On the Storage Capacity of Rewritable Memories

Christoph Bunte and Amos Lapidoth
Signal and Information Processing Laboratory
ETH Zurich

***Abstract*—We study the capacity of rewritable storage cells with discrete memoryless write channels. We consider a constraint on the maximum number of rewrites and give a closed-form expression for the capacity of cells with symmetric write channels. For binary symmetric write channels this result is extended to the case where only noisy observations of the state are available to the encoder and decoder. We also consider a constraint on the expected number of rewrites and derive a lower bound on the capacity. This bound coincides with the capacity for cells with binary symmetric write channels.**

I. Introduction

A rewritable memory consists of a number of cells, each of which carries information in the physical state it assumes (e.g., an electrical resistance). Thus, storing information amounts to setting the states of the cells to a desired configuration, and retrieving information amounts to observing the states. The former is commonly referred to as "writing" to memory, and the latter as "reading" from memory. Both of these processes may be subjected to noise. In other words, the true state of a cell may differ from the state intended by the person storing information, and it may also differ from the state observed by the person wishing to retrieve the information.

The capacity of a rewritable memory, i.e., the amount of information that can be stored and retrieved reliably, depends on the statistical properties of the write and read mechanisms. An obvious way to compensate for uncertainties in the write mechanism is to observe the state of a cell after each write operation and rewrite it in case the outcome is undesired. This effectively increases the reliability of the write mechanism and hence the capacity. The increase in capacity, however, comes at a cost (e.g., higher energy consumption, longer encoding delays, et cetera). Thus, it seems reasonable to impose certain restrictions on the number of rewrite operations performed on each cell.

In [1]–[6], Mittelholzer et al. studied the capacity of cells with continuous-alphabet write channels subject to constraints on the number of rewrites. Assuming an ideal (i.e., noiseless) read mechanism, reasonably well-behaved write noise (e.g., additive Gaussian or uniform noise), one expects the storage capacity to be an unbounded function of the allowed (average) number of rewrites. In fact, the authors of [1]–[6] show that, for a certain class of cells, the capacity grows logarithmically in the expected number of rewrite iterations.

In this paper we consider cells with *discrete* write channels, which implies that the number of different states a cell can assume is finite. In this case the capacity is always upper bounded by the logarithm of the cardinality of the state-alphabet. What follows is an attempt to study the capacity of such cells subject to constraints on the maximum and the expected number of rewrites.

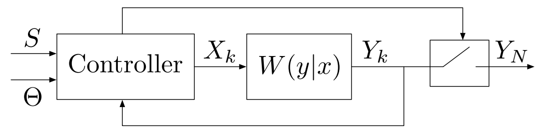

Fig. 1. A cell with controller.

II. The Model

In this section we give basic definitions and establish the mathematical model of a cell.

Definition 1. *Let μ and ν be positive integers. A cell is a discrete memoryless channel (DMC) with input alphabet $\mathcal{X} = \{0, \ldots, \mu - 1\}$, output alphabet $\mathcal{Y} = \{0, \ldots, \nu - 1\}$, and transition law*

$$W(y|x) = \mathbb{P}[Y = y|X = x], \quad x \in \mathcal{X}, y \in \mathcal{Y}.$$

The channel input X corresponds to the write stimulus, and the channel output Y corresponds to the state of the cell induced by X. Thus, every use of the channel corresponds to a write operation, and the transition law W models the uncertainty in the write mechanism. We refer to $\mathcal{Y}$ as the state-alphabet of the cell. In light of Definition 1 we can identify a cell with its transition law W.

Definition 2. *A* strategy *is a sequence $\{(\sigma_k, \mathcal{T}_k)\}$ of pairs of mappings*[1]

$$\sigma_k \colon \mathcal{Y}^{k-1} \times \mathcal{U} \to \mathcal{X},$$
$$\mathcal{T}_k \colon \mathcal{Y}^{k-1} \times \mathcal{U} \to 2^{\mathcal{Y}}, \quad k = 1, 2, \ldots$$

The set of all strategies will be denoted by $\mathcal{S}$.

