Capacities of Erasure Networks

by

Brian Smith, S.B., M.Eng.

Qualifying Proposal
Communications, Networks, and Systems (CommNetS) Area
Department of Electrical and Computer Engineering
The University of Texas at Austin

Committee Members
Prof. Ari Arapostathis
Dr. Piyush Gupta
Prof. Robert Heath
Prof. Sanjay Shakkottai
Prof. Sriram Vishwanath (Supervisor)

The University of Texas at Austin

Feb 2007

Capacities of Erasure Networks

Brian Smith

The University of Texas at Austin, 2007

Supervisor: Sriram Vishwanath

We propose to investigate, in multiple senses, the “capacity” of several models of erasure networks. The defining characteristic of a memoryless erasure network is that each channel between any two nodes is an independent erasure channel. The models that we explore differ in the absence or presence of interference at either the transmitters, the receivers, or both; and in the availability of feedback at the transmitters. We will also investigate different performance measures for these networks: traditional information capacity, multicast capacity, and transport capacity.

Contents

List of Figures

Chapter 1

Introduction

Network information theory is a subset of information theory which deals with the transfer of information over a networked combination of channels, rather than across a single channel. There may be either a single, or more than one, source or destination, which require the transfer of the same, or perhaps unique, pieces of data. Large wireless systems, those with many users which often lack the ability for coordinated, universal control, are particularly valuable to study, both for practical and theoretical reasons. These kinds of systems are becoming ubiquitous: cellular networks continually must support increasing numbers of customers; more people expect higher-speed broadband access over their laptops, personal organizers, and phones; military applications which keep track of the course of battle can help save lives.

While both theoretical limits and many practical methods are known for single-user point-to-point communication models, basic understanding of these complex multiple-user systems is still lacking. There are several very natural, easy to formulate, of practical use, and still open problems in the field of multi-user information theory. For example, the precise capacity regions of the broadcast channel (where one transmitter wishes to communicate two independent data sources to two different receivers) and the interference channel (where two transmitters send independent data to two different receivers) are, in general, unknown. In these cases, it is important to do as much as possible to characterize the behavior of these multi-user channels. We can do so in several ways - by providing upper and lower bounds on the general case, by choosing simple but representative models that are more amenable to analytic description, and by studying the asymptotic behavior of such systems. My research focuses primarily on the last two methods.

We choose the erasure channel as a simple, yet wide-reaching, model of wireless packet communication, and then study both traditional information capacity and asymptotic transport and throughput capacities.

If we take a directed graph as the basis of our network model, then the distinguishing feature of an erasure network is that the edges each represent erasure channels: Either the symbol transmitted across an edge is received correctly at the endpoint, or a distinct "Error" symbol is received. The model results from viewing a system from the *network layer*, rather than the physical layer. We assume that some mode of error detection coding has already been performed, so that when a packet is received, we can be assured either that its contents are correct, or that we have no information about what the contents might actually be. Pioneering work on networks of erasure channels was done in [2] and [5]. We expand this work by examining more general models of erasure networks, involving different types of transmitter and receiver interference and feedback availability. My primary research aim is to understand the fundamental mathematical limits which govern communication in different models of erasure networks, in determining the information theoretic capacity when possible, and in other cases determining the properties that exact solutions to problems involving both systems that increase in size and density must have.

1.1 Information Capacity and Transport Capacity

The information capacity of a single channel is the maximum rate at which reliable communication is possible; it allows us to answer the question "In n uses, how many unique inputs can the output distinguish amongst with high probability?" (The number of distinguishable signals grows exponentially in n, with the capacity as the exponent. [13]) Information capacity is a theoretical limit on reliable communication; it ignores practical consideration by allowing arbitrarily long delays and unlimited computational abilities as the transmitter and receivers. Nonetheless, it is important to understand and be aware of the fundamental limits of communication, in an attempt to approach them with real-world algorithms and approximations.

When dealing with more than a single source and a single destination, the information capacity of a system becomes a capacity region, rather than the maximum of a scalar rate. The convex region is a set of all rate vectors $(R_1, R_2, ...)$ where the R_i are rates corresponding to the different source-destination pairs. Even the simple act of describing such a region, let along computing it, becomes exponentially more difficult as the number of possible sources and destinations increases.

One example of a multiple-source multiple-destination network problem is "multicast," which refers to the case when a number of different destinations all require the information generated by a single source. Alternatively, "multiple unicast" is the situation when several source-destination pairs all communication unique data.

Even with a single source-destination pair, networks composed of more than two nodes and more than a single channel also become problematic: these networks are generally called relay networks (since the intermediate nodes have no data of their own to transmit, and have the purpose of aiding the source in relaying its data to the destination), and the general capacity of such networks is unknown. There are several achievable strategies for the three-terminal discrete memoryless relay channel [4] dating back twenty-five years or more, but these only correspond to best known upper-bound in a particular subset of cases (the physically degraded and deterministic relay channels, for example).

Because there are only a relatively few cases in which the information capacity region of multiple-source and/or multiple-destination networks is known, different approaches must be considered. In order to still obtain a descriptive measure of the communication capability of networks, other tools are available. Even though information capacity is seen by some as the apex of information theory, it does not always tell the whole story. In addition to the difficulty of actually computing the capacity region, it reflects nothing about the ease of implementation or the delay requirement of a particular coding scheme. Random coding, the stock tool for demonstrating the achievability of a particular capacity value, is largely not considered for actual use in any real (or even simulated) system - the computation required and delay incurred would be astronomically huge (doubly exponential growth, with respect to desired probability of error, is hardly unknown).

All of this makes looking at other measures of the capabilities of any network a valuable exercise. One popular quantity is the *transport capacity*. Transport capacity is the distance-weighted sum of rates for a network, and provides a convenient scalar description of the amount of traffic that a network can support. The notion was introduced by Gupta and Kumar [10] to study the capabilities of wireless networks with additive Gaussian noise. Later, Xie and Kumar [11] provided a information-theoretic scaling law that shows, under certain high-attenuation conditions, the transport capacity of the additive Gaussian interference network can grow no faster than linearly in the number of nodes in the network. Franceschetti et al. [12] demonstrate that linear growth in the number of nodes is achievable in a random network with an additive Gaussian noise model using routing alone (and no network coding). The work is often done in two different network settings: First, the case of *dense* scaling, where more and more nodes are placed in a fixed area. This is the sense of Gupta and Kumar's original paper. [10] Alternatively, the case of an *expansive* network is studied, where the physical size of the network grows and the density of nodes remains a constant as the total number of nodes increases. These views are unified for the Gaussian interference network in [1], where it is shown that (arbitrarily close to) linear growth can be achieved in a dense network using a cooperative, multi-layered MIMO technique.

We use the descriptive tool of transport capacity to investigate the asymptotic capabilities of several different models of erasure networks.

