Design of a Coded Modulation for Deep Space Optical Communications

Bruce Moision, Jon Hamkins and Michael Cheng
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109 USA
Email: {bmoision,hamkins,mkcheng}@jpl.nasa.gov

***Abstract*— We show how to design a modulation and error correcting code (ECC) to achieve near-capacity performance on a deep space optical link. We begin by comparing the capacity of various modulations, given the power and bandwidth constraints of the channel and show that pulse position modulation (PPM) is near-optimal for a nominal mission scenario to Mars. We present a coded PPM solution comprised of the serial concatenation of a convolutional code, interleaver, accumulator, and PPM. This code is amenable to a high-speed, low-complexity hardware implementation and has performance with about 1 dB of capacity and several dB better than competing coded PPM designs.**

I. Introduction

All of NASA's previous deep space missions have relied upon Radio Frequency (RF) communications for their telemetry links. Such links, particularly at S-band and X-band, are well understood, with operational flight experience spanning several decades. The additive white Gaussian noise channel model for these RF links is both accurate and one of the simplest to analyze, providing a happy convergence of academic research and practical communications link design. Theoretical breakthroughs in ECC design regularly make their way onto deep space missions. The most common signal format used today by deep space missions is binary phase-shift keying (BPSK) with either a concatenated Reed-Solomon (RS) and convolutional code, or a turbo code. Low density parity-check codes are being standardized for space missions as well.

RF communications, however, will not always meet NASA's needs. Scientists and mission planners continue to demand ever increasing data rates from deep space, to support hyperspectral imagers, high-resolution stereo cameras, and other data-intensive instruments. Optical communications is attractive for these emerging mission scenarios because of advantages in transmitter mass, beam divergence, and bandwidth, compared to RF communications. This has led to the development of the Mars Laser Communications Demonstration (MLCD) project, scheduled to launch in 2009 and achieve data rates up to 50 Mbps from Mars using a 5 W laser. NASA strategic redirection has led to the cancellation of MLCD, but before this unfortunate news, techniques were developed to design a near-capacity signal format for deep-space optical links. This paper summarizes the design process.

The paper is organized as follows. In section II, we discuss the channel model, power and bandwidth constraints, and show how to select a modulation. In Section III, we compare RS, turbo, serially concatenated convolutional, and low-density parity-check code classes. In IV, we discuss the design of the serially concatenated code, including component code design, block length, interleaver, and bit-to-symbol mapping.

II. Modulation

In an optical deep space link, the received signal, a Poisson point process, is focused by a telescope onto a photodetector that produces a pulse in response to each incident photon. For our purposes, we assume this impulse stream is efficiently mapped back to the incident Poisson process, so that the entire channel (transmitter, free-space channel, photodetector, receiver) is modeled as Poisson. We first address the choice of modulation for this channel.

The modulation must satisfy the following physical constraints: a maximum average signal photon transmission rate of P_{av}, a maximum peak signal photon transmission rate of P_{pk}, and a minimum pulsewidth constraint of T_p. Wyner[1] showed that under peak and average power constraints there is negligible loss in capacity by restricting the transmitted signal to binary values over discrete time increments (slots). Hence, we restrict ourselves to binary, slotted modulations.

Letting n_b be the mean noise photons in a slot and n_s the mean signal photons in a pulsed slot, the channel then reduces to a binary-input, discrete-output Poisson channel, with signal and noise slot mass functions

$$p_1(y) = \frac{(n_s+n_b)^y e^{-(n_s+n_b)}}{y!}$$
$$p_0(y) = \frac{(n_b)^y e^{-(n_b)}}{y!}$$

Letting $1/M$ be the probability of a pulse transmission, or duty cycle, we may restate the peak and average power constraints as

$$n_s \le P_{pk}$$
$$\frac{n_s}{M} \le P_{av}$$

Parameterized by the duty cycle, no modulation for this channel can exceed the capacity of on-off-keying

$$C_{OOK} = \frac{1}{M} E_{Y_1} \log \frac{p_1(Y_1)}{p(Y_1)} + \frac{M-1}{M} E_{Y_0} \log \frac{p_0(Y_0)}{p(Y_0)} \text{bits/slot} \quad (1)$$

where Y_1 has distribution p_1, Y_0 has distribution p_0, and $p(y) = p_1(y)/M + p_0(y)(M-1)/M$.

