

Capacity and Zero-Error Capacity of the Chemical Channel with Feedback

Haim Permuter, Paul Cuff, Benjamin Van Roy and Tsachy Weissman

Department of Electrical Engineering

Stanford University

Stanford, CA

Email: {haim1, pcuff, bvr, tsachy } at stanford.edu

Abstract—We consider a family of channels, collectively referred to as the ‘chemical channel’, which generalizes the trapdoor channel. We show that the feedback capacity of the chemical channel can be cast as the solution to a dynamic programming (DP) problem. We obtain numerical values for the feedback capacity of the chemical channel by approximating the solution of the DP problem using value iteration. For the special case of the trapdoor channel, by solving the DP problem analytically, we prove that the feedback capacity of the trapdoor channel is the logarithm of the golden ratio. Further, we describe a simple scheme that achieves the capacity of the trapdoor channel. The scheme has zero probability of error, which allows us to conclude that the logarithm of the golden ratio is also the zero error capacity of the chemical channel.

I. INTRODUCTION

The trapdoor channel, depicted in Figure 1, was introduced by Blackwell in 1961 [1] as a “simple two-state channel”. This channel is discussed in detail in the book by Ash [2], which features it on its cover.

The channel behaves as follows. Balls labeled ‘0’ or ‘1’ are used to communicate through the channel. The channel starts with a ball already in it. To use the channel, a ball is inserted into the channel by the transmitter, and the receiver receives one of the two balls in the channel with equal probability. The ball that does not exit the channel remains inside for the next channel use. Simple as it may be to describe, the capacity of the trapdoor channel remains an open problem after more than 45 years [1].




Fig. 1. The trapdoor(chemical) channel.

The *chemical channel* has the same structure as the trapdoor channel but, rather than being equal to $1/2$, the exit probability of a ball depends on whether it was the first of the two balls to enter the channel, and on its color (a precise description of this channel is given in the next section).

The name “chemical channel” suggests a physical system in which the concentrations of chemicals are used to communicate, such as might be the case in some cellular

biological systems as shown by Berger [3]. The transmitter adds molecules to the channel, and the receiver samples molecules randomly from the channel. The trapdoor channel is the most basic realization of this type of channel: it has only two types of molecules, all the molecules in the channel are equally likely to exit, and the concentration can only be one of $\{0, 0.5, 1\}$ or, equivalently, only one molecule remains in the channel between uses.

In this paper we study the capacity of the chemical channel with perfect feedback. After concretely describing the channel model and some preliminaries in Section II, we develop in Section III a dynamic programming (DP) framework to evaluate the capacity. We then use this framework in Section IV to obtain numerical values for the feedback capacity of the chemical channel by approximating the solution of the DP problem using value iteration. For the special case of the trapdoor channel, we solve the DP problem analytically in Section V, thereby proving that the feedback capacity of the trapdoor channel is the logarithm of the golden ratio. In Section VI we describe a simple scheme that achieves the capacity of the trapdoor channel. The scheme has zero probability of error, which we argue in Section VII to imply that the logarithm of the golden ratio is also the zero error capacity of the chemical channel. We conclude in Section VIII with a summary of our findings and a related future research direction.

A preliminary account of some of our results was given in [4], where we have shown that the feedback capacity of the trapdoor channel is the logarithm of the golden ratio. The novelty of the present paper is on several levels. For one thing, we extend the scope of consideration to the chemical channel, of which the trapdoor channel is a special case. No less importantly, unlike the scheme in [4], the capacity achieving scheme for the trapdoor channel that we describe here is error free (zero probability of decoding error), while neither encoder nor decoder need to know the initial state. This scheme allows us establish the logarithm of the golden ratio as the zero error capacity not only of the trapdoor channel, but of the chemical channel as well.

