The Anytime Reliability of Constrained Packet Erasure Channels with Feedback

Anant Sahai
Department of Electrical Engineering and Computer Science
University of California
Berkeley, CA 94720-1770
sahai@eecs.berkeley.edu

Qing Xu
Department of Mechanical Engineering
University of California
Berkeley, CA 94720-1770
qingxu@me.berkeley.edu

Abstract

We study the anytime reliability of packet erasure channels where the erasures are iid and known to both the transmitter and receiver. The packets are allowed to carry a variable number of bits within them although only one packet may be transmitted at any given time step. For any moment constraint on the size of the packets, we show that the anytime reliability is constant at all rates up to the Shannon capacity of the channel and that this constant is essentially the logarithm of the probability of erasure. For cases where there is both a first moment constraint and a constraint on the maximum number of bits that a packet can carry, we show that the optimum anytime reliability is determined primarily by the peak-packet size, but drops abruptly to zero above the Shannon capacity. In order to show achievability of these reliabilities, we give schemes involving FIFO queues where the server can adjust its service rate based on the number of bits awaiting transmission. When many bits are waiting, larger packets are used.

I. Introduction

Packet erasure channels (Figure 1) are a natural model for communication systems where the underlying physical layer communication is already partially specified. They emerge naturally in both wired and wireless communication systems where packet-oriented designs abound, often from an implicit desire to achieve statistical multiplexing among many different users. Erasures generally model two types of events: an unfortunate noise sequence that the underlying error correcting code could not correct[1] or "collisions" at either an intermediate node in a network (where it leads to a packet drop) or over the shared communication medium [13][2] (where it leads to an undecodeable reception of the packet).[2] Knowledge of the erasure comes back to the transmitter through either an acknowledgment packet or by the transmitter observing the packet getting mangled over the link.

Variable size packets can occur in cases where the transmitter has a choice of modulation to use in transmitting the packet. For example, as shown in Figure 2, we can encode 1, 2, or 4 bits in one channel symbol by using different constellations. In systems with access to wide bandwidths, a multi-tone transmitter could use more frequencies to carry a larger amount of data within the same packet transmission time.[3]

Since they are a useful abstraction that is actually supported by today's communication infrastructure, it is natural to consider using such packet erasure channels within control systems. [12] [11] In most prior work [4], the issue of rate is ignored by considering packets to be infinite precision or

[1] We will not consider the case of uncorrected and undetected errors in this paper. That is studied from first principles for communication over an underlying AWGN+erasure channel with feedback in [15] and [10].

[2] Generally, interference is reduced by Medium Access Control (MAC) protocols (e.g. the 802.11 MAC protocol [1]), which controls the transmitters so that they are not likely to transmit at the same time. However, in our simplified model we can view a denial by the MAC layer as an erasure since the packet was unable to be sent at the requested time.

[3] The case where the packet is simply made longer in time is not covered by the analysis given in this paper since we assume that the time between allowed packet transmissions is always constant.

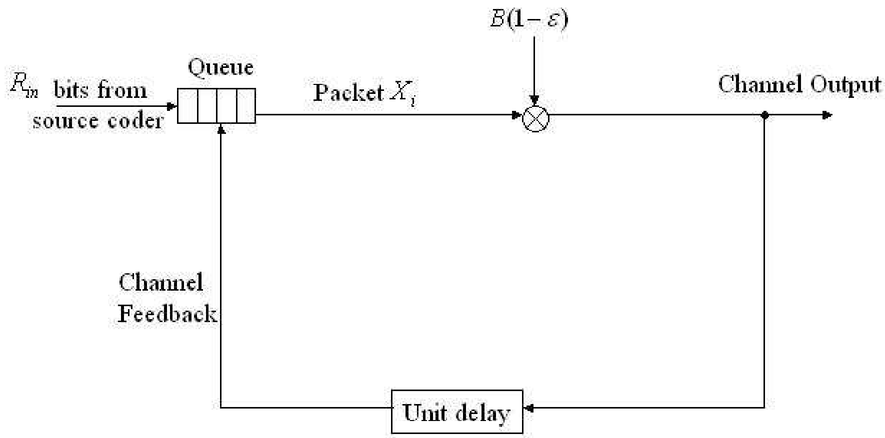

Fig. 1. Packet erasure channel with feedback, fed by a queue

at least to have a high enough rate that the dynamic range is not an issue. But if we really want to leverage a packet-oriented communication for the control of an unstable dynamic system, then [6], [7], and [8] tell us that we require enough anytime capacity to be able to meet our performance objectives.[4]

The anytime capacity of a channel relates the bit error with the delay in a system where we require every bit to get through eventually. We review the definition of anytime capacity: [7]

