Hybrid Packet FEC and Retransmission-based Erasure Recovery Mechanisms (HARQ) for Lossy Networks: Analysis and Design

Vijaynarayanan Subramanian, ECSE Dept, RPI, subrav@rpi.edu
Shivkumar Kalyanaraman, ECSE Dept, RPI, kalyas@rpi.edu
K.K. Ramakrishnan, AT&T Labs Research, kkrama@research.att.com

***Abstract*— With increasing dependence on wireless networks as an integral part of the communication infrastructure, it is critical that data link and transport layer protocols perform reasonably under potentially severe lossy conditions. A key strategy is to use Hybrid ARQ (HARQ) with erasure codes (a.k.a. Forward Error Correction or FEC) sent both proactively and reactively in response to feedback about dynamic loss statistics. A challenge is to design HARQ to satisfy multiple objectives such as high goodput, low latency and negligible residual loss rate. In this paper, we analyze the performance benefits and trade-offs of these reliability strategies (Hybrid ARQ+FEC). We derive expressions for the expected goodput (and overhead in terms of FEC wastage), latency, and residual loss for a given raw erasure loss process (e.g. uniform and bursty loss models). We show how the analysis can be used to explain and provide specialized design guidance for link-layer HARQ that is subject to tight delay constraints and a recently designed transport layer HARQ scheme (called Loss-Tolerant TCP). We validate our analysis by comparing the predictions with values obtained from simulations performed on the link and transport layer HARQ strategies with *ns-2*. We believe that such an analysis could also have value for other adaptive protocols using network coding and incremental redundancy techniques.**

I. Introduction

Diversity techniques are popular at several layers of the networking stack to combat variability in performance and provide reliability over unreliable time-varying channels. When the channel performance is good, we want to send at a higher information rate, and when it is bad, we want to add more reliability and send at a lower information rate. Sophisticated diversity schemes use more diversity modes (e.g. time, frequency, space, etc.) and use each mode efficiently (e.g., through coding and adaptive targeting of overhead).

In this paper we analyze hybrid ARQ/FEC approaches, a class of time-diversity techniques, in the context of erasure channels at the link and transport layers. Our analysis provides guidance in designing these schemes to be efficient. We validate the analysis using specific protocols we have recently designed for the link layer (LL-HARQ/FEC) and transport layer (loss-tolerant TCP (LT-TCP)).

Schemes designed for error protection need to deal with the trade-off between error protection and overhead. To improve end-end performance, we have proposed a set of mechanisms at the transport and link layers that attempt to achieve this balance.

This work was funded by AT&T Labs Research, MIT Lincoln Labs/AFOSR Grant Letter No. 14-S-06-0206 and NSF-ITR 0313095.

We take into account the following measures: goodput (both link level and end-to-end), residual error rate at the link level, latency and wastage of added reliability information. Our analysis examines the trade-off between these measures to guide us in the amount of error protection provided at each layer, and to tune it appropriately for each layer.

In the case of HARQ/FEC involving erasure channels, the trade-off is between goodput (or information rate), residual erasure rate and latency. In particular, indiscriminately adding erasure coding (FEC) reduces goodput unnecessarily, while possibly reducing the latency for block recovery more than needed, and bringing the residual erasure rate down to very low levels (see [1]). On the other hand, if we depend upon ARQ with a high degree of persistence (e.g., at the link layer or end-to-end like TCP) and do not provision any error protection (e.g., FEC), we would have to pay a price in terms of latency under high and/or volatile erasure conditions. The goodput and residual erasure rate will depend upon the specific schemes adopted.

Trade-offs differ when we consider the functionality needed at the link-layer versus the transport layer. If a wireless link is to be used as a building block for multiple applications, and as one link in a multi-hop wireless network (e.g., meshed wireless), there is a need to balance residual erasure rate with goodput and latency. Data-oriented interactive applications such as gaming, ssh, Voip, videoIM will use such networks and cannot tolerate high latencies or poor goodput. Clearly, residual erasure rates need to be reduced with link-level HARQ; otherwise they will accumulate end-to-end to yield unacceptable values. However, attempting to reduce it to zero, especially in the context of multi-hop wireless hops, will require either a severe goodput penalty or a severe latency penalty on an end-to-end basis. It may be desirable to complement a link layer's limited attempts to bring the residual erasure rate down to a reasonable level while limiting the impact on latency with an end-to-end transport layer mechanism that overcomes this residual erasure rate with even a small amount of FEC for error protection.

In this paper, we investigate designs for the link layer to see how effective we can be in reducing the residual erasure rate even if we limit the link level retransmissions to at most one ARQ attempt. The associated goal is to still achieve high goodput, and low residual erasure rates despite bursty losses. Such links can then be used as building blocks for a multi-hop wireless network. The end-to-end cumulative residual erasure rate can be handled by our recent enhancements to TCP to make it more loss-tolerant (LT-TCP) that integrates adaptive HARQ/FEC, and uses ECN to distinguish between congestion and erasures. A detailed overview and performance evaluations of LT-TCP can be found in [10]. This proposed balance of functionality between the link and transport layers delivers a better trade-off than today's solution of using a brittle loss-intolerant TCP which imposes on the link layer the responsibility to effectively achieve a zero residual erasure rate. Interestingly, the designs of both the link and transport layer adaptive HARQ/FEC schemes can be guided by a single analysis developed below.

We validate the results of the analysis through simulation of a link-layer HARQ scheme and a transport layer scheme, loss-tolerant TCP (LT-TCP). We have used the analysis to guide the design of these schemes, though there are special considerations to specialize the adaptive HARQ ideas in the context of TCP.

The rest of this paper is organized as follows. Section II briefly describes hybrid ARQ/FEC for erasure channels. Section III provides an overview of the general design using an abstract model of a lossy channel, analyzes the components of the design and provides insight into the structuring of the various building blocks. Section IV compares the analytical predictions with the simulation results and discusses the comparisons in both the link and transport contexts. Section V summarizes our paper and presents our conclusions.

