Fibonacci Solitaire

Alexander Gnedin
University of Newcastle

and

Sergei Kerov
Steklov Mathematical Institute at St. Petersburg

Abstract. Fibonacci Solitaire is a combinatorial algorithm which associates with a permutation of $[n] = \{1, \ldots, n\}$ a partition of $[n]$ into couples and singletons. We study the output configuration when the algorithm is applied to a random permutation, with emphasis on the large n-asymptotics. We show that the set of singletons, properly scaled, resembles a familiar 'stick-breaking' Poisson configuration, while the configuration of couples becomes close to uniform.

1 Introduction. Fibonacci Solitaire is a combinatorial algorithm invented in [1] as a method to convert random permutations into Fibonacci words. The algorithm owes its name to a simple chance game which can be played with real cards.

To play the game a deck of n cards with values $1, 2, \ldots, n$ is taken. The deck is shuffled, cards are turned up one at a time and dealt into or removed from a row on the table. Each time a card, say A, is turned we compare it with the lowest card on the table (if any), say B, and the cards are moved according to the rules

(S): if A is lower than B, card A is added to the row on the table in the leftmost position,

(C): if A is higher than B both cards are removed.

Card A is dealt into the row also in the case when at the moment there are no other cards on the table, in particular when A is the first card turned.

For example, suppose a shuffled deck of 8 cards is in the order

position	1	2	3	4	5	6	7	8
value	6	2	4	8	3	7	5	1

then the player sees the rows

$$\begin{array}{cc} & 6 \\ 2 & 6 \\ & 6 \\ & \emptyset \\ & 3 \\ & \emptyset \\ & 5 \\ 1 & 5 \end{array}$$

where $\emptyset$ means that the row is empty after the 4th and the 6th moves.

For n even, the player might count the game a success if it ends with no cards on the table. A result from [1] which strongly motivated the present work says that the probability of a successful game is equal to the quantity

$$\binom{n}{n/2}\frac{1}{2^n} \tag{1}$$

best known as the return probability for the simple random walk. (This gives an approximate 11% probability of winning for the standard deck of 52 cards.) Extending this result, we will show that the number of cards which remain on the table has the same probability law as the number of odd-size cycles in a random permutation of $[n] = \{1, 2, \dots, n\}$.

When the rule (C) applies we say that the cards A and B are *coupled*. At the end of the game some cards are coupled while the others are not, thus we can regard the final configuration of cards as an *involution*, i.e. a partition of the deck into couples and singletons. In our example the resulting involution is

$$(1) \qquad (2,4) \qquad (3,7) \qquad (5) \qquad (6,8)$$

with the blocks written in the order of increase of their least elements.

The nature of this involution is a central theme of this paper. We find the exact probability law and study asymptotic proprties of the involution for large n. We show that the asymptotic configuration of the singletons resembles a Poisson configuration appearing via a familiar stick-breaking construction (Theorem 7), and that the configuration of couples becomes close to uniform (Theorem 16). We also discuss an interacting particle process and various combinatorial objects associated with the algorithm.

2 Involutions. An *involution* on a finite set L of different real numbers is a partition of L into blocks of size one or two, singletons and couples. Of two mates in a couple the smaller is called a *bottom* and the larger is called a *top*. If the involution has at least one singleton the smallest of them is called the *base*. For $n = |L|$ even we define a *perfect matching* as involution without singletons, and for n odd as involution with exactly one singleton which is larger than all bottoms (our definition of perfect matching differs from the conventional graph-theoretic one).

An arrangement of an n-card deck can be identified with a permutation $v_1, \dots, v_n$ of $[n]$ by saying that v_p is the value of the card at position p. Fibonacci Solitaire is a sequential construction of an involution on $[n]$ from a permutation.

The algorithm starts with the single involution on $L_1 = \{v_1\}$ and at each step $p = 2, \ldots, n$ yiels an involution on $L_p = \{v_1, \ldots, v_p\}$ according to the following rules. Given an involution on L_{p-1}, the value v_p becomes the base of L_p if L_{p-1} has no singletons smaller than v_p (rule (S)), or v_p is coupled with the base of L_{p-1} otherwise (rule (C)).

Note that in case (C) v_p becomes a top, the base of L_{p-1} becomes a bottom and the second smallest singleton of L_{p-1} (if any) becomes the base of L_p. Figure 1 shows the course of the algorithm applied to the same permutation as in Introduction, with the arc systems representing involutions arising at intermediate steps.

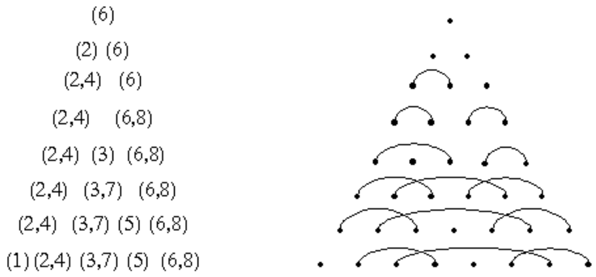

Figure 1: Fibonacci Solitaire

The procedure becomes more transparent when a permutation is represented in the matrix form, via marking the squares with coordinates (v_p, p) of a $n \times n$ chessboard ($(1,1)$ is at the top left). The chessboard is scanned row-wise. At step p the mark in (v_p, p) is linked with the leftmost yet uncoupled mark located to the north-west of (v_p, p) (if any), as in Figure 2. The final involution on $[n]$ is obtained by coupling values v an w if the marks in the columns v and w get linked.

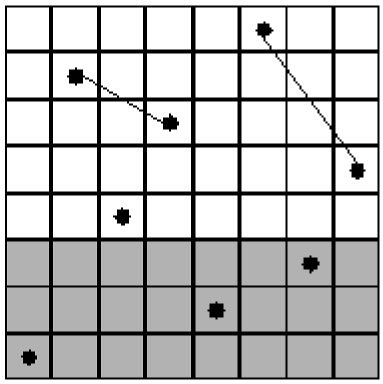

Figure 2: Step 5 in the matrix representation

3 Probability law. We are interested in the involution arising when the algorithm is applied to a random permutation π_n chosen uniformly from the set of all $n!$ permutations of $[n]$. Determining the probability law of the output amounts to counting the number of permutations which yield a particular involution.