II. CHANNEL MODELS AND PRELIMINARIES

The chemical channel can be represented as a finite state channel (FSC). An FSC [5, ch. 4] is a channel that, for each

time index, has one of a finite number of possible states, s_{t-1} , and has the property that $p(y_t, s_t | x_t, s_{t-1}, y^{t-1}) = p(y_t, s_t | x_t, s_{t-1})$. The state of the chemical channel, which is shown in Figure 1, is the ball, 0 or 1, that is in the channel before the transmitter transmits a new ball. Let $x_t \in \{0, 1\}$ be the ball that is transmitted at time t and $s_{t-1} \in \{0, 1\}$ be the state of the channel when ball x_t is transmitted. The probability of the output y_t given the input x_t and the state of the channel s_{t-1} is shown in table I. For the case that $p_1 = p_2 = 0.5$ the chemical channel is the trapdoor channel as defined by Blackwell [1].

TABLE I

THE PROBABILITY OF THE OUTPUT y_t IN THE CHEMICAL CHANNEL GIVEN THE INPUT x_t AND THE STATE s_{t-1} .

x_t	s_{t-1}	$p(y_t = 0 x_t, s_{t-1})$	$p(y_t = 1 x_t, s_{t-1})$
0	0	1	0
0	1	p_1	$1 - p_1$
1	0	p_2	$1 - p_2$
1	1	0	1

The chemical channel has the property that the next state s_t is a deterministic function of the state s_{t-1} , the input x_t , and the output y_t . For a feasible tuple, (x_t, y_t, s_{t-1}) , the next state is given by the equation

$$s_t = s_{t-1} \oplus x_t \oplus y_t, \quad (1)$$

where \oplus denotes the binary XOR operation.

Because of this property, we consider only the family of FSCs known as unifilar channels, as considered by Ziv [6]. A *unifilar* FSC has the property that the state s_t is deterministic given (s_{t-1}, x_t, y_t) . We also define a *strongly connected* FSC, as follows.

Definition 1: A finite state channel is *strongly connected* if for any state s there exists an integer T and an input distribution of the form $\{p(x_t | s_{t-1})\}_{t=1}^T$ such that the probability that the channel reaches state s from any starting state s' , in less than T time-steps, is positive. I.e.

$$\sum_{t=1}^T \Pr(S_t = s | S_0 = s') > 0, \quad \forall s \in \mathcal{S}, \forall s' \in \mathcal{S}. \quad (2)$$

We assume a communication setting that includes feedback. The transmitter (encoder) knows at time t the message m and the feedback samples y^{t-1} . The output of the encoder at time t is denoted by x_t and is a function of the message and the feedback. The channel is a unifilar FSC and the output of the channel y_t enters the decoder (receiver). The encoder receives the feedback sample with one unit delay.

III. FEEDBACK CAPACITY AND DYNAMIC PROGRAMMING

Based on the upper and lower bounds on the feedback capacity of FSCs derived in [7], the following is proved in [8].

Theorem 1: The feedback capacity of a strongly connected unifilar FSC when initial state s_0 is known at the encoder and decoder can be expressed as

$$C_{FB} = \sup_{\{p(x_t | s_{t-1}, y^{t-1})\}_{t \geq 1}} \liminf_{N \rightarrow \infty} \frac{1}{N} \sum_{t=1}^N I(X_t, S_{t-1}; Y_t | Y^{t-1}) \quad (3)$$

where $\{p(x_t | s_{t-1}, y^{t-1})\}_{t \geq 1}$ denotes the set of all distributions such that $p(x_t | y^{t-1}, x^{t-1}, s^{t-1}) = p(x_t | s_{t-1}, y^{t-1})$ for $t = 1, 2, \dots$.

By using this theorem we can formulate the feedback capacity problem as an average-reward dynamic program.

A. Dynamic Programs

Here we introduce a formulation for average-reward dynamic programs. Each problem instance is defined by a septuple $(\mathcal{Z}, \mathcal{U}, \mathcal{W}, F, P_z, P_w, g)$. We will explain the roles of these parameters.

We consider a discrete-time dynamic system evolving according to $z_t = F(z_{t-1}, u_t, w_t)$, $t \geq 1$, where each *state* z_t takes values in a Borel space \mathcal{Z} , each *action* u_t takes values in a compact subset \mathcal{U} of a Borel space, and each *disturbance* w_t takes values in a measurable space \mathcal{W} . The initial state z_0 is drawn from a distribution P_z . Each disturbance w_t is drawn from a distribution $P_w(\cdot | z_{t-1}, u_t)$ which depends only on the state z_{t-1} and action u_t . All functions considered in this paper are assumed to be measurable, though we will not mention this each time we introduce a function or set of functions.

