Proving Disjointness with Foliated Patches

Richard Evan Schwartz *

February 4, 2000

1 Introduction

Let X be an n-dimensional homogeneous space, with an analytic Lie group action $G : X \to X$. We will assume that $X \subset \boldsymbol{R}^N$ for some $N \geq n$. The Euclidean metric on $\boldsymbol{R}^N$ induces a metric on X, but typically G does not act isometrically in this (or any) metric. The purpose of this paper is to describe an algorithm which sometimes helps prove that two sets in X are disjoint, by making only finitely many computations in $\boldsymbol{R}^N$. In special cases, this algorithm is entirely practical; it is meant to be implemented on a computer.

Let us narrow our task a bit. We say that two subsets $S, T \subset X$ are *gently separated* if there is a function $f : X \to \boldsymbol{R}$ which is contracting on $S \cup T$, such that $\sup_S f < \inf_T f$. More generally, we say that a *piecewise gentle separation* of S and T is a decomposition $S = S_1 \cup ... \cup S_n$ and $T = T_1 \cup ... \cup T_m$, such that S_i is gently separated from T_j, for each pair (i, j). Obviously, if S and T have a piecewise gentle separation, they are disjoint.

This paper does not have anything to say about finding piecewise gentle separations for S and T. In low dimensions, such an object probably would be found by looking at a lot of computer plots. At any rate, we imagine that the *search* for a piecewise gentle separation would be done experimentally, perhaps without any regard for rigor. The work here is on the other side of the fence. We will suppose that a *suspected* piecewise gentle separation for S and T has been found. Our main task is to prove that the given object is actually a piecewise gentle separation. This amounts to proving (finitely many times) that certain functions are positive on certain subsets of X.

*Supported by N.S.F. Research Grant

To put this paper in context, suppose that $\Gamma \subset G$ is a countable subgroup which seems to act properly discontinuously on an open subset $\Delta \subset X$, with some subset $F \subset \Delta$ as a fundamental polyhedron. We think of F has having a cellular subdivision, in which the individual cells are, say, pieces of algebraic varieties. A big step in proving that F is a fundamental domain for Γ is showing that various cells of F, together with Γ-translates of these cells, are disjoint from each other.

If X is a fairly exotic space, the equations for the simplest useful building blocks for fundamental polyhedra might be very complicated. It might be impossible to determine *algebraically* if a given function is positive on a set. To give a computational proof, we could parametrize our subset, obtain Lipschitz bounds on the parametrization, and evaluate the function on sufficiently many points. However, this approach might lead to an impractical computation.

For us, the motivating example is that of $PU(2,1)$ acting on S^3, the ideal boundary of the complex hyperbolic plane. The basic sets with which we work are *hybrid cones*. We developed hybrid cones to solve the *Goldman-Parker conjecture*, a question [**GP**] about triangle groups acting on the complex hyperbolic plane, and have already had to perform extensive numerical analysis on them. See [**S1**].

We found that the straightforward methods for dealing with hybrid cones lead to computations which would take millions of years to complete on today's computers. We had to use methods which, while much more complicated, lead to feasible calculations. We think that hybrid cones will find wide application in complex hyperbolic geometry, and it seems useful to have an efficient and robust general framework for treating them numerically. The method here is more natural, and also faster computationally, than the rather *ad hoc* method used in [**S1**].

The algorithm we present here, which we call *the method of foliated patches*, takes into account the geometry of X, and obtains *a posteriori* estimates on the fly, removing the need for *a priori* bounds on the parametrizations. The cost of this is that our algorithm requires some differential geometric information about specific curve families in X. In the special cases, such information can be obtained, once and for all, without too much trouble.

Our point of view is that one can construct useful sets in X by starting with some "geometrically simple" examples and considering all G-translates of them. The geometry of the G-translates might be much less tractible than that of the original examples. However, we will construct convenient "measuring devices", which we call *foliated patch clusters*, for the original examples. The foliated patch clusters, which are adapted to the action of G, are transported by the G-action to foliated patch clusters for the translates. These new foliated patch clusters, which are still "geometrically simple", serve as useful measuring devices for the more complicated translates. We will make these terms precise in the body of paper.

We will describe our algorithms under the assumption that we have a perfectly accurate computing maching at our disposal. In practice, one would have to use interval arithmetic, or some other guarantee of sufficiently high precision, when actually implementing our algorithm.

In §2 we will describe the method of foliated patches in general. While reading §2, the reader should be aware that we are not presenting an algorithm which is guaranteed to work. All of the elements described in §2 must combine, in the exactly the right way, to lead to a proof. It may appear that such a collection of coincidences will never occur in practice. This is not so. At the end of §2 we will give a heuristic argument that our method is actually bound to succeed, given a sensible design. In §3 we adapt the general framework to the motivating example discussed above. The background material can be found, in more detail, in [**E**], [**G**], [**S1**], or [**S2**].

I would like to thank Todd Drumm, Bill Goldman, and Justin Wyss-Gallifent for encouraging me to write this paper. I would like to to thank my wife, Brienne Brown, for computer help, and Jeong-Hwan Jang for helpful conversations about curvature.

