Transforming Inductive Proofs to Bijective Proofs

Nathaniel Shar
Stanford University
nshar@stanford.edu

June 2010

1 Introduction

A bijective proof of an identity $A = B$ is a pair of sets S, T and a bijection $\phi : S \to T$, along with a proof that $|S| = A$ and $|T| = B$.

Bijective proofs are prized by many mathematicians because they often provide natural explanations for combinatorial identites. For example, the identity

$$\sum_{k=0}^{n} \frac{n!}{k!(n-k)!} = 2^n \tag{1}$$

is not obvious as a relation among the integers, but has a natural bijective explanation. Namely, let S_k be the set of k-subsets of $[n]$. (Here, $[n]$ denotes the set of positive integers less than or equal to n.) Then $|S_k| = \frac{n!}{k!(n-k)!}$, so the left side of the identity is $\sum_{k=0}^{n} |S_k|$. Since the sets S_k are disjoint, this is equal to $|\bigcup_{k=0}^{n} S_k| = |\mathcal{P}([n])| = 2^n$, which completes the proof.

Of course, there are multiple proofs of (1), many of which are often claimed to be the most natural explanation. For example, if one desires a proof using generating functions, one can expand $(x+1)^n$, and then set $x = 1$. This is typical; it is rare for the only proof of an identity to be bijective. In fact, the Wilf-Zeilberger (WZ) method provides an algorithm for proving a wide variety of combinatorial identities, meaning that the first proof of an identity is almsot never bijective. Nevertheless, bijective proofs are often more elegant than proofs by analytic methods.

Many familiar proofs are simply bijective proofs in disguise. Consider, for example, the famous "proof without words" of the fact

$$\sum_{k=1}^{n} (2n-1) = n^2 \tag{2}$$

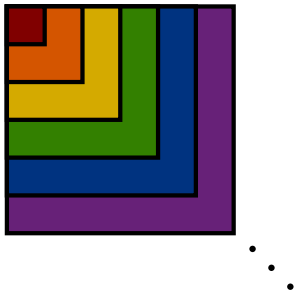

Figure 1: A "proof without words" that n^2 is the sum of the first n positive odd integers.

depicted in Figure 1. If you pay attention to the colors, the figure appears to be the disjoint union S of sets S_k each containing $(2k-1)$ unit squares, where k ranges over the elements of $[n]$. If you ignore the colors, the figure appears to be a set T of n^2 unit squares. There is clearly a bijection between these sets: namely, $\phi : S \to T$ is the function that forgets colors. Unfortunately, to actually show that ϕ is a bijection — for example, by exhibiting its inverse — is quite involved. This is a common feature of bijective proofs: it is often much easier to "see" that the function is a bijection than it is to formally prove it.

In Section 2, we construct a bijective proof of (2) to illustrate the use of the "translation method," a relatively mechanical method for producing bijective proofs of certain identities. In Section 3, we apply the translation method to produce the first known bijective proof of a Fibonacci number identity of Vajda. Then, in Section 4, we explain why the translation method is not successful at producing a bijective proof of a formula counting binary De Bruijn cycles. Finally, in Section 5, we describe a recently-published bijective proof (not using the translation method) of thesame formula counting binary De Bruijn cycles.

2 Induction and bijections

The identity (2) is often proved by induction. As we will see, the inductive proof is closely related to the bijective proof. First, let us review the inductive proof carefully, writing out every step. In the base case, the identity reduces to $1 = 1^2$. For the inductive case, if we assume that the identity

holds for n, then

$$\begin{aligned}\sum_{k=1}^{n+1}(2n-1) &= \sum_{k=1}^{n}(2n-1)+(2n+1)\\ &= n^2+(2n+1)\\ &= (n+1)^2,\end{aligned}$$

which completes the proof.

Now, let us pretend that the geometric intuition of Figure 1 has escaped us. In an attempt to generate a bijective proof of (2) anyway, let us view each of the steps of this inductive proof as a statement about sets. We first introduce some notation. For any two sets, let $A \sim B$ denote the statement "A bijection exists between A and B," and let $A \sim_\phi B$ denote the statement "ϕ is a bijection between A and B."

Now, the base case of our induction is the statement that $\{1\} \sim (\{1\} \times \{1\})$, which is obvious. For the inductive case, we begin with the assumption that a bijection ϕ_n exists from $[1] \sqcup [3] \sqcup \cdots \sqcup [2n-1]$ to $[n] \times [n]$, where $A \sqcup B$ denotes the disjoint union of sets A and B. Then, in the first equality,

$$\sum_{k=1}^{n+1}(2n-1) = \sum_{k=1}^{n}(2n-1)+(2n+1),$$

we note that

$$[1] \sqcup [3] \sqcup \cdots \sqcup [2n+1] \sim_\zeta ([1] \sqcup [3] \sqcup \cdots \sqcup [2n-1]) \sqcup [2n+1],$$

where ζ is the obvious bijection. In the second equality, we use the existence of ϕ_n to say that

$$([1] \sqcup [3] \sqcup \cdots \sqcup [2n-1]) \sqcup [2n+1] \sim_\psi ([n] \times [n]) \sqcup [2n+1].$$

Here, ψ maps any element of $[1] \sqcup [3] \sqcup \cdots \sqcup [2n-1]$ to an element of $[n] \times [n]$ by applying ϕ_n to it, while it leaves any element of $[2n+1]$ alone. Finally, the last equality states that

$$([n] \times [n]) \sqcup [2n+1] \sim_\xi [n+1] \times [n+1].$$

Here, ξ is defined as follows:

$$\xi(s) = \begin{cases} (x,y) & \text{if } [n] \times [n] \text{ and } s = \phi_n^{-1}(x \times y) \\ (n+1,s) & \text{if } s \in [2n+1] \text{ and } s \leq n \\ (n+1,n+1) & \text{if } s \in [2n+1] \text{ and } s = n+1 \\ (2n+2-s,n+1) & \text{if } s \in [2n+1] \text{ and } s \geq n+2. \end{cases}$$

If we let $\phi_{n+1} = \zeta \circ \psi \circ \xi$, then $\phi_{n+1} : [1] \sqcup [3] \sqcup \cdots \sqcup [2n+1] \to [n+1] \times [n+1]$. This completes the induction, and shows the existence of bijections ϕ_n for all n.

