Optimal FEC Strategies in Connections with Large Delay-Bandwidth Products

Lavy Libman*
National ICT Australia
Bay 15, Australian Technology Park
Eveleigh, NSW 1430, Australia
Lavy.Libman@nicta.com.au

Ariel Orda
Department of Electrical Engineering
Technion – Israel Institute of Technology
Haifa 32000, Israel
ariel@ee.technion.ac.il

***Abstract*— We study the problem of optimal packet coding for connections with large delay-bandwidth products. Generally, for a given loss rate, using a higher coding redundancy achieves a higher average throughput, but also incurs higher transmission costs (e.g. in terms of energy of a wireless device) and creates a higher load on the network. We define an optimal coding strategy as one that minimizes the expected cost/throughput ratio, for a connection that has a cost per unit time and a cost per transmitted packet. We present an algorithm for computing the optimal strategy and study its properties. We demonstrate that the cost/throughput ratio can be significantly better than with simple retransmission schemes, showing, in particular, that it strongly depends on the decoding buffer size, and obtain several asymptotic bounds on the optimal strategy performance for both unlimited and fixed-size decoding buffers.**

I. Introduction

There are two families of techniques commonly used for achieving reliable communication over an unreliable network connection: *Automatic Repeat reQuest (ARQ)* and *Forward Error Correction (FEC)*. With ARQ, packets that are corrupted or lost in transit are retransmitted upon receiving a negative acknowledgment, or after a timeout during which a positive acknowledgment fails to arrive. With FEC, extra redundancy, in the form of coding, is added at the sender, allowing the receiver to extract the information even in the event of some packets arriving corrupted or lost. The choice of the technique depends both on the application and the network connection parameters; generally, it can be said that FEC is most appropriate for applications requiring a high throughput, and for connections with a high *delay-bandwidth product*, i.e., with a round-trip delay that is significantly higher than a packet transmission time. A good example is a geostationary satellite link, with a round-trip propagation delay of roughly 0.25 seconds, used within a high-speed connection where a packet transmission typically takes a fraction of a millisecond; the delay-bandwidth product is then measured in thousands. In such connections, the overhead required to add the extra redundancy is small compared to the saving of a round-trip wait, which would be required in case of a packet failure. We note that, strictly speaking, to achieve 100% reliability, FEC must still be supplemented by ARQ to cater for the events when even the extra redundant coding fails to recover the information and requires it to be resent; however, proper FEC deployment can reduce the occurence rate of such events by several orders of magnitude.

*This research was performed while L. Libman was with the Dept. of Electrical Engineering, Technion – Israel Institute of Technology.

FEC coding mechanisms can be found at several layers. They are frequently employed at the physical layer (*channel coding*), where they are typically fixed and independent of the applications running above, and can be considered part of the channel specifications. More importantly, they can also be used by the connections, that is, at the data link or the end-to-end transport layers (*source coding*); in that case, they can be more flexible, e.g., the code used and the amount of redundancy added can depend on the application or connection parameters. Indeed, proper selection of the FEC coding mechanism involves a nontrivial tradeoff: too much redundancy may simply increase the overhead without achieving any significant improvement of the probability of successful arrival at the receiver, and may actually reduce the resulting throughput [1], [2]. Furthermore, transmitting too much redundant information has other undesirable effects, in terms of energy consumption (especially important for wireless mobile devices) and of contributing to the network load. To quantify this tradeoff, we associate with the connection a 'cost' per unit time and a 'cost' per packet transmission, and define the optimal coding strategy as one that minimizes the average cost/throughput ratio over time, or, in other words, the average cost per successfully communicated packet. We emphasize that the optimal strategy depends on such parameters as the loss rate and end-to-end delay of the connection.

Forward error correction in a similar framework has recently been the subject of [3], which studied an extension to the classic "sliding windows" flow control mechanism, allowing packet retransmissions to be made before their timeout elapses. It was shown that this extension – essentially, a primitive form of FEC – suffices to reduce the cost/throughput considerably; in particular, it was demonstrated that, when using the optimal retransmission strategy, the average cost per successful packet increases merely *logarithmically* in the price per unit time, rather than *linearly* as is the case for classic sliding-windows implementations.

