Optimal Codes for the Burst Erasure Channel

Jon Hamkins*

We make the simple observation that the erasure burst correction capability of any (n, k) code can be extended to arbitrary lengths above n with the use of a block interleaver, and discuss nuances of this property when channel symbols are over GF(p) and the code is defined over GF(p^J), $J > 1$. The results imply that maximum distance separable codes (e.g., Reed-Solomon) offer optimal burst erasure protection with linear complexity, and that the optimality does not depend on the length of the code.

I. Introduction

This article makes the fairly straightforward observation that block interleaving a code to depth I results in an approximate I-fold increase in burst erasure correction capability, which implies that short, low-complexity maximum distance separable (MDS) codes together with long interleavers maximize the length of resolvable erasure bursts among signaling schemes of the same rate and total transmission duration.

This observation is simple enough that perhaps this article is unnecessary. Indeed, this basic property of block interleaving on bursty channels has been made before, e.g., on the related burst error channel [9] (the present article addresses the burst erasure channel). However, the impact of the result has not been fully appreciated in the design of practical systems, and the literature contains several imprecise or incorrect statements and is continuing down a path of much more complicated designs that do not match the performance of simpler systems.

This article aims to present practical linear-complexity designs for the burst erasure channel, with the goal of maximizing the resolvable erasure burst length of a transmission. Along the way, we provide an explicit accounting of the relationship between the length of a channel symbol erasure burst and a code symbol erasure burst, both with and without an interleaver. This leads to a theorem on the power of block interleavers to increase the resolvable length of erasure bursts.

---

*Communications Architectures and Research Section.

The research described in this publication was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

II. Channel model and previous work

An erasure channel is one in which each transmitted symbol is either received correctly or is corrupted so badly as to be considered erased. When the erasures are clustered together, as shown in Fig. 1, we refer to the channel as a *burst erasure channel*.

In typical space communications systems currently deployed, incorrectly decoded frames of data are flagged, either by a forward error correction (FEC) decoder or through a subsequent cyclic redundancy check (CRC) code [2, 4]. Such damaged frames are removed from the data stream before being handed up to higher layers, and the result is a stream of correctly decoded data containing bursts of erasures. Note that in this application, there are no other impairments other than the bursts of erasures.

A burst erasure channel model is also appropriate for several other communication scenarios: in Ka-band space-Earth links, which experience weather-induced outages [14]; in optical and magnetic storage (CD, DVD, hard disk drive, etc.), in which thermal asperity, physical defects from scratches, or other impairments can erase a contiguous block of symbols [15, 19]; in free space optical links operating in a low background light regime, where detecting a background photon is a rare event and the primary channel impairment is atmospheric events that prevent the detection of any signal photons for a duration of time [6]; during operational outages at the receiver due to loss of carrier lock or frame lock, or on-the-fly changes of code rate or modulation [4]; and in channels dynamically impaired by large interferers [16].

The need for efficient solutions to the burst erasure problem for the space application has led the Consultative Committee for Space Data Systems (CCSDS) to study it, primarily with emphasis on the design of long, iteratively decoded LDPC codes or related structures, including irregular repeat-accumulate (IRA) codes, generalized IRA codes, Tornado codes, and protograph-based codes [4, 5].

The computational complexity of decoding (n, k) Reed-Solomon codes is $O(n \log^2 n)$ [7], and popular practical algorithms decode in $n - k$ clock cycles and use less than $6(n - k)$ multipliers, for an area-time complexity of $O((n - k)^2)$ [13]. Because of this quadratic complexity, Reed-Solomon (RS) codes have been dismissed as a viable solution, and "the impossibility of exploiting long codeword lengths represents a limit to the performance achievable" [4]. In this article, our observations will show that a simple block interleaver overcomes these complexity and length issues, making RS codes not only viable but an optimal solution for the noiseless burst erasure channel application.

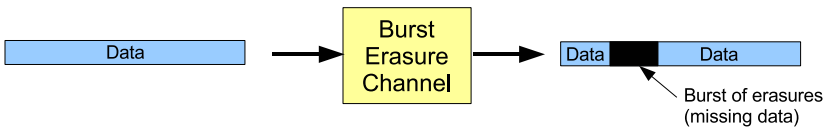

Figure 1. Burst erasure channel.

Yang and Ryan [19] noted the power of block interleaved RS codes to resolve long bursts of erasures (although the result was slightly misstated, as we will discuss later), and compared several codes of similar lengths and rates on the basis of both coding gain and resolvable erasure burst length. The channel model in that case was burst erasure plus additive white Gaussian noise (AWGN), and the comparison included RS codes, an LDPC code of Mackay , Euclidean geometry LDPC codes, extended IRA codes, and array codes (see [19] for references). The authors concluded that "LDPC codes are very effective against noise bursts and, in fact, are superior to RS codes in the regime simulated." This conclusion is appropriate for the burst erasure plus AWGN channel considered, where a hybrid metric of both good erasure burst protection and coding gain are desired. A similar conclusion was reached when comparing nonbinary LDPC codes and RS codes [15]. However, the fact is that neither LDPC codes nor RS codes are best at both metrics – LDPC codes provide better coding gain, and as we shall see, interleaved RS codes provide unbeatable burst erasure protection.

The length of burst erasure protection that an arbitrary linear block code has is bounded by a certain value, called the zero-covering span, of its parity-check matrix [17]. This has led to comparisons of quasi-cyclic LDPC codes to earlier designs [19], and later, to a constructive technique for LDPC codes that approach maximum efficiency for asymptotically long codeword lengths [8]. In another work [12], it was shown that codes designed for the burst erasure channel achieve essentially the same performance on the Rayleigh fading channel, which means that to design good codes on the Rayleigh channel, the simpler process of designing codes for the burst erasure channel may be employed.