1.2 Erasure Network Models

There is a growing interest in the study of the capacity of erasure networks with constraints that reflect the underlying physical layer [5]. One of the primary techniques used to study such networks is network coding. Network coding was first used to achieve the multicast capacity of deterministic wireline networks [6]. It has since been put to a variety of purposes, including wireless erasure networks with broadcast constraints (but no interference constraints) which were studied in [5]. This work studied systems with independent erasures between nodes using both random coding and random linear encoding techniques.

Erasure networks generate interest for two main reasons: First, they can be a reasonable model for a packetized network which uses error correction and detection coding. In fact, some practical in-use communication protocols, such as ethernet, use a check to decide whether to accept or reject a packet. Secondly, erasure networks are one of the few multiple-terminal networks which offer themselves up to analysis. By comparison, much recent work has been concentrated on network coding models in zero-error, wired networks without any kind of interference whatsoever [6]. Work involving network coding with interference is now beginning to be more common, for example [7–9].

We look at several different models of erasure networks, in order to cover multiple physical phenomenon, and to gain as much understanding as to the underlying fundamental dynamics of networks.

The simplest of these is the non-interference network model introduced in [2]. It represents a network by a directed acyclic graph, where each edge is an independent erasure channel, with a dedicated input and output for each edge. When the erasure probability is a constant ϵ on each edge, then it was shown that the single-source single destination capacity of the network is the traditional min-cut value (the sum of the number of edges crossing the cut) multiplied by $\bar{\epsilon}$.

The broadcast nature of wireless networks is accounted for by a new erasure network model in [5]. In this model, all of the edges that depart any given node are required to carry an identical symbol in any given timeslot, just as a wireless antenna transmits just one signal to all the antennas which may be receiving at that time. This *wireless erasure network* model assumes some time-sharing or other interference mitigation scheme, so that the symbols along all incoming edges to a node are received without interference, as in a vector. It has been demonstrated that, if the final destinations all know the positions of any erasures in the network, then the information theoretic cut-set upper-bound (a modification of the traditional min-cut bound for flow networks) is indeed achievable for multicasting.

This result capacity-achieving result inspires the question, what about a network with no broadcast constraint, but a multiple-access interference constraint instead? Indeed, for any wireless erasure network with only a broadcast constraint, one can create a "dual" multiple-access constraint network: Swap the source and destination nodes, and reverse the direction of every edge. Eliminate the broadcast constraint, and institute an additive finite-field multiple access channel at every receiver node. The cut-set upper bound for this network is then identical to that of the wireless broadcast erasure network.

Further, we can generalize and create a model which takes into account all of the cases so far described: In this model, each node is allowed multiple outputs, and each output is connected to one or more nodes' receivers by erasure channels. Each node has at least one receiver, which obtains the finite-field sum of all the unerased inputs to that receiver. The cut-set upper bound is easy to derive for this global model, but whether that bound is achievable is still unknown. Further, if we wish to investigate the transport capacity of any of the above networks, it is necessary to have a model which translates the geographic distance between any two nodes and the probability of erasure along the channel between them. We investigate three separate models: a threshold model (where perfect reception is assumed for distances smaller than a certain value, and no reception for greater distances), an exponential decay model, and a polynomial decay model.

1.3 Completed Work

Our research can be categorized under two main headings: Work in the field of information capacity, and work in transport capacity.

For a single-source single-destination erasure network, with additive-finite field receiver interference but no broadcast requirement at the transmitter, we have shown that the min-cut max-flow capacity is in fact achievable when side information detailing the erasure locations along all the links is available to the final destination.

Under the heading of transport capacity, we have shown that with or without receiver interference, the wireless erasure network with a broadcast transmit constraint has a transport capacity which grows no faster than linearly in the number of nodes when a minimum node separation constraint is enforced. This result holds for both one- and two-dimensional networks when the probability of a successful packet transmission decays exponentially with increasing distance between two nodes. We have also shown that it holds for a one-dimensional network with a polynomially decay, as long as the decay exponent β is greater than 3. Further, it has been shown that linear growth in the number of nodes is achievable by routing only, without any need for network coding at the intermediate nodes.

The ways in which we expect to extend these results are detailed in the following section.

1.4 Expected Contributions

(i) Characterize the multi-cast capacity region for a general erasure network with finite-field additive interference.

The multicast region of erasure networks with no interference, and with the wireless erasure network (with broadcast interference only), have been previously known. In our work, we have shown the single-user capacity of the erasure network with additive finite-field receiver interference, only. It is our goal to prove that the min-cut max-flow cut-set upper-bound is achievable both for multi-cast and single-user cases in the general model where both multi-antenna transmitter interference and multi-antenna receiver with additive finite-field interference are present. This work should also then allow us to describe more precisely the duality relationship between broadcast-constraint only erasure networks, and multiple-access constraint only erasure networks.

(ii) Describe the benefits of feedback to erasure networks.

We have demonstrated with a simple counter-example how, if the min-cut max-flow rate is to be achievable for the multicast wireless erasure network of [5], even with complete instantaneous feedback available to the transmitters, some sort of network coding is required. We propose to demonstrate that, for the single-source, single-destination model of this network, even limited feedback allows the maximum capacity to be achieved in with a simple strategy involving no coding of packets.

(iii) Provide new bounds on the capacity of Gaussian networks.

We propose to put new upper and lower bounds on models of Gaussian interference networks which have similar characteristics to our models of erasure networks: Solely a broadcast constraint at the transmitters, or solely an additive multiple-access constraint at the receivers.

(iv) Upper bound the transport capacity in erasure networks with polynomial decay models.

We have shown for a one-dimensional network, the total transport capacity is linearly bounded in the number of nodes for erasure networks under a polynomial decay model. We will demonstrate that the same claim holds for two-dimensional networks.

(v) Provide an achievable scheme for super-linear growth in an erasure model with a slowly decaying probability of error function.

The linear upper bound on total transport capacity is only valid for large enough decay factors in the polynomial decay model. We propose to find a scheme which manages to achieve an actual super-linear total transport capacity growth.

1.5 Organization

In this proposal, we first discuss what is known about the exact information capacity of different erasure network models in Chapter 2. This chapter continues by describing interesting extensions to the different erasure network models, and discussing the impact that these modifications may make. Chapter 3 describes the concept and history of transport capacity, and specifically details the contributions to the transport capacity of erasure networks. Both Chapter 2 and Chapter 3 conclude with a description of the proposed work. Finally, we summarize and present a schedule for the proposed work in Chapter 4.

Chapter 2

Information Capacity

Here we look at the capacity of erasure networks in the traditional information theory sense. We are interested in both single-source, single-destination capacity and multicast capacity. The distinguishing feature of erasure networks is that each edge in the directed graph describing the network topology represents an independent erasure channel. For simplicity, in this section we will consider networks whose transmit alphabet is limited to the two symbols $\{0, 1\}$, but all results hold for larger alphabets(say, of size q), with the trivial modification of multiplying the rate values by $\lg q$.

2.1 Prior Work

David Julian help pioneer the concept of the erasure network in [2]. This original model is loosely based on the concept of flow networks [14]: each link in the network is an independent erasure channel with an identical erasure probability. Further, each node in the network sees all of its incoming and outgoing edges as distinct: it is a completely interference-free network, as if it were a wireline network.

