8-1-2010

Combinatorially Composing Chebyshev Polynomials

Arthur T. Benjamin
Harvey Mudd College

Daniel Walton '07
Harvey Mudd College

Journal of Statistical Planning and Inference 140 (2010) 2161–2167

Contents lists available at ScienceDirect

Journal of Statistical Planning and Inference

journal homepage: www.elsevier.com/locate/jspi

Combinatorially composing Chebyshev polynomials

Arthur T. Benjamin [a,*], Daniel Walton [b]

[a] Department of Mathematics, Harvey Mudd College, Claremont, CA 91711, USA
[b] Department of Atmospheric and Oceanic Sciences, UCLA, Los Angeles, CA 90095-1565, USA

ARTICLE INFO

Available online 20 January 2010

Keywords:
Chebyshev polynomials
Combinatorial proof
Tiling

ABSTRACT

We present a combinatorial proof of two fundamental composition identities associated with Chebyshev polynomials. Namely, for all $m,n \geq 0$, $T_m(T_n(x)) = T_{mn}(x)$ and $U_{m-1}(T_n(x))U_{n-1}(x) = U_{mn-1}(x)$.

1. Introduction

The *Chebyshev polynomials of the first kind* are defined by $T_0(x) = 1$, $T_1(x) = x$, and for $n \geq 2$,

$$T_n(x) = 2xT_{n-1}(x) - T_{n-2}(x).$$

The next few polynomials are $T_2(x) = 2x^2 - 1$, $T_3(x) = 4x^3 - 3x$, $T_4(x) = 8x^4 - 8x^2 + 1$.

The *Chebyshev polynomials of the second kind* differ only in the initial conditions. They are defined by $U_0(x) = 1$, $U_1(x) = 2x$, and for $n \geq 2$,

$$U_n(x) = 2xU_{n-1}(x) - U_{n-2}(x).$$

The next few polynomials are $U_2(x) = 4x^2 - 1$, $U_3(x) = 8x^3 - 4x$, $U_4(x) = 16x^4 - 12x^2 + 1$.

Shapiro (1981) showed that Chebyshev polynomials have a simple combinatorial structure, and exploited it to prove some Chebyshev polynomial identities. Several more Chebyshev polynomial identities are given combinatorial proofs in Benjamin et al. (2010), Benjamin and Walton (2009), and Walton (2007). Continuing in that spirit, we present a combinatorial proof of two fundamental composition identities associated with Chebyshev polynomials. Namely, for all $m,n \geq 0$,

$$T_m(T_n(x)) = T_{mn}(x)$$

and

$$U_{m-1}(T_n(x))U_{n-1}(x) = U_{mn-1}(x).$$

These identities are well-known (see, for instance Rivlin, 1990), and have elementary algebraic proofs. Nevertheless, their combinatorial structure compels us to seek a direct combinatorial proof.

2. Combinatorial interpretations

The fundamental combinatorial objects of this article are *n-tilings*, which are sequences of light squares, dark squares, and dominoes that cover a total length of n cells, where squares cover one cell and dominoes cover two cells. A *restricted*

* Corresponding author.
E-mail address: benjamin@hmc.edu (A.T. Benjamin).

0378-3758/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.jspi.2010.01.012

n-tiling is an n-tiling that is not allowed to start with a dark square, and we let $\mathcal{T}_n$ denote the set of restricted n-tilings. For example, if we let the symbols a, b, and D, respectively, denote light squares, dark squares, and dominoes, then $\mathcal{T}_3$ comprises the tilings

aaa aab aba abb aD Da Db

We define the *weight* of an n-tiling to be $(-1)^k x^{n-2k}$, where k is the number of dominoes in the n-tiling (and therefore $n-2k$ is the number of squares). You can think of each domino contributing a multiplicative weight of -1 and each square contributing a multiplicative weight of x. Among the restricted 3-tilings, there are four tilings of weight x^3 and three tilings of weight $-x$, giving a total weight of $4x^3-3x$ which, not coincidentally, is $T_3(x)$.

