

MARKOV CHAINS ON WEYL GROUPS FROM THE GEOMETRY OF THE FLAG VARIETY

PERSI DIACONIS AND CALDER MORTON-FERGUSON

ABSTRACT. This paper studies a basic Markov chain, the Burnside process, on the space of flags G/B with $G = \mathrm{GL}_n(\mathbb{F}_q)$ and B its upper triangular matrices. This gives rise to a shuffling: a Markov chain on the symmetric group realized via the Bruhat decomposition $S_n \cong B \backslash G/B$. Actually running and describing this Markov chain requires understanding Springer fibers and the Steinberg variety. The main results give a practical algorithm for all n and q and determine the limiting behavior of the chain when q is large. In describing this behavior, we find interesting connections to the combinatorics of the Robinson–Schensted correspondence and to the geometry of orbital varieties. The construction and description is then carried over to finite Chevalley groups of arbitrary type, describing a new class of Markov chains on Weyl groups.

CONTENTS

1. Introduction	1
2. The Burnside process on G/B	3
3. A sampling algorithm	7
4. GL_n behavior: Tableaux and the Robinson–Schensted correspondence	10
5. General type behavior: orbital varieties and Steinberg cells	16
6. Consequences for mixing times	18
7. Examples and simulations	19
8. Final comments	21
References	24

1. INTRODUCTION

A basic group theory question is to understand the decomposition of a group G into double cosets $H \backslash G/K$. For G a finite group and H, K subgroups, one may ask:

- How many double cosets are there?
- What are typical sizes?
- Do the double cosets have “nice names”?
- Do they fit together into some sort of understandable moduli space?

A famous example is the Bruhat decomposition

$$W = U \backslash G/B,$$

MSC 2020: Primary: 60J10, 14M15. Secondary: 05E10, 05A17, 20G40.

with $G = \mathrm{GL}_n(\mathbb{F}_q)$ the invertible matrices over a finite field $\mathbb{F}_q$, B the upper triangular matrices, U the unipotent matrices in B , and $W = S_n$. This decomposition is realized by Gaussian elimination, and generalizes to the case of G a finite Chevalley group and W its Weyl group. Going back to the case of a general finite group G , the Burnside process is a stochastic algorithm for picking a double coset at random. This gives an effective way to get approximate answers to the first two questions above.

This paper began with the practical question: how would one actually carry out the Burnside process for the action of U on G/B , for $G = \mathrm{GL}(\mathbb{F}_q)$? In this case G/B is the space of flags F of vector spaces in $\mathbb{F}_q^n$. The process entails two steps:

- (1) From some $F \in G/B$, choose $u \in U$ with $F = uF$, uniformly at random.
- (2) From this $u \in U$, choose $F' \in G/B$ with $F' = uF'$, uniformly at random.

The chain moves from F to F' with probability given by

$$(1.1) \quad P(F, F') = \frac{1}{|\mathrm{stab}_U(F)|} \sum_{g \in \mathrm{stab}_U(F) \cap \mathrm{stab}_U(F')} \frac{1}{|\mathrm{Fix}_X(g)|},$$

where Section 2 explains the notation. This gives a Markov chain on G/B with stationary distribution

$$\pi(x) = \frac{1}{n!|O_F|},$$

where O_F is the U -orbit of F . Thus simply reporting the orbit after each step gives a method for making a uniform choice of orbit. The orbits are indexed by permutations, so a random choice of $w \in S_n$ results.

It turns out that even in this special case, practical implementation of steps (1)–(2) requires an understanding of the geometry and combinatorics of flag varieties and Springer fibers. Fixing $u \in U$, the set

$$\{F : F^u = F\}$$

is the Springer fiber: a variety whose cohomology underlies the representation theory of W . Running the Burnside process involves understanding the size of Springer fibers, while computing explicit transition probabilities involves understanding the intersection of Springer fibers with a fixed double coset UwB . This is a difficult task in general. For GL_n , we are able to give an explicit implementation of the Burnside process using connections between Springer fibers and the combinatorics of Hall–Littlewood and Green polynomials.

Section 2 gives background on the Burnside process and the geometry of flag space required to carry out steps (1) and (2) for $G = \mathrm{GL}_n(\mathbb{F}_q)$. Section 3 spells out the details of a full sampling algorithm for this case. This algorithm has been implemented, and actual runs of the Markov chain are presented and discussed in Section 7. This may be read now for further motivation.

Hands-on understanding of the Markov chain, even for $W = S_n$, is difficult. However, when q is large and n fixed, we show that there is a simple description of its behavior in terms of Young tableaux and the Robinson–Schensted correspondence. This is developed in Section 4 with examples in Section 7. This analysis gives a useful lower bound on the convergence of the Markov chain to its stationary distribution whose details are in Section 6. We also discuss a more precise analysis and set of bounds in the case of $G = \mathrm{GL}_3(\mathbb{F}_q)$ in Proposition 7.1.

Section 5 develops the same theme for the more general case of finite Chevalley groups in any type. It has a self-contained introduction and analogues of the main results of Section 4. Now the Young tableaux and their associated combinatorics must be replaced by more sophisticated tools related to Steinberg and orbital varieties. A practical version of our algorithm from Section 3 in this general case remains for the future. That being said, in all types we are still able to describe large q behavior in terms of interesting combinatorial data of general Weyl groups called Steinberg cells.

The topics developed may interest researchers in disparate fields (probability, combinatorics, geometric representation theory). Thus we have tried to write a broadly accessible account.

Acknowledgments. We thank Michael Howes, Matthew Nicoletti, and Arun Ram for very helpful comments. P.D. was supported in part by NSF grant 1954042.

2. THE BURNSIDE PROCESS ON G/B

2.1. The Burnside process for the action of a finite group.

2.1.1. *General setup.* The Burnside process for the action of a finite group G on a set X is a Markov chain on X built from the following two steps. From a given element $x \in X$ one proceeds as follows:

- (1) Choose an element $g \in \text{stab}_G(x)$ uniformly at random.
- (2) Choose an element $y \in \text{Fix}_X(g)$ uniformly at random.

Here $\text{stab}_G(x) \subset G$ is the stabilizer of x and $\text{Fix}_X(g) \subset X$ is the set of elements fixed by g . The transition probability $P(x, y)$ of moving from x to y is then given by the formula in (1.1). It is explained in [7] that P is an ergodic, reversible Markov kernel with stationary distribution given by

$$\pi(x) = \frac{|\Theta|^{-1}}{|O_x|},$$

where Θ is the set of G -orbits on X and O_x is the orbit $G \cdot x$.

One can then consider the Burnside process on the set Θ of G -orbits, obtained by remembering only the G -orbit at each step of the Burnside process. One then gets a reversible Markov chain on orbits whose stationary distribution is always uniform; we call its Markov kernel P_Θ .

The Burnside process has been used to generate and count “Pólya trees” (shapes of rooted labeled trees) on n vertices for n of order 10^8 [1], to generate random partitions of n for n of order 10^8 and random contingency tables ($i \times j$ arrays of non-negative integers with given row and column sums) [7] and for counting the number of conjugacy classes of the uni-upper triangular matrices in $\text{GL}_n(\mathbb{F}_q)$ [10]. In all of these applications, as in the present paper, there was substantial mathematical and computational effort required to get the two steps up and running. Convergence estimates such as a spectral gap for the underlying Markov chain were left for future work. There have been successful running time estimates for the Burnside process in special cases; see [9], [19], [20]. We hope to do this for the walks studied here.

2.1.2. *The Burnside process on the flag variety for $\mathrm{GL}_n(\mathbb{F}_q)$.* Now let $G = \mathrm{GL}_n(\mathbb{F}_q)$, invertible $n \times n$ matrices with entries in the finite field $\mathbb{F}_q$. Let B be its subgroup of upper triangular matrices and U the set of unipotent matrices in B . We then let $X = G/B$; this is the *flag variety* of $\mathrm{GL}_n(\mathbb{F}_q)$ over the finite field $\mathbb{F}_q$. It is the set of *complete flags*

$$0 = V_0 \subset V_1 \subset \cdots \subset V_n = \mathbb{F}_q^n$$

of vector subspaces of $\mathbb{F}_q^n$ where $\dim(V_k) = k$ for all $1 \leq k \leq n$; we write $V_\bullet$ for convenience. It admits a left action of the group U . We now recall the Bruhat decomposition, which describes the orbits of the action of U (or equivalently, of B) on X . It gives that the orbits of this action are labelled by elements of the symmetric group S_n . We write ℓ for the length function on S_n , which counts the number of inversions of a permutation.

Proposition 2.1. *There is a decomposition*

$$G/B = \coprod_{w \in S_n} BwB = \coprod_{w \in S_n} UwB,$$

where for each $w \in S_n$, $UwB = BwB$ is considered as a subset of G/B , and $|UwB| = q^{\ell(w)}$.

In the present paper we study the Burnside process for the action of U on X . Given $x, y \in X$, we will write $P(x, y)$ for the transition probabilities, and since by Proposition 2.1 the orbits for this action are indexed by S_n , for $w, z \in W$ we will use $P_{S_n}(w, z)$ to denote the transition probabilities for the corresponding Burnside process on the set of orbits. We refer to it as the q -Burnside process when emphasizing the dependence on q .

2.2. **Stabilizers in U .** Given some choice of $x \in X$, the first step of the Burnside process requires uniform sampling from $\mathrm{stab}_U(x)$. Using Proposition 2.1 to write $x = uwB$ for some $u \in U$, $w \in W$, we note that

$$\mathrm{stab}_U(x) = u \cdot \mathrm{stab}_U(wB) \cdot u^{-1},$$

so we now recall an explicit description of $\mathrm{stab}_U(wB)$ for any $w \in W$.