Julian proves, among other statements concerning the erasure channel, that the capacity of such a network is equal to the capacity of the equivalent flow network, multiplied by $\bar{\epsilon}$, where ϵ is the uniform erasure probability.

The concept of an erasure network was expanded to provide applicability to wireless networks in [5]. In this work, a model which accounts for the broadcast nature of the wireless medium was proposed: A *wireless erasure network* is still represented by a directed graph, but each node is required to transmit an identical symbol down each of its outgoing links.

This modification requires a change in the evaluation of the cut-set bound. While in flow networks and in the network of [2], the cut-capacity is simply the sum of the capacities of all edges that cross the cut, the broadcast requirement of [5] introduces a new formula. We take ϵ_{ij} as the erasure probability across the edge connecting node i to node j; $\mathcal{V}_s$ is a partition of the nodes such that source s and destination d satisfy $s \in \mathcal{V}_s$ and $d \in \mathcal{V}_d = \mathcal{V}_s^C$, and let $[\mathcal{V}_s, \mathcal{V}_d]$ represent the set of all forward edges which traverse the cut. Then the communication rate is upper-bounded by

$$R \leq \min_{\mathcal{V}_s : \mathcal{V}_s \text{is an s-d cut}} \sum_{i \in \mathcal{V}_s} \left(1 - \prod_{j:(i,j) \in [\mathcal{V}_s, \mathcal{V}_d]} \epsilon_{ij} \right). \tag{2.1}$$

Intuitively, we can understand the rate bound (2.1) as follows: Given a cut, every node which can possibly transmit a symbol across the cut, contributes to the sum a quantity equal to the probability that the symbol transmitted along *at least one of the outgoing edges* is successfully received across the cut. Intuitively, as long as any of the symbols crossing the cut from that node are not erased, the transmission is accounted successful.

The main result of [5] is that the cut-set rate of Equation (2.1) is indeed achievable, for multicast as well as single-destination networks, under one additional assumption: The destinations are all aware, as side-information, of the positions of all the erasures on each link of the network. This assumption is reasonable, as packets usually carry headers, and the extra amount of information required does not increase with the packet size.

To demonstrate the achievability of (2.1), the authors of [5] make use of an extremely interesting proof technique. Random coding is employed over B blocks, each of size n, at all nodes in the network. Given the locations of all erasures, the operation of the network is a *deterministic function* of any input. The destinations simply simulate the operation of the network, and an error only occurs if two different input sequences (corresponding to two different messages) produce identical output sequences at the destination.

Because there is no interference between incoming edges at any node, there is no ambiguity as to which time block any symbol corresponds. (Each node waits until all symbols corresponding to a given message block arrive before calculating its output function for that particular message block). It is therefore possible to follow the evolution of a message block as it traverses the network. As noted above, an error occurs when there exists at least one input sequence (corresponding to a message other than the message actually sent) which produces exactly the same output at the destination node as the actual message. The critical insight of the proof is in representing this error event (call it E) as the union of error events E_x. Let w_0 be the actual message transmitted, and w_1 be an alternative message. If $\mathcal{V}_x$ is a cut, then E_x is the event that, for all nodes in $\mathcal{V}_x$, the inputs w_0 and w_1 produce distinct outputs, but for all nodes in $\mathcal{V}_x^C$, the received symbols for that block are identical. Note that the E_x (corresponding to the different cut-sets) are disjoint, and their union is the event E that the destination has an identical block of symbols for both w_o and w_1.

Using the union bound, the probability of E is less than or equal to the sum of the probabilities of E_x, each of which decays exponentially in n. It is therefore the event E_x with the minimum exponent (corresponding to the min-cut) which dominates the sum and governs the achievable rate.

The key to this proof is the fact the we can follow the progression of a single message through the network - that there is no “mixing” between symbols corresponding to different blocks. As soon as receiver interference is introduced to the network, however, this simplification is no longer available. For example, in the simple relay network of Figure 2.1, the symbols transmitted by the relay in the second time block (corresponding to the message of the first time block) become entwined with the symbols which are simultaneously being transmitted by the source (corresponding to the message of the second time block).

In the next section, we will detail how we overcame this problem to show the

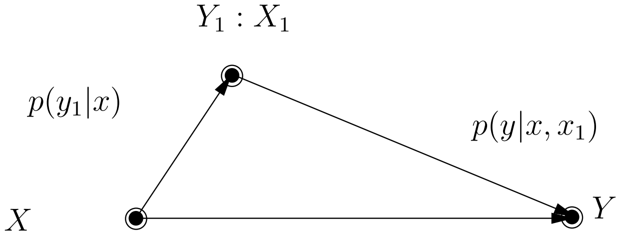

Figure 2.1: Relay Network with Interference

achievability of the cut-set bound in a network with additive finite-field interference, only.

2.2 Erasure Networks with Additive Finite-Field MAC Constraint

Consider a network model very similar to that of [5], but with a dual multiple-access channel constraint, instead of the broadcast channel constraint. That is, each node now can send a distinct symbol down all outgoing edges. Instead of a vector of incoming symbols, however, each node receives the finite-field sum of the unerased symbols along all incoming edges (erasures are treated as zeros in the sum). We show that the cut-set upperbound for a given MAC erasure network is identical to the upperbound obtained for the wireless erasure network in which

- The source and destination nodes of the MAC network are interchanged, and
- the direction of every edge in the network is reversed, but
- the erasure probabilities along each edge remain unchanged.

Further, we can show that this cut-set bound is achievable. The details of the model, and a discussion of our results, follow.

2.2.1 Notation

A directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ has vertex set $\mathcal{V} = \{1, 2, ..., |\mathcal{V}|\}$ and edge set $\mathcal{E} \subseteq (\mathcal{V} \times \mathcal{V})$. For a vertex $v \in \mathcal{V}$, let $\mathcal{N}_I(v)$ and $\mathcal{N}_O(v)$ be the sets of edges entering and leaving the vertex v, respectively. That is,

$$\mathcal{N}_I(v) = \{(u, v) : (u, v) \in \mathcal{E}\}$$
$$\mathcal{N}_O(u) = \{(u, v) : (u, v) \in \mathcal{E}\}$$

2.2.2 System Model

We investigate a single-source (denoted by s), single-destination (denoted by d) network, modeled by by an acyclic graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Let L be the length of the longest path between source and destination. Each node i has $|\mathcal{N}_O(i)|$ outputs and exactly one input, and at time t transmits the $|\mathcal{N}_O(i)|$-length vector of symbols $X_{ij}(t)$, where each symbol $X_{ij}(t)$, $(i, j) \in \mathcal{N}_O(i)$, is chosen from the alphabet $\{0, 1\}$. The symbol $X_{ij}(t)$ can depend on inputs to the node i from times $1..t-1$ and, if i is the source node, the current message.

