Priority Encoding Transmission*

Andres Albanese Johannes Blömer Jeff Edmonds Michael Luby
Madhu Sudan

Abstract

We introduce a novel approach for sending messages simultaneously over high and low-speed networks. The new method, called Priority Encoding Transmission, allows a user to specify a different priority on each part the message to be transmitted over a lossy heterogeneous medium. Based on the priorities, the sender uses the system to encode the message into packets for transmission. The system ensures a smooth rate of recovery of the message (in order of priority) by receivers as a function of the received amount of the encoding transmitted.

We define the *rate* of a set of priorities, which is a natural measure of how much information about the message the priorities dictate must be contained in each bit of the encoding. We develop systems for implementing any set of priorities with rate one. We also give an information-theoretic proof that there is no system that can implement a set of priorities with rate greater than one.

This work has immediate applications to multi-media and high speed networks applications, especially in applications with multiple receivers with heterogeneous capabilities. The systems show promise of being practical.

*The first four authors work at the International Computer Science Institute, Berkeley, California. The fourth author works at IBM Yorktown Heights, and was involved in this research while visiting the International Computer Science Institute.

1 Introduction

In many multi-media applications, long messages are to be transmitted in real time across a multiple network links. A message is not sent as one unit, but broken into small indivisible packets that are sent through the medium. Some of the packets arrive quickly, some are delayed indefinitely, and others are lost completely. At some point in time, the receiver cannot wait for packets any longer and must recover as much of the original message as possible from the packets received.

In an idealized world, this message could be directly partitioned into packets that would then be sent using a stream media. A *stream media* is a media where no packets are lost and packets are delivered in the same order as sent. The only irregularity the stream media may have is that the rate at which packets are delivered may vary over time. When the receiver is to display an image, the best image possible is reconstructed from the packets delivered up till that time.

However, the actual media through which the packets are sent, e.g., connection-less networks (IP) and connection-oriented networks (ATM), do not yield stream media. In connection-less networks, in addition to delivering packets at rates that may vary over time, the packets don't necessarily arrive in the same order as sent because the media may send packets over a different routes with varying rates. In connection-oriented networks, an arbitrary subsets of packets may be lost or delayed beyond the point where they are useful due to congestion, interference with other users, hardware failures, and etc. We hereafter call such a media *lossy*, emphasizing the point that each packet is either delivered intact or lost. A lossy media is indiscriminating, in the sense that no packet has priority over any other packet. If lossy media is used to send message in a direct way there is no guarantee that a prefix of the message can be recovered, even if most of the packets are delivered.

There are encoding systems which allow recovery of the message from only a subset of the encoding, e.g., forward error correction suggested by [8, McAuley] and empirically evaluated by [4, Biersack]. However, these systems allow only one priority level for the entire message. The systems we introduce in this paper can be thought of as extensions of such systems.

In this paper, we describe the notion of a Priority Encoding Transmission (PET) system. PET can be viewed as a transmission system for a lossy media that simulates a stream media. A user specified priority is associated with each part of the message to be sent, where the priority specifies the amount of the encoding sufficient to decode that part of the message. The encoding/decoding systems are designed to adhere to these priorities, i.e., the system guarantees that a part of the message can be reconstructed depending only on the quantity (determined by the priority of the part), and not the particular subset, of packets received.

[10, Shacham] also suggests methods for sending prioritized messages over networks. However, those methods require computation of channel capacities from the sender to each receiver, which may be impractical for very large networks with capacities that vary quickly because of congestion. Furthermore, this work does not handle packet losses.

Section 2 describes potential applications of the PET system to transmit multicast video images over heterogeneous networks, to encode IP packets for the recovery the header and control information from a partial delivery of ATM cells, and to increase the quality of service (only the packet loss) provided by the network layer to an application.

Section 3 gives a definition of PET systems and it describes properties for deterministic and probabilistic models. A deterministic PET system is described in Section 4, and a probabilistic system is described Section 5. Section 6 defines a geometric measure of information and gives an information-theoretic proof that the rate of any PET system is at most 1.

2 Applications

Present networks use multiple window protocols to retransmit missing information for communicating with multiple receivers. Consequently, the information rate is determined by the worst case receiver. Future information highways will provide an ever wider range of performance due to the proliferation of wide area networks and broadband technology. Information from a sender must be received by all users participating in the multicast session. Furthermore, each user should be free to select among the available transport services and receiving stations.