II. Discussion of Hybrid ARQ/FEC For Erasure Channels

Reliability schemes typically involve provisioning redundancy, and matching it with the diversity resources of the underlying unreliable channel. The simplest scheme is to send duplicates of packets (i.e. using repetition coding): the average probability of packet loss can be reduced from p for a single transmission to p^m for m duplicate transmissions. This open-loop repetition method can be made more efficient by using closed-loop ARQ and selective retransmission, i.e. send duplicate transmissions only when a packet loss has been detected, and feedback precisely identifies the sequence number of the packet needing retransmission. There is an implicit trade-off of latency to achieve better goodput. ARQ works well when the loss probability p is small, as in wired networks, and does not require the source to predict when to provision redundancy.

However, when the loss rate becomes high and bursty, more efficient erasure coding can be used instead of repetition coding in reliability schemes. Maximum Distance Separable (MDS) codes like Reed-Solomon (RS codes) allow the source to add $N - K$ erasure coded parity packets to protect K data packets [6]. As shown in Figure 1, as long as at least any K packets arrive at the receiver, it can recover the original K data packets to achieve an information rate or *goodput* of $\frac{K}{N}$. The *open-loop* coded redundancy transmission (called *proactive FEC* or PFEC) can again be made more efficient by complementing it with closed-loop ARQ feedback, and sending coded redundancy in response (called *reactive FEC* or RFEC). This hybrid of ARQ and FEC is an example of a Type II HARQ scheme (see [2] for an overview of HARQ schemes).

At the transport layer, TCP currently uses ARQ techniques only. However, Rizzo showed the feasibility of RS coding in software at high speeds [7]. Proactive/reactive FEC have been proposed for multi-cast erasure channels at the transport layer [8], [3], [5]. Moreover, new approaches such as LT-codes[4] and Raptor codes[9] (so-called rateless codes) that have very good encoding and decoding efficiencies have made it easier to integrate FEC into transport protocols.

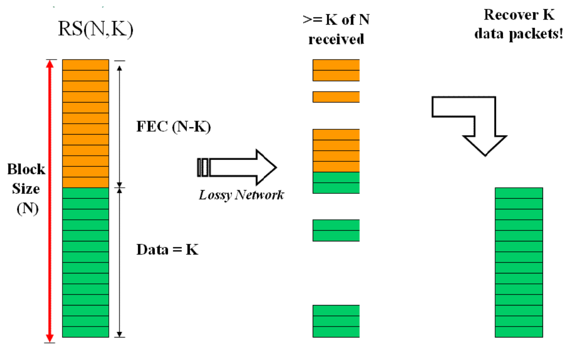

Fig. 1. Operation of Forward Error Correction: RS coding has the property that any FEC unit can stand in for a lost data unit. As long as K out of N units are received, the original data bytes can be recovered.

In addition to coding gains, HARQ can take advantage of the properties of erasure codes. For example, with repetition-coding, the source requires selective ACKs to precisely target retransmission. In contrast, with erasure coded packets, the source only needs to know the number of packets required to reach the threshold of K for decoding. ARQ feedback status can also be simplified because of this *sequence agnostic* property of FEC. The receiver needs to only accumulate enough packets in the first and subsequent ARQ-triggered attempts to reach the threshold of K for decoding. The overall structure of HARQ with PFEC and RFEC is shown in Figure 2.

In this paper, we focus on the following HARQ design questions and develop analytical guidance for them by posing the following questions.

1) How should we set the block size for coding? Given a set of data bytes (at TCP) or a packet at the link-layer, how should this be granulated so as to maximize performance?
2) How should we set PFEC overhead adaptively? How does it depend upon short term loss rate statistics (the mean and variance)? FEC coding can be compared to taking insurance: too much PFEC will lead to wastage and too little will lead to insufficient protection. An inappropriate PFEC protection policy will lead to sub-optimal goodput.
3) What is the need for RFEC? How should we set RFEC overhead adaptively so that residual loss rate *and* RFEC wastage are minimized?
4) What is the balance between PFEC and RFEC overhead provisioning? Should we be more aggressive with PFEC or RFEC? What are the impacts of different choices?

III. Analytical Model:

We will develop our analytical model for a bursty ON-OFF erasure channel. Consider M blocks sent through an ON-OFF erasure channel which has, on average, equal short and alternating ON-OFF periods (shorter compared to the measurement of per-unit loss rate). We refer to this as the ON-OFF model for convenience because it resembles a commonly used 2-state

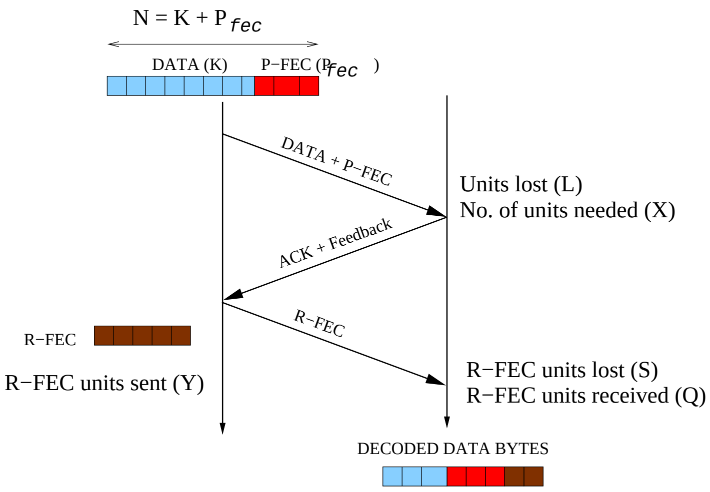

Fig. 2. Time diagram showing the use of Proactive and Reactive FEC. The initial transmission (Phase 1) consists of data and PFEC units for the block. Acks contain feedback that trigger transmission of RFEC segments.

