Shuffling cards

Chapter 31

> *How often does one have to shuffle a deck of cards until it is random?*

The analysis of random processes is a familiar duty in life ("How long does it take to get to the airport during rush-hour?") as well as in mathematics. Of course, getting meaningful answers to such problems heavily depends on formulating meaningful questions. For the card shuffling problem, this means that we have

- to specify the size of the deck ($n = 52$ cards, say),
- to say how we shuffle (we'll analyze top-in-at-random shuffles first, and then the more realistic and effective riffle shuffles), and finally
- to explain what we mean by "is random" or "is close to random."

So our goal in this chapter is an analysis of the riffle shuffle, due to Edgar N. Gilbert and Claude Shannon (1955, unpublished) and Jim Reeds (1981, unpublished), following the statistician David Aldous and the former magician turned mathematician Persi Diaconis according to [1]. We will not reach the final precise result that 7 riffle shuffles *are* sufficient to get a deck of $n = 52$ cards very close to random, while 6 riffle shuffles do not suffice — but we will obtain an upper bound of 12, and we will see some extremely beautiful ideas on the way: the concepts of stopping rules and of "strong uniform time," the lemma that strong uniform time bounds the variation distance, Reeds' inversion lemma, and thus the interpretation of shuffling as "reversed sorting." In the end, everything will be reduced to two very classical combinatorial problems, namely the coupon collector and the birthday paradox. So let's start with these!

Persi Diaconis' business card as a magician. In a later interview he said: "If you say that you are a professor at Stanford people treat you respectfully. If you say that you invent magic tricks, they don't want to introduce you to their daughter."

The birthday paradox

Take n random people — the participants of a class or seminar, say. What is the probability that they all have different birthdays? With the usual simplifying assumptions (365 days a year, no seasonal effects, no twins present) the probability is

$$p(n) \quad = \quad \prod_{i=1}^{n-1} \Big(1 - \frac{i}{365}\Big),$$

M. Aigner, G. M. Ziegler, *Proofs from THE BOOK*, https://doi.org/10.1007/978-3-662-57265-8_31

which is smaller than $\frac{1}{2}$ for $n = 23$ (this is the “birthday paradox”!), less than 9 percent for $n = 42$, and exactly 0 for $n > 365$ (the “pigeon-hole principle,” see Chapter 28). The formula is easy to see — if we take the persons in some fixed order: If the first i persons have distinct birthdays, then the probability that the $(i+1)$-st person doesn't spoil the series is $1 - \frac{i}{365}$, since there are $365 - i$ birthdays left.

Similarly, if n balls are placed independently and randomly into K boxes, then the probability that no box gets more than one ball is

$$p(n,K) \;=\; \prod_{i=1}^{n-1}\Big(1 - \frac{i}{K}\Big).$$

The coupon collector

Children buy photos of pop stars (or soccer stars) for their albums, but they buy them in little nontransparent envelopes, so they don't know which photo they will get. If there are n different photos, what is the expected number of pictures a kid has to buy until he or she gets every motif at least once?

Equivalently, if you randomly take balls from a bowl that contains n distinguishable balls, and if you put your ball back each time, and then again mix well, how often do you have to draw on average until you have drawn each ball at least once?

If you already have drawn k distinct balls, then the probability not to get a new one in the next drawing is $\frac{k}{n}$. So the probability to need exactly s drawings for the next new ball is $(\frac{k}{n})^{s-1}(1 - \frac{k}{n})$. And thus the expected number of drawings for the next new ball is

$$\sum_{s\geq 1}\Big(\frac{k}{n}\Big)^{s-1}\Big(1 - \frac{k}{n}\Big)s \;=\; \frac{1}{1 - \frac{k}{n}},$$

as we get from the series in the margin. So the expected number of drawings until we have drawn *each* of the n different balls at least once is

$$\begin{aligned}\sum_{s\geq 1} x^{s-1}(1-x)s \;&=\\ &= \sum_{s\geq 1} x^{s-1}s - \sum_{s\geq 1} x^s s\\ &= \sum_{s\geq 0} x^s(s+1) - \sum_{s\geq 0} x^s s\\ &= \sum_{s\geq 0} x^s \;=\; \frac{1}{1-x},\end{aligned}$$

where at the end we sum a geometric series (see page 48).

$$\sum_{k=0}^{n-1}\frac{1}{1-\frac{k}{n}} \;=\; \frac{n}{n} + \frac{n}{n-1} + \cdots + \frac{n}{2} + \frac{n}{1} \;=\; nH_n \;\approx\; n\log n,$$

with the bounds on the size of harmonic numbers that we had obtained on page 13. So the answer to the coupon collector's problem is that we have to expect that roughly $n \log n$ drawings are necessary.