Let Θ be a random variable taking value in a set $\mathcal{U}$. The k-th channel output (state of the cell) Y_k is induced by the k-th channel input (write stimulus) $X_k = \sigma_k(Y^{k-1}, \Theta)$. We define the k-th target set as $T_k = \mathcal{T}_k(Y^{k-1}, \Theta)$. Let N be the smallest positive integer such that $Y_N \in T_N$. Then N is a stopping time and we call Y_N the *final* state of the cell.

[1]For notational convenience we define $Y^k = (Y_1, \ldots, Y_k)$ for $k \geq 1$. We also define $Y^0 = 0$, $Y_0 = 0$, and $\mathcal{Y}^0 = \{0\}$. The power set of $\mathcal{Y}$, i.e., the collection of all subsets of $\mathcal{Y}$, is denoted by $2^{\mathcal{Y}}$.

978-1-4244-8682-3/10/$26.00

The purpose of introducing Θ is to allow strategies that make nondeterministic decisions.

One can think of the cell as being connected to a controller, which, at time k, has access to $Y_1, \ldots, Y_{k-1}$ and Θ. The controller is fed the strategy and produces the channel input $X_k = \sigma_k(Y^{k-1}, \Theta)$ accordingly. It stops after the k-th write iteration if, and only if, $Y_k \in T_k$. Hence N equals the number of writes performed on the cell (and $N-1$ equals the number of rewrites). Let S denote the strategy. We then have a memoryless channel connecting $S \in \mathcal{S}$ with $Y_N \in \mathcal{Y}$ as illustrated in Figure 1. This channel is specified by the conditional probabilities

$$\mathbb{P}\left[Y_N = y | S = s\right], \quad s \in \mathcal{S}, y \in \mathcal{Y}.$$

Since different strategies may induce the same conditional distribution of Y_N, it is convenient to put them into equivalence classes. The equivalence class $[s]$ of $s \in \mathcal{S}$ is the set

$$\{s' \in \mathcal{S} : \mathbb{P}\left[Y_N = y | S = s'\right] = \mathbb{P}\left[Y_N = y | S = s\right], \forall y \in \mathcal{Y}\}.$$

Then $[\mathcal{S}] = \{[s] : s \in \mathcal{S}\}$ is the collection of all equivalence classes. Instead of considering $\mathcal{S}$ it suffices to consider a subset $\tilde{\mathcal{S}} \subset \mathcal{S}$ that contains exactly one representative of every[2] element of $[\mathcal{S}]$. Then $\tilde{\mathcal{S}}$ has the same cardinality as $[\mathcal{S}]$. If a certain cost is associated with each strategy $s \in \mathcal{S}$, as is done in subsequent sections, the representative of each equivalence class should be chosen to have the lowest cost among all members of its class.

Consider an array of n identical and statistically independent cells, and let N_i denote the stopping time (number of writes) for the i-th cell.

Definition 3. *A rate-R blocklength-n code consists of a message set $\mathcal{M} = \{1, \ldots, 2^{nR}\}$, an encoder mapping $f\colon \mathcal{M} \to \mathcal{S}^n$, and a decoder mapping $g\colon \mathcal{Y}^n \to \mathcal{M}$.*

To store the message $m \in \mathcal{M}$, the encoder produces a sequence of strategies $S^n = f(m)$; one for every cell. The decoder observes the final state of each cell to obtain the sequence Y_N^n from which it produces the estimate $\hat{M} = g(Y_N^n)$. Associated with every code is the maximum probability of error

$$P_e = \max_{m \in \mathcal{M}} \mathbb{P}\left[\hat{M} \neq M | M = m\right].$$

Definition 4. *A rate R (in bits per cell) is said to be achievable (subject to a constraint), if there exists a sequence $\{(f_n, g_n)\}$ of rate-R blocklength-n codes (satisfying the constraint) such that $P_e \to 0$ as $n \to \infty$.*

III. A Constraint on the Maximum Number of Rewrites

We begin our discussion of capacity by considering a constraint on the maximum number of rewrites performed on each cell.