Definition 1: $C_{anytime}(\alpha)$, the α-anytime capacity, is the max rate at which the channel can be used to communicate with a bit error probability that drops with delay exponentially at a rate of α.

$$C_{anytime}(\alpha) = sup\{R|\exists \mathscr{E}^{\mathscr{R}}, K > 0, \forall N, \exists \mathscr{D}^{\mathscr{R}}_{\mathscr{N}}, P_{error}(\mathscr{E}^{\mathscr{R}}, \mathscr{D}^{\mathscr{R}}_{\mathscr{N}}) < K2^{-\alpha N}\} \quad (1)$$

In above definition $\mathscr{E}^{\mathscr{R}}$ is the anytime encoder, $\mathscr{D}^{\mathscr{R}}_{\mathscr{N}}$ is the decoder, and N is the delay that a bit experiences in units of channel uses. The parameter α is called the *anytime reliability.*[5] Since the probability of error on every single bit goes to zero with increasing delay, it is clear that the anytime capacity is always less than or equal to the classical Shannon capacity.

Fundamentally, what we have is a region of achievable (α, R) pairs — the region between the α axis and the anytime capacity curve. Whether we choose to look at maximizing R as a function of α or as maximizing α as a function of R is a matter of convenience for the problem at hand. Once we know the anytime capacity we know the anytime achievable region, and vice versa.

In this submission, we study the anytime reliability of packet erasure channels by first reviewing the results for the case of unconstrained packet sizes. There, it is obvious that the optimal transmission strategy is just to transmit all the bits awaiting transmission. Next, we review the case of fixed-size packets. For this, the optimal strategy is a first-in first-out (FIFO) queue with bits being removed from the queue when a packet is successfully received. With these results in hand, we consider what happens if we impose a moment constraint on the packet size. We consider a general moment but show that there is something fundamentally different about first-and-higher moments. To show achievability of

[4]Essentially, to hold the η-moment of the state of an unstable plant finite, it is necessary and sufficient for the feedback channel's anytime capacity evaluated corresponding to anytime-reliability $\alpha = \eta \log_2 \lambda$ to be greater than $\log_2 \lambda$ where λ is the unstable eigenvalue of the plant. As such, understanding the anytime capacity under various models is essential if we are able to evaluate engineering questions like whether or not we can have a control system share the communication channel with an existing packet system or whether we need to have a reserved channel for just the control application. Alternatively, sometimes it is possible to schedule the transmissions from nodes in a multiple access network such that interfering nodes share the channel and transmit in turn — e.g. in the wireless token ring protocol[5] — but unless we can evaluate the anytime capacities, it is hard to evaluate whether this additional complexity is justified.

[5]The anytime reliability with noiseless feedback is fundamentally different from both the classical error exponents of Gallager and the exponents for variable delay decoding given by Burnashev[3]. See [9] for more discussion on this.

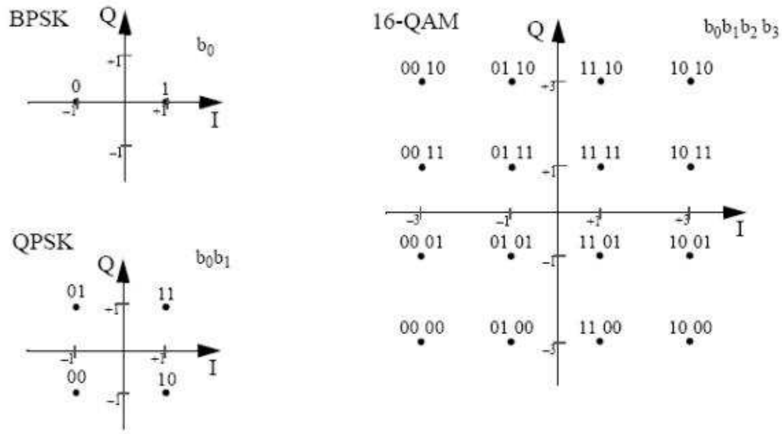

Fig. 2. BPSK, QPSK, and 16-QAM constellation bit encoding

the optimum anytime reliability, we use longer packets when the queue is long.[6] Finally, we consider what happens when we have a peak packet-size constraint in addition to a first moment constraint.

II. Unconstrained Packets

In this section we analyze the anytime capacity of packet erasure channels with no constraint on the packet size. We assume that the channel noise is zero or the transmit power is unbounded so we can encode arbitrarily large numbers of bits in one packet. When the packet is received, the bit error is negligible.