Definition 2.2. For any $w \in W$, let

$$U_w = \mathrm{stab}_U(wB) \subset U.$$

Setting $U^w = wUw^{-1}$, one then clearly has $U_w = U \cap wUw^{-1} = U \cap U^w$.

Proposition 2.3 (see Proposition 5.2). *For any $w \in S_n$, let*

$$\mathrm{Inv}(w) = \{(i, j) \mid i < j, w(i) > w(j)\}$$

denote the set of inversions of w .

Then

$$U_w = \{(u_{ij}) \in U \mid u_{ij} = 0 \text{ for } (i, j) \in \mathrm{Inv}(w)\}.$$

In other words, U_w consists of upper triangular unipotent matrices in which the entries corresponding to inversions of w are forced to vanish. In particular, for any $w \in W$ the cardinality of the indexing set in the product which appears in Proposition 2.3 is exactly $\binom{n}{2} - \ell(w)$, so

$$(2.1) \quad |U_w| = q^{\binom{n}{2} - \ell(w)}.$$

Example 2.4. Suppose $n = 3$. In this case, $W = \{e, s_1, s_2, s_2s_1, s_1s_2, w_0\} \cong S_3$, where $s_1 = (12)$, $s_2 = (23)$ and $w_0 = s_1s_2s_1$.

$$\begin{aligned} \text{stab}_U(s_1B) &\cong \left\{ \begin{pmatrix} 1 & 0 & * \\ 0 & 1 & * \\ 0 & 0 & 1 \end{pmatrix} \right\} & \text{stab}_U(s_2B) &\cong \left\{ \begin{pmatrix} 1 & * & * \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \right\}, \\ \text{stab}_U(s_2s_1B) &\cong \left\{ \begin{pmatrix} 1 & * & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \right\} & \text{stab}_U(s_1s_2B) &\cong \left\{ \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & * \\ 0 & 0 & 1 \end{pmatrix} \right\} \end{aligned}$$

while $\text{stab}_U(eB) \cong U$ and $\text{stab}_U(w_0B) \cong \{1\}$.

2.3. Fixed-point sets and Springer fibers. Now given an arbitrary $u \in U$, the second step in the Burnside process calls for an understanding of $\text{Fix}_X(u)$, which we refer to as X_u . This is the set of flags $V_\bullet$ such that

$$u(V_k) \subset V_k$$

for all $1 \leq k \leq n$.

This is the *Springer fiber* associated to the unipotent matrix u . We will explain the relationship between the geometry of X_u to the transition probabilities in the Burnside process on S_n . A starting point for a discussion of these combinatorics is an understanding of the cardinality $|X_u|$ itself.

2.3.1. Green polynomials. In [18, III.3], the polynomials

$$X_\mu^\lambda(q)$$

associated to partitions μ, λ of n are defined as the coefficients expressing the power-sum symmetric functions in terms of Hall–Littlewood polynomials:

$$p_\mu(x) = \sum_{\lambda \vdash n} X_\mu^\lambda(q) P_\lambda(x; q).$$

Here μ and λ are partitions of n , p_μ is the power-sum symmetric function, and $P_\lambda(x; q)$ is the Hall–Littlewood polynomial. Thus the $X_\mu^\lambda(q)$ form the entries of the transition matrix from the power-sum basis $\{p_\mu\}$ to the Hall–Littlewood basis $\{P_\lambda\}$.

The *Green polynomials* $Q_\mu^\lambda(q)$ are related to these coefficients by the reciprocity identity

$$(2.2) \quad q^{-b(\lambda)} Q_\mu^\lambda(q) = X_\mu^\lambda(1/q),$$

where

$$b(\lambda) = \sum_{i \geq 1} (i-1)\lambda_i = \deg(X_\mu^\lambda) = \deg(Q_\mu^\lambda).$$

We note that in [18] it is referred to as $n(\lambda)$, a notation we avoid here so as not to confuse it with the n in S_n .

We are primarily interested in the case $\mu = (1^n)$; in this case, the polynomials

$$Q_{(1^n)}^\lambda(q)$$

encode the cardinalities of the Springer fibers X_u as follows.

Proposition 2.5 (see [18, Ch. III, §7]). *Let $u \in \mathrm{GL}_n(\mathbb{F}_q)$ be a unipotent matrix with Jordan type $\lambda \vdash n$. Then*

$$|X_u| = Q_{(1^n)}^\lambda(q),$$

where $Q_{(1^n)}^\lambda(q)$ is the corresponding Green polynomial.

Although here we consider X_u only as a finite set, the Springer fiber $\mathbf{X}_u$ can be similarly defined and considered as a variety over $\overline{\mathbb{F}}_q$. We note that by Proposition 2.5, $X_{(1^n)}^\lambda(0)$, which is the leading coefficient of $Q_{(1^n)}^\lambda(q)$, can be viewed as the number of irreducible components of $\mathbf{X}_u$. In [23], it is explained that when $u \in \mathrm{GL}_n(\mathbb{F}_q)$ has Jordan type $\lambda \vdash n$, these irreducible components are in bijection with the set of standard Young tableaux of shape λ . One can also directly deduce that $X_{(1^n)}^\lambda(0)$ is the number of such tableaux using the following combinatorial argument.

Proposition 2.6. *For any partition λ of n ,*

$$X_{(1^n)}^\lambda(0) = f^\lambda,$$

where f^λ is the number of standard Young tableaux of shape λ .

Proof. By definition, $X_{(1^n)}^\lambda(q)$ is a coefficient of the transition matrix between power-sum products and Hall-Littlewood functions. This means

$$p_{(1^n)}(x) = \sum_{\mu} X_{(1^n)}^\mu(q) P_\mu(x; q).$$

We note that $p_{(1^n)}(x) = p_1(x)^n$, where $p_1(x) = x_1 + \cdots + x_n$, while $P_\mu(x; 0) = s_\mu(x)$, the Schur function corresponding to μ .

Noting that $p_1(x)$ is the Frobenius characteristic of the regular representation of S_n , we have

$$p_1(x)^n = \sum_{\mu} f^\mu s_\mu(x),$$

since f^μ is the dimension of the irreducible S_n -representation indexed by μ . Comparing coefficients on the Schur polynomial $s_\lambda(x)$ gives $X_{(1^n)}^\lambda(0) = f^\lambda$. $\square$

2.3.2. Springer fibers and Bruhat cells. We now explain why an understanding of transition probabilities for our Burnside process reduces to an understanding of the cardinalities of the sets

$$(2.3) \quad X_u \cap BwB$$

for all $w \in W$.

The following comes from combining formulas (1.1) and (2.1).

Proposition 2.7. *The transition probability $P_{S_n}(w, z)$ for the q -Burnside process on S_n is given by*

$$P_{S_n}(w, z) = \frac{1}{q^{\binom{n}{2} - \ell(w)}} \sum_{u \in U_w(\mathbb{F}_q)} \frac{1}{Q_{(1^n)}^\lambda(q)} |X_u \cap ByB|$$

This means that an understanding of the quantities $|X_u \cap ByB|$ for all $u \in U$ and $y \in S_n$ would yield a precise understanding of the transition probabilities in our Markov matrix P_{S_n} . In this section, we give a brief overview of what is known of these quantities. The upshot is that for certain “carefully-chosen” u , there is a very nice combinatorial description of $|X_u \cap ByB|$ in terms of y and the Jordan type

of u , whereas determining the cardinalities $|X_u \cap ByB|$ for arbitrary u has been shown to be very difficult. (More precisely, [25, Section 6] indicates that necessary and sufficient conditions on u for the following combinatorial description to hold are not known, but a sufficient condition is explained in detail in [26, Section 4].)

It is always true (for any u) that $\mathbf{X}_u$ has an *affine paving* (a filtration by closed subvarieties whose successive differences are isomorphic to affine spaces) which can be used to give a geometric explanation for an expression of $|X_u|$ as a positive linear combination of monomials q^d . If u is “carefully chosen” in the sense of [26], then the resulting partition of X_u is given by $\{X_u \cap ByB\}_{y \in W}$, and therefore each quantity $|X_u \cap ByB|$ is equal to q^d for d the dimension of the corresponding affine space. The dimension d can then be computed combinatorially in terms of y and the Jordan type of u in terms of the number of *row-strict fillings* of a certain Young diagram; this combinatorial description is explained in [25, Section 3.2].

It is explained in [25] that little is known about the intersections $|X_u \cap ByB|$ for other u other than the fact that they can be very complicated; examples where these intersections deviate from the combinatorial formula previously explained are exhibited in [16] and [21]. It is known (c.f. [25]) to be a difficult open problem to produce a formula for these quantities in more generality.

We view this as evidence toward our belief that the precise transition probabilities which appear in our Burnside process have the potential to be extremely complicated, and we do not believe it is likely that they can be precisely characterized in any neat algebraic form. Despite this, we are able to analyze leading order terms in q of these transition probabilities to develop a complete understanding of the approximate behavior of the Markov chain up to “error terms” of order $O(1/q)$. Such a characterization is the main result of Sections 4 and 5.

Remark 2.8. Despite how complicated we expect general entries of the transition matrix $\{P_{S_n}(w, z)\}_{w, z \in S_n}$ to be, we note that the single column $\{P_{S_n}(w_0, w)\}_{w \in S_n}$ (where w_0 is the unique “longest permutation” with $\binom{n}{2}$ inversions) is very simple. In fact, for any $w \in S_n$,

$$P_{S_n}(w_0, w) = \frac{q^{\ell(w)}}{|X|}$$

This means that starting at w_0 and running the Burnside process for one step is equivalent to sampling from the *Mallows measure*, a well-studied probability measure on S_n . This is because U_{w_0} is trivial, and therefore this step of the Burnside process samples uniformly from X . The Bruhat cell corresponding to any $w \in S_n$ has cardinality $q^{\ell(w)}$, from which this observation follows.

