

Algorithmically Distinguishing Irreducible Characters of the Symmetric Group

Timothy Y. Chow Jennifer Paulhus

May 2020

Abstract

Suppose that χ_λ and χ_μ are distinct irreducible characters of the symmetric group S_n . We show how to construct explicitly a permutation $\pi \in S_n$ such that $\chi_\lambda(\pi)$ is provably different from $\chi_\mu(\pi)$. In fact, we show a little more. Suppose f is an irreducible character of S_n , but we do not know which one it is, and we are given only oracle access to f . We give an algorithm that, using only polynomially many queries to f , determines which irreducible character f is.

1 Introduction

This paper is motivated by the following question. Suppose that we are given two distinct irreducible characters χ_λ and χ_μ of the symmetric group S_n . How hard is it to find a permutation $\pi \in S_n$ such that $\chi_\lambda(\pi) \neq \chi_\mu(\pi)$?

Surprisingly, this simple and natural question does not seem to have been considered before in the literature, perhaps because there is a sense in which the answer is, “not hard.” Namely, if one simply tries various permutations—especially permutations with a lot of fixed points—then empirically it seems to take at most a few tries to find a π such that $\chi_\lambda(\pi) \neq \chi_\mu(\pi)$.

However, proving that the above heuristic procedure always works does not seem to be easy. Computing $\chi_\lambda(\pi)$ for π with many fixed points leads to the consideration of hook-length formulae, and we find ourselves needing to show that the numbers that arise cannot satisfy an unexpected linear relation. It seems hard to rigorously rule out the possibility of “accidental”

linear relations. Known results on character values (e.g., [1]) do not seem to be strong enough to yield an immediate answer to our question.

Nevertheless, in this paper we describe an algorithm that provably solves the following slightly harder problem. We are given a positive integer n as well as *oracle access* to a function f on the symmetric group S_n , meaning that the only way we can obtain information about f is to submit a *query* (i.e., an input value that we are free to choose) $\pi \in S_n$ to an oracle, which then truthfully tells us the value of $f(\pi)$. We know that f is an irreducible character of S_n , but we do not know which one. Our job is to determine, via a sequence of queries to the oracle, which irreducible character f is. Our queries are allowed to be *adaptive*; that is, we may examine the results of previous queries when deciding which query to submit next.

Theorem 1. *There is a deterministic algorithm that, given oracle access to a function f that is promised to be an irreducible character of S_n , determines which irreducible character it is, using a number of queries that is polynomial in n . The running time of the algorithm is also polynomial in n .*

An interesting feature of our algorithm is that, instead of querying permutations with many fixed points, it mainly queries permutations with rather long cycles and very few fixed points; the possible border-strip tableaux are thereby severely constrained, allowing us to enumerate them explicitly and prove the inequalities we want.

As will become clear when we describe the algorithm, it can be readily adapted to the problem we stated initially, of finding π such that $\chi_\lambda(\pi) \neq \chi_\mu(\pi)$. One simply simulates the algorithm using $f = \chi_\lambda$ until one reaches a query that rules out the possibility that $f = \chi_\mu$.

After some necessary preliminaries in Section 2, we describe the overall structure of our algorithm in Section 3. There is one step of the algorithm that, as far as we can see, requires a complicated case analysis; this is carried out in Section 4.

2 Background

In this section we review some standard material. It turns out that in the case of the symmetric group, there is a natural bijection between irreducible characters and conjugacy classes, and conjugacy classes are naturally indexed

by partitions (the lengths of the cycles of the permutation). So the first order of business is to review some of the combinatorics of partitions.

2.1 Partitions and Young Diagrams

Let n be a positive integer. A *composition of n* is defined to be a sequence $\alpha = (\alpha_1, \dots, \alpha_\ell)$ of positive integers such that $\sum_{i=1}^{\ell} \alpha_i = n$. A *partition of n* is a composition $\lambda = (\lambda_1, \dots, \lambda_\ell)$ of n such that

$$\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_\ell.$$

A partition of n may be visualized as a *Young diagram*, which is a left-justified grid of boxes having λ_i boxes in row i . See Figure 1 for an example. The boxes of a Young diagram are coordinatized in the same way that matrix entries are coordinatized; i.e., box (i, j) is the j th box from the left in the i th row from the top, where $i \geq 1$ and $j \geq 1$.

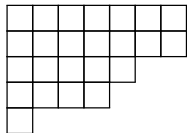


Figure 1: Young Diagram of the Partition $(7, 7, 5, 4, 1)$

Following standard terminology [4, Section 7.2], we define the *conjugate λ'* of a partition λ to be the sequence of column lengths of the Young diagram of λ . For example, the conjugate of $(7, 7, 5, 4, 1)$ is $(5, 4, 4, 4, 3, 2, 2)$. A partition λ is *self-conjugate* if $\lambda = \lambda'$.

The *principal diagonal* of a Young diagram is the set of boxes with coordinates (i, i) for some i .

For the purposes of this paper, it will be convenient to think of Young diagrams in a slightly nonstandard manner, namely as a nested sequence of *principal hooks*.

Definition 1. The *i th principal hook* of a Young diagram D is the set

$$H_i := \{(i, j) \in D : j \geq i\} \cup \{(j, i) \in D : j \geq i\}.$$

The *i th principal hook length h_i* is the area (i.e., cardinality) of H_i .

For example, Figure 2 illustrates the (nonempty) principal hooks of the Young diagram from Figure 1, where we have colored H_1 red, H_2 orange, H_3 yellow, and H_4 green.

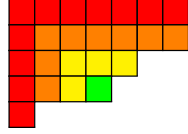


Figure 2: Principal Hook Decomposition of $(7, 7, 5, 4, 1)$

The horizontal part of a principal hook is known as its *arm* and the vertical part is known as its *leg*. Note that the arm of H_i extends at least as far to the right as the arm of H_{i+1} , and may extend farther. If it extends farther, we refer to the extra boxes as the *i th arm overhang* (and similarly for the legs). More formally, we have the following definitions.

Definition 2. Suppose that a Young diagram has exactly k nonempty principal hooks. Then for $1 \leq i \leq k$, the *i th arm overhang* is

$$\{(i, j) \in H_i : j > i \text{ and } (i + 1, j) \notin H_{i+1}\}.$$

Similarly, the *i th leg overhang* is

$$\{(j, i) \in H_i : j > i \text{ and } (j, i + 1) \notin H_{i+1}\}.$$

If the *i th arm overhang* and the *i th leg overhang* have different cardinalities, then we call the smaller one the *i th short overhang* and we let a_i be its cardinality; similarly we call the longer one the *i th long overhang* and we let b_i be its cardinality. If the *i th arm overhang* and the *i th leg overhang* have the same cardinality then we set both a_i and b_i equal to that cardinality.

In our running example, the 1st leg overhang and the 2nd and 3rd arm overhangs are nonempty, as indicated by the colored boxes in Figure 3, and $a_1 = a_2 = a_3 = a_4 = 0$ and $b_1 = 1, b_2 = 2, b_3 = 1, b_4 = 0$.

Lemma 1. *If there are more than k principal hooks, then $h_k - a_k \geq 3$.*

Proof. By assumption, the $(k + 1)$ st principal hook has at least one box, so the k th principal hook must have at least three boxes that do not overhang the $(k + 1)$ st principal hook. \square

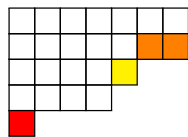


Figure 3: Nonempty Overhangs of $(7, 7, 5, 4, 1)$

The next definition plays a key role in this paper.

Definition 3. The *doppelgänger* $\hat{\lambda}$ of a partition λ is the partition whose Young diagram \hat{D} is the same as the Young diagram D of λ except that the cardinality of the 1st arm overhang of $\hat{\lambda}$ equals the cardinality of the 1st leg overhang of λ , and vice versa.

For example, the doppelgänger of $(7, 7, 5, 4, 1)$ is $(8, 7, 5, 4)$, obtained by swapping the 1st leg overhang (whose cardinality is 1) with the 1st arm overhang (whose cardinality is 0). Note that λ and $\hat{\lambda}$ have the same values of h_i , a_i , and b_i for all i .

The reason we have chosen the term “doppelgänger” is that it turns out to be surprisingly tricky to find a permutation π such that $\chi_\lambda(\pi)$ is provably different from $\chi_{\hat{\lambda}}(\pi)$.

2.2 Border-Strip Tableaux

In the introduction, we mentioned but did not define irreducible characters of S_n . The only fact about irreducible characters of S_n that we need in this paper (besides the fact that they are indexed by partitions of n) is a famous result known as the *Murnaghan–Nakayama rule*, which gives a combinatorial rule for computing them. In this section, we give a complete statement of the Murnaghan–Nakayama rule, so the reader unfamiliar with the concept of an irreducible character may take the Murnaghan–Nakayama rule as a definition. For more details, including a proof of the Murnaghan–Nakayama rule, the interested reader can consult Sagan [3, Section 4.10] or Stanley [4, Chapter 7].