2 The General Framework

We continue the notation from the introduction. Let $f : X \to \boldsymbol{R}$ be a function which is contracting on $S \subset X$. We suppose that $S \subset \Phi(Q_0)$, where Q_0 is the unit k-cube, and $\Phi : Q_0 \to X$ is a continuous map. We will explain how to prove that $f|_S > 0$, using only finitely many computations.

An alternate situation is that f is contracting in a region which we suspect contains S, and part of our proof amounts to showing that S is contained in this region. The methods we explain here sometimes apply to a case such as this though, for the sake of brevity, we will not go into details.

2.1 The Subdivision Algorithm

A *k-rectangle* is a product of the form $I_1 \times ... \times I_k$, where each of the I_j are intervals. Let I_j^1 and I_j^2 be the left and right halves of I_j. If Q is a k-rectangle, let $c(Q)$ be the center of Q and $v(Q)$ be the vertex set of Q. Also, for $j = 1, ..., k$ and $i = 1, 2$ we define

$$Q_{ji} = I_1 \times ... \times I_{j-1} \times I_j^i \times I_{j+1} \times ... \times I_k.$$

Let $\mathcal{Q}$ denote the set of k-rectangles which are contained in $Q_0 = [0,1]^k$. Suppose that $\boldsymbol{A} : \mathcal{Q} \to \{0, 1, ..., k\}$ is some function. We think of 0 as the *true state* and the values $1, ..., k$ as *false states*. A *partition* of Q_0 is a subdivision $Q_0 = Q_1 \cup ... \cup Q_m$, where the interiors of the smaller rectangles are pairwise disjoint. We say that such a partition is *true* if $\boldsymbol{A}(Q_j) = 0$ for $j = 1, ..., m$.

Here is an algorithm which searches for a true partition:

1. Let U be a list of elements of $\mathcal{Q}$, having Q_0 as its only member. Let V be the empty list.

2. Let Q be the last member of U. If $\boldsymbol{A}(Q) = 0$ transfer Q from U to V. If U is now empty, halt. If U is not now empty then repeat this step.

3. If $\boldsymbol{A}(Q) = j > 0$ then delete Q from U and append to U the two rectangles Q_{j1} and Q_{j2}. Go to Step 2.

If the algorithm halts, the list of rectangles in V is a true partition. We call this algorithm the *subdivision algorithm*. It is a depth-first search over the tree of dyadic subcubes of Q_0.

2.2 Foliated Patches

Our method is centered around a G-invariant, finite dimensional family $\mathcal{F}$ of closed analytic arcs in X. For instance, one could take $\mathcal{F}$ to consist of the arcs of the form $\bigcup_{t\in I} h_t(x)$, where $I \subset \boldsymbol{R}$ is a closed interval, and $\{h_t|\ t \in \boldsymbol{R}\}$ is a 1-parameter subgroup of G. As a simpler example, one could take $\mathcal{F}$ to be the union of all G-translates of some finite list of anaytic arcs. This is what we will do in §4. The purpose of choosing a family such as $\mathcal{F}$ is twofold:

1. The finite dimensionality condition means, in principle, that one could estimate the geometry of arcs in $\mathcal{F}$ by sampling a uniformly bounded number of points on each arc.

2. The arcs in $\mathcal{F}$ have analytic continuations into larger curves. In carrying out Item 1, one could possibly sample points not just on the arcs, but on the continuations.

We suppose that each arc $\gamma \in \mathcal{F}$ has a pre-chosen analytic continuation $\hat{\gamma}$. We define $\mathcal{F}(a,b)$ to be the set of triples of the form (γ, A, B). Here A is a collection of a ordered points chosen from $\gamma - \partial\gamma$ and B is a collection of b ordered points chosen from $\hat{\gamma} - \gamma$. In addition to these $a+b$ points, we wish to distinguish the endpoints of γ. Thus, elements of $\mathcal{F}(a,b)$ consist of curves $\hat{\gamma}$ with $\nu = a+b+2$ distinguished points. Associated to each $X \in \mathcal{F}(a,b)$ is the $\nu \times \nu$ symmetric matrix $M(X)$ encoding the Euclidean distances between these points. The differential geometric input of our method amounts to bounding the curvature and diameter of arcs in $\mathcal{F}$ based on the matrix described above. This is further described in §2.4

Say that an *n-dimensional $\mathcal{F}$-patch* is an embedding $\psi : Q \to X$, all of whose maximal coordinate curves are elements of $\mathcal{F}$. We will use the words *foliated patch* and *$\mathcal{F}$-patch* as synonyms. Here Q is an n-dimensional rectangle. Note the following property: If $g \in G$ and ψ is an $\mathcal{F}$-patch, then so is $g \circ \psi$.