For an involution on a finite set L and a bottom b we define the *depth* $\delta(b)$ to be one plus the total cardinality of blocks which are completely larger than b. Denoting by T, B ans S the sets of tops, bottoms and singletons, respectively, we have by the definition

$$\delta(b) := 1 + |\{s \in S : s > b\}| + 2|\{b' \in B : b' > b\}|. \tag{2}$$

We define a *rigging* to be a correspondence $b \mapsto \iota(b)$ which attributes to each bottom an integer *index* in the range $1 \leq \iota(b) \leq \delta(b)$.

Fibonacci Solitaire involution on $[n]$ comes up with the *canonical* rigging which assigns to each bottom b the index $\iota(b)$ equal to the depth of $\delta_p(b)$ in L_p, at step p when b gets coupled (the upper bound for the index cannot be broken after the coupling because $\delta_p(b)$ is nondecreasing in p). In our running example there are three bottoms $2, 3$ and 6 and their depthes are $\delta_8(2) = 6$, $\delta_8(3) = 4$ and $\delta_8(6) = 1$. The canonical rigging is $\iota(2) = 2$, $\iota(3) = 3$ and $\iota(6) = 1$.

Our first result says that rigging is exactly the information necessary to re-construct a permutation from the output involution. The idea to keep in mind is that the canonical rigging is also defined for any partial involution arising at an intermediate step of the algorithm.

Theorem 1 *Fibonacci Solitaire yields a bijection between permutations and rigged involutions on $[n]$.*

Proof. The algorithm yields a sequence of canonically rigged involutions in such a manner that at step $p > 1$ a rigged involution I_{p-1} supported by L_{p-1} and $v_p \notin L_{p-1}$ are transformed into a rigged involution I_p on $L_p = L_{p-1} \cup \{v_p\}$. Let us denote this operation by

$$I_p = FS(I_{p-1}, v_p) \tag{3}$$

and examine it in detail for arbitrary I_{p-1}. When the rule (C) applies, v_p couples with the base $b^* \in L_{p-1}$, which becomes a bottom with the index

$$\iota(b^*) = \delta_p(b^*) \tag{4}$$

and has the properties (a): b^* is the least bottom in I_p satisfying (4), and (b): b^* is smaller than any singleton of I_p. While (b) is obvious, (a) follows by noting that

$$\delta_p(b') = \delta_{p-1}(b') + 1 > \iota(b') \tag{5}$$

for all bottoms $b' < b^*$ in L_p. When the rule (S) applies, there is no b^* meeting (a) and (b), because v_p becomes the base and (5) holds for all bottoms $b' \in L_p$ less than v_p.

We see that if (3) holds then either (c): there exists b^* satisfying (a) and (b) and I_p is obtained from I_{p-1} by coupling b^* with v_p and setting (4), or (s): such b^* does not exist,

v_p is the base of I_p and the rule (S) applies. On the other hand, each of the conditions (c) and (s) is sufficient for (3), whence (3) is equivalent to '(c) or (s)'.

We argue now that the transform can be uniquely inverted: for arbitrary rigged involution I_p on L_p we can find unique I_{p-1} and v_p such that (3) holds. Indeed, if $b^* \in L_p$ satisfying (a) and (b) exists, we determine v_p to be the top mate of b^* and I_{p-1} to be the restriction of I_p on $L_p \setminus \{v_p\}$. Restricting the rigging we will not violate the constraints $\delta_{p-1}(b) \leq \iota(b)$ because the depths of bottoms $b' > b^*$ remain unchanged, while the depths of bottoms $b' < v_p$ satisfy (5). The data I_p, I_{p-1}, v_p agree with (c), whence (3).

If such b^* does not exist, we let v_p to be the base of I_p and also define I_{p-1} by restricting I_p to $L_p \setminus \{v_p\}$. Then we are in case (s) and again (3) follows. The uniqueness claim follows because condition '(c) or (s)' is necessary.

Starting with arbitrary I_n and inverting (3) for $p = n, n-1, \ldots, 2$ we arrive at a unique permutation $v_1, \ldots, v_n$. Therefore the correspondence is bijective. □

The probability law of the Fibonacci Solitaire involution is an immediate consequence of the bijection.

Theorem 2 *A particular involution on $[n]$ appears as the output of the algorithm with probability*

$$\frac{1}{n!} \prod_b \delta(b) \tag{6}$$

where the product is over all bottoms of the involution.

Proof. For a given involution the product in (6) is the number of ways to select a rigging, and each choice corresponds to a permutation which yields this involution.□

Recall that the double factorial is defined as $n!! = n(n-2)(n-4)\ldots$.

Corollary 3 *The Fibonacci Solitaire involution resulting from a random permutation π_n has the following properties.*

(i) *Each of $(n-1)!!$ particular perfect matchings of $[n]$ occurs with probability $(n!!)^{-1}$. The total probability of perfect matchings amounts to $(n-1)!!/n!!$, which is (1) for n even.*

(ii) *Perfect matchings are the most likely involutions. The least probable is the pure singleton involution, which has probability $(n!)^{-1}$.*

Proof. (i) To compute the number of perfect matchings on n points we will examine configurations spanned on the k largest non-top values.

Consider first odd n. The largest non-top value must be a singleton, say s. The second largest non-top value is a bottom b which we put to the left of s. Now for the relative location of the top mate t of this b there are two possibilities bts and bst. Similarly, the top mate of the next largest bottom can be inserted in 4 different ways, etc. It follows that the number of perfect matchings is $2 \cdot 4 \cdot \ldots = n - 1!!$. The number of riggings for a perfect matching is also $(n-1)!!$, because the depths from the right to the left are also $2, 4, \ldots$.

Similarly, for n even, both sequences are $1, 3, \ldots, n-1$ and we can apply (6).