In all of our analyses, we use a FIFO queue to achieve the anytime capacity. Bits enter the queue at a steady rate R_{in}. Whenever there is a non-erasure in the channel in time step i, a packet X_i is transmitted with no error. By feedback, the queue knows whether an erasure happened in the previous time step. The size of the packet X_i is denoted as $S(X_i)$.

When there is no constraint on the packet size, the anytime capacity is easy to obtain. At every time step, the channel attempts to transmit all the bits in the queue. This channel is discussed as the real erasure channel in [7] and it is clear that the encoding strategy is optimal.

Theorem 1: The anytime capacity of the real erasure channel is

$$C_{anytime}(\alpha) = \begin{cases} \infty & \text{if } 0 \le \alpha \le -\log_2 \varepsilon \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

or expressed in terms of the anytime reliability, we have for all rates R:

$$\alpha^*(R) = -\log_2 \varepsilon \tag{3}$$

Proof: There is an bit error only when there is a sequence of consecutive erasures from time t through $t+d$, and the bit error probability is $\varepsilon^d = 2^{d \log_2 \varepsilon}$. Therefore the anytime reliability α must be less than or equal to $-\log_2 \varepsilon$, but the incoming rate can be as high as we would like. ■

III. Fixed packet sizes

When $S(X_i)$ is fixed to be R_{out}, the channel is an R_{out}-bit erasure channel with feedback. The feedback anytime capacity of the binary erasure channel was derived in [6] and [7]. By a simple change of units argument, it is clear that the anytime capacity of the R_{out}-bit channel is the anytime capacity of the BEC (1 bit) channel scaled by R_{out}.

[6]Similar in spirit, Tse et.al. study the statistical multiplexing of multiple time-scale Markov streams [14]. They model each stream with a singularly perturbed Markov-modulated process with some state transitions occurring less frequently than others. They estimate the buffer overflow probabilities in various asymptotic regimes in the buffer size, rare transition probabilities, and the number of streams.

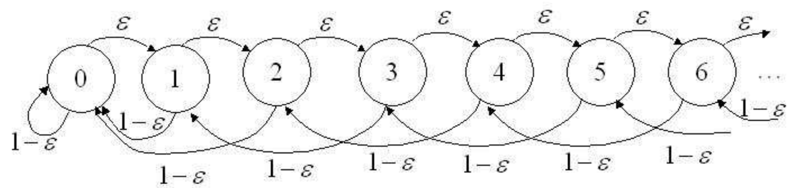

Fig. 3. Markov chain of queue length when the packet size is fixed: $R = 1/3$

Theorem 2: The anytime capacity of a packet erasure channel with packet size constrained to be R_{out} bits, and encoder having access to noiseless feedback, is given by:

$$C_{anytime}(\alpha) = R_{out} \frac{\alpha}{\alpha + \log_2\left(\frac{1-\varepsilon}{1-2^{\alpha}\varepsilon}\right)} \tag{4}$$

where ε is the erasure probability.

Alternatively, we can vary an arbitrary positive parameter η to see that the anytime reliability $\alpha^* = \eta - \log_2(1+\varepsilon(2^{\eta}-1))$ corresponds to the rate of $R_{out}\left(1 - \frac{1}{\eta}\log_2(1+\varepsilon(2^{\eta}-1))\right)$.

Although the optimality proof for (4) involved a control system [7], it should be clear that the optimal anytime encoder is just a FIFO queue with a server that tries to make R_{out} sized packets and send them out. See Figure 3 for an example of the resulting Markov chain.

IV. CONSTRAINT ON THE j-TH MOMENT OF PACKET SIZE

The unconstrained packet size case in section II imposes no constraint while in many applications, the rigid packet-size constraint of section III is too rigid and does not adequately capture the flexibility that might exist. One way of constraining the communication in a flexible way is to impose an average constraint on the j-th moment of the packet size. This ensures that large packets are unlikely and by increasing j, we increase the relative cost of larger packets as compared to smaller ones. $j = 1$ corresponds to the most natural constraint on the average packet size.

To prove achievability, we will use the following type of policy that changes the nature of the packet sizes at a certain critical queue length L_c.

- When the queue length l is smaller than or equal to a critical length L_c, as both measured in units[7] of R_{in}, the system attempts to transmit $R_{out} = R_1$ bits in each packet.
- When the queue length l is larger than L_c, the system transmits a giant packet to reduce the queue length back down to L_c upon successful reception.[8]

To study the anytime capacity of this system we need the following lemma.

Lemma 1: Let the input rate be R_{in} and the system have this queuing rule:

- *Short queue mode*: When the queue length is smaller than L_c, the system transmits packet of size $R_1 > \frac{R_{in}}{(1-\varepsilon)}$
- *Long queue mode*: When the queue length is larger than L_c, the system transmits packets of size larger than R_1.