PET is a new method for sending information messages through an unreliable transmission system to multiple receivers. Using a PET system, the user can collect together a set of related information objects to form a message and assign different priority levels to the objects within the message. PET encodes the message into multiple packets to be transmitted over the unreliable media. A crucial property of PET is that an object can be recovered by receivers based only on its priority level and on the number of packets received, where higher priority objects can be recovered from less packets than lower priority objects.

PET can be applied in several of the network layers of the protocol architecture. This section describes two possible applications: multicasting of video images over heterogeneous networks, and assembly layers in ATM networks.

2.1 Video Multicasting

High quality images may consist of as much as 96 Mbits per image, and images may be sent at the rate of 30 per second. If some of the packets containing the image are delayed or lost, the receiver cannot delay displaying the video image. An unfavorable scheme would be to physically partition the image into small regions and place in each packet the information about a single region. The resulting displayed image could be visually displeasing, consisting of a patchwork of high resolution regions corresponding to received packets intermixed with blank regions corresponding to lost packets.

Using JPEG or MPEG, a discrete cosine transform can be applied to a video image to produce what is hereafter called a message [13, Wallace], [6, Le Gall]. Besides allowing a highly compressed representation of the image, this message has a nice property. Consider ordering the information in the message so that the lowest frequency coefficients come first followed by successively higher frequency coefficients. The nice property is that the quality of the image that can be reconstructed from a prefix of this ordered message improves gracefully as a function of the length of the prefix. A PET system can be used to send this real-time prioritized information over a media that sometimes loses and/or inordinately delays delivery of packets.

2.2 IP and ATM Applications

Broadband communication systems using ATM techniques can carry IP packets. We propose the encoding of multimedia applications using IP packets to guarantee the delivery of highest-priority data or the timely recovery of real-time data when IP packets are lost. Voice and video can be coded into the same packet with different priorities to guarantee the quality of service required by each media. Another possible application is in the ATM adaptation layer to encode IP packets into ATM cells in such a way that the IP packet headers and other control information in the packet are recovered with higher priority when ATM cells are lost.

3 Definitions of PET systems

The basic scenario is that a sender encodes a message to be transmitted over a lossy medium in such a way that a receiver is able to recover parts of the original message from parts of the encoding. Moreover, the sender should be able to assign different priorities to different pieces of the message and the receiver should be able to recover the pieces of the message in order of their priority.

3.1 Definition of a Deterministic PET System

Let M be a message of length m to be sent over a lossy network.

Definition 3.1 (priority function) *The importance of receiving individual message bits is specified by* priority function β *that maps* $[1..m]$ *to integral multiples of the packet size* ℓ*. The intended meaning of* β_i *is that message bit* M_i *should be recoverable from any subset of packets that contain in total at least* β_i *bits. We assume without loss of generality that the priority function is monotonically increasing, i.e.,* $\beta_1 \leq \beta_2 \leq \cdots \leq \beta_m$*. Thus,* β_i *can also be thought of as the number of packet bits needed to recover the first* i *bits of the message.*

Definition 3.2 (PET system) *A PET system with message length* m*, packet size* ℓ*,* n *packets, and encoding length* $e = n\ell$ *consists of the following: An encoding function* E *that maps a message* M *of length* m *onto* n *packets; A decoding function* D *that maps a set of at most* n *packets onto a bit string of length* m*; A priority function* β *that maps* $[1..m]$ *to integral multiples of* ℓ*. The decoding guarantee of the system is that, for all messages* M *of length* m*,* D *decodes the first* i *bits of the message from any subset of packets that contain in total at least* β_i *bits.*

An important information-theoretical measure for a PET system is how much information each bit in the encoding contains about the message.

Definition 3.3 (Rate of a priority function/PET system) *For a function β mapping $[1..m]$ to the natural numbers, the* rate *of β is* $\text{rate}_\beta = \sum_{i\in[1..m]} 1/\beta_i$. *The rate of a PET system is the rate of its priority function.*

