

Fault Tolerant Network Coding

Ravishankar Krishnaswamy

ravishankar.k@gmail.com

Guide: Pandurangan Chandrasekaran

rangan@shiva.cs.iitm.ernet.in

Indian Institute of Technology Madras

May 15, 2007

Abstract

Consider a communication network in which a source node wishes to multicast information to some sink nodes on the network. In the traditional setting, every node can only pass on the data it receives from incoming links to the links leaving it. We consider an extended setting which gives the intermediate nodes more “power”. In our model, each node may send any *linear combination* of the received data on the outgoing links. Such protocols, known as *Linear Network Coding* schemes, have been proved to guarantee optimal multicast throughput (that is, the maximum number of different data packets the source can transmit to all the sinks in one execution of the protocol matches the theoretical bound). In this work, we address the problem of fault tolerance in network coding. More specifically, we are interested in obtaining reliable coding schemes which guarantee optimal (in a sense we shall explain later) throughput even when some edges stop functioning, or are corrupt by noise.

We modify an existing algorithm to provide a centralized scheme for finding such codes that tolerate any failure pattern from a polynomially bounded set of possible failure patterns in a network. A failure pattern is a set of edges that fail (transmits a zero message). When an edge fails, it is assumed to fail throughout the execution of the protocol. Our algorithm is rate optimal in the asymptotic limit with respect to the message packet size, in the sense that it sustains a flow rate equal to the minimum of the *max-flows* from the source to the sinks with the failed edges being deleted. Such codes could be utilized in networks where certain edges are prone to failure, and yet optimal throughput is expected regardless of the edges failing. We also present simple extensions which optimally tolerate corruption of edges by white noise (they transmit a purely random message).

1 Introduction

The concept of *Network Coding* was first introduced in the landmark paper [1] by Ahlswede et al. The authors prove bounds on the capacity of network coding and show that it is possible to obtain optimal throughput using network coding. The implication in a multicast setting is that one can achieve a throughput equal to the minimum of the *max flows* from the source to each of the sinks - the theoretical upper bound on network throughput. In [2], Jaggi et al obtain a polynomial centralized algorithm that identifies the throughput maximizing linear network code for a directed acyclic graph.

Our Focus: We primarily focus on the *robustness* (tolerance to the failure/corruption of certain edges) of network coding. We are interested in efficiently designing network codes that tolerate polynomially many failure patterns. A similar setting is also considered by [2] and [3], but they assume that the sinks have knowledge about the edges that failed. This assumption, however, may not be valid in real-life networks. To the best of our knowledge, this is the first centralized algorithm for network coding when edges fail, without any assumptions about apriori information of the specific failures.

Preliminaries: All networks considered are directed acyclic graphs $G = (V, E)$. Let $s \in V$ be the source node and $T = \{t_1, t_2, \dots, t_{|T|}\} \subseteq V \setminus \{s\}$ be the set of sinks. The source s wishes to multicast h elements x_1, x_2, \dots, x_h where $x_i \in \mathbb{F}_{q^i}$. Each edge $e \in E$ can transmit a single *packet* of information. If an edge were to have a higher capacity k , we could replace it with k links of unit capacity. In our framework, a packet p that is transmitted on an edge e is an $h + 1$ tuple $(m, b_1, b_2, \dots, b_h)$ where $m \in \mathbb{F}_{q^i}$ and $b_j \in \mathbb{F}_q$. The vector $b(e) = (b_1, b_2, \dots, b_h)$ is called the *global encoding vector* of edge e . Also $b(e)$ is such that $m = \sum_{i=1}^h b_i x_i$.

2 Tolerance to Edge Failures

This section deals with network codes that can tolerate failure of some edges. We first consider codes that tolerate single edge failures, and later extend it. We now provide an upper bound on the throughput an algorithm can achieve under an edge failure when network coding is and is not employed.