To state the Murnaghan–Nakayama rule, we must define border strips (also known as rim hooks or ribbons) and border-strip tableaux.

Definition 4. A *border strip* is a finite set of boxes such that in each row, the boxes in that row are contiguous, and except for the top row, the *rightmost*

box in each row lies directly underneath the *leftmost* box in the row above it. The *area* or *size* of a border strip is the total number of boxes. If B is a border strip, its *height* $h(B)$ is the number of rows of B minus 1.

In Figure 4, the four border strips have areas 11, 7, 1, 8 and heights 3, 2, 0, 3 respectively. Note that the definition of a border strip ensures that the boxes are orthogonally connected, and that a border strip never contains two boxes (i, j) and (i', j') with $i < i'$ and $j < j'$.

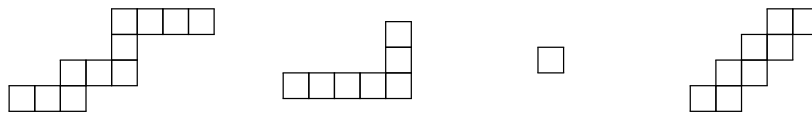


Figure 4: Four Examples of Border Strips

Definition 5. Let n be a positive integer, let λ be a partition of n , and let α be a composition of n . A *border-strip tableau (BST)* of shape λ and type α is a tiling of the Young diagram of λ with border strips such that

1. the area of the i th border strip is α_i , and
2. if the number i is written in each box of the i th border strip, then the numbers weakly increase across every row and down every column.

Figure 5 shows an example of a BST where we have colored each border strip with a different color as a visual aid.



Figure 5: Border-Strip Tableau of Shape $(5,4,4,1)$ and Type $(4,5,1,1,3)$

In Section 4.2, we will want to consider *partial border-strip tableaux*. We will not give a completely formal definition, but the idea is that we take a BST and consider only the placement of the first few border strips, ignoring the placement of the remaining border strips.

We are now ready to state the Murnaghan–Nakayama rule.

Theorem 2. *Let λ be a partition of n , and let χ_λ be the irreducible character of S_n indexed by λ . If $\pi \in S_n$ and (α_i) is the sequence of cycle lengths of π , then*

$$\chi_\lambda(\pi) = \sum_T \prod_{B \in T} (-1)^{h(B)}, \quad (1)$$

where the sum is over all BSTs T of shape λ and type α , and the product is over the border strips B that tile T .

The expression $\prod_{B \in T} (-1)^{h(B)}$ appearing in Equation (1) is called the *sign* of the BST T .

The alert reader may notice that the above statement of Theorem 2 speaks of “the” sequence of cycle lengths of π , but there is no canonical ordering on the set of cycle lengths of a permutation. It is a remarkable and nontrivial fact that Theorem 2 remains true no matter what ordering is chosen. For example, let $\lambda = (5, 4, 2)$ and let π be a permutation with cycle lengths 6, 3, and 2. If we let $\alpha = (6, 3, 2)$ then there are two BSTs of shape λ and type α , as shown in Figure 6. One of these BSTs contributes $+1$ to the sum in Equation (1) and the other contributes -1 , so $\chi_\lambda(\pi) = 0$.



Figure 6: Border-Strip Tableaux of Shape $(5,4,2)$ and Type $(6,3,2)$

On the other hand, the reader can check that if we let $\alpha = (6, 2, 3)$ then there are *no* BSTs of shape λ and type α , so again Equation (1) tells us that $\chi_\lambda(\pi) = 0$, but for a “different reason” combinatorially.

Since the value of $\chi_\lambda(\pi)$ depends only on the sequence of cycle lengths of π , for the rest of this paper we will usually regard χ_λ as a function of a composition of n , and our queries will be compositions rather than permutations.

We conclude this section with a couple of simple but useful facts about BSTs.

Lemma 2. *A border strip of a BST cannot contain more than one box on the principal diagonal.*

Proof. Let i be the smallest positive integer such that the border strip contains the box of the Young diagram with coordinates (i, i) . This box—call

it x —appears in some row, say the j th row, of the border strip. By the definition of a border strip, no box in the border strip in any row below the j th can appear further to the right than x does, so in particular no other box on the principal diagonal can appear in the border strip. \square

Lemma 3. *In any BST, for all k , the first k border strips fit into the first k principal hooks.*

Proof. We use induction on k . The case $k = 0$ is vacuously true. Suppose the claim is true for some $k \geq 0$. For the claim to fail for $k + 1$, the $(k + 1)$ st border strip must contain a box x in the i th principal hook for some $i \geq k + 2$. Then since the box $(k + 2, k + 2)$ lies (weakly) above and to the left of x , the weakly increasing property of a BST forces $(k + 2, k + 2)$ to belong to border strip j for some $j \leq k + 1$. But by the induction hypothesis, $(k + 2, k + 2)$ does not belong to any of the first k border strips, so $j = k + 1$. The same argument proves that the box $(k + 1, k + 1)$ must also belong to border strip $k + 1$. This contradicts Lemma 2. \square

3 High-Level Structure of the Algorithm

The algorithm consists of a *forward pass* followed by a *backward pass*. As we explain in Section 3.1, during the forward pass, we determine the principal hook lengths h_i one at a time, in the order h_1, h_2, h_3, \dots . As we explain in Section 3.2, during the backward pass, we determine the principal hooks (i.e., their actual shapes, not just their lengths) in *reverse* order, starting with the last principal hook and working backward. It turns out that for the backward pass, it suffices to show how to recover the 1st principal hook given that we know all the later hooks, and we accomplish this in two steps:

1. Determine the overhang lengths a_1 and b_1 . We give the details in Section 4.1. This reduces the problem to distinguishing between λ and its doppelgänger $\hat{\lambda}$.
2. Distinguish between λ and $\hat{\lambda}$. This step is surprisingly complicated, and subdivides into several cases, which we explain in detail in Section 4.2.

3.1 The Forward Pass

In the forward pass, we determine the principal hook lengths of λ . To determine h_1 , we define, for each $i \in \{0, 1, \dots, n - 1\}$, the composition $\alpha^{(i)}$ of n

by

$$\alpha_j^{(i)} = \begin{cases} n - i, & \text{if } j = 1; \\ 1, & \text{if } 2 \leq j \leq i + 1. \end{cases}$$

We submit the queries $\alpha^{(0)}, \alpha^{(1)}, \alpha^{(2)}, \dots$ successively to the oracle, stopping as soon as we encounter a nonzero value.

The key observation is that by Lemma 3 (with $k = 1$), there cannot exist a BST of shape λ and type $\alpha^{(i)}$ if $\alpha_1^{(i)} > h_1$, because the 1st border strip would be simply too large to fit inside the 1st principal hook. So if $\alpha_1^{(i)} > h_1$, then $\chi_\lambda(\alpha^{(i)}) = 0$. Conversely, if $\alpha_1^{(i)} = h_1$, then the BSTs of shape λ and type $\alpha^{(i)}$ are obtained by letting the 1st border strip cover the entire 1st principal hook, and then arranging the numbers from 2 through $i + 1$ in the rest of λ so that they increase across rows and down columns. There is always at least one such BST; if there are more, then they all have the same sign, and hence in particular, $\chi_\lambda(\alpha^{(i)}) \neq 0$. Therefore, if i_0 is the smallest i such that $\chi_\lambda(\alpha^{(i)}) \neq 0$, then $h_1 = \alpha_1^{(i_0)} = n - i_0$.

Once we know h_1 , we can determine h_2 by a similar procedure. If $h_1 = n$ then we are done. Otherwise, we fix the size of the 1st border strip at h_1 , and find, by repeated “guessing and checking,” the largest possible size of the 2nd border strip; this is h_2 . More formally, we define compositions $\beta^{(i)}$ of n as follows:

$$\beta_j^{(i)} = \begin{cases} h_1, & \text{if } j = 1; \\ \min(h_1 - 2, n - h_1) - i, & \text{if } j = 2; \\ 1, & \text{if } 3 \leq j \leq n - \beta_1^{(i)} - \beta_2^{(i)} + 2. \end{cases}$$

We submit the queries $\beta^{(0)}, \beta^{(1)}, \beta^{(2)}, \dots$ successively to the oracle, stopping as soon as we encounter a nonzero value. Since $\beta_1^{(i)} = h_1$, Lemma 3 tells us that in any BST of shape λ and type $\beta^{(i)}$, the 1st border strip must entirely cover the 1st principal hook. If $\beta_2^{(i)} > h_2$, then there cannot be any BSTs of shape λ and type $\beta^{(i)}$, because the 2nd border strip would be too large to fit inside the 2nd principal hook. Conversely, if $\beta_2^{(i)} = h_2$, then the BSTs of shape λ and type $\beta^{(i)}$ are obtained by letting the 1st and 2nd border strips cover the entire 1st and 2nd principal hooks respectively, and then arranging the numbers from 3 to $n - \beta_1^{(i)} - \beta_2^{(i)} + 2$ in the rest of λ so that they increase across rows and down columns. As before, there is always at least one BST of this form, and if there are more, then they all have the same sign. So if i_0 is the smallest i such that $\chi_\lambda(\beta^{(i)}) \neq 0$, then $h_2 = \beta_2^{(i_0)}$.

