Final solutions

Problem F.1 (70 points)

In this problem we will consider coded modulation schemes based on a one-to-one mapping $t : \mathbb{F}_3 \to \mathcal{A}$ from the finite field $\mathbb{F}_3$ to a 3-simplex signal set $\mathcal{A}$ in $\mathbb{R}^2$ with energy $E(\mathcal{A})$ per symbol. The symbols from $\mathcal{A}$ will be transmitted by QAM modulation over a passband AWGN channel with single-sided power spectral density N_0. In everything that follows, we assume that the receiver performs optimal detection.

The amount of information that can be conveyed in one ternary symbol will be called one trit. *We will normalize everything "per information trit;" i.e.,*

- *we will use E_t/N_0 as our normalized signal-to-noise ratio, where E_t is the average energy per information trit;*
- *we will define the nominal spectral efficiency ρ_t as the number of information trits per two dimensions (t/2D) conveyed by a given transmission scheme; and*
- *we will define $P_t(E)$ as the probability of error per information trit.*

(a) What is the ultimate Shannon limit on E_t/N_0 in dB?

The amount of information conveyed by one equiprobable ternary symbol is $\log_2 3$ bits. Thus one trit is equal to $\log_2 3 = \log_{10} 3/\log_{10} 2 = 4.77/3.01 = 1.58$ bits. Alternatively, we can just take logarithms to the base 3 to measure information-theoretic quantities directly in trits; *i.e.,* the amount of information conveyed by one equiprobable ternary symbol is $\log_3 3 = 1$ trit.

The capacity of an AWGN channel is thus

$$C_t = \frac{\log_2(1+\mathrm{SNR})}{\log_2 3} = \log_3(1+\mathrm{SNR}) \quad \text{t/2D}.$$

The signal energy per two dimensions is $\rho_t E_t$, so $\mathrm{SNR} = \rho_t E_t/N_0$. Thus for reliable transmission

$$\rho_t \le \log_3(1+\rho_t E_t/N_0),$$

which is equivalent to

$$E_t/N_0 \ge \frac{3^{\rho_t}-1}{\rho_t}.$$

As $\rho_t \to 0$, we have $3^{\rho_t} = \exp(\rho_t \ln 3) \to 1+\rho_t \ln 3$, so this lower bound decreases monotonically to

$$\ln 3 = \frac{\log_{10} 3}{\log_{10} e} = \frac{4.77}{4.34} = 1.10 \text{ (0.41 dB)}.$$

Alternatively, since the ultimate Shannon limit on E_b/N_0 is $\ln 2$ (-1.59 dB), the ultimate Shannon limit on E_t/N_0 is $(\ln 2)(\log_2 3) = \ln 3 = 1.10$ (0.41 dB).

(b) What is the baseline performance ($P_t(E)$ vs. E_t/N_0) of the signal set $\mathcal{A}$?

A 3-simplex signal set $\mathcal{A}$ may be constructed by starting with a 3-orthogonal signal set $\mathcal{A}'$ and subtracting out the mean $\mathbf{m}(\mathcal{A}')$: $\mathcal{A} = \mathcal{A}' - \mathbf{m}(\mathcal{A}')$. Then $d^2_{\min}(\mathcal{A}) = d^2_{\min}(\mathcal{A}')$, and because $\mathbf{m}(\mathcal{A}) = \mathbf{0}$, we have $E(\mathcal{A}') = E(\mathcal{A}) + ||\mathbf{m}(\mathcal{A}')||^2$, or $E(\mathcal{A}) = E(\mathcal{A}') - ||\mathbf{m}(\mathcal{A}')||^2$.

Take $\mathcal{A}' = \{(\alpha, 0, 0), (0, \alpha, 0), (0, 0, \alpha)\}$; then $d^2_{\min}(\mathcal{A}') = 2\alpha^2$, $E(\mathcal{A}') = \alpha^2$, and $\mathbf{m}(\mathcal{A}') = (\alpha, \alpha, \alpha)/3$. Thus $\mathcal{A} = \mathcal{A}' - \mathbf{m}(\mathcal{A}')$ has $d^2_{\min}(\mathcal{A}) = d^2_{\min}(\mathcal{A}') = 2\alpha^2$, and $E(\mathcal{A}) = E(\mathcal{A}') - ||\mathbf{m}(\mathcal{A}')||^2 = 2\alpha^2/3$. We conclude that $d^2_{\min}(\mathcal{A}) = 3E(\mathcal{A})$.