The *history* $h_t = (z_0, w_0, \dots, w_{t-1})$ summarizes information available prior to selection of the t th action. The action u_t is selected by a function μ_t which maps histories to actions. In particular, given a policy $\pi = \{\mu_1, \mu_2, \dots\}$, actions are generated according to $u_t = \mu_t(h_t)$.

We consider an objective of maximizing average reward, given a bounded reward function $g : \mathcal{Z} \times \mathcal{U} \rightarrow \mathbb{R}$. The average reward for a policy π is defined by

$$\rho_\pi = \liminf_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}_\pi \left\{ \sum_{t=0}^{N-1} g(Z_t, \mu_{t+1}(h_{t+1})) \right\},$$

where the subscript π indicates that actions are generated by the policy $\pi = (\mu_1, \mu_2, \dots)$. The optimal average reward is defined by $\rho^* = \sup_\pi \rho_\pi$.

B. The Bellman Equation

An alternative characterization of the optimal average reward is offered by the Bellman Equation. This equation offers a mechanism for verifying that a given level of average reward is optimal. The following encapsulation of the Bellman equation, which can be found in [9], will be of later use.

Theorem 2: If $\rho \in \mathbb{R}$ and a bounded function $h : \mathcal{Z} \mapsto \mathbb{R}$ satisfy for all $z \in \mathcal{Z}$

$$\rho + h(z) = \sup_{u \in \mathcal{U}} \left(g(z, u) + \int P_w(dw | z, u) h(F(z, u, w)) \right) \quad (4)$$

then $\rho = \rho^*$. Further, if there is a function $\mu : \mathcal{Z} \mapsto \mathcal{U}$ such that $\mu(z)$ attains the supremum for each z then $\rho_\pi = \rho^*$ for $\pi = (\mu_0, \mu_1, \dots)$ with $\mu_t(h_t) = \mu(z_{t-1})$ for each t .

It is convenient to define a dynamic programming operator T by

$$(Th)(z) = \sup_{u \in \mathcal{U}} \left(g(z, u) + \int P_w(dw|z, u) h(F(z, u, w)) \right),$$

for all functions h . Then, Bellman's equation can be written as $\rho \mathbf{1} + h = Th$.

C. Feedback Capacity as a Dynamic Program

We will now formulate a dynamic program such that the optimal average reward equals the feedback capacity of a unifilar channel as presented in Theorem 1. This entails defining the septuple $(\mathcal{Z}, \mathcal{U}, \mathcal{W}, F, P_z, P_w, g)$ based on properties of the unifilar channel and then verifying that the optimal average reward is equal to the capacity of the channel.

Let β_t denote the $|\mathcal{S}|$ -dimensional vector of channel state probabilities given information available to the decoder at time t . In particular, each component corresponds to a channel state s_t and is given by $\beta_t(s_t) \triangleq p(s_t|y^t)$. We take states of the dynamic program to be $z_t = \beta_t$. Hence, the state space \mathcal{Z} is the $|\mathcal{S}|$ -dimensional unit simplex. Each action u_t is taken to be the matrix of conditional probabilities of the input x_t given the previous state s_{t-1} of the channel. Hence, the action space \mathcal{U} is the set of stochastic matrices of dimension $|\mathcal{S}| \times |\mathcal{X}|$. The disturbance w_t is taken to be the channel output y_t . The disturbance space \mathcal{W} is the output alphabet Y .

The initial state distribution P_z is concentrated at the prior distribution of the initial channel state s_0 . Note that the channel state s_t is conditionally independent of the past given the previous channel state s_{t-1} , the input probabilities u_t , and the current output y_t . Hence, $\beta_t(s_t) = p(s_t|y^t) = p(s_t|\beta_{t-1}, u_t, y_t)$.

