

Green Codes: Energy-Efficient Short-Range Communication

Pulkit Grover and Anant Sahai
Wireless Foundations, Department of EECS
University of California at Berkeley, CA-94720, USA
{pulkit, sahai}@eecs.berkeley.edu

Abstract—A green code attempts to minimize the total energy per-bit required to communicate across a noisy channel. The classical information-theoretic approach neglects the energy expended in processing the data at the encoder and the decoder and only minimizes the energy required for transmissions. Since there is no cost associated with using more degrees of freedom, the traditionally optimal strategy is to communicate at rate zero.

In this work, we use our recently proposed model for the power consumed by iterative message passing. Using generalized sphere-packing bounds on the decoding power, we find lower bounds on the total energy consumed in the transmissions and the decoding, allowing for freedom in the choice of the rate. We show that contrary to the classical intuition, the rate for green codes is bounded away from zero for any given error probability. In fact, as the desired bit-error probability goes to zero, the optimizing rate for our bounds converges to 1.

I. INTRODUCTION

With the development of billion transistor chips, the range of communication has come down dramatically from hundreds of kilometers (e.g. deep space communication) to a few meters (e.g. ad-hoc wireless networks) or a few millimeters or even less (e.g. on chip communication). To communicate over smaller distances, the transmit power required is much smaller. At these distances, the energy used in transmissions can be comparable to that expended by the system processes. The small size limits the ability of these chips to dissipate heat. Further, the chip might be battery operated, imposing stringent constraints on its energy usage. It is therefore of interest to design coding techniques that minimize the *total energy* consumed, which includes the transmission energy as well as the processing energy. We refer to the coding techniques that minimize the total energy as *green codes*.

The classical information theoretic approach finds the minimum *transmission* energy required to communicate reliably across the channel. The approach is motivated by long-range communication, that corresponds to power constrained channels. Shannon [1] first characterized the minimum energy required to communicate across a channel with fixed rate. The resulting bounds are expressed using ‘waterfall’ curves that convey the revolutionary idea that unboundedly low probabilities of bit-error are attainable using only finite transmit power. This characterization raises a natural question: what is the minimum energy required for communication that is free of a rate constraint? The classical approach [2] [3] gives the minimum transmission energy required (on average) to

communicate one bit reliably across the channel. For example, for an AWGN channel of noise variance 1, this minimum energy is

$$\lim_{P_T \rightarrow 0} \frac{P_T}{C(P_T)} = 2 \ln(2) \text{ Joules.} \quad (1)$$

Since there is no penalty associated with lower rates, it is good to use as many degrees of freedom as are available, and the optimal transmission rate is zero.

The problem of minimizing combined transmission and processing energy is well studied in networks. The common thread in [4], [5], [6], [7], [8], [9] is that the energy consumed in processing the signals can be a substantial fraction of the total power. In [7], an information-theoretic formulation is considered. The authors model the processing energy by a constant ϵ per unit time when the transmitter is transmitting (and hence, is in the ‘on’ state). A total of r channel uses are allowed, and the total energy available is $r\mathcal{E}$, where \mathcal{E} is a constant. Let P_i be the transmit power at i -th time instant, and let $C(P_i)$ be the capacity of the corresponding channel. Then the problem is to transmit maximum number of bits with the total power less than $r\mathcal{E}$. That is,

$$\max \sum_{i=1}^r 1_i C(P_i) \quad (2)$$

$$\text{subject to } \sum_{i=1}^r 1_i \times (P_i + \epsilon) \leq r\mathcal{E} \quad (3)$$

where $1_i = 1$ if a symbol is transmitted in the i -th channel use, and is 0 otherwise. This is equivalent to dividing the channel into r sub-channels, with independent coding on each sub-channel. Since the capacity function $C(P)$ is concave in its argument, for maximizing the total number of information bits communicated, the transmission energy P_i should be equal for all i where $1_i = 1$. Without accounting for the energy consumed by the system processes, the optimal strategy would be to use all the r parallel channels, and share the energy equally amongst them. However, the energy consumed by the system processes imposes a fixed penalty on each channel use. The authors quantify this tension by measuring ‘burstiness’ Θ of signaling defined as $\Theta = \frac{1}{r} \sum_{i=1}^r 1_i$.