Unfortunately, the bijections ϕ_n (except, of course, for ϕ_1) are recursively defined in terms of ϕ_{n-1}. As a result, this proof of (2), while technically a "bijective proof," lacks the desirable qualities of bijective proofs, such as elegance and explanatory power. Nevertheless, the definition of ϕ_n is perfectly comprehensible to a computer. This means that, for example, we would have no trouble writing down the value of, say, ϕ_{100}, or its inverse, at any desired point. Now, by looking at these values, we can hope to come up with a non-recursive definition of ϕ_n or its inverse. In this case, we get lucky, because if we write down the values of ϕ_4^{-1}, we see the following:

$$[\phi_4^{-1}(i,j)] = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 3 & 2 & 2 & 2 \\ 5 & 4 & 3 & 3 \\ 7 & 6 & 5 & 4 \end{bmatrix}.$$

Seeing the obvious pattern, we conjecture that

$$\phi_n^{-1}(i,j) = \begin{cases} i & \text{if } i \leq j \\ 2i - j & \text{if } i > j. \end{cases}$$

This conjecture is easily proved by induction, so we now have a direct bijection.

The method we have used here is called the *translation method* and was described by Wood and Zeilberger [5]. The translation method can be summarized as a sequence of steps:

1. Prove the given identity by induction.
2. Fix n.
3. View one side (without loss of generality the left side) of the identity as counting the number of elements in a set S_0.
4. For the ith step in the inductive proof, create a corresponding set S_i and a bijection $\psi_{n,i}$ from S_{i-1} to S_i. When the step uses the induction hypothesis, it is permissible to recursively define $\psi_{n,i}$ in terms of functions ϕ_m, where $m < n$. If there are r steps in the induction, then the right side of the identity should count S_r.
5. Compose all the bijections $\psi_{n,i}$ to get ϕ_n.

6. Use any method to explore the behavior of ϕ_n in hopes of describing it without resorting to recursion.

Unsurprisingly, the last step is often the most difficult part of the process; it generally involves an insight or a stroke of luck. These steps can be carried out with the aid of computers; using a computer to evaluate the bijections ϕ_n and display the results is often the easiest way to explore the behavior of ϕ_n in step 6.

In step 4, a variety of tricks, familiar to combinatorialists, can be used create the sets S_i. For example, an identity of the form $\sum x_j$ corresponds to counting the set $\bigsqcup X_j$, where $|X_j| = x_j$, and an identity $\prod_{j=1}^{n} x_j$ counts the set $X_1 \times \cdots \times X_n$.

3 Application: Fibonacci number identities[1]

The Fibonacci numbers F_n are defined for $n \geq 0$ by the recurrence

$$F_n = \begin{cases} F_0 = 0 \\ F_1 = 1 \\ F_{n+2} = F_{n+1} + F_n \quad n \geq 0 \end{cases}.$$

Many surprising identities involving the Fibonacci numbers are known. In [1], Benjamin and Quinn provide bijective proofs of many such identities. They also provide a list of identities for which bijective proofs are not known. Many of these identities were drawn from a list created by Vajda [4], including the following identity:

$$-1 + \sum_{k=2}^{n} \frac{1}{F_{2^k}} = -\frac{F_{2^n-1}}{F_{2^n}} \qquad (V90).$$

An equivalent form, which we find more convenient to prove, is the following, which results from clearing the denominators, multiplying both sides by F_2 (which is equal to 1), and rearranging the result slightly:

$$F_{2^n-1} \prod_{1 \leq j \leq n-1} F_{2^j} + \sum_{k=1}^{n-1} \prod_{\substack{1 \leq j \leq n \\ j \neq k+1}} F_{2^j} = \prod_{1 \leq j \leq n} F_{2^j}, \qquad (1)$$

Identity (V90) is easily proved by induction. Using the techniques from [5] for translating proofs by induction into bijective proofs, we were able

[1] The contents of this chapter have been submitted for publication.

to find a bijective proof of (V90). In Section 2.1, we introduce sets whose cardinalities are Fibonacci numbers. In Section 2.2, we describe a "tail-swap maneuver," inspired by the tail-swapping bijections highlighted in [1], which serves as the building block of the bijective proof. In Section 2.3, we describe the bijection itself. Finally, in Section 2.4, we describe a family of identities which we believe to be new, and which can be proved by bijections very similar to those which we used to prove (V90).

3.1 Combinatorial interpretation of Fibonacci numbers

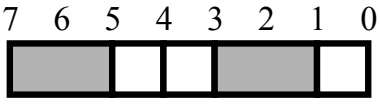

Figure 2: A tiling of length 7

Following [1], we note that the set of tilings, using only squares and dominoes, of a 1-by-n strip, has cardinality F_{n+1}. We will write T_n to denote the set of such tilings, and also write $f_n = |T_n|$. Such a tiling (where $n = 7$) is illustrated in Figure 2.

We generally identify the rightmost end of a tiling as "position 0," with positions $1, \ldots, n$ numbered from right to left. (This is different notation than is used in our sources, but it is more convenient for our purposes.) We say that a tile is at position k if its rightmost edge is at position k. For example, in Figure 2 there are dominoes at positions 1 and 5, and squares at positions 0, 3, and 4. We say that a tiling has a *fault* at position k, with $0 \leq k \leq n$, unless it has a domino at position $k-1$. For example, the tiling in Figure 2 has faults at all positions except 2 and 6.

When we consider multiple tilings at once, we consider them to be aligned vertically in some way. In case their rightmost edges are not aligned, will consider position 0 to be the rightmost edge of the tiling with the rightmost edge located furthest to the right. Given two tilings written one above the other, we say they have a *common fault* at position k unless at least one of them has a domino at position $k-1$.

3.2 The tail-swap maneuver

In our bijective proof of (1) we make use of a *tail-swap maneuver* (see [1]), illustrated in Figure 3.

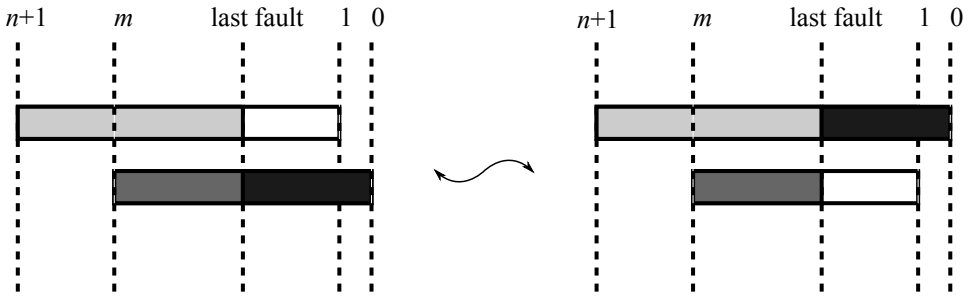

Figure 3: The tail-swap maneuver

Given integers $m, n \geq 0$, define

$$\tau_{m,n} : T_m \times T_n \to (T_m \cup T_{m+1}) \times (T_n \cup T_{n-1})$$

as follows. Given two tilings t_1 and t_2, of lengths m and n respectively, we write t_1 above t_2 so that the rightmost block of t_2 protrudes 1 unit beyond the rightmost block of t_1. (Thus, the rightmost edge of t_1 is at position 1, and the rightmost edge of t_2 is at position 0.) Then we locate the rightmost fault common to both tilings, if it exists. Any tiles to the right of this fault are swapped to the other tiling. After this process, we have a tiling t_1' of length $m+1$ and a tiling t_2' of length $n-1$. We write $\tau_{m,n}(t_1, t_2) = (t_1', t_2')$ and say that the tail-swap succeeded. If no common fault exists, we write $\tau_{m,n}(t_1, t_2) = (t_1, t_2)$. In this case, we say that the tail-swap failed.