Then the probability of having a queue of length $L > dR_{in}$ bits is upper-bounded by

$$P(L > dR_{in}) \leq T_1 2^{-\alpha_1^* d}$$

where α_1^* is the feedback anytime reliability of the R_1-sized packet erasure channel corresponding to a rate of R_{in} from Theorem 2 and T_1 is some positive constant.

[7]Measuring the queue in units of R_{in} is not strictly correct since R_{out} need not be an integer multiple of R_{in}. However, dealing with the general case involves extra notation and provides no additional understanding and so we avoid the issue in this submission.

[8]Note we are using potentially unboundedly large packets.

Proof: The key is to realize that for all possible realizations of erasures, the queue length can only be shorter than the queue length if we used packets of fixed size R_1. So we can bound the probability of a large queue length by the corresponding probability for the fixed-size packet system. Whenever an erasure happens, the symbol is resent until it is successfully received. The only possible bit error event is that the bit is still in the queue and has not been sent. The queue must therefore contain at least dR_{in} bits and thus the anytime reliability of this channel therefore bounds the tail distribution of the queue. How to calculate α_1^* graphically from the binary erasure channel anytime curve is illustrated in Figure 4. The constant T_1 is playing the role of the constant K in (1). ■

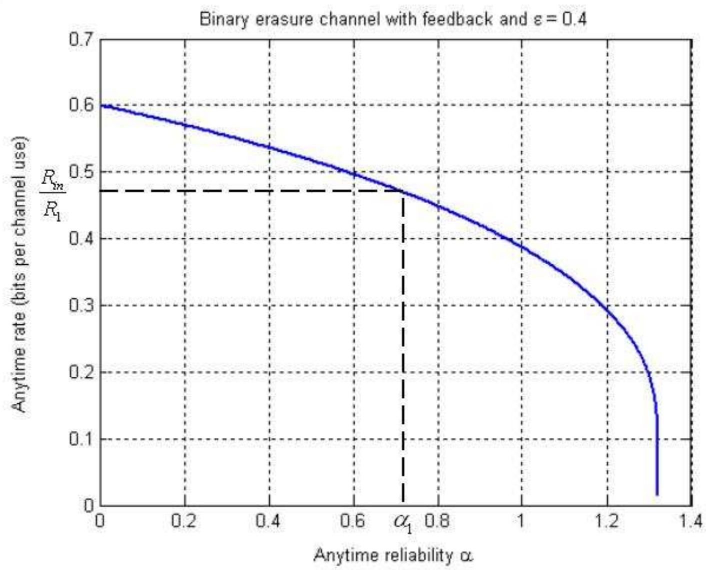

Fig. 4. Finding α_1^* for given R_1 and R_{in} using the anytime capacity curve of binary erasure channel

With Lemma 1, we can prove the following theorems for the anytime capacity of the constrained packet erasure channel. It turns out that the anytime capacity for the case when $j \geq 1$ and $j < 1$ are quite different. We discuss them separately.

Theorem 3: For $j \geq 1$, when the j-th ($j \geq 1$) moment of packet size is constrained $E\{[S(X_i)]^j\} \leq \bar{S}_j$, the anytime capacity of the packet erasure channel is

$$C_{anytime}(\alpha) = \begin{cases} (1-\varepsilon)\sqrt[j]{\bar{S}_j} & \text{if } 0 \leq \alpha \leq -\log_2 \varepsilon \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

Or when viewed in terms of anytime reliability:

$$\alpha^*(R) = \begin{cases} -\log_2 \varepsilon & \text{if } R < (1-\varepsilon)\sqrt[j]{\bar{S}_j} \\ 0 & \text{otherwise} \end{cases}$$

Proof: Clearly we can only do worse for this channel than the unconstrained packet-erasure channel and hence the upper bound on α is an immediate corollary to Theorem 1. Moreover, the anytime capacity can never exceed the Shannon capacity and $(1-\varepsilon)\sqrt[j]{\bar{S}_j}$ is clearly the Shannon capacity of the constrained channel since $\sqrt[j]{\bar{S}_j}$ is the induced constraint on the average number of bits per packet that comes from the moment constraint and the factor of $(1-\varepsilon)$ represents the independent probability that a packet gets through.

For achievability, we use the following queuing rule illustrated in Figure 5.