The framework of our current work is closely related to that of [3], and uses the same model parameters and assumptions, with one exception: instead of simple retransmissions, we allow the connection endpoints to use general FEC packet-level codes. Our analysis is absolutely indifferent to the specific coding mechanism used, and merely assumes that the code is capable of converting k data packets into $n > k$ code packets, so that any k thereof received successfully allow the original data to be reconstructed. Following common terminology, we say that the connection endpoints communicate with (n, k)-*codes*, and we study the properties of optimal strategies available to them that use such codes.[†]

Since decoding an (n, k) code block requires a buffer space of k packets (at least), it is evident that the decoding buffer size available at the receiver plays an important part in the performance that can be attained by the optimal coding strategy. In this paper, we consider the cases of both an unlimited buffer, which allows any amount k of data packets to be encoded at a time, and a limited buffer of size $K_{\max}$, which limits the number of packets that the strategy is allowed to encode in a single code block to $K_{\max}$ at most. We show that, for an unlimited buffer size, an optimal coding strategy technically does not exist: the cost/throughput ratio can be reduced arbitrarily by using ever-larger code blocks, up to an absolute limit that is due only to the cost of transmitting a packet. For a limited buffer size, we present an algorithm for finding the optimal coding strategy, based on a dynamic programming approach. We analyze the asymptotic dependence of the optimal strategy performance on the instance parameters, and show, in particular, that the 'contribution' of the price per unit time to the average cost per packet is proportional to at most the inverse square root of the decoding buffer size.

The special concerns raised by connections with large delay-bandwidth products in general, and satellite links in particular, have attracted considerable research in recent years. Most of these studies are in the context of the widely-used TCP protocol and propose how to improve its performance, either by tuning the parameters of existing features like extended windows, slow-start, and congestion avoidance [5], or by introducing extensions, such as explicit congestion notifications [6]. On the other hand, there exists a vast amount of research on FEC coding, including considerable attention devoted to the bandwidth tradeoff it introduces [2], [7], and to coding schemes that are able to adapt to higher-layer applications and protocols, e.g. in the context of multimedia applications with real-time requirements [8] or in conjunction with TCP [9]. In this paper, we perform optimal FEC analysis with a specific focus on large delay-bandwidth product connections, where the key feature is that a virtually unlimited amount of coding redundancy can be introduced with a negligible overhead; hence, the decision on the optimal coding strategy is not due to a bandwidth tradeoff, but concerns such as energy consumption, which, as expained above, are captured by a price factor per transmitted code packet.

The rest of the paper is structured as follows. Section II describes the model and formally defines the underlying optimization problem. Section III studies the unlimited buffer case, showing that the strategy performance can be improved indefinitely by using ever-larger code blocks, and that an optimal strategy therefore does not exist. A solution algorithm is presented in section IV for the limited-buffer case, and its asymptotic properties are analyzed in Section V. Finally, section VI concludes with a discussion of our methodology and possible extensions that remain for further research.

[†]It is important to observe that, for any $n > k$, an (n, k)-code exists and can be calculated in a straightforward fashion; see, e.g., the discussion on frequency-domain Reed-Solomon codes in [4].

II. Model and Problem Formulation

A. The model

With the exception of FEC coding instead of retransmissions, our current setting is similar to that of [3], and the model assumptions are repeated here for convenience. We are interested in network connections with high delay-bandwidth products; to capture the essence of such connections, we take the packet transmission time to be zero, which implies that the number of packets that can be transmitted within a round-trip period is unlimited. Furthermore, we assume no other factors interfere with the packet communication; e.g., the receiving application is able to process the arriving packets instantly. We assume packets must be delivered at the receiving end in order only; thus, a code block carrying packets other than the next expected ones is discarded, even if received successfully.

We denote the packet loss rate in the connection's path by L, and assume that losses are independent, as is the case, e.g., for white noise or a random discard policy such as RED [10]. Thus, the probability of an (n, k) code block to be successful, i.e. contain at least k successful packets, is

$$P(n, k) \triangleq \sum_{i=k}^{n} \binom{n}{i} (1 - L)^i L^{n-i}. \tag{1}$$

In addition, we neglect the loss rate of acknowledgments, since they are, typically, much shorter than data packets, and therefore suffer less from noise and their paths are often less congested. Consequently, for every code block, the sender knows whether it was successfully received after a round-trip time, which we denote by T.