An LDPC code designed for the erasure channel can be block interleaved to achieve excellent performance on the burst erasure channel [5]. Alternatively, without lengthening the code, one may permute the variable nodes of an LDPC code in order to maximize the maximum resolvable erasure burst length, either with a structured algorithm [10] or by simulated annealing [16].

III. Preliminaries

This article assumes the channel transmits symbols from GF(p) and that J uses of the channel are used to transmit a symbol from GF(p^J). A linear (n, k) code over GF(p^J) takes each k information symbols in GF(p^J) and encodes them into a codeword of n symbols in GF(p^J), $n \geq k$. We say a code symbol in GF(p^J) is *fully erased* if all J of its channel symbols are erased, and *partially erased* if at least one but fewer than J of its channel symbols are erased. In either case, such a code symbol is typically considered to be erased by a conventional decoder operating over GF(p^J). For example, if $p = 2$ and $J > 1$, then a single binary channel symbol erasure would lead to the erasure of a J-bit code symbol. Following, e.g., [19, 10, 12], the *maximum resolvable erasure burst length*, $L^{ch}_{\max}$, is defined as the maximum number of consecutive erased channel symbols that the code is guaranteed to correct regardless of where the burst begins. In the present article, we insert the superscript *ch* to indicate that the length is measured in channel symbol erasures. We

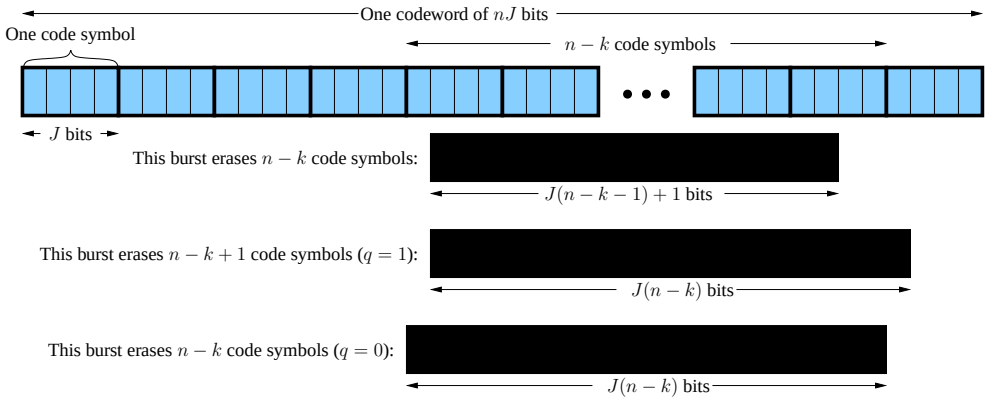

Figure 2. A channel symbol erasure burst of length $J(n-k-1)+1$ erases $n-k$ code symbols, regardless of its alignment. A burst of length $J(n-k)$ may erase $n-k$ or $n-k+1$ code symbols, depending on its alignment.

use $L^c_{\max}$ to denote the maximum resolvable erasure burst length in code symbols.

The Singleton bound states that the minimum distance of any linear (n,k) code satisfies $d \leq n-k+1$ [3], which is a consequence of the facts that every linear code is equivalent to a systematic linear code and that changing one systematic code symbol can change at most $n-k$ parity symbols. Any code that satisfies the Singleton bound with equality is *maximum distance separable* (MDS). It is well known that Reed-Solomon (RS) codes, for example, are MDS.

For an MDS code, the maximum resolvable code symbol erasure burst length is $L^c_{\max} = n-k$, since an MDS code satisfies the Singleton bound with equality, $d = n-k+1$. This length is the fraction $(n-k)/n = 1-r$ of the symbols in the codeword where $r = k/n$ is the code rate, and no linear (n,k) code has a larger $L^c_{\max}$.

It has been noted that an (n,k) RS code over GF(2^J) has maximum resolvable channel symbol erasure burst length $L^{ch}_{\max} = JL^c_{\max} = J(n-k)$ (e.g., [19]), but this assessment glosses over the impact of misalignment between channel symbols and code symbols. As illustrated in Fig. 2, some bursts of length $J(n-k)$ result in $n-k+1$ code symbol erasures, which are uncorrectable in general. A moment's reflection reveals that any MDS code over GF(p^J) has maximum resolvable erasure-burst length

$$L^{ch}_{\max} = J(n-k-1)+1, \tag{1}$$

and when a burst is aligned to the J-bit code symbols the code can correct up to length $J(n-k)$.

More generally, for an arbitrary code over GF(p^J), not necessarily MDS, we may replace $n-k$ above with the maximum resolvable code symbol erasure burst length $L^c_{\max}$, and conclude that

$$L^{ch}_{\max} = J(L^c_{\max}-1)+1. \tag{2}$$

Note that for $J>1$ we have $L^{ch}_{\max} = 1 \mod J$, a fact that we will use later in the article.

When $J = 1$, (2) simplifies to $L^{ch}_{\max} = L^{c}_{\max}$, as expected.

This article considers the block-interleaved coded system shown in Fig. 3. For each labeled point in Figure 3, the format of the signal at that point is shown in Figure 4. For ease of exposition, Figure 4 refers to "bits" (i.e., channel symbols from GF(2)) and code symbols from GF(2^J), although the description applies equally to a channel transmitting symbols from GF(p) and a code over GF(p^J) by replacing "bit" with "p-ary symbol." When an (n, k) code over $GF(2^J)$ is used, messages at the input to the encoder are grouped into blocks of kJ bits, as shown in Fig. 4(a), and then encoded to form codewords of nJ bits as shown in Fig. 4(b). A block interleaver [1] can be viewed as a function that takes codewords written in rows and reads them out by column, one code symbol (J bits) at a time. The transmission order is shown in Fig. 4(c), with each small colored section on the right-hand side representing J bits. The entire interleaved block comprises I codewords of n symbols, each of J bits, for a total of InJ bits. After going through the channel, a number of erasure bursts occur, one of which is shown in Fig. 4(d). After de-interleaving, a single burst has the form shown in Fig. 4(e). If the burst is not too long, then the erasures are corrected and the original message is recovered, as shown in Fig. 4(f).

