Desargues Theorem, Dynamics, and Hyperplane Arrangments

Richard Evan Schwartz *

June 18, 2001

1 Introduction

Desargues' theorem is a familiar and basic theorem in projective geometry. It concerns the configuration in Figure 1.1, and says that the points x_1, x_2, and x_3 are collinear, no matter how the other points and lines are moved around, subject to the incidence relations implied by the diagram.

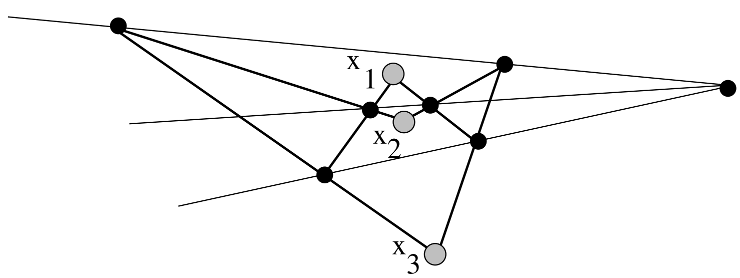

Figure 1.1

In this paper we will generalize Desargues' theorem in the direction of dynamical systems. The generalization comprises an infinite family of planar configurations. There is a sense in which the configurations, and the statements made about them, are "unboundedly intricate". Before stating the main result, we will explain some special cases, which may give the reader a feel for the situation in general.

* Supported by a Sloan Research Fellowship

The starting point for our result is a different point of view on Desargues' theorem. Let H and H' be the hexagons shown in Figure 1.2. H' is obtained from H by drawing in the "diagonals" of H and then intersecting them.

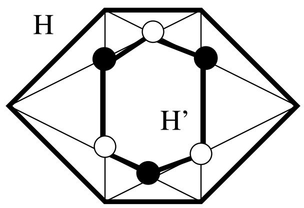

Figure 1.2

This construction is easiest to draw when H is convex. But now imagine deforming H until the sides of H are alternately parallel to the x and y axes, as shown in the Figure 1.3. It turns out, in this case, that the points of H' lie on a pair of lines, provided that these points are defined. (Generically they are defined.) This fact follows from two applications of Desargues' theorem.

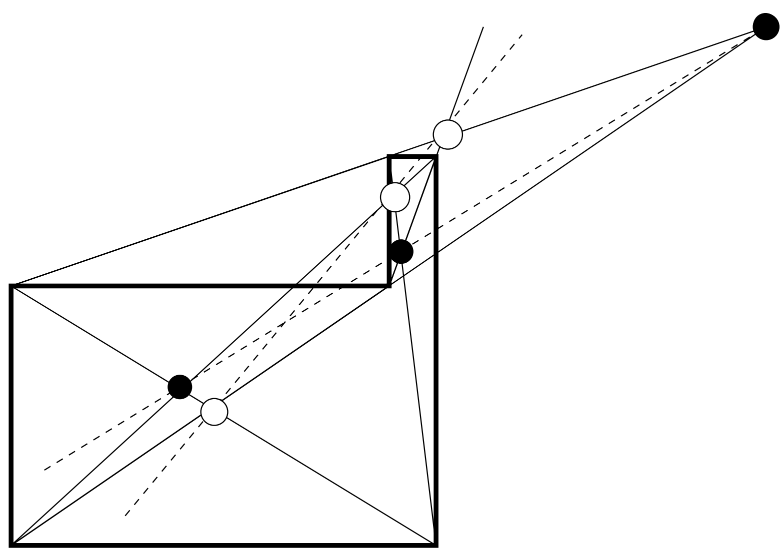

Figure 1.3

Now begin with an octagon O, and apply the above construction twice, to obtain O''. If we deform O until its sides are alternately parallel to the x and y axes, then the eight points of O'' lie on a pair of lines. This is shown, somewhat schematically in Figure 1.3. This phenomenon continues. If D is a decagon whose sides are alternately parallel to the x and y axes, then the 10 points of D''' lie on a pair of lines. And so on. (The same comments on definedness apply here as above.) The suggested infinite sequence of statements is in turn a special case of Theorem 1.1.

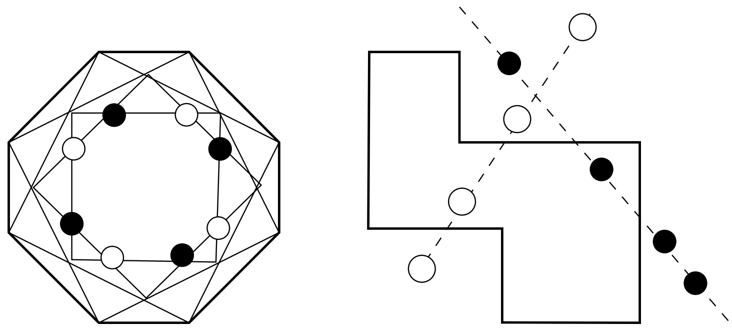

Figure 1.3

In order to present Theorem 1.1 in a natural and comprehensive manner, we need some unusual terminology. Say that a *PolyPoint* is a cyclically ordered, finite collection of points in the projective plane. (For convenience, we will work in the complex projective plane.) We say that an *n-Point* is a PolyPoint consisting of n points. Dually, we say that a *PolyLine* is a cyclically ordered finite collection of lines in the projective plane. We say that an *n-Line* is a PolyLine consisting of n lines. We will use the more familiar term, *polygon*, to refer to an object which could either be a PolyPoint or a PolyLine.

Let P be an n-Point and suppose k is not a multiple of n. A k-diagonal of P is a line determined by the jth point of P and the $(j+k)$th point of P, for some j. Indices are taken mod n. The k-diagonals of P inherit a cyclic order from P. Taken all together they comprise an n-Line L. Dually, let L be an n-Line. A *k-diagonal* of L is a point obtained by intersecting the jth and $(j+k)$th lines of L. The k-diagonals of L inherit a cyclic ordering, and thus comprise an n-Point.

We will denote the two maps described above by δ_k. When X is an n-Point, $\delta_k(X)$ is an n-Line, and vice versa. In fact, δ_k^2 is the identity, both on PolyPoints and on PolyLines. Figure 1.4 shows the construction just described, for $n = 7$ and $k = 3$.

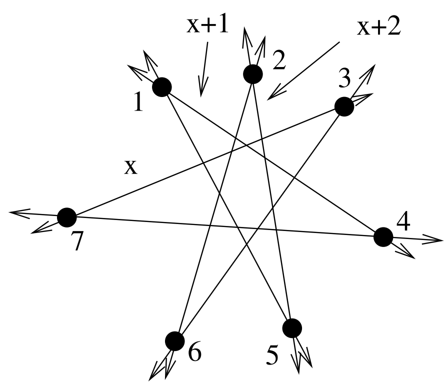

Figure 1.4

To get some nontrivial dynamics, we consider the group generated by δ_p and δ_q, for distinct p and q. That is, we take a polygon P, and form the sequence

$$...\delta_q\delta_p(P) \leftrightarrow \delta_p(P) \leftrightarrow P \leftrightarrow \delta_q(P) \leftrightarrow \delta_p\delta_q(P)...$$

Typically, this sequence is bi-infinite. However, the sequence could terminate, at either end, due to a singularity of some kind which prevents the definition of the relevant map. Our result involves cases where the sequence terminates, at one end, in a highly symmetric singularity.

Say that an np-Point C is an *n-cover* if the jth and $(j+p)$th points of C coincide, for all j. Our terminology is chosen so that C is an n-fold covering of a p-Point. (It is useful to think of n large and p small.) We say that an np-Line S is an (n,p)-*satellite* if $C = \delta_p(S)$ is an n-cover. If $S = \{s_1, s_2, ..., s_{np}\}$, then the sets $\{s_j, s_{j+p}, s_{j+2p}, ...\}$ consist of coincident lines. The same definitions may be made with the words *Point* and *Line* interchanged, using projective duality.

Figure 1.5 shows examples for $(n,p) = (4,3)$. The examples H, O, and D, considered above, are respectively $(3,2)$-, $(4,2)$-, and $(5,2)$-satellites. To give a general example, define $P_{n,p} = \{p_1, p_2, ..., p_{np}\}$, where

$$p_j = \begin{bmatrix} \cos(\frac{2\pi}{p}) & \sin(\frac{2\pi}{p}) \\ -\sin(\frac{2\pi}{p}) & \cos(\frac{2\pi}{p}) \end{bmatrix} \begin{bmatrix} \exp(\frac{2\pi i}{n}) \\ 0 \end{bmatrix}.$$

When n and p are relatively prime, $P_{n,p}$ is an (n,p)-satellite.

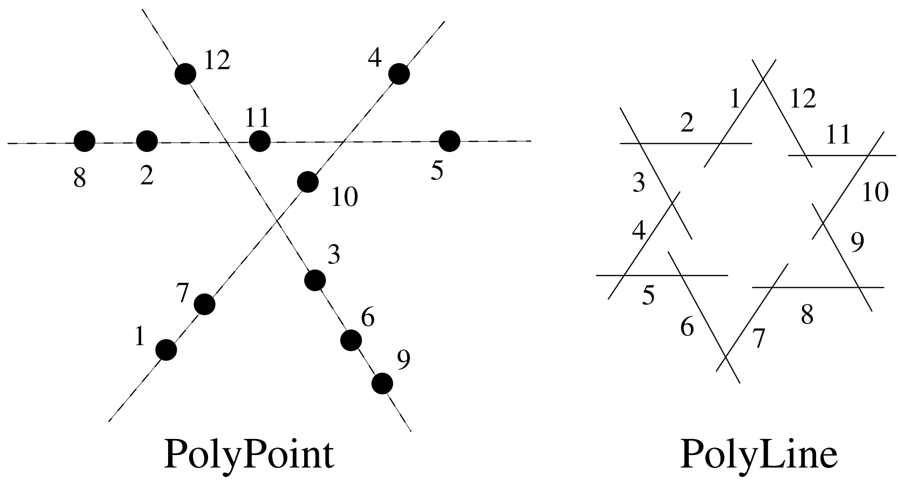

Figure 1.5

Now, suppose that S is an (n,p)-satellite and $C = \delta_p(S)$ is an n-cover, just as above. A sequence of the form

$$(*) \quad C = \delta_p(S) \leftrightarrow S \leftrightarrow \delta_q(S) \leftrightarrow \delta_p\delta_q(S) \leftrightarrow \delta_q\delta_p\delta_q(S)...$$

cannot be continued much further to the left, because δ_p is not defined on $\delta_q(C)$. On the other hand, we can ask what happens as the maps proceed towards the right.