The same conclusion could be reached by taking $\mathcal{A}'$ to be the vertices of an equilateral triangle in $\mathbb{R}^2$ centered on the origin, or from our general formulas for the inner products of an M-simplex signal set, namely $||\mathbf{a}_j||^2 = E(\mathcal{A})$; $\langle \mathbf{a}_j, \mathbf{a}_{j'} \rangle = -E(\mathcal{A})/(M-1)$ if $j \neq j'$.

Since the energy per symbol or per trit is $E_t = E(\mathcal{A})$, and each signal in $\mathcal{A}$ has $K_{\min}(\mathcal{A}) = 2$ nearest neighbors, the union bound estimate (UBE) of the probability of error per symbol (or per trit) is

$$P_t(E) \approx 2Q\sqrt{\left(\frac{d^2_{\min}(\mathcal{A})}{2N_0}\right)} = 2Q\sqrt{\left(\frac{3}{2}E_t/N_0\right)}.$$

(c) How far is this baseline performance from the ultimate Shannon limit at $P_t(E) \approx 10^{-5}$?

The baseline ternary curve of part (b) may by obtained by moving the baseline binary curve $P_b(E) = Q\sqrt{(2E_b/N_0)}$ of Figure 1 to the right by the "coding loss" of $\frac{3}{4}$ (-1.25 dB) and up by a factor of 2, which costs about 0.2 dB at $P_t(E) \approx 10^{-5}$. Thus we obtain $P_t(E) \approx 10^{-5}$ when $E_t/N_0 \approx 9.6 + 1.25 + 0.2 \approx 11$ dB. This is about 10.6 dB from the ultimate Shannon limit on E_t/N_0 of 0.4 dB.

Let $\mathcal{C}$ be the $(4, 2, 3)$ linear "tetracode" over $\mathbb{F}_3$, and let $t(\mathcal{C})$ be the Euclidean image of $\mathcal{C}$ under the map $t : \mathbb{F}_3 \to \mathcal{A}$.

(d) What are the state and branch complexities of a minimal trellis for $\mathcal{C}$?

The tetracode $\mathcal{C}$ meets the Singleton bound $d + k \leq n + 1$ with equality, and therefore is MDS. Its trellis-oriented generator matrix thus must have the following form:

$$\begin{bmatrix} xxx0 \\ 0xxx \end{bmatrix}.$$

From this matrix we see that the state complexity profile of a minimal trellis for $\mathcal{C}$ is $\{1, 3, 9, 3, 1\}$, and the branch complexity profile is $\{3, 9, 9, 3\}$.

(e) What is the performance ($P_t(E)$ vs. E_t/N_0) of the signal set $t(\mathcal{C})$?

We first note that the minimum Hamming distance of $\mathcal{C}$ is 3, and that all 8 nonzero codewords have weight 3, since $\mathcal{C}$ is MDS and thus $N_d = \binom{4}{3}(3-1) = 8$.

The minimum squared distance of $t(\mathcal{C})$ is therefore $3d^2_{\min}(\mathcal{A})$, since every sequence in $t(\mathcal{C})$ differs from every other by $d^2_{\min}(\mathcal{A})$ in 3 places. The number of nearest neighbors is $K_{\min}(t(\mathcal{C})) = 8$. (In fact, $t(\mathcal{C})$ is a 9-simplex.)

Finally, $E_t = 4E(\mathcal{A})/2 = 2E(\mathcal{A})$, and $P_t(E) = \frac{1}{2}\Pr(E)$. The union bound estimate (UBE) of the probability of error per information trit is thus

$$P_t(E) \approx \frac{1}{2} K_{\min}(t(\mathcal{C})) Q\sqrt{\left(\frac{3d^2_{\min}(\mathcal{A})}{2N_0}\right)} = 4Q\sqrt{\left(\frac{9}{4}E_t/N_0\right)}.$$

In other words, the nominal coding gain over the baseline curve $P_t(E) \approx 2Q\sqrt{(\frac{3}{2}E_t/N_0)}$ is $\gamma_c(\mathcal{C}) = kd/n = 3/2$ (1.76 dB). Because of the doubling of the error coefficient, the effective coding gain at $P_t(E) \approx 10^{-5}$ is about 0.2 dB less, or about 1.55 dB.

Now let $\mathcal{C}'$ be a linear rate-1/2 convolutional code over $\mathbb{F}_3$ with generator 2-tuple $\mathbf{g}(D) = (1 + D, 1 + 2D)$, and let $t(\mathcal{C}')$ be the Euclidean image of $\mathcal{C}'$ under the map t.