IV. Burst Erasures on Block Interleaved Codes

A burst of binary channel symbol erasures results in a corresponding burst of code symbol erasures in GF(p^J). We now discuss the relationship between the lengths of these bursts, and the effect of interleaving. Throughout the section, we assume an (n, k) code over GF(p^J) is block interleaved to depth I. To keep the accounting straight, we let B_{ch} be the length of a burst of channel symbol erasures, let B be the length of the corresponding code symbol burst, and let $B(i)$ denote the number of code symbols erased in the ith row (the ith codeword) of the block interleaver.

The detailed counting done in this section can be seen graphically in Fig. 5. Each of the I codewords in the interleaver is labeled $1, 2, \ldots, I$, and contains n symbols of J bits each. As seen in the burst on the left, a channel symbol burst of length $IJL^{c}_{\max} - J + 1$ gives rise to exactly $L^{c}_{\max}$ code symbol erasures in each codeword, and this count is unaffected by the relative alignment of the channel symbol erasure burst to the code symbols. If an erasure burst is aligned to the code symbols, as it is for the burst on the right, then a channel

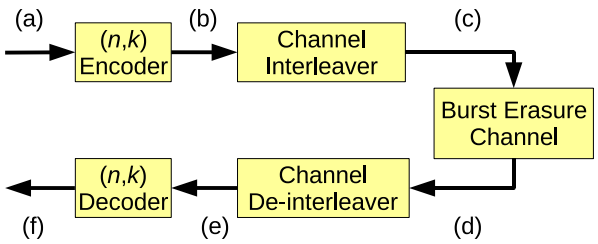

Figure 3. Block interleaved RS coding on burst erasure channel. Letters refer to the formats shown in Fig. 4.

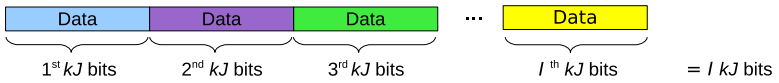

(a) Information messages.

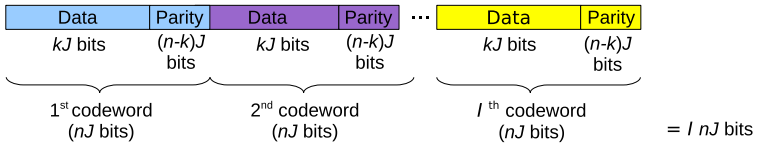

(b) Codewords.

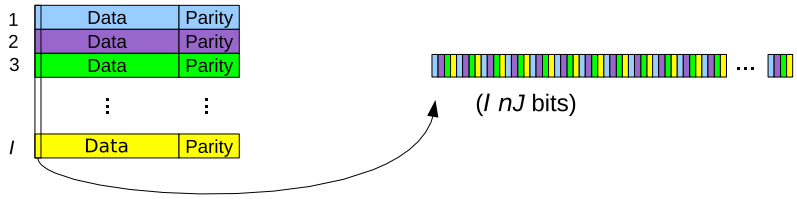

(c) Codewords in rows, and corresponding interleaver output.

(d) An erasure burst.

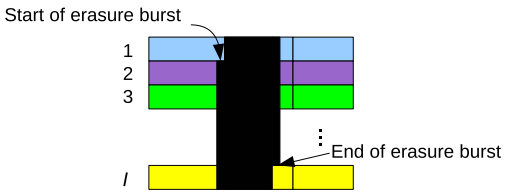

(e) An erasure burst after de-interleaving.

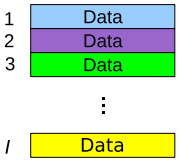

(f) Decoder output.

Figure 4. The data format in positions (a) - (f) of Figure 3, when $p = 2$.

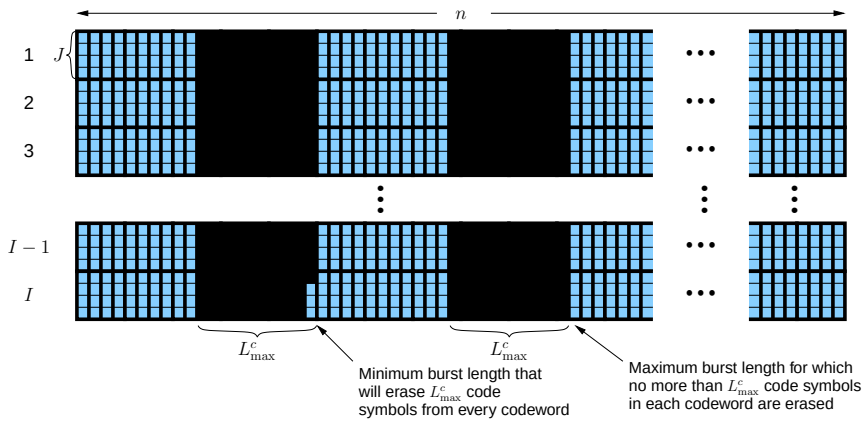

Figure 5. A graphical view of burst erasures.

symbol burst of length $IJL^c_{\max}$ gives rise to $L^c_{\max}$ code symbol erasures in each codeword.

A. Length of de-interleaved code symbol erasure bursts

For each i, $1 \leq i \leq I$, we have

$$B(i) \in \left\{ \left\lfloor \frac{B}{I} \right\rfloor, \left\lceil \frac{B}{I} \right\rceil \right\}. \tag{3}$$

That is, a code interleaved to depth I that experiences a burst of B code symbol erasures results in each codeword experiencing a code symbol erasure burst of the same length, within one. This can be proved by induction by noting that when $B < I$, and regardless of which row the erasure begins, each codeword in the interleaved block experiences either 0 or 1 code symbol erasure, i.e., $B(i) \in \{0, 1\}$, and each additional length of exactly I code symbol erasures results in each codeword seeing one additional erasure. This equal-erasures effect is illustrated in Figure 4(e), where $B(1) = B(I) = \lfloor B/I \rfloor$ and $B(2) = \cdots = B(I-1) = \lceil B/I \rceil$.