The estimate that we need in the following is for the probability that you need significantly more than $n \log n$ trials. If V_n denotes the number of drawings needed (this is the random variable whose expected value is $E[V_n] \approx n \log n$), then for $n \geq 1$ and $c \geq 0$, the probability that we need more than $m := \lceil n \log n + cn \rceil$ drawings is

$$\text{Prob}\big[V_n > m\big] \;\leq\; e^{-c}.$$

Indeed, if A_i denotes the event that the ball i is not drawn in the first m drawings, then

$$\begin{aligned}\text{Prob}\big[V_n > m\big] &= \text{Prob}\Big[\bigcup_i A_i\Big] \;\le\; \sum_i \text{Prob}\big[A_i\big] \\ &= n\Big(1 - \frac{1}{n}\Big)^m \;<\; ne^{-m/n} \;\le\; e^{-c}.\end{aligned}$$

A little calculus shows that $\big(1 - \frac{1}{n}\big)^n$ is an increasing function in n, which converges to $1/e$. So $\big(1 - \frac{1}{n}\big)^n < \frac{1}{e}$ holds for all $n \ge 1$.

Now let's grab a deck of n cards. We number them 1 up to n in the order in which they come — so the card numbered "1" is at the top of the deck, while "n" is at the bottom. Let us denote from now on by $\mathfrak{S}_n$ the set of all permutations of $1, \dots, n$. *Shuffling* the deck amounts to the application of certain *random permutations* to the order of the cards. Ideally, this might mean that we apply an arbitrary permutation $\pi \in \mathfrak{S}_n$ to our starting order $(1, 2, \dots, n)$, each of them with the same probability $\frac{1}{n!}$. Thus, after doing this just once, we would have our deck of cards in order $\pi = (\pi(1), \pi(2), \dots, \pi(n))$, and this would be a perfect random order. But that's not what happens in real life. Rather, when shuffling only "certain" permutations occur, perhaps not all of them with the same probability, and this is repeated a "certain" number of times. After that, we expect or hope the deck to be at least "close to random."

"Top-in-at-random"

Top-in-at-random shuffles

These are performed as follows: you take the top card from the deck, and insert it into the deck at one of the n distinct possible places, each of them with probability $\frac{1}{n}$. Thus one of the permutations

$$\tau_i \;=\; \big(2, 3, \dots, i, \overset{\overset{\textstyle i}{\downarrow}}{1}, i{+}1, \dots, n\big)$$

is applied, $1 \le i \le n$. After one such shuffle the deck doesn't look random, and indeed we expect to need lots of such shuffles until we reach that goal. A typical run of top-in-at-random shuffles may look as follows (for $n = 5$):

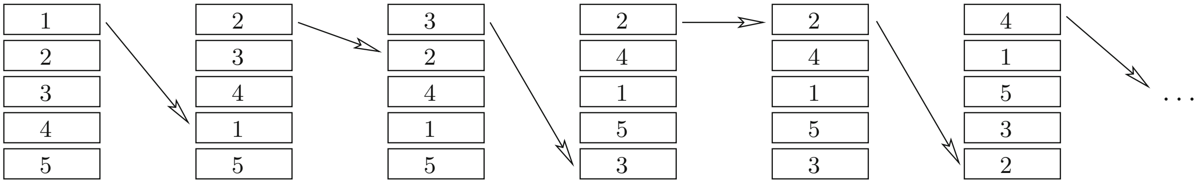

How should we measure "being close to random"? Probabilists have cooked up the "variation distance" as a rather unforgiving measure of randomness: We look at the probability distribution on the $n!$ different orderings of our deck, or equivalently, on the $n!$ different permutations $\sigma \in \mathfrak{S}_n$ that yield the orderings.

