RC24653 (W0804-110) April 30, 2008
Mathematics

IBM Research Report

Rewritable Storage Channels

L. A. Lastras-Montaño M. Franceschini, T. Mittelholzer*, M. Sharma
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598
USA

*IBM Zurich Research Laboratory
Säumerstrasse 4
CH-8803 Rüschlikon
Switzerland

International Symposium on Information Theory and its Applications, ISITA2008
Auckland, New Zealand, 7-10, December, 2008

Rewritable storage channels

L. A. Lastras-Montaño†, M. Franceschini†, T. Mittelholzer‡ and M. Sharma†

† IBM T.J. Watson Research Center
1101 Kitchawan Rd, Yorktown Heights, NY, 10598
e-mail: lastrasl, franceschini, mxsharma@us.ibm.com

‡ IBM Zurich Research Laboratory
Säumerstrasse 4
CH-8803 Rüschlikon, Switzerland
e-mail: tmi@zurich.ibm.com

Abstract

We introduce a general model of the rewritable storage channel in which rewriting can be used to improve storage capacity at the cost of encoding time. We then focus on a particular rewritable storage channel and show that if the average number of writes increases by a factor of $\boldsymbol{c}$, storage capacity increases by an additional $\mathbf{log}\,\boldsymbol{c}$ bits; we conjecture that capacity gains due to iterations for other scenarios will follow a similar trend. We will also discuss a number of open problems with practical significance in the context of rewritable storage channels.

1. Introduction

What limits the information storage capacity of a memory that can take on analog values? It may be that reading the memory with great accuracy is very difficult. It may also be that after writing with high precision a value to a memory, its contents degrade over time, either as a result of external effects or due to the nature of the physical medium itself. Or perhaps, intersymbol interference prevents an encoder from writing arbitrary values to nearby locations. Alternatively, it is possible that the writing mechanisms available in a memory have uncertainties that place limits on how well one may approach a desired analog level. In the latter, such uncertainties may be related to the fact that the input/output physical behavior of the medium may possess fundamentally random properties. Even if such behavior is deterministic, it may be governed by parameters that are unknown to the writing mechanism.

In a rewritable memory, the ability to read and possibly rewrite immediately creates a very flexible feedback path that one may exploit in order to improve storage capacities in those technologies which suffer from write mechanisms with some degree of uncertainty in their outcome. This kind of feedback is different from the classical notion of feedback introduced by Shannon [1] in the context of communication systems, since overwritten data never reaches an external read request.

Practitioners in the storage field have long recognized the advantage of "rewrites". For example, error correction protected memory is sometimes periodically read, decoded, and written back (a process called scrubbing) in order to avoid the buildup of errors in the memory. An example with a different flavor and much closer to the model we study can be found in Flash memory in which "write-and-verify" techniques are employed for attaining reliable multibit cell storage [2]. In these techniques a target level is approached through a number of write steps with intervening reads that provide feedback on the state of the Flash cell after the last write. Such feedback allows the device to get an improved control over the final analog level reached and also diminish potential dependencies on Flash cell variability.

The theoretical study of memories has a long distinguished tradition in information theory and computer science. As examples, we point to the original work of Kuznetsov and Tsybakov on memories with defective cells [3] and Rivest and Shamir's work on Write Once Memories [4]. The memory models in these groundbreaking articles have been studied in more depth, generalized, and also modified to adapt to other relevant settings. Examples of these developments can be found in Heegard and El Gamal [5], Fiat and Shamir [6], Wolf, Wyner, Ziv and Körner [7], Heegard [8], Simonyi [9], Kuznetsov and Vinck [10], Ahlswede and Zhang [11], Fu and Vinck [12] and Fu and Yeung [13].

A common thread in these works is the idea that the existing contents on a memory can either restrict or assign a cost to what can be written next. The problem we consider in this work is of an entirely different nature, as the focus in our model is a write mechanism that is inherently unreliable. The goal for this article is to draw to the reader's attention the role that information theory can play in understanding how the ability

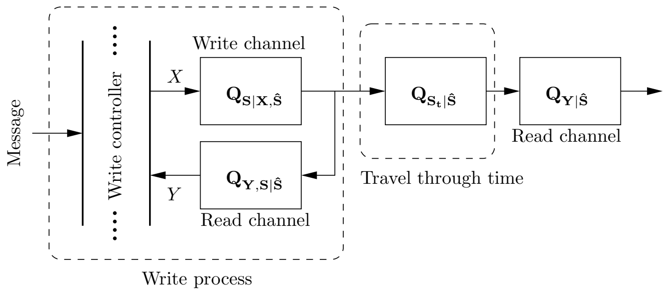

Figure 1: Schematic of a model for the rewritable storage channel.

of rewriting in a memory may impact its ultimate storage capabilities. In order to sharpen our exposition, we focus on the scenario where a single message is transferred to the memory, leaving the treatment of multiple write requests for future work.

