Algebraic Soft-Decision Decoding of Reed–Solomon Codes

Ralf Koetter, *Member, IEEE,* and Alexander Vardy, *Fellow, IEEE*

Abstract—A polynomial-time soft-decision decoding algorithm for Reed–Solomon codes is developed. This list-decoding algorithm is algebraic in nature and builds upon the interpolation procedure proposed by Guruswami and Sudan for hard-decision decoding. Algebraic soft-decision decoding is achieved by means of converting the probabilistic reliability information into a set of interpolation points, along with their multiplicities. The proposed conversion procedure is shown to be asymptotically optimal for a certain probabilistic model. The resulting soft-decoding algorithm significantly outperforms both the Guruswami-Sudan decoding and the generalized minimum distance (GMD) decoding of Reed–Solomon codes, while maintaining a complexity that is polynomial in the length of the code. Asymptotic analysis for a large number of interpolation points is presented, leading to a geometric characterization of the decoding regions of the proposed algorithm. It is then shown that the asymptotic performance can be approached as closely as desired with a list size that does not depend on the length of the code.

Index Terms—Berlekamp–Welch algorithm, Guruswami–Sudan algorithm, list decoding, polynomial interpolation, Reed–Solomon codes, soft-decision decoding.

I. INTRODUCTION

GURUSWAMI and Sudan [27], [14] have recently achieved a breakthrough in algebraic decoding of Reed–Solomon codes. A long-standing open problem in hard-decision decoding of Reed–Solomon codes was that of decoding beyond the error-correction radius. The algorithm of Guruswami and Sudan [14] corrects any fraction of $\tau \leqslant 1 - \sqrt{R}$ erroneous positions for a Reed–Solomon code of rate R. Thus the error-correction capability of this algorithm exceeds the packing radius bound $(1 - R)/2$ for all rates in the interval $[0, 1)$.

Soft-decision decoding of Reed–Solomon codes is, however, an entirely different matter. Although the decoder can be often supplied with reliable soft-decision data relatively easily [5], the high complexity of optimal soft-decision decoding makes full utilization of such data prohibitive. Indeed, all the available *optimal* soft-decoding algorithms for Reed–Solomon codes, such

as [29] and its modifications [7], [23], [24], are nonalgebraic and run in time that scales *exponentially* with the length of the code. This makes the use of such algorithms generally infeasible in practice. An alternative approach to the problem of efficient soft decoding is the generalized minimum distance (GMD) decoding of [10], [11]. While the complexity of GMD decoding is moderate—ultimately, of the same order as that of hard-decision decoding [3], [16], [17], [26], the gains that can be realized by GMD decoding are also moderate (cf. Fig. 1). Thus in light of the ubiquity of Reed–Solomon codes, efficient soft-decision decoding of Reed–Solomon codes is one of the most important problems in coding theory and practice.

Our goal in this paper is to present a polynomial-time soft-decision decoding algorithm for Reed–Solomon codes. This algorithm significantly outperforms both the Guruswami-Sudan list decoding [14] and the GMD-based decoding methods. For example, Fig. 1 shows the performance of the three algorithms for a simple coding scheme: codewords of the $(255, 144, 112)$ Reed–Solomon code over GF(256) are modulated using a 256-QAM signal constellation and transmitted over an AWGN channel. We note that similar coding schemes, although with higher rate Reed–Solomon codes, are in use today on satellite communication channels.

The proposed algorithm is based on the algebraic interpolation techniques developed by Guruswami and Sudan [14], [27]. Guruswami and Sudan also present a weighted version of their list-decoding algorithm in [14]. As pointed out by a referee, this version can be viewed as a soft-decoding algorithm, assuming the interpolation weights can be set "appropriately" given the channel observations. However, the referee also points out that Guruswami and Sudan [14] do assume that interpolation weights are somehow magically given to the algorithm. The main contribution of the present paper is this: we show how the soft-decision reliability information provided by the channel should be translated into algebraic interpolation conditions. Once this is done, we appeal to the interpolation-based techniques of [14].

Specifically, given the channel output vector $(y_1, y_2, \ldots, y_n)$ and the *a posteriori* transition probabilities $\Pr(c_i | y_i)$, we iteratively compute a set of interpolation points along with their multiplicities. We show that, at each step of the computation, this choice of interpolation points is optimal, in a certain precise sense defined in Section IV.

Notably, the algebraic interpolation and factorization techniques of Guruswami and Sudan [27], [14] can be implemented efficiently in polynomial time [1], [8], [9], [21], [20], [25], [32]. Our soft-decision decoding procedure inherits these properties of Guruswami-Sudan decoding. In addition, one of the most

R. Koetter is with the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801 USA (e-mail: koetter@csl.uiuc.edu).

A. Vardy is with the Department of Electrical and Computer Engineering, the Department of Computer Science, and the Center for Wireless Communications, University of California, San Diego, La Jolla, CA 92093 USA (e-mail: vardy@kilimanjaro.ucsd.edu).

Fig. 1. Performance comparison for a simple coding scheme. Codewords of the $(255, 144, 112)$ Reed–Solomon code are modulated using a 256-QAM signal constellation and transmitted over an additive white Gaussian noise (AWGN) channel. At the channel output, soft decisions are quantized to 8 bits. The different curves correspond to the performance achieved by two hard-decision decoding algorithms and two soft-decision decoding algorithms. The two hard-decision algorithms are the conventional Berlekamp-Welch [30] decoding up to half the minimum distance and the list-decoding algorithm by Guruswami and Sudan [14]. For the latter, asymptotic performance is shown, assuming that the multiplicity of interpolation points tends to infinity (cf. Theorem 2). The two soft-decision algorithms are Forney's GMD decoding [10] and the algebraic soft-decision list-decoding algorithm developed here. The curve marked $\triangledown$ describes asymptotic performance for a large number of interpolation points, and hence large list size. However, the curve marked $\circ$ shows that the asymptotic performance can be closely approached with a finite list that is guaranteed to have at most 32 codewords (cf. Section VI).

useful characteristics of our algorithm is a complexity–performance tradeoff that can be chosen freely. In particular, the complexity can be adjusted to any required level of performance within a certain fundamental bound (cf. Theorem 12).

The rest of this paper is organized as follows. The next section contains a brief overview of the Guruswami-Sudan list-decoding algorithm [14]. Section III then sets the ground for algebraic soft-decision decoding of Reed–Solomon codes. In particular, we define the concepts of *score* and *cost* associated with each possible set of interpolation points. We then give a sufficient condition for successful list decoding in terms of these concepts. The core of our soft-decoding algorithm is developed in Section IV, which deals with the computation of (the multiplicities of) the interpolation points. In particular, we show how to iteratively compute the interpolation *multiplicity matrix* so as to maximize the expected score in a certain probabilistic model. Thus Section IV contains the new algorithmic component of this paper. Section V presents an asymptotic performance analysis for our algorithm as the the number of interpolation points approaches infinity. The analysis leads to a simple geometric characterization of the (asymptotic) decoding regions of our algorithm. In Section VI, we show that the asymptotic performance can be approached arbitrarily closely with a list size that depends on the rate but *not* on the the length of the code at hand. We also present simulation results for various list sizes, for both half-rate and high-rate Reed–Solomon codes. In Section VII, we consider channels with memory, and show how the soft-decision decoder should be adapted to deal with such channels. Finally, we conclude with a brief discussion in Section VIII.

II. THE GURUSWAMI-SUDAN LIST-DECODING ALGORITHM

We first set up some of the notation that will be used throughout this work. Let $\mathbb{F}_q$ be the finite field with q elements. The ring of polynomials over $\mathbb{F}_q$ in a variable X is denoted $\mathbb{F}_q[X]$. Reed–Solomon codes are obtained by evaluating certain subspaces of $\mathbb{F}_q[X]$ in a set of points $\mathcal{D} = \{x_1, x_2, \ldots, x_n\} \subseteq \mathbb{F}_q$. Specifically, the Reed–Solomon code $\mathbb{C}_q(n, k)$ of length n and dimension k is defined as follows:

$$\mathbb{C}_q(n, k) \overset{\text{def}}{=} \{ (f(x_1), \ldots, f(x_n)) : x_1, \ldots, x_n \in \mathcal{D},$$

$$f(X) \in \mathbb{F}_q[X], \ \deg f(X) < k \}. \quad (1)$$

The set of polynomials of degree less than k in $\mathbb{F}_q[X]$ is a linear space, which together with the linearity of the evaluation map (1) establishes that $\mathbb{C}_q(n, k)$ is a linear code. The minimum Hamming distance of $\mathbb{C}_q(n, k)$ is $d = n - k + 1$, which follows from the fact that any nonzero polynomial of degree less than k evaluates to zero in less than k positions.

Given an arbitrary vector $\underline{y} \in \mathbb{F}_q^n$, the **hard-decision decoding task** consists of finding the codeword $\underline{c} \in \mathbb{C}_q(n, k)$ such that the Hamming weight $\text{wt}(\underline{e})$ of the error vector $\underline{e} = \underline{y} - \underline{c}$ is minimized. The Berlekamp–Welch algorithm [30] is a well-known algorithm that accomplishes this task, provided $\text{wt}(\underline{e}) < d/2$. Generalizing upon Berlekamp–Welch [30], Sudan [27] as well as Guruswami and Sudan [14] derived a polynomial-time algorithm that achieves error correction substantially beyond half the minimum distance of the code. In the remainder of this section, we describe the essential elements of this algorithm.

Fig. 2. A bound on the number of monomials of (w_X, w_Y)-weighted degree at most δ. Here, $N_{w_X, w_Y}(\delta)$ is the area under the solid line, with each monomial $X^a Y^b$ represented by the unit square whose lower left corner is at the point (a, b). It is easy to see that the triangle of area $\delta^2/2w_X w_Y$ (bounded by the dashed line) is completely enclosed by the solid line.

Definition 1: Let $\mathcal{A}(X,Y) = \sum_{i=0}^{\infty} \sum_{j=0}^{\infty} a_{i,j} X^i Y^j$ be a bivariate polynomial over $\mathbb{F}_q$, and let w_X, w_Y be real numbers. Then, the (w_X, w_Y)-weighted degree of $\mathcal{A}(X,Y)$, denoted $\deg_{w_X, w_Y} \mathcal{A}(X,Y)$, is the maximum over all numbers $iw_X + jw_Y$ such that $a_{i,j} \neq 0$.

The $(1,1)$-weighted degree is simply the *degree* of a bivariate polynomial. The number of monomials of (w_X, w_Y)-weighted degree at most δ is denoted $N_{w_X, w_Y}(\delta)$. Thus

$$N_{w_X, w_Y}(\delta) \overset{\text{def}}{=} \left| \left\{ X^i Y^j : i, j \geqslant 0 \text{ and } iw_X + jw_Y \leqslant \delta \right\} \right|.$$

The following lemma provides a closed-form expression for $N_{w_X, w_Y}(\delta)$ for the case $w_X = 1$. Similar statements can be found in [14], [20], [25], [27], and other papers.

Lemma 1:

$$N_{1,k}(\delta) = \left\lceil \frac{\delta+1}{k} \right\rceil \left(\delta - \frac{k}{2} \left\lfloor \frac{\delta}{k} \right\rfloor + 1 \right).$$

The lemma follows by a straightforward counting of monomials; for a proof, see [14, Lemma 6]. The exact expression in Lemma 1 can be converted into a simple lower bound

$$
\begin{aligned}
N_{1,k-1}(\delta) \\
&= \frac{(\delta+1)^2}{2(k-1)} + \frac{k-1}{2} \left(\left\lceil \frac{\delta+1}{k-1} \right\rceil - \left(\left\lceil \frac{\delta+1}{k-1} \right\rceil - \frac{\delta+1}{k-1} \right)^2 \right) \\
&> \frac{\delta^2}{2(k-1)}.
\end{aligned}
\tag{2}
$$

This is, in fact, a special case of the more general lower bound $N_{w_X, w_Y}(\delta) > \delta^2/2w_X w_Y$. The latter bound can be easily proved using geometric arguments, as shown in Fig. 2.

Given the channel output vector $\underline{y} = \underline{c} + \underline{e} = (y_1, y_2, \ldots, y_n)$ and the corresponding point set $\mathcal{D} = \{x_1, x_2, \ldots, x_n\}$, we consider the set of pairs $\mathcal{P} = \{(x_1, y_2), (x_2, y_2), \ldots, (x_n, y_n)\}$ as *points* in a two-dimensional affine space. Given a point (α, β) and a bivariate polynomial $\mathcal{A}(X,Y)$, we say that (α, β) *lies on* $\mathcal{A}(X,Y)$ if $\mathcal{A}(\alpha, \beta) = 0$. Equivalently, we say that $\mathcal{A}(X,Y)$ *passes through* the point (α, β). Here, we will be interested in bivariate polynomials that not only pass through all the points in $\mathcal{P}$ but do so with high multiplicities.

Definition 2: A bivariate polynomial $\mathcal{A}(X,Y)$ is said to pass through a point (α, β) with multiplicity m if the shifted polynomial $\mathcal{A}(X + \alpha, Y + \beta)$ contains a monomial of degree m and does not contain a monomial of degree less than m. Equivalently, the point (α, β) is said to be a zero of multiplicity m of the polynomial $\mathcal{A}(X,Y)$.

Using a well-known explicit relation (cf. [14]) between the coefficients of a bivariate polynomial $\mathcal{A}(X,Y)$ and the coefficients of the shifted polynomial, we find that Definition 2 imposes the following linear constraints:

$$\sum_{i=k}^{\infty} \sum_{j=l}^{\infty} \binom{i}{k}\binom{j}{l} \alpha^{i-k}\beta^{j-l} a_{i,j} = 0,$$
$$\forall\, k, l \geqslant 0 \text{ such that } k + l < m \tag{3}$$

on the coefficients $a_{i,j}$ of $\mathcal{A}(X,Y)$. Thus $\mathcal{A}(X,Y)$ passes through a given point with multiplicity at least m if and only if $a_{i,j}$ satisfy the $m(m+1)/2$ constraints specified by (3). We are now ready to formulate the first step of the Guruswami-Sudan [14], [27] algorithm.

Interpolation step: Given the set $\mathcal{P}$ points in $\mathbb{F}_q \times \mathbb{F}_q$ and a positive integer m, compute a nontrivial bivariate polynomial $\mathcal{Q}_{\mathcal{P}}(X,Y)$ of minimal $(1, k-1)$- weighted degree that passes through all the points in $\mathcal{P}$ with multiplicity at least m.