Remark 1. For brevity, when the lengths of the tilings are known, we write τ in place of $\tau_{m,n}$. We will also refer to the two tilings that are elements of $\tau(t_1, t_2)$ as the first and second *resultant tilings* of the tail-swap.

Proposition 1. *Suppose t_1 and t_2 are tilings of lengths m and n, respectively, with $m \leq n$. If a tail swap of t_1 and t_2 fails, then m is even and t_1 consists only of dominoes. Furthermore, the rightmost $m/2+1$ tiles of t_2 are all dominoes.*

Proof. If t_1 contains a square, then it has two adjacent faults. Tiling t_2 must have a fault at one of these positions (if it did not, then it would have dominoes at two adjacent positions, which is impossible), so t_1 and t_2 have a common fault. Therefore, a tail swap of t_1 and t_2 does not fail. This proves that if a tail swap of t_1 and t_2 fails, then t_1 consists only of dominoes. In particular, m is even. Furthermore, there are no common faults between t_1 and t_2, so the rightmost $m/2+1$ tiles of t_2 must also be dominoes. □

The following proposition, which is apparent from Figure 3 , states that when a tail-swap succeeds, it can be reversed by performing a second tail-swap.

Proposition 2. *Given two tilings t_1 and t_2, if $\tau(t_1, t_2) = (t_1', t_2')$ is a successful tail swap, then $\tau(t_2', t_1') = (t_2, t_1)$.*

Proof. Because the tail-swap succeeded, after the tails of t_1 and t_2 are swapped, the result is that tiling t_1' is aligned 1 block to the right of t_2'. The rightmost common fault is still at the same location, so performing $\tau(t_2', t_1')$ simply swaps the tails back to their original positions. □

3.3 The bijection

We begin by describing sets whose cardinalities are the right- and left-hand-sides of (1). Let

$$S_1 = \prod_{1 \le j \le n} T_{2^j - 1};$$

then the cardinality of S is clearly

$$\prod_{1 \le j \le n} F_{2^j},$$

the right-hand side of (1). Let

$$S_2 = T_{2^n - 2} \times \prod_{2 \le j \le n-1} T_{2^j - 1} \cup \bigcup_{k=2}^{n} \prod_{\substack{2 \le j \le n \\ j \ne k}} T_{2^j - 1}. \tag{2}$$

Note that all of the unions in (2) are disjoint. Clearly the cardinality of S_2 is the left-hand-side of (1). We will establish a bijection between S_1 and S_2.

Before describing the bijection, we describe an iterative version of the tail swap process, called a multiple tail-swap, that takes place on a sequence of $k \ge 2$ tilings $t_1, \ldots, t_k$ of lengths $l_1, \ldots, l_k$. Define $\tau(t_1, \ldots, t_k)$ as follows. First find $\tau(t_1, t_2)$. If this tail swap fails, let $\tau(t_1, \ldots, t_k)$ be $(t_1, \ldots, t_k)$ If it succeeds, replace t_1 and t_2 with the first and second resultant tilings, respectively. Then find $\tau(t_2, t_3)$. If this tail-swap fails, let $\tau(t_1, \ldots, t_k)$ be $(t_1, \ldots, t_k)$. If it succeeds, replace t_2 and t_3 with the first and second resultant tilings and find $\tau(t_3, t_4)$. Continue in this way until every tiling has participated in a tail-swap, or until one of the tail-swaps has failed. If any tail-swap $\tau(t_r, t_{r+1})$ fails, then we say that the *failure index* of the multiple

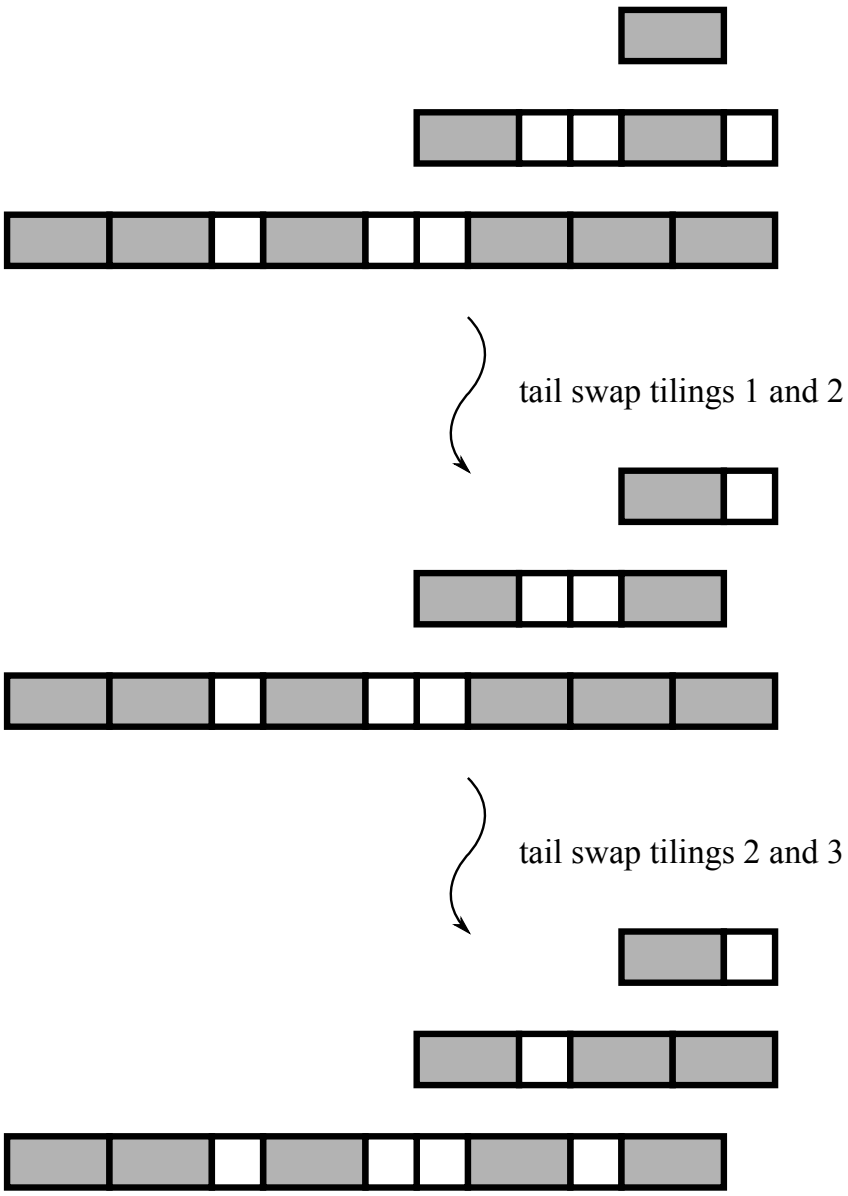

Figure 4: A multiple tail-swap on 3 tilings, in which the failure index is 3.

tail-swap is r and the process stops. If all of the successive tail-swaps succeed, the failure index is k. Notice that the failure index can be determined by comparing the lengths of the resultant tilings compared to their original lengths. In particular, if $2 \leq i \leq k$ and t_i has been shortened by one unit, then the failure index is i, and if all tilings still have their original lengths, then the failure index is 1.