The remaining material in this paper is organized as follows: in Section 2 we introduce a general model for the rewritable channel which is intended to convey to the reader the richness of this relevant problem space. In Section 3, we will present upper and lower bounds on rewritable storage channel capacity of an interesting problem instance. In Section 4 we discuss this result along with numerical results for the Gaussian rewritable channel model. Acknowledgements are found in Section 5; the Appendix (Section 6) contains the mathematical details needed to complete the proofs in Section 3.

2. General problem setup

We propose a model of a memory that consists of discrete memory parts which we call cells. In response to a stimulus, a cell takes on a new internal state in a manner dependent on its current state and according to a statistical relation $Q_{S|X,\hat{S}}(s|x,\hat{s};\boldsymbol{\theta})$, which we call the *write channel*, where $x \in \mathcal{X}$ denotes the write stimulus (or input) to the cell, $s \in \mathcal{S}$ denotes its final state after the write stimulus is applied, $\hat{s} \in \mathcal{S}$ denotes the state at the time x is applied and $\boldsymbol{\theta}$ denotes a vector of potentially unknown parameters for the cell.

The *read channel* is modeled as $Q_{Y,S|\hat{S}}(y,s|\hat{s};\boldsymbol{\theta})$, where $y \in \mathcal{Y}$, which allows for possible state changes as a result of the read operation. These changes are generally an undesired side effect of the act of reading. We assume that the state of a cell is observable only through Y in the read channel defined above. The alphabets $\mathcal{X}$, $\mathcal{S}$ and $\mathcal{Y}$ will be subsets of the real line in the particular examples examined in detail in this article, but of course can be modified to fit particular storage medium characteristics. These concepts are illustrated in Figure 1, where the read and write channels are placed in the context of a write process managed by a *write controller* that in general operates on multiple cells.

We shall be interested in storing information into a group of n cells. The role of the write controller is to accept a message $W \in \{1,\cdots,2^{nR}\}$, where $R \in (0,1]$ is a rate parameter, to be encoded in the memory and to provide input signals to be applied to the memory so that a future requestor can retrieve the intended message with very high probability. The write controller plays a role similar to that of channel encoder in traditional transmission schemes. Nonetheless, an extra degree of freedom is allowed: the write controller may read after writing and decide to write again an arbitrary subset of the n cells. Physically this increases the time for writing and also increases the amount of energy spent for encoding. For memories that degrade after a large number of write/read cycles, the lifetime of the memory may also be affected. *Nonetheless, a write controller may find advantageous to exercise the rewrite ability in order to increase the storage capacity and/or reduce decoding error rates for the underlying memory system.* The set of rules that the write controller uses to coordinate its actions is termed a *write policy*. This policy in general can be adaptive - future write actions may be determined not only by the message one is attempting to encode, but also by the results of previous write attempts.

After the write process, the state of a cell may change with time. This is modeled with a channel $Q_{S_t|\hat{S}}(s_t|\hat{s};t,\boldsymbol{\theta})$, as shown in Figure 1, which is parametrized by the amount of time t elapsed between the finish of the write process and its first subsequent read.

We assume that each individual cell can have a parameter vector $\boldsymbol{\theta}_i,\ i \in \{1,\cdots,n\}$, although some of these parameters may be common to all n cells. The cells are statistically independent. In order to address this and related concepts it will be useful to introduce some notation. We denote vectors with n entries with capital bold letters. Individual entries in a vector are denoted using the same letter without the bold font and with a subscript denoting the index within the vector. Suppose that $Q_{A|B}$ is a conditional density for random variables taking on an alphabet $\mathcal{A}$ where the conditioning random variable is in an alphabet $\mathcal{B}$. The notation $\mathbf{A} \overset{Q_{A|B}}{\longleftarrow} \mathbf{B}$ defines $\mathbf{A}$ to be a random vector with n entries, each entry obtained by passing the corresponding entry in $\mathbf{B}$ through the channel $Q_{A|B}$ in a stochastically independent manner. Also, the notation $Q_{A|B} \circ Q_{B|C}$ is the channel $Q_{A|C}(a|c) = \sum_b Q_{A|B}(a|b)Q_{B|C}(b|c)$.