Intuitively, for a PET system with priority function β, any β_i bits of the code E must determine the i-th message bit M_i. Hence every bit of the code must contain $1/\beta_i$ bits "about" M_i. Therefore, every bit of the code must contain $\text{rate}_\beta = \sum_{i\in[1..m]} 1/\beta_i$ bits in total "about" the message. However, a single bit of the code can contain at most one bit of information. Hence, it is reasonable to expect that such a system is possible only if $\text{rate}_\beta \leq 1$, and we prove this to be the case in Section 6 (Theorem 6.5). On the other hand, we show in Section 4 (Theorem 4.3) that, for a given priority function β with $\text{rate}_\beta = 1$, a PET system with priority function γ can be constructed such that $\text{rate}_\gamma = 1$ and such that, for all $i \in [1..m]$,

$$\gamma_i \leq (1 + 5/\alpha) \cdot \beta_i.$$

Here, $\alpha \geq 3$ is an adjustable parameter that balances the tradeoff between the closeness of the approximation of γ to β, the total encoding length, and the packet size.

3.2 Definition of a Probabilistic PET System

We only highlight the differences between probabilistic and deterministic PET systems. The main difference is that there is a random string R shared by both the sender and receiver that is used to encode and decode messages. Once the string R has been selected the encoding and decoding is deterministic. There is a strictly decreasing confidence function p that quantifies the probabilistic decoding guarantee. The decoding guarantee is that, for any fixed subset of packets that contain in total at least $(1+\delta)\beta_i$ bits, the probability that the decoding function decodes for all messages M the i-th message bit from this subset of packets is at least $1 - p(\delta)$. The probability is with respect to the uniform distribution on R. We stress that this probability is not over a particular distribution over the messages. For any fixed value of R, an encoding/decoding pair succeeds or fails on certain subsets of packets, independent of the message. The confidence function p should be a rapidly decreasing function like $p(\delta) = \exp(-\delta b)$, for some parameter b.

In later sections we describe theorems for probabilistic PET systems that are analogous to those already described for deterministic PET systems.

4 A PET System

We describe a general method that takes any given priority function β and produces a PET system which has a priority function that closely approximates β. The method works by first partitioning the message into blocks based on the priority function β, and then using the partition to implement a PET system based on erasure codes.

In the first subsection we describe erasure codes. In the second subsection, we assume we have the partitioned message and show how to implement a PET system based on erasure codes. Finally, we describe an algorithm that accepts the description of an arbitrary priority function β and produces a partitioned message. The PET system that results from combining these parts has a priority function which closely approximates β.

4.1 Erasure Codes: A Basic Encoding System

An erasure code is specified by a triple $\langle b, n, w\rangle$, where $n \geq b$. It encodes a message M of length $m = wb$ into a code E of length $e = wn$. Both the message and the code consists of words of length w each. The code has the property that all b words of M can be recovered from any b words of E.

Descriptions of erasure codes can be found, for example, in [3, Berlekamp] and [9, Rabin]. One implementation of erasure codes is the following. The b words of M are viewed as the coefficients of univariate polynomial of degree $b-1$ over $\mathrm{GF}[2^w]$. Call this polynomial G. The j^{th} word of the code consists of the value of the polynomial G evaluated at the field element $j \in \mathrm{GF}[2^w]$. Since G is of degree $b-1$, any b words (together with the indices of the words) uniquely determine G. The message M, i.e., the coefficients of G, can be recovered from any b words by interpolation.

This implementation requires that $n \leq 2^w$, or equivalently that

$$w \geq \log(n). \tag{1}$$

This ensures that there are at least n different elements in the field $\mathrm{GF}[2^w]$ on which to evaluate the polynomial.

4.2 Block Systems

The first step in constructing a PET system given a priority function β is to partition the message into blocks based on β. This first step is described in the next subsection. In this subsection, we show how to implement a PET system given a partition of the message. In this and all subsequent constructions, we ignore small round-off errors.

Definition 4.1 (m-partition) *An m-*partition *consists of a sequence of positive integers*

$$\langle m_1, \ldots, m_d \rangle$$

such that

$$\sum_{j \in [1..d]} m_j = m.$$

Lemma 4.2
Given an m-partition $\langle m_1, \ldots, m_d \rangle$, a PET system with priority function γ can be constructed with the following properties:

(i) For all $j \in [1..d]$, $\gamma_j = dm_j$.

(ii) $\mathrm{rate}_\gamma = 1$.

(iii) The encoding length is $e = \max_{j \in [1..d]}\{\gamma_j\}$.