B. Relationship to channel symbol erasure bursts

We now relate the p^J-ary code symbol erasures in (3) to the number of p-ary channel symbol erasures. When $B \geq 2$, all J channel symbols of the middle $B-2$ code symbols are erased and at least 1 channel symbol, and up to all J channel symbols, of each of the two code symbols on the ends are erased, and thus it follows that

$$(B-2)J + 2 \leq B_{ch} \leq BJ, \tag{4}$$

which can also be seen to apply when $B < 2$. Rearranging (4), we have

$$\frac{B_{ch}}{J} \leq B \leq \frac{B_{ch}}{J} + 2 - \frac{2}{J}. \tag{5}$$

Since $B_{ch}/J < \lfloor B_{ch}/J \rfloor + 1$, we have

$$\frac{B_{ch}}{J} \leq B < \left\lfloor \frac{B_{ch}}{J} \right\rfloor + 3 - \frac{2}{J}. \tag{6}$$

Table 1. Allowable values of s.

$B_{ch} \mod J$	Permissible s values
0	0, 1
1	1
2 to $J-1$	1,2

The difference of the last two terms of (6) is at least one and strictly less than three, and since B is an integer, we have only the three possibilities

$$B = \left\lfloor \frac{B_{ch}}{J} \right\rfloor + s, \quad s \in \{0, 1, 2\}. \tag{7}$$

Note that for fixed B_{ch} and J, some values of s may be impossible, depending on the relative alignment of the binary channel symbols to the code symbols. For example, when $J \mid B_{ch}$, we see from (5) that $s \neq 2$. The possible values of s, as a function of $B_{ch} \mod J$, are shown in Table 1. This allows the three possibilities to be reduced to two, as

$$B = \left\lfloor \frac{B_{ch}}{J} \right\rfloor + I_{\{J \nmid B_{ch}\}} + qI_{\{B_{ch} \neq 1 \mod J\}}, \quad q = \{0, 1\} \tag{8}$$

where $I_{\{c\}}$ is the indicator function, equal to 0 if c is false, and 1 if c is true. In (8), q may be 0 or 1 for any values of B_{ch} and $J \geq 2$, and reflects an additional code symbol erasure that may result from the misalignment of the binary channel symbol burst erasure to the code symbols. The exception to this independence of q is that when $J = 1$, we can see from (5) that $q = 0$, so that $B = B_{ch}$. This alignment-dependency is illustrated in Figure 2. Plugging (8) into (3), we have for each i, $1 \leq i \leq I$,

$$B(i) \in \left\{ \left\lfloor \frac{\left\lfloor \frac{B_{ch}}{J} \right\rfloor + I_{\{J \nmid B_{ch}\}} + qI_{\{B_{ch} \neq 1 \mod J\}}}{I} \right\rfloor, \left\lceil \frac{\left\lfloor \frac{B_{ch}}{J} \right\rfloor + I_{\{J \nmid B_{ch}\}} + qI_{\{B_{ch} \neq 1 \mod J\}}}{I} \right\rceil \right\}. \tag{9}$$

In particular, the scenario in which all code symbols involved in the erasure burst are fully erased corresponds to $J \mid B_{ch}$ and $q = 0$, and (9) reduces to

$$B(i) \in \left\{ \left\lfloor \frac{B_{ch}}{JI} \right\rfloor, \left\lceil \frac{B_{ch}}{JI} \right\rceil \right\}, \tag{10}$$

and if in addition JI divides B_{ch}, we have simply

$$B(i) = \frac{B_{ch}}{JI} \tag{11}$$

for all i.

C. Implications for interleaved codes

We are now ready to state the main result of the article, the effect of interleaving on the maximum resolvable erasure burst length.

Theorem 1. *When an (n, k) code over GF(p^J) with maximum resolvable channel symbol erasure burst length $L^{ch}_{\max}$ is block interleaved to depth I, the overall transmission has maximum resolvable channel symbol erasure burst length $L^{ch}_{\max}(I) = L^{ch}_{\max} I + (I-1)(J-1)$ channel symbols, and some channel symbol erasure bursts as long as $L^{ch}_{\max} I + I(J-1)$ are correctable.*

Proof. First consider $J = 1$. We must show that all channel symbol erasure bursts of length $B_{ch} \le L^{ch}_{\max} I$ are correctable. By (3), we have

$$B(i) \le \left\lceil \frac{B}{I} \right\rceil = \left\lceil \frac{B_{ch}}{I} \right\rceil \le L^{ch}_{\max} = L^{c}_{\max}$$

By definition, the code can correct the $L^c_{\max}$ code symbol erasures in each codeword.

In the remainder, we assume $J > 1$. Let $B_{ch} = L^{ch}_{\max} I + I(J-1)$. Using (2), we have

$$B_{ch} = (JL^c_{\max} - (J-1))I + I(J-1) = L^c_{\max} IJ.$$

This is divisible by IJ and if the erasure burst alignment corresponds to $q = 0$ in (9) as well, then (11) holds and we have $B(i) = L^c_{\max}$ for all i. By definition, the code can correct the $L^c_{\max}$ code symbol erasures in each codeword.