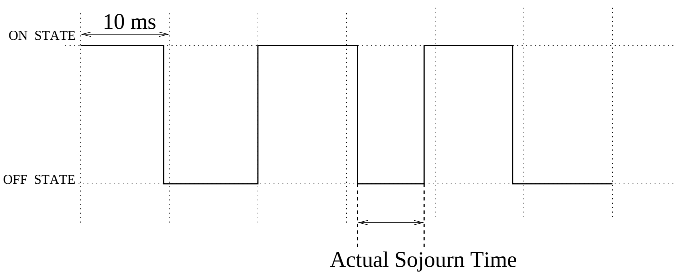

Fig. 3. The behavior of the bursty error process is as shown. The PER in the OFF state is lower than that in the ON state. For the uniform error process the PER in the 2 states is equal. In our simulations, the actual sojourn time in a state is randomized between 9 and 11 ms.

Symbol	Meaning
M	No. of blocks
N	No. of units in a block
$P_{fec} = N - K$	No. of PFEC units
X	No. of units needed after first attempt
$R_{fec} = Y$	No. of RFEC units sent
S	No. of RFEC units lost
L	No. of units lost during first attempt
Q	No. of RFEC units that survive
p	Average Packet Error Rate (PER)
q, r	ON-OFF Model State-Loss probabilities

TABLE I
THE VARIOUS SYMBOLS USED IN THE ANALYSIS AND THEIR MEANINGS ARE AS SHOWN IN THIS TABLE.

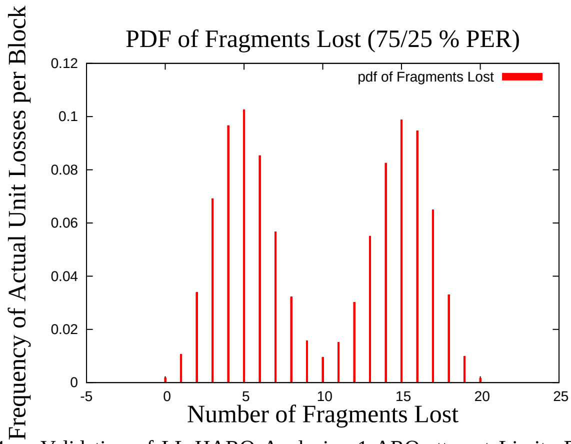

Fig. 4. Validation of LL-HARQ Analysis. 1-ARQ-attempt Limit: Bimodal Distribution of Lost-Units-Per-Block (prior to ARQ attempt). This reflects ON-OFF PER model with Avg PER = 50%, and binomial modes for 75% ON-PER and 25% OFF-PER.

Markov-model. We also assume that each block is sent either during an ON or an OFF period with equal probability. Each block is divided into N units. We assume that the average erasure rate (PER) p refers to the long-run average probability of a unit being lost. During the ON periods, there is a higher erasure rate of $q = 1.5 * p$ (with $p \leq 2/3$, so that $q \leq 1$) and during the OFF period, there is a lower erasure rate $r = 0.5 * p$. For p = 50% these ON-OFF period loss rates are 75% and 25% respectively. This bursty ON-OFF model can be simplified to the uniform per-packet erasure model by setting $q = r = p$. The average ON and OFF *sojourn* periods are the same and are randomized to be between 9 and 11 ms in our simulations. The error model is depicted in Figure 3.

The probability that L, the number of erased units in each block equals i erasures is therefore a bimodal-binomial distribution:

$$P(L = i) = \binom{N}{i} x^i (1-x)^{(N-i)} \text{ where}$$

$$x = \begin{cases} q & \text{if ON period;} \\ r & \text{if OFF period;} \end{cases}$$

A simulation of the bimodal-binomial behavior is shown in Figure 4 which plots the frequencies of actual unit losses per block prior to the ARQ attempt (p = 50%, q = 75% and r = 25%); 10 flows, 1-link case. We capture burstiness as unmeasured bimodality in the underlying loss distribution; the two modes (e.g. 75%-25%) are equidistant from the nominal mode (e.g. 50%); and the distance between the two modes is equal to the original mode value.

The number of blocks that experience i unit erasures is $M * P(L = i)$. The above model makes a simplifying assumption that a block is either in the ON period or OFF period, but does not straddle both. The bursty ON-OFF time series of erasures when viewed as a distribution of erasures per block becomes bimodal-binomial. When simplified as a uniform erasure model, this distribution simplifies to a binomial distribution (illustrated below in Figure 5 for different loss rates):

Recall some useful properties of the binomial distribution. Its mean $m = Np$, Variance = $s^2 = Npq$, coefficient of variation (C.o.V) = $s/m = \sqrt{(pq/N)}$. The scaling of the C.o.V means that the distribution is relatively *tighter* and relatively more concentrated around the mean with increasing N. Also the $\frac{1}{\sqrt{(N)}}$ scaling means that the C.o.V reduction effect is biggest for small N (e.g. as N increases from 1-10). For $Npq >> 1$, the binomial is approximated by a normal distribution near the mean and has $O(e^{-x^2})$ decay of the pdf tails. In other words, the distribution tends to concentrate around the mean and decays very rapidly in

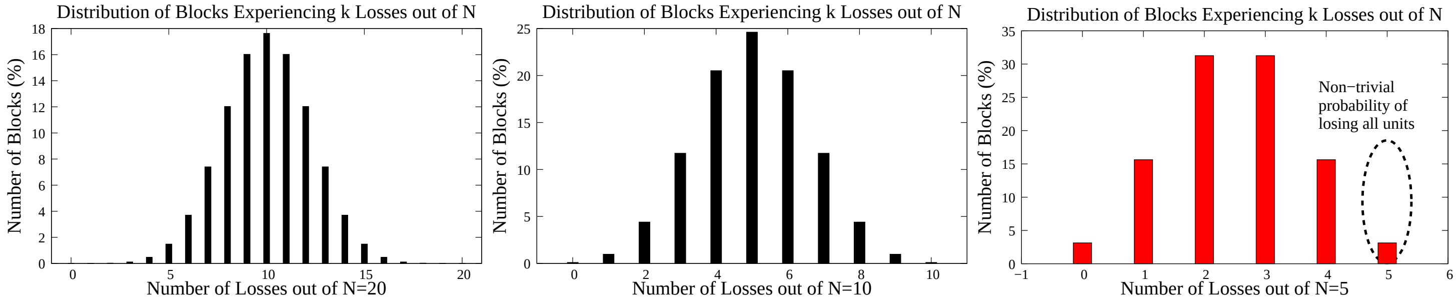

(a) Loss Distribution with N=20

(b) Loss Distribution with N=10

(c) Loss Distribution with N=5

Fig. 5. The binomial distribution of blocks experiencing k losses out of N is shown for different values of N. We note that the key metric i.e. probability of losing all N units is non-trivial for N = 5 and is very small for N=10.

either direction. These properties will have important implications in our modeling of FEC overheads, latency and goodput.