[2]The existence of such a set is guaranteed by the Axiom of Choice.

Definition 5. *For every positive integer η we define the capacity C_η as the supremum of the set of rates achievable subject to the constraint $\mathbb{P}\left[N_i \leq \eta\right] = 1$ for $i = 1, \ldots, n$.*

The set of all strategies permissible under the given constraint is $\mathcal{S}_\eta = \{s \in \mathcal{S} : \mathbb{P}\left[N \leq \eta | S = s\right] = 1\}$. Thus,

$$C_\eta = \max_{S \in \mathcal{S}_\eta} I(S; Y_N), \quad \eta = 1, 2, \ldots,$$

where the maximum is taken over all distributions of S on $\mathcal{S}_\eta$. Since the channel defining a cell is memoryless, one can show that it suffices to consider the set $\mathcal{S}'_\eta \subset \mathcal{S}_\eta$ of all strategies $s \in \mathcal{S}_\eta$ such that $\sigma_k(Y^{k-1}, \Theta) = \sigma'_k(Y_{k-1})$ and $\mathcal{T}_k(Y^{k-1}, \Theta) = \mathcal{T}'_k(Y_{k-1})$ for suitable mappings σ'_k and $\mathcal{T}'_k$. That is, Θ may be chosen to be deterministic, and only the most recent state is relevant for determining the input and target set. Note that $[\mathcal{S}'_\eta]$ is a finite set. Using convexity arguments one can prove the following lemma.

Lemma 1. *Let $\mathcal{S}'_\eta$ and $\mathcal{S}_\eta$ be as above. Then*

$$\max_{S \in \mathcal{S}'_\eta} I(S; Y_N) = \max_{S \in \mathcal{S}_\eta} I(S; Y_N), \quad \eta = 1, 2, \ldots,$$

where the maximization on the left is over all distributions of S on $\mathcal{S}'_\eta$, and the maximization on the right is over all distributions of S on $\mathcal{S}_\eta$.

A. Some General Observations

Let C_{DMC} denote the channel capacity of the DMC defining a cell W. The following proposition summarizes some basic properties of the sequence C_η.

Proposition 1.

1) *$C_\eta = C_{\text{DMC}}$ for $\eta = 1$;*
2) *C_η is monotonically nondecreasing;*
3) *if for every $y \in \mathcal{Y}$ there exists some $x \in \mathcal{X}$ such that $W(y|x) > 0$, then $\lim_{\eta \to \infty} C_\eta = \log \nu$.*

The next proposition shows that a single rewrite suffices to obtain positive capacity for all "nondegenerate" cells.

Proposition 2. *The following two conditions are equivalent.*

1) *If $\eta \geq 2$, then $C_\eta > 0$.*
2) *There exist two distinct symbols $y_1 \in \mathcal{Y}$ and $y_2 \in \mathcal{Y}$ such that $\max_{x \in \mathcal{X}} W(y_1|x) > 0$ and $\max_{x \in \mathcal{X}} W(y_2|x) > 0$.*

B. Binary State-Alphabets

For $\mathcal{Y} = \{0, 1\}$ it can be shown that, without loss of generality, one can restrict the set of strategies to $\{s_0, s_1\} \subset \mathcal{S}'_\eta$, where s_0 and s_1 are defined by

$$X_k(s_i) = \min\{j \in \mathcal{X} : W(i|j) \geq W(i|j') \text{ for all } j' \in \mathcal{X}\},$$

$$T_k(s_i) = \begin{cases} \{i\}, & k < \eta, \\ \{0, 1\}, & k \geq \eta, \end{cases}$$

for $i \in \{0, 1\}$ and $k = 1, 2, \ldots$ The capacity can then be computed by maximizing $I(S; Y_N)$ with respect to the parameter $p = \mathbb{P}\left[S = s_0\right] = 1 - \mathbb{P}\left[S = s_1\right]$.