To prove the guaranteed-correctability part of the theorem, we consider the two largest values of B_{ch} first, then prove it for smaller values. If $B_{ch} = L^{ch}_{\max} I + (I-1)(J-1)$, then by (2) we have $B_{ch} = L^c_{\max} IJ - J + 1$. Since $B_{ch} = 1 \mod J$, the indicator functions in (8) are 1 and 0, respectively. If $B_{ch} = L^c_{\max} IJ - J$, then $J \mid B_{ch}$ and the indicator functions in (8) are 0 and 1, respectively. In either case, at most one indicator function contributes to B, regardless of q, and using (9) we have, for all i,

$$B(i) \le \left\lceil \frac{\left\lfloor \frac{L^c_{\max} IJ - J + 1}{J} \right\rfloor + 1}{I} \right\rceil = \left\lceil \frac{L^c_{\max} I - 1 + \left\lfloor \frac{1}{J} \right\rfloor + 1}{I} \right\rceil = L^c_{\max}.$$

If $B_{ch} \le L^c_{\max} IJ - J - 1$, then assuming both indicator functions in (8) are 1 yields the bound

$$B(i) \le \left\lceil \frac{\left\lfloor \frac{L^c_{\max} IJ - J - 1)}{J} \right\rfloor + 2}{I} \right\rceil = \left\lceil \frac{L^c_{\max} I - 1 + \left\lfloor -\frac{1}{J} \right\rfloor + 2}{I} \right\rceil = L^c_{\max}.$$

Thus, for any $B_{ch} \le L^{ch}_{\max} IJ - J + 1$, every codeword has at most $L^c_{\max}$ code symbol erasures, and such erasures are always correctable. □

When Theorem 1 is applied to MDS codes, we have the following.

Corollary 1. *An (n, k) MDS code over GF(p^J) with block interleaving to depth I has maximum resolvable channel symbol erasure burst length $L^{ch}_{\max} = (n-k)IJ - J + 1$ channel symbols, and can correct some erasure bursts as long as $(n-k)IJ$ bits.*

Proof. From (1), an MDS code has $L^{ch}_{\max} = J(n-k) - J + 1$. Plugging this into Theorem 1 gives the result. □

Corollary 1 implies that a channel symbol erasure burst with length between $(n-k)IJ - J + 2$ and $(n-k)IJ$ may or may not be correctable by an MDS code, depending on it's alignment with the code symbols. We are able to strengthen Corollary 1, using an impressive theorem due to Xu:

Theorem 2. *[18] An arbitrary (n, k) GF(p^J)-linear MDS code can be modified so that in GF(p) it can correct all burst erasures of length up to $J(n-k)$, while it is still MDS in GF(p^J).*

This result is achieved by demonstrating that with careful ordering of the bit representations of code symbols in GF(p^J), it follows that every consecutive $J(n-k)$ channel symbols in the equivalent (nJ, kJ) code over GF(p) forms a basis for the code. Xu showed that when this modification is performed for RS codes, the result is a generalized RS code, in which each of the n coordinates of a RS code are multiplied by a (possibly non-distinct) nonzero element of GF(p^J). Note that achieving the full burst erasure correcting capability requires a decoder that operates in GF(p), not the higher field GF(p^J).

Using Theorem 2, we have the following corollary to Theorem 1:

Corollary 2. *An arbitrary (n, k) GF(p^J)-linear MDS code can be modified so that when block interleaved to depth I its maximum resolvable channel symbol erasure burst length is $L_{\max}(I) = (n-k)IJ$ bits.*

Proof. When $B_{ch} = (n-k)IJ$, exactly $(n-k)J$ channel symbols are erased in each codeword. □

The Singleton bound implies that no (nIJ, kIJ) code over GF(p) has $L^{ch}_{\max} > (n-k)IJ$. This means that Corollary 2 shows that a block interleaved MDS code, suitably modified using Xu's technique, provides optimal burst erasure protection among all codes of the same rate and total transmission length.

V. Code Comparisons

A. Efficiency

The burst erasure *efficiency* η [19] of an (n, k) code as the ratio of the maximum resolvable burst erasure length divided by $n-k$, the maximum indicated by the Singleton bound. An (n, k) MDS code over GF(2^J) interleaved to depth I, then, is an (nIJ, kIJ) code over GF(2) with efficiency

$$\eta = \frac{(n-k)IJ - J + 1}{(n-k)IJ} = 1 - \frac{J-1}{(n-k)IJ} = 1 - \frac{J-1}{nIJ(1-r)}. \tag{12}$$

The form of (12) suggests defining the *inefficiency* of a code by $\mu = 1 - \eta$, which allows comparisons of codes on a log scale. An inefficiency $\mu = 0$ would correspond to an optimum, MDS code.

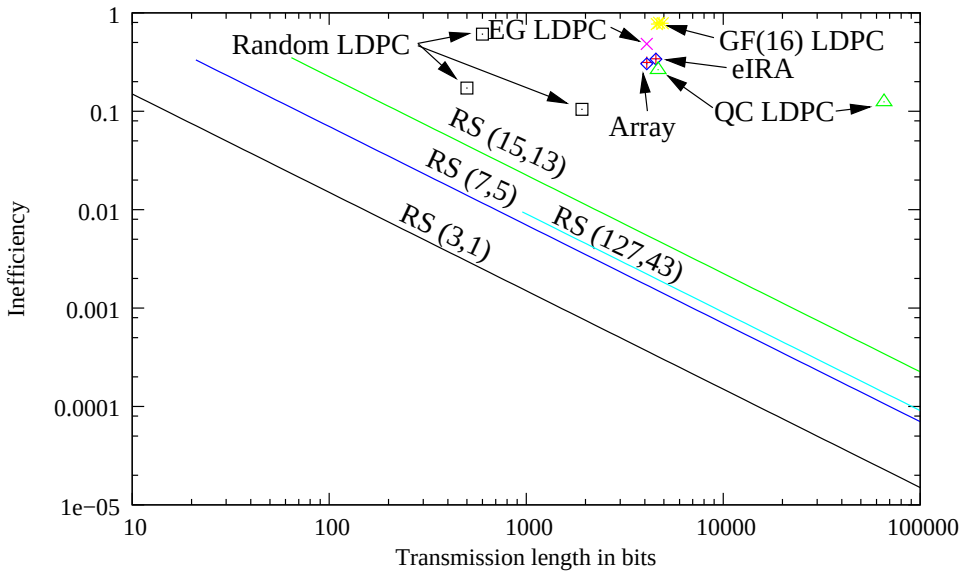

Figure 6. Comparison of inefficiency of various LDPC and interleaved RS codes.