Theorem 1.1 below says (in the cases covered) that generically $(*)$ continues to the right as

$$(**) \quad C = \delta_p(S) \leftrightarrow S \leftrightarrow \delta_q(S) \leftrightarrow ... \leftrightarrow \delta_q(S^*) \leftrightarrow S^* \leftrightarrow \delta_p(S^*) = C^*.$$

Both C and C^* are n-covers of p-gons, and the total number of arrows is exactly $2n-1$. If C is a PolyPoint, then C^* is a PolyLine, and vice versa. One is tempted to call C^* the *inverse* of C.

Let $P(n,p)$ and $L(n,p)$ denote the space of (n,p)-satellites which respectively are PolyPoints and PolyLines. By *generic set* we mean, roughly, that the complement of the set is a lower dimensional subvariety. See §2.4 for a precise definition.

Theorem 1.1 *Suppose that n,p,q are positive integers, pairwise relatively prime, such that $\min(n,p) \geq 3$. Then there are generic subsets $GP \subset P(n,p)$ and $GL \subset L(n,p)$, dual to each other, such that $(\delta_q \circ \delta_p)^{n-2} \circ \delta_q$ is defined on, and interchanges, GP and GL.*

As we will see in §2, the case $(3,p,q)$ is equivalent to Desargues' theorem, regardless of the values p and q. Theorem 1.1 grows more global and dynamical as $n \to \infty$. The most dramatic cases are when n is large in comparison to p and q. The drama is that the singularity at one end of $(**)$ reappears at other end, after (say) trillions of iterations of the basic construction.

For the sake of exposition, we have not tried to present the most general version of Theorem 1.1. It is worth mentioning, however, that Theorem 1.1 holds for $(n,p,q) = (n,2,1)$, for $n = 3,4,5....$ These cases correspond to the specific examples described at the beginning of this introduction, and in certain coordinates are quite closely related to C.L. Dodgson's celebrated *method of condensation* for computing determinants.

Aside from the connection to determinants, we are not sure how Theorem 1.1 fits into the general scheme of things. We discovered the result by computer experimentation. The main motivation was to understand the singularities of the kind of birational mapping described above. (We studied a case of this mapping in [**S**].) The motivating idea, however, does not find its way into this paper.

Theorem 1.1 is a direct consequence of Lemma 2.1, which we prove in §2-3 and Lemma 4.1, which we prove in §4-5. Lemma 2.1 says that the sequence $(*)$ generically has the structure of $(**)$, where the total number of arrows is $\xi + 2$, for some $\xi \leq 2n-3$. Just as the classical proof of Desargues theorem involves a certain arrangement of planes in three space, our proof of Lemma 2.1 involves certain hyperplane arrangements in n-space.

Lemma 4.1 says that that $\xi(n,p,q) = 2n-3$, when (n,p,q) satisfy the hypotheses of Theorem 1.1. The proof of Lemma 4.1 is quite delicate (and intricate) in places. Some of this subtlety is to be expected, because there are choices of (n,p,q), such as $(6,2,3)$, for which $\xi(n,p,q) < 2n-3$. Lemma 4.1 has a trivial computational proof for each fixed choice of (n,p,q). Let $\rho = (\delta_q \circ \delta_p)^{n-2} \circ \delta_q$. One simply has to exhibit a single element $P \in P(n,p)$ such that ρ is defined on P and on $\rho(P)$. Call P a *good example*. A good example is incompatible with $\xi < 2n-3$.

To the reader who wishes to avoid the technicalities of §4-5 we say: In any given case, Theorem 1.1 follows from Lemma 2.1 and from a single computer experiment, which exhibits a good example. Indeed, ample experimental evidence suggests that the element $P_{n,p}$, described above, is always a good example.

I would like to thank Peter Doyle, John Millson, and Madhav Nori, for helpful, interesting and enthusiastic conversations about this work.

2 Upper Bound

2.1 Overview

Suppose that (n,p,q) is chosen. Recall that $P(n,p)$ and $L(n,p)$ are spaces of (n,p)-satellites. Let ϕ_0 be the identity map. For $i \geq 1$ odd, define $\phi_i = \delta_q \circ \phi_{i-1}$. For $j \geq 2$ even, define $\phi_j = \delta_p \circ \phi_{j-1}$. The goal of §2-3 is to prove:

Lemma 2.1 (Upper Bound) *Suppose that (n,p,q) satisfy the hypotheses of Theorem 1.1. There is an odd positive integer $\xi \leq 2n-3$ and generic subsets $GP \subset P(n,p)$ and $GL \subset L(n,p)$ such that ϕ_ξ is defined on, and interchanges, GP and GL.*

In this chapter, we will deduce Lemma 2.1 from:

Lemma 2.2 *Suppose that $(n,p,q) \geq (3,2,1)$ componentwise, and p,q are relatively prime. Suppose that $L \in L(n,p)$ and $L^* = \phi_{2n-3}(L)$ is defined. Then $L^* \in P(n,p)$.*

Lemma 2.2 is, in turn, a consequence of the completely geometrical Theorem 2.3, described below and proved in §3 via hyperplane arrangements.

2.2 Mating PolyPoints

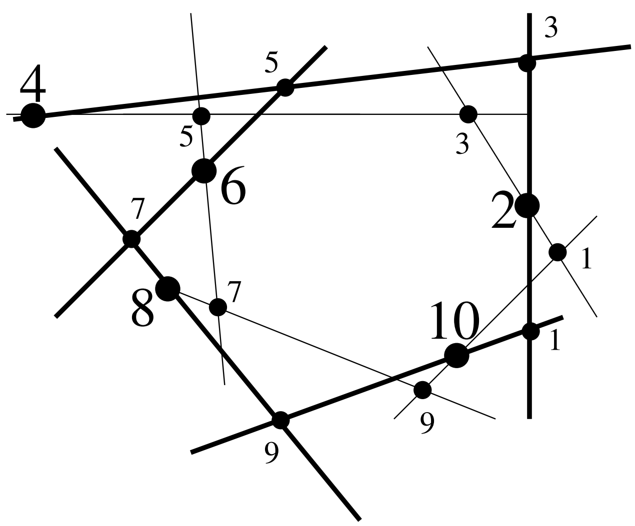

Figure 2.1

We fix $n \geq 3$. So that our definitions below make linguistic sense, we will say that *chain* and *n-Point* are synonyms. Let $X = (x_1, x_3, ..., x_{2n-1})$ and $Y = (y_1, y_3, ..., y_{2n-1})$ be two n-Points. We construct $Z = (z_2, z_4, ... z_{2n})$ by defining

$$z_j = \overline{x_{j-1}x_{j+1}} \cap \overline{y_{j-1}y_{j+1}}.$$

This is shown in Figure 2.1.

Of course, we only make this construction if the relevant lines and intersection points are defined. Assuming Z exists, we call $Z = X * Y$ the *offspring* of X and Y. If X and Y are labelled by even integers, we can make the same definition, taking indices mod $2n$.

We say that a sequence of n-Points $A_1 = (A_{1,1}, A_{1,3}, ..., A_{1,2m-1})$ is a *PolyChain*, and more specifically, an *m-Chain*. If $A_j = (A_{j,j}, A_{j,j+2}, ..., A_{j,2m-j})$ exists, we write

$$A_{j+1} = (A_{j+1,j+1}, ..., A_{j+1,2m-(j+1)}); \qquad A_{j+1,k} = A_{j,k-1} * A_{j,k+1},$$

provided this is defined. We write $A_1 \to A_2 \to ... \to A_g$, if all these PolyChains exist. We call this progression the *mating process*. We will say that the mating process is *well defined* on A_1 if $A_1 \to ... \to A_m$. We call the $A_{m,m}$ the *final offspring* of A_1. It is convenient to represent the mating process as a triangular array. For instance, if $m = 4$, we could write

$$\begin{array}{ccccccc} A_{1,1} & & A_{1,3} & & A_{1,5} & & A_{1,7} \\ & A_{2,2} & & A_{2,4} & & A_{2,6} & \\ & & A_{3,3} & & A_{3,5} & & \\ & & & A_{4,4} & & & \end{array}$$

provided that the mating process was defined.

Suppose X and Y are the two n-Points above. We say that X and Y are *compatible* if the n lines $\{\overline{x_j y_j}\}$ are concurrent−i.e. have a common intersection point. We will sometimes write $X \leftrightarrow Y$ in this case. We say that the PolyChain A_1, described above, is a *special PolyChain* if $A_{1,j-1} \leftrightarrow A_{1,j+1}$ for all j. Let $\mathcal{D}_n$ denote the space of special $(n-1)$-Chains. In §3 we will prove:

Theorem 2.3 *Let $n \geq 3$. If the mating process is defined on $A \in \mathcal{D}_n$, then the final offspring of A consists of collinear points.*

The case $n = 3$ of Theorem 2.3 involves 6 points, arranged into 2 triangles, and is a restatement of Desargues' theorem. In general, Theorem 2.3 involves $n(n-1)$ points, arranged into a sequence of n-Points. For a schematic representation (in the case $n = 4$), we could write

$$\begin{array}{ccccccc} A_{1,1} & \leftrightarrow & A_{1,3} & \leftrightarrow & A_{1,5} & \leftrightarrow & A_{1,7} \\ & A_{2,2} & & A_{2,4} & & A_{2,6} & \\ & & A_{3,3} & & A_{3,5} & & \\ & & & !A_{4,4} & & & \end{array}$$

The exclamation point means that the points of $A_{4,4}$ all lie on a single line.

2.3 Deriving PolyChains from PolyLines

In this section, we explain how to deduce Lemma 2.2 from Theorem 2.3. This deduction is just a matter of carefully keeping track of labellings.