(ii) Fix an involution and consider a fragment squeezed between two singletons like $s_1, b_1, b_2, \ldots, b_j, s_2$ where we write the elements in increasing order and skip the tops. Denoting by d the depth of b_1, the contribution of the fragment into the product in (6) is $d(d-2)\ldots(d-2j+2)$. If we modify the involution by coupling s_1 with s_2, then s_1 will be a bottom and s_2 a top, and the fragment will contribute then the larger value $(d+1)(d-1)(d-3)\ldots(d-2j+1)$. Because other factors remain unchanged, the modified involution is more likely. Continuing coupling in this manner we will always increase probability. For n even coupling of singletons ends with a perfect matching. For n odd coupling continues until exactly one singleton remains, and among such compositions the most likely are those which have the singleton larger than all bottoms (a direct check).

Finally, it is immediate from (6) that the pure-singleton involution is the least probable. The only permutation which yields this involution is $(v_1, \ldots, v_n) = (n, n-1, \ldots, 1)$. $\Box$

4 The dual algorithm. The exact formula (6) seems to be of limited value to determine distributions and moments of most interesting characteristics of the involution. Fortunately, there is an alternative construction which leads to recurrent formulas.

The following *dual* algorithm constructs the involution by processing the inverse permutation $p_1, \ldots, p_n$, with p_v being the position of value v. At any step of the algorithm each of the positions can be *occupied*, *free* or *reserved*. Initially, all n positions are free. At step 1 position p_1 becomes occupied, and we determine value 1 to be a singleton when $p_1 = n$, or otherwise value 1 becomes a bottom and position $p_1 + 1$ becomes reserved. In the second case the top mate of value 1 will appear at one of the forthcoming steps in the reserved position.

Inductively, upon completing step $v-1$ of the algorithm the set of values $\{1, \ldots, v-1\}$ is comprised of complete blocks and bottoms, the occupied positions are $p_1, \ldots, p_{v-1}$, each reserved position is attributed to one of the bottoms and the remaining positions are free.

The piece of information added at step v is p_v. If p_v is one of the formerly reserved positions, v gets coupled with the corresponding bottom. If p_v is the highest of currently free positions, the value v is determined to be a singleton. If p_v is a free position but not the highest free, then v is determined to be a bottom and the lowest of those free positions which are higher than p_v becomes reserved for the top mate of v. In any case, p_v becomes occupied.

To illustrate, the inverse permutation in our running example is

value	1	2	3	4	5	6	7	8
position	8	2	5	3	7	1	6	4

and the dual algorithm yields

$$\begin{array}{c}
(1) \\
(1) \quad (2, v_3) \\
(1) \quad (2, v_3) \quad (3, v_6) \\
(1) \quad (2, 4) \quad (3, v_6) \\
(1) \quad (2, 4) \quad (3, v_6) \quad (5) \\
(1) \quad (2, 4) \quad (3, v_6) \quad (6, v_4) \\
(2, 4) \quad (3, 7) \quad (5) \quad (6, v_4) \\
(1) \quad (2, 4) \quad (3, 7) \quad (5) \quad (6, 8)
\end{array}$$

For instance, at step 3 the value 3 is determined to be a bottom and position 6 is reserved for the mate of 3; this mate $v_6 = 7$ becomes definite at step 7.

When the matrix representation of permutations is used, the dual algorithm proceeds by column-wise scanning of the chessboard. Figure 3 shows the configuration after step 4, with the shadowed boxes in rows 3 and 6 demonstrating reservations at steps 2 and 3.

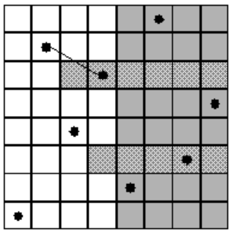

Figure 3: Step 4 of the dual algorithm

We wish to show that the dual algorithm and Fibonacci Solitaire produce the same involution on $[n]$. The argument to follow will be based on an important reducibility property.

Proposition 4 *Let R be a block of the involution appearing when the solitaire rules are applied to the sequence $x_1, \ldots, x_n$. The restricted involution on $\{x_1, \ldots, x_n\} \setminus R$ coincides with the involution which appears when the algorithm is applied to the reduced sequence, with elements of R removed.*

Proof. Suppose first that block R is a couple (b, t) and write the input sequence as L', b, L'', t, L''', where any of the subsequences L', L'', L''' can be empty. Immediately as L' gets processed the base goes to b and then neither exceeds b nor disappears until t gets processed. Therefore, the restriction of the involution onto L'' is a perfect matching. It is easily seen that processing L'' immediately after L' does not alter the set of singletons which remain after L' gets processed, hence elements of L' can be coupled only with elements of L'''. The same coupling is achieved when the elements b and t are removed from the input sequence.

If R is a singleton block $\{s\}$ we represent the input sequence as L', s, L''. Suppose that s', the base of the involution arising when L' just has been processed, is defined. Then at the next step the base moves to s and never returns to s' thereafter. Each time the base is at s the next element (if any) becomes a new base. Therefore, the elements of L' do not couple with the elements of L''. But removing s leads to the same outcome, because s' never gets coupled and therefore the elements of L' do not couple with the elements of L'' as well.

A similar argument also applies in the case when $|L'|$ is even and the restriction of the involution onto L' is a perfect matching.□

Now the dual algorithm is easy to justify. Indeed, coupling of the value 1 agrees with the solitaire rules. Removing the block containing value 1 from the sequence and using reducibility we see that coupling of the minimal element of the reduced sequence is the same for both principal and dual algorithms. Proceeding by induction we see that both algorithms result in the same involution.

5 Counting singletons. Let S_n be the set of singleton values resulting from processing random permutation π_n and

$$\sigma_n(k) := P(|S_n| = k)$$

be the probability distribution of the number of singletons.