If $\deg_{1,k-1} \mathcal{Q}\mathcal{P}(X,Y) = \delta$, then $\mathcal{Q}\mathcal{P}(X,Y)$ may have up to $N_{1,k-1}(\delta)$ nonzero coefficients. These coefficients should be chosen so as to satisfy the $nm(m+1)/2$ linear constraints of type (3), imposed by the interpolation step. This produces a system of $nm(m+1)/2$ linear equations (not all of them necessarily linearly independent) over $\mathbb{F}_q$ in $N_{1,k-1}(\delta)$ unknowns. It is clear that this system has a nonzero solution as long as

$$N_{1,k-1}(\delta) > \frac{nm(m+1)}{2}. \tag{4}$$

For efficient algorithms to solve such a system of linear equations and, hence, accomplish the interpolation step, we refer the reader to [9], [16], [20], [21], [25], [32].

The idea of the Guruswmi-Sudan algorithm [14], [27] is that, under certain constraints on the weight of the error vector, we can read off a list of decoding decisions as factors of $\mathcal{Q}_{\mathcal{P}}(X,Y)$ of type $Y - f(X)$. Thus the second (and last) step of the algorithm is as follows.

Factorization step: Given the bivariate polynomial $\mathcal{Q}_{\mathcal{P}}(X,Y)$, identify all the factors of $\mathcal{Q}_{\mathcal{P}}(X,Y)$ of type $Y - f(X)$ with $\deg f(X) < k$. The output of the algorithm is a list of the codewords that correspond to these factors.

Notice that full factorization of $\mathcal{Q}_{\mathcal{P}}(X,Y)$ is not required to find all the factors of type $Y - f(X)$ with $\deg f(X) < k$. Efficient algorithms to accomplish such "partial factorization" are given in [1], [8], [9], [20], [32]. The eventual decoder output can be taken as the codeword on the list produced at the factorization step that is closest to the received vector $\underline{y}$.

The fundamental question is under which conditions can one guarantee that the correct decoding decision is found on the list. The answer to this question is given in Theorem 2.

Theorem 2: Suppose that a vector $\underline{y}$ and a positive integer m are given. Then the factorization step produces a list that contains all codewords of $\mathbb{C}_q(n, k)$ at distance less than

$$t = n - \left\lfloor \frac{\delta}{m} \right\rfloor > \left\lfloor n\left(1 - \sqrt{R\frac{m+1}{m}} \right) \right\rfloor \quad (5)$$

from $\underline{y}$, where δ is the smallest integer such that $N_{1,k-1}(\delta) > nm(m+1)/2$ and $R = k/n$.

For a proof of (5), see [14], [20]. The inequality in (5) follows from (2). Here, we observe that Theorem 2 is a special case of Theorem 3, which we prove in the next section. Theorem 2 is the main result of Guruswami and Sudan in [14]. The theorem shows that as $m \to \infty$, the algorithm of [14] corrects any fraction of $\tau \leqslant 1 - \sqrt{R}$ erroneous positions.

III. ALGEBRAIC SOFT-DECISION DECODING

In many situations [5], [29], the decoder can be supplied with probabilistic reliability information concerning the received symbols. A decoding algorithm that utilizes such information is generally referred to as a *soft-decision* decoding algorithm. We now specify this notion precisely, in the context of the present paper. First, we define a *memoryless channel*, or simply a *channel*, as a collection of a finite input alphabet $\mathfrak{X}$, an output alphabet $\mathfrak{Y}$, and $|\mathfrak{X}|$ functions

$$f(\cdot|x) : \mathfrak{Y} \to \mathbb{R} \qquad \text{for all } x \in \mathfrak{X} \quad (6)$$

that are assumed to be known to the decoder. We think of channel input and output as random variables $\mathcal{X}$ and $\mathcal{Y}$, respectively, and assume that $\mathcal{X}$ is uniformly distributed over $\mathfrak{X}$. If the channel is continuous (e.g., Gaussian), then $\mathcal{Y}$ is continuous and the $f(\cdot|x)$ are probability-density functions, while if the channel is discrete then $\mathcal{Y}$ is discrete and the $f(\cdot|x)$ are probability-mass functions. In either case, the decoder can easily compute the probability that $\alpha \in \mathfrak{X}$ was transmitted given that $y \in \mathfrak{Y}$ was observed as follows:

$$\Pr\left(\mathcal{X}=\alpha \mid \mathcal{Y}=y\right) = \frac{f(y|\alpha)\Pr(\mathcal{X}=\alpha)}{\sum_{x\in\mathcal{X}} f(y|x)\Pr(\mathcal{X}=x)} = \frac{f(y|\alpha)}{\sum_{x\in\mathcal{X}} f(y|x)} \quad (7)$$

where the second equality follows from the assumption that $\mathcal{X}$ is uniform. For Reed–Solomon codes, the input alphabet is always $\mathfrak{X} = \mathbb{F}_q$. Let $\alpha_1, \alpha_2, \ldots, \alpha_q$ be a fixed ordering of the elements of $\mathbb{F}_q$; this ordering will be implicitly assumed in the remainder of this paper. Given the vector $\underline{y} = (y_1, y_2, \ldots, y_n) \in \mathfrak{Y}^n$ observed at the channel output, we compute

$$\pi_{i,j} \stackrel{\text{def}}{=} \Pr\left(\mathcal{X}=\alpha_i \mid \mathcal{Y}=y_j\right),$$
$$\text{for } i = 1, 2, \ldots, q \text{ and } j = 1, 2, \ldots, n \quad (8)$$

according to the expression in (7). Let Π be the $q \times n$ matrix with entries $\pi_{i,j}$ defined in (8). We will refer to Π as the *reliability matrix* and assume that Π is the input to a soft-decision decoding algorithm. For notational convenience, we will sometimes write $\Pi(\alpha, j)$ to refer to the entry found in the jth column of Π in the row indexed by $\alpha \in \mathbb{F}_q$.

We note that in some applications [5], [31], it is the reliability matrix Π rather than the vector $\underline{y} \in \mathfrak{Y}^n$ that is directly available at the channel output. In many other cases, the channel output alphabet $\mathfrak{Y}$ is quite different from $\mathbb{F}_q$. Thus the first step in

hard-decision decoding is the construction of the **hard-decision vector** $\underline{u} = (u_1, u_2, \ldots, u_n) \in \mathbb{F}_q^n$, where

$$u_j \stackrel{\text{def}}{=} \arg\max_{\alpha \in \mathbb{F}_q} \Pi(\alpha, j), \qquad \text{for } j = 1, 2, \ldots, n. \quad (9)$$

This hard-decision vector is then taken as the channel output $\underline{c} + \underline{e}$ (cf. Section II), thereby converting the channel at hand into a hard-decision channel.

On the other hand, a soft-decision decoder works directly with the probabilities compiled in the reliability matrix Π. If the decoder is algebraic, it must somehow convert these probabilities into algebraic conditions. The algebraic soft-decision Reed–Solomon decoder developed in this paper converts the reliability matrix Π into a *choice of interpolation points and their multiplicities* in the Guruswami–Sudan [14] list-decoding algorithm.

A convenient way to keep track of the interpolation points and their multiplicities is by means of a multiplicity matrix. A *multiplicity matrix* is a $q \times n$ matrix M with nonnegative integer entries $m_{i,j}$. Thus the first step of our decoding algorithm consists of computing the multiplicity matrix M from the reliability matrix Π. This step is discussed in detail in the next section. From there, the soft-decision decoder proceeds as in the Guruswami-Sudan [14] algorithm. In particular, the second step consists of the following.

> **Soft interpolation step:** Given the point set $\mathcal{D}$ and the multiplicity matrix $M = [m_{i,j}]$, compute a nontrivial bivariate polynomial $\mathcal{Q}_M(X, Y)$ of minimal $(1, k-1)$-weighted degree that has a zero of multiplicity at least $m_{i,j}$ at the point (x_j, α_i) for every i, j such that $m_{i,j} \neq 0$.

The third and final step of the algorithm is the factorization step, which is identical to the factorization step of the Guruswami-Sudan algorithm, described in the previous section.

In the remainder of this section, we characterize the conditions under which the decoder will produce the transmitted codeword $\underline{c} \in \mathbb{C}_q(n, k)$, for a *given choice* of interpolation points and their multiplicities (that is, for a given multiplicity matrix M).

Definition 3: Given a $q \times n$ matrix M with nonnegative integer entries $m_{i,j}$, we define the cost of M as follows:

$$\mathcal{C}(M) \stackrel{\text{def}}{=} \frac{1}{2} \sum_{i=1}^{q} \sum_{j=1}^{n} m_{i,j}(m_{i,j} + 1).$$

It is easy to see that the computation of $\mathcal{Q}_M(X, Y)$ is equivalent to solving a system of linear equations of type (3). Since a given zero of multiplicity m imposes $m(m+1)/2$ linear constraints on the coefficients of $\mathcal{Q}_M(X, Y)$, the cost $\mathcal{C}(M)$ is precisely the total number of linear equations. As in (4), we can always find a nonzero solution $\mathcal{Q}_M(X, Y)$ to the soft interpolation task if the $(1, k-1)$-weighted degree δ is large enough, namely, if

$$N_{1,k-1}(\delta) > \mathcal{C}(M) \quad (10)$$

so that the number of degrees of freedom is greater than the number of linear constraints. Thus we define the function

$$\Delta_{w_X, w_Y}(\nu) \stackrel{\text{def}}{=} \min\left\{ \delta \in \mathbb{Z} : N_{w_X, w_Y}(\delta) > \nu \right\}. \quad (11)$$

Observe that $\Delta_{1,k-1}(\nu) < \sqrt{2(k-1)\nu}$ in view of (2). Next, given two $q \times n$ matrices A and B over the same field, we define the inner product

$$\langle A, B \rangle \overset{\text{def}}{=} \text{trace}(AB^T) = \sum_{i=1}^{q} \sum_{j=1}^{n} a_{i,j} b_{i,j}.$$

Finally, it will be convenient to think of the codewords of the Reed–Solomon code $\mathbb{C}_q(n,k)$ as $q \times n$ matrices over the reals. Specifically, any vector $\underline{v} = (v_1, v_2, \ldots, v_n)$ over $\mathbb{F}_q$ can be represented by the $q \times n$ real-valued matrix $[\underline{v}]$ defined as follows: $[\underline{v}]_{i,j} = 1$ if $v_j = \alpha_i$, and $[\underline{v}]_{i,j} = 0$ otherwise. With this notation, we have the following definition.

Definition 4: The score of a vector $\underline{v} = (v_1, v_2, \ldots, v_n)$ over $\mathbb{F}_q$ with respect to a given multiplicity matrix M is defined as the inner product $\mathcal{S}_M(\underline{v}) = \langle M, [\underline{v}] \rangle$.

The following theorem characterizes the set of codewords produced by our soft-decision decoding algorithm for a given multiplicity matrix. Notice that Theorem 2 follows as a special case of Theorem 3 for the multiplicity matrix $M = m [\underline{y}]$.

Theorem 3: Let $\mathcal{C}$ be the cost of a given multiplicity matrix M. Then the polynomial $\mathcal{Q}_M(X, Y)$ has a factor $Y - f(X)$, where $f(X)$ evaluates to a codeword $\underline{c} \in \mathbb{C}_q(n,k)$, if the score of $\underline{c}$ is large enough compared to $\mathcal{C}$, namely, if

$$\mathcal{S}_M(\underline{c}) > \Delta_{1,k-1}(\mathcal{C}). \tag{12}$$

Proof: Let $\underline{c} = (c_1, \ldots, c_n)$ be a codeword of $\mathbb{C}_q(n,k)$, and let $f(X)$ be the polynomial that evaluates to $\underline{c}$. That is, $f(x_j) = c_j$ for all $x_j \in \mathcal{D}$, where $\mathcal{D}$ is the set of points that specify $\mathbb{C}_q(n,k)$ as in (1). Given $\mathcal{Q}_M(X, Y)$, we define the polynomial $g(X) \in \mathbb{F}_q[X]$ as follows:

$$g(X) \overset{\text{def}}{=} \mathcal{Q}_M(X, f(X)).$$

It would clearly suffice to prove that (12) implies that $g(X)$ is the all-zero polynomial, since then $\mathcal{Q}_M(X, Y)$ must be divisible by $Y - f(X)$. To prove that $g(X) \equiv 0$, we will show that $\deg g(X) \leqslant \Delta_{1,k-1}(\mathcal{C})$ and yet $g(X)$ has a factor of degree $\mathcal{S}_M(\underline{c})$. We write

$$\mathcal{S}_M(\underline{c}) = \langle M, [\underline{c}] \rangle = m_1 + m_2 + \cdots + m_n.$$

Thus the polynomial $\mathcal{Q}_M(X, Y)$ passes through the point (x_j, c_j) with multiplicity at least m_j, for $j = 1, 2, \ldots, n$. We will make use of the following lemma.

Lemma 4: Suppose that a bivariate polynomial $\mathcal{Q}(X, Y)$ passes through a point (α, β) in $\mathbb{F}_q \times \mathbb{F}_q$ with multiplicity at least m, and let $p(X)$ be any polynomial in $\mathbb{F}_q[X]$ such that $p(\alpha) = \beta$. Then $\mathcal{Q}(X, p(X))$ is divisible by $(X - \alpha)^m$.

This lemma is identical to [14, Lemma 4], and we omit the proof. Since $f(x_j) = c_j$ for $j = 1, 2, \ldots, n$, it follows from Lemma 4 and the fact that $x_1, x_2, \ldots, x_n$ are all distinct that the polynomial $g(X) = \mathcal{Q}_M(X, f(X))$ is divisible by the product

$$(X - x_1)^{m_1} (X - x_2)^{m_2} \cdots (X - x_n)^{m_n}$$

whose degree is $\mathcal{S}_M(\underline{c})$. We conclude that either $\deg g(X) \geqslant \mathcal{S}_M(\underline{c})$ or $g(X) \equiv 0$. Since $\deg f(X) \leqslant k - 1$, it is easy to see that the degree of $g(X) = \mathcal{Q}_M(X, f(X))$ cannot exceed the $(1, k-1)$-weighted degree of $\mathcal{Q}_M(X, Y)$. Yet it follows from

(10) and (11) that $\deg_{1,k-1} \mathcal{Q}_M(X, Y) \leqslant \Delta_{1,k-1}(\mathcal{C})$. Thus if $g(X) \not\equiv 0$ then $\mathcal{S}_M(\underline{c}) \leqslant \deg g(X) \leqslant \Delta_{1,k-1}(\mathcal{C})$. Hence, (12) implies that $g(X) \equiv 0$. $\square$

Corollary 5: Let $\mathcal{C} = \mathcal{C}(M)$ be the cost of a given multiplicity matrix M. Then, $\mathcal{Q}_M(X, Y)$ has a factor $Y - f(X)$, where $f(X)$ evaluates to $\underline{c} \in \mathbb{C}_q(n,k)$, if $\mathcal{S}_M(\underline{c}) \geqslant \sqrt{2(k-1)\mathcal{C}}$.