At each time t, the node j will receive the single symbol $Y_j(t)$, where

$$Y_j(t) = \sum_{(i,j) \in \mathcal{N}_I(j)} \gamma_{ij}(t) X_{ij}(t) \tag{2.2}$$

and the $\gamma_{ij}(t)$ are all independent (over both time and edge indices) Bernoulli random variables which take the value 0 with probability ϵ_{ij}. We assume that each node knows the state of each of the channels incoming to that node, as well as the value of $Y_j(t)$.

In summary, each node has a single input and multiple outputs, which are connected to other nodes' inputs via erasure channels. At each timestep, every node receives an input $Y_j(t)$, the value of which is the finite-field sum of all the non-erased symbols transmitted by exactly the edges which are connected that node. This is illustrated in Figure 2.2. The probability that a symbol transmitted from an output of node i is successfully received (and therefore added to the sum) of the

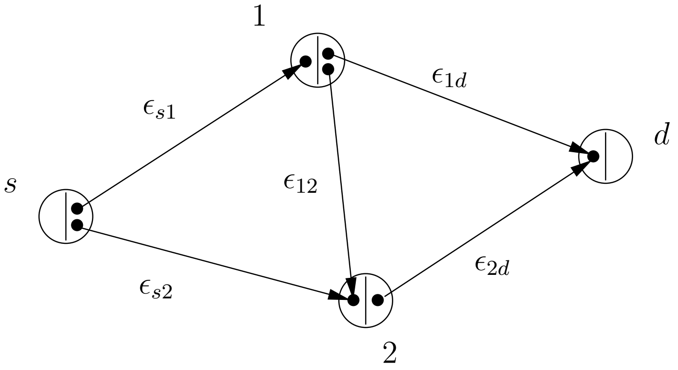

Figure 2.2: Erasure Network With MAC Constraint

input of node j is then $1-\epsilon_{ij}$. Equation (2.3) represents the relationships between the input and outputs of each node in Figure 2.2.

$$\begin{aligned} y_1 =& \gamma_{s1}x_{s1} \\ y_2 =& \gamma_{s2}x_{s2} + \gamma_{12}x_{12} \\ y_d =& \gamma_{1d}x_{1d} + \gamma_{2d}x_{2d} \end{aligned} \tag{2.3}$$

2.2.3 Results

The single-source single-destination rate must satisfy the upperbound

$$R \leq \min_{\mathcal{V}_s : \mathcal{V}_s \text{is an s-d cut}} \sum_{j \in \mathcal{V}_d} \left(1 - \prod_{i:(i,j) \in [\mathcal{V}_s, \mathcal{V}_d]} \epsilon_{ij} \right). \tag{2.4}$$

Further, the rate of Equation (2.4) is achievable, when side information on the locations of all erasures is available to the destination node.

2.2.4 Discussion

The proof again involves random coding over B blocks, each of length n, but instead of choosing B messages randomly from a set of size 2^{nR} (one for each block), we choose one message from a set of 2^{nRB}. The encoding functions at each intermediate node, however, remain functions of the length-n channel outputs of a single block. That is, each node i other than the destination randomly generates the $m = |\mathcal{N}_O(i)|$ different functions

$$f_{i^m} : \{0,1\}^n \times \{0,\epsilon\}^{n|\mathcal{N}_I(i)|} \rightarrow \{0,1\}^n.$$

Since we assume each node knows the channel state of all incoming edges, the inputs to each function are the received symbols in each time block and those erasure locations.

The error event E is no longer the simple union of error events associated with each cut $\mathcal{V}_s$, because we work over $B+L$ blocks (where L is the length of the longest path between source and destination). Here, let the event E_x^b be the event that, after the b^{th} time block, the outputs received at all nodes in $\mathcal{V}_x$ are distinct, but the outputs received at all nodes in $\mathcal{V}_x^C$ are identical. The error event E (that after all time blocks, the destination receives identical outputs for messages w_0 and w_1) can be described as

$$E = \bigcap_{b \in 1..B+L} \bigcup_{\mathcal{V}_x \text{is an s-d cut}} E_x^b. \tag{2.5}$$

Let $\{\mathcal{V}_x^{(b)}\}$ be a $B+L$ length sequence of s-d cuts $\mathcal{V}_x$. Interchanging union and intersection results in the error event expression

$$E = \bigcup_{\text{All sequences of s-d cuts}\{\mathcal{V}_x^{(b)}\}} \bigcap_{b \in 1..B+L} E_{x^{(b)}}^b, \tag{2.6}$$

which we can evaluate by properly bounding the probabilities of each possible sequence of error events. The calculation is complicated by the fact that the error events E_x^b are not independent across time. However, it can be shown that

$$P(E_{x^{(b)}}^b | E_{x^{(b-1)}}^{b-1}, E_{x^{(b-2)}}^{b-2}, ..., E_{x^{(1)}}^1) = P(E_{x^{(b)}}^b | E_{x^{(b-1)}}^{b-1}),$$

i.e. that the dependence is only over one previous time block.

This tool, with some further calculation, allows us to bound the probability of mistaking any two input messages for each other, and complete the achievability proof.

It still remains to prove that this cut-set rate is achievable in the multi-cast, as well as single-source single-destination, version of the problem.

Further, recall that the cut-set bound for a given MAC erasure network is identical to that of a wireless erasure network, with edge directions reversed and source and destination nodes swapped. We propose that there is a relationship between coding schemes for these two, dual networks. That is, if a network code for a given MAC erasure network is valid, then it can be used to construct an equivalent network code for the corresponding wireless erasure network.

2.3 Erasure Networks with Arbitrary Interference

The two erasure network models described so far each only consider a single type of interference - at the transmitter, only, or alternatively, at the receiver, only. The erasure network with arbitrary interference is a fairly general network description which encompasses both of the above models, incorporating both broadcast and receiver interference.

In this model, each node is allowed multiple outputs, and each output is connected to one or more nodes' receivers by erasure channels. Each node has at least one receiver, which obtains the finite-field sum of all the unerased inputs to that receiver. Specifically,

- Each node i is allowed to have multiple inputs (receivers) and multiple outputs (transmitters).
- Each edge in the network is a connection between an output i^m of one node and an input j_n of a different node. Each edge acts as an independent erasure channel.
- Each output i^m of a node is constrained to send the same symbol along all outgoing edges, but different outputs of a the same node may transmit different symbols.

- Each input j_n to a node receives the finite-field sum of all the non-erased symbols along incoming edges.

We must begin with a description of notation used to represent such a model.

2.3.1 Notation and Preliminaries

We modify the notation for directed graphs in order to allow for a broad variety of access and broadcast constraints on a network. A directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ has vertex set $\mathcal{V} = \{1, 2, ..., |\mathcal{V}|\}$ and edge set $\mathcal{E} \subseteq (\mathcal{V} \times \mathbb{Z}^+) \times (\mathcal{V} \times \mathbb{Z}^+)$. An edge $((i, m), (j, n)) \in \mathcal{E}$ will be abbreviated as (i^m, j_n). Let M_i and N_i denote the number of outputs and inputs a node i has (which is a distinct concept from the total number of edges entering or leaving the vertex - each edge connects one particular output i^m of a vertex i to one particular input j_n of a different vertex j).