(f) What are the state and branch complexities of a minimal trellis for $\mathcal{C}'$?

The encoder for $\mathcal{C}'$ has one memory element storing one trit. It therefore has 3 states. There is a 3-way branch out of every state, so its branch complexity is 9. (The VA decoding complexity of $\mathcal{C}'$ is very nearly the same as that of $\mathcal{C}$.)

(g) What is the performance ($P_t(E)$ vs. E_t/N_0) of $t(\mathcal{C}')$?

We will first establish that the minimum Hamming distance of $\mathcal{C}'$ is 4, and that there are $K_t = 2$ error events of weight 4 per unit time (per information trit).

In the trellis diagram of $\mathcal{C}'$, the branch from the zero state to the zero state is labelled 00. The branches leaving the zero state to nonzero states are labelled 11 and 22, and the branches arriving at the zero state from nonzero states are labelled 12 and 21. The labels of the 9 branches run through the 9 ternary linear combinations of 11 and 12, which comprise all of $(\mathbb{F}_3)^2$ since 11 and 12 are linearly independent. The labels of the 4 branches from nonzero states to nonzero states therefore have Hamming weight 1. Thus every nonzero trellis path from the zero state to the zero state has Hamming weight 2 in its first branch, 2 in its last branch, and 1 in every branch in between. We conclude that the minimum Hamming weight is 4, and that only the 2 error events of length 2 (*i.e.,* $\mathbf{g}(D)$ and $2\mathbf{g}(D)$) have the minimum weight.

The minimum squared distance of $t(\mathcal{C}')$ is therefore $4d^2_{\min}(\mathcal{A})$, since every sequence in $t(\mathcal{C}')$ differs from every other by $d^2_{\min}(\mathcal{A})$ in at least 4 places. The number of nearest neighbors per information trit is $K_t = 2$. The energy per information trit is $E_t = 2E(\mathcal{A})$.

The union bound estimate (UBE) of the probability of error per information trit is thus

$$P_t(E) \approx K_t Q\sqrt{\left(\frac{4d^2_{\min}(\mathcal{A})}{2N_0}\right)} = 2Q\sqrt{(3E_t/N_0)}.$$

In other words, the nominal coding gain over the baseline curve $P_t(E) \approx 2Q\sqrt{(\frac{3}{2}E_t/N_0)}$ is $\gamma_c(\mathcal{C}) = kd/n = 2$ (3.01 dB). Because the error coefficient is the same, the effective coding gain is also 3 dB at all $P_t(E)$.

In summary, even though the block and convolutional codes have about the same VA decoding complexity and the block code is as good as possible (MDS), the effective coding gain of the convolutional code is about 1.5 dB greater.

Problem F.2 (50 points)

Consider the $(16, 7, 6)$ *binary linear block code* $\mathcal{C}$ *generated by the following generator matrix:*

$$\begin{bmatrix} 1111 & 1100 & 0000 & 0000 \\ 0101 & 1011 & 1000 & 0000 \\ 1100 & 1001 & 0110 & 0000 \\ 1001 & 1111 & 0101 & 0000 \\ 1010 & 0001 & 0100 & 1100 \\ 1100 & 0101 & 0000 & 1010 \\ 0011 & 0010 & 0100 & 1001 \end{bmatrix}.$$

(a) It is known that $k_{\max}(n, 6) = \{0, 0, 0, 0, 0, 1, 1, 1, 2, 2, 3, 4, 4, 5, 6, 7\}$ *for* $1 \leq n \leq 16$. *Show that there exist shortened codes of* $\mathcal{C}$ *that meet this bound for every* $n \leq 16$.

By inspection, for $1 \leq n \leq 16$, the first $k_{\max}(n, 6)$ generators of $\mathcal{C}$, shortened to the first n coordinates, generate an $(n, k_{\max}(n, 6))$ binary linear block code. Since these codes are each a shortened code of $\mathcal{C}$, which has minimum distance $d = 6$, each shortened code must have minimum distance at least 6.

(b) Give the state complexity profile and the branch complexity profile of a 16-section minimal trellis for $\mathcal{C}$.

We first reduce the generator matrix above to trellis-oriented form, obtaining:

$$\begin{bmatrix} 1111 & 1100 & 0000 & 0000 \\ 0101 & 1011 & 1000 & 0000 \\ 0011 & 0101 & 0110 & 0000 \\ 0000 & 1101 & 1011 & 0000 \\ 0000 & 0110 & 1100 & 1100 \\ 0000 & 0001 & 1101 & 1010 \\ 0000 & 0000 & 0011 & 1111 \end{bmatrix}.$$

Note that this generator matrix is symmetrical, and that the stopping times $(6, 9, 11, 12, 14, 15, 16)$ of the generators are the same as in the original generator matrix. The starting times are symmetrical, $(1, 2, 3, 5, 6, 8, 11)$, and thus the shortened codes generated by the last k generators also meet the bound of part (a).