The transmissions should not be too bursty because of the law of diminishing returns associated with the $\log(\cdot)$ function. On the other hand, the transmission strategy should not make use of all degrees of freedom either, since there is an ϵ cost

associated with the use of each degree of freedom. The authors conclude that for minimum total energy, $0 < \Theta < 1$. Contrary to conventional information theoretic wisdom, it is no longer optimal to use all available degrees of freedom. Consequently, the optimal rate that minimizes the total energy consumption is bounded away from zero. That is, *if processing energy is taken into account, green codes may not communicate at zero rate!*

The objective in [7] [5] [9] is to reduce the energy consumption for wireless devices that consume energy continuously when operating e.g. hand-held computers, high-end laptops, etc. Energy consumption per unit time for such devices is indeed well modeled by a constant possibly independent of the coding strategy being used. In this paper, we are interested in the energy expended by the decoding process itself. The decoding circuit requires some non-zero energy to perform each operation. As opposed to energy consumed by system processes in [7], [5], [9], the decoding energy depends significantly on the code construction, the rate and the desired error probability, and therefore needs more careful modeling.

In this work, we study explicit models of energy expended at the decoder. Owing to their low implementation complexity, and hence low energy consumption, we concentrate on the message passing decoder. For this decoder, we derive lower bounds on the combined transmission and decoding energy, with no constraint on the rate. We show that the optimizing rate for green codes based on message passing decoding is indeed bounded away from zero. As the error probability decreases to zero, the optimizing rate increases. In a result that is qualitatively different from those in [7], we show that there is no advantage in increasing the rate beyond 1. Therefore, as the error probability converges to zero, the optimizing rate converges to 1!

The organization of the paper is as follows : In Section II, we introduce the channel model, the decoder model, and the energy model. In Section III, we summarize some of our results in [10]. In Section IV, we build on the results in [10] to find bounds on the minimum total energy required to communicate across a channel, with no rate constraint, taking into account the decoding energy as well. We conclude in Section V.

II. SYSTEM MODEL

Consider a point-to-point communication link. An information sequence \mathbf{B}_1^k is encoded into 2^{mR} codeword \mathbf{X}_1^m , using a possibly randomized encoder. The observed channel output is \mathbf{Y}_1^m . The information sequences are assumed to consist of iid fair coin tosses and hence the rate of the code is $R = k/m$.

The channel model considered is an average power constrained AWGN channel of noise variance $\sigma_{\mathcal{P}}^2$. We also obtain some results for the BSC arising from performing hard-decision on BPSK symbols transmitted over an AWGN channel. The true channel is denoted by \mathcal{P} . The channel capacity is denoted by $C_{\sigma^2}(P_T)$, where σ^2 is the noise variance, and P_T is the average power constraint. We drop σ^2 from this notation when no ambiguity is created in doing so.

For maximum generality, we do not impose any *a priori* structure on the code itself. Instead, inspired by [11], [12], [13], we focus on the parallelism of the decoder and the energy consumed within it. We assume that the decoder is physically made of computational nodes that pass messages to each other in parallel along physical (and hence unchanging) wires. A subset of nodes are designated ‘message nodes’ in that each is responsible for decoding the value of a particular message bit. Another subset of nodes (not necessarily disjoint), called the ‘observation nodes’ has members that are each initialized with at most one observation of the received channel output symbols. There may be additional computational nodes to merely help in decoding.

The implementation technology is assumed to dictate that each computational node is connected to at most $\alpha + 1 > 2$ other nodes¹ with bidirectional wires. No other restriction is assumed on the topology of the decoder. In each iteration, each node sends (possibly different) messages to all its neighboring nodes. **No restriction is placed on the size or content of these messages except for the fact that they must depend only on the information that has reached the computational node in previous iterations.** If a node wants to communicate with a more distant node, it has to have its message relayed through other nodes. The neighborhood size at the end of l iterations is denoted by $n \leq \alpha^{l+1}$. Each computational node is assumed to consume a fixed E_{node} joules of energy at each iteration.