For fixed n and k, the efficiency can be made as close to one as desired simply by using an MDS code and increasing the interleaver length sufficiently. Furthermore, if J is an increasing function of n, then codes are more efficient with *smaller* code lengths. For example, among RS codes with a fixed rate r and fixed total transmission length nIJ, larger values of n in (12) reduces the efficiency, since $J = \log_2(n+1)$ and the denominator of (12) is fixed. This result is consistent with the fact that larger field sizes allow a single channel bit erasure (at the tail of a burst) to affect a larger code symbol.

In Figure 6, we compare the inefficiency of several RS codes with LDPC codes [19] designed for the burst erasure channel. The inefficiency of a block interleaved (3,1) RS code over GF(4)— the shortest nontrivial RS code— indicates an inefficiency about two orders of magnitude better than those of rate 1/3 LDPC codes in [16], which were modified for improved performance on the burst erasure channel. The inefficiencies of a block interleaved (7,5) RS code over GF(8) and a block interleaved (15,13) RS code over GF(15) indicate that increasing the code length does not improve efficiency. Lowering the rate does improve efficiency, as seen with a block interleaved (127,43) RS code over GF(128).

We see that very short RS codes are less than 1% inefficient at many lengths of practical interest.

B. Coding Gain on Noisy Channel

This article only addresses the burst erasure channel, when no other channel impairments are present. This is directly applicable to scenarios of interest to NASA, where an inner FEC code corrects errors and leaves erasures for codewords that could not be corrected. When such a code is not present, an erasures plus noise channel is more appropriate, and

in those cases, block interleaved RS codes are substantially outperformed by LDPC or other codes [19, 5, 10, 16].

C. Capacity bounds

A binary-input *memoryless* (not bursty) erasure channel with erasure rate ϵ has capacity [3]

$$C = 1 - \epsilon \text{ bits per channel use,} \tag{13}$$

where ϵ is the probability of erasure. Define a *genie-aided* encoder as one with perfect a priori knowledge of the positions of erasures. Codes designed for such an (impossible) encoder could not improve upon a non-genie aided encoder utilizing a capacity-approaching code, because according to (13) erasures cause a loss in capacity exactly equal to the fraction of symbols erased. That is, side information about the positions of erasures does not increase capacity. Since the position of the erasures does not affect capacity, it follows that the capacity of the burst erasure channel is the same as the memoryless erasure channel, if the overall probability of erasures is the same.

If a probability of bit error P_b is acceptable, then by the converse to the channel coding theorem for discrete memoryless channels, rates up to

$$r(P_b) = \frac{C}{1 - \mathcal{H}(P_b)} \text{ bits / channel use}$$

are achievable, where $\mathcal{H}(x) = -x \log_2 x - (1-x)\log_2(1-x)$ is the binary entropy function. Equivalently, when coding at rate $r > C$, we have this bound on the probability of bit error:

$$P_b \geq \begin{cases} \mathcal{H}^{-1}\left(1 - \frac{1-\epsilon}{r}\right) & r > 1 - \epsilon \\ 0 & r \leq 1 - \epsilon \end{cases} \tag{14}$$

The weak converse channel coding theorem also bounds the block decoding error (the codeword error rate) of a binary-input erasure channel as [3]

$$P_w \quad \geq \quad \begin{cases} \frac{1}{2}\left(r - C - \frac{1}{n}\right) & r > 1 - \epsilon \\ 0 & r \leq 1 - \epsilon \end{cases} \tag{15}$$

where n is the length of the codewords. This bounds P_w away from zero for sufficiently large n, and hence for all values of n, because if $P_w = 0$ for small n we could use a repetition code to achieve $P_w = 0$ for large n. This results in the simple bound $P_w \geq \frac{1}{2}(r - C)$, which corresponds to guessing, with 50% chance of success, at all bits transmitted at a rate above capacity.

D. Memoryless Erasure Channel

A convenient feature of burst erasures is that they fully erase code symbols over GF(p^J), except possibly at the tails of the burst. This allows the code symbol erasure rate to be approximately equal to the channel symbol erasure rate.

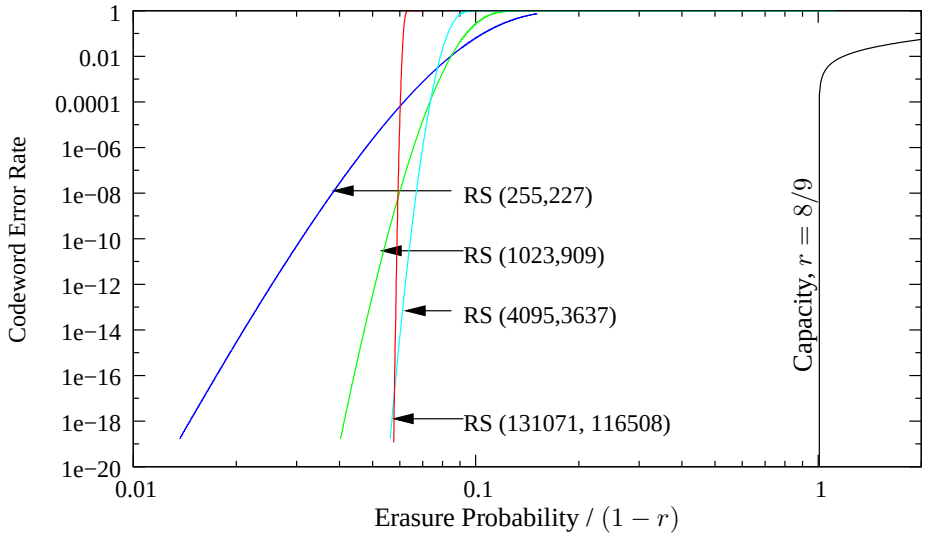

Figure 7. RS codes of rate 8/9, compared to capacity, on a memoryless erasure channel.