In general, we have time-varying channels whose average erasure rate changes over time. For example we could imagine a doubly stochastic situation where every unit i experiences loss rate with an unknown, random probability p_i (known as Poisson trials). If we assume a quasi-static behavior (where these p_i are correlated over a short period), we can approximate Poisson trials as Bernoulli trials, and estimate the average per-unit erasure rate through simple estimation techniques such as exponentially weighted moving average (EWMA). We then assume that the binomial loss count behavior at the block level holds if the blocks are smaller in time-scale than the time-scale for the average erasure rate to change significantly.

We now move on to the design questions raised in the previous section, the first being how to choose the block size for coding.

A. *Adaptive Granulation: Choice of Block Size N*

Block coding schemes such as RS codes require a block size parameter N units, out of which K are data units, and $N - K$ parity check units. In LT-TCP, the block size is based upon a window comprising data bytes and the desired inventory of PFEC. In LL-HARQ, a PDU from the higher layer is fragmented into K units and $N - K$ parity units are added. In both cases, the block size determines the size of each unit (e.g. a TCP segment or link-layer fragment), and the granulation of the set of data bytes (e.g. TCP's window granulation into segments, or the link-layer PDU's granulation into fragments).

In general, we prefer a larger block size N, to sharply reduce the risk of all the units in the block being lost. When all units of the block are lost, it signals the failure of feedback-based HARQ/FEC method, and there tend to be significant protocol penalties. For example, in LT-TCP this will lead to a retransmission timeout, and associated loss in terms of goodput via window reductions, *ssthresh* reduction and idle times. The penalties at the link layer may vary depending upon whether stop-and-wait is used or other pipelining methods are used. If stop-and-wait is used, a link-level timeout is necessary, but the penalties are not as severe as the transport layer.

Smaller transmission units offer some auxiliary benefits especially at the transport layer. TCP tends to be very brittle and susceptible to timeout when its window is small, even if not all segments are lost. This occurs because of the interaction between retransmission and congestion control mechanisms, and the use of a threshold (e.g. 2 or 3 segments) to trigger the efficient selective/fast retransmission procedures. Granulation of the window would reduce the risk of such small-window brittleness phenomena. Further, with more units in flight, TCP methods such as self-clocking lead to smoother transmission and offer more measurement samples and feedback opportunities (robustness to a lossy feedback channel).

We now revisit our binomial distribution model of the number of losses per block (L). For binomial parameters $(N, p, 1-p)$, i.e block size N and per-unit loss rate p, we ask: what is the probability of losing all units? Figure5 illustrates the distribution of lost packets for $p = 50\%$, with N = 20, 10 and 5. For p = 50%, N = 5 we have a non-trivial 3.175% probability of losing all units. Increasing N to 10 or 20 dramatically reduces this probability. If the loss distribution per-block (L) is not binomial (as our scheme assumes), but bursty on small time-scales (i.e. multi-modal), the larger value of N is even more important. For example, for our ON-OFF error model, with block sizes of N = 5, 10 and 20, the probabilities of losing all the units are 12%, 2.8% and 0.1% respectively. So, a block size of at least 10 gives us robustness and timeout risk 'reduction even with a bursty underlying loss process. Given the penalty associated with the event of losing all units , the risks (i.e. probability multiplied by event penalty) are high, and this suggests a guideline that we set N above 5 definitely and at least 10 if at all possible.

However, our enthusiasm for higher block granulations ($N > 10$) is tempered by the per-packet overheads that affect goodput. TCP/IP overhead of at least 40 bytes/segment increases sharply as the segment size drops to 400 bytes. Other forms of per-packet overhead include wireless MAC overhead. For example, the 802.11b MAC sends MAC-ACK per-PDU at the lowest transmission rate. The 802.16e MAC has similar overheads sent at lower bit rate for robustness.

If the target is to keep the overhead to less than 10% , then a granulation of $N = 10$ is reasonable for medium-speed wireless links (10 Mb/s, shared as will be seen with WiFi and Wimax meshes). For LL-HARQ, the link-layer per-unit overheads are much smaller. Therefore, we target a granulation level (N) of 10 units/block for TCP and 20 units/block for the link-layer HARQ scheme.

B. Adaptive PFEC Sizing:

We set the level of PFEC as a function of the mean ($\mu = Np$) and standard deviation ($\sigma = \sqrt{Npq}$) of the assumed binomial distribution of number of losses in a block, where p is the estimated per-packet loss rate. We examine the following choices of PFEC per block: ($\mu, \mu + \sigma$ and $\mu - \sigma$). The reason is that the binomial distribution is well concentrated around the mean (and indeed approximated by normal distribution for the special cases when $Npq >> 1$). Setting PFEC either too high or too low relative to the mean would lead to heavy PFEC wastage or little PFEC impact respectively. This is modeled as follows:

1) PFEC Wastage: PFEC units are wasted (i.e. more than necessary for block decoding) when more than K units arrive, or equivalently, $L < P_{fec}$. The number of wasted PFEC is $P_{fec} - L$ in such blocks. The total number of wasted PFEC is calculated as a weighted sum as L goes from zero to P_{fec}:

$$\text{\# PFEC Wasted} = \sum_{i=0}^{P_{fec}} (P_{fec} - i) * P(L = i) * M \quad (1)$$