It is well known that M-ary pulse-position-modulation (PPM), in which $\log_2(M)$ bits map to an M-slot symbol with a single pulse in one slot, with capacity

$$C = \frac{1}{M} E_{Y_1,\ldots,Y_M} \log_2 \left[\frac{ML(Y_1)}{\sum_{j=1}^{M} L(Y_j)} \right] \text{ bits/slot}$$

where $L(y) = p_1(y)/p_0(y)$, Y_1 is distributed as p_1 and $Y_j, j \neq 1$ is distributed as p_0, achieves capacity close to that of (1) at low duty cycles, e.g., [2], [3]. Current link budgets for deep-space links show that the optimum duty cycle is less than $1/32$ for the entirety of the (in orbit) mission [4], [5]. For these duty cycles, we see a negligible loss in choosing PPM.

Only when the optimum duty cycle is large, say greater than $1/16$, does PPM begin to suffer significant losses relative to OOK [6]. In this region, alternatives to PPM should be pursued, e.g., [7]. However, this occurs only for high average powers, which are not typical for current deep space missions.

A. Capacity of PPM

We describe a few properties of the capacity of PPM that will be relevant to the design. Expanding C in a Taylor series we have [5]

$$C = \left(\frac{M-1}{2M \ln 2} \right) \frac{n_s^2}{n_b} + O(n_s^3) \text{ bits/slot}$$

Hence, for small n_s, the capacity goes as n_s^2/n_b. To see the behavior as a function of slotwidth T_s, write the capacity as a function of the incident signal and noise rates

$$\lambda_s = \frac{n_s}{MT_s} \quad \text{photons/sec}$$
$$\lambda_b = \frac{n_b}{T_s} \quad \text{photons/sec}$$

and convert units to bits/sec yielding

$$\frac{C}{T_s} \approx \left(\frac{M(M-1)}{2 \ln 2} \right) \frac{\lambda_s^2}{\lambda_b} \text{ bits/sec.}$$

Hence the capacity is invariant to the slot width for small n_s. This is illustrated in Fig. 1, where λ_b is held constant and the capacity is plotted as a function of λ_s with slotwidths varying by powers of 10.

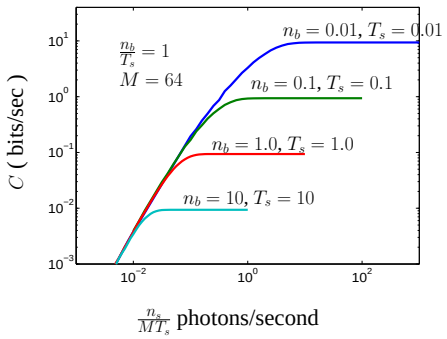

Fig. 1. Illustrating invariance of PPM capacity to slotwidth

Finally, the optimum PPM order in the presence of peak and average power constraints is always less than or equal to the peak to average power constraint ratio [5]. Hence the peak power constraint translates into an order constraint. For the lasers in consideration for MLCD, the imposed order constraint was $M \leq 64$.

B. Partitioning the received signal power/noise power space

The throughput of the coded modulation is given by

$$\text{Throughput } = \frac{\log_2 M}{M} R \frac{1}{T_s} \text{ bits/sec}$$

where R is the rate of the error correction code (ECC). The received signal and noise powers allow capacities that range from 0.1 Mbps to over 50 Mbps. A system designer can vary M, R and T_s, to tile this range of capacities. How should one accomplish this with limited complexity? Each parameter has physical constraints and complexity costs. As noted above, we should choose $M \leq P_{pk}/P_{av}$. We are also bounded by $T_s \geq T_p$. Finally, the required speed of the hardware for a fixed user data rate goes as $1/R$.