An example of a multiple tail-swap, with $k = 3$, can be seen in Figure 4.

Notice that if $k = 2$, this is merely the original tail-swap operation. We now consider the reversibility of the multiple tail-swap.

Proposition 3. *If $\tau(t_1, \ldots, t_n) = (t'_1, \ldots, t'_n)$ is a multiple tail-swap with failure index $i > 1$, then $\tau(t'_i, \ldots, t'_1) = (t'_i, \ldots, t'_1)$ is a successful multiple tail-swap.*

Proof. Because the failure index is i, the first $i-1$ tail-swaps of the multiple tail-swap succeed, and then the ith either fails or, if $i = n$, is not performed because the last tiling has been swapped. Failed tail-swaps do not change any of the tilings, so the original multiple tail swap leaves the tilings in the state they are in after the $(i-1)$th tail-swap. Performing $\tau(t'_i, \ldots, t'_1)$ reverses these $(i-1)$ tail-swaps, which restores the sequence $(t_i, \ldots, t_1)$ by Proposition 2. □

Remark 2. Proposition 3 shows that

$$\tau : \prod_{k=1}^{n} T_{l_k} \to \left(\prod_{k=1}^{n} T_{l_k} \right) \cup \bigcup_{I=2}^{n-1} \left(T_{l_k+1} \times \prod_{1 \leq k \leq I-1} T_{l_k} \times T_{l_I-1} \times \prod_{I+1 \leq k \leq n} T_{l_k} \right)$$

is injective.

We next describe $\mathrm{img}(\tau)$ when the lengths of the tilings satisfy a certain condition.

Theorem 3. *Suppose l_1 is even and $l_2, \ldots, l_n$ are all odd, and that $l_1, \ldots, l_n$ is increasing. Given*

$$(t'_1, \ldots, t'_n) \in \left(\prod_{k=1}^{n} T_{l_k} \right) \cup \bigcup_{I=2}^{n-1} \left(T_{l_k+1} \times \prod_{1 \leq k \leq I-1} T_{l_k} \times T_{l_I-1} \times \prod_{I+1 \leq k \leq n} T_{l_k} \right),$$

where $t'_1, \ldots, t'_n$ have lengths $l'_1, \ldots, l'_n$ respectively, there exists $(t_1, \ldots, t_n) \in \prod_{k=1}^{n} T_{l_k}$ such that $\tau(t_1, \ldots, t_n) = (t'_1, \ldots, t'_n)$ if and only if one of the following is true:

Case 1: All the l'_i are equal to the l_i, and t'_1 consists entirely of dominoes, and the rightmost $l'_1/2$ tiles of t'_2 are dominoes, or

Case 2: $l'_I = l_I - 1$ for some $I \geq 2$, t'_I consists entirely of dominoes, and the rightmost $(l'_I)/2$ tiles of t'_{I+1} are dominoes.

Proof. If $(t_1, \ldots, t_n)$ exist, then consider the failure index of $\tau(t_1, \ldots, t_n)$. If the failure index is 1, then Case 1 holds, by Proposition 1. If the failure index is $I \geq 2$, then Case 2 holds by Proposition 1. So one of the three cases must hold.

Now, suppose that one of the cases holds. We will show that the appropriate choice of $(t_1, \ldots, t_n)$ exists. First, suppose that Case 1 holds. Then $(t_1, \ldots, t_n) = (t'_1, \ldots, t'_n)$ is an appropriate choice for $(t_1, \ldots, t_n)$.

Otherwise, Case 2 holds. Note that $l'_1, \ldots, l'_{I-1}$ are all odd, and l'_I is even. Also, we have $l'_I \geq l'_{I-1} \geq \cdots \geq l_1$. Therefore, by Proposition 1, the tail-swap

$$\tau(t'_I, t'_{I-1}, \ldots, t'_1) = (t_I, t_{I-1}, \ldots, t_1)$$

succeeds. Now, by Proposition 3, $\tau(t_1, \ldots, t_I) = (t'_1, \ldots, t'_I)$. Furthermore, if we let $t_{I+1} = t'_{I+1}$, then $\tau(t_I, t_{I+1})$ fails. So

$$\tau(t_1, \ldots, t_I, t'_{I+1}, t'_{I+2}, \ldots, t'_n) = (t'_1, \ldots, t'_n),$$

as desired. □

Corollary 4. *The tail-swap procedure τ is bijective from* $\mathrm{dom}(\tau)$ *to* $\mathrm{img}(\tau)$*, where* $\mathrm{img}(\tau)$ *is characterized in Theorem 1.*

Proof. By Propositions 3 and 4 (see Remark 2), τ is injective. Therefore, it is bijective onto its image. □

Remark 5. The condition on the lengths $l_1, \ldots, l_n$ is satisfied by the sequence $0, 3, 7, \ldots, 2^n - 1$. The sequence of the lengths of the tilings in S_1 (for the Vajda identity) is $1, 3, 7, \ldots, 2^n - 1$. However, the number of tilings of length 0 is the same as the number of tilings of length 1 (namely, there is one of each). Thus, we prove Vajda's identity while assuming that the length of the first tiling is 0, rather than 1.

We are now ready to describe the bijection $f : S_1 \to S_2$. Let $\mathbf{t} \in S_1$ be $(t_1, t_2, \ldots, t_n)$. The lengths of these tilings are $2^1 - 1, 2^2 - 1, \ldots, 2^n - 1$. Let $(t'_1, \ldots, t'_n) = \tau(\emptyset, t_2, t_3, \ldots, t_n)$, and let i be the failure index. Now, if $i = 0$, then let $f(\mathbf{t}') = (t'_1, \ldots, t'_n)$. Note that $f(\mathbf{t}) \in S_2$ in this case, because the lengths of $t'_1, t'_2, \ldots, t'_n$ are $2^1 - 1, 2^2 - 1, \ldots, 2^{n-1} - 1, 2^n - 2$, respectively.