Proof: Follows immediately from Theorem 3 and the fact that $\Delta_{1,k-1}(\mathcal{C}) < \sqrt{2(k-1)\mathcal{C}}$ by Lemma 1. $\square$

IV. FROM POSTERIOR PROBABILITIES TO INTERPOLATION POINTS

This section develops the core contribution of our paper: an algorithm that converts posterior probabilities derived from the channel output into a choice of interpolation points and their multiplicities. More specifically, given a reliability matrix Π, as defined in (8), we compute the multiplicity matrix M that serves as input to the soft interpolation step. Let $\mathfrak{M}_{q,n}$ denote the set of all $q \times n$ matrices with nonnegative integer entries $m_{i,j}$, and let $\mathfrak{M}(\mathcal{C})$ be the finite set of all matrices in $\mathfrak{M}_{q,n}$ whose cost is equal to $\mathcal{C}$. Thus

$$\mathfrak{M}(\mathcal{C}) \overset{\text{def}}{=} \left\{ M \in \mathfrak{M}_{q,n} : \frac{1}{2} \sum_{i=1}^{q} \sum_{j=1}^{n} m_{i,j}(m_{i,j}+1) = \mathcal{C} \right\}.$$

In view of Theorem 3, we would like to choose $M \in \mathfrak{M}(\mathcal{C})$ so as to maximize the score of the transmitted codeword $\underline{c} \in \mathbb{C}_q(n,k)$. However, this codeword is obviously unknown to the decoder; only some stochastic information about $\underline{c}$ is available through the observation of the channel output $(y_1, y_2, \ldots, y_n) \in \mathfrak{Y}^n$ and the knowledge of the transition probabilities $\Pr(\mathcal{X} = \alpha \mid \mathcal{Y} = y)$. In fact, as far as the decoder is concerned, the transmitted codeword may be thought of as a random vector $\boldsymbol{\mathcal{X}} = (\mathcal{X}_1, \mathcal{X}_2, \ldots, \mathcal{X}_n)$. We shall specify the distribution of $\boldsymbol{\mathcal{X}}$ shortly (see (15) on the next page and the discussion immediately following this equation). For the time being, observe that, for a given multiplicity matrix M, the score of the transmitted codeword is a function of $\boldsymbol{\mathcal{X}}$ given by $\mathcal{S}_M(\boldsymbol{\mathcal{X}}) = \langle M, [\boldsymbol{\mathcal{X}}] \rangle$.

Thus $\mathcal{S}_M(\boldsymbol{\mathcal{X}})$ is a random variable, and the question is: what is the best choice of a multiplicity matrix $M \in \mathfrak{M}(\mathcal{C})$ in this probabilistic setting? We choose to compute the matrix $M \in \mathfrak{M}(\mathcal{C})$ that maximizes the expected value of $\mathcal{S}_M(\boldsymbol{\mathcal{X}})$. This choice is based on the following considerations. First, this is a reasonable optimization criterion for the probabilistic setup which is the focus of this paper. The obvious alternative is to compute $M \in \mathfrak{M}(\mathcal{C})$ that directly maximizes $\Pr\{\mathcal{S}_M(\boldsymbol{\mathcal{X}}) > \Delta_{1,k-1}(\mathcal{C})\}$. However, this computation appears to be extremely difficult, except for certain special cases of simple channels.

The second reason is this. Theorem 14 of Section V-B shows that this criterion is asymptotically optimal in the following sense. Let P_e denote the probability of decoding failure, defined as the probability that the transmitted codeword $\underline{c}$ is not on the list produced by the soft-decoder. Theorem 14 implies that for every $\varepsilon > 0$, we have $P_e \leqslant \varepsilon$ provided

$$\sqrt{R} \leqslant \frac{\mathsf{E}\{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\}}{\sqrt{2n\mathcal{C}}} - \frac{1}{\sqrt{\varepsilon n}} \tag{13}$$

where $R = k/n$ is the rate of the Reed–Solomon code and $\mathsf{E}\{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\}$ is the expected value of the score for a given multiplicity matrix $M \in \mathfrak{M}(\mathcal{C})$. Thus at least asymptotically for $n \to \infty$, maximizing the expectation of the score allows for reliable transmission at the highest possible rate. One might argue that such asymptotic reasoning has little meaning for Reed–Solomon codes, since $n \leqslant q$. However, the proposed soft-decoding algorithm can be generalized to algebraic-geometry codes, so that $n \to \infty$ makes sense for a fixed q. More importantly, the bound in (13) essentially follows from the fact that the random variable $\mathcal{S}_M(\boldsymbol{\mathcal{X}})$ concentrates about its expected value as n becomes large. We have observed that n does not have to be very large for this concentration to take place.

To proceed, let us define the ***expected score*** with respect to a probability distribution $P(\cdot)$ on the random vector $\boldsymbol{\mathcal{X}} = (\mathcal{X}_1, \mathcal{X}_2, \ldots, \mathcal{X}_n)$ as follows:

$$\mathsf{E}_P\{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\} \stackrel{\text{def}}{=} \sum_{\underline{x} \in \mathfrak{X}^n} \mathcal{S}_M(\underline{x})P(\underline{x}) = \sum_{\underline{x} \in \mathbb{F}_q^n} \sum_{j=1}^{n} M(x_j, j)P(\underline{x})$$

(14)

where $M(x_j, j)$ denotes the entry found in the jth column of M in the row indexed by x_j. It remains to specify $P(\cdot)$. For this purpose, we adopt the product distribution determined by the channel output $(y_1, y_2, \ldots, y_n) \in \mathfrak{Y}^n$, namely

$$P(x_1, x_2, \ldots, x_n) \stackrel{\text{def}}{=} \prod_{j=1}^{n} \Pr\left(\mathcal{X}_j = x_j \mid \mathcal{Y}_j = y_j\right)$$
$$= \prod_{j=1}^{n} \Pi(x_j, j)$$

(15)

where Π is the reliability matrix defined in (8). It is easy to see that this would be the *a posteriori* distribution of $\boldsymbol{\mathcal{X}}$ given the channel observations, if the *a priori* distribution of $\boldsymbol{\mathcal{X}}$ were uniform over the space $\mathbb{F}_q^n$. However, the decoder knows that $\boldsymbol{\mathcal{X}}$ was drawn *a priori* from the code $\mathbb{C}_q(n, k)$ rather than from the entire space $\mathbb{F}_q^n$, and hence the probability model in (15) is suboptimal. Taking this into account results in the probability model given in (47). This model is optimal in that it reflects precisely *all* the information available to the decoder. Unfortunately, this model leads to an intractable optimization problem, as shown in the Appendix (cf. Theorem 20). Thus the remainder of this section is concerned with the computation of the matrix $M(\Pi, \mathcal{C})$ defined as follows:

$$M(\Pi, \mathcal{C}) \stackrel{\text{def}}{=} \operatorname{argmax}_{M \in \mathfrak{M}(\mathcal{C})} \mathsf{E}_P\{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\}$$

(16)

where the expectation is taken with respect to the probability distribution $P(\cdot)$ in (15). We start with the following lemma, which gives a useful expression for the expected score.

Lemma 6: The expected score with respect to the probability distribution in (15) is equal to the inner product of the multiplicity matrix and the reliability matrix, namely

$$\mathsf{E}_P\{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\} = \langle M, \Pi \rangle.$$

Proof: It is easy to see that if $\boldsymbol{\mathcal{X}}$ is distributed according to (15), then Π is precisely the componentwise expected value of $[\boldsymbol{\mathcal{X}}]$. The lemma now follows by linearity of expectation:

$$\mathsf{E}_P\{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\} = \mathsf{E}_P\{\langle M, [\boldsymbol{\mathcal{X}}]\rangle\} = \langle M, \mathsf{E}_P\{[\boldsymbol{\mathcal{X}}]\}\rangle = \langle M, \Pi \rangle.$$

More explicitly, we have

$$\mathsf{E}_P\{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\} = \sum_{\underline{x} \in \mathbb{F}_q^n} \langle M, [\underline{x}]\rangle P(\underline{x})$$
$$= \sum_{\underline{x} \in \mathbb{F}_q^n} \langle M, [\underline{x}]P(\underline{x})\rangle$$
$$= \left\langle M, \sum_{\underline{x} \in \mathbb{F}_q^n} [\underline{x}]P(\underline{x})\right\rangle = \langle M, \Pi \rangle. \qquad \square$$

We will construct $M(\Pi, \mathcal{C})$ iteratively, starting with the all-zero matrix and increasing one of the entries in the matrix at each iteration. Referring to Lemma 6, we see that increasing $m_{i,j}$ from 0 to 1 increases the expected score by $\pi_{i,j}$ while increasing the cost by 1. If we require that $\mathcal{Q}_M(X, Y)$ passes through the same point again (that is, increase $m_{i,j}$ from 1 to 2), then the expected score again grows by $\pi_{i,j}$, but now we have to "pay" two additional linear constraints. In general, increasing $m_{i,j}$ from a to $a + 1$ always increases the expected score by $\pi_{i,j}$ while introducing $a + 1$ additional constraints of type (3). These observations lead to the following algorithm, which greedily maximizes the ratio of the increase in the expected score to the increase in cost at each iteration.

Algorithm A

Input: Reliability matrix Π and a positive integer s, indicating the total number of interpolation points.

Output: Multiplicity matrix M.

Initialization step: Set $\Pi^* := \Pi$ and $M := \text{all}-\text{zero}$ matrix.

Iteration step: Find the position (i, j) of the largest entry $\pi_{i,j}^*$ in Π^*, and set

$$\pi_{i,j}^* := \frac{\pi_{i,j}}{m_{i,j} + 2}$$
$$m_{i,j} := m_{i,j} + 1$$
$$s := s - 1$$

Control step: If $s = 0$, return M; otherwise go to the iteration step.

Let $\mathcal{M}(\Pi, s)$ denote the multiplicity matrix produced by Algorithm A for a given reliability matrix Π and a given number of interpolation points s (counted with multiplicities). The following theorem shows that this matrix is optimal.

Theorem 7: The matrix $\mathcal{M}(\Pi, s)$ maximizes the expected score among all matrices in $\mathfrak{M}_{q,n}$ with the same cost. That is, if $\mathcal{C}$ is the cost of $\mathcal{M}(\Pi, s)$, then

$$\mathcal{M}(\Pi, s) = \operatorname{argmax}_{M \in \mathfrak{M}(\mathcal{C})} \langle M, \Pi \rangle.$$

Proof: With each position (i, j) in the reliability matrix Π, we associate an infinite sequence of rectangles $\mathcal{B}_{i,j,1}, \mathcal{B}_{i,j,2}, \ldots$ indexed by the positive integers. Let $\mathfrak{B}$ denote the set of all such rectangles. For each rectangle $\mathcal{B}_{i,j,l} \in \mathfrak{B}$, we define its $\texttt{length}(\mathcal{B}_{i,j,l}) = l$, $\texttt{height}(\mathcal{B}_{i,j,l}) = \pi_{i,j}/l$, and

$$\texttt{area}(\mathcal{B}_{i,j,l}) = \texttt{length}(\mathcal{B}_{i,j,l}) \cdot \texttt{height}(\mathcal{B}_{i,j,l}) = \pi_{i,j}.$$

For a multiplicity matrix $M \in \mathfrak{M}_{q,n}$, we define the corresponding set of rectangles $\mathfrak{S}(M)$ as

$$\mathfrak{S}(M) \stackrel{\text{def}}{=} \{\mathcal{B}_{i,j,l} : 1 \leqslant i \leqslant q, \ 1 \leqslant j \leqslant n, \ 1 \leqslant l \leqslant m_{i,j}\}. \tag{17}$$

Observe that the number of rectangles in $\mathfrak{S}(M)$ is given by $\sum_{i=1}^{q} \sum_{j=1}^{n} m_{i,j}$, which is precisely the total number of interpolation points imposed by the multiplicity matrix M (counted with multiplicities). Furthermore

$$\mathcal{C}(M) = \sum_{\substack{i=1 \\ j-1}}^{q,n} \frac{m_{i,j}(m_{i,j}+1)}{2} = \sum_{\substack{i=1 \\ j-1}}^{q,n} \sum_{l=1}^{m_{i,j}} l$$

$$= \sum_{\substack{i=1 \\ j-1}}^{q,n} \sum_{l=1}^{m_{i,j}} \texttt{length}(\mathcal{B}_{i,j,l}) = \sum_{\mathcal{B} \in \mathfrak{S}(M)} \texttt{length}(\mathcal{B})$$

$$\langle M, \Pi \rangle = \sum_{\substack{i=1 \\ j-1}}^{q,n} m_{i,j} \cdot \pi_{i,j} = \sum_{\substack{i=1 \\ j-1}}^{q,n} \sum_{l=1}^{m_{i,j}} \pi_{i,j}$$

$$= \sum_{\substack{i=1 \\ j-1}}^{q,n} \sum_{l=1}^{m_{i,j}} \texttt{area}(\mathcal{B}_{i,j,l}) = \sum_{\mathcal{B} \in \mathfrak{S}(M)} \texttt{area}(\mathcal{B}).$$

Thus the cost of M is the total length of all the rectangles in $\mathfrak{S}(M)$ and the expected score $\langle M, \Pi \rangle$ is the total area of all the rectangles in $\mathfrak{S}(M)$. It is intuitively clear that to maximize the total area for a given total length, one has to choose the highest rectangles. This is precisely what Algorithm A does: the algorithm constructs the matrix $\mathcal{M}(\Pi, s)$ that corresponds to the set of s highest rectangles in $\mathfrak{B}$. Indeed, it is easy to see that the ratios $\pi_{i,j}^*$ with which Algorithm A operates are precisely the heights of the rectangles. The algorithm "removes" from $\mathfrak{B}$ and puts in $\mathfrak{S}(M)$ the highest rectangle available at each iteration. It is now obvious that if the *s highest* rectangles in $\mathfrak{B}$ have total length $\mathcal{C}$, then no collection of rectangles of total length at most $\mathcal{C}$ can have a larger total area. $\square$

Although Algorithm A produces an optimal multiplicity matrix $\mathcal{M}(\Pi, s)$ for an arbitrary number of interpolation points s, it cannot be used to solve the optimization problem (16) for an arbitrary value of the cost $\mathcal{C}$. The algorithm computes a solution to (16) only for those costs that are expressible as the total length of the s highest rectangles in $\mathfrak{B}$ for some s. In other words, if $M(\Pi, 1), M(\Pi, 2), M(\Pi, 3), \ldots$ is the infinite sequence of matrices defined by (16) for $\mathcal{C} = 1, 2, 3, \ldots$, then $\mathcal{M}(\Pi, 1), \mathcal{M}(\Pi, 2), \mathcal{M}(\Pi, 3), \ldots$ is a subsequence of this sequence. This subsequence will generally suffice for our purposes.