Theorem 1. *For $n \geq 0$, $T_n(x)$ is the total weight of all restricted n-tilings.*

Proof. The proof is by induction on n. When $n=0$, the empty tiling has weight 1, and when $n=1$, the tiling consisting of a single light square has weight x. For $n \geq 2$, by induction, the total weight of all restricted n-tilings that end with a square (of either color) is $2xT_{n-1}(x)$ and the total weight of all restricted n-tilings that end with a domino is $(-1)T_{n-2}(x)$, for a combined total weight of $2xT_{n-1}(x)-T_{n-2}(x) = T_n(x)$, as desired. □

If we remove the initial tile restriction, and allow n-tilings to start with a dark square, then more tilings are counted. We let $\mathcal{U}_n$ denote the set of unrestricted n-tilings. For example, when $n=3$, in addition to the tilings of $\mathcal{T}_3$, $\mathcal{U}_3$ also contains

baa bab bba bbb bD

whose total weight is $4x^3-x$. Adding this to the previous tilings gives us a total weight of $8x^3-4x = U_3(x)$. By the exact same reasoning as before, we obtain

Theorem 2. *For $n \geq 0$, $U_n(x)$ is the total weight of all (unrestricted) n-tilings.*

Viewed combinatorially, many Chebyshev identities become transparent. For example, for $n \geq 1$, $U_n(x) = xU_{n-1}(x)+T_n(x)$ is simply stating that an unrestricted n-tiling either begins with a dark square (where the dark square has weight x) followed by an unrestricted tiling of length $n-1$ or else the n-tiling satisfies the leading tile restriction. Here is another one. By considering whether a restricted n-tiling begins with a square or a domino, we get: for $n \geq 2$, $T_n(x) = xU_{n-1}(x)-U_{n-2}(x)$.

To prove the composition identities we will need the technique of *tailswapping*, explored extensively in Benjamin and Quinn (2003). Before we do this, we must introduce the concept of a tiling being *breakable*.

Definition 1. We say that a tiling is *m-breakable* or *breakable at cell m* if there is no domino covering both cells m and $m+1$.

The intuitive idea is that an m-breakable tiling can be separated into two distinct tilings with the separation occurring between cells m and $m+1$. Note that the weight of all m-breakable $(m+n)$- tilings is $U_m(x)U_n(x)$. If the tilings are restricted, then the total weight is $T_m(x)U_n(x)$. The total weight of m-unbreakable tilings is $-U_{m-1}(x)U_{n-1}(x)$, since there is a domino of weight -1 covering cells m and $m+1$. Likewise the total weight of restricted m-unbreakable tilings is $-T_{m-1}(x)U_{n-1}(x)$. Observing that each length $(m+n)$- tiling is either breakable or unbreakable at cell m immediately gives us the following two identities.

Identity 1. *For integers $m, n \geq 1$,*

$$U_{m+n}(x) = U_m(x)U_n(x)-U_{m-1}(x)U_{n-1}(x)$$

and

$$T_{m+n}(x) = T_m(x)U_n(x)-T_{m-1}(x)U_{n-1}(x).$$

This next theorem uses the idea of breakability as well as the more advanced technique of tailswapping.

Identity 2. *For $n \geq 1$, $U_n^2(x)-U_{n+1}(x)U_{n-1}(x) = 1$.*

Proof. Here we describe a weight preserving correspondence between $\mathcal{U}_n \times \mathcal{U}_n$ and $\mathcal{U}_{n+1} \times \mathcal{U}_{n-1}$ that is *almost* a bijection. Let $C = (A, B)$ be an element of $\mathcal{U}_n \times \mathcal{U}_n$, where A is an n-tiling occupying cells 1 through n, B is also an n-tiling, but it is offset so that it occupies cells 2 through $n+1$. We define the weight of C to be the product of the weight of its tiles, i.e., $w(C) = w(A)w(B)$. We say that C has a *fault* at cell j (where $1 \leq j \leq n$) if A and B are both breakable at cell j. (Note that A is considered breakable at cell n and B is considered breakable at cell 1.) Now suppose C has its rightmost fault at cell k. Say $A = A_1A_2$ where A_1 is the subtiling covering cells 1 through k, and A_2 is the subtiling covering cells $k+1$ through n, and similarly $B = B_1B_2$, where B_1 covers cells 2 through k, and B_2 covers cells $k+1$ through $n+1$. Then by *tailswapping*, we map C to $C' = (A', B')$, where $A' = A_1B_2$ and $B' = B_1A_2$. Note that C' is in $\mathcal{U}_{n+1} \times \mathcal{U}_{n-1}$ and has the same weight as C. Also C' has the same rightmost fault at cell k, so tailswapping C' produces C again. See Fig. 1.