The capacity increases with the inverse of the slot width. Hence, an optimum solution is to fix $T_s = T_p$, choose the optimum M as a function of λ_s for the given n_b, and select a code rate of $R = CM/(\log_2(M)T_s)$ (assuming a close-to-capacity achieving code). However, this would require a wide range of changes to the code rate, which is prohibitively complex. Noting that the capacity is roughly invariant to the slotwidth, one can vary the throughput by varying the slotwidth and PPM order, while keeping the rate fixed. A rate of $1/2$ was chosen to tradeoff power and bandwidth efficiency.

Due to fluctuations of the received signal and noise power and errors in the estimation of these parameters, it would be impractical to have partitions of the signal and noise power space to a resolution greater than a few tenths of a dB. Figure 2 illustrates partitions chosen for MLCD, optimized under the constraints $M \leq 64$ and T_s a multiple of 1.6 ns, and overlaid with a representative range of received signal powers. Incident signal and noise powers are scaled by photodetector/receiver losses and a system margin, so that photon rates at the input to the decoder are $\lambda_s - 8.23$ dB and $\lambda_b - 5.23$ dB.

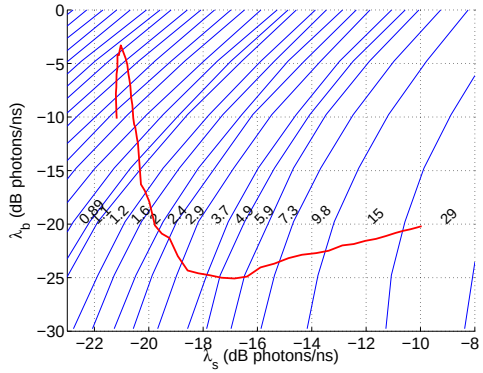

Fig. 2. Incident signal and noise powers and partitioning of the space. Partitions labeled with achievable throughput in Mbps

III. Code Selection

We now address the selection of an appropriate ECC. We consider four classes of codes: Reed-Solomon, parallel concatenated convolutional, serially concatenated convolutional, and low-density parity-check.

A. Serially concatenated convolutional code

We were able to obtain the best performance by treating PPM as an inner code in a serial concatenation with a short-constraint length outer convolutional code, a structure illustrated in Fig. 3 (the cyclic-redundancy-check (CRC) is used for error detection and as a stopping rule). Massey [8] first investigated the serial concatenation of a convolutional code with PPM with non-iterative decoding. Iteratively decoding the same concatenation, as done in [9] for convolutional codes, provides remarkably good performance–a gain of 4.8 dB over non-iterative decoding, within less than 1 dB from capacity over the wide range of operating points for the deep-space Mars link–with moderate encoding and decoding complexity. We denote this combination as SCPPM, and will go into detail on the design of the components in the following section. Figure 4 illustrates performance of SCPPM and several considered alternatives for a nominal operating point. We briefly discuss other considered alternatives in the following sections.

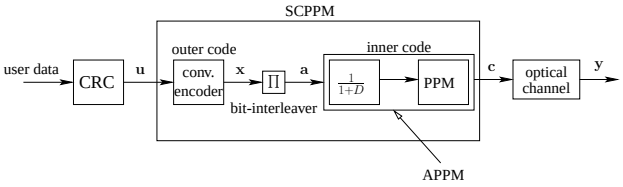

Fig. 3. Serially concatenated convolutionally coded PPM

B. Reed-Solomon codes

Reed-Solomon codes were first proposed for use on the PPM channel by choosing RS symbols to correspond to PPM symbols [10]. The convention of mapping code symbols to PPM symbols can be extended by grouping multiple PPM symbols to form a single RS symbol [11], yielding some improvement in performance. At nominal background rates, RS coded PPM suffers a 3 dB loss relative to soft-decision-codes (This is for hard-decision decoding–some of this loss would be recovered by soft-decision decoding).