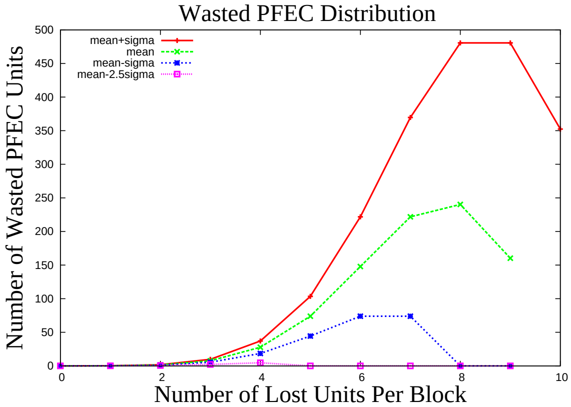

Fig. 6. PFEC Wastage with different PFEC protection levels: (Uniform Loss Scenario with Prob(Unit loss) = 50 %) The number of wasted PFEC units (for a total of 1000 blocks transmitted) is plotted against the number of lost units per block. It can be seen that as the PFEC protection increases, so does the wastage but the need for an ARQ phase decreases.

Recall that the above analysis holds for both the bursty and uniform models. For example, consider the special case of the uniform per-packet erasure model and let $N = 20$ and $p = 50\%$; the distribution of lost units per block has a mean of 10 and σ of 2.2 units. The three choices of PFEC above would imply 8, 10 or 12 PFEC units per block respectively (after rounding off). When PFEC = 8 out of 20 units (i.e. $\mu - \sigma$), there is little PFEC wastage (1%), but 75 % of blocks need another round of transmission placing a large burden on RFEC and increasing timeout risk. When PFEC = 10 units (i.e. μ), there is still low PFEC wastage (4.4%) and 41% need further retransmission. When PFEC = 12 units (i.e. $\mu+\sigma$), the PFEC wastage increases to 11%, but only 13% of blocks need further retransmission and experience timeout risk. These effects are shown in Figure 6. Beyond one standard deviation from the mean, PFEC over-provisioning almost eliminates the burden on future rounds, however, it leads to heavy PFEC waste and lost goodput.

This simple analysis suggests a PFEC-per-block choice of $\mu + \sigma$ to balance the goodput penalty against latency and dependence on retransmission and risk of timeouts. It points to the need to estimate the short-run first and second order statistics of the erasure process in order to efficiently provision FEC, and the huge goodput penalties of over-provisioning PFEC beyond a point.

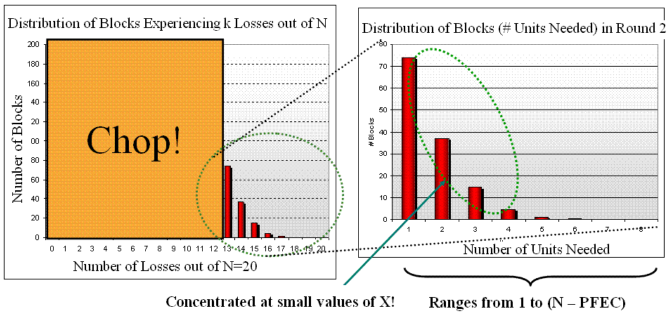

Fig. 7. Chopped Binomial Distribution of # Units Required for Recovery (X) in Round 2. We also see that X is concentrated at small values and decays exponentially.

By setting PFEC-per-block to $\mu+\sigma$, we accept a higher upfront goodput penalty, to reduce the risk of timeouts (for TCP) or residual loss rate (for the link layer). Our analysis also implies that setting PFEC to a very low or very high constant relative to the mean would lead to a poorer trade-off among expected goodput, residual loss (or timeout risk) and expected latency (number of recovery rounds). In fact, when we set PFEC =25% (i.e. $\mu - 2.5\sigma$ in the above case), almost 98% of blocks were un-recovered after the first round! Observe that from the narrow perspective of PFEC waste, this setting is very efficient, but it has little effect in terms of the number of blocks recovered and transfers almost all the burden of recovery to future rounds, increasing latency and timeout risks.

The over-provisioning of PFEC by σ is beneficial in the ON-OFF bursty erasure case, especially during the ON period when the actual erasure rate spikes up. If these ON-OFF periods are less than the measurement time-scales, then we have unmeasured bimodality. During the OFF period, the overhead is also higher due to estimation errors. However, the residual loss rates are smaller because of our conservative PFEC provisioning and aggressive RFEC provisioning (discussed in the next section).

The shape of the wasted PFEC distribution is also revealing: most of the PFEC is wasted around the mean (m) due to the bell-shaped nature of the binomial distribution around the mean (which is the distribution of number of lost units in the block).

C. Adaptive RFEC Sizing:

If retransmission/ARQ is required (i.e., $L > P_{fec}$), the HARQ scheme sends RFEC units. In this case, we have lost L units, received $N - L$ units prior to ARQ, and FEC decoding requires at least K out of N units. Therefore if ARQ is required, we still need $K - (N - L)$ units to reconstruct the block. Let $X = K - N + L$, the minimum number of units needed for reconstruction. The value of X is fed-back from the receiver and is used to tune the RFEC strategy.

The effect of using PFEC on a binomial distribution is to effectively "chop" (see Figure 7) the first part of distribution upto the level of PFEC sizing leaving a residual distribution of the units still required for recovery ($X = N - P_{fec} - L$). This tail decays rapidly (O(e^{-x^2})) due to the normal-like tendencies of the original binomial distribution. The distribution also displays an significant concentration at the lowest values of $X = 1, 2, 3$.

1) RFEC Wastage: The total number of RFEC wasted can be computed as a sum over the wasted RFEC for each case when the ARQ is successful. We had sent Y units of RFEC, Q of which arrived, but we needed only X and so $(Q - X)$ represents wasted RFEC.

Now, $[Q - X]^+$ is the positive projection of $(Q - X)$:

$$[Q - X]^+ = \begin{cases} (Q - X) & \text{if } (Q - X) > 0 \\ 0 & \text{if } (Q - X) \leq 0 \end{cases}$$

The total number of wasted RFEC is therefore:

$$\sum_{L=P_{fec}+1}^{N} P(L) * M * P(S = (Y - Q)) * [Q - X]^+ \qquad (2)$$

The above expression sums up $[Q - X]^+$ with the appropriate probability weights, over the M original blocks.