For a vertex $v \in \mathcal{V}$, let $\mathcal{N}_I(v)$ and $\mathcal{N}_O(v)$ be the sets of edges entering and leaving the vertex v, respectively. That is,

$$\mathcal{N}_I(v) = \{(u^m, v_n) : (u^m, v_n) \in \mathcal{E}\}$$
$$\mathcal{N}_O(v) = \{(v^m, u_n) : (v^m, u_n) \in \mathcal{E}\}$$

We define similar notation for the set of edges entering and leaving any particular input or output of the vertex i, as well: $\mathcal{N}_I(v_n)$ and $\mathcal{N}_O(v^m)$.

An s-d cut for $s, d \in \mathcal{V}$ is a partition of $\mathcal{V}$ into two subsets $\mathcal{V}_s$ and $\mathcal{V}_d$ $(= \mathcal{V}_s^C)$ such that $s \in \mathcal{V}_s$ and $d \in \mathcal{V}_d$. The cut-set $[\mathcal{V}_s, \mathcal{V}_d]$ is the set of edges from the s-set to the d-set,

$$[\mathcal{V}_s, \mathcal{V}_d] = \{(u^m, v_n) : (u^m, v_n) \in \mathcal{E}, u \in \mathcal{V}_s, v \in \mathcal{V}_d\}$$

Further, for any s-d cut $\mathcal{V}_s$, define $\mathcal{V}_s^*$ and $\mathcal{V}_d^*$ as

$$\mathcal{V}_s^* = \{i | \exists j \text{ s.t.} (i, j) \in [\mathcal{V}_s, \mathcal{V}_d]\} \tag{2.7}$$
$$\mathcal{V}_d^* = \{j | \exists i \text{ s.t.} (i, j) \in [\mathcal{V}_s, \mathcal{V}_d]\} \tag{2.8}$$

so that $\mathcal{V}_s^*$ is the set of nodes in the s-set that has at least one of its outgoing edges in the cut-set, and $\mathcal{V}_d^*$ is the set of nodes in the d-set that has at least one of its incoming edges in the cut-set.

2.3.2 System Model

We consider a single-source (denoted by s), single-destination (denoted by d) network, modeled by by an acyclic graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Each node i has M_i outputs and N_i inputs, and at time t transmits the M_i-length vector of symbols $X_i(t)$, where each symbol $X_i^m(t)$, $m \in 1..M_i$, is chosen from the alphabet $\{0, 1\}$.

At each time t, the node j will receive the vector of symbols $Y_j(t)$ (of length N_j), where

$$Y_j^n(t) = \sum_{(i^m, j_n) \in \mathcal{N}_I(j^n)} \gamma_{j,n}^{i,m}(t) X_i^m(t) \tag{2.9}$$

and the $\gamma_{j,n}^{i,m}(t)$ are all independent (over both time and edge indices) Bernoulli random variables which take the value 0 with probability $\epsilon_{j,n}^{i,m}$. We assume that each node knows the state of each of the channels incoming to that node, as well.

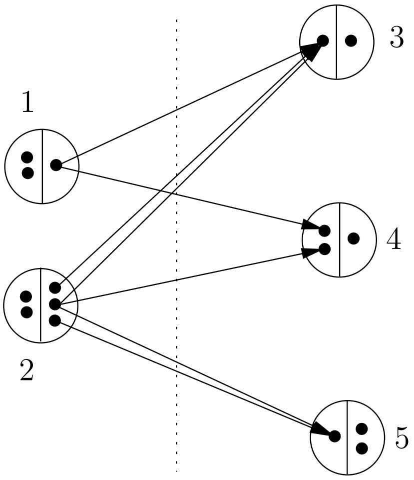

Figure 2.3: Detail of a General Erasure Network With Interference

In summary, each node has multiple inputs and outputs, which are connected to other nodes' inputs and outputs via erasure channels. At each timestep, every node receives a vector of symbols, where each element of the vector corresponds to one of that node's inputs. The value of each input's symbol is the finite-field sum of all the non-erased symbols transmitted by exactly the outputs which are connected that input. This is illustrated in Figure 2.3. The probability that a symbol transmitted from the m^{th} output of node i is successfully received (and therefore added to the sum) of the n^{th} input of node j is then $1 - \epsilon_{j,n}^{i,m}$.

2.3.3 Results

As of this point, we have only been able to prove an upper-bound on the possible rate for a single-source single-destination network.

For any $s - d$ cut, define the matrix A to be a random matrix with $0 - 1$ entries of size

$$\sum_{j \in \mathcal{V}_d^*} N_j \times \sum_{i \in \mathcal{V}_s^*} M_i.$$

Each column represents an output i^m of a node in $\mathcal{V}_s^*$, (i.e. an input into an erasure channel in the cut-set) and each row represents an input j_n of a node in $\mathcal{V}_d^*$ (i.e. the output of a channel, on the right side of the cut). For every edge (i^m, j_n) in the cut-set, there is an entry in (the appropriate row and column of) A_s which takes the value 0 with probability $\epsilon_{j,n}^{i,m}$ and 1 with the probability $1 - \epsilon_{j,n}^{i,m}$. Every entry in A_s which does have a corresponding edge in $\mathcal{G}$ will be zero with probability one. The matrix A_s then acts as a transfer matrix between the outputs of $\mathcal{V}_s$ and the inputs of $\mathcal{V}_d$: if all nodes on each of the two sides of the cut could cooperate perfectly, then we could collect all the outputs of the nodes in $\mathcal{V}_s^*$ into the vector X^* and the part of the inputs of the nodes in $\mathcal{V}_d^*$ that depends only on outputs from the s side of the $s - d$ cut into a vector Y^*. The relationship $Y^*(t) = A_s(t)X^*(t)$ describes the transfer of information across the cut.

Theorem 2.3.1. *The rate of reliable communication between the source s and the destination d in an erasure network, as defined in Section 2.3.2, is upperbounded as*

$$R \leq \min_{\mathcal{V}_s : \mathcal{V}_s \textit{is an s-d cut}} E[rank(A_s)]. \tag{2.10}$$

For example, the matrix A_s for the s-d cut $\mathcal{V}_s = \{1, 2\}$ illustrated in Figure 2.3 is demonstrated as

$$\begin{bmatrix} y_3^1 \\ y_4^1 \\ y_4^2 \\ y_5^1 \end{bmatrix} = \begin{bmatrix} \gamma_{3,1}^{1,1} & \gamma_{3,1}^{2,1} & \gamma_{3,1}^{2,2} & 0 \\ \gamma_{4,1}^{1,1} & 0 & 0 & 0 \\ 0 & 0 & \gamma_{4,2}^{2,2} & 0 \\ 0 & 0 & \gamma_{5,1}^{2,2} & \gamma_{5,1}^{2,3} \end{bmatrix} \begin{bmatrix} x_1^1 \\ x_2^1 \\ x_2^2 \\ x_2^3 \end{bmatrix}.$$

2.3.4 Discussion

We propose that this upper-bound is again achievable, when the destination nodes have side information on the positions of all erasures in the network, in both the single-source single-destination and multicast network cases. The proof will require some use of matrix theory.