Two examples are our starting distribution E, which is given by

$$\begin{aligned} \mathsf{E}(\mathrm{id}) &= 1, \\ \mathsf{E}(\pi) &= 0 \quad \text{otherwise}, \end{aligned}$$

and the uniform distribution U given by

$$\mathsf{U}(\pi) = \tfrac{1}{n!} \quad \text{for all } \pi \in \mathfrak{S}_n.$$

The *variation distance* between two probability distributions Q_1 and Q_2 is now defined as

$$\|\mathsf{Q}_1 - \mathsf{Q}_2\| := \tfrac{1}{2} \sum_{\pi \in \mathfrak{S}_n} |\mathsf{Q}_1(\pi) - \mathsf{Q}_2(\pi)|.$$

By setting $S := \{\pi \in \mathfrak{S}_n : \mathsf{Q}_1(\pi) > \mathsf{Q}_2(\pi)\}$ and using $\sum_\pi \mathsf{Q}_1(\pi) = \sum_\pi \mathsf{Q}_2(\pi) = 1$ we can rewrite this as

$$\|\mathsf{Q}_1 - \mathsf{Q}_2\| = \max_{S \subseteq \mathfrak{S}_n} |\mathsf{Q}_1(S) - \mathsf{Q}_2(S)|,$$

with $\mathsf{Q}_i(S) := \sum_{\pi \in S} \mathsf{Q}_i(\pi)$. Clearly we have $0 \leq \|\mathsf{Q}_1 - \mathsf{Q}_2\| \leq 1$. In the following, "being close to random" will be interpreted as "having small variation distance from the uniform distribution." Here the distance between the starting distribution and the uniform distribution is very close to 1:

$$\|\mathsf{E} - \mathsf{U}\| = 1 - \tfrac{1}{n!}.$$

After one top-in-at-random shuffle, this will not be much better:

$$\|\mathsf{Top} - \mathsf{U}\| = 1 - \tfrac{1}{(n-1)!}.$$

For card players, the question is not "exactly how close to uniform is the deck after a million riffle shuffles?", but "is 7 shuffles enough?"

(Aldous & Diaconis [1])

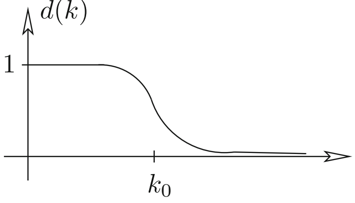

The probability distribution on $\mathfrak{S}_n$ that we obtain by applying the top-in-at-random shuffle k times will be denoted by Top^{*k}. So how does $\|\mathsf{Top}^{*k} - \mathsf{U}\|$ behave if k gets larger, that is, if we repeat the shuffling? And similarly for other types of shuffling? General theory (in particular, Markov chains on finite groups; see e. g. Behrends [3]) implies that for large k the variation distance $d(k) := \|\mathsf{Top}^{*k} - \mathsf{U}\|$ goes to zero exponentially fast, but it does not yield the "cut-off" phenomenon that one observes in practice: After a certain number k_0 of shuffles "suddenly" $d(k)$ goes to zero rather fast. Our margin displays a schematic sketch of the situation.

Strong uniform stopping rules

The amazing idea of strong uniform stopping rules by Aldous and Diaconis captures the essential features. Imagine that the casino manager closely watches the shuffling process, analyzes the specific permutations that are applied to the deck in each step, and after a number of steps that depends on the permutations that he has seen calls "STOP!". So he has a *stopping rule* that ends the shuffling process. It depends only on the (random) shuffles that have already been applied. The stopping rule is *strong uniform* if the following condition holds for all $k \geq 0$:

> ***If** the process is stopped after exactly k steps, **then** the resulting permutations of the deck have uniform distribution (exactly!).*

Let T be the number of steps that are performed until the stopping rule tells the manager to cry "STOP!"; so this is a random variable. Similarly, the ordering of the deck after k shuffles is given by a random variable X_k (with values in $\mathfrak{S}_n$). With this, the stopping rule is strong uniform if for all feasible values of k,

$$\text{Prob}\big[X_k = \pi \mid T = k\big] \;=\; \frac{1}{n!} \qquad \text{for all } \pi \in \mathfrak{S}_n.$$