Remark: Algorithm A can be also used to generate a sequence of multiplicity matrices indexed by a bound on the size of the list produced by the soft-decision decoder. Clearly, the number of factors of $\mathcal{Q}_M(X, Y)$ of type $Y - f(X)$ is bounded by $\deg_{0,1} \mathcal{Q}_M(X, Y)$, and

$$\deg_{0,1} \mathcal{Q}_M(X, Y) \leqslant \left\lfloor \frac{\deg_{1,k-1} \mathcal{Q}_M(X, Y)}{k-1} \right\rfloor \leqslant \left\lfloor \frac{\Delta_{1,k-1}(\mathcal{C})}{k-1} \right\rfloor. \tag{18}$$

Thus given a bound on the desired list size, all one has to do is to keep track of the total cost $\mathcal{C}$, and stop Algorithm A just before the right-hand side of (18) exceeds this bound.

V. ASYMPTOTIC PERFORMANCE ANALYSIS

In the next subsection, we investigate the multiplicity matrix $\mathcal{M}(\Pi, s)$ produced by Algorithm A as $s \to \infty$. We shall see that for $s \to \infty$ this matrix becomes proportional to Π. Based on this result, we derive an asymptotic condition for successful list decoding, and provide a geometric characterization of the asymptotic decoding regions of our algorithm. In a subsequent subsection, we focus instead on long codes—that is, we study the limiting performance of our algorithm as the code length n approaches infinity.

A. Asymptotic Analysis for Large Costs

We start with two simple lemmas. In all of the subsequent analysis, we keep the reliability matrix $\Pi = [\pi_{i,j}]$ fixed, while s ranges over the positive integers. For convenience, we define $\Phi = \{1, 2, \ldots, q\} \times \{1, 2, \ldots, n\}$. Let $\chi(\Pi)$ denote the set of all $(i, j) \in \Phi$ such that $\pi_{i,j} \neq 0$. Let $m_{i,j}(s)$ denote the entries in the matrix $\mathcal{M}(\Pi, s)$ produced by Algorithm A.

Lemma 8: As $s \to \infty$, every nonzero entry in $\mathcal{M}(\Pi, s)$ grows without bound. In other words $m_{i,j}(s) \to \infty$ when $s \to \infty$ for all $(i, j) \in \chi(\Pi)$.

Proof: Define

$$m_{\max}(s) = \max_{(i,j) \in \Phi} m_{i,j}(s)$$

and

$$m_{\min}(s) = \min_{(i,j) \in \chi(\Pi)} m_{i,j}(s).$$

Clearly, it would suffice to show that $m_{\min}(s) \to \infty$ for $s \to \infty$. Notice that

$$s = \sum_{i=1}^{q} \sum_{j=1}^{n} m_{i,j}(s) \leqslant qn \, m_{\max}(s). \tag{19}$$

It follows from (19) that $m_{\max}(s) \to \infty$ as $s \to \infty$. Hence, there exists an infinite integer sequence $s_1, s_2, s_3, \ldots$ defined by the property that $m_{\max}(s_r) = r$ and $m_{\max}(s_r + 1) = r+1$. The iterative nature of Algorithm A implies that for all $s \geqslant 1$, there is exactly one position (i_0, j_0) such that

$$m_{i_0,j_0}(s+1) = m_{i_0,j_0}(s) + 1.$$

We say that (i_0, j_0) is the position ***updated at iteration*** $s+1$ of Algorithm A. This position is distinguished by the property that

$$\frac{\pi_{i_0,j_0}}{m_{i_0,j_0}(s)+1} \geqslant \frac{\pi_{i,j}}{m_{i,j}(s)+1}, \qquad \text{for all } (i,j) \in \Phi. \tag{20}$$

For $r = 1, 2, \ldots$, let (i_r, j_r) denote the position updated at iteration $s_r + 1$ of Algorithm A. Then it follows from (20) and the definition of s_r that

$$m_{i,j}(s_r) + 1 \geqslant \frac{\pi_{i,j}}{\pi_{i_r,j_r}} \Big(m_{\max}(s_r) + 1 \Big)$$

$$\geqslant \frac{\pi_{\min}}{\pi_{\max}} r + \frac{\pi_{\min}}{\pi_{\max}}, \qquad \text{for all } (i,j) \in \chi(\Pi)$$

where $\pi_{\max} = \max_{(i,j) \in \Phi} \pi_{i,j}$ and $\pi_{\min} = \min_{(i,j) \in \chi(\Pi)} \pi_{i,j}$. Denoting by ρ the ratio $\pi_{\min}/\pi_{\max}$, we conclude from the above

that $m_{\min}(s_r) \geqslant \rho r + \rho - 1$. Since ρ is a positive constant while $r \to \infty$ as $s \to \infty$, it follows that $m_{\min}(s)$ grows without bound for $s \to \infty$. $\square$

Henceforth, let (i_s, j_s) denote the position updated at iteration s of Algorithm A, and consider the sequence of ratios of the increase in the expected score to the increase in cost at successive iterations of Algorithm A, namely

$$\theta_1 \stackrel{\text{def}}{=} \frac{\pi_{i_1, j_1}}{m_{i_1, j_1}(1)}, \ \theta_2 \stackrel{\text{def}}{=} \frac{\pi_{i_2, j_2}}{m_{i_2, j_2}(2)}, \ \theta_3 \stackrel{\text{def}}{=} \frac{\pi_{i_3, j_3}}{m_{i_3, j_3}(3)}, \ldots.$$

It follows from (20) that the sequence $\theta_1, \theta_2, \ldots$ is nonincreasing. (Indeed, this was our goal in the design of Algorithm A.) Clearly, $\theta_1 = \pi_{\max}$ while $\lim_{s \to \infty} \theta_s = 0$ by Lemma 8.

Lemma 9: For every positive integer s, there exists a positive constant $K = K(s) \leqslant \pi_{\max}$, such that

$$K\big(m_{i,j}(s) + 1\big) \geqslant \pi_{i,j} \geqslant K m_{i,j}(s), \quad \text{for all } (i, j) \in \Phi. \quad (21)$$

Conversely, for every positive constant $K \leqslant \pi_{\max}$, there exists a positive integer $s = s(K)$ such that (21) holds.

Proof: Given s, we choose $K = K(s)$ so that

$$\theta_{s+1} \leqslant K \leqslant \theta_s$$

which is always possible as the sequence $\theta_1, \theta_2, \ldots$ is nonincreasing. To prove the first inequality in (21), observe that

$$K \geqslant \theta_{s+1} = \frac{\pi_{i_{s+1}, j_{s+1}}}{m_{i_{s+1}, j_{s+1}}(s+1)}$$

$$= \frac{\pi_{i_{s+1}, j_{s+1}}}{m_{i_{s+1}, j_{s+1}}(s) + 1}$$

$$\geqslant \frac{\pi_{i,j}}{m_{i,j}(s) + 1}, \quad \text{for all } (i, j) \in \Phi$$

where the last inequality follows from (20). The second inequality in (21) holds vacuously if $m_{i,j}(s) = 0$, so assume that $m_{i,j}(s) \geqslant 1$. This assumption implies that position (i, j) was updated at least once, and we let $s^* \leqslant s$ denote the number of the *most recent* iteration of Algorithm A at which position (i, j) was updated. Then

$$K \leqslant \theta_s \leqslant \theta_{s^*} = \frac{\pi_{i,j}}{m_{i,j}(s^*)} = \frac{\pi_{i,j}}{m_{i,j}(s)}$$

where the last equality follows from the fact that position (i, j) was *not* updated since iteration s^*. Finally, given $0 < K \leqslant \pi_{\max}$, we choose $s = s(K)$ so that $\theta_{s+1} < K \leqslant \theta_s$ once again. This choice is possible because the sequence $\theta_1, \theta_2, \ldots$ is nonincreasing, $\theta_1 = \pi_{\max}$, and $\lim_{s \to \infty} \theta_s = 0$. The proof then remains exactly the same, except that the first inequality in (21) can be now strengthened to a strict inequality. $\square$

Since (21) holds for all $(i, j) \in \Phi$, both inequalities in (21) remain valid under summation over all $(i, j) \in \Phi$. Thus it follows from Lemma 9 that

$$\sum_{(i,j) \in \Phi} K\big(m_{i,j}(s) + 1\big) \geqslant \sum_{(i,j) \in \Phi} \pi_{i,j} \geqslant \sum_{(i,j) \in \Phi} K m_{i,j}(s). \quad (22)$$

These inequalities lead to upper and lower bounds on the constant $K(s)$ in Lemma 9. Since $\sum_{(i,j) \in \Phi} m_{i,j}(s) = s$ while $\sum_{(i,j) \in \Phi} \pi_{i,j} = n$, we conclude from (22) that

$$\frac{n}{s} \geqslant K(s) \geqslant \frac{n}{s + qn}. \quad (23)$$

Next, we define the ***normalized multiplicity matrix*** $\mathcal{M}'(\Pi, s) = [\mu_{i,j}(s)]$ and the ***normalized reliability matrix*** $\Pi' = [\pi'_{i,j}]$ as follows: $\mu_{i,j}(s) = m_{i,j}(s)/s$ and $\pi'_{i,j} = \pi_{i,j}/n$ for all $(i, j) \in \Phi$. It is clear from these definitions that $\langle \mathcal{M}', \mathbf{1} \rangle = \langle \Pi', \mathbf{1} \rangle = 1$, where $\mathbf{1}$ denotes the all-one matrix. The following theorem is the key result of this subsection: the theorem shows that the optimal multiplicity matrix $\mathcal{M}(\Pi, s)$ becomes proportional to Π as $s \to \infty$.

Theorem 10: As $s \to \infty$, the normalized multiplicity matrix converges to the normalized reliability matrix

$$\mathcal{M}'(\Pi, s) \longrightarrow_{s \to \infty} \Pi'.$$

In other words, for every $\varepsilon > 0$, there exists an s_0 such that for all $s \geqslant s_0$ we have

$$\left| \pi'_{i,j} - \mu_{i,j}(s) \right| = \left| \frac{\pi_{i,j}}{n} - \frac{m_{i,j}(s)}{s} \right| \leqslant \varepsilon, \quad \text{for all } (i, j) \in \Phi. \quad (24)$$

Proof: It follows from Lemma 9 that for all s, there exists a constant $K(s)$ such that $1 \geqslant \pi_{i,j}/K(s) - m_{i,j}(s) \geqslant 0$ for all $(i, j) \in \Phi$. Dividing this inequality by s, we obtain

$$\frac{1}{s} \geqslant \frac{\pi_{i,j}}{s K(s)} - \mu_{i,j}(s) \geqslant 0. \quad (25)$$

From the bounds on $K(s)$ in (23), we conclude that

$$\pi'_{i,j} \leqslant \pi_{i,j}/s K(s) \leqslant \pi'_{i,j} + q \pi_{i,j}/s.$$

Combining this with (25), we get

$$\frac{1}{s} \geqslant \pi'_{i,j} - \mu_{i,j}(s) \geqslant -\frac{q \pi_{\max}}{s}. \quad (26)$$

It follows that for all $s \geqslant \max\{1/\varepsilon, \pi_{\max} q/\varepsilon\} = \pi_{\max} q/\varepsilon$, the bound in (24) holds for all $(i, j) \in \Phi$. Thus $s_0 = \lceil \pi_{\max} q/\varepsilon \rceil$. $\square$

Asymptotically, for a large number s of interpolation points (and, hence, for a large cost), a constraint on the cost

$$\mathcal{C}(M) = (\langle M, M \rangle + \langle M, \mathbf{1} \rangle)/2$$

is equivalent to a constraint on the L_2-norm $\sqrt{\langle M, M \rangle}$ of the multiplicity matrix. Obviously, for a fixed norm $\sqrt{\langle M, M \rangle}$, maximizing the expected score $\langle M, \Pi \rangle$ is equivalent to maximizing the correlation between M and Π, which is clearly achieved by letting M be proportional to Π. This intuition confirms the result established in Theorem 10.

Remark: Finding the optimal multiplicity matrix $\mathcal{M}(\Pi, s)$ can be viewed as a gambling problem. Assume that a gambler has a certain wealth in the form of a maximal number of linear constraints the gambler can satisfy. The matrix Π provides all the information the gambler can use in order to place bets on interpolation points with the goal of maximizing the return, which is the score of the transmitted codeword. In this context, Theorem 10 shows that proportional betting is the asymptotically optimal gambling strategy. Proportional betting is known [6] to be the optimal strategy in the context of a *fair horse race*. However, these results do not appear to be related to Theorem 10 in an obvious way.

We conclude this section with a geometric characterization of the (asymptotic) decoding regions of our soft-decision decoding algorithm. To start with, the following simple lemma essentially recasts Theorem 3 in slightly different terms.

Lemma 11: For a given multiplicity matrix M, the algebraic soft-decision decoding algorithm outputs a list that contains a codeword $\underline{c} \in \mathbb{C}_q(n, k)$ if

$$\frac{\langle M, [\underline{c}] \rangle}{\sqrt{\langle M, M \rangle + \langle M, \mathbf{1} \rangle}} \geqslant \sqrt{k-1}. \tag{27}$$

The lemma follows from Corollary 5 by observing that

$$2\mathcal{C}(M) = \langle M, M \rangle + \langle M, \mathbf{1} \rangle.$$

Theorem 10 and Lemma 11 now lead to a precise characterization of the performance limits of our algorithm as the number of interpolation points approaches infinity. In the following theorem and its corollaries, $o(1)$ denotes a function of s that tends to zero as $s \to \infty$.