Note that we have a dilemma in choosing the value Y (see Table I)not directly captured in the above analysis. If we send RFECs using a simple scaling factor based upon the per-unit loss rate p, i.e. ($Y = \frac{X}{(1-p)}$), the total number of recovery units (Y) sent in this round could be small. As discussed earlier, $Y > 5$, and ideally $Y >= 10$ is desired for minimizing the timeout risk, especially when the underlying loss distribution could be bursty or multi-modal. We propose a conservative provisioning of Y as a function of σ, X and p: $Y = \frac{(X+3\sigma)}{(1-p)}$ as depicted in Figure 8.

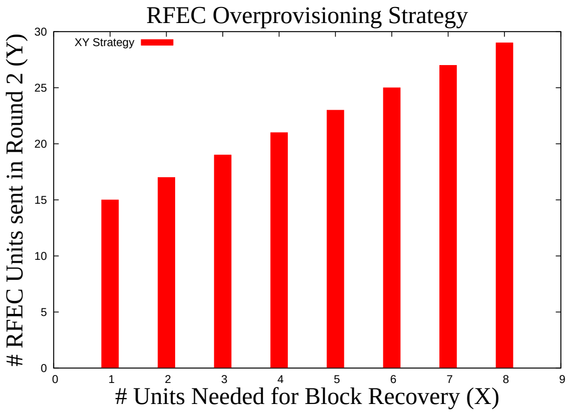

Fig. 8. The RFEC over-provisioning strategy is as shown. We send a larger relative number of RFEC units for small X to counter the *small* N *binomial effect*. As X increases, our RFEC response is less aggressive. This strategy in conjunction with slightly over-provisioned PFEC leads to the most favorable performance. The PER is 50%, with N set to 20.

Observe that, the excess RFEC $(Y - X - pY)$, a relative measure of RFEC wastage, could be large. For example, with a Uniform per-packet erasure model, for $X = 1, p = 0.5, N = 20$, Y is provisioned as 15 units and excess RFEC provisioned is 6.5 units out of Y =15 units. In other words, a large number of RFECs could be wasted relative to RFEC sent (43% in this case). However, since the use of higher PFEC means that a smaller fraction of blocks require RFEC (e.g. 13% of blocks for PFEC = $\mu + \sigma$, when $p = 0.5, N = 20$), and the highest relative RFEC wastage occurs primarily when X is 1, the *total* RFEC wastage is small (3.8%) compared to the *total* PFEC wastage (11%). This total RFEC wastage (assuming a negligible residual loss rate after the ARQ attempt) is a key metric, because it affects goodput. This analysis leads to a non-intuitive insight: high RFEC over-provisioning can still lead to negligible negative impact on goodput, if PFEC is mildly over-provisioned as well.

The RFEC over-provisioning policy is also crucial for the bursty/ON-OFF loss process. For example, if the average PER p is 0.5 under the bursty model, our total PFEC and RFEC wastage values grow to 12% and 10% respectively. The higher level of total FEC wastage occurs because of unpredictability and higher-than-expected variance in the underlying distribution. The residual loss rate in this case grows from almost zero to about 3.6%. The residual loss rate is far worse without the conservative over-provisioning. For example if we use only 1σ instead of 3σ in the formula for Y, the residual loss rate jumps to 12%, even though RFEC wastage reduces to 5%.

D. *Residual Losses:*

Residual loss occurs when the total number of units after ARQ is still short of K, i.e., $(N - L) + Q < K$, where Q out of Y RFEC units arrive. In this case, we are interested in the probability of residual loss, i.e. residual loss rate. We can rearrange the residual erasure inequality condition above as $Q < (K - N + L)$, or simplified as $Q < X$.

The number of lost RFEC units $S = (Y - Q)$ is distributed bimodal-binomial like the r.v. L seen earlier, but with the conditional variable Y replacing the block size N. Given this, the probability of residual loss $P(Q < X)$ is equal to $P(Y - S) < X)$. This can be rearranged as $P(S > (Y - X))$, which is a CCDF value (complementary CDF value, or tail probability sum) of the bimodal-binomial distribution of r.v. S, the number of lost RFEC units.

The number of residual losses can be computed as the sum over the r.v. $V(L)$, as L ranges from $(N - K) + 1$ to N. $V(L)$ is the number of residual losses when L units were lost in the first round prior to the ARQ attempt. $V(L)$ is computed as $V(L) = M * P(L) * P(Q < X)$, where $P(Q < X)$ is the probability of residual loss as explained above.

$$\text{\# Residual Losses} = \sum_{L=P_{fec}+1}^{N} V(L)$$
$$\text{where } V(L) = M * P(L) * P(Q < X) \qquad (3)$$

The percentages (residual loss rates, % RFEC wasted, % PFEC wasted etc..) can then be computed by multiplying equations (1, 2 and 3) by $100/M$. Though this analysis involves several cases and doesn't admit a simple closed form, it is easy to perform numerically for specific values of p. In summary, this analysis points to a trade-off between estimation/prediction errors, FEC-targeting efficiency/goodput, residual loss rate and latency which are validated using the simulations described in the next section.