Tailswapping is well-defined, provided that at least one fault exists. Hence the quantity $U_n^2(x)-U_{n+1}(x)U_{n-1}(x)$ is the total weight of *fault-free* tilings of $\mathcal{U}_n \times \mathcal{U}_n$ minus the weight of the fault-free tilings of $\mathcal{U}_{n+1} \times \mathcal{U}_{n-1}$.

When n is even, the only fault-free tiling of $\mathcal{U}_n \times \mathcal{U}_n$ is the all-domino tiling $(A, B) = (D^{n/2}, D^{n/2})$ which has weight $(-1)^n = 1$, and $\mathcal{U}_{n+1} \times \mathcal{U}_{n-1}$ has no fault free tilings since $\mathcal{U}_{n-1}$ contains at least one square which will generate a fault.

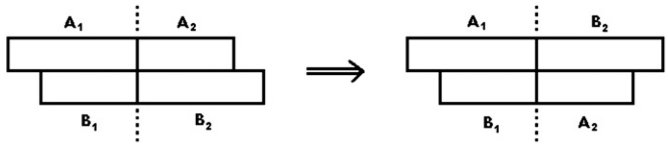

Fig. 1. Tailswapping: line up n-tilings A and B so that the right end of B extends one space past the right end of A. Locate the right most fault line, then swap the portions of A and B that are to the right of this line.

When n is odd, then $\mathcal{U}_n \times \mathcal{U}_n$ has no fault-free tilings, but $\mathcal{U}_{n+1} \times \mathcal{U}_{n-1}$ has one fault-free tiling, the all-domino tiling $(D^{(n+1)/2}, D^{(n-1)/2})$ with weight $(-1)^n = -1$. Thus, regardless of the parity of n, the difference of the weights of the fault-free tilings is 1, as desired. □

3. A composition formula for $T_{mn}(x)$

We now offer a bijective proof of the main identity of this paper.

Identity 3. *For* $n \geq 0$, $T_{mn}(x) = T_m(T_n(x))$.

From Theorem 1, $T_{mn}(x)$ is the total weight of all restricted mn-tilings. We denote the set of restricted mn-tilings by $\mathcal{T}_{mn}$.

But what does $T_m(T_n(x))$ count? On first inspection, it is the total weight of all restricted m-tilings, where dominoes have weight -1 and squares have weight $T_n(x)$. But we can decompose $T_n(x)$ as $w_1(x) + w_2(x) + \cdots + w_p(x)$, where p is the number of restricted n-tilings and $w_i(x)$ is the weight of the i th restricted n-tiling. Thus $T_m(T_n(x))$ can be thought of as the sum of the weights of all restricted m-tilings where a domino has weight -1 and each square has the weight of a restricted n-tiling. We summarize this by saying that $T_m(T_n(x))$ is the total weight of all length m restricted *metatilings*, in which each square has a weight given by a length n restricted *minitiling*, where the restriction is that no metatiling nor minitiling may begin with a dark square. Each domino in the metatiling has weight -1. We call the set of such objects $\mathcal{T}_m(\mathcal{T}_n)$.

For example, $T_3(T_2(x))$ is the sum of the length 3 metatilings in which squares have weights corresponding to length 2 minitilings. Fig. 2 shows three examples of such tilings.

We now prove Identity 3, by exhibiting a weight preserving bijection between $\mathcal{T}_{mn}$ and $\mathcal{T}_m(\mathcal{T}_n)$. First we consider the case where n is odd.

Case 1: *n is odd*. Consider a tiling $\sigma \in \mathcal{T}_{mn}$. Write σ as m rows of n-tilings, stacked on top of each other (called an $m \times n$ board), where the first row (at the top) consists of the cells 1 through n, the second row consists of cells $n+1$ through $2n$, ..., the m th row (at the bottom) consists of cells $(m-1)n+1$ through mn. If a domino starts in the last cell of row and ends in the first cell of the next row, we say that such a domino is *out of phase*. Fig. 3 shows an example of how the length 18 tiling *DababDDDbDabb* can be turned into a 6×3 board. Notice that the out of phase domino starting on row 3 and ending on row 4 is denoted by the dashed lines.

