Séminaire Lotharingien de Combinatoire **89B** (2023)
Article #89, 12 pp.

Proceedings of the 35th Conference on Formal Power Series and Algebraic Combinatorics (Davis)

Interpolating between classic and bumpless pipe dreams

Allen Knutson*[1] and Gabe Udell†[1]

[1]*Department of Mathematics, Cornell University*

Abstract. We introduce *hybrid pipe dreams,* in which some rows are specified to contain classic pipe dream tiles (plus a couple new ones) and some to contain bumpless pipe dream tiles (albeit upside down). For each of the 2^n possibilities of row specification, there is a sum over hybrid pipe dreams to compute each double Schubert polynomial. We give a bijection to swap the specifications on two adjacent rows, demonstrating the equivalence of the many formulæ.

Keywords: Schubert polynomials, pipe dreams

1 Hybrid pipe dreams compute Schubert polynomials

1.1 Hybrid pipe dreams and the formula

Define the **classic pipe dream tiles** to be

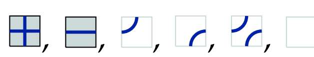 but not ,

the **(upside down) bumpless pipe dream tiles** to be

 but not ,

and in each case call the shaded tile(s) **weighty**.

Given $\pi \in S_n$ a permutation, and $\tau \in \{C, B\}^n$ a word in **C**lassic and **B**umpless, define a **τ-hybrid pipe dream for** π to be an $n \times n$ square made of tiles, such that

- if two tiles share an edge, they must agree as to whether it is blank or not. Hence the pipes end only on the boundary of the square
- the ith row is called "a classic row" or "a bumpless row" depending on τ_i, and uses only tiles from one set or the other
- each classic row has a number on the left edge, and is blank on the right edge, with the opposite being true for bumpless rows; this number is the **row label**

*allenk@math.cornell.edu. AK was partially supported by NSF Grant DMS–1953948.
†gru5@cornell.edu.

- the top edges are labeled with $\pi^{-1}(1), \pi^{-1}(2), \ldots, \pi^{-1}(n)$ from left to right
- the bottom edges are blank
- the labeled boundary edges on the left and right are labeled $1 \ldots n$ in order, counterclockwise starting from the Northwest corner
- the numbers at the two ends of a pipe agree, and no two pipes cross twice.

Some examples are below. Our first main theorem is the following:

Theorem 1. *For $\pi \in S_n$ a permutation, and any word $\tau \in \{C, B\}^n$, the double Schubert polynomial $S_\pi(x_1, \ldots, x_n, y_1, \ldots, y_n)$ can be computed as a sum over the τ-hybrid pipe dreams δ for π:*

$$S_\pi = \sum_\delta \prod \{x_i - y_j \ : \ i, j \in [n] \text{ and the tile in } \textbf{physical} \text{ column } j \text{ and row } \textbf{labeled } i \text{ is weighty}\}$$

In the case $\tau = C^n$, this reproduces the pipe dream formula in [3, Prop. 6.2] (and [1] in the $y_i \equiv 0$ setting). In the case $\tau = B^n$, this reproduces the bumpless pipe dream formula in [7] (albeit upside down) and is equivalent to the alternating sign matrix formula from [9, 8]. The pipes are labeled *differently here (by π^{-1}) than they are in [1, 3, 6, 7], but the placements of tiles are exactly the same.*

The number of terms in this sum is $S_\pi|_{x_i \equiv 1, y_j \equiv 0}$, independent of τ. Our other main theorem (in Section 2), which implies this one, will be a bijective correspondence.

Here are the hybrid pipe dreams for $\pi = 132$ and various choices of τ.

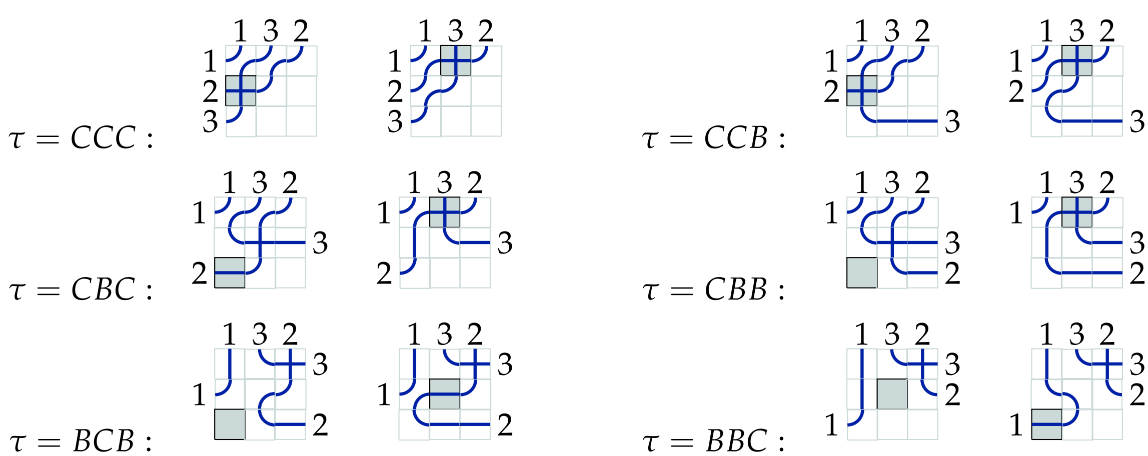

1.2 The all-classic case $\tau = C^n$

The boundary conditions mean that every edge on North and West is crossed by a pipe, whereas every edge on South and East is blank.

Proposition 1. *A τ-hybrid pipe dream for $\tau = C^n$ has*

- *only* `j`*-tiles on the antidiagonal, and only blank tiles below the antidiagonal;*
- *only crosses and elbows above the antidiagonal.*

In particular the and `r`*-tile , despite being allowed, are never used.*

Proof. Consider the rightmost column, from top to bottom. The Northeast tile has a number on its North edge and blank on its East edge, and cannot be a -tile (as those are forbidden in classic rows), so must be a -tile. The next tile below it is blank on North and East, so must be , and the same argument holds all the way down the rightmost column.

Now consider the first row, from right to left, starting just West of the Northeast tile. The next tile has pipes crossing its North and East side, so must be a or an . The same argument holds all the way to the left in the first row.

Remove the last column and first row and do induction on the remaining $(n-1) \times (n-1)$ square. □

This picture is slightly different *mathematically* from the classic pipe dreams in, say, [1] which fill the fourth quadrant and have only crossing and elbow tiles. That formulation is based on the forward-stability of Schubert polynomials, $S_\pi = S_{\pi \oplus I_m}$ (where the $\oplus$ is on the permutation matrices). However, even in [1] they omit "the 'sea' of wavy strands" and obtain as pictures *exactly* the C^n-hybrid pipe dreams here.

1.3 The all-bumpless case $\tau = B^n$

There is less to say here – our definition of B^n-hybrid pipe dreams exactly matches that of [7], except for being upside down and for the shift in labeling (they have $\pi(1)\ldots\pi(n)$ along the East side, where we have $n \ldots 1$).

1.4 Repeating letters

We inverted the labeling principally to simplify the description of the labels on the boundary (as $1 \ldots n$ counterclockwise). There is another benefit (though not one we will pursue here); in the case that π has descents in at most the positions $k_1, k_1 + k_2, \ldots, k_1 + k_2 + k_3 + \ldots + k_d$ it is natural to write π^{-1} as a **string** with content $1^{k_1} 2^{k_2} \cdots d^{k_d} (d+1)^{n - \sum k_i}$. Put another way, π is the shortest representative inside a coset $S_n / \prod_i S_{k_i}$ (where $k_{d+1} := n - \sum_{i=1}^{d} k_i$), which is an equivalence class within which the *positions* are ambiguous, and its inverse has ambiguous *values*. If one allows in this way multiple pipes to have the same label, one must forbid those pipes from crossing even *once*.

Placing the ambiguity on values is more natural in these scattering-diagram-inspired contexts, in that physics is comfortable with indistinguishable particles but not with ambiguous positions.

2 Equivalence of the 2^n formulæ

In this section we give combinatorial proofs of the equivalence of the 2^n formulæ. For this purpose, we introduce **decorated hybrid pipe dreams**, as hybrid pipe dreams with decorations of x or $-y$ on the weighty tiles. Following the conventions in [5] (which introduced this notion in the bumpless case), we write the set of decorated τ-hybrid pipe dreams for π as

$$\widetilde{\mathrm{HPD}}(\tau,\pi) := \{(D,f) \mid D \in \mathrm{HPD}(\tau,\pi),\ f : \mathrm{weighty}(D) \to \{x,-y\}\}$$

where $\mathrm{HPD}(\tau,\pi)$ is the set of τ-hybrid pipe dreams for π, and $\mathrm{weighty}(D)$ is the set of weighty tiles in D. Define the **monomial of a decorated hybrid pipe dream** to be

$$\mathrm{mon}(D,f) := \prod_{\substack{(i,j)\in\mathrm{weighty}(D)\\ f(i,j)=x}} x_i \prod_{\substack{(i,j)\in\mathrm{weighty}(D)\\ f(i,j)=-y}} (-y_j), \qquad \text{so } S_\pi = \sum_{(D,f)\in\widetilde{\mathrm{HPD}}(\tau,\pi)} \mathrm{mon}(D,f)$$

as follows from Theorem 1.

We prove the equivalence of these 2^n formulæ by constructing a bijection

$$\phi : \widetilde{\mathrm{HPD}}(\tau_1,\pi) \to \widetilde{HPD}(\tau_2,\pi)$$

for any $\tau_1,\tau_2 \in \{C,B\}^n$ while making sure that $\mathrm{mon}(\phi(D,f)) = \phi(\mathrm{mon}(D,f))$ for all $(D,f) \in \widetilde{\mathrm{HPD}}(\tau_1,\pi)$. We will do this in two steps (each then iterated many times). First, we handle the case where τ,τ' differ only in the last letter. That is, we describe how to change the last row from classic to bumpless, and in fact, we do this in such a way that $\phi(D,f)$ differs from (D,f) only in its final row. Second, we handle the case where τ' is the same as τ but with some (consecutive) substring CB replaced by BC. That is, we describe how to switch the order of two adjacent rows, one bumpless and one classic. We do this in such a way that $\phi(D,f)$ differs from (D,f) only in the rows that were swapped. The composite bijections ϕ are automatically coherent in the sense that any triangle $\widetilde{\mathrm{HPD}}(\tau_1,\pi) \to \widetilde{\mathrm{HPD}}(\tau_2,\pi) \to \widetilde{\mathrm{HPD}}(\tau_3,\pi)$ commutes.

2.1 Monk's rule and the Gao–Huang bijection

Hybrid pipe dreams were created to help understand the connection between pipe dreams and bumpless pipe dreams, but they are not the first tool for doing so. Another approach to this connection involves Monk's rule.

Theorem 2 (Monk's rule for double Schubert polynomials)**.** *For π a permutation, α a positive integer, $t_{a,b}$ the transposition of a and b, and $\gtrdot$ the Bruhat covering relation,*

$$(x_\alpha - y_{\pi(\alpha)})S_\pi + \sum_{\substack{s<\alpha \\ \pi t_{s,\alpha} \gtrdot \pi}} S_{\pi t_{s,\alpha}} = \sum_{\substack{l>\alpha \\ \pi t_{\alpha,l} \gtrdot \pi}} S_{\pi t_{\alpha,l}}.$$

Monk's rule for single Schubert polynomials can be proven bijectively using either pipe dreams [1, 2, 5] or bumpless pipe dreams [5] by using, for either $\tau = C^n$ or $\tau = B^n$, maps $x_\alpha \rightsquigarrow$: $\mathrm{HPD}(\tau, \pi) \to \bigcup_{\substack{\alpha<l \\ \pi t_{\alpha,l} \gtrdot \pi}} \mathrm{HPD}(\tau, \pi t_{\alpha,l})$ and $m_{s,\beta} : \mathrm{HPD}(\tau, \pi t_{s,\beta}) \to \bigcup_{\substack{l>\beta \\ \pi t_{\beta,l} \gtrdot \pi}} \mathrm{HPD}(\tau, \pi t_{\beta,l})$ when $s < \beta, \pi t_{s,\beta} \gtrdot \pi$. In [4], Gao and Huang constructed a bijection ϕ_{GH} between bumpless and classic pipe dreams preserving $\mathrm{mon}(D, f)$ when all decorations are x. Furthermore, they proved their bijection to be canonical in the sense that it commutes with $x_\alpha \rightsquigarrow$ and $m_{s,\beta}$. When $\tau_1 = B^n$, $\tau_2 = C^n$ and either all decorations being x or all decorations being $-y$, we get specializations ϕ_x, ϕ_y of the hybrid pipe dream bijection. We conjecture that $\phi_{GH}(D) = \phi_x(D) = (\phi_y(D^\top))^\top$ for any bumpless pipe dream D, where $D^\top$ denotes the transpose of a pipe dream or bumpless pipe dream.

2.2 Changing the bottom row

A single picture shows the full gamut of possibilities for the bottom rows, in the C and B cases:

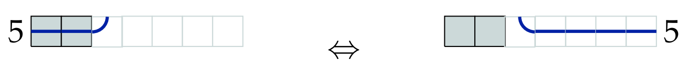

Note that this correspondence preserves the positions of the weighty tiles, and hence the associated monomials (assuming we keep the same decoration).
We give an in-context example of the correspondence:

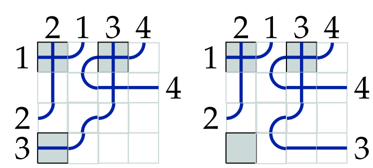

2.3 Swapping a classic row with an adjacent bumpless row

Define a $\substack{C \\ B}$ **section** to be a tiled $2 \times k$ rectangle that could appear in rows i and $i+1$ (and k contiguous columns) of some τ-hybrid pipe dream, where $\tau_i = C$ and $\tau_{i+1} = B$. Define $\substack{B \\ C}$ **sections** similarly. We say the pipe connecting i on the top of a hybrid pipe dream to i on a side is **labeled** i or is **the** i **pipe**. Where a pipe enters/exits an $\substack{C \\ B}$ section ($\substack{B \\ C}$ section) we will place a number indicating the label of the pipe. These numbers are just to distinguish which pipes connect to each other outside the $\substack{C \\ B}$ section ($\substack{B \\ C}$ section). That is, if we take an $\substack{C \\ B}$ section ($\substack{B \\ C}$ section) and apply a permutation to each number on its boundary, we will not distinguish this from the original.

Theorem 3. *Assuming all the decorations are $-y$s, there exists* uniquely *a monomial-preserving bijection* $\phi_y : \widetilde{HPD}(\tau_1 CB\tau_2, \pi) \to \widetilde{HPD}(\tau_1 BC\tau_2, \pi)$ *changing only those two rows, using the replacements below:*

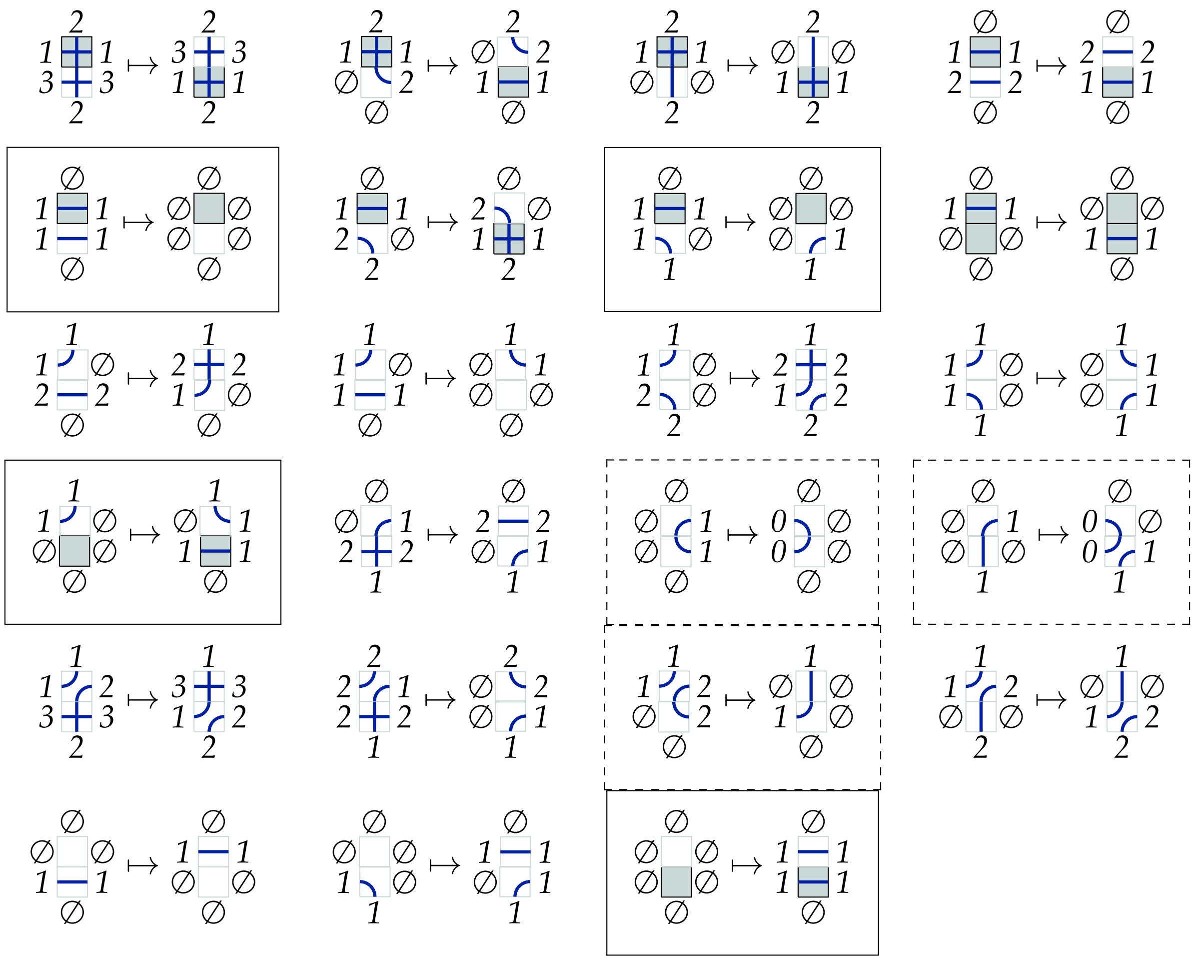

The proof is a not-too-laborious case check, determining all the ways to glue a classic tile atop a bumpless tile (or vice versa), then to specify the possible connectivities; finally to match the 23 of each, preserving the net connectivity (up to potential pipes entering and leaving on the same side!) and the number of weighty tiles. It takes some additional analysis to confirm that the replacement tiles glue together. Uniqueness follows (slightly non-trivially) from the rest of the lemmas and theorems of this section and the observation that each of the 23 pairs of tiles has a distinct connectivity (i.e. no two have the same labels around the outside). The four boxed replacements are the ones that would not preserve the monomial if the weighty tiles were labeled x rather than $-y$. We spend the rest of the abstract dealing with them (and bringing the dash-boxed replacements along for the ride).

When a $\genfrac{}{}{0pt}{}{C}{B}$ section or $\genfrac{}{}{0pt}{}{B}{C}$ section is of (full) width n, call it **full**. The major property that full $\genfrac{}{}{0pt}{}{C}{B}$ sections ($\genfrac{}{}{0pt}{}{B}{C}$ sections) have, that non-full ones need not, is that a full $\genfrac{}{}{0pt}{}{C}{B}$ section ($\genfrac{}{}{0pt}{}{B}{C}$ section) has blank labels on the right edge of its top row and the left edge of its bottom row (left edge of its top row and right edge of its bottom row) and blank labels on its other vertical edges, whereas an arbitrary $\genfrac{}{}{0pt}{}{C}{B}$ section ($\genfrac{}{}{0pt}{}{B}{C}$ section) may have any combination of blank and non-blank labels on its vertical edges.

Two $\genfrac{}{}{0pt}{}{C}{B}$ sections ($\genfrac{}{}{0pt}{}{B}{C}$ sections) are said to be the same if they have the same tiles and the same connectivity. As in Theorem 3, we put labels on $\genfrac{}{}{0pt}{}{C}{B}$ sections ($\genfrac{}{}{0pt}{}{B}{C}$ sections) in order to show their connectivity and not to indicate the labels the pipes may have had within some larger hybrid pipe dream.

We attach decorations f to $\genfrac{}{}{0pt}{}{C}{B}$ sections ($\genfrac{}{}{0pt}{}{B}{C}$ sections) in the same way that we did with hybrid pipe dreams. To assign weights to $\genfrac{}{}{0pt}{}{C}{B}$ sections ($\genfrac{}{}{0pt}{}{B}{C}$ sections) in the context of a complete hybrid pipe dream we need to know the labels on the two rows. However, for the purpose of verifying that our bijection preserves weights, we will say the classic row is labeled c and the bumpless row is labeled b so that we attach the weight x_c to weighty tiles decorated with x in the classic row and x_b to weighty tiles decorated with x in the bump less row.

Although the bijection in Theorem 3 is not weight-preserving for arbitrary decorated $\genfrac{}{}{0pt}{}{C}{B}$ section, most of the replacement rules in the bijection described in Theorem 3 will still be a part of the more general bijection we present in Theorem 5. The leftover columns will be the ones appearing in the seven $\boxed{\text{boxed}}$ and dash-boxed rules in Theorem 3. To define the row swapping bijection for arbitrary decorated $\genfrac{}{}{0pt}{}{C}{B}$ sections we will (1) organize these leftover tiles into chunks which are described in Definition 1, (2) define the bijection on these chunks in Theorem 4, and finally (3) show that everything outside those chunks can be dealt with using the local replacement rules from Theorem 3.

Definition 1. *We say that a* $\genfrac{}{}{0pt}{}{C}{B}$ *section is* **obstructive** *if it can be made with the following non-deterministic finite state machine (on the left) by starting at one of the leftmost 5 nodes and ending at one of the rightmost 5 modes. In the finite state machine j and k mean to repeat that node j or k times respectively where j and k may be any non-negative integer including 0.*

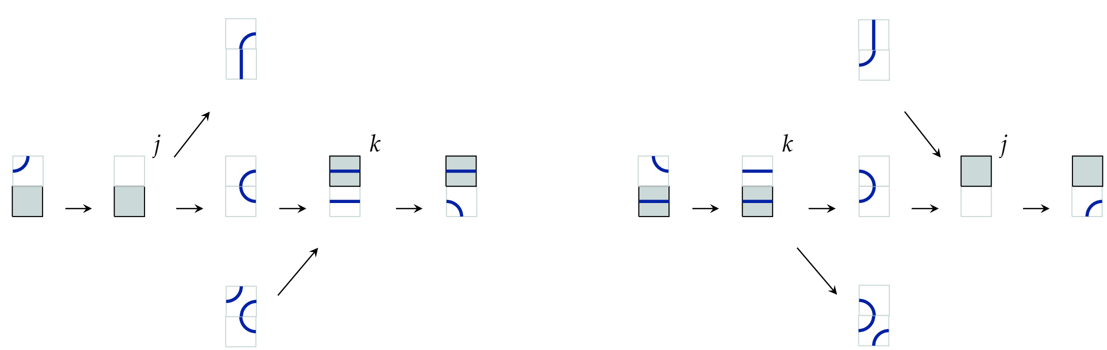

The state machine diagram defining **obstructive** ${}^{B}_{C}$ *sections is the one on the right.*

Theorem 4. *There is a unique weight preserving bijection between decorated obstructive* ${}^{C}_{B}$ *sections and decorated obstructive* ${}^{B}_{C}$ *sections such that the labels along the top and bottom edges are preserved.*

Proof. Every obstructive ${}^{C}_{B}$ section (${}^{B}_{C}$ section) begins with a possibly empty sequence of columns with weighty tiles in the bumpless (classic) row and non-weighty tiles in the classic (bumpless) row. After this, there will be a single column with no weighty tiles. Finally, the ${}^{C}_{B}$ section (${}^{B}_{C}$ section) ends with a possibly empty sequence of columns with weighty tiles in the classic (bumpless) row and non-weighty tiles in the bumpless (classic) row. An obstructive ${}^{C}_{B}$ section (${}^{B}_{C}$ section) is completely determined by the labels on its top and bottom edges together with the position of its single non-weighty column. Further, for each possible boundary of an obstructive ${}^{C}_{B}$ section (${}^{B}_{C}$ section) and each i between 1 and the width of the ${}^{C}_{B}$ section (${}^{B}_{C}$ section), there exists an obstructive ${}^{C}_{B}$ section (${}^{B}_{C}$ section) with those boundary labels that has a non-weighty ith column.

Let D be an obstructive ${}^{C}_{B}$ section (${}^{B}_{C}$ section) with width k and let f be some decoration on it. Then, $\mathrm{mon}(D,f) = x_c^a x_b^{k-1-a-m} \prod_{j=1}^{m} -y_{i_j}$ for some $0 \leq m \leq k-1$, some $0 \leq a \leq k-1-m$,and some sequence $i_1, i_2, \ldots, i_m \in \{1, 2, \ldots, k\}$. We claim that mon restricts to a bijection from ${}^{C}_{B}$ sections (${}^{B}_{C}$ sections) with fixed labels along their top and bottom edges to monomials of the form in the previous sentence. Once proven, this claim implies that there is a unique weight preserving bijection between decorated obstructive ${}^{C}_{B}$ sections and decorated obstructive ${}^{B}_{C}$ sections. Suppose we have some $i_1, \ldots, i_m \in [k]$ with $0 \leq m \leq k-1$, and some $0 \leq a \leq k-1-m$. Given labels along the top and bottom edges we want to show there is a unique obstructive decorated ${}^{C}_{B}$ section (D, f) and a unique obstructive decorated ${}^{B}_{C}$ section (D', f') with those labels and with $\mathrm{mon}(D,f) = \mathrm{mon}(D',f') = x_c^a x_b^{k-1-a-m} \prod_{j=1}^{m} -y_{i_j}$. To achieve that monomial, there needs to be a weighty tile in both D and D' decorated with $(-y)$ in the (i_j)th column for each $j = 1, \ldots, m$. The a factors of x_c in the monomial must come from the rightmost (leftmost) a columns of D (D') which are not already reserved for giving $-y_{i_j}$ weights. The $k-1-a-m$ factors of x_b must come from the leftmost (rightmost) $k-1-a-m$ columns of D (D') which are not already reserved for giving $-y_{i_j}$ weights. We have thus determined the weights we'll need to have coming from $m + a + (k-1-a-m) = k-1$ of the k columns of D (D'). Hence, the unique remaining column of D (D') will be the unique non-weighty column of D (D′). Setting a column as the non-weighty column uniquely determines what D and D' should each be based on the labels along the top and bottom edges. In the process of determining where the non-weighty column of D and D' had to be to achieve the monomial $x_c^a x_b^{k-1-a-m} \prod_{j=1}^{m} -y_{i_j}$, we specified which columns would contribute which weights. This tells us exactly how to decorate D and D' because each column of an obstructive ${}^{C}_{B}$ section (${}^{B}_{C}$ section) has at most one weighty tile. We have thus shown that mon gives a weight preserving bijection between obstructive decorated $\substack{C\\B}$ sections ($\substack{B\\C}$ sections) with set labels on the top and bottom edges, and monomial weights of the form $x_c^a x_b^{k-1-a-m} \prod_{j=1}^m -y_{i_j}$ for appropriate a, m, i_j. We conclude that for any decorated $\substack{C\\B}$ section (D, f), there is a unique decorated $\substack{B\\C}$ section (D', f') such that the the top and bottom edge labels of D match those of D' and $\mathrm{mon}(D, f) = \mathrm{mon}(D', f')$. □

Columns of a full $\substack{C\\B}$ section ($\substack{B\\C}$ section) which aren't contained in some obstructive $\substack{C\\B}$ section ($\substack{B\\C}$ section) turn out to be completely determined by the connectivity of the full $\substack{C\\B}$ section ($\substack{B\\C}$ section). To aid in the discussion of the connectivity of hybrid pipe dreams we introduce the following definitions.

Definition 2. *Define the ith edge content $C_i(D)$ of either a $\substack{C\\B}$ section or $\substack{B\\C}$ section D to be (x, y) where x is the label of the pipe passing from the tile in the $(i+1)$st column of the classic row of D to the ith column of the classic row of D (or $\emptyset$ if there is no such pipe) and y is the label of the pipe passing from the ith column of the bumpless row of D to the $(i+1)$st column of the bumpless row of D (and $\emptyset$ if there is no such pipe).*

Define the boundary ∂D of D to be the list of labels along the top, bottom, left, and right boundary edges of D where D may be either a $\substack{C\\B}$ section or a $\substack{B\\C}$ section.

If D is a $\substack{C\\B}$ section ($\substack{B\\C}$ section) such that for all $\substack{C\\B}$ sections ($\substack{B\\C}$ sections) D' such that $\partial D' = \partial D$, $C_i(D') = (x, y)$ we call i **separating** *for D.*

Lemma 1. *If D is a $\substack{C\\B}$ section ($\substack{B\\C}$ section) and $C_i(D) = (x, y)$ with $x \neq y$ then i is separating for D and in fact $C_i(D') = (x, y)$ for any $\substack{C\\B}$ section ($\substack{B\\C}$ section) D' with $\partial D = \partial D'$. In particular, if we introduce the following conditions:*

- *(i) the label x appears on the top edge of D in a column right of the ith column (or on the right edge of the classic row)*
- *(ii) the label x appears on the bottom edge of D in a column left of the ith column (or on the right edge of the classic row)*
- *(iii) the label y appears on the top edge of D in the ith column or to the left of it (or on the left edge of the bumpless row)*
- *(iv) the label y appears on the bottom edge of D in a column right of the ith column (or on the right edge of the bumpless row),*

then

1. *$C_i(D) = (x, y)$ with $x \neq y$ and $x \neq \emptyset \neq y$ iff conditions (i),(ii),(iii), and (iv) all hold.*
2. *$C_i(D) = (x, \emptyset)$ with $x \neq \emptyset$ iff conditions (i) and (ii) hold and for all y, either condition (iii) does not hold or condition (iv) does not hold.*

3. *$C_i(D) = (\emptyset, y)$ with $y \neq \emptyset$ iff conditions (iii) and (iv) hold, and for all x, either condition (i) does not hold or condition (ii) does not hold.*

Lemma 2. *If the label $x \neq \emptyset$ appears on the top edge of the ith column of D and the bottom edge of the jth column of D for some $j \geq i$ or the right edge of the bumpless row of D when D has width j, then $i-1, i, \ldots, j$ are all separating for D. Additionally, if the label $x \neq \emptyset$ appears on the bottom edge of the ith column and the left edge of the bumpless row of D then $1, 2, \ldots, i$ are all separating for D.*

Definition 3. *Let D be a ${}^{C}_{B}$ section (${}^{B}_{C}$ section) of width n. The interval $[a, b]$ is irreducible in D if no $i \in [a, b-1]$ is separating for D. An irreducible interval $[a, b]$ in D is called maximal if $a-1$ and b are separating for D.*

Lemma 3. *Let F be a full ${}^{C}_{B}$ section (${}^{B}_{C}$ section) and let D be the ${}^{C}_{B}$ section (${}^{B}_{C}$ section) sitting in columns a through b of F where $b > a$. If $[a, b]$ is a maximal irreducible interval for F then D is an obstructive ${}^{C}_{B}$ section (${}^{B}_{C}$ section).*

Proof. Suppose that $[a, b]$ is a maximal irreducible interval for F.

Say a pipe p enters D from the top edge of the ith column (with $i \in [a, b]$ of course). If p exits D from the bottom edge of a column right of i (or exits from the bottom of i) then $i-1$ and i are both separating for D by Lemma 2. In that case, irreducibility implies that $i-1 = a-1$ and $i = b$ which would then contradict the assumption that $b > a$. We conclude that p must exit D either to the left of the ith column or through a vertical edge of D. If p exits through the right edge of the bumpless row of D then $b-1$ is separating for D by Lemma 2 so either D is not irreducible or $b = a+1$. Either case is a contradiction so we conclude that p exits through the left edge of the classic row of D or somewhere left of the ith column. If p exits D in the jth column for some $a \leq j < i$ then $i-1$ will be separating for D by Lemma 1 parts 1 and 2. Irreducibility then implies $i-1 = a-1$ which contradicts our assumption that $a \leq j < i$. Since p can't exit D in the bottom of column i and it can't exit D in a column $j < i$ or a column $j > i$, the only remaining possibility is that p exits D through the left edge of the classic row of D and $i = a$. In summary, if there is a pipe entering D from a top edge, it must enter from the top edge of column a and exit through the left edge of the classic row of D.

Suppose a pipe p exits D from the bottom edge of the ith column of D. By the preceding paragraph, p cannot have entered D from a top edge, so p must have entered D from the right edge of the classic row of D or the left edge of the bumpless row of D. Suppose towards a contradiction that p entered through the left edge of the bumpless row. Then, a will be separating for D by Lemma 2. Since $b-1 > a$, that contradicts irreducibility. Therefore, we conclude that p can only have entered D from the right edge of the classic row. If i is less than b, then Lemma 1 shows $b-1$ to be separating for D. Thus, we conclude $i = b$ by irreducibility. Our takeaway is that if a pipe p exits

D from a bottom edge, it must be entering from the right edge of the classic row and exiting from the bottom edge of column b.

In addition to what we know already, Lemma 1 tells us that for all $i = a, \ldots, b-1$, $C_i(D) = (\emptyset, \emptyset)$ or $C_i(D) = (x, x)$ for some x. With all the restrictions on the contents of D we can now go through the list of 2×1 $\substack{C\\B}$ section columns and and find that if D is a $\substack{C\\B}$ section, the possible columns are: 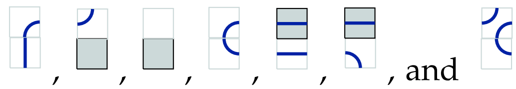 . If D is a $\substack{B\\C}$ section, the possible tiles will be: 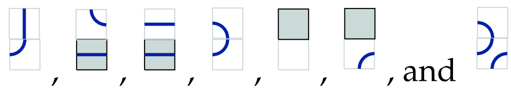 . This does not yet tell us that D is an obstructive $\substack{C\\B}$ section ($\substack{B\\C}$ section). The problem is that D could start with a column which is not a start state of the finite state machine or end with a column that is not an end state of the finite state machine. Call a column of a $\substack{C\\B}$ section ($\substack{B\\C}$ section) problematic if it is included in the finite state machine for obstructive $\substack{C\\B}$ sections ($\substack{B\\C}$ sections). In the finite state machine for $\substack{C\\B}$ sections ($\substack{B\\C}$ sections), the valid start states are exactly the columns which can be preceded by some non-problematic tile and the valid end states are exactly the columns that can be followed by some non-problematic tile. The columns which are not valid start states cannot appear at the beginning of a full $\substack{C\\B}$ section ($\substack{B\\C}$ section). The columns which are not valid end states cannot appear at the end of a full $\substack{C\\B}$ section ($\substack{B\\C}$ section). Hence, if D starts with an invalid start state/ends with an invalid end state, it will be preceded/followed by a sequence of problematic tiles beginning/ending with a valid start/end column. This larger $\substack{C\\B}$ section ($\substack{B\\C}$ section) will be obstructive. If this containment is proper, then either $a-1$ or b won't be separating. That would contradict maximality, so we conclude that D is an obstructive $\substack{C\\B}$ section ($\substack{B\\C}$ section). □

Theorem 5. *There is a weight preserving bijection between decorated full $\substack{C\\B}$ sections and decorated full $\substack{B\\C}$ sections.*

Proof. To map a $\substack{C\\B}$ section D to a $\substack{B\\C}$ section we first split it up into smaller pieces, splitting it up between columns i and $i+1$ whenever i is separating for D. This leaves us with maximal irreducible intervals which by Lemma 3 will either be single columns or obstructive $\substack{C\\B}$ sections. Obstructive $\substack{C\\B}$ sections (including ones which are single columns) are mapped to obstructive $\substack{B\\C}$ sections in the manner described in Theorem 4 and all other single columns are mapped to single columns in the manner described in Theorem 3. The 23 replacement rules in Theorem 3 are characterized by the property that
$$C_i(\phi_y(D)) = \begin{cases} (\emptyset, \emptyset) & C_i(D) = (x, x) \text{ for some } x \\ (x, x) \text{ for some } x & C_i(D) = (\emptyset, \emptyset) \\ C_i(D) & \text{else.} \end{cases}$$
This property also holds on the leftmost and rightmost boundary edges for the bijection between decorated obstruc-

tive ${}^{C}_{B}$ sections (${}^{B}_{C}$ sections). Thus, when we apply our construction on each maximal irreducible interval, the results will fit together. This construction sends maximal irreducible intervals to maximal irreducible intervals so it is indeed a bijection. □

As an example of Theorem 5, the following ${}^{C}_{B}$ section and ${}^{B}_{C}$ section are mapped to one another when all decorations are x.

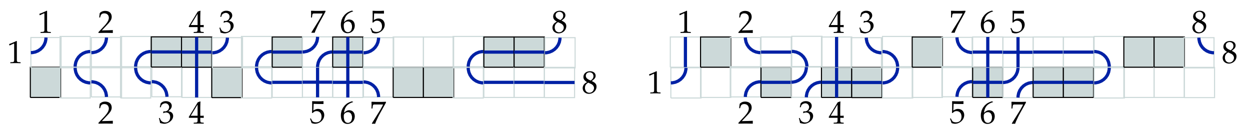

Acknowledgements

We thank Daoji Huang and Paul Zinn-Justin for many pertinent conversations. We are grateful to Zachary Hamaker and Adam Gregory for writing some of the LaTeX macros used here. Finally, we thank the reviewers for improvement of the exposition.

References

[1] N. Bergeron and S. Billey. "RC-graphs and Schubert polynomials". *Experimental Mathematics* **2**.4 (1993), pp. 257–269.

[2] S. C. Billey, A. E. Holroyd, and B. J. Young. "A bijective proof of MacdonaldâĂŹs reduced word formula". *Algebraic Combinatorics* **2**.2 (2019), pp. 217–248.

[3] S. Fomin and A. N. Kirillov. "The Yang-Baxter equation, symmetric functions, and Schubert polynomials". *Discrete Mathematics* **153**.1-3 (1996), pp. 123–143.

[4] Y. Gao and D. Huang. "The canonical bijection between pipe dreams and bumpless pipe dreams". 2021. arXiv:2108.11438.

[5] D. Huang. "Bijective Proofs of Monk's rule for Schubert and Double Schubert Polynomials with Bumpless Pipe Dreams". 2020. arXiv:2010.15048.

[6] A. Knutson and E. Miller. "Gröbner geometry of Schubert polynomials". *Annals of Mathematics* (2005), pp. 1245–1318.

[7] T. Lam, S. J. Lee, and M. Shimozono. "Back stable Schubert calculus". *Compositio Mathematica* **157**.5 (2021), pp. 883–962.

[8] A. Lascoux. "Chern and Yang through ice". 2002. Link.

[9] A. Weigandt. "Bumpless pipe dreams and alternating sign matrices". *Journal of Combinatorial Theory, Series A* **182** (2021), p. 105470.