A write process is associated with a cost, generally related to the number of iterations required to finish it.

We may wish to place a limit on the maximum number iterations since in some settings the time for encoding a message into the n cells is dominated by the worst time across all cells. We may also place a limit on the average number of iterations since one expects it is related to energy consumption during a write process.

Formally, a write controller is given by

1. a sequence of *stimuli generator* functions

$$f^i(w, \mathbf{y}^1 \cdots \mathbf{y}^{i-1}) : \{1, ..., 2^{nR}\} \times \mathcal{Y}^n \times \cdots \times \mathcal{Y}^n \to \mathcal{X}^n$$

where i is the iteration count,

2. a stopping set for each cell $j \in \{1, \cdots, n\}$ on each iteration i called $D^i_j(w) \subset \mathcal{Y}$.

The first argument in f^i (the only argument to D^i_j) is the message we wish to encode in the memory. The notation $\mathbf{y}^{i-1}$ denotes the read contents of the n cells at the time prior to applying the ith write iteration. We assume that every write is followed by a read. The first write, indexed by 1, is a special case as it could be aided by a previous read or not. Accordingly, the f^1 function may or may not have a dependence on a preliminary read $\mathbf{y}^0$; in this article we assume that there is no such dependence, and hence no read prior to the first write. One may also consider policies in which a final write is not followed by a "confirming" read operation. We do not consider this type of policies in this work.

Let $\mathbf{S}^0$ be the state of the n cells prior to the first write; this random variable is distributed according to a law $P_{\mathbf{S}^0}$. We assume that the message W to encoded in the memory is uniformly distributed over the range $\{1, \cdots, 2^{nR}\}$. The write process is defined by

$$\begin{aligned} \mathbf{X}^i &= f^i(W, \mathbf{Y}^1 \cdots \mathbf{Y}^{i-1}) \\ (\mathbf{Y}^i, \mathbf{S}^i) &\overset{Q_{Y,S|\hat{S}} \circ Q_{S|X,\hat{S}}}{\longleftarrow} (\mathbf{X}^i, \mathbf{S}^{i-1}) \end{aligned}$$

for $i = 1, 2, \cdots$. The stopping set $D^i_j(W)$ is used to determine when it is that a cell has an acceptable content, *after* the ith write is applied and the resulting contents are observed through a read operation. Define the stopping time for the jth cell as $L_j = \min\{i \geq 1 : Y^i_j \in D^i_j(W)\}$.

Physically, once a cell has met its stopping condition, the write controller will cease writing or reading to that cell. In a situation in which all cells are written at the same time, the write controller finishes the entire write operation at time $\max_{1 \leq j \leq n} L_j$. The state of the cells after the write process is given by $(S^{L_1}_1, \cdots S^{L_n}_n)$. If the memory is then accessed at time t for retrieval of the stored data, the resulting observations are given by $\mathbf{V}_t \overset{Q_{Y|S} \circ Q_{S_t|\hat{S}}}{\longleftarrow} (S^{L_1}_1, \cdots S^{L_n}_n)$. Assuming that each write attempt has cost equal to one, the average cost associated with this write controller is given by $\frac{1}{n}\sum_{j=1}^n EL_j$. In this article we will wish to place an upper constraint on this average cost and we shall use the letter κ to denote this constraint. As we discussed earlier, other cost metrics may be employed; we choose the one above in the basis of simplicity. The reading mechanism in the memory is given by a decoding function $g : \mathcal{Y}^n \to \{1, \cdots, 2^{nR}\}$.