The basic strategy of the bijection is as follows. The bijection "tries to" map the k-th row of the $m \times n$ board is to a square occupying cell k of the metatiling, and the associated minitiling has the same weight as (and is often the same as) row k. If it cannot map the k-th row in that fashion, then it performs a "tailswap" with a later row j, allowing rows k through j to be mapped in a natural way. If k and j cannot be tailswapped, a rare event, then these rows are brought together and mapped to a domino. The details are spelled out in the cases that follow.

Given an $m \times n$ board representing an element of $\mathcal{T}_{mn}$, we use the following algorithm to generate its corresponding element of $\mathcal{T}_m(\mathcal{T}_n)$. The algorithm starts from the first row of the $m \times n$ board and proceeds downwards, and illustrated in Figs. 4–6.

Case 1a: (Row has no out of phase dominoes. Starts a or D.) Suppose the given row (call it row k) begins with a light square or a domino. Then the board is mapped to a metatiling with a light square at cell k whose embedded minitiling is the same as the tiling of row k. For example, in Fig. 4 we see that the first row (*Da*) begins with a domino, so the corresponding metatiling has a light square in the first cell with embedded minitiling *Da*, the same as the first row of the board. The mapping of row 6 in Fig. 4 is another example of this case.

Case 1b: (Row has no out of phase dominoes. Starts b.) Suppose row k does not contain part of an out of phase domino and begins with a dark square. Here the board is mapped to a metatiling with a dark square at cell k whose embedded minitiling is the same as the tiling of row k, except the initial dark square is changed to a light square. (This color swap is made so that the resulting minitiling is an element in $\mathcal{T}_n$.) For example, in Fig. 4 we see that the second row (*bab*) begins with a dark square. Thus the corresponding metatiling has a dark square in the second cell and the embedded minitiling is nearly the same (*aab*), except the initial square has changed color from dark to light. Note that the weight of the minitiling *bab* has the same weight as *aab*.

Case 2a: (Row ends with out of phase dominoes. Tailswappable.) Suppose row k is the first row to contain part of an out of phase domino. Since it is the first such row, it must have the out of phase domino starting at the last cell. Since row k contains an out of phase domino, it cannot be mapped directly to cell k of the metatiling (and its embedded minitiling).

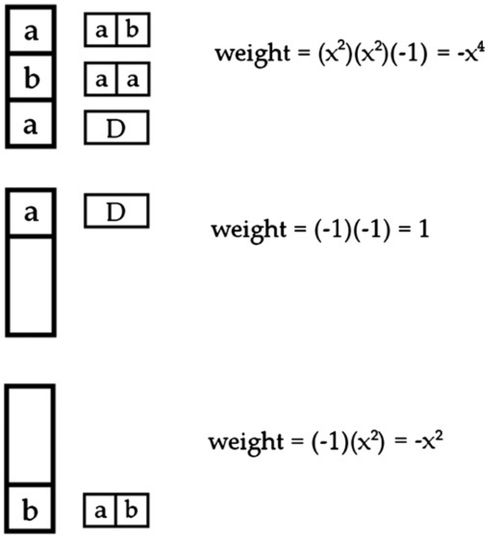

Fig. 2. Here are three examples of metatilings in the set $\mathcal{T}_3(\mathcal{T}_2)$. Notice that each square in each length 3 metatiling has an embedded length 2 minitiling. The weight of the square in the metatiling is the same as the weight of the embedded minitiling. Dominoes do not have embedded minitilings and always have weight -1.

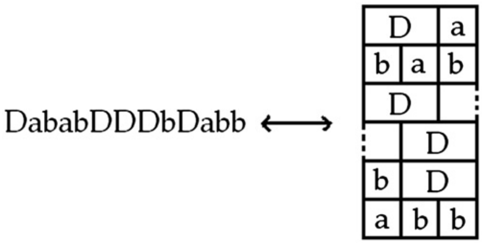

Fig. 3. This is an example of how to convert an 18-tiling to 6×3 board.

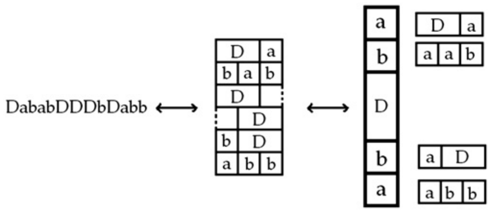

Fig. 4. In this example, the bijection creates a domino metatile.