Upper Bound: Consider a network with one source, two sinks and n intermediate nodes. The source is connected to the n intermediate nodes, each of which has an outgoing arc to both the sinks. We classify the edges connecting the source and the intermediate nodes to be in *level 1*, and the edges connecting the intermediate nodes to the sinks to be in *level 2*. Clearly, the minimum of the *max-flows* from the source to each of the sinks in this setting is n . When we wish to tolerate an edge that can fail, we notice that this value drops to $n - 1$. This is the bound that our routing protocols should try to achieve. If we do not employ network coding, a rate of $n - 1$ cannot be achieved. This is because, if the source were to send $n - 1$ packets of information to the intermediate nodes, it could repeat only one of these packets in the remaining *level 1* edge. If $n > 1$ and a *level 1* edge not carrying a packet that is repeated fails, there is no way by which that packet could be retrieved. With network coding, however, the source could send packets x_1, x_2, \dots, x_{n-1} along the first $n - 1$ *level 1* edges and the packet $x_1 + x_2 + \dots + x_{n-1}$ to the last intermediate node. Note that any $n - 1$ out of the n vectors that are transmitted along *level 1* edges are *linearly independent*. Therefore, even when any one edge fails, the sink will definitely receive $n - 1$ linearly independent equations in the variables x_1, x_2, \dots, x_{n-1} . A similar argument works when up to k edges can fail in the network. We can show while the upper bound is $n - k$, without network coding, the maximum possible throughput is $\lfloor n/(k + 1) \rfloor$.

Lower Bound: Let G have a max flow of n from the source to each of the sinks. Let $carried(e_i, k)$ stand for the global encoding vector on edge e_i in the occurrence of failure on edge e_k . We first identify n edge disjoint paths from the source to each sink t . At any instant we consider the set B_t^k of such global encoding vectors on one edge from each path. As we traverse the edges, we keep updating these sets with the goal that eventually there would exist $n - 1$ linearly independent vectors in each set. $sinks(e_j)$ is the set of sinks in whose edge disjoint paths edge e_j occurs. a represents the pseudo-inverse

of b just like in [2]. We also maintain the $b - a$ table with values of the pseudo-inverse vectors (a) for $b \in B_t^k$ with respect to each subset of B_t^k of size $n - 1$ containing b (such a set would form a basis for \mathbb{F}_q^{n-1} due to the invariant maintained).

Theorem 2.1 *Given any graph G whose minimum max flow from s to the sinks is n , there exists a centralized network coding scheme that guarantees a throughput of $n - 1$ under any single edge failure.*

1. The flow decomposition procedure ([4]) is executed to obtain the n edge disjoint paths from s to each sink t . The imaginary edge g_i is made to precede the i^{th} such path.
2. For all k , Set $B_t^k = \mathcal{B}$ where $\mathcal{B} = \{(1, 0, 0, \dots, 0), (0, 1, 0, \dots, 0), \dots, (0, 0, \dots, 1), (1, 1, \dots, 1)\}$. Let each vector be the global encoding vector on each of the n imaginary edges leading to s . That is, for all $k = 1, 2, \dots, |E|$, we set $carried(g_r, k) = (0, 0, \dots, 1, 0, \dots, 0)$ - the r^{th} identity basis vector in \mathbb{F}_q^{n-1} for $1 \leq r < n$, and $carried(g_n, k) = (1, 1, \dots, 1)$. The $b - a$ table is then populated accordingly.
3. Traverse the edges in the topological order e_1 to $e_{|E|}$. It is to be noted that when traversing e_i , all incoming edges to $head(e_i)$ have already been traversed. Let e_j be the edge that is currently traversed. Perform steps 4 to 21
4. For $k = 0$ to $j - 1$
5. Initialize the list $U = \phi$.
6. For $t \in sinks(e_j)$
7. Obtain from the $b - a$ table for sink t , all values of (b_i, a_j) such that $b_i = carried(prev(e_j, t), k)$ and add the (b_i, a_j) tuple to the list U .
8. End For
9. Remove all redundant entries from U .
10. Using the replacement vector algorithm ([2]), find a vector b such that $b \cdot a_j \neq 0$ for all a_j such that $(b_i, a_j) \in U$.
11. Set $carried(e_j, k) = b$.
12. For $t \in sinks(e_j)$
13. Mark B_t^k for updating to $B_t^k \setminus \{carried(prev(e_j, t), k)\} \cup \{b\}$. Also run the pseudo-inverse procedure ([2]) to obtain the changes to be done to the $b - a$ table. Mark these changes as well.
14. End For
15. End For
16. Set $carried(e_j, j) = 0$.
17. For $t \in sinks(e_j)$
18. Set $B_t^k = B_t^k \setminus \{carried(prev(e_j, t), k)\} \cup \{0\}$.
19. Update the j^{th} group of the $b - a$ table such that only the b, a values corresponding to the non-zero b vectors in B_t^k remain (there will be $n - 1$ such vectors).
20. End For
21. Update all the changes in the sets C_t, B_t^k and the $b - a$ table that were marked for updating in step 13.
22. End Traversal