If ψ is an $\mathcal{F}$-patch, we define $c(\psi) = \psi(c(Q))$ and $v(\psi) = \psi(v(Q))$. Note that $c(\psi)$ is a single point and $v(\psi)$ is a set of 2^n points. We say that ψ is *guaranteed* if

$$|Q| := \sup_{x\in\psi(Q)} \rho(c(\psi), x) = \max_{v\in v(\psi)} \rho(c(\psi), v).$$

We have the implication

$$f(c(\phi)) - |Q| > 0 \quad \Longrightarrow \quad f|_{\psi(Q)} > 0. \tag{1}$$

If Q is a k-rectangle for $k \leq n$, we define $Q\!\uparrow = Q \times [0,1]^{n-k}$. Let S and Φ be as above. Let $\mathcal{X}$ denote the set of finite lists of $\mathcal{F}$-patches. Given an element $A \in \mathcal{X}$, we will let A_1 be the first element of A. We think of A_1 as being the main $\mathcal{F}$-patch in A, and the other elements on the list as playing supporting roles. This point of view is discussed more fully in §2.6.

We say that a *patch map dominating* Φ is a map $\Psi : \mathcal{Q} \to \mathcal{X}$ such that, for all $Q \in \mathcal{Q}$,

$$\Phi(Q) \subset \psi(Q\!\uparrow); \qquad \psi = (\Psi(Q))_1.$$

Here $Q\!\uparrow$ is the domain of ψ.

We say that a finite computer program $\boldsymbol{G}$ is a *guarantee test* for Ψ if there is an integer M such that $\boldsymbol{G}$ accepts as input any element of $A \in \mathcal{X}$ and returns, after at most M basic computations, an integer $\boldsymbol{G}(A) \in \{0, 1, ..., k\}$ such that $\boldsymbol{G}(A) = 0$ only if X_1 is guaranteed. A return of values greater than 0 indicates how the computer thinks A_1 should be subdivided so as to most increase the chances that the subdivided pieces will be guaranteed. We will not discuss how $\boldsymbol{G}$ selects its nonzero values. The basic idea is to arrange the subdivision so that the images of rectangles, under the patch map, remain fairly cube-like. In practice, the best method would probably be found by trial and error.

To prove that f is positive on S, we define a function $\boldsymbol{A}$ as follows: For each $Q \in \mathcal{Q}$, we have $\boldsymbol{A}(Q) = 0$ iff $\boldsymbol{G} \circ \Psi(Q) = 0$ and the left hand of Equation 1 is true. Once we know that $(\Psi(Q))_1$ is guaranteed, the left hand side of Equation 1 can be checked in finitely many computations. We run the subdivision algorithm, with respect to $\boldsymbol{A}$. If this algorithm halts, the information constitutes a proof that f is positive on S.

The computer program $\boldsymbol{G}$ has three ingredients

1. Curvature and diameter guarantees for $\mathcal{F}$

2. The Straight Lemma.

3. Foliated patch clusters.

We will explain these ingredients in the next three sections. Following this, we will explain the operation of $\boldsymbol{G}$.

2.3 Diameter Guarantees

For any arc γ, define

$$\lambda(\gamma) = \sup_{x \in \gamma} \frac{\|x - \partial\gamma\|}{\|\partial_1\gamma - \partial_2\gamma\|}.$$

Here $\partial_1\gamma$ and $\partial_2\gamma$ are the endpoints of γ. A *diameter guarantee* is a map $\Delta : \mathcal{F}(a,b) \to [1,\infty]$ such that $\lambda(F(A)) \leq \Delta(A)$. Here $F : \mathcal{F}(a,b) \to \mathcal{F}$ is the forgetful map.

Here is a simple example. Let $A \in \mathcal{F}(1,0)$ and $\gamma = F(A)$, so that $A = (\gamma, p)$, where $p \in \hat{\gamma} - \gamma$. We define $\Delta(A) = \sqrt{2}/2$ if p is further from either endpoint of γ than these endpoints are from each other. In this case, γ is contained in a semicircle. Otherwise, we simply make no conclusion, and set $\Delta(A) = \infty$.

2.4 Curvature Guarantees

Let M_ν denote the set of $\nu \times \nu$ nonnegative symmetric matrices. We have a map $M : \mathcal{F}(a,b) \to M_\nu$, obtained by evaluating the distances between distinguished points. Here $\nu = a + b + 2$. Recall that the space curvature of a curve $\gamma \subset \boldsymbol{R}^N$ is the quantity $\|\gamma''(s)\|$, where $s \to \gamma(s)$ is a unit speed parameterization. We define $\kappa(\gamma)$ to be the the maximum space curvature of γ. We say that a *curvature guarantee* is a map $K : M_\nu \to [0,\infty]$ with the following property: for any $A \in \mathcal{F}(a,b)$ we have $\kappa(F(A)) \leq K \circ M(A)$.

We will actually use a slightly fancier object than K. Say that a *block neighborhood* of a matrix $M \in M_\nu$ is any set of matrices having the form

$$N_M = \{M' \in M_\nu : \ |M'_{ij} - M_{ij}| < \epsilon_{ij}\}.$$

Let BM_ν denote the set of block neighborhoods of elements of M_ν. We say that a *block version* of K is a map $\hat{K} : BM_\nu \to [0,\infty]$ such that

$$\sup_{M' \in N_M} K(M') \leq \hat{K}(N_M).$$

$\hat{K}$ can obtained from K if one knows how K varies over block neighborhoods.