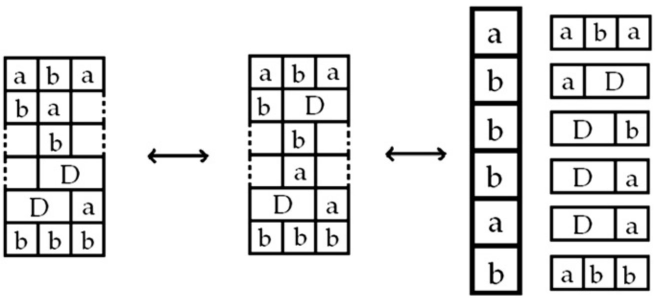

Fig. 5. In this tiling, the bijection performs a tail-swap.

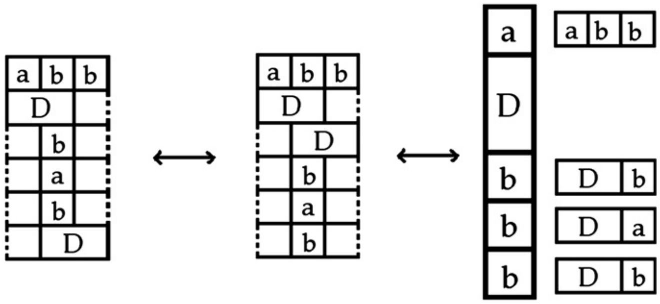

Fig. 6. In this example, the bijection creates a domino, causing some rows to shift down.

Instead, we must first tailswap row k with the first row after k that does not end with an out of phase domino (call this row j). If rows k and j are tailswappable, tailswap them. For an example of tailswapping see Fig. 5.

Once this tailswap has occurred, row k no longer contains an out of phase domino. Therefore, apply case 1 to obtain cell k of the corresponding metatiling and its corresponding minitiling.

In contrast, each of the rows $k+1, k+2, \ldots, j-1, j$ has part of an out of phase domino in cell 1 and in cell n. Each row is mapped to a dark square in the metatiling where the embedded minitiling is the tiling of the row, except the two out of phase domino pieces are put together to form a domino at the beginning of the minitiling.

In Fig. 5, $k=2$ and $j=4$. So rows 2 and 4 are tailswapped. After tailswapping, new row 2 is mapped using case 1(b), after which new rows 3 and 4 are mapped as just described.

Case 2b: (Row ends with out of phase dominoes. NOT tailswappable.) Again, suppose row k is the first row to contain part of an out of phase domino in cell n and that row j is the first row after k that does not end in an out of phase domino. Unlike case 2(a), suppose now that rows k and j cannot be tailswapped. Since n is odd, the only situation where this is possible is when rows j and k contain only dominoes. In this case, insert the tiles in row j between rows k and $k+1$, effectively shifting rows $k+1$ through $j-1$ down one row. The cells k and $k+1$ of the corresponding metatiling are covered by a domino of weight -1. Note that since n is odd, $(-1)^n=-1$, so the metadomino has the same weight as the n dominoes that it represents.

For a simple example, in Fig. 4, $k=3$ and $j=4$. Since these rows are already next to each other, $k+1=j$, so no shifting needs to be done. For a more detailed example, see Fig. 6. Here $k=2$ and $j=6$, so row 6 becomes row 3 and rows 3–5 are shifted down to 4–6.

Applying the algorithm above to each row, starting at row 1 and working down, each element in $\mathcal{T}_{mn}$, thought of as an $m \times n$ board, is sent to a unique element of $\mathcal{T}_m(\mathcal{T}_n)$. To show that this algorithm produces a bijection is straightforward. However, we will now point out some of the subtleties of the process.

First, notice that every image of a board in $\mathcal{T}_{mn}$ is in fact an element of $\mathcal{T}_m(\mathcal{T}_n)$. Since the first row of the board cannot start with an out of phase domino or dark square, the first tile of the metatiling must be either a light square or a domino, which fits the restriction on the metatilings. Furthermore, any minitiling embedded in any square of the metatiling cannot start with a dark square, since in case 1(b), if row k of the board starts with a dark square, cell k is a dark square, but the color of the first tile in the embedded minitiling is switched from dark to light. So the image of each board satisfies the restrictions of $\mathcal{T}_m(\mathcal{T}_n)$.