IV. Comparison of Analysis and Simulation

In this section, we look at the performance of both the link-layer protocol (LL-HARQ) and the transport-layer protocol (LT-TCP). We will see that the LL-HARQ performance is modeled very accurately by the analysis above. The LT-TCP transport performance is captured to a large extent by the analysis though the complexities of the transport protocol cause deviations from the model due to reasons discussed below. We consider both

Error Model	Uniform model (p=0.5)		Uniform model, (p=0.1)		ON-OFF model, (p=0.75/0.25)	
	Analysis	Simulation	Analysis	Simulation	Analysis	Simulation
Link Goodput (Mb/s)	**3.61**	**3.59**	**8.15**	**8.08**	**2.88**	**2.74**
Residual Error Rate (%)	0.0	0.034	0.53	0.0	4.2	4.02
PFEC waste (Mb/s)	1.0	0.99	0.57	0.59	1.13	1.47
RFEC waste (Mb/s)	0.39	0.39	0.28	0.26	0.994	0.662
Avg no. of rounds	1.13	1.12	1.11	1.11	1.90	1.27

TABLE II

LL-HARQ COMPARISON OF ANALYTIC PREDICTIONS WITH THE VALIDATING SIMULATION RESULTS [UNIFORM CASES: 10% AND 50%; BURSTY (ON-OFF) CASE: 75%-25%]. WE ASSUME N TO BE 20 UNITS.

the Uniform and the ON-OFF loss processes (described in Section III). We validated our analysis for two Uniform per-packet erasure model scenarios (with $p = 0.1, 0.5$) and one bursty (ON-OFF) error model scenario (with $p = 0.5, q = 0.75, r = 0.25$). The simulations were run over a single-bottleneck topology (10 Mb/s, 20 ms delay, 10 flows). Six simulation runs were used with each simulation lasting 300 seconds.

Our metrics of interest at the link and transport layers include the obtained goodput, the proactive and reactive FEC usage and wastage, average number of rounds needed per packet. Finally, we also measure the residual loss rate (at the link layer) and frequency of timeout at the transport layer. We will see that in spite of the differences in the LL-HARQ and LT-TCP mechanisms (for example in scheduling RFECs), our analysis explains a large part of the behavior.

A. Link-level Protocol (LL-HARQ)

The link-layer scheme LL-HARQ scheme is captured well by the analysis. As mentioned earlier in Section III, we set $N = 20$ with PFEC set to $\mu+\sigma$ and RFEC set as $Y = \frac{(X+3\sigma)}{(1-p)}$. The comparisons between the analysis and simulations at the link layer can be seen in Table II. We see a close match between all the parameters measured. In particular, for the Uniform case with $p = 0.5$, the goodput is nearly identical (around 3.61 Mb/s). The average number of rounds needed is around 1.13 which validates our earlier expectation that around 13% of the packets will need a RFEC phase.

The results also show that FEC wastage due to over-provisioning is a cause for the loss of goodput, but this is a necessary trade-off to have low residual loss rate after just one round of HARQ. Total PFEC wasted is larger than RFEC waste despite substantial relative over-provisioning of the latter. Goodput is lower for the bursty case because of inaccuracies in targeting FEC to the loss process.

B. Transport-level Protocol (LT-TCP)

At the transport layer, modelling LT-TCP is much more challenging. Among the factors that preclude a close match are:

LT-TCP's scheduling: LT-TCP mechanisms where we send either a new data / PFEC / RFEC or data retransmission packet in response to an incoming ack are not modeled in the analysis.

Self-clocking/ Timer Behavior/RTT Variations: Our model does not capture TCP characteristics such as self-clocking, timer granulation and behavior and the impact of RTT and RTT-sampling. These TCP idiosyncrasies have myriad interactions that impact the obtained performance.

Partial Feedback: LT-TCP acks each packet immediately upon receipt with an ack or dupack. In contrast, at the link-layer each incoming fragment is not acked but a single ack is sent for the entire packet (of N fragments). LT-TCP encodes a partial or current estimate of X (number of units still needed)in the acks. This partial-X increases monotonically and indicates the number of missing units currently seen at the receiver.

Since X is not explicitly known, LT-TCP responds to incoming acks by scheduling RFEC based on this partially known X. Moreover, since LT-TCP may also send RFECs in response to X = 0, we set PFEC = μ for LT-TCP (instead of $\mu + \sigma$ as in the LL-HARQ scheme). Effective PFEC (by including RFEC sent when X = 0) is a little more than ($\mu + \sigma$). Figure 9 shows the behavior of LT-TCP in sending RFEC units in response to the incoming partial-X information. We see that while LT-TCP can only schedule RFEC units in response to incomplete information, it still comes close to the actual value of Y called for in the simulations.

Table III shows the comparison of LT-TCP metrics for the analysis and simulation. We see a good match for some of the metrics such as goodput. For example, under the Uniform loss model, the goodputs from analysis and simulation match exactly for $p = 0.5$ (2.85 Mb/s) and are close for $p = 0.1$. The average number of rounds also matches closely (e.g. 1.37 at $p = 0.5$, Uniform loss). However, percentage of blocks timing out in the ON-OFF case is lower (2.07%) compared to our simulation result of 4.87%. However, the model does not capture several key LT-TCP mechanisms as explained earlier and in the worst cases (ON/OFF scenario with p=0.5), the measured values differ from the model.

V. CONCLUSIONS

Wireless networks can suffer from high and bursty losses which can have serious impact on the performance obtained at the individual link and end-to-end transport layers. Common mitigation techniques include the use of ARQ and FEC schemes. One such approach to providing loss and erasure protection is the use of hybrid ARQ/FEC techniques. We have proposed an HARQ protocol based on adaptive granulation, loss estimation and FEC provisioned in a proactive and reactive manner. To be widely applicable, including the use of the link layer mechanisms in long-delay links (such as airborne or satellite links) it is desirable to constrain the level of persistence in the ARQ mechanism. For example, in our proposed HARQ link layer protocol, we limit ARQ to one retransmission attempt. In this paper, we looked at the trade-offs between latency (ARQ persistence), goodput and residual loss rate. By studying these trade-offs while considering