The pattern should now be clear. Given that we know the first few principal hook lengths, we fix the corresponding border-strip sizes to be equal to the known principal hook lengths, and use the next border strip to “guess” the size of the next principal hook length, starting with the largest conceivable value and working our way downward. If our guess is too large, then the oracle will return zero. As soon as the oracle returns a nonzero value for χ_λ , that tells us that our guess for the size of the next principal hook length is correct. In this way, we can recover all the principal hook lengths.

3.2 The Backward Pass

Knowing the principal hook lengths of λ does not, in general, determine λ uniquely, because there can be many different hooks with the same hook length. Our overall strategy for determining λ will be to recover the principal hooks themselves (not just their lengths) inductively, starting with the *last* (or innermost) principal hook, and working backward one principal hook at a time. Each principal hook will be recovered in two steps; first, we will recover the overhang lengths a_i and b_i , and then (if $a_i \neq b_i$) we will recover which of a_i and b_i is the arm overhang length and which is the leg overhang length.

In this backward pass, we claim that we may assume without loss of generality that we know λ completely except that we are unsure about the shape of the 1st principal hook. To see this, suppose that we know only the shapes of the j th principal hooks of λ for j greater than some value j_0 , and we want to recover the j_0 th principal hook of λ . Consider what happens if $\alpha_i = h_i$ for all $i < j_0$. If we query the value of $\chi_\lambda(\alpha)$, then by the same kind of argument we gave in Section 3.1, each of the first $j_0 - 1$ border strips is forced to cover the corresponding principal hook entirely. In effect, we are querying a smaller shape μ —one that has been obtained from λ by “stripping off” the first $j_0 - 1$ principal hooks, and that we know completely except for its 1st principal hook. In other words, given a sequence of queries for μ that recover its 1st principal hook, we can simply prepend $\alpha_i = h_i$ for all $i < j_0$ to each of these queries; this will give us a sequence of queries for λ that allow us to recover its j_0 th principal hook.

Our description of the backward pass, and our proof of its correctness, proceeds by induction on the number of principal hooks of λ . The induction step—which assumes that the exact shapes of all the principal hooks H_i with $i \geq 2$ are known, and shows how to recover the principal hook H_1 —is the

topic of Section 4.

Let us now establish the base case, when we know that λ is a hook, and we know its length h_1 , but we do not know its precise shape. Let λ_1 denote the number of boxes in the first row of λ . If we let $\alpha_i = 1$ for $1 \leq i \leq n$, then every BST of shape λ and type α has positive sign, and Theorem 2 implies that

$$\chi_\lambda(\alpha) = \binom{n-1}{\lambda_1-1},$$

because we may choose any $\lambda_1 - 1$ of the numbers in the set $\{2, 3, \dots, n\}$ to place in the first row of our BST, and everything else about the BST is uniquely determined. For fixed n , the binomial coefficients $\binom{n-1}{r}$ are distinct, except that $\binom{n-1}{r} = \binom{n-1}{n-1-r}$. Therefore once we know $\chi_\lambda(\alpha)$, we know that our shape is either λ or its doppelgänger $\hat{\lambda}$. If $\lambda = \hat{\lambda}$ (i.e., λ is self-conjugate) then we are done. Otherwise, we define

$$\beta_i := \begin{cases} 2, & \text{if } i = 1; \\ 1, & \text{if } 2 \leq i \leq n-1. \end{cases}$$

We query the value of $\chi_\lambda(\beta)$. There are two types of BSTs of shape λ and type β , depending on whether the first border strip is arranged horizontally or vertically; in the former case, the first border strip has height 0 so the BST has positive sign, whereas in the latter case, the first border strip has height 1 so the BST has negative sign. A straightforward application of Theorem 2 yields

$$\chi_\lambda(\beta) = \binom{n-2}{\lambda_1-2} - \binom{n-2}{\lambda_1-1}.$$

On the other hand, $\chi_{\hat{\lambda}}(\beta) = -\chi_\lambda(\beta)$. Therefore, as long as $\chi_\lambda(\beta) \neq 0$, we can distinguish between λ and $\hat{\lambda}$ just by examining whether the answer to our query is positive or negative. But the only way that $\chi_\lambda(\beta)$ can be zero is if $(\lambda_1 - 2) + (\lambda_1 - 1) = n - 2$, which can happen only if λ is self-conjugate—a case that we already dealt with above.

4 Recovering the First Principal Hook

For the remainder of this paper, we make the following standing assumptions: λ is a partition with principal hook lengths h_i and overhang lengths a_i and b_i ,

where $a_i \leq b_i$. We assume that λ has exactly $k+1$ nonempty principal hooks, with $k \geq 1$. We further assume that we know all the h_i as well as the exact shape formed by all the principal hooks H_i with $i \geq 2$.

4.1 Determining the Overhang Lengths a_1 and b_1

Our goal in this section is to explain how to determine the overhang lengths a_1 and b_1 . We define compositions $\alpha^{(i)}$ of n as follows:

$$\alpha_j^{(i)} = \begin{cases} h_1 - i, & \text{if } j = 1; \\ h_2 + i, & \text{if } j = 2; \\ h_j, & \text{if } j \geq 3. \end{cases}$$

We submit the queries $\alpha^{(1)}, \alpha^{(2)}, \alpha^{(3)}, \dots$ successively to the oracle, stopping as soon as we encounter a nonzero value.

Suppose that for some $i \geq 1$, there exists a BST T of shape λ and type $\alpha^{(i)}$. Lemma 3 implies that the 2nd border strip must be contained within the first two principal hooks, and because $\alpha_2^{(i)} > h_2$, the 2nd border strip must contain some—in fact, exactly i —boxes from the 1st principal hook. Let x be the box in the 1st principal hook adjacent to the overhang of length a_1 , and let y be the box in the 1st principal hook adjacent to the overhang of length b_1 . Then the 1st border strip cannot cover *both* x and y , because then the 2nd border strip (which must in particular contain the box $(2, 2)$) would not be able to contain any boxes from the 1st principal hook. See for example Figure 7, where the 1st border strip (which we have colored red, omitting the 1s in the boxes to avoid clutter) contains both x and y , and thus blocks the 2nd border strip (in orange) from containing any boxes from the 1st principal hook.

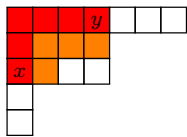


Figure 7: 1st Border Strip Contains Both x and y

But if $i \leq a_1$ then the 1st border strip is so long that it is forced to cover both x and y . Hence if $i \leq a_1$, then $\chi_\lambda(\alpha^{(i)}) = 0$.

Conversely, suppose that $i = a_1 + 1$. Then there is a BST of shape λ and type $\alpha^{(i)}$ in which

- the 1st border strip covers everything in the 1st principal hook except x and the overhang adjacent to x ;
- the 2nd border strip covers x , the overhang adjacent to x , and the entire 2nd principal hook; and
- for $j > 2$, the j th border strip covers the entire j th principal hook.

See for example Figure 8.

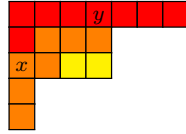


Figure 8: Border-Strip Tableau of Type $\alpha^{(a_1+1)}$

Moreover, if $a_1 < b_1$, then this is the *only* BST of shape λ and type $\alpha^{(a_1+1)}$, because any other placement of the 1st border strip would cover both x and y , and once the 1st border strip is placed, the areas of the border strips (together with the constraint imposed by Lemma 3) force a unique placement of the remaining border strips. Therefore, $\chi_\lambda(\alpha^{(a_1+1)}) = \pm 1$, and $i = a_1 + 1$ is the smallest value of i such that $\chi_\lambda(\alpha^{(i)}) \neq 0$.

On the other hand, if $a_1 = b_1$, then there is a second BST of shape λ and type $\alpha^{(a_1+1)}$, where the 1st border strip covers everything in the 1st principal hook except y and the overhang adjacent to y . In the case of one of these two BSTs, the heights of the first two border strips are $\lambda'_1 - 1$ and $\lambda'_1 - a_1 - 1$, and in the other case, the heights of the first two border strips are $\lambda'_1 - a_1 - 2$ and $\lambda'_1 - 2$. Therefore the two BSTs have the same sign, so $\chi_\lambda(\alpha^{(a_1+1)}) = \pm 2$, and again $i = a_1 + 1$ is the smallest value of i such that $\chi_\lambda(\alpha^{(i)}) \neq 0$.

Once the value of a_1 is determined, the value of b_1 is also determined, since $h_1 = h_2 + a_1 + b_1 + 2$.

4.2 Distinguishing Doppelgänger

Our goal in this section is to construct a query α that distinguishes between λ and its doppelgänger; i.e., such that $\chi_\lambda(\alpha) \neq \chi_{\hat{\lambda}}(\alpha)$. We assume that $a_1 \neq b_1$, since otherwise $\lambda = \hat{\lambda}$.

Definition 6. The *second imbalance* of λ , denoted $I(\lambda)$, is the smallest integer $i \geq 2$ such that $a_i \neq b_i$. If $a_i = b_i$ for $2 \leq i \leq k + 1$, then we set $I(\lambda) := \infty$.

The reason for calling $I(\lambda)$ the “second” imbalance is that the first imbalance is always 1, since we have assumed that $a_1 \neq b_1$. Figure 9 illustrates a partition with $k = 4$ and $I(\lambda) = 5$.

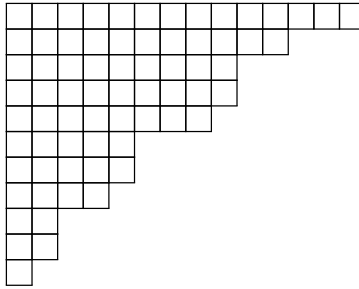


Figure 9: Partition with $k = 4$ and $I(\lambda) = 5$

Throughout this section, our queries will always start off the same way, as described in the following crucial definition.

Definition 7. Let $m := \min(I(\lambda), k + 1) - 1$, or in other words let $m = k$ if $I(\lambda) = \infty$ and $m = I(\lambda) - 1$ if $I(\lambda) < \infty$. Define

$$\alpha_i(\lambda) = \begin{cases} h_1 - a_1 - 1, & \text{if } i = 1; \\ h_i - a_i + a_{i-1}, & \text{if } 2 \leq i \leq m. \end{cases} \quad (2)$$

4.2.1 The Greedy Arrangement

As we now explain, there is always a way to create a partial border-strip tableau of shape λ in which the i th border strip has area $\alpha_i(\lambda)$, for $1 \leq i \leq m$. Namely, let us place the 1st border strip so that it entirely covers the 1st long overhang. The area of the 1st border strip is $\alpha_1(\lambda) = h_1 - a_1 - 1$, which is less than h_1 , so it is short enough to fit inside the 1st principal hook. We must also check that it is long enough to reach the box in the $(1, 1)$ position, since in any BST, the 1st border strip must cover $(1, 1)$. Because its area is $h_1 - a_1 - 1$, what the 1st border strip leaves uncovered (inside the 1st principal hook) is the 1st short overhang and the box adjacent to the 1st

short overhang—call this box x_1 . Since λ has at least 2 nonempty principal hooks, x_1 is *not* in the $(1, 1)$ position, so the 1st border strip does indeed cover $(1, 1)$. See the red border strip in Figure 10 for an example.

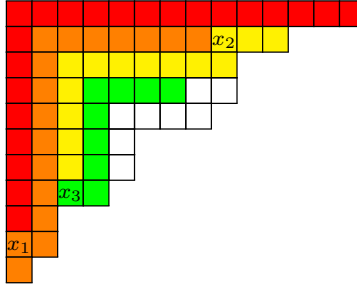


Figure 10: Greedy Arrangement

If $m \geq 2$, then let us place the 2nd border strip so that it entirely covers the 1st short overhang. Because x_1 is *not* in the 1st short overhang, it is adjacent to some box in the 2nd principal hook; indeed, it is adjacent to one end of the 2nd principal hook. Thus the 2nd border strip can “spill over” into the 2nd principal hook; indeed, it spills over by precisely

$$\alpha_2(\lambda) - (a_1 + 1) = h_2 - a_2 + a_1 - (a_1 + 1) = h_2 - a_2 - 1$$

boxes. Now since $m \geq 2$, the 2nd arm overhang and the 2nd leg overhang both have length a_2 , and there are at least 3 nonempty principal hooks. Arguing as we did for the 1st border strip, we conclude that what the 2nd border strip leaves uncovered (in the 2nd principal hook) is the overhang at the opposite end of the 2nd principal hook from x_1 , plus the box (call it x_2) adjacent to that overhang. Moreover, positioning the 2nd border strip in this fashion does indeed cover the $(2, 2)$ box, as it is required to do. See the orange border strip in Figure 10.

The placements of the remaining border strips (if they exist) follow the same recipe. The i th border strip is positioned to cover the boxes in the $(i - 1)$ st principal hook that are *not* covered by the $(i - 1)$ st border strip, namely the uncovered overhang and the box x_{i-1} adjacent to that overhang. The border strip then “spills over” into the i th principal hook, covering everything except the overhang at the opposite end of the i th principal hook and the box x_i adjacent to that overhang. This recipe works even when an overhang is empty, as illustrated by the green border strip in Figure 10. We formalize this construction with the following definition.

Definition 8. Let $(\alpha_i(\lambda))_{i=1,\dots,m}$ be given by Equation (2). A BST of shape λ is called *greedy* if the first m border strips have areas $(\alpha_i(\lambda))_{i=1,\dots,m}$ and are arranged in the manner described in the preceding paragraphs. It is called *non-greedy* if the first m border strips have areas $(\alpha_i(\lambda))_{i=1,\dots,m}$ but some of those first m border strips are arranged differently.

In Sections 4.2.2 to 4.2.4, we deal with the case when $I(\lambda) < \infty$. In a nutshell, our algorithm will submit queries α whose first few values are given by Equation (2), and we will argue that, for the values of α that we choose,

- the number of non-greedy BSTs of shape λ and type α equals the number of non-greedy BSTs of shape $\hat{\lambda}$ and type α , and
- the number of greedy BSTs of shape λ and type α has the opposite parity from the number of greedy BSTs of shape $\hat{\lambda}$ and type α .

Given these two facts, it follows that, regardless of the signs of the BSTs, $\chi_\lambda(\alpha) \not\equiv \chi_{\hat{\lambda}}(\alpha) \pmod{2}$, and therefore $\chi_\lambda(\alpha) \neq \chi_{\hat{\lambda}}(\alpha)$.

Finally, in Section 4.2.5, we deal with the case $I(\lambda) = \infty$.

4.2.2 Non-Greedy Border-Strip Tableaux

In what follows, we will overload the notation α_i , using it to mean both the i th part of a composition α , as well as the function given by Equation (2). There should be no confusion because we will only be considering compositions α whose first few parts coincide with the values given in Equation (2).

Let us make some general observations about non-greedy BSTs. Let α be a composition of n with $\alpha_i = \alpha_i(\lambda)$ for $1 \leq i \leq m$. Let T be a non-greedy BST of shape λ and type α . Let $i \leq m$ be the smallest number such that the i th border strip is *not* positioned as the greedy arrangement dictates. We claim that the only other allowable positions of the i th border strip are *slides* of its greedy position, meaning that we remove some boxes from one end of the border strip and append the same number of boxes to the other end; if we move $j \geq 1$ boxes from one end to the other then we call the resulting border strip the j th *slide*. For example, if $\lambda = (14, 11, 9, 9, 8, 5, 4, 4, 2, 2, 1)$, then Figure 11 illustrates the two possible slides of the 1st border strip if $i = 1$.

Figure 12 illustrates, for the same λ , the 1st and 3rd slides of the 2nd border strip if $i = 2$.

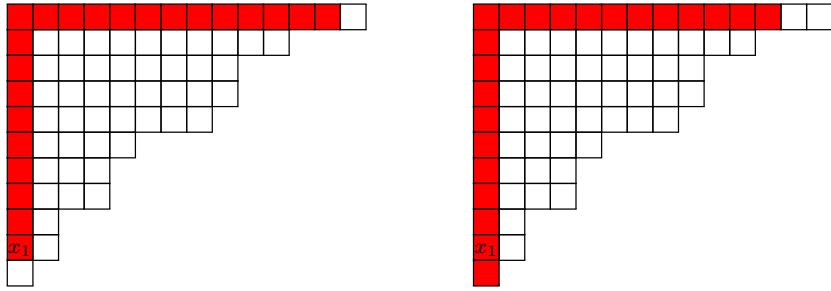


Figure 11: Non-Greedy Arrangements of 1st Border Strip

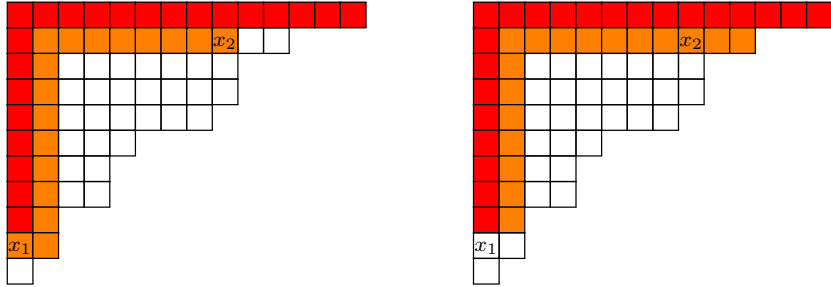


Figure 12: 1st and 3rd Slides of 2nd Border Strip

To see why slides are the only allowable positions of the i th border strip, recall that Lemma 3 implies that the first i border strips must lie within the first i principal hooks. The only available boxes for the i th border strip are in the $(i - 1)$ st and i th principal hooks, because the greedy arrangement entirely fills all the earlier principal hooks.

We now claim that the only boxes in the i th principal hook of T that are not covered by the i th border strip must lie in the i th arm overhang or the i th leg overhang.

To see why the claim is true if $i = 1$, note that every slide of the 1st border strip necessarily covers the box x_1 , and any uncovered boxes in the 1st principal hook lying beyond x_1 must be in the 1st short overhang. On the other hand, even if the 1st border strip is slid as far as possible, so that it covers the entire 1st short overhang, it will leave uncovered at most $a_1 + 1 \leq b_1$ boxes, which must therefore all lie in the 1st long overhang.

A similar argument establishes the claim if $i > 1$. For example, if $i = 2$, then every slide of the 2nd border strip covers x_2 , and any uncovered boxes

in the 2nd principal hook lying beyond x_2 must be in an overhang. On the other hand, even if the 2nd border strip is slid as far as possible, it will leave uncovered

$$h_2 - (h_2 - a_2 + a_1) = a_2 - a_1 \leq a_2$$

boxes at one end of the 2nd principal hook, so these boxes must all lie in an overhang.

In fact, we can say a little more. Consider the set S of boxes in the first i principal hooks that are *not* covered by the first i border strips. There are j boxes in S that were covered in the greedy arrangement and that are now uncovered as a result of the slide; call this set of boxes the *outer island* (see for example the blue boxes in Figure 13). The outer island comprises some boxes in an $(i - 1)$ st overhang, plus possibly x_{i-1} and some boxes in the adjacent i th overhang. The remaining boxes in S lie in the other i th overhang, at the other end of the i th border strip; call this set of boxes the *inner island* (see for example the violet boxes in Figure 13; note that the inner island may be empty). The reason we call these sets “islands” is that they are disconnected from each other and from all later principal hooks; i.e., no border strip can contain both a box from an (inner or outer) island and a box from the i' -th principal hook for $i' > i$, because they are disconnected from each other by the i th border strip. Finally, note also that neither island can contain a box from the principal diagonal; boxes in an overhang can never lie on the principal diagonal, so the only worry is that x_{i-1} might lie on the principal diagonal, but this can happen only if x_{i-1} is on the innermost hook, which is not possible because x_{i-1} is on the $(i - 1)$ st principal hook and there are at least i principal hooks.

Lemma 4. *If α is a composition of n with $\alpha_i = \alpha_i(\lambda)$ for $1 \leq i \leq m$, then the number of non-greedy BSTs of shape λ and type α is equal to the number of non-greedy BSTs of shape $\hat{\lambda}$ and type α .*

Proof. We describe a bijection between non-greedy BSTs of shape λ and type α and non-greedy BSTs of shape $\hat{\lambda}$ and type α . As above, if T is a non-greedy BST of shape λ and type α , we let i be the smallest number such that the i th border strip is *not* positioned greedily; the i th border strip in T is then the j th slide for some j . To construct the corresponding BST \hat{T} of shape $\hat{\lambda}$, we proceed as follows:

1. Arrange the first $i - 1$ border strips of \hat{T} greedily.

2. Put the i th border strip of \hat{T} in the j th slide position.
3. For each remaining border strip, if, in T , it lies in the inner or outer island, then in \hat{T} , put it in the *transposed* location (explained below). Otherwise, put it in the same place that it appears in T .

Step 3 requires further explanation. Suppose without loss of generality that the 1st long overhang of λ is its 1st arm overhang (as opposed to the 1st leg overhang), so that the 1st long overhang of $\hat{\lambda}$ is its 1st leg overhang. Then in the greedy arrangement of the first $i - 1$ border strips of T , an extreme end of the 1st, 2nd, 3rd, 4th, \dots border strips will be at the end of an arm, leg, arm, leg, \dots respectively, whereas in the greedy arrangement of the first $i - 1$ border strips of \hat{T} , an extreme end of the 1st, 2nd, 3rd, 4th, \dots border strips will be at the end of a leg, arm, leg, arm, \dots . Therefore if the i th border strip of T is slid by j from the end of an arm, then the i th border strip of \hat{T} will be slid by j from the end of a leg, and vice versa. The inner and outer islands in T and \hat{T} will therefore be identical except that they will be transposed—the $(i - 1)$ st and i th rows in T will become the $(i - 1)$ st and i th columns in \hat{T} , and vice versa. See Figure 13 for an example with $i = 2$, where the outer and inner islands are colored blue and violet respectively.

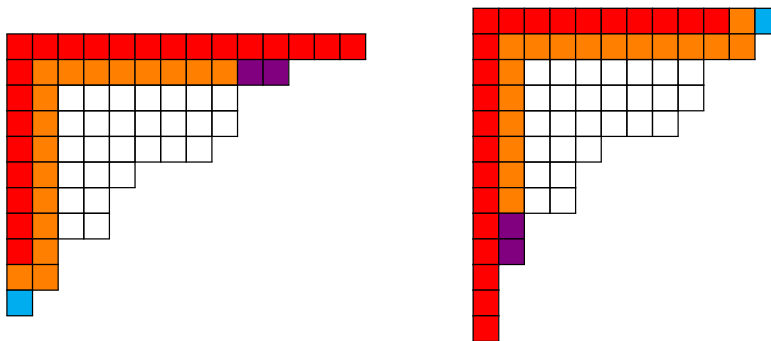


Figure 13: T and \hat{T} with $i = 2$ and $j = 1$

This is what we mean in Step 3 above by “transposed”; the border strips that ultimately cover the inner and outer islands in T should be used to cover the inner and outer islands in the transposed positions in \hat{T} . This procedure will never cause a violation of the weakly-increasing property of a BST, because the inner and outer islands will be covered by border strips beyond the i th.

Verifying that our putative bijection is indeed a bijection is straightforward and we will not belabor the details. \square

4.2.3 The Case $I(\lambda) \leq k$

Lemma 5. *Assume that $I := I(\lambda) \leq k$. For $1 \leq i < I$, let $\alpha_i = \alpha_i(\lambda)$ as defined by Equation (2). Let*

$$\alpha_i := \begin{cases} h_i - a_i + a_{i-1}, & \text{if } i = I; \\ h_i + a_{i-1} + 1, & \text{if } i = I + 1; \\ h_i, & \text{if } I + 1 < i \leq k + 1. \end{cases}$$

Then either there is exactly one greedy BST of shape λ and type α and no greedy BST of shape $\hat{\lambda}$ and type α , or there is exactly one greedy BST of shape $\hat{\lambda}$ and type α and no greedy BST of shape λ and type α .

Proof. Informally, the main idea is that, depending on whether the I th principal hook covers the I th long overhang or the I th short overhang (and this in turn depends on whether the 1st long overhang is the arm overhang or the leg overhang), either there will be a unique way to place all the border strips, or the I th border strip will disconnect the remainder of the shape in such a way that it cannot be completed to a full BST. See for example Figure 14, where $I = 4$.

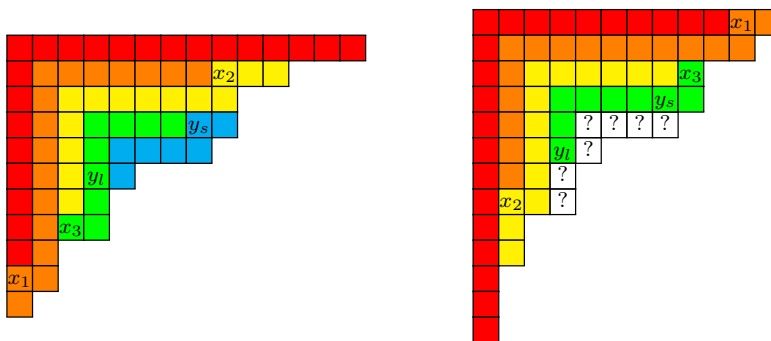


Figure 14: Greedy Arrangement

On the left, the 4th (green) border strip covers the 4th long overhang and there is a unique way to complete the BST, whereas on the right, the 4th

border strip disconnects the remainder of the shape, rendering it impossible to cover with the single remaining border strip of area 7.

More formally, in the Young diagrams of both λ and $\hat{\lambda}$, let y_s be the box in the I th principal hook that is adjacent to the I th short overhang, and let y_l be the box in the I th principal hook that is adjacent to the I th long overhang. In one of the two cases—the left-hand diagram in Figure 14—there will be a way to place the I th border strip so that it covers everything in the I th principal hook up to but not including y_s . In effect, this is a “greedy” placement of the I th border strip. Given this placement, there is a unique way to place the $(I + 1)$ st border strip, namely in such a way that it covers the I th short overhang, the box y_s , and the entire $(I + 1)$ st principal hook. All subsequent border strips are then uniquely forced to entirely cover the corresponding principal hooks. On the other hand, if we attempt to place the I th border strip “non-greedily” then it must cover y_s , but this will create an outer island that is disconnected from the $(I + 1)$ st principal hook, and there will be nowhere to place the $(I + 1)$ st border strip.

In the other case—the right-hand diagram in Figure 14—the I th border strip will cover both y_s and y_l , as well as $b_I - a_I - 1$ boxes of the I th long overhang, leaving $a_I + 1 > 0$ boxes in the I th long overhang uncovered. We are now in trouble with the $(I + 1)$ st border strip, because it must cover these $a_I + 1$ boxes as well as the entire $(I + 1)$ st principal hook, but this is impossible because these two nonempty regions are connected only via the box y_l , which is occupied by the I th border strip. So there are no BSTs in this case.

Note that it is in this final step of the argument that we use the assumption that $I(\lambda) \leq k$; this implies that the $(I + 1)$ st principal hook is nonempty, which we need for our contradiction. \square

4.2.4 The Case $I(\lambda) = k + 1$

As we observed in the proof of Lemma 4, if we compare the greedy arrangements of λ and $\hat{\lambda}$, then in one case an extreme end of the 1st, 2nd, 3rd, 4th, \dots border strips will be at the end of an arm, leg, arm, leg, \dots respectively, whereas in the other case an extreme end of the 1st, 2nd, 3rd, 4th, \dots border strips will be at the end of a leg, arm, leg, arm, \dots respectively. Since $I(\lambda) = k + 1$, this process will continue until all that is left uncovered is the final $(k + 1)$ st principal hook, the box x_k , and the adjacent k th overhang. In one case, x_k and the k th overhang will be attached to the arm of the $(k + 1)$ st

principal hook, and in the other case, they will be attached to the leg of the $(k + 1)$ st principal hook. Motivated by this situation, let us define an *arm extension* of a hook to be the shape obtained by attaching an additional row of boxes to the hook, with the leftmost box of the new row directly above the box at the top right corner of the hook. Let us also define a *leg extension* of a hook to be the shape obtained by attaching an additional column of boxes to the hook, with the topmost box of the new column directly to the left of the box at the bottom left corner of the hook. See Figure 15 for an example, where the original hook is white and the additional row or column has been colored red for visual clarity.



Figure 15: Arm and Leg Extensions of a Hook

Arm and leg extensions are defined even for *improper* hooks; i.e., hooks that consist of just a single row or a single column. See Figure 16 for an example of the arm and leg extensions of an improper hook.



Figure 16: Arm and Leg Extensions of an Improper Hook

The observations above show that comparing the number of greedy BSTs of type α and shapes λ and $\hat{\lambda}$ reduces to comparing the number of BSTs of type α^- of the arm and leg extensions¹ of the $(k + 1)$ st principal hook, where α^- consists of the parts α_i for $i > k$. Since $I(\lambda) = k + 1$, we may further assume that the $(k + 1)$ st principal hook is *not* self-conjugate.

Lemma 6. *Assume that $I(\lambda) = k + 1$. For $1 \leq i \leq k$, let $\alpha_i = \alpha_i(\lambda)$ as*

¹We have not formally defined a BST of a shape that is not a Young diagram, but the definition is the obvious generalization.

defined by Equation (2). Let

$$\alpha_i := \begin{cases} h_{k+1} - a_{k+1} + a_k, & \text{if } i = k + 1; \\ a_{k+1} + 1, & \text{if } i = k + 2. \end{cases}$$

Then the number of greedy BSTs of shape λ and type α has the opposite parity from the number of greedy BSTs of shape $\hat{\lambda}$ and type α .

Proof. Let H denote the $(k+1)$ st principal hook. Without loss of generality, we may assume that the length of the first row of H exceeds the height of its first column. For both the arm extension and the leg extension of H , we want to count the number of BSTs of type

$$(h_{k+1} - a_{k+1} + a_k, a_{k+1} + 1).$$

In each case, the answer will be at most 2, since the 2nd border strip must be placed at either the upper right end or the lower left end of the arm/leg extension; it may be less than 2 because one or both of these putative placements may violate the weakly increasing condition of a BST.

Note first that placing the 2nd border strip at the lower left end of the arm extension of H is illegal, because then it occupies precisely the entire 1st column of H ; in particular, the top box of the 1st column will contain a 2, and there will be boxes to its right that contain a 1. On the other hand, placing the 2nd border strip at the upper right end of the leg extension of H is legal; since $a_{k+1} + 1 < b_{k+1} + 1$, the 2nd border strip will occupy a proper subset of the first row of H . See Figure 17 for an example using the arm and leg extensions of Figure 15.



Figure 17: Illegal (left) and Legal (right) Tableaux

Now we split into two cases, depending on whether $a_k = a_{k+1}$. Suppose first that $a_k = a_{k+1}$. Then placing the 2nd border strip at the upper right end of the arm extension of H is illegal, because then it occupies precisely the boxes in the additional row; in particular, the leftmost additional box will

contain a 2, and there will be a box just below it containing a 1. Similarly, placing the 2nd border strip at the lower left end of the leg extension of H is illegal, because then it occupies precisely the boxes in the additional column; in particular, the topmost additional box will contain a 2, and there will be a box just to the right of it containing a 1. See Figure 18.



Figure 18: Illegal Tableaux

Combining this observation with the previous observation, we see that the arm extension admits 0 BSTs while the leg extension admits 1 BST, and 0 and 1 have opposite parity, as required.

Now suppose that $a_k \neq a_{k+1}$. Then placing the 2nd border strip at the upper right end of the arm extension will be legal, because it will occupy either a proper subset of the additional row, or it will occupy the entire additional row plus a proper subset of the first row of H (it cannot occupy the entire first row of H because the first row of H has $b_{k+1} + 1 > a_{k+1} + 1$ boxes, which is the area of the 2nd border strip, and the 2nd border strip must cover at least one box in the additional row), and either way, there is no violation of the weakly increasing property. Similarly, placing the 2nd border strip at the lower left end of the leg extension will be legal. So in this case, the arm extension admits 1 BST while the leg extension admits 2 BSTs, and 1 and 2 have opposite parity, as required. \square

4.2.5 The Case $I(\lambda) = \infty$

If $I(\lambda) = \infty$, then there is no second imbalance, and it is not hard to see that this implies that if the 1st principal hook is deleted from λ , then what remains is self-conjugate. In particular, $\lambda' = \hat{\lambda}$. We need the following standard fact.

Lemma 7. *For any λ and π ,*

$$\chi_\lambda(\pi) = (\text{sgn } \pi)\chi_{\lambda'}(\pi).$$

Proof. Textbooks typically prove this fact by noting that taking the conjugate of a partition corresponds to tensoring with the sign representation,

but as pointed out to us by Richard Stanley, it can be easily proved directly from the Murnaghan–Nakayama rule as follows. It is not hard to see that a border strip of even area must have either an odd number of rows and an even number of columns, or an even number of rows and an odd number of columns; similarly, a border strip with odd area must have either an odd number of rows and an odd number of columns, or an even number of rows and an even number of columns. It follows that if we transpose a BST of shape λ , then in the resulting BST of shape λ' , the signs of the border strips with even area will reverse while the signs of the border strips with odd area will remain the same. Now, in the disjoint-cycle decomposition of π , a cycle of odd length is an even permutation, and a cycle of even length is an odd permutation. Thus the overall sign of the BST will reverse if and only if π has an odd number of cycles of even length; i.e., if and only if π is an odd permutation. Since this argument applies for every BST, regardless of type, the lemma follows. \square

Lemma 7 implies that any odd permutation π such that $\chi_\lambda(\pi) \neq 0$ has the desired property that $\chi_\lambda(\pi) \neq \chi_{\hat{\lambda}}(\pi)$. From this point on, we focus on finding such a π .

Lemma 8. *Assume that $I(\lambda) = \infty$. For $1 \leq i \leq k$, let $\alpha_i = \alpha_i(\lambda)$ as defined by Equation (2), and let $\alpha_{k+1} = h_{k+1} + a_k + 1$. Then there is exactly one BST of shape λ and type α .*

Proof. There certainly exists a unique greedy BST of shape λ and type α , whose $(k + 1)$ st border strip covers all the boxes not covered by the first k border strips. The point is that there cannot be any non-greedy BSTs, because then the first k border strips would leave uncovered at least two disconnected components—an outer island, and some boxes in the $(k + 1)$ st principal hook—which therefore cannot both be covered by the single remaining border strip. \square

If the α described in Lemma 8 is the cycle type of an odd permutation, then we are done, so let us assume the contrary. Note that breaking any single cycle of an even permutation into two nonempty cycles yields an odd permutation. In particular, any composition β with exactly $k + 2$ nonzero parts where $\beta_i = \alpha_i(\lambda)$ (as defined by Equation (2)) for $i \leq k$, and whose last two parts β_{k+1} and β_{k+2} sum to $h_{k+1} + a_k + 1$, is the cycle type of an odd permutation. If we can show that the number of BSTs of type β is odd, then that will imply that $\chi_\lambda(\beta) \neq 0$, regardless of the signs of the BSTs.

The following lemma tells us that there are some strong constraints on what a non-greedy BST can look like.

Lemma 9. *Assume that $I(\lambda) = \infty$. For $1 \leq i \leq k$, let $\alpha_i = \alpha_i(\lambda)$ as defined by Equation (2), and assume that α has exactly $k+2$ nonzero parts. Let T be a non-greedy BST of shape λ and type α . Let $i \leq k$ be the smallest number such that the i th border strip is not positioned greedily. Then for $i < i' \leq k$,*

$$h_{i'} - a_{i'} + a_{i'-1} > a_i + 1 \quad (3)$$

and

$$a_{i'} = a_{i'-1}. \quad (4)$$

Moreover, $\alpha_{k+1} = h_{k+1}$ or $\alpha_{k+2} = h_{k+1}$.

Proof. The i th border strip of T is the j th slide for some $j > 0$; we use j for this number in the rest of this proof.

By Lemma 2, a border strip cannot contain more than one box on the principal diagonal. Here we have $k+2$ border strips and $k+1$ boxes on the principal diagonal, so there can be at most one *non-principal* border strip (i.e., a border strip that does not contain a box on the principal diagonal). In particular, there can be at most one border strip that lies entirely in an overhang.

It follows that the inner island of T must be empty, or else the inner and outer islands would both contain non-principal border strips. Thus the i th border strip must be slid as far as possible. i.e., $j = a_i + 1$.

For brevity, call the outer island S . The area of S is $j = a_i + 1$. Equations (3) and (4) are vacuously true if $i = k$, so assume that $i < k$. We prove the two equations jointly, by induction on i' . The base case is $i' = i + 1$. By Lemma 1,

$$h_{i+1} - a_{i+1} \geq 3 > 1,$$

so $h_{i+1} - a_{i+1} + a_i > a_i + 1$, proving Equation (3) for $i' = i + 1$. But what Equation (3) says is that the $(i+1)$ st border strip is too large to fit inside S ; therefore, it must lie inside the $(i+1)$ st principal hook. In fact, it must entirely fill the $(i+1)$ st principal hook; the reason is that the area of the $(i+1)$ st border strip is

$$h_{i+1} - a_{i+1} + a_i \geq h_{i+1} - a_{i+1},$$

so any boxes in the $(i+1)$ st principal hook not covered by the $(i+1)$ st border strip must lie in an $(i+1)$ st overhang, creating another region that must be

covered by a non-principal border strip. Therefore $h_{i+1} - a_{i+1} + a_i = h_{i+1}$, proving $a_{i'} = a_{i'-1}$ for $i' = i + 1$.

The proof of the induction step is similar to the proof of the base case. Lemma 1 implies $h_{i'} - a_{i'} > 1$, and by induction we may assume that $a_{i'-1} = a_i$, so this proves Equation (3). The i' th border strip cannot fit inside S and hence must lie inside the i' th principal hook. Furthermore, by the same reasoning as we gave above, it must entirely fill the i' th principal hook, proving that $a_{i'} = a_{i'-1}$.

We have now almost completely specified the structure of T . The border strips before the i th are placed greedily; the i th border strip is slid as far as possible; the remaining border strips up to the k th border strip each entirely occupy the respective principal hook. The only boxes not covered by the first k border strips are in two disconnected components, namely S and the $(k + 1)$ st principal hook. Therefore α_{k+1} and α_{k+2} must specify the areas of these two regions, and in particular one of them must equal h_{k+1} . \square

It turns out that for most shapes λ with $I(\lambda) = \infty$, we can find an α for which there exists a unique greedy BST of shape λ and type α and no non-greedy BST of shape λ and type α ; see Lemma 10. However, there are a few exceptions when the innermost hooks are very small; these are dealt with in Lemma 11.

Lemma 10. *Assume that $I(\lambda) = \infty$. For $1 \leq i \leq k$, let $\alpha_i = \alpha_i(\lambda)$ as defined by Equation (2). Let*

$$\alpha_{k+2} := \begin{cases} 2, & \text{if } a_k \geq 2 \text{ and } a_{k+1} = 0; \\ a_{k+1} + 1, & \text{if } a_k \neq a_{k+1} \text{ and } a_{k+1} \geq 1; \\ a_{k+1} + 2, & \text{if } a_k = a_{k+1} \geq 2. \end{cases}$$

Let $\alpha_{k+1} := h_{k+1} + a_k + 1 - \alpha_{k+2}$. Then there is exactly 1 BST of shape λ and type α , and this BST is greedy.

Proof. First, it is easy to verify that neither α_{k+1} nor α_{k+2} is equal to h_{k+1} , so by Lemma 9, there can be no non-greedy BST of shape λ and type α . So we just need to verify that there is exactly one greedy BST of shape λ and type α . To do this, it suffices to consider the residual shape left uncovered by the greedy arrangement, and show that there is a unique BST of this residual shape of type $(\alpha_{k+1}, \alpha_{k+2})$.

In the case that $a_k \geq 2$ and $a_{k+1} = 0$, the residual shape is an arm extension (or a leg extension, but we may assume without loss of generality that it is an arm extension; this is also true for the remaining cases, so we will assume arm extensions in the rest of this proof without further comment) of a single box, where the additional row has length $a_k + 1 \geq 3$. In this case, $\alpha_{k+2} = 2$, so putting the 2nd border strip at the lower left end is illegal, while putting the 2nd border strip at the upper right end is legal (since the first row has length at least 3). See Figure 19.

2	1	1
2		

1	2	2
1		

Figure 19: Illegal (left) and Legal (right) Tableaux

In the case that $a_k \neq a_{k+1}$ and $a_{k+1} \geq 1$, the residual shape is an arm extension of a proper self-conjugate hook H , where the number of boxes in the additional row is nonzero and is different from the number of boxes in the first row (equivalently, the first column) of H . In this case, $\alpha_{k+2} = a_{k+1} + 1$, which is exactly the number of boxes in the first column of H . Therefore putting the 2nd border strip at the lower left end is illegal, since it would fill up the entire first column of H , resulting in a 2 to the left of a 1 in the first row of H . On the other hand, putting the 2nd border strip at the upper right end is legal, because the 2nd border strip will either be a proper subset of the additional row, or it will contain the entire additional row as well as a nonzero number of boxes at the right end of the first row of H (but not all of the first row of H), and either way, the danger of having a 2 above a 1 will be averted. See Figures 20 and 21 for examples in which $a_k < a_{k+1}$ and $a_{k+1} < a_k$ respectively.

			1	1	1
2	1	1	1		
2					
2					

				2	2	2
1	1	1	2			
1						
1						

Figure 20: Illegal (left) and Legal (right) Tableaux

There remains the case that $a_k = a_{k+1} \geq 2$, which is similar to the preceding case except that now the additional row is exactly the same length

			1	1	1	1
2	1	1				
2						
2						

			1	2	2	2
1	1	1				
1						
1						

Figure 21: Illegal (left) and Legal (right) Tableaux

as the first row of H . In this case, $\alpha_{k+1} = a_{k+1} + 2$, which is one more than the height of the first column of H . Putting the 2nd border strip at the lower left end fills up the first column of H as well as an extra box in the first row of H ; since the first row of H has length at least 3, there will be a 1 in the first row of H to the right of a 2, which is illegal. On the other hand, putting the 2nd border strip at the upper right end fills up the additional row as well as the rightmost box in the first row of H , and this is legal. See Figure 22.

			1	1	1
2	2	1			
2					
2					

			2	2	2
1	1	2			
1					
1					

Figure 22: Illegal (left) and Legal (right) Tableaux

This completes the proof. □

Finally, we deal with the exceptional cases not covered by Lemma 10.

Lemma 11. *Assume that $I(\lambda) = \infty$. For $1 \leq i \leq k$, let $\alpha_i = \alpha_i(\lambda)$ as defined by Equation (2), and let $\alpha_{k+1} = h_{k+1}$ and $\alpha_{k+2} = a_k + 1$. Let G be the number of greedy BST of shape λ and type α and let N be the number of non-greedy BST of shape λ and type α .*

1. *If $a_k = 0$ and $a_{k+1} = 0$ then $G = 1$ and $N = 2$.*
2. *If $a_k = 1$ and $a_{k+1} = 0$ then $G = 0$ and $N = 1$.*
3. *If $a_k = 1$ and $a_{k+1} = 1$ then $G = 0$ and $N = 1$.*

Proof. Let us consider Case 1 first. The greedy arrangement leaves just two adjacent boxes uncovered, and there is obviously only one way to cover them with the (singleton) $(k + 1)$ st and $(k + 2)$ nd border strips. See Figure 23 for an example.

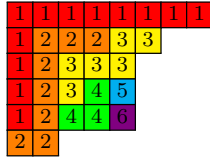


Figure 23: Case 1: $h_{k+1} = 1$ and $a_k = 0$

Now consider the non-greedy BSTs. Let i' be the smallest integer with the property that $a_{i'} = a_l = 0$ for all $i' \leq l \leq k$, and let the i th border strip be the first border strip that is not positioned greedily. We claim that $i = i'$. Note first that Lemma 9 implies that $i \geq i'$. If $i > i'$, then the outer island consists of the singleton box x_{i-1} , which cannot be covered by the $(k+1)$ st or $(k+2)$ nd border strip since it would violate the weakly increasing property. See Figure 24 for an example with $4 = i > i' = 3$.

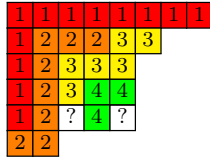


Figure 24: Case 1: Illegal Non-Greedy BST

Conversely, if $i = i'$, then since $a_i = 0$, the only possible slide of the i th border strip is the 1st slide, so the outer island consists of a singleton box at the extreme end of an $(i-1)$ st overhang. This box and the (singleton) $(k+1)$ st hook can be covered by the $(k+1)$ st and $(k+2)$ nd border strip in either order, while the intervening border strips are uniquely forced to entirely cover their corresponding principal hooks. We obtain a total of 2 non-greedy BSTs. See Figure 25 for an example.



Figure 25: Case 1: Two Non-Greedy BSTs

Now consider Case 2. The greedy arrangement leaves uncovered a set of

three boxes (two in one row and one in another), and there is no legal way to cover this with a $(k + 1)$ st border strip of area 1 and a $(k + 2)$ nd border strip of area 2. See Figure 26 for an example.

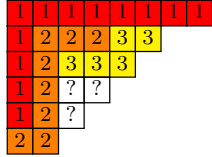


Figure 26: Case 2: Illegal Greedy BST

Now consider the non-greedy BSTs. Let i' be the smallest integer with the property that $a_{i'} = a_l = 1$ for all $i' \leq l \leq k$, and let the i th border strip be the first border strip that is not positioned greedily. The i th border strip must be slid as far as possible, since otherwise we would have both an outer island and an inner island, creating more disconnected components than we have border strips to cover them. We claim that $i = i'$. The argument is similar to the argument in Case 1. Lemma 9 implies that $i \geq i'$. If $i > i'$, then because $a_i = 1$, the outer island consists of two boxes, one of which is x_{i-1} , which cannot be covered by the $(k + 1)$ st or $(k + 2)$ nd border strip without violating the weakly increasing property. See Figure 27 for an example with $3 = i > i' = 2$.

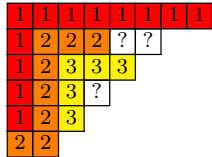


Figure 27: Case 2: Illegal Non-Greedy BST

Conversely, if $i = i'$ then there exists a unique non-greedy BST, because the outer island has area 2 and the innermost principal hook has area 1, so the outer island must be covered by the $(k + 2)$ nd border strip and the innermost principal hook must be covered by the $(k + 1)$ st border strip. It is readily checked that covering the outer island with the $(k + 2)$ nd border strip does not violate the weakly increasing property, regardless of whether $a_{i-1} = 0$ or $a_{i-1} > 1$; see Figure 28 for diagrams of these two subcases.

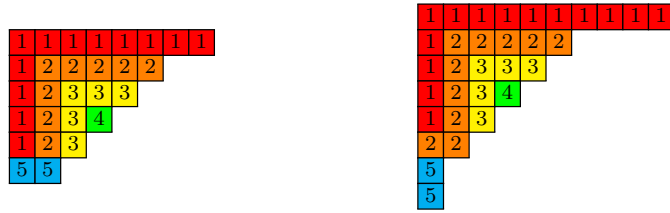


Figure 28: Case 2: $h_{k+1} = 1$ and $a_k = 1$

Finally, there is Case 3, which is very similar to Case 2, so we will just sketch the argument. The greedy arrangement leaves 5 boxes uncovered, which cannot be covered by a $(k + 1)$ st border strip of area 3 and a $(k + 2)$ nd border strip of area 2; see Figure 29.

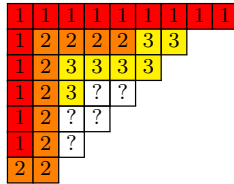


Figure 29: Case 3: Illegal Greedy BST

For the non-greedy BSTs, we define i and i' as in Case 2, and argue as before that we must have $i = i'$; see Figure 30.

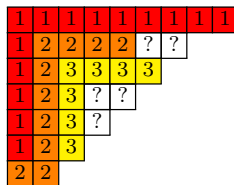


Figure 30: Case 3: Illegal Non-Greedy BST

If $i = i'$ then as in Case 2 we argue that there is a unique non-greedy BST, with the $(k + 1)$ st border strip covering the innermost principal hook and the $(k + 2)$ nd border strip covering the outer island; see Figure 31.

In all cases, we have found an odd number of BSTs, so the character value cannot be zero. \square

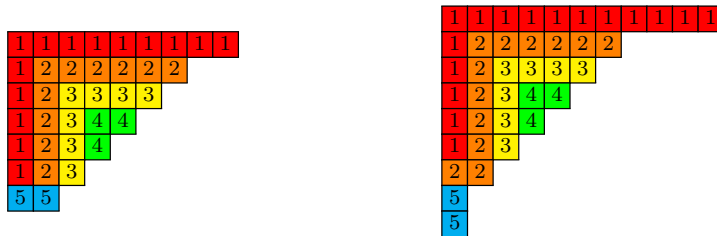


Figure 31: Case 3: $h_{k+1} = 3$ and $a_k = 1$

5 Concluding Remarks

A crude upper bound on the number of queries needed for our algorithm is $O(n^{3/2})$. In the forward pass, there are at most $n^{1/2}$ principal hooks, and the size of each principal hook can be determined in at most n queries. In the backward pass, there are again at most $n^{1/2}$ principal hooks, and for each principal hook, determining a_i and b_i requires at most n queries, and distinguishing doppelgänger requires a constant number of queries. We expect that a more careful analysis, which we have not carried out, will show that the required number of queries is (approximately) linear in n , because if there are a lot of principal hooks then the number of queries per principal hook will be reduced.

It is natural to ask if a more efficient algorithm can be found. As we mentioned in the Introduction, empirically it seems that permutations that consist mostly of fixed points are good distinguishers. Enumerating the corresponding BSTs naturally leads to enumerating skew tableaux, for which there exist formulae such as the Naruse hook-length formula (see [2] for a readable description and proof of the Naruse hook-length formula). Although it does not seem easy to prove that various alternating sums of hook-length formulae cannot coincide in value, perhaps it can be done. If so, the number of queries needed could conceivably be drastically reduced.

References

- [1] Michael Larsen and Aner Shalev. Characters of symmetric groups: sharp bounds and applications. *Inventiones Mathematicae*, 174(3):645–687, 2008.

- [2] Alejandro H. Morales, Igor Pak, and Greta Panova. Hook formulas for skew shapes I. q -analogues and bijections. *Journal of Combinatorial Theory Series A*, 154:350–405, February 2018.
- [3] Bruce E. Sagan. *The Symmetric Group: Representations, Combinatorial Algorithm, and Symmetric Functions*. Graduate Texts in Mathematics. Springer, second edition, 2001.
- [4] Richard P. Stanley. *Enumerative Combinatorics*, volume 2. Cambridge University Press, 1999.