Continuing with our simple example, $M(A)$ is a 3×3 matrix. One can interpret $M(A)$ as a Euclidean triangle, having side lengths E_1, E_2, E_3. By convention, we take E_1 to be the distance between the endpoints of γ. Let $\mathcal{F}$ be the family of circular arcs in $X = \boldsymbol{R}^N$. We define $K(M) = 2/s$, where $s = \max(E_1, E_2, E_3)$. One can easily construct a block version in this case.

2.5 The Straight Lemma

Let $C \subset X$ be a compact connected set. Given a continuous family $\Gamma = \{\gamma_t | \ t \in C\}$ of arcs in $\mathcal{F}$, we define $\kappa(\Gamma) = \sup_{t \in C} \kappa(\gamma_t)$. Note that $t \to \kappa(\gamma_t)$ need not be a continuous function.

Given a ball B, and an arc γ we say that γ is *straight enough for B* if $\partial\gamma \subset B$ and

$$\kappa(\gamma)\rho(B) < 1; \qquad 2\lambda(\gamma) < \frac{1}{\rho(B)\kappa(\gamma)} - 1.$$

Here $\rho(B)$ is the radius of B. This is a scale-invariant notion.

Suppose that $\phi : Q \to X$ is a k-dimensional $\mathcal{F}$-patch. We say that two $(k-1)$-dimensional $\mathcal{F}$-patches ϕ_1 and ϕ_2 *sandwich* ϕ if there are two parallel $(k-1)$-faces $Q_1, Q_2 \in \partial Q$ such that $\phi_j = \phi|_{Q_j}$. In this case, we say that *arcs connecting the sandwich* are the maximal coordinate arcs of ϕ whose endpoints are contained in $\phi_1(Q_1) \cup \phi_2(Q_2)$.

Inductively, we say that a k-dimensional $\mathcal{F}$-patch ϕ is straight enough for B if

1. ϕ has a sandwich $(\phi_1, \ \phi_2)$, such that ϕ_j is straight enough for B, for $j = 1, 2$.

2. If Γ is the family of connecting arcs for the sandwich then $\kappa(\Gamma)\rho(B) < 1$.

3. Some arc $\gamma \in \Gamma$ is straight enough for B.

Lemma 2.1 (Straight) *If ϕ is an $\mathcal{F}$-patch which is straight enough for B then $\phi \subset B$.*

Proof: We translate and scale so that B is the unit ball. Thus $\rho(B) = 1$. Our proof goes by induction. First, we claim that $\gamma \subset B$ if γ is an arc which is straight enough for B. Let $R = 1/\kappa(\gamma)$. Let $t \to \gamma(t)$ be a parametrization of γ. Here $t \in [0, 1]$. Consider the function $f(t) = \|\gamma(t)\|$. If f has a local maximum, at some $t \in (0, 1)$, then the curvature of γ at $f(t)$ is at least $1/f(t)$. Hence, $f(t) \geq R$. Since $f(0), f(1) \leq 1$, we have $\lambda(\gamma) > (R-1)/2$. This a contradiction. Hence, f does not have a maximum at any point in $(0, 1)$. In short, $f(t) \leq 1$ for all $t \in [0, 1]$. Thus, $\gamma \subset B$, as claimed.

For the induction step, we have a continuous family $\Gamma = \{\gamma_t | \ t \in C\}$. such that $\kappa(\Gamma) < 1$ and $\partial\gamma_s \subset B$ for all $s \in C$. Also $\gamma_t \subset B$ for some $t \in C$. We wish to prove that $\gamma_s \subset B$ for all $s \in C$.

We first make the following observation: Suppose $s \in C$, and $p \in \gamma_s$, and U is the unit vector tangent to γ at $\gamma(p)$. Let $s_n \in K$ be a sequence such that $s_n \to s$. Let $p_n \in \gamma_{s_n}$ be a sequence of points such that $p_n \to p$. Let U_n be the unit vector tangent to γ_{s_n} at $\gamma_{s_n}(p_n)$. Then $U_n \to \pm U$. This standard result follows from the uniform curvature bound we have on all arcs of Γ.

To complete our proof, it suffices to prove that $H = \{t \in C | \ \gamma_t \subset B\}$ is open, since H is clearly closed and nonempty and C is connected. Suppose that $\gamma_t \subset B$. The curvature bound implies that γ cannot be tangent to B at any point. There are three cases, either $\partial\gamma \cap \partial B$ is 0, 1, or 2 points. In the first case, γ_t is contained in the interior of B, and the result is obvious. The second case is like the third, so we will only consider the third case

In the third case, γ_t is transverse to ∂B at both endpoints. From the convergence of tangent vectors noted above, the part of γ_s near $\partial\gamma_s$ makes a definite angle with ∂B, once s is sufficiently close to t. Thus, there is a neighborhood N_s of $\partial\gamma_s$, such that $\gamma_s \cap N_s \subset B$, for s close to t. The curvature bound implies that the diameter of N_s can be taken independent of s, as long as s is sufficiently close to t. There is a positive lower bound on the distance from $\gamma_t - N_t$ to ∂B. By continuity, the same is true for $\gamma_s - N_s$, once s is sufficiently close to t. Thus, $\gamma_s - N_s \subset B$. In short, $\gamma_s \subset B$, and the set H is open. ♠

2.6 Foliated Patch Clusters

Here is an entertaining analogy which gives a good mental picture of the construction in this section. Imagine that you are the captain of a space ship, and you want to defend yourself against enemies. Some of your space ship is heavily shielded. However, there are some parts of your ship which are vulnerable to attack. So, you command several smaller spaceships to guard the vulnerable spots on your ship. Once again, these smaller ships are mainly shielded, but have some spots of vulnerability. Each smaller ship employs yet smaller ships to help it in the same way. This situation continues to a finite extent, the smallest ships in the fleet being too small to attack.