The distribution of the disturbance w_t is $p(w_t|z^{t-1}, w^{t-1}, u^t) = p(w_t|z_{t-1}, u_t)$. Conditional independence from z^{t-2} and w^{t-1} given z_{t-1} is due to the fact that the channel output is determined by the previous channel state and current input.

We consider a reward of $I(Y_t; X_t, S_{t-1}|y^{t-1})$. Note that the reward depends only on the probabilities $p(x_t, y_t, s_{t-1}|y^{t-1})$ for all x_t, y_t and s_{t-1} . Further,

$$p(x_t, y_t, s_{t-1}|y^{t-1}) = \beta_{t-1}(s_{t-1}) u_t(s_{t-1}, x_t) p(y_t|x_t, s_{t-1})$$

Recall that $p(y_t|x_t, s_{t-1})$ is given by the channel model. Hence, the reward depends only on β_{t-1} and u_t .

Given an initial state z_0 and a policy $\pi = (\mu_1, \mu_2, \dots)$, u_t and β_t are determined by y^{t-1} . Further, (X_t, S_{t-1}, Y_t) is conditionally independent of y^{t-1} given β_{t-1} as shown in (5). Hence,

$$g(z_{t-1}, u_t) = I(Y_t; X_t, S_{t-1}|y^{t-1}) = I(X_t, S_{t-1}; Y_t|\beta_{t-1}, u_t). \quad (5)$$

It follows that the optimal average reward is

$$\rho^* = \sup_{\pi} \liminf_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}_{\pi} \left[\sum_{t=1}^N I(X_t, S_{t-1}; Y_t|Y^{t-1}) \right] = C_{FB}.$$

The dynamic programming formulation that is presented here is an extension of the formulation presented in [10] by Yang, Kavčić and Tatikonda. In [10] the formulation is for channels with the property that the state is deterministically determined by the previous inputs, whereas here we allow the state to be determined, in addition to the previous inputs, also by the previous outputs.

IV. COMPUTING THE FEEDBACK CAPACITY

In this section we briefly describe some results of computing the feedback capacity of the chemical channel by solving the dynamic program using value iteration. This computation generates a sequence of iterates according to

$$J_{k+1} = TJ_k, \quad (6)$$

initialized with $J_0 = 0$. For each k and z , $J_k(z)$ is the maximal expected reward over k time periods given that the system starts in state z . Since rewards are positive, for each z , $J_k(z)$ grows with k . For each k , we define a differential reward function $h_k(z) \triangleq J_k(z) - J_k(0)$. These functions capture differences among values $J_k(z)$ for different states x .

Figure 2 shows the capacity of the chemical channel as a function of two parameters. The left plot depicts the capacity as a function of the parameter p_{zero} , which is the probability that the ball of type '0' will exit if in the channel the other ball is of type '1'. The parameter p_{zero} implies that the probabilities p_1, p_2 in Table I both equal p_{zero} .




Fig. 2. Capacity of the chemical channel as function of two parameters.

The right plot depicts the capacity as a function of the parameter p_{switch} , which is the probability that the ball that just entered will exit. The parameter p_{switch} implies that the probabilities p_1, p_2 satisfy $p_1 = 1 - p_2 = p_{switch}$. If we interpret the chemical channel as a permutation channel [11], as seen in Figure 3, then p_{switch} is the probability of switching two adjacent balls.




Fig. 3. The chemical channel as a permuting channel. Going from left to right, there is a probability p_{switch} that two adjacent bits switch places.

V. FEEDBACK CAPACITY OF THE TRAPDOOR CHANNEL

For the case where $p_{zero} = p_{switch} = 0.5$, which is the trapdoor channel, the results in the previous section, obtained from value iteration, gave an approximate value of 0.694, which is suspiciously close to $\log \phi \approx 0.6942$, where $\phi = \frac{1+\sqrt{5}}{2}$ is the golden ratio [12]. One context in which the golden ratio comes up is the entropy rate of the binary Markov chain with transition matrix

$$P = \begin{bmatrix} 1-p & p \\ 1 & 0 \end{bmatrix}. \quad (7)$$

The entropy rate of this chain is given by $\frac{H(p)}{1+p}$. It is readily checked that $\max_{0 \leq p \leq 1} \frac{H(p)}{1+p} = \log \phi$, the maximizing value being $p = \frac{3-\sqrt{5}}{2}$.