We assume the connection incurs a cost composed of a 'price' of a per unit of time and b per transmitted packet, and define an optimal strategy as one that minimizes the cost/throughput ratio over time; as explained in the Introduction, these prices can have generic interpretations, e.g. in terms of energy. This linear cost structure is appropriate for a variety of scenarios and cost interpretations [3], [11]. A different (nonlinear) cost structure may be used instead, provided that the cost of transmitting a code block (or a sequence thereof) depends only on its total number of packets, and not on their identities, contents, or the actual number successfully received. Such a different cost structure can affect only the analytical results, e.g. the asymptotic dependence of the optimal strategy performance on the prices, whereas the actual algorithm for finding it remains intact.

The computation of the optimal strategy from the connection parameters (L,T,a,b) implicitly assumes that they are known; therefore, they must either remain constant or change quasi-statically, allowing the strategy to adapt after a change is detected. If any parameter, e.g. the round-trip time, changes quickly and unpredictably, it should be modeled by a random variable (e.g., as in [11]) rather than a constant value. We point out, however, that this is not typical of the kind of network connections that are the subject of this study: e.g., for satellite links, the round-trip time is dominated by the propagation delay, which can be considered essentially constant.

The above assumptions readily imply two fundamental properties. First, in the optimal strategy, packets are transmitted in 'bursts' only at multiples of T. Indeed, suppose that a sequence of code blocks is sent at time $t = 0$; then, their acknowledgments arrive at $t = T$, and until that time no further information is available to the sender. Thus, any transmissions between $0 < t < T$ can instead be made at $t = 0$, without any adverse effect on the performance of the strategy; the same consideration can be repeated inductively for all multiples of T. Second, once a sequence of packets is sent at time t and the acknowledgments arrive back at $t + T$, the index of the last packet to have arrived successfully and in order is known, so the strategy simply restarts ('slides') at the subsequent packet. Consequently, the description of a strategy consists of a single vector, specifying the configuration of code blocks to be sent at every multiple of T relative to the next-expected packet index. Our purpose subsequently will be to find the optimal such vector and its properties.

B. Problem formulation

Consider a vector $\vec{n} = \langle (n_1, k_1), \ldots, (n_i, k_i), \ldots \rangle$, where n_i, k_i are whole and non-negative and $n_i \geq k_i$, and define a random variable S to be the number of in-order successful packets at the receiver if the sender transmits an (n_1, k_1)-code of the first k_1 packets, followed by the (n_2, k_2)-code of the following k_2 packets, etc. We define the *score* of $\vec{n}$, denoted by $\phi(\vec{n})$, to be the expected value of S; thus

$$\phi(\vec{n}) \triangleq \mathrm{E}[S] = \sum_{j=1}^{\infty} k_j \cdot \prod_{i=1}^{j} P(n_i, k_i), \tag{2}$$

where $P(n_i, k_i)$, the individual probability of the (n_i, k_i) code block to arrive successfully, is given by (1).† We seek the vector $\vec{n} = \langle (n_1, k_1), \ldots, (n_i, k_i), \ldots \rangle$ that minimizes

$$\frac{a \cdot T + b \cdot \sum_{i=1}^{\infty} n_i}{\phi(\vec{n})}. \tag{3}$$

The above expression describes the cost/throughput ratio attained by the strategy $\vec{n}$ over time. The numerator is the fixed cost of a period of T, during which one 'burst' (code block sequence) is transmitted, and the denumerator is the expected number of packets successfully communicated in that period.

We note that the problem tackled in [3] is merely a special case in our current terms, with the extra constraint $k_i = 1$ for all i; indeed, retransmitting n_i copies of a packet is equivalent to coding it with the code $(n_i, 1)$. We subsequently show that the major properties of the optimal strategy, including the complexity of its solution and the asymptotic dependence of its cost/throughput performance on the time price, remain similar to the case of simple retransmissions [3] as long as there exists a bound of $k_i \leq K_{\max}$. However, if k_i are not constrained, the problem possesses entirely different properties; in particular, there is not even a finite optimal vector then.

†Expression (2) is correct under the assumption that a failed code block (i.e., one without a sufficient number of successful packets) is worthless. Some codes may allow recovery of partial information in that case, which increases the score expression somewhat; however, this does not have a significant effect on our subsequent analysis.