The probability of error of this write controller and decoding function for reading at time t is given by $P(g(\mathbf{V}_t) \neq W)$. We say that the rate R is achievable at cost κ with initial state distribution $P_{\mathbf{S}^0}$ and reading time t if for every $\epsilon > 0$ there is a write controller $\{f^i\}_{i \geq 1}$ and read mechanism g with rate R for a sufficiently large number of cells n, such that

$$P(g(\mathbf{V}_t) \neq W) \leq \epsilon, \qquad \frac{1}{n}\sum_{j=1}^{n} EL_j \leq \kappa.$$

The largest such R for a fixed cost κ is the capacity of this rewritable storage channel and we use the notation $\mathsf{C}(\kappa, P_{\mathbf{S}^0}, t)$ to denote it. If this capacity does not depend on $P_{\mathbf{S}^0}$ nor t, as in the main contribution of this paper, we will use the simpler notation $\mathsf{C}(\kappa)$ instead. The dependence of capacity on the initial distribution $P_{\mathbf{S}^0}$ will allow future work to study several important notions of capacity. For example, one may consider writing different messages to a memory repeatedly, the writing of a message determining the initial state for the writing of the next message. In this case it is sensible to relate $P_{\mathbf{S}^0}$ to a stationary behavior of the write controller (assuming such behavior exists). Or one may possess a mechanism for initializing the cells to a known state in which case it is interesting to study capacity when the intial state distribution places unit mass on this known state.

3. Analysis of a simple model

In the following we will focus on a model of a rewritable storage channel that although simple, captures the essence of this article's view of the role of rewriting in memories.

3.1. A model with uniform noise

We shall assume that the state of the cell is directly observable by a read, without any noise and without effecting any new state changes on the cell. We shall also assume that there are no dependencies with respect to the current state of the cell $\hat{s}$ when we obtain a next state s due to an input x ($Q_{S|X,\hat{S}} = Q_{S|X}$). We shall also assume that there are no state changes due to the passage of time after the write process is finished.

Within these constraints, the basic building block is a channel $Q_{Y|X} : \mathcal{Y} \times \mathcal{X} \to [0,1]$ that describes the statistical relation between the input and the output of an individual cell. We assume that $\mathcal{Y} = [-a/2, 1 + a/2]$ and that $\mathcal{X} = [0,1]$, where a is a real number whose role will be evident shortly. We assume that $Q_{Y|X}$ is such that $Y = x + N$ where N is a random variable uniformly distributed in the interval $[-a/2, a/2]$. Thus if $\frac{1+a}{a}$ is an integer, then it is clear that one can store $\log \frac{1+a}{a}$ bits in a single cell with no coding and zero probability of error; more generally for any $a \leq 1$ there is a simple strategy for encoding $\log\lfloor \frac{1+a}{a} \rfloor$ bits.

3.2. The result and its proof

In our capacity result for this channel model we further limit the class of stimuli generator functions to depend solely on the message w. Thus for every iteration i, $f^i(w) : \{1, \cdots, 2^{nR}\} \to \mathcal{X}^n$.

Within this restriction, our main result is as follows:

Theorem 1

$$\log \left\lfloor \frac{1+a}{a}\kappa \right\rfloor \leq \mathsf{C}(\kappa) \leq \log \left(\frac{1+a}{a}\kappa \right). \qquad (1)$$

3.2.1. Proof of the lower bound

The basic idea is to use rewrites for shaping the noise to be uniformly distributed with a smaller range. To this end, let $0 < b < a$. We construct from $\mathcal{Y}$ disjoint open intervals each of length b. Then one can obtain a cell storage capacity of $\log\lfloor \frac{1+a}{b} \rfloor$ bits by selecting as input the center of any of the $\lfloor \frac{1+a}{b} \rfloor$ intervals[1], and then attempting as many writes as necessary in order to fall within the desired interval of length b. The average number of iterations is then a/b.

3.2.2. Proof of the upper bound

Define $\Delta_j^i = D_j^i \cap [X_j^i - a/2, X_j^i + a/2]$. The significance of Δ_j^i is that it is the *effective* stopping set for cell j at iteration i, since given an input X_j^i, one may only reach values in the interval $[X_j^i - a/2, X_j^i + a/2]$. If $A \subset \mathcal{Y}$, let $|A|$ denote the total length of the subset; formally if 1_A denotes the indicator function for set A, $|A| = \int 1_A(x)dx$ where dx is the Lebesgue measure. Thus for example $|\Delta_j^i| \leq |D_j^i|$. During this proof, we shall make use of the following facts, the proof of which can be found in the Appendix:

Lemma 1 *For any* $j \in \{1, \cdots, n\}$,

1. $h(Y_j^{L_j} | D_j^{L_j}, X_j^{L_j}) = E \log |\Delta_j^{L_j}|$,
2. $(Y_1^{L_1} \cdots Y_{j-1}^{L_{j-1}}, W) \to (X_j^{L_j}, D_j^{L_j}) \to Y_j^{L_j}$,
3. $E\left[\frac{a}{|\Delta_j^{L_j}|} \right] = EL_j$.