From this trellis-oriented generator matrix, we find that the state complexity profile is

$$\{1, 2, 4, 8, 8, 16, 16, 16, 32, 16, 16, 16, 8, 8, 4, 2, 1\},$$

and that the branch complexity profile is

$$\{2, 4, 8, 8, 16, 32, 16, 32, 32, 16, 32, 16, 8, 8, 4, 2\}.$$

(c) From the information given, is it possible to say whether another coordinate ordering might give a less complex trellis for $\mathcal{C}$?

The past subcode at time n is generated by the generators that stop by time n, and the future subcode at time n is generated by the generators that start after time n. By part (a), the past subcode always has the largest dimension $k_{\max}(n, 6)$ that it could have, and by part (b) the same is true of the future subcodes. Therefore the Muder bound $\dim S_n \geq \dim \mathcal{C} - k_{\max}(n, 6) - k_{\max}(17 - n, 6)$ is met at all times, so no $(16, 7, 6)$ code could have a better state or branch complexity profile with any coordinate ordering.

(d) Find the sectionalization that gives the minimum number of sections without increasing the maximum branch complexity. Give the state complexity profile and the branch complexity profile of the resulting trellis.

Using the heuristic clustering rule of Chapter 10 (or the LV rule), the section at time 8 may be extended back to time 7 and forward to time 10 before meeting the first or last generator, giving a central section of length 4 with branch complexity 32. Similarly, the section at time 6 may be extended back to the beginning and the symmetrical section at time 11 may be extended to the end, giving first and last sections of length 6 with branch complexity 32. In short, this sectionalization gives a 3-section trellis with state complexity profile $\{1, 16, 16, 1\}$ and branch complexity profile $\{32, 32, 32\}$.

(e) Count the number of arithmetic operations required by decoding using a straightforward Viterbi algorithm of the trellises of parts (b) and (d). Which is less complex?

The 16 branch types, sizes and corresponding number of add and compare operations in a standard Viterbi algorithm decoding of the 16-section trellis are as follows

time	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
type	$<$	$<$	$<$	$=$	$<$	$\bowtie$	$=$	$<$	$>$	$=$	$\bowtie$	$>$	$=$	$>$	$>$	$>$
size	2	4	8	8	16	32	16	32	32	16	32	16	8	8	4	2
adds	0	4	8	8	16	32	16	32	32	16	32	16	8	8	4	2
comps	0	0	0	0	0	16	0	0	16	0	16	8	0	4	2	1

There are therefore a total of 234 additions (of two variables) and 63 2-way comparisons.

The three-section trellis has 32 branches of length 6 in the first section, which require $32 \times 5 = 160$ additions of single-symbol metrics in a straightforward implementation. At the end of the first section, 16 2-way comparisons are required. In the second section there are 32 branches of length 4, which require $32 \times 4 = 128$ additions, followed by 16 2-way comparisons. The final section has 32 branches of length 6, which require $32 \times 6 = 192$ additions, followed by a single 32-way comparison, equivalent to 31 2-way comparisons. Thus the total number of 2-way comparisons is 63, the same as for the unsectionalized trellis, but the total number of additions is 480, about twice as much. Evidently the unsectionalized trellis organizes the metric additions more efficiently.

However, the metric additions could be organized in a similar way in the sectionalized trellis, and the sectionalized trellis logic is less complex.

Problem F.3 (40 points)

For each of the propositions below, state whether the proposition is true or false, and give a proof of not more than a few sentences, or a counterexample. No credit will be given for a correct answer without an adequate explanation.

(a) There exist sequences of Reed-Muller codes which can approach the Shannon limit arbitrarily closely, but the trellis complexity of such a sequence of codes necessarily grows without limit.

TRUE. For $m \geq 0$, the Euclidean image of the first-order Reed-Muller code $\mathrm{RM}(1,m) = (2^m, m+1, 2^{m-2})$ is a 2^{m+1}-biorthogonal signal set. It is known that as $M \to \infty$ the probability of decoding error with M-biorthogonal signal sets goes to zero whenever $E_b/N_0 > \ln 2$ (-1.59 dB), the ultimate Shannon limit on E_b/N_0.