Conditional probabilities

The *conditional probability*

$$\text{Prob}[A \,|\, B]$$

denotes the probability of the event A under the condition that B happens. This is just the probability that both events happen, divided by the probability that B is true, that is,

$$\text{Prob}[A \,|\, B] \;=\; \frac{\text{Prob}[A \wedge B]}{\text{Prob}[B]}.$$

Three aspects make this interesting, useful, and remarkable:

1. Strong uniform stopping rules exist: For many examples they are quite simple.
2. Moreover, these can be analyzed: Trying to determine $\text{Prob}[T > k]$ leads often to simple combinatorial problems.
3. This yields effective upper bounds on the variation distances such as $d(k) = \|\mathsf{Top}^{*k} - \mathsf{U}\|$.

For example, for the top-in-at-random shuffles a strong uniform stopping rule is

> "STOP after the original bottom card (labelled n) is first inserted back into the deck."

Indeed, if we trace the card n during these shuffles,

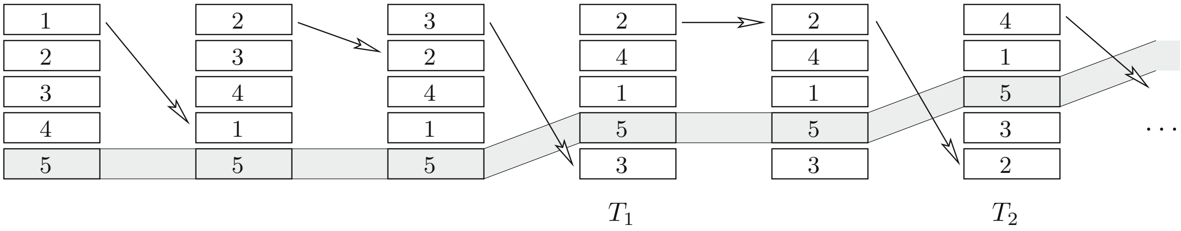

we see that during the whole process the ordering of the cards below this card is completely uniform. So, after the card n rises to the top and then is inserted at random, the deck is uniformly distributed; we just don't know when precisely this happens (but the manager does).

Now let T_i be the random variable which counts the number of shuffles that are performed until for the first time i cards lie below card n. So we have to determine the distribution of

$$T \;=\; T_1 + (T_2 - T_1) + \cdots + (T_{n-1} - T_{n-2}) + (T - T_{n-1}).$$

But each summand in this corresponds to a coupon collector's problem: $T_i - T_{i-1}$ is the time until the top card is inserted at one of the i possible places below the card n. So it is also the time that the coupon collector takes from the $(n-i)$-th coupon to the $(n-i+1)$-st coupon. Let V_i be the number of pictures bought until he has i different pictures. Then

$$V_n \;=\; V_1 + (V_2 - V_1) + \cdots + (V_{n-1} - V_{n-2}) + (V_n - V_{n-1}),$$

and we have seen that $\text{Prob}[T_i - T_{i-1} = j] = \text{Prob}[V_{n-i+1} - V_{n-i} = j]$ for all i and j. Hence the coupon collector and the top-in-at-random shuffler perform equivalent sequences of independent random processes, just in the opposite order (for the coupon collector, it's hard at the end). Thus we know that the strong uniform stopping rule for the top-in-at-random shuffles takes more than $k = \lceil n \log n + cn \rceil$ steps with low probability:

$$\text{Prob}\big[T > k\big] \;\leq\; e^{-c}.$$

And this in turn means that after $k = \lceil n \log n + cn \rceil$ top-in-at-random shuffles, our deck is "close to random," with

$$d(k) \;=\; \|\mathsf{Top}^{*k} - \mathsf{U}\| \;\leq\; e^{-c},$$

due to the following simple but crucial lemma.