We will now show that $\log \phi$ is indeed the capacity of the trapdoor channel by analytically solving the dynamic program. Consider the following policy associated with the Markov chain in (7): The state of the Markov process indicates whether the input to the channel will be the same or different from the state of the channel. In other words, if at time t the binary Markov sequence is '0' then the input to the channel is equal to the state of the channel, i.e. $x_t = s_{t-1}$. Otherwise, the input to the channel will be the complement to the state of the channel, i.e. $x_t = s_{t-1} \oplus 1$. This scheme uniquely defines the distribution $p(x_t | s_{t-1}, y^{t-1})$:

$$p(X_t = s_{t-1} | s_{t-1}, y_{t-1}) = \begin{cases} 1-p & \text{if } s_{t-1} = y_{t-1}, \\ 1 & \text{if } s_{t-1} \neq y_{t-1}. \end{cases} \quad (8)$$

This policy led us to the construction of the function $\tilde{h}(z)$ for $z \in [\sqrt{5}-2, 3-\sqrt{5}]$ described as follows:

$$\begin{aligned} \tilde{h}(z) &= H(z) - \tilde{\rho}z + c_2, & \sqrt{5}-2 \leq z \leq \frac{3-\sqrt{5}}{2} \\ \tilde{h}(z) &= 1, & \frac{3-\sqrt{5}}{2} \leq z \leq \frac{\sqrt{5}-1}{2} \\ \tilde{h}(z) &= H(z) + \tilde{\rho}z + c_1, & \frac{\sqrt{5}-1}{2} \leq z \leq 3-\sqrt{5} \end{aligned} \quad (9)$$

where $c_1 = \log(3-\sqrt{5})$, $c_2 = \log(\sqrt{5}-1)$ and $\tilde{\rho} = \log \frac{1+\sqrt{5}}{2}$. The function \tilde{h} is continuous and symmetric around $z = 0.5$.

A. Verification

According to Theorem 2, to verify that the trapdoor feedback capacity is indeed $\log \phi$, it will suffice to establish the existence of $h(z)$ that solves the Bellman equation (4) for the trapdoor channel with $\rho^* = \log \phi$. The main obstacle for making this verification is the fact that $\tilde{h}(z)$ exists only for $z \in [\sqrt{5}-2, 3-\sqrt{5}]$, and in order to verify Bellman equation we need a candidate function defined for the whole range $z \in [0, 1]$.

We solve the problem by constructing a sequence of functions $h_k(z)$, $k = 0, 1, 2, \dots$ that has the property that $h_k(z) = \tilde{h}(z)$ for all $k \geq 0$, $z \in [\sqrt{5}-2, 3-\sqrt{5}]$, and that converges uniformly to a function $h^*(z)$ that solves the Bellman equation. Let $h_0(z)$ be the pointwise maximum among concave

functions satisfying $h_0(z) = \tilde{h}(z)$ for $z \in [\sqrt{5}-2, 3-\sqrt{5}]$, where $\tilde{h}(z)$ is defined in (9). Let $h_{k+1}(z) = (Th_k)(z) - \tilde{\rho}$, and

$$h^*(z) \triangleq \limsup_{k \rightarrow \infty} h_k(z). \quad (10)$$

It can be shown that the limit supremum is in fact a limit and that the convergence is uniform. These facts, along with the definition of \tilde{h} , allow to verify the following.

Theorem 3: The function h^* and scalar $\tilde{\rho}$ satisfy $\tilde{\rho}1 + h^* = Th^*$, where $\tilde{\rho} = \log \phi$.

By Theorem 2, this proves that $\log \phi$ is the trapdoor channel feedback capacity.

VI. ERROR-FREE CAPACITY-ACHIEVING SCHEME

We begin with a description of a simple encoder and decoder pair that provides error-free communication through the trapdoor channel with feedback and *known initial state*, at any rate below capacity. We will then enhance it to a zero error scheme where neither encoder nor decoder need to know the initial state.