III. The Unlimited Buffer Case

We begin by showing that the problem as described by (3), without any further constraints on n_i and k_i, does not have a finite solution.

Lemma 1. *There does not exist a finite coding strategy that achieves a cost/throughput ratio of* $\frac{b}{1-L}$ *or less.*

Proof. Consider a strategy $\vec{n}$ that transmits a total of $N = \sum_{i=1}^{\infty} n_i$ packets per period. Its score $\phi(\vec{n})$ cannot be higher than the expected number of individually successful packets, which is $N(1-L)$ (even if the in-order arrival requirement is disregarded). Thus, $\frac{aT+bN}{\phi(\vec{n})} \geq \frac{aT+bN}{N(1-L)} > \frac{b}{1-L}$. □

Lemma 2. *For any* $\epsilon > 0$*, there exists a coding strategy that achieves a cost/throughput ratio of less than* $\frac{b(1+\epsilon)}{1-L}$.

Proof. Consider a vector $\vec{n} = \left\langle \left(N, \frac{1-L}{1+0.25\epsilon}N\right), (0,0), \ldots \right\rangle$. By the law of large numbers, $\lim_{N\to\infty} P\left(N, \frac{1-L}{1+0.25\epsilon}N\right) = 1$; hence, there exists some N_1 such that, for any $N > N_1$, $P\left(N, \frac{1-L}{1+0.25\epsilon}N\right) > \frac{1+0.25\epsilon}{1+0.5\epsilon}$, and, therefore, $\phi(\vec{n}) > \frac{1+0.25\epsilon}{1+0.5\epsilon} \cdot \frac{1-L}{1+0.25\epsilon}N = \frac{(1-L)N}{1+0.5\epsilon}$. Also, denote $N_2 = \frac{aT(1+0.5\epsilon)}{b \cdot 0.5\epsilon}$. Now, choose some $N > \max(N_1, N_2)$. Then, the cost/throughput ratio attained by the strategy $\vec{n}$ is $\frac{aT+bN}{\phi(\vec{n})} < \frac{aT+bN}{\frac{(1-L)N}{1+0.5\epsilon}} < 0.5\epsilon\frac{b}{1-L} + \frac{b(1+0.5\epsilon)}{1-L} = \frac{b(1+\epsilon)}{1-L}$. □

The following theorem is a direct corollary of lemmas 1–2.

Theorem 1. *For unlimited* k_i*, there does not exist a finite optimal coding strategy.*

IV. Limited Buffer: Solution Algorithm

The previous section showed that there is no 'optimal' buffer size at the receiver; the higher it is, the lower the cost/throughput ratio that can be attained. In practice, however, the decoding buffer size available at the receiving side is limited; this corresponds to an extra constraint on the possible strategy vectors, namely, the strategy must satisfy $k_i \leq K_{\max}$ for all i, where $K_{\max}$ denotes the buffer size. We now present a solution for the case when this constraint is present. Our method is based on a dynamic programming approach.

Consider expression (3), and note that, for any N, all the vectors with $\sum_{i=1}^{\infty} n_i = N$ attain the same numerator value; hence, the comparison among them is based merely on their

Initialization: Set $\vec{n}(0) = \langle(0,0),(0,0),\ldots\rangle$, $N \leftarrow 0$, $E_L(0) \leftarrow 0$, $Best\_CTR \leftarrow \infty$
Loop until $Best\_CTR$ has not decreased for several iterations:
Set $N \leftarrow N+1$
Set $E_L(N) \leftarrow \max_{\substack{1\le n_1 \le N \\ 1 \le k_1 \le K_{\max}}} P(n_1,k_1)\,[k_1 + E_L(N-n_1)]$, where $P(n_1,k_1)$ is given by (1)
Set (n_1^*, k_1^*) to the arguments that achieved the maximum in the previous line
Set $\vec{n}(N)$ to the concatenation of $\langle(n_1^*, k_1^*)\rangle$ and $\vec{n}(N-n_1^*)$
Set $CTR \leftarrow$ the cost/throughput ratio for $\vec{n}(N)$
If $CTR < Best\_CTR$
Set $Best\_CTR \leftarrow CTR$, $N^* \leftarrow N$

Fig. 1. Algorithm Dynamic-Limited Buffer (**D-LB**).