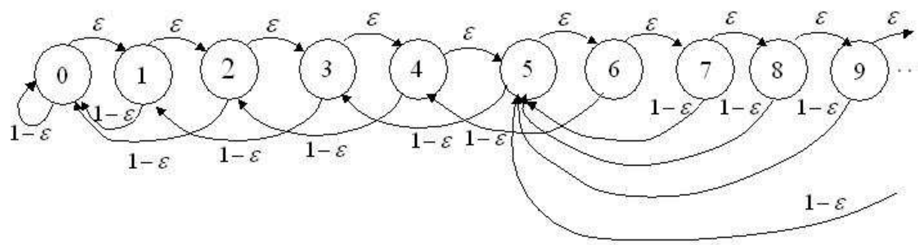

Fig. 5. Markov chain of queue length when $E\{S^j\} < \bar{S_j}$ and $j \geq 1$: L_c=4, $R = 1/3$ when $l < L_c$

- When the queue length l, in units of R_{in}, is smaller than or equal to L_c, also in units of R_{in}, the system takes out $R_{out}(l) = R_1$ bits for every non-erasure. We require $R_1 = \sqrt[j]{\bar{S_j}} - \varepsilon_1$, with ε_1 being an arbitrary small positive number. Notice to make the system stable we need $R_{in} < (1-\varepsilon)R_1$.
- When the queue length l is larger than L_c, the system transmits a giant packet to reduce the queue length to L_c.

We now bound the probability of error. A bit error implies the bit is not transmitted. If the delay is larger than L_c and the bit is still not transmitted, the queue must be longer than $L_c R_{in}$ bits. In such situations, any non-erasure will reduce the queue length to $L_c R_{in}$ bits and so there must have been a string of consecutive erasures since the instant the queue length grew longer than $L_c R_{in}$ bits. Thus $P_{error} < \varepsilon^{(d-L_c)} = (\varepsilon^{-L_c})2^{-(\log_2 \varepsilon)d}$ as delay d goes large. Since we can make ε_1 arbitrarily small, we can obtain the anytime capacity as stated in (5).

We only need to prove that the constraint on the j-th moment of queue length is met. From Lemma 1 we know that the the queue length distribution satisfies $P(l > L_c) = P(lR_{in} > L_c R_{in}) = P(L > L_c R_{in}) \leq T_1 2^{-\alpha_1^* L_c}$, with T_1 and α_1 being positive constants. Hence we have:

$$\begin{aligned} E\{S^j\} &= \sum_{i=0}^{\infty} R_{out}^j(i) P(l=i) \\ &= P(l \leq L_c) R_1^j + P(l > L_c) \sum_{i=L_c+1}^{\infty} (i-L_c)^j R_{in}^j \cdot P(l=i | l > L_c) \\ &\leq R_1^j + T_1 2^{-\alpha_1 L_c} \sum_{i=L_c+1}^{\infty} (i-L_c)^j R_{in}^j \cdot \varepsilon^{i-L_c} \\ &= R_1^j + T_1 R_{in}^j 2^{-\alpha_1 L_c} \sum_{i=1}^{\infty} i^j \varepsilon^i \end{aligned} \tag{6}$$

Since $\sum_{i=1}^{\infty} i^j \varepsilon^i$ converges and is independent of L_c, the second term is an exponentially decreasing function of L_c. Furthermore, R_1^j is selected to be smaller than $\bar{S_j}$. Therefore we can always find L_c large enough such that the constraint on the j-th moment of the packet size is met. ■

In effect, having a constraint on the high moments of the packet size just imposes a bound on the anytime capacity, not the anytime reliability.

Theorem 4: When the j-th ($j < 1$) moment of packet size is constrained by $E\{[S(X_i)]^j\} \leq \bar{S_j}$, the anytime capacity of the packet erasure channel is

$$C_{anytime}(\alpha) = \begin{cases} \infty & \text{if } 0 \leq \alpha \leq -\log_2 \varepsilon \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

or equivalently for anytime reliability:

$$\alpha^*(R) = -\log_2 \varepsilon \tag{8}$$

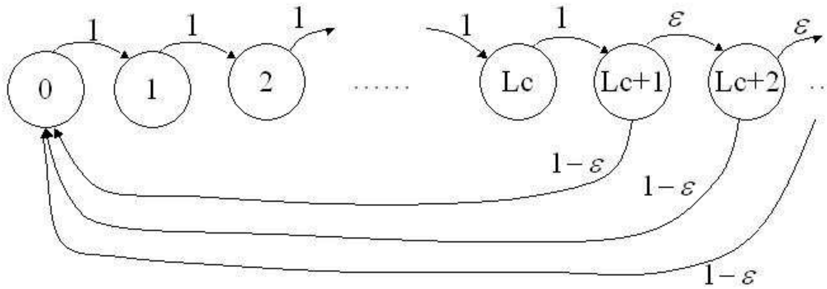

Fig. 6. Markov chain of queue length when $E\{S^j\} < \bar{S}_j$ and $j < 1$, labels in units of R_{in}