Let the average probability of bit error of a code be denoted by $\langle P_e \rangle$ when it is used over channel \mathcal{P} . The main tool is a lower bound on the neighborhood size n as a function of $\langle P_e \rangle$ and R . This then translates into a lower bound on the number of iterations that can in turn be used to lower bound the required decoding power.

Throughout this paper, we allow the encoding and decoding to be randomized with all computational nodes allowed to share a pool of common randomness. We use the term ‘average probability of error’ to refer to the probability of bit error averaged over the channel realizations, the messages, the encoding, and the decoding.

III. LOWER BOUNDS ON THE DECODING COMPLEXITY AND TOTAL ENERGY

In this section we summarize our results for lower bounds on decoding complexity for an AWGN channel from [10]. The main bounds are given by theorems that capture a local sphere-packing effect. These can be turned around to give a family of lower bounds on the neighborhood size n as a function of $\langle P_e \rangle$ and R . Using a simple lower bound on the number of iterations, $l \geq \frac{\log(n)}{\log(\alpha)} - 1$, we get a lower bound² on complexity. The family of lower bounds is indexed by the choice of a hypothetical channel \mathcal{G} and the bounds can be optimized numerically for any desired set of parameters.

¹In practice, this limit could come from the number of metal layers on a chip. $\alpha = 1$ would just correspond to a big ring of nodes and is therefore uninteresting.

²We approximate this by $l \geq \frac{\log(n)}{\log(\alpha)}$ for the rest of the paper.

Theorem 3.1: For the AWGN channel and the decoder model in Section II, let n be the maximum size of the decoding neighborhood of any individual message bit. The following lower bound holds on the average probability of bit error.

$$\langle P_e \rangle \geq \sup_{\sigma_G^2: C_{\sigma_G^2}(P_T) < R} \frac{h_b^{-1}(\delta(\sigma_G^2))}{2} \exp\left(-nD(\sigma_G^2 || \sigma_P^2)\right) - \sqrt{n} \left(\frac{3}{2} + 2 \ln \left(\frac{2}{h_b^{-1}(\delta(\sigma_G^2))} \right) \right) \left(\frac{\sigma_G^2}{\sigma_P^2} - 1 \right), \quad (4)$$

where $\delta(\sigma_G^2) = 1 - C_{\sigma_G^2}(P_T)/R$, the capacity $C_{\sigma_G^2}(P_T) = \frac{1}{2} \log_2 \left(1 + \frac{P_T}{\sigma_G^2} \right)$, and the KL divergence $D(\sigma_G^2 || \sigma_P^2) = \frac{1}{2} \left[\frac{\sigma_G^2}{\sigma_P^2} - 1 - \ln \left(\frac{\sigma_G^2}{\sigma_P^2} \right) \right]$.

Proof: See [10]. There is a better bound in [10] as well. This bound is presented here for ease of exposition. ■

Observe that the required value of n increases as $\langle P_e \rangle$ decreases. Taking log on both sides of (4), it is evident that for small $\langle P_e \rangle$, the term $nD(\sigma_G^2 || \sigma_P^2)$ dominates the other terms in the RHS. For small $\langle P_e \rangle$, σ_G^2 can be taken close to σ_G^{*2} that satisfies $C_{\sigma_G^{*2}}(P_T) = R$. Neglecting the other two terms, we get

$$n \gtrsim \frac{\log(1/\langle P_e \rangle)}{D(\sigma_G^{*2} || \sigma_P^2)}. \quad (5)$$

IV. MINIMIZATION OF TOTAL ENERGY BY OPTIMIZING OVER THE RATE AND TRANSMIT POWER.

Consider the total energy spent in transmission. For transmitting k bits at rate R , the number of channel uses is $m = k/R$. If each transmission has power $\xi_T P_T$, the total energy used in the transmissions is $\xi_T P_T m$.