When erasures occur independently the code symbol erasure rate may be as much as J times the channel symbol erasure rate. That worst case corresponds to a single channel symbol erasure within each code symbol comprising J channel symbols. This worst case is the typical behavior when $\epsilon \ll 1/J$, and it severely weakens the ability of a nonbinary code, even if MDS, to correct erasures. Hence, we expect that MDS codes over GF(p^J) would not be appropriate for the memoryless erasure channel. This remains true for long codeword lengths, where the value of J may be even higher.

More explicitly, the erasure rate of code symbols in GF(p^J) on a GF(p) erasure channel with erasure probability ϵ is given by

$$P_s = 1 - (1-\epsilon)^J, \tag{16}$$

and the codeword error rate of an MDS code over GF(p^J) is

$$P_w = \sum_{i=n-k+1}^{n} \binom{n}{i} P_s^i (1-P_s)^{n-i} \tag{17}$$

$$= \sum_{i=n-k+1}^{n} \binom{n}{i} \left(1-(1-\epsilon)^J\right)^i (1-\epsilon)^{J(n-i)} \tag{18}$$

The performance of several RS codes of approximately rate 8/9 is shown in Figure 7, along with the capacity of the erasure channel restricted to rate 8/9. As can be seen, RS codes can handle less than 10% of the erasure rate that is possible with a capacity-approaching code. For this channel, LDPC codes discussed elsewhere [19, 5, 10, 16] would be a better alternative.

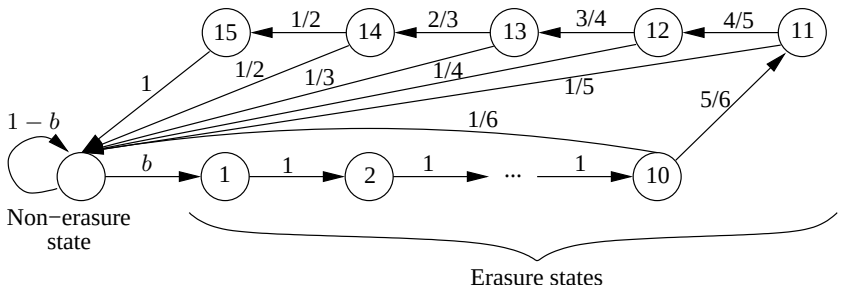

Figure 8. Channel model.

E. Complexity

It has been noted that while the burst erasure efficiency of an (n, k) Reed-Solomon (RS) code is one, its encoding complexity grows quadratically in n, which places a "severe limitation to the codeword length n of practical use" [11]. This sentiment is echoed in [4]. What has not been appreciated, however, is that low, fixed values of (n, k) may be used together with a long interleaver. This results in a near optimal efficiency, as discussed above, and an encoding and decoding complexity that grows linearly with the transmission length nIJ, since with n fixed each codeword will takes a constant complexity to encode. For example, among rate 6/7 codes with total transmission length approximately 5000, an (7,6) RS code with a block interleaver of depth 238 results in a total transmission length of 4998 binary channel symbols and an efficiency of $\eta = 0.997$.

VI. Design Example

We now revisit a design example proposed by Calzolari [4], motivated by an application involving a multiple Mbps deep space link in which in which 8 kbit packets are transferred on a channel that produces erasure bursts with lengths between 10 and 15 packets. The packet erasures result from undecodable frames of FEC encoded data, such as would happen with a weather outage or with loss of synchronization by the receiver.

The channel is modeled as a Gilbert-like channel with one non-erasure (good) state and 15 erasure (bad) states. For each frame, the probability of remaining in the good state is $1 - b$, and the probability of moving to the first bad state is b. Once in the first bad state, a return to the good state occurs in B frames, with B uniformly distributed between 10 and 15. This is modeled by the Markov chain shown in Figure 8. Of interest are codes of rate at least 9/10 and total transmission length of 2 to 4 million. The performance of LDPC codes designed for this application [4] are shown in Figure 9.

To compare this to a RS code of the same rate and overall transmission length, we begin with a (16,15) RS code over GF(16) (i.e., $J = 4$) and shorten it to a (10,9) code. This is then block interleaved to depth 50,000 to achieve a transmission length of two million, and to depth 100,000 to achieve a transmission length of four million. Aligning the packets to the 4-bit code symbols is easy to do, so we do not need to consider that misalignment here.

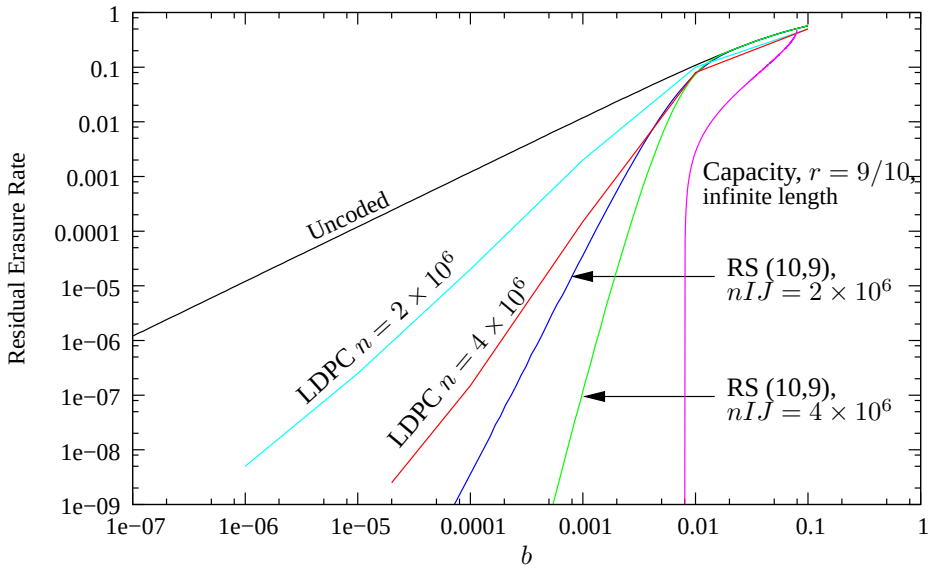

Figure 9. Comparison of interleaved (10,9) RS code to LDPC codes of length 2 million and 4 million.