score. Consequently, let us define

$$\mathrm{E}_L(N) \triangleq \max_{\substack{n_1,k_1,n_2,k_2,\ldots \\ s.t. \sum_i n_i = N}} \left\{ \sum_{j=1}^{\infty} k_j \prod_{i=1}^{j} P(n_i,k_i) \right\}; \tag{4}$$

then, the minimum attained by (3) for a given N is

$$\frac{a \cdot T + b \cdot N}{\mathrm{E}_L(N)}. \tag{5}$$

Furthermore, consider the score expression (2) and rearrange it as follows:

$$\sum_{j=1}^{\infty} k_j \prod_{i=1}^{j} P(n_i,k_i) = P(n_1,k_1) \left[k_1 + \sum_{j=2}^{\infty} k_j \prod_{i=2}^{j} P(n_i,k_i) \right]. \tag{6}$$

Thus, if (n_1,k_1) is fixed, the dependence of the vector's overall score on the other elements (code blocks) is only through the score of the subvector that begins with the second element. Consequently, the following relation holds:

$$\mathrm{E}_L(N) = \max_{\substack{1\le n_1 \le N \\ 1 \le k_1 \le K_{\max}}} P(n_1,k_1)\left[k_1 + E_L(N-n_1)\right]. \tag{7}$$

Relation (7) suggests that the optimal score for a certain N (and the vector that achieves it) can be found from the optimal scores of lesser N by dynamic programming. The corresponding algorithm, termed **D-LB** (for "Dynamic-Limited Buffer"), is stated formally in Figure 1.[†] It computes the optimal scores for ever-increasing N, until the cost/throughput ratio decreases no more. We point out that the termination condition is purposely left somewhat vague, and only state that the search should not be terminated prematurely at the first minimum, which may be a 'false' local rather than a global one. Figure 2, which plots the cost/throughput ratio as a function of N for $a = 1$, $T = 1$, $b = 1$, $L = 0.1$, $K_{\max} = 8$, exhibits that the dependence can be quite erratic; indeed, one can clearly observe the local minima at $N = 4$ and $N = 12$, before the true global minimum at $N = 23$, corresponding to the optimal strategy of $\langle(11,8),(10,8),(1,1),(1,1),(0,0),\ldots\rangle$, attaining a cost/throughput of 1.43602.[‡] Finding an efficient search termination condition, which would ensure reaching

[†]We point out that this algorithm can also be used in the context of [3] to find an *exact* optimal retransmission strategy (as opposed to the *approximate* algorithms suggested there), by setting $K_{\max} = 1$.

[‡]Incidentally, the optimal strategy for the same parameters except $K_{\max} = 1$ (i.e. using retransmissions only) reaches a cost/throughput of 3.0951.

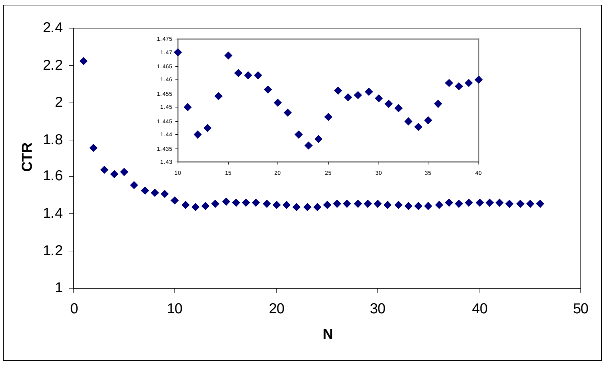

Fig. 2. The ratio $\frac{a \cdot T + b \cdot N}{\mathrm{E}_L(N)}$ as a function of N, for $a = 1$, $T = 1$, $b = 1$, $L = 0.1$, $K_{\max} = 8$. The inset 'zooms in' on $10 \le N \le 40$.

the global optimum without continuing the search too much beyond the optimal N, remains a subject for further study.

V. Limited Buffer: Asymptotic Properties

In this section, we are concerned with the asymptotic dependence of the cost/throughput performance of the optimal coding strategy on the problem parameters.