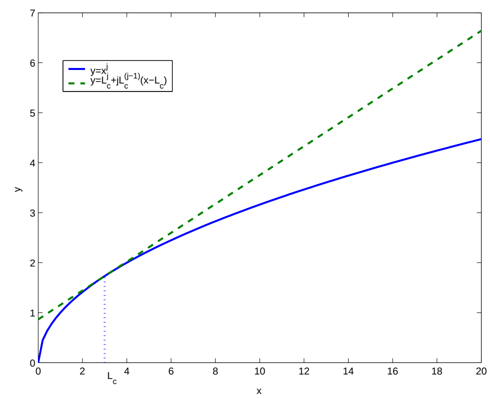

Fig. 7. Sub-linearity of function $y = x^j$ when $j = 0.5$, $L_c = 3$

Proof: We now change the policy to be as illustrated in Figure 6:

- When the queue is shorter than a critical length $L_c R_{in}$ bits, no bits are sent: the size is zero.
- When the queue is longer than $L_c R_{in}$ bits, empty the queue with every non-erasure.

As before, a bit error can happen when bit is held in the queue. When the delay is d (large enough), the probability of bit error is the probability of the queue being longer than dR_{in}, which is upper-bounded as before by $P(L > d) < \varepsilon^{d-L_c}$.

To verify the queue length constraint, we look at stationary distribution of the Markov chain in Figure 6. Obviously, $\pi_0 = \pi_1 = \cdots = \pi_{L_c}$, and $\pi_0 = (1-\varepsilon)\sum_{i=1}^{\infty} \pi_{L_c+i}$. Therefore $\pi_0 = \pi_1 = \cdots = \pi_{L_c} = \frac{1}{L_c+1+(1-\varepsilon)} < \frac{1}{L_c}$, and $P(l > L_c) = \sum_{i=1}^{\infty} \pi_{L_c+i} = \frac{\pi_0}{(1-\varepsilon)} < \frac{1}{L_c(1-\varepsilon)}$.

Hence we have

$$E\{S^j\} < \frac{R_{in}^j}{L_c(1-\varepsilon)} \sum_{i=1}^{\infty} (L_c+i)^j \varepsilon^i \tag{9}$$

As is illustrated in Figure 7, due to the sub-linearity of the function $y = x^j$ when $j < 1$ and $x > 0$, we have

$$\begin{aligned} y &< L_c^j + \left(\frac{dy}{dx}\right)_{x=L_c} \cdot (x - L_c) \\ &= L_c^j + jL_c^{j-1}(x - L_c) \end{aligned}$$

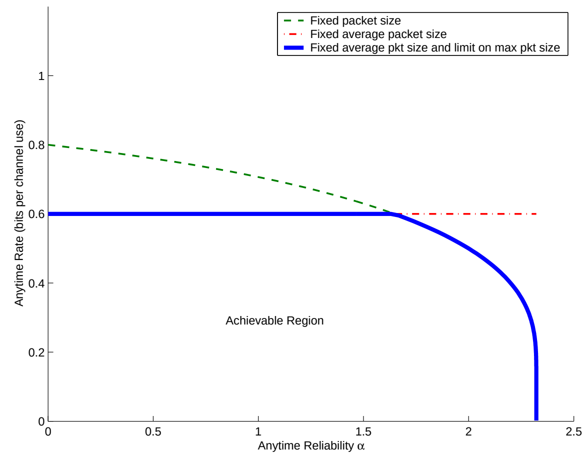

Fig. 8. Anytime achievable region of packet erasure channel with both average and peak packet size constraints

Therefore continuing from equation (9) we have

$$\begin{aligned} E\{S^j\} &< \frac{R_{in}^j}{L_c(1-\varepsilon)}\sum_{i=1}^{\infty}(L_c^j + jL_c^{j-1}i)\varepsilon^i \\ &= \frac{R_{in}^j L_c^j \varepsilon}{L_c(1-\varepsilon)^2} + \frac{R_{in}^j jL_c^{j-1}}{L_c(1-\varepsilon)^3} \\ &= \frac{R_{in}^j \varepsilon}{L_c^{1-j}(1-\varepsilon)^2} + \frac{R_{in}^j j}{L_c^{2-j}(1-\varepsilon)^3} \end{aligned} \quad (10)$$

Since $j < 1$, both terms in (10) decay to zero with L_c. Thus for every R_{in}, we can choose a large enough L_c to satisfy the constraint on the j-th moment of the packet size. ■

So, having a constraint on a moment less than the first moment is as good as having no constraint at all, at least when it comes to the anytime reliability!