C. Parallel concatenated convolutional codes

With the advent of turbo-codes [12], several authors investigated the application of parallel-concatenated-convolutional codes to the PPM channel. In [13], [14], the modulation was not iteratively decoded. In [15], which uses PPM on a discrete-time memoryless Rayleigh-fading channel where a small duty cycle is also optimal, PPM is iteratively decoded, providing performance close to capacity. However, we find that we have the same or better performance with lower complexity by using a lower complexity outer code.

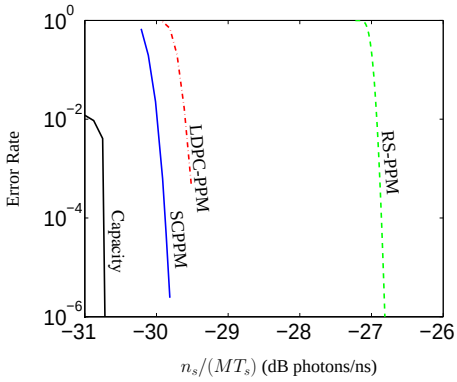

Fig. 4. SCPPM, LDPC-PPM and RS-PPM performance, $M = 64$, $n_b = 0.2$, $T_s = 32$ ns

D. Low-density parity-check codes

Low-density parity-check (LDPC) codes have proven to have performance comparable or better than turbo-like codes on the Gaussian channel while having a decoder that is more amenable to high-speed implementation. We carried out some designs [16], using the approach in [17] to derive a good degree distribution, and the approach in [18] to implement the degree distribution as a protograph. We demonstrated performance about 0.5 dB worse than that for the SCPPM code. We expect that with additional effort in design, LDPC-coded PPM performance will match or exceed that of SCPPM. However, for the deep space channel, the onus on complexity lies with the space-based encoder more so than with the single, Earth-based, decoder, and an LDPC encoder may have higher complexity than the SCPPM encoder. At the time of selection, performance, maturity of technology, and complexity factors favored the SCPPM code.

IV. SCPPM Code Design

In the following sections we describe the design of the SCPPM code: component codes, block length, interleaver, and bit-to-symbol mapping.

A. Component code selection

We utilize EXIT chart analysis [19] to select inner and outer codes. We first examine the chart for PPM. Let $L(A), L_e(A)$ be the apriori and extrinsic log-likelihoods of A and A, X be defined by Fig. 3. We assume that $L_e(A)$ is Gaussian and symmetric, i.e, $l = \log(f_{L_e}(l)/f_{L_e}(-l))$. The initial mutual information is given by

$$I(A; L_e(A)) = 1 - E\left(\log_2\left(\frac{\sum_{i=1}^{M}(1+n_s/n_b)^{Y_i}}{\sum_{i=1}^{M/2}(1+n_s/n_b)^{Y_i}}\right)\right)$$

where Y_1 is distributed as p_1 and $Y_j, j = 2, \ldots, M$ are distributed as p_0. In the limit of large mean $L(A)$, the mean extrinsic output of the inner code will approach

$$E[L_e(A)] \underset{E[L(A)]\to\infty}{\longrightarrow} n_s \log\left(1 + \frac{n_s}{n_b}\right)$$

and the mutual information approaches

$$I(A;L_e(A)) \underset{E[L(X)]\to\infty}{\longrightarrow} 1 - E\left[\log_2\left(1+\left(1+\frac{n_s}{n_b}\right)^{(Y_2-Y_1)}\right)\right]$$

where Y_1 is distributed as p_1 and Y_2 as p_0. Hence, the information we can extract from demodulating PPM saturates at a constant independent of the order. The relative gain from demodulation increases with the PPM order, with no gain when $M = 2$.

By preceding the PPM mapping with a recursive $1/(1+D)$ binary accumulator, a combination referred to as $APPM$, we prevent the saturation constraint. Fig. 5 illustrates EXIT charts for $M = 64$-ary PPM and APPM with $n_b = 0.2$ and $n_s = 2.0$. The figure illustrates the well-known tradeoff of adding a recursive element to the inner code–a transfer function that goes to one at the cost of a slightly lower threshold.