On the other hand, if $i < n$, then the ith tail swap failed. Therefore, by Proposition 1, t'_{i+1} ends with 2^{i-2} dominoes. Let t^*_{i+1} be the result of removing those 2^{i-2} dominoes from t^*_i. Then let $f(\mathbf{t}) = (t'_1, \ldots, t'_{i-1}, t^*_{i+1}, t'_{i+2}, t'_{i+3}, \ldots, t'_n)$. Note that the lengths of the tilings here are $2^1 - 1, 2^2 - 1, \ldots, 2^{i-1} - 1, 2^i - 1, 2^{i+2} - 1, 2^{i+3} - 1, \ldots, 2^n - 1$, so $f(\mathbf{t})$ is once again in S_2.

Proposition 4. *The map f is injective.*

Proof. If $f(t_1, \ldots, t_n) = (u_1, \ldots, u_n)$ (that is, if the result has n tilings), then f was just a multiple tail-swap, so by Corollary 1, it is injective.

On the other hand, suppose $f(t_1, \ldots, t_n) = (u_1, \ldots, u_{n-1})$ (that is, there are only $n - 1$ tilings), then the set of the lengths of $u_1, \ldots, u_{n-1}$ is equal to $\{1, 3, \ldots, 2^n - 1\} \setminus \{2^k - 1\}$ for some $k \geq 2$. How did this happen? Well, first of all, we know that the $(k-1)$th tail swap failed. Using the notation from above, this means that t'_{k-1} had length $2^{k-1} - 2$. Furthermore, by Proposition 1, it consisted only of dominoes. Also, we know that $u_{k-1} = t^*_k$, so we can reconstruct t'_k by appending $2^{k-2} - 1$ dominoes to the end of u_{k-1}. Finally, for $1 \leq j \leq k-2$, we have $u_j = t'_j$, while for $k \leq j \leq n-1$, we have $u_j = t'_{j+1}$. Thus, we can reconstruct all the values $t'_1, t'_2, \ldots, t'_n$. By Corollary 1, the values $t_1, t_2, \ldots, t_n$ are uniquely determined. □

Remark 6. The above proposition demonstrates the existence of an inverse map $f^{-1} : S_2 \to S_1$ for f.

Proposition 5. *The map f^{-1} is injective.*

Proof. Suppose $f^{-1}(u_1, \ldots, u_l) = f^{-1}(v_1, \ldots, v_m)$, where $(u_1, \ldots, u_l)$ and $(v_1, \ldots, v_m)$ are both in S_2 and $l, m \in \{n-1, n\}$.

Case 1: $l = m = n$. In this case, f^{-1} is just a multiple tail swap, which is a bijection by Corollary 1. In particular, it must be the case that $(u_1, \ldots, u_l) = (v_1, \ldots, v_m)$.

Case 2: $l = n-1$ and $m = n$, or vice versa. Without loss of generality, let $l = n-1$, and let the lengths of the tilings $u_1, \ldots, u_l$ be $\{1, 3, \ldots, 2^n - 1\} \setminus \{2^k - 1\}$. Let $u'_1, \ldots, u'_n$ be the result of the reconstruction process described above; that is,

$$u'_j = \begin{cases} u_j & 1 \leq j \leq k-2 \\ 2^{n-2} - 1 \text{ dominoes} & j = k-1 \\ u_{k-1} \text{with } 2^{n-2} - 1 \text{ dominoes appended} & j = k \\ u_{k-1} & k+1 \leq j \leq n \end{cases}.$$

Now, $u'_1, \ldots, u'_n$ is the result of a multiple tail-swap with failure index $k-1$. Note that $v_1, \ldots, v_m$ is the result of a multiple tail-swap with failure index n (since $m = n$). Therefore, applying τ^{-1} to $u'_1, \ldots, u'_n$ and $v_1, \ldots, v_m$ yields different results, by Corollary 1. But this is a contradiction, since

$$\tau^{-1}(u'_1, \ldots, u'_n) = f^{-1}(u_1, \ldots, u_l) = f^{-1}(v_1, \ldots, v_m) = \tau^{-1}(v_1, \ldots, v_m).$$

So case 2 cannot occur.

Case 3: $l = m = n-1$. Let the lengths of $u_1, \ldots, u_{n-1}$ be $\{1, 3, \ldots, 2^n - 1\} \setminus \{2^{k_u} - 1\}$, and let the lengths of $v_1, \ldots, v_{n-1}$ be $\{1, 3, \ldots, 2^n - 1\} \setminus \{2^{k_v} - 1\}$. As above, let $u'_1, \ldots, u'_n$ be the result of the reconstruction process applied to $u_1, \ldots, u_{n-1}$, and let $v'_1, \ldots, v'_n$ be the reconstruction of $v_1, \ldots, v_{n-1}$. Now, $u'_1, \ldots, u'_n$ is the result of a multiple tail-swap with failure index $k_u - 1$, and $v_1, \ldots, v'_n$ is the result of a multiple tail-swap with failure index $k_v - 1$. As above, if $k_u \neq k_v$, we have a contradiction, so we can assume that $k_u = k_v$. Then we have

$$(u'_1, \ldots, u'_n) = \tau(f^{-1}(u_1, \ldots, u_l)) = \tau(f^{-1}(v_1, \ldots, v_m)) = (v'_1, \ldots, v'_n).$$

But if $(u'_1, \ldots, u'_n) = (v'_1, \ldots, v'_n)$, then $(u_1, \ldots, u_{n-1}) = (v_1, \ldots, v_{n-1})$, as claimed.

So f^{-1} is injective. □

Theorem 7. *The function f is a bijection.*

Proof. Propositions 5 and 6 show that both f and f^{-1} are injective. □

In our proof, we used the form (1) of the identity rather than the form (V90). There are three differences between (1) and (V90). The first is that both sides have been multiplied by F_2. This makes no combinatorial difference, but merely clarifies the bijection. The second difference is that some terms have been moved from one side to the other. Again, this does not make any combinatorial difference. The second is that both sides have been multiplied by $F_4 F_8 \cdots F_{2^n}$. Therefore, to understand the identity in the form (V90), we use the following probabilistic reasoning. Suppose we are given a sequence of tilings

$$(t_1, \ldots, t_n) \in S_1$$

Replace the tiling of length 1 with the empty tiling and then perform a multiple tail-swap. There are several mutually exclusive possible outcomes: we

could have any failure index in $\{0, 1, \ldots, n-1\}$. Our bijection demonstrates that the number of $t \in S_1$ with failure index $k \geq 1$ is exactly

$$\prod_{\substack{1 \leq j \leq n-1 \\ j \neq k+1}} F_{2^j},$$

and hence the probability of a randomly chosen element of S_1 having failure index $k \geq 1$ is $\frac{1}{F_{2^{n+1}}}$. Similarly, the probability of having failure index 0 is $\frac{F^{2^n-1}}{F_{2^n}}$. Since these are the only possible outcomes and they are mutually exclusive, we have

$$\sum_{k=1}^{n-1} \frac{1}{F_{2^{k+1}}} + \frac{F_{2^n-1}}{F_{2^n}} = 1,$$

which is (V90). So the identity (V90) has both a *probabilistic form* and a *combinatorial form*.