2.4 Feedback

For a point-to-point channel, it is well known that feedback cannot increase the communication capacity. However, feedback can often be used to improve the actual mechanism of transmitting the information. For example, the Schalkwijk-Kailath coding scheme can substantially improve the error exponent.

For a single erasure *channel*, feedback allows an extremely simple communication scheme, requiring no coding over the bits whatsoever. The transmitter simply repeats each bit in an input sequence until it is alerted that the receiver successfully received that bit, and then moves on to the next bit.

For the wireless erasure network model of [5], feedback may also allow communication at the min-cut rate without the requirement of coding. However, for a wireless erasure network that wishes to multicast messages to multiple destinations, coding is required to achieve the cut-set bound. We exhibit a simple counter-example to demonstrate that some sort of coding will be required to achieve the multicast cut-set bound when only causal feedback (i.e. the transmitter obtains knowledge of the channel state only after it has already chosen which symbol it desires to send in a given timeslot) is available.

Alternatively, however, we propose to show that for a single-destination, feedback allows a very simple, randomized strategy, which allows communication at capacity with virtually no coordination an no knowledge of the network topology at any intermediate relay node.

2.4.1 Multicasting Counter-Example

Consider the simple three-node network of one source and two destinations in Figure 2.4. Each destination desires to have all bits, and each link has an erasure probability of 1/2. The cut-set rate to each destination is therefore 1/2.

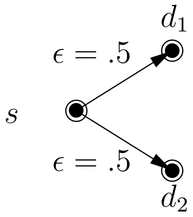

Figure 2.4: Simple Wireless Erasure Network

Without coding, the source has very few choices of transmit policies. In fact, they all are equivalent to the policy of the point-to-point erasure channel: transmit the first available bit until it is received by both destinations. After it has been received by both destinations, transmit the next bit. All possible policies are equivalent to this policy because of the restriction on any coding - without the ability to mix bits at any point, the transmitter must repeat the bit until it is received at both destinations. The transmitter may alter the order of the bits it sends (for example, attempt to send bits 1 through 10, then repeat the unsuccessful bits) but any permutation has no effect because the channel is memoryless.

Calculation of the Expected Number of Timeslots Required to Transmit One Bit: In the first timeslot, one of three things can happen.

- With probability 1/4, both destinations successfully receive the bit.

- With probability 1/2, exactly one destination receives the bit.
- With probability 1/4, neither destination receives the bit.

Let T be the expected number of timeslots required for both destinations to successfully receive the bit. Then,

$$T = \frac{1}{4}(1) + \frac{1}{2}(1+2) + \frac{1}{4}(1+T)$$

This equation is interpreted as follows: The expected number of timeslots required is the probability that both destinations get the bit in the first timeslot multiplied by 1, plus the probability that one of the destinations gets the bit in the first timeslot multiplied by 1 plus the expected number of timeslots required for the other destination to receive the bit, given that it did not in the first timeslot(2), plus the probability that neither destination received the bit multiplied by 1 plus the expected number of timeslots required until they both will receive the bit.

Solving the equation yields $T = 8/3 > 2$, so the multicast rate of this network is $3/8 < 1/2$.

Note that if the transmitter had *acausal* knowledge of the state, this network could achieve the multicast cut-set rate of 1/2. We know that 1/4 of the time, it can send a bit to both destinations successfully, 1/4 it can send a bit to destination d_1 successfully, and 1/4 of the time it can send a bit to destination d_2 successfully. If it knows beforehand when to schedule bits that need to go to each of the destinations individually, it could do so. (This would also require either packet headers, because the bits would then arrive out of order at the destinations, or it would also be possible if the destinations get feedback about the state of the other channel - knowing the scheduling algorithm, they would then be able to appropriately re-order the bits.)

2.4.2 Single Destination Wireless Erasure Network with Feedback

I propose to show that it is possible to communicate data, as the cut-set rate, at the cut-set upper bound rate when feedback is available to each node in the network. In fact, it seems to be the case that *only* feedback from the final destination, to each node in the network, is required.

The policy, which appears to be throughput optimal, has several nice features:

- No node needs to know anything about the topology of the network. (The only thing that the user needs to know is the value of the min-cut).
- The only feedback that is necessary is from the destination node, to all other nodes in the network. The destination notifies every other node whenever it successfully receives a packet.

This policy requires that packets be labeled with packet numbers, since they will arrive completely out of order at the destination.

Let each node, including the source but excluding the destination, have a queue. New packets will be injected into the source's queue with probability λ, where λ is less than the min-cut of the network (as defined in [5]).

At each time-step, any node with at least one packet in its queue will *randomly* choose a packet to broadcast. The node will not necessarily delete that packet from its queue, however. A packet is only removed from a node's queue when it is received at the final destination, at which point it is removed from all queues in the network.

Intuitively, imagine first a network consisting of just two channels in series. If the first channel is the bottleneck, then the second queue will remain relatively short. The probability of the first node having to resend a packet that it has already sent to the second node is small - it is more likely that that packet will go on, be received by the destination, and be deleted from the first queue, then it is that the packet will be randomly selected again from the (relatively long) first node's queue.

All of the nodes on the source-side of the min-cut will have relatively long queues, while all of the nodes on the destination side of the min-cut will have relatively short queues. We have to show, however, that all the queues are stable, which is tough, because the rate of innovative arrivals to any queue is dependent on the current length of that queue.

The policy has been simulated in a four node network, with values of ϵ_{ij} corresponding the min-cuts corresponding to each of the different cutsets, for a sanity check.

2.5 Expected Contributions

The expected contributions related to the information capacity of erasure networks are summarized below.

- Multicast for the MAC Erasure Network

  We have shown that the single-source single-destination cut-set rate is achievable for the MAC Erasure Network, but it remains to show that this rate is achievable in a multicast network. This work would be subsumed by the successful completion of the next proposed item.

- General Case (Multicast and Single-Source Single-Destination)

  We have computed an upper-bound for the erasure network with arbitrary combinations of transmitter and receiver interference. We propose to show that this bound is indeed achievable.

- Feedback

  We propose that, for a single-source single-destination wireless erasure network, no network coding is required to achieve the cut-set rate when a limited amount of feedback, from the final destination only, is available at all intermediate nodes.

- Duality

  We propose to show that the construction of a code for a MAC erasure network will lead directly to a valid code for the dual wireless erasure network.

2.6 Gaussian Interference Networks

Much of the work on the capacity of wireless network has focused on the physical layer, see for example the seminal work of [10]. A physical layer Gaussian interference network can be modeled on a directed graph in an intuitive manner: Each node has a single transmit and a single receive antenna. Nodes transmit a real signal under a transmit power constraint

$$E[X_i(t)] \leq P$$

and receive the signal

$$Y_j(t) = \sum_{i \in \mathcal{N}_I(j)} \sqrt{h_{ij}} X_i(t) + Z_j(t)$$

where h_{ij} are constant power decay coefficients and $Z_j(t)$ are white Gaussian noise with power σ_j.