At the decoder, let the number of iterations be l . Assume that each node consumes E_{node} joules of energy in each iteration. The number of computational nodes can be lower bounded by m , the number of received channel outputs, and also by k , the number of bits to be decoded. We lower bound by the maximum of the two³

$$E_{dec} \geq E_{node} \times \max\{k, m\} \times l. \quad (6)$$

There is no lower bound on the encoding complexity and so the encoding is considered free. For m transmissions with average power P_T , we require mP_T joules of energy. This results in the following bound for the weighted total energy⁴

$$E_{total} \geq \xi_T m P_T + \xi_D E_{node} \max\{k, m\} \times l. \quad (7)$$

³A lower bound of $m + k$ would not allow for node sharing between the set of observation nodes and the message nodes.

⁴The parameters ξ_T and ξ_D are weights assigned to the transmit and the decoding energy respectively. ξ_T depends on the path-loss across the channel. ξ_D indicates the relative importance of decoding energy. For example, if the energy use at the decoder is severely constrained, ξ_D would be large.

Using $l \geq \frac{\log(n)}{\log(\alpha)}$,

$$E_{total} \geq m \xi_T P_T + \frac{\xi_D E_{node} \max\{k, m\} \log(n)}{\log(\alpha)} \propto \frac{m P_T}{\sigma_P^2} + \gamma \max\{k, m\} \log(n), \quad (8)$$

where $\gamma = \frac{\xi_D E_{node}}{\sigma_P^2 \xi_T \log(\alpha)}$ is a constant that summarizes all the technological and environmental terms. The expression in (8) gives the normalized total energy, normalized by the noise variance σ_P^2 . Figure 1 provides example⁵ behavior of γ with distance. The neighborhood size n itself can be lower bounded by plugging the desired average probability of error into Theorem 3.1.




Fig. 1. The plot shows the behavior of γ with distance d for path loss $\frac{1}{\xi_T} = \frac{1}{d^2}$ for $d > 0.1\text{mm}$ (and path loss 1 for smaller d). E_{node} is 1pJ, $\alpha = 4$, $\xi_D = 1$, and $\sigma_P^2 = 4 \times 10^{-21}\text{J}$. The energy per bit is normalized by σ_P^2 .

We thus obtain the following expression for the minimum normalized total energy,

$$E_{\text{per bit}} = \min_{P_T, R} \frac{1}{R} \frac{P_T}{\sigma_P^2} + \frac{1}{R} \gamma \max \left\{ \frac{1}{R}, 1 \right\} \log(n). \quad (9)$$

Observe that in (9), the decoding energy increases as the error probability decreases for constant transmit power and rate. This behavior is not reflected by using the model inspired from [7] for decoding energy. The bounds in [7] are for error probability converging to zero. To compare our bounds with the black-box model of [7], in Appendix I we derive bounds for non-zero error probability based on the model in [7]. We plot the two bounds against each other in Figure 2 for $k = 10,000$ bits.

We choose $\epsilon = 4$, for which the total energy per bit for the black-box model equals the energy per bit for $\gamma = 0.2$ for our bound for $\langle P_e \rangle = 10^{-13}$. The figure shows that for

⁵The energy cost of one iteration at one node $E_{node} \approx 1$ pJ is arrived at by an optimistic extrapolation from the reported values in [14], [15], thermal noise energy per sample $\sigma_P^2 \approx 4 \times 10^{-21}\text{J}$ from kT with T around room temperature.

$\langle P_e \rangle$ smaller than this threshold, the model inspired from [7] underestimates the total energy. It is because this model treats the decoder as a black-box where ϵ does not change with error probability or rate.

It is interesting to observe what values of R optimize (9). Under the small $\langle P_e \rangle$ approximation in (5), we now heuristically argue that the optimal rate R_{opt} should converge to 1 as $\langle P_e \rangle \rightarrow 0$.