3.4 Further identities

There is nothing unique about the values $\left\{2^k\right\}_{k=1}^n$. Given any increasing sequence of even integers $\{a_n\}_{k=1}^n$ with $a_1 = 2$, we canuse the multiple tail-swap procedure to give a bijective proof of an identity about products of Fibonacci numbers:

$$\prod_{k=1}^{n} F_{a_k} = F_{a_n-1} \prod_{k=1}^{n-1} F_{a_k} + \sum_{1 \leq k \leq n-1} F_{a_{k+1}-a_k} \prod_{\substack{j \notin \{k,k+1\} \\ 1 \leq j \leq n}} F_{a_j}. \qquad (*)$$

The proof of this general formula proceeds exactly as in the proof of the specific case of Vajda's 90th identity. Like Vajda's 90th identity, this general identity has a nice probabilistic form:

$$\sum_{1 \leq k \leq n-1} \frac{F_{a_{k+1}-a_k}}{F_{a_k} F_{a_{k+1}}} + \frac{F_{a_n-1}}{F_{a_n}} = 1. \qquad (**)$$

Notice that when $a_n = 2^n$, the values $F_{a_{k+1}-a_k}$ in the numerator and F_{a_k} in the denominator cancel, accounting for the elegance of (V90).

As an example, we provide the special cases of $(*)$ for the sequences $\{2k\}_{k=1}^n$, $\{4k-2\}_{k=1}^n$, and $\left\{3^k-1\right\}_{k=1}^n$. Using these sequences in $(*)$, we get the identities

$$\prod_{k=1}^{n} F_{2k} = F_{2n-1} \prod_{k=1}^{n-1} F_{2k} + \sum_{1\le k\le n-1} \prod_{\substack{j\notin\{k,k+1\}\\ 1\le j\le n}} F_{2j}$$

$$\prod_{k=1}^{n} F_{4k-2} = F_{4n-3} \prod_{k=1}^{n-1} F_{4k-2} + 3 \sum_{1\le k\le n-1} \prod_{\substack{j\notin\{k,k+1\}\\ 1\le j\le n}} F_{4j-2}$$

$$\prod_{k=1}^{n} F_{3^k-1} = F_{3^k-2} \prod_{k=1}^{n-1} F_{3^k-1} + F_{2\cdot 3^k} \prod_{\substack{j\notin\{k,k+1\}\\ 1\le j\le n}} F_{3^k-1}.$$

The equivalent probabilistic forms, given by $(**)$, are

$$\sum_{k=1}^{n-1} \frac{1}{F_{2k}F_{2k+2}} + \frac{F_{2n-1}}{F_{2n}} = 1,$$

$$\sum_{k=1}^{n-1} \frac{3}{F_{4k-2}F_{4k+2}} + \frac{F_{4n-3}}{F_{4n-2}} = 1,$$

and

$$\sum_{k=1}^{n-1} \frac{F_{2\cdot 3^k}}{F_{3^k-1}F_{3^{k+1}-1}} + \frac{F_{3^n-2}}{F_{3^n-1}} = 1.$$

Some of these identities are reasonably attractive; in particular, when $F_{a_{k+1}-a_k}$ is constant, the probabilistic form of the identity is relatively simple. This happens when the sequence $\{a_k\}$ is arithmetic.

4 Counting binary De Bruijn cycles

A *binary De Bruijn cycle* of order n is a cyclic sequence of 0s and 1s such that every string of 0s and 1s of length n occurs exactly once as a conesecutive substring. For example, 00010111 is a binary De Bruijn cycle of order 3. We adopt the convention of treating this as the same cycle as its cyclic rearrangements (such as 00101110), though in other works these are treated as different objects.

The following striking result was first proved by Sainte-Marie [3]:

Theorem 8. *There are $2^{2^{n-1}-n}$ binary De Bruijn cycles of order n.*

Before proving this theorem, we introduce some useful concepts and some simple results. In what follows, let Σ^n denote the set of $\{0,1\}$-strings of length n.

Definition 9. A *2-in 2-out directed graph* is a directed graph in which every vertex has indegree and outdegree 2.

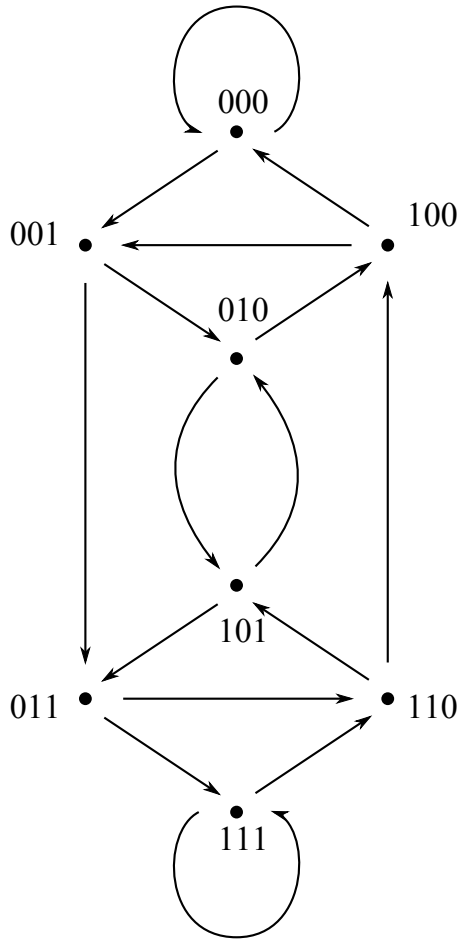

Figure 5: The graph B_3.

Definition 10. The De Bruijn graph B_n of order n is the directed graph whose vertex set is Σ^n and whose edges are the pairs of strings (a, b) such that for $i \in [n-1]$, the ith character of b is the $(i+1)$th character of a. For example, Figure 5 shows the De Bruijn graph B_3.

Definition 11. Given a directed graph $G = (V, E)$, the *line graph* $L(G)$ of G is the graph whose vertices are the edges of G and whose edges are the set of $(e, f) \in E \times E$ such that there exists $v \in V$ where e is directed to v and f is directed from v.