Determining the capacity of the network type described above would be a fiendishly difficult problem: it includes, for example, the Gaussian relay channel as a subproblem. However, it can be instructive to look at simpler problems. In a follow-on to his work on the erasure network [3], Julian studied the case of parallel Gaussian channels composed with erasures. As we have had some success by simplifying the erasure network with arbitrary interference model into models with broadcast interference, only, or receiver interference, only, we propose to investigate Gaussian interference network from the same viewpoints. Because the Gaussian interference network has historically been such a difficult problem, we expect modest results. We hope for new bounds on the capacity of these network, but believe that regardless of the results, the investigation of such problems will prove illuminating.

Chapter 3

Transport Capacity

In Section 2.6, we described the Gaussian interference network model, and noted the difficulty of finding the capacity of the network. It is possible to continue to enhance and complicate the system model by allowing an arbitrary number of source-destination node pairs. Finding the exact capacity region of such a network is far beyond the capabilities of the current research. Nonetheless, it is still possible to provide an informative measure of the amount of information traffic a Gaussian interference network can support, regardless of the number of sources and destinations. The concept of transport capacity was introduced in [10] in order to do just that. In this chapter, we describe how we adopt the notion of transport capacity and apply it to different kinds of erasure networks.

3.1 Prior Work

Transport capacity, introduced in [10], was first used to describe the asymptotic capabilities of a Gaussian interference network on a fixed area, as the number of nodes in the network n (and therefore the node density of the network) scales larger and larger. Transport capacity is defined as the maximum, over all possible sets of feasible rate vectors, of the distance-weighted rate-sum, i.e.

$$T = \sup_{feasible\{R_1,\ldots,R_L\}} \sum_{l=1}^{L} R_l d(l) \tag{3.1}$$

where the supremum is taken over all possible sets of source-destination pairs, there are L source-destination pairs, $d(l)$ be the distance between the source node s_l and destination node t_l, and R_l is the rate associated with the l^{th} pair.

Making transport capacity meaningful requires that increasing physical distance separating two nodes to have a degrading effect on the communication ability between those nodes. Many authors [10],[11],[12],[1] invoke a polynomially or exponentially decaying power law to accomplish this.

Under the outage capacity and routing-only constraints of the model, [10] showed that the total transport capacity of the Gaussian interference network is upper-bounded as $O(\sqrt{n})$, and that at least $\Omega(\sqrt{n}/\lg n)$ growth is achievable by routing alone. Franceschetti et. al [12] sharpen the lower bound and demonstrate that linear growth in the number of nodes (removing the $\lg n$ term) is achievable in a random network with an additive Gaussian noise model using routing alone.

In an alternative to the fixed area, dense network scaling of [10], some authors [11] [12], [1] study a network whose geographic expanse increases with the number of nodes, so that the density remains constant. Xie and Kumar [11] provide an information theoretic upperbound of linear transport capacity growth for a sufficiently high-attenuation power law. Finally, Tse et. al demonstrate in [1] a hierarchical cooperative MIMO scheme which achieves linear transport capacity growth for both the dense and expansive network cases.

We have investigated the capabilities of erasure networks from a transport capacity viewpoint, and present our results in the following sections.

3.2 Transport Capacity of Erasure Networks

We provide a linear upper-bound (in terms of the number of nodes n) to the transport capacity in both the wireless erasure network (broadcast constraint, but no receiver interference) and the broadcast erasure network with interference (a single transmit and single receive antenna at each node), under a variety of decay models. Further, we have demonstrated that in a random wireless erasure network with interference, routing alone (no network coding) can achieve this order-wise optimal linear bound.

3.2.1 Models

In demonstrating the upper bound, we use three different models to describe how the probability of an unerased transmission decays with distance:

In the threshold model, over distances less than or equal to d^*, there is never an erasure, but for distances greater than d^*, an erasure will always occur. This model is motivated by straightforward signal-to-noise ratio considerations.

In the exponential model, the probability of successful transmission decays exponentially with increasing distance:

$$\epsilon\left(d\right) = 1 - \alpha^d = 1 - e^{-d/d^*},$$

where the decay parameter α satisfies $0 < \alpha < 1$. The exponential model is motivated by the error rates for coding over long blocks with a signal-to-noise ratio which decays with increasing distance according to a power-law.

In the polynomial model, the probability of successful transmission follows a power-law decay:

$$\epsilon\left(d\right) = 1 - \frac{1}{1 + d^\beta},$$

where $\beta > 0$.

Note that for all of the models, there is no erasure when $d = 0$ and a zero probability of successful transmission when $d = \infty$.

Because of the broadcasting requirement of the wireless medium, we can consider the network to be a complete graph: There exists some probability (non-zero, in all but the threshold model) that any node will successfully receive the symbol transmitted by any node in the network. Therefore, we no longer have a directed *acyclic* graph, as per the system models in Chapter 2. We do prove, however, that back-edges cannot increase the cut-set bound on capacity.

We also work in the expansive network model, by imposing a minimum separation distance constraint on the node placement: For all nodes i and j, $d_{ij} \geq d_{min}$ for some d_{min} which is independent of n.

In the wireless erasure network model, we make the same assumption of no receiver interference as [5]. For the broadcast erasure network with interference, we assume that in each timeslot, the node j receives the sum $Y_j(t)$, where

$$Y_j(t) = \sum_{i \neq j, i=1}^{n+1} h_{ij}(t) X_i(t) \tag{3.2}$$

and $h_{ij}(t)$ are independent (over both indices and timeslots) zero-one binary random variables that take the value 0 with probability ϵ_{ij}.

3.3 Results and Discussion

While the cut-set bounds for the wireless erasure network and broadcast erasure network with interference differ in form, we demonstrate that they both, in fact, are upper-bounded by a single expression.

This fact allows us to prove a common linear upper bound to transport capacity in a one-dimensional network, for all three (threshold, exponential, and polynomial with $\beta > 3$) decay models. When the nodes are distributed over a two-dimensional plane, however, the current bounds are only tight enough to show a linear upper-bound to transport capacity growth for the threshold and exponential decay models.

These proofs are, in general, accomplished by the technique of examining a set of individual network cuts, and *squeezing* nodes not involved in the particular cuts of interest as close as possible to each other, while still obeying the minimum separation constraint.

This proof technique suggest how to achieve *superlinear* transport capacity growth, if we remove the minimum inter-node separation constraint. Two clusters of nodes, each of size $n/2$, which are separated from each other by a distance of $\lg n$, can achieve at least $\Omega(n \lg n)$ total transport capacity.

3.4 Proposed Work

The expected contributions related to the transport capacity of erasure networks are summarized below.