Constraining the selection of outer code to rate $1/2$ convolutional codes, we found the constraint length 3, $(5,7)$ (octal taps), convolutional code to be a good fit. The joint EXIT chart with a simulated transfer function for this combination is illustrated in Fig. 6.

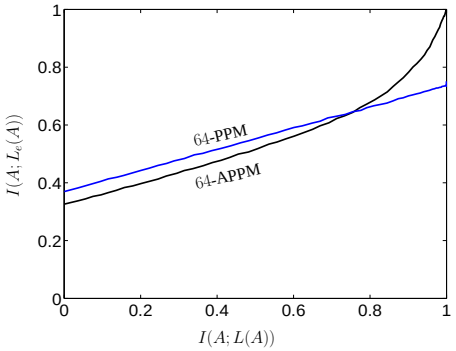

Fig. 5. EXIT charts for PPM and APPM, $M = 64$, $n_s = 2$, $n_b = 0.2$

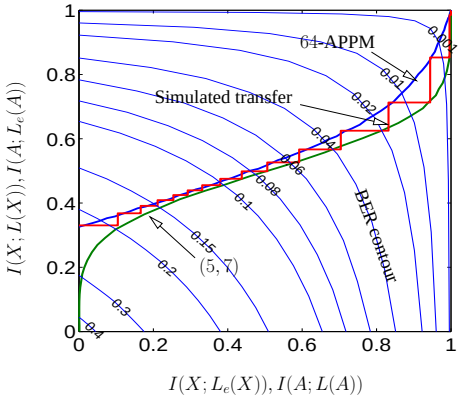

Fig. 6. EXIT chart for SCPPM

B. Block Length–Error Rate Bounds

Let $\mathcal{C}$ be the collection of SCPPM codewords $\mathbf{c}$ and $\mathcal{C}_h(\mathbf{c})$ the collection of codewords at distance h from $\mathbf{c}$. The word error rate is given by

$$P_w = \frac{1}{2^k}\sum_{\mathbf{c}\in\mathcal{C}}\sum_{h=2}^{2n/\log_2(M)} P[e_h] \tag{2}$$

where n is the length of $\mathbf{x}$ and $P[e_h]$ is the probability of a distance h event given $\mathbf{c}$ was transmitted. The pairwise union-Bhattacharyya bound may be tightened below the critical signal power using a similar approach to that used for the Gaussian channel, e.g., [20]. Let $\mathcal{R}$ be the region $\mathcal{R} = \{\mathbf{y} : \langle \mathbf{y}, \mathbf{x}\rangle \geq \epsilon\}$, with $\langle \mathbf{y}, \mathbf{x}\rangle$ the inner product, and expand $P[e_h]$ as

$$\begin{aligned} P[e_h] &= P[e_h, \mathbf{y} \notin \mathcal{R}] + P[e_h, \mathbf{y} \in \mathcal{R}] \\ &\leq P[\mathbf{y} \notin \mathcal{R}] + P[e_h, \mathbf{y} \in \mathcal{R}] \end{aligned} \tag{3}$$

The first term is given by

$$P[\mathbf{y} \notin \mathcal{R}] = F_S(\epsilon - 1)$$

where F_S is a cumulative distribution and S is Poisson with mean $\lambda_S = (n/\log_2(M))(n_s + n_b)$. The second term may be bounded via an application of the union bound

$$P[e_h, \mathbf{y} \in \mathcal{R}] \leq \sum_{\hat{\mathbf{c}}\in\mathcal{C}_h(\mathbf{c})} P[\epsilon - U \leq T \leq N] \tag{4}$$

where U is the inner product of $\mathbf{y}$ and the pulse positions common to $\mathbf{x}$ and $\hat{\mathbf{x}}$ so that U, T, N are independent Poisson with means

$$\begin{aligned} \lambda_T &= \frac{h}{2}(n_s + n_b) \\ \lambda_U &= \lambda_S - \lambda_T \\ \lambda_N &= \frac{h}{2}n_b \end{aligned}$$