Our construction is based on $\mathcal{F}(a,b)$. The utility of the construction depends entirely on the space X, on the family $\mathcal{F}$, and on length and curvature guarantees associated to $\mathcal{F}(a,b)$. A general description of what features make for a useful computational tool is beyond us, though the heuristic argument in §2.8 offers some insights.

We define inductively the space $\mathcal{X}(m;a,b)$ of *m-dimensional standard $\mathcal{F}$-patch clusters*. Elements of $\mathcal{X}(0;a,b)$ are single points. For $k \geq 1$ an element of $\mathcal{X}(m;a,b)$ is a quintuple having the form $A = (\psi, a_1, a_2, \alpha, \beta)$, where

1. ψ is a m-dimensional $\mathcal{F}$-patch.

2. $a_1 = (\psi_1, ...)$ and $a_2 = (\psi_2, ...)$ are elements of $\mathcal{X}(m-1;a,b)$.

3. (ψ_1, ψ_2) sandwiches ψ.

4. α is a collection $\alpha_1, .., \alpha_a$, where each α_j is an element of $\mathcal{X}(m_j;a,b)$, and $m_j < m$.

5. β is a collection $\beta_1, .., \beta_b$, where each β_j is an element of $\mathcal{X}(n_j;a,b)$, and $n_j < m$.

6. If γ is a connecting arc of ψ, then $\gamma - \partial\gamma$ has nontrivial intersection with each α_i.

7. If γ is a connecting arc of ψ, then $\hat{\gamma} - \gamma$ has nontrivial intersection with each β_j.

By construction, $\mathcal{X}(1;a,b)$ and $\mathcal{F}(a,b)$ coincide.

We will call ψ the *main patch* of A. The other patches we call *auxilliary patches*. To say a bit more about the structure of A, the foliated patches in A are naturally organized into various binary trees. For instance, ψ is sandwiched by the foliated patches ψ_1 and ψ_2. Each ψ_j is in turn sandwiched by ψ_{j1} and ψ_{j2}. And so on. We call the union of all these patches the *main sandwich tree* in the patch cluster. The main patches in the α-clusters and the main patches in the β-clusters are at the pinnacles of *auxilliary sandwich trees*. And so on. All the patches in a sandwich tree are all contained in the single patch at the pinnacle of the tree. One can think of A as a tree of sandwich trees.

Our massive cluster collapses into a manageable number in low dimensional examples. Our purpose here is to explain one general *theoretical* approach for constructing the computer program $\boldsymbol{G}$.

2.7 The Guarantee Algorithm

We will describe a finite computer program $\boldsymbol{G}$, whose input is an element $A \in \mathcal{X}$ and whose output is 0 only if A_1 is guaranteed. We will automatically define $\boldsymbol{G}(A)$ to be nonzero if A does not belong to $\mathcal{X}(m;a,b)$, for some $m \leq n$. In practice, $\boldsymbol{G}$ will only be evaluated on elements of $\mathcal{X}(m;a,b)$, so the provision just made is a red herring. The operation of $\boldsymbol{G}$ depends on the existence of a diameter guarantee Δ and a block version $\hat{K}$ of the curvature guarantee K.

Let $A = (\psi, ...)$, as above. In particular, $A_1 = \psi$. Here $\psi : Q \to X$ is the main patch of the patch cluster. Let B_0 be the smallest ball centered at $c(\psi)$ and containing all of $v(\psi)$. Our computer program $\boldsymbol{G}$ can verify that $\psi(Q)$ is straight enough for any given ball B. If we run $\boldsymbol{G}$ for the ball $B = B_0$ then $\boldsymbol{G}$ can verify that ψ is guaranteed.

$\boldsymbol{G}$ operates in an inductive fashion. Suppose that $\boldsymbol{G}$ has already verified that every auxilliary patch in A is guaranteed and that every auxilliary patch in the main sandwich tree is straight enough for B. Let Γ be the family of connecting arcs for the sandwich (ψ_1, ψ_2). Using the Straight Lemma, $\boldsymbol{G}$ can complete its task by checking that $\kappa(\Gamma)\rho(B) < 1$ and by exhibiting an arc of Γ which is straight enough for B. In the second case, $\boldsymbol{G}$ can just pick some arc γ and use the diameter guarantee.

The first case is the interesting one. Let $B(\alpha_j)$ be the smallest ball centered at $c(\alpha_j)$ and containing all of $v(\alpha_j)$. We know, from the guarantees on all the auxilliary patches, that the image of α_j is contained in $B(\alpha_j)$. The same statement can be made for the β_j patches. Finally, the same statement can be made for the balls $B(\psi_j)$.