Theorem 2. *The maximum score attained by a vector of size N satisfying $k_i \le K_{\max}$ for all i is $O\left(K_{\max} \frac{N}{\log_{1/L} N}\right)$.*

Proof. Given a strategy $\vec{n} = \langle(n_1,k_1),\ldots,(n_i,k_i),\ldots\rangle$, consider the following vector of simple retransmissions: $\vec{n'} = \langle(n_1,1),\ldots,(n_i,1),\ldots\rangle$. Since $P(n_i,k_i) \le P(n_i,1)$, we have $\phi(\vec{n}) = \sum_{j=1}^{\infty} k_j \cdot \prod_{i=1}^{j} P(n_i,k_i) \le K_{\max} \sum_{j=1}^{\infty} \prod_{i=1}^{j} P(n_i,1) = K_{\max}\phi(\vec{n'})$. Thus, the optimal score attained by a coding vector of size N cannot be more than $K_{\max}$ times the optimal score attained by a retransmission strategy of the same size, which was shown in [3] to be $O\left(\frac{N}{\log_{1/L} N}\right)$. Therefore, we have $E_L(N) = O\left(K_{\max} \frac{N}{\log_{1/L} N}\right)$. □

Theorem 3. *As $a \to \infty$ (for fixed values of T, b, L, $K_{\max}$), the cost/throughput attained by the optimal vector is logarithmic in a.*[§]

Proof. In light of theorem 2, the best cost/throughput attained by a strategy of size N is $O\left(\frac{aT+bN}{K_{\max} N/\log_{1/L} N}\right)$, which, for a fixed value of $K_{\max}$, is simply $O\left(\frac{aT+bN}{N/\log_{1/L} N}\right)$, i.e., exactly as for a retransmission strategy of size N. Consequently, the asymptotic dependence of the optimal cost/throughput on a is also the same; as shown in [3], this dependence is logarithmic. □

Finally, our last theorem provides a quantitative evaluation of the tradeoff between the receiving buffer size and the cost/throughput attainable by the optimal coding strategy.

[§]For comparison, in the 'classic' sliding-window scheme, which does not use encoding or retransmissions until a timeout elapses, the cost/throughput is linear in the time price [3].

Theorem 4. *For fixed a, T, b, L, and $K_{\max} \to \infty$, the optimal cost/throughput is not worse than* $\frac{b}{1-L} + O\left(\frac{1}{\sqrt{K_{\max}}}\right)$.

Proof. It suffices to present one strategy that achieves the claimed cost/throughput. Indeed, consider the strategy $\vec{n} = \langle (N(K_{\max}), K_{\max}), (0,0), \ldots \rangle$, where $N(K) = \frac{K}{1-L}\left(1 + K^{-0.5}\right)$. Let X be the random variable denoting the number of successful packets among $N(K)$. Observe that its mean is $\mathrm{E}[X] = K\left(1 + K^{-0.5}\right)$, and its variance $\sigma_X^2 = L \cdot K\left(1 + K^{-0.5}\right)$. Therefore, by Chebyshev's inequality,

$$Prob\{X < K\} \leq Prob\left\{\frac{|X - \mathrm{E}[X]|}{\sigma_X} > \frac{1}{\sqrt{L\left(1 + K^{-0.5}\right)}}\right\} \leq L\left(1 + K^{-0.5}\right).$$

Thus,

$$\phi(\vec{n}) = K_{\max} \cdot P(N(K_{\max}), K_{\max}) \geq K_{\max} \cdot \left[1 - L\left(1 + K_{\max}^{-0.5}\right)\right],$$

so, finally,

$$\frac{aT + bN(K_{\max})}{\phi(\vec{n})} \leq \frac{aT + \frac{b}{1-L}K_{\max}\left(1 + K_{\max}^{-0.5}\right)}{K_{\max}\left[1 - L\left(1 + K_{\max}^{-0.5}\right)\right]} = aT \cdot O\left(\frac{1}{K_{\max}}\right) + \frac{b}{1-L}\left[1 + \cdot O\left(K_{\max}^{-0.5}\right)\right] = \frac{b}{1-L} + O\left(\frac{1}{\sqrt{K_{\max}}}\right).$$

□

It remains an open question whether the bound presented by Theorem 4 is tight. Finding a more precise characterization of the asymptotic dependence between the optimal cost/throughput and the receiver's decoding buffer size, as well as an analytic representation of the optimal strategy itself, is left for further investigation.