- Two-Dimensional Linear Transport Capacity Upper-Bound

  We have shown that, for the wireless erasure network and the erasure network with additive finite-field interference, the total transport capacity is linearly bounded in the number of nodes when the network is one-dimensional, or when the network is two-dimensional, and the probability of a successful transmission decays exponentially with distance. The transport capacity for the two-dimensional case, when the probability decays as a polynomial function of distance, must be shown to also have a linear bound, for some sufficiently large decay factor β, likely $\beta > 4$.

- Low Polynomial Decay Factor

  We have shown that, when the minimum node-separation distance requirement is relaxed, greater than linear growth in total transport capacity is achievable. Our linear bound on total transport capacity, in the polynomial decay case, also requires sufficiently high attenuation. We propose to show a super-linear achievability scheme for the case of a smaller, attenuation factor.

Chapter 4

Summary and Plan of Work

4.1 Summary

The erasure channel is an appropriate model to describe the action of channel-coded communication from the network-layer point of view. It has the additional advantage of analytic tractability. Investigation of the operation of networks of such channels has prompted a number of compelling and intellectually interesting system level questions. For example, why do multiple-access constraint, only, and broadcast constraint, only, networks have a duality in their capacities? Can we extended these simplified network models to Gaussian interference networks? What are the similarities and critical differences between the erasure model and the Gaussian model that effect the asymptotic transport capacity results?

By exploiting several different analytical tools (for example, matrix theory, optimization theory, and real analysis), by cooperation with peers and researchers at the University of Texas at Austin and elsewhere, I hope to uncover more knowledge related to the common understanding of multi-user networks, in general.

4.2 Expected Time-line

(i) Ph.D. Proposal - February 2007.

(ii) **Spring 2007:** Complete information capacity of erasure networks proof, submit work to Information Theory Workshop (ITW). Travel to California Institute of Technology to work with Prof. B. Hassibi on feedback in wireless erasure networks.

(iii) **Summer and Fall 2007:** Complete transport capacity proofs for two-dimensional network case. Develop modified Gaussian interference network models, and define what side-information, if any, would help in achieving greater rates. Determine better bounds on capacity.

(iv) **Spring 2008:** Write Dissertation and achieve Graduation.

Appendices

Graduate Coursework

Major Coursework

Course ID	Course Name	Instructor	Grade
EE381J	Probability and Stochastic Processes	de Veciana	A
EE381K-8	Digital Signal Processing	Bovik	A
EE381K-2	Digital Communications	Andrews	A
EE381K-7	Information Theory	Vishwanath	A
EE381K-9	Advanced Signal Processing	Heath	A
EE381V	Channel Coding	Vishwanath	A
EE380N	Optimization in Engineering Systems	Baldick	A
EE381K-13	Communications Networks: Analysis and Design (Queuing Theory)	Shakkottai	A
EE381K-5	Adv. Telecommunication Networks	Shakkottai	A

Supporting Coursework

Course ID	Course Name	Instructor	Grade
CS388G	Algorithms: Technique and Theory	Plaxton	A
M381C	Real Analysis	Beckner	A
CS388C	Combinatorics and Graph Theory	Zimmerman	B+
M385C	Theory of Probability	Zitkovic	CR
M393C	Statistical Physics	Radin	A

Publications

Published work related to proposal

1. "Transport Capacity of Wireless Erasure Networks," B. Smith, S. Vishwanath. In *Proceedings of The 44th Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Sep. 2006
2. "Network-Coding in Interference Networks," B. Smith, S. Vishwanath. In *Proceedings of 2005 Conference on Information Sciences and Systems*, Baltimore, MD, Mar. 2005.
3. "Asymptotic Transport Capacity of Erasure Networks," Under preparation for submission to *IEEE Transactions on Information Theory*

Work under review/preparation related to proposal

1. "Routing is Order-Optimal in Erasure Networks with Interference," B. Smith, P. Gupta, S. Vishwanath. Submitted to the *2007 IEEE ISIT*.
2. "Capacity of MAC Erasure Networks," B. Smith, S. Vishwanath. Under preparation for submission.

Previous work

1. "Cooperative Communication in Sensor Networks: Relay Channel with Correlated Sources," In *Proceedings of The 42nd Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Oct. 2004
2. "Capacity Analysis of the Relay Channel with Correlated Sources," *IEEE Transactions on Information Theory*, in revision

Bibliography

[1] A. Ozgur, O. Leveque, and D. Tse, "Hierarchical cooperation achieves optimal capacity scaling in ad hoc networks," in *Proc. 44th Allerton Conf. on Communication, Control, and Computing*, Monticello, IL, Oct 2006.

[2] D. Julian, "Erasure networks," in *Proc. IEEE ISIT 2002*, Lausanne, Switzerland, Jul 2002.

[3] D. Julian, "Dependent and Gaussian erasure networks," in *Proc. IEEE ISIT 2003*, Yokohama, Japan, Jul 2003.

[4] T. M. Cover and A. A. El Gamal,"Capacity theorems for the relay channel," *IEEE Trans. Inform. Theory*, vol. IT-25, No.5, pp. 572-584, Sep 1979.

[5] A. Dana, R. Gowaikar, R. Palanki, B. Hassibi, and M. Effros, "Capacity of wireless erasure networks," *IEEE Trans. Inform. Theory*, vol. 52, no. 3, pp. 789-804, Mar 2006.

[6] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, "Network information flow," *IEEE Trans. Inform. Theory*, vol. 46, no. 4, pp. 1204-1216, Jul 2000.

[7] N. Ratnakar and G. Kramer, "The multicast capacity of deterministic relay networks with no interference," *IEEE Trans. Inform. Theory*, vol. 53, no. 6, pp. 2425-2432, June 2006.

[8] B. A. Nazer and M. Gastpar, "Computing over multiple-access channels with connections to wireless network coding," in *Proc. IEEE ISIT 2006*, Seattle, WA, Jul 2006.

[9] S. Bhadra, P. Gupta, and S. Shakkottai, "On network coding for interference networks," in *Proc. IEEE ISIT 2006*, Seattle, WA, Jul 2006.

[10] P. Gupta and P. R. Kumar, “The capacity of wireless networks,” *IEEE Trans. Inform. Theory*, vol. 46, no. 3, pp. 388–404, Mar 2000.

[11] L. Xie and P. R. Kumar, “A network information theory for wireless communication: scaling laws and optimal operation,” *IEEE Trans. Inform. Theory,* vol. 50, no. 5, pp. 748-767, May 2004.

[12] M. Franceschetti, O. Dousse, D. Tse, and P. Thiran, “On the throughput capacity of random wireless networks,” *IEEE Trans. Inform. Theory*, submitted for publication.

[13] T. Cover and J. Thomas, *Elements of Information Theory.* New York: Wiley, 1991.

[14] T. Cormen, C. Leiserson, R. Rivest, and C. Stein, *Introduction to Algorithms, Second Edition.* Cambridge: MIT Press, 2001.

[15] C. Cooper, “On the distribution of rank of a random matrix over a finite field,” *Random Structures and Algorithms*, vol. 17, no. 3-4, pp. 197-212, Nov 2000.