(iv) The packet size is $\ell = dw$, where $w = \log(e)$ *is the word size.*

Proof of Lemma 4.2: Let $B_1, \ldots, B_d$ be the blocks of M, and thus the length of B_j is m_j. The basic idea is to use a separate erasure code for each of the d blocks of the message. The j^{th} erasure code is used to encode B_j into a code E_j consisting of n words, each of length w, where n and w are fixed below. Thus, B_j consists of $b_j = m_j/w$ words. The entire encoding E consists of n packets of size $\ell = wd$ each, where the k^{th} packet consists of the concatenation, for $j \in [1..d]$, of the k^{th} word from the code E_j. Thus, the code length is $e = \ell n$. The decoding works in the obvious way.

Since we use an erasure code for each block, all bits in the same block have the same priority. Any b_j words of the code E_j suffice to recover block B_j. Since there is one such word in each packet, it follows that b_j packets of E are sufficient to recover B_j. Thus, the priority of all bits in block B_j is

$$\gamma_j = \ell b_j = dm_j. \tag{2}$$

This proves item (i). Note that

$$\mathrm{rate}_\gamma = \sum_{j \in [1..d]} m_j/\gamma_j = 1. \tag{3}$$

This proves item (ii). To ensure that the entire message can be recovered from all the packets, we need

$$n \geq \max_{j \in [1..d]}\{b_j\}. \tag{4}$$

With the number of packets set to make this an equality, the total encoding length is $e = \ell n = \max_{j \in [1..d]}\{\gamma_j\}$. This proves item (iii). To use the implementation of erasure codes described in Section 4.1, we need the word length w to be at least $\log(n)$ from Inequality (1). Thus, we can set $w = \log(e) \geq \log(n)$. This proves item (iv). ∎

In the system described above, each packet needs to contain an identifier which is interpreted as the field element value at which the d message blocks considered as polynomials are evaluated. Although this is part of the packet, we did not include it in the packet size. The overhead per packet because of this is at most w bits.

4.3 Partitioning a Message

In this subsection, we show how to construct an m-partition based on a message length m, a priority function β with $\text{rate}_\beta = 1$, and a parameter $\alpha \geq 3$. When this m-partition is used to construct a PET system as described in Lemma 4.2, the priority function γ of the system is a close approximation of β. The parameter α is used to balance the tradeoff between the closeness of the approximation of γ to β, the total encoding length, and the packet size.

We first give the main theorem (Theorem 4.3) and then give the partitioning lemma (Lemma 4.4) upon which the theorem is based.

Theorem 4.3 *Let β be a priority function with* $\text{rate}_\beta = 1$ *for messages of length m. There is an efficient algorithm that, on input β, m, and a value $\alpha \geq 3$, produces a PET system with priority function γ with the following properties:*

(i) The encoding length e is at most $3\alpha m$.

(ii) The packet size ℓ is at most $\alpha^2 \log^2(3\alpha m)$.

(iii) For all $i \in [1..m]$, $\gamma_i \leq (1 + 5/\alpha) \cdot \beta_i$.

(iv) $\text{rate}_\gamma = 1$.

Proof of Theorem 4.3: The first step is to partition M based on β, m and α as described in Lemma 4.4. We then use Lemma 4.2 to get the PET system. It is easy to verify that it has claimed properties. ∎

Note that as α increases the closeness of the approximation of γ to β improves whereas the encoding length and the packet size both increase. Little attempt is made in the theorem to optimize the minimal value $\alpha = 3$ for which the result holds or the other absolute constants associated with α.

In practice, the transmission medium dictates the size ℓ_{medium} of a packet. If $\ell < \ell_{\text{medium}}$ then the packet size of the PET system can be easily scaled up to ℓ_{medium}.

Lemma 4.4 *Let β be a priority function with* $\text{rate}_\beta = 1$ *for messages of length m. There is an efficient algorithm that, on input β, m, and a value $\alpha \geq 3$, produces an m-partition $\langle m_1, \ldots, m_d \rangle$ that satisfies the following properties:*

(i) $\max_{j \in [1..d]}\{m_j\} \leq 3\alpha m/d$.

(ii) For all indices i in the j^{th} block B_j, $m_j \leq (1 + 5/\alpha) \cdot \beta_i/d$.

(iii) $d = \alpha^2 \log(2\alpha m)$.