Observe that for $R < 1$,

$$\begin{aligned} E_{\text{per bit}} &= \frac{P_T}{\sigma_p^2 R} + \frac{\gamma}{R} \log_2(n) \\ &= \frac{P_T}{\sigma_p^2 R} + \frac{\gamma}{R} \log_2 \left(\log_2 \left(\frac{1}{\langle P_e \rangle} \right) \right) - \log_2(D(\sigma_G^{*2} || \sigma_P^2)) \end{aligned}$$

As $\langle P_e \rangle \rightarrow 0$, $n \rightarrow \infty$. Therefore, the decoding energy increases to infinity. Increasing the rate R at the cost of increasing P_T offsets the increasing decoding costs. However, for $R \geq 1$,

$$E_{\text{per bit}} \gtrsim \frac{P_T}{\sigma_p^2 R} + \gamma \log_2 \left(\frac{\log_2 \left(\frac{1}{\langle P_e \rangle} \right)}{D(\sigma_G^{*2} || \sigma_P^2)} \right), \quad (10)$$

which indicates there is no advantage in increasing rate beyond $R = 1$, since it no longer decreases the decoding energy.

Evidently, for finite $\langle P_e \rangle$, there exists an optimal rate $R_{\text{opt}} > 0$ that minimizes the combined energy consumed. Using numerical evaluation of the bound (9), we plot the behavior of the optimal rate with $\langle P_e \rangle$ in Figure 5. The plots demonstrate that the optimal rate indeed converges to 1.

Figure 3 shows the behavior of our lower bound on sum energy with $\langle P_e \rangle$ for various values of γ . Figure 4 shows that similar behavior also holds for a BSC arising from performing hard-decision on BPSK symbols transmitted over an AWGN channel. The optimal rate for this channel also converges to 1 as $\langle P_e \rangle \rightarrow 0$. Due to lack of space, we omit the plots.

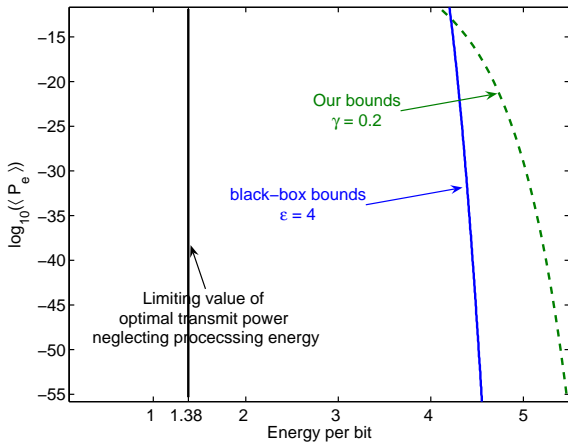


Fig. 2. The plot shows the comparison of lower bounds on the minimum normalized energy for $k = 10,000$ bits. The ‘black-box bounds’ plot is based on the model in [7], where the details of the processor are ignored. Our bounds take into account the decoder structure as well.

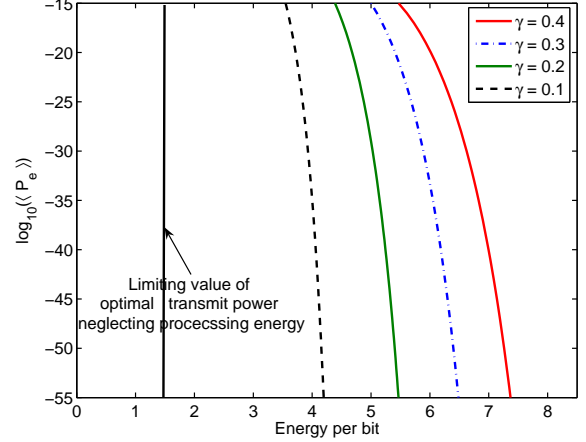


Fig. 3. The plot shows the behavior of lower bound on the normalized sum energy with $\langle P_e \rangle$ for various values of γ . The sum energy goes to infinity as $\langle P_e \rangle \rightarrow 0$.