Proposition 6. *The De Bruijn graph is a 2-in 2-out digraph.*

Proposition 7. *The binary De Bruijn cycles of order n are in bijection with the Eulerian cycles of B_{n-1}.*

Proof. We label the edges of B_{n-1} as follows: label the edge $(a_1a_2\cdots a_{n-1}, b_1b_2\cdots b_{n-1})$ with the string $a_1a_2a_3\cdots a_{n-1}b_{n-1}$. It is evident that each element in Σ^n appears exactly once as a label. Futhermore, an edge e is adjacent to an edge f if and only if the label of f can succeed the label of e in a De Bruijn cycle. Therefore, a De Bruijn cycle of order n corresponds to a cyclic list of 2^n distinct edges, each adjacent to its successor; that is, a cycle of length 2^n. Because B_{n-1} has exactly 2^n edges, the cycle is Eulerian. □

Proposition 8. $L(B_n) = B_{n+1}$.

Proof. Using the same labeling as above, the result is immediate. □

Corollary 12. *The binary De Bruijn cycles of order n are in bijection with the Hamiltonian cycles of B_n.*

Proposition 9. *If G is a 2-in, 2-out digraph with k Eulerian cycles and n vertices, then $L(G)$ has $2^{n-1}k$ Eulerian cycles.*

Proof. We induct on n. In the base case $n = 1$, there is only one 2-in, 2-out digraph, and the theorem is immediate.

Suppose the theorem holds for digraphs with n vertices. Let G have $n+1$ vertices and k_G Eulerian cycles and choose one vertex arbitrarily. For the moment, assume that this vertex is not adjacent to a loop (the case where it is adjacent to a loop is much easier and is left to the reader). Now, "splice out" that vertex in two different ways to get two new graphs H_1 and H_2 with n vertices. The "splicing out" operation consists of deleting the vertex and all edges adjacent to it, and then connecting each former in-neighbor to a former out-neighbor. Since there are two in-neighbors and two out-neighbors, there are always two ways to do this splice. Observe that any (Eulerian) cycle in G corresponds to a cycle in H_1 or a cycle in H_2 (but not both), so if k_{H_1} denotes the number of cycles in H_1 and k_{H_2} the number of cycles in H_2, then $k = k_{H_1} + k_{H_2}$.

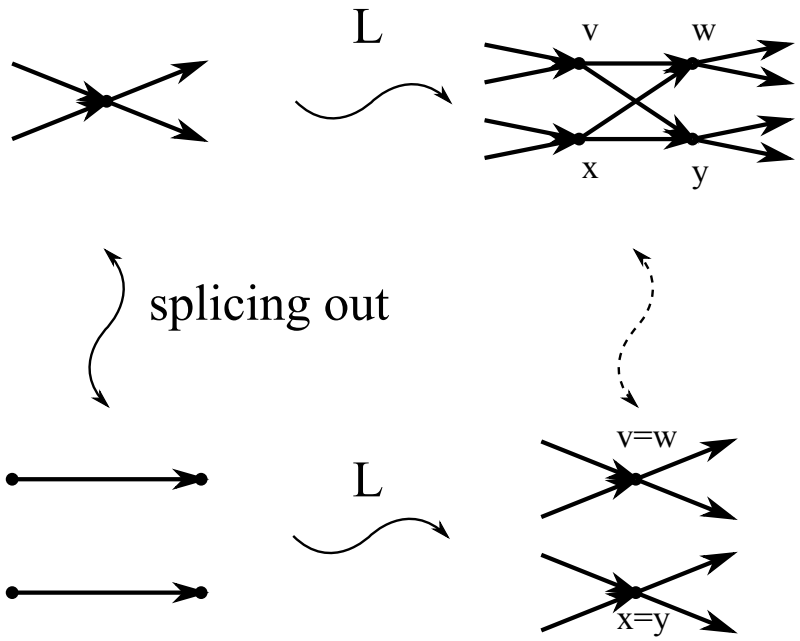

Figure 6: Splicing out vertices in G and $L(G)$.

Referring to Figure 6, we observe that for every cycle in $L(H_i)$, there are 2 corresponding cycles in $L(G)$. To see this, consider the two paths leaving x and returning to v or w. Either both go to p, or both go to q, or one goes to p and the other to q. In each case, there are 4 cycles in $L(G)$, and two cycles in H_1 and H_2. So $k_{L(G)} = 2(k_{L(H_1)} + k_{L(H_2)})$. By the induction hypothesis, $k_{L(H_i)} = 2^{n-1}k_{H_i}$, so

$$k_{L(G)} = 2(2^{n-1}k_{H_1} + 2^{n-1}k_{H_2}) = 2^n(k_{H_1} + k_{H_2}) = 2^n k_H.$$

This is what we wanted. □

Corollary 13. *For $n \geq 1$, B_{n-1} has $2^{2^{n-1}-n}$ Eulerian cycles. Therefore, there are $2^{2^{n-1}-n}$ De Bruijn cycles of order n.*

Proof. Induction on n. When $n = 1$, the result is immediate because $2^{2^{1-1}-1} = 1$, and B_0 does indeed have 1 Eulerian cycle. Assume that B_{n-1} has $2^{2^{n-1}-n}$ Eulerian cycles. Then $B_n = L(B_{n-1})$ has $2^{2^{n-1}-1} \cdot 2^{2^{n-1}-n} = 2^{2^n-(n+1)}$ Eulerian cycles, by the previous proposition. This completes the proof. □

Note that this counting result has been proved by induction. This suggests that we could use the translation method to convert the inductive proof into a bijective proof. However, this approach runs into trouble because the inductive proof here is significantly different from the previous inductive proofs we have encountered. Namely, in the previous proofs, there was only one natural way to move from the $(n+1)$-case to the n-case.

For example, when we considered the sum of the $n+1$ terms

$$1 + 3 + \cdots + (2n-1) + (2n+1),$$

we wanted to use the assumption that $1 + 3 + \cdots + (2n-1) = n^2$. The assumption was specific to the first $(2n-1)$ terms of the summation, so there was no doubt about the next step in the proof.

However, in this problem, the induction step involves picking an arbitrary vertex of the graph B_n and splicing it out. There are 2^n ways to pick this vertex, and although symmetries cause some of these choices to be equivalent, this still means there are on the order of 2^n different potential bijections to consider. Even worse, after picking this vertex, we are not done, because now we have a graph with $2^n - 1$ vertices. But B_{n-1} has 2^{n-1} vertices, which is vastly less. So we have to pick one of the remaining $2^n - 1$ vertices to splice out, to get down to a graph with $2^n - 2$ vertices, and so on, with each of these choices again yielding a different bijection. Overall, there are approximately $\frac{2^n!}{2^{n-1}!}$ potential bijections between $\{0,1\}^{2^{n-1}-1} \times B_{n-1}$ and B_n.

Remember that we are seeking an elegant bijection that can be expressed concisely. With so many choices to make, we have little hope of going through with the final step of the translation method: rewriting our bijection in an insightful, direct form. The only chance is to stumble upon a fruitful rule for choosing which point to splice out at each step, since we know of no systematic way for coming up with such a rule.