Thus, the shorter code can correct a channel erasure burst $IJ(n-k) = 200,000$ bits long, and the longer code can correct a burst $400,000$ bits long. Erasures using the channel model above were simulated, and result in the RS performance shown in Figure 9. It can be seen that block interleaving a simple 40-bit RS code yields better performance than long LDPC codes designed for the application, by more than an order of magnitude in b.

Also included in Figure 9 is the channel capacity of rate 9/10 codes, with no constraint on the length of the code. Since the fraction of time in the non-erased state is approximately $12.5b$, we set $\epsilon = 12.5b$ in (14) to determine the capacity.

VII. Conclusions

Block interleaved RS codes are optimal for the burst erasure channel. This optimality does not depend on the length of the RS code, which suggests that short RS codes with long interleavers are a good solution for this problem, as they outperform long LDPC codes designed for this channel and are simpler to decode. When other impairments are on the channel, the low coding gain of RS codes makes them a poor choice compared to other codes.

Acknowledgment

The author thanks Sam Dolinar and Dariush Divsalar for helpful discussions, and reviewer Kar-Ming Cheung for suggesting Fig. 5.

References

[1] Telemetry synchronization and channel coding. Consultative Committee for Space Data Systems (CCSDS) 131.0-B-1. Blue Book. Issue 1. September 2003. http://public.ccsds.org/publications/archive/131x0b1.pdf.

[2] K.S. Andrews, D. Divsalar, S. Dolinar, J. Hamkins, C.R. Jones, and F. Pollara. The development of turbo and LDPC codes for deep-space applications. *Proceedings of the IEEE*, 95(11):2142–2156, Nov. 2007.

[3] Richard E. Blahut. *Theory and Practice of Error Control Codes*. Addison Wesley, Reading, Massachusetts, 1984.

[4] G.P. Calzolari, M. Chiani, F. Chiaraluce, R. Garello, and E. Paolini. Channel coding for future space missions: New requirements and trends. *Proceedings of the IEEE*, 95(11):2157–2170, Nov. 2007.

[5] D. Divsalar, S. Dolinar, and C. Jones. Protograph LDPC codes over burst erasure channels. *Military Communications Conference, 2006. MILCOM 2006*, pages 1–7, Oct. 2006.

[6] Hamid Hemmati, editor. *Deep-Space Optical Communications*. John Wiley & Sons, New York, 2003.

[7] J. Justesen. On the complexity of decoding Reed-Solomon codes (corresp.). *Information Theory, IEEE Transactions on*, 22(2):237–238, Mar 1976.

[8] Lan Lan, Lingqi Zeng, Y.Y. Tai, Lei Chen, Shu Lin, and K. Abdel-Ghaffar. Construction of quasi-cyclic LDPC codes for AWGN and binary erasure channels: A finite field approach. *Information Theory, IEEE Transactions on*, 53(7):2429–2458, July 2007.

[9] S. Lin and D. J. Costello Jr. *Error Control Coding: Fundamentals and Applications*. Prentice-Hall, New Jersey, 1983.

[10] E. Paolini and M. Chiani. Improved low-density parity-check codes for burst erasure channels. *Communications, 2006. ICC '06. IEEE International Conference on*, 3:1183–1188, June 2006.

[11] Enrico Paolini and Marco Chiani. Long erasure correcting codes: the new frontier for zero loss in space applications? *SpaceOps*, pages 1–12, 2006.

[12] Fei Peng, M. Yang, and W.E. Ryan. Simplified eIRA code design and performance analysis for correlated rayleigh fading channels. *Wireless Communications, IEEE Transactions on*, 5(4):720–725, April 2006.

[13] D.V. Sarwate and N.R. Shanbhag. High-speed architectures for Reed-Solomon decoders. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 9(5):641–655, Oct 2001.

[14] S. Shambayati. Ka-band telemetry operations concept: A statistical approach. *Proceedings of the IEEE*, 95(11):2171–2179, Nov. 2007.

[15] Hongxin Song and J.R. Cruz. Reduced-complexity decoding of Q-ary LDPC codes for magnetic recording. *Magnetics, IEEE Transactions on*, 39(2):1081–1087, Mar 2003.

[16] Gokul Sridharan, Abishek Kumarasubramanian, Andrew Thangaraj, and Srikrishna Bhashyam. Optimizing burst erasure correction of LDPC codes by interleaving. *Information Theory, 2008. ISIT 2008. IEEE International Symposium on*, pages 1143–1147, July 2008.

[17] Y.Y. Tai, L. Lan, L. Zeng, S. Lin, and K.A.S. Abdel-Ghaffar. Algebraic construction of quasi-cyclic LDPC codes for the AWGN and erasure channels. *Communications, IEEE Transactions on*, 54(10):1765–1774, Oct. 2006.

[18] L. Xu. Maximizing burst erasure-correction capability of MDS codes. *Communications, IEEE Transactions on*, 54(11):1901–1904, Nov. 2006.

[19] M. Yang and W.E. Ryan. Performance of efficiently encodable low-density parity-check codes in noise bursts on the EPR4 channel. *Magnetics, IEEE Transactions on*, 40(2):507–512, March 2004.