A. Encoding

Each message is mapped to a unique binary sequence of N actions, \tilde{x}_N , that ends with 0 and has no occurrences of two 1's in a row. The input to the channel is derived from the action and the state as $x_k = \tilde{x}_k \oplus s_{k-1}$.

B. Decoding

The channel outputs are recorded differentially as, $\tilde{y}_k = y_k \oplus y_{k-1}$, for $k = 2, \dots, N$. Decoding of the action sequence is accomplished in reverse order, beginning with $\tilde{x}_N = 0$ by construction. It can be verified, by exhaustively considering all the possibilities, that if \tilde{x}_{k+1} is known to the decoder, \tilde{x}_k can be correctly decoded. Since \tilde{x}_N is known by the decoder (to be 0), this implies that, going in reverse, the decoder obtains the whole channel input sequence error free.

Decoding example. Table II shows an example of decoding a sequence of length $N = 4$.

TABLE II
DECODING EXAMPLE

Variable	Value	Reason
y_n	1110	Channel output
\tilde{y}_n	*001	Differential output
\tilde{x}_n	0	Given
	10	$\tilde{x}_4 = 0$ and $\tilde{y}_4 = 1$
	010	$\tilde{x}_3 \neq 0$ (or $\tilde{y}_3 \neq 1$)
	0010	$\tilde{y}_2 \neq 1$

C. Rate

Evidently, the number of sequences between which this scheme can distinguish error free is the number of binary sequences of length $N-1$ that have no repeating 1's. Exponentially, this is readily seen to be equivalent to ϕ^N , where ϕ is the golden ratio. In other words, rates arbitrarily close to $\log \phi$ are achievable without error.

D. Unknown Initial State

We now modify the scheme to establish error-free communication that still achieves capacity without knowledge of the initial state.

1) *Encoding*: Encoding is done in three phases. Phase 1 uses the channel N times, phase 2 uses the channel three times, and phase 3 uses the channel an additional (deterministic) number of times $\theta(N)$, where $\theta(N)$ is derived from the Fibonacci sequence and is roughly $\frac{1}{\log \phi} \log N$ (i.e. $\lim_{N \rightarrow \infty} \frac{\theta(N)}{\frac{1}{\log \phi} \log N} = 1$), where ϕ is the golden ratio. Each message is mapped to a unique action sequence, \tilde{x}_k , that will be used in phase 1.

Phase 1. The encoder assumes the initial state of the channel is 0 and communicates according to the scheme in section VI-A. At some point during phase 1 it is possible that the encoder will discover that the initial state assumption was wrong. This will happen if a ball labeled ‘1’ leaves the channel when according to the encoders calculation there were two balls labeled ‘0’ in the channel. We will call this event a ‘contradiction.’ If a contradiction occurs then the correct state is known to the encoder from that point on and the encoding continues based on the correct state.

Phase 2. If a contradiction did not occur in phase 1, then the encoder sends three 0’s. Otherwise, the encoder sends three 1’s.

Phase 3. If a contradiction did not occur in phase 1, then the encoder sends all zeros in phase 3. If a contradiction did occur in phase 1, then the encoder knows the initial state for phase 3, so it uses the scheme from Section VI-A to communicate the index of the contradiction. The indices 1 through N are assigned to unique admissible action sequences.

2) *Decoding*: Decoding begins with phase 2 and finishes with phase 1, where the message is discovered. In the case of no contradiction, phase 3 is skipped and phase 1 is decoded according to section VI-B.

Phase 2. The decoder notes that a contradiction occurred (three 1’s were sent) if and only if two or more 1’s were received at the output in phase 2.

Phase 3. If a contradiction occurred, then the index of the contradiction is decoded as explained in section VI-B.

Phase 1. It is not hard to show that if a contradiction occurred at index m , then \tilde{x}_k can be correctly decoded by first changing y_m to its complement and then decoding according to Section VI-B.

3) *Rate*: The number of unique input sequences distinguishable with zero-error is as in the original scheme. The number of channel uses is now $N + 3 + \frac{1}{\log \phi} \log N$, which is still $N + o(N)$ so this scheme too asymptotically achieves the rate $\log \phi$.