Proof of Lemma 4.4: To satisfy part (i) of the lemma, we first introduce an intermediate priority function β'. For all $i \in [1..m]$, let $\beta'_i = c' \cdot \min\{\beta_i, \alpha m\}$, where $1 \leq c' \leq 1 + 1/\alpha$ is a small normalizing constant that makes $\text{rate}_{\beta'} = 1$. Note that for all i,

$$\beta'_i \leq c'\beta_i. \tag{5}$$

Furthermore,

$$\beta'_m \leq c'\alpha m \leq 2\alpha m. \tag{6}$$

We set $d = \alpha^2 \log(2\alpha m)$, which satisfies part (iii) of the lemma. We also set two intermediate parameters $c = 1 + 1/\alpha$ and $k' = \alpha \log(2\alpha m)$ and define $k = d - k'$. These parameters are set so as to satisfy the following:

$$c^{k'} \geq 2\alpha m. \tag{7}$$

$$k = d(1 - 1/\alpha). \tag{8}$$

Inequality (7) holds for any $\alpha \geq 2$. Based on these settings of parameters, we then iteratively cut the message into blocks $B_1, \ldots, B_d$ as follows, where i_j denotes the first index in block B_j and m_j is the length of B_j. Suppose that indices $i_1, \ldots, i_j$ have already been set. Then i_{j+1} is set to be the smallest index greater than i_j that satisfies at least one of the following two conditions:

Condition 1: $\beta'_{i_{j+1}} > c\beta'_{i_j}$.

Condition 2: $\sum_{i=i_j}^{i_{j+1}} 1/\beta'_i > 1/k$.

We first verify that the entire message is completely partitioned into the d blocks. From $\beta'_1 \geq 1$, from Inequality (6), and from Inequality (7), Condition (1) can happen at most k' times. Because $\text{rate}_{\beta'} = 1$, Condition (2) can happen at most k times. Thus, the total number of blocks used to partition the entire message is at most $k + k' = d$.

We now derive an upper bound on the number m_j of bits in block B_j. By Condition (1), for all $i \in B_j$, $\beta'_i \leq c\beta'_{i_j}$. By considering the worst case, i.e., when this is equality for all $i \in B_j$, and using Condition (2), it follows that $m_j \leq c\beta'_{i_j}/k$. From this and from Inequality (5), it follows that, for all indices i in block B_j,

$$\begin{aligned} m_j &\leq c \cdot (d/k) \cdot \beta'_{i_j}/d \leq c \cdot (d/k) \cdot \beta'_i/d \\ &\leq c' \cdot c \cdot (d/k) \cdot \beta_i/d. \end{aligned} \tag{9}$$

Note that $c' \leq 1 + 1/\alpha$, $c = 1 + 1/\alpha$, and $d/k \leq 1/(1 - 1/\alpha)$, It can be easily shown that, for all $\alpha \geq 3$,

$$(1 + 1/\alpha)^2/(1 - 1/\alpha) \leq 1 + 5/\alpha. \tag{10}$$

From this inequality, and from Inequality (9), it can be seen that for all indices i in B_j, $m_j \leq (1 + 5/\alpha) \cdot \beta_i/d$. This satisfies part (ii) of the lemma.

Because $m_j \leq c \cdot (d/k) \cdot \beta'_{i_j}/d$ from Inequality (9), and because $\beta'_{i_j} \leq \beta'_m \leq c'\alpha m$ from Inequality (6), it follows that

$$\max_{j \in [1..d]} \{m_j\} \leq c' \cdot c \cdot (d/k) \cdot \alpha m/d.$$

Thus, from Inequality (10), and because $\alpha \geq 3$ implies that $1 + 5/\alpha \leq 3$, it follows that $\max_{j \in [1..d]}\{m_j\} \leq 3\alpha m$, proving part (i) of the lemma. ◻

5 A probabilistic PET system

Erasure codes, as described in Section 4.1 and used in the deterministic PET system in Section 4.2, are specified by a triple $\langle b, n, w \rangle$. Recovering the message M of length $m = wb$ requires the interpolation of polynomial of degree $b - 1$ over GF$[2^w]$, and for large values of b and w this may turn out to be impractical. In this section we briefly describe a probabilistic PET system that is based on probabilistic erasure codes. These codes allow a smaller word size and smaller degree polynomials. The idea is to break the message into fixed size pieces, called *bundles*, of $t < b$ words each. The encoding is probabilistic in the sense that given any $(1 + \delta)b$ words of the code a bundle of the message can be decoded with some probability depending on δ. However, the decoding of a bundle involves only the interpolation of a degree $t - 1$ polynomial over GF$[2^w]$.