Theorem 12: The algebraic soft-decision decoding algorithm outputs a list that contains a codeword $\underline{c} \in \mathbb{C}_q(n, k)$ if

$$\frac{\langle \Pi, [\underline{c}] \rangle}{\sqrt{\langle \Pi, \Pi \rangle}} \geqslant \sqrt{k-1} + o(1). \tag{28}$$

Proof: Substituting the optimal multiplicity matrix $\mathcal{M}(\Pi, s)$ in (27) and normalizing (dividing by s the numerator and the denominator), we obtain the equivalent condition

$$\frac{\langle \mathcal{M}'(\Pi, s), [\underline{c}] \rangle}{\sqrt{\langle \mathcal{M}'(\Pi, s), \mathcal{M}'(\Pi, s) \rangle + \frac{1}{s}}} \geqslant \sqrt{k-1}. \tag{29}$$

It follows from Theorem 10 that for $s \to \infty$, one can replace $\mathcal{M}'(\Pi, s)$ in (29) by Π', which upon renormalization yields (28). More explicitly, we have

$$\frac{\langle \mathcal{M}'(\Pi, s), [\underline{c}] \rangle}{\sqrt{\langle \mathcal{M}'(\Pi, s), \mathcal{M}'(\Pi, s) \rangle + \frac{1}{s}}}$$

$$\geqslant \frac{\langle \Pi, [\underline{c}] \rangle - \frac{n^2}{s}}{\sqrt{\langle \Pi, \Pi \rangle + \frac{(2q\pi_{\max}+1)n^2}{s} + \frac{q^3 n^3 \pi_{\max}^2}{s^2}}}$$

$$= \frac{\langle \Pi, [\underline{c}] \rangle}{\sqrt{\langle \Pi, \Pi \rangle}} + o(1)$$

where the first inequality follows from (26) after some straightforward manipulations. In conjunction with (29), this completes the proof. $\square$

Note that decoding under the asymptotic condition (28) can be directly achieved by choosing large integer multiplicities that are in proportion to the entries of the reliability matrix Π. What is particularly interesting about Theorem 12 is that it shows that the greedy iterative algorithm (Algorithm A) asymptotically achieves this limiting performance.

Finally, Theorem 12 has an especially nice interpretation if the reliability matrix Π and the codeword $[\underline{c}]$ are viewed as vectors in the qn-dimensional Euclidean space $\mathbb{R}^{qn}$.

Corollary 13: Let $\beta = \beta(\Pi, \underline{c})$ denote the angle in $\mathbb{R}^{qn}$ between Π and $[\underline{c}]$. Then the algebraic soft-decision decoding algorithm outputs a list that contains $\underline{c}$ if

$$\cos \beta \geqslant \sqrt{R} + o(1).$$

Proof: Follows directly from Theorem 12 and the identity $\langle \Pi, [\underline{c}] \rangle = \sqrt{n \langle \Pi, \Pi \rangle} \cos \beta$. $\square$

Thus the asymptotic decoding regions of our algorithm are *spherical cones* in the Euclidean space $\mathbb{R}^{qn}$, extending from the origin to the surface of a sphere $\mathbf{S}$ of radius $\sqrt{n}$. The code-word $[\underline{c}]$ is a point of $\mathbf{S}$, and the line connecting the origin to this point constitutes the central axis of the spherical cone. The angle of each spherical cone is $\cos^{-1} \sqrt{R}$. Notice that since the algorithm is a list-decoding algorithm, its decoding regions are not disjoint: the spherical cones of angle $\cos^{-1} \sqrt{R}$ are overlapping. Also notice that we are concerned only with the positive 2^{qn}-part of the Euclidean space $\mathbb{R}^{qn}$ (which consists of points with all coordinates nonengative), since all the entries of both Π and $[\underline{c}]$ are nonnegative.

It follows from Theorem 2 that the asymptotic (for $m \to \infty$) decoding regions of the Guruswami-Sudan [14] algorithm are *spherical caps* on the surface of $\mathbf{S}$ of the same spherical angle $\cos^{-1} \sqrt{R}$, but the decoding process involves projecting Π onto a point $[y]$ on the surface of $\mathbf{S}$ in a nonlinear fashion, according to (9). Finally, the decoding regions of conventional Berlekamp-Welch [30] hard-decision decoding are also spherical caps on the surface of $\mathbf{S}$ and the same nonlinear projection is employed, but the spherical angle of these caps is only $\cos^{-1}(\frac{1+R}{2})$, and they are nonoverlapping.

B. Asymptotic Analysis for Long Codes

As noted in Section IV, from the point of view of the receiver, the transmitted codeword is a random vector $\boldsymbol{\mathcal{X}}$ whose *a posteriori* probability distribution is given by (15). For notational convenience, let us introduce two random variables

$$\mathcal{Z} \stackrel{\text{def}}{=} \langle M, [\boldsymbol{\mathcal{X}}] \rangle \quad \text{and} \quad \mathcal{Z}^* \stackrel{\text{def}}{=} \frac{\langle M, [\boldsymbol{\mathcal{X}}] \rangle}{\sqrt{n \langle M, M \rangle + n \langle M, \mathbf{1} \rangle}}.$$

The key result of this subsection is the following theorem which shows that as $n \to \infty$, the random variable $\mathcal{Z}^*$ converges to its expected value.

Theorem 14: Suppose that a $q \times n$ reliability matrix Π is given, and let M be an arbitrary $q \times n$ multiplicity matrix. Then for any $\varepsilon > 0$, we have

$$\Pr\{|\mathcal{Z}^* - \mathsf{E}\{\mathcal{Z}^*\}| \geqslant \varepsilon\} \leqslant \frac{1}{n\varepsilon^2}. \tag{30}$$

Proof: Consider the random variable

$$\mathcal{Z} = \langle M, [\boldsymbol{\mathcal{X}}] \rangle = M(\mathcal{X}_1, 1) + \cdots + M(\mathcal{X}_n, n)$$

and define $\mathcal{Z}_j = M(\mathcal{X}_j, j)$ for $j = 1, 2, \ldots, n$. Thus $\mathcal{Z}_j$ is the entry found in the jth column of M in the row indexed by $\mathcal{X}_j$. The distribution of $\mathcal{X}_j$, computed by marginalizing the distribution of $\boldsymbol{\mathcal{X}}$ in (15), is given by

$$\Pr\left(\mathcal{X}_j = \alpha_i\right) = \pi_{i,j}, \text{ for } i = 1, 2, \ldots, q \text{ and } j = 1, 2, \ldots, n.$$

Using this distribution, we find that $\mathsf{E}\{\mathcal{Z}_j\} = \sum_{i=1}^{q} m_{i,j} \pi_{i,j}$ and $\mathsf{E}\{\mathcal{Z}_j^2\} = \sum_{i=1}^{q} m_{i,j}^2 \pi_{i,j}$. The key observation is this: since the random variables $\mathcal{X}_1, \mathcal{X}_2, \ldots, \mathcal{X}_n$ are independent, so are $\mathcal{Z}_1, \mathcal{Z}_2, \ldots, \mathcal{Z}_n$. Hence,

$$\text{Var}(\mathcal{Z}) = \sum_{j=1}^{n} \text{Var}(\mathcal{Z}_j)$$

$$= \sum_{j=1}^{n} \left(\sum_{i=1}^{q} m_{i,j}^2 \pi_{i,j} - \left(\sum_{i=1}^{q} m_{i,j} \pi_{i,j} \right)^2 \right)$$

$$\leqslant \sum_{j=1}^{n} \sum_{i=1}^{q} m_{i,j}^2 = \langle M, M \rangle.$$

The theorem now follows by a straightforward application of the Chebycheff inequality [22, p. 193] to the random variable $\mathcal{Z}^*$. Thus

$$
\Pr\left\{|\mathcal{Z}^* - \mathsf{E}\{\mathcal{Z}^*\}| \geqslant \varepsilon\right\} \leqslant \frac{\mathrm{Var}(\mathcal{Z}^*)}{\varepsilon^2}
$$
$$
= \frac{\mathrm{Var}(\mathcal{Z})}{\varepsilon^2\left(n\langle M, M\rangle + n\langle M, \mathbf{1}\rangle\right)} \leqslant \frac{1}{n\varepsilon^2}. \qquad \square
$$

Remark: The proof of Theorem 14 is essentially similar to a well-known proof of the weak law of large numbers. We note that using the strong law of large numbers, it is possible to show that as $n \to \infty$, the random variable $\mathcal{Z}^*$ equals its expectation with probability 1.

We can use Theorem 14 to derive a relationship between the probability P_e of list-decoding failure, the expected score, and the rate R of the Reed–Solomon code. Indeed, we have

$$
P_e \leqslant \Pr\left\{\mathcal{Z} \leqslant \Delta_{1,k-1}(\mathcal{C})\right\} \tag{31}
$$

since $\mathcal{Z} = \mathcal{S}_M(\boldsymbol{\mathcal{X}})$ by defintion, and the "score condition" $\mathcal{S}_M(\underline{c}) > \Delta_{1,k-1}(\mathcal{C})$ is sufficient for successful decoding by Theorem 3. It follows from (31) that

$$
P_e \leqslant \Pr\left\{\mathcal{Z} \leqslant \sqrt{2k\mathcal{C}}\right\} = \Pr\left\{\mathcal{Z}^* \leqslant \sqrt{R}\right\} \tag{32}
$$

since $\Delta_{1,k-1}(\mathcal{C}) < \sqrt{2(k-1)\mathcal{C}} < \sqrt{2k\mathcal{C}}$ in view of Lemma 1 and (2). In conjuction with (32), Theorem 14 immediately implies the following. For all $\varepsilon > 0$, if

$$
\sqrt{R} \leqslant \mathsf{E}\{\mathcal{Z}^*\} - \frac{1}{\sqrt{\varepsilon n}}
$$

then $P_e \leqslant \varepsilon$. This is precisely condition (13) relating the expected score to the probability of decoding failure, discussed in the fourth paragraph of Section IV.

VI. PERFORMANCE ANALYSIS FOR A FIXED LIST SIZE

In this section, we study the performance achieved by our soft-decoding algorithm under a constraint which guarantees that the number of codewords on the list produced by the decoder does not exceed a given bound L. The key analytical result in this section is Theorem 17. This theorem extends Theorem 12 by providing a bound on how quickly the decoding algorithm converges to the asymptotic performance as a function of L. The analytical results are confirmed by simulations for both high-rate and low-rate codes.

We start with two lemmas. As observed in Section IV, the number of codewords on the list produced by the soft-decision decoder is upper-bounded by $\deg_{0,1}\mathcal{Q}_M(X,Y)$, where $\mathcal{Q}_M(X,Y)$ is the interpolation polynomial. This leads to the following lemma.

Lemma 15: The number of codewords on the list produced by the soft-decision decoder for a given multiplicity matrix M does not exceed

$$
\mathcal{L}_k(M) \stackrel{\mathrm{def}}{=} \frac{\sqrt{\langle M, M\rangle + \langle M, \mathbf{1}\rangle}}{\sqrt{k-1}}. \tag{33}
$$

Proof: The size of the list is at most $\deg_{0,1}\mathcal{Q}_M(X,Y)$. By the definition of weighted-degree, we have

$$
\deg_{0,1}\mathcal{Q}_M(X,Y) \leqslant \frac{\deg_{1,k-1}\mathcal{Q}_M(X,Y)}{k-1}.
$$

Now

$$
\frac{\deg_{1,k-1}\mathcal{Q}_M(X,Y)}{k-1} \leqslant \frac{\Delta_{1,k-1}(\mathcal{C})}{k-1} \leqslant \frac{\sqrt{\langle M, M\rangle + \langle M, \mathbf{1}\rangle}}{\sqrt{k-1}}
$$

where the first inequality follows from (10) and (11), while the second inequality follows from Lemma 1, the definition of the cost $\mathcal{C} = \mathcal{C}(M)$, and (11). $\qquad \square$

Let Π be a given reliability matrix, and let $\mathcal{M}(\Pi, s)$ be the corresponding multiplicity matrix produced by Algorithm A. For convenience, we define $\mathcal{M}(\Pi, 0)$ as the all-zero $q \times n$ matrix. Let $\mathcal{J}$ denote a $q \times n$ matrix all of whose entries are nonnegative real numbers not exceeding 1. We write $\mathcal{J}^*$ instead of $\mathcal{J}$ if all the entries in the matrix are strictly less than 1.

Lemma 16: For every positive real number λ, there exists a nonnegative integer s, such that the matrix $\mathcal{M}(\Pi, s)$ can be written as

$$
\mathcal{M}(\Pi, s) = \lfloor \lambda\Pi \rfloor \stackrel{\mathrm{def}}{=} \lambda\Pi - \mathcal{J}^*. \tag{34}
$$

Conversely, for every nonnegative integer s, there exists a positive real λ such that (34) holds, possibly with $\mathcal{J}^*$ replaced by $\mathcal{J}$.

Proof: As before, let $\pi_{\max}$ be the largest entry in Π. If $\lambda < \pi_{\max}^{-1}$, then $\mathcal{M}(\Pi, 0)$ satisfies (34). Otherwise, set $K = \lambda^{-1}$, so that $0 < K \leqslant \pi_{\max}$. We know from Lemma 9 and its proof that there exists a positive integer s, such that

$$
\frac{\pi_{i,j}}{K} - 1 < m_{i,j}(s) \leqslant \frac{\pi_{i,j}}{K}, \qquad \text{for all } (i,j) \in \Phi. \tag{35}
$$

It follows from (35) that $\mathcal{M}(\Pi, s)$ is of the form (34). Conversely, given $\mathcal{M}(\Pi, s)$, we take $\lambda = K^{-1}$, where

$$
K = K(s) \leqslant \pi_{\max}
$$

is the constant derived in Lemma 9. $\qquad \square$

Given $M = \mathcal{M}(\Pi, s)$, let λ be a positive real constant such that $M = \lambda\Pi - \mathcal{J}$. Such a constant exists by Lemma 16. Then

$$
\langle M, M\rangle + \langle M, \mathbf{1}\rangle = \lambda^2\langle \Pi, \Pi\rangle - \lambda\langle \Pi, 2\mathcal{J}-\mathbf{1}\rangle + \langle \mathcal{J}, \mathcal{J}-\mathbf{1}\rangle. \tag{36}
$$

We can use (36) and (33) to derive an expression for λ in terms of Π, $\mathcal{J}$, and $\mathcal{L}_k(M)$. Equating the right-hand side of (36) to $(k-1)\mathcal{L}_k^2(M)$, we obtain a quadratic equation in λ. Since $\langle \Pi, \Pi\rangle > 0$ and $\langle \mathcal{J}, \mathcal{J}-\mathbf{1}\rangle \leqslant 0$, this equation has one positive root and one negative root. Solving for the unique positive root yields

$$
\lambda = \frac{\langle \Pi, 2\mathcal{J}-\mathbf{1}\rangle}{2\langle \Pi, \Pi\rangle}
$$
$$
+ \sqrt{\frac{\langle \Pi, 2\mathcal{J}-\mathbf{1}\rangle^2}{4\langle \Pi, \Pi\rangle^2} + \frac{\langle \mathcal{J}, \mathbf{1}-\mathcal{J}\rangle}{\langle \Pi, \Pi\rangle} + \frac{(k-1)\mathcal{L}_k^2(M)}{\langle \Pi, \Pi\rangle}}. \tag{37}
$$

Suppose now that we are given a positive integer L and would like to guarantee that the number of codewords on the list produced by the soft-decision decoder does not exceeed L. In view

of Lemma 15, we can do so by computing $\mathcal{L}_k(M)$ at each iteration of Algorithm A, and stopping the algorithm just before $\mathcal{L}_k(M)$ equals or exceeds $L+1$. At this point

$$L \leqslant \mathcal{L}_k(M) < L+1 \qquad (38)$$

and since the number of codewords on the list produced by the decoder is an integer not exceeding $\mathcal{L}_k(M)$, this number is at most L. We will refer to this decoding procedure[1] as *algebraic soft-decoding with list size limited to L.*

Theorem 17: Algebraic soft-decoding with list size limited to L produces a list that contains a codeword $\underline{c} \in \mathbb{C}_q(n,k)$ if

$$\frac{\langle \Pi, [\underline{c}] \rangle}{\sqrt{\langle \Pi, \Pi \rangle}} \geqslant \frac{\sqrt{k-1}}{1 - \frac{1}{L}\left(\frac{1}{R^*} + \frac{\sqrt{q}}{2\sqrt{R^*}}\right)} = \sqrt{k-1} \cdot \left(1 + O\left(\frac{1}{L}\right)\right) \qquad (39)$$

where Π is the reliability matrix derived from the channel output, $R^* = (k-1)/n$ is the rate of $\mathbb{C}_q(n,k-1)$, and the constant in $O(\cdot)$ depends only on R^* and q.