Second, the map is surjective. Every element in $\mathcal{T}_m(\mathcal{T}_n)$ has a preimage, which can be found by applying the inverse of this algorithm, starting at row m and working up the rows.

Third, the map is injective. Suppose two $m \times n$ boards have the same image. Since the image is created by working linearly down the rows (or a block of contiguous rows if tailswapping occurs) of the board, the only way two boards could have images with identical metatilings and embedded minitilings is if each of their rows were the same.

Finally, notice that the map preserves weights. For each row being mapped to a square metatile, the associated minitiling has the same weight as that row. And as noted earlier, each pair of rows mapped to a domino (necessarily the pair contained n dominoes) is mapped to a domino with weight -1.

Now we consider the case where n is even.

Case 2: *n is even*. Here we use the same mapping used in the odd case to map elements from $\mathcal{T}_{mn}$ to $\mathcal{T}_m(\mathcal{T}_n)$. However, notice that because n is even, tailswapping is always possible, eliminating the need for case 2(b). It follows that no image of an $m \times n$ board contains a domino. Furthermore, no image of an $m \times n$ board contains the metatiling ab with each embedded minitiling containing $n/2$ dominoes ($D^{n/2}$). We call this an *all-domino ab*. We call the set of metatilings with a domino or an all-domino ab, the exceptional tilings of $\mathcal{T}_m(\mathcal{T}_n)$. Now we show that the sum of the weights of the exceptional tilings of $\mathcal{T}_m(\mathcal{T}_n)$ is zero.

Suppose $\sigma \in \mathcal{T}_m(\mathcal{T}_n)$ is an exceptional tiling. Let k be the smallest number such that metatiling cells k and $k+1$ are either covered by a domino or by an all-domino ab. If k and $k+1$ are covered by a domino, then define $f(\sigma)$ to be the tiling where that domino (of weight -1) is replaced by an all-domino ab, and all other metatiles and minitiles are unchanged.

Note that since n is even, the all-domino minitiling associated with ab has weight $(-1)^n = 1$. On the other hand, if k and $k+1$ are covered by an all-domino ab, then we let $f(\sigma)$ be the tiling where that ab has been replaced by a domino. Hence f is an involution, and $w(f(\sigma)) = -w(\sigma)$. By pairing up each element with its image under f, we see that combined weight of all exceptional tilings is zero.

Hence, in the case where n is even, while our weight-preserving map is no longer a bijection between $\mathcal{T}_{mn}$ and all of $\mathcal{T}_m(\mathcal{T}_n)$, we do have a weight preserving bijection from $\mathcal{T}_{mn}$ to a subset of $\mathcal{T}_m(\mathcal{T}_n)$, where the sum of the weights of elements not hit by the bijection is 0.

Therefore,

$$T_m(T_n(x)) = T_{mn}(x).$$

4. A composition formula for $U_{mn-1}(x)$

Next we will prove the related composition identity for the Chebyshev polynomials of the second kind. The proof uses the same weight preserving bijection as the previous proof, but with a few minor changes. First, we make two quick definitions that will be useful in the next proof.

Definition 2. A row of a board is *closed on the left* if its first cell does not contain half of an out of phase domino. Likewise, a row is *closed on the right* if its last cell does not contain half of an out of phase domino.

Definition 3. A row of a board is *open on the left* if its first cell contains half of an out of phase domino. Similarly, a row is *open on the right* if its last cell contains half of an out of phase domino.

Identity 4. If n, m are nonnegative integers, then

$$U_{m-1}(T_n(x))U_{n-1}(x) = U_{mn-1}(x).$$

Proof. Consider the tilings in the set $\mathcal{U}_{mn-1}$. Write each tiling as an $m \times n$ board with the first cell removed, referred to as a *notched board*. Our goal is to convert each notched board into an unrestricted regular tiling of length $n-1$ and an unrestricted length $m-1$ metatiling with restricted length minitilings of length n. Hence we want a weight preserving bijection taking $\mathcal{U}_{mn-1}$ to $\mathcal{U}_{n-1} \times \mathcal{U}_{m-1}(\mathcal{T}_n)$. The overall idea is that the first row of the notched board corresponds to the unrestricted $(n-1)$- tiling in $\mathcal{U}_{n-1}$. The remaining $m-1$ by n board corresponds to the unrestricted length $m-1$ metatiling with restricted length n minitilings in $\mathcal{U}_{m-1}(\mathcal{T}_n)$.