A straightforward method to do this is to choose the encoding such that for all $j \in [1..n]$ with probability $1/b$ the j^{th} word of the encoding is the i^{th} word of the message, i.e., $t = 1$.

This method contains some ideas and features of the probabilistic erasure code eventually developed. For example, the expected number of encoding words necessary to get the i^{th} message word is b. However, it has several flaws including the following two related drawbacks. With probability $(1 - 1/b)^{(1+\delta)b} \approx \exp(-(1 + \delta))$ more than $(1 + \delta)b$ encoding words are necessary to get the i^{th} message word. Hence the variance is high and the probability of not getting the i^{th} message word drops to $1/b$ only after $\Omega(b \log b)$ encoding words have been received.

Secondly, the case that all message words are received corresponds exactly to the classical coupon collecting problem. Hence the expected number of encoding words necessary to receive all message words is $\Omega(b \log b)$, i.e., the encoding must have length $\Omega(m \log m)$ instead of linear length.

To overcome these problems we combine this method with erasure codes. Let $t > 1$ be the size of a bundle and let $U_1, \ldots, U_{b/t}$, be the partition of the message M into bundles. A bundle U_i is viewed as the coefficients of a polynomial G_i of degree $t-1$ over $\mathrm{GF}[2^w]$. The r^{th} word of the code E is chosen as follows. An index $i_r \in \{1, \ldots, b/t\}$ and an element $s_r \in \mathrm{GF}[2^w]$ are chosen uniformly at random, and code word E_r is set to $G_{i_r}(s_r)$.

Given a set of $(1+\delta)b \le n$ words of E, a fixed bundle U_i can be recovered from this set if it contains the value of G_i at t different elements of $\mathrm{GF}[2^w]$. Let $\mu = 2^w/t \ge 8$. Using Chernoff-bounds (e.g., see [2, Alon, Spencer]), the following lemma is easy to prove.

Lemma 5.1 *Let $\delta > 2/(\mu-1)$. For all messages M, any fixed bundle U_i and any fixed set of $(1+\delta)b \le n$ words of the encoding, with probability at least $1-\exp(-\delta^2 t/8)$ the bundle U_i of M can be recovered from these code words. The probability is over the random choices of the bundles and field elements.*

Choosing for each code word a bundle and a field element uniformly at random requires a lot of truly random bits. Using the construction of (γ, k)-independent random variables given in [1, AGHP], and using the analysis given in [5, EGLNV], the number of random bits required by the probabilistic erasure codes can be reduced significantly. The details of this method are described in the full paper.

Replacing deterministic erasure codes by probabilistic erasure codes, we obtain a probabilistic version of Lemma 4.2. Combining this with Lemma 4.4, we obtain the following probabilistic version of Theorem 4.3.

Theorem 5.2 *Let β be a priority function with $\mathrm{rate}_\beta = 1$ for messages of length m. There is an efficient algorithm that, on input β, m, a pair of integers $\langle w, t\rangle$, satisfying $\mu = 2^w/t \ge 8$, and values $\alpha \ge 3$, $\epsilon > 0$, produces a PET system with priority function γ with the following properties:*

(i) *The encoding length e is at most $3(1+\epsilon)\alpha m$.*

(ii) *The packet size ℓ is at most $\alpha^2 \log(2\alpha m) w$.*

(iii) *For all $i \in [1..m]$, $\gamma_i \le (1+5/\alpha) \cdot \beta_i$.*

(iv) $\mathrm{rate}_\gamma = 1$.

(v) *For any $\delta \in (2/(\mu-1), \epsilon]$ the confidence function p satisfies $p(\delta) \le \exp(-\delta^2 t/8)$.*

As in Theorem 4.3 the parameter α balances the tradeoff between the closeness of the priority function γ to β, the total encoding length, and the packet size. The parameter ϵ balances the tradeoff between the total encoding length and the range for which a precise statement about the confidence function can be made. The pair of parameters $\langle w, t\rangle$ balance a tradeoff between the efficiency of the encoding and decoding processes and how fast the confidence function p decreases with increasing δ.