Proof: Writing $M = \lambda \Pi - \mathcal{J}$ as in Lemma 16 and using the definition of $\mathcal{L}_k(M)$ in (33), we can recast the sufficient condition (27) of Lemma 11 in the following way:

$$\frac{\langle \Pi, [\underline{c}] \rangle}{\sqrt{\langle \Pi, \Pi \rangle}} \cdot \left(\lambda - \frac{\langle \mathcal{J}, [\underline{c}] \rangle}{\langle \Pi, [\underline{c}] \rangle}\right) \cdot \frac{\sqrt{\langle \Pi, \Pi \rangle}}{\mathcal{L}_k(M)\sqrt{k-1}} \geqslant \sqrt{k-1}. \quad (40)$$

Using the expression for λ in (37), we now express the factor multiplying $\langle \Pi, [\underline{c}] \rangle / \sqrt{\langle \Pi, \Pi \rangle}$ on the left-hand side of (40) as $\mathcal{F}_1(\Pi, L) - \mathcal{F}_2(\Pi, L) - \mathcal{F}_3(\Pi, \underline{c}, L)$, where

$$\mathcal{F}_1(\Pi, L)$$
$$\stackrel{\text{def}}{=} \sqrt{1 + \frac{\langle \Pi, 2\mathcal{J}-\mathbf{1}\rangle^2}{4 \langle \Pi, \Pi \rangle \mathcal{L}_k^2(M)(k-1)} + \frac{\langle \mathcal{J}, \mathbf{1}-\mathcal{J}\rangle}{\mathcal{L}_k^2(M)(k-1)}} \geqslant 1 \qquad (41)$$

$$\mathcal{F}_2(\Pi, L) \stackrel{\text{def}}{=} \frac{1}{\mathcal{L}_k(M)\sqrt{k-1}} \cdot \frac{\langle \Pi, \mathbf{1}-2\mathcal{J}\rangle}{2\sqrt{\langle \Pi, \Pi \rangle}}$$
$$\leqslant \frac{1}{\mathcal{L}_k(M)\sqrt{k-1}} \cdot \frac{n}{2\sqrt{n/q}} \qquad (42)$$

and

$$\mathcal{F}_3(\Pi, \underline{c}, L) \stackrel{\text{def}}{=} \frac{1}{\mathcal{L}_k(M)\sqrt{k-1}} \cdot \frac{\langle \mathcal{J}, [\underline{c}] \rangle \sqrt{\langle \Pi, \Pi \rangle}}{\langle \Pi, [\underline{c}] \rangle}$$
$$\leqslant \frac{1}{\mathcal{L}_k(M)\sqrt{k-1}} \cdot \frac{n}{\sqrt{k-1}}. \qquad (43)$$

To obtain the inequality in (42), we have used the fact that

$$\langle \Pi, \mathbf{1}-2\mathcal{J}\rangle \leqslant \langle \Pi, \mathbf{1}\rangle = n \quad \text{and} \quad \langle \Pi, \Pi \rangle \geqslant n/q.$$

To obtain the inequality in (43), we made use of the following two observations. First, we have $\langle \mathcal{J}, [\underline{c}] \rangle \leqslant \langle \mathbf{1}, [\underline{c}] \rangle = n$. Second, if Π and $\underline{c}$ are such that (39) holds, then *a fortiori* $\langle \Pi, [\underline{c}] \rangle / \sqrt{\langle \Pi, \Pi \rangle} \geqslant \sqrt{k-1}$. Since $L \leqslant \mathcal{L}_k(M)$ by (38), it follows from (42) and (43), respectively, that

$$\mathcal{F}_2(\Pi, L) \leqslant \sqrt{q}/2L\sqrt{R^*} \quad \text{and} \quad \mathcal{F}_3(\Pi, \underline{c}, L) \leqslant 1/LR^*.$$

[1] In practice, algebraic soft-decoding with list size limited to L almost always produces lists with much less than L codewords, most often a single codeword.

In conjunction with (41) and (40), this completes the proof of the theorem. $\qquad \square$

We observe that Theorem 17 is a very loose bound. The actual performance of algebraic soft-decoding with list size limited to L is usually orders of magnitude better than that predicted by (39). In the proof of Theorem 17, we have used the inequality $\langle \Pi, \Pi \rangle \geqslant n/q$, which is a weak lower bound since $\langle \Pi, \Pi \rangle \simeq n$ for signal-to-noise ratios (SNRs) of practical interest. Replacing n/q on the right-hand side of (42) by the actual value $\langle \Pi, \Pi \rangle$, we obtain a somewhat stronger bound, which guarantees that $\underline{c} \in \mathbb{C}_q(n,k)$ is on the list produced by the decoder, provided

$$\frac{\langle \Pi, [\underline{c}] \rangle}{\sqrt{\langle \Pi, \Pi \rangle}} \geqslant \frac{\sqrt{k-1}}{1 - \frac{1}{L}\left(\frac{1}{R^*} + \frac{1}{2\sqrt{R^*}} \cdot \frac{\sqrt{n}}{\sqrt{\langle \Pi, \Pi \rangle}}\right)}$$
$$\simeq \frac{\sqrt{k-1}}{1 - \frac{1}{L}\left(\frac{1}{R^*} + \frac{1}{2\sqrt{R^*}}\right)}. \qquad (44)$$

This works well for large L, although (44) is still a loose bound for moderate list sizes. Nevertheless, the significance of Theorem 17 is that it proves convergence to the asymptotic performance *at least* as fast as $O(1/L)$. Furthermore, the theorem shows that the size of the list required to approach the asymptotic performance within any given constant does not depend on the length of the code.

In addition to the analysis of Theorem 17, we have performed extensive simulations of algebraic soft-decoding with list size limited to L for various Reed–Solomon codes over GF(256). As the running channel model, we have assumed an AWGN channel with a 256-QAM signal constellation. The 256 constellation points were matched to the 256 elements of GF(256) in an arbitrary manner. The reliability matrix Π was computed by measuring the distance from the channel output to the four nearest constellation points. Thus only four entries in each column of Π were nonzero. Moreover, all the entries in Π were normalized and quantized to 8 bits of precision. The performance curves were obtained by running Algorithm A as discussed in the remark at the end of Section IV, then interpolating and factoring as discussed in Section III. We note that the same curves result by simply evaluating the sufficient condition of Theorem 3: the difference between the two error rates was in the second or third significant digit at all SNRs, in all cases we have simulated.

Simulation results for the $(255, 144, 112)$ Reed–Solomon code of rate ~ 0.56 are summarized in Fig. 3. One can see from Fig. 3 that at codeword error rates of 10^{-5} and lower, algebraic soft-decision decoding provides a coding gain of about 1.5 dB, whereas GMD decoding and Guruswami–Sudan decoding achieve coding gains of about 0.2 and 0.4 dB, respectively, compared to conventional hard-decision decoding. Although the 1.5-dB coding gain corresponds to asymptotic performance (cf. Theorem 12), it is evident from Fig. 3 that most of this gain can be obtained with very small list sizes. A list of size $L = 4$ already outperforms both GMD and Guruswami–Sudan decoding by a substantial margin, while a list of size $L = 32$ approaches the asymptotic performance to within 0.1 dB.

Fig. 3. Performance of algebraic soft-decision decoding for the $(255, 144, 112)$ Reed–Solomon code with 256-QAM modulation on an AWGN channel.


Fig. 4. Performance of algebraic soft-decoding for the $(204, 188, 17)$ shortened Reed–Solomon code with 256-QAM modulation on an AWGN channel.

Simulation results for the $(204, 188, 17)$ shortened Reed–Solomon code of rate ~ 0.92 are presented in Fig. 4. We observe that this code, in conjunction with a 256-QAM signal constellation, is implemented today in certain satellite communications systems. Here, algebraic soft-decision decoding provides an ultimate coding gain of about 0.75 dB. The fact that the asymptotic coding gain decreases with the rate of a code is to be expected since list decoding, in general, is less effective for high-rate codes. In fact, the asymptotic performance of Guruswami–Sudan list decoding coincides with that of the conventional Berlekamp–Wech decoding for the $(204, 188, 17)$

code: the Guruswami–Sudan decoder finds all codewords within Hamming distance of

$$\left\lfloor 204(1 - \sqrt{0.92}) \right\rfloor = \lfloor 8.16 \rfloor = (17-1)/2$$

from the (hard-decision) channel output (cf. Theorem 2). In contrast, soft-decision list decoding does provide a significant coding gain. As in the case of half-rate codes, most of this gain can be achieved with small list sizes. Moreover, one can see from Fig. 4 that the coding gain grows with SNR. Extrapolating the simulation results to error rates of about 10^{-10} (that

are of interest for many applications), one should expect coding gains in excess of about 1.0 dB for high-rate as well as low-rate Reed–Solomon codes.

VII. CHANNELS WITH MEMORY AND CONCATENATED CODING

Throughout this paper, we have assumed that the channel input comes from a product distribution, and that the channel is memoryless. In other words, if $\boldsymbol{\mathcal{X}} = (\mathcal{X}_1, \mathcal{X}_2, \ldots, \mathcal{X}_n)$ and $\boldsymbol{\mathcal{Y}} = (\mathcal{Y}_1, \mathcal{Y}_2, \ldots, \mathcal{Y}_n)$ denote the random vectors at the channel input and output, respectively, we have assumed that $\mathcal{X}_1, \mathcal{X}_2, \ldots, \mathcal{X}_n$ are independent and identically distributed (i.i.d.) and that $\mathcal{Y}_j$ depends only on $\mathcal{X}_j$, which makes the random variables $\mathcal{Y}_1, \mathcal{Y}_2, \ldots, \mathcal{Y}_n$ i.i.d. as well. These assumptions are reflected in our definition of the reliability matrix in (7) and (8) and of the expected score in (14) and (15). These are the basic concepts that underlie our soft-decoding algorithm.

While the preceding assumptions are justified in a variety of contexts, there are important applications of Reed–Solomon and algebraic-geometry codes where these assumptions are not valid. In practice, the most consequential of such applications is the use of Reed–Solomon (and algebraic-geometry) codes as outer codes in concatenated coding schemes [11], [15].

In the most general setup, we have to assume that the channel input $\boldsymbol{\mathcal{X}}$ and output $\boldsymbol{\mathcal{Y}}$ are governed by a $2n$-dimensional joint probability distribution $P_{\boldsymbol{\mathcal{X}}, \boldsymbol{\mathcal{Y}}}(\underline{x}, \underline{y})$. This setup encompasses arbitrary channels with memory and allows for arbitrary distributions on the channel input. For $\underline{y} \in \mathfrak{Y}^n$, let

$$P_{\boldsymbol{\mathcal{X}}|\boldsymbol{\mathcal{Y}}}(\underline{x}\,|\,\underline{y}) = \frac{P_{\boldsymbol{\mathcal{X}}, \boldsymbol{\mathcal{Y}}}(\underline{x}, \underline{y})}{P_{\boldsymbol{\mathcal{Y}}}(\underline{y})}$$

be the conditional joint distribution on the channel input $(\mathcal{X}_1, \mathcal{X}_2, \ldots, \mathcal{X}_n)$ given that $\boldsymbol{\mathcal{Y}} = \underline{y}$. Then, in the most general case, given the vector $\underline{y} \in \mathfrak{Y}^n$ observed at the channel output and a $q \times n$ multiplicity matrix M, we need to compute the expected score with respect to $P_{\boldsymbol{\mathcal{X}}|\boldsymbol{\mathcal{Y}}}(\cdot\,|\,\underline{y})$, namely

$$\mathsf{E}_{P_{\boldsymbol{\mathcal{X}}|\boldsymbol{\mathcal{Y}}}} \{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\} \stackrel{\text{def}}{=} \sum_{\underline{x} \in \mathfrak{X}^n} \mathcal{S}_M(\underline{x}) P_{\boldsymbol{\mathcal{X}}|\boldsymbol{\mathcal{Y}}}(\underline{x}\,|\,\underline{y})$$
$$= \sum_{\underline{x} \in \mathbb{F}_q^n} \langle M, [\underline{x}] \rangle\, P_{\boldsymbol{\mathcal{X}}|\boldsymbol{\mathcal{Y}}}(\underline{x}\,|\,\underline{y}). \quad (45)$$

We then need to find a multiplicity matrix $\mathcal{M} \in \mathfrak{M}(\mathcal{C})$ that maximizes this expected score. In what follows, we show that the decoding procedure of Section IV can be easily modified (in an optimal way) to accommodate channels with memory in the general setup of (45).

To this end, we introduce a **generalized reliability matrix** $\Pi^\star$, which reduces to the reliability matrix Π for memoryless channels and product input distributions. In the general case, $\Pi^\star = [\pi_{i,j}^\star]$ is a $q \times n$ matrix whose entries are defined as follows:

$$\pi_{i,j}^\star \stackrel{\text{def}}{=} \Pr\left(\mathcal{X}_j = \alpha_i \mid \boldsymbol{\mathcal{Y}} = \underline{y}\right),$$
$$\text{for } i = 1, 2, \ldots, q \text{ and } j = 1, 2, \ldots, n \quad (46)$$

where $\underline{y}$ is the observed channel output, $\boldsymbol{\mathcal{X}} = (\mathcal{X}_1, \mathcal{X}_2, \ldots, \mathcal{X}_n)$ is the channel input, and α_i is the ith element of the input alphabet $\mathfrak{X} = \mathbb{F}_q$. The following lemma is the counterpart of

Lemma 6 for channels with memory and/or non-product input distributions.