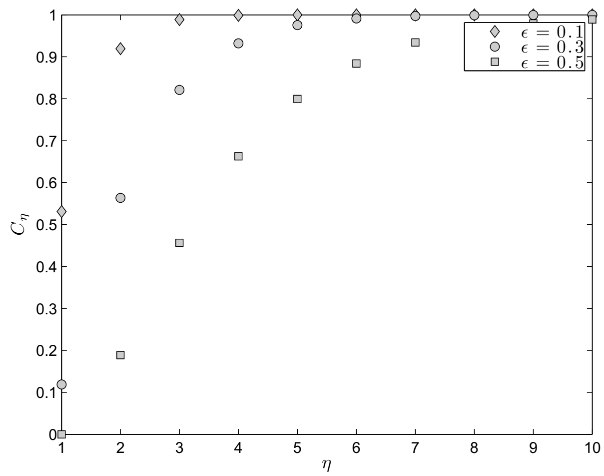

Fig. 2. Capacity of the BSC cell as a function of η for different values of ϵ.

C. The Binary Symmetric Channel Cell

Consider the cell with $\mathcal{X} = \mathcal{Y} = \{0, 1\}$ and transition law

$$W(y|x) = \begin{cases} 1-\epsilon, & x = y, \\ \epsilon, & x \neq y, \end{cases} \quad x \in \mathcal{X}, y \in \mathcal{Y},$$

for some $\epsilon \in [0, 1/2]$. For obvious reasons we refer to it as the binary symmetric channel (BSC) cell of parameter ϵ. Based on the observation in Section III-B it is not difficult to prove the following proposition.

Proposition 3. *For the BSC cell of parameter ϵ,*

$$C_\eta = 1 - H_b(\epsilon^\eta), \quad \eta = 1, 2, \ldots$$

Moreover, capacity is achieved by a uniform distribution of S on the two strategies s_0 and s_1 defined in Section III-B.

D. Symmetric Cells

The following definition is based on the definition of symmetry for DMCs as given in [7].

Definition 6. *We say that a cell is symmetric if W is the transition law of a symmetric DMC.*

Thus, if W is a symmetric cell and $\mathrm{W} \in \mathbb{R}^{\nu \times \mu}$ is the corresponding stochastic matrix with entries[3] $\mathrm{W}_{i,j} = W(i-1|j-1)$ for $i = 1, \ldots, \nu$ and $j = 1, \ldots, \mu$, then the rows of W are permutations of each other, and the columns of W are permutations of each other. Using the notion of Schur convexity (see [8]) one can prove the following theorem.

Theorem 1. *Let* W *be the stochastic matrix corresponding to a symmetric cell. Let* $\mathbf{w}$ *be a column of* W *and let* $k_0 \in \{1, \ldots, \nu\}$ *be an index of a largest entry of* $\mathbf{w}$. *Define the k-th column* $\mathbf{a}^{(k)}$ *of the matrix* $\mathrm{A} \in \mathbb{R}^{\nu \times \nu}$ *as*

$$\mathbf{a}^{(k)} = \begin{cases} \mathbf{w}, & k \neq k_0, \\ \mathbf{I}^{(k_0)}, & k = k_0, \end{cases} \quad k = 1, \ldots, \nu,$$

where $\mathbf{I}^{(k_0)}$ *denotes the* k_0*-th column of the* $\nu \times \nu$ *identity matrix. Then*

$$C_\eta = \log \nu - H\left(\mathrm{A}^{\eta-1}\mathbf{w}\right), \quad \eta = 1, 2, \ldots$$

[3] $\mathrm{W}_{i,j}$ denotes the entry in the i-th row and j-th column of W.

E. The BSC Cell with Impaired Read Mechanisms

We now extend the capacity formula for the BSC cell (of parameter ϵ) to the case where the read and feedback mechanisms are corrupted by independent BSCs. More precisely, at every epoch k the encoder can only observe the output of a BSC with crossover probability $\delta \in [0, 1/2]$ whose input is Y_k. Likewise, the decoder can observe the final state Y_N only through a BSC with crossover probability $\gamma \in [0, 1/2]$. The necessary modifications to the definitions of "strategy" and "code" are obvious.