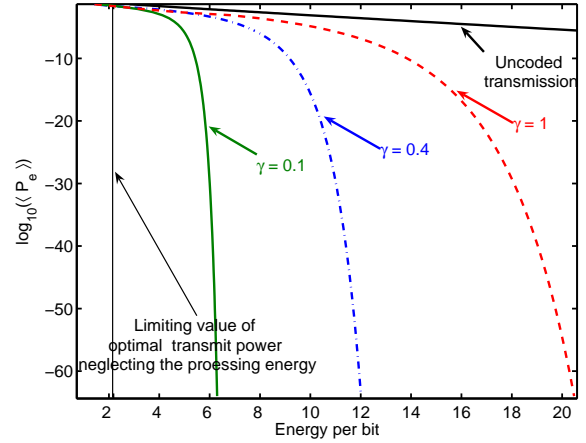


Fig. 4. The plot shows the behavior of lower bound on normalized sum energy with $\langle P_e \rangle$ for various values of γ for a BSC arising from performing hard-decision on BPSK symbols transmitted over an AWGN channel. The optimizing rate converges to 1 as $\langle P_e \rangle \rightarrow 0$. Even so, this plot shows that the optimal strategy is not uncoded transmission at low $\langle P_e \rangle$ since coded transmission outperforms uncoded transmission at small $\langle P_e \rangle$.

V. DISCUSSIONS AND CONCLUSIONS

In this work, we derived lower bounds on the combined transmission and decoding energy for iterative decoding with unconstrained rates. It is important to note that these are lower bounds, and the actual energy consumption would only be higher. An interesting feature of our bounds is that the optimizing rate for green codes is bounded away from zero, and, in fact, converges to 1 as the error probability converges to zero. This is qualitatively different from a pure black-box modeling of the decoding process, where energy consumption is independent of the desired error probability and the rate. In that case, as observed in [7], the optimal rate is a constant that can be greater than 1.




Fig. 5. Optimal value of rate vs error probability: As $\langle P_e \rangle$ converges to 0, the optimizing rate converges extremely slowly to 1.

For an AWGN channel, the value 1 for optimal rate is a result of a bit-wise representation of the information at the decoder. If, however, the message nodes represent the information in base M then the optimizing rate would converge to $\log_2(M)$.

For the BSC arising from performing hard-decision on BPSK symbols transmitted over an AWGN channel, the optimal rate still converges to 1. The rate is upper bounded by 1 because the channel has binary input alphabet, and thus this case might seem somewhat uninteresting. However, uncoded transmission over BSC also corresponds to rate 1, which might falsely suggest that uncoded transmission is asymptotically optimal for minimizing the total energy. We observe that despite the optimal rate approaching 1, coded transmission attains the same error probability with much smaller total energy than uncoded transmission.

We note that the total energy per-bit required to communicate at arbitrarily low error probability increases to infinity for the message passing decoder. This is in contrast to the classical information-theoretic result for transmit power, which shows that the transmit power is bounded even as $\langle P_e \rangle \rightarrow 0$. Based on results in [10], the total energy per bit increases to infinity for most known codes and decoding algorithms. It would be interesting to extend this result to all possible codes and decoding algorithms. An approach based on laws of physics is suggested in [10] for the fixed rate problem. The approach might yield results here as well.

APPENDIX I

BOUNDS IN [7] FOR NON-ZERO ERROR PROBABILITY

Observe that the results in [7] are for $\langle P_e \rangle \rightarrow 0$ and infinitely many information bits. Parallel to our analysis for message passing decoding, in this appendix, we build on the analysis in [7] to derive bounds on the minimum energy required for communicating with a non-zero error probability $\langle P_e \rangle$ and finite information bits.