We have

$$\sigma_n(k) = \frac{1}{n}\sigma_{n-1}(k-1) + \frac{n-1}{n}\sigma_{n-2}(k). \tag{7}$$

(with $\sigma_n(-1) = 0$). To see this suppose $|S_n| = k$ and consider the block R containing value 1. If value 1 is a singleton, the involution on $[n] \setminus R$ has $k-1$ singletons, otherwise this number is k. Clearly, $P(1 \in S_n) = 1/n$ and the conditional probability that the restricted involution has $k-1$ singletons is $\sigma_{n-1}(k-1)$, as follows from Proposition 4 and the familiar *sampling property* of π_n: given that the values $u_1, \ldots, u_j$ occupy positions $q_1, \ldots, q_j$, any arrangement of the remaining values into the remaining positions has probability $1/(n-j)!$. The second term in (7) is derived similarly, when the value 1 is a bottom.

The recursion is easy to apply. For instance, the perfect matching formula (1) follows as a particular case $k = 0$ (n even). Averaging in (7) we see that the expected number of singletons is

$$E|S_n| = E|S_{n+1}| = 1 + \frac{1}{3} + \frac{1}{5} + \ldots + \frac{1}{n} \qquad (n \ \text{ odd}).$$

We establish next that the number of singletons has the same distribution as the number of odd-size cycles in a random permutation. To this end, consider a permutation with k odd cycles and let $s_n(k)$ be the number of such permutations. If (1) is a one-element cycle the remaining $n-1$ elements must comprise a permutation with $k-1$ odd cycles. Otherwise discarding 1 and its follower we obtain a partition with k odd cycles, and there are $(n-1)^2$ ways to choose a follower of 1 along with their position in a cycle. This leads to

$$s_n(k) = s_{n-1}(k-1) + (n-1)^2 s_{n-2}(k),$$

and dividing by $n!$ we get a recursion coinciding with (7).

The coincidence allows to translate known results on cycle structure of permutations (see[7], [8]) into facts about the involution. Adapting a result from [7], we have a generating function

$$\sum_{k=0}^{n} \sigma_n(k)x^k = \sum_{k=0}^{n} \binom{(x-1)/2}{n-k}\binom{(x-1)/2+k}{k}$$

and a central limit theorem:

Theorem 5 *As $n \to \infty$ both the variance and the expectation of $|S_n|$ are asymptotic to $(\log n)/2$. The probability distribution of the scaled number of singletons*

$$\frac{|S_n| - \frac{1}{2}\log n}{\sqrt{\frac{1}{2}\log n}}$$

approaches the standard normal distribution.

6 Singleton values. The number of singletons does not tell much about their values. One can ask about the maximum, the minimum or a typical singleton value. In this section we view S_n as a random subset of $[n]$ and seek to study its asymptotic pattern.

The exact probability law of S_n is complicated, because the events $\{v \in S_n\}$ are not independent for different v. (One obvious obstruction for independence is that $|S_n|$ must have the same parity as n.) Nevertheless, we will see that independence appears asymptotically, opening the way to Poisson approximation.

Our method is to compare the set of singleton values with another random set $\widehat{S}_n \subset [n]$ which has the events $\{v \in \widehat{S}_n\}$ independent exactly and satisfies

$$P(v \in \widehat{S}_n) = (n - v + 1)^{-1}.$$

Different realizations of $\widehat{S}_n$ have appeared in the literature in connection with random permutations, see [3] for interpretations as record times, record values, new-cycle indices and others. It is known that the distribution of $|\widehat{S}_n|$ approaches the normal distribution with expectation and variance $\log n$. Another kind of asymptotics appears if the random set $\widehat{S}_n$ is scaled by the factor n^{-1}, then the probability law approaches that of a Poisson point process $\mathcal{P}$ which lives on $[0, 1[$ and has intensity measure $\mathrm{d}x/(1-x)$.

For future reference we recall how the points of $\mathcal{P}$ can be represented via a classical stick-breaking procedure. The minimum point is obtained by breaking the unit 'stick' $[0, 1]$ in two pieces at a uniform random location, then the left piece remains indivisible, and the right piece is broken at a uniform random location again, etc. The divisible part becomes smaller after each step, so that the points of $\mathcal{P}$ accumulate at 1.

We will find it convenient to deal with a special realization of $\widehat{S}_n$ associated with the dual algorithm. Let ν_n be the random value defined by setting $\nu_n = v$ if at step v the set of free positions becomes empty for the first time. (Clearly, ν_n is a stopping time adapted to $p_1, \ldots, p_n$.) Define further $\widehat{S}_n$ by setting $v \in \widehat{S}_n$ if either $v \le \nu_n$ and p_v is the highest free position remaining after step $v-1$ of the dual algorithm, or $v > \nu_n$ and p_v is the highest reserved position remaining after step $v-1$. Recalling the course of the algorithm we see that

(8) $$v \in S_n \Longleftrightarrow v \in \widehat{S}_n \text{ and } v \le \nu_n$$

that is, the two sets coincide until there is at least one free position available. On completing step ν_n of the dual algorithm the uncertainty is largely reduced: all values $v = \nu_n + 1, \ldots, n$ are tops which should be attributed to the relevant bottoms.

From the sampling property of π_n mentioned in Section 5 follows that $\widehat{S}_n$ has the required probability law. A key fact is that ν_n and $\widehat{S}_n$ are independent as follows by noting that the event $\{v \in \widehat{S}_n\}$ occurs if at step v a certain position is chosen at random out of $n - v + 1$ possibilities, whatever the allocation of free positions. Therefore (8) implies

$$P(v \in S_n) = (n - v + 1)^{-1} P(\nu_n \geq v). \tag{9}$$

The events $\{v \in S_n\}$ are independent exactly for $v \leq n/2$ if n is even, and for $v \leq (n+1)/2$ if n is odd. For the 'second half' of values v the distortion is caused by the shortage of free positions. In view of this it is important to have information about typical values of ν_n. Remarkably, there is a simple formula for the distribution

$$q_n(v) := P(\nu_n = v).$$

Proposition 6 *For n even we have*

$$q_n(v) = q_{n+1}(v) = \binom{v}{n/2} 2^{-v} \qquad n/2 \leq v \leq n. \tag{10}$$