Suppose that (n, p, q) are as in Lemma 2.2. Let L be any np-Line, not necessarily a satellite, such that $\delta_q(L)$ exists. Let X and Y be the n-Points above. We write $X \bowtie_p L$ if $X \subset \delta_q(L)$, and the points of X are evenly spaced in the cyclic ordering on $\delta_q(L)$. (This is to say that x_{j+1} succeeds x_{j-1} by p.) If $X, Y \bowtie_p L$, we write $X \prec_q Y$ if, for all relevant indices, the jth point of Y succeeds the jth point of X by q, in the cyclic order of $\delta_q(L)$. We say that the chain A_1 is (p, q)-*derived* from L if, for all relevant indices, $A_{1,j} \bowtie_p L$ and $A_{1,j-1} \prec_q A_{1,j+1}$.

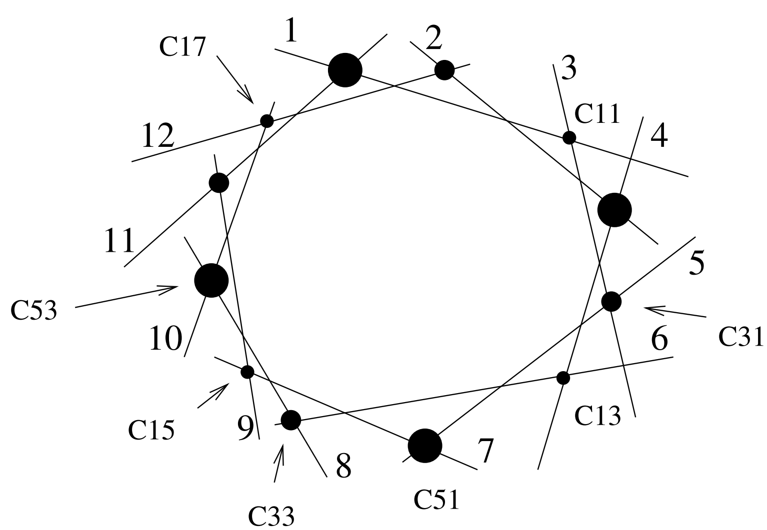

Figure 2.2

Figure 2.2 shows a 3-Chain (C_1, C_3, C_5) which is $(3,2)$-derived from a 12-Line. The points of C_j are, successively $C_{j,1}$, $C_{j,3}$, $C_{j,5}$, and $C_{j,7}$.

Lemma 2.4 *For m and v, let $g \in [1, \min(m, v+1)]$. Let L be an np-Line, such that $(\delta_q \circ \delta_p)^v \circ \delta_q$ is defined on L. Let A_1 be an m-Chain, (p,q)-derived from L. Then $A_1 \to \ldots \to A_g$ and A_g is (p,q)-derived from $(\delta_p \circ \delta_q)^{g-1}(L)$.*

Proof: In general, for $k \geq 2$, let α be a k-Chain (p,q)-derived from an np-Line Ω. If $\delta_q \circ \delta_p \circ \delta_q(\Omega)$ is defined, it follows straight from the definitions that $\alpha \to \beta$, where β is a $(k-1)$-Chain which is (p,q)-derived from $\delta_p \circ \delta_q(\Omega)$. Our lemma follows from this observation and induction. ♠

Now, assume that $(\delta_p \circ \delta_q)^{n-2} \circ \delta_q$ is defined on $L \in L(n,p)$. Suppose that A_1 is an $(n-1)$-Chain which is (p,q)-derived from L. It follows straight from the definitions that $A_1 \in \mathcal{D}_n$. From Lemma 2.4, the mating process is well defined on A_1, and $A_{n-1,n-1} \bowtie_p M = (\delta_p \circ \delta_q)^{n-2}(L)$. By Theorem 2.3, the n points of $A_{n-1,n-1}$ are collinear. Cyclically relabelling, we see that this is true for all A such that $A \bowtie_p M$. This is the same as saying that $L' = \delta_q(M) = (\delta_q \circ \delta_p)^{n-2} \circ \delta_q(L)$ belongs to $P(n,p)$.

2.4 Complex Algebraic Manifolds

What remains in this chapter is to deduce Lemma 2.1 from Lemma 2.2. We begin by discussing a particular kind of complex manifold.

A map $\phi : \boldsymbol{C}^k \to \boldsymbol{C}^k$ is *birational* if both ϕ and ϕ^{-1} are rational. We say that a manifold M is a *complex algebraic manifold* if there is an atlas of local coordinate charts $M \to \boldsymbol{C}^k$, whose transition functions are restrictions of birational maps.

A subset $S \subset M$ is *singular* if, in each local chart, S is a subset of a lower dimensional complex affine variety. We say that $U \subset M$ is *generic* if $M - U$ is singular. Finite intersections and unions of generic subsets are generic. A map $f : M \to N$ between complex algebraic manifolds is said to be *rational* iff f is rational when considered in local coordinates. Typically, such a map will only be defined on a generic subset of M.

In what follows, we attach a linear order to our polygons which is compatible with our cyclic order. This amounts to choosing a distinguished point for each polygon. Once we do this, every configuration space in sight is a complex algebraic manifold and all maps in sight are rational.

2.5 Proof of Lemma 2.1

Say that the map ϕ_h is *bi-defined* on some subset $S \subset P(n,p)$ if ϕ_h is defined and injective on S, and ϕ_h^{-1} is defined and injective on $\phi_h(S)$. We will say that S is *h-good* if ϕ_h is bi-defined on S. Since our maps are rational maps, the existence of a nonempty h-good subset implies the existence of a generic h-good subset. In fact, the set of all h-good elements is generic, if nonempty. Let ξ denote the largest value of h such that there exists a nonempty h-good subset. It follows immediately from Lemma 2.2 that $\xi \leq 2n-3$. Let $GP \subset P(n,p)$ and $GL \subset L(n,p)$ denote the generic sets of ξ-good elements.

Sub-Lemma 2.5 *$\phi_{\xi+1}$ is defined on a generic subset $V \subset P(n,p)$.*

Proof: It suffices to show that $\phi_{\xi+1}$ is defined on some element of $P(n,p)$. Suppose this is false. Suppose also that ξ is odd. The case when ξ is even, which does not actually arise, has a similar proof. Let $X \in GP$, and let $Y = \phi_\xi(X)$. We claim that $Y = (y_1, y_2, ..., y_{np})$ is an n-cover of a p-Point. If not, then $y_j \neq y_{j+p}$ for some j. This means that this equality fails on a generic set. Cyclically relabelling, and intersecting, we see that this equality fails, for all j, on a generic set $S \subset GP$. But $\phi_{\xi+1} = \delta_p \circ \phi_\xi$ is defined for $X \in S$, a contradiction. So, assuming $\phi_{\xi+1}$ is never defined, we see that $\phi_\xi(X)$ is an n-cover of a p-Point. By choice of X, the PolyPoint $\delta_q(\phi_\xi(X)) = \phi_{\xi-1}(X)$ is defined. This PolyLine is clearly the n-cover of a p-Line. Hence, δ_p is not defined on $\phi_{\xi-1}(X)$. This implies that $\phi_{\xi-1}^{-1}$ is not defined on $\phi_{\xi-1}(X)$, a contradiction to $X \in GP$. ♠

Let $V \subset P(n,p)$ be the generic subset on which $\phi_{\xi+1}$ is defined. $\phi_{\xi+1}^{-1}$ is never defined on $\phi_{\xi+1}(V)$, for otherwise, we could increase the value of ξ. The same argument as above says that $\phi_{\xi+1}(X)$ is an n-cover, for any $X \in V$. Taking limits, we see that this is also true for $X \in GP$. Hence, by definition $\phi_\xi(GP) \subset L(n,p)$. Since ϕ_ξ^{-1} is defined on $\phi_\xi(GP)$, we have $\phi_\xi(GP) \subset GL$. Dually, $\phi_\xi(GL) \subset GP$. Since ξ is odd, ϕ_ξ is an involution on the above domains. Therefore, $\phi_\xi(GP) = GL$ and $\phi_\xi(GL) = GP$.

If ξ is even, then $\phi_{\xi+1}(X)$ is a $\frac{np}{q}$-cover of a q-gon, for $X \in V$. This is impossible since (by hypothesis) q does not divide n. Hence ξ cannot be even. Since the even case does not occur, the proof is complete.

3 Hyperplane Arrangements

In this chapter, we will prove Theorem 2.3. For $d \in [1, n-1]$, say that a *d-flat* is a copy of $\boldsymbol{C}^d$ in $\boldsymbol{C}^n$. As usual, we call the extreme cases *hyperplanes* and *lines*. Define $\pi(z_1, ..., z_n) = (z_1, z_2)$.

3.1 Joints and Tubes

A *joint* is a sequence $J = (p_1, p_3, ..., p_{2n-1})$ of n general position points in $\boldsymbol{C}^n$. Let $|J|$ denote the hyperplane spanned by the points of J. A *tube* is sequence of n lines, having the form $qJ = (\overline{q}p_1, ..., \overline{q}p_{2n-1})$, where $q \in \boldsymbol{C}^n - |J|$. The complex lines of qJ inherit labellings from J.

Given joints J_1 and J_3, such that $J_1 \cap J_3 = \emptyset$, let $\overline{J_1 J_3}$ be the collection of n lines determined by corresponding points of J_1 and J_3. We write $J_1 \leftrightarrow J_3$ if and only if $\overline{J_1 J_3}$ is a tube. A *PolyJoint* (or *m-Joint*) is a sequence of joints $\Omega = (J_1 \leftrightarrow J_3 \leftrightarrow ... \leftrightarrow J_{2m-1})$. By definition, Ω determines a sequence of (typically distinct) tubes $(T_2, T_4, ...T_{2m-2})$, where $T_{2k} = \overline{J_{2k-1} J_{2k+1}}$.

We call $\pi\Omega = (\pi(J_1), ..., \pi(J_{2m-1}))$ the *projection* of Ω. We call Ω a *cover* of $\pi\Omega$. By construction, $\pi\Omega$ is a special PolyChain. Conversely, let $A = (A_1, A_3, ..., A_{2m-1})$ be a special PolyChain. Let L_j denote the collection of lines determined by correspondingly labelled points of A_{j-1} and A_{j+1}. We say that A is *coverable* if $\forall j$ the points of A_j are in general position and belong to $\boldsymbol{C}^2$ and the lines of L_j are pairwise distinct, and nonparallel in $\boldsymbol{C}^2$.

Lemma 3.1 (Lifting) *Coverable PolyChains have PolyJoint covers.*

Proof: Let $A = (A_1, A_3 ..., A_{2m-1})$ be a coverable m-Chain. We will induct on m. Certainly, we may find a joint J_1 such that $\pi(J_1) = A_1$. Suppose, by induction, we have a PolyJoint $(J_1, ..., J_{2k-1})$ which projects to $(A_1, ..., A_{2k-1})$. Let $x_{2k} \in \boldsymbol{C}^2$ be the unique common point of the lines in L_{2k}. Since the points of A_{2k-1} are in general position, $\pi(|J_{2k-1}|) = \boldsymbol{C}^2$. Therefore, $|J_{2k-1}|$ is transverse to the level sets of π. This means that there is some $\tilde{x}_{2k} \in \pi^{-1}(x_{2k}) - |J_{2k-1}|$. Since $\pi(\tilde{x}_{2k} J_{2k-1}) = L_{2k}$, and the points of A_{2k+1} are contained, one per line, in the lines of L_{2k}, the set $J_{2k+1} = \pi^{-1}(A_{2k+1}) \cap \tilde{x}_{2k} J_{2k-1}$ is a joint. By construction, $J_{2k+1} \leftrightarrow J_{2k-1}$ and $\pi(J_{2k+1}) = A_{2k+1}$. This completes the induction step. ♠

3.2 Cyclic Skeleton

Let $T = (t_1(1), t_1(3)..., t_1(2n-1))$ be a tube. We set $\Sigma_1 T = T$. Inductively, we set $\Sigma_k T = (..., t_k(j), ...)$, where $t_k(j)$ is the flat spanned by the flats $t_{k-1}(j-1)$ and $t_{k-1}(j+1)$. Note that $t_k(*)$ is spanned by k consecutive lines of T, and hence has dimension k. The n flats of $\Sigma_k T$ inherit their cyclic order from the cyclic order of T. When k is odd (resp. even), the elements of $\Sigma_k T$ are labelled by odd (resp. even) integers. Note also that $t_k(j) = t_{k-1}(j-1) \cap t_{k-1}(j+1)$. We call the union $\Sigma_* T$ the *cyclic skeleton* of T.

Suppose that W is a codimension j flat. Suppose that T is a tube. We say that W is *slices* T if $W \cap \Sigma_j T$ consists of n distinct points and (provided $j \leq n-2$) $W \cap \Sigma_{j+1} T$ consists of n distinct lines. If W slices T, we define $W_T = \pi(W \cap \Sigma_j T)$. The points of W_T are labelled by odd or even integers, depending on the parity of j. Let W and W' be distinct flats. We say that the pair (W, W') *slices* T if $\dim(W \cap W') + 1 = \dim(W) = \dim(W')$ and if W, W' and $W \cap W'$ all slice T. The following simple result is the key to our entire proof:

Lemma 3.2 (Mating) *Suppose the pair (V, V') slices T. If the offspring $(V)_T * (V')_T$ is defined, then $(V \cap V')_T = (V)_T * (V')_T$.*

Proof: Let d be the common codimension of V and V'. Assume that d is even. The odd case is the same, except for the parity of the labellings. $V \cap \Sigma_d$ consists of n points $p_0, p_2, ..., p_{2n-2}$ and $V \cap \Sigma_{d+1} T$ consists of n distinct lines, $L_1, L_3..., L_{2n-1}$. (likewise for V'.) Let $q_j = (V \cap V') \cap t_{d+1}(j)$. We claim that

$$(1) \qquad L_j = \overline{p_{j-1} p_{j+1}}; \qquad L'_j = \overline{p'_{j-1} p'_{j+1}} \qquad q_j = L_j \cap L'_j.$$

By hypothesis, $p_i \in t_d(i)$. Since $t_d(j \pm 1) \subset t_{d+1}(j)$, we have $p_{j-1}, p_{j+1} \in V \cap t_{d+1}(j)$. Since these points are distinct, $\overline{p_{j-1} p_{j+1}}$ is uniquely defined and $\overline{p_{j-1} p_{j+1}} \subset V \cap t_{d+1}(j)$. Since $V \cap t_{d+1}(j) = L_j$, we get the first part of (1). The second part of (1) follows from the trivial fact $(V \cap t_{d+1}(j)) \cap (V' \cap t_{d+1}(j)) = (V \cap V') \cap t_{d+1}(j)$. By definition,

$$(2) \quad [V]_T = (..., \pi(p_i)...); \ [V']_T = (..., \pi(p'_i)...); \ [V \cap V']_T = (..., \pi(q_j), ...).$$

Our lemma follows from projecting (1) into the plane, and using (2). ♠

3.3 Full Transversal

Suppose that Ω is the m-Joint described in §3.1. Define $H_{1,k} = |J_k|$, and

$$H_{g,k} = H_{g-1,k-1} \cap H_{g-1,k+1}; \qquad g \in [2,m], \qquad k \in [g, 2m-g].$$

In all, we get $H_{g,g}, H_{g,g+2}, ..., H_{g,2m-g}$, for $g \in [1,m]$. We do not rule out the possibility that some H_* is a point or the emptyset.

Define $k \hookleftarrow h$ iff h is even and $|h-k| \leq 1$. If $s < \min(m,n)$, we say that Ω is s-*sliced* if the following holds for all relevant indices: If $g \leq s$ and $k \hookleftarrow h$, then $H_{g,k}$ has codimension g and $H_{g,k}$ slices T_h. In this case, and for these indices, define

$$A_{g,k} = \pi(H_{g,k} \cap \Sigma_k T_h); \qquad A_g = (A_{g,g}, A_{g,g+2}, ..., A_{g,2m-g}).$$

Note that $A_1 = \pi\Omega$.

Lemma 3.3 *Suppose Ω is d-sliced, and the mating process is d-defined on $A_1 = \pi\Omega$. Then $A_1 \to ... \to A_d$.*

Proof: We just have to show that $A_{g+1,k}$ is the offspring of $A_{g,k-1}$ and $A_{g,k+1}$, for $g \leq d-1$, and $k \in \{g+1, ..., 2m-g-1\}$. Suppose first that g is odd and k is even. Then

$$A_{g,k-1} = \pi(H_{g,k-1} \cap \Sigma_g T_k); \qquad A_{g,k+1} = \pi(H_{g,k+1} \cap \Sigma_g T_k);$$

$$A_{g+1,k} = \pi(H_{g+1,k} \cap \Sigma_{g+1} T_k); \qquad H_{g+1,k} = H_{g,k-1} \cap H_{g,k+1}.$$

Applying the Mating Lemma to the pair $(H_{g,k-1}, H_{g,k+1})$, and the tube T_k, we get the desired result in this case. If g is even and k is odd,

$$A_{g,k-1} = \pi(H_{g,k-1} \cap \Sigma_g T_{k-1}); \qquad A_{g,k+1} = \pi(H_{g,k+1} \cap \Sigma_g T_{k+1});$$

$$A_{g+1,k} = \pi(H_{g+1,k} \cap \Sigma_{g+1} T_{k+1}); \qquad H_{g+1,k} = H_{g,k-1} \cap H_{g,k+1}.$$

By construction, $\Sigma_* T_{k-1} \cap H_{1,k} = \Sigma_* T_{k+1} \cap H_{1,k}$. Since $2 \leq g < k$, we have $H_{g,k-1} \subset H_{g-1,k} \subset H_{1,k}$. Putting the two facts together, we have

$$H_{g,k-1} \cap \Sigma_g T_{k-1} = H_{g,k-1} \cap \Sigma_g T_{k+1}.$$

Therefore,

$$A_{g,k-1} = \pi(H_{g,k-1} \cap \Sigma_g T_{k+1});$$

Applying the Mating Lemma to the pair $(H_{g,k-1}, H_{g,k+1})$, and the tube T_{k+1}, we get the desired result in this case. ♠

3.4 Proof of Theorem 2.2

Let $\mathcal{B}$ be the space of coverable $(n-1)$-Chains, such that each chain is an n-Point. Thus $\mathcal{B}$ is a subset of $\mathcal{D}_n$, which appears in Theorem 2.3. Let $\mathcal{E}$ be the set of n-Tubes which cover $\mathcal{B}$. Let $\mathcal{S} \subset \mathcal{E}$ be the set of $(n-1)$-sliced elements.

Lemma 3.4 *$\mathcal{S}$ is open dense in $\mathcal{E}$.*

Proof: By stringing out the coordinates of the points of an element of $\mathcal{B}$ and of $\mathcal{E}$, it is easy to see that both these spaces can be expressed as $V - W$, where V and W are complex affine varieties. It is easy to see that $\mathcal{B}$ is connected, since we are working over $\boldsymbol{C}$. Given an element of $\mathcal{B}$, one finds an element of $\mathcal{E}$ which covers it by the procedure of the Lifting Lemma. The proof given there shows that $\mathcal{E}$ fibers over $\mathcal{B}$. The fiber again has the form $V - W$, where V and W are complex affine varieties. Since we are working over $\boldsymbol{C}$, this fiber is connected. In summary, $\mathcal{B}$ and $\mathcal{E}$ are both connected complex algebraic manifolds.

$\mathcal{S}$ is determined by the nonvanishing of certain rational functions. Thus, $\mathcal{S}$ is open and dense provided it is nonempty. Let $F_1, ..., F_{n-1}$ be a collection of general position hyperplanes. Let Z be any tube. By transversality, every F_j will be transverse to every flat in the cyclic skeleton of Z', where Z' is a suitable isometric copy of Z. Let $\Theta = (F_1 \cap Z', ..., F_{n-1} \cap Z')$. By construction, $\Theta \in \mathcal{S}$. ♠

The preceding Lemma implies that $\mathcal{B}$ is dense in $\mathcal{D}_n$, which is fairly obvious anyhow. Suppose that the mating process is defined on $A_1 \in \mathcal{D}_n$. By continuity, the mating process is defined on all elements of $\mathcal{D}_n$ sufficiently close to A_1. By perturbation, we can assume that $A_1 \in \mathcal{B}$. In other words, $A_1 = \pi\Omega$, where $\Omega \in \mathcal{S}$. Let $A_1 \to ... \to A_{n-1}$ be the mating process for A_1. In the notation of §3.3, we have $A_{n-1,n-1} = \pi(H_{n-1,n-1} \cap T_v)$, where $n-1 \hookleftarrow v$. Since Ω is $(n-1)$-sliced, $H_{n-1,n-1}$ is a 1-flat. This is to say that the points of $A_{n-1,n-1}$ are collinear.

4 Lower Bound

Theorem 1.1 follows from Lemma 2.1 and from

Lemma 4.1 (Lower Bound) *Let (n, p, q, ξ) be as in Lemma 2.1. Then $\xi = 2n - 3$.*

4.1 Proof Modulo Transversality

In this section we prove Lemma 4.1, modulo a transversality lemma.

Let $\underline{h} = \frac{1}{2}(\xi + 1)$, so that $\delta_q \circ (\delta_p \circ \delta_q)^{\underline{h}-1} = \phi_\xi$. Let $N > np$ be any integer, fixed once and for all. Say that an element $L \in L(n, p)$ is *good* if there is an N-joint Ω such that $\pi\Omega$ is (p, q)-derived from L, and (see §3.3)

1. Ω is $\underline{h}$-sliced,
2. $H_{d,k} \neq \emptyset$ for $d \in [1, n]$ and for all relevant k.
3. $H_{g,k} \cap \Sigma_g T_z$ spans $H_{g,k}$, for $g \in [1, \underline{h}]$ and $k \hookleftarrow z$.

In this case, we will say that Ω is *associated to* L. Let $\mathcal{G} \subset L(n, p)$ be the set of good elements. Assume for now that:

Lemma 4.2 *The closure of $\mathcal{G}$ in $L(n, p)$ has nonempty interior.*

Let $\mathcal{P} \subset GP$ denote those elements Y such that the np points of Y are distinct, and the p lines of $\delta_p(Y)$ are in general position. By Theorem 2.1, the map ϕ_ξ is a birational involution swapping GP and GL. Hence, ϕ_ξ takes generic sets to generic sets. $\mathcal{P}$ is clearly generic, and hence so is $\phi_\xi(\mathcal{P})$. By 4.2 there exists some $X \in \phi_\xi(\mathcal{P}) \cap \mathcal{G}$. Let $Y = \phi_\xi(X)$, and let Ω be an N-Joint associated to X. Pick $k \in [2n, 2N - 2n]$. By Lemma 2.4 and Lemma 3.3 combined, the 3-Chain $(A_{\underline{h},k-2}, A_{\underline{h},k}, A_{\underline{h},k+2})$ is (p, q)-derived from $(\delta_p \circ \delta_q)^{\underline{h}-1}(X)$. The points of $A_{\underline{h},v}$ belong to Y and are evenly spaced in the cyclic ordering on Y, for $v = k-2, k, k+2$. Since $Y \in \mathcal{P}$, the span L_v of $A_{\underline{h},v}$ is a line, for our choices of v. From Property 3 above, the span of $H_{g,v} \cap \Sigma_g T_w$ projects to the span of $A_{g,v} = \pi(H_{g,v} \cap \Sigma_g T_w)$, provided $v \hookleftarrow w$. Therefore, $L_v = \pi(H_{\underline{h},v})$. If $\underline{h} \leq n - 2$, then

$$L_{k-2} \cap L_k \cap L_{k+2} \supset \pi(H_{\underline{h},k-2} \cap H_{\underline{h},k} \cap H_{\underline{h},k+2}) = \pi(H_{\underline{h}+2,k}) \neq \emptyset,$$

from Property 2 above. The above three lines are consecutive in $\delta_p(Y)$. Since $p \geq 3$, their concurrence contradicts $Y \in \mathcal{P}$. Hence $\underline{h} = n-1$ and $\xi = 2n-3$.

4.2 Configuration Spaces

Let $\mathcal{B} = \mathcal{B}(N; n, p, q)$ be the subspace of coverable N-Chains (p, q)-derived from elements of $L(n, p)$. Let $\mathcal{E} = \mathcal{E}(N; n, p, q)$ be the space of N-Joints which cover elements $\mathcal{B}(N; n, p, q)$.

Lemma 4.3 *$\mathcal{B}$ and $\mathcal{E}$ are connected algebraic manifolds.*

Proof: The following construction reveals that $\mathcal{B}$ is smooth and connected. Choose p distinct points in $\boldsymbol{C}^2$. Choose np distinct lines, which cyclically contain the chosen points. This determines an element $L \in L(n, p)$. Now take the m-Chain A_1 which is (p, q)-derived from L, such that the first point of $A_{1,1}$ is the first point of L. By consruction, $A_1 \in \mathcal{B}$. Since $N > np$, different choices of L determine different elements of $\mathcal{B}$.

Let $\Omega \in \mathcal{E}$. Let $A_1 = \pi\Omega \in \mathcal{B}$. Any lift of A_1 can be constructed from the following *lifting data*: A joint J_1 such that $\pi J_1 = A_{1,1}$, and a sequence of points $\tilde{x}_{2k}$, such that $\pi(\tilde{x}_{2k}) = x_{2k}$. Define $\psi(z_1, ..., z_n) = (z_3, z_4, ..., z_n)$. Each point in the lifting data is determined by the corresponding point in A_1 and by its image under ψ. Since $\mathcal{B}$ is a smooth manifold, our construction shows that there is a "product neighborhood" of Ω, diffeomorphic to a ball, which fibers over a neighborhood of A_1 in $\mathcal{B}$. The connectivity of $\mathcal{E}$ follows from the connectivity of $\mathcal{B}$, and the (obvious) connectivity of the fibers. ♠

4.3 Transversality Proof Overview

In this section we outline the proof of Lemma 4.2. Let $\Omega \in \mathcal{E}$. We adopt the notation of §3.3. Given a tube T, let $\Sigma_r^j T$ denote the jth flat of $\Sigma_r T$, when these flats are ordered in a way which is compatible with their cyclic order.

Say that Ω is *cyclically invariant* if the function $j \to \dim(H_{g,k} \cap \Sigma_r^j T_m)$ is independent of j for any choice of indices (g, k, r, m). Let $\mathcal{CI}$ be the set of such Ω. In §4.5 we will show that $\mathcal{CI}$ is open and dense in $\mathcal{E}$.

Say that a PolyChain $A_k = (..., A_{k,i}, ...)$ is *affine* if, for all relevant indices α and β the following is true: $A_{k,\alpha-1} \cup A_{k,\alpha+1}$ consists of $2n$ distinct points of $\boldsymbol{C}^2$ and $\delta_1(A_{k,\beta})$ consists of n distinct lines, all of which intersect $\boldsymbol{C}^2$. Say that Ω is $\underline{h}$-*affine* if the mating process is $\underline{h}$-defined on A_1, and A_j is an affine chain, for $j = 1, ..., \underline{h}$. Let $\mathcal{AF}$ denote the set of such Ω. In §4.6 we will show that $\mathcal{AF}$ is open dense in $\mathcal{E}$.

Say that Ω is *transverse invariant* if $\exists\gamma \in [1,n]$ such that $\dim(H_{g,k}) = \max(n-\gamma, n-g)$, for $g \in [1,n]$. Let $\mathcal{TI}$ be the set of such Ω. In §4.7 we will show that $\mathcal{TI}$ is open dense in $\mathcal{E}$.

Let $\mathcal{U} = \mathcal{CI} \cap \mathcal{TI} \cap \mathcal{AF}$. From the statements above, $\mathcal{U}$ is open dense in $\mathcal{E}$. In §4.8 we will show that any element of $\mathcal{U}$ satisfies the three properties listed in §4.1. Since $\pi : \mathcal{E} \to \mathcal{B}$ is continuous, $\pi(\mathcal{U})$ is dense in $\mathcal{B}$. Now, $\mathcal{B}$ contains an open subset of the space $\mathcal{C}$ of all N-Chains which are (p,q)-derived from elements of $L(n,p)$. This is to say that the closure of $\pi(\mathcal{U})$ in $\mathcal{C}$ has nonempty interior. Since $N > np$, an element of $\mathcal{C}$ determines, and is determined by, the corresponding element of $L(n,p)$. The spaces $\mathcal{C}$ and $L(n,p)$ are thus canonically homeomorphic. This homeomorphism carries $\mathcal{G} \subset L(n,p)$ to $\pi(\mathcal{U}) \subset \mathcal{C}$. Hence the closure of $\mathcal{G}$ has nonempty interior in $L(n,p)$, as desired.

4.4 The Basic Set

Given a flat V, let $\overline{V}$ denote the completion of V in $\boldsymbol{CP}^n$. Let $V^\infty = \overline{V} - V$. Let $\Theta = (V_1, ..., V_k)$ be a family of hyperplanes comprising part of Ω. We think of Θ as specifying combinatorially which hyperplanes to take. The actual hyperplanes will vary with the choice of $\Omega \in \mathcal{E}$.

Let $\overline{\Theta} = (\overline{V}_1, ..., \overline{V}_k)$, and let $\Theta^\infty = (V_1^\infty, ..., V_k^\infty)$. Define $\cap\Theta$ to be the intersection of all the hyperplanes in Θ. Define $\cap\overline{\Theta}$ and $\cap\Theta^\infty$ similarly. There are only finitely many distinct choices for Θ. Let $\mathcal{M} \subset \mathcal{E}$ denote the subset on which the functions $\dim(\cap\overline{\Theta})$ and $\dim(\cap\Theta^\infty)$ attain their minimum value, for each and every choice of Θ. From Lemma 4.3, we know $\mathcal{E}$ is a connected complex algebraic manifold. The complement of $\mathcal{M}$ is described by the vanishing of various rational functions. Since $\mathcal{M}$ is nonempty, it is generic.

4.5 Cyclic Invariance

In this section we prove that $\mathcal{M} \subset \mathcal{CI}$, which implies that $\mathcal{CI}$ is open dense.

Given a joint $J = (p_1, p_3, ..., p_{2n-1})$, let $\rho(J) = (p_3, p_5, ..., p_{2n-1}, p_1)$ be the joint obtained by cyclically relabelling the points of J. Given a PolyJoint $\Omega = (J_1, ..., J_{2N-1})$, let $\rho(\Omega) = (\rho(J_1), ..., \rho(J_{2N-1}))$. Note that $\mathcal{E}$ is clearly invariant under ρ. By construction $\rho(\mathcal{M}) = \mathcal{M}$. This implies that $\mathcal{M} \subset \mathcal{CI}$.

4.6 Affine PolyJoints

In this section we prove that $\mathcal{AF}$ is open dense. By algebraicity, it suffices to show that $\mathcal{AF} \neq \emptyset$. Let GL and GP be the sets in Lemma 2.1. Suppose we could find some $L \in GL$ such that the n lines/points of $\phi_i(L)$ are all distinct, for $i = 1, ..., \xi$. Since ϕ_i commutes with complex projective transformations, we could perturb L by a complex projective transformation so that all relevant points belong to $\boldsymbol{C}^2$ and all relevant lines intersect $\boldsymbol{C}^2$. Any N-chain A_1 derived from this perturbation would satisfy the conditions used to define *affine*. By the lifting lemma, there is some Ω such that $\pi\Omega = A_1$. By construction, Ω would be $\underline{h}$-affine.

If no element such as L exists, there is some index $j \in [1, \xi]$ such that the ath and bth points/lines of $\phi_j(L)$ coincide, for all $L \in GP$. Relabelling, we see that the $(a+k)$th and $(b+k)$th points/lines coincide for all k. This property is inherited by $\phi_{j+j'}(L)$, for $j' = 1, 2, ...$, so that the points of $\phi_\xi(L)$ are never entirely distinct. This contradicts the fact that $\phi_\xi(GL) = GP$, because elements in GP generically consist of np-distinct points.

4.7 Transverse Invariance

In this section we prove, modulo a technical detail, that $\mathcal{M} \subset \mathcal{TI}$, which implies that $\mathcal{TI}$ is open dense.

As a bit of notation, a *projective flat* in $\boldsymbol{CP}^n$ is the projectivization of a linear subspace of $\boldsymbol{C}^{n+1}$.

Lemma 4.4 *Let $P_1, P_2, P_3 ...$ be a sequence of codimension one projective flats. Suppose that $d_r = d_r(j) = \dim(P_i \cap ... \cap P_{i+r})$ depends only on r. Then $d_r - d_{r+1} \leq 1$. If $d_{r+1} = d_r \geq 0$, for some $r \leq n$ then $P_{i+1} \cap ... \cap P_{i+n+1}$ is independent of i. In particular, $\cap P_j \neq \emptyset$.*

Proof: The first statement is standard. (See [**K**, Cor 3.21].) If $d_r(j) = d_{r+1}(j)$, then $P_j \cap ... \cap P_{j+r} = P_{j+1} \cap ... \cap P_{j+r+1}$. This is independent of j, and gives the second statement. ♠

Let $\Theta(g,k)$ denote the collection of hyperplanes $\{H_{1,j}\}$ which intersect to give $H_{g,k}$. So, $\Theta(g,k)$ consists of g consecutive hyperplanes in the list. The choice of k determines which hyperplanes to take. Define

$$\overline{d}(g,k) = \dim(\cap\overline{\Theta}(g,k)); \quad d^\infty(g,k) = \dim(\cap\Theta^\infty(g,k)).$$

By definition, $\overline{d}(g,k)$ and $d^\infty(g,k)$ are constant functions on $\mathcal{M}$.

Lemma 4.5 *The functions* $\overline{d}(g,k)$ *and* $d^\infty(g,k)$ *are independent of* k *on* $\mathcal{M}$.

Proof: $\Theta(g,k) = \{H_{1,k-g+1},...,H_{k+g-1}\}$. Let $\Omega[g,k] = (J_{k-g+1},...,J_{k+g-1})$ be the corresponding g-joint. Let $\mathcal{E}[g,k] = \{\Omega[g,k] |\ \Omega \in \mathcal{E}\}$. The argument in the Lifting Lemma implies that $\mathcal{E}[g,k]$ consists exactly in those g-Joints associated to elements of $L(n,p)$. The point is that any such g-Joint can be prolonged to an N-Joint, by the procedure in the Lifting Lemma. In particular, $\mathcal{E}[g,k]$ is independent of k. Since our two functions only depend on $\Omega[g,k]$, they are independent of k on $\mathcal{M}$. ♠

Given a PolyJoint Ω, define $\overline{I}(\Omega) = \cap \overline{H}_{1,j}$, where the intersection is taken over all possible indices j. Recall that $\pi : \boldsymbol{C}^n \to \boldsymbol{C}^2$ is our projection. Let $A = \overline{\pi^{-1}(0,0)} - \boldsymbol{C}^n$. Let $B = \overline{\boldsymbol{C}^2}$. In §5 we will prove

Lemma 4.6 (Existence) $\overline{I}(\Omega_0) \cap A = \emptyset$ *for some* $\Omega_0 \in \mathcal{E}$.

Corollary 4.7 $\exists \Omega_1 \subset \mathcal{M}$ *such that either* $\overline{I}(\Omega_1) = \emptyset$ *or* $\overline{I}(\Omega_1) \cap \boldsymbol{C}^n \neq \emptyset$.

Proof: We will base our example on Ω_0. We can perturb Ω_0 so that $\Omega_0 \in \mathcal{M}$. The only case we need to consider is when $I(\Omega_0) \neq \emptyset$. Since $\mathcal{M}$ is open and PolyJoints are compact, $T(\Omega_0) \in \mathcal{M}$ for all T sufficiently close to the identity. The subgroup of projective transformations which stabilize the pair (A,B) acts transitively on $\boldsymbol{CP}^n - A - B$, and has the identity as limit point. Hence $\Omega_1 = T(\Omega_0)$ intersects $\boldsymbol{C}^n$ for a suitable choice of T. ♠

Suppose, with a view towards getting a contradiction, that $\Omega_2 \in \mathcal{M} - \mathcal{TI}$. Let $\overline{d}_g = \overline{d}(g,k)$ and $d_g^\infty = d^\infty(g,k)$. If $\forall g\ \dim(H_{g,k}) = \overline{d}_k$ then Lemma 4.4 implies $\Omega_2 \in \mathcal{TI}$. Hence $\dim(H_{u,k}) \neq \overline{d}_u$ for some $u \in [1,n]$. Since $\overline{H}_{u,k}$ is a projective flat, we have $\overline{d}_u = d_u^\infty$. This implies that $\overline{I}(\Omega) \cap \boldsymbol{C}^n = \emptyset$ for all $\Omega \in \mathcal{M}$.

Now, $d_1^\infty = \overline{d}_1 - 1$. Also, these functions decrease by at most 1 as g increases by 1. Therefore, $\overline{d}_v = \overline{d}_{v+1} \geq 0$ for some $v \in [1,n]$. This fact implies that $\overline{I}(\Omega) \neq \emptyset$ for all $\Omega \in \mathcal{M}$. The two conclusions about Ω contradict the existence of Ω_1 above.

4.8 Dimension and Slicing

In this section we prove that any element of $\mathcal{CI} \cap \mathcal{TI} \cap \mathcal{AF}$ satisfies the three properties of §4.1. First of all, Property 2 holds for any element of $\mathcal{TI}$, by definition. We will show in Lemma 4.9 that Property 1 implies Property 3, so that we only really need to establish Property 1.

Let T be a tube, as usual.

Lemma 4.8 *If $V \cap \Sigma_k T$ consists of n distinct points, and $V \cap \Sigma_{k-1} T = \emptyset$. Then* $\dim(V) \geq n - k$.

Proof: We may normalize by a complex affine map so that $T = qJ$, where $q = \{0, ..., 0\}$, and $J = \{e_1, e_2, ..., e_n\}$. Here e_i is the standard basis vector. Let $p_j = V \cap t_k(j)$. The hypotheses imply that (after suitably labelling), the jth and $(j + k - 1)$st coordinates of p_j are nonzero, for $j = 1, 2, ..., n - k + 1$. These $n - k + 1$ vectors are independent, and all contained in V. ♠

Lemma 4.9 *Suppose that W is a codimension d flat, slicing a tube T. Every hyperplane of W intersects $\Sigma_{d+1} T$, and $W \cap \Sigma_d T$ spans W.*

Proof: Let $L_1, L_3, ..., L_{2n-1}$ be the lines of $W \cap \Sigma_{d+1} T$. If the first statement of this lemma is false, then some hyperplane of W is disjoint from $\cup L_i$. This in turn implies that $\cup L_i \subset W'$, for hyperplane W' of W. By transversality, and the distinctness of the lines L_i, there is a hyperplane $W'' \subset S'$ such that the intersections $q_j = W'' \cap L_j$ are all distinct, and also disjoint from Σ_d. Lemma 4.8 says that $\dim(W'') \geq n - d - 1$, which is a contradiction.

The lines L_j are each determined by pairs of points in $W \cap \Sigma_d T$. If the second statement of this lemma is false, then the lines L_j are contained in a hyperplane of W, which contradicts the first statement of the lemma. ♠

By construction, $H_{1,k}$ is in general position with respect to $\Sigma_* T_z$, so $H_{1,k}$ slices T_z whenever $k \hookleftarrow z$. Suppose, by induction, that Ω is d-sliced, for some $d < \underline{h}$. We now show that $H_{d+1,k}$ slices T_z, when $k \hookleftarrow z$. For ease of exposition, we assume that d and k are even.

Suppose that $\dim(H_{d+1,k}) = \dim(H_{d,k\pm 1})$. Consider the smallest d for which this equality holds. The above equality implies that $H_{d,k-1} = H_{d,k+1}$. By induction and the Mating Lemma, $A_{d,k-1} = A_{d,k+1}$. Since A_d is an affine chain, this is only possible if $\underline{h} < d$. Therefore, $H_{d+1,k}$ is a hyperplane of $H_{d,k\pm1}$.

By induction, $H_{d,k\pm1}$ slice T_k, since $(k\pm1) \hookleftarrow k$. Let $L_{\pm,1}, L_{\pm,3}, ..., L_{\pm,2n-1}$ be the n distinct lines of $H_{d,k\pm1} \cap \Sigma_{d+1} T_k$. Let $p_{\pm,1}, p_{\pm,3}, ..., p_{\pm,2n-1}$ be the n distinct points of $H_{d,k\pm1} \cap \Sigma_d T_k$. It follows from the Mating Lemma that

1. $\pi(\{p_{\pm,j}\})$ consists of the n-distinct points of the chain $A_{d,k\pm1}$.

2. $\pi(\{L_{\pm,j}\})$ consists of the n distinct lines of $\delta_1(A_{d,k\pm1})$.

3. $\pi(L_{-,j}) \cap \pi(L_{+,j})$ is the jth point of $A_{d+1,k}$.

Let $I_j = H_{d+1,k} \cap \Sigma^j_{d+1} T_k$. We have

$$I_j = (H_{d,k-1} \cap \Sigma^j_{d+1} T) \cap (H_{d,k+1} \cap \Sigma^j_{d+1} T) = L_{-,j} \cap L_{+,j}.$$

By Lemma 4.9, we see that $H_{d+1,k} \cap \Sigma_{d+1} T_k \neq \emptyset$. Hence, $\dim(I_j) \geq 0$ for some j. Since the mating process is defined on the chain A_d, and A_d is affine, the lines $L_{-,j}$ and $L_{+,j}$ project to distinct lines in the plane. Therefore, these lines are distinct in $\boldsymbol{C}^n$. Hence, $\dim(I_j) \leq 0$. In short I_j is a point, for some j. Since $\Omega \in \mathcal{CI}$, this fact is true for all j. Since the chain A_{d+1} is affine, the points I_j project to distinct points in the plane, and hence are distinct in $\boldsymbol{C}^n$. If $d = n-2$, we are done.

Suppose $d \leq n-3$. Let $Q_i = H_{d+1,k} \cap \Sigma^i_{d+2} T_k$. Each Q_i contains points $I_{i\pm1}$. Hence $\dim(Q_i) = q_i \geq 1$. Suppose that $q_i \geq 2$. Define

$$V_1 = \Sigma^{i+1}_{d+1} T_k; \quad V_2 = Q_i; \quad W = \Sigma^i_{d+2} T_k.$$

Since

$$Q_i \cap \Sigma^{i+1}_{d+1} T_k = (H_{d+1,k} \cap \Sigma^i_{d+2} T_k) \cap \Sigma^{i+1}_{d+1} T_k = H_{d+1,k} \cap \Sigma^{i+1}_{d+1} T_k = I_{i+1},$$

we have $V_1 \cap V_2 = I_{i+1}$. In particular, $\dim(V_1 \cap V_2) = 0$. This contradicts the fact that $V_1, V_2 \subset W$, and $\dim(V_1) + \dim(V_2) \geq \dim(W) + 1$. Hence, $q_i = 1$. The lines $Q_2, Q_4, ..., Q_{2n-2}$ project to the lines of $\delta_1(A_{d+1,k})$. Since A_{d+1} is affine these lines are all distinct in the plane. Therefore, $Q_1, Q_3, ..., Q_{2n-1}$ are all distinct. Putting everything together, we see that $H_{d+1,k+1}$ slices T_k.

5 Proof of Lemma 4.6

5.1 Abstract Formulation

Recall that $\pi(z_1, ..., z_n) = (z_1, z_2)$ and that $\psi(z_1, ..., z_n) = (z_3, ..., z_n)$. If $Z = \phi^{-1}(0,0)$ then $A = \overline{Z} - Z$, where $\overline{Z}$ is the closure of Z in $\boldsymbol{CP}^n$.

Given a point $a = (a_1, ..., a_n)$ and a sequence of vectors $Q = \{q_1, ..., q_n\} \subset \boldsymbol{C}^k$, define $a \cdot Q = \sum_{\alpha=1}^n a_\alpha q_\alpha$. Suppose that $C \subset \boldsymbol{C}^n$ is an n-Point. Let J_C be the space of n-Joints J such that $\pi(J) = C$. Given $J \in J_C$, we define the linear map $\Psi_J : \boldsymbol{C}^n \to \boldsymbol{C}^{n-2}$ by the formula $\Psi_J(v) = v \cdot \psi(J)$.

Let H be a hyperplane. We say that a flat F is *adapted to the pair* (J, H) if $\Psi_J(F) = \psi(H \cap Z)$. Suppose that $H()$ is a *hyperplane function*. That is $H(J)$ is a hyperplane, for each $J \in J_C$. Say that a flat F is *adapted to the pair* $(C, H())$ if F is adapted to the pair $(J, H(J))$ for a generic choice of $J \in J_C$.

Given a hyperplane $V \subset \boldsymbol{C}^n$, let V' be the linear subspace which is parallel to V. We say that the collection of hyperplane functions $\{H_1(), ..., H_n()\}$ is *ample* if there exist a collection of codimension 3 flats $\{F_1, ..., F_n\}$ and a collection of hyperplanes $\{V_1, ..., V_n\}$ such that $V_1', ..., V_n'$ are in general position, and for all j $F_j \subset V_j$ and F_j is adapted to the pair $(C, H())$.

Lemma 5.1 *Suppose that $\{H_1(), ..., H_n()\}$ is an ample collection of adapted hyperplane functions. Then $A \cap (\cap \overline{H}_j(J)) = \emptyset$ for generic $J \in J_C$.*

Here is how we will use Lemma 5.1. After constructing a particular PolyChain Γ, we define, for each $J_1 \in J_{C_1}$, a canonical PolyJoint $\Omega(\Gamma, J_1) = \{J_1, J_3, ..., J_{2N-1}\}$ which covers Γ. Next, for $j = 1, ..., n$ we define hyperplane functions $\tilde{H}_j(J_1) = |J_{1+2(j-1)p}|$. Finally, we unravel the definitions to prove that these hyperplane functions are ample. Lemma 5.1 then says that $\Omega(\Gamma, J_1)$ satisfies the conclusion of Lemma 4.6 for generic choice of $J_1 \in J_C$, and all we really need is one example.

The rest of this section is devoted to the proof of Lemma 5.1.

Sub-Lemma 5.2 *Let $W_1, ..., W_n$ be general position codimension one linear subspaces. For $j = 1, ..., n$, let $G_j \subset W_j$ be linear subspace which has codimension 2 in W_j. Then the generic linear map $\Psi : \boldsymbol{C}^n \to \boldsymbol{C}^{n-2}$ has the property that $\Psi(G_1) \cap ... \cap \Psi(G_n) = \{0\}$*

Proof: It suffices, by algebraicity, to find a single linear map which has the desired property. Equip $\boldsymbol{C}^n$ with the usual Hermitian inner product. Our conditions imply that 3 dimensional subspaces $Q_j = \{G_j^{\perp}\}$ are such that $\text{span}(Q_1, ..., Q_k)$ has dimension at least k. We will show, inductively, that there is a dimension $d-2$ linear subspace R_d, such that $\{R_d \cap Q_1, ..., R_d \cap Q_d\}$ spans R_d. Certainly, this is true for $d = 2, 3$. Suppose it is true for $d = k$. Since R_k has dimension $k-2$, and the span of $Q_1, ..., Q_{k+1}$ has dimension at least $k+1$, there is some $j \in [1, k+1]$ such that $Q_j \not\subset R_k$. Let $v \in Q_j - R_k$. If we take R_{k+1} to be the span of R_k and v, then R_{k+1} has the desired properties. By induction, $R = R_n$ exists. It follows from duality that orthogonal projection onto R has the desired properties. ♠

The collection of vectors in $\psi(J)$ can be chosen generically. Hence, the generic projection $\boldsymbol{C}^n \to \boldsymbol{C}^{n-2}$ arises as a map of the form Ψ_J. Let F_j' be the linear subspace parallel to F_j. Note that $F_j' \subset V_j'$. SubLemma 5.2 and the hypotheses imply that $\Psi_J(F_1') \cap ... \cap \Psi_J(F_n') = \{0\}$ for the generic choice of $J \in J_C$. Hence, there is no nontrivial flat in $\boldsymbol{C}^{n-2}$ which is simultaneously parallel to all of $\Psi_J(F_j)$. Since $\psi(Z \cap H_j(J)) = \Psi_J(F_j)$ there there is no flat in $\boldsymbol{C}^{n-2}$ which simultaneously parallel to all of $\psi(Z \cap H_j(J))$. Such a flat would exist, however, if $A \cap (\cap \overline{H}_j(J)) \neq \emptyset$.

5.2 Separation and General Position

Say that a vector $(a_1, ..., a_n)$ is *separated* if $|a_1| > 1$ and if $|a_j| > n! \| a_1^n ... a_{j-1}^n|$ for all j. Say that a vector $\Lambda = (\lambda_1, \lambda_3, ..., \lambda_{2n-1})$ is *inverse separated* if the associated sequence $(a_1, ..., a_n)$ is separated. Here $a_j = \lambda_{2j-1}^{-1}$. (The reason for the strange labelling of Λ will be clear soon.)

Define the linear map $T_{\Lambda,q}(e_j) = a_j e_{j+q}$. Here e_j is the standard basis vector. Let $W_0 = \{x_1, ..., x_n | \sum x_i = 0\}$. For $j \geq 1$ define $W_j = T_{\Lambda,q}^{j-1}(W_0)$.

Lemma 5.3 *If Λ is inverse separated then $W_1, ..., W_n$ are in general position.*

Proof: Given vectors X and Y, let XY be the vector obtained by componentwise multiplication. Taking indices mod n, let $V_h = (a_{h+1}, ..., a_{h+n})$. Let $W_{q,1} = (1, ..., 1)$. Inductively, define $W_{q,k} = W_{q,k-1} V_{kq}$. To prove this lemma, it suffices to show that the vectors $W_{q,1}, ..., W_{q,n}$ are linearly independent. Let $P_{n,q}(a_1, ..., a_n) = \det(M_{n,q})$, where $M_{n,q}$ is the matrix whose columns are these vectors. It suffices to prove that $X = P_{n,q}(a_1, ..., a_n) \neq 0$.