Lemma 18: The expected score with respect to the probability distribution $P_{\boldsymbol{\mathcal{X}}|\boldsymbol{\mathcal{Y}}}(\cdot\,|\,\underline{y})$ is equal to the inner product of $\Pi^\star$ and the multiplicity matrix, namely

$$\mathsf{E}_{P_{\boldsymbol{\mathcal{X}}|\boldsymbol{\mathcal{Y}}}} \{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\} = \langle M, \Pi^\star \rangle.$$

Proof: Let $\mathfrak{P}$ denote the expected value of $[\boldsymbol{\mathcal{X}}]$ with respect to the distribution $P_{\boldsymbol{\mathcal{X}}|\boldsymbol{\mathcal{Y}}}(\cdot\,|\,\underline{y})$, namely

$$\mathfrak{P} = \sum_{\underline{x} \in \mathbb{F}_q^n} [\underline{x}] P_{\boldsymbol{\mathcal{X}}|\boldsymbol{\mathcal{Y}}}(\underline{x}\,|\,\underline{y}).$$

Since $[\underline{x}]_{i,j} = 1$ if $x_j = \alpha_i$, and $[\underline{x}]_{i,j} = 0$ otherwise, the entry found in row i and column j of the matrix $\mathfrak{P}$ is given by

$$p_{i,j} = \sum_{\substack{\underline{x} \in \mathbb{F}_q^n \\ x_j = \alpha_i}} P_{\boldsymbol{\mathcal{X}}|\boldsymbol{\mathcal{Y}}}(\underline{x}\,|\,\underline{y}) = \Pr\left(\mathcal{X}_j = \alpha_i \mid \boldsymbol{\mathcal{Y}} = \underline{y}\right) = \pi_{i,j}^\star.$$

Thus $\mathfrak{P}$ is precisely the generalized reliability matrix $\Pi^\star$. The theorem now follows by the linearity of expectation

$$\mathsf{E}_{P_{\boldsymbol{\mathcal{X}}|\boldsymbol{\mathcal{Y}}}} \{\langle M, [\boldsymbol{\mathcal{X}}] \rangle\} = \langle M, \mathsf{E}_{P_{\boldsymbol{\mathcal{X}}|\boldsymbol{\mathcal{Y}}}} \{[\boldsymbol{\mathcal{X}}]\} \rangle = \langle M, \Pi^\star \rangle. \quad \square$$

The result of Lemma 18 is of exactly the same form as that of Lemma 6. This makes it possible to apply Algorithm A, without change, to the generalized reliability matrix $\Pi^\star$ to compute a multiplicity matrix that maximizes the expected score in (45).

Corollary 19: Let $\mathcal{M}(\Pi^\star, s)$ denote the multiplicity matrix produced by Algorithm A for a given generalized reliability matrix $\Pi^\star$ and a given number of interpolation points s. Let $\mathcal{C}$ be the cost of $\mathcal{M}(\Pi^\star, s)$. Then

$$\mathcal{M}(\Pi^\star, s) = \arg\max_{M \in \mathfrak{M}(\mathcal{C})} \mathsf{E}_{P_{\boldsymbol{\mathcal{X}}|\boldsymbol{\mathcal{Y}}}} \{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\}.$$

Corollary 19 follows immediately from Theorem 7 and Lemma 18, and provides the basis for soft-decision decoding on channels with memory: all the results of Sections III–VI apply, with the reliability matrix Π replaced by the generalized reliability matrix $\Pi^\star$.

The remaining problem is how to compute $\Pi^\star$ given the channel observations. Fortunately, a computation of this kind is very common in communication systems.

Given a joint distribution $P_{\boldsymbol{\mathcal{X}}, \boldsymbol{\mathcal{Y}}}(\underline{x}, \underline{y})$ on the channel input $\boldsymbol{\mathcal{X}}$ and output $\boldsymbol{\mathcal{Y}}$ together with a specific observation $\underline{y} = (y_1, y_2, \ldots, y_n)$, we have to compute the conditional probabilities $\pi_{i,j}^\star = \Pr(\mathcal{X}_j = \alpha_i | \boldsymbol{\mathcal{Y}} = \underline{y})$ for all $\alpha_1, \alpha_2, \ldots, \alpha_q$ and all positions $j = 1, 2, \ldots, n$. This is precisely the task known as maximum *a posteriori* (MAP) symbol-by-symbol decoding. General algorithms for MAP symbol-by-symbol decoding, such as the sum–product algorithm or the forward–backward algorithm, are well known [18].

In particular, if the channel is a finite-state machine with a moderate number of states, then the generalized reliability matrix $\Pi^\star$ can be computed with the Bahl–Cocke–Jelinek–Raviv (BCJR) forward–backward algorithm [2]. Important special

(a)

(b)

Fig. 5.   Algebraic soft-decision coding gains on a fast Rayleigh-fading channel. Codewords of the Reed–Solomon code are binary phase-shift keying (BPSK) modulated to produce a sequence of $255 \cdot 8$ bits $(b_1, b_2, \ldots, b_{2040})$. This sequence is multiplied componentwise by the vector $(\alpha_1, \alpha_2, \ldots, \alpha_{2040})$, where α_i are i.i.d. Rayleigh random variables with unit mean. The channel output is given by $\alpha_1 b_1 + z_1, \alpha_2 b_2 + z_2, \ldots, \alpha_{2040} b_{2040} + z_{2040}$, where z_i are zero-mean i.i.d. Gaussian random variables. In computing the reliability matrix Π, the channel states α_i are assumed to be unknown to the receiver. Large cost (high interpolation multiplicity) asymptotic performance is shown for both the Guruswami–Sudan decoder and the algebraic soft-decision decoder. (a) Performance of the $(255, 144, 112)$ Reed–Solomon code. (b) Performance of the $(255, 191, 65)$ Reed–Solomon code.

cases include intersymbol interference (ISI) channels and outer channels in a concatenated coding scheme, whose memory derives from an inner convolutional (or block) code. If the trellis complexity of the inner code is moderate (as is the case in practice), then the BCJR [2] algorithm is usually quite efficient . Thus one of the key conclusions of this section is as follows: in the context of concatenated coding, the BCJR algorithm turns out to be an efficient means for converting a channel with memory into a "memoryless" channel for the purposes of algebraic soft-decision decoding.

Note Added in Proof: We point out that the coding gains due to algebraic soft-decision decoding of Reed–Solomon codes on certain important channels (with or without memory) turn out to be much higher than the corresponding coding gains on a memoryless AWGN channel. For example, simulation results for a fast Rayleigh-fading channel are presented in Fig. 5. We see from Fig. 5(a) that algebraic soft-decoding of the $(255, 144, 112)$ Reed–Solomon code provides a coding gain of about 3.0 dB over hard-decision decoding, whereas the corresponding gain on the AWGN channel is 1.5 dB (cf.

Fig. 3). In Fig. 5, we assume that channel state information is unknown to the receiver. Given the channel states, one can do even better. For example, for the $(255, 191, 65)$ Reed–Solomon code, we observe a soft-decision coding gain of about 2.5 dB in Fig. 5(b), whereas Gross *et al.* [12], [13] assume perfect state information and obtain a soft-decision coding gain of 3.5 dB for the same code.

VIII. CONCLUSION

We have shown that interpolation-based decoding can be used to devise an efficient soft-decision decoding algorithm for Reed–Solomon codes. The soft-decoding algorithm outperforms both GMD decoding and Guruswami–Sudan list decoding by a substantial margin.

The focus of this paper has been the performance achievable in a probabilistic setting, where the channel output is characterized in terms of *a posteriori* probabilities rather than error patterns. This is quite different from several recent papers [15], [19] which focus on a combinatorial setting, and provide guarantees on the number (and type) of errors that can be corrected on certain hard-decision channels. In particular, for long codes, the criterion derived here for the computation of a multiplicity matrix allows for reliable transmission at the highest possible rate, although this is not necessarily the criterion that maximizes the number of errors that one can guarantee to correct.

The asymptotic performance of the proposed soft-decoding algorithm for a large number of interpolation points or, equivalently, for large lists has been characterized in terms of simple geometric conditions. Moreover, it has been shown that that the asymptotic performance can be approached arbitrarily closely with list sizes that are bounded by a constant, even as the length of a code grows beyond all bounds.

APPENDIX
ON THE UNDERLYING PROBABILISTIC MODEL

In Section IV, in order to convert posterior probabilities (the reliability matrix Π) into interpolation points (the multiplicity matrix M), we regard the transmitted codeword as a random vector $\boldsymbol{\mathcal{X}} = (\mathcal{X}_1, \mathcal{X}_2, \ldots, \mathcal{X}_n) \in \mathfrak{X}^n$ and use the following probability distribution:

$$P(x_1, x_2, \ldots, x_n) = \prod_{j=1}^{n} \Pr\left(\mathcal{X}_j = x_j \mid \mathcal{Y}_j = y_j\right)$$
$$= \prod_{j=1}^{n} \Pi(x_j, j)$$

where $\underline{y} = (y_1, y_2, \ldots, y_n) \in \mathfrak{Y}^n$ is the vector observed at the channel output (cf. (15) of Section IV). Recall that this distribution corresponds to the following scenario: a vector $\boldsymbol{\mathcal{X}}$ is drawn uniformly at random from the space $\mathbb{F}_q^n$ and transmitted over a memoryless channel characterized by (6); thereupon the vector $\underline{y} \in \mathfrak{Y}^n$ is observed at the channel output. Up to certain natural assumptions, this is indeed what happens, *except* that the transmitted codeword $\boldsymbol{\mathcal{X}}$ is drawn uniformly at random from the code $\mathbb{C}_q(n, k)$ rather than the entire space $\mathbb{F}_q^n$. Thus the *a priori* distribution of $\boldsymbol{\mathcal{X}}$ is $\Pr(\boldsymbol{\mathcal{X}} = \underline{x}) = \mathcal{I}_{\mathbb{C}}(\underline{x})/q^k$, where

$\mathcal{I}_{\mathbb{C}}(\underline{x}) : \mathbb{F}_q^n \to \{0, 1\}$ is the indicator function for $\mathbb{C}_q(n, k)$ defined by

$$\mathcal{I}_{\mathbb{C}}(\underline{x}) \stackrel{\text{def}}{=} \begin{cases} 1, & \text{if } \underline{x} \in \mathbb{C}_q(n, k) \\ 0, & \text{otherwise.} \end{cases}$$

Given the channel observations $\underline{y} = (y_1, y_2, \ldots, y_n) \in \mathfrak{Y}^n$ one can easily compute the true posterior probability distribution of $\boldsymbol{\mathcal{X}}$ as follows:

$$P^*(x_1, x_2, \ldots, x_n) \stackrel{\text{def}}{=} \Pr\left(\boldsymbol{\mathcal{X}} = \underline{x} \mid \boldsymbol{\mathcal{Y}} = \underline{y}\right)$$
$$= \gamma \mathcal{I}_{\mathbb{C}}(\underline{x}) \prod_{j=1}^{n} \Pi(x_j, j). \quad (47)$$

The normalization constant γ in (47) is given by

$$\gamma \stackrel{\text{def}}{=} \frac{1}{\sum_{\underline{x} \in \mathbb{C}} \prod_{j=1}^{n} \Pi(x_j, j)} = q^{n-k} \frac{\prod_{j=1}^{n} f_{\mathcal{Y}_j}(y_j)}{f_{\boldsymbol{\mathcal{Y}}}(\underline{y})} \quad (48)$$

where $f_{\boldsymbol{\mathcal{Y}}}(\cdot)$ is the probability density function of the channel output $\boldsymbol{\mathcal{Y}} = (\mathcal{Y}_1, \mathcal{Y}_2, \ldots, \mathcal{Y}_n)$ (we assume without loss of generality that $\boldsymbol{\mathcal{Y}}$ is continuous), and $f_{\mathcal{Y}_j}(\cdot)$ are the marginal probability densities derived from $f_{\boldsymbol{\mathcal{Y}}}(\cdot)$. The expression in (47) follows by repeated application of the Bayes rule, first to $\Pr(\boldsymbol{\mathcal{X}} = \underline{x} \mid \boldsymbol{\mathcal{Y}} = \underline{y})$ and then to $\Pr(\mathcal{Y}_j = y_j \mid \mathcal{X}_j = x_j)$. Hence, the precise optimization problem we would like to solve is

$$M_{\text{opt}}(\Pi, \mathcal{C}) \stackrel{\text{def}}{=} \text{argmax}_{M \in \mathfrak{M}(C)} \mathsf{E}_{P^*} \{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\} \quad (49)$$

where, in contrast to (16), the expectation $\mathsf{E}_{P^*}\{\cdot\}$ is taken with respect to the true posterior distribution (47). While (49) gives a natural optimality criterion for the computation of the multiplicity matrix, we shall see that the computation itself is likely to be intractable.

There are two sources of difficulty in performing the maximization in (49). One of these has to do with the fact that computing $P^*(\underline{x})$ is difficult, even for a single input vector $\underline{x} = (x_1, x_2, \ldots, x_n) \in \mathbb{F}_q^n$. While $\mathcal{I}_{\mathbb{C}}(\underline{x})$ and $\prod_{j=1}^{n} \Pi(x_j, j)$ are easy to evaluate, it can be shown that computing γ in (48) for an arbitrary reliability matrix Π and an arbitrary linear code $\mathbb{C}$ is NP-hard. This difficulty, however, can be avoided as follows. Let

$$\Psi(\underline{x}) \stackrel{\text{def}}{=} \frac{P^*(\underline{x})}{\gamma} = \mathcal{I}_{\mathbb{C}}(\underline{x}) \prod_{j=1}^{n} \Pi(x_j, j) \quad (50)$$

be a density function. Given a multiplicity matrix M, let us formally define the expected score with respect to $\Psi(\cdot)$ as follows:

$$\mathsf{E}_{\Psi}\{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\} \stackrel{\text{def}}{=} \sum_{\underline{x} \in \mathfrak{X}^n} \mathcal{S}_M(\underline{x})\Psi(\underline{x})$$
$$= \sum_{\underline{x} \in \mathbb{F}_q^n} \sum_{j=1}^{n} M(x_j, j)\Psi(\underline{x}). \quad (51)$$

Then, it is easy to see from (50) and (51) that $\mathsf{E}_{P^*}\{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\}$ and $\mathsf{E}_{\Psi}\{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\}$ differ by a factor of γ that does not depend on the

multiplicity matrix M. Thus the knowledge of γ is not essential for the computation of argmax in (49), and we have

$$M_{\text{opt}}(\Pi, \mathcal{C}) \overset{\text{def}}{=} \text{argmax}_{M \in \mathfrak{M}(\mathcal{C})} \mathsf{E}_{P^*}\{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\}$$
$$= \text{argmax}_{M \in \mathfrak{M}(\mathcal{C})} \mathsf{E}_\Psi\{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\}. \quad (52)$$

Unfortunately, the second difficulty in the optimization of (49) and (52) is inherent in the presence of the indicator function $\mathcal{I}_{\mathbb{C}}(\cdot)$ in both $P^*(\cdot)$ and $\Psi(\cdot)$. Specifically, we now show that given a polynomial-time algorithm for the computation of $M_{\text{opt}}(\Pi, \mathcal{C})$ in (52), one could devise a polynomial-time algorithm for maximum-likelihood hard-decision decoding of $\mathbb{C}_q(n, k)$. If $\mathbb{C}_q(n, k)$ is a general linear code, the latter task is known [4] to be NP-hard.