6 Inherent Limits of PET Systems

In this section we describe some of the inherent limitations of PET systems. We start with a sketch of the bound on the rate of any PET system. Let $\mathcal{E} \subseteq \{E(M) \mid M \in \{0,1\}^m\}$ denote an arbitrary subset of the encodings sent by the system. A PET encoding $E = \langle E_1, \ldots, E_n\rangle \in \left[\{0,1\}^\ell\right]^n$ has length e and is broken into $n = e/\ell$ packets of size ℓ. Hence, it can be viewed as a point in the n dimensional lattice $\mathbf{Z}^n$ where each coordinate lies between 0 and $2^\ell - 1$. The set of encodings $\mathcal{E}$ can be viewed as a set of such points.

Definition 6.1 *For each $q \in [1..n]$ we define the following measure of $\mathcal{E}$:*

$$V_q(\mathcal{E}) = \left(\prod_{\vec{t} \in \binom{n}{q}} |\mathcal{E}_{\vec{t}}| \right)^{\frac{1}{\binom{n-1}{q-1}}}.$$

Here $\mathcal{E}_{\vec{t}}$ denotes the projection of $\mathcal{E}$ onto the dimensions $\vec{t}$, where $\vec{t} \in \binom{n}{q}$ is any q of the n dimensions.

Lemma 6.2

$$a^n \geq V_1(\mathcal{E}) \geq V_2(\mathcal{E}) \geq \ldots \geq V_n(\mathcal{E}) = |\mathcal{E}|.$$

This lemma is an easy generalization of a result given in [7, Loomis, Whitney].

Definition 6.3 *Let $\mathcal{E} = \mathcal{E}^0 \dot{\cup} \mathcal{E}^1$ and let $b = 0$ on $\mathcal{E}^0$ and $b = 1$ on $\mathcal{E}^1$. For $\vec{t} \in \binom{n}{q}$, we say that the coordinates $\vec{t}$* determines *b if $\mathcal{E}^0_{\vec{t}} \dot{\cup} \mathcal{E}^1_{\vec{t}}$ is a partition of $\mathcal{E}_{\vec{t}}$.*

Lemma 6.4 *If every q coordinates determines the bit b, then there is a setting of $b \in \{0,1\}$ for which*

$$V_q(\mathcal{E}^b) \leq 2^{-n/q} \cdot V_q(\mathcal{E}).$$

We now have the necessary tools for proving the lower bound.

Theorem 6.5 *For any PET with priority function β, $\text{rate}_\beta \leq 1$.*

Proof of Theorem 6.5: Let

$$\mathcal{E} \subseteq \{E(M) \mid M \in \{0,1\}^m\}$$

denote the set of encodings associated with messages sent by the PET system. For all $i \in [1..m]$, let $q_i = \beta_i/\ell$ denote the number of packets needed to determine the message bit M_i in the PET system. For any sequence of $\langle b_1 \cdots b_i \rangle$ of i bits, let $\mathcal{E}^{b_1 \cdots b_i}$ denote the set of encodings possible subject to $M_1 = b_1, \ldots, M_i = b_i$. Let $\vec{t} \in \binom{n}{q_i}$ denote any q_i of the n packets. Let $\mathcal{E}^{b_1 \cdots b_{i-1} 0} \dot{\cup} \mathcal{E}^{b_1 \cdots b_{i-1} 1}$ be the partition of $\mathcal{E}^{b_1 \cdots b_{i-1}}$ based on whether M_i is 0 or 1. Since the value of M_i is determined by the values of the packets in $\vec{t}$, it follows that $\mathcal{E}^{b_1 \cdots b_{i-1} 0}_{\vec{t}} \dot{\cup} \mathcal{E}^{b_1 \cdots b_{i-1} 1}_{\vec{t}}$ is a partition of $\mathcal{E}^{b_1 \cdots b_{i-1}}_{\vec{t}}$.

Applying Lemma 6.2 followed by Lemma 6.4 in sequence m times, and using $q_1 \leq q_2 \leq \ldots \leq q_m$ and $n/q_i = e/\beta_i$, it follows that there is a setting $\langle b_1 \cdots b_m \rangle$ for the message M such that

$$\begin{aligned}
2^e &\geq V_{q_1}(\mathcal{E}) \\
&\geq 2^{e/\beta_1} \cdot V_{q_1}(\mathcal{E}^{b_1}) \\
&\geq 2^{e/\beta_1} \cdot V_{q_2}(\mathcal{E}^{b_1}) \\
&\geq 2^{e/\beta_2} 2^{e/\beta_1} \cdot V_{q_2}(\mathcal{E}^{b_1 b_2}) \\
&\vdots \\
&\geq 2^{\sum_{i \in [1..m-1]} e/\beta_i} \cdot V_{q_m}(\mathcal{E}^{b_1 b_2 \ldots b_{m-1}}) \\
&\geq 2^{\sum_{i \in [1..m]} e/\beta_i} \cdot V_{q_m}(\mathcal{E}^{b_1 b_2 \ldots b_m}) \\
&\geq 2^{\sum_{i \in [1..m]} e/\beta_i} \cdot |\mathcal{E}^{b_1 b_2 \ldots b_m}|
\end{aligned}$$