VII. ZERO-ERROR CAPACITY OF THE CHEMICAL CHANNEL

The zero error feedback capacity is the same for all chemical channels with parameters $0 < p_1 < 1$ and $0 < p_2 < 1$, since they all share the same set of channel realizations that

have positive probability. Denote this capacity by C_{ZE} . Note that on the one hand $C_{ZE} \leq \log \phi$ since $\log \phi$ was seen in Section V to be the Shannon (as opposed to zero-error) feedback capacity of the trapdoor channel which cannot be smaller than its zero error capacity. On the other hand, the scheme of the previous Section achieves rates arbitrarily close to $\log \phi$ with zero probability of error, implying $C_{ZE} \geq \log \phi$. We conclude then that $C_{ZE} = \log \phi$.

VIII. CONCLUSION AND FURTHER WORK

The feedback capacity of the chemical channel was shown to be expressible as an average-reward dynamic program, which allowed us to evaluate it by using a value iteration algorithm. For the trapdoor channel, which is a special case of the chemical channel, by analytically solving the dynamic programming problem we showed that the feedback capacity is the logarithm of the golden ratio. A simple scheme that achieves this capacity with zero decoding error probability was described, and allowed us to deduce that the logarithm of the golden ratio is also the zero-error feedback capacity of the chemical channel.

Perhaps most imminent among the several directions in which this work can be extended is that of finding an analytic solution for the Shannon feedback capacity of the chemical channel at values of (p_1, p_2) other than $(1/2, 1/2)$.

ACKNOWLEDGMENT

We are grateful to Tom Cover for introducing us to the chemical channel, and for very helpful discussions. This work was supported by the National Science Foundation through the grants CCR-0311633, CCF-0515303, IIS-0428868 and CAREER.

REFERENCES

- [1] D. Blackwell, “Information theory,” *Modern mathematics for the engineer: Second series*, pp. 183–193, 1961.
- [2] R. Ash, *Information Theory*. New York: Wiley, 1965.
- [3] T. Berger, “The generalized shannon-blackwell billiard ball channel,” in *Lecture 2 of CSL Distinguished Visiting Professorship, University of Illinois, Urbana, IL, 22 April 2002*. (“Information Theory Invades Biology”, 22 April 2002.
- [4] H. H. Permuter, P. W. Cuff, B. Van-Roy, and T. Weissman, “Capacity of the trapdoor channel with feedback,” in *Allerton*, September 2006.
- [5] R. C. Gallager, *Information Theory and Reliable Communication*. New York: Wiley, 1968.
- [6] J. Ziv, “Universal decoding for finite-state channels,” *IEEE Trans. Inform. Theory*, vol. 31, no. 4, pp. 453–460, 1985.
- [7] H. H. Permuter, T. Weissman, and A. J. Goldsmith, “Finite state channels with time-invariant deterministic feedback,” Sep 2006, submitted to IEEE Trans. Inform. Theory. Available at arxiv.org/pdf/cs.IT/0608070.
- [8] H. H. Permuter, P. W. Cuff, B. Van-Roy, and T. Weissman, “Capacity of the trapdoor channel with feedback,” Sep 2006, submitted to IEEE Trans. Inform. Theory. Available at arxiv.org/pdf/cs.IT/0610047.
- [9] A. Arapostathis, V. S. Borkar, E. Fernandez-Gaucherand, M. K. Ghosh, and S. Marcus, “Discrete time controlled Markov processes with average cost criterion - a survey,” *SIAM Journal of Control and Optimization*, vol. 31, no. 2, pp. 282–344, 1993.
- [10] A. K. S. Yang and S. Tatikonda, “Feedback capacity of finite-state machine channels,” *IEEE Trans. Inform. Theory*, pp. 799–810, 2005.
- [11] T. W. Benjamin, “Coding for a noisy channel with permutation errors,” Ph.D. Dissertation, Cornell Univ., Ithaca, NY, 1975.
- [12] L. Mario, *The golden ratio : the story of phi, the world’s most astonishing number*. New York: Broadway Books, 2002.