V. Constraints on both average and peak packet sizes

Usually, it will be unrealistic to allow unboundedly large packet sizes. Seeing as how having a higher moment constraint seems to only effect the anytime capacity and reliability through the induced constraint on the first moment, it is natural to consider the case where not only is $E\{S\} \leq \overline{S}$, but $S < S_{max}$ as well where $S_{max} > \overline{S}$ to avoid triviality. To study the anytime capacity we need the following lemma which refines Lemma 1 further:

Lemma 2: Let the input rate be R_{in} and the system have this queuing rule:

- *Short queue mode*: When the queue length is not larger than L_c, in units of R_{in}, the system transmits packet of fixed size $R_1 > \frac{R_{in}}{(1-\varepsilon)}$ bits.
- *Long queue mode*: When the queue length is larger than L_c, in units of R_{in}, the system transmits packet of the fixed size $R_2 > R_1$ bits.

Then when queue length $L > L_c R_{in}$ bits, the probability of having a queue of length $L > dR_{in}$ is upper-bounded as

$$P(L > dR_{in}) \leq T_2' 2^{-\alpha_2^* d}$$

where α_2^* is the feedback anytime reliability corresponding to a rate of R_{in} of the R_2-size packet erasure channel from Theorem 2 and T_2' is some positive constant.

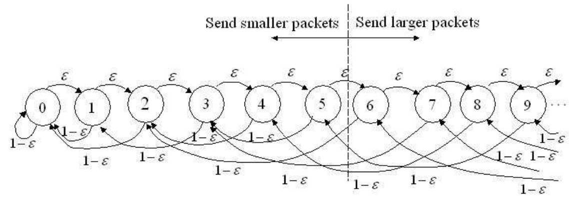

Fig. 9. An example of Markov chain of queue length when both average packet size and peak packet size are constrained: $L_c = 5$, $R_{in} = 1$, $R_{out} = 3$ when $l \le L_c$, and $R_{out} = S_{max} = 4$ when $l > L_c$

Proof: This is similar to the proof of Lemma 1. Since $R_2 > R_1$, we have $\alpha_2^* > \alpha_1^*$. When $L > L_c R_{in}$, we consider the part of queue that is over $L_c R_{in}$. In this part corresponding to R_2 we have anytime exponent α_2^*. Let $i = L - L_c R_{in}$, we can bound $P(i \ge dR_{in} - L_c R_{in})$ by $T_2 2^{-\alpha_2^*(d-L_c)}$.

$$\begin{aligned} P(L > dR_{in}) &= P(L > L_c R_{in}) P(L - L_c R_{in} > dR_{in} - L_c R_{in} | L > L_c) \\ &\le T_1 2^{-\alpha_1^* L_c} T_2 2^{-\alpha_2^*(d-L_c)} \\ &= T_1 T_2 2^{(\alpha_2^* - \alpha_1^*) L_c} 2^{-\alpha_2^* d} \\ &= T_2' 2^{-\alpha_2^* d} \end{aligned}$$

since we can bound the queue length of this queuing system by L_c plus the length of one where only R_2 sized packets are used. ■

Using Lemmas 1 and 2 we can prove the following theorem on the anytime capacity of the packet erasure channel when both the average packet size and the peak packet size are constrained.

Theorem 5: The anytime capacity of the packet erasure channel, when both the average packet size and the peak packet size are constrained, i.e., $E\{S(X_i)\} \le \bar{S}$ and $\max(S(X_i)) < S_{max}$, is the following: $C_{anytime}(\alpha)$

$$= \begin{cases} \min\left((1-\varepsilon)\bar{S}, \frac{\alpha}{\alpha + \log_2\left(\frac{1-\varepsilon}{1-2^{\alpha}\varepsilon}\right)} S_{max} \right) & \text{if } 0 \le \alpha \le -\log_2 \varepsilon \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

for $\eta > 0$.

Expressed in terms of anytime reliability,

$$\alpha^* = \eta - \log_2(1 + \varepsilon(2^\eta - 1))$$

whenever $R = S_{max}\left(1 - \frac{1}{\eta}\log_2(1 + \varepsilon(2^\eta - 1))\right) < (1-\varepsilon)\bar{S}$, and 0 otherwise, where η is a positive parameter.

Proof: The achievable anytime region of this system should be contained in the achievable regions with only one of the constraints: peak or average. Therefore the achievable region lies in their intersection. We achieve every (α, R) point in this intersection by using:

- When the queue length l is smaller than or equal to L_c, both in units of R_{in}, the system uses packets of size $R_{out} = R_1$ bits. We require $R_1 = \bar{S} - \varepsilon_1$, with ε_1 being an arbitrary small positive number. Notice to make the system stable we need $R_{in} < (1-\varepsilon)R_1$.
- When the queue length l is larger than L_c, the system transmits packets with size S_{max}, until the queue is shorter than L_c again.