Note that $|\mathcal{E}^{b_1 b_2 \ldots b_m}| \geq 1$, because there is an encoding $E \in \mathcal{E}$ sent when the message is fixed to $M = \langle b_1 b_2 \ldots b_m \rangle$. This gives

$$e \geq \sum_{i \in [1..m]} e/\beta_i = e \cdot \text{rate}_\beta$$

□

The lower bound for the probabilistic case and its proof are similar.

Theorem 6.6 *For any probabilistic PET with priority function β and failure probability p, $\text{rate}_\beta \leq 1/(1-p)$.*

6.1 A Lower Bound on the Packet Size

Theorem 6.7 *Consider a system that encodes messages of length b into codes of length e with packet size ℓ, so that any β/ℓ of the packets determines the entire message. Then $\ell \geq \log(e/\beta) - O((e/\beta)2^{-b})$.*

There are two corollaries of Theorem 6.7. For erasure codes, $\beta = b$ and $\ell = w$, because any b/ℓ of the packets determines the entire message and the word length is w. Therefore, for any erasure code (not just those using polynomials), the word length w must be at least $\log(e/b) - O((e/b)2^{-b})$.

A reasonable priority function β has the property that $e^{1-\epsilon}$ bits of the code determines at least $\log(e)$ bits of the message, i.e. $\beta_{\log(e)} \leq e^{1-\epsilon}$. The second corollary is that a PET system with a reasonable priority function β requires packet size at least $\Omega(\log(e))$. Setting $b = \log(e)$ and $\beta = e^{1-\epsilon}$ in the lower bound yields a packet size of at least

$$\ell \geq \log(e/\beta) - O((e/\beta)2^{-b}) \geq \epsilon \log(e) - 1.$$

7 Conclusions

Algorithms based on erasure codes have been described for implementing a PET system for arbitrary priority functions. PET is a new approach that requires a feasibility demonstration, benchmark evaluation for different software/firmware platforms, and opens the search for fast implementations of the algorithm in software and firmware.

References

[1] N. Alon, O. Goldreich, J. Hastad, R. Peralta, *Simple constructions of almost k-wise independent random variables*, Random Structures and Algorithms, 3(3) (1992), pp. 289-304.

[2] N. Alon, J. H. Spencer, *The probabilistic method*, John Wiley & Sons, Inc., New York, 1992.

[3] E. R. Berlekamp, *Algebraic Coding Theory*, McGrawhill, New York, 1968.

[4] E. W. Biersack, *Performance evaluation of forward error correction in ATM networks*, in Proceedings of SIGCOMM'92, Baltimore, 1992.

[5] G. Even, O. Goldreich, M. Luby, N. Nisan, B. Veličković, *Approximations of general independent distributions*, in Proc. 24^{th} Symposium on Theory of Computing (STOC), 1992, pp. 10-16.

[6] D. Le Gall, *MPEG: A Video Compression Standard for Multimedia Applications*, CACM, Vol. 34, No. 4, April 1991, pp. 47-58.

[7] L. H. Loomis, H. Whitney, *An inequality related to the isoperimetric inequality*, Bulletin of the American Mathematical Society, 55(7) (1949), pp. 961-962.

[8] A. J. McAuley, *Reliable broadband communication using a burst erasure correcting code*, in Proceedings SIGCOMM'90, Philadelphia, 1990.

[9] M. O. Rabin, *Efficient Dispersal of Information for Security, Load Balancing, and Fault Tolerance*, J. ACM, Vol. 36, No. 2, April 1989, pp. 335-348.

[10] N. Shacham, *Multicast Routing of Hierarchical Data,* Proceedings of ICC'92, Chicago, 1992.

[13] G. K. Wallace, *The JPEG Still Picture Compression Standard*, CACM, Vol. 34, No. 4, April 1991, pp. 30-44.