

COMPUTATION AND SAMPLING FOR SCHUBERT SPECIALIZATIONS

DAVID ANDERSON, GRETA PANOVA, AND LEONID PETROV

ABSTRACT. We present computational results related to principal specializations of the Schubert polynomials $\mathfrak{S}_w(1^n)$ for permutations $w \in S_n$. Equivalently, these specializations count reduced pipe dreams (and reduced bumpless pipe dreams – RBPd) with boundary conditions determined by w . We find the first counterexample, at $n = 17$, to the conjecture of Merzon–Smirnov [MS16] that the maximal value of $\mathfrak{S}_w(1^n)$ is obtained at a layered permutation. However, the simulations suggest that $\lim_{n \rightarrow \infty} \log(\max_{w \in S_n} \mathfrak{S}_w(1^n))/n^2$ is the same constant arising for layered permutation from [MPP19]. Simultaneously, we explore the typical permutation obtained from uniformly random RBPds, i.e. drawn from the distribution proportional to $\mathfrak{S}_w(1^n)$. Our simulations reveal a permuton-like asymptotic behavior similar to the one derived for the analogous problem for Grothendieck polynomials in [MPPY25].

We implement and compare the performance of three recurrence relations for computing $\mathfrak{S}_w(1^n)$: the descent formula of Macdonald, the transition formula of Lascoux–Schützenberger, and the cotransition formula of Knutson. We investigate Markov chain algorithms for sampling uniformly random reduced bumpless pipe dreams (whose number is $\sum_{w \in S_n} \mathfrak{S}_w(1^n)$). We prove a negative result: the global constraint of reducedness breaks the sublattice property of the underlying alternating sign matrix (ASM) lattice, which prevents the use of the standard monotone Coupling From The Past (CFTP), and leads to false coalescence of the extremal chains. To bypass this, we develop a highly efficient Markov chain Monte Carlo (MCMC) sampler, augmented with macroscopic “droop” updates to guarantee state space connectivity and accelerate mixing. Our implementations enable computation of $\mathfrak{S}_w(1^n)$ for permutations w in S_n up to $n \sim 20$ on a personal computer (and beyond on a computing cluster), as well as uniform sampling of reduced bumpless pipe dreams up to $n \sim 60$ on a personal computer (and $n \sim 100$ on a cluster).

CONTENTS

1. Introduction	2
2. Computation of Schubert specializations	6
3. Performance comparison	9
4. Sampling: failure of CFTP with local flips	12
5. Sampling: MCMC for reduced bumpless pipe dreams	17
6. Conjectures and open problems	23
Appendix A. Background on Schubert polynomials	30
References	31

Date: March 20, 2026.

DA was partially supported by NSF grant DMS-1945212 and by a Membership at the Institute for Advanced Study funded by the Charles Simonyi Endowment. GP was partially supported by NSF grant CCF:AF-2302174. LP was partially supported by NSF grant DMS-2153869 and by the Simons Foundation Travel Support for Mathematicians Awards. Part of this research was performed while GP and LP were visiting the Institute for Pure and Applied Mathematics (IPAM), supported by NSF Grant No. DMS-1925919. The authors acknowledge Research Computing at the University of Virginia for providing computational resources and technical support (<https://rc.virginia.edu>).

1. INTRODUCTION

It is difficult to find a black cat in a dark room, especially if the cat is not there. Likewise, it is difficult to prove a conjecture, especially if it is wrong. Here we investigate the saga of the maximal principal specializations of Schubert polynomials, initiated by Stanley in the aptly named “shenanigans” paper [Sta17] which asked for the leading term asymptotics of the maximal principal specialization of a Schubert polynomial. Asymptotically, this is the same as the number of reduced bumpless pipe dreams (RBDs), a six-vertex model with long-range interaction. Merzon and Smirnov [MS16] conjectured that the maximum of principal specializations is achieved by a layered permutation, and Morales, Pak, and Panova [MPP19] determined the optimal layered permutation and its asymptotics. For $n \leq 13$, the maximizer is indeed a layered permutation.

In this paper we discover a slightly larger Schubert polynomial for $n = 17$, thereby disproving the conjectured maximal family. At the same time, we explore the typical permutations sampled from uniformly random RBDs, that is, permutations w drawn with probabilities proportional to the principal specialization $\mathfrak{S}_w(1^n)$ of their Schubert polynomials. We experimentally derive the conjectured permutation and limit shapes for the RBD. The computational data resembles the behavior shown to hold for the analogous Grothendieck polynomials in [MPPY25] and lends support to the conjectured asymptotic behavior. Obtaining these experimental results is a significant computational challenge which we tackle in this paper. It is still not practically possible to search for the permutation maximizing the Schubert polynomial for values of $n > 20$, but the counterexamples and evidence we produce show that such permutations may not have a simple defining structure, yet the asymptotic growth of the principal Schubert specializations is the same as for layered.

1.1. Background and definitions. Let S_n denote the group of permutations of n elements. Schubert polynomials $\mathfrak{S}_w$, indexed by $w \in S_n$ and depending on n variables $x_1, \dots, x_n$, are fundamental objects in algebraic combinatorics and algebraic geometry which give a basis for the homology of flag varieties. For a thorough survey of the role of Schubert polynomials in enumerative geometry and algebraic combinatorics, see [BGP25].

Schubert polynomials admit two combinatorial interpretations of significance to statistical mechanics, namely through tilings with certain global constraints.

The first, introduced by Bergeron and Billey [BB93] and Fomin and Kirillov [FK97], represents Schubert polynomials as generating functions over *reduced pipe dreams* (also called RC-graphs). These are configurations of crossing and elbow tiles in a staircase shape, satisfying the constraint that each pair of pipes crosses at most once (the *reduced* condition). To each reduced pipe dream one can associate its *boundary permutation* $w \in S_n$ by tracing where the pipes of all colors exit at the top. See Figure 1 for an illustration.

More recently, Schubert polynomials were interpreted as generating functions over objects of another type, *reduced bumpless pipe dreams* (RBD, for short). This interpretation was promoted by Lam, Lee, and Shimozono [LLS21], and further developed by Weigandt [Wei21], who also observed a connection with earlier work by Lascoux [Las02]. Bumpless pipe dreams are tilings of an $n \times n$ square grid with six types of tiles, subject to the same global *reduced* condition that each pair of pipes crosses at most once. Without the reducedness constraint, bumpless pipe dreams are the same as configurations of the six-vertex (square ice) model with domain wall boundary conditions. These configurations are well-known to be in bijection with alternating sign matrices, see, e.g., Bressoud [Bre99] and Zinn-Justin [ZJ09].




FIGURE 1. Pipe dreams for $n = 4$. Left: A reduced pipe dream for $w = 1432$ where each pair of pipes crosses at most once. Right: A non-reduced (forbidden) pipe dream where pipes 3 and 4 cross twice. Color is added merely as a visual aid to highlight the pipes' paths.

Similarly to the ordinary pipe dreams in the staircase, to each RBPD one can associate its *boundary permutation* $w \in S_n$ by tracing where the pipes of all colors exit through the right side. See Figure 2 for an illustration.





FIGURE 2. **Top:** The six tile types used in bumpless pipe dreams. **Bottom:** A reduced bumpless pipe dream for $n = 4$ corresponding to the permutation $w = 2143$.

Schubert polynomials $\mathfrak{S}_w(x_1, \dots, x_n)$ are defined by attaching variables x_i to the rows of these combinatorial objects. (More generally, one can assign two sets of variables x_i, y_j in a suitable way to get double Schubert polynomials, see Lascoux-Schützenberger [LS85a], Macdonald [Mac91].) Since here we focus only on the principal specializations, we simply define

$$\begin{aligned} \Upsilon_w &:= \mathfrak{S}_w(1^n) = \mathfrak{S}_w(\underbrace{1, 1, \dots, 1}_n) \\ &:= \#\{\text{reduced pipe dreams of size } n \text{ with boundary permutation } w\} \\ &= \#\{\text{reduced bumpless pipe dreams of size } n \text{ with boundary permutation } w\}. \end{aligned} \tag{1.1}$$

We recall the full definition and relevant background results on Schubert polynomials in Appendix A.

1.2. Asymptotics of Schubert polynomial evaluations. Stanley [Sta17] posed the following fundamental question about the asymptotic behavior of the principal specializations (1.1): does the limit

$$\lim_{n \rightarrow \infty} \frac{1}{n^2} \log_2 \max_{w \in S_n} \Upsilon_w \quad (1.2)$$

exist, and if so, what is its value and for which permutations w is the maximum value of Υ_w achieved? This question remains open. Morales, Pak, and Panova [MPP19] established a lower bound of approximately 0.29 for this limit by using *layered permutations* (permutations whose diagram consists of consecutive decreasing blocks, e.g., 32165487). An upper bound of approximately 0.37 follows from a connection with Alternating Sign Matrices and the six-vertex model. We refer to Morales, Panova, Petrov, and Yeliussizov [MPPY25, Section 6] for further discussion.

In [MS16] Merzon-Smirnov observed that the maximum of Υ_w is achieved by layered permutations for values $n \leq 10$ and this formed the basis for the working conjecture:

Conjecture 1.1 ([MS16]). *The maximal value of Υ_w for $w \in S_n$ is achieved when w is a layered permutation.*

This conjecture was exhaustively verified by one of us (DA) for $n \leq 13$ in February 2025. The maximal values for $n \leq 12$ were reported in [OEI25, sequence A284661]. For $n = 13$, the tools developed in the present paper greatly improve the speed of previous computations (and confirm their reliability); this is the first report of the verified maximal value of Υ_w for $w \in S_{13}$ (see Proposition 3.6).

In May 2025, Adam Zsolt Wagner (along with DA and Alejandro Morales) deployed Google DeepMind’s FunSearch [RPBN+24] to seek counterexamples to Conjecture 1.1. For $n \leq 16$ the heuristics found by the model did not uncover any counterexamples, providing weak evidence in favor of the conjecture in this range. (For larger n , time constraints limited the power of this method.)

At $n = 17$, however, using our improved computational methods together with the simple idea of looking at permutations that are “close” to layered, we disprove Conjecture 1.1.

Theorem 1.2. *Conjecture 1.1 is false. The non-layered permutation*

$$w^* = (1, 3, 2, 7, 6, 5, 17, 4, 16, 15, 14, 13, 12, 11, 10, 9, 8) \in S_{17}, \quad (1.3)$$

obtained from the optimal layered permutation $w(1, 2, 4, 10)$ by a single adjacent transposition (s_7), satisfies

$$\Upsilon_{w^*} = 3,272,424,600,397,137,120,000,$$

exceeding the layered maximum

$$\Upsilon_{w(1,2,4,10)} = 3,050,684,475,186,219,300,000$$

by about 7%. The permutation

$$u^* = (1, 3, 2, 8, 6, 5, 17, 4, 16, 15, 14, 13, 12, 11, 10, 9, 7) \in S_{17}, \quad (1.4)$$

obtained by two transpositions from $w(1, 2, 4, 10)$, satisfies

$$\Upsilon_{u^*} = 3,528,445,515,842,977,489,500,$$

exceeding $\Upsilon_{w(1,2,4,10)}$ by about 15.6%.

Proof. Direct computation using the (co)transition formula (Sections 2.3–2.4). □

While layered permutations fail to achieve the absolute maximum value of Υ_w for large n , the general bounds found in Section 6.4 (Proposition 6.7 and Corollary 6.8) imply that the maximal asymptotic behavior appears to be the same, whether one considers all permutations or only layered ones.

1.3. Computation and sampling. Motivated by these open questions, we develop computational tools for exploring principal specializations of Schubert polynomials and the underlying combinatorial structures. We address two main questions:

- How to **efficiently compute** Υ_w for permutations $w \in S_n$ when n is large?
- What does a **typical reduced bumpless pipe dream** of size n look like, when sampled uniformly at random from all RBPDs (regardless of the corresponding permutation w)?

The lower bound of about 0.29 for the (conjectural) limit (1.2) was obtained in [MPP19] by using an explicit product formula for Υ_w when w is a layered permutation. This formula enables efficient optimization over layered permutations. For general permutations, by contrast, even if one can compute Υ_w efficiently for any given w , exhaustively searching over all $n!$ permutations in S_n to find the maximum is infeasible even for moderate n (say, $n > 13$). This motivated us to study *typical* permutations (in the sense of random reduced bumpless pipe dreams) and, in particular, compare the values of Υ_w on them with the maximum over layered permutations known from [MPP19].

Despite the counterexamples, the simulations and observed *typical* behavior suggest that the asymptotic maximum is still the same as for layered, so we expect that

$$\limsup_{n \rightarrow \infty} \frac{1}{n^2} \log_2 \max_{w \in S_n} \Upsilon_w < 0.3$$

and that the actual limit exists.

Remark 1.3 (Typical vs. maximal). The RBPD-typical permutation, whose limiting shape is conjecturally described by the permuton in Conjecture 6.1, does not resemble a layered permutation at large scale, and its Υ_w may be significantly smaller than the maximum.

For sampling, we first attempted to apply the standard monotone Coupling From The Past (CFTP) algorithm. However, we show that reduced bumpless pipe dreams do not form a sublattice of the alternating sign matrix lattice (Section 4). This breaks the monotone coupling required by CFTP: the internal rejection scheme needed to enforce reducedness can cause the two extremal chains to cross, invalidating the sandwiching argument. Instead, we implement a specialized Markov chain Monte Carlo (MCMC) algorithm with macroscopic block updates (Section 5).

Based on extensive simulations, we pose **Conjecture 6.1 on the limiting permuton** structure and other properties of typical Schubert permutations and **Conjecture 6.2 on the limit shape and arctic curves of uniformly random reduced bumpless pipe dreams** in Section 6. The Schubert and Grothendieck permutons show striking similarities as pictured in Figure 3. The Grothendieck permuton arises from a non-reduced BPD model and is analyzed rigorously using integrable probability techniques coming from TASEP in [MPY25]. The requirement in the Schubert model that no pipes cross more than once makes it a nonlocal model and not amenable to integrable techniques at present¹.

¹The Schubert model can also be interpreted as a colored vertex model removing the “long-range interaction” condition. However, then the number of colors is n and asymptotic analysis cannot be performed.




FIGURE 3. The simulated density histogram for $n = 100$ for the Schubert (left) and Grothendieck (right) random permutations, displaying striking similarities in behavior. The permuton limit for the Grothendieck case is proven in [MPPY25, Theorem 1.5], while the Schubert case is open.

Acknowledgements. We are grateful to the many colleagues in the algebraic combinatorics and integrable probability communities who have expressed interest in this topic. In particular, we benefited from conversations with Amol Aggarwal, William Fulton, Vadim Gorin, Daoji Huang, Allen Knutson, David Speyer, and Adam Zsolt Wagner; from Igor Pak’s rightful suspicion of Conjecture 1.1; and from fruitful discussions with Richard Stanley about (1.2) and related questions.

A special acknowledgement goes to Alejandro Morales for his earlier work and computations with principal specializations, and for popularizing this topic; as well as to Linus Setiabrata for insightful attempts at proving Conjecture 1.1. We heartily thank both of them for many enjoyable and inspiring conversations.

Note on AI tools. We used Claude Code as a coding assistant for optimizing the Schubert evaluation and sampling code, and then subjected our code to thorough testing and verification across multiple implementations (C++ and Julia). After we found the first permutation w^* (1.3) violating Conjecture 1.1, the second example u^* (1.4) was obtained with the help of ChatGPT-5.4 Pro, in the process of verifying the computation of Υ_{w^*} .

2. COMPUTATION OF SCHUBERT SPECIALIZATIONS

We use three recurrence formulas for computing Υ_w . In this section, we state these formulas and discuss their implementation details. Performance benchmarks and comparisons of the implementations are presented in Section 3.

2.1. Notation for permutations. We fix some standard notation: $\ell(w)$ denotes the number of inversions of w , $s_i = (i, i + 1)$ are adjacent transpositions, $t_{i,j} = (i, j)$ are not necessarily adjacent transpositions, $\text{Des}(w) = \{i : w(i) > w(i + 1)\}$ is the descent set, $e = (1, 2, \dots, n)$ is the identity, and $w_0 = (n, n - 1, \dots, 1)$ is the longest permutation. We use right multiplication on one-line notation: $w \cdot (a, b)$ is obtained from w by swapping the entries in positions a, b . A word

$(a_1, \dots, a_k)$ with $1 \leq a_i < n$ is called a *reduced word* for a permutation $w \in S_n$ if $w = s_{a_1} s_{a_2} \cdots s_{a_k}$ and $k = \ell(w)$ (that is, this expression for w in elementary transpositions is of minimal possible length).

The permutation matrix associated to w has 1's in positions $(i, w(i))$ and 0's elsewhere. Often we indicate the positions of 1's by a dot, leaving the 0's blank.

2.2. Descent formula. The descent recurrence formula [Mac91], [FS94], [BHY19] expresses Υ_w in terms of values at permutations of smaller length, where the sum is over the descents:

Theorem 2.1 (Descent formula). *For $w \in S_n$ with $w \neq e$, we have:*

$$\Upsilon_w = \sum_{i \in \text{Des}(w)} \frac{i}{\ell(w)} \cdot \Upsilon_{w \cdot s_i} = \sum_{i \in \text{Des}(w^{-1})} \frac{i}{\ell(w)} \cdot \Upsilon_{s_i \cdot w}. \quad (2.1)$$

The base case is $\Upsilon_e = 1$.

Note that each permutation $w \cdot s_i$ occurring on the RHS of (2.1) has length $\ell(w) - 1$, so the recurrence terminates at $\Upsilon_e = 1$.

Proof of Theorem 2.1. From [Mac91, (6.11)], we have

$$\Upsilon_w = \frac{1}{\ell(w)!} \sum_{(a_1, \dots, a_{\ell(w)}) \in R(w)} a_1 a_2 \cdots a_{\ell(w)}, \quad (2.2)$$

where $R(w)$ is the set of reduced words for w . Splitting the sum according to the descents of w (or w^{-1} for the second sum) and using (2.2) for the resulting shorter permutations $w \cdot s_i$ yields the desired recurrence. $\square$

Let us discuss the implementation of the descent formula of Theorem 2.1.

- 1. Permutation encoding:** Permutations are encoded as 128-bit integers using 5 bits per element, supporting $n \leq 25$, which enables fast hashing and efficient memory usage. Adjacent transposition $w \mapsto w \cdot s_i$ is performed via bitwise operations directly on the packed representation. For enumeration over S_n (Section 3.3), we use a more compact 64-bit encoding (4 bits per element) for $n \leq 16$, halving memory usage.
- 2. Preprocessing.** Before running the recurrence, we strip trailing fixed points of w (suffix positions where $w(k) = k$), since this does not change Υ_w , to reduce the effective size of the problem, if possible.
- 3. Memoization and length decrement.** We store computed values in a hash table keyed by the packed permutation, with a hard cap on the number of entries ($3 \cdot 2^{26} \approx 192\text{M}$ for double, $2^{27} \approx 128\text{M}$ for exact arithmetic) to keep the table cache-friendly. The length-decreasing nature of the recursion lets us pass $\ell(w)$ as a parameter and decrement it at each step, avoiding the $O(n^2)$ cost of recomputing length from scratch.
- 4. BFS sort-reduce evaluation.** We evaluate the descent recurrence level-by-level using a breadth-first search (BFS) sort-reduce algorithm. Starting from the target w at level $\ell(w)$, we process all permutations at a given length simultaneously. For each permutation v at the current level and each descent $i \in \text{Des}(v)$, we emit the child $v \cdot s_i$ (at length $\ell(v) - 1$) with its weighted contribution $\frac{i}{\ell(v)} \cdot \Upsilon_v$. Distinct parents may produce the same child. We *sort* the children by their packed key and *reduce* by summing contributions for each distinct key, yielding Υ for every permutation at the new level. The previous level is then discarded.
- 5. Arithmetic precision.** Formula (2.1) involves division by $\ell(w)$, requiring either rational or floating-point arithmetic. We provide two implementations:

- (a) *Double-precision arithmetic* (64-bit): Provides approximately 15 significant digits, which becomes imprecise for large values of Υ_w .
- (b) *Rational arithmetic* (exact): Uses a custom implementation of rational numbers with 128-bit integers for numerator and denominator and GCD reduction after each operation. Produces exact integer results (when they fit into 128 bits) at slower speed.

Since Υ_w is always an integer (counting reduced pipe dreams), the rational implementation serves as ground truth for validation.

2.3. Transition formula. The transition formula of Lascoux-Schützenberger [LS85b] (and its refinement by Fan-Guo-Sun [FGS18]) can be specialized to Υ_w as follows. Given $w = w_1 w_2 \cdots w_n$, let r be the largest index such that w_r appears as the “3” in a 132 pattern, i.e., there exist $i < r < s$ with $w_i < w_s < w_r$. If no such index exists, then w is dominant and $\Upsilon_w = 1$ [Mac91, (4.7)]. Otherwise, we have the following recurrence:

Theorem 2.2 (Transition formula [LS85b]). *For non-dominant $w \in S_n$, let r be as above, and let $s > r$ be the largest index such that $w_s < w_r$ and there exists $i < r$ with $w_i < w_s$. Then*

$$\Upsilon_w = \Upsilon_v + \sum_{i < r : \ell(v \cdot (i, r)) = \ell(w)} \Upsilon_{v \cdot (i, r)}, \quad (2.3)$$

where $v = w \cdot (r, s)$.

By the choice of r and s , we have $\ell(v) = \ell(w) - 1$.

The implementation of Theorem 2.2 uses the same packed permutation representation as in Section 2.2, and is based on iterative depth-first search (DFS) with memoization.

- 1. Packed representation and memoization.** States are stored as 128-bit packed permutations (5 bits per entry, $n \leq 25$) and memoized in hash tables keyed by the packed code. As in Section 2.2, we first strip trailing fixed points of w , and use the same memoization hard caps.
- 2. Transition index computation and base case.** For each state, we compute (r, s) directly on the packed representation. If no such pair exists (equivalently, the permutation is dominant), we apply the base case $\Upsilon_w = 1$.
- 3. Child generation and depth-first search state.** For non-dominant w , we first form $v = w \cdot (r, s)$ and then generate all terms $v \cdot (i, r)$ satisfying $\ell(v \cdot (i, r)) = \ell(w)$. We then follow one branch of this recursion tree to a base case before backtracking to process remaining children. The stack frame stores this child list explicitly; the first child has length $\ell - 1$, while the remaining children stay at length ℓ . Because of these same-level dependencies, this implementation uses depth-first search (DFS) rather than a level-by-level BFS variant.
- 4. Arithmetic precision.** We provide two precision versions: one with double precision, and another one with exact integer arithmetic based on GMP’s `mpz_class`. Since (2.3) uses only addition, no rational division is required.

2.4. Cotransition formula. The cotransition formula of Knutson [Knu19] provides yet another recurrence that uses only integer addition. Recall that w_0 denotes the longest permutation.

Theorem 2.3 (Cotransition formula [Knu19]). *For $w \in S_n$ with $w \neq w_0$, let us denote*

$$i := \min\{j : j + w(j) \leq n\}. \quad (2.4)$$

Then we have:

$$\Upsilon_w = \sum_{v > w, v(i) \neq w(i)} \Upsilon_v. \quad (2.5)$$

Here $v \succ w$ means that v covers w in Bruhat order (see Lemma 2.6 below). The base case is $\Upsilon_{w_0} = 1$.

Remark 2.4. The cotransition formula (2.5) proceeds *toward* the longest permutation w_0 , increasing the length at each step. Like the transition formula, (2.5) also involves only addition, and thus can be implemented using exact integer arithmetic.

Remark 2.5. From (2.5) it is not hard to see that there exists $w \in S_n$ with the maximal Υ_w such that $w(1) = 1$.

The following well-known characterization of the Bruhat order is useful for implementation:

Lemma 2.6 ([BB05, Lemma 2.1.4]). *A permutation v covers w in Bruhat order if and only if $v = w \cdot (a, b)$ for some $a < b$ such that $w(a) < w(b)$, and there is no k with $a < k < b$ and $w(a) < w(k) < w(b)$.*

Let us discuss the implementation details of the cotransition formula of Theorem 2.3.

1. **Permutation encoding.** As in Sections 2.2 and 2.3, we first strip trailing fixed points of w before evaluation. Permutations are then handled in the same packed form, 64-bit for $n \leq 16$ and 128-bit for $n \leq 25$.
2. **BFS sort-reduce evaluation.** The cotransition evaluator runs level-by-level using breadth-first search (BFS). At each level, for each permutation w on the current frontier with value Υ_w , we enumerate all Bruhat covers $v \succ w$ satisfying $v(i) \neq w(i)$ and emit the pair (v, Υ_w) . Since distinct permutations at the same level may share a common cover v , we apply sort-reduce as in Section 2.2: sort emitted pairs by packed key and sum the values for each distinct v . This computes $\Upsilon_w = \sum_v \Upsilon_v$ per (2.5) for all permutations at the new level. This collapses repeated states and keeps memory usage bounded.²
3. **Bruhat cover enumeration.** For each permutation w , we enumerate all Bruhat covers $v \succ w$ as in Lemma 2.6. We filter to those covers where $v(i) \neq w(i)$ at the cotransition index i , see (2.4). We store the cover pairs in a fixed-size array (not heap-allocated), eliminating memory allocation overhead in the inner loop.
4. **Arithmetic precision.** Since formula (2.5) involves only addition, exact integer arithmetic is natural. We provide two implementations:
 - (a) *Exact arithmetic* uses the GMP library (`mpz_class`) to compute exact integer values regardless of magnitude. This is essential for large n , where Υ_w can exceed 50 digits.
 - (b) *Double-precision* (64-bit) is typically faster, but may become imprecise for large Υ_w .

Remark 2.7. A folklore rule of thumb holds that the transition formula is the most efficient means of computing Schubert polynomials. (See discussion and debate in [MPY22, §1.2].) As we will see in the next section, this expectation is close to true in our simplified setting of principal evaluations, especially as n grows. The similar cotransition formula performs equally well or better for smaller n .

3. PERFORMANCE COMPARISON

We present a comparison of the performance of the descent, transition, and cotransition formulas from Section 2. Our comparisons are anchored by known closed formulas for *layered permutations* and, in particular, the permutations maximizing Υ_w over the layered ones [MPP19] (which

²An older implementation used DFS with memoization but it turned out to be much slower, as reported in Remark 3.4 below.

are *not* the absolute maximizers over all permutations, see Theorem 1.2 for a counterexample). Most computations were performed on a laptop with Apple M2 Pro chip and 16GB RAM, and some of the heavier computations (where indicated) used a more powerful desktop computer with an AMD Ryzen 9 processor with 64GB RAM.

3.1. Layered permutations. For a composition $(b_1, \dots, b_k)$ of n we recursively define the *layered permutation* $w(b_1, \dots, b_k) := (w(b_1, \dots, b_{k-1}), n, n-1, \dots, n-b_k+1)$, where $w(b) = (b, b-1, \dots, 1)$. In other words, the layered permutation $w(b_1, \dots, b_k)$ is obtained by splitting $1, \dots, n$ into blocks of sizes $b_1, b_2, \dots, b_k$ and reversing each block.

It is straightforward to compute the length of a layered permutation:

Lemma 3.1. *The length of a layered permutation $w(b_1, \dots, b_k)$ is equal to $\sum_{i=1}^k \binom{b_i}{2}$.*

From [MPP19], we explicitly know the permutations on which Υ_w achieves the maximum among the layered permutations. Let $w^*(n)$ denote this permutation, i.e., $w^*(n)$ is the layered permutation maximizing Υ_w over all layered permutations in S_n .

Lemma 3.2. *We have*

$$\lim_{n \rightarrow \infty} \frac{\ell(w^*(n))}{\binom{n}{2}} = \frac{1-\alpha}{1+\alpha} \approx 0.3955, \quad (3.1)$$

where $\alpha \approx 0.4331818312$ is the constant from [MPP19].

Proof. By [MPP19, Theorem 1.1], the optimal block sizes satisfy $b_i \sim \alpha^{i-1}(1-\alpha)n$ as $n \rightarrow \infty$. By Lemma 3.1,

$$\frac{\ell(w^*(n))}{\binom{n}{2}} \sim \frac{\sum_i b_i^2}{n^2} = (1-\alpha)^2 \sum_{i=0}^{\infty} \alpha^{2i} = \frac{(1-\alpha)^2}{1-\alpha^2} = \frac{1-\alpha}{1+\alpha}. \quad \square$$

As the first benchmark, we compute $\Upsilon_{w^*(n)}$ for the maximizing layered permutations, for n up to 17, using all three formulas with double-precision and exact arithmetic implementations. Table 1 summarizes the timing results.

For layered permutations, cotransition is the fastest method overall in double precision. The transition formula has larger constant overhead for small n (≈ 0.1 seconds already for $n = 8-11$), but scales comparably to cotransition for larger n . In exact arithmetic, transition is already faster than cotransition at $n = 16, 17$. The descent formula remains competitive only at very small n and then grows much more steeply; its exact (rational) variant hits the 180s timeout by $n = 16$.

Remark 3.3 (Precision). At $n = 15$, the maximal value of Υ_w is $\approx 2.3 \times 10^{16}$, which exceeds $2^{53} \approx 9 \times 10^{15}$, the threshold beyond which double-precision floating-point arithmetic loses integer precision. The descent formula with double precision yields 23,399,330,089,073,392, while exact arithmetic gives 23,399,330,089,073,400 — a discrepancy of 8 in the units digit. The cotransition and transition formulas with double precision still produce the correct integer value at $n = 15$. At $n \geq 16$, both cotransition and transition in double precision lose integer accuracy.

3.2. Random permutations from the RBDP sampler. We also test the performance on typical permutations obtained from the reduced bumpless pipe dream (RBDP) sampler described in Section 5 below. We pick $n = 15$ and sample 200 permutations. Table 2 summarizes the timing.

The cotransition formula is by far the fastest on these RBDP-typical permutations.

Remark 3.4. For larger n , the computation time grows substantially. For example, the RBDP-typical permutation

$$w = (1, 3, 11, 2, 8, 4, 13, 7, 21, 9, 6, 19, 17, 18, 5, 16, 15, 14, 20, 10, 12) \in S_{21}$$

TABLE 1. Timing comparison for maximal layered permutations from [MPP19]. Times are in seconds. The “Layers” column shows the block structure of an optimal layered permutation (for example, for $n = 8$ the permutation is $w(1, 2, 5) = 13287654$). Asterisks indicate results from the range where integer accuracy is lost. The descent rational variant exceeded the 180s timeout at $n \geq 16$. The row $n = 300$ shows the largest value computed in [MPP19] for comparison (no timing data).

n	Layers	$\ell(w)$	$\frac{\log_2 \Upsilon_w}{n^2}$	Descent		Cotransition		Transition	
				double	rational	double	exact	double	exact
8	(1, 2, 5)	11	0.206	0.0001	0.0002	0.0001	0.0002	0.101	0.075
9	(1, 2, 6)	16	0.214	0.0002	0.0017	0.0005	0.0005	0.085	0.088
10	(1, 3, 6)	18	0.221	0.0006	0.0059	0.0009	0.0013	0.112	0.108
11	(1, 3, 7)	24	0.227	0.0069	0.0320	0.0040	0.0084	0.094	0.104
12	(1, 3, 8)	31	0.230	0.0373	0.2770	0.0168	0.0255	0.147	0.127
13	(1, 1, 3, 8)	31	0.234	0.0444	0.2823	0.0370	0.0582	0.185	0.139
14	(1, 1, 4, 8)	34	0.237	0.1491	1.1652	0.0784	0.1328	0.287	0.339
15	(1, 1, 4, 9)	42	0.242	1.4425*	14.5971	0.4032	0.7926	0.902	1.108
16	(1, 1, 4, 10)	51	0.244	20.5393*	>180	2.4683*	5.6003	3.367*	4.316
17	(1, 2, 4, 10)	52	0.247	93.5795*	>180	9.7029*	17.6501	9.779*	12.966
300	(1, 2, 6, 14, 32, 74, 171)	17839	0.290	—	—	—	—	—	—

TABLE 2. Performance statistics for 200 random permutations from the RBPB sampler at $n = 15$.

Statistic	Descent		Cotransition		Transition	
	double	rational	double	exact	double	exact
Mean time (s)	0.213	2.262	0.017	0.030	0.314	0.365
Median time (s)	0.105	0.974	0.012	0.021	0.271	0.312
Max time (s)	2.961	34.007	0.099	0.189	1.242	1.601
Total time (s)	42.7	452.3	3.5	5.9	62.8	73.1

with $\ell(w) = 67$ requires about 29 seconds using the BFS-mode exact-precision cotransition formula (the DFS mode with memoization takes several dozen minutes, and the descent formula does not terminate within hours). For a RBPB-typical permutation

$$w = (2, 1, 4, 13, 10, 3, 7, 19, 6, 23, 22, 12, 17, 21, 5, 16, 20, 11, 18, 14, 15, 9, 8) \in S_{23}$$

with $\ell(w) = 98$, the BFS-mode cotransition formula completes in about 9 minutes, while the DFS mode did not finish within a week.

Remark 3.5. While the cotransition formula is much faster on RBPB-typical permutations, the descent formula wins decisively on permutations close to the identity. For example, consider

$$w = (3, 4, 5, 1, 2, 8, 9, 10, 6, 7, 13, 14, 15, 11, 12, 18, 19, 20, 16, 17, 22, 21, 23, 24, 25) \in S_{25}$$

with $\ell(w) = 25$. After stripping trailing fixed points (n reduced to 22), the descent formula computes Υ_w in about 0.005 seconds, while the cotransition formula takes about 0.03 seconds.

3.3. Search for maximum. We search for the permutation maximizing Υ_w over all $w \in S_n$ and not just over the layered ones.

In the *full search for maximum* up to $n = 13$, we start two parallel threads: the descent thread computes Υ_w starting from the identity and proceeding upward in length, while the cotransition thread starts from w_0 and proceeds downward. Each thread uses the sort-reduce approach described in Sections 2.2 and 2.4. This keeps only the current level in memory. The two threads use *dynamic meeting*: each checks whether it has crossed the other thread's current level, and stops when the levels cross. This approach balances the workload between threads, as the descent and cotransition frontiers grow at different rates. Both threads track the maximum Υ_w encountered at each level, guaranteeing that the global maximum is found. We run the full search up to $n = 12$, where it completes in about 1 minute. For $n = 13$, this implementation runs out of the 16GB RAM available, but was completed in about 15 minutes on a machine with 64GB RAM. During some of the steps, it had exhausted all of the available RAM and used the Linux swap space.

We have the following experimental result:

Proposition 3.6. *For $n \leq 13$, the permutation maximizing Υ_w over all $w \in S_n$ is the same as the one maximizing over layered permutations, as given in [MPP19].*

The prior state of the art was exhaustive verification for $n \leq 10$, as reported in [MS16]. The cases $n = 11, 12$ were subsequently verified and reported in [OEI25] (sequence A284661), and the present work provides the first published record of exhaustive search through $n = 13$. For $14 \leq n \leq 16$, exhaustive search is infeasible, but no counterexample was found within Cayley distance 4 of the optimal layered permutation. At $n = 17$, however, a permutation at Cayley distance 1 from the optimal layered permutation exceeds the layered maximum (Theorem 1.2), along with similar counterexamples for $18 \leq n \leq 20$ (Section 6.3).

4. SAMPLING: FAILURE OF CFTP WITH LOCAL FLIPS

In this section, we analyze the state space of reduced bumpless pipe dreams (RBPDs) under local tile moves. A natural approach for exact sampling of uniformly random RBPDs is to embed them within the larger space of alternating sign matrices (ASMs) and use Coupling From The Past (CFTP) [PW98]. We explain why this naive approach fails.

4.1. Rothe diagram and Rothe BPD. The *Rothe diagram* of $w \in S_n$ is

$$D(w) := \{(i, j) \in [n] \times [n] : w(i) > j \text{ and } w^{-1}(j) > i\}. \quad (4.1)$$

Equivalently, $D(w)$ is obtained by placing a dot at $(i, w(i))$ for each i and removing all cells weakly right of or below each dot (in the same row and column, respectively). The number of cells is equal to the length: $|D(w)| = \ell(w)$. See Figure 4 for an example.

Interpreting the shaded cells of the diagram $D(w)$ as empty, and the rays deleting the boxes as pipes, we obtain a reduced bumpless pipe dream $\mathbf{b}_w$ for w , called the *Rothe BPD*. This is an RBPD canonically associated to w .

4.2. Height functions and local flips. Recall that a BPD uses the six tile types shown in Figure 2. Disregarding the reducedness constraint, the set of all BPDs of size n is in bijection with configurations of the six-vertex model with domain wall boundary conditions, or equivalently, ASMs of size n [Kup96], [Bre99]. This space is equipped with a natural lattice structure governed by height functions:




FIGURE 4. Rothe diagram $D(w)$ for $w = (3, 1, 4, 2)$. Dots mark positions $(i, w(i))$; shaded cells form $D(w)$. (Here $|D(w)| = 3 = \ell(w)$.)

Definition 4.1 (Height function and partial order). For an $n \times n$ BPD D , the *height function* $h_D(i, j) \in \mathbb{Z}_{\geq 0}$ is defined on the dual-lattice vertices $0 \leq i, j \leq n$ so that six-vertex lines are level lines of h_D .

We define a partial order on all BPDs by $D \leq D'$ if and only if $h_D(v) \leq h_{D'}(v)$ pointwise for all vertices v of the dual square lattice. Under this order, ASMs form a distributive lattice with a unique minimum $\mathbf{b}_{w_0}$ and maximum $\mathbf{b}_{\text{id}}$ (the Rothe BPDs for the longest and the identity permutation, respectively). We refer to Figure 5 for an illustration.





FIGURE 5. Height functions on $n = 4$ BPDs. Values of $h_D(i, j)$ appear at dual-lattice vertices; six-vertex lines are the level lines of h_D .

A natural dynamical system on the space of the ASMs proceeds by *flips* on 2×2 windows of tiles: whenever a 2×2 block admits an alternative valid tiling with the same boundary conditions, the two configurations can be exchanged. Each flip shifts the height function at the shared interior vertex by ± 1 , while preserving the height values everywhere else. Figure 6 displays all possible flips, organized into three types based on the presence of cross tiles:

- (a) *Drips*: no cross tile is involved, and the boundary permutation is preserved.
- (b) *Cross creation/annihilation*: a cross tile appears or disappears, either at the top-left corner (upper sub-row) or at the bottom-right corner (lower sub-row).
- (c) *Cross relocation*: the cross tile shifts from one corner to the opposite.

Flips of types (b) and (c) alter the global pipe crossing pattern and may change the boundary permutation w . In particular, a flip can turn a reduced BPD into a non-reduced one, or vice versa.

FIGURE 6. All flips on 2×2 windows of BPD tiles.

These flips define a Markov chain on the ASM lattice: at each step, choose an interior dual-lattice vertex (i, j) with $1 \leq i, j \leq n - 1$ uniformly at random and a direction (up or down) with equal probability, then apply the flip to the four cells sharing vertex (i, j) if one exists (otherwise the state is unchanged). It is not hard to see that flips connect the state space: any two BPDs (or equivalently, any two ASMs) are related by a sequence of flips. This chain is ergodic and reversible with respect to the uniform measure on all $n \times n$ ASMs.

Restricting to the subset of *reduced* BPDs, we conjecture that the same 2×2 flips suffice to connect the state space:

Conjecture 4.2 (Connectivity of the RBPD graph). *For every n , the set of reduced bumpless pipe dreams of size n is connected under the 2×2 flips (Figure 6) that preserve reducedness.*

We have verified Conjecture 4.2 computationally for $n \leq 8$ by checking that the graph whose vertices are RBPDs and edges are flips that preserve reducedness is connected. The remainder of this section assumes the conjecture holds. We discuss obstructions to “easy” proofs of this conjecture in Section 5.1 below.

4.3. Failure of monotone CFTP. We now explain why the standard Propp–Wilson monotone Coupling From The Past (CFTP) algorithm [PW98] cannot be used to sample uniformly from RBPDs of size n . We begin with a concrete counterexample at $n = 4$, then state the general obstruction.

At $n = 4$, the ASM lattice has 42 elements. Exactly 41 are reduced BPDs and one is not. The non-reduced element $\mathbf{b}^*$ is the unique 4×4 BPD in which a single pair of pipes crosses twice (Figure 7, right). Recall that the meet $A \wedge B$ in the ASM lattice is the BPD whose height function is the pointwise minimum $h_{A \wedge B}(v) = \min(h_A(v), h_B(v))$. One can check that among the $\binom{41}{2}$ pairs of RBPDs, there exist nine pairs $\{A, B\}$ whose meet $A \wedge B$ equals $\mathbf{b}^*$. One such pair is shown in Figure 7, left and center. Thus, we have the following result:

Proposition 4.3 (Failure of the sublattice property). *For $n \geq 4$, the set of reduced bumpless pipe dreams does not form a sublattice of the ASM lattice.*

Proof. The $n = 4$ case is verified directly as described above. For $n \geq 5$, the same pair of RBPDs embeds into the bottom-right 4×4 block of the $n \times n$ grid; the remaining $n - 4$ pipes simply

travel up and then right without interacting with the 4×4 block. The meet of this embedded pair remains non-reduced. $\square$




FIGURE 7. Failure of the sublattice property for RBPDs at $n = 4$.

The failure of the sublattice property has a direct consequence for CFTP. The Propp–Wilson algorithm [PW98] requires the random update map Φ to be *monotone*: if $X \leq Y$ in the partial order, then $\Phi(X) \leq \Phi(Y)$ with probability one. When this holds, it suffices to track the two extremal chains starting from the unique minimum and maximum elements of the state space, $\mathbf{b}_{w_0}$ and $\mathbf{b}_{\text{id}}$ (see Figure 5). The coalescence of these two extremal chains (together with the time-doubling procedure of the CFTP algorithm) guarantees that all intermediate chains have also coalesced (the “sandwiching” property), yielding an exact uniform sample.

To restrict the flip dynamics on the ASM lattice to RBPDs, one must reject any flip that produces a non-reduced BPD. One may consider two natural rejection schemes which still deal with two extremal chains X and Y starting from $\mathbf{b}_{w_0}$ and $\mathbf{b}_{\text{id}}$, respectively:

1. **Internal (per-chain) rejection.** Each chain independently checks whether the proposed flip preserves reducedness, and rejects it if not. This scheme preserves the marginal evolution of X and Y as the original flip dynamics on RBPDs, but breaks monotonicity: a shared flip may be accepted by one chain (which remains reduced) and rejected by the other (which would become non-reduced), causing the ordering $X \leq Y$ to be violated after the update. See Figure 8 for an example.
2. **Coupled rejection.** Reject the flip for both chains X, Y whenever it would make either chain non-reduced. This preserves monotonicity, but the transition probabilities of each chain now depend on the state of the other chain. The resulting coupling from the past process no longer targets the uniform distribution on RBPDs.

We see that neither scheme produces a valid monotone CFTP sampler.

Because reducedness depends on the global pipe crossing pattern, the height function at the flip site cannot detect whether a distant cross will cause a double crossing. The direct enumeration, presented in Table 3, shows that these monotonicity violations appear as soon as $n = 4$ and grow with n .

4.4. False coalescence. The internal rejection mechanism preserves the correct stationary distribution but breaks monotonicity. Here we directly test whether the “naive” version of the CFTP algorithm that tracks only the extremal chains X and Y (started from $\mathbf{b}_{w_0}$ and $\mathbf{b}_{\text{id}}$) can still be used to sample from the uniform distribution on RBPDs. Indeed, one might hope that the extremal chains rarely cross, and their coalescence may be close to the universal coalescence.




FIGURE 8. Monotonicity violation at $n = 4$ under internal rejection. States $X \leq Y$ have identical 2×2 windows around $v = (3, 3)$ (dashed box). The flip is accepted by Y but rejected by X due to the extra cross outside the window (highlighted in X). This yields $X' \not\leq Y'$ since $h_{X'}(3, 3) = 3$ but $h_{Y'}(3, 3) = 2$.

TABLE 3. Monotonicity violations under internal rejection. For each ordered pair $X \leq Y$ of RBPDs, all $2(n-1)^2$ flips are tested ($(n-1)^2$ interior vertices $\times$ 2 directions). A violation occurs when a shared flip is accepted by one chain but rejected by the other, breaking the order $X' \leq Y'$.

n	RBPDs	Ordered pairs	Flip checks	Violations
3	7	26	208	0
4	41	618	11,124	16
5	393	39,302	1,257,664	2,259

For each trial, we run the standard backward CFTP protocol with time doubling [PW98]: for $T = 1, 2, 4, 8, \dots$, we draw fresh random updates for times $-2T, \dots, -(T+1)$, place them before the existing updates for $-T, \dots, -1$, and apply the combined sequence of Markov steps in the two extremal chains, doubling T until the extremal chains coalesce at time 0. We record the final update sequence. We then replay *the same sequence* starting from *every* RBPd in the state space, and check whether all chains arrive at the same final state. We say that *false coalescence* occurs when $X(0) = Y(0)$, but for some intermediate chain Z started from a different RBPd, $Z(0) \neq X(0)$. The simulation results (Table 4) confirm that this is not a rare event: at $n = 5$, over 18% of CFTP terminations are false.

To confirm that this sampling algorithm produces a biased output distribution, we compare the naive CFTP output at $n = 4$ to the target distribution on permutations. A uniform distribution on RBPDs induces a distribution on permutations $w \in S_n$ with probability proportional to

TABLE 4. False coalescence rates. Each trial runs backward CFTP until the extremal chains coalesce, then replays from all starting states. A failure means the extremal chains agreed but at least one intermediate chain did not.

n	RBPDs	Trials	False coalescences	Rate
3	7	1,000	0	0.0%
4	41	500,000	37,476	7.5%
5	393	50,000	9,361	18.7%

$\Upsilon_w = \mathfrak{S}_w(1^n)$ (the number of reduced BPDs for w). Over 500,000 naive CFTP samples, the permutation frequencies deviate significantly from this target: a Pearson χ^2 test yields $\chi^2 = 60.7$ with 23 degrees of freedom ($p \approx 3 \times 10^{-5}$). Table 5 shows the per-permutation breakdown, grouped by Υ_w . The per-permutation relative deviations are around 2%, but the bias is statistically significant.

TABLE 5. Naive CFTP output distribution on permutations at $n = 4$ (500,000 trials). Permutations are grouped by Υ_w . Under an unbiased sampler, the ratios of observed/expected probability should be much closer to 1.00. Expected counts are rounded from $500,000 \cdot \Upsilon_w / 41$.

Υ_w	w	Expected	Observed	Obs/Exp	$(O-E)^2/E$
5	(1, 4, 3, 2)	60,976	61,607	1.010	6.5
3	(1, 2, 4, 3)	36,585	36,093	0.987	6.6
	(1, 3, 4, 2)	36,585	36,669	1.002	0.2
	(1, 4, 2, 3)	36,585	36,848	1.007	1.9
	(2, 1, 4, 3)	36,585	36,545	0.999	0.0
2	(1, 3, 2, 4)	24,390	23,924	0.981	8.9
	(2, 4, 1, 3)	24,390	24,489	1.004	0.4
	(2, 4, 3, 1)	24,390	24,536	1.006	0.9
	(3, 1, 4, 2)	24,390	24,669	1.011	3.2
	(4, 1, 3, 2)	24,390	24,546	1.006	1.0
1	14 permutations (expected 12,195 each):				
	max obs/exp	12,195	12,447	1.021	5.2
	min obs/exp	12,195	11,905	0.976	6.9

We therefore abandon exact sampling via CFTP and instead, in Section 5 below, develop a Markov chain Monte Carlo (MCMC) sampler for uniformly random RBPDs.

5. SAMPLING: MCMC FOR REDUCED BUMPLESS PIPE DREAMS

We turn to a Markov Chain Monte Carlo (MCMC) approach for sampling uniformly random RBPDs: a random walk on the state space of RBPDs whose stationary distribution is the uniform measure. In Section 5.1, we further discuss Conjecture 4.2 (connectivity of the state space under the 2×2 flips that preserve reducedness), and present examples of “traps” that require non-monotone paths to escape. In Section 5.2, we introduce droops and undroops — rectangular

moves from [LLS21] that bypass the traps — and prove that the combined move set makes the state space connected, establishing the correctness of the MCMC sampler.

5.1. Traps for local flips. We verified Conjecture 4.2 for $n \leq 8$ by direct exploration of the state space. This exploration also revealed the presence of “traps” in the state space that require non-monotone paths to escape.

Definition 5.1. Recall the 2×2 flips in Figure 6. Call a flip *up* if it *increases* the height function at the flip site by one (bringing the RBPd closer to the maximum $\mathbf{b}_{\text{id}}$), and *down* otherwise (bringing it closer to the minimum $\mathbf{b}_{w_0}$). In Figure 5, up flips are replacing the left local configuration with the right one, and down flips are the reverse.

One might hope that any RBPd can be transformed into $\mathbf{b}_{\text{id}}$ by a sequence of up flips, but this is not the case. We call a RBPd $\mathbf{b} \neq \mathbf{b}_{\text{id}}$ *stuck* if it does not admit any up flips (that preserve reducedness).

For $n = 8$, we found exactly ten such stuck RBPDs, all exhibiting the same mechanism, the *mutual frustration* of two j-elbows (Figure 9, left). Each of the two j-elbows is locked by a triple crossing which prevents it from being flipped up. Moreover, every box-cross annihilation near one of the locked j-elbows would create a double crossing, which is forbidden. The state space remains connected via non-monotone paths: escaping the trap requires down flips. Another example, with three j-elbows at $n = 9$, is shown in Figure 9, right.




FIGURE 9. Examples of stuck RBPDs for $n = 8$ (left) and $n = 9$ (right).

5.2. Droops, undroops, and connectivity. The traps of Section 5.1 show that establishing connectivity of the RBPd graph under the 2×2 flips alone (Conjecture 4.2) appears difficult. While we still believe this conjecture is true, we establish a weaker connectivity result by supplementing the 2×2 flips with *droops* and *undroops*, rectangular moves introduced in [LLS21]. We show that the Markov chain with this larger move set is ergodic and preserves the uniform distribution on the set of all RBPDs, thus providing a valid MCMC sampler.

Definition 5.2 (Droop and undroop [LLS21]). Let $R = [i_1, i_2] \times [j_1, j_2]$ be a rectangle in an $n \times n$ BPD with $i_1 < i_2$ and $j_1 < j_2$. Call R *droopable* if the NW corner (i_1, j_2) is an r-elbow, the SE corner (i_2, j_1) is empty, and R has no elbows except possibly at the four corners (i.e., every non-corner tile of R is empty, cross, vertical, or horizontal).

The *droop* at R replaces tiles on the boundary of R as follows: the NW corner becomes empty, the SE corner becomes a j-elbow, and the NE and SW corners both become r-elbows. On the four borders (excluding corners): the north and west borders *retract* (horizontal $\rightarrow$ empty, cross $\rightarrow$ vertical on north; vertical $\rightarrow$ empty, cross $\rightarrow$ horizontal on west), while the south and east borders *extend* (empty $\rightarrow$ horizontal, vertical $\rightarrow$ cross on south; empty $\rightarrow$ vertical, horizontal $\rightarrow$ cross on east). Interior tiles are unchanged.

The *undroop* is the inverse: if the NW corner is empty, the SE corner is a j-elbow, the NE and SW corners are r-elbows, and there are no interior elbows, then the undroop restores the r-elbow at NW and the empty tile at SE. A 3×3 droop is shown in Figure 10.




FIGURE 10. A 3×3 droop.

Droops and undroops preserve reducedness and the boundary permutation w of an RBPB. Moreover, by [LLS21, Proposition 5.3], for any $w \in S_n$, every w -RBPB can be obtained from the Rothe RBPB $\mathbf{b}_w$ (see Section 2.1 for the definition) by a sequence of droops.

To connect RBPBs with *different* boundary permutations, we use the 2×2 flips:

Lemma 5.3. *For $w \neq \text{id}$, the Rothe RBPB $\mathbf{b}_w$ admits a 2×2 flip of type (b) (Figure 6) that produces an RBPB with boundary permutation w' satisfying $\ell(w') = \ell(w) - 1$.*

Proof. The Rothe RBPB $\mathbf{b}_w$ has no j-elbows: its empty tiles are exactly the cells of the Rothe diagram $D(w)$, and all other non-cross tiles are r-elbows, vertical, or horizontal. Since $w \neq \text{id}$, the diagram $D(w)$ is nonempty. Choose a cross in position (i, j) , in the leftmost column where crosses occur (so j is minimal), and such that the cell $(i-1, j-1)$ is empty. (For instance, the topmost cross in this leftmost column works.) The vertical pipe of this cross comes from column $b = j$ on the bottom (south) edge and exits at row $w(b)$ on the right (east) edge. The horizontal pipe of the cross comes from column $a = w^{-1}(i)$ on the bottom (south) edge and exits at row $i = w(a)$ on the right (east). Our labels chosen such that $a < b$ and (since the pipes cross) $w(a) > w(b)$. From our choice of cross, and the fact that there are no j-elbows, for any c with $a < c < b$, the pipe beginning at column c on the south edge must exit at row $w(c) > w(a) > w(b)$. It follows that $\ell(w \cdot (a, b)) = \ell(w) - 1$. See Figure 11.

The 2×2 configuration bounding around our chosen cross at (i, j) and empty tile at $(i-1, j-1)$ looks like one of the following configurations:



(Highlighted in red in Figure 11.) These are exactly the configurations before flips in Figure 6, namely, in the lower row in type (b). The corresponding type (b) flip annihilates the chosen cross at (i, j) , and the resulting permutation is $w \cdot (a, b)$, which is of length one less than w , as observed above. $\square$




FIGURE 11. Schematic view of a cross in the Rothe BPD

Theorem 5.4 (Connectivity). *For every n , the set of all $n \times n$ RBPDs is connected under 2×2 flips (Figure 6) that preserve reducedness, together with droops and undroops.*

Proof. By [LLS21, Proposition 5.3], any RBPD with boundary permutation w can be undrooped to the Rothe RBPD $\mathbf{b}_w$, and by Lemma 5.3, a type-(b) flip then produces an RBPD whose boundary permutation has length $\ell(w) - 1$. This new RBPD can again be undrooped to its Rothe RBPD, and the length can be reduced again, until reaching the maximal RBPD $\mathbf{b}_{\text{id}}$. Since all moves are reversible, any two RBPDs can be connected by a sequence of 2×2 flips and droops/undroops. $\square$

Theorem 5.4 is weaker than Conjecture 4.2, which asserts connectivity under 2×2 flips alone.

Corollary 5.5. *Consider a random walk on $n \times n$ RBPDs that at each step attempts one of the $2(n-1)^2$ possible 2×2 flips (rejecting those that break reducedness) or one of the $\binom{n}{2}^2$ possible droop/undroop rectangles (the count corresponds to choosing row and column coordinates for the rectangle), such that a move and its inverse are attempted with equal probability (but this probability may depend on the move). Then the walk is ergodic on the space of all RBPDs, and its stationary distribution is uniform.*

Proof. Connectivity follows from Theorem 5.4, and aperiodicity from the positive rejection probability. For detailed balance with respect to the uniform distribution, note that each 2×2 flip is self-inverse, and each droop at a rectangle R is paired with the undroop at the same R , so the symmetry condition ensures equal transition rates in both directions. $\square$

Remark 5.6. As discussed in Section 4, already the 2×2 flips alone preclude a valid CFTP scheme. Adding droops and undroops to the move set only compounds the difficulty. Thus, we do not pursue exact sampling via CFTP, and instead rely on the approximate MCMC sampler based on Corollary 5.5.

5.3. MCMC parameters and simulation results. The move set of the MCMC sampler developed in Section 5.2 above consists of:

- **Local 2×2 flips** (Section 4.2): there are $(n-1)^2$ possible 2×2 windows in an $n \times n$ grid, each with two directions (flip up or down), giving $2(n-1)^2$ flip moves in total.

- **Droops and undroops** (Definition 5.2): there are $\binom{n}{2}^2$ candidate rectangles in total. Depending on the tile types at the corners of a rectangle, each rectangle can lead to either a droop, an undroop, or a non-move (when the rectangle is not droopable or undroopable).

To run the chain, we must specify a *proposal rule* that determines how the chain selects a candidate move at each step. There is considerable freedom in this choice (probabilities of each move type, distribution over rectangle sizes, etc.); by Corollary 5.5, any symmetric rule with full support preserves the uniform stationary distribution, and different rules affect only the mixing rate. We use the following:

Definition 5.7 (MCMC proposal rule). At each step, with probability $\frac{3}{4}$ we attempt a 2×2 flip (choosing a uniformly random 2×2 window and a random direction, up or down, as in Section 4.2). This rule is symmetric in the sense that up and down flips are attempted with equal probability.

With the complementary probability $\frac{1}{4}$ we attempt a droop or undroop. We pick a rectangle R from the $\binom{n}{2}^2$ candidates (the specific distribution over rectangle sizes is discussed below) and check the corner tiles: if the northwest corner is an r-elbow and the southeast corner is empty, we perform a droop at R ; if the northwest corner is empty and the southeast corner is a j-elbow, we perform an undroop; otherwise, the move is rejected. The rule is symmetric: each rectangle is selected with the same probability regardless of whether it leads to a droop or undroop.

Reducedness of the BPD is maintained at each step by tracking the pipe colors throughout the $n \times n$ grid (see Figure 2, bottom, for an example) and incrementally updating the boundary permutation and its inversion count. A BPD is reduced if and only if the number of cross tiles equals the number of inversions of the boundary permutation. For a 2×2 flip, the changed tiles can redirect pipes passing through the 2×2 block; we trace each affected pipe to the grid boundary in $O(n)$ time, updating edge colors along the way. The flip is accepted if reducedness is preserved. Droops and undroops automatically preserve reducedness (they do not change the boundary permutation or the cross count), so no rejection is needed. After an accepted droop/undroop, we update edge colors incrementally within the affected rectangle, and propagate changes outward, in time $O(k^2 + kn)$, where k is longest side length of the rectangle.

Let us now discuss the distribution over rectangle sizes for droop proposals. We pick the southeast corner (i, j) of the rectangle uniformly at random from $\{2, \dots, n\}^2$, and then pick the offsets $\delta_i \in \{1, \dots, i-1\}$, $\delta_j \in \{1, \dots, j-1\}$ to determine the northwest corner $(i-\delta_i, j-\delta_j)$ of the rectangle (it must be inside the $n \times n$ grid). For exploratory diagnostics, we tested the following distributions for the offsets (here, $\delta_{\max}$ stands for $i-1$ or $j-1$, depending on the coordinate):

1. *Geometric*: $\mathbb{P}(\delta = k) \propto 2^{-k}$, $1 \leq k \leq \delta_{\max}$.
2. *Uniform*: $\mathbb{P}(\delta = k)$ constant on $\{1, \dots, \delta_{\max}\}$.
3. *Log-uniform (reciprocal)*: $\mathbb{P}(\delta = k) \propto 1/k$, $1 \leq k \leq \delta_{\max}$.
4. *Reverse log-uniform (reverse reciprocal)*: $\mathbb{P}(\delta = k) \propto 1/(\delta_{\max} + 1 - k)$, $1 \leq k \leq \delta_{\max}$.

The diagnostic experiments are performed for $n = 60$. Each experiment consists of initial *burn-in* steps, followed by *collection* steps, where we record samples at regular intervals (thinning). We started from the two extreme pipe dreams, $\mathbf{b}_{\text{id}}$ and $\mathbf{b}_{w_0}$ (cf. Figure 5). Figure 12 shows the plot of $\ell(w)$ over 10^9 steps for each of the four rectangle size distributions, starting from both states. The dashed line marks $\ell \approx 0.396 \binom{n}{2}$, the expected length of the optimal layered permutation (Lemma 3.2). For all four distributions of the rectangle sizes, the chain started from $\mathbf{b}_{w_0}$ (red) reaches this stationary range quickly, with no apparent difference between the distributions. The chain starting from $\mathbf{b}_{\text{id}}$ (blue) rises much more slowly, and systematically stays below the




FIGURE 12. Plots of $\ell(w)$ (number of crosses in a BPD, equivalently, length of the boundary permutation) for the four rectangle size distributions at $n = 60$, starting from $\mathbf{b}_{\text{id}}$ (blue; starting from $\ell = 0$) and $\mathbf{b}_{w_0}$ (red; starting from $\ell = \binom{n}{2} = 1770$).

stationary value, even after 10^9 steps. For this initial condition, the average permutation matrix and an example of an RBPD (starting from $\mathbf{b}_{\text{id}}$) display characteristic “stuck” patterns which seem to prevent faster mixing, see Figure 13 for an illustration. We conclude that **starting from $\mathbf{b}_{w_0}$ mixes more reliably**.

Let us discuss the choice of the rectangle-size distribution for the *collection* phase (after burn-in). We compared the four distributions (geometric, uniform, log-uniform, and reverse log-uniform) by measuring the lag-1 autocorrelation $\text{Corr}(\ell(w_t), \ell(w_{t+1}))$ of the $\ell(w)$ statistic. The geometric distribution (with parameter $1/2$, as above) achieves the lowest autocorrelation. We attribute this to its acceptance rate: geometric droops (small rectangles preferred) are accepted 100 to 1000 times more often than for other distributions, which apparently leads to faster decorrelation between successive samples.

For the simulations with $n = 100$, we use the geometric rectangle-size distribution for droop proposals, both in burn-in and collection phases. We verified (by checking the equilibration of the $\ell(w)$ statistic) that a burn-in of 10^{10} steps is sufficient for $n = 100$. The $B = 10,000$ samples are then collected with thinning by 5×10^8 steps. We used 40 CPU cores and ran 40 Markov chains in parallel on the Rivanna HPC cluster at the University of Virginia, each collecting 250




FIGURE 13. A “stuck” chain started from $\mathbf{b}_{\text{id}}$ at $n = 60$ after 10^9 steps. Left: average permutation matrix (5.1) over 500 samples (thinning 10^6 steps). Right: a single RBPd with colored pipes. Two values ($\sigma(44) = 58$ and $\sigma(54) = 59$) are frozen across all 500 samples (dark dots), visible in the RBPd as vertical gaps of horizontal tiles in the bottom-right corner. These samples were done with uniform droop proposals, but the same phenomenon appears with other distributions.

samples after an independent burn-in. The total runtime for the burn-in and collection phases was about 3 hours.

Each sample records the boundary permutation w , from which we accumulate the *average permutation matrix*

$$(M_{ij})_{i,j=1}^n, \quad M_{ij} := \frac{1}{B} \sum_{t=1}^B \mathbf{1}_{w_t(i)=j}, \quad (5.1)$$

see Figure 14 for the result. The matrix M serves as a histogram approximating the hypothetical limiting *permuton* [HKM⁺13], [Gru24] of uniform random RBPds: if a limit shape exists as $n \rightarrow \infty$, the rescaled matrix M converges to the density of that permuton.

We also record the height function h_t for each sample, and compute the average height function

$$\bar{h}(x, y) := \frac{1}{B} \sum_{t=1}^B h_t(x, y), \quad 0 \leq x, y \leq n. \quad (5.2)$$

From the height functions, we extract two quantities. First, the discrete mixed derivative $\Delta_x \Delta_y \bar{h}$ outlines the “*liquid region*” (called by analogy with [KOS06]), the part of the grid where the average height function is not linear (where all six tile types coexist), see Figure 15, left. The complement to the liquid region consists of *frozen regions* consisting entirely of empty, cross, or vertical/horizontal tiles. Second, we record the height function fluctuations $h_t - \bar{h}$, which are nontrivial inside the liquid region and vanish in the frozen regions, see Figure 16.

6. CONJECTURES AND OPEN PROBLEMS

6.1. Schubert measure on permutations and the limiting permuton. The uniform measure on RBPds of size n weights each permutation $w \in S_n$ proportionally to its number of reduced




FIGURE 14. Average permutation matrix (5.1) at $n = 100$ from 10,000 MCMC samples. The support is contained inside a cone, with a singular (delta-measure) component visible along the southeast boundary curve.




FIGURE 15. Left: the discrete mixed derivative $\Delta_x \Delta_y \bar{h}$. Right: plot of the difference of $\Delta_x \Delta_y \bar{h}$ and the average permutation matrix M (normalized as probability distributions) showing that the southeast boundary curves of M and of the liquid region coincide (see Conjecture 6.2 and Proposition 6.3).

bumpless pipe dreams, which equals the principal specialization Υ_w . We may call the resulting probability measure on S_n ,

$$\mathbb{P}(w) = \frac{\Upsilon_w}{\sum_{v \in S_n} \Upsilon_v}, \quad (6.1)$$

the *Schubert measure* on permutations.

A measure on permutations similar to (6.1) was studied in [MPPY25], when the Schubert polynomials are replaced by Grothendieck polynomials with the Grothendieck parameter specialized




FIGURE 16. A single MCMC sample at $n = 100$. Left: height function fluctuation $h_t - \bar{h}$; red indicates $h_t > \bar{h}$, blue indicates $h_t < \bar{h}$. Right: a random reduced bumpless pipe dream. One can see that the region adjacent to the southeast corner is saturated with cross tiles.

at $\beta = 1$. As $n \rightarrow \infty$, these Grothendieck random permutations converge [MPY25, Theorem 1.5] to a deterministic permuton supported inside a cone, with a singular component concentrated on a quarter circle forming the southeast boundary of the support. We refer to [HKM⁺13], [Gru24] for background on the permuton convergence (which, in particular, is equivalent to the convergence of all finite pattern densities).

The average permutation matrix we observe for the Schubert measure is visually similar (Figure 14), motivating us to make the following conjecture:

Conjecture 6.1 (Convergence to the Schubert permuton). *The random permutations (6.1) converge, as $n \rightarrow \infty$, to a deterministic permuton μ on $[0, 1]^2$ (which we call the Schubert permuton) supported inside a cone and having a singular component along the southeast boundary curve.*

For the Schubert permuton, the boundary curve does not appear to be a quarter circle.

6.2. Limit shape and arctic curve. The discrete mixed derivative $\Delta_x \Delta_y \bar{h}$ (Figure 15, left) and fluctuations of $h_t - \bar{h}$ (Figure 16, left) detect the *liquid region* of the uniformly random RBPB, where all six tile types coexist with positive density. The liquid region is symmetric about the main diagonal (reflecting the $w \mapsto w^{-1}$ symmetry of the model), but does *not* appear to have quarter-turn symmetry, in contrast with the Grothendieck case [MPY25] (equivalent to domino tilings of the Aztec diamond [JPS98], [CEP96]) and the uniformly random six-vertex model with domain wall boundary conditions [CP10b], [Agg20], whose limit shapes are invariant under quarter-turn rotation of the square.

Conjecture 6.2 (Limit shape and arctic curve). *As $n \rightarrow \infty$, the height function of a uniformly random RBPB of size n converges to a deterministic limit shape, consisting of four frozen regions adjacent to the corners of the domain (where the height function is linear) and a liquid region (where it is curved), separated by a deterministic arctic curve. The southeast arc of the*

RBPD arctic curve coincides with the support of the singular component of the Schubert permuton (Conjecture 6.1).

We prove the following coincidence of the southeast component of the arctic curve and the singular component of the permuton, assuming that the frozen region in the RBPDs exists with an exponential rate of convergence:

Proposition 6.3. *Let $R_n \subset \{1, \dots, n\}^2$ be a SE-justified region in the BPD grid, that is, if $(i, j) \in R_n$, then $(i', j') \in R_n$ for all $i' \geq i$ and $j' \geq j$.³ Furthermore, assume that R_n does not touch the north or west boundaries of the grid. Suppose that the RBPDs exhibit the frozen region: for some constant $c > 0$ not depending on n , every tile $(i, j) \in R_n$ of a uniformly random RBPD of size n is a cross with probability at least $1 - e^{-cn}$. Then there exists a constant $c' > 0$ such that with probability $1 - e^{-c'n}$ as $n \rightarrow \infty$, the region R_n in the permutation matrix of a Schubert random permutation w of size n does not contain any points, that is, $(w(k), k) \notin R_n$ for all $k \in \{1, \dots, n\}$.*

Proof. Condition on the event that every tile in R_n is a cross (by a union bound over the $O(n^2)$ tiles in R_n , this has probability $1 - e^{-c'n}$ for some $c'' > 0$). Recall that the pipe k enters the grid from the south at column k and exits on the east boundary at row $w(k)$.

Assume first that $(n, k) \in R_n$. Then the cells of R_n in column k form a contiguous block $\{i_k^*, i_k^* + 1, \dots, n\}$ for some i_k^* . Pipe k enters from the south at (n, k) and passes straight up through R_n , exiting column k at row $\leq i_k^* - 1$. Since after that the pipe must continue traveling only north and east, its exit row $w(k)$ must satisfy $w(k) \leq i_k^* - 1 < i_k^*$, so $(w(k), k) \notin R_n$.

It remains to consider the case that $(n, k) \notin R_n$. Then $(i, k) \notin R_n$ for all i . So $(w(k), k) \notin R_n$ regardless of $w(k)$. This completes the proof. $\square$

Heuristically, in Figures 14 to 16, we see that some mass from the southeast frozen region in the RBPD grid is concentrated along the southeast boundary curve in the permuton, but Proposition 6.3 does not rule out the possibility that the permuton may be absolutely continuous with respect to the Lebesgue measure in $[0, 1]^2$.

The height function fluctuations $h_t - \bar{h}$ are nontrivial inside the liquid region and vanish in the frozen regions (Figure 16). For dimer models, such fluctuations converge to the Gaussian Free Field (GFF) [Ken01], [Pet15], [BG18], [CLR21] (see [She07] for the standard background on the GFF). For the non-free-fermion six-vertex model with domain wall boundary conditions, the fluctuation structure is far less understood.⁴

Question 6.4. *Do the height function fluctuations of uniformly random RBPDs (see Figure 16) converge to the GFF inside the liquid region?*

In support of the GFF behavior, let us note that our MCMC samples at $n = 100$ exhibit fluctuations of order $\sqrt{\log n}$: the maximum of $|h_t - \bar{h}|$ is approximately 2.2–2.9 across independent samples, consistent with $\sqrt{\log 100} \approx 2.15$.

6.3. Merzon–Smirnov conjecture. The Merzon–Smirnov conjecture (Conjecture 1.1) holds for $n \leq 13$ by exhaustive search (Proposition 3.6), extending prior verification for $n \leq 10$ in [MS16]. At $n = 17$ the conjecture fails (Theorem 1.2). The counterexample w^* (1.3) differs from the optimal layered permutation $w(1, 2, 4, 10)$ by a single adjacent transposition. We computed Υ_w for all permutations within Cayley distance 4 of the optimal layered permutation for each

³The grid coordinates are as in Figure 4, with $(1, 1)$ in the northwest corner and (n, n) in the southeast corner.

⁴Even the full limit shape picture is not fully developed, cf. [CP10b], [CP10a], [dGKW21], [Agg20].

$n \leq 16$, and found no counterexamples; at $n = 17$, the counterexample w^* at Cayley distance 1 is the unique immediate neighbor of $w(1, 2, 4, 10)$ that exceeds the layered maximum. We also found counterexamples at $n = 18, 19$, and 20 : in each case, the Cayley distance 1 neighborhood of the optimal layered permutation contains a unique permutation exceeding the layered maximum, obtained by swapping the last entry of the penultimate block with the first entry of the last block. These are s_7 applied to $w(1, 2, 4, 11)$ at $n = 18$, s_8 applied to $w(1, 2, 5, 11)$ at $n = 19$, and s_8 applied to $w(1, 2, 5, 12)$ at $n = 20$, exceeding the layered maximum by about 5%, 9%, and 8%, respectively. The $n = 20$ computation required a 40-core cluster node (University of Virginia's Rivanna HPC) due to the size of the BFS frontiers.⁵ However, applying the same swap to the optimal layered permutation for $n \leq 16$ always *decreases* Υ_w , so the pattern does not extend downward.

Searching further from the optimal layered permutation reveals even larger values. The permutation u^* (1.4), obtained from w^* by additionally transposing the entries in positions 4 and 17, exceeds the layered maximum by about 16% (compared to the 7% excess in Theorem 1.2). The analogous permutation

$$(1, 3, 2, 8, 6, 5, 18, 4, 17, 16, 15, 14, 13, 12, 11, 10, 9, 7) \in S_{18} \quad (6.2)$$

exceeds the layered maximum by about 12%. We illustrate the maximal layered and the newly discovered permutation in Figure 17.




FIGURE 17. Permutation matrices for the $n = 18$ examples: the optimal layered permutation $w(1, 2, 4, 11)$ (left), and two permutations exceeding the layered maximum.

Remark 6.5. It now becomes an even more challenging problem to characterize the permutations maximizing Υ_w for general n : the newly found examples show that the characteristics would likely not be structural (e.g. pattern avoidance, block structure etc).

6.4. Bounds and asymptotics. Despite the failure of Conjecture 1.1, we expect that the asymptotic growth rate in Stanley's problem (1.2) is determined by layered permutations:

⁵While the benchmarks for layered permutations (Section 3.1) show the transition formula beating cotransition for $n = 16, 17$, for larger n the caching required for this speedup becomes infeasible, so we default to cotransition and BFS rather than DFS; see Remark 3.4.

Conjecture 6.6. *The maximum of Υ_w over layered permutations has the same exponential growth rate as the maximum over the whole S_n , that is,*

$$\lim_{n \rightarrow \infty} \frac{1}{n^2} \log_2 \max_{w \in S_n} \Upsilon_w = \lim_{n \rightarrow \infty} \frac{1}{n^2} \log_2 \max_{w \text{ layered}} \Upsilon_w \approx 0.29,$$

where the layered limit was computed in [MPP19]. In other words, any potential improvement over layered permutations is subexponential in n^2 .

In support of this we observe the asymptotic behavior of nearby permutations.

Proposition 6.7 (Bounds). *Suppose that $u = w \cdot s_i$ or $u = s_i \cdot w$ with $\ell(u) = \ell(w) + 1$. Then*

$$\frac{1}{n-i} \leq \frac{\Upsilon_u}{\Upsilon_w} \leq i. \quad (6.3)$$

Proof. Suppose that $u = w \cdot s_i$. The case $u = s_i \cdot w$ follows from it by observing that $u^{-1} = w^{-1} s_i$ and that $\Upsilon_u = \Upsilon_{u^{-1}}$ for every permutation.

For the lower bound, we use the divided difference formula, $\mathfrak{S}_w(x_1, \dots, x_n) = \partial_i \mathfrak{S}_u(x_1, \dots, x_n)$. (See Appendix A.) Every monomial of $\mathfrak{S}_u$ is of the form $x_1^{a_1} \cdots x_{n-1}^{a_{n-1}}$ with $a_i \leq n-i$ (as can be easily seen from the PD interpretation). We have

$$\partial_i \mathbf{x}^{\mathbf{a}} = \begin{cases} \sum_{j=0}^{a_i - a_{i+1} - 1} x_i^{a_i - j - 1} x_{i+1}^{a_{i+1} + j} \prod_{k \neq i, i+1} x_k^{a_k}, & \text{for } a_i > a_{i+1}, \\ - \sum_{j=0}^{a_{i+1} - a_i - 1} x_i^{a_i - j - 1} x_{i+1}^{a_{i+1} + j} \prod_{k \neq i, i+1} x_k^{a_k}, & \text{for } a_i < a_{i+1}, \\ 0, & \text{for } a_i = a_{i+1}. \end{cases}$$

Setting all x variables equal to 1, we find that the specialization of $\partial_i \mathbf{x}^{\mathbf{a}}$ is bounded above by $|a_i - a_{i+1}| \leq n - i$. Applying this bound to every monomial in $\mathfrak{S}_u$, we obtain

$$\Upsilon_w \leq (n - i) \Upsilon_u,$$

which gives the desired lower bound.

To establish the upper bound, we use Monk's formula for multiplying Schubert polynomials; see Equation (A.2). Evaluating at 1, this says

$$i \Upsilon_w = \sum_{r \leq i < s} \Upsilon_{w \cdot (r, s)},$$

the sum over transpositions (r, s) such that $\ell(w \cdot (r, s)) = \ell(w) + 1$. Evidently from the assumption that $u = w \cdot (i, i+1)$ and $\ell(u) = \ell(w) + 1$, u is among the permutations appearing on the RHS. This gives

$$i \Upsilon_w \geq \Upsilon_u$$

as claimed. $\square$

Corollary 6.8. *Suppose that two permutations u and w differ by at most a linear number of adjacent transpositions, i.e. $u = \sigma w \pi$, where $\ell(\sigma), \ell(\pi) \leq an$ for some constant a . Then*

$$2^{-2an \log(n)} \leq \frac{\Upsilon_u}{\Upsilon_w} \leq 2^{2an \log(n)}$$

Proof. Follows by presenting σ and π as products of at most an simple transpositions each, and repeatedly applying Proposition 6.7. $\square$

Remark 6.9. In particular, Corollary 6.8 implies that if a permutation w is close to maximal, i.e. $\Upsilon_w = 2^{cn^2+o(n^2)}$, then a permutation u differing from w by a linear number of transpositions will have the same leading term asymptotics. Thus, the counterexamples found in Theorem 1.2 do not contradict the conjectured leading order asymptotics.

To characterize the maximal permutations, we discuss some consequences of the descent formula (Theorem 2.1).

Proposition 6.10. *Suppose that $w \in S_n$ is such that Υ_w is maximal among all specializations for permutations in S_n . Then $\text{maj}(w) \geq \ell(w)$, where $\text{maj}(w) := \sum_{i \in \text{Des}(w)} i$ is the major index of the permutation w .*

Proof. We have that $\Upsilon_{w \cdot s_i} \leq \Upsilon_w$, so (2.1) gives

$$\Upsilon_w = \sum_{i \in \text{Des}(w)} \frac{i}{\ell(w)} \cdot \Upsilon_{w \cdot s_i} \leq \sum_{i \in \text{Des}(w)} \frac{i}{\ell(w)} \cdot \Upsilon_w.$$

After cancellation, we get $\ell(w) \leq \sum_{i \in \text{Des}(w)} i$, as desired. $\square$

Remark 6.11. One can directly check that Proposition 6.10 holds for all layered permutations.

6.5. Connectivity and exact sampling. We bypassed the issue of connectivity of RBPDs under local flips in 2×2 windows by adding droop moves to the MCMC sampler, but the question of connectivity under local flips alone remains open (Conjecture 4.2). Neither 2×2 flips nor droops/undroops seem to be appropriate ingredients for an exact sampling algorithm with CFTP [PW98], and so the question of exact sampling of uniform RBPDs also remains open. We remark that sometimes exact sampling can be achieved by a non-local Markov chain utilizing algebraic/combinatorial structure of the model, for example, in shuffling algorithms for domino [Pro03] and lozenge [BGR10] tilings.

6.6. Mixing time. We have no rigorous mixing time bounds for the MCMC chain of Section 5. Our experiments (Figure 12) suggest that with sufficiently large rectangular moves (such as the reverse log-uniform distribution on rectangle sizes), the chain mixes from all starting states in polynomial time. On the other hand, the chain restricted to 2×2 flips alone appears to mix exponentially slowly (due to “stuck” states as in Figure 13). For comparison, Glauber dynamics for the six-vertex model (which performs the same 2×2 flips as in Section 4.2, but with probabilities depending on the six-vertex weights) mixes exponentially slowly for certain parameter regimes [FR19]. Establishing a polynomial/exponential distinction rigorously for the RBPD chain is an important open problem.

6.7. Other specializations and measures. One can extend our methods beyond the principal specialization $\Upsilon_w = \mathfrak{S}_w(1^n)$ in several directions. The q -specialization $\mathfrak{S}_w(1, q, q^2, \dots)$ refines Υ_w , and one can readily assign a q -weight to each RBPD, and perform a similar MCMC sampling investigation. The double Schubert polynomials $\mathfrak{S}_w(x; y)$ (Appendix A) also enjoy transition and cotransition formulas [Knu19], [LS85b]. For them, one can adapt our algorithms described in Section 2 and compute $\mathfrak{S}_w(x; y)$ at various specialization points. The Grothendieck polynomials $\mathfrak{S}_w^{(\beta)}(x)$ deform the Schubert polynomials by adding a parameter β (setting $\beta = 0$ recovers the Schubert case). For all $\beta > 0$, the bumpless pipe dreams are allowed to be non-reduced, but their enumeration involves the factors $\beta^{-\ell(w)}$. For $\beta = 1$, these factors disappear, and the bumpless pipe dream picture becomes equivalent to domino tilings of the Aztec diamond [MPPY25, Section 6].

The permuton picture was also developed for the $\beta = 1$ case in [MPPY25]. It would be interesting to explore MCMC sampling and principal specializations for all $0 < \beta < 1$.

Another natural relative of the Schubert measure is the *symmetrized product measure* $\mathbb{P}(w) \propto \Upsilon_w \cdot \Upsilon_{w_0 w}$. Its expected length can be computed exactly from the Cauchy identity for Schubert polynomials [Mac91]:

$$\sum_{w \in S_n} \mathfrak{S}_w(x) \mathfrak{S}_{w_0 w}(y) = \prod_{i+j \leq n} (x_i + y_j). \quad (6.4)$$

Setting $x_i = t$ and $y_j = 1$ and using $\mathfrak{S}_w(t^n) = t^{\ell(w)} \Upsilon_w$ gives $\sum_w t^{\ell(w)} \Upsilon_w \Upsilon_{w_0 w} = (1+t)^{\binom{n}{2}}$. Taking $\frac{d}{dt} \Big|_{t=1}$ of both sides, divided by the value at $t = 1$, yields $\mathbb{E}[\ell(w)] = \binom{n}{2}/2$. This contrasts with the Schubert measure (6.1), which (conjecturally) concentrates on permutations of length $\approx 0.396 \binom{n}{2}$, see Lemma 3.2.

6.8. Code availability. The code accompanying this paper is available at <https://github.com/lenis2000/schubert-computations-sampling>. The repository includes:

- `schubert.cpp` — descent, cotransition, and transition formulas for computing Υ_w , together with exact and heuristic max search.
- `bpd_mcmc.cpp` — MCMC sampler for reduced bumpless pipe dreams.
- `bpd_cftp_sampler.cpp` — backward-CFTP sampler used for the negative results of Section 4.
- Diagnostic and validation tools for CFTP failure analysis and RBPD connectivity.

APPENDIX A. BACKGROUND ON SCHUBERT POLYNOMIALS

We collect here the standard definitions underlying the recurrences used in Sections 2.2 and 2.4. Standard references include Macdonald [Mac91] and Manivel [Man01].

For $1 \leq i \leq n-1$, the *divided difference operator* ∂_i acts on polynomials $f \in \mathbb{Z}[x_1, \dots, x_n]$ by

$$\partial_i f := \frac{f - s_i \cdot f}{x_i - x_{i+1}}, \quad (\text{A.1})$$

where s_i acts on f by permuting x_i and x_{i+1} . These operators satisfy the nilCoxeter relations: $\partial_i^2 = 0$, $\partial_i \partial_j = \partial_j \partial_i$ for $|i - j| \geq 2$, and $\partial_i \partial_{i+1} \partial_i = \partial_{i+1} \partial_i \partial_{i+1}$.

Schubert polynomials are recursively determined by the following conditions:

- For the longest permutation $w_0 = (n, n-1, \dots, 2, 1)$, we have $\mathfrak{S}_{w_0} = x_1^{n-1} x_2^{n-2} \dots x_{n-1}$.
- For all $w \in S_n$ and $i = 1, \dots, n-1$ such that $\ell(ws_i) = \ell(w) + 1$, we have $\mathfrak{S}_w = \partial_i \mathfrak{S}_{ws_i}$.

This recursive definition was introduced by Bernstein, Gelfand, and Gelfand [BGG73] and Demazure [Dem74] in the context of Schubert classes in the cohomology of flag varieties, and made explicit by Lascoux and Schützenberger [LS82].

Lascoux and Schützenberger [LS85a] also introduced *double Schubert polynomials*

$$\mathfrak{S}_w(x_1, \dots, x_n; y_1, \dots, y_n),$$

depending on two sets of variables. They are defined by the same divided difference recursion (A.1), with ∂_i acting on the x -variables only, and the modified base case

$$\mathfrak{S}_{w_0}(x; y) = \prod_{i+j \leq n} (x_i - y_j).$$

The ordinary Schubert polynomial is recovered by setting $y = 0$: $\mathfrak{S}_w(x) = \mathfrak{S}_w(x; 0, \dots, 0)$. The transition formula of Lascoux-Schützenberger [LS85b] (Theorem 2.2, see also [KV97, Wei21]) and the cotransition formula of Knutson [Knu19] (Theorem 2.3) hold more generally for double

Schubert polynomials; our specialization $\Upsilon_w = \mathfrak{S}_w(1^n)$ corresponds to $x_i = 1$, $y_j = 0$ for all i, j . The transition and cotransition formulas follow quickly from *Monk's formula*, which describes how to multiply by a linear Schubert polynomial $\mathfrak{S}_{s_k}$. Restricting attention to single Schubert polynomials, the formula is

$$(x_1 + \cdots + x_k) \cdot \mathfrak{S}_w = \sum_{i \leq k < j} \mathfrak{S}_{w \cdot (i,j)}, \quad (\text{A.2})$$

the sum over transpositions (i, j) such that $\ell(w \cdot (i, j)) = \ell(w) + 1$. This is equivalent to

$$x_r \mathfrak{S}_w = - \sum_{i < r} \mathfrak{S}_{w \cdot (i,r)} + \sum_{j > r} \mathfrak{S}_{w \cdot (r,j)}, \quad (\text{A.3})$$

where both sums run over only those terms of length equal to $\ell(w) + 1$. Transition is the case where the second sum consists of exactly one term (move the first sum to the other side to obtain a positive formula); this can occur for several choices of r . Cotransition is the case where the first sum is empty; this only occurs for r as in Theorem 2.3.

Schubert polynomials are independent of the ambient symmetric group: if $\iota: S_n \hookrightarrow S_{n+1}$ denotes the natural inclusion (fixing $n+1$), then

$$\mathfrak{S}_w(x_1, \dots, x_n) = \mathfrak{S}_{\iota(w)}(x_1, \dots, x_{n+1}). \quad (\text{A.4})$$

In particular, $\mathfrak{S}_{\iota(w)}$ does not depend on x_{n+1} , so the principal specialization satisfies $\Upsilon_w = \mathfrak{S}_w(1^n) = \mathfrak{S}_{\iota(w)}(1^{n+1})$. This means Υ_w depends only on the permutation w belonging to the infinite symmetric group S_∞ (i.e., the inductive limit of the S_n 's), provided that the number of ones in the Schubert polynomial is large enough to accommodate w .

REFERENCES

- [Agg20] A. Aggarwal, *Arctic Boundaries of the Ice Model on Three-Bundle Domains*, Invent. Math. **220** (2020), 611–671. arXiv:1812.03847 [math.PR]. [↑25](#), [26](#)
- [BB05] A. Björner and F. Brenti, *Combinatorics of Coxeter Groups*, Grad. Texts Math., vol. 231, Springer, Berlin, Heidelberg, 2005. [↑9](#)
- [BB93] N. Bergeron and S. Billey, *RC-graphs and Schubert polynomials*, Experiment. Math. **2** (1993), no. 4, 257–269. MR1281474 [↑2](#)
- [BG18] A. Bufetov and V. Gorin, *Fluctuations of particle systems determined by Schur generating functions*, Adv. Math. **338** (2018), 702–781. arXiv:1604.01110 [math.PR]. [↑26](#)
- [BGG73] I.N. Bernstein, I.M. Gelfand, and S.I. Gelfand, *Schubert cells, and the cohomology of the spaces G/P* , Russian Math. Surveys **28** (1973), no. 3, 1–26. [↑30](#)
- [BGP25] S. C. Billey, Y. Gao, and B. Pawłowski, *Introduction to the Cohomology of the Flag Variety*, arXiv preprint (2025). arXiv:2506.21064 [math.CO]. [↑2](#)
- [BGR10] A. Borodin, V. Gorin, and E. Rains, *q -Distributions on boxed plane partitions*, Selecta Math. **16** (2010), no. 4, 731–789. arXiv:0905.0679 [math-ph]. [↑29](#)
- [BHY19] S. C. Billey, A. E. Holroyd, and B. Young, *A bijective proof of Macdonald's reduced word formula*, Algebr. Comb. **2** (2019), no. 2, 217–248. arXiv:1702.02936 [math.CO]. [↑7](#)
- [Bre99] D. M. Bressoud, *Proofs and Confirmations: The Story of the Alternating-Sign Matrix Conjecture*, Cambridge University Press, 1999. [↑2](#), [12](#)
- [CEP96] H. Cohn, N. Elkies, and J. Propp, *Local statistics for random domino tilings of the Aztec diamond*, Duke Math. J. **85** (1996), no. 1, 117–166. arXiv:math/0008243 [math.CO]. [↑25](#)
- [CLR21] D. Chelkak, B. Laslier, and M. Russkikh, *Bipartite dimer model: perfect t -embeddings and Lorentz-minimal surfaces*, arXiv preprint (2021). arXiv:2109.06272 [math-ph]. [↑26](#)
- [CP10a] F. Colomo and A. Pronko, *The arctic curve of the domain-wall six-vertex model*, J. Stat. Phys **138** (2010), no. 4-5, 662–700. arXiv:0907.1264 [math-ph]. [↑26](#)
- [CP10b] F. Colomo and A. G. Pronko, *The limit shape of large alternating sign matrices*, SIAM J. Discrete Math. **24** (2010), no. 4, 1558–1571. arXiv:0803.2697 [math-ph]. [↑25](#), [26](#)

- [Dem74] M. Demazure, *Désingularisation des variétés de Schubert généralisées*, Ann. Sci. École Norm. Supér. **7** (1974), no. 1, 53–88. [↑30](#)
- [dGKW21] J. de Gier, R. Kenyon, and S.S. Watson, *Limit shapes for the asymmetric five vertex model*, Commun. Math. Phys. **385** (2021), 793–836. arXiv:1812.11934 [math.PR]. [↑26](#)
- [FGS18] N.J.Y. Fan, P.L. Guo, and S.C.C. Sun, *Bumpless pipedreams, reduced word tableaux and Stanley symmetric functions*, arXiv preprint (2018). arXiv:1810.11916 [math.CO]. [↑8](#)
- [FK97] S. Fomin and A.N. Kirillov, *Reduced words and plane partitions*, J. Algebraic Combin. **6** (1997), no. 4, 311–319. MR1471891 [↑2](#)
- [FR19] M. Fahrback and D. Randall, *Slow Mixing of Glauber Dynamics for the Six-Vertex Model in the Ordered Phases*, Approx/random 2019, 2019, pp. 37:1–37:20. arXiv:1904.01495 [cs.DS]. [↑29](#)
- [FS94] S. Fomin and R. P. Stanley, *Schubert polynomials and the nilCoxeter algebra*, Adv. Math. **103** (1994), no. 2, 196–207. [↑7](#)
- [Gru24] R. Gruebel, *Ranks, copulas, and permutons*, Metrika **87** (2024), 155–182. arXiv:2206.12153 [math.ST]. [↑23](#), [25](#)
- [HKM⁺13] C. Hoppen, Y. Kohayakawa, C. G. Moreira, B. Ráth, and R. M. Sampaio, *Limits of permutation sequences*, J. Comb. Theory Ser. B **103** (2013), no. 1, 93–113. arXiv:1103.5844 [math.CO]. [↑23](#), [25](#)
- [JPS98] W. Jockusch, J. Propp, and P. Shor, *Random Domino Tilings and the Arctic Circle Theorem*, arXiv preprint (1998). arXiv:math/9801068 [math.CO]. [↑25](#)
- [Ken01] R. Kenyon, *Dominos and the Gaussian Free Field*, Ann. Probab. **29** (2001), no. 3, 1128–1137. arXiv:math-ph/0002027. [↑26](#)
- [Knu19] A. Knutson, *Schubert polynomials, pipe dreams, equivariant classes, and a co-transition formula*, arXiv preprint (2019). arXiv:1909.13777 [math.CO]. [↑8](#), [29](#), [30](#)
- [KOS06] R. Kenyon, A. Okounkov, and S. Sheffield, *Dimers and amoebae*, Ann. Math. **163** (2006), 1019–1056. arXiv:math-ph/0311005. [↑23](#)
- [Kup96] G. Kuperberg, *Another proof of the alternative-sign matrix conjecture*, Intern. Math. Research Notices **1996** (1996), no. 3, 139–150. [↑12](#)
- [KV97] A. Kohnert and S. Veigneau, *Using Schubert basis to compute with multivariate polynomials*, Adv. Appl. Math. **19** (1997), no. 1, 45–60. [↑30](#)
- [Las02] A. Lascoux, *Chern and Yang through ice*, 2002. Unpublished manuscript. [↑2](#)
- [LLS21] T. Lam, S. J. Lee, and M. Shimozono, *Back stable Schubert calculus*, Compos. Math. **157** (2021), no. 5, 883–962. arXiv:1806.11233 [math.CO]. [↑2](#), [18](#), [19](#), [20](#)
- [LS82] A. Lascoux and M.-P. Schützenberger, *Polynômes de Schubert*, C. R. Acad. Sci. Paris Sér. I Math. **294** (1982), 447–450. [↑30](#)
- [LS85a] A. Lascoux and M.-P. Schützenberger, *Interpolation de Newton à plusieurs variables*, Séminaire d’algèbre Paul Dubreil et Marie-Paule Malliavin, 36ème année (Paris, 1983–1984), 1985, pp. 161–175. [↑3](#), [30](#)
- [LS85b] A. Lascoux and M.-P. Schützenberger, *Schubert polynomials and the Littlewood-Richardson rule*, Lett. Math. Phys. **10** (1985), no. 2–3, 111–124. [↑8](#), [29](#), [30](#)
- [Mac91] I.G. Macdonald, *Notes on Schubert Polynomials*, Université du Québec à Montréal, 1991. [↑3](#), [7](#), [8](#), [30](#)
- [Man01] L. Manivel, *Symmetric functions, Schubert polynomials and degeneracy loci*, SMF/AMS Texts and Monographs, vol. 6, American Mathematical Society, Providence, RI; Société Mathématique de France, Paris, 2001. Translated from the 1998 French original by John R. Swallow, Cours Spécialisés, 3. [Specialized Courses]. MR1852463 [↑30](#)
- [MPP19] A.H. Morales, I. Pak, and G. Panova, *Asymptotics of principal evaluations of Schubert polynomials for layered permutations*, Proc. Amer. Math. Soc. **147** (2019), 1377–1389. arXiv:1805.04341 [math.CO]. [↑1](#), [2](#), [4](#), [5](#), [9](#), [10](#), [11](#), [12](#), [28](#)
- [MPPY25] A.H. Morales, G. Panova, L. Petrov, and D. Yeliussizov, *Grothendieck Shenanigans: Permutons from pipe dreams via integrable probability*, Adv. Math. **480** (2025), 110510. arXiv:2407.21653 [math.PR]. [↑1](#), [2](#), [4](#), [5](#), [6](#), [24](#), [25](#), [29](#), [30](#)
- [MPY22] C. Monical, B. Pankow, and A. Yong, *Reduced word enumeration, complexity, and randomization*, Electron. J. Combin. **29** (2022), no. 2, P2.46. [↑9](#)
- [MS16] G. Merzon and E. Smirnov, *Determinantal identities for flagged Schur and Schubert polynomials*, Eur. J. Math. **2** (2016), 227–245. arXiv:1410.6857 [math.CO]. [↑1](#), [2](#), [4](#), [12](#), [26](#)
- [OEI25] OEIS Foundation Inc., *The On-Line Encyclopedia of Integer Sequences*, 2025. Published electronically at <https://oeis.org>. [↑4](#), [12](#)

- [Pet15] L. Petrov, *Asymptotics of Uniformly Random Lozenge Tilings of Polygons*. *Gaussian Free Field*, Ann. Probab. **43** (2015), no. 1, 1–43, available at [1206.5123](#). arXiv:1206.5123 [math.PR]. [↑26](#)
- [Pro03] J. Propp, *Generalized domino-shuffling*, Theoretical Computer Science **303** (2003), no. 2-3, 267–301. arXiv:math/0111034 [math.CO]. [↑29](#)
- [PW98] J. Propp and D. Wilson, *Coupling from the past: a user’s guide*, Microsurveys in discrete probability (Princeton, NJ, 1997), 1998, pp. 181–192. [↑12](#), [14](#), [15](#), [16](#), [29](#)
- [RPBN⁺24] Bernardino Romera-Paredes, Mohammadamin Barekatali, Alexander Novikov, Matej Balog, M. Pawan Kumar, Emilien Dupont, Francisco J. R. Ruiz, Jordan S. Ellenberg, Pengming Wang, Omar Fawzi, Pushmeet Kohli, and Alhussein Fawzi, *Mathematical discoveries from program search with large language models*, Nature **625** (2024), no. 7995, 468–475. [↑4](#)
- [She07] S. Sheffield, *Gaussian free fields for mathematicians*, Probab. Theory Relat. Fields **139** (2007), no. 3-4, 521–541. arXiv:math/0312099 [math.PR]. [↑26](#)
- [Sta17] R. Stanley, *Some Schubert shenanigans*, arXiv preprint (2017). arXiv:1704.00851 [math.CO]. [↑2](#), [4](#)
- [Wei21] A. Weigandt, *Bumpless pipe dreams and alternating sign matrices*, J. Comb. Theory Ser. A **182** (2021), 105470. arXiv:2003.07342 [math.CO]. [↑2](#), [30](#)
- [ZJ09] P. Zinn-Justin, *Six-Vertex, Loop and Tiling models: Integrability and Combinatorics*, arXiv preprint (2009). arXiv:0901.0665 [math-ph]. [↑2](#)

OHIO STATE UNIVERSITY, COLUMBUS, OH

Email address: anderson.2804@math.osu.edu

UNIVERSITY OF SOUTHERN CALIFORNIA, LOS ANGELES, CA

Email address: gpanova@usc.edu

UNIVERSITY OF VIRGINIA, CHARLOTTESVILLE, VA

Email address: lenia.petrov@gmail.com