below). For details, see [4].

II. COMPARISON

Recently, Ahlswede *et al.* [5] show that the *multicast capacity*, which is the maximum rate that a sender can communicate common information to a set of receivers, is given by the minimum $C = \min_{t \in T} C_t$ of max-flows $C_t = \text{maxflow}(s, t)$ between the sender and each receiver. Moreover, they show that while the multicast capacity cannot be achieved in general by routing, it can be achieved by network coding.

Using the network topologies of six commercial Internet service providers, we compare the achievable rate for the greedy tree packing algorithms with the multicast capacity and the rate achieved with a distributed practical network coding system [6][4]. The achievable throughput for both greedy tree packing and practical network coding is observed to be reasonably close to the multicast capacity. It is also observed that packing multiple distribution trees offers a significant gain in throughput compared to using only one distribution tree (as in traditional IP multicast, for example).

Nevertheless, network coding offers additional benefits, including fewer network resources consumed, ease of management, and robustness to both ergodic losses (e.g., packet losses) and non-ergodic failures (e.g., node and link failures). In particular, network coding can provide a given rate at the minimum cost of consumed resources. If the cost is linear and additive in the used rates on the edges, the minimum cost can be found with a linear program that assigns the *union of flows* while minimizing the total cost; the solution is called the *minimum-cost union of flows*. For details, see [4].

REFERENCES

- [1] K. Jain, M. Mahdian, M.R. Salavatipour, “Packing Steiner trees,” *14th ACM-SIAM Symp. on Discrete Algorithms*, 2003.
- [2] L. Lovasz, “On two minimax theorems in graph theory,” *J. Combin. Theory B* 21, pp. 96-103, 1976.
- [3] J. Edmonds, “Edge-disjoint branchings,” in: *Combinatorial Algorithms*, ed. R. Rustin, pp. 91-96, Academic Press, NY, 1973.
- [4] Y. Wu, P. A. Chou, and K. Jain, “Practical network coding,” Technical report, Microsoft Research. In preparation.
- [5] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, “Network information flow,” *IEEE Trans. Information Theory*, IT-46(4):1204-1216, Jul. 2000.
- [6] P. A. Chou, Y. Wu, and K. Jain, “Practical network coding,” *51st Allerton Conf. Comm., Control and Comp.*, Oct. 2003.