The entropy efficiency of point-push mapping classes on the punctured disk

PHILIP BOYLAND

JASON HARRINGTON

We study the maximal entropy per unit generator of point-push mapping classes on the punctured disk. Our work is motivated by fluid mixing by rods in a planar domain. If a single rod moves among N fixed obstacles, the resulting fluid diffeomorphism is in the point-push mapping class associated with the loop in $\pi_1(D^2 - \{N \text{ points}\})$ traversed by the single stirrer. The collection of motions where each stirrer goes around a single obstacle generate the group of point-push mapping classes, and the entropy efficiency with respect to these generators gives a topological measure of the mixing per unit energy expenditure of the mapping class. We give lower and upper bounds for $\mathrm{Eff}(N)$, the maximal efficiency in the presence of N obstacles, and prove that $\mathrm{Eff}(N) \to \log(3)$ as $N \to \infty$. For the lower bound we compute the entropy efficiency of a specific point-push protocol, HSP_N, which we conjecture achieves the maximum. The entropy computation uses the action on chains in a $\mathbb{Z}$–covering space of the punctured disk which is designed for point-push protocols. For the upper bound we estimate the exponential growth rate of the action of the point-push mapping classes on the fundamental group of the punctured disk using a collection of incidence matrices and then computing the generalized spectral radius of the collection.

1 Introduction

The topological entropy $h(f)$ of a map f is a standard measure of its dynamical complexity. The entropy of an isotopy class is the infimum of all the entropies of the homeomorphisms in the class. If a surface isotopy class has positive entropy, a theorem of Thurston [29] guarantees the existence of an (almost) canonical map in the isotopy class which realizes the infimum and whose dynamics are described by a finite collection of mixing subshifts of finite type (see also Fathi, Laudenbach and Poenaru [12]). Further, Handel's theorem [19] shows that these dynamics are present up to semiconjugacy in every homeomorphism in the mapping class (see Boyland [5]).

The prototypical positive entropy mapping class is of pseudo-Anosov type and the Thurston representative in the class is a pseudo-Anosov map ϕ. In this case, $\lambda = \exp(h(\phi))$ is the dilation of the pseudo-Anosov map, an algebraic integer with a wide variety of geometric and topological interpretations. As a consequence of these interpretations it is natural to seek the pseudo-Anosov map with least dilation within semigroups of isotopy classes (see, for example, Hironaka and Kin [20], Farb, Leininger and Margalit [10] or Lanneau and Thiffeault [24]). From a dynamics perspective, maximizing complexity is also of interest. Since entropy is multiplicative under iteration, $h(f^n) = nh(f)$, the entropy over a semigroup is always unbounded, and thus one needs a normalization. In many cases the semigroup of interest has a natural set of generators, for example, a collection of Dehn twists or motions of stirring rods as discussed below. In these cases following Finn and Thiffeault [14; 28] and Moussafir [26], we consider the *entropy efficiency* of a mapping class which is the topological entropy per unit generator. The entropy efficiency of the semigroup is then the supremum of the efficiencies of its members.

As a consequence of Thurston's theorem, the entropy of an isotopy class is equal to the exponential growth rate of words under iteration by the induced map on the fundamental group of the surface. One may also define this growth rate and thus the entropy of any automorphism of a finitely generated group G. Given a finite collection of automorphisms $\mathcal{A} = \{a_1, \ldots, a_n\}$ of the group, we define the *entropy efficiency* of the semigroup generated by $\mathcal{A}$ as the maximal entropy per unit generator (see Section 2.2). If G is a free abelian group, the automorphisms are linear isomorphisms and the entropy efficiency of $\mathcal{A}$ is the log of the generalized spectral radius of the collection of linear transformations (Remark 2.1). The generalized and joint spectral radii are very useful and much-studied linear objects, and the entropy efficiency is their nonlinear analog.

Our main motivation for studying the entropy efficiency is fluid mixing via the motion of stirring rods in a planar domain (see Boyland, Aref and Stremler [7], Finn and Thiffeault [14; 28] and Boyland [6]). In the simplest situation the mixing is accomplished by n rods which after time T return to their original positions. Such a stirring protocol is naturally described as a braid with n–strings. When the rods move through the fluid there is a resulting motion of the entire fluid region, and the evolution of a fluid particle from its initial position to its position after time T defines the *fluid motion diffeomorphism*. The exact nature of this diffeomorphism depends on the physical properties of the fluid and the geometry of the fluid region and rod paths. However, as long as the fluid does not tear, all possible fluid motion diffeomorphisms arising from the same stirring protocol are in the same mapping class, and this class is the one usually associated with the braid (see Section 2.3). Thus if we can determine

the entropy of this isotopy class we have a lower bound for the complexity of the evolution of any fluid motion resulting from the stirring protocol. The topological entropy has a particular physical significance because it measures, among other things, the exponential growth rate of material lines in the fluid under the action of the protocol. Exponential stretching of material lines does not guarantee good mixing, but in many situations this rate of stretching is a good topological "first order" indicator.

The other consideration in implementing a stirring protocol is its energy cost. Protocols which maximize the mixing per unit expenditure of energy could have great practical value. The energy required to implement a stirring protocol certainly depends on the geometry of the path of the stirrer, but since we are characterizing protocols just by their braids, we seek a topological measure of effort which just depends on the braid. Thus we pick some class of fundamental motions of the stirrers, for example, the switching of a pair of rods, and treat these as generators of a semigroup of mapping classes. The corresponding maximal entropy efficiency gives a rough notion of the maximal mixing per unit expenditure of energy possible in the class of protocols. A protocol which realizes the maximum would be optimal in a precise mathematical sense, but perhaps not in any true physical sense; see [14].

The entropy efficiency of the full braid group with respect to the standard generators was studied by Finn and Thiffeault in [14]. They show that the maximal efficiency is bounded above by $\log\left((1 + \sqrt{5})/2\right)$. Further, when $n > 4$, this bound is strict, and when $n = 3$ and 4 the bound is realized by $\sigma_1\sigma_2^{-1}$ and $\sigma_1^{-1}\sigma_2\sigma_3^{-1}\sigma_2$, respectively. They also studied a generating set consisting of words of the form $\sigma_{i_1}^{\pm 1}\cdots\sigma_{i_k}^{\pm 1}$, where for all j, $|i_j - i_{j+1}| > 1$. For this class they obtained $\log(1 + \sqrt{2})$ as an upper bound for the maximal efficiency. This class of generators models motions of the collection of rods in which the entire motion can be done simultaneously.

Here we study the class of stirring protocols in which there are N fixed obstacles and a single stirrer moves. Thus the stirring protocols are determined by the path of the single stirrer in the N–punctured disk. This path can be considered an element of the fundamental group and so these stirring motions are called π_1–*stirring protocols*. Each π_1–stirring protocol determines a braid on $(N+1)$ strings and an isotopy class on the $(N+1)$–punctured disk. The mapping classes which result from π_1–stirring protocols are called *point-push* in the mathematics literature (see Birman [4], Kra [23], Dowdall [9] and Farb and Margalit [11]). Thus the collection of π_1–stirring protocols, denoted here as $\mathcal{P}_N$, is the subgroup of the braid group on $(N+1)$ strings corresponding to the point-push isotopy classes in the N–punctured disk. To maintain the traditional mathematical terminology while incorporating the physical intuition and motivation of our problem, we will usually use *point-push protocol* to refer to elements of $\mathcal{P}_N$.

We use the analog of the standard generators of the fundamental group of the N–punctured disk as our fundamental stirring motions. Thus for $j = 1, \ldots, N$, let α_j correspond to moving the stirrer counterclockwise around the j–th obstacle (Figure 1), and we consider the entropy efficiency of a point-push protocol with respect to the generating set $\{\alpha_1, \ldots, \alpha_N\}$ of $\mathcal{P}_N$. The maximal entropy efficiency over $\mathcal{P}_N$ is denoted $\mathrm{Eff}(N)$. Our main theorem is:

Theorem 1.1 *For $N \geq 2$ there are explicitly defined matrices $H^{(N)}$ and $\hat{H}^{(N)}$ such that the maximal efficiency $\mathrm{Eff}(N)$ of a point-push protocol in the presence of N obstacles satisfies*

$$\frac{\log(3^N - 3N - 1) - \log(N)}{N} \leq \frac{\log(\rho(H^{(N)}))}{N} \leq \mathrm{Eff}(N) \leq \frac{\log(\rho(\hat{H}^{(N)}))}{N} \leq \frac{\log(3^N - 2)}{N},$$

and so $\mathrm{Eff}(N) \to \log(3)$ as $N \to \infty$.

Note that the lower bound has no content when $N = 2$. The most efficient protocol with two obstacles is given in Theorem 4.7.

The proof of Theorem 1.1 follows directly from the lower bound given in Theorem 4.5 and the upper bound given in Theorem 5.9. Both the upper and lower bound proceed by reducing the problem to the computation of the spectral radius of a specific product of generating matrices, but the reduction in the two cases is quite different. In both cases the required computations are tractable because of the simple forms of actions made possible by the adapted set of generators shown in Figure 2.

Since the maximal efficiency is the supremum over the efficiencies of all point-push protocols, the efficiency of any single protocol serves as a lower bound. In Section 4.1 we define a point-push protocol HSP_N which numerically appears to give the maximal efficiency (Figure 3). The invariant train track for HSP_N is rather simple, but the corresponding Markov transition matrix has very large entries. Thus the computation or even estimation of the dilation of HSP_N as its spectral radius appears to be quite difficult. Instead we use homological methods developed by Fried [16; 15] (see also Band and Boyland [1]). In short, the idea is to lift the mapping class to finite covers where the dilation is computable as the spectral radius of the action on homology. This cover is found as the quotient of a $\mathbb{Z}$–covering space whose homology is treated as a free $\mathbb{Z}[t^{\pm 1}]$–module. The action of the mapping class is then a matrix over $\mathbb{Z}[t^{\pm 1}]$, and substitutions of $t = \exp(2\pi i p/q)$ give pieces of the action of the mapping class on the homology of the corresponding $\mathbb{Z}_q$–covering space.

As noted above, the entropy of a point-push protocol is equal to the exponential growth rate of words under iteration by the induced map on the fundamental group of the

punctured disk, $\pi_1(D_{N+1})$. For each generator α_j of $\mathcal{P}_N$, we find the incidence matrix $A^{(j)}$ of its action on $\pi_1(D_{N+1})$. The incidence matrix just counts the number of occurrence of a generator in the image of another generator. As also remarked above, the entropy efficiency is the nonlinear analog of the generalized spectral radius and in Proposition 5.1 we show that for a finitely generated semigroup of free group automorphisms, the log of the generalized spectral radius of the incidence matrices of the generators is an upper bound for the entropy efficiency of the semigroup. In general, finding a generalized spectral radius is difficult, but because of the particularly simple form of the action of the α_j on the generators in Figure 2, in our case we can determine the exact product which yields the generalized spectral radius and estimate its spectral radius (Proposition 5.3).

We conjecture that the protocol HSP_N which is used for the lower bound actually realizes the maximal efficiency of a point-push protocol in the presence of N obstacles when $N \geq 3$. Figure 3(right) shows a symmetric view of the path of the single stirrer of HSP_4. The path lies on a hypotrochoid which lead to the terminology hypotrochoid stirring protocol (see Kobayashi [22] and Binder [3]). The fact that the path of the single stirrer can be generated by a disk rolling inside a circle has practical implications; it implies that the stirring protocol may be implemented using a fairly simple system of gears to move the stirrer [22; 14; 3]. It is also interesting to note that the spectral radii of the matrices used in Theorem 1.1 appear to be algebraic units of a special type. Numerical computations for $N \leq 20$ indicate that each $\rho(H^{(N)})$ is a Salem number and each $\rho(\hat{H}^{(N)})$ is a Pisot number. These classes of algebraic numbers arise in a variety of topological and dynamical circumstances, especially in optimization problems (see, for example, Ghate and Hironaka [17] or Lanneau and Thiffeault [24]). The generator matrices whose product defines $\hat{H}^{(N)}$ are the absolute values of those which define $\hat{H}^{(N)}$; the connection of this fact to the algebraic type of their spectral radii warrants further investigation.

2 Preliminaries

2.1 Linear algebra and the generalized spectral radius

We recall some basic facts and establish a few conventions. The identity matrix of the appropriate dimension in a given context is denoted I. An inequality of a pair of

vectors means that the inequality holds in each component; the analogous convention holds for inequalities of matrices of the same dimensions.

For a square matrix M, $\rho(M)$ will denote its spectral radius and for $1 \leq p \leq \infty$, $\|M\|_p$ is the matrix norm induced by the vector norm $\|\vec{v}\|_p$. Recall the basic facts that $\rho(M) \leq \|M\|_p$ and $\|M\|_2 = (\rho(MM^T))^{1/2}$. Gelfand's formula says that $\rho(M) = \lim_{n\to\infty} \|M^n\|_p^{1/n}$. This implies that if M and M' are nonnegative, square matrices with $M \geq M'$, then $\rho(M) \geq \rho(M')$.

For a given matrix M, we let $c(M)$ be the row vector of its column sums, and so $c_j(M) = \sum_i M_{ij}$. An easy computation yields that for any pair of square matrices, $c(MM') = c(M)M'$. Also note that for a positive matrix M, $\max_j(c_j(M)) = \|M\|_1$.

We shall also work with free modules over the ring of all Laurent polynomials in the variable t with integer coefficients, $\mathbb{Z}[t^{\pm 1}]$. The collection of $(n \times n)$–matrices over $\mathbb{Z}[t^{\pm 1}]$ is denoted $\mathrm{Mat}(n, \mathbb{Z}[t^{\pm 1}])$.

We recall the definition of the generalized spectral radius. For more information, see Jungers [21]. For a finite set of matrices $\mathcal{M} = \{M_1, \ldots, M_n\}$ and a $k > 0$ let

$$(2\text{-}1) \qquad \rho_k(\mathcal{M}) = \max\{\rho(M_{i_1} M_{i_2} \cdots M_{i_k})^{1/k} : M_{i_j} \in \mathcal{M}\}.$$

Since norms are submultiplicative,

$$(2\text{-}2) \qquad \rho_k(\mathcal{M}) \leq \max\{\|M_i\|_p : 1 \leq i \leq n\}.$$

Thus the *generalized spectral radius* defined as

$$\rho(\mathcal{M}) = \limsup_{k\to\infty} \rho_k(\mathcal{M})$$

is always finite. Note that since the spectral radius is multiplicative, for all $n > 0$, $\rho_{nk}(\mathcal{M}) \geq \rho_k(\mathcal{M})$, and thus

$$(2\text{-}3) \qquad \rho(\mathcal{M}) = \sup\{\rho_k(\mathcal{M}) : k \in \mathbb{Z}^+\}.$$

2.2 Word growth and the efficiency of automorphisms

Let G be a finitely generated group with generators $x_1, \ldots, x_m$. For $g \in G$, let $\ell(g)$ be the minimum word length in the generators. If $a\colon G \to G$ is an automorphism, its *entropy* is defined as

$$(2\text{-}4) \qquad h(a) = \max_i \limsup_{n\to\infty} \frac{\log(\ell(a^n(x_i)))}{n}.$$

Note that $h(a)$ does not depend on the choice of generators of G. If $\mathcal{A} \subset \mathrm{Aut}(G)$ is a semigroup generated by $\{a_1, \ldots, a_k\}$, for $a \in \mathcal{A}$ define its *efficiency* with respect to the given generating set as

$$\mathrm{eff}(a; \{a_1, \ldots, a_k\}) = \frac{h(a)}{\mathcal{L}(a)},$$

where $\mathcal{L}(a)$ is the minimal word length of a in the generators $\{a_i\}$. Finally, the *maximal efficiency* of $\mathcal{A}$ with respect to the given generating set is

$$\mathrm{Eff}(\mathcal{A}, \{a_1, \ldots, a_k\}) = \sup\{\mathrm{eff}(a; \{a_1, \ldots, a_k\}) : a \in \mathcal{A}\}.$$

Note that $\mathrm{eff}(a)$ and $\mathrm{Eff}(\mathcal{A})$ usually *do* depend on the choice of generators of $\mathcal{A}$.

Remark 2.1 If the group G is free abelian, then an automorphism A is a linear transformation and treating $g \in G$ as a vector $\vec{g}$, its word length is $\ell(g) = \|\vec{g}\|_1$. It then follows that $h(A) = \log(\rho(A))$. Further, as a consequence of (2-3), the efficiency of the semigroup $\mathcal{A}$ generated by a collection of linear isomorphisms $\{A_1, \ldots, A_n\}$ is the log of their generalized spectral radius,

$$\mathrm{Eff}(\mathcal{A}; \{A_1, \ldots, A_n\} = \log(\rho(\{A_1, \ldots, A_n\})).$$

Thus the entropy efficiency may be viewed as the nonlinear analog of the generalized spectral radius.

2.3 Braids

There are three of the many interpretations of braids which will be used here: as the motion of a collection of points in the disk, as physical braids, and as elements of a mapping class group. We describe each of these briefly and informally. For more information; see Birman [4] and Farb and Margalit [11]. Fix a model of the two-dimensional disk D^2 and n distinguished points $\{p_1, \ldots, p_n\} \subset \mathrm{Int}(D^2)$. The disk punctured at the distinguished points is $D_n := D^2 - \{p_1, \ldots, p_n\}$.

A motion of the distinguished points is described by a family of paths $\eta_i \colon [0, 1] \to D^2$ for $i = 1, \ldots, n$ with $\eta_i(t) \neq \eta_j(t)$ for all $t \in [0, 1]$ and $i \neq j$. To find the corresponding physical braid define arcs $\hat{\eta} \colon [0, 1] \to D^2 \times [0, 1]$ by $\hat{\eta}_i(t) = (\eta_i(t), 1-t)$. The physical braid corresponding to the motion of points is the image of $(\hat{\eta}_1, \ldots, \hat{\eta}_n)$. We can define the product of two point motions by following one by the other. This corresponds to stacking the associated physical braids one below the other.

The elements of Artin's braid group on n–strings, denoted B_n, correspond to equivalence classes of these objects. It can be considered as homotopy classes of point

paths or else of physical braids. There are $(n-1)$–standard generators of B_n with σ_i corresponding to switching the i–th and $(i+1)$–st strings (or points) counterclockwise. There are two types of relations which define the braid group: $\sigma_i \sigma_{i+1} \sigma_i = \sigma_{i+1} \sigma_i \sigma_{i+1}$ for all i, and $\sigma_i \sigma_j = \sigma_j \sigma_i$ for $|i - j| > 1$.

For the connection to a mapping class group, we let $\mathrm{MCG}(D_n)$ be the collection of isotopy classes of homeomorphisms of D_n, where the homeomorphisms and isotopies are required to fix the boundary of the disk pointwise. The collection of mapping classes $\mathrm{MCG}(D_n)$ becomes a group under composition. For each generator σ_i of B_n, let f_i be a homeomorphism that switches p_i and p_{i+1} counterclockwise. The map $\sigma_i \mapsto [f_i] \in \mathrm{MCG}(D_n)$ then yields an isomorphism $B_n \cong \mathrm{MCG}(D_n)$.

The *pure braid group* on n strings is the subgroup P_n of B_n corresponding to paths η_j with $\eta_j(0) = \eta_j(1)$. Thus the corresponding physical braids do not permute their endpoints and the mapping classes in D_n fix all the punctures.

Conventions There are many possible conventions for the braid group and its various interpretations. Since there is not much standardization, we summarize our conventions. Braid words are written from left to right, and so $\sigma_1 \sigma_2$ means perform σ_1 and then σ_2. In the corresponding physical braid, we draw σ_2 below σ_1. We maintain the same conventions when treating braids as mapping classes, and so $\sigma_1 \sigma_2$ means switch the first two punctures from the left counterclockwise and then switch the second two punctures. Thus the convention is the opposite of what one would use when composing functions. To be consistent with these conventions, when we have a matrix representation of the braid group, the matrices will act on a row vector from the right. Thus the matrices are the transpose of the most common convention. Specifically, the components of the image of the i–th basis vector gives the i–th row of the matrix.

2.4 Thurston–Nielsen theory and the entropy of a mapping class

Thurston's theorem classifies mapping classes on a surface M into three classes: pseudo-Anosov, finite order and reducible (see Thurston [29], Fathi, Laudenbach and Poenaru [12], Casson and Bleiler [8] and Farb and Margalit [11]). In each case the class contains what is now called the Thurston representative which has the same name, so pseudo-Anosov classes contain pseudo-Anosov homeomorphisms, etc. We define the (topological) entropy of a surface mapping class χ as

$$(2\text{-}5) \qquad h(\chi) = \inf\{h_{\mathrm{top}}(f) : f \in \chi\}.$$

Part of Thurston's classification theorem says that this infimum is achieved by the Thurston representative in the mapping class. In addition, the exponential word growth

under iteration by the mapping class on $\pi_1(M)$ is equal to $h(\chi)$. Specifically, if M is a surface with chosen basepoint $Q \in M$, pick a homeomorphism $f \in \chi$ with $f(Q) = Q$ and let $f_\sharp$ denote its action on $\pi_1(M, Q)$. Then treating $f_\sharp$ as an automorphism of $\pi_1(M, Q)$ and defining $h(f_\sharp)$ as in (2-4), Thurston's theorem says that $h(f_\sharp) = h(\chi)$. For a pseudo-Anosov class χ, the number $\lambda = \exp(h(\chi)) > 1$ is called the *dilation*.

We call a braid β pseudo-Anosov, finite order, or reducible according to whether its corresponding mapping class has that type, and we define its entropy $h(\beta)$ as in (2-5), again using its corresponding mapping class.

3 Point-push protocols and their topological efficiency

In accordance with the motivation from fluid mechanics we will sometimes use the terminology "stirring protocol" to refer to a braid or equivalently, a mapping class on the punctured disk. We study a special class of stirring protocols here in which just one stirrer moves. The other N stirrers (or punctures) are stationary and will be referred to as the *obstacles*. The initial and final point of the path of the single stirrer is denoted P, and so the protocol is given by a path $\Gamma : [0, 1] \to D_N$ with $\Gamma(0) = \Gamma(1) = P$. We may extend Γ to an isotopy f_t on D_N so that $f_t(P) = \Gamma(t)$ for all $t \in [0, 1]$. We further puncture D_N at P forming D_{N+1}, and the class of f_1 in MCG(D_{N+1}) is the point-push class corresponding to Γ. Since Γ is a closed loop starting and ending at P it gives an element of the fundamental group with basepoint P, $\pi_1(D_N, P)$. A theorem of Kra [23] and Birman [4] says that the point-push mapping class is uniquely determined by this element in the fundamental group and so elements of this class of stirring motions are sometimes called π_1–*protocols*.

3.1 The subgroup of point-push protocols

Each point-push protocol with N obstacles yields a mapping class that fixes the punctures and so they correspond to elements of the pure braid group on $(N+1)$ strings. The collection of such classes is a subgroup we denote $\mathcal{P}_N$. If $h : P_{N+1} \to P_N$ is the homomorphism which deletes the left-most string of the pure braid, then $\mathcal{P}_N = \ker(h)$. Equivalently, the point-push protocols correspond to mapping classes on D_{N+1} with the property that when the left-most puncture is filled in, the resulting class in MCG(D_N) is the identity.

To compute the entropy efficiency we need a specified set of generators of $\mathcal{P}_N$. For $i = 1, \ldots, N$, let

$$
(3\text{-}1) \qquad\qquad \alpha_i = \sigma_1 \sigma_2 \cdots \sigma_{i-1} \sigma_i^2 \sigma_{i-1}^{-1} \cdots \sigma_2^{-1} \sigma_1^{-1},
$$

where the σ_i are the standard generators of B_{N+1}. Note that we have adopted the convention that the stirrer P is the left-most puncture and the obstacles are the next N punctures. Thus in terms of a stirring motion, the generator α_i corresponds to the stirrer passing in front of the first $(i-1)$ obstacles from the left, going around the i–th obstacle (which is the $(i+1)$–st puncture) once counterclockwise, and then returning to its initial position passing once again in front of the first $(i-1)$ obstacles (see Figure 1). The generators α_j are the members of the standard generating set of the pure braid group P_{N+1} in which only the left-most strand does not go straight down. The next proposition follows from the theorem of Kra [23] and Birman [4] described above.

Proposition 3.1 *The group $\mathcal{P}_N$ of point-push protocols with N obstacles is a free group on the generators $\{\alpha_1, \ldots, \alpha_N\}$ and is naturally isomorphic to $\pi_1(D_N, P)$.*
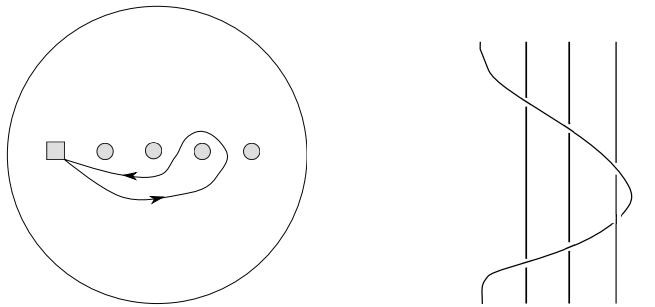

Figure 1: $N = 4$: Left, the path of α_3; Right, the braid of α_3

3.2 An adapted set of generators of $\pi_1(D_{N+1})$

In working with point-push protocols with N obstacles it will be frequently useful to use a collection of loops in $\pi_1(D_{N+1})$ which generate and are adapted to the protocols in the sense that the generators α_j of $\mathcal{P}_N$ act particularly simply on them. The case $N = 4$ is shown in Figure 2. The square represents the stirrer P. Fix a basepoint Q, and for $i = 1, \ldots, N$, the loop δ_i begins at Q and goes counterclockwise around the i–th obstacle. The last generator δ_{N+1} encloses the stirrer and all the obstacles going *clockwise* around the disk.

The reason for the last convention is to get cleaner formulas for the action of the α_j on the δ_i. We may assume that the generators $\{\alpha_1, \ldots, \alpha_N\}$ of $\mathcal{P}_N$ act while fixing the basepoint Q. Figure 2 right shows the action of α_3 when $N = 4$. Writing $\alpha_i(\delta_j)$ for
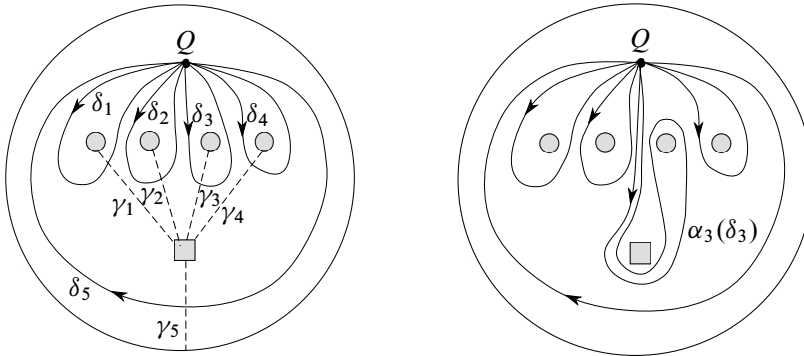
Figure 2: Left, the adapted generators δ_j for $\pi_1(D_5)$. The arcs γ_j are used in the construction of the covering spaces in Section 4.3. Right, the action of α_3

the action of α_i on the loop δ_i, we have

$$
\alpha_j(\delta_j) = (\delta_{j-1}^{-1}\delta_{j-2}^{-1}\cdots\delta_1^{-1}\delta_{N+1}^{-1}\delta_N^{-1}\cdots\delta_{j+1}^{-1})
$$

(3-2)
$$
\cdot\delta_j\,(\delta_{j+1}\cdots\delta_N\delta_{N+1}\delta_1\cdots\delta_{j-2}\delta_{j-1}),
$$

$$
\alpha_i(\delta_j) = \delta_j \quad \text{for } j \neq i.
$$

3.3 Topological efficiency of point-push protocols

For a point-push protocol $\chi \in \mathcal{P}_N$, pick a homeomorphism $h \in \chi$ with $h(Q) = Q$, and let $\chi_\sharp$ be the induced automorphism of $\pi_1(D_{N+1}, Q) \cong F_{N+1}$. Now the efficiency $\mathrm{eff}(\chi_\sharp; \{\alpha_1^{\pm 1}, \ldots, \alpha_N^{\pm 1}\})$ is independent of the choice of h and Q, and we use the notation

$$
\mathrm{eff}(\chi) := \mathrm{eff}(\chi_\sharp; \{\alpha_1^{\pm 1}, \ldots, \alpha_N^{\pm 1}\}).
$$

The maximal efficiency over the group $\mathcal{P}_N$ of point-push protocols with N obstacles is denoted here as

$$
\mathrm{Eff}(N) := \sup\{\mathrm{eff}(\chi) : \chi \in \mathcal{P}_N\}.
$$

Remark 3.2 To avoid potential confusion we emphasize that a point-push protocol with N obstacles is determined by an element of the fundamental group of the $N-$ punctured disk, $\pi_1(D_N)$, but after also puncturing at the stirrer P it yields a mapping class on the $(N+1)$–punctured disk or a braid on $(N+1)$ strings.

4 The lower bound

Since we are seeking $\mathrm{Eff}(N)$, the maximal efficiency of a point-push protocol with N obstacles, the efficiency of any single $\chi \in \mathcal{P}_N$ gives a lower bound. In the first

subsection we define $\mathrm{HSP}_N \in \mathcal{P}_N$ which numerically appears to give the maximal efficiency, and we shall see that its topological efficiency converges as $N \to \infty$ to the upper bound given in Section 5. Most of this section is concerned with the computation of the topological entropy of HSP_N.

4.1 The hypertrophied stirring protocol

Using the generators α_i of $\mathcal{P}_N$ we define the *hypotrochoid stirring protocol* with N obstacles as

$$\mathrm{HSP}_N := \alpha_1 \alpha_2 \cdots \alpha_N.$$

As with any braid (see Section 2.3), in various contexts we treat HSP_N as a physical braid, a motion of a collection of points (in this case just the single stirrer is moving), and a mapping class. Figure 3 left shows the corresponding motion of the stirrer when $N = 4$. Figure 3 right gives a topologically equivalent but more symmetric view in which the path of the stirrer is a hypotrochoid. The fluid mixing induced by physical realizations of HSP_N has been studied by Finn, Cox and Byrne [13], Kobayashi [22], Finn and Thiffeault [14] and Binder [3].

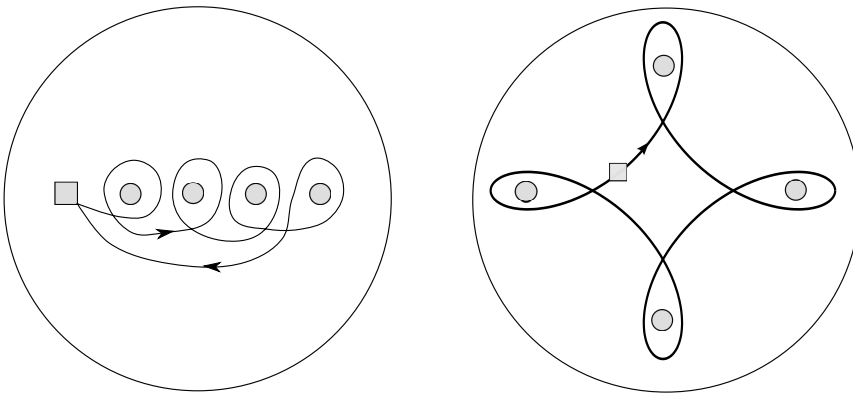Figure 3: HSP_4: Left, in the braid configuration; Right, on a hypotrochoid

The train track of a mapping class is a collapsed version the unstable foliation (see Penner and Harer [27], Casson and Bleiler [8] and Farb and Margalit [11]). Given combinatorial data about a mapping class, it is always computable in finite time by the Bestvina–Handel algorithm [2]. The edges of the train track give a Markov partition for a pseudo-Anosov representative ϕ in the class, and the induced transition on the edges from the action of ϕ gives the corresponding transition matrix. Figure 4 shows the invariant train track for HSP_4. Its invariance under the action of HSP_4 as well as the

fact that its transition matrix is strictly positive can be confirmed by a rather tedious and long graphical calculation or with a computer implementation of the Bestvina–Handel algorithm (see Hall [18]). For other values of N the repeated regularity of HSP_N ensures that the train tracks in those cases is constructed by a repetition of the simple one-prongs connected by a single edge seen on the last three right-hand punctures of HSP_4. The track demonstrates the important fact that the only singularities of the invariant foliations of HSP_N are one-prongs at the punctures and a singularity on the outer boundary. Also note that the train track and its transition matrix prove that each HSP_N is a pseudo-Anosov mapping class, though this is verifiable by elementary means or from Kra's Theorem [23].
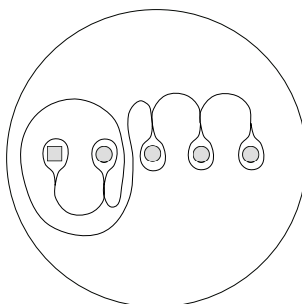

Figure 4: The invariant train track for HSP_4

4.2 The method for computing $h(\mathrm{HSP}_N)$

In principle for any N one can use the invariant train track of HSP_N to compute a Markov transition matrix for ϕ. The dilation of HSP_N is then the spectral radius of this matrix. However, this computation appears intractable and so we resort to homological methods.

As noted in Section 2.4, the topological entropy of a pseudo-Anosov map ϕ can be computed as the exponential growth rate of words in the fundamental group under the induced action of ϕ. In practise, this computation is very difficult. The corresponding computation in homology is much simpler, but usually just gives a lower bound (see Manning [25]). Fried [16] develops the idea that the action of the induced map on homology when lifted to finite covers (or equivalently, on twisted homology) is usually computable, often gives a better lower bound, and in certain cases, the bound is sharp.

To implement this program, rather than working with individual finite covers it is simpler to construct a $\mathbb{Z}$–covering space $\widetilde{M}$ and study its finite quotients. The first homology of $\widetilde{M}$ is treated as a module over $\mathbb{Z}[t^{\pm 1}]$ with t representing the generator

of the deck group. If the given map f (or the mapping class) lifts to $\widetilde{M}$ and commutes with the deck transformation T, then the lifted action of f is given by a matrix $B(t) \in \mathrm{Mat}(m, \mathbb{Z}[t, t^{-1}])$, where m is the rank of $H_1(\widetilde{M})$ as a $\mathbb{Z}[t^{\pm 1}]$ module. One may then show (see [1, Lemma 3.2]) that a substitution of $t = \exp(2\pi i p/q)$ into $B(t)$ gives the action of f on a piece of the first homology of the $\mathbb{Z}_q$–cover constructed by taking the quotient of $\widetilde{M}$ by T^q.

In the case of the n–punctured disk, there is a natural $\mathbb{Z}$–cover to which every mapping class lifts and commutes with the deck transformations. The matrix $B(t)$ giving the action of the mapping class (or braid) is the Burau representation and indeed, this is now the most common descriptions of the topological meaning of the Burau representation. Computing entropy of braids using substitutions into the Burau representation is considered in [1], and that paper is a good source for more information on the techniques used in this section.

4.3 The $\mathbb{Z}$–covering spaces

The mapping classes representing point-push protocols are rather special and so the general cover used in Burau representation is not the best way to proceed. Instead we construct a pair of $\mathbb{Z}$–covering spaces, $\widetilde{Z}_N$ and $\widetilde{Z}'_N$, which are adapted to the particular nature of the action of point-push protocol. Both $\widetilde{Z}_N$ and $\widetilde{Z}'_N$ are covering spaces of the $(N+1)$–punctured disk, D_{N+1}. The constructions are based on the generating loops $\{\delta_j\}$ for $\pi_1(D_{N+1})$ given in Section 3.2. We shall work with simplicial one-chains and give the relation to homology in Proposition 4.2.

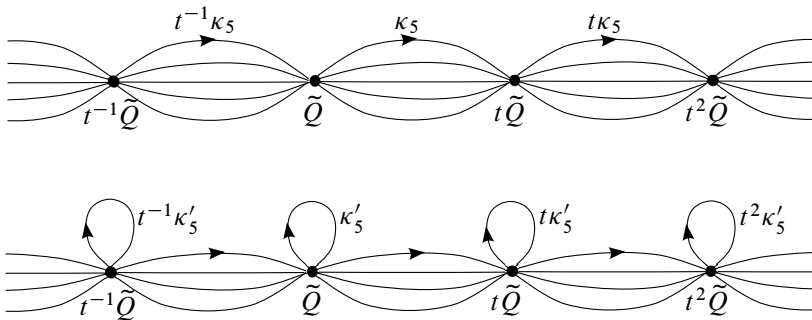Figure 5: Top, the spine of the covering space $\widetilde{Z}_4$; bottom, the spine of the covering space $\widetilde{Z}'_4$. Only selected simplices are labeled and have their orientations shown.

Geometrically, both $\mathbb{Z}$–covering spaces are constructed by cutting along arcs "dual" to the generating loops $\delta_1, \dots, \delta_{N+1}$ of $\pi_1(D_{N+1}, Q)$. The arcs γ_j for $j = 1, \dots, N$

connect the j–th obstacle to the stirrer, and γ_{N+1} connects the stirrer to the outer boundary (see Figure 2). For the cover $\widetilde{Z}_N$ one cuts along all these arcs and for $\widetilde{Z}'_N$ one cuts along all of the arcs except γ_{N+1}. After cutting along the arcs, copies of resulting space are then stacked and pasted into a $\mathbb{Z}$–cover.

For algebraic descriptions, recall that a covering space with abelian deck group of a manifold M corresponds to a subgroup $H \subset H_1(M; \mathbb{Z})$, and the deck group of the resulting cover is $H_1(M; \mathbb{Z})/H$. Thus a $\mathbb{Z}$–covering space of D_{N+1} is given by an epimorphism $\iota\colon H_1(D_{N+1}; \mathbb{Z}) \to \mathbb{Z}$. To define this epimorphism for our two covering spaces, for $j = 1, \ldots, N+1$, we let Δ_j be the element of $H_1(D_{N+1}; \mathbb{Z})$ corresponding to the generating loop δ_j. Thus the set $\{\Delta_1, \ldots, \Delta_{N+1}\}$ gives a basis for $H_1(D_{N+1}; \mathbb{Z})$. The covering spaces $\widetilde{Z}_N$ and $\widetilde{Z}'_N$ are then respectively defined by epimorphisms

$$(4\text{-}1) \qquad \iota\left(\sum_{i=1}^{N+1} c_i \Delta_i\right) = \sum_{i=1}^{N+1} c_i \quad \text{and} \quad \iota'\left(\sum_{i=1}^{N+1} c_i \Delta_i\right) = \sum_{i=1}^{N} c_i.$$

In both cases the generator of the deck group will be denoted T.

We now give the description of the one-chains in $\widetilde{Z}_N$ as a free $\mathbb{Z}[t^{\pm 1}]$ module. The wedge of $(N+1)$ circles given by the union of the generating loops $\{\delta_j\}$ of $\pi_1(D_{N+1}, Q)$ is denoted Λ_N and it is a spine of D_{N+1}. We work with the simplicial homology of the one-skeleton Λ_N and its lift to $\widetilde{Z}_N$, denoted $\widetilde{\Lambda}_N$. The simplicial one-chains over $\mathbb{Z}$ are denoted $C_1(\Lambda_N)$ and $C_1(\widetilde{\Lambda}_N)$, respectively. Fix a lift $\widetilde{Q}$ of the basepoint Q. For $j = 1, \ldots, N+1$, let κ_j be the oriented arc (or one-chain) in $\widetilde{Z}_N$ which is a lift of δ_i and connects $\widetilde{Q}$ to $T(\widetilde{Q})$. The collection $\{\kappa_1, \ldots, \kappa_{N+1}\}$ will be our basis for $C_1(\widetilde{\Lambda}_N)$, and for $n \in \mathbb{Z}$, the formal expression $t^n \kappa_j$ refers to the one-chain $T_*^n(\kappa_i)$, where $T_*\colon C_1(\widetilde{\Lambda}_N; \mathbb{Z}) \to C_1(\widetilde{\Lambda}_N; \mathbb{Z})$ is the induced action of the deck transformation T. The one-chains in $\widetilde{\Lambda}_N$ are thus the module $C_1(\widetilde{Z}_N; \mathbb{Z}) \cong \mathbb{Z}[t^{\pm 1}]^{N+1}$, where the vector of Laurent polynomials $(q_1, \ldots, q_{N+1})$ with each $q_j(t) \in \mathbb{Z}[t, t^{-1}]$ is identified with the one-chain in $\widetilde{\Lambda}_N$ given by $\sum q_j(T_*)\kappa_j$.

The one-chains in the spine $\widetilde{\Lambda}'_N$ of $\widetilde{Z}'_N$ have a similar description. However in this case δ_{N+1} doesn't lift to an arc but rather lifts to a loop based at $\widetilde{Q}$ which is denoted κ'_{N+1}. We then use the simplicial one-chains $\{\kappa_1, \ldots, \kappa_N, \kappa'_{N+1}\}$ as a basis for $C_1(\widetilde{\Lambda}'_N)$ over $\mathbb{Z}[t^{\pm 1}]$ and proceed as with $\widetilde{Z}_N$.

4.4 The lifted action of a point-push protocol

We next consider the lifted action of a point-push protocol $\chi \in \mathcal{P}_N$. Pick a homeomorphism $h \in \chi$ which fixes the basepoint Q and let h^Δ be its induced action on

the spine Λ_N. More precisely, let $h^\Delta\colon \Lambda_N \to \Lambda_N$ be a simplicial map homotopic to $r \circ h|_{\Lambda_N}$, where $r\colon D_{N+1} \to \Lambda_N$ is a deformation retract. Since h fixes all the punctures in D_{N+1}, it acts as the identity on $H_1(D_{N+1};\mathbb{Z})$. Thus h lifts to $\widetilde{Z}_N$ and commutes with all deck transformations. Let $\widetilde{h}$ be the lift with $\widetilde{h}(\widetilde{Q}) = \widetilde{Q}$. The simplicial map h^Δ also lifts to $\widetilde{h}^\Delta\colon \widetilde{\Lambda}_N \to \widetilde{\Lambda}_N$. Since $\widetilde{h}^\Delta$ commutes with the deck transformation, $\widetilde{h}^\Delta_*\colon C_1(\widetilde{\Lambda}_N) \to C_1(\widetilde{\Lambda}_N)$ can be represented by a matrix $C(t) \in \mathrm{Mat}(N+1, \mathbb{Z}[t^{\pm 1}])$. The assignment $\chi \mapsto C(t)$ is a representation of the group of point-push protocols with N obstacles which we denote $\Phi\colon \mathcal{P}_N \to \mathrm{Mat}(N+1, \mathbb{Z}[t, t^{-1}])$. It is clear that the construction does not depend on the choice of $h \in \chi$ or its retract h^Δ as long as $\widetilde{h}^\Delta(\widetilde{Q}) = \widetilde{Q}$. If $\widetilde{h}^\Delta(\widetilde{Q}) = T^n(\widetilde{Q})$, then the corresponding matrix is multiplied by t^n. Thus we adopt the notation $\widetilde{\chi}_* = \widetilde{h}^\Delta_*$ with the understanding that it is always computed using an $\widetilde{h}^\Delta$ with $\widetilde{h}^\Delta(\widetilde{Q}) = \widetilde{Q}$. The same considerations for the cover $\widetilde{Z}'_N$ yields a representation $\Phi'\colon \mathcal{P}_N \to \mathrm{Mat}(N+1, \mathbb{Z}[t^{\pm 1}])$, and we sometimes write $\Phi'(\chi) = C'(t)$ as an abbreviated notation.

Using their action on the generators δ_j it is straightforward now to compute the action of the lift of generators α_j of $\mathcal{P}_{N+1}$ on the chosen basis of simplicial one-chains in the $\mathbb{Z}$–covering space. This procedure is described in [16]. We let $\widetilde{\alpha}_j(\kappa_j)$ denote the lifted action of α_j on the simplex κ_j.

Example 4.1 We illustrate the general result by computing the lifted action of α_3 on $\widetilde{\Lambda}_4$. The action on π_1 from (3-2) is

$$(4\text{-}2) \qquad \alpha_3(\delta_3) = \delta_2^{-1}\delta_1^{-1}\delta_5^{-1}\delta_4^{-1}\delta_3\delta_4\delta_5\delta_1\delta_2,$$

and $\alpha_3(\delta_j) = \delta_j$ for $j \neq 3$. For the lifted action to the generators of the simplicial chains in $\widetilde{\Lambda}$, it is immediate that $\widetilde{\alpha}_3(\kappa_j) = \kappa_j$ for $j \neq 3$, so we need only compute $\widetilde{\alpha}_3(\kappa_3)$. Since our convention is to represent mapping classes with homeomorphisms whose lifts fix $\widetilde{Q}$, the image of κ_3 must be an arc beginning at $\widetilde{Q}$. Since from (4-2) $\alpha_3(\delta_3)$ begins with δ_2^{-1}, the beginning of $\widetilde{\alpha}_3(\kappa_3)$ must traverse a copy of κ_2 in the negative direction and thus must in fact be $-t^{-1}\kappa_2$. The image of δ_3 continues with δ_1^{-1} and so the image of κ_3 continues with $-t^{-2}\kappa_1$. Continuing in this manner the next stage of the image $\widetilde{\alpha}_3(\kappa_3)$ is always forced by the next generator (or its inverse) in the image $\alpha_3(\delta_3)$. The full image is then

$$\widetilde{\alpha}_3(\kappa_3) = -t^{-1}\kappa_2 - t^{-2}\kappa_1 - t^{-3}\kappa_5 - t^{-4}\kappa_4 + t^{-4}\kappa_3 + t^{-3}\kappa_4 + t^{-2}\kappa_5 + t^{-1}\kappa_1 + \kappa_2$$

$$= \kappa_2 + t^{-1}(\kappa_1 - \kappa_2) + t^{-2}(\kappa_5 - \kappa_1) + t^{-3}(\kappa_4 - \kappa_5) + t^{-4}(\kappa_3 - \kappa_4).$$

The process of deriving the formulas for the action of α_3 lifted to $\widetilde{\Lambda}'_4$ is similar, but now instead of traversing the arcs corresponding to lifts of δ_5 to $\widetilde{\Lambda}_4$, one goes around the loops $t^n\kappa'_5$ which are the lifts of δ_5 to $\widetilde{\Lambda}'_4$.

The same procedure in general yields the following.

$$(4\text{-}3) \quad \begin{cases} \widetilde{\alpha}_1(\kappa_1) = \kappa_{N+1} + t^{-1}(\kappa_N - \kappa_{N+1}) + \cdots + t^{-N}(\kappa_1 - \kappa_2), \\ \widetilde{\alpha}_j(\kappa_j) = \kappa_{j-1} + t^{-1}(\kappa_{j-2} - \kappa_{j-1}) & \text{when } j > 1, \\ \qquad + \cdots + t^{-(j-2)}(\kappa_1 - \kappa_2) + t^{-(j-1)}(\kappa_{N+1} - \kappa_1) \\ \qquad + t^{-j}(\kappa_N - \kappa_{N+1}) + t^{-(j+1)}(\kappa_{N-1} - \kappa_N) \\ \qquad + \cdots + t^{-N}(\kappa_j - \kappa_{j+1}) \\ \widetilde{\alpha}_i(\kappa_j) = \kappa_j & \text{when } j \neq i. \end{cases}$$

Letting $\widetilde{\alpha}_i'$ the lifted action of α_i on one-chains in $\widetilde{\Lambda}'$ we have in general

$$(4\text{-}4) \quad \begin{cases} \widetilde{\alpha}_1'(\kappa_1) = \kappa_1 + \kappa_{N+1} + t^{-1}(\kappa_N - \kappa_1 - \kappa_{N+1}) \\ \qquad + t^{-2}(\kappa_{N-1} - \kappa_N) + \cdots + t^{-(N-1)}(\kappa_2 - \kappa_3), \\ \widetilde{\alpha}_2'(\kappa_2) = t\kappa_{N+1} + (\kappa_N - \kappa_{N+1}) \\ \qquad + t^{-1}(\kappa_{N-1} - \kappa_N) + \cdots + t^{-(N-1)}(\kappa_1 - \kappa_2), \\ \widetilde{\alpha}_j'(\kappa_j) = \kappa_{j-1} + t^{-1}(\kappa_{j-2} - \kappa_{j-1}) & \text{when } j > 1, \\ \qquad + \cdots + t^{-(j-3)}(\kappa_2 - \kappa_3) + t^{-(j-2)}(\kappa_1 - \kappa_2 + \kappa_{N+1}) \\ \qquad + t^{-(j-1)}(\kappa_N - \kappa_1 - \kappa_{N+1}) + t^{-j}(\kappa_{N-1} - \kappa_N) \\ \qquad + \cdots + t^{-(N-1)}(\kappa_j - \kappa_{j+1}) \\ \widetilde{\alpha}_i'(\kappa_j) = \kappa_j & \text{when } j \neq i. \end{cases}$$

4.5 The double cover

The two-fold covering spaces formed by taking the quotient of $\widetilde{Z}_N$ and $\widetilde{Z}_N'$ by the square of the deck transformation will be central to our calculations. We recall some basic facts concerning the connection of various objects in the full $\mathbb{Z}$–covering spaces, the base space and this two-fold cover, and then compute the action on homology in this cover. We will initially just discuss $\widetilde{Z}_N$ and comment on $\widetilde{Z}_N'$ afterwards. To simplify notation fix N and drop it from the notation, so now $\widetilde{Z} = \widetilde{Z}_N$.

Recall that the $\mathbb{Z}$–covering space $p: \widetilde{Z} \to D_{N+1}$ corresponds to an epimorphism $\iota: H_1(D_{N+1}; \mathbb{Z}) \to \mathbb{Z}$. The two-fold covering space $p^{(2)}: \widetilde{Z}^{(2)} \to D_{N+1}$ corresponds to the epimorphism $\iota^{(2)}: H_1(D_{N+1}; \mathbb{Z}) \to \mathbb{Z}_2$ where $\iota^{(2)} := \pi^{(2)} \circ \iota$, with $\pi^{(2)}: \mathbb{Z} \to \mathbb{Z}_2 = \mathbb{Z}/(2\mathbb{Z})$ the projection. Thus $\widetilde{Z}^{(2)} = \widetilde{Z}/T^2$, and we denote the deck transformation of $\widetilde{Z}^{(2)}$ as $T^{(2)}$. The corresponding one-skeleton of $\widetilde{Z}^{(2)}$ is denoted $\widetilde{\Lambda}^{(2)}$.

If τ is an indeterminate with $\tau^2 = 1$, then $C_1(\widetilde{\Lambda}^{(2)})$ is a $\mathbb{Z}[\tau]$ module with $C_1(\widetilde{\Lambda}^{(2)}) \cong \mathbb{Z}[\tau]^{N+1}$. The action of $\widetilde{\alpha}_i$ is denoted as $\widetilde{\alpha}_i^{(2)}$ and using (4-3) we get

$$(4\text{-}5) \quad \begin{cases} \widetilde{\alpha}_j^{(2)}(\kappa_j) = (-1)^j (1-\tau)\kappa_1 + (-1)^{j+1}(1-\tau)\kappa_2 & \text{when } N \text{ is even,} \\ \qquad + \cdots + (1-\tau)\kappa_{j-1} + \kappa_j - (1-\tau)\kappa_{j+1} \\ \qquad + \cdots + (-1)^{(j+1)}(1-\tau)\kappa_{N+1} \\ \widetilde{\alpha}_j^{(2)}(\kappa_j) = (-1)^j (1-\tau)\kappa_1 + (-1)^{j+1}(1-\tau)\kappa_2 & \text{when } N \text{ is odd,} \\ \qquad + \cdots + (1-\tau)\kappa_{j-1} - \tau\kappa_j + (1-\tau)\kappa_{j+1} \\ \qquad + \cdots + (-1)^{(j+1)}(1-\tau)\kappa_{N+1} \\ \widetilde{\alpha}_i^{(2)}(\kappa_j) = \kappa_j & \text{for all } N \text{ when } j \neq i. \end{cases}$$

Similarly, $\widetilde{Z}'^{(2)} = \widetilde{Z}'/T^2$ is defined by the epimorphism $\iota^{(2)} := \pi^{(2)} \circ \iota'$. Its one-skeleton is $\widetilde{\Lambda}'^{(2)}$ and $C_1(\widetilde{\Lambda}'^{(2)}) \cong \mathbb{Z}[\tau]^{N+1}$ and using (4-4) the action of $\widetilde{\alpha}_i$ is

$$(4\text{-}6) \quad \begin{cases} \widetilde{\alpha}_j'^{(2)}(\kappa_j) = (-1)^j (1-\tau)\kappa_1 + (-1)^{j+1}(1-\tau)\kappa_2 & \text{when } N \text{ is even,} \\ \qquad + \cdots + (1-\tau)\kappa_{j-1} - \tau\kappa_j + (1-\tau)\kappa_{j+1} \\ \qquad + \cdots + (-1)^j (1-\tau)\kappa_{N+1} \\ \widetilde{\alpha}_j'^{(2)}(\kappa_j) = (-1)^j (1-\tau)\kappa_1 + (-1)^{j+1}(1-\tau)\kappa_2 & \text{when } N \text{ is odd,} \\ \qquad + \cdots + (1-\tau)\kappa_{j-1} + \kappa_j - (1-\tau)\kappa_{j+1} \\ \qquad + \cdots + (-1)^j (1-\tau)\kappa_{N+1} \\ \widetilde{\alpha}_i'^{(2)}(\kappa_j) = \kappa_j & \text{for all } N \text{ when } j \neq i. \end{cases}$$

Let χ be a point-push protocol treated as a mapping class. The projection of the action $\widetilde{\chi}_*$ on $C_1(\widetilde{\Lambda})$ to one on $C_1(\widetilde{\Lambda}^{(2)})$ is denoted $\widetilde{\chi}_*^{(2)}$. The action of χ on $H_1(\widetilde{\Lambda}^{(2)}; \mathbb{Z}) \cong H_1(\widetilde{Z}^{(2)}; \mathbb{Z})$ is written $\widetilde{\chi}_\star^{(2)}$. The same notation is used for $\widetilde{Z}'^{(2)}$, but with a prime attached.

The next proposition gives the connection of the substitution $t = -1$ into the matrices $C(t)$ and $C'(t)$ and the growth rate of the action of χ on the first homology of the two-fold cover $\widetilde{Z}^{(2)}$. It follows from [1, Lemma 3.2] and the fact that the growth on chains is the same as that on homology. For sake of having the paper self-contained, we give the elementary proof.

Proposition 4.2 *Assume that $\chi \in \mathcal{P}_N$ is a given point-push protocol. As constructed above, $\widetilde{Z}^{(2)}$ and $\widetilde{Z}'^{(2)}$ are double covers and $\widetilde{\chi}_\star^{(2)}$ and $\widetilde{\chi}_\star'^{(2)}$ are the induced actions on $H_1(\widetilde{Z}^{(2)})$ and $H_1(\widetilde{Z}'^{(2)})$, respectively, and Φ and Φ' are representations of $\mathcal{P}_N$. Letting $C(t) = \Phi(\chi)$ and $C'(t) = \Phi'(\chi)$, then*

$$\rho(\widetilde{\chi}_\star^{(2)}) = \rho(C(-1)) \quad \text{and} \quad \rho(\widetilde{\chi}_\star'^{(2)}) = \rho(C'(-1)).$$

Proof We will initially just discuss $\widetilde{Z}_N$ and comment on $\widetilde{Z}'_N$ afterwards.

Since $(T^{(2)})^2 = \mathrm{id}$, the eigenvalues of its action on simplicial one-chains in $\widetilde{\Lambda}^{(2)}$ are ± 1, and we denote the corresponding eigenspace decomposition as $C_1(\widetilde{\Lambda}^{(2)}) = V_+ \oplus V_-$. The subspaces $V_\pm$ are also characterized as $V_+ = \{c + T_*^{(2)}(c) : c \in C_1(\widetilde{\Lambda}^{(2)})\}$ and $V_- = \{c - T_*^{(2)}(c) : c \in C_1(\widetilde{\Lambda}^{(2)})\}$. Since $\widetilde{\chi}_*^{(2)}$ commutes with $T_*^{(2)}$, we have that $\widetilde{\chi}_*^{(2)}$ preserves the decomposition $C_1(\widetilde{\Lambda}^{(2)}) = V_+ \oplus V_-$.

Treating $C_1(\widetilde{\Lambda}^{(2)}) \cong \mathbb{Z}[\tau]^{N+1}$, the eigensubspaces have alternative descriptions as $V_+ = \{c + \tau c : c \in C_1(\widetilde{\Lambda}^{(2)})\}$ and $V_- = \{c - \tau c : c \in C_1(\widetilde{\Lambda}^{(2)})\}$. The homomorphism $\widetilde{\chi}_*^{(2)}$ acts on $C_1(\widetilde{\Lambda}^{(2)})$ as the matrix $C(\tau)$, and we will write $C(\tau) = B_0 + \tau B_1$ with $B_i \in \mathrm{Mat}(N+1, \mathbb{Z})$. Thus acting on an element of V_+ we have $(c + \tau c)C(\tau) = (cB_0 + cB_1) + \tau(cB_1 + cB_0) = (c + \tau c)(B_0 + B_1) = (c + \tau c)C(1)$ Similarly acting on an element of V_- we have $(c - \tau c)C(\tau) = (c - \tau c)C(-1)$. So $\widetilde{\chi}_*^{(2)}$ acts on V_+ by $C(1)$ and on V_- by $C(-1)$. Substituting $\tau = 1$ into (4-5), we find $C(1) = I$ (this also follows from the fact that $p_*^{(2)} : C_1(\widetilde{\Lambda}^{(2)}) \to C_1(D_{N+1})$ is an isomorphism when restricted to V_+). Since $\det(C(-1)) = 1$, $\rho(C(-1)) \geq 1$ and so $\rho(\widetilde{\chi}_*^{(2)}) = \rho(C(-1))$.

To find out what this means for homology we go back to working over $\mathbb{Z}$. For $1 \leq i \leq N+1$, let $\overline{\kappa}_i = T_*^{(2)}(\kappa_i)$, and so $\{\kappa_i, \overline{\kappa}_i\}$ forms a basis for $C_1(\widetilde{\Lambda}^{(2)}; \mathbb{Z})$. Now for $1 \leq i \leq N$, let $\Gamma_i = \kappa_i - \kappa_{i+1}$ and $\overline{\Gamma}_i = \overline{\kappa}_i - \overline{\kappa}_{i+1}$. Letting $\widehat{\Gamma} = \kappa_{N+1} + \overline{\kappa}_{N+1}$ it is easy to check that $\{\Gamma_1, \ldots, \Gamma_N, \overline{\Gamma}_1, \ldots, \overline{\Gamma}_N, \widehat{\Gamma}, \kappa_{N+1}\}$ is also a basis for $C_1(\widetilde{\Lambda}^{(2)}; \mathbb{Z})$ and we define V as the linear span $V = \langle \Gamma_1, \ldots, \Gamma_N, \overline{\Gamma}_1, \ldots, \overline{\Gamma}_N, \widehat{\Gamma} \rangle$ Thus $C_1(\widetilde{\Lambda}^{(2)}; \mathbb{Z}) = V \oplus \langle \kappa_{N+1} \rangle$. Now $\widehat{\Gamma}$ and all of the Γ_i and $\overline{\Gamma}_i$ are cycles because there are no two cells in $\widetilde{\Lambda}^{(2)}$. In fact, V is a concrete realization of the simplicial homology $H_1(\widetilde{\Lambda}^{(2)}; \mathbb{Z})$. Either by direct calculation using (4-5), or from general principles, $\widetilde{\chi}_*^{(2)}(V) = V$. Also from (4-5), $\widetilde{\alpha}_j^{(2)}(\kappa_{N+1}) = \kappa_{N+1}$, and so we see that the spectral radius of $\widetilde{\chi}_*^{(2)}$ acting on all of $C_1(\widetilde{\Lambda}^{(2)}; \mathbb{Z})$ is the same as the spectral radius of its restriction to V. But $\widetilde{\chi}_*^{(2)}$ restricted to V is exactly $\widetilde{\chi}_\star^{(2)}$, the induced action of $\widetilde{h}^{(2)}$ on $H_1(\widetilde{Z}^{(2)}; \mathbb{Z})$ and so by the conclusion of the previous paragraph, $\rho(\widetilde{\chi}_\star^{(2)}) = C(-1)$.

For the other cover $\widetilde{Z}'^{(2)}$, the proof on the chain level that $\rho(\widetilde{\chi}'^{(2)}_*) = \rho(C'(-1))$ is exactly the same as for $\widetilde{Z}^{(2)}$. A slight alteration is required when passing to homology because δ_{N+1} lifts to a loop in $\widetilde{Z}'^{(2)}$: for $1 \leq i \leq N-1$, let $\Gamma_i = \kappa_i - \kappa_{i+1}$ and $\overline{\Gamma}_i = \overline{\kappa}_i - \overline{\kappa}_{i+1}$, and finally let $\widehat{\Gamma} = \kappa_1 + \overline{\kappa}_1$. It is easy to check that $\{\Gamma_1, \ldots, \Gamma_{N-1}, \kappa_{N+1}, \overline{\Gamma}_1, \ldots, \overline{\Gamma}_N, \widehat{\Gamma}, \kappa_1\}$ is a basis for $C_1(\widetilde{\Lambda}'^{(2)}; \mathbb{Z})$, and we let $V' = \langle \Gamma_1, \ldots, \Gamma_{N-1}, \kappa_{N+1}, \overline{\Gamma}_1, \ldots, \overline{\Gamma}_N, \widehat{\Gamma} \rangle$. So $C_1(\widetilde{\Lambda}'^{(2)}; \mathbb{Z}) = V' \oplus \langle \kappa_1 \rangle$, and V' is a concrete realization of $H_1(\widetilde{\Lambda}'^{(2)}; \mathbb{Z})$ with $\widetilde{\chi}'^{(2)}_*(V') = V'$. From (4-6), $\widetilde{\alpha}'^{(2)}_1(\kappa_1) = \tau\kappa_1 + \ldots$ when N is even, and $\widetilde{\alpha}'^{(2)}_1(\kappa_1) = \kappa_1 + \ldots$ when N is odd, and $\widetilde{\alpha}'^{(2)}_j(\kappa_1) = \kappa_1$ for $j \neq 1$, and so the spectral radius of $\widetilde{\chi}'^{(2)}_*$ acting on all of $C_1(\widetilde{\Lambda}'^{(2)}; \mathbb{Z})$ is the same as the spectral radius of its restriction to V'. But recall that

$\widetilde{\chi}_*'^{(2)}$ restricted to V' is exactly $\widetilde{\chi}_\star'^{(2)}$, the induced action of $\widetilde{\chi}'^{(2)}$ on $H_1(\widetilde{\Lambda}'^{(2)}; \mathbb{Z})$. Therefore by the conclusion of the first paragraph, $\rho(\widetilde{\chi}_\star'^{(2)}) = C'(-1)$. $\qquad\qquad\qquad\qquad\square$

4.6 Another representation

Proposition 4.2 shows the significance of the substitution $t = -1$ into the matrices $C(t)$ and $C'(t)$. It will turn out that the matrix of significance for computing $h(\mathrm{HSP}_N)$ will be $C(t)$ when N is even and $C'(t)$ when N is odd. Conveniently, the substitution of $t = -1$ into (4-3) (or $\tau = -1$ into (4-5)) when N is even yields an identical formula to the result of substituting $t = -1$ into (4-4) when N is odd.

Now note that for all j, $\widetilde{\alpha}_j^{(2)}(\kappa_{N+1}) = \kappa_{N+1}$, and so for any product of the $\widetilde{\alpha}_j^{(2)}$, κ_N is an eigenvector with eigenvalue 1. Since we are interested in exponential growth rates, we may ignore κ_{N+1}. For $1 \leq k \leq N$, define the $(N \times N)$–matrix $T^{(k)}$ by

$$(4\text{-}7) \qquad \begin{cases} T_{kj}^{(k)} = 2 \cdot (-1)^{k-j+1} & \text{for } 1 \leq j \leq k-1, \\ T_{kk}^{(k)} = 0, \\ T_{kj}^{(k)} = 2 \cdot (-1)^{j-k} & \text{for } k+1 \leq j \leq N, \\ T_{ij}^{(k)} = 0 & \text{for } i \neq k, \end{cases}$$

and let

$$E^{(k)} = I + T^{(k)}.$$

Extend the assignment $\widehat{\alpha}_k \mapsto E^{(k)}$ to a homomorphism $\Psi \colon \mathcal{P}_N \to \mathrm{SL}(N, \mathbb{Z})$. Thus $\Psi(\chi)$ is the matrix $C(-1)$ with the last row and column deleted for N even and $\Psi(\chi)$ is the matrix $C'(-1))$ with the last row and column deleted for N odd.

4.7 A trace computation

Fix an $N \geq 2$. Recall that $\mathrm{HSP}_N = \alpha_1 \cdots \alpha_N$. We will compute $\mathrm{trace}(\Psi(\mathrm{HSP}_N))$ inductively using the initial segments of the product $\alpha_1 \cdots \alpha_k$ for $k = 1, \ldots, N$. These products have the matrix representation $H^{(k)} := \Psi(\alpha_1 \cdots \alpha_k) = E^{(1)} E^{(2)} \cdots E^{(k)}$.

Lemma 4.3 *With Ψ and HSP_N as defined above, $\mathrm{trace}(\Psi(\mathrm{HSP}_N)) = -3^N + 3N + 1$.*

Proof It follows easily from the definitions that for any $1 \leq a, b \leq N$,

$$(4\text{-}8) \qquad H_{a,b}^{(k+1)} = H_{a,b}^{(k)} + H_{a,k+1}^{(k)} \cdot T_{k+1,b}^{(k+1)}.$$

We now prove by induction on $k < N$ that for $m = 1, \ldots, N - k$,

(4-9)
$$\begin{cases} H^{(k)}_{i,k+m} = 2 \cdot 3^{k-i}(-1)^{i+k+m} & 1 \le i \le k, \\ H^{(k)}_{i,k+m} = 0 & k+1 \le i \le N \text{ and } i \ne k+m, \\ H^{(k)}_{k+m,k+m} = 1. \end{cases}$$

The case $k = 1$ is trivial. So assume (4-9) is true for k. Thus using (4-8). for $1 \le i \le k$ and $m = 2, \ldots, N - k$,

$$\begin{aligned} H^{(k+1)}_{i,k+m} &= H^{(k)}_{i,k+m} + H^{(k)}_{i,k+1} T^{(k+1)}_{k+1,k+m} \\ &= 2 \cdot 3^{k-i}(-1)^{i+k+m} + 2 \cdot 3^{k-i}(-1)^{i+k+1} \cdot 2(-1)^{k+m-(k+1)} \\ &= 2 \cdot 3^{k+1-i}(-1)^{i+k+1+m}, \end{aligned}$$

as required. The other values of i are special cases easily checked, completing the proof of (4-9).

Now we prove that for $1 \le k \le N$, $\text{trace}(H^{(k)}) = -3^k + 2k + N + 1$, again by induction on k, with the case $k = 1$ being trivial. Using (4-8) and the formulas (4-9),

$$\begin{aligned} \text{trace}(H^{(k+1)}) &= \text{trace}(H^{(k)}) + \sum_{a=1}^{N} H^{(k)}_{a,k+1} T^{(k+1)}_{k+1,a} \\ &= \text{trace}(H^{(k)}) + \sum_{a=1}^{k} 2 \cdot 3^{k-a}(-1)^{a+k+1} 2(-1)^{k-a} \\ &= -3^{k+1} + 2(k+1) + N + 1. \end{aligned}$$

The statement in the lemma now follows by letting $k = N$. □

4.8 The orientation double cover of a pseudo-Anosov map

A pseudo-Anosov map ϕ is characterized by the existence of a pair of transverse invariant measured foliations. If these foliations are oriented, it is now standard that the dilation of the pseudo-Anosov map is the spectral radius of the action of ϕ on first homology. For pseudo-Anosov maps on the punctured disk and in many other cases, a pseudo-Anosov map's foliations are not oriented. However, one can lift ϕ and its foliations to a new pseudo-Anosov map $\vec{\phi}$ on a branched cover $\vec{M}$ where there foliations are oriented, and thus the dilation is detectable using homology. The construction of this cover requires a fair amount of a priori information about ϕ, and the exact connection between ϕ and $\vec{\phi}$ can be difficult to ascertain, so in many situations the orientation double cover is not a useful tool. However, because the family HSP_N

lives on a genus zero surface and its foliations are of a particularly simple type, it will turn out that its orientation double cover always corresponds to one of the two-fold finite covers constructed by taking quotients of the $\mathbb{Z}$–covering spaces $\widetilde{Z}^{(2)}$ and $\widetilde{Z}'^{(2)}$ which were considered in Section 4.5.

A few definitions are required to give a characterization applicable to the case at hand. Assume ϕ is a pseudo-Anosov map on the n–punctured disk D_n with unstable foliation $\mathcal{F}^u$. Let S be the collection of singularities of $\mathcal{F}^u$ which are of odd order and contained in the interior of D_n. In addition, P denotes the collection of punctures of D_n and $S' \subset P$ is the subset consisting of those punctures at which the singularity of $\mathcal{F}^u$ is of odd order. The orientation double cover of ϕ is the two-fold cover branched over $S \cup S'$. More formally, first further puncture D_n at S to form $M' = D_n - S$. For each $p_i \in P \cup S$, let Γ_i be a small loop going counterclockwise around p_i, and so the homology classes $\{[\Gamma_i]\}$ are a basis for $H_1(M'; \mathbb{Z})$. If I is the set of indices i so that $p_i \in S' \cup S$, define the epimorphism $H_1(M'; \mathbb{Z}) \to \mathbb{Z}_2$ by

$$\sum c_i[\Gamma_i] \mapsto \sum_{i \in I} c_i \bmod 2.$$

The corresponding two-fold cover over M' is denoted $\widehat{M}'$. Finally, fill in the punctures in $\widehat{M}'$ that project to points in S to form $\widehat{M}''$ which is a branched covering space over D_n. The following is standard: for example, it follows from in [1, Lemma 2.2].

Proposition 4.4 *Let ϕ be pseudo-Anosov map on the n–punctured disk D_n and $\widehat{M}''$ the branched double cover of D_n formed by branching over the singularities of the foliations of ϕ which have odd order as was just constructed. Then $\widehat{M}''$ is the orientation double cover of ϕ.*

4.9 The entropy of HSP$_N$ and the lower bound

We now have all the ingredients needed to obtain our lower bound.

Theorem 4.5 *The topological entropy of the hypotrochoid stirring protocol with N obstacles is given by the log of the spectral radius of the matrix $H^{(N)}$ defined in the previous subsection,*

(4-10)
$$h(\mathrm{HSP}_N) = \log(\rho(H^{(N)})).$$

As a consequence,

$$\log\left(\frac{3^N - 3N - 1}{N}\right) \le h(\mathrm{HSP}_N),$$

and thus the maximal efficiency of a point-push protocol with N obstacles satisfies

$$\frac{\log(3^N - 3N - 1) - \log(N)}{N} \leq \mathrm{Eff}(N).$$

Proof The main observation we need is that the double covers $\tilde{Z}^{(2)}$ and $\tilde{Z}'^{(2)}$ from Section 4.5 are identical to the orientation double covers of a pseudo-Anosov map $\phi \in \mathrm{HSP}_N$ when N is even and odd, respectively. Once again we discuss the case of $\tilde{Z}^{(2)}$ first. Recall that $\tilde{Z}^{(2)}$ corresponds to the epimorphism $\rho^{(2)}\colon H_1(D_{N+1}; \mathbb{Z}) \to \mathbb{Z}_2$ given by $\rho^{(2)}(\sum_{i=1}^{N+1} c_i \Delta_i) \equiv \sum_{i=1}^{N+1} c_i \mod 2$ with Δ_i the basis given in Section 4.3. Let Ω be the homology class of a small loop going once counterclockwise around the stirrer P. We then have $\Omega = -\sum_{i=1}^{N+1} \Delta_i$ and so $\rho^{(2)}(\Omega) \equiv 1 \mod 2$ when N is even. As noted in Section 4.1, the foliations of the pseudo-Anosov representative $\phi \in \mathrm{HSP}_N$ have their only singularities at the punctures and outside boundary, and the singularities at the punctures all have odd order. Thus by Proposition 4.4, the orientation double cover of ϕ is defined by the epimorphism, $\bar{\rho}(\Delta_i) \equiv 1 \mod 2$ for $i = 1, \dots, N$ and $\bar{\rho}(\Omega) \equiv 1 \mod 2$. Thus the orientation double cover of HSP_N is exactly $\tilde{Z}^{(2)}$ when N is even.

Again for even N, letting $C(t) = \Phi(HSP_N)$, Proposition 4.2 says that $\rho(\tilde{\phi}_\star^{(2)}) = \rho(C(-1))$, where $\tilde{\phi}^{(2)}$ is the lift of the pseudo-Anosov $\phi \in \mathrm{HSP}_N$ to $\tilde{Z}^{(2)}$. But since $\tilde{Z}^{(2)}$ is also the orientation double cover for ϕ, we have $h(\phi^{(2)}) = \log(\rho(\tilde{\phi}_\star^{(2)}))$. Since entropy is preserved by finite covers, we also have $h(\phi) = \log(\rho(\tilde{\phi}_\star^{(2)}))$. Finally, as remarked in Section 4.6, $\rho(H^{(N)}) = \rho(C(-1))$, finishing the proof when N is even.

To obtain the result when N is odd, we use the two-fold cover $\tilde{Z}'^{(2)}$. It corresponds to the epimorphism $\rho'^{(2)}(\sum_{i=1}^{N+1} c_i \Delta_i) \equiv \sum_{i=1}^{N} c_i \mod 2$. Thus when N is odd, $\rho'^{(2)}(\Omega) \equiv 1 \mod 2$, and once again $\tilde{Z}'^{(2)}$ is the orientation double cover of the pseudo-Anosov $\phi \in \mathrm{HSP}_N$. The rest of the argument yielding (4-10) is the same as when N is even.

Since any $(N \times N)$–matrix C satisfies $\rho(C) \geq \mathrm{trace}(C)/N$, the last two statements of the theorem follow from Lemma 4.3. $\qquad\square$

4.10 The case $N = 2$

The case of two obstacles requires separate consideration. While Theorem 4.5 is true for $N = 2$, the lower bound is simply $0 \leq h(\mathrm{HSP}_2)$ and indeed, it is easy to check that $\rho(\mathrm{HSP}_2) = 1$. Fortunately, there are well-known, simple methods for analyzing the required mapping classes. The Euler–Poincaré formula [12] implies that any pseudo-Anosov on the thrice-punctured disk has invariant foliations whose only singularities are one-prongs at the punctures and a one-prong on the outer boundary. As in the proof of

Theorem 4.5, this implies that for any point-push protocol $\chi \in \mathcal{P}_2$, $h(\chi) = \log(\rho(\Psi(\chi))$ with $\Psi(\chi)$ defined in Section 4.6.

For $j = 1, 2$, let $N_j = \Psi(\alpha_j)$. Now χ is a word in the generators $\alpha_1^{\pm 1}, \alpha_2^{\pm 1}$, and thus with exactly the same indices, $\Psi(\chi)$ is a product of matrices from the set $\mathcal{N} := \{N_1^{\pm 1}, N_2^{\pm 1}\}$. Finding the maximal efficiency in $\mathcal{P}_2$ is then exactly the same as finding the generalized spectral radius of $\mathcal{N}$ (see Remark 2.1). From (4-7),

$$N_1^{\pm 1} = \begin{pmatrix} 1 & 2 \cdot (-1)^{\pm 1} \\ 0 & 1 \end{pmatrix}, \quad N_2^{\pm 1} = \begin{pmatrix} 1 & 0 \\ 2 \cdot (-1)^{\mp 1} & 1 \end{pmatrix}.$$

The following proposition is no doubt well known; we include the simple proof.

Proposition 4.6 If $\mathcal{M} = \{M_1, \ldots, M_n\}$ and the matrix M_1 has the property that $\|M_1\|_2 = \max\{\|M_i\|_2 : 1 \le i \le n\}$ and further, $M_1^T \in \mathcal{M}$, then the generalized spectral radius of the collection $\mathcal{M}$ is $\rho(\mathcal{M}) = \rho_2(\mathcal{M}) = \|M_1\|_2 = (\rho(M_1 M_1^T))^{1/2}$.

Proof By (2-2), $\rho(\mathcal{M}) \le \|M\|_2$, and since $\rho(M_1 M_1^T)^{1/2} = \|M\|_2$, the upper bound is realized by $M_1 M_1^T$. □

Theorem 4.7 The maximal entropy efficiency for point-push protocol with two obstacles is realized by the protocol $\alpha_1 \alpha_2^{-1}$, and so has value

$$\text{Eff}(2) = \frac{\log(3 + 2\sqrt{2})}{2} = \log(1 + \sqrt{2}).$$

Proof The set $\mathcal{N}$ is closed under taking transposes and all its elements have the same two-norm, namely $(3 + 2\sqrt{2})^{1/2} = 1 + \sqrt{2}$, and so the theorem follows from Proposition 4.6. □

The same argument obviously also shows that the maximal efficiency is also realized by $\alpha_1^{-1}\alpha_2$ or $\alpha_2\alpha_1^{-1}$ or $\alpha_2^{-1}\alpha_1$. Also notice that in terms of the usual generators of the braid group on three strings, $\alpha_1\alpha_2^{-1} = \sigma_1^2 \sigma_1 \sigma_2^{-2} \sigma_1^{-1}$. Conjugating by σ_1 we get $\sigma_1^2 \sigma_2^{-2}$, and so the path of the stirrer is a figure eight with an obstacle inside each loop. This protocol is called the "taffy puller" because it is used in making certain kinds of candy (cf[14, Section II]).

5 The upper bound

The main tool for finding an upper bound to the entropy efficiency is the incidence matrix of the action of a point-push protocol on the fundamental group of the punctured

disk. Proposition 5.1 says that the generalized spectral radius of the collection of incidence matrices gives an upper bound for the entropy efficiency. The computation of the appropriate generalized spectral radius is possible here because of the particularly simple form of the action of the generating automorphisms α_k on the generating loops δ_i of the free group $\pi_1(D_{N+1}, Q)$ given in (3-2).

5.1 Incidence matrices and the generalized spectral radius

We start with a general result that says that the generalized spectral radius of the collection of incidence matrices is an upper bound for the efficiency of a collection of automorphisms of a free group.

Assume now that the group $G = F_m$, the free group on the generators $x_1, \ldots, x_m$. Let $\mathbf{n}_i(g)$ be the number of occurrences of the generator x_i or its inverse x_i^{-1} in the reduced form of $g \in G$. We denote the row vector $\mathbf{n}(g) = (\mathbf{n}_1(g), \ldots, \mathbf{n}_m(g))$, and so $\|\mathbf{n}(g)\|_1 = \ell(g)$, the word length of g. The *incidence matrix* of the automorphism $a \colon F_m \to F_m$ is the $(m \times m)$–matrix A with $A_{ij} = \mathbf{n}_j(a(x_i))$.

Proposition 5.1 *If $\{a_1, \ldots, a_k\}$ is a collection of generators of the semigroup $\mathcal{A} \subset$ Aut(F_n) whose incidence matrices are $\{A_1, \ldots, A_k\}$, then the entropy efficiency of $\mathcal{A}$ is bounded above by the generalized spectral radius of the collection of incidence matrices or*

$$\mathrm{Eff}(\mathcal{A}; \{a_1, \ldots, a_n\}) \leq \rho(\{A_1, \ldots, A_n\}).$$

Proof Let $a \in \mathcal{A}$ have incidence matrix A. By construction, $\mathbf{n}(a(g)) \leq \mathbf{n}(g)A$, and iterating we have for $n > 0$,

$$\mathbf{n}(a^n(g)) \leq \mathbf{n}(g)A^n.$$

Taking one-norms, n–th roots, and then using Gelfand's formula

$$\limsup_{n \to \infty} \|\mathbf{n}(a^n(g))\|_1^{1/n} \leq \limsup_{n \to \infty} \|\mathbf{n}(g)\|_1^{1/n} \|A^n\|_1^{1/n} = \rho(A).$$

Since this is true for all $g \in \mathbb{F}_m$, using the definition in (2-4) and taking the logarithm we have $h(a) \leq \log(\rho(A))$. Now treat a as a product of the generators in reduced form, $a = a_{i_1} \cdots a_{i_k}$. Since incidence matrices ignore cancellations, $A \geq A_{i_1} \cdots A_{i_k}$, and thus because the matrices are positive, $\rho(A) \geq \rho(A_{i_1} \cdots A_{i_k})$. Therefore for all a of word length k,

$$\frac{h(a)}{k} \leq \log(\rho(A_{i_1} \cdots A_{i_k})^{1/k}),$$

and the result follows. $\qquad\square$

5.2 The incidence matrices and the form of the upper bound

The action of α_k on the generating loops δ_i of the free group $\pi_1(D_{N+1}, Q)$ as given in (3-2) has $((N+1)\times(N+1))$–dimensional incidence matrix $\bar{A}^{(k)} = I + S^{(k)}$ with $S^{(k)}$ defined by $S_{kk}^{(k)} = 0$, $S_{kj}^{(k)} = 2$ for $j \neq k$, and $S_{ij}^{(k)} = 0$ otherwise. As in Section 4.6, for all k, the last row of $\bar{A}^{(k)}$ is $(0\,0\cdots 0,\,,1)$. Thus if we define the *reduced incidence matrix* $A^{(k)}$ as the matrix obtained by deleting the last row and column of $\bar{A}^{(k)}$, then

$$(5\text{-}1) \qquad \rho(\{A^{(1)}, \dots, A^{(N)}\}) = \rho(\{\bar{A}^{(1)}, \dots, \bar{A}^{(N)}\}).$$

Thus we may restrict attention to the $A^{(k)}$. We still have $A^{(k)} = I + S^{(k)}$, but now $S^{(k)}$ is $N \times N$. We could have alternatively defined $A^{(k)}$ as the absolute value of the matrix $\Psi(\alpha_k) = E^{(k)}$ giving the action of α_k on lifted chains to the two-fold cover from Section 4.6.

It is easy to check that $\bar{A}^{(k)}$ is also the incidence matrix for the functional inverse, α_k^{-1} of α_k and further that $\|A^{(k)}\|_1 = 3$ for all k. Thus by Proposition 5.1, (2-2) and (5-1) we have:

Proposition 5.2 *If $\{A^{(1)}, \dots, A^{(N)}\}$ is the set of reduced incidence matrices of the standard set of generators $\{\alpha_1^{\pm 1}, \dots, \alpha_N^{\pm 1}\}$ of the group of point-push protocols $\mathcal{P}_N$, then the maximal entropy efficiency in the presence of N obstacles satisfies*

$$\mathrm{Eff}(N) \leq \rho(\{A^{(1)}, \dots, A^{(N)}\}) \leq \log(3),$$

where $\rho(\{A^{(1)}, \dots, A^{(N)}\})$ is the generalized spectral radius of the set of incidence matrices.

5.3 Computing the generalized spectral radius

We can improve the upper bound of $\log(3)$ in Proposition 5.2 by computing the generalized spectral radius $\rho(\{A^{(1)}, \dots, A^{(N)}\})$, at least in the sense of showing that it is realized by an explicit matrix.

Proposition 5.3 *If $\hat{H}^{(N)} := A^{(1)} \cdots A^{(N)}$, then $\rho(\hat{H}^{(N)})^{1/N} = \rho(\{A^{(1)}, \dots, A^{(N)}\})$.*

We shall see in Remark 5.7 that the product realizing the generalized spectral radius in Proposition 5.3 is not unique.

The computation of the generalized spectral radius makes frequent use of the fact that multiplication by an incidence matrix $A^{(k)}$ changes the column sum in a fairly simple fashion. Recall that $c(M)$ denotes the column sums of M. We then have

$$(5\text{-}2) \quad c_j(M A^{(k)}) = c_j(M) + 2c_k(M) \text{ for } j \neq k \quad \text{and} \quad c_k(M A^{(k)}) = c_k(M).$$

Now the maximal column sum of a positive matrix is its one-norm, $\|M\|_1 = \max c(M)$. Thus if we can determine a sequence of multiplications by incidence matrices which maximizes the column sums, then we know the sequence which maximizes the one-norm. An argument using Gelfand's formula then completes the proof.

Finding the proper sequence is not difficult, again because of the simple action of the $A^{(k)}$ on column sums. An example will make it clear.

Example 5.4 We begin with the identity matrix I and $c(I) = (1, 1, \ldots, 1)$. Since all the entries are equal, it doesn't matter which $A^{(k)}$ we choose to multiply by, and so for simplicity we choose $A^{(1)}$, and we get $c(IA^{(1)}) = (1, 3, 3, 3, \ldots, 3)$. Now to maximize the column sums in the next product, (5-2) indicates we should choose an index k for which the corresponding entry in $c(IA^{(1)})$ is maximal. Again, we may choose any $k > 1$, but for simplicity we choose $A^{(2)}$ yielding $c(IA^{(1)}A^{(2)}) = (4, 3, 6, 6, \ldots, 6)$. Continuing we see that at the j–th step, the j–th entry of the current column sum will be among the maximal values and so we choose to multiply by $A^{(j)}$. It is more or less clear then that after N steps we will get a maximal column sum by using the sequence prescribed in $\hat{H}^{(N)} := A^{(1)} \cdots A^{(N)}$

We need a few definitions and preliminary lemmas to formalize the insight of the example. Fix $N \geq 3$ and let $\mathcal{O}_N$ be all N–tuples of positive integers given in nonincreasing order, so

$$\mathcal{O}_N = \{\vec{a} \in (\mathbb{Z}^+)^n : a_1 \geq a_2 \geq \cdots \geq a_N > 0\}.$$

The *sorting function* is the map $S \colon (\mathbb{Z}^+)^n \to \mathcal{O}_n$ which sends a row vector of positive integers to the row vector of its entries sorted in nonincreasing order. In keeping with our conventions we will have S act from the right, and so $\vec{b} S$ is the sorted version of $\vec{b} \in (\mathbb{Z}^+)^n$.

It is easy to check that for each j and each $\vec{a} \in (\mathbb{Z}^+)^N$, $\vec{a} A_j \in (\mathbb{Z}^+)^N$, and so we may define $W_j \colon \mathcal{O}_N \to \mathcal{O}_N$ by $\vec{a} W_j := \vec{a} A_j S$. Using the fact that all entries of $\vec{a}$ are strictly positive we get the formulas

$$
\begin{aligned}
\text{(5-3)} \quad &\vec{a} W_1 = (a_2 + 2a_1, \ldots, a_n + 2a_1, a_1), \\
&\vec{a} W_j = (a_1 + 2a_j, \ldots, a_{j-1} + 2a_j, a_{j+1} + 2a_j, a_n + 2a_j, a_j) \quad \text{when } j > 1.
\end{aligned}
$$

The next lemma collects some useful facts about the action of the W_j and the usual partial order on $(\mathbb{Z}^+)^N$. Its proof is an elementary exercise using (5-3).

Lemma 5.5 *Assume $\vec{a}, \vec{b} \in (\mathbb{Z}^+)^N$.*

(a) *If $\vec{a} \geq \vec{b}$, then $\vec{a}\, W_j \geq \vec{b}\, W_j$ for all j.*

(b) *If $\vec{a} > \vec{b}$, then $\vec{a}\, W_j > \vec{b}\, W_j$ for all j.*

(c) *If $j < k$ and $a_j = a_k$, then $\vec{a}\, W_j = \vec{b}\, W_k$.*

(d) *If $j < k$ and $a_j > a_k$, then $\vec{a}\, W_j > \vec{b}\, W_k$.*

(e) *If $j \leq k$ and $\vec{a} > \vec{b}$, then $\vec{a}\, W_j > \vec{b}\, W_k$.*

The next step is to formalize the various sequences of multiplying the matrices in our family. A *strategy* is a list $\underline{s} = (s_1, s_2, \ldots, s_k)$ with $1 \leq s_i \leq N$ for all i. Each strategy $\underline{s}$ defines a map $W_{\underline{s}}: \mathcal{O}_N \to \mathcal{O}_N$ given by $\vec{a}\, W_{\underline{s}} := \vec{a}\, W_{s_1} W_{s_2} \cdots W_{s_k}$.

The next definition gives a canonical form for a strategy which will facilitate comparing the results of two strategies. Given $\vec{a} \in \mathcal{O}_N$ and a strategy $\underline{s}$, the *standardization* of $\underline{s}$ with respect to $\vec{a}$ is the strategy $\underline{s}' = (s_1', s_2', \ldots, s_k')$ defined inductively as follows. Let

$$s_1' = \min\{i : a_i = a_{s_1}\},$$
$$s_m' = \min\{i : b_i^{(m)} = b_{s_m}^{(m)}\} \quad \text{for } m > 1,$$

where
$$b^{(m)} = \vec{a}\, W_{s_1'} \cdots W_{s_{m-1}'}.$$

Note that it follows easily that for all m, $b^{(m)} = \vec{a}\, W_{s_1} \cdots W_{s_{m-1}}$, and so in particular,

$$(5\text{-}4) \qquad\qquad\qquad\qquad \vec{a}\, W_{\underline{s}} = \vec{a}\, W_{\underline{s}'}.$$

If the strategy $\underline{s}$ is equal to its standardization with respect to $\vec{a}$, then $\underline{s}$ is said to be *standard* for $\vec{a}$. In light of (5-4) we henceforth restrict consideration to standard strategies.

The simplest strategy is to repeatedly use the first place, which means using the largest entry; for each k, let $\underline{1}^{(k)}$ be the k–tuple of all ones, $\underline{1}^{(k)} = (1, 1, \ldots, 1)$. The next proposition collects some useful facts about strategies and products of the matrices from $\mathcal{M}$. The first part says that $\underline{1}^{(k)}$ is always the best strategy for maximizing the end result. We adopt the simplified notation that given an k–tuple $\underline{\ell} = (\ell_1 \cdots \ell_k)$, $A_{\underline{\ell}} = A^{(\ell_1)} \cdots A^{(\ell_k)}$.

Lemma 5.6 *Assume $\vec{a} \in \mathcal{O}_N$.*

(a) *If $\underline{s} = (s_1, \ldots, s_k)$ is standard strategy for $\vec{a}$ and $\underline{s} \neq \underline{1}^{(k)}$, then $\vec{a}\, W_{\underline{1}^{(k)}} > \vec{a}\, W_{\underline{s}}$.*

(b) *If $\underline{s} = (s_1, \ldots, s_k)$ is any strategy, then $\vec{a}\, W_{\underline{1}^{(k)}} \geq \vec{a}\, W_{\underline{s}}$.*

(c) *Given any k–tuple $\underline{\ell} = (\ell_1, \ldots, \ell_k)$, there exits a strategy $\underline{s} = (s_1, \ldots, s_k)$ with $\vec{a}\, A_{\underline{\ell}}\, S = \vec{a}\, W_{\underline{s}}$.*

(d) *If $\underline{\ell}$ is the k–tuple $\underline{\ell} = (1, 2, \ldots, k)$ and $\vec{a} = (1, 1, \ldots, 1) \in (\mathbb{Z}^+)^N$, then $\vec{a}\, A_{\underline{\ell}}\, S = \vec{a}\, W_{\underline{1}^{(k)}}$.*

Proof Since by hypothesis, $\underline{s} \neq \underline{1}^{(k)}$, there is a smallest index i with $s_i > 1$. If $i = 1$, let $\vec{b} = \vec{a}$ and if $i > 1$, let $\vec{b} = \vec{a}\, W_{(s_1,\dots,s_{i-1})} = W_{\underline{1}^{(i-1)}}(\vec{a})$. Since $\underline{s}$ is standard for $\vec{a}$, when $i = 1$, $a_{s_i} < a_1$ and so $b_{s_i} < b_1$, and when $i > 1$, again since $\underline{s}$ is standard for $\vec{a}$, $b_{s_i} < b_{s_i-1} \leq b_1$. Thus in either case by Lemma 5.5(d), $\vec{b}\, W_1 > \vec{b}\, W_{s_i}$ which is to say that $\vec{a}\, W_{\underline{1}^{(i)}} > \vec{a}\, W_{s_1,\dots,s_i}$. Since $s_n \geq 1$ for all $n > i$, by Lemma 5.5(e) $\vec{a}\, W_{\underline{1}^{(k)}} > \vec{a}\, W_{\underline{s}}$, as required for (a).

For (b), if $\underline{1}^{(k)}$ is the standardization of $\underline{s}$, then $\vec{a}\, W_{\underline{1}^{(k)}} = \vec{a}\, W_{\underline{s}}$, otherwise use (a). We prove (c) by induction on k, with $k = 1$ being trivial. The inductive step is based on the easily checked observation that for any (perhaps not ordered vector) $\vec{d} \in (\mathbb{Z}^+)^N$, and for any j, if we let j' be the index with $(\vec{d}\,S)_{j'} = d_j$, then $\vec{d}\,A_j\,S = \vec{d}\,SA_{j'}\,S$.

Now assume that the result is true for a k–tuple $\underline{\ell}$. Given $\underline{\ell}^+ = (\ell_1, \dots, \ell_k, \ell_{k+1})$, using the observation of the previous paragraph with $\vec{d} = \vec{a}\,A_{\underline{\ell}}$, there is a j' with $\vec{a}\,A_{\underline{\ell}}\,A_{\ell_{k+1}}\,S = \vec{a}\,A_{\underline{\ell}}\,S\,A_{j'}\,S = \vec{a}\,W_{\underline{s}}\,A_{j'}\,S$, using the inductive hypothesis. Thus $\vec{a}\,A_{\underline{\ell}^+}\,S = \vec{a}\,W_{\underline{s}^+}$ with $\underline{s}^+ = (s_1, \dots, s_k, j')$, completing the induction.

The proof of (d) is a simple calculation. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

Proof of Proposition 5.3 Given a $k > 0$, write $k = mN + j$ with $0 \leq j < N$ and let $B^{(k)} = (\hat{H}^{(N)})^m A^{(1)} \cdots A^{(j)}$. As a first step we show that for any product of k matrices $P := A^{(j_1)} \cdots A^{(j_k)}$, $\|B^{(k)}\|_1 \geq \|P\|_1$.

Let $\vec{a} = (1, 1, \dots, 1) \in (\mathbb{Z}^+)^N$. By Lemma 5.6(c), there is a strategy $\underline{s}$ which we may assume to be standard so that $\vec{a}\,P\,S = \vec{a}\,W_{\underline{s}}$. Further, by Lemma 5.6(d), $\vec{a}\,B^{(k)}\,S$ is produced by the strategy $\underline{1}^{(k)}$, so $\vec{a}\,B^{(k)}\,S = \vec{a}\,W_{\underline{1}^{(k)}}$. Thus by Lemma 5.6(b), $\vec{a}\,B^{(k)}\,S \geq \vec{a}\,P\,S$. But for any positive matrix M, $\max(\vec{a}\,M) = \|M\|_1$ and so $\|B^{(k)}\|_1 \geq \|P\|_1$, finishing the proof of the first step.

Since $B^{(k)} = (\hat{H}^{(N)})^m A^{(1)} \cdots A^{(j)}$, the first step implies that $\|P\|_1 \leq c\|(\hat{H}^{(N)})^m\|_1$, where

$$c = \max\{\|A^{(1)} \cdots A^{(j)}\|_1 : 1 \leq j < N\}.$$

But for any matrix M, $\rho(M) \leq \|M\|_1$, and so for each $k = mN + j$ letting $\mathcal{M} = \{A^{(1)}, \dots, A^{(N)}\}$ we have

$$\rho_k(\mathcal{M}) \leq \|(\hat{H}^{(N)})^m\|_1^{1/k} c^{1/k} \leq \|(\hat{H}^{(N)})^m\|_1^{1/(mN)} c^{1/k},$$

and so using Gelfand's formula, $\rho(\mathcal{M}) \leq \rho(\hat{H}^{(N)})^{1/N}$. On the other hand, for any $m > 0$ $\rho((\hat{H}^{(N)})^m) = \rho(\hat{H}^{(N)})^m$, and so $\rho_{mN}(\mathcal{M}) \geq \rho(\hat{H}^{(N)})^{1/N}$ and thus $\rho(\mathcal{M}) \geq \rho(\hat{H}^{(N)})^{1/N}$, finishing the proof. $\qquad\qquad\qquad\square$

Remark 5.7 The product of matrices in Proposition 5.3 is by no means unique. If $\underline{\ell}$ is any N–tuple which is a cyclic permutation of $(1, 2, \ldots, N)$ and $\vec{a} = (1, 1, \ldots, 1) \in (\mathbb{Z}^+)^N$, then it is easy to check that $\vec{a}\, A_{\underline{\ell}}\, S = \vec{a}\, W_{1^{(k)}}$. Thus the joint spectral radius of the collection $\{A^{(1)}, \ldots, A^{(N)}\}$ it is also realized by any product of N of the $A^{(k)}$ in which each $A^{(k)}$ occurs exactly once.

5.4 Properties of $\widehat{H}^{(N)}$ and the upper bound

Recall that $c_j(M)$ denotes the sum of the entries in the j–th column of the matrix M.

Proposition 5.8 The product of incidence matrices $\widehat{H}^{(N)} := A^{(1)} \cdots A^{(N)}$ satisfies $c_j(\widehat{H}^{(N)}) = 3^N - 2 \cdot 3^{j-1}$ and $\operatorname{trace}(\widehat{H}^{(N)}) = 3^N - N - 1$, and so

$$\frac{3^N - N - 1}{N} \le \rho(\widehat{H}^{(N)}) \le 3^N - 2.$$

Proof Fix N and for $1 \le k \le N$, let $\widehat{H}^{(k)} = A^{(1)} A^{(2)} \cdots A^{(k)}$. An easy induction yields that $c_j(\widehat{H}^{(k)}) = 3^k - 2 \cdot 3^{j-1}$ for $j \le k$ and $c_j(\widehat{H}^{(k)}) = 3^k$ for $j > k$. The column sum formula then follows by letting $k = N$.

For any matrix M, it follows from the definition of $A^{(j)}$ that $\operatorname{trace}(MA^{(j)}) = \operatorname{trace}(M) + 2c_j(M) - 2M_{jj}$. Since for any k, it is easy to check that $\widehat{H}^{(k-1)}_{k,k} = 1$, and we have just seen that $c_k(\widehat{H}^{(k-1)}) = 3^{k-1}$, so another simple induction gives $\operatorname{trace}(\widehat{H}^{(k)}) = N + 3^k - 2k - 1$, and we again let $k = N$.

The estimates on the spectral radius then follow from the general fact that for any $(N \times N)$–matrix M, $|\operatorname{trace}(M)|/N \le \rho(M) \le \|M\|_1$. $\square$

As a consequence of Propositions 5.2, 5.3 and 5.8 we have our upper bound for the entropy efficiency.

Theorem 5.9 The maximal efficiency $\operatorname{Eff}(N)$ of a point-push protocol in the presence of N obstacles satisfies

$$\operatorname{Eff}(N) \le \frac{\log(\rho(\widehat{H}^{(N)}))}{N} \le \frac{\log(3^N - 2)}{N}.$$

References

[1] **G Band**, **P Boyland**, *The Burau estimate for the entropy of a braid*, Algebr. Geom. Topol. 7 (2007) 1345–1378 MR2350285

[2]   **M Bestvina**, **M Handel**, *Train-tracks for surface homeomorphisms*, Topology 34 (1995) 109–140   MR1308491

[3]   **B J Binder**, *Ghost rods adopting the role of withdrawn baffles in batch mixer designs*, Phys. Letters A 374 (2010) 3483–3486

[4]   **J S Birman**, *Braids, links, and mapping class groups*, Annals of Math. Studies 82, Princeton Univ. Press (1974)   MR0375281

[5]   **P Boyland**, *Isotopy stability of dynamics on surfaces*, from: "Geometry and topology in dynamics (Winston-Salem, NC, 1998/San Antonio, TX, 1999)", (M Barge, K Kuperberg, editors), Contemp. Math. 246, Amer. Math. Soc. (1999) 17–45   MR1732369

[6]   **P Boyland**, *Dynamics of two-dimensional time-periodic Euler fluid flows*, Topology Appl. 152 (2005) 87–106   MR2160808

[7]   **P Boyland**, **H Aref**, **M A Stremler**, *Topological fluid mechanics of stirring*, J. Fluid Mech. 403 (2000) 277–304   MR1742169

[8]   **A J Casson**, **S A Bleiler**, *Automorphisms of surfaces after Nielsen and Thurston*, London Math. Soc. Student Texts 9, Cambridge Univ. Press (1988)   MR964685

[9]   **S Dowdall**, *Dilatation versus self-intersection number for point-pushing pseudo-Anosov homeomorphisms*   `arXiv:1004.3936`

[10]   **B Farb**, **C Leininger**, **D Margalit**, *Small dilatation pseudo-anosovs and* 3*–manifolds* `arXiv:0905.0219v1`

[11]   **B Farb**, **D Margalit**, *A primer on mapping class groups*, Princeton Math. Ser. 49, Princeton Univ. Press (2011)

[12]   **A Fathi**, **F Laudenbach**, **V Poénaru**, **editors**, *Travaux de Thurston sur les surfaces*, Astérisque 66–67, Soc. Math. France, Paris (1979)   MR568308   Séminaire Orsay, In French with an English summary

[13]   **M D Finn**, **S M Cox**, **H M Byrne**, *Mixing measures for a two-dimensional chaotic Stokes flow*, J. Engrg. Math. 48 (2004) 129–155   MR2033473

[14]   **M D Finn**, **J-L Thiffeault**, *Topological optimisation of rod-stirring devices*, to appear in SIAM Review   `arXiv:1004.0639v1`

[15]   **D Fried**, *Periodic points and twisted coefficients*, from: "Geometric dynamics (Rio de Janeiro, 1981)", (J Pallis, editor), Lecture Notes in Math. 1007, Springer, Berlin (1983) 261–293   MR730272

[16]   **D Fried**, *Entropy and twisted cohomology*, Topology 25 (1986) 455–470   MR862432

[17]   **E Ghate**, **E Hironaka**, *The arithmetic and geometry of Salem numbers*, Bull. Amer. Math. Soc. (N.S.) 38 (2001) 293–314   MR1824892

[18]   **T Hall**, *Software to compute train tracks of surface homeomorphisms* (2005) Available at `http://www.liv.ac.uk/maths/PURE/MIN_SET/CONTENT/members/T_Hall.html`

[19] **M Handel**, *Global shadowing of pseudo-Anosov homeomorphisms*, Ergodic Theory Dynam. Systems 5 (1985) 373–377   MR805836

[20] **E Hironaka**, **E Kin**, *A family of pseudo-Anosov braids with small dilatation*, Algebr. Geom. Topol. 6 (2006) 699–738   MR2240913

[21] **R Jungers**, *The joint spectral radius: Theory and applications*, Lecture Notes in Control and Information Sci. 385, Springer, Berlin (2009)   MR2507938

[22] **T Kobayashi**, **S Umeda**, *A design for pseudo-Anosov braids using hypotrochoid curves*, Topology Appl. 157 (2010) 280–289   MR2556106

[23] **I Kra**, *On the Nielsen–Thurston–Bers type of some self-maps of Riemann surfaces*, Acta Math. 146 (1981) 231–270   MR611385

[24] **E Lanneau**, **J-L Thiffeault**, *On the minimum dilatation of pseudo-anosov homeomorphisms on surfaces of small genus* (2011) Available at `http://aif.cedram.org/item?id=AIF_2011__61_1_105_0`

[25] **A Manning**, *Topological entropy and the first homology group*, from: "Dynamical systems—Warwick 1974 (Proc. Sympos. Appl. Topology and Dynamical Systems, Univ. Warwick, Coventry, 1973/1974; presented to E C Zeeman on his fiftieth birthday)", (A Manning, editor), Lecture Notes in Math. 468, Springer, Berlin (1975) 185–190   MR0650661

[26] **J-O Moussafir**, *On computing the entropy of braids*, Funct. Anal. Other Math. 1 (2006) 37–46   MR2381961

[27] **R C Penner**, **J L Harer**, *Combinatorics of train tracks*, Annals of Math. Studies 125, Princeton Univ. Press (1992)   MR1144770

[28] **J-L Thiffeault**, **M D Finn**, *Topology, braids and mixing in fluids*, Philos. Trans. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci. 364 (2006) 3251–3266   MR2317905

[29] **W P Thurston**, *On the geometry and dynamics of diffeomorphisms of surfaces*, Bull. Amer. Math. Soc. (N.S.) 19 (1988) 417–431   MR956596

Department of Mathematics, University of Florida
Gainesville FL 32611-8105, USA

`boyland@ufl.edu, mathguy@ufl.edu`

`http://www.math.ufl.edu/~boyland/, http://www.math.ufl.edu/~mathguy/`