More precisely, let q be a fixed prime power and let $d(\cdot, \cdot)$ denote the Hamming distance. Then the following decision problem:

Problem: MAXIMUM-LIKELIHOOD DECODING
Instance: Positive integers n, k, t, an $(n-k) \times n$ matrix H over $\mathbb{F}_q$, and a vector $\underline{y} \in \mathbb{F}_q^n$.
Question: Is there a vector $\underline{c} \in \mathbb{F}_q^n$ such that $d(\underline{c}, \underline{y}) \leqslant t$ and $H\underline{c}^T = \mathbf{0}$?

was shown to be NP-complete by Berlekamp, McEliece, and van Tilborg [4]. Let $\mathbb{Q}$ denote the field of rational numbers. In this appendix, we exhibit a polynomial transformation from MAXIMUM-LIKELIHOOD DECODING to the following decision problem:

Problem: OPTIMAL MULTIPLICITY MATRIX
Instance: Positive integers n, k, and $\mathcal{C}$, an $(n-k) \times n$ matrix H over $\mathbb{F}_q$ which defines a code $\mathbb{C}_q(n, k)$, a $q \times n$ reliability matrix Π over $\mathbb{Q}$, and a rational number β.
Question: Is there a matrix $M \in \mathfrak{M}(\mathcal{C})$ such that $\mathsf{E}_\Psi\{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\} \geqslant \beta$?

It is easy to see that OPTIMAL MULTIPLICITY MATRIX is just a reformulation of the optimization problem (52) as a decision problem. Notice that this decision problem is not necessarily in NP, since given a putative solution $M \in \mathfrak{M}(\mathcal{C})$, there is no obvious way to verify that $\mathsf{E}_\Psi\{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\} \geqslant \beta$ in polynomial time.

Theorem 20: OPTIMAL MULTIPLICITY MATRIX is NP-hard.

Proof: We reduce from MAXIMUM-LIKELIHOOD DECODING. Given an instance $\{H, \underline{y}, t\}$ of MAXIMUM-LIKELIHOOD DECODING, we generate an instance of OPTIMAL MULTIPLICITY MATRIX as follows. Fix a rational number ϵ such that $1 > \epsilon > q^k/(q^k + q - 1)$ and let $\delta = (1 - \epsilon)/(q - 1)$. This choice of ϵ and δ ensures that $\epsilon + (q - 1)\delta = 1$ and $\epsilon/\delta > q^k$. In terms of ϵ, δ, and $\underline{y}$, we set

$$\Pi = \epsilon[\underline{y}] + \delta(1 - [\underline{y}]).$$

The fact that $\epsilon + (q - 1)\delta = 1$ implies that Π is a valid reliability matrix. We take $\beta = n\epsilon^{n-t}\delta^t$. Finally, we use the same parity-check matrix H, and set $\mathcal{C} = n$. This completes the mapping of $\{H, \underline{y}, t\}$ onto an instance $\{H, \Pi, \mathcal{C}, \beta\}$ of OPTIMAL MULTIPLICITY MATRIX.

Suppose that $\{H, \underline{y}, t\}$ is a "YES" instance of MAXIMUM-LIKELIHOOD DECODING. Then there exists a codeword

$\underline{c} \in \mathbb{C}_q(n, k)$ such that $d(\underline{c}, \underline{y}) \leqslant t$. Let $M = [\underline{c}]$. It is easy to see that $\mathcal{C}(M) = n$, and so $M \in \mathfrak{M}(\mathcal{C})$ for $\mathcal{C} = n$. Furthermore

$$\mathsf{E}_\Psi\{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\} = \sum_{\underline{x} \in \mathfrak{X}^n} \langle M, [\underline{x}] \rangle \Psi(\underline{x})$$
$$= \sum_{\underline{x} \in \mathbb{C}_q(n,k)} \langle [\underline{c}], [\underline{x}] \rangle \prod_{j=1}^n \Pi(x_j, j)$$
$$\geqslant n \prod_{j=1}^n \Pi(c_j, j)$$

where the inequality follows by retaining a single term in the summation over $\underline{x} \in \mathbb{C}_q(n, k)$ that corresponds to $\underline{x} = \underline{c}$. With the reliability matrix given by $\Pi = \epsilon[\underline{y}] + \delta(1 - [\underline{y}])$, we further conclude that

$$\mathsf{E}_\Psi\{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\} \geqslant n \prod_{j=1}^n \Pi(c_j, j)$$
$$= n\,\epsilon^{n-d(\underline{c},\underline{y})}\delta^{d(\underline{c},\underline{y})} \geqslant n\,\epsilon^{n-t}\delta^t = \beta$$

where the last inequality follows from $d(\underline{c}, \underline{y}) \leqslant t$ and $\delta \leqslant \epsilon$. Therefore, if $\{H, \underline{y}, t\}$ is a "YES" instance of MAXIMUM-LIKELIHOOD DECODING then $\{H, \Pi, \mathcal{C}, \beta\}$ is also a "YES" instance of OPTIMAL MULTIPLICITY MATRIX.

Now let $\{H, \underline{y}, t\}$ be a "NO" instance of MAXIMUM-LIKELIHOOD DECODING. Then, $d(\underline{x}, \underline{y}) \geqslant t + 1$ for all $\underline{x} \in \mathbb{C}_q(n, k)$. Observe that for any matrix $M \in \mathfrak{M}(n)$ and any vector $\underline{x} \in \mathbb{F}_q^n$, we have $\langle M, [\underline{x}] \rangle \leqslant \langle M, \mathbf{1} \rangle \leqslant \mathcal{C}(M) = n$. It follows that

$$\mathsf{E}_\Psi\{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\} = \sum_{\underline{x} \in \mathbb{C}_q(n,k)} \langle M, [\underline{x}] \rangle \prod_{j=1}^n \Pi(x_j, j)$$
$$\leqslant \sum_{\underline{x} \in \mathbb{C}_q(n,k)} n\,\epsilon^{n-t-1}\delta^{t+1}$$
$$= q^k\left(\frac{\delta}{\epsilon}\right)\beta < \beta$$

for any $M \in \mathfrak{M}(n)$. Hence, if $\{H, \underline{y}, t\}$ is a "NO" instance of MAXIMUM-LIKELIHOOD DECODING then $\{H, \Pi, \mathcal{C}, \beta\}$ is a "NO" instance of OPTIMAL MULTIPLICITY MATRIX. $\qquad\square$

It follows from Theorem 20 that solving the optimization problem (52) for an arbitrary linear code $\mathbb{C}_q(n, k)$ and an arbitrary cost $\mathcal{C}$ is NP-hard. It is possible to argue that the original optimization problem (16) might be also NP-hard for arbitrary costs; nevertheless, Algorithm A solves this problem for certain specific costs. However, in contrast to (16), the optimization in (52) remains NP-hard even if we restrict the cost to $\mathcal{C} = n$. Furthermore, as can be seen from the proof of Theorem 20, maximizing $\mathsf{E}_{P^*}\{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\}$ over all multiplicity matrices M such that $\langle M, \mathbf{1} \rangle = n$ (this is equivalent to selecting n interpolation points regardless of the cost) is still NP-hard. The analogous problem for $\mathsf{E}_P\{\mathcal{S}_M(\boldsymbol{\mathcal{X}})\}$, where $P(\cdot)$ is the distribution in (15) is trivial: it is solved by allocating all the n points at the position of the largest entry in Π.

Finally, one might argue that while the OPTIMAL MULTIPLICITY MATRIX problem has to do with arbitrary linear codes over $\mathbb{F}_q$, the codes involved in the optimization task (52) are Reed–Solomon codes and thus have a lot of structure. In this context, Theorem 20 shows that the computation of

$M_{\mathrm{opt}}(\Pi, \mathcal{C})$ in (52) subsumes maximum-likelihood hard-decision decoding of Reed–Solomon codes. No polynomial-time algorithm for maximum-likelihood hard-decision decoding of Reed–Solomon codes is presently known [28], and the problem is generally considered to be hard.

REFERENCES

[1] D. Augot and L. Pecquet, "A Hensel lifting to replace factorization in list-decoding of algebraic-geometric and Reed–Solomon codes," *IEEE Trans. Inform. Theory*, vol. 46, pp. 2605–2614, Nov. 2000.

[2] L. R. Bahl, J. Cocke, F. Jelinek, and J. Raviv, "Optimal decoding of linear codes for minimizing symbol error rate," *IEEE Trans. Inform. Theory*, vol. IT-20, pp. 284–287, Mar. 1974.

[3] E. R. Berlekamp, "Bounded distance+1 soft-decision Reed–Solomon decoding," *IEEE Trans. Inform. Theory*, vol. 42, pp. 704–721, May 1996.

[4] E. R. Berlekamp, R. J. McEliece, and H. C. A. van Tilborg, "On the inherent intractability of certain coding problems," *IEEE Trans. Inform. Theory*, vol. IT-24, pp. 384–386, May 1978.

[5] E. R. Berlekamp, R. E. Peile, and S. P. Pope, "The application of error control to communications," *IEEE Commun. Mag.*, vol. 25, pp. 44–57, Jan. 1987.

[6] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley Interscience, 1991.

[7] A. B. Cooper 3rd, "Soft-decision decoding of Reed-Solomon codes," in *Reed-Solomon Codes and their Applications*, S. B. Wicker and V. K. Bhargava, Eds. New York: IEEE Press, 1994, pp. 108–124.

[8] G.-L. Feng, "A fast special factorization algorithm in the Sudan decoding procedure," in *Proc. 31st Allerton Conf. Communications, Control, and Computing*, Oct. 2000, pp. 593–602.

[9] G.-L. Feng and X. Giraud, "Fast algorithms in Sudan decoding procedure for Reed–Solomon codes," *IEEE Trans. Inform. Theory*, submitted for publication.

[10] G. D. Forney Jr., "Generalized minimum distance decoding," *IEEE Trans. Inform. Theory*, vol. IT-12, pp. 125–131, Apr. 1966.

[11] ——, *Concatenated Codes*. Cambridge, MA: MIT Press, 1966.

[12] W. J. Gross, F. R. Kschischang, R. Koetter, and P. G. Gulak, "A VLSI architecture for interpolation in soft-decision list decoding of Reed-Solomon codes," *Proc. IEEE Workshop on Signal Processing Systems*, pp. 39–44, Oct. 2002.

[13] ——, "Applications of algebraic soft-decision decoding of Reed-Solomon codes," *IEEE Trans. Commun.*, submitted for publication.

[14] V. Guruswami and M. Sudan, "Improved decoding of Reed–Solomon and algebraic-geometric codes," *IEEE Trans. Inform. Theory*, vol. 45, pp. 1757–1767, Sept. 1999.

[15] ——, "List decoding algorithms for certain concatenated codes," *IEEE Trans. Inform. Theory*, submitted for publication.

[16] R. Koetter, "On Algebraic Decoding of Algebraic-Geometric and Cyclic codes," Ph.D. dissertation, Univ. Linköping, Linköping, Sweden, 1996.

[17] ——, "Fast generalized minimum distance decoding of algebraic geometric and Reed–Solomon codes," *IEEE Trans. Inform. Theory*, vol. 42, pp. 721–738, May 1996.

[18] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Trans. Inform. Theory*, vol. 47, pp. 498–519, Feb. 2001.

[19] R. R. Nielsen, "Decoding concatenated codes with Sudan's algorithm," *IEEE Trans. Inform. Theory*, submitted for publication.

[20] R. R. Nielsen and T. Høholdt, "Decoding Reed-Solomon codes beyond half the minimum distance," preprint, 1998.

[21] V. Olshevsky and A. Shokrollahi, "A displacement structure approach to efficient decoding of Reed-Solomon and algebraic-geometric codes," in *Proc. 31th ACM Symp. Theory of Computing (STOC)*, Atlanta, GA, May 1999, pp. 235–244.

[22] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*. New York: McGraw-Hill, 1984.

[23] V. Ponnampalam and B. S. Vucetic, "Soft decision decoding of Reed-Solomon codes," in *Proc. 13th Symp. Applied Algebra, Algebraic Algorithms, and Error-Correcting Codes*, Honolulu, HI, Nov. 1999.

[24] S. Ray-Chaudhuri and A. H. Chan, "Bit-level parallel decoding of Reed-Solomon codes," in *Proc. 31th Allerton Conf. Communications, Control, and Computing*, Sept. 1993.

[25] R. M. Roth and G. Ruckenstein, "Efficient decoding of Reed–Solomon codes beyond half the minimum distance," *IEEE Trans. Inform. Theory*, vol. 46, pp. 246–258, Jan. 2000.

[26] U. Sorger, "A new Reed-Solomon decoding algorithm based on Newton's interpolation," *IEEE Trans. Inform. Theory*, vol. 39, pp. 358–365, Mar. 1993.

[27] M. Sudan, "Decoding of Reed-Solomon codes beyond the error correction bound," *J. Complexity*, vol. 12, pp. 180–193, 1997.

[28] A. Vardy, "Algorithmic complexity in coding theory and the minimum distance problem," in *Proc. 29th Symp. Theory of Computing*, El Paso, TX, 1997, pp. 92–109.

[29] A. Vardy and Y. Be'ery, "Bit-level soft-decision decoding of Reed-Solomon codes," *IEEE Trans. Commun.*, vol. 39, pp. 440–445, Mar. 1991.

[30] L. R. Welch and E. R. Berlekamp, "Error Correction for Algebraic Block Codes," U.S. Patent 633 470, Dec. 30, 1986.

[31] W. W. Wu, D. Haccoun, R. E. Peile, and Y. Hirata, "Coding for satellite communication," *IEEE J. Select. Areas Commun.*, vol. 5, pp. 724–785, May 1987.

[32] X.-W. Wu and P. H. Siegel, "Efficient root-finding algorithm with application to list decoding of algebraic-geometric codes," *IEEE Trans. Inform. Theory*, vol. 47, pp. 2579–2587, Sept. 2001.