3. A SAMPLING ALGORITHM

In this section, we continue with the case of $G = \mathrm{GL}_n(\mathbb{F}_q)$. In this case, the Burnside process gives rise to a *shuffling*, i.e. a Markov chain on S_n . In the present section, we describe an explicit sampling algorithm which one can use to implement the shuffling arising from this Burnside process in practice. This algorithm can be implemented with code (we implemented it in SageMath to run the simulations explained in Section 7) and runs efficiently even in cases when a brute-force enumeration of flags is impractical.

3.1. Some partition combinatorics. Suppose that $\lambda = (\lambda_1, \dots, \lambda_k)$ is a partition of n ; we freely identify it with its corresponding Young diagram. When describing such a Young diagram (and in the examples in Section 7) we will always use English notation, so the row lengths are weakly decreasing when reading from top to bottom. We can then make the following definitions.

Definition 3.1. Let $R(\lambda)$ be the set of partitions corresponding to those Young diagrams obtained by removing a single box from the Young diagram of λ . In other words, partitions which arise from decrementing a single part of λ by 1.

Given $\lambda' \in R(\lambda)$, we write $r(\lambda, \lambda')$ to denote the index of the part of λ which was decremented by 1 to form λ' . By convention, $r(\lambda, \lambda')$ is always maximal among those j for which $\lambda_j = \lambda_{r(\lambda, \lambda')}$.

Finally, let $m(\lambda, \lambda')$ denote the multiplicity of $\lambda_{r(\lambda, \lambda')}$ in λ .

Definition 3.2. Let $I(\lambda) \subset \{1, \dots, n\}$ be the k -element set consisting of 1 along with all integers of the form

$$(3.1) \quad \lambda_1 + \lambda_2 + \dots + \lambda_i + 1$$

for $1 \leq i \leq k-1$.

For any partition λ of n , we write J_λ for the unipotent matrix in Jordan canonical form with its ordered set of k Jordan blocks having sizes $(\lambda_1, \dots, \lambda_k)$.

The following is immediate.

Lemma 3.3. *The k -dimensional subspace of $\mathbb{F}_q^n$ consisting of eigenvectors of J_λ is*

$$(3.2) \quad E(\lambda) = \text{span}\{e_i\}_{i \in I(\lambda)}.$$

If $v \in E(\lambda)$, then the linear transformation $\overline{J_\lambda}$ is well-defined on the quotient $\mathbb{F}_q^n / \text{span}(v)$.

Definition 3.4. For $\lambda' \in R(\lambda)$, let $E^{\lambda'}(\lambda)$ be the subset of $E(\lambda)$ consisting of eigenvectors v of J_λ for which the transformation $\overline{J_\lambda}$ has Jordan type λ' on $\mathbb{F}_q^n / \text{span}(v)$.

Lemma 3.5. *Suppose $\lambda' \in R(\lambda)$. Then $E^{\lambda'}(\lambda)$ is exactly the set of $(v_i)_{i=1}^n$ such that*

- $v_i = 0$ unless $i \in I(\lambda)$,
- $v_i = 0$ for $i \geq \lambda_1 + \dots + \lambda_{r(\lambda, \lambda')} + 1$,
- $v_i \neq 0$ for at least one i satisfying

$$\lambda_1 + \dots + \lambda_{r(\lambda, \lambda') - m(\lambda, \lambda')} + 1 \leq i < \lambda_1 + \dots + \lambda_{r(\lambda, \lambda')} + 1.$$

This gives an explicit bijection

$$(3.3) \quad E^{\lambda'}(\lambda) \cong \mathbb{F}_q^{r(\lambda, \lambda') - m(\lambda, \lambda')} \times (\mathbb{F}_q^{m(\lambda, \lambda')} \setminus \{0\})$$

for any $\lambda' \in R(\lambda)$, which leads to the following result.

Corollary 3.6. *Suppose $\lambda' \in R(\lambda)$. Then*

$$(3.4) \quad |E^{\lambda'}(\lambda)| = q^{r(\lambda, \lambda') - m(\lambda, \lambda')} (q^{m(\lambda, \lambda')} - 1) = q^{r(\lambda, \lambda')} - q^{r(\lambda, \lambda') - m(\lambda, \lambda')}$$

3.2. The sampling algorithm. Suppose that $V_\bullet = 0 \subset V_1 \subset \dots \subset V_n = \mathbb{F}_q^n$ is a complete n -step flag of subspaces in $\mathbb{F}_q^n$. We now describe an explicit algorithm for sampling another flag of subspaces according to the Burnside process for the action of U on G/B .

3.2.1. *Bruhat decomposition via Gaussian elimination.* Choose a basis $\{v_i\}_{i=1}^n$ for $\mathbb{F}_q^n$ such that for all $1 \leq k \leq n$,

$$(3.5) \quad V_k = \text{span}(v_1, \dots, v_k).$$

Then $V_\bullet$ is the flag corresponding to the point $gB \in G/B$, where $g \in G$ is the matrix with v_i (written with respect to the standard basis $\{e_j\}_{j=1}^n$) as its i th column. This identification is independent of the choice of basis.

Bruhat decomposition guarantees that one can then apply Gaussian elimination (represented as multiplication by a unipotent upper triangular matrix on the left) to reduce g to a permutation matrix. This gives an element $u \in U$ and $w \in S_n$ such that

$$uwB = gB \in G/B.$$

A very clear exposition of Bruhat decomposition as Gaussian elimination appears in Roger Howe's expository work [14].

3.2.2. *Sampling from the stabilizer.* First we explain that uniformly sampling from $\text{stab}_U(uwB) \subset U$ is straightforward. Indeed, note that

$$(3.6) \quad \text{stab}_U(uwB) = u \cdot \text{stab}_U(w) \cdot u^{-1} = uU_w u^{-1},$$

so it is enough to sample uniformly from U_w and then conjugate the resulting element by u .

Since by Proposition 2.3,

$$(3.7) \quad U_w = \{(u_{ij}) \in U \mid u_{ij} = 0 \text{ for } (i, j) \in \text{Inv}(w)\} \cong \mathbb{F}_q^{\ell(w_0) - \ell(w)},$$

it is enough to sample uniformly from $\mathbb{F}_q^{\ell(w_0) - \ell(w)}$. Passing the result through the bijection (3.7) and conjugating by u , we have chosen an element $a \in \text{stab}_U(uwB)$ uniformly at random. We then write $a = QJQ^{-1}$ for some $Q \in G$ with J a unipotent upper triangular matrix in Jordan form with Jordan blocks of non-increasing size corresponding to a permutation $\lambda = (\lambda_1, \dots, \lambda_k)$.

3.2.3. *Weighted sampling on $R(\lambda)$.* The next step of our algorithm is a weighted sampling among elements of $R(\lambda)$. For any $\lambda' \in R(\lambda)$, let

$$\text{wt}(\lambda') = (q^{r(\lambda, \lambda')} - q^{r(\lambda, \lambda') - m(\lambda, \lambda')}) Q_{(1^n)}^{\lambda'}(q).$$

We then perform a weighted sampling among elements $\lambda' \in R(\lambda)$ with probabilities proportional to the weights $\text{wt}(\lambda')$. We write μ for the resulting element of $R(\lambda)$.

3.2.4. *Sampling from the Springer fiber.* By means of the bijection

$$E^\mu(\lambda) \cong \mathbb{F}_q^{r(\lambda, \mu) - m(\lambda, \mu)} \times (\mathbb{F}_q^{m(\lambda, \mu)} \setminus \{0\})$$

described in Lemma 3.5, one then samples an element of $E^\mu(\lambda)$ uniformly at random, calling it v .

Since by Corollary 3.6 there are $q^{r(\lambda, \mu)} - q^{r(\lambda, \mu) - m(\lambda, \mu)}$ vectors in $E^\mu(\lambda)$. This means that these two steps (weighted sampling among elements of $R(\lambda)$ followed by uniform sampling among $E^\mu(\lambda)$) give a proper weighted sampling among all vectors in $E(\lambda)$ with $v \in E^\mu(\lambda)$ having probability proportional to $Q_{(1^n)}^\mu(q)$.

With the $v \in E^\mu(\lambda)$ chosen above, one then notes that the transformation $\bar{J}$ on $\mathbb{F}_q^n/\text{span}(v)$ has Jordan type μ . Inductively applying this algorithm, one obtains an n -step flag fixed by J which we call $V'_\bullet$, writing

$$0 \subset V'_1 = \text{span}(v) \subset V'_2 \subset \cdots \subset V'_n = \mathbb{F}_q^n.$$

By Proposition 2.5 applied to $\bar{J}$, there are $Q_{(1^n)}^\mu(q)$ such flags fixed by $\bar{J}$ whose first step V_1 is $\text{span}(v)$. For the sake of induction, we assume for now that each of these flags is chosen uniformly at random.

Since our algorithm chooses v with probability proportional to $Q_{(1^n)}^\mu(q)$, which is the number of possible extensions of $\text{span}(v)$ into an n -step flag which are fixed by J , we then get that our algorithm chooses uniformly at random among flags fixed by J . The base case where $n = 1$ is clearly a uniform sampling.