Proof. Let R be the block containing value 1. Suppose $\nu_n = v$, then the values $v+1, \ldots, n$ are tops. If $1 \in S_n$ then $p_1 = n$ and $v+1, \ldots, n$ must be tops of the restricted involution on $[n] \setminus R$ as well. If value 1 is a bottom then there are $n-1$ choices for p_1 and $n-1$ choices for the mate top which can be among the values $w > v$ or not; in the first case the reduced involution ends with $n - v - 1$ tops and in the second with $n - v$. Counting possibilities and arguing as in Section 5 we get

$$q_n(v) = \frac{1}{n} q_{n-1}(v-1) + \frac{n-v}{n} q_{n-2}(v-1) + \frac{v-2}{n} q_{n-2}(v-2)\,.$$

A direct check shows that (10) solves the recursion and the claim follows by the uniqueness of solution. □

Remark. Observe that $\nu_n = n$ holds when the value n is a singleton. Thus (10) implies that for n even $P(n \in S_n)$ is given by (1) (this probability is really ubiquitous!). We can compute this probability differently, by noting that the value n can be a singleton if and only if its position p_n is odd, and the values in lower positions $1, 2, \ldots, p_n - 1$ are partitioned into couples. Recalling that the value n is equally likely to appear in any of n positions and using (1) and (10) we obtain as a by-product the identity

$$\sum_{k=0}^{m} \binom{2k}{k} 2^{-2k} = (2m+1) \binom{2m}{m} 2^{-2m}$$

(where we substituted odd $n = 2m + 1$).

Applying Stirling's formula to (10) it is not hard to show that $(n-\nu_n)n^{-1/2} \xrightarrow{d} \xi$ where ξ is a random variable with density

$$P(\xi \in x) = 2\phi(x)\mathrm{d}x, \quad x > 0 \tag{11}$$

and ϕ is the standard normal density. On the other hand, for the scaled set $n^{-1}\widehat{S}_n = \{v/n : v \in \widehat{S}_n\}$ we have a weak limit $n^{-1}\widehat{S}_n \xrightarrow{d} \mathcal{P}$, which taken together with the approximation to ν_n results in the convergence of scaled S_n.

Theorem 7 *As $n \to \infty$ we have*

$$n^{-1}S_n \xrightarrow{d} \mathcal{P}. \tag{12}$$

Proof. In view of (8), for any $x \in [0,1[$ we have $S_n \cap [0, xn] = \widehat{S}_n \cap [0, xn]$ provided $\nu_n > xn$, and by (11) $\nu_n/n \to 1$ (in probability), whence these sets coincide with probability going to one, and their weak limits on the $1/n$-scale must be the same. Letting $x \to 1$ completes the proof.□

From the stick-breaking interpretation we know that the minimum point of $\mathcal{P}$ is uniformly distributed, thus the theorem implies that the asymptotic distribution of the least singleton is uniform.

Approximation (12) discards a major distinction between S_n and $\widehat{S}_n$ in the range of values $n - v \asymp n^{1/2}$. To stress the difference we need another scaling. To this end, let $\mathcal{P}_+$ be a Poisson process on $]0, \infty[$ with intensity measure $x^{-1}\mathrm{d}x$. Using independence properties of $\widehat{S}_n$ and examining its intensity measure we get $n^{-1/2}(n - \widehat{S}_n) \xrightarrow{d} \mathcal{P}_+$. The corresponding scaled limit of S_n is a thinned version of $\mathcal{P}_+$.

Theorem 8 *As $n \to \infty$ we have*

$$n^{-1/2}(n - S_n) \xrightarrow{d} \mathcal{P}_+ \cap [\xi, \infty[$$

where random variable ξ is independent of $\mathcal{P}_+$ and has distribution (11).

The result says that the limit configuration coincides with $\mathcal{P}_+$ to the right of the random cut ξ, and has no points to the left of ξ. The proof follows at once from the convergence of S_n and independence of ν_n and S_n.

We turn now to the distribution of the maximum singleton. The minimum point of $\mathcal{P}_+ \cap [\xi, \infty[$ has a distribution with density $x^{-2}(2/\pi)^{1/2}(1 - \exp(-x^2/2)), x > 0$, as derived from (11) and the fact that $\mathcal{P}_+$ has no points within $[x, y]$ with probability x/y. By Theorem 8 this is also the limit density of $n^{-1/2}(n - \max S_n)$.

The connections between $\widehat{S}_n$ and S_n become clear. These sets coincide within $[1, \nu_n]$, where a typical difference $n - \nu_n$ is of the order of $n^{1/2}$, and within the range $v > \nu_n$ the set $\widehat{S}_n$ has about $(\log n)/2$ elements, while S_n is empty. This explains why the cardinality of $\widehat{S}_n$ is about twice the cardinality of S_n.

7 Motzkin path. With involution on $[n]$ we associate a lattice *Motzkin path* $M_n(v), v = 0, 1, \ldots, n$ which starts at $(0, 0)$, ends at $(n, 0)$ and has n links. The vth increment of the path is $(1, 1), (1, 0)$ or $(1, -1)$ if value v is a bottom, a top or a singleton, respectively. Figure 4 shows the Motzkin path for our running example.

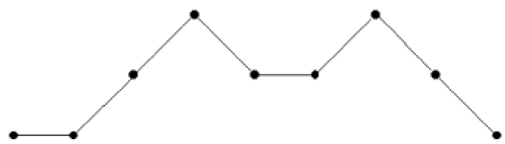

Figure 4: Motzkin path

$M_n(v)$ is equal to the number of couples (b, t) which are separated by v in the sense that $b \leq v < t$. Remarkably, the Motzkin path associated with our random involution is a Markov chain.

Proposition 9 *$M_n(v)$ is a time-inhomogeneous Markov chain in $v = 0, 1, \ldots, n$, with $M_n(0) = 0$ and $M_n(n) = n$. Given $M_n(v) = h$ with $h < n - v$, the path is incremented at step step $v + 1$ by $(1, 1), (1, 0)$ or $(1, -1)$ with probability*

$$\frac{n-h-v-1}{n-v}\,, \quad \frac{1}{n-v}\,, \quad \text{or} \quad \frac{h}{n-v}$$

respectively. For $h = n - v$ the only possible increment is $(1, -1)$.