Theorem 2. *Let ϵ, δ, and γ have the above meaning. Then*

$$C_\eta = 1 - H\left(\mathrm{B}\mathrm{E}^{\eta-1}\mathbf{p}\right), \quad \eta = 1, 2, \ldots,$$

where

$$\mathrm{B} = \begin{pmatrix} 1-\gamma & \gamma \\ \gamma & 1-\gamma \end{pmatrix}, \quad \mathrm{E} = \begin{pmatrix} 1-\epsilon\delta & (1-\delta)(1-\epsilon) \\ \epsilon\delta & \epsilon(1-\delta)+\delta \end{pmatrix},$$

and $\mathbf{p} = (1-\epsilon, \epsilon)^T$. *Moreover,*

$$\lim_{\eta \to \infty} C_\eta = 1 - H_b\left(\frac{(1-\epsilon)(1-\delta)(1-\gamma) + \epsilon\delta\gamma}{(1-\epsilon)(1-\delta) + \epsilon\delta}\right).$$

IV. A Constraint on the Average Number of Rewrites

In this section we discuss the capacity of cells subject to a constraint on the expected number of rewrites.

Definition 7. *For every $\zeta \geq 1$ we define the capacity $C(\zeta)$ as the supremum of the set of rates achievable subject to the constraint $\sum_{i=1}^{n} E[N_i] \leq n\zeta$.*

Thus,

$$C(\zeta) = \max_{S: E[N] \leq \zeta} I(S; Y_N), \quad \zeta \geq 1,$$

where the maximum is taken over all distributions of S on $\mathcal{S}$ such that $E[N] \leq \zeta$. It is clear that $C(1) = C_{\mathrm{DMC}}$. Furthermore, assuming that for every $y \in \mathcal{Y}$ there exists some $x \in \mathcal{X}$ such that $W(y|x) > 0$,

$$C(\zeta) = \log \nu, \quad \text{if } \zeta \geq \beta \triangleq \frac{1}{\nu} \sum_{i=0}^{\nu-1} \frac{1}{\max_{j \in \mathcal{X}} W(i|j)}.$$

For $\zeta \geq \beta$ capacity is achieved by a uniform distribution of S on $\{s_0, \ldots, s_{\nu-1}\} \subset \mathcal{S}$, where s_i is defined by

$$X_k(s_i) = \min\{j \in \mathcal{X} : W(i|j) \geq W(i|j') \text{ for all } j' \in \mathcal{X}\},$$
$$T_k(s_i) = \{i\}, \quad k = 1, 2, \ldots,$$

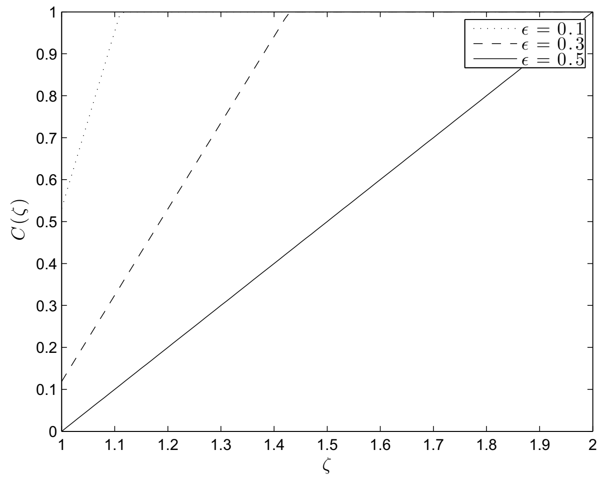

Fig. 3. Capacity of the BSC cell as a function of ζ for different values of ϵ.

for $i = 0, \ldots, \nu - 1$. Since $C(\zeta)$ is a concave function of ζ, we immediately get the lower bound:

Proposition 4.

$$C(\zeta) \geq \left(1 - \frac{\zeta - 1}{\beta - 1}\right) C_{\text{DMC}} + \frac{\zeta - 1}{\beta - 1} \log \nu, \quad 1 \leq \zeta \leq \beta.$$

For the BSC cell, this lower bound coincides with the capacity:

Proposition 5. *For the BSC cell of parameter ϵ,*

$$C(\zeta) = \begin{cases} 1 - \frac{1-\zeta(1-\epsilon)}{\epsilon} H_b(\epsilon), & 1 \leq \zeta < \frac{1}{1-\epsilon}, \\ 1, & \zeta \geq \frac{1}{1-\epsilon}. \end{cases}$$

Moreover, for $1 \leq \zeta < 1/(1-\epsilon)$, capacity is achieved by a distribution of S on $\{s_0^\infty, s_1^\infty, s_0^1, s_1^1\} \subset \mathcal{S}$ with

$$\mathbb{P}\left[S = s_0^1\right] = \mathbb{P}\left[S = s_1^1\right] = \frac{1 - \zeta(1-\epsilon)}{2\epsilon},$$

and

$$\mathbb{P}\left[S = s_0^\infty\right] = \mathbb{P}\left[S = s_1^\infty\right] = \frac{1}{2} - \mathbb{P}\left[S = s_0^1\right],$$

where the strategies $s_0^\infty, s_1^\infty, s_0^1, s_1^1$ are defined by

$$X_k(s_i^\infty) = i, \quad T_k(s_i^\infty) = \{i\},$$

and

$$X_k(s_i^1) = i, \quad T_k(s_i^1) = \{0, 1\},$$

for $i \in \{0, 1\}$ and $k = 1, 2, \ldots$

V. Conclusion

We have presented a closed-form expression for the capacity of symmetric cells as a function of the maximum number of rewrites. Some properties of the capacity for general cells have been discussed; in particular, we showed that every cell with two distinct reachable states yields positive capacity provided that at least one rewrite is allowed. We extended the capacity formula for the BSC cell to the case where only noisy versions of the state are available to the encoder and the decoder. Finally, a lower bound on the capacity of cells subject to a constraint on the average number of rewrites was derived.

Acknowledgment

The authors thank Ligong Wang for helpful discussions.

References

[1] L. Lastras-Montano, M. Franceschini, T. Mittelholzer, and M. Sharma, "Rewritable storage channels," *Proc. ISITA08*, pp. 7–10.

[2] T. Mittelholzer, M. Franceschini, L. Lastras-Montano, I. Elfadel, and M. Sharma, "Rewritable channels with data-dependent noise," in *2009 International Conference on Communications*, 2009, pp. 1–6.

[3] M. Franceschini, L. Lastras-Montano, T. Mittelholzer, and M. Sharma, "The role of feedback in rewritable storage channels [Lecture Notes," *IEEE Signal Processing Magazine*, vol. 26, pp. 190–194, 2009.

[4] L. Lastras-Montano, T. Mittelholzer, and M. Franceschini, "Superposition coding in rewritable channels," in *Information Theory and Applications workshop (ITA2010)*.

[5] L. Lastras-Montano, M. Franceschini, and T. Mittelholzer, "The capacity of the uniform noise rewritable channel with average cost," in *Information Theory Proceedings (ISIT), 2010 IEEE International Symposium on*. IEEE, 2010, pp. 201–205.

[6] T. Mittelholzer, L. Lastras-Montano, M. Sharma, and M. Franceschini, "Rewritable storage channels with limited number of rewrite iterations," in *Information Theory Proceedings (ISIT), 2010 IEEE International Symposium on*. IEEE, 2010, pp. 973–977.

[7] T. Cover and J. Thomas, *Elements of information theory*. wiley, 2006.

[8] A. Marshall and I. Olkin, "Inequalities: Theory of Majorization and Its Applications, Mathematics in Science and Engineering, Vol 143," 1979.