3.2.5. *Conclusion.* We then know that the n -step flag

$$(3.8) \quad 0 \subset QV'_1 \subset \cdots \subset QV'_n = \mathbb{F}_q^n$$

has been chosen uniformly at random among flags fixed by $a = QJQ^{-1}$. Let $g'B$ be the corresponding point in G/B , with g' constructed as in Section 3.2.1. By the method described therein, we can also then write $g' = u'w'B$ for some $u' \in U$ and $w' \in S_n$.

Starting with the original point gB , we have just explained how to sample an element a uniformly from its U -stabilizer, and then an element of G/B uniformly among a -fixed points. Therefore this algorithm indeed samples according to the Burnside process for the U -action on G/B .

Writing $gB = uwB$ and $g'B = u'w'B$ and recording only the permutations w and w' , this therefore also describes a sampling algorithm for the lumped Burnside process orbits, i.e. the corresponding Markov chain on S_n .

4. GL_n BEHAVIOR: TABLEAUX AND THE ROBINSON–SCHENSTED CORRESPONDENCE

In this section, we prove our main result for $\text{GL}_n(\mathbb{F}_q)$, in the form of Theorem 4.12, describing the large- q limiting behavior of the Markov chain on S_n arising from the Burnside process on the flag variety in terms of the combinatorics of the Robinson–Schensted correspondence.

Before doing so, we now describe this behavior informally. We will explain that to any permutation S_n , one can associate a standard Young tableau $P(w)$ with n boxes. We will show that when q is large, then with very high probability, the Markov chain sends any $w \in S_n$ to a uniform choice of $z \in S_n$ subject to the condition that $P(z) = P(w)$. Thus the Markov chain tends to cluster into “buckets” labelled by Young tableaux which are very difficult to exit. Despite this, over a very long period of time, the Markov chain still eventually finds its way from bucket to bucket so as to converge to the uniform distribution.

In Section 5, we will show that the same behavior holds in the more general setting when $\text{GL}_n(\mathbb{F}_q)$ is replaced by a finite Chevalley group $G(\mathbb{F}_q)$ of any type, thereby obtaining a Markov chain on an arbitrary Weyl group W . We will then show that the same behavior roughly holds, but the “buckets” from the previous paragraph can no longer be labeled by Young tableaux. Instead, they will be indexed by *Steinberg cells*, a certain partition of W which arises from the geometry of orbital varieties.

4.1. The Robinson–Schensted correspondence. The Robinson–Schensted correspondence is a bijection between permutations $w \in S_n$ and pairs of standard Young tableaux (P, Q) of the same shape $\lambda \vdash n$. Given a permutation w , the correspondence produces two tableaux:

- $P(w)$, called the *insertion tableau*, constructed by successively inserting the entries of w into a tableau by the row-insertion algorithm described in [22].
- $Q(w)$, called the *recording tableau*, which records the order in which boxes are added during the insertion process.

We refer the reader to [11] for a clear summary and exposition of the algorithm in question.

Both $P(w)$ and $Q(w)$ are standard Young tableaux of the same shape λ . We let $T(w) = \lambda$ be the corresponding partition of n . We now recall the following basic property of the Robinson–Schensted correspondence.

Proposition 4.1 (Lemma 7 in [22]). *For any $w \in S_n$,*

$$(4.1) \quad P(w^{-1}) = Q(w).$$

Definition 4.2. For a partition $\lambda \vdash n$, recall that f^λ is the number of standard Young tableaux of shape λ . Equivalently, f^λ is the dimension of the irreducible representation of the symmetric group S_n corresponding to λ . The hook length formula gives

$$f^\lambda = \frac{n!}{\prod_{(i,j) \in \lambda} h_{ij}},$$

where h_{ij} is the hook length of the box (i, j) in the Young diagram of λ .

4.2. Underlying geometry. Recall that in Proposition 2.3, we described the subspace $U_w = U \cap wUw^{-1}$ of matrices fixing $wB \in G/B$. It will be convenient to instead work with $\mathfrak{n}$, the set of strictly upper-triangular (nilpotent) $n \times n$ matrices. We let $\mathfrak{n}^w = w\mathfrak{n}w^{-1}$ and write $\mathfrak{n}_w = \mathfrak{n} \cap \mathfrak{n}^w$. We note that

$$\mathfrak{n}_w = \{u - \text{id}_{n \times n} : u \in U_w\}.$$

Since a flag $V_\bullet$ is fixed by $u \in U_w$ if and only if it is fixed by $x = u - \text{id}_{n \times n} \in \mathfrak{n}_w$, we note that U_w and $\mathfrak{n}_w$ can and will be used interchangeably for the purposes of the Burnside process. Note that $\mathfrak{n}$ and its subspaces admit an action by G under conjugation. We note that $\mathfrak{n}$, its subspaces, and G -orbits therein all naturally inherit the structures of varieties over $\overline{\mathbb{F}}_q$, which we will use in what follows to discuss notions of genericity and dimension.

Definition 4.3. For any $w \in W$, let

$$V_w = U \cdot \mathfrak{n}_w.$$

Note that when considered as an algebraic variety over $\overline{\mathbb{F}}_q$, the closure $\overline{V_w}$ is an *orbital variety*: an irreducible component of the intersection of a nilpotent orbit with $\mathfrak{n}$. The following is a restatement of Steinberg’s beautiful geometric realization of the Robinson–Schensted correspondence which appears in [24].

Proposition 4.4 (Steinberg, [24]). *For a pair of permutations $w, z \in S_n$, $P(w) = P(z)$ if and only if*

$$(4.2) \quad U \cdot \mathfrak{n}_w = U \cdot \mathfrak{n}_z.$$

An alternative way to see this result is as follows. In the forthcoming Definition 5.3, we recall the definition of *right Steinberg cells* in general Weyl groups and explain that by definition, w and y are in the same right Steinberg cell if and only if (4.2) holds. Proposition 4.4 then asserts that for S_n , right Steinberg cells are determined by the P -symbol of a permutation. Indeed, in [2], it is shown that for S_n , right Steinberg cells agree with right Kazhdan–Lusztig cells; our result then follows from Kazhdan–Lusztig’s original work [15] where right Kazhdan–Lusztig cells in S_n are indexed by standard Young tableaux, and membership of an element w in a cell indexed by P is decided by the condition that $P(w) = P$.

Lemma 4.5. *When q is large, a generic element of $\mathfrak{n}_w$ (considered as a variety over $\mathbb{F}_q$) has Jordan type given by $T(w)$, the shape of $P(w)$.*

Further, any arbitrary element of $\mathfrak{n} \cap \mathfrak{n}^w$ has Jordan type μ for some $\mu \leq T(w)$ in the dominance order for partitions.

Proof. By definition, a generic element of $\mathfrak{n} \cap \mathfrak{n}^w$ has a zero in entry (i, j) if (i, j) is not in the inversion set of w and a generic nonzero element of $\mathbb{F}_q$ if (i, j) is in the inversion set of w for $i < j$. When q is large, one can choose generic nonzero elements such that the Jordan type of the corresponding matrix has the same Jordan type as a generic element (over $\mathbb{C}$) of the subspace of $\mathfrak{n}$ corresponding to the *inversion digraph* of w as in [12].

Gansner’s theorem, explained in loc. cit., gives a partition from any generic nilpotent matrix with support determined by a digraph by taking the Jordan type. The inversion digraph is precisely the complement of the poset of allowed entries for $\mathfrak{n} \cap \mathfrak{n}^w$. Gansner explains that when this procedure is applied to a generic nilpotent matrix whose support is determined by the the inversion digraph of a permutation, one obtains the transpose of the underlying Young diagram obtained from applying Robinson–Schensted algorithm, as detailed in [22], to that same permutation. This relationship is established using Greene’s theorem from [13]. Thus applying it to the complement (the non-inversion digraph), we obtain simply $T(w)$ without a transpose; this is exactly the first claim in the lemma.

Now note that the same logic combined with [12, Lemma 4.1] implies that any matrix (generic or not) in $\mathfrak{n} \cap \mathfrak{n}^w$ must have Jordan type dominated by $T(w)$ in the dominance order. $\square$

Definition 4.6. For any $w \in S_n$, let $d(w) = \dim V_w$ as a variety over $\overline{\mathbb{F}}_q$.

The following explicit computation for $d(w)$ follows from [2].

Proposition 4.7. *There exists some generic element $x \in \mathfrak{n}_w$ such that*

$$(4.3) \quad d(w) = \frac{1}{2} \dim \mathbb{O}_x,$$

where $\mathbb{O}_x$ is the nilpotent orbit $G \cdot x$.

We note that the nilpotent orbit $\mathbb{O}_x$ depends only on the Jordan type $\lambda(x)$, so it makes sense to simply refer to a nilpotent orbit $\mathbb{O}_\lambda$ for λ a partition of n . Applying Lemma 4.5 to the dimension formula in Proposition 4.7, we get the following.

Corollary 4.8. *For any $w \in S_n$, let $\lambda = T(w)$. Then $d(w)$ depends only on λ , and we have*

$$(4.4) \quad d(w) = \frac{\dim \mathbb{O}_\lambda}{2}$$

The dimension $\dim \mathbb{O}_\lambda$ can be described combinatorially as follows (see, e.g., [6, III.7]).

Lemma 4.9. *For any partition λ of n , the dimension of the nilpotent orbit $\mathbb{O}_\lambda$ is*

$$(4.5) \quad \dim \mathbb{O}_\lambda = n^2 - \sum_i (\lambda'_i)^2,$$

where λ' is the transpose (conjugate partition) of λ .

4.3. Combinatorics of Green polynomials. Recall that for any $\lambda \vdash n$ with conjugate partition λ' ,

$$(4.6) \quad b(\lambda) = \deg Q_{(1^n)}^\lambda(q) = \sum_{i \geq 1} (i-1)\lambda_i = \sum_j \binom{\lambda'_j}{2}.$$

Comparing this with Lemma 4.9 gives us the following comparison between $b(\lambda)$ and $d(w)$, whenever $P(w)$ has shape λ .

Lemma 4.10. *For any partition λ of n and any $w \in S_n$ such that $P(w)$ has shape λ , we have*

$$(4.7) \quad b(\lambda) = \binom{n}{2} - d(w).$$

Proof. Let $\lambda \vdash n$ with conjugate partition λ' . Consider the nilpotent orbit $\mathbb{O}_\lambda \subset \mathfrak{gl}_n$ of Jordan type λ . By Proposition 4.9, its dimension is given by

$$(4.8) \quad \dim \mathbb{O}_\lambda = n^2 - \sum_i (\lambda'_i)^2.$$

Let $\mathbb{O}_{\text{reg}}$ denote the regular nilpotent orbit (corresponding to the Jordan type $(1, 1, \dots, 1)$), which has dimension

$$\dim \mathbb{O}_{\text{reg}} = n^2 - n.$$

Then we can rewrite

$$(4.9) \quad \sum_i (\lambda'_i)^2 = n + 2 \sum_j \binom{\lambda'_j}{2} = n + 2b(\lambda),$$

so that

$$(4.10) \quad \dim \mathbb{O}_\lambda = n^2 - \sum_i (\lambda'_i)^2 = n^2 - n - 2b(\lambda) = \dim \mathbb{O}_{\text{reg}} - 2b(\lambda).$$

By Corollary 4.8, $d(w) = \frac{1}{2} \dim \mathbb{O}_\lambda$. Combining the above, we get

$$(4.11) \quad b(\lambda) = \frac{\dim \mathbb{O}_{\text{reg}}}{2} - d(w).$$

Noting that $\binom{n}{2} = \frac{n^2 - n}{2} = \frac{\dim \mathbb{O}_{\text{reg}}}{2}$, we obtain

$$(4.12) \quad b(\lambda) = \ell(w_0) - d(w).$$

□

By the formula (4.6), one can directly check the following.

Lemma 4.11. *For any $\mu \leq \lambda$ in the dominance order we have $b(\mu) \leq b(\lambda)$*

4.4. The $q \rightarrow \infty$ limit. We arrive at our main theorem, which is the following combinatorial explanation of the behavior of the Markov chain on S_n for large q .

Theorem 4.12. *For $w, z \in S_n$, writing λ for the shape of $P(w)$, we have*

$$(4.13) \quad \lim_{q \rightarrow \infty} P_{S_n}(w, z) = \begin{cases} \frac{1}{f^\lambda} & P(w) = P(z), \\ 0 & P(w) \neq P(z). \end{cases}$$

In fact,

$$P_{S_n}(w, z) = \frac{1}{f^\lambda} \delta_{P(w), P(z)} + O(1/q).$$

In this section, we develop some intermediate results which will be required for the proof of Theorem 4.12.

Lemma 4.13. *For any $z \in S_n$, the cardinality of the stabilizer of $\mathfrak{n}_z$ under the adjoint action of U on the Grassmannian of $\binom{n}{2} - \ell(z)$ -dimensional linear subspaces of $U \cdot \mathfrak{n}^z$ is*

$$(4.14) \quad \frac{|U| |\mathfrak{n}_z|}{|V_z|}$$

Proof. Consider the action of U by conjugation on subspaces of V_z . We note that $|V_z| = |\mathfrak{n}_z| |\mathbb{O}|$, where $\mathbb{O}$ is the orbit of U on $\mathfrak{n} \cap \mathfrak{n}^z$ as a subset of the Grassmannian. This means $|\mathbb{O}| = |V_z|/|\mathfrak{n}_z|$. On the other hand, by the orbit-stabilizer theorem, the stabilizer has cardinality $|U|/|\mathbb{O}|$, giving the result. $\square$

Proposition 4.14. *For any $w \in S_n$, we have*

$$(4.15) \quad \lim_{q \rightarrow \infty} \frac{1}{|U_w|} \sum_{a \in U_w} \frac{q^{b(\lambda)}}{|X_a|} = \frac{1}{f^\lambda},$$

where λ is the shape of $P(w)$.

Proof. Let $w \in S_n$ and let λ be the shape of $P(w)$. We note that the statement remains identical if U_w is replaced by $\mathfrak{n}_w = U_w - \text{id}_{n \times n}$. Let $\mathfrak{n}_w^\circ$ be the generic subset of $\mathfrak{n}_w$ consisting of elements with Jordan type λ .

First note that for generic $a \in V_w$, we have by Proposition 2.5 and Lemma 4.5 that $|X_a| = Q_{(1^n)}^\lambda(q)$, which has degree $b(\lambda)$ by Lemma 4.10. Further, by Lemmas 4.5 and 4.11, for any $a' \in V_w$ not generic in this sense, we have that $|X_a|$ is a polynomial in q with degree at least $b(\lambda)$. This means that, up to some $O(1/q)$ terms, we have

$$(4.16) \quad \frac{1}{|\mathfrak{n}_w|} \sum_{a \in U_w} \frac{q^{b(\lambda)}}{|X_a|} \approx \frac{1}{|\mathfrak{n}_w|} \sum_{a \in \mathfrak{n}_w^\circ} \frac{q^{b(\lambda)}}{|X_a|}$$

$$(4.17) \quad = \frac{|\mathfrak{n}_w^\circ|}{|\mathfrak{n}_w|} \cdot \frac{q^{b(\lambda)}}{Q_{(1^n)}^\lambda(q)}$$

$$(4.18) \quad \approx \frac{q^{b(\lambda)}}{Q_{(1^n)}^\lambda(q)}$$

with the last step holding since the number of elements of U_w° , an open subset of the irreducible variety V_w , is $q^{d(w)}$ up to a linear combination of smaller powers of q . Recalling as in (2.2) that

$$q^{-b(\lambda)} Q_{(1^n)}^\lambda(q) = X_{(1^n)}^\lambda(1/q),$$

by the Taylor series expansion of (4.18) at $q = \infty$, we see that its leading (degree zero) term is exactly $(X_{(1^n)}^\lambda(0))^{-1}$. Since by Proposition 2.6, $X_{(1^n)}^\lambda(0) = f^\lambda$, putting this together gives that

$$\begin{aligned} \lim_{q \rightarrow \infty} \frac{1}{q^{d(w)-b(\lambda)}} \sum_{a \in V_w} \frac{1}{|X_a|} &= \lim_{q \rightarrow \infty} \frac{1}{q^{-b(\lambda)} Q_{(1^n)}^\lambda(q)}, \\ &= \frac{1}{X_{(1^n)}^\lambda(0)} \\ &= \frac{1}{f^\lambda}, \end{aligned}$$

as claimed in the proposition. $\square$

We now use this estimate to prove Theorem 4.12.

Proof of Theorem 4.12. Suppose that $w, z \in W$ are such that $V_w = V_z$. Then by Lemma 4.13, Lemma 4.10, and Proposition 4.14 applied in succession, we have

$$\begin{aligned} P_{S_n}(w, z) &= \frac{1}{|U_z|} \sum_{u \in U} P(w, uz) \\ &= \frac{1}{|U_z||U_w|} \sum_{u \in U} \sum_{a \in U_w \cap uU_z u^{-1}} \frac{1}{|X_a|} \\ &= \frac{1}{|\mathbf{n}_z||\mathbf{n}_w|} \sum_{u \in U} \sum_{a \in \mathbf{n}_w \cap u \cdot \mathbf{n}_z} \frac{1}{|X_a|} \\ &\approx \frac{1}{|\mathbf{n}_z||\mathbf{n}_w|} \cdot \frac{|U||\mathbf{n}_z|}{|\mathbf{n}_z|} \sum_{a \in \mathbf{n}_w \cap V_z} \frac{1}{|X_a|} \\ &= \frac{1}{|\mathbf{n}_w|} \cdot \frac{|U|}{|V_z|} \sum_{a \in \mathbf{n}_w \cap V_w} \frac{1}{|X_a|} \\ &\approx \frac{1}{|\mathbf{n}_w|} \sum_{a \in \mathbf{n}_w} \frac{q^{\ell(w_0)-d(w)}}{|X_a|} \\ &= \frac{1}{|\mathbf{n}_w|} \sum_{a \in U_w} \frac{q^{b(\lambda)}}{|X_a|} \\ &\approx \frac{1}{f^\lambda}, \end{aligned}$$

up to terms of order $O(1/q)$, for λ the shape of $P(w)$. A similar direct computation shows that

$$(4.19) \quad P_{S_n}(w, z) = O(1/q)$$

whenever $P(w) \neq P(z)$. Alternatively, we note that for any fixed $w \in S_n$, the above implies that

$$(4.20) \quad \lim_{q \rightarrow \infty} \sum_{\substack{z \in S_n \\ P(w)=P(z)}} P_{S_n}(w, z) = f^\lambda \cdot \frac{1}{f^\lambda} = 1,$$

and so we must have $P_{S_n}(w, z) = 0$ for all z for which $P(w) \neq P(z)$. $\square$

5. GENERAL TYPE BEHAVIOR: ORBITAL VARIETIES AND STEINBERG CELLS

Although our result for $\mathrm{GL}_n(\mathbb{F}_q)$ can be stated in terms of the combinatorics of the Robinson–Schensted correspondence, the behavior we will describe in Theorem 5.9 is a phenomenon which occurs in general type, for G a finite Chevalley group. To state and prove it, we introduce the notion of Steinberg cells, which will serve as analogues of the subsets of S_n determined by a fixed choice of $P(w)$ or $Q(w)$, but for general-type Weyl groups W . Once we set up the geometric analogues of the combinatorial objects studied in Section 4 in general type, we can immediately use the same analysis as in the proof of Theorem 4.12 to establish its general-type analogue in the form of Theorem 5.9.

5.1. Setup for Chevalley groups of general type. One of the great achievements of twentieth-century mathematics is the unified treatment of a vast class of finite simple groups as “finite groups of Lie type.” Of course $\mathrm{GL}_n(\mathbb{F}_q)$ is among them (it has a simple quotient) as are the other classical orthogonal, unitary and symplectic groups over $\mathbb{F}_q$. These along with with twisted variants and the exceptional groups G_2, F_4, E_6, E_7, E_8 have appearances and applications in all corners of mathematics. A readable introduction to all that is needed for the present paper is in [3, §1]. We now recall how to construct such a group over $\mathbb{F}_q$ in any Dynkin type, and we will see that our results concerning all of these groups can be treated at once with unified arguments.

Let $\mathfrak{g}$ be a semisimple Lie algebra over $\mathbb{C}$. A construction of Chevalley [5] ensures that one can construct an affine group scheme $\mathbf{G}$ over $\mathbb{Z}$ corresponding to $\mathfrak{g}$. We fix a Cartan subgroup $\mathbf{T}$, a Borel subgroup $\mathbf{B} \subset \mathbf{G}$ containing $\mathbf{T}$, and we let $\mathbf{U} \subset \mathbf{G}$ denote its unipotent radical. These have corresponding Lie subalgebras $\mathfrak{h}, \mathfrak{b}, \mathfrak{n}$ of $\mathfrak{g}$.

Associated to $\mathfrak{g}$ is a root system Φ ; we write Φ^+ for its positive roots. Let $W = N_{\mathbf{G}}(\mathbf{T})/\mathbf{T}$ be the associated Weyl group, and let $\ell : W \rightarrow \mathbb{Z}_{\geq 0}$ be the length function. Let $w_0 \in W$ be the longest element. For any $\alpha \in \Phi^+$, there is a corresponding 1-parameter subgroup $\mathbf{U}_\alpha$ of $\mathbf{U}$; similarly let $\mathfrak{g}_\alpha$ denote the corresponding root subspace of $\mathfrak{g}$.

Fix once and for all a prime power q , and let $G = \mathbf{G}(\mathbb{F}_q)$, $B = \mathbf{B}(\mathbb{F}_q)$, $U = \mathbf{U}(\mathbb{F}_q)$, and $U_\alpha = \mathbf{U}_\alpha(\mathbb{F}_q)$ for any $\alpha \in \Phi^+$.

5.1.1. The Burnside process on the flag variety. Letting

$$X = \mathbf{G}(\mathbb{F}_q)/\mathbf{B}(\mathbb{F}_q) = G/B,$$

we continue to call this the *flag variety* of $\mathbf{G}$ over the finite field $\mathbb{F}_q$. In this setting we again have the Bruhat decomposition.

Proposition 5.1 (see [3], §1.10). *There is a decomposition*

$$X = \coprod_{w \in W} BwB = \coprod_{w \in W} UwB,$$

where for each $w \in W$, $UwB \cong \mathbb{A}_{\mathbb{F}_q}^{\ell(w)}$ is thought of as a subset of X and so $|UwB| = q^{\ell(w)}$.

We now study the Burnside process for the action of U on X in this more general setup. We again will write $P(x, y)$ for the transition probabilities, and since by Proposition 2.1 the orbits for this action are indexed by W , for $w, z \in W$ we will

use $P_W(w, z)$ to denote the transition probabilities for the corresponding Burnside process on the Weyl group.

5.2. Stabilizers in U . We now give a more general description of $\text{stab}_U(wB)$ for any $w \in W$, generalizing Proposition 2.3. Given $w \in W$, we again set $U^w = wUw^{-1} \subset G$, $\mathfrak{n}^w = w\mathfrak{n}w^{-1} \subset \mathfrak{g}$, and $U_w = \text{stab}_U(wB) = U \cap U^w$.

Proposition 5.2 (see [3, §2.5]). *For any $w \in W$, the stabilizer U_w is the subgroup of U generated by those root subgroups U_α with $\alpha \in \Phi^+$ such that $w^{-1}\alpha \in \Phi^+$. Equivalently,*

$$U_w = \prod_{\substack{\alpha \in \Phi^+ \\ w^{-1}\alpha \in \Phi^+}} U_\alpha.$$

Similarly,

$$\mathfrak{n}_w = \mathfrak{n} \cap \mathfrak{n}^w = \bigoplus_{\substack{\alpha \in \Phi^+ \\ w^{-1}\alpha \in \Phi^+}} \mathfrak{g}_\alpha.$$

In particular, since the cardinality of the indexing set in the product and direct sum which appear in Proposition 5.2 is exactly $\ell(w_0) - \ell(w) = \ell(w_0w)$, we will use that for any $w \in W$,

$$(5.1) \quad |U_w| = q^{\ell(w_0) - \ell(w)}.$$

5.3. General type behavior. As in Section 4, for any $w \in W$ let $V_w = U \cdot \mathfrak{n}_w$.

Definition 5.3. For $w \in W$, the *right Steinberg cell* $\mathcal{S}(w)$ is the set of all $z \in W$ such that

$$V_z = V_w$$

That is, $\mathcal{S}(w)$ consists of all $z \in W$ for which the U -orbit of $\mathfrak{n} \cap \mathfrak{n}^z$ for the adjoint action coincides with that of $\mathfrak{n} \cap \mathfrak{n}^w$.

Definition 5.4. In [2], it is explained that $\mathfrak{n}_w$ has a dense open subset consisting of elements x such that $\overline{V_w}$, considered now as a variety over $\overline{\mathbb{F}}_q$ is an irreducible component of $\overline{\mathbb{O}_x \cap \mathfrak{n}}$, where $\mathbb{O}_x$ is the nilpotent orbit $G \cdot x$. Let $\mathfrak{n}_w^\circ$ be this set.

Definition 5.5. Suppose $u \in \mathfrak{n}_w^\circ$. Let $b(w) = \dim X_u$, and let f^w be the number of irreducible components of X_u .

The following result, expressing a relationship between Springer fibers for $u \in U_w$ and the geometry of the varieties $V_w = U \cdot \mathfrak{n} \cap \mathfrak{n}^w$, is explained in [2, Appendix B].

Proposition 5.6. *For $u \in \mathfrak{n}_w^\circ$, the number of irreducible components of X_u is equal to $|\mathcal{S}(w)|$. Further, for any $u \in \mathfrak{n}_w$,*

$$\dim X_u \leq b(w).$$

Example 5.7. In Type A , we note that $b(w) = b(\lambda)$. Further, our comments after Proposition 2.5 guarantee that in this case, the irreducible components of X_u for $u \in \mathfrak{n}_w^\circ$ is equal to $|\mathcal{S}(w)| = f^\lambda$ for λ the shape of $P(\lambda)$.

The following is then a direct consequence of Proposition 5.6.

Corollary 5.8. *For any $w \in W$,*

$$(5.2) \quad \lim_{q \rightarrow \infty} \frac{1}{|U_w|} \sum_{u \in U_w} \frac{q^{b(w)}}{|X_u|} = \frac{1}{f^w}$$

The very same proof we used to establish Theorem 4.12 can then be used to show the following generalization.

Theorem 5.9. *For any $w, z \in W$,*

$$\lim_{q \rightarrow \infty} P_W(w, z) = \begin{cases} \frac{1}{f^{\mathcal{S}(w)}} & z \in \mathcal{S}(w) \\ 0 & z \notin \mathcal{S}(w) \end{cases}$$

In fact,

$$P_W(w, z) = \frac{1}{f^\lambda} \delta_{\mathcal{S}(w), \mathcal{S}(z)} + O(1/q).$$

6. CONSEQUENCES FOR MIXING TIMES

We now explain that the leading-order behavior of the Burnside process described in Sections 4 and 5 gives a lower bound on the mixing time.

For a given Weyl group W , let $\mathcal{S}_1, \dots, \mathcal{S}_k$ be an enumeration of the right Steinberg cells. By the analysis done in Section 5, there exists some constant C' not depending on q such that for any w, z with $\mathcal{S}(w) \neq \mathcal{S}(z)$,

$$(6.1) \quad P_W(w, z) \leq \frac{C'}{q}.$$

In other words, there is a constant C not depending on q such that for any $w \in W$,

$$(6.2) \quad P_W(\mathcal{S}(w), W - \mathcal{S}(w)) \leq \frac{C}{q}.$$

We will use this to prove the following lower bound on the mixing time of the Markov chain.

Theorem 6.1. *For any $\varepsilon > 0$, the ε -mixing time $\tau_{\text{mix}}(\varepsilon)$ of the Burnside process on W satisfies*

$$\tau_{\text{mix}}(\varepsilon) \geq \frac{q}{2C} \log \left(\frac{1}{2\varepsilon} \right) = \Omega(q \log(1/\varepsilon)).$$

Proof. Pick any Steinberg cell S and note that $\pi(S) \leq 1/2$; since the stationary distribution of the chain is uniform, we have $\pi(x) = 1/|W|$ for all $w \in W$ and hence $\pi(S) = |S|/|W|$. For any $w \in S$, we know the probability $P(w, S^c)$ of leaving S in one step is bounded by C/q . Now consider the stationary flow out of S :

$$Q(S, S^c) = \sum_{w \in S} \pi(w) P(w, S^c) \leq \frac{C}{q} \pi(S).$$

Therefore the conductance of S satisfies

$$\Phi(S) = \frac{Q(S, S^c)}{\pi(S)} \leq \frac{C}{q}.$$

Since $\Phi = \min_{T: 0 < \pi(T) \leq 1/2} \Phi(T)$, we conclude $\Phi \leq C/q$.

By the standard conductance lower bound (see, e.g., [17, Theorem 7.4]),

$$\tau_{\text{mix}}(\varepsilon) \geq \frac{1}{2\Phi} \ln \frac{1}{2\varepsilon}.$$

Plugging in $\Phi \leq C/q$ yields

$$\tau_{\text{mix}}(\varepsilon) \geq \frac{q}{2C} \ln \frac{1}{2\varepsilon} = \frac{q}{2C(|W| - 1)} \ln \frac{1}{2\varepsilon}.$$

In particular, $\tau_{\text{mix}}(\varepsilon) = \Omega(q \log(1/\varepsilon))$ as claimed. $\square$

Theorem 6.1 guarantees that, by choosing q large, the Burnside process on W defined in the present paper provides a natural family of Markov chains on the Weyl group which can be made to mix “arbitrarily slowly.”

7. EXAMPLES AND SIMULATIONS

7.1. Explicit transition probabilities for $\mathrm{GL}_3(\mathbb{F}_q)$. Using a computer, we produced the transition matrix for $G = \mathrm{GL}_3(\mathbb{F}_q)$ for arbitrary q , expressing each transition probability as a rational function in q . In this case, $W \cong S_3$, and we again label its elements in the order

$$\{e, s_1, s_2, s_2s_1, s_1s_2, w_0\}.$$

The transition matrix is as follows:

$$\begin{pmatrix} \frac{q^3}{q^3+2q^2+2q+1} & \frac{q^3+2q^2+q-1}{2q^4+5q^3+6q^2+4q+1} & \frac{q^3+2q^2+q-1}{2q^4+5q^3+6q^2+4q+1} & \frac{q^3+q^2+1}{2q^4+5q^3+6q^2+4q+1} & \frac{q^3+q^2+1}{2q^4+5q^3+6q^2+4q+1} & \frac{1}{q^3+2q^2+2q+1} \\ \frac{q^3+2q^2+q-1}{2q^4+5q^3+6q^2+4q+1} & \frac{q^4+2q^3+q^2-q}{2q^4+5q^3+6q^2+4q+1} & \frac{q^3+q^2+1}{2q^4+5q^3+6q^2+4q+1} & \frac{q^4+q^3+q}{2q^4+5q^3+6q^2+4q+1} & \frac{1}{q^3+2q^2+2q+1} & \frac{q}{q^3+2q^2+2q+1} \\ \frac{q^3+2q^2+q-1}{2q^4+5q^3+6q^2+4q+1} & \frac{q^3+q^2+1}{2q^4+5q^3+6q^2+4q+1} & \frac{q^4+2q^3+q^2-q}{2q^4+5q^3+6q^2+4q+1} & \frac{1}{q^3+2q^2+2q+1} & \frac{q^4+q^3+q}{2q^4+5q^3+6q^2+4q+1} & \frac{q}{q^3+2q^2+2q+1} \\ \frac{q^3+q^2+1}{2q^4+5q^3+6q^2+4q+1} & \frac{q^4+q^3+q}{2q^4+5q^3+6q^2+4q+1} & \frac{1}{q^3+2q^2+2q+1} & \frac{q^4+q^3+2q^2-1}{2q^4+5q^3+6q^2+4q+1} & \frac{q}{q^3+2q^2+2q+1} & \frac{q^2}{q^3+2q^2+2q+1} \\ \frac{q^3+q^2+1}{2q^4+5q^3+6q^2+4q+1} & \frac{1}{q^3+2q^2+2q+1} & \frac{q^4+q^3+q}{2q^4+5q^3+6q^2+4q+1} & \frac{q}{q^3+2q^2+2q+1} & \frac{q^4+q^3+2q^2-1}{2q^4+5q^3+6q^2+4q+1} & \frac{q^2}{q^3+2q^2+2q+1} \\ \frac{1}{q^3+2q^2+2q+1} & \frac{q}{q^3+2q^2+2q+1} & \frac{q}{q^3+2q^2+2q+1} & \frac{q^2}{q^3+2q^2+2q+1} & \frac{q^2}{q^3+2q^2+2q+1} & \frac{q^3}{q^3+2q^2+2q+1} \end{pmatrix}.$$

Even in this small example, we see that the transition probabilities are difficult to understand algebraically, and exhibit some of the “wild” behavior described in Section 2.3.2. However, we note that looking at their leading terms reveals the combinatorial patterns described in Section 4. Indeed, the $q \rightarrow \infty$ limit of this matrix is the matrix

$$\begin{array}{c} e \quad s_1 \quad s_2 \quad s_2s_1 \quad s_1s_2 \quad w_0 \\ \begin{array}{c} e \\ s_1 \\ s_2 \\ s_2s_1 \\ s_1s_2 \\ w_0 \end{array} \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} & 0 & 0 \\ 0 & 0 & \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} & 0 & 0 \\ 0 & 0 & \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \end{array}$$

Unlike the wild entries in the first matrix, this limiting matrix can be explained by Theorem 4.12 and the fact that under the Robinson–Schensted correspondence,

$$(7.1) \quad P(e) = \begin{array}{|c|c|c|} \hline 1 & 2 & 3 \\ \hline \end{array}$$

$$(7.2) \quad P(s_1), P(s_2s_1) = \begin{array}{|c|c|} \hline 1 & 3 \\ \hline 2 & \\ \hline \end{array}$$

$$(7.3) \quad P(s_2), P(s_1s_2) = \begin{array}{|c|c|} \hline 1 & 2 \\ \hline 3 & \\ \hline \end{array}$$

$$(7.4) \quad P(s_1s_2s_1) = \begin{array}{|c|} \hline 1 \\ \hline 2 \\ \hline 3 \\ \hline \end{array}.$$

The beautiful mathematical structure underlying the transition matrix above suggests that representation theory might offer a way diagonalizing these Markov chains. Indeed, after a lot of massaging, this happened for the simplest case of the Burnside process [9]. One hallmark of a “nice diagonalization” is “nice eigenvalues” (something like rational numbers). Alas, for the example above, when $q = 2$, the eigenvalues are $0, 1, \frac{2}{5}$, and the three irrational roots of the cubic polynomial

$$525x^3 - 315x^2 + 50x - 2.$$

These don’t inspire hope.

That being said, inspection of the matrix for general q when $n = 3$ reveals that the second-largest eigenvalue is $\frac{2q-2}{2q+1}$ with (right) eigenvector $(0, 1, -1, 1, -1, 0)$. Indeed, once observed, it is easy to check with a computer that this is an eigenvalue-eigenvector pair. There is an eigenvalue 0 with eigenvector $(0, 1, -1, -1, 1, 0)$, and of course 1 is an eigenvalue with eigenvector $(1, 1, 1, 1, 1, 1)$. This gives three roots of the characteristic polynomial. The remaining eigenvalues are roots of a cubic and standard analysis shows that for any q , each such root is smaller than $\frac{2q-2}{2q+1}$. All eigenvalues are non-negative because of the structure of these two step chains. Using these observations, we can give sharp bounds on the rate of convergence of the chain to stationarity on S_3 for general q . The argument shows that order $2q+1$ steps are necessary and sufficient for convergence, from any start.

Proposition 7.1. *For the Burnside process on flags in $\text{GL}_3(\mathbb{F}_q)$, the Markov chain on S_3 after l steps, satisfies*

$$\frac{1}{2} \left(1 - \frac{3}{2q+1}\right)^l \leq \|K^l(x, \cdot) - \pi(\cdot)\| \leq \sqrt{\frac{3}{2}} \left(1 - \frac{3}{2q+1}\right)^l$$

Above, the total variation distance is

$$\sum_{w \in S_3} \left| K^l(x, w) - \frac{1}{6} \right|.$$

The upper bound holds for all $x \in S_3$ and all $l \geq 1$. The lower bound holds for all x except id and w_0 and all $l \geq 1$.

Proof. The upper bound follows from the general upper bound (for any reversible chain on a finite state space with stationary distribution $\pi(y)$)

$$4\|K^l(x, \cdot) - \pi(\cdot)\|^2 \leq \frac{1}{\pi_*} \beta_*^{2l}$$

Here, $\pi_* = \min_{y \in W} \pi(y)$ and β_* is the absolute value of the second eigenvalue. In the present case, $\pi_* = \frac{1}{6}$ and $\beta_* = \frac{2q-2}{2q+1}$. The claimed upper bound follows by elementary manipulation.

For the lower bound, use [8, Lemma 2.1] which states if $f(x)$ is an eigenvector with eigenvalue beta of a Markov chain with $\sup |f(x)| = f_*$, then, for any x and l ,

$$\|K^l(x, \cdot) - \pi(\cdot)\| \geq \frac{f(x)}{f_*} |\beta|^l.$$

In the present case, the eigenfunction for $\frac{2q-2}{2q+1}$ takes values in $\{0, 1, -1\}$ with mean zero. $\square$

We remain curious about whether clean formulas for the second eigenvalue and its corresponding eigenvector may exist in other types.

7.2. Simulations in larger types. Using the algorithm described in Section 3, we implemented the Burnside process in Type A and ran some simulations for $\mathrm{GL}_4(\mathbb{F}_q)$ and $\mathrm{GL}_5(\mathbb{F}_q)$, some of which we include here to better illustrate the behavior we describe in Section 4. All of the figures in this sections are histograms depicting the number of times the Markov chain visited a given element in S_n . In these figures, we let t be the number of steps for which the simulation was run.

First, we note that when q is small, we expect sampling from the Burnside process to be very close to sampling from the uniform distribution on S_n , as shown in Figures 1 and 2. If q is instead made to be large (compared to the number of steps), we then expect the uniform-like sampling property to disappear, instead being replaced by uniform sampling on the Steinberg cell, which in this case is the set of permutations with the same P -symbol as the initial element; this can be seen in Figures 3 and 4.

Finally, Figure 5 shows an example where q and t have the same order of magnitude. Here we see that after remaining in a certain Steinberg cell for many steps, sometimes the chain exits and moves to a different Steinberg cell.

FIGURE 1. The histogram for $q = 3$, $n = 4$, $t = 1000$, with the chain starting at the permutation 3 2 1 4.





FIGURE 2. The histogram for $n = 5$, $q = 5$, $t = 1000$, with the chain starting at the permutation 3 2 1 4.



8. FINAL COMMENTS

There is much more to say about the Burnside process on G/B . First of all, we are interested in explicit algorithms as in Section 3 for finite Chevalley groups outside of Type A . We are hopeful that a similar approach as in the present paper

FIGURE 3. The histogram for $q = 1997, n = 4, t = 1000$, starting at 3 2 1 4. We note that 3 elements share the same P -symbol with 3 2 1 4, so the behavior here is described by Theorem 4.12.





FIGURE 4. The histogram for $n = 5, q = 20011, t = 1000$, starting at 3 2 1 4 5. We note that 4 elements share the same P -symbol with 3 2 1 4 5 in this case.




could accomplish this for classical types, and curious as to whether such an explicit algorithm is possible to describe in any exceptional types (even, for example, for G_2). We note that an algorithm implementing the first step of the Burnside process is easy in any type, while the real work is in uniform sampling from Springer fibers.

In any type where such an algorithm can be described, one can then apply Theorem 5.9 as follows. This theorem says that by setting q to be very large, such an algorithm would effectively allow for near-uniform sampling from any Steinberg cell. In [4], it was shown that access to a uniform sampling gives allows for many ways to obtain good estimates for the size of a set. Thus, such an algorithm would provide a direct computational way to answer combinatorial questions about Steinberg cells, e.g. how large is a given cell, and how many total cells are there? For example, this idea applied to our algorithm in Type A provides a (redundant) way to estimate the quantities f^λ , which of course already have precise formulas; this may be less redundant in other types where the combinatorics of Steinberg cells is much more complicated.

In other examples of the Burnside process, there have sometimes been “interpretable” descriptions for the Markov chain, see e.g. the examples in [7]. We are curious as to whether, even in Type A , such an interpretable description could be provided for the Burnside process on G/B studied in the present paper. Further, there is surely more to say about the mixing times for these chains than the lower bound we give in Section 6, and a more refined estimate would be interesting.

FIGURE 5. The histogram for $n = 5$, $t = 10000$, $q = 20011$, starting at $3\ 2\ 1\ 4\ 5$. In this example, the chain happened to visit four different Steinberg cells, staying in each one for a certain number of steps as pictured. The four “buckets” into which these data cluster are labelled by the pictured Young tableaux according to their color.



REFERENCES

- [1] Laurent Bartholdi and Persi Diaconis. An algorithm for uniform generation of unlabeled trees (Pólya trees), with an extension of Cayley’s formula, 2024. To appear, Forum Math. Sigma. URL: <https://arxiv.org/abs/2411.17613>, [arXiv:2411.17613](https://arxiv.org/abs/2411.17613).
- [2] W. Borho and J.-L. Brylinski. Differential operators on homogeneous spaces. III. Characteristic varieties of Harish-Chandra modules and of primitive ideals. *Invent. Math.*, 80(1):1–68, 1985. [doi:10.1007/BF01388547](https://doi.org/10.1007/BF01388547).
- [3] Roger W. Carter. *Finite groups of Lie type*. Wiley Classics Library. John Wiley & Sons, Ltd., Chichester, 1993.
- [4] Sourav Chatterjee, Persi Diaconis, and Susan Holmes. Estimating the size of a set using cascading exclusion, 2025. URL: <https://arxiv.org/abs/2508.05901>, [arXiv:2508.05901](https://arxiv.org/abs/2508.05901).
- [5] Claude Chevalley. *Classification des groupes algébriques semi-simples*. Springer-Verlag, Berlin, 2005. Collected works. Vol. 3, Edited and with a preface by P. Cartier, With the collaboration of Cartier, A. Grothendieck and M. Lazard.
- [6] David H. Collingwood and William M. McGovern. *Nilpotent orbits in semisimple Lie algebras*. Van Nostrand Reinhold Mathematics Series. Van Nostrand Reinhold Co., New York, 1993.
- [7] Persi Diaconis and Michael Howes. Random sampling of contingency tables and partitions: two practical examples of the Burnside process. *Stat. Comput.*, 35(6):Paper No. 181, 20, 2025. [doi:10.1007/s11222-025-10708-5](https://doi.org/10.1007/s11222-025-10708-5).
- [8] Persi Diaconis, Kshitij Khare, and Laurent Saloff-Coste. Gibbs sampling, exponential families and orthogonal polynomials. *Statist. Sci.*, 23(2):151–178, 2008. With comments and a rejoinder by the authors. [doi:10.1214/07-STS252](https://doi.org/10.1214/07-STS252).
- [9] Persi Diaconis, Andrew Lin, and Arun Ram. A curiously slowly mixing Markov chain, 2025. to appear.
- [10] Persi Diaconis and Chenyang Zhong. Hahn polynomials and the Burnside process. *Ramanujan J.*, 61(2):567–595, 2023. [doi:10.1007/s11139-021-00482-z](https://doi.org/10.1007/s11139-021-00482-z).
- [11] William Fulton. *Young tableaux*, volume 35 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1997. With applications to representation theory and geometry.
- [12] Emden R. Gansner. Acyclic digraphs, Young tableaux and nilpotent matrices. *SIAM J. Algebraic Discrete Methods*, 2(4):429–440, 1981. [doi:10.1137/0602046](https://doi.org/10.1137/0602046).
- [13] Curtis Greene. Some partitions associated with a partially ordered set. *J. Combinatorial Theory Ser. A*, 20(1):69–79, 1976. [doi:10.1016/0097-3165\(76\)90078-9](https://doi.org/10.1016/0097-3165(76)90078-9).
- [14] Roger Howe. A century of Lie theory. In *American Mathematical Society centennial publications, Vol. II (Providence, RI, 1988)*, pages 101–320. Amer. Math. Soc., Providence, RI, 1992.
- [15] David Kazhdan and George Lusztig. Representations of Coxeter groups and Hecke algebras. *Invent. Math.*, 53(2):165–184, 1979. [doi:10.1007/BF01390031](https://doi.org/10.1007/BF01390031).
- [16] Bertram Kostant. Flag manifold quantum cohomology, the Toda lattice, and the representation with highest weight ρ . *Selecta Math. (N.S.)*, 2(1):43–91, 1996. [doi:10.1007/BF01587939](https://doi.org/10.1007/BF01587939).
- [17] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov chains and mixing times*. American Mathematical Society, Providence, RI, 2009. With a chapter by James G. Propp and David B. Wilson. [doi:10.1090/mbk/058](https://doi.org/10.1090/mbk/058).
- [18] I. G. Macdonald. *Symmetric functions and Hall polynomials*. Oxford Mathematical Monographs. The Clarendon Press, Oxford University Press, New York, second edition, 1995. With contributions by A. Zelevinsky, Oxford Science Publications.
- [19] J. E. Paguyo. Mixing times of a burnside process Markov chain on set partitions, 2023. URL: <https://arxiv.org/abs/2207.14269>, [arXiv:2207.14269](https://arxiv.org/abs/2207.14269).
- [20] John Rahmani. Mixing times for the commuting chain on CA groups. *J. Theoret. Probab.*, 35(1):457–483, 2022. [doi:10.1007/s10959-020-01049-2](https://doi.org/10.1007/s10959-020-01049-2).
- [21] Konstanze Rietsch. Totally positive Toeplitz matrices and quantum cohomology of partial flag varieties. *J. Amer. Math. Soc.*, 16(2):363–392, 2003. [doi:10.1090/S0894-0347-02-00412-5](https://doi.org/10.1090/S0894-0347-02-00412-5).
- [22] C. Schensted. Longest increasing and decreasing subsequences. *Canadian J. Math.*, 13:179–191, 1961. [doi:10.4153/CJM-1961-015-3](https://doi.org/10.4153/CJM-1961-015-3).
- [23] Nicolas Spaltenstein. *Classes unipotentes et sous-groupes de Borel*, volume 946 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin-New York, 1982.

- [24] Robert Steinberg. An occurrence of the Robinson-Schensted correspondence. *J. Algebra*, 113(2):523–528, 1988. doi:[10.1016/0021-8693\(88\)90177-9](https://doi.org/10.1016/0021-8693(88)90177-9).
- [25] Julianna Tymoczko. The geometry and combinatorics of Springer fibers. In *Around Langlands correspondences*, volume 691 of *Contemp. Math.*, pages 359–376. Amer. Math. Soc., Providence, RI, 2017. doi:[10.1090/conm/691/13903](https://doi.org/10.1090/conm/691/13903).
- [26] Julianna S. Tymoczko. Linear conditions imposed on flag varieties. *Amer. J. Math.*, 128(6):1587–1604, 2006. URL: http://muse.jhu.edu/journals/american_journal_of_mathematics/v128/128.6tymoczko.pdf.