Proof Sketch: Though the way we assign the global encoding vector $carried(e, k)$ makes it appear that it depends on the edge e_k that has failed, we note that the global encoding vector is actually just a function of the incoming packets' global encoding vectors. This is so because, for two different sets of incoming global encoding vectors $\{b_1, b_2, \dots, b_i\}$ and $\{b'_1, b'_2, \dots, b'_i\}$, where i is the in-degree of the current node we are considering, the list U that is chosen would be the same if $b_j = b'_j$. Therefore, the outgoing global encoding vector would also be the same, regardless of which edge has failed. The fact that the above procedure keeps as many vectors linearly independent as it can is because of step 10.

Polynomial Edge Failure Patterns: We now extend the above result to an algorithm allowing failure in any one of polynomially many edge patterns. A failure pattern is a set of edges that fail (rather than just a single edge).

Theorem 2.2 *Given any graph G and a polynomial set \mathcal{FP} of failure patterns, suppose the minimum of the max-flow from s to each sink t is at least h even after deleting the failed edges of any pattern $\mathbf{p} \in \mathcal{FP}$. There exists a centralized network coding scheme that guarantees a throughput of h under any edge failure pattern.*

The basic idea is to generate a set of n vectors such that any h of them are linearly independent. We observe that a purely random generation of n vectors satisfies this requirement with sufficient probability so long as the field size is large. We then run the centralized algorithm trying to maintain one set of h linearly independent vectors for every failure pattern.

Noisy Edges: To handle a scenario where a failed edge transmits a perfectly random message (instead of $\mathbf{0}$), we add some redundancy to the message packet transmitted. If we modify the transmitted packet on an edge e to an $h + 2$ tuple $(m, b_1, b_1, b_2, \dots, b_h)$ instead of the $h + 1$ tuple $(m, b_1, b_2, \dots, b_h)$, we could detect white noise on edges with probability $1/q$ by performing the check that the 2^{nd} and 3^{rd} elements of the received packet are the same. We then treat these edges as failed and proceed.

References

- [1] R. Ahlswede, N. Cai, S-Y.R. Li, and R.W. Yeung. Network information flow. *IEEE Transactions on Information Theory*, 46(4):1204–1216, 2000.
- [2] S. Jaggi, P. Sanders, P.A. Chou, M. Effros, S. Egner, K. Jain, and L. Tolhuizen. Polynomial time algorithms for multicast network code construction. *IEEE Transactions on Information Theory*, 51(6):1973–1982, 2005.
- [3] R. Koetter and M. Medard. Beyond routing: An algebraic approach to network coding. In *Proceedings of the 21st Annual Joint Conference of the IEEE Computer and Communications Societies(INFOCOM)*, pages 122–130, 2002.
- [4] L.R. Ford Jr. and D.R. Fulkerson. Maximal flow through a network. *Canadian Journal of Mathematics*, 8:399–404, 1956.