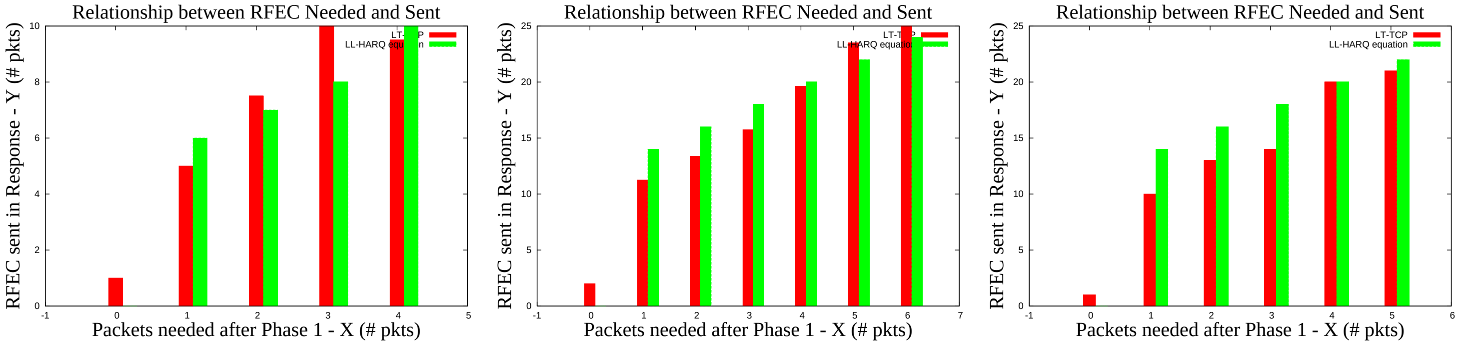

(a) Uniform Loss, PER=10% (b) Uniform Loss, PER=50% (c) ON-OFF Loss, Avg PER=50%

Fig. 9. This figure shows the comparisons between the calculated and empirical RFEC response (Y) for RFEC units needed (X) at the transport layer (LT-TCP). It can be seen that the number of RFEC units sent by LT-TCP is close to what the analysis calls for. It should be noted that LT-TCP response is based on partial feedback. For example, at X =0, RFEC units are sent by LT-TCP when the analysis predicts none are needed.

Error Model	Uniform model (p=0.5)		Uniform model, (p=0.1)		ON-OFF model, (p=0.75/0.25)	
	Analysis	Simulation	Analysis	Simulation	Analysis	Simulation
LT-TCP Goodput (Mb/s)	**2.85**	**2.85**	**7.48**	**7.07**	**2.30**	**2.61**
PFEC waste (Mb/s)	1.02	0.64	1.05	0.40	1.29	1.02
RFEC waste (Mb/s)	0.70	0.84	0.24	0.81	1.41	0.70
Avg no. of rounds	1.37	1.37	1.01	1.07	1.28	1.43
Blocks Timing out(%)	0.01	0.42	0.0	0.02	2.07	4.87

TABLE III

LT-TCP COMPARISON OF ANALYTIC PREDICTIONS WITH THE VALIDATING SIMULATION RESULTS [UNIFORM CASES: 10% AND 50%; BURSTY (ON-OFF) CASE: 75%-25%]. WE ASSUME N TO BE 10 UNITS.

the operation over an abstract lossy channel, we gained insights into the appropriate and most favorable structuring of the building blocks for high-performance link and transport protocols that achieve a high level of loss/erasure tolerance.

The novelty of our protocols lies in way we dynamically adapt our FEC protection based on the estimated loss rate. Both the proactive and the reactive protection levels are tuned based on the feedback and the currently estimated loss rates. This adaptation enables us to minimize the wastage while providing low residual loss rate, high goodput and controlled latency. We then compared the analytical model with simulation results, where the protocols were tailored for the link and transport layers. We find a close match at the link layer in almost all the metrics considered including obtained goodput and FEC wastages. At the transport layer, we have a reasonable match even though loss-tolerant TCP (LT-TCP) complexities such as packet scheduling, incomplete feedback information, variable feedback latencies etc. are not captured by the analytical model. Overall, the analysis captures the significant aspects that influence the performance at the link and transport layers.

In summary, we provided insights into the structuring of the building blocks to create protocols with attractive trade-offs at the link and transport layers. The protocols achieve superior performance in terms of latency, goodput and residual loss rate even under high loss and bursty loss conditions. There are several recently proposed approaches to loss and error correction that use techniques such as network coding and incremental redundancy. We believe that our analysis can be useful in analyzing and structuring such approaches to provision the redundancy needed for error and erasure/loss protection appropriately and get optimal performance.

REFERENCES

[1] C. Barakat and A. Fawal. Analysis of Link-level Hybrid FEC/ARQ-SR for Wireless Links and Long-lived TCP Traffic. volume 57, pages 453–476, Amsterdam, The Netherlands, The Netherlands, 2004. Elsevier Science Publishers B. V.

[2] E. Berlekamp, R. Peile, and S. Pope. The Application of Error Control to Communications. *Communications Magazine, IEEE*, 25(4), April 1987.

[3] J. Byers, M. Luby, M. Mitzenmacher, and A. Rege. A Digital Fountain Approach to Reliable Distribution of Bulk Data. In *Proceedings of the ACM SIGCOMM '98 conference on Applications, technologies, architectures, and protocols for computer communication (SIGCOMM '98)*, pages 56–67, New York, NY, USA, 1998. ACM Press.

[4] M. Luby. LT-Codes. In *Proceedings of the 43rd Annual IEEE Symposium on the Foundations of Computer Science (STOC)*, pages 271–280, November 2002.

[5] J. Nonnenmacher and E. Biersack. Reliable Multicast: Where to use FEC. In *Protocols for High-Speed Networks*, pages 134–148, 1996.

[6] I. Reed and G. Solomon. Polynomial Codes Over Certain Finite Fields. *Journal of the Society for Industrial and Applied Mathematics*, 8:300–304, June 1960.

[7] L. Rizzo. On the Feasibility of Software FEC. Technical Report DEIT Technical Report LR-97013, University of Pisa, 1997. Available at http://www.iet.unipi.it/ luigi/softfec.ps.

[8] D. Rubenstein, J. Kurose, and D. Towsley. A Study of Proactive Hybrid FEC/ARQ and Scalable Feedback Techniques for Reliable, Real-time Multicast. *Computer Communications*, 24(5–6):563–574, 2001.

[9] A. Shokrollahi. Raptor Codes. Technical report, Digital Fountain, Inc, June 2003. Technical Report DF2003-06-001,.

[10] V. Subramanian, S. Kalyanaraman, and K.K.Ramakrishnan. An End-to-End Transport Protocol for Extreme Wireless Network Environments . In *Proceedings of MILCOM 06, IEEE Military Communications Conference*, Washington D.C., USA, October 2006.