5 A bijective counting of the binary De Bruijn cycles

In this section we present a bijection, constructed by Bidkhori and Kishore [2], between the binary De Bruijn cycles of order n and the binary strings of length $2^{n-1} - n$. We will slightly modify their terminology in what follows.

Given a digraph G, a *spanning tree* of G is a subgraph T of G such that every vertex of G is in T, and such that there exists a vertex v such that

from every vertex w of T, there is a unique path in T from w to v. (In particular, T is a tree.) The vertex v is called the *root* of T.

Given a digraph G, define a *tree array* τ of $G = (V, E)$ with $V = \{v_1, \ldots, v_n\}$ to be a $(n+1)$-tuple $(T, l_{v_1}, \ldots, l_{v_n})$, where T is a spanning tree of G, and l_{v_i} is an ordered list of $\text{indeg}(v) - 1$ edges (not necessarily distinct), all with source v_i.

Assume that E is ordered (arbitrarily, if necessary). Define a bijection σ from the tree arrays of G to the spanning trees of $L(G)$. Let $\tau = (T, l_{v_1}, \ldots, l_{v_n})$ be a tree array of G. To calculate $\sigma(\tau)$, we use the following algorithm:

1. Let $\mathcal{T}$ be the empty subgraph of $L(G)$.
2. Let
$$S = \{e \in E : e \text{ is not in any } l_{v_i},\ e \notin T, \text{ and } \text{outdeg}_{\mathcal{T}}(e) = 0\}.$$
3. Let f be the smallest element of S under the ordering of E. Let v be the target of f in G.
4. If l_v has any elements, let g be the first element. Remove g from l_v, then add the edge (f, g) to $\mathcal{T}$. Then go to Step 3.
5. Otherwise, if v is the root of T, stop immediately and let $\sigma(\tau) = \mathcal{T}$.
6. Otherwise, let g be the edge in T with source v. Add the edge (f, g) to $\mathcal{T}$. Then go to Step 3.

The inverse σ^{-1} can be defined by the following algorithm:

1. Let l_v be empty for each $v \in V$, and let T be empty.
2. Take the leaf f of $\mathcal{T}$ which is least under the ordering of E.
3. If f is not the root of $\mathcal{T}$, then let g be the target of the only edge in $\mathcal{T}$ with source f. Delete f from $\mathcal{T}$. Let v be the target of f in G. Append g to l_v, and repeat this step.
4. If f is the root of $\mathcal{T}$, then for each $v \in V$, remove the last element of l_v and add it to T. Let $\sigma^{-1}(\mathcal{T}) = (T, l_{v_1}, \ldots, l_{v_n})$.

The reader should refer to [2] for proofs that these maps are indeed what they claim to be (that is, maps from tree arrays of G to spanning trees of $L(G)$, and vice versa).

Now we can describe a bijection between De Bruijn cycles and binary strings of length $2^{n-1} - n$. Let a De Bruijn cycle be given. Take the corresponding Hamiltonian circuit in B_n, and delete the edge with source $111\cdots1$ to turn this circuit into a Hamiltonian path.

Think of this Hamiltonian path as a spanning tree T_n of B_n. Let T_{n-1} be the tree from the tree array $\sigma^{-1}(T_n)$, and so on; that is, let T_k be the tree from the tree array $\sigma^{-1}(T_{k+1})$ for all $1 \leq k \leq n-1$. In addition, for $1 \leq k \leq n-1$, denote by A_k the tree array $\sigma^{-1}(T_{k+1})$.

Let $1 \leq k \leq n-2$. For each $A_k = (T_k, l_{v_1}^{(k)}, \ldots, l_{v_{2^k}}^{(k)})$, where the vertices are ordered lexicographically (recall that A_k is a tree array for B_k, and the vertices of B_k are binary strings) we create a ternary string with 2^k characters. To find the ith character s_i of this string, take the list $l_{v_i}^{(k)}$. If it has any elements, take the first element. It goes from vertex v_i to some other vertex w. Since w is a vertex of B_k, it is a binary string. Let s_i be the last character of w. On the other hand, if $l_{v_i}^{(k)}$ does not have any elements and v_i is not the root of T_k, then take the unique edge in T_k that has source v_i. Let the target of this edge be w, and let s_i be the last character of w. If v_i is the root of T_k, let $s_i = 2$. It is easy to see this case happens if $i = 2^{k+1} - 1$.

Now, concatenate the ternary strings associated to $A_1, A_2, \ldots, A_{n-2}$. This concatenation has $2^{n-1} - 1$ characters. However, some of them are twos: these are at indices $1, 2, 4, 8, \ldots, 2^{n-2}$. If we eliminate these $(n-1)$ twos, we have a binary string of length $2^{n-1} - n$, and this is our answer.

To show that this is a bijection, we describe the inverse map. Given a binary sequence $s_1 s_2 \ldots s_{2^{n-1}-n}$, we begin by inserting 2s at indices $1, 2, 4, 8, \ldots, 2^{n-2}$.

We now construct tree arrays $A_1, A_2, \ldots, A_{n-2}, A_{n-1}$. Let $A_1 = (T, (), ())$, and let T be the tree rooted at 0 in B_1 if $s_1 = 0$, and the tree rooted at 1 if $s_1 = 1$. For $2 \leq k \leq n-1$, let i range from 1 to 2^k. Let the vertices (in lexicographic order) of B_k be $v_1, \ldots, v_{2^k}$. Add the "zero edge" with source v_i to $l_{v_i}^{(k)}$ if $s_{2^k+i-1} = 0$, and add the "one edge" with source v_i if $s_{2^k+i-1} = 1$. (If $s_{2^k+i-1} = 2$, do not add any edge.) Then let $T_k = \sigma(A_{k-1})$, and write $A_k = (T_k, l_{v_1}^{(k)}, \ldots, l_{v_{2^k}}^{(k)})$.

Now, take $\sigma(A_{n-1}) = T_n$ to get a Hamiltonian path in B_n. Finally, add the remaining edge, which closes the Hamiltonian path, to get a Hamiltonian cycle, and read off the corresponding De Bruijn cycle. This map is easily seen to be the inverse of the map from De Bruijn cycles to binary strings, so this is indeed a bijection.

References

[1] A. Benjamin and J. Quinn, *Proofs that really count*, The Mathematical Association of America, 2003.

[2] H. Bidkhori and S. Kishore, *A bijective proof of the theorem of knuth*, Combinatorics, Probability and Computing **20** (2011), 11–25.

[3] C. F. Sainte-Marie, *Solution to question nr. 48*, l'Intermédiaire des Mathématiciens **1** (1894), 107–110.

[4] S. Vajda, *Fibonacci and lucas numbers, and the golden section: Theory and applications*, Dover Publications, 2007.

[5] P. M. Wood and D. Zeilberger, *A translation method for finding combinatorial bijections*, Annals of Combinatorics **13** (2009), 383–402.