Assume k bits are to be transmitted across the channel, with desired error probability $\langle P_e \rangle$. In [7], the authors maximize the information bits communicated under a total energy constraint. Turning around the problem in [7], we can instead minimize the total energy consumed given the number of bits transmitted. Now we can add an error probability constraint to the bits transmitted. Assume that a block code is used to communicate across the channel. The corresponding error exponent is bounded by the sphere-packing bound [16]. Assuming optimistically that the code actually achieves the sphere-packing bound in the exponent,

$$\langle P_e \rangle \leq P_{e,block} \approx e^{-m E_{sp}(P_T, R)}$$

where $E_{sp}(P_T, R)$ is the sphere-packing bound at rate R and transmit power P_T . The objective, therefore, is

$$\begin{aligned} & \min_{P_T, m} \quad m \times (P_T + \epsilon) \\ & \text{subject to} \quad m \times E_{sp}\left(P_T, \frac{k}{m}\right) = \ln\left(\frac{1}{\langle P_e \rangle}\right). \quad (11) \end{aligned}$$

REFERENCES

- [1] C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, pp. 379–423, 623–656, Jul./Oct. 1948.
- [2] J. Pierce, "Optical channels: Practical limits with photon counting," *IEEE Transactions on Communications*, pp. 1819–1821, Dec. 1978.
- [3] S. Verdú, "On channel capacity per unit cost," *IEEE Trans. Inform. Theory*, vol. 36, pp. 1019–1030, 1990.
- [4] P. Agrawal, "Energy efficient protocols for wireless systems," in *IEEE International Symposium on Personal, Indoor, Mobile Radio Communication*, 1998, pp. 564–569.
- [5] S. Cui, A. J. Goldsmith and A. Bahai, "Energy Constrained Modulation Optimization," *IEEE Trans. Wireless Commun.*, vol. 4, no. 5, pp. 1–11, 2005.
- [6] A. J. Goldsmith and S. B. Wicker, "Design challenges for energy constrained ad hoc wireless networks," *IEEE Trans. Wireless Commun.*, pp. 8–27, 2002.
- [7] P. Massaad, M. Medard, and L. Zheng, "Impact of Processing Energy on the Capacity of Wireless Channels," in *International Symposium on Information Theory and its Applications (ISITA)*, 2004.
- [8] S. Vasudevan, C. Zhang, D. Goeckel, and D. Towsley, "Optimal power allocation in wireless networks with transmitter-receiver power tradeoffs," *Proceedings of the 25th IEEE International Conference on Computer Communications INFOCOM*, pp. 1–11, Apr. 2006.
- [9] R. Kravertz and P. Krishnan, "Power management techniques for mobile communications," *The Fourth Annual ACM/IEEE International Conference on Mobile Computing and Networking*, 1998.
- [10] A. Sahai and P. Grover, "The price of certainty : "waterslide curves" and the gap to capacity," *Submitted to IEEE Transactions on Information Theory*, available online at <http://arXiv.org/abs/0801.0352v1>, Dec. 2007.
- [11] P. P. Sotiriadis, V. Tarokh, and A. P. Chandrakasan, "Energy reduction in VLSI computation modules: an information-theoretic approach," *IEEE Trans. Inform. Theory*, vol. 49, no. 4, pp. 790–808, Apr. 2003.
- [12] N. Shanbhag, "A mathematical basis for power-reduction in digital VLSI systems," *IEEE Trans. Circuits Syst. II*, vol. 44, no. 11, pp. 935–951, Nov. 1997.
- [13] T. Koch, A. Lapidoth, and P. P. Sotiriadis, "A channel that heats up," in *Proceedings of the 2007 IEEE Symposium on Information Theory*, Nice, France, 2007.
- [14] M. Bhardwaj and A. Chandrakasan, "Coding under observation constraints," in *Proceedings of the Allerton Conference on Communication, Control, and Computing*, Monticello, IL, Sep. 2007.
- [15] S. L. Howard, C. Schlegel, and K. Iniewski, "Error control coding in low-power wireless sensor networks: when is ECC energy-efficient?" *EURASIP Journal on Wireless Communications and Networking*, pp. 1–14, 2006.
- [16] C. E. Shannon, "Probability of error for optimal codes in a gaussian channel," *Bell System Technical Journal*, vol. 38, pp. 611–656, 1959.