Note that (4) is a function only of h and invariant to $\mathbf{c}$. Hence, although PPM is nonlinear, we may drop the dependence on $\mathbf{c}$ and choose a convenient value to simplify analysis. Combining terms yields the bound

$$P_w \leq \sum_{h=2}^{2n/M}\left(F_S(\epsilon-1) + \sum_{\hat{\mathbf{c}}\in\mathcal{C}_h(\mathbf{c})} P[\epsilon - U \leq T \leq N]\right) \tag{5}$$

In order to find the optimum ϵ to minimize (3), we restrict the search to integer ϵ and expand the second term as

$$\begin{aligned} P[\epsilon - U \leq T \leq N] &= \sum_{u=0}^{\infty}\sum_{n=\epsilon-u}^{\infty}\sum_{t=\epsilon-u}^{n} f_U(u)f_N(n)f_T(t) \\ &= P[T \leq N]P[U \geq \epsilon] + \\ &\quad \sum_{u=0}^{\epsilon-1}\sum_{n=\epsilon-u}^{\infty}\sum_{t=\epsilon-u}^{n} f_U(u)f_N(n)f_T(t) \end{aligned}$$

which lends itself to efficient numerical computation by pre-computing certain terms.

Fig. 7 illustrates union-Bhattacharyya bounds, simulation results and expurgated versions of (5) [21] for the SCPPM code with blocklengths of 516, 1026 and 2052 bits. We see that block lengths greater than 2 Kbits have an error floor below 10^{-8}, with negligible impact for our purposes.

A interleaver length of $15120 = 2^4 \times 3^3 \times 5 \times 7$ was chosen for MLCD as a good tradeoff between threshold and complexity. This value is divisible by $\log_2 M$ for a large range of M, allowing easy integration with various PPM orders.

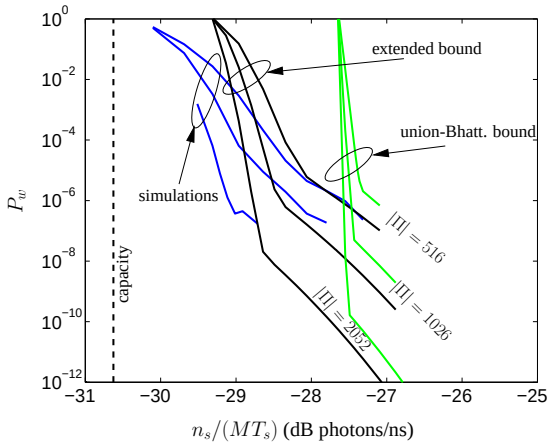

Fig. 7. Word error rates: simulations, capacity, union-Bhattacharyya bound and extended bound, $M = 64$, $n_b = 0.2$

C. Interleaver

Good performance has been demonstrated for long block-length codes with pseudo-random spread interleavers. However, pseudo-random functions may be complex to implement in hardware. In [22], quadratic permutation polynomials are proposed for use as an interleaver and we adopted a permutation polynomial for SCPPM. We observed no loss in performance relative to spread-interleavers with the polynomial interleaver $f(x) = 11x + 210x^2$, with inverse $f^{-1}(x) = 7331x + 7770x^2$, which is conveniently also quadratic.

D. Bit-to-symbol Mapping

On a truly memoryless channel, the bit-to-PPM-symbol mapping has no impact on performance. However, in practice, finite detector pulsewidths and timing jitter may add appreciable inter-slot-interference (ISI). If the detector/receiver pulsewidth is on the order of a slotwidth, the effect of ISI will be large. We found the best performance in the presence of ISI with an anti-Gray mapping–anti-Gray in the sense that we map pairs of $\log_2 M$ bits at the input to the accumulator with distance $\log_2(M) - 1$ to PPM symbols with pulses in adjacent slots–those symbols closest in distance in the presence of ISI.

ACKNOWLEDGMENT

The design of the coded modulation for MLCD was done in collaboration with Richard Barron and Don Boroson of Lincoln Laboratories.

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

REFERENCES

[1] A. D. Wyner, "Capacity and error exponent for the direct detection photon channel–Part I," *IEEE Transactions on Information Theory*, vol. 34, pp. 1449–1461, Nov. 1988.