Here comes the punchline. The continuation $\hat{\gamma}$ of each arc $\gamma \in \Gamma$ intersects both $B(\alpha_i)$ and $B(\beta_j)$ for all indices i and j. Also, one point of $\partial\gamma$ is contained in $B(\psi_1)$ and the other point of $\partial\gamma$ is contained in $B(\psi_2)$. If all these balls are small, then we have tight control on a list of $a+b+2$ points on $\hat{\gamma}$. Applying our block version of the curvature guarantee, $\boldsymbol{G}$ verifies *simultaneously* that $\kappa(\gamma)\rho(B) < 1$ for all $\gamma \in \Gamma$. The point is that all these arcs intersect the same little balls, and so their associated matrices all lie in the same block neighborhood of a given matrix.

An obvious generalization of our algorithm is that we have a decomposition of $\mathcal{F}$ into several families of arcs $\mathcal{F}_1$,...,$\mathcal{F}_m$, and we have curvature and diameter guarantees for each subfamily.

2.8 Heuristic Argument

Here is an argument that $\boldsymbol{G}$ will *eventually* work, when combined with the subdivision algorithm, and a sensible design of a foliated patch map. Our argument is that each time $\boldsymbol{G}$ fails, and then reconsiders a patch bluster built around a smaller cube, the chance is better that it will succeed on the next attempt. Thus, long chains of failures are not possible.

We will design our patch map so that it is an *approximation* of the map $\Phi : Q_0 \to X$. That is, the main patch in the patch cluster $\Psi(Q)$ is small if Q is a small cube. Likewise for the auxilliary patches. On the other hand, we arrange that the centers of the auxilliary patches in the cluster $\Psi(Q')$ are roughly in the same places as the centers of the auxilliary patches in the cluster $\Psi(Q)$, once Q is small. Here Q' is a subdivision of Q.

As the subdivision algorithm proceeds without success, $\boldsymbol{G}$ looks at increasingly small block neighborhoods of essentially the same finite set of matrices. The curvature and diameter guarantees obtained by $\boldsymbol{G}$ during this long chain of failures does not change much, and perhaps even improves. On the other hand, the balls considered by $\boldsymbol{G}$ get vanishingly small during this losing streak, making the "straight enough" hypotheses increasingly easy to actually verify. Eventually, $\boldsymbol{G}$ succeeds and a little cube is taken off the list of cubes which need to be checked.

3 An Example

3.1 The Homogeneous Space

$\boldsymbol{C}^{2,1}$ is a copy of the vector space $\boldsymbol{C}^3$ equipped with the Hermitian form

$$\langle u, v\rangle = u_1\overline{v}_1 + u_2\overline{v}_2 - u_3\overline{v}_3 \tag{2}$$

The map

$$\Theta(v_1, v_2, v_3) = \left(\frac{v_1}{v_3}, \frac{v_2}{v_3}\right) \tag{3}$$

Takes the spaces

$$N_- = \{v \in \boldsymbol{C}^{2,1} | \ \langle v, v\rangle < 0\}; \qquad N_0 = \{v \in \boldsymbol{C}^{2,1} | \ \langle v, v\rangle = 0\} \tag{4}$$

respectively to the open unit ball and unit sphere S^3 in $\boldsymbol{C}^2$. As usual, $\boldsymbol{CH}^2$, the complex hyperbolic plane, is identified with the open unit ball. In this way, S^3, the space we are actually interested in, appears as the ideal boundary of $\boldsymbol{CH}^2$.

The group $G = PU(2,1)$ is the group of complex projective transformations which preserve S^3. These maps are precisely projectivizations of $\langle , \rangle$ preserving complex linear tranfrormtions. This group acts transitively on G. The distribution $\mathcal{E}$ of complex lines tangent to S^3 is invariant under G.

S^3 is a lie group itself. The group law is given by:

$$(z_1, w_1) \cdot (z_2, w_2) = (z_1 w_1 - z_2\overline{w}_2, z_1 w_2 + z_2\overline{w}_1). \tag{5}$$

The right multiplication map $R_X(Y) = Y \cdot X$ is an element of $PU(2,1)$ which acts isometrically on S^3, in the round metric, and moves all points by the same amount. The point here is that R_X commutes with the complex structure, which is left multiplication by i.

We call $\mathcal{H}= \boldsymbol{C} \times \boldsymbol{R}$ *Heisenberg space*. Given $p \in S^3$, a *Heisenberg stereographic projection from* p is a transformation $\boldsymbol{B} : S^3 - \{p\} \to \mathcal{H}$ of the form $\boldsymbol{B} = \pi \circ \beta$, where $\pi(z, w) = (z, \Im(w))$ and β is a complex projective transformation of $\boldsymbol{CP}^2$ which identifies $\boldsymbol{CH}^2$ with the *Siegal domain* $\{(z, w)| \ \Re(w) > |z|^2\}$. We write $\boldsymbol{B}(p) = \infty$ in this case.

The map

$$\boldsymbol{B}_0(z, w) = \left(\frac{w}{1+z}, \Im\frac{1-z}{1+z}\right) \tag{6}$$

is an example of Heisenberg stereographic projection from $(-1, 0)$.