The proof starts with a variant of the familiar argument based on Fano's inequality:

$$\begin{aligned} nR &= H(W) \\ &= I(W; Y_1^{L_1} \cdots Y_n^{L_n}) + H(W | Y_1^{L_1} \cdots Y_n^{L_n}) \\ &\leq I(W; Y_1^{L_1} \cdots Y_n^{L_n}) + 1 + nR\epsilon. \end{aligned}$$

We continue the proof as follows:

$$\begin{aligned} & I(W; Y_1^{L_1} \cdots Y_n^{L_n}) \\ &= h(Y_1^{L_1} \cdots Y_n^{L_n}) - h(Y_1^{L_1} \cdots Y_n^{L_n} | W) \\ &\overset{(a)}{\leq} n \log(1+a) - h(Y_1^{L_1} \cdots Y_n^{L_n} | W) \\ &= n \log(1+a) - \sum_{j=1}^n h(Y_j^{L_j} | Y_1^{L_1} \cdots Y_{j-1}^{L_{j-1}}, W) \\ &\overset{(b)}{\leq} n \log(1+a) \\ &\quad - \sum_{j=1}^n h(Y_j^{L_j} | D_j^{L_j}, X_j^{L_j}, Y_1^{L_1} \cdots Y_{j-1}^{L_{j-1}}, W) \\ &\overset{(c)}{=} n \log(1+a) - \sum_{j=1}^n h(Y_j^{L_j} | D_j^{L_j}, X_j^{L_j}) \end{aligned}$$

where (a) follows from the fact that the maximum differential entropy distribution with support $[-a/2, 1 + a/2]$ is $\log(1+a)$, (b) follows from the fact that conditioning does not increase differential entropy and (c) follows from Lemma 1. We then write

$$\begin{aligned} & I(W; Y_1^{L_1} \cdots Y_n^{L_n}) \\ &\overset{(d)}{\leq} n \log(1+a) - \sum_{j=1}^n E \log |\Delta_j^{L_j}| \\ &= n \log \frac{1+a}{a} + \sum_{j=1}^n E \log \frac{a}{|\Delta_j^{L_j}|} \\ &\overset{(e)}{\leq} n \log \frac{1+a}{a} + \sum_{j=1}^n \log E \frac{a}{|\Delta_j^{L_j}|} \\ &\overset{(f)}{=} n \log \frac{1+a}{a} + n \sum_{j=1}^n \frac{1}{n} \log EL_j \\ &\overset{(g)}{\leq} n \log \frac{1+a}{a} + n \log \left(\frac{1}{n} \sum_{j=1}^n EL_j \right) \\ &= n \log \left(\frac{1+a}{a}\kappa \right) \end{aligned}$$

[1] An adjustment to this description is needed at the borders.

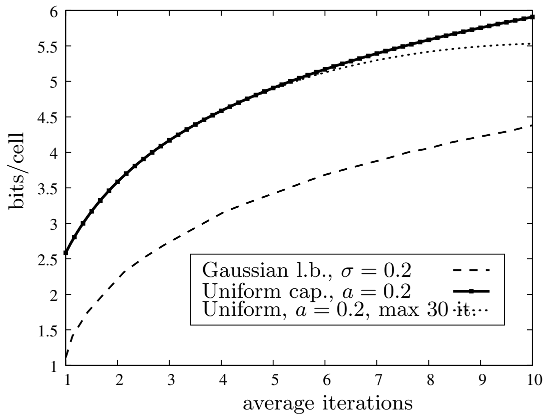

Figure 2: Capacity bounds for two models.

where (d), (f) follow from Lemma 1 and (e), (g) follow from Jensen's inequality.

4. Discussion and Concluding Remarks

The additive logaritmic increase with the average number of iterations is illustrated in Figure 2, in which the solid line is the capacity upper bound in (1), the square dots correspond to the points in which the lower bound and the upper bound touch. The dotted line represents an *upper bound* on capacity when in addition to an average number of iterations, a constraint on the maximum number of iterations is specified (max = 30 in this case); we expect these type of bounds to be more relevant in practice. The dashed line is a lower bound to the capacity of an additive white Gaussian write noise channel, with writing signal constrained to be in the interval $[0, 1]$ and impaired by a zero mean Gaussian write noise with standard deviation $\sigma = 0.2$.

The qualitative behavior of storage capacity as a function of the average number of iterations, in the case with arbitrary write noise models, is an open problem, although numerical evidence suggests that the logarithmic behavior is likely to be common. Determining the fundamental tradeoff between write cost and storage capacity in presence of unknown, possibly random parameters, is also an open problem that we are currently addressing.

We believe that in the long term, development of the simple ideas presented in this paper can have significant implications in the design of real memory/storage systems.

5. Acknowledgments

We acknowledge the deeply insightful discussions that we have had with Stefanie R. Chiras and John P. Karidis during the course of our work. We also acknowledge the guidance that we have received on memory technologies by our colleagues Matt Breitwisch, Roger Cheek, Bipin Rajendran and Chung Lam.

6. Appendix: Proof of Lemma 1.

To prove item 1, we omit the dependence of the various random variables on j to simplify the notation. Write

$$h(Y^L|D^L, X^L) = \int h(Y^L|D^L = d, X^L = x) d\mu(x, d)$$

where μ is the probability law governing (D^L, X^L). The insight is that Y^L, under the conditioning above, is uniformly distributed over $d \cap [x - a/2, x + a/2]$. The proof is as follows:

$$\begin{aligned} & P(Y^L \le y | D^L = d, X^L = x) \\ &\quad = \sum_{i\ge 1} P(Y^L \le y, L = i | D^L = d, X^L = x) \\ &\quad = \sum_{i\ge 1} P(L = i | D^L = d, X^L = x) \times \\ &\qquad\quad P(Y^i \le y | D^L = d, X^L = x, L = i) \end{aligned}$$

and note that by definition of L and the assumption on the probability law governing $(D^1, Y^1, D^2, Y^2, \cdots)$,

$$\begin{aligned} & P(Y^i \le y | D^i = d, X^i = x, L = i) \\ &\quad = P(Y^i \le y | D^i = d, X^j = x, Y^1 \notin D^1, \cdots \\ &\qquad Y^{i-1} \notin D^{i-1}, Y^i \in D^i) \\ &\quad = P(Y^i \le y | X^i = x, Y^i \in d). \end{aligned}$$

From this it is easy to see that Y^i, conditioned on $X^i = x, Y^i \in d$, is uniformly distributed over $d \cap [x - a/2, x + a/2]$, which is also Δ^i under the same conditioning. The proof follows immediately. To demonstrate the second item, let

$$\mathcal{I} = \Big\{ i \ge 1 : P(L_j = i, D_j^{L_j} = d, X_j^{L_j} = x, Y_1^{L_1} = y_1 \\ , \cdots Y_{j-1}^{L_{j-1}} = y_{j-1}, W = w) > 0 \Big\}$$

and let $d, x, y_1, \cdots, y_{j-1}, w$ be any choice for which the set above is nonempty. Now write

$$
\begin{aligned}
& P(Y_j^{L_j} \leq y | D_j^{L_j} = d, X_j^{L_j} = x, Y_1^{L_1} = y_1, \cdots \\
& \quad Y_{j-1}^{L_{j-1}} = y_{j-1}, W = w) \\
= & \sum_{i \in \mathcal{I}} P(Y_j^{L_j} \leq y, L_j = i | D_j^{L_j} = d, \\
& X_j^{L_j} = x, Y_1^{L_1} = y_1, \cdots Y_{j-1}^{L_{j-1}} = y_{j-1}, W = w) \\
= & \sum_{i \in \mathcal{I}} P(Y_j^i \leq y | L_j = i, D_j^i = d, X_j^i = x) \times \\
& P(L_j = i | D_j^{L_j} = d, X_j^{L_j} = x, Y_1^{L_1} = y_1 \\
& , \cdots Y_{j-1}^{L_{j-1}} = y_{j-1}, W = w).
\end{aligned}
$$

From the proof of item 1, we know that $P(Y_j^i \leq y | L_j = i, D_j^i = d, X_j^i = x) = P(Y_j^i \leq y | Y_j^i \in d, X_j^i = x)$. Note that the latter has the same value for all iterations $i \in \mathcal{I}$. Pick an $i^* \in \mathcal{I}$, then

$$
\begin{aligned}
& P(Y_j^{L_j} \leq y | D_j^{L_j} = d, X_j^{L_j} = x, Y_1^{L_1} = y_1, \cdots \\
& \quad Y_{j-1}^{L_{j-1}} = y_{j-1}, W = w) \\
= & \quad P(Y_j^{i^*} \leq y | Y_j^{i^*} \in d, X_j^{i^*} = x).
\end{aligned}
$$

The latter has no dependence on $y_1, \cdots, y_{i-1}$ nor w, which proves the desired Markov chain property.

To prove the third item, we shall condition on W and focus on the conditioned terms:

$$
\begin{aligned}
E\left[\frac{a}{|\Delta_j^{L_j}|}\right] &= \sum_w P(W = w) E\left[\frac{a}{|\Delta_j^{L_j}|} | W = w\right] \\
EL_j &= \sum_w P(W = w) E[L_j | W = w].
\end{aligned}
$$

Define $p_i = P(Y_j^i \in D_j^i | Y_j^1 \notin D_j^1, \cdots Y_j^{i-1} \notin D_j^{i-1}, W = w)$ so that $P(L_j = i | W = w) = (1-p_1) \cdots (1-p_{i-1}) p_i$. Note that $p_i = \frac{|\Delta_j^i(w)|}{a}$ because conditioned on W, both X_j^i and D_j^i are fully known. One then writes

$$
\begin{aligned}
& E\left[\frac{a}{|\Delta_j^{L_j}|} | W = w\right] = \sum_{i \geq 1} (1-p_1) \cdots (1-p_{i-1}) \frac{a p_i}{|\Delta_j^i(w)|} \\
= & \sum_{i \geq 1} (1-p_1) \cdots (1-p_{i-1}) = \sum_{i \geq 1} P(L_j \geq i | W = w) \\
= & E[L_j | W = w]
\end{aligned}
$$

where the last step is a standard equality for nonnegative integer valued random variables. This finishes the proof of the Lemma.

References

[1] C.E. Shannon. The zero-error capacity of a noisy channel. *IRE Transactions on Information Theory*, IT-2:8–19, 1956.

[2] B. Ricco M. Grossi, M. Lanzoni. Program schemes for multilevel flash memories. *Proceedings of the IEEE*, 91(4), April 2003.

[3] A. V. Kuznetsov and B.S. Tsybakov. Coding in a memory with defective cells. *Problemy Peredachi Informatsii*, 10(2):52–60, April-June 1974.

[4] R. L. Rivest and A. Shamir. How to reuse a write-once memory. *Inform. and Control*, 55:1–19, 1982.

[5] C. Heegard and A.A. El Gamal. On the capacity of computer memory with defects. *IEEE Transactions on Information Theory*, 29(5):731–739, September 1983.

[6] A. Fiat and A. Shamir. Generalized "write-once" memories. *IEEE Transactions on Information Theory*, 30:470–480, May 1984.

[7] J.K. Wolf, A.D. Wyner, J. Ziv and J. Körner. Coding for a write-once memory. *AT&T Bell Laboratories Technical Journal*, 63(6):1089–1112, July-August 1984.

[8] C. Heegard. On the capacity of permanent memory. *IEEE Transactions on Information Theory*, 31:34–42, January 1985.

[9] G. Simonyi. On write-unidirectional memory codes. *IEEE Trans. Inform. Theory*, 35(3):663–667, May 1989.

[10] A. V. Kuznetsov and A.J. Han Vinck. On the general defective channel with informed encoder and capacities of some constrained memories. *IEEE Transactions on Information Theory*, 40(6):1866–1871, November 1994.

[11] R. Ahlswede and Z. Zhang. Coding for write-efficient memory. *Inform. and Comput.*, 83(1):80–97, 1989.

[12] F-W Fu and A.J. Han Vinck. On the capacity of generalized write-once memory with state transitions described by an arbitrary directed acyclic graph. *IEEE Transactions on Information Theory*, 45(1):308–313, January 1999.

[13] F-W Fu and R. Yeung. On the capacity and error correcting codes for write efficient memories. *IEEE Transactions on Information Theory*, 46(7):2299–2314, November 2000.

[14] S.W. McLaughlin. Signal processing and coding for m-ary optical storage. *Proceedings of the 2005 International Conference on Acoustics, Speech and Signal Processing*, 5:18–23, March 2005.