[2] J. R. Pierce, "Optical channels: Practical limits with photon counting," *IEEE Transactions on Communications*, vol. COM-26, pp. 1819–1821, Dec. 1978.

[3] R. G. Lipes, "Pulse-position-modulation coding as near-optimum utilization of photon counting channel with bandwidth and power constraints," *DSN Progress Report*, vol. 42, pp. 108–113, Apr. 1980.

[4] A. Biswas and S. Piazzolla, "Deep-space optical communications downlink budget: Mars system parameters," *IPN Progress Report*, July 2003. to appear.

[5] B. Moision and J. Hamkins, "Deep-space optical communications downlink budget: Modulation and coding," in *IPN Progress Report*, vol. 42-154, Aug. 2003.

[6] J. Hamkins and B. Moision, "Multipulse pulse position modulation on discrete memoryless channels." submitted to IEEE Transactions on Communications, Nov. 2005.

[7] R. Barron, "Binary shaping for low duty-cycle communications," in *Proceedings of IEEE International Symposium on Information Theory*, p. 513, 2004.

[8] J. L. Massey, "Capacity, cutoff rate, and coding for a direct-dectection optical channel," *IEEE Transactions on Communications*, vol. COM-29, pp. 1615–1621, Nov. 1981.

[9] S. Benedetto, D. Divsalar, G. Montorsi, and F. Pollara, "Serial concatenation of interleaved codes: Performance analysis, design and iterative decoding," *Telecommunications and Data Acquisition Progress Report*, vol. 42-126, Aug. 1996.

[10] R. J. McEliece, "Practical codes for photon communication," *IEEE Transactions on Information Theory*, vol. IT-27, pp. 393–398, July 1981.

[11] J. Hamkins and B. Moision, "Performance of long blocklength Reed-Solomon codes with low-order pulse position modulation." submitted to IEEE Transactions on Communications, Nov. 2005.

[12] C. Berrou and A. Glavieux, "Near optimum error correcting coding and decoding: Turbo-codes," *IEEE Transactions on Communications*, vol. 44, pp. 1276–1271, Oct. 1996.

[13] K. Kiasaleh, "Turbo-coded optical PPM communication systems," *Journal of Lightwave Technology*, vol. 16, pp. 18–26, Jan. 1998.

[14] J. Hamkins, "Performance of binary turbo coded 256-PPM," *TMO Progess Report*, vol. 42, pp. 1–15, Aug. 1999.

[15] M. Peleg and S. Shamai, "Efficient communication over the discrete-time memoryless Rayleigh fading channel with turbo coding/decoding," *European Transactions on Telecommunications*, vol. 11, pp. 475–485, September-October 2000.

[16] M. Barsoum and B. Moision, "An LDPC code for optical space communication," tech. rep., Jet Propulsion Laboratory, Oct. 2005. Inter-Office-memorandum.

[17] S. ten Brink, G. Kramer, and A. Ashikhmin, "Design of low-density parity-check codes for modulation and detection," *IEEE Transactions on Communications*, vol. 52, pp. 670–678, April 2004.

[18] J. Thorpe, K. Andrews, and S. Dolinar, "Methodologies for designing LDPC codes using protographs and circulants," in *isit*, p. 236, 2004.

[19] S. Brink, "Convergence of iterative decoding," *Electronics Letters*, vol. 35, pp. 806–808, May 1999.

[20] D. Divsalar and E. Biglieri, "Upper bounds to error probabilities of coded systems beyond the cutoff rate," *IEEE Transactions on Communications*, vol. 51, pp. 2011–2018, December 2003.

[21] B. Moision and J. Hamkins, "Error rate bounds for coded PPM on the Poisson channel." submitted to IEEE Transactions on Communications, Nov 2005.

[22] J. Sun and O. Y. Takeshita, "Interleavers for turbo codes using permutation polynomials over integer rings," *IEEE Transactions on Information Theory*, vol. 51, pp. 101–119, Jan. 2005.