This queuing rule is illustrated in Figure 9.

Similar to the proof of Theorem 3, a bit error implies the bit is still awaiting transmission. If the delay is larger than L_c and the bit is still not transmitted, the queue must be longer than L_c. From Lemma 2 we know that the queue distribution, and therefore the probability of bit error, is bounded by $T_2' 2^{-\alpha_2^*(R_{in})d}$, where $\alpha_2^*(R_{in})$ is the feedback anytime reliability for the S_{max}-size erasure channel corresponding to anytime rate R_{in}. Since we can make ε_1 arbitrarily small, we have the anytime reliability as stated in the theorem. To check the average packet size constraint:

$$\begin{aligned} E\{S\} &= \sum_{i=0}^{\infty} R_{out}(i) P(l = i) \\ &= R_1 P(l \leq L_c) + S_{max} P(l > L_c) \\ &\leq R_1 + T_2' S_{max} 2^{-\alpha_1 L_c} \end{aligned} \quad (12)$$

Since $R_1 < \bar{S}$ and the second term is an exponential function of L_c, we can always select large enough L_c such that $E\{S\} \leq \bar{S}$. ∎

VI. Conclusion and Future Work

By using variable sized packets, we can substantially increase the anytime reliability of a packet erasure channel. This can be achieved by using bigger packets when we have many bits awaiting transmission. When the constraint on the packet size takes the form of a moment constraint, then the anytime reliability is effectively the same as that of the unconstrained channel for all rates up to capacity. When there is a peak-size constraint as well, then the peak-size constraint dominates the anytime reliability at all rates up to the Shannon capacity of the channel.

In [15] and [10], we extend this style of analysis to the AWGN channel with erasures and noiseless feedback. It turns out there that the doubly-exponential vanishing of the probabilities of error with delay lets us effectively conceptualize the channel as a noiseless packet erasure channel, though there are many technical steps along the way. We believe that this style of analysis can be extended to give us the anytime reliabilities of many Gaussian channels with noiseless feedback and channel state side information available at both the transmitter and receiver.

References

[1] Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications. *IEEE Standard 802.11-1997*, 1997.
[2] N. Abramson. Multiple access in wireless digital networks. *Proceedings of IEEE*, 82:1360–1369, 1994.
[3] M. V. Burnashev. Data transmission over a discrete channel with feedback, random transmission time. *Problemy Perdachi Informatsii*, 12(4):10–30, 1976.
[4] N. Elia. Feedback stabilization in the presence of fading channels. *Proceedings of the American Control Conference*, pages 4438–4443, June 2003.
[5] D. Lee, A. Attias, A. Puri, R. Sengupta, S. Tripakis, and P. Varaiya. A wireless token ring protocol for intelligent transportation systems. *IEEE 4th International Conference on Intelligent Transportation Systems, Oakland, CA*, 2001.
[6] A. Sahai. Evaluating channels for control: Capacity reconsidered. *American Control Conference*, 4:2358–2362, June 2000.
[7] A. Sahai. *Anytime Information Theory*. PhD thesis, Massachusetts Institute of Technology, February 2001.
[8] A. Sahai. The necessity and sufficiency of anytime capacity for scalar control over a noisy communication link. *To appear at the CDC*, 2004.
[9] A. Sahai and T. Şimşek. On the reliability function of discrete memoryless channels with feedback and variable-delay decoding. *To appear in the Information Theory Workshop*, 2004.
[10] A. Sahai and Q. Xu. The anytime capacity of AWGN+erasure channel with feedback. *Also appeared at the Allerton Conference*, 2004.
[11] P. Seiler. *Coordinated Control of Unmanned Aeiral Vehicles*. PhD thesis, University of California, Berkeley, 2001.
[12] P. Seiler and R. Sengupta. Analysis of communication losses in vehicle control problems. *American Control Conference*, 2:1491–1496, June 2001.
[13] F. Tobagi and L. Kleinrock. Packet switching in radio channels: Part I- carrier sense multiple-access modes and their throughput/delay characteristics. *IEEE Trans. Comm.*, COM-23(12):1400–1416, December 1975.
[14] D. Tse, R. Gallager, and J. Tsitsiklis. Statistical multiplexing of multiple time-scale markov streams. *IEEE Journal on Selected Areas in Communications*, 13(6):1028–1038, August 1995.
[15] Q. Xu. The anytime capacity of AWGN+erasure channel with feedback. Master's thesis, University of California at Berkeley, Fall 2004.