Lemma. *Let $\mathsf{Q} : \mathfrak{S}_n \longrightarrow \mathbb{R}$ be any probability distribution that defines a shuffling process Q^{*k} with a strong uniform stopping rule whose stopping time is T. Then for all $k \geq 0$,*

$$\|\mathsf{Q}^{*k} - \mathsf{U}\| \;\leq\; \text{Prob}\big[T > k\big].$$

■ **Proof.** If X is a random variable with values in $\mathfrak{S}_n$, with probability distribution Q, then we write $\mathsf{Q}(S)$ for the probability that X takes a value in $S \subseteq \mathfrak{S}_n$. Thus $\mathsf{Q}(S) = \text{Prob}[X \in S]$, and in the case of the uniform distribution $\mathsf{Q} = \mathsf{U}$ we get

$$\mathsf{U}(S) \;=\; \text{Prob}\big[X \in S\big] \;=\; \frac{|S|}{n!}.$$

For every subset $S \subseteq \mathfrak{S}_n$, we get the probability that after k steps our deck is ordered according to a permutation in S as

$$\begin{aligned}
\mathsf{Q}^{*k}(S) \;&=\; \text{Prob}[X_k \in S] \\
&=\; \sum_{j \leq k} \text{Prob}[X_k \in S \;\wedge\; T = j] \;+\; \text{Prob}[X_k \in S \;\wedge\; T > k] \\
&=\; \sum_{j \leq k} \mathsf{U}(S)\, \text{Prob}[T = j] \;+\; \text{Prob}[X_k \in S \,|\, T > k] \cdot \text{Prob}[T > k] \\
&=\; \mathsf{U}(S)\,(1 - \text{Prob}[T > k]) \;+\; \text{Prob}[X_k \in S \,|\, T > k] \cdot \text{Prob}[T > k] \\
&=\; \mathsf{U}(S) \;+\; \big(\text{Prob}[X_k \in S \,|\, T > k] - \mathsf{U}(S)\big) \cdot \text{Prob}[T > k].
\end{aligned}$$

This yields

$$|\mathsf{Q}^{*k}(S) - \mathsf{U}(S)| \;\leq\; \text{Prob}[T > k]$$

since

$$\text{Prob}[X_k \in S \,|\, T > k] \;-\; \mathsf{U}(S)$$

is a difference of two probabilities, so it has absolute value at most 1. □

This is the point where we have completed our analysis of the top-in-at-random shuffle: We have proved the following upper bound for the number of shuffles needed to get "close to random."

Theorem 1. *Let $c \geq 0$ and $k := \lceil n \log n + cn \rceil$. Then after performing k top-in-at-random shuffles on a deck of n cards, the variation distance from the uniform distribution satisfies*

$$d(k) \;:=\; \|\mathsf{Top}^{*k} - \mathsf{U}\| \;\leq\; e^{-c}.$$

One can also verify that the variation distance $d(k)$ stays large if we do significantly fewer than $n \log n$ top-in-at-random shuffles. The reason is that a smaller number of shuffles will not suffice to destroy the relative ordering on the lowest few cards in the deck.

Of course, top-in-at-random shuffles are extremely inefficient — with the bounds of our theorem, we need more than $n \log n \approx 205$ top-in-at random shuffles until a deck of $n = 52$ cards is mixed up well. Thus we now switch our attention to a much more interesting and realistic model of shuffling.

Riffle shuffles

This is what dealers do at the casino: They take the deck, split it into two parts, and these are then interleaved, for example by dropping cards from the bottoms of the two half-decks in some irregular pattern.

Again a riffle shuffle performs a certain permutation on the cards in the deck, which we initially assume to be labelled from 1 to n, where 1 is the top card. The riffle shuffles correspond exactly to the permutations $\pi \in \mathfrak{S}_n$ such that the sequence

$$(\pi(1), \pi(2), \;\ldots\;, \pi(n))$$

consists of two interlaced increasing sequences (only for the identity permutation it is one increasing sequence), and that there are exactly $2^n - n$ distinct riffle shuffles on a deck of n cards.

“A riffle shuffle”

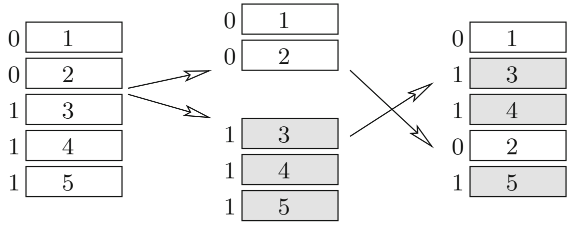

In fact, if the pack is split such that the top t cards are taken into the right hand $(0 \leq t \leq n)$ and the other $n-t$ cards into the left hand, then there are $\binom{n}{t}$ ways to interleave the two hands, all of which generate distinct permutations — except that for each t there is one possibility to obtain the identity permutation.

Now it’s not clear which probability distribution one should put on the riffle shuffles — there is no unique answer since amateurs and professional dealers would shuffle differently. However, the following model, developed first by Edgar N. Gilbert and Claude Shannon in 1955 (at the legendary

Bell Labs "Mathematics of Communication" department at the time), has several virtues:

- it is elegant, simple, and seems natural,
- it models quite well the way an amateur would perform riffle shuffles,
- and we have a chance to analyze it.

Here are three descriptions — all of them describe the same probability distribution Rif on $\mathfrak{S}_n$:

1. $\mathsf{Rif} : \mathfrak{S}_n \longrightarrow \mathbb{R}$ is defined by

$$\mathsf{Rif}(\pi) \ := \ \begin{cases} \frac{n+1}{2^n} & \text{if } \pi = \text{id}, \\ \frac{1}{2^n} & \text{if } \pi \text{ consists of two increasing sequences}, \\ 0 & \text{otherwise.} \end{cases}$$

2. Cut off t cards from the deck with probability $\frac{1}{2^n}\binom{n}{t}$, take them into your right hand, and take the rest of the deck into your left hand. Now when you have r cards in the right hand and ℓ in the left, "drop" the bottom card from your right hand with probability $\frac{r}{r+\ell}$, and from your left hand with probability $\frac{\ell}{r+\ell}$. Repeat!

3. An *inverse shuffle* would take a subset of the cards in the deck, remove them from the deck, and place them on top of the remaining cards of the deck — while maintaining the relative order in both parts of the deck. Such a move is determined by the subset of the cards: Take all subsets with equal probability.

   Equivalently, assign a label "0" or "1" to each card, randomly and independently with probabilities $\frac{1}{2}$, and move the cards labelled "0" to the top.

The inverse riffle shuffles correspond to the permutations $\pi = (\pi(1), \dots, \pi(n))$ that are increasing except for at most one "descent." (Only the identity permutation has no descent.)

It is easy so see that these descriptions yield the same probability distributions. For (1) $\Longleftrightarrow$ (3) just observe that we get the identity permutation whenever all the 0-cards are on top of all the cards that are assigned a 1.

This defines the model. So how can we analyze it? How many riffle shuffles are needed to get close to random? We won't get the precise best-possible answer, but quite a good one, by combining three components:

(1) We analyze inverse riffle shuffles instead,

(2) we describe a strong uniform stopping rule for these,

(3) and show that the key to its analysis is given by the birthday paradox!

Theorem 2. *After performing k riffle shuffles on a deck of n cards, the variation distance from a uniform distribution satisfies*

$$\|\mathsf{Rif}^{*k} - \mathsf{U}\| \ \leq \ 1 - \prod_{i=1}^{n-1}\left(1 - \frac{i}{2^k}\right).$$

■ **Proof.** (1) We may indeed analyze inverse riffle shuffles and try to see how fast they get us from the starting distribution to (close to) uniform. These inverse shuffles correspond to the probability distribution that is given by $\overline{\mathsf{Rif}}(\pi) := \mathsf{Rif}(\pi^{-1})$.

Now the fact that every permutation has its unique inverse, and the fact that $\mathsf{U}(\pi) = \mathsf{U}(\pi^{-1})$, yield

$$\|\mathsf{Rif}^{*k} - \mathsf{U}\| \;=\; \|\,\overline{\mathsf{Rif}}^{*k} - \mathsf{U}\|.$$

(This is Reeds' inversion lemma!)

(2) In every inverse riffle shuffle, each card gets associated a digit 0 or 1:

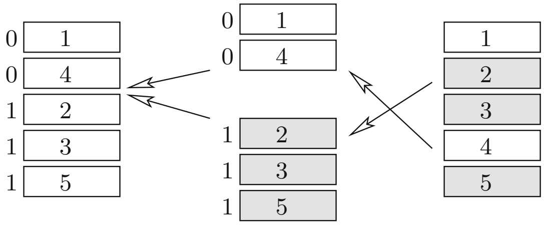

If we remember these digits — say we just write them onto the cards — then after k inverse riffle shuffles, each card has gotten an ordered string of k digits. Our stopping rule is:

> "STOP as soon as all cards have distinct strings."

When this happens, the cards in the deck are *sorted* according to the binary numbers $b_k b_{k-1} \dots b_2 b_1$, where b_i is the bit that the card has picked up in the i-th inverse riffle shuffle. Since these bits are perfectly random and independent, this stopping rule is strong uniform!

In the following example, for $n = 5$ cards, we need $T = 3$ inverse shuffles until we stop:

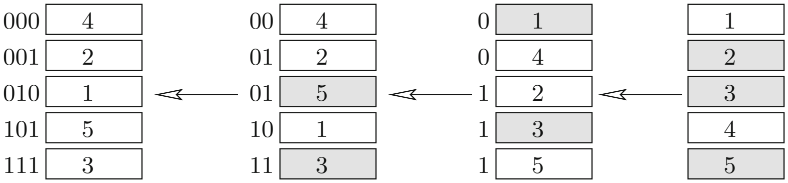

(3) The time T taken by this stopping rule is distributed according to the birthday paradox, for $K = 2^k$: We put two cards into the same box if they have the same label $b_k b_{k-1} \dots b_2 b_1 \in \{0,1\}^k$. So there are $K = 2^k$ boxes, and the probability that some box gets more than one card ist

$$\mathrm{Prob}[T > k] \;=\; 1 - \prod_{i=1}^{n-1}\left(1 - \frac{i}{2^k}\right),$$

and as we have seen this bounds the variation distance $\|\mathsf{Rif}^{*k} - \mathsf{U}\| = \|\,\overline{\mathsf{Rif}}^{*k} - \mathsf{U}\|$. □

k	$d(k)$
1	1.000
2	1.000
3	1.000
4	1.000
5	0.952
6	0.614
7	0.334
8	0.167
9	0.085
10	0.043

The variation distance after k riffle shuffles, according to [2]

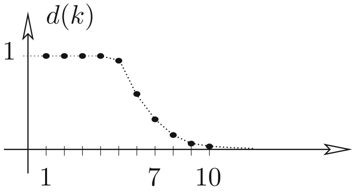

So how often do we have to shuffle? For large n we will need roughly $k = 2\log_2(n)$ shuffles. Indeed, setting $k := 2\log_2(cn)$ for some $c \geq 1$ we find (with a bit of calculus) that $P[T > k] \approx 1 - e^{-\frac{1}{2c^2}} \approx \frac{1}{2c^2}$.
Explicitly, for $n = 52$ cards the upper bound of Theorem 2 reads $d(10) \leq 0.73$, $d(12) \leq 0.28$, $d(14) \leq 0.08$ — so $k = 12$ should be "random enough" for all practical purposes. But we don't do 12 shuffles "in practice" — and they are not really necessary, as a more detailed analysis shows (with the results given in the margin). The analysis of riffle shuffles is part of a lively ongoing discussion about the right measure of what is "random enough." Diaconis [4] is a guide to recent developments.

Indeed, does it matter? Yes, it does: Even after three good riffle shuffles a sorted deck of 52 cards looks quite random ... but it isn't. Martin Gardner [5, Chapter 7] describes a number of striking card tricks that are based on the hidden order in such a deck!

References

[1] D. ALDOUS & P. DIACONIS: *Shuffling cards and stopping times,* Amer. Math. Monthly **93** (1986), 333-348.

[2] D. BAYER & P. DIACONIS: *Trailing the dovetail shuffle to its lair,* Annals Applied Probability **2** (1992), 294-313.

[3] E. BEHRENDS: *Introduction to Markov Chains,* Vieweg, Braunschweig/ Wiesbaden 2000.

[4] P. DIACONIS: *Mathematical developments from the analysis of riffle shuffling,* in: "Groups, Combinatorics and Geometry. Durham 2001" (A. A. Ivanov, M. W. Liebeck and J. Saxl, eds.), World Scientific, Singapore 2003, pp. 73-97.

[5] M. GARDNER: *Mathematical Magic Show,* Knopf, New York/Allen & Unwin, London 1977.

[6] E. N. GILBERT: *Theory of Shuffling,* Technical Memorandum, Bell Laboratories, Murray Hill NJ, 1955.

"Random enough?"