Proof. In terms of the dual algorithm $M_n(v)$ is the number of reserved positions and $n - v$ is the joint number of free and reserved positions. The Markov property follows from the sampling property of π_n, which implies that the conditional distribution of p_v given $p_1, \ldots, p_{v-1}$ is uniform within the set of unoccupied positions.

In the event $M_n(v) = h$ there are h reserved and $n - h - v$ free positions. In case $v + h < n$ at least one position is free and $v + 1$ is a singleton with probability $1/(n - v)$, a top with probability $(n - h - v - 1)/(n - v)$ or a bottom with probability $h/(n - v)$. In case $v + h = n$ all positions are either occupied or reserved and a top is the only possibility. □

We show next that the scaled Motzkin path has a limiting parabolic shape. For unified treatment we define $M_n(v)$ also for noninteger $v \in [0, n]$ by linear interpolation from the integer values.

Theorem 10 *For $n \to \infty$ and arbitrary $\epsilon > 0$*

$$P\left(\sup_{x \in [0,1]} |n^{-1} M_n(nx) - x(1 - x)| < \epsilon\right) \to 1.$$

Proof. Viewing v and h as 'time' and 'state' variables, the phase space of the process scaled by the factor $1/n$ becomes a lattice with step $1/n$ within the triangle with vertices $(0,0),(1/2,1/2)$ and $(1,0)$. The increment of the scaled process over the time interval $[v/n,(v+1)/n]$ is at most 1 and the average increment is, by Proposition 9,

$$\frac{1}{n}\left(1-\frac{2y}{1-x}\right)+o\left(\frac{1}{n}\right),$$

provided the process starts at $x=v/n, y=h/n$. The o-term in this formula is uniform in x,y for $x<1-\delta$.

By Theorem 4.7 of [5] the path of the scaled process is with high probability uniformly close to a solution of the differential equation

$$y'=1-\frac{2y}{1-x}, \qquad x\in[0,1-\delta].$$

To check formally the assumptions of the theorem we should take $\alpha_n=n,\epsilon_n=1/n, F(x,y)=1-2y/(1-x)$, in the notation of the cited paper. The function F satisfies the required Lipshitz condition in the indicated range.

Instead of giving technical details we comment briefly on the nature of this result (see also [4], p. 517). For continuity reasons the distribution of the increment over a $1/n$-interval is not sensitive to small variation of the initial point in the phase space, thus locally the process has almost i.i.d. increments. These increments are of the order of $1/n$ and their variance is of the order of $1/n^2$, whence locally the path (comprised of, say ϵn, increments) with high probability must be close to a straight line (the law of large numbers). Putting the local approximations together shows that the path approaches an appropriate integral curve of the differential equation.

The unique solution of the differential equation with initial condition $y(0)=0$ is $y(x)=x(1-x)$, therefore this is also the limit shape of the scaled path. Letting $\delta\to 0$ we obtain convergence in the full range $x\in[0,1]$.□

8 Interacting particle system. We switch to another kind of model which is based on an infinite sequence of random values $X_1,X_2,\dots$ sampled independently from the uniform distribution on the unit interval. Applying the solitaire rules we obtain a consistent sequence of involutions on finite sets $\{X_1,\dots,X_n\}$. The continuous model provides a unified framework for treatment of the involutions of different size and allows study of asymptotic properties which hold with probability one.

It is very intuitive to think of the set of singletons as a system of particles which interact in the sense that the probability of birth or death of a particle at a certain site is determined by the current configuration. At time n there is a configuration of particles $\mathcal{S}_n$ on the unit interval, with the *base* $Z_n=\min\mathcal{S}_n$ (for convenience, the base of the empty configuration $\emptyset$ is defined to be 1). the process starts with some configuration $\mathcal{S}_0$ and the solitaire rules become the rules for evolution of the particle system. At time n a new particle appears at a random location X_n and one of the following changes can occur:

(S): the particle joins the existing population if $X_n<Z_{n-1}$,

(C): or the particle gets coupled with the base and at the same time both particles disappear if $X_n > Z_{n-1}$.

We are interested in the long time behaviour of the system. Because the same transition mechanism is used at each step, the particle system is a time-homogeneous Markov chain. Basically we are interested in the system which starts with the empty configuration and has therefore at most n particles at time n, but it will be necessary also to consider some infinite configurations which have an accumulation point at 1.

Dynamics of the system are determined essentially by fluctuations of the base. When the base is at z the configuration to the left of z is empty, and at the next step no particles to the right of z can appear or die. A related distinguished feature which makes the system amenable to study is what we call the *one-sided interaction*: the probability that at the next time a particle appears or dies at a given location depends solely on the partial configuration to the left of this location.

The chain has one recurrent state $\emptyset$, all other states are transient.

Proposition 11 *The empty configuration $\emptyset$ is a null-recurrent state of the chain.*

Proof. Returns to $\emptyset$ constitute a discrete renewal process. On the other hand, the event $\{\mathcal{S}_n = \emptyset\}$ is just another expression for a perfect matching, whence its probability is given by (1). This is the same as the return probability for the simple random walk, whence the recurrence time distributions are the same. The result follows by recalling that the random walk is null-recurrent.□

(One curious phenomenon associated with the infinite expectation is that the population which starts with $\mathcal{S}_0 = \emptyset$ dies out from time to time despite of the fact that the average number of particles grows as $(\log n)/2$).

For any state of the chain the number of returns to this state and the recurrence time depend only on the location of the base, as it follows by virtue of the one-sided interaction. If the base is initially at $z \in [0, 1[$ then the number of returns to the initial configuration is geometrically distributed with parameter $1-z$. And the recurrence time is the *life-time at z* defined as

$$\rho(z) = \min\{n : Z_n > z\}$$

(given $Z_0 = z$). The life-times are nondecreasing in z and have the following probability distribution.