Let $\omega = \exp(2\pi i/n)$. We have $P_{n,q}(\omega, ..., \omega^n) = \omega^b P_{n,0}(\omega, ..., \omega^n)$. Here b depends on n and q. Now, $P_{n,0}$ is the usual Vandermonde determinant: $P_{n,0}(\omega, ..., \omega^n) = \prod_{i<j}(\omega^i - \omega^j) \neq 0$. Hence, $P_{n,q}$ is not the zero polynomial. The coefficients of the monomials in X are all ± 1. The separation hypothesis says that one of these monomials has larger norm than the sum of all the other norms. ♠

5.3 Separated PolyChains

We will use the notation of §2. Suppose that $\Gamma = (C_1, C_3, ..., C_{2N-1})$ is a PolyChain, (p, q) derived from an element of $L(n, p)$. Let c_{ij} be the points of C_i. Also, let

$$x_{2k} = \overline{c_{2k-1,i}c_{2k+1,i}} \cap \overline{c_{2k-1,i}c_{2k+1,i}}.$$

This point is independent of i. The points $c_{2k-1,i}$, x_{2k}, and $c_{2k+1,i}$ all lie in the same complex line, and hence the following is well defined:

$$\beta_{2k,i} = \frac{c_{2k+1,i} - x_{2k}}{c_{2k-1,i} - x_{2k}}; \qquad \lambda_i = \prod_{k=1}^{p} \beta_{2k,i}.$$

We call $\Lambda = \{\lambda_1, \lambda_3, ..., \lambda_{2n-1}\}$ the *return sequence* for Γ. We say that Γ is *separated* if Λ is inverse separated, as defined above.

Lemma 5.4 *Suppose (n, p, q) are as in Theorem 1.1. Then some element of $\mathcal{B}(N; n, p, q)$ is separated.*

Proof: Let $\omega = \exp(2\pi i/n)$. For any complex number z, let $\underline{z}$ be the point in $\boldsymbol{R}^2$ which represents z. Choose generic points $x_{2p-2}, x_{2p} \notin U = \cup\overline{0\underline{\omega}^j}$. Let $2 > r_1 > r_3 > ... > r_{2p-5} = 1$. Let $s_1, s_3, ..., s_{2n-1} < 1$ be a sequence of small numbers. Let ρ_q be the cyclic permutation: $\rho_q(z_1, z_2, ..., z_n) = (z_{q+1}, z_{q+2}, ..., z_n, z_1, ..., z_q)$. (Our notation is a bit redundant when $p = 3, 4$.) As shown in figure 5, define

$$C_i = \{r_i\underline{\omega}, ..., r_i\underline{\omega}^n\}; \qquad i = 1, 2, ..., 2p-5.$$

$$C_{2p-3} = \{s_1\underline{\omega}, s_3\underline{\omega}^2, s_5\underline{\omega}^3, ..., s_{2n-1}\underline{\omega}^n\}.$$

$$C_{2p-1} = (c_{2p-1,1}, c_{2p-1,3}, ..., c_{2p-1,2n-1}); \qquad c_{2p-1,j} = \overline{x_{2p-2}c_{2p-3,j}} \cap \overline{x_{2p}c_{1,j+2q}}.$$

$$C_{2p+j} = \rho^q(C_j) \qquad j = 1, 3, ..., 2N-2p-1$$

By construction, $\{C_1, C_3, ..., C_{2N-1}\}$ is a PolyChain which is (p,q)-derived from an element of $L(n,p)$.

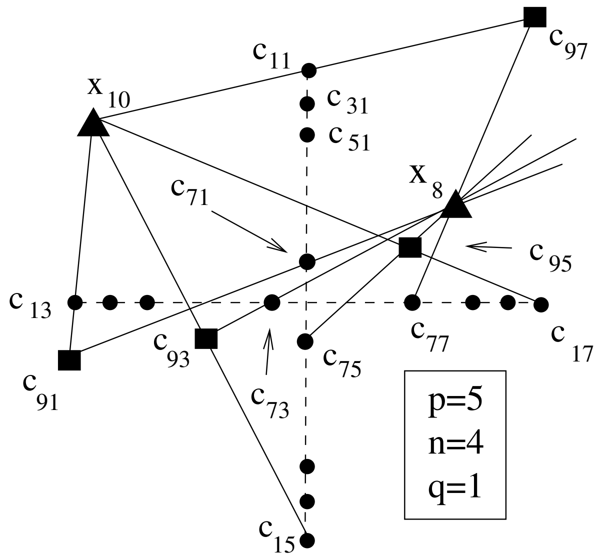

Figure 5

To estimate the return sequence note that

$$\beta_{2,j}, \beta_{4,j}, ..., \beta_{2p-6,j} \in [1,2]; \qquad \beta_{2p-4} = s_j.$$

As $s_1, ..., s_n \to 0$, the points of C_{2p-1} converge to U. Since $x_{2p-2}, x_{2p} \notin U$, there are constants $\epsilon, K > 0$ such that

$$\max(s_j) < \epsilon \quad \Rightarrow \quad \beta_{2p-2,j}, \beta_{2p,j} \in [1/K, K],$$

All in all, $\lambda_j = K_j s_j$, where K_j lies in a compact subset of $\boldsymbol{R}^+$ which is independent of the choice of $\{s_j\}$, and we can prescribe this sequence as we like.

5.4 Canonical Lifts

Let Γ be any separated PolyChain, as constructed above. Perturbing if necessary, we can assume that Γ is coverable. Recall from the Lifting Lemma that a PolyJoint lift of a coverable chain Γ can be specified by choosing a joint J_1 such that $\pi(J_1) = C_1$, and points $\tilde{x}_2, \tilde{x}_4, ..., \tilde{x}_{2N-2}$ such that $\pi\tilde{x}_j = x_j$. Let J_1 be a joint. The J_1-*canonical lift* of Γ is the object $\Omega(\Gamma, J_1)$ whose lifting data is given by $(J_1; x_2, x_4, ..., x_{2N-2})$.

Lemma 5.5 *For generic $J_1 \in J_C$, the object $\Omega(\Gamma, J_1)$ is a PolyTube and belongs to $\mathcal{E}(N; n, p, q)$.*

Proof: By induction, there is a generic set of choices for J_3 so that the J_3-canonical lift $(C_3, C_3, ..., C_{2N-1})$ is a $(N-1)$-Joint. By transversality, there is a generic set of choices of J_3 so that $x_2 \notin |J_3|$. Finally, J_3 depends birationally on J_1. Hence, there is a generic set of choices of J_1 which guarantee that J_3 has the above two properties. For such choices of J_1, the J_1-canonical lift of $(C_1, C_3, ..., C_{2N-1})$ will be an N-Joint, and hence will belong to $\mathcal{E}$. (Note that we are using that fact that $x_2 \notin J_3$ implies $x_2 \notin J_1$.) ♠

We write $\Omega = \Omega(J_1, \Gamma)$, and as in §3 write $\Omega = \{J_1, J_3, ..., J_{2N-1}\}$. As in §3, let $H_{1,j} = |J_j|$, as usual. The functions $H_{1,j}$ are all hyperplane functions of J_1. We are mainly interested in $\tilde{H}_j() = H_{1+2(j-1)p}()$. All that remains is to prove

Lemma 5.6 *The collection $\{\tilde{H}_1(), ..., \tilde{H}_n()\}$ is ample.*

Proof: We will use the notation $F \cdot Q = \{a \cdot q |\ a \in F\}$ for any flat F.

Define $V_0 = \{(x_1, ..., x_n) | \sum x_i = 1\}$. Let F_1 be the maximal subflat of V_0 such that $F_1 \cdot C_1 = \{(0\}$. Note that $V_0 \cdot J_1 = \tilde{H}_1$. Since the points of C span $\boldsymbol{C}^2$, we see that $\mathrm{codim}(F_1) = 3$. Hence $F_1 \cdot J_1 = Z \cap \tilde{H}_1$. In particular, F_1 is adapted to the pair $(\tilde{H}_1(), C)$.

Define $F_j = T^{j-1}_{\Lambda,q}(F_1)$. Also define $V_j = T^{j-1}_{\Lambda,q}(V_1)$ and $W_j = V_j'$. Obviously, F_j jas codimension 3 and V_j has codimension 1. It follows from Lemma 5.3 that $W_1, ..., W_n$ are in general position. To finish the entire proof, all we need to show is that F_j is adapted to the pair $(C_1, \tilde{H}_j())$. We will give a proof for $j = 2$. The general case follows from the iteration of our argument.

Let c_{ij} be the ith point of C_j. Let $\tilde{c}_{ij}$ be the ith point of J_j. Let $\hat{c}_{ij} = \psi(\tilde{c}_{ij})$. By construction $\pi(\tilde{c}_{ij}) = c_{ij}$. Recall that $Z = \pi^{-1}(0,0)$. Z has codimension 2. We also use the constants β_{ij} defined in the preceding section.

First, we claim that $\hat{c}_{2k+1,i} = \beta_{2k,i}\ \hat{c}_{2k-1,i}$. None of the above values changes if we translate $\boldsymbol{C}^n$ in such a way as to preserve $\boldsymbol{C}^2$. Hence, we may assume $x_{2k} = (0,0)$. In this case, the complex linear map of $\boldsymbol{C}^n$ defined by $S(x) = \beta x$ takes $c_{2k-1,i}$ to $c_{2k+1,i}$. Since S preserves the lines through the origin, $S(\tilde{c}_{2k-1,i}) = \tilde{c}_{2k+1,i}$. Since S commutes with ψ, we establish our claim. Iterating, and using the fact that Γ is (p,q)-derived from an element of $L(n,p)$, we have

$$(*) \quad \hat{c}_{1+2p,i} = \lambda_i \hat{c}_{1,i}; \quad c_{1+2p,i} = c_{1,i+2q}; \quad i = 1,3,...,2n-1$$

Given the standard basis $\{e_j\}$ of $\boldsymbol{C}^n$, define maps $R(e_i) = e_{i+2q}$ and $D(e_i) = \lambda_{i+2j} e_i$. We have

$$(**) \quad \Psi_{J_1}(F_2) = \psi(F_2 \cdot J_1) = \psi(D(R(F_1)) \cdot J_1) = \psi(R(F_1) \cdot J_{1+2p}).$$

The only nontrivial equality is the last one, which follows from $(*)$. Since $R(V_0) = V_0$, we have

$$R(V_0) \cdot J_{1+2p} = V_0 \cdot J_{2p+1} = H_{1+2p} = \tilde{H}_2.$$

From equation $(*)$, we have $R(F_1) \cdot J_{2p+1} = Z$. A dimension count now implies that $R(F_1) \cdot J_{1+2p} = Z \cap \tilde{H}_2$. Combining this with $(**)$ shows that $\Psi_{J_1}(F_2) = \psi(Z \cap H_2)$, as desired. ♠

6 References

[**K**] K. Kendig, *Elementary Algebraic Geometry*, Springer-Verlag G.T.M. **44**, 1977.

[**S**] R. Schwartz, *The Pentagram Map*, J. Experimental Math., 1992.