A minimal trellis for the $\mathrm{RM}(1,m) = (2^m, m+1, 2^{m-2})$ code has 2^{m-1} states at the central state space, since the $|u|u+v|$ construction gives the optimal coordinate ordering for RM codes, and under the $|u|u+v|$ construction the dimension of the central state space S is

$$\dim S = \dim \mathrm{RM}(1,m-1) - \dim \mathrm{RM}(0,m-1) = m-1.$$

Thus the trellis complexity of the $\mathrm{RM}(1,m)$ codes goes to ∞ as $m \to \infty$.

It is possible that there exist other sequences of RM codes whose performance approaches the relevant Shannon limit. However, the length of these codes must go to infinity in order to approach the Shannon limit arbitrarily closely, and the trellis complexity of an RM code $\mathrm{RM}(r,m)$ other than the universe, SPC and repetition codes (whose performance does not approach the Shannon limit) is lowerbounded by the trellis complexity of the $\mathrm{RM}(1,m)$ code.

(b) Let G be a finite abelian group of order $|G|$, and let X and N be independent random variables defined on G, where the probability distribution of N is uniform:

$$p_N(n) = 1/|G|, \forall n \in G.$$

Then $Y = X + N$ is uniformly distributed over G and independent of X, regardless of the distribution of X.

TRUE. The conditional probability distribution of Y given x is then uniform:

$$p_{Y|X}(y \mid x) = p_N(y-x) = 1/|G|, \forall x, y \in G;$$

i.e., $p_{Y|X}(y \mid x)$ is uniform independent of x. Thus Y is uniform:

$$p_Y(y) = \sum_{x \in G} p_{Y|X}(y \mid x) p_X(x) = 1/|G|, \forall y \in G.$$

Moreover, since $p_{Y|X}(y \mid x) = p_Y(y), \forall x \in G$, Y is independent of X.

This is the principle of the "one-time pad" in cryptography, which ensures that the encrypted text Y is independent of the plaintext X. This principle is also the basis of various "scrambling" and "dither" processes used in data communications to ensure that the transmitted signal is quasi-random, regardless of the actual data.

(c) There exists no MDS binary linear block code with block length greater than 3.

FALSE. An (n, k, d) linear code is MDS if it meets the Singleton bound, $d + k \leq n + 1$. The $(n, n, 1)$ binary universe code, the $(n, n-1, 2)$ binary single-parity-check code, and the $(n, 1, n)$ binary repetition code are thus all MDS for any $n \geq 1$.

(d) Given an (n, k, d) linear block code over a finite field $\mathbb{F}_q$ and optimal erasure correction:

(i) up to $d-1$ erasures can always be corrected;

(ii) up to $n-k$ erasures may be able to be corrected;

(iii) more than $n-k$ erasures can never be corrected.

TRUE. Optimal erasure correction is a matter of finding a codeword $\mathbf{c}$ that agrees with the received word $\mathbf{r}$ in the set $\mathcal{J}$ of unerased places; *i.e.,* such that the projections onto $\mathcal{J}$ agree: $\mathbf{c}_{|\mathcal{J}} = \mathbf{r}_{|\mathcal{J}}$. Since the transmitted codeword $\mathbf{c}$ always satisfies this condition, optimal erasure correction fails only if there is some other codeword $\mathbf{c}'$ such that $\mathbf{c}'_{|\mathcal{J}} = \mathbf{c}_{|\mathcal{J}}$, in which case there is no way to choose between $\mathbf{c}$ and $\mathbf{c}'$. By linearity, such a codeword $\mathbf{c}'$ exists if and only if there exists a nonzero codeword $\mathbf{c}'' = \mathbf{c} - \mathbf{c}'$ such that $\mathbf{c}''_{|\mathcal{J}} = \mathbf{0}_{|\mathcal{J}}$.

In case (i), if there are fewer than d erasures, then no such ambiguous case can arise, since the minimum Hamming weight of any nonzero codeword is d.

In case (ii), erasure correction is possible if and only if the projection of the code onto $\mathcal{J}$ is one-to-one; *i.e.,* if and only if $\mathcal{J}$ includes an information set for the code. (In the case of MDS codes every set of size k is an information set, but on the other hand $n-k = d-1$.)

In case (iii), unambiguous erasure correction is never possible, since the dimension of the code is k, so if there are fewer than k unerased places in $\mathcal{J}$ then the projection of the code onto $\mathcal{J}$ cannot possibly be one-to-one; *i.e.,* a set $\mathcal{J}$ of size less than k cannot possibly include an information set.