3.2 Foliated Patches

A $\boldsymbol{C}$-*circle* (also known as a *chain*) is the intersection of a complex line with S^3, provided that this intersection is more than a point. $\boldsymbol{C}$-circles are all round circles, and G transitively permutes the $\boldsymbol{C}$-circles. A $\boldsymbol{C}$-*arc* is a nontrivial arc of a $\boldsymbol{C}$-circle. Given two points $p \neq q \in S^3$, there is a unique $\boldsymbol{C}$-circle containing p and q. The image of a $\boldsymbol{C}$-circle in $\mathcal{H}$ is either an ellipse, which projects to a round circle in $\boldsymbol{C}$, or the one point completion of a vertical line $\{z\} \times \boldsymbol{R}$.

Any G-translate of the circle $\boldsymbol{R}^2 \cap S^3$ is called an $\boldsymbol{R}$-*circle*. Not all $\boldsymbol{R}$-circles are round circles. An $\boldsymbol{R}$-*arc* is a nontrivial arc of an $\boldsymbol{R}$-circle. There is more than one $\boldsymbol{R}$-arc joining two points in S^3. The image of an $\boldsymbol{R}$-circle in $\mathcal{H}$ is either a curve which projects to a lemniscate in $\boldsymbol{C}$, or the one point completion of a nonvertical line. In this case, the line is horizontal if and only if it intersects $\{0\} \times \boldsymbol{R}$. We will call these horizontal lines *level* $\boldsymbol{R}$-*circles*.

Let $\mathcal{F}_{\boldsymbol{C}}$ and $\mathcal{F}_{\boldsymbol{R}}$ respectively be the sets of $\boldsymbol{C}$-arcs and $\boldsymbol{R}$-arcs. Let $\mathcal{F}$ be the union of $\mathcal{F}_{\boldsymbol{C}}$ and $\mathcal{F}_{\boldsymbol{R}}$. Our foliated patches in S^3 are based on $\mathcal{F}$.

3.3 Curvature Guarantees

We use the notation of §2.4. Since $\boldsymbol{C}$-arcs are arcs of round circles, we use the curvature guarantee $K_{\boldsymbol{C}} = K$, constructed in §2.4, as a curvature guarantee for $\mathcal{F}_{\boldsymbol{C}}$. For $\mathcal{F}_{\boldsymbol{R}}$, we use $K_{\boldsymbol{R}}(M) = ((6/s)^2 - 8)^{1/2}$. The lemma to follow justifies this choice.

Lemma 3.1 *Let γ be an $\boldsymbol{R}$-circle of Euclidean diameter δ. The maximum space curvature of γ is $\sqrt{(6/\delta)^2 - 8}$.*

Proof: Modulo isometry of $\boldsymbol{C}^2$, there is a one parameter family $\{f_v | \ v \geq 0\}$ of $\boldsymbol{R}$-circles:

$$f_v(t) = \Big(\frac{u\cos(t) - iv}{u + iv\cos(t)}, \frac{\sin(t)}{u + iv\cos(t)}\Big); \qquad u^2 - v^2 = 1.$$

$f_0 = S^3 \cap \boldsymbol{R}^2$ is a great circle. f_v shrinks to a point as $v \to \infty$. Clearly, $f_v \subset S^3$. One can show that the *angular invariant* [**G** p. 210] vanishes on triples of points on f_v, which suffices to prove that f_v is an $\boldsymbol{R}$-circle.

We claim that the Euclidean diameter of f_v is $\delta_v = 2/\sqrt{1+v^2}$. First of all, an easy calculation gives $\|f_v(\pi/2) - f_v(3\pi/2)\| = \delta_v$. To prove our equality, we will show that every point of f_v is within $\delta_v/2$ of $P_v = (-iv/u, 0)$. By symmetry, it suffices to consider $f_v(\theta)$ for $\theta \in [0, \pi]$. For this range, we set $\cos(t) = r/\sqrt{1+r^2}$ and $\sin(t) = 1/\sqrt{1+r^2}$ and compute:

$$g(r) = \|f_v(r) - P_v\|^2 = \frac{1+r^2+v^2}{(1+v^2)(1+r^2+v^2+2r^2v^2)}.$$

It is easy to check that $dg/dr = 0$ iff $r = 0$ and that 0 is a local maximum. Plugging in $r = 0$, and using the triangle inequality, we get our result.

Once we prove that the maximum space curvature of f_v is $\sqrt{1+9v^2}$, our lemma follows from basic algebra. Once again, it suffices to consider $f_v(\theta)$ for $\theta \in [0, \pi]$. Again, we use the substitution above. Let $K_v(r)$ be the curvature of f_v at r. We have

$$K_v(r) = \frac{|f_v'(r) \wedge f_v''(r)|}{|f_v'(r)|^3}.$$

Here $|A \wedge B|$ is the area of the parallelogram spanned by A and B. Let $\langle, \rangle$ be the definite Hermitian inner product on $\boldsymbol{C}^2$. Using the equations

$$|A \wedge B|^2 = |A|^2|B|^2 - (\Re\langle A, B\rangle)^2$$

we compute that

$$K_v^2 = \frac{1+r^2+10v^2+2r^2v^2+9v^4}{1+r^2+v^2+2r^2v^2},$$

and that $dK_v^2/dr = 0$ iff $r = 0$, and that 0 is a local maximum. Plugging in $r = 0$ gives us our curvature bound. ♠

Remark: We define the *radius* of a set to be half its diameter. Being a round circle, a $\boldsymbol{C}$-circle has the property that its geodesic curvature is $\tan(\theta)$ provided that its radius is $\cos(\theta)$. The calculation above says that an $\boldsymbol{R}$-circle has maximum geodesic curvature $3\tan(\theta)$, provided that its radius is $\cos(\theta)$.

