

A curved Brunn-Minkowski inequality for the symmetric group

Weerachai Neeranartvong

Department of Mathematics
Massachusetts Institute of Technology
Cambridge, MA.

weeracha@mit.edu

Jonathan Novak

Department of Mathematics
University of California
San Diego, CA

jinovak@ucsd.edu

Nat Sothanaphan

Department of Mathematics
Massachusetts Institute of Technology
Cambridge, MA

nsothana@mit.edu

Submitted: Jul 22, 2015; Accepted: Jan 28, 2016; Published: Feb 5, 2016

Mathematics Subject Classifications: 05A19, 05A20

Abstract

In this paper, we construct an injection $A \times B \rightarrow M \times M$ from the product of any two nonempty subsets of the symmetric group into the square of their midpoint set, where the metric is that corresponding to the conjugacy class of transpositions. If A and B are disjoint, our construction allows to inject two copies of $A \times B$ into $M \times M$. These injections imply a positively curved Brunn-Minkowski inequality for the symmetric group analogous to that obtained by Ollivier and Villani for the hypercube. However, while Ollivier and Villani's inequality is optimal, we believe that the curvature term in our inequality can be improved. We identify a hypothetical concentration inequality in the symmetric group and prove that it yields an optimally curved Brunn-Minkowski inequality.

Keywords: Brunn-Minkowski inequality; discrete curvature

1 Introduction

The classical Brunn-Minkowski inequality may be formulated as follows: given two compact nonempty sets $A, B \subset \mathbb{R}^n$, one has

$$\log |M_t| \geq (1-t) \log |A| + t \log |B|$$

for any $0 \leq t \leq 1$, where

$$M_t = \{(1-t)a + tb : a \in A, b \in B\}$$

is the set of t -midpoints of A and B , and $|\cdot|$ is Lebesgue measure. If \mathbb{R}^n is replaced by a smooth complete Riemannian manifold with positive Ricci curvature bounded below by $K > 0$, the Brunn-Minkowski inequality can be strengthened to

$$\log |M_t| \geq (1-t) \log |A| + t \log |B| + \frac{K}{2} t(1-t) d(A, B)^2, \quad (1)$$

where d is the Hausdorff distance and M_t the set of points of the form $\gamma(t)$ with γ a geodesic in X such that $\gamma(0) \in A$ and $\gamma(1) \in B$, see [4] and references therein. The heuristic here is that midpoint sets are larger in positively curved space than in flat space, and the degree of this distortion is controlled by Ricci curvature.

Y. Ollivier and C. Villani [4] have shown that a curved Brunn-Minkowski inequality analogous to (1) holds for the discrete hypercube \mathbb{Z}_2^N equipped with the Hamming distance. While the definition of t -midpoints in discrete space is somewhat messy, in the case $t = \frac{1}{2}$, at least, it is reasonable to define the midpoint set $M = M_{1/2}$ of $A, B \subseteq \mathbb{Z}_2^N$ to be the collection of $m \in \mathbb{Z}_2^N$ which satisfy $d(a, m) + d(m, b) = d(a, b)$ and $d(a, m) = d(m, b) + \varepsilon$, with $\varepsilon \in \{-1, 0, 1\}$, for some $(a, b) \in A \times B$. Adopting these definitions, Ollivier and Villani proved the following curved Brunn-Minkowski inequality for the hypercube [4, Theorem 1].

Theorem 1. *For any nonempty sets $A, B \subseteq \mathbb{Z}_2^N$,*

$$\log |M| \geq \frac{1}{2} \log |A| + \frac{1}{2} \log |B| + \frac{K}{8} d(A, B)^2,$$

where $K = \frac{1}{2N}$.

Ollivier and Villani moreover verify that the dependence of K on N in their result is optimal. As discussed in [4], this result supports the statement that the “discrete Ricci curvature” of \mathbb{Z}_2^N is of order N^{-1} .

For any $n \geq 2N$, there is an injective group homomorphism

$$\mathbb{Z}_2^N \longrightarrow S(n)$$

from the N -dimensional hypercube into the symmetric group of rank n determined by

$$\begin{aligned} e_1 &\mapsto (1 \mapsto 2) \\ e_2 &\mapsto (3 \mapsto 4) \\ &\vdots \\ e_N &\mapsto (2N-1 \mapsto 2N), \end{aligned}$$

where $e_i \in \mathbb{Z}_2^N$ is the bitstring in which all bits are zero except the i th bit, and $(2i-1 \mapsto 2i) \in S(n)$ is the transposition which swaps $2i$ and $2i-1$. If one equips $S(n)$ with the

metric induced by the word norm corresponding to the conjugacy class of transpositions (see Section 2 below for a detailed construction), this injection is an isometric embedding of \mathbb{Z}_2^N in $S(n)$. It is thus natural to seek an extension of Theorem 1 to the symmetric group, viewed as a metric space in this way. In this paper, we prove the following curved Brunn-Minkowski inequality for $S(n)$.

Theorem 2. *For any nonempty sets $A, B \subseteq S(n)$,*

$$\log |M| \geq \frac{1}{2} \log |A| + \frac{1}{2} \log |B| + \frac{K}{8} d(A, B)^2,$$

where $K = \frac{4 \log 2}{(n-1)^2}$.

The Brunn-Minkowski inequality presented in Theorem 2 is only slightly curved, and we believe that Theorem 2 in fact holds with $K = \frac{c}{n-1}$, c a positive constant. Although we do not prove this, we identify a hypothetical concentration inequality in the symmetric group which generalizes the hypercube concentration inequality of Ollivier and Villani [4, Corollary 6], and demonstrate that it implies an optimally curved Brunn-Minkowski inequality for the symmetric group.

2 Symmetric Group Basics

In this section we fix basic notation and terminology concerning the symmetric group $S(n)$. We identify $S(n)$ with its right Cayley graph as generated by the conjugacy class of transpositions. Thus $a, b, c, \dots \in S(n)$ are the vertices of our graph, and $\{a, b\}$ is an edge if and only if $a^{-1}b$ fixes all but two points of $\{1, \dots, n\}$. In this way $S(n)$ becomes a *graded graph*: it decomposes as the disjoint union

$$S(n) = \bigsqcup_{r=0}^{n-1} L_r,$$

where L_r is the set of permutations which factor into exactly $n - r$ disjoint cycles; each L_r is an independent set (contains no edges); and, finally, there exists an edge between L_r and $L_{r'}$ if and only if $|r - r'| = 1$. Figure 1 shows the case $n = 4$.

Each level L_r of $S(n)$ further decomposes as the disjoint union

$$L_r = \bigsqcup_{\substack{\lambda \vdash n \\ \ell(\lambda) = n-r}} C_\lambda,$$

where the union is over partitions λ of n with $n - r$ parts, and C_λ is the set of permutations with cycle type λ . The sets C_λ are the conjugacy classes of $S(n)$.

In this paper we make use of a decomposition of $S(n)$ which is finer than the usual decomposition into conjugacy classes. Given $p \in S(n)$, factor p into disjoint cycles, and present each cycle so that its leftmost element is its minimal element. That is, each cycle of p is presented in the form

$$(i_1 \mapsto i_2 \mapsto i_3 \mapsto \dots),$$




Figure 1: The Cayley graph of $S(4)$.

where $i_1 < \min\{i_2, i_3, \dots\}$. Next, list the cycles of p from left to right in increasing order of their minimal elements. Thus p is presented in the form

$$p = (i_1 \mapsto i_2 \mapsto i_3 \mapsto \dots)(j_1 \mapsto j_2 \mapsto j_3 \mapsto \dots)(k_1 \mapsto k_2 \mapsto k_3 \mapsto \dots) \dots,$$

where $i_1 < j_1 < k_1 < \dots$. We call this the *ordered cycle factorization* of p . Figure 1 displays the elements of $S(4)$ using their ordered cycle factorizations. We refer to the vector

$$(i_1, j_1, k_1, \dots)$$

as the *sequence of cycle minima* of p . The vector

$$(\mu_1, \mu_2, \mu_3, \dots)$$

of cycle lengths

$$p = \underbrace{(i_1 \mapsto i_2 \mapsto i_3 \mapsto \dots)}_{\text{length } \mu_1} \underbrace{(j_1 \mapsto j_2 \mapsto j_3 \mapsto \dots)}_{\text{length } \mu_2} \underbrace{(k_1 \mapsto k_2 \mapsto k_3 \mapsto \dots)}_{\text{length } \mu_3} \dots,$$

in the ordered cycle factorization of p is a composition μ of n which we call the *ordered cycle type* of p . We denote by $\ell(\mu)$ the number of parts of μ , so that $r = n - \ell(\mu)$ if $p \in L_r$. Given two permutations p, p' of the same ordered cycle type, there is a *unique* permutation u which both conjugates p into p' and transforms the sequence of cycle minima of p into that of p' .

Given a composition $\mu \vDash n$, we denote by \vec{C}_μ the set of all permutations whose ordered cycle type is μ . Then each conjugacy class C_λ in $S(n)$ decomposes as the disjoint union

$$C_\lambda = \bigsqcup_{\mu} \vec{C}_\mu,$$

where μ ranges over all compositions obtained by permuting the parts of λ . For the symmetric group $S(4)$, the successive decompositions we have discussed are

$$\begin{aligned} S(4) &= L_0 \sqcup L_1 \sqcup L_2 \sqcup L_3 \\ &= C_{(1,1,1,1)} \sqcup C_{(2,1,1)} \sqcup C_{(3,1)} \sqcup C_{(2,2)} \sqcup C_{(4)} \\ &= \vec{C}_{(1,1,1,1)} \sqcup \vec{C}_{(2,1,1)} \sqcup \vec{C}_{(1,2,1)} \sqcup \vec{C}_{(1,1,2)} \sqcup \vec{C}_{(3,1)} \sqcup \vec{C}_{(1,3)} \sqcup \vec{C}_{(2,2)} \sqcup \vec{C}_{(4)}. \end{aligned}$$

Each class \vec{C}_μ contains a canonical permutation p_μ , which acts by cyclically permuting the first μ_1 positive integers in the canonical way, cyclically permuting the next μ_2 positive integers in the canonical way, and so on. Given $p \in \vec{C}_\mu$, we denote by $u_p \in S(n)$ the unique permutation which both conjugates p_μ to p and transforms the sequence of cycle minima of p_μ into that of p .

We equip $S(n)$ with the graph theory distance d . Thus level L_r in the Cayley graph coincides with the sphere of radius r centred at the identity permutation $e \in S(n)$. The following properties of d are easily checked:

$$\begin{aligned} d(a, b) &= d(pap^{-1}, pbp^{-1}) \\ &= d(ab^{-1}, e) \\ &= d(e, a^{-1}b) \\ &= d(ap, bp) \\ &= d(pa, pb). \end{aligned}$$

In particular, the diameter of the Cayley graph is

$$\max\{d(a, b) : a, b \in S(n)\} = \max\{d(e, p) : p \in S(n)\} = n - 1.$$

We have already mentioned the fact that the set of permutations which lie on a geodesic path from the identity permutation e to an involution v is isometrically isomorphic to a hypercube whose dimension is half the size of the support of v . We will also make use of the fact that a permutation lies on a geodesic path from e to a forward cycle f if and only if it is a product of forward cycles which together induce a noncrossing partition of the support of f . A proof of this folklore result may be found in [3, Lecture 23].

3 Midpoint Calculus

In this section, we generalize the encoding/decoding formalism of Ollivier and Villani from the hypercube to the symmetric group. Where possible, we try to be consistent with the notation and terminology of [4].

3.1 Crossovers

Let $(a, b) \in S(n) \times S(n)$ be a pair of permutations, and let $M(a, b)$ be the corresponding midpoint set. Our first observation is that $M(a, b)$ is the isometric image of a “standard” set of midpoints. More precisely, let μ be the ordered cycle type of $a^{-1}b$, and let $\text{Cr}(\mu)$ denote the set of midpoints of e , the identity permutation, and p_μ , the canonical permutation of ordered cycle type μ . Adopting the terminology of [4], we call the elements of $\text{Cr}(\mu)$ *crossovers*, or μ -*crossovers* to be precise. Consider the function

$$S(n) \longrightarrow S(n)$$

defined by

$$x \mapsto au_{a^{-1}b}xu_{a^{-1}b}^{-1},$$

where $u_{a^{-1}b}$ is the unique permutation which conjugates p_μ to $a^{-1}b$ and transforms the sequence of cycle minima of p_μ into that of $a^{-1}b$. This function is an isometry, being composed of rotation by $u_{a^{-1}b}$ followed by translation by a . Moreover, under this mapping

$$e \mapsto a, \quad p_\mu \mapsto b.$$

Thus the mapping restricts to a bijection

$$\text{Cr}(\mu) \longrightarrow M(a, b).$$

We write

$$\varphi_c(a, b) = au_{a^{-1}b}cu_{a^{-1}b}^{-1},$$

and view

$$\varphi_c(a, b), \quad c \in \text{Cr}(\mu)$$

as a parameterization of the locus $M(a, b)$ by a “standard” set of midpoints. Following the terminology of [4], we call $\varphi_c(a, b) \in M(a, b)$ the midpoint of a and b *encoded* by the crossover $c \in \text{Cr}(\mu)$.

Here is an example to help solidify these definitions. Consider $a = (1)(23)(4)$ and $b = (1324)$ as elements of $S(4)$. We have

$$M(a, b) = \{(132)(4), (1)(243), (14)(23)\}.$$

Since $a^{-1}b = (124)(3)$, the ordered cycle type of $a^{-1}b$ is $(3, 1)$. So $p_\mu = (123)(4)$ and the set of μ -crossovers is

$$\text{Cr}(\mu) = \{(12)(3)(4), (13)(2)(4), (1)(23)(4)\}.$$

Since $u_{a^{-1}b} = (1)(2)(34)$, we obtain a parametrization of $M(a, b)$ as

$$\begin{aligned} \varphi_{(12)(3)(4)}(a, b) &= (132)(4), \\ \varphi_{(13)(2)(4)}(a, b) &= (14)(23), \\ \varphi_{(1)(23)(4)}(a, b) &= (1)(243). \end{aligned}$$

Thus, for example, $(132)(4)$ is the midpoint of a and b encoded by the crossover $(12)(3)(4)$.

3.2 Duality

There is a natural geometric operation on crossovers: we view a crossover as the lower half of a geodesic path from the identity to the canonical permutation with a given ordered cycle type, and map it to the corresponding upper half. More precisely, given $c \in \text{Cr}(\mu)$, its *dual* c^\vee is defined by

$$c^\vee := c^{-1}p_\mu.$$

We now establish some technical properties of the operation $c \mapsto c^\vee$ which will be needed below. First is the basic but important closure property.

Proposition 3. *$\text{Cr}(\mu)$ is closed under taking duals.*

Proof. Let c be a midpoint of e and p_μ ; we have to check that c^\vee is a midpoint of e and p_μ . We have

$$d(e, c^\vee) = d(e, c^{-1}p_\mu) = d(c, p_\mu)$$

and

$$d(c^\vee, p_\mu) = d(c^{-1}p_\mu, p_\mu) = d(c^{-1}, e) = d(e, c).$$

Thus

$$d(e, c^\vee) + d(c^\vee, p_\mu) = d(e, c) + d(c, p_\mu) = d(e, p_\mu),$$

and

$$d(e, c^\vee) - d(c^\vee, p_\mu) = d(c, p_\mu) - d(e, c) = \varepsilon$$

with $\varepsilon \in \{-1, 0, 1\}$, as required. \square

Note that while the map $\text{Cr}(\mu) \rightarrow \text{Cr}(\mu)$ defined by $c \mapsto c^\vee$ is bijective, it is not involutive: the dual of the dual of c is c conjugated by p_μ^{-1} . For future use, we extend the duality operation from points to sets: given $C \subseteq \text{Cr}(\mu)$, we define

$$C^\vee := \{c^\vee : c \in C\}.$$

Next is the following important property of crossover duals.

Lemma 4. *If $c \in \text{Cr}(\mu)$, then $c^{-1}c^\vee$ has both the same ordered cycle type and the same sequence of cycle minima as p_μ .*

Proof. First note that $c^{-1}c^\vee = c^{-2}p_\mu$. Since c lies on a geodesic path linking e to p_μ , each cycle of c is a subcycle of some cycle of p_μ . Since the cycles of p_μ are intervals, i.e. consist of consecutive numbers, our task reduces to proving the following general statement: whenever $c_1c_2\dots$ is a product of forward cycles which induce a noncrossing partition of $\{1, \dots, k\}$, the product

$$c_1^{-2}c_2^{-2}\dots(1 \mapsto \dots \mapsto k)$$

is a cyclic permutation of the numbers $1, \dots, k$. If this statement holds, then left multiplication of p_μ by c^{-2} will change neither the cycle structure nor cycle minima of p_μ .

Note that we can assume each cycle c_i is of length at least two. Suppose that the first cycle, c_1 , is

$$c_1 = (i_1 \mapsto i_2 \mapsto \dots \mapsto i_{k_1}),$$

where $1 \leq i_1 < i_2 < \dots < i_{k_1} \leq k$. Let us write the full forward k -cycle in the form

$$(1 \mapsto \dots \mapsto k) = (i_1 \mapsto I_1 \mapsto i_2 \mapsto I_2 \mapsto \dots \mapsto i_{k_1} \mapsto I_{k_1}),$$

where the I_* 's are intervals. Since

$$(i_1 \mapsto I_1 \mapsto i_2 \mapsto I_2 \mapsto \dots \mapsto i_{k_1} \mapsto I_{k_1}) = (i_1 \mapsto \dots \mapsto i_{k_1})(i_1 \mapsto I_1)(i_2 \mapsto I_2) \dots (i_{k_1} \mapsto I_{k_1}),$$

we have

$$\begin{aligned} (i_1 \mapsto \dots \mapsto i_{k_1})^{-2}(1 \mapsto \dots \mapsto k) &= (i_1 \mapsto \dots \mapsto i_{k_1})^{-1}(i_1 \mapsto I_1)(i_2 \mapsto I_2) \dots (i_{k_1} \mapsto I_{k_1}) \\ &= (i_{k_1} \mapsto \dots \mapsto i_1)(i_{k_1} \mapsto I_{k_1}) \dots (i_2 \mapsto I_2)(i_1 \mapsto I_1) \\ &= (i_{k_1} \mapsto I_{k_1} \mapsto \dots \mapsto i_2 \mapsto I_2 \mapsto i_1 \mapsto I_1). \end{aligned}$$

Now, since the cycles c_1, c_2, \dots induce a noncrossing partition of $\{1, \dots, k\}$, the cycle c_2 is contained in one of the intervals I_1, \dots, I_{k_1+1} . Thus the same argument applies to compute

$$c_2^{-2}(i_{k_1} \mapsto I_{k_1} \mapsto \dots \mapsto i_2 \mapsto I_2 \mapsto i_1 \mapsto I_1). \quad \square$$

3.3 Encoding

The duality operation has been introduced, and its basic properties developed, in order to make available a structured means of encoding *pairs* of midpoints using crossovers. To this end, we introduce the mapping

$$\text{Cr}(\mu) \longrightarrow M(a, b) \times M(a, b)$$

defined by

$$c \mapsto \Phi_c(a, b) := (\varphi_c(a, b), \varphi_{c^\vee}(a, b)).$$

Thus $(x, y) = \Phi_c(a, b)$ is the pair of midpoints of a and b encoded by c and c^\vee . The duality relationship between c and c^\vee induces algebraic and geometric relations between the pairs (a, b) and (x, y) which may be collectively called *duality of midpoints*.

Proposition 5. *If $(x, y) = \Phi_c(a, b)$, then:*

1. $x^{-1}y$ has the same ordered cycle type and sequence of cycle minima as $a^{-1}b$;
2. a and b are midpoints of x and y ;
3. $u_{x^{-1}y} = u_{a^{-1}b}u_{c^{-1}c^\vee}$.

Proof. Let us prove these assertions in order.

1. First, we have

$$\begin{aligned} x^{-1}y &= (\varphi_c(a, b))^{-1}(\varphi_{c^\vee}(a, b)) \\ &= (au_{a^{-1}b}cu_{a^{-1}b}^{-1})^{-1}(au_{a^{-1}b}c^\vee u_{a^{-1}b}^{-1}) \\ &= u_{a^{-1}b}c^{-1}c^\vee u_{a^{-1}b}^{-1}. \end{aligned}$$

Let μ be the ordered cycle type of $a^{-1}b$. By definition, $u_{a^{-1}b}$ conjugates p_μ into $a^{-1}b$ and transforms the sequence of cycle minima of p_μ into that of $a^{-1}b$. By Lemma 4, $c^{-1}c^\vee$ has both the same ordered cycle type and sequence of cycle minima as p_μ . Thus, $u_{a^{-1}b}c^{-1}c^\vee u_{a^{-1}b}^{-1}$ has both the same ordered cycle type and sequence of cycle minima as $a^{-1}b$.

2. We now show that a and b are midpoints of x and y . Since

$$a^{-1}x = a^{-1}(au_{a^{-1}b}cu_{a^{-1}b}^{-1}) = u_{a^{-1}b}cu_{a^{-1}b}^{-1}$$

and

$$\begin{aligned} y^{-1}b &= (au_{a^{-1}b}c^\vee u_{a^{-1}b}^{-1})^{-1}b \\ &= (au_{a^{-1}b}c^{-1}p_\mu u_{a^{-1}b}^{-1})^{-1}b \\ &= (au_{a^{-1}b}c^{-1}u_{a^{-1}b}^{-1}u_{a^{-1}b}p_\mu u_{a^{-1}b}^{-1})^{-1}b \\ &= (au_{a^{-1}b}c^{-1}u_{a^{-1}b}^{-1}a^{-1}b)^{-1}b \\ &= (b^{-1}au_{a^{-1}b})c(b^{-1}au_{a^{-1}b})^{-1}, \end{aligned}$$

both $a^{-1}x$ and $y^{-1}b$ are conjugates of c . By the conjugation invariance of d we thus have

$$d(e, a^{-1}x) = d(e, c) = d(e, y^{-1}b),$$

whence

$$d(a, x) = d(e, c) = d(y, b).$$

Similarly, since

$$a^{-1}y = a^{-1}(au_{a^{-1}b}c^\vee u_{a^{-1}b}^{-1}) = u_{a^{-1}b}c^\vee u_{a^{-1}b}^{-1}$$

and

$$\begin{aligned} x^{-1}b &= (au_{a^{-1}b}cu_{a^{-1}b}^{-1})^{-1}b \\ &= u_{a^{-1}b}c^{-1}u_{a^{-1}b}^{-1}a^{-1}b \\ &= u_{a^{-1}b}c^\vee p_\mu^{-1}u_{a^{-1}b}^{-1}a^{-1}b \\ &= u_{a^{-1}b}c^\vee (u_{a^{-1}b}^{-1}(a^{-1}b)^{-1}u_{a^{-1}b})u_{a^{-1}b}^{-1}a^{-1}b \\ &= u_{a^{-1}b}c^\vee u_{a^{-1}b}^{-1}, \end{aligned}$$

both $a^{-1}y$ and $x^{-1}b$ are conjugates of c^\vee , and we have

$$d(a, y) = d(e, c^\vee) = d(x, b).$$

Now, since $x^{-1}y$ has the same ordered cycle type as $a^{-1}b$, we have

$$d(x, y) = d(a, b) = d(e, p_\mu).$$

Putting this all together, we have

$$\begin{aligned} d(x, a) + d(a, y) &= d(e, c) + d(e, c^\vee) \\ &= d(e, c) + d(c, p_\mu) \\ &= d(e, p_\mu) \\ &= d(x, y), \end{aligned}$$

and

$$\begin{aligned} d(x, a) - d(a, y) &= d(e, c) - d(e, c^\vee) \\ &= d(e, c) - d(c, p_\mu) \\ &= \varepsilon, \end{aligned}$$

where $\varepsilon \in \{-1, 0, 1\}$ because c is a midpoint of e and p_μ . Thus, a is a midpoint of x and y . The proof that b is a midpoint of x and y is just the same.

3. Finally, we prove the identity

$$u_{x^{-1}y} = u_{a^{-1}b}u_{c^{-1}c^\vee},$$

which will be needed below in the proof of Proposition 6. Since

$$\begin{aligned} x^{-1}y &= (au_{a^{-1}b}cu_{a^{-1}b}^{-1})^{-1}(au_{a^{-1}b}c^\vee u_{a^{-1}b}^{-1}) \\ &= u_{a^{-1}b}c^{-1}c^\vee u_{a^{-1}b}^{-1} \\ &= u_{a^{-1}b}u_{c^{-1}c^\vee}p_\mu u_{c^{-1}c^\vee}^{-1}u_{a^{-1}b}^{-1} \\ &= (u_{a^{-1}b}u_{c^{-1}c^\vee})p_\mu(u_{a^{-1}b}u_{c^{-1}c^\vee})^{-1}, \end{aligned}$$

we have that $u_{a^{-1}b}u_{c^{-1}c^\vee}$ conjugates p_μ into $x^{-1}y$. This is one of the two properties that uniquely defines the permutation $u_{x^{-1}y}$, the other being that it transforms the sequence of cycle minima of p_μ into the sequence of cycle minima of $x^{-1}y$. Since $u_{c^{-1}c^\vee}$ stabilizes the sequence of cycle minima of p_μ (by Lemma 4), conjugation of p_μ by $u_{a^{-1}b}u_{c^{-1}c^\vee}$ produces a permutation whose sequence of cycle minima coincides with that of $a^{-1}b$, and hence with that of $x^{-1}y$ by Part (1). \square

3.4 Decoding

Our constructions so far may be thought of in cryptographic terms, as follows. Alice and Bob wish to transmit messages to one another across an insecure channel. They meet at a secure location, and agree on a composition $\mu \vDash n$ and a crossover $c \in \text{Cr}(\mu)$ to be used as a secret encryption key. Alice and Bob then return to their respective locations on opposite ends of the channel. The plaintext messages to be transmitted are pairs

$(a, b) \in S(n) \times S(n)$ such that $a^{-1}b \in \vec{C}_\mu$. To send the message (a, b) to Bob, Alice computes the ciphertext $(x, y) = \Phi_c(a, b)$ and transmits it to Bob across the channel. Bob receives the ciphertext (x, y) , and wishes to recover the plaintext message. Our next result proves that there is a well-defined decryption key $\delta(c)$ such that $(a, b) = \Phi_{\delta(c)}(x, y)$ — in fact, the proof explains how to compute $\delta(c)$.

Proposition 6. *For each composition $\mu \vDash n$, there exists a function*

$$\delta_\mu : \text{Cr}(\mu) \longrightarrow \text{Cr}(\mu)$$

such that

$$\Phi_{\delta_\mu(c)}(\Phi_c(a, b)) = (a, b)$$

holds for all $c \in \text{Cr}(\mu)$ and each $(a, b) \in S(n) \times S(n)$ verifying $a^{-1}b \in \vec{C}_\mu$. Moreover, this function is an involution.

We call the function δ_μ of Proposition 6 the *decoding function* of type μ . In order to lighten the notation, we will henceforth omit the dependence of δ on μ .

Proof. Fix a composition $\mu \vDash n$. We claim that the corresponding decoding function is given by

$$\delta(c) := u_{c^{-1}c^\vee}^{-1}c^{-1}u_{c^{-1}c^\vee}.$$

First, let us check that the codomain of δ is indeed $\text{Cr}(\mu)$, i.e. that $\delta(c)$ is in fact a midpoint of e and p_μ . We have

$$\begin{aligned} d(e, \delta(c)) &= d(e, u_{c^{-1}c^\vee}^{-1}c^{-1}u_{c^{-1}c^\vee}) \\ &= d(e, c^{-1}) \\ &= d(e, c) \end{aligned}$$

and

$$\begin{aligned} d(\delta(c), p_\mu) &= d(u_{c^{-1}c^\vee}^{-1}c^{-1}u_{c^{-1}c^\vee}, p_\mu) \\ &= d(c^{-1}, u_{c^{-1}c^\vee}p_\mu u_{c^{-1}c^\vee}^{-1}) \\ &= d(c^{-1}, c^{-1}c^\vee) \\ &= d(e, c^\vee) \\ &= d(e, c^{-1}p_\mu) \\ &= d(c, p_\mu). \end{aligned}$$

Thus, since c is a midpoint of e and p_μ , we have

$$d(e, \delta(c)) + d(\delta(c), p_\mu) = d(e, c) + d(c, p_\mu) = d(e, p_\mu)$$

and

$$d(e, \delta(c)) - d(\delta(c), p_\mu) = d(e, c) - d(c, p_\mu) = \varepsilon$$

with $\varepsilon \in \{-1, 0, 1\}$.

Next, let $(a, b) \in S(n) \times S(n)$ be a valid plaintext, i.e. a pair of permutations such that $a^{-1}b \in \tilde{C}_\mu$, and let $(x, y) = \Phi_c(a, b)$ be the corresponding ciphertext. We then have

$$\begin{aligned} x &= \varphi_c(a, b) = au_{a^{-1}b}cu_{a^{-1}b}^{-1} \\ y &= \varphi_c(a, b) = au_{a^{-1}b}c^\vee u_{a^{-1}b}^{-1}, \end{aligned}$$

By Proposition 5, Part (1), we may re-encode (x, y) using $\delta(c)$ as an encryption key, arriving at a new ciphertext $(x', y') = \Phi_{\delta(c)}(x, y)$. We claim that this re-encoding decrypts (x, y) , i.e. that $(x', y') = (a, b)$. Indeed,

$$\begin{aligned} x' &= \varphi_{\delta(c)}(x, y) \\ &= xu_{x^{-1}y}\delta(c)u_{x^{-1}y}^{-1} \\ &= (au_{a^{-1}b}cu_{a^{-1}b}^{-1})(u_{a^{-1}b}u_{c^{-1}c^\vee})(u_{c^{-1}c^\vee}^{-1}c^{-1}u_{c^{-1}c^\vee})(u_{a^{-1}b}u_{c^{-1}c^\vee})^{-1} \\ &= a, \end{aligned}$$

where we made use of Proposition 5, Part (3). Similarly, we have

$$\begin{aligned} y' &= \varphi_{\delta(c)^\vee}(x, y) \\ &= xu_{x^{-1}y}\delta(c)^\vee u_{x^{-1}y}^{-1} \\ &= (au_{a^{-1}b}cu_{a^{-1}b}^{-1})(u_{a^{-1}b}u_{c^{-1}c^\vee})(u_{c^{-1}c^\vee}^{-1}c^{-1}u_{c^{-1}c^\vee})^\vee (u_{a^{-1}b}u_{c^{-1}c^\vee})^{-1} \\ &= (au_{a^{-1}b}cu_{a^{-1}b}^{-1})(u_{a^{-1}b}u_{c^{-1}c^\vee})(u_{c^{-1}c^\vee}^{-1}c^\vee u_{c^{-1}c^\vee}p_\mu)(u_{a^{-1}b}u_{c^{-1}c^\vee})^{-1} \\ &= au_{a^{-1}b}c^2u_{c^{-1}c^\vee}p_\mu u_{c^{-1}c^\vee}^{-1}u_{a^{-1}b}^{-1} \\ &= au_{a^{-1}b}c^\vee u_{a^{-1}b} \\ &= au_{a^{-1}b}p_\mu u_{a^{-1}b} \\ &= b. \end{aligned}$$

It remains to show that

$$\delta : \text{Cr}(\mu) \longrightarrow \text{Cr}(\mu)$$

is an involution. Let c be a μ -crossover, let (a, b) be a valid message, and consider the triple encoding

$$\Phi_{\delta^2(c)}(\Phi_{\delta(c)}(\Phi_c(a, b))).$$

Since $\delta(c)$ is the decryption key corresponding to the encryption key c , we have

$$\Phi_{\delta^2(c)}(\Phi_{\delta(c)}(\Phi_c(a, b))) = \Phi_{\delta^2(c)}(a, b).$$

On the other hand, since $\delta^2(c)$ is the decryption key corresponding to the encryption key $\delta(c)$, we also have

$$\Phi_{\delta^2(c)}(\Phi_{\delta(c)}(\Phi_c(a, b))) = \Phi_c(a, b).$$

Thus $\Phi_{\delta^2(c)}(a, b) = \Phi_c(a, b)$, which readily implies $\delta^2(c) = c$. □

4 The Brunn-Minkowski inequality

4.1 Without a curvature term

We now put our encoding-decoding formalism to work. We begin by proving the “flat” Brunn-Minkowski inequality for $S(n)$.

Theorem 7. *For any nonempty $A, B \subseteq S(n)$, we have*

$$\log |M| \geq \frac{1}{2} \log |A| + \frac{1}{2} \log |B|,$$

where M is the midpoint set of A and B .

Proof. Suppose that there exists an injection

$$\Phi : A \times B \longrightarrow M \times M.$$

Then

$$|M|^2 \geq |A||B|,$$

and taking logs this becomes

$$\log |M| \geq \frac{1}{2} \log |A| + \frac{1}{2} \log |B|.$$

To construct such an injection, let μ_1, \dots, μ_m be an enumeration of the ordered cycle types of the permutations

$$a^{-1}b, \quad (a, b) \in A \times B.$$

For each $1 \leq i \leq m$, choose an encryption key $c_i \in \text{Cr}(\mu_i)$, and let $d_i = \delta(c_i)$ be the corresponding decryption key.

Partition $A \times B$ into classes

$$C_i = \{(a, b) \in A \times B : a^{-1}b \in \vec{C}_{\mu_i}\}, \quad 1 \leq i \leq m,$$

and consider the map

$$\Phi : A \times B \longrightarrow M \times M$$

whose restriction to each C_i is given by encryption using key c_i :

$$\Phi(a, b) = \Phi_{c_i}(a, b).$$

We claim that Φ is injective. Indeed, if

$$\Phi(a, b) = \Phi(a', b')$$

for some $(a, b), (a', b') \in A \times B$, then we must have $(a, b), (a', b') \in C_i$ for some $1 \leq i \leq m$ by Part 1 of Proposition 5. Hence

$$(a, b) = \Phi_{d_i}(\Phi_{c_i}(a, b)) = \Phi_{d_i}(\Phi_{c_i}(a', b')) = (a', b'),$$

by Proposition 6. □

4.2 With a curvature term

Theorem 7 above was proved by injecting $A \times B$ into $M \times M$. The injection used was a straightforward application of the formalism of encoding and decoding developed above. We now obtain the curved Brunn-Minkowski inequality stated as Theorem 2 in the introduction by injecting two copies of $A \times B$ into $M \times M$. The construction of the required injection involves a subtler use of encoding/decoding: the proof requires choosing two parallel sequences of encryption keys which are coupled in a special way.

Theorem 8. *For any nonempty $A, B \subseteq S(n)$, we have*

$$\log |M| \geq \frac{1}{2} \log |A| + \frac{1}{2} \log |B| + \frac{K}{8} d(A, B)^2$$

where M is the midpoint set of A and B and

$$K = \frac{4 \log 2}{(n-1)^2}.$$

Proof. First note that if $d(A, B) = 0$, the claimed inequality degenerates to the flat Brunn-Minkowski inequality, which we have already proved. We thus assume that A and B are disjoint.

Suppose that there exists an injection

$$\Phi : A \times B \times \{0, 1\} \longrightarrow M \times M.$$

Then

$$|M|^2 \geq 2|A||B|,$$

and taking logs this becomes

$$\log |M| \geq \frac{1}{2} \log |A| + \frac{1}{2} \log |B| + \frac{\log 2}{2}.$$

The desired inequality now follows from the fact that the diameter of $S(n)$ is $n - 1$.

In order to construct Φ as required, let μ_1, \dots, μ_m and C_1, \dots, C_m be as in the proof of Theorem 7. Choose a system of encryption keys

$$c_i \in \text{Cr}(\mu_i), \quad 1 \leq i \leq m,$$

and consider a second system of encryption keys obtained from the first system by setting

$$\tilde{c}_i := \delta(\delta(c_i)^\vee), \quad 1 \leq i \leq m.$$

The keys \tilde{c}_i are defined in this way so that, by the involutive property of δ , their corresponding decryption keys are the duals of the decryption keys of the first system:

$$\delta(\tilde{c}_i) = \delta(c_i)^\vee. \tag{2}$$

We build Φ from the above data as follows. First, partition $A \times B \times \{0, 1\}$ into the $2m$ sets

$$C_i^{(j)} = \{(a, b, j) : a^{-1}b \in C_i\}, \quad 1 \leq i \leq m, \quad 0 \leq j \leq 1.$$

We define Φ by declaring its restriction to $C_i^{(j)}$ to be

$$\Phi(a, b, j) := \begin{cases} \Phi_{c_i}(a, b), & \text{if } j = 0 \\ \Phi_{\bar{c}_i}(a, b), & \text{if } j = 1 \end{cases}.$$

We now prove that Φ so defined is an injection. Suppose

$$(x, y) \in M \times M, \quad (a, b, j) \in C_i^{(j)}, \quad (a', b', j') \in C_{i'}^{(j')}$$

are such that

$$(x, y) = \Phi(a, b, j) = \Phi(a', b', j').$$

Then, by the first part of Proposition 5, we must have $i = i'$. We claim that also $j = j'$. If not, then (relabelling if necessary) we have

$$(x, y) = \Phi(a, b, 0) = \Phi(a', b', 1),$$

so that

$$(x, y) = \Phi_{c_i}(a, b) = \Phi_{\bar{c}_i}(a', b').$$

Decoding, we obtain

$$(a, b) = \Phi_{\delta(c_i)}(x, y) = (\varphi_{\delta(c_i)}(x, y), \varphi_{\delta(c_i)^\vee}(x, y))$$

and, using (2),

$$(a', b') = \Phi_{\delta(\bar{c}_i)}(x, y) = (\varphi_{\delta(\bar{c}_i)^\vee}(x, y), \varphi_{\delta(\bar{c}_i)}(x, y)).$$

This implies $a' = b$, which is impossible since A and B are disjoint.

There are now two cases: $j = j' = 0$ and $j = j' = 1$. In the first case, we have

$$\Phi_{c_i}(a, b) = \Phi(a, b, 0) = \Phi(a', b', 0) = \Phi_{c_i}(a', b'),$$

whence

$$(a, b) = \Phi_{\delta(c_i)}(\Phi_{c_i}(a, b)) = \Phi_{\delta(c_i)}(\Phi_{c_i}(a', b')) = (a', b').$$

In the second case, we have

$$\Phi_{\bar{c}_i}(a, b) = \Phi(a, b, 1) = \Phi(a', b', 1) = \Phi_{\bar{c}_i}(a', b'),$$

whence

$$(a, b) = \Phi_{\delta(\bar{c}_i)}(\Phi_{\bar{c}_i}(a, b)) = \Phi_{\delta(\bar{c}_i)}(\Phi_{\bar{c}_i}(a', b')) = (a', b'). \quad \square$$

4.3 With an optimal curvature term

We conjecture that the curved Brunn-Minkowski inequality proved in Theorem 8 can be improved to the following optimal statement.

Conjecture 9. For any nonempty $A, B \subseteq S(n)$, we have

$$\log |M| \geq \frac{1}{2} \log |A| + \frac{1}{2} \log |B| + \frac{K}{8} d(A, B)^2$$

where M is the midpoint set of A and B ,

$$K = \frac{c}{n-1},$$

and c is a positive constant.

Note that this statement is indeed best possible, by Theorem 1. While we are at present unable to prove Conjecture 9, we show here that it is implied by the following conjectural concentration inequality.

Conjecture 10. There exists a positive constant $\varepsilon > 0$ such that, for any $\mu \models n$, we have

$$d(C, C^\vee) \geq r \implies \mathbb{P}(C) \leq e^{-\frac{\varepsilon r^2}{n-\ell(\mu)}},$$

where \mathbb{P} is the uniform probability measure on $\text{Cr}(\mu)$.

Remark 11. In the case where $\mu = (\mu_1, \mu_2, \dots)$ satisfies $\mu_i \in \{1, 2\}$, Conjecture 10 is true — via the embedding of the hypercube described in the introduction, it is equivalent to Corollary 6 in [4].

We now explain how Conjecture 9 can be deduced from Conjecture 10. This argument, which lifts the proof of [4, Theorem 1] from the hypercube to the symmetric group, differs substantially from the proofs of Theorems 7 and 8. To prove these coarser results, we used static encoding to construct our injections, i.e. a predetermined list of encryption keys. Here, we use an adaptive coding scheme in which all crossovers associated to $A \times B$ are employed.

Theorem 12. *Suppose that Conjecture 10 is true. Then, for any nonempty $A, B \subseteq S(n)$, we have*

$$\log |M| \geq \frac{1}{2} \log |A| + \frac{1}{2} \log |B| + \frac{K}{8} d(A, B)^2$$

where M is the midpoint set of A and B and

$$K = \frac{4\varepsilon}{n-1}.$$

Proof. Let μ_1, \dots, μ_m and C_1, \dots, C_m be as in the proof of Theorem 7. Put

$$U_i := C_i \times \text{Cr}(\mu_i), \quad 1 \leq i \leq m.$$

These sets are pairwise disjoint, but note that they are not subsets of $A \times B$. Rather,

$$U := \bigsqcup_{i=1}^m U_i$$

is an “enriched” version of $A \times B$ in which each pair $(a, b) \in C_i$ appears with multiplicity $|\text{Cr}(\mu_i)|$.

Consider the map

$$\Phi : U \longrightarrow M \times M$$

defined by

$$\Phi(a, b, c) := \Phi_c(a, b).$$

By Proposition 5, we know that the image of U_i under Φ is contained in

$$V_i := \{(x, y) \in M \times M : x^{-1}y \in \vec{C}_{\mu_i}\}.$$

The sets V_1, \dots, V_m are disjoint, and

$$\sum_{i=1}^m |V_i| \leq |M \times M|.$$

Fix an arbitrary $i \in \{1, \dots, m\}$, and an arbitrary pair of midpoints $(x, y) \in V_i$. By definition, $(a, b, c) \in U_i$ maps to (x, y) under Φ if and only if

$$\Phi_c(a, b) = (x, y),$$

which by Proposition 6 is equivalent to

$$\Phi_{\delta(c)}(x, y) = (a, b).$$

Thus the cardinality of the fibre of Φ over (x, y) agrees with that of the set

$$D(x, y) := \{d \in \text{Cr}(\mu_i) : \Phi_d(x, y) \in C_i\}.$$

Now let $d_1, d_2 \in D(x, y)$. Then, by definition,

$$\varphi_{d_1}(x, y) \in A, \quad \varphi_{d_2^\vee}(x, y) \in B.$$

Because the map

$$d \mapsto \varphi_d(x, y)$$

is an isometry, this implies

$$d(d_1, d_2^\vee) = d(\varphi_{d_1}(x, y), \varphi_{d_2^\vee}(x, y)) \geq d(A, B),$$

and hence

$$d(D(x, y), D(x, y)^\vee) \geq d(A, B).$$

Hence, assuming Conjecture 10, we have that

$$|D(x, y)| \leq e^{-\frac{\varepsilon d(A, B)^2}{n - \ell(\mu_i)}} |\text{Cr}(\mu_i)| \leq e^{-\frac{\varepsilon d(A, B)^2}{n-1}} |\text{Cr}(\mu_i)|.$$

Summing over all $(x, y) \in V_i$ we thus obtain

$$|U_i| \leq e^{-\frac{\varepsilon d(A, B)^2}{n-1}} |\text{Cr}(\mu_i)| |V_i|.$$

Since $|U_i| = |C_i| |\text{Cr}(\mu_i)|$, this implies

$$|C_i| \leq e^{-\frac{\varepsilon d(A, B)^2}{n-1}} |V_i|.$$

Summing over $1 \leq i \leq m$, we get

$$\sum_{i=1}^m |C_i| \leq e^{-\frac{\varepsilon d(A, B)^2}{n-1}} \sum_{i=1}^m |V_i|,$$

whence

$$|A \times B| \leq e^{-\frac{\varepsilon d(A, B)^2}{n-1}} |M \times M|.$$

Taking logs, we obtain the curved Brunn-Minkowski inequality with

$$K = \frac{4\varepsilon}{n-1}. \quad \square$$

Acknowledgements

We thank Prasad Tetali for introducing us to the subject of discrete curvature, and suggesting the problem of lifting the curved Brunn-Minkowski inequality from the hypercube to the symmetric group. We are grateful to Professor Tetali for communicating to us a set of notes [5] on this topic compiled by a group of researchers following an AIM workshop on discrete curvature, and for pointing out the papers [1, 2]. We thank Mike Lacroix for creating Figure 1.

References

- [1] M. Erbaar, J. Maas, P. Tetali, *Discrete Ricci curvature bounds for the Bernoulli-Laplace and random transposition models*, [arXiv:1409.8605v1](https://arxiv.org/abs/1409.8605v1)
- [2] B. Klartag, G. Kozma, P. Ralli, P. Tetali, *Discrete curvature and abelian groups*, [arXiv:1501.00516v2](https://arxiv.org/abs/1501.00516v2)
- [3] A. Nica, R. Speicher, *Lectures on The Combinatorics of Free Probability*, London Mathematical Society Lecture Note Series, No. 335, 2006.

- [4] Y. Ollivier, C. Villani, *A curved Brunn-Minkowski inequality on the hypercube, or: What is the Ricci curvature of the discrete hypercube?*, SIAM J. Discrete Math. **26** (2012), 983-996.
- [5] K. Costello, N. Gozlan, J. Melbourne, W. Perkins, C. Roberto, P.-M. Samson, P. Tetali, *Notes on displacement convexity on the symmetric group*, unpublished.