VI. Conclusion

We have investigated optimal FEC coding strategies for network connections where the packet transmission time is negligible compared to the round-trip delay. We associated a cost per unit of time and per packet transmission with the connection, and defined the optimal strategy as one that minimizes the expected cost/throughput ratio, i.e., the average cost per successfully communicated packet. We proposed a algorithm, based on a dynamic programming approach, to find the optimal strategy, and studied the analytic properties of the result. We showed that, using a sufficiently large coding buffer, the impact of the time price on the average packet cost can be made arbitrarily small – more precisely, proportional to at most the inverse square root of the buffer size. Additionally, for a fixed (limited) buffer size, we proved that the cost/throughput ratio of the optimal strategy increases logarithmically in the time price. These results provide a significant insight on the buffer requirements for reliable connections in networks with large delay-bandwidth products and transmission costs, and demonstrate the tradeoff that exists between buffer size and the cost of communication. Our approach was demonstrated to attain a significantly lower cost/throughput ratio than both 'classic' sliding windows, where a packet is not retransmitted until after a timeout or negative acknowledgment, and the approach of [3], which used only simple retransmissions rather than genuine FEC coding.

The strategies discussed in this paper make transmission bursts at multiples of the round-trip time, and wait for all the acknowledgments from a previous transmission to arrive before making the next one. As explained in the model description in section II, this behavior is optimal if the packet transmission time is neglected, and is quite adequate if the connection's delay-bandwidth product is large (i.e. a packet transmission time is negligible compared to the connection round-trip time). Otherwise, i.e. if a packet transmission takes a sizeable fraction of the round-trip time, it may be better not to wait for all acknowledgments from the previous burst, but, rather, proceed with transmission with only a partial information on previous successes and losses. Then, a strategy is no longer described by a vector applied at every multiple of the round-trip time, but a rule applied after every packet transmission, specifying the packet or code-word most worthwhile to transmit next (if at all), according to the information available up to that moment. The investigation of optimal strategies and their properties in this framework, for FEC coding and even for simple retransmission schemes, is left as a subject for future work.

References

[1] A. Guha and T.-S. Chang. The effectiveness of cell-level FEC for packet delivery in ATM networks. In *Proc. 22nd IEEE Conference on Local Computer Networks (LCN)*, pages 264–273, Minneapolis, MN, November 1997.

[2] C. Barakat and E. Altman. Bandwidth tradeoff between TCP and link-level FEC. *Computer Networks*, 39(2):133–150, June 2002.

[3] L. Libman and A. Orda. Optimal sliding-window strategies in networks with long round-trip delays. In *Proc. IEEE Infocom*, San Francisco, CA, April 2003.

[4] A.D. Houghton. *The engineer's error coding handbook*. Chapman & Hall, London, UK, 1997.

[5] M. Allman, C. Hayes, H. Kruse, and S. Ostermann. TCP performance over satellite links. In *Proc. 5th International Conference on Telecommunication Systems*, pages 456–469, Nashville, TN, March 1997.

[6] D. Katabi, M. Handley, and C. Rohrs. Internet congestion control for future high bandwidth-delay product environments. In *Proc. ACM SIGCOMM*, Pittsburgh, PA, August 2002.

[7] Q. Zhao, P.C. Cosman, and L.B. Milstein. Optimal bandwidth allocation for source coding, channel coding and spreading in a CDMA system. In *Proc. 3rd IASTED Conference on Wireless and Optical Communications (WOC)*, pages 463–468, Banff, Alberta, Canada, July 2003.

[8] K. Park and W. Wang. AFEC: An adaptive forward error correction protocol for end-to-end transport of real-time traffic. In *Proc. 7th International Conference on Computer Communications and Networks (ICCCN)*, pages 196–205, Lafayette, LA, October 1998.

[9] B. Liu, D.L. Goeckel, and D. Towsley. TCP-cognizant adaptive forward error correction in wireless networks. In *Proc. IEEE Globecom*, pages 2128–2132, Taipei, Taiwan, November 2002.

[10] S. Floyd and V. Jacobson. Random early detection gateways for congestion avoidance. *IEEE/ACM Transactions on Networking*, 1(4):397–413, August 1993.

[11] L. Libman and A. Orda. Optimal timeout and retransmission strategies for accessing network resources. *IEEE/ACM Transactions on Networking*, 10(4):551–564, August 2002.