3.4 Diameter Guarantees

For the case of $\mathcal{F}_{\boldsymbol{C}}$ we will use the same diameter guarantee as constructed in §2.3. The case of $\mathcal{F}_{\boldsymbol{R}}$ is based on an idea from [**S1**].

Define

$$U(\epsilon) = \{(z, w) \in S^3 |\ |z+1| \leq \epsilon\}. \tag{7}$$

Let $\pi_{\boldsymbol{C}} : \mathcal{H} \to \boldsymbol{C}$ be projection. Let l be the arc-length function. Let $\boldsymbol{B}_0$ be the map from Equation 6.

Lemma 3.2 *Suppose $\gamma \in S^3$ is a curve which is everywhere tangent to $\mathcal{E}$. If $\gamma \subset U(\epsilon)$ then $l(\gamma) \leq \epsilon\ l(\pi_{\boldsymbol{C}} \circ \boldsymbol{B}_0(\gamma))$.*

Proof: It suffices to prove the following infinitesimal version of this result: If V is a unit tangent vector to S^3, based at a point $p \in U(\epsilon)$, and tangent to $\mathcal{E}$, then $\|d(\pi_{\boldsymbol{C}} \circ \boldsymbol{B}_0)(V)| \geq 1/\epsilon$. Define

$$\nu = \frac{(1+z)^2}{1+\overline{z}} \begin{bmatrix} -\overline{w} \\ \overline{z} \end{bmatrix}$$

Since $\langle \nu, (z,w) \rangle = 0$, the vector ν is tangent to $\mathcal{E}$. Also $\|\nu\| \leq \epsilon$. Let $b = \pi_{\boldsymbol{C}} \circ \boldsymbol{B}_0$. We compute

$$db = \Big(\frac{-w}{(1+z)^2}, \frac{1}{1+z} \Big).$$

From this, we get $db(\nu) = 1$. This lemma now follows from the fact that b is a holomorphic map, and V is a complex multiple of ν. ♠

Given an element $(\gamma, p) \in \mathcal{F}(1,0)$, the $\boldsymbol{R}$-circle $-R_p^{-1}(\hat{\gamma})$ contains the point $(-1,0)$. Hence, $\boldsymbol{B}_0(-R_p^{-1}(\hat{\gamma}))$ is the one point completion of a straight line. We define

$$S(\gamma, p) = \pi_{\boldsymbol{C}} \circ \boldsymbol{B}_0(-R_p^{-1}(\gamma)). \tag{8}$$

By construction $S(\gamma, p)$ is a line segment in $\boldsymbol{C}$.

Definitions: Let $\epsilon_1(\gamma, p)$ be half the length of the line segment $S(\gamma, p)$. Let $\epsilon_2(\gamma, p)$ be the smallest ϵ such that $-R_p^{-1}(\partial\gamma) \in U(\epsilon)$. Let $\epsilon_3(\gamma)$ be the Euclidean distance between the endpoints of γ.

Since γ is integral to $\mathcal{E}$, Lemma 3.2 says that $l(-R_p^{-1}\gamma)) \leq \epsilon_1(\gamma, p)$. Hence,

$$\gamma \subset U(\epsilon); \qquad \epsilon = \epsilon_1(\gamma, p) + \epsilon_2(\gamma, p).$$

Applying Lemma 3.2, and using the fact that R_p^{-1} is an isometry, we get

$$l(\gamma) \leq 2(\epsilon_1(\gamma, p) + \epsilon_2(\gamma, p))\epsilon_1(\gamma, p). \tag{9}$$

The formula

$$\Delta(\gamma, p) = \frac{\epsilon_1(\gamma, p)(\epsilon_1(\gamma, p) + \epsilon_2(\gamma, p))}{\epsilon_3(\gamma)} \tag{10}$$

serves as a diameter guarantee for $\mathcal{F}_{\boldsymbol{R}}$.

3.5 Foliated Patch Clusters

Let $\Sigma = \boldsymbol{R}^+ \times [0, 2\pi) \times \boldsymbol{R}$. Consider the map $\phi : \Sigma \to \mathcal{H}$, given by

$$Z(r, \theta, t) = (r \exp(i\theta), t).$$

The r-curves of ϕ are $\boldsymbol{R}$-arcs. The θ-curves and the t-curves of ϕ are both $\boldsymbol{C}$-circles. Thus, the restriction of X to rectangle in Σ is a $\mathcal{F}$-patch.

We will only will consