Combinatorics, Probability and Computing (2013) 22, 118–132. © Cambridge University Press 2012
doi:10.1017/S0963548312000478

Improved Mixing Time Bounds for the Thorp Shuffle

BEN MORRIS[†]

Department of Mathematics, University of California, Davis, CA 95616, USA
(e-mail: morris@math.ucdavis.edu)

Received 29 December 2010; revised 12 September 2012

E. Thorp introduced the following card shuffling model. Suppose the number of cards is even. Cut the deck into two equal piles, then interleave them as follows. Choose the first card from the left pile or from the right pile according to the outcome of a fair coin flip. Then choose from the other pile. Continue this way, flipping an independent coin for each pair, until both piles are empty.

We prove an upper bound of $O(d^3)$ for the mixing time of the Thorp shuffle with 2^d cards, improving on the best known bound of $O(d^4)$. As a consequence, we obtain an improved bound on the time required to encrypt a binary message of length d using the Thorp shuffle.

AMS 2010 *Mathematics subject classification*: Primary 60J10

1. Introduction

1.1. The Thorp shuffle

Thorp [17] introduced the following card shuffling model in 1973. Assume that the number of cards, n, is even. Cut the deck into two equal piles, then interleave them as follows. Choose the first card from the left pile or from the right pile according to the outcome of a fair coin flip. Then choose from the other pile. Continue this way, flipping an independent coin for each pair, until both piles are empty. This is one time step.

1.2. Connection to cryptography

We now briefly describe the connection between the Thorp shuffle and cryptography. For a more comprehensive account, see [10]. Let x be a binary string of length d. Think of x as the binary representation of the initial position of a card in the deck (where the positions range from 0 at the bottom to $2^d - 1$ at the top), and let $\mathrm{Thorp}(x, t)$ be the position of the card at time t. The behaviour of $\mathrm{Thorp}(x, t)$ for a given x is easy to describe (see, *e.g.*, [8]). One can obtain $\mathrm{Thorp}(x, t + 1)$ recursively from $\mathrm{Thorp}(x, t)$ as follows:

[†] Research partially supported by NSF grant DMS-1007739.

https://doi.org/10.1017/S0963548312000478 Published online by Cambridge University Press

1 do a circular shift of the bits left by one position, and then
2 flip the bit now in the rightmost position with probability $\frac{1}{2}$ (according to the outcome of the corresponding coin flip at time t)

Note that one can trace the route of any given card in the deck without attending to the remaining cards. Moni Naor called this property *obliviousness*.

It was Naor [11, p. 62] who first noticed potential cryptographic applications of the Thorp shuffle. (See also Rudich [16, p.17].) An analysis of the Thorp shuffle was open at that time, but it is now better understood. Recently, in [10], the author, P. Rogaway and T. Stegers have used the Thorp shuffle as the basis for an encryption algorithm. This algorithm, which is briefly described in the following subsection, is contained in a proposal currently under consideration by the US National Institute of Standards and Technology [1].

1.2.1. Format-preserving encryption. Suppose one wants to encipher a 10-digit social security number so that its encoding is also a 10-digit number, or a 16-digit credit card number, the result being of the same form. This would allow the encryption of the entries of a database without affecting its operation. Rising concerns about identity theft and new banking regulations that prohibit storage of unencrypted credit card numbers in databases give this problem practical importance.

The mathematical setup for this problem is as follows. A *random function* is a function whose argument is a binary string and whose output is a long block of random bits. For example, if $\{0,1\}^*$ denotes the set of all binary strings, then $g : \{0,1\}^* \to \{0,1\}^{128}$ is a random function, if the random variables $\{g(x) : x \in \{0,1\}^*\}$ are independent uniform random samples from $\{0,1\}^{128}$. Let $\mathcal{M} = \{1,\ldots,N\}$ be a set of messages. The problem is to turn a random function g into an algorithm that returns $f(x)$ on input x, where f is a uniform random permutation $f : \mathcal{M} \to \mathcal{M}$. (More precisely, f must be *pseudorandom*, meaning that an adversary who can query f and f^{-1} cannot distinguish it from a uniform permutation within a practical number of steps.)

A naïve solution would be to use the random function to generate random numbers for each value of $x \in \mathcal{M}$ and then define f according to the induced ordering. The problem with this scheme is that it requires time and space on the order of N (*e.g.*, 10^{16} in the case of credit card numbers). What is desired is an algorithm to encode messages in time that is polynomial in the *length* of the message, *i.e.*, polylogarithmic in the number of messages.

1.2.2. Feistel networks. A classical method for format-preserving encryption is a balanced Feistel network. Suppose that the set of messages is $\mathcal{M} = \{0,1\}^d$ with d even, so that the number of messages $N = 2^d$. A Feistel network consists of m rounds (*i.e.*, shuffles), where for $1 \leqslant i \leqslant m$, round i makes the transformation $(L,R) \to (R, L \oplus F_i(R))$. Here $\oplus$ denotes xor, the F_i are independent random functions from $\{0,1\}^{d/2}$ to $\{0,1\}^{d/2}$, and L and R denote the leftmost and rightmost $d/2$ bits of the input, respectively.

A long sequence of papers, starting with the classic result of Luby and Rackoff (see [4], [5], [11], [6], [12], [13] and [14]) have analysed balanced Feistel networks. These culminated in Patarin's result, which says that an adversary must make on the order of (roughly) $2^{d/2}$ queries in order to distinguish a balanced Feistel network from a uniform permutation. (This is essentially

https://doi.org/10.1017/S0963548312000478 Published online by Cambridge University Press

the best bound of this type possible unless the number of rounds is exponential in d: see [10] for a discussion.)

1.2.3. Thorp encryption. When the number of bits d is only moderately large (*e.g.*, $20 \leqslant d \leqslant 50$), such as in the case of a social security number or credit card number, then the $2^{d/2}$ bound might be considered insufficient. In this case, an improved guarantee can be obtained using the Thorp shuffle.

In Thorp encryption, the encoding of the binary string x is $\mathrm{Thorp}(x, t)$ for some specified value of t. In [10], it is shown that for any $\epsilon > 0$ there is a constant C such that if the number of shuffles $t = Cd$, then the Thorp shuffle is indistinguishable from a uniform permutation unless the number of queries is at least $2^{(1-\epsilon)d}$. The proof is based on bounding the time required for the Thorp shuffle to mix up $2^{(1-\epsilon)d}$ of the cards.

The mixing time Even better security is obtained when all the cards are mixed. Let τ_{mix} be the *mixing time* of the Thorp shuffle, that is, the number of shuffles necessary to mix up the whole deck (see Section 2 for a precise definition). Note that if we perform $t = \tau_{\mathrm{mix}}$ shuffles, that is, we encode x by $\mathrm{Thorp}(x, \tau_{\mathrm{mix}})$, then an adversary cannot distinguish the shuffle from a uniform permutation after any number of queries, since the ordering of the deck has an almost uniform distribution. This motivates determining the mixing time of the Thorp shuffle.

An analysis of the Thorp shuffle was a long-standing open problem until the author, in [8], proved a mixing time bound of $O(d^{44})$. This was subsequently improved to $O(d^{29})$ [7] and then to $O(d^4)$ [9]. In the present paper, we improve the bound to $O(d^3)$. This is close to being practical for applications, and not too far off the trivial lower bound, which is on the order of d.

Our analysis builds on the work in [9], but is more complex. Note that the effect of each coin flip in the Thorp shuffle is to interchange a pair of cards with probability $\frac{1}{2}$. In [9], such an operation is called a *collision*. The proof in [9] proceeds by identifying a large set of collisions that occur over a period of d shuffles and finding their effect on the relative entropy. Crucial to the approach is the fact that when collisions involve distinct cards, the order in which they are performed does not affect the resulting distribution. In the present paper, we dispense with the requirement that the collisions involve distinct cards. This allows for a larger set of collisions to be considered and gives an improved bound on the decay of relative entropy and the mixing time. However, the analysis becomes quite a bit more delicate.

2. Background

In this section we give some basic definitions and extract what is necessary from [9]. Let $p(x, y)$ be transition probabilities for a Markov chain on a finite state space V with a uniform stationary distribution. For probability measures μ and ν on V, define the *total variation distance*

$$\|\mu - \nu\| = \tfrac{1}{2} \sum_{x \in V} |\mu(x) - \nu(x)|,$$

and define the mixing time

$$T_{\mathrm{mix}} = \min\{n : \|p^n(x, \cdot) - \mathcal{U}\| \leqslant \tfrac{1}{4} \text{ for all } x \in V\}, \tag{2.1}$$

where $\mathcal{U}$ denotes the uniform distribution.

https://doi.org/10.1017/S0963548312000478 Published online by Cambridge University Press

For a probability distribution $\{p_i : i \in V\}$, define the (*relative*) *entropy* of p by

$$\mathrm{ENT}(p) = \sum_{i \in V} p_i \log(|V| p_i),$$

where we define $0 \log 0 = 0$. The following well-known inequality (see, *e.g.*, [3]) links relative entropy to total variation distance. We have

$$\|p - \mathcal{U}\| \leqslant \sqrt{\tfrac{1}{2}\mathrm{ENT}(p)}. \tag{2.2}$$

If X is a random variable taking finitely many values, define $\mathrm{ENT}(X)$ as the relative entropy of the distribution of X. For reasons that will become clear later, we shall think of $\mathcal{S}_n$ as the set of permutations on $\{0, \ldots, n-1\}$. We shall think of the distribution of a random permutation in $\mathcal{S}_n$ as a sequence of probabilities of length $n!$, indexed by permutations in $\mathcal{S}_n$. If $\mathcal{F}$ is a sigma field, then we shall write $\mathrm{ENT}(X \mid \mathcal{F})$ for the relative entropy of the conditional distribution of X given $\mathcal{F}$. Note that $\mathrm{ENT}(X \mid \mathcal{F})$ is a random variable. If π is a random permutation in $\mathcal{S}_n$, let $\mathcal{H}_n$ be the trivial sigma field and for $0 \leqslant k \leqslant n-1$, define $\mathcal{H}_k = \sigma(\pi^{-1}(k), \ldots, \pi^{-1}(n-1))$. For $0 \leqslant k \leqslant n-1$, define $\mathrm{ENT}(\pi, k) = \mathbb{E}(\mathrm{ENT}(\pi^{-1}(k) \mid \mathcal{H}_{k+1}))$, where we think of the conditional distribution of $\pi^{-1}(k)$ given $\mathcal{H}_{k+1}$ as being a sequence of length $k+1$. The standard entropy chain rule (see, *e.g.*, [2]) can be used to derive the following proposition regarding relative entropy.

Proposition 2.1. *For any $i \leqslant n-1$, we have*

$$\mathrm{ENT}(\pi) = \mathbb{E}\big(\mathrm{ENT}(\pi \mid \mathcal{H}_i)\big) + \sum_{k=i}^{n-1} \mathrm{ENT}(\pi, k). \qquad \square$$

To compute the relative entropy in the first term on the right-hand side, we think of the distribution of π given $\mathcal{H}_i$ as a sequence of probabilities of length $i!$.

Remark. Since $\mathrm{ENT}(\pi \mid \mathcal{H}_1) = 0$, substituting $i = 1$ into the formula gives

$$\mathrm{ENT}(\pi) = \sum_{k=1}^{n-1} \mathrm{ENT}(\pi, k).$$

If we think of π as representing the order of a deck of cards, with $\pi(i) =$ location of card i, then this allows us to think of $\mathrm{ENT}(\pi, k)$ as the portion of the overall entropy $\mathrm{ENT}(\pi)$ that is attributable to the location k.

We will also need the following easy proposition.

Proposition 2.2. *Suppose that $\mathcal{G}$ and $\mathcal{F}$ are independent sigma fields, and $\mathcal{G}'$ is a sigma field with $\mathcal{G}' \subset \mathcal{G}$. Then $\mathcal{G}$ and $\mathcal{F}$ are conditionally independent given $\mathcal{G}'$.*

Proof. We need to show that for any events $A \in \mathcal{G}$ and $B \in \mathcal{F}$, we have

$$\mathbb{E}(\mathbf{1}_A \mathbf{1}_B \mid \mathcal{G}') = \mathbb{E}(\mathbf{1}_A \mid \mathcal{G}') \mathbb{E}(\mathbf{1}_B \mid \mathcal{G}'). \tag{2.3}$$

https://doi.org/10.1017/S0963548312000478 Published online by Cambridge University Press

By the tower property of expectation,

$$\begin{aligned}\mathbb{E}(\mathbf{1}_A\mathbf{1}_B \mid \mathcal{G}') &= \mathbb{E}\big(\mathbb{E}(\mathbf{1}_A\mathbf{1}_B \mid \mathcal{G}) \mid \mathcal{G}'\big)\\ &= \mathbb{E}\big(\mathbf{1}_A\mathbb{E}(\mathbf{1}_B \mid \mathcal{G}) \mid \mathcal{G}'\big)\\ &= \mathbb{E}\big(\mathbf{1}_A\mathbb{P}(B) \mid \mathcal{G}'\big)\\ &= \mathbb{E}(\mathbf{1}_A \mid \mathcal{G}')\mathbb{P}(B), \qquad (2.4)\end{aligned}$$

where the second line holds because $A \in \mathcal{G}$ and the third line holds because B is independent of $\mathcal{G}$. Finally, note that since B is independent of $\mathcal{G}'$ we have $\mathbb{E}(\mathbf{1}_B \mid \mathcal{G}') = \mathbb{P}(B)$. Combining this with (2.4) yields (2.3) and verifies the proposition. □

We will also need the following definition.

Definition. For $p, q \geqslant 0$, define

$$d(p,q) = \tfrac{1}{2}p\log p + \tfrac{1}{2}q\log q - \frac{p+q}{2}\log\left(\frac{p+q}{2}\right).$$

Finally, we will need the following proposition, which is easily verified using calculus.

Proposition 2.3 ([9]). *Fix $p \geqslant 0$. The function $d(p,\cdot)$ is convex.* □

Observe that $d(p,q) \geqslant 0$, with equality if and only if $p = q$ by the strict convexity of the function $x \to x\log x$. If $p = \{p_i : i \in V\}$ and $q = \{q_i : i \in V\}$ are both probability distributions on V, then we can define the 'distance' $d(p,q)$ between p and q, by $d(p,q) = \sum_{i\in V} d(p_i, q_i)$. (We use the term *distance* loosely and do not claim that $d(\cdot,\cdot)$ satisfies the triangle inequality.) Note that $d(p,q)$ is the difference between the average of the entropies of p and q and the entropy of the average (*i.e.*, an even mixture) of p and q.

We will use the following projection lemma.

Lemma 2.4 ([9]). *Let X and Y be random variables with distributions p and q, respectively. Fix a function g and let P and Q be the distributions of $g(X)$ and $g(Y)$, respectively. Then $d(p,q) \geqslant d(P,Q)$.* □

Let $\mathcal{U}$ denote the uniform distribution on V. Note that if μ is an arbitrary distribution on V, then $\mathrm{ENT}(\mu)$ and $d(\mu,\mathcal{U})$ are both notions of a distance from μ to $\mathcal{U}$. The following lemma relates the two.

Lemma 2.5 ([9]). *For any distribution μ on V, we have*

$$d(\mu,\mathcal{U}) \geqslant \frac{c}{\log|V|}\mathrm{ENT}(\mu),$$

for a universal constant $c > 0$. □

https://doi.org/10.1017/S0963548312000478 Published online by Cambridge University Press

3. Thorp shuffle

Recall that the Thorp shuffle has the following description. Assume that the number of cards, n, is even. Cut the deck into two equal piles, then interleave them as follows. Choose the first card from the left pile or from the right pile according to the outcome of a fair coin flip. Then choose from the other pile. Continue this way, flipping an independent coin for each pair, until both piles are empty. This is one time step.

We will actually work with the time reversal of the Thorp shuffle, which has the same mixing time (since the Thorp shuffle is a random walk on a group: see [15]). In this paper, we assume that $n = 2^d$ is a power of two. By writing the position of each card, from the bottom card (0) to the top card ($2^d - 1$), in binary, we can view the positions as elements of $\{0,1\}^d$. The reverse Thorp shuffle can then be constructed in the following way (see, for example, [10]). Let

$$Z = \{Z(l,t) : l \in \{0,1\}^{d-1}, t \in \{0,1,\dots\}\}$$

be a collection of i.i.d. Bernoulli(1/2) random variables, hereafter referred to as *coin flips*. Note that $x \in \{0,1\}^d$ can be written as $x = (L(x), R(x))$, where $L(x)$ and $R(x)$ are the leftmost $d-1$ bits and rightmost bit, respectively, of x. The reverse Thorp shuffle starts with $X_0 = \mathrm{id}$ and has the following transition rule. If the state at time t is $X_t = \pi$, then the next state $X_{t+1} = v \circ \pi$, where v is the permutation that sends

$$(L, R) \to (R \oplus Z(L,t), L).$$

Thus the position of each card undergoes a 'cyclic bit shift right' followed by a possible flipping of the first bit. For example, in a step of the reverse Thorp shuffle with 2^4 cards, the card in position 1001 moves to either 1100 or 0100, with probability $\frac{1}{2}$ each.

We are now ready to state the main technical result of this paper.

Lemma 3.1. *Let X_t be the reverse Thorp shuffle with 2^d cards. There is a universal constant c such that if μ is a random permutation which is independent of $\{X_t\}$, then*

$$\mathrm{ENT}(X_d \circ \mu) \leqslant (1 - c/d)\mathrm{ENT}(\mu).$$

Before proving this lemma we show how it gives the desired mixing time bound.

Theorem 3.2. *The mixing time of the reverse Thorp shuffle with 2^d cards is $O(d^3)$.*

Proof. Repeated applications of Lemma 3.1 give

$$\begin{aligned}\mathrm{ENT}(X_{kd}) &\leqslant (1 - c/d)^k \mathrm{ENT}(\mathrm{id}) \\ &\leqslant e^{-ck/d} \log(2^d!) \\ &\leqslant e^{-ck/d} d 2^d.\end{aligned}$$

Now let α be large enough so that $d\left(\frac{2}{e^\alpha}\right)^d \leqslant 1/8$ for all d. Then if $k = \lceil \alpha d^2/c \rceil$ we have

$$\mathrm{ENT}(X_{kd}) \leqslant e^{-ck/d} d 2^d \leqslant \tfrac{1}{8},$$

and hence $\|X_{kd} - \mathcal{U}\| \leqslant \frac{1}{4}$ by (2.2). The theorem follows since k is $O(d^2)$. □

https://doi.org/10.1017/S0963548312000478 Published online by Cambridge University Press

We will now give an informal description of our strategy for proving Lemma 3.1. Define the sigma field $\mathcal{F} = \sigma(Z(l,t) : l \in \{0,1\}^{d-1}, 0 \leqslant t < d\}$, that is, $\mathcal{F}$ is the sigma field generated by the coin flips used to make X_d. Since X_d is $\mathcal{F}$-measurable and μ is independent of $\mathcal{F}$, we have

$$\mathrm{ENT}(X_d \circ \mu \,|\, \mathcal{F}) = \mathrm{ENT}(\mu \,|\, \mathcal{F}) = \mathrm{ENT}(\mu). \tag{3.1}$$

Now suppose that we initially know the outcomes of all the coin flips, but then 'cover them up' one by one until none are known. Each cover-up changes the relative entropy of the conditional distribution of $X_d \circ \mu$, and by equation (3.1), the sum of the expected changes in relative entropy is $\mathrm{ENT}(X_d \circ \mu) - \mathrm{ENT}(\mu)$. Our aim will be to show that for each value of j there is a special coin flip such that the expected change of relative entropy when this coin flip is covered up is $\leqslant -\frac{c}{d}\mathrm{ENT}(\mu, j)$, for a universal constant c. Summing up these expected changes gives

$$-\frac{c}{d}\sum_j \mathrm{ENT}(\mu, j) = -\frac{c}{d}\mathrm{ENT}(\mu),$$

as desired.

We now give the formal proof.

Proof of Lemma 3.1. Suppose that the number of cards is $n = 2^d$. For integers j with $1 \leqslant j < n$, define $R_j = \max\{i : j_i \neq 0\}$, where $j_{d-1}j_{d-2}\cdots j_0$ is the binary representation of j. Thus $R_j + 1$ is the number of bits required to write j in binary (if leading zeros are omitted). Let T be a geometric(1/2) random variable on $\{0, 1, \dots\}$, and define $T_j = R_j - T$. Note that T_j is non-decreasing in j. Let $\tilde{Z}^j$ be obtained from Z by flipping the value of $Z(L(X_{T_j}(j)), T_j)$ if $T_j \geqslant 0$. More precisely, for $l \in \{0,1\}^{d-1}$ and $0 \leqslant t < d$, define

$$\tilde{Z}^j(l,t) = \begin{cases} 1 - Z(l,t) & \text{if } T_j = t \text{ and } L(X_{T_j}(j)) = l, \\ Z(l,t) & \text{otherwise.} \end{cases}$$

Note that if $T_j < 0$ then $\tilde{Z}^j$ is the same as Z. Let $\{\tilde{X}^j_t : t \geqslant 0\}$ be the reverse Thorp shuffle process defined by using $\tilde{Z}^j$ instead of Z. Define $X(j) = (X_1(j), \dots, X_d(j))$, with a similar definition for $\tilde{X}^j(j)$. For j with $1 \leqslant j \leqslant n$, define

$$\begin{aligned} \mathcal{F}_j &= \sigma(X(k), \tilde{X}^k(k) : k \geqslant j) \\ &= \sigma(X_t(k), \tilde{X}^k_t(k) : k \geqslant j, 0 \leqslant t \leqslant d). \end{aligned}$$

We think of μ as a random initial ordering of a deck of cards, and X_d as d reverse Thorp shuffles. We write 'card j' for the card initially in position j. (So if the initial ordering is μ, then card j is the card with the label $\mu^{-1}(j)$.) Thus, $X(j)$ is the trajectory of card j, and $\mathcal{F}_j$ is the sigma field that corresponds to knowing, for each $k \geqslant j$,

(1) the trajectory of card k, and
(2) what its trajectory would have been if the coin flip at time T_k had been different (see the schematic diagram in Figure 1).

https://doi.org/10.1017/S0963548312000478 Published online by Cambridge University Press

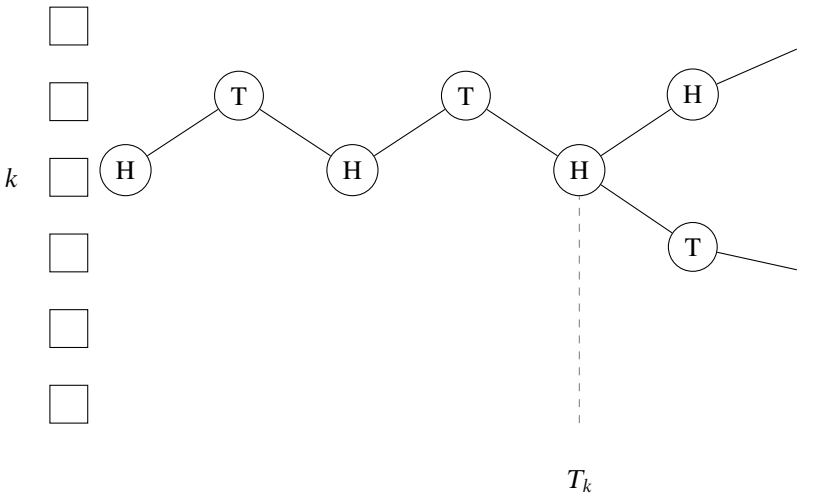


Figure 1. $\mathcal{F}_j$ corresponds to knowing, for all $k \geqslant j$, the coin flips pictured above.

Since $\mathcal{F}_n$ is trivial and X_d is $\mathcal{F}_1$-measurable (because the trajectories for cards $k \geqslant 1$ determine that of card 0), we have

$$\begin{aligned}\mathrm{ENT}(X_d \circ \mu) - \mathrm{ENT}(\mu) &= \mathrm{ENT}(X_d \circ \mu \,|\, \mathcal{F}_n) - \mathrm{ENT}(X_d \circ \mu \,|\, \mathcal{F}_1) \\ &= \sum_{j=1}^{n-1} \big(\mathrm{ENT}(X_d \circ \mu \,|\, \mathcal{F}_{j+1}) - \mathrm{ENT}(X_d \circ \mu \,|\, \mathcal{F}_j)\big). \end{aligned} \tag{3.2}$$

We claim that for all j with $1 \leqslant j < n$ we have

$$\mathbb{E}\big(\mathrm{ENT}(X_d \circ \mu \,|\, \mathcal{F}_{j+1}) - \mathrm{ENT}(X_d \circ \mu \,|\, \mathcal{F}_j)\big) \leqslant -\mathrm{ENT}(\mu, j)\frac{c}{d}, \tag{3.3}$$

for a universal constant $c > 0$. Combining this with equation (3.2) gives

$$\mathrm{ENT}(X_d \circ \mu) - \mathrm{ENT}(\mu) \leqslant -\frac{c}{d}\sum_{j=1}^{n-1} \mathrm{ENT}(\mu, j) = -\frac{c}{d}\mathrm{ENT}(\mu),$$

which proves the lemma. It remains to verify equation (3.3).

Say that cards i and j are *adjacent at time t* if $L(X_t(i)) = L(X_t(j))$, that is, if they use the 'same coin' at time t. For j with $1 \leqslant j < n$, define

$$a(j) = \begin{cases} \text{the card adjacent to } j \text{ at time } T_j & \text{if } T_j \geqslant 0, \\ j & \text{otherwise,} \end{cases}$$

For j with $1 \leqslant j < n$, define $\widehat{\mathcal{F}}_j = \sigma(\mathcal{F}_{j+1}, \{X(j), \tilde{X}^j(j)\})$, that is, the sigma field generated by $\mathcal{F}_{j+1}$ and the random *unordered set* $\{X(j), \tilde{X}^j(j)\}$. See the schematic diagram in Figure 2.

The sigma field $\widehat{\mathcal{F}}_j$ is almost the same as $\mathcal{F}_j$. However, we will see later that the coin flip $Z(L(X_{T_j}(j)), T_j)$ is not $\widehat{\mathcal{F}}_j$-measurable because if $a(j) < j$ then flipping the value of $Z(L(X_{T_j}(j)), T_j)$ does not change $\{X(j), \tilde{X}^j(j)\}$ or $(X(k), \tilde{X}^k(k))$ for $k > j$. (On the other hand,

https://doi.org/10.1017/S0963548312000478 Published online by Cambridge University Press

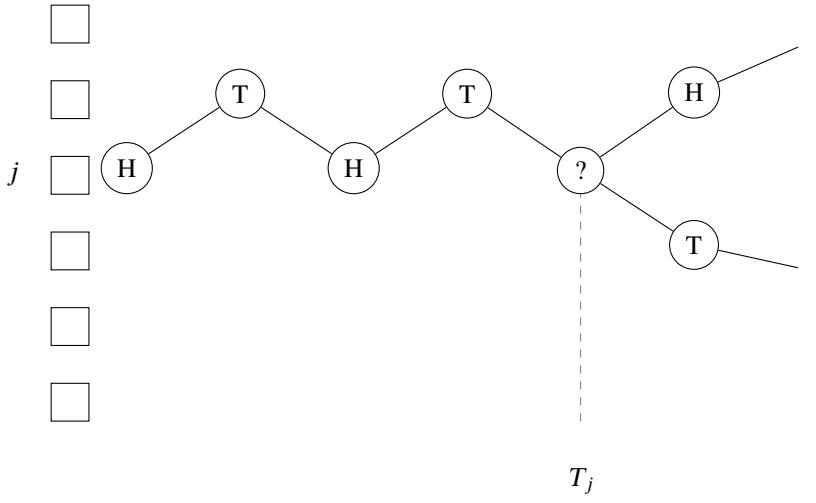


Figure 2. $\widehat{\mathcal{F}}_j$ is like $\mathcal{F}_j$, except that the coin flip for card j at time T_j might not be known.

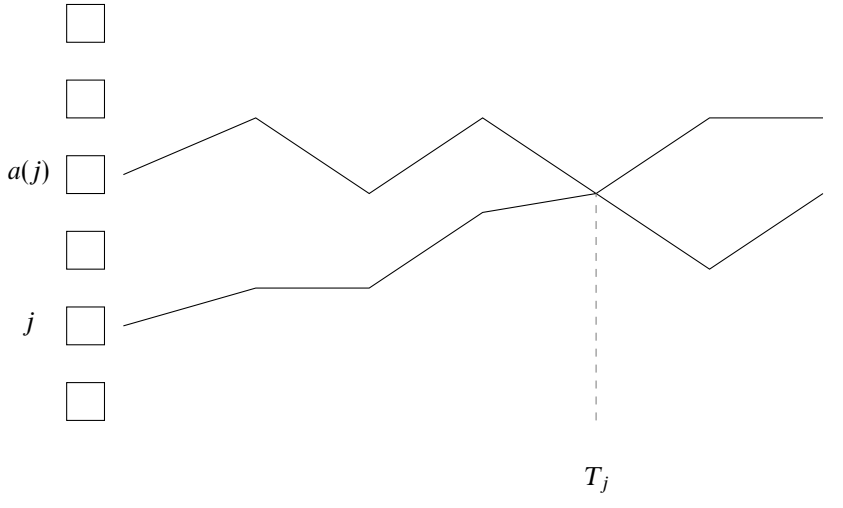


Figure 3. If $a(a) > j$, then the outcome of this flip can be inferred from the coin flips corresponding to card $a(j)$.

if $a(j) > j$ then the outcome of this flip can be inferred from the coin flips corresponding to card $a(j)$: see Figure 3.)

Since $\widehat{\mathcal{F}}_j \supset \mathcal{F}_{j+1}$, we have

$$\mathbb{E}(\mathrm{ENT}(X_d \circ \mu \,|\, \mathcal{F}_{j+1})) \leqslant \mathbb{E}(\mathrm{ENT}(X_d \circ \mu \,|\, \widehat{\mathcal{F}}_j))$$

by convexity of $p \to \mathrm{ENT}(p)$ and Jensen's inequality. Thus, it is enough to show that for all $j \geqslant 1$ we have

$$\mathbb{E}(\mathrm{ENT}(X_d \circ \mu \,|\, \widehat{\mathcal{F}}_j)) - \mathbb{E}(\mathrm{ENT}(X_d \circ \mu \,|\, \mathcal{F}_j)) \leqslant -\frac{c}{d}\mathrm{ENT}(\mu, j), \tag{3.4}$$

https://doi.org/10.1017/S0963548312000478 Published online by Cambridge University Press

for a universal constant c. (Thus the coin flip for card j at time T_j is the 'special' coin flip referred to earlier; revealing it changes the expected entropy by at least $\frac{c}{d}\mathrm{ENT}(\mu, j)$.)

Fix j with $1 \leqslant j < n$. Define $\mathcal{G}_{j+1} = \sigma(\mu^{-1}(j+1), \ldots, \mu^{-1}(n-1))$. Since the random variable $(\mu^{-1}(j+1), \ldots, \mu^{-1}(n-1))$ is $\sigma(X_d \circ \mu, \mathcal{F}_j)$-measurable, a straightforward calculation using the entropy chain rule shows that

$$\mathrm{ENT}(X_d \circ \mu \,|\, \mathcal{F}_j) = \mathrm{ENT}(\mu^{-1}(j+1), \ldots, \mu^{-1}(n-1) \,|\, \mathcal{F}_j) + \mathbb{E}(\mathrm{ENT}(X_d \circ \mu \,|\, \mathcal{F}_j, \mathcal{G}_{j+1})), \tag{3.5}$$

and a similar equation holds with $\widehat{\mathcal{F}}_j$ replacing $\mathcal{F}_j$. The random permutation μ is independent of $\{X_t : t \geqslant 0\}$ and T, hence independent of $\mathcal{F}_j$ and $\widehat{\mathcal{F}}_j$. Thus the first term on the right-hand side of equation (3.5) does not change if we replace $\mathcal{F}_j$ by $\widehat{\mathcal{F}}_j$. It follows that

$$\begin{aligned} &\mathrm{ENT}(X_d \circ \mu \,|\, \widehat{\mathcal{F}}_j) - \mathrm{ENT}(X_d \circ \mu \,|\, \mathcal{F}_j) \\ &\quad = \mathbb{E}\big(\mathrm{ENT}(X_d \circ \mu \,|\, \widehat{\mathcal{F}}_j, \mathcal{G}_{j+1}) - \mathrm{ENT}(X_d \circ \mu \,|\, \mathcal{F}_j, \mathcal{G}_{j+1})\big). \end{aligned} \tag{3.6}$$

Recall that for j with $1 \leqslant j < n$, we denote by $a(j)$ the card adjacent to j at time T_j.

Lemma 3.3. *For j with $1 \leqslant j < n$, the conditional distribution of $a(j)$ given $a(j) < j$ is uniform over k with $0 \leqslant k < j$. Furthermore, the probability that $a(j) < j$ is at least $1/2$.*

Proof. Fix j with $1 \leqslant j < n$. Recall that $R_j + 1$ is the number of bits required to write j in binary if leading zeros are omitted. We claim that for every k with $k < j$ we have

$$\mathbb{P}(a(j) = k) = \left(\frac{1}{2}\right)^{R_j+1}.$$

To see this, suppose that $k < j$, let $k_{d-1}k_{d-2}\cdots k_0$ and $j_{d-1}j_{d-2}\cdots j_0$ be the binary representations of k and j, respectively, and let $D(k,j) = \max\{l : k_l \neq j_l\}$. Note that $D(k,j)$ is the only time before time d that k and j can be adjacent. Thus $a(j) = k$ only if $T_j = D(k,j)$. Since the probability that k and j are adjacent after $D(k,j)$ steps is $\left(\frac{1}{2}\right)^{D(k,j)}$, it follows that

$$\begin{aligned} \mathbb{P}(a(j) = k) &= \left(\frac{1}{2}\right)^{D(k,j)} \mathbb{P}(T_j = D(k,j)) \\ &= \left(\frac{1}{2}\right)^{D(k,j)} \mathbb{P}(R_j - T_j = R_j - D(k,j)) \\ &= \left(\frac{1}{2}\right)^{D(k,j)} \left(\frac{1}{2}\right)^{R_j - D(k,j)+1} \\ &= \left(\frac{1}{2}\right)^{R_j+1}, \end{aligned} \tag{3.7}$$

where the second line follows from the fact that $R_j - T_j = T$ has the geometric($\frac{1}{2}$) distribution on $\{0, 1, 2, \ldots\}$. Since the probability in (3.7) does not depend on k, this implies that the conditional distribution of $a(j)$ given $a(j) < j$ is uniform over k with $k < j$. To see the second part of

https://doi.org/10.1017/S0963548312000478 Published online by Cambridge University Press

the lemma, note that $j \geqslant 2^{R_j}$ and hence

$$\begin{aligned}\mathbb{P}(a(j) < j) &= \sum_{k<j} \mathbb{P}(a(j) = k) \\ &= j\left(\frac{1}{2}\right)^{R_j+1} \\ &\geqslant 2^{R_j}\left(\frac{1}{2}\right)^{R_j+1} \qquad (3.8) \\ &= \frac{1}{2}. \qquad \square\end{aligned}$$

For k with $1 \leqslant k \leqslant n$, let S^k denote the set of pairs (l,t), with $l \in \{0,1\}^{d-1}$ and $0 \leqslant t < d$, such that either $L(X_t(k)) = l$ or $L(\tilde{X}_t^k(k)) = l$, and let $S = \cup_{k>j} S^k$. Thus S indexes the coin flips used to generate $X(k)$ and $\tilde{X}^k(k)$ for $k > j$.

Note that $\mathcal{F}_{j+1} = \sigma(S, Z(l,t) : (l,t) \in S)$, that is, $\mathcal{F}_{j+1}$ is the sigma field generated by the random set S and $Z(l,t)$ evaluated for $(l,t) \in S$. For a sigma field $\mathcal{F}$ and random variable W, we write $\mathcal{L}(W \mid \mathcal{F})$ for the conditional distribution of W given $\mathcal{F}$. Note that $\mathcal{L}(X_d \circ \mu \mid \widehat{\mathcal{F}}_j) = \mathcal{L}(X_d \circ \mu \mid \mathcal{F}_j)$ whenever $(L(X_{T_j}(j)), T_j) \in S$, and hence in this case the expression on the left-hand side of (3.6) is 0.

However, we claim that if $a(j) < j$ then $(L(X_{T_j}(j)), T_j) \notin S$. (That is, the coin flip for card j at time T_j is a 'fresh' coin flip.) To see this, suppose that $(L(X_{T_j}(j)), T_j) \in S$. Then $(L(X_{T_j}(j)), T_j) \in S^k$ for some $k > j$, so at least one of the following holds:

(1) $L(X_{T_j}(k)) = L(X_{T_j}(j))$,
(2) $L(\tilde{X}^k_{T_j}(k)) = L(X_{T_j}(j))$

But since $X_t(k) = \tilde{X}_t^k(k)$ for $t = 0, \ldots, T_k$ and $T_j \leqslant T_k$, we have $X_{T_j}(k) = \tilde{X}^k_{T_j}(k)$. Thus, if condition (2) above holds then condition (1) must also hold. But if condition (1) holds, then $a(j) = k > j$. This verifies the claim.

Let A be the event that $a(j) < j$. Note that $A \in \widehat{\mathcal{F}}_j$. Define

$$Z_j = \begin{cases} Z(L(X_{T_j}(j)), T_j) & \text{if } T_j \geqslant 0, \\ 0 & \text{otherwise.} \end{cases}$$

Since $\sigma(X_d, T)$ and $\mathcal{G}_{j+1}$ are independent and $\widehat{\mathcal{F}}_j \subset \sigma(X_d, T)$, Proposition 2.2 implies that $\sigma(X_d, T)$ and $\mathcal{G}_{j+1}$ are conditionally independent given $\widehat{\mathcal{F}}_j$. Furthermore, Z_j is measurable with respect to $\sigma(X_d, T)$, and hence conditionally independent of $\mathcal{G}_{j+1}$ given $\widehat{\mathcal{F}}_j$. Combining this with the claim gives that on the event A, the conditional distribution of $Z(L(X_{T_j}(j)), T_j)$ given $\sigma(\widehat{\mathcal{F}}_j, \mathcal{G}_{j+1})$ is uniform over $\{0,1\}$. It follows that

$$\mathcal{L}(X_d \circ \mu \mid \widehat{\mathcal{F}}_j, \mathcal{G}_{j+1})\mathbf{1}_A = \tfrac{1}{2}\mathcal{L}(X_d \circ \mu \mid \mathcal{F}_j, \mathcal{G}_{j+1})\mathbf{1}_A + \tfrac{1}{2}\mathcal{L}(\tilde{X}_d^j \circ \mu \mid \mathcal{F}_j, \mathcal{G}_{j+1})\mathbf{1}_A. \qquad (3.9)$$

(Note that there is no hat on the $\mathcal{F}_j$ on the right-hand side of this equation.) For a sigma field $\mathcal{F}$ and random variables W_1, W_2, we let $d(W_1, W_2 \mid \mathcal{F})$ denote the random variable

$$d(\mathcal{L}(W_1 \mid \mathcal{F}), \mathcal{L}(W_2 \mid \mathcal{F})).$$

https://doi.org/10.1017/S0963548312000478 Published online by Cambridge University Press

Equation (3.9) implies that

$$\begin{aligned}
&\mathrm{ENT}(X_d \circ \mu \,|\, \widehat{\mathcal{F}}_j, \mathcal{G}_{j+1})\mathbf{1}_A \\
&\quad = \big(\tfrac{1}{2}\mathrm{ENT}(X_d \circ \mu \,|\, \mathcal{F}_j, \mathcal{G}_{j+1}) + \tfrac{1}{2}\mathrm{ENT}(\tilde{X}_d^j \circ \mu \,|\, \mathcal{F}_j, \mathcal{G}_{j+1}) \\
&\qquad - d(X_d \circ \mu, \tilde{X}_d^j \circ \mu \,|\, \mathcal{F}_j, \mathcal{G}_{j+1})\big)\mathbf{1}_A \\
&\quad = \big(\mathrm{ENT}(X_d \circ \mu \,|\, \mathcal{F}_j, \mathcal{G}_{j+1}) - d(X_d \circ \mu, \tilde{X}_d^j \circ \mu \,|\, \mathcal{F}_j, \mathcal{G}_{j+1})\big)\mathbf{1}_A, \qquad (3.10)
\end{aligned}$$

where the final equality follows from the fact that

$$\mathrm{ENT}(X_d \circ \mu \,|\, \mathcal{F}_j, \mathcal{G}_{j+1}) = \mathrm{ENT}(\tilde{X}_d^j \circ \mu \,|\, \mathcal{F}_j, \mathcal{G}_{j+1}),$$

which holds because $\tilde{X}_d^j \circ \mu$ is obtained from $X_d \circ \mu$ by swapping the cards in the ($\mathcal{F}_j$-measurable) positions $X_d(j)$ and $\tilde{X}_d^j(j)$. By the projection lemma, the distance between the conditional distributions of $X_d \circ \mu$ and $\tilde{X}_d^j \circ \mu$ can be bounded by the corresponding distance when we evaluate each permutation's inverse at a given point. Note that $(X_d \circ \mu)^{-1}(X_d(j)) = \mu^{-1}(j)$, and since $(\tilde{X}_d^j)^{-1}(X_d(j)) = a(j)$, we have $(\tilde{X}_d^j \circ \mu)^{-1}(X_d(j)) = \mu^{-1}(a(j))$. Thus the projection lemma gives

$$d\big(X_d \circ \mu, \tilde{X}^j \circ \mu \,|\, \mathcal{F}_j, \mathcal{G}_{j+1}\big)\mathbf{1}_A \geqslant d\big(\mu^{-1}(j), \mu^{-1}(a(j)) \,|\, \mathcal{F}_j, \mathcal{G}_{j+1}\big)\mathbf{1}_A. \qquad (3.11)$$

We claim that

$$\mathrm{ENT}(X_d \circ \mu \,|\, \widehat{\mathcal{F}}_j, \mathcal{G}_{j+1}) \leqslant \mathrm{ENT}(X_d \circ \mu \,|\, \mathcal{F}_j, \mathcal{G}_{j+1})$$

everywhere. To see this, note that $\mathcal{L}(X_d \circ \mu \,|\, \widehat{\mathcal{F}}_j, \mathcal{G}_{j+1})$ is always either $\mathcal{L}(X_d \circ \mu \,|\, \mathcal{F}_j, \mathcal{G}_{j+1})$ (when $(L(X_{T_j}(j)), T_j) \in S$) or else it is $\frac{1}{2}\mathcal{L}(X_d \circ \mu \,|\, \mathcal{F}_j, \mathcal{G}_{j+1}) + \frac{1}{2}\mathcal{L}(\tilde{X}_d^j \circ \mu \,|\, \mathcal{F}_j, \mathcal{G}_{j+1})$ (otherwise). In the second case, both conditional distributions have the same entropy, so the claim follows from the fact that $p \to \mathrm{ENT}(p)$ is convex and Jensen's inequality. Since

$$\mathrm{ENT}(X_d \circ \mu \,|\, \widehat{\mathcal{F}}_j, \mathcal{G}_{j+1}) \leqslant \mathrm{ENT}(X_d \circ \mu \,|\, \mathcal{F}_j, \mathcal{G}_{j+1})$$

everywhere, equations (3.6), (3.10) and (3.11) imply that in order to verify (3.4), it is enough to show that

$$\mathbb{E}\big(d\big(\mu^{-1}(j), \mu^{-1}(a(j)) \,|\, \mathcal{F}_j, \mathcal{G}_{j+1}\big)\mathbf{1}_A\big) \geqslant \frac{c}{\log n}\mathrm{ENT}(\mu, j), \qquad (3.12)$$

for a universal constant c.

Let $\tilde{U}$ be a uniform($\{0, 1, \ldots, j-1\}$) random variable which is independent of everything else. For l with $0 \leqslant l < n$, define the random variable $W_l = \mathbb{P}(\mu^{-1}(a(j)) = l \,|\, \mathcal{F}_j, \mathcal{G}_{j+1})\mathbf{1}_A$. We claim that for l with $0 \leqslant l < n$ we have

$$\mathbb{E}(W_l \,|\, \mathcal{G}_{j+1}) = \mathbb{P}(A)\mathbb{P}(\mu^{-1}(\tilde{U}) = l \,|\, \mathcal{G}_{j+1}). \qquad (3.13)$$

Before proving this, let us first show that the result follows from it. Note that equation (3.13) can be restated as follows. If $\nu = \mathcal{L}(\mu^{-1}(a(j)) \,|\, \mathcal{F}_j, \mathcal{G}_{j+1})$, then

$$\mathbb{E}(\nu \,|\, \mathcal{G}_{j+1}, A) = \mathcal{L}(\mu^{-1}(\tilde{U}) \,|\, \mathcal{G}_{j+1}). \qquad (3.14)$$

Now consider the quantity

$$\mathbb{E}\big(d\big(\mu^{-1}(j), \mu^{-1}(a(j)) \,|\, \mathcal{F}_j, \mathcal{G}_{j+1}\big)\mathbf{1}_A \,|\, \mathcal{G}_{j+1}\big). \qquad (3.15)$$

https://doi.org/10.1017/S0963548312000478 Published online by Cambridge University Press

Since $\mathcal{F}_j$ and μ are independent and $\mathcal{G}_{j+1} \subset \sigma(\mu)$, Proposition 2.2 implies that the conditional distribution of $\mu^{-1}(j)$ given $\mathcal{G}_{j+1}$ is independent of $\mathcal{F}_j$. That is, the first argument of $d(\cdot, \cdot)$ in the quantity (3.15) is conditionally independent of $\mathcal{F}_j$ given $\mathcal{G}_{j+1}$. Furthermore, the conditional expectation of the second argument given $\mathcal{G}_{j+1}$ and A is $\mathcal{L}(\mu^{-1}(\tilde{U}) \,|\, \mathcal{G}_{j+1})$. Note also that Proposition 2.3 implies that if W_1 and W_2 are random variables and $\mathcal{H}$ is a sigma field that is independent of W_1, then $\mathbb{E}(d(W_1, W_2 \,|\, \mathcal{H})) \geqslant d(W_1, W_2)$. It follows that

$$\begin{aligned}\mathbb{E}\big(d\big(\mu^{-1}(j), \mu^{-1}(a(j)) \,|\, \mathcal{F}_j, \mathcal{G}_{j+1}\big)\mathbf{1}_A \,|\, \mathcal{G}_{j+1}\big) &\geqslant d\big(\mu^{-1}(j), \mu^{-1}(\tilde{U}) \,|\, \mathcal{G}_{j+1}\big)\mathbb{P}(A \,|\, \mathcal{G}_{j+1}) \\ &\geqslant \tfrac{1}{2} d\big(\mu^{-1}(j), \mu^{-1}(\tilde{U}) \,|\, \mathcal{G}_{j+1}\big), \qquad (3.16)\end{aligned}$$

where the last line holds because A is independent of $\mathcal{G}_{j+1}$ and by Lemma 3.3. Let U be an independent uniform$(\{0, 1, \ldots, j\})$ random variable. Note that $\mathcal{L}(\mu^{-1}(U) \,|\, \mathcal{G}_{j+1})$ can be written as a mixture as follows:

$$\mathcal{L}(\mu^{-1}(U) \,|\, \mathcal{G}_{j+1}) = \tfrac{j}{j+1}\mathcal{L}(\mu^{-1}(\tilde{U}) \,|\, \mathcal{G}_{j+1}) + \tfrac{1}{j+1}\mathcal{L}(\mu^{-1}(j) \,|\, \mathcal{G}_{j+1}).$$

Thus convexity implies

$$\begin{aligned}d\big(\mu^{-1}(j), \mu^{-1}(U) \,|\, \mathcal{G}_{j+1}\big) &\leqslant \tfrac{j}{j+1} d(\mu^{-1}(j), \mu^{-1}(\tilde{U}) \,|\, \mathcal{G}_{j+1}) + \tfrac{1}{j+1} d(\mu^{-1}(j), \mu^{-1}(j) \,|\, \mathcal{G}_{j+1}) \\ &= \tfrac{j}{j+1} d(\mu^{-1}(j), \mu^{-1}(\tilde{U}) \,|\, \mathcal{G}_{j+1}).\end{aligned}$$

Since $\frac{j}{j+1} < 1$, we have

$$d\big(\mu^{-1}(j), \mu^{-1}(\tilde{U}) \,|\, \mathcal{G}_{j+1}\big) \geqslant d\big(\mu^{-1}(j), \mu^{-1}(U) \,|\, \mathcal{G}_{j+1}\big),$$

and hence the quantity (3.16) is at least

$$\tfrac{1}{2} d\big(\mu^{-1}(j), \mu^{-1}(U) \,|\, \mathcal{G}_{j+1}\big) \;\geqslant \frac{c}{2 \log n} \mathrm{ENT}(\mu^{-1}(j) \,|\, \mathcal{G}_{j+1})$$

for a universal constant c by Lemma 2.5. Incorporating a factor of $\frac{1}{2}$ into the constant c and then taking expectations gives

$$\mathbb{E}\big(d\big(\mu^{-1}(j), \mu^{-1}(a(j)) \,|\, \mathcal{F}_j, \mathcal{G}_{j+1}\big)\mathbf{1}_A\big) \geqslant \frac{c}{\log n} \mathrm{ENT}(\mu, j),$$

which verifies (3.12) and completes the proof, assuming (3.13). It remains to verify (3.13). Recall that for l with $0 \leqslant l < n$, the random variable W_l is defined by

$$W_l = \mathbb{P}(\mu^{-1}(a(j)) = l \,|\, \mathcal{F}_j, \mathcal{G}_{j+1})\mathbf{1}_A.$$

Note that

$$W_l = \sum_{k<j} \mathbb{P}(\mu(l) = k, a(j) = k \,|\, \mathcal{F}_j, \mathcal{G}_{j+1}). \qquad (3.17)$$

Proposition 2.2 implies that since $\sigma(X_d, T)$ and μ are independent, and $\mathcal{G}_{j+1} \subset \sigma(\mu)$, the random variable $\mu(l)$ is conditionally independent of $\sigma(X_d, T)$ given $\mathcal{G}_{j+1}$. Furthermore, since both $\sigma(a(j))$ and $\mathcal{F}_j$ are contained in $\sigma(X_d, T)$, we have $\sigma(\mathcal{F}_j, a(j)) \subset \sigma(X_d, T)$. Hence the quantity

https://doi.org/10.1017/S0963548312000478 Published online by Cambridge University Press

(3.17) is

$$\begin{aligned}\sum_{k<j} \mathbb{P}(\mu(l) = k \mid \mathcal{F}_j, \mathcal{G}_{j+1}, a(j) = k)\mathbb{P}(a(j) = k \mid \mathcal{F}_j, \mathcal{G}_{j+1}) \\ = \sum_{k<j} \mathbb{P}(\mu(l) = k \mid \mathcal{G}_{j+1})\mathbb{P}(a(j) = k \mid \mathcal{F}_j, \mathcal{G}_{j+1}) \\ = \sum_{k<j} \mathbb{P}(\mu(l) = k \mid \mathcal{G}_{j+1})\mathbb{P}(a(j) = k \mid \mathcal{F}_j), \end{aligned} \tag{3.18}$$

where the first line follows from conditioning, the second line follows from the fact that $\mu(l)$ is conditionally independent of $\sigma(\mathcal{F}_j, a(j))$ given $\mathcal{G}_{j+1}$, and the third line follows from Proposition 2.2, since $a(j)$ is measurable with respect to $\sigma(X_d, T)$, which is independent of $\mathcal{G}_{j+1}$, and $\mathcal{F}_j \subset \sigma(X_d, T)$. Taking the conditional expectation given $\mathcal{G}_{j+1}$ of the left-hand side of (3.17) and the right-hand side of (3.18) gives

$$\begin{aligned}\mathbb{E}(W_l \mid \mathcal{G}_{j+1}) &= \sum_{k<j} \mathbb{P}(\mu(l) = k \mid \mathcal{G}_{j+1})\mathbb{E}\big(\mathbb{P}(a(j) = k \mid \mathcal{F}_j) \mid \mathcal{G}_{j+1}\big) \\ &= \sum_{k<j} \mathbb{P}(\mu(l) = k \mid \mathcal{G}_{j+1})\mathbb{P}(a(j) = k) \\ &= \mathbb{P}(A) \cdot \frac{1}{j} \sum_{k<j} \mathbb{P}(\mu(l) = k \mid \mathcal{G}_{j+1}) \\ &= \mathbb{P}(A)\mathbb{P}(\mu^{-1}(\tilde{U}) = l \mid \mathcal{G}_{j+1}), \end{aligned} \tag{3.19}$$

where the second line holds because $\mathcal{F}_j$ is independent of $\mathcal{G}_{j+1}$, the third line follows from Lemma 3.3, and the last line holds because the sum in (3.19) is $\mathbb{P}(\mu(l) < j \mid \mathcal{G}_{j+1})$. This verifies (3.13). □

Acknowledgements

I am grateful to O. Blumberg, D. Romik and the referees for numerous helpful suggestions.

References

[1] Bellare, M., Rogaway, P. and Spies, T. (2010) The FFX mode of operation and format preserving encryption (draft 1.1). NIST submission.

[2] Cover, T. and Thomas, J. (1991) *Elements of Information Theory*, Wiley.

[3] Jonasson, J. (2011) Mixing time bounds for overlapping cycles shuffles. *Electron. J. Probab.* **16** 1281–1295.

[4] Luby. M. and Rackoff, C. (1988) How to construct pseudorandom permutations from pseudorandom functions. *SIAM J. Comput.* **17** 373–386.

[5] Maurer, U. (1993) A simplified and generalized treatment of Luby–Rackoff pseudorandom permutation generators. In *Advances in Cryptology: EUROCRYPT 1992*, Vol. 658 of *Lecture Notes in Computer Science*, Springer, pp. 239–255.

[6] Maurer, U. and Pietrzak, K. (2003) The security of many-round Luby–Rackoff pseudorandom permutations. In *Advances in Cryptology: EUROCRYPT 2003*, Vol. 2656 of *Lecture Notes in Computer Science*, Springer, pp. 544–561.

https://doi.org/10.1017/S0963548312000478 Published online by Cambridge University Press

[7] Montenegro, R. and Tetali, P. (2006) *Mathematical Aspects of Mixing Times in Markov Chains*, Foundations and Trends in Theoretical Computer Science, Now Publishers.
[8] Morris, B. (2008) The mixing time of the Thorp shuffle. In *STOC 2005*, *SIAM J. Comput.* **38** 484–504.
[9] Morris, B. (2009) Improved mixing time bounds for the Thorp shuffle and L-reversal chain. *Ann. Probab.* **37** 453–477.
[10] Morris, B., Rogaway, P. and Stegers, T. (2009) How to encipher messages on a small domain: Deterministic encrption and the Thorp shuffle. In *Advances in Cryptology: CRYPTO 2009*, Vol. 5677 of *Lecture Notes in Computer Science*, Springer, pp. 286–302.
[11] Naor, M. and Reingold, O. (1999) On the construction of pseudo-random permutations: Luby–Rackoff revisited. *J. Cryptology* **12** 29–66.
[12] Patarin, J. (1991) New results on pseodorandom permutation generators based on the DES scheme. In *Advances in Cryptology: CRYPTO 1991*, Vol. 576 of *Lecture Notes in Computer Science*, Springer, pp. 301–312.
[13] Patarin, J. (1998) About Feistel schemes with six (or more) rounds. In *Fast Software Encryption 1998*, Vol. 1372 of *Lecture Notes in Computer Science*, Springer, pp. 103–121.
[14] Patarin, J. (2004) Security of random Feistel schemes with 5 or more rounds. In *Advances in Cryptology: CRYPTO 2004*, Vol. 3152 of *Lecture Notes in Computer Science*, Springer, pp. 106–122.
[15] Saloff-Coste, L. (2004) Random walks on finite groups. In *Probability on Discrete Structures*, Vol. 110 of *Encyclopedia of Mathematical Sciences* (H. Kesten, ed.), Springer, pp. 263–346.
[16] Rudich, S. (1989) Limits on the provable consequences of one-way functions. PhD Thesis, UC Berkeley.
[17] Thorp, E. (1973) Nonrandom shuffling with applications to the game of Faro. *J. Amer. Statist. Assoc.* **68**, 842–847.

https://doi.org/10.1017/S0963548312000478 Published online by Cambridge University Press