Suppose that row 1 of the notched board is closed on the right. Then row 1, of length $n-1$, is mapped directly to the unrestricted $n-1$ tiling. Then we are left with an $m-1$ by n board that is closed on the left of its first row. See Fig. 7 for an example of this case.

On the other hand, suppose that row 1 of the notched board is open on the right. To be able to map row 1 to the unrestricted $(n-1)$- tiling, we need it to be closed on the right. So we will tailswap it with the first available row that is closed on the right. Since cell 1 has been removed, the first row is always breakable after cell 1. Hence, we are guaranteed that we can tailswap row 1 with the first row that is closed on the right. Once the tailswap has been performed, we map the new row 1 (now closed on the right) directly to the unrestricted $n-1$ tiling. We are then left with an $m-1$ by n board that is open on the left of its first row. See Fig. 8 for an example of this case.

Next, we convert the $m-1$ by n board to an unrestricted $m-1$ metatiling with embedded restricted n tilings, in $\mathcal{U}_{m-1}(\mathcal{T}_n)$.

When n is odd, we proceed almost exactly as in the proof of Identity 3. The only difference is that because our $m-1$ by n board can be open on the left of its first row (which could not happen in the previous proof), the corresponding $m-1$ metatilings are no longer restricted. For instance, in Fig. 8, the first row of the 5 by 3 board is open on the left, and therefore the first cell in the length 5 metatiling is a dark square, which could not have happened in the previous proof. Thus it makes sense that our $m-1$ by n boards get mapped to *unrestricted* length $m-1$ metatilings. Thus we have a weight preserving bijection between $\mathcal{U}_{nm-1}$ and $\mathcal{U}_{m-1}(\mathcal{T}_n) \times \mathcal{U}_{n-1}$.

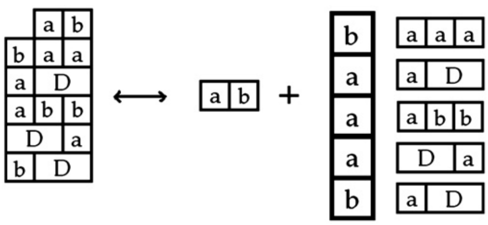

Fig. 7. A notched board that does not need tailswapping.

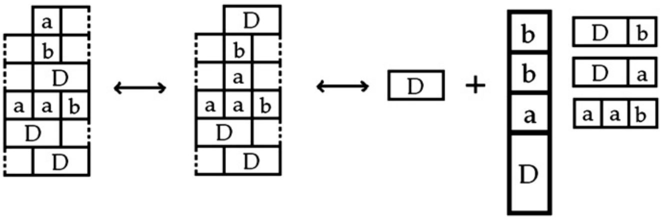

Fig. 8. A notched board that does require tailswapping.

When n is even, then we follow the same proof strategy incorporated in Identity 3, again noting that because the first row of the board can be open on the left, our metatilings are unrestricted. Like the last proof, this mapping is a weight preserving injective map, where the exceptional elements without preimages have total weight zero.

It follows that, regardless of the parity of n,

$$U_{m-1}(T_n(x))U_{n-1}(x) = U_{mn-1}(x). \qquad \square$$

References

Benjamin, A.T., Quinn, J.J., 2003. Proofs That Really Count: The Art of Combinatorial Proof. Mathematical Association of America, Washington, DC.

Benjamin, A.T., Ericksen, L., Jayawant, P., Shattuck, M., 2010. Combinatorial trigonometry with Chebyshev polynomials. Journal of Statistical Planning and Inference, this issue, doi:10.1016/j.jspi.2010.01.011.

Benjamin, A.T., Walton, D., 2009. Counting on Chebyshev polynomials. Mathematics Magazine 82, 117–126.

Rivlin, T.J., 1990. Chebyshev Polynomials: From Approximation Theory to Algebra and Number Theory. Wiley, New York.

Shapiro, L.W., 1981. A combinatorial proof of a Chebyshev polynomial identity. Discrete Mathematics 34, 203–206.

Walton, D., 2007. A tiling approach to Chebyshev polynomials, Senior Thesis, Harvey Mudd College, Claremont, CA.