Proposition 12 *The life-time at z attains only odd values and their probabilities are*

$$r_{2k+1}(z) := P(\rho(z) = 2k+1) = (1-z)\left(\frac{z}{2}\right)^k \left(1 - \frac{z}{2}\right)^k \binom{2k}{k} \tag{13}$$

The expected life-time in z is $E\rho(z) = (1-z)^{-2}$.

Proof. Suppose $Z_0 = z$. If the first particle to appear comes to the right of z, then $\rho(z) = 1$, and thus $r_1(z) = 1 - z$ in full agreement with (13).

If the new particle appears at some location $y < z$, the base moves to y. Having spent some time in $[0, y]$ the base returns to z, whence the life-time at z can be decomposed as $1 + \rho_1(y) + \rho_2(z)$ where the ρ_j's are two independent replicas of $\rho(\cdot)$. Averaging we obtain

$$r_{2k+1}(z) = \int_0^z \sum_{\substack{i,j \geq 0 \\ i+j=k-1}} r_{2i+1}(z) r_{2j+1}(y) \mathrm{d}y. \tag{14}$$

Assuming by induction that (13) holds for all $j < k$ we have

$$\int_0^z r_{2j+1}(y)\mathrm{d}y = \frac{2}{j+1}\binom{2j}{j}\left(\frac{z}{2}\right)^{j+1}\left(1-\frac{z}{2}\right)^{j+1}$$

as can be seen by taking derivatives and checking for $z = 0$. Substituting this equation into (14) and using the identity

$$\sum_{j=0}^{k} \frac{2}{j+1}\binom{2j}{j}\binom{2k-2j}{k-j} = \binom{2k+2}{k+1}$$

(which is the first recurrence time decomposition of the random walk) we get (13).

Putting $s = 2(z/2)^{1/2}(1 - z/2)^{1/2}$ into the binomial identity

$$(1-s^2)^{-1/2} = \sum_{k=0}^{\infty}\binom{2k}{k}\left(\frac{s}{2}\right)^{2k}$$

we see that the probabilities sum up to unity and $\rho(z)$ is therefore finite with probability one. The expectation can be computed directly, but it is more instructive to make use of the above decomposition to derive an integral equation for $g(z) = E\rho(z)$,

$$g(z) = (1-z) + \int_0^z (g(z) + g(y) + 1)dy$$

whose unique solution is $g(z) = (1 - z)^{-2}$. $\square$

From Proposition 12 follows that the state $\emptyset$ is accessible with probability one from any finite configuration. Indeed, the particles disappear in the increasing order, and since the interaction is only one-sided the total life-time of a finite configuration is the sum of independent individual life-times over the particles comprising the configuration.

Remark. None of the singletons on $]z, 1[$ can become a bottom as long as the base walks within $[0, z]$, therefore the quantity

$$\left(\frac{z}{2}\right)^k\left(1-\frac{z}{2}\right)^k\binom{2k}{k},$$

appearing in (13) is the probability that the algorithm yields a perfect matching on $2k$ random points, with all bottoms smaller than z. As $z \to 1$ we get the unconstrained perfect matching probability (1).

9 Equilibrium. A random configuration $\mathcal{S}_0$ is *equlibrium* if all $\mathcal{S}_n$ have the same distribution as $\mathcal{S}_0$, which means that this distribution is a finite invariant measure for the Markov chain. A general result on Markov chains (see [6], p. 234) implies that there is no equilibrium with finitely many particles, because the chain has a null-recurrent state $\emptyset$ accessible from all finite configurations. We turn therefore to infinite configurations. In view of Theorem 7 one could guess that the Poisson configuration $\mathcal{P}$ is equilibrium.

Theorem 13 *The Poisson point process $\mathcal{P}$ on $[0,1[$ with intensity measure $\mathrm{d}x/(1-x)$ is the unique equilibrium configuration.*

Proof. A complete proof of the equilibrium property appears in [2] and need not be reproduced here. We briefly illustrate the idea by showing that for $\mathcal{S}_0 = \mathcal{P}$ one iteration does not alter the distribution of the base. Recalling the stick-breaking construction of $\mathcal{P}$, the distribution of its base Z_0 is standard uniform, and given $Z_0 = z$ the distribution of the second smallest particle, say Y, is uniform on $[z,1]$. Now, if the rule (S) applies the base goes to a random uniform point in $[0,z]$, and if (C) it goes to Y. Whatever z, the conditional distribution of Z_1 is comprised of two uniform pieces, whence it is standard uniform. This implies that the unconditional distribution is also uniform, since the value of z does not matter.

Proving uniqueness is based on specific features of our particle system. For $x \in]0,1[$ let $\mathcal{S}_n^{(x)} = \mathcal{S}_n \cap [0,x]$ be the partial configuration within $[0,x]$. By virtue of the one-sided interaction the sequence $(\mathcal{S}_n^{(x)})$ is itself a Markov chain. Let $\emptyset^{(x)}$ be the empty configuration on $[0,x]$ and $\mathcal{P}^{(x)} = \mathcal{P} \cap [0,x]$ be the reduced Poisson configuration.

Using the same argument which led us to accessibility of $\emptyset$ (for $\mathcal{S}_n$) we conclude that any particle configuration on $[0,x]$ disappears in a finite time. Moreover, this time has a finite expectation, because it is the sum of individual life-times, whose expectations are finite, by Proposition 12. It follows that the reduced chain $\mathcal{S}_n^{(x)}$ has the *positive*-recurrent state $\emptyset^{(x)}$, accessible from everywhere with probability one. By the cited general result (see [6], p. 234) the reduced chain must have a unique equilibrium configuration. Obviously, this is $\mathcal{P}^{(x)}$, because truncating $\mathcal{P}$ retains the equilibrium property.

The proof is finished by observing that the unique configuration on $[0,1[$ compatible with all $\mathcal{P}^{(x)}$'s is $\mathcal{P}$, because the probability law of a configuration on $[0,1[$ is uniquely determined by the laws of reduced configurations on the intervals $[0,x]$ for $x < 1$.□

10 Laws of large numbers for involutions. We proceed with the continuous model assuming $\mathcal{S}_0 = \emptyset$.

Proposition 14 *For any $x \in]0,1[$ the total variation distance between $\mathcal{S}_n^{(x)}$ and $\mathcal{P}^{(x)}$ goes to zero. Hence $\mathcal{S}_n \xrightarrow{d} \mathcal{P}$.*

Proof. We wish to apply an ergodic theorem for Markov chains. The state $\emptyset^{(x)}$ has a $\mathcal{P}^{(x)}$-positive measure, and any $\mathcal{P}^{(x)}$-positive set of configurations with k particles on $[0,x]$ can be accessed from $\emptyset^{(x)}$ in k steps with positive probability. It follows that the chain $(\mathcal{S}_n^{(x)})$ is Harris-positive recurrent. Also, this chain is aperiodic, since the transition from $\emptyset^{(x)}$ to itself occurs with positive probability $1-x$. The convergence in total variation now follows from [6], p. 323.

The second statement follows from the first and the definition of weak convergence of point processes on $[0,1[$ (which amounts to convergence of joint distribution of particle counts within any finite system of nonintersecting intervals bounded from 1).□

(The convergence in total variation does not hold for $(\mathcal{S}_n)$ starting with a finite $\mathcal{S}_0$; in this case the event $|\mathcal{S}_n| < \infty$ has probability one while its $\mathcal{P}$-probability is zero.)

To describe the asymptotic pattern of the couples configuration of the involution on $\{X_1, \dots, X_n\}$ we represent each couple (X_i, X_j) by a point in $\nabla = \{(x,y) : 0 < x < y < 1\}$ and consider the 'empirical measure'

$$\mu_n(Q) = n^{-1} |\{i, j \leq n : (X_i, X_j) \in Q \text{ and they are coupled}\}|, \qquad Q \subset \nabla$$

which counts the couples at step n.

Proposition 15 *For any measurable $Q \subset \nabla$, $\mu_n(Q)$ converge to the Lebesgue measure of Q, with probability one.*

Proof. It is sufficient to consider Q bounded away from the vertex $(1,1)$. At step n a point is added if $(Z_{n-1}, X_n) \in Q$. Under the invariant measure (Z_{n-1}, X_n) is uniform in the unit square, whence the statement follows by applying the ergodic theorem to the Markov chain $(\mathcal{S}_n^{(x)}, X_{n+1})$ (as in [6], p. 417), with x selected sufficiently close to 1. □

A link between the continuous model and the random permutation model is established by using the ranks

$$v_{k,n} = |\{j \leq k : X_j \leq X_k\}|, \qquad k = 1, \dots, n$$

as substitutes for the X_k's. Since ties have only probability zero there is no ambiguity in the definition of ranks and the exchangeability of X_k's implies that $v_{1,n}, \dots, v_{n,n}$ is a uniform random permutation. Passing to ranks commutes with the algorithm, thus the involution on $[n]$ induced by $\{X_1, \dots, X_n\}$ is the same as when the algorithm is applied to the related sequence of ranks $v_{1,n}, \dots, v_{n,n}$.

A major advantage of introducing random permutations in this manner is that we have now involutions on $[n]$ defined simultaneously for all n on a single probability space. Asymptotic results for these involutions follow from a consequence of the Glivenko–Cantelli theorem which states that the difference between the X_k's and the scaled ranks becomes uniformly small:

$$\max_{k \leq n} |v_{k,n}/n - X_k| \to 0, \quad \text{as} \quad n \to \infty, \tag{15}$$

with probability one. In this framework, Theorem 7 on the configuration of singletons can be seen as a consequence of Proposition 14. Our second principal result on the configuration of couples follows from (15) and Proposition 15.

Theorem 16 *Let $\widehat{\mu}_n$ be a random measure on ∇ obtained by placing a mass $1/n$ at $(u/n, v/n)$ for each couple (u, v) of the involution on $[n]$. The measures $\widehat{\mu}_n$ converge weakly to the Lebesgue measure on ∇, with probability one.*

We can also treat the associated Motzkin paths as defined for all n on a single probability space. In this setting a stronger version of Theorem 10 is valid: the uniform convergence of $n^{-1}M_n(nx)$ to $h(x) = x(1-x)$ holds with probability one. This follows by noting that a couple (b, t) is separated by $v = nx$ if its representing point $(b/n, t/n)$ is in the rectangle $\{(\beta, \tau) \in \nabla : \beta \leq x < \tau\}$, whose area is $h(x)$.

Acknowledgement. The second author appreciates the hospitality of the Institute of Mathematical Stochastics, University of Göttingen, during the fall semester of 1999, when this paper was initiated. At that time S. Kerov held the Mercator Professorship supported by the Deutsche Forschungsgemeinschaft. Helpful comments of the referees are greatfully acknowledged.

References

[1] Gnedin, A. and Kerov, S. (2000) The Plancherel measure of the Young–Fibonacci graph. *Math. Proc. Camb. Phil. Soc.* **129**, 433-446.

[2] Gnedin, A. and Kerov, S. (2001) A characterization of GEM distributions. *Combinatorics, Probability and Computing* (to appear).

[3] Goldie, Ch. (1989) Records, permutations and the greatest convex minorants. *Math. Proc. Camb. Phil. Soc.* **106**, 169-177.

[4] Jacod, J. and Shiryaev, A.N. *Limit Theorems for Stochastic processes.* Springer, NY 1987.

[5] Kurtz, T. (1970) Solutions of ordinary differential equations as limits of pure jump Markov processes. *J. Appl. Prob.* **7**, 49-58.

[6] Meyn, S.P. and Tweedie, R.L. *Markov Chains and Stochastic Stability.* Springer, NY 1996.

[7] Sachkov, V. *Probabilistic Methods in Combinatorial Analysis*, Cambridge Univ. Press, 1997.

[8] Wilf, H. *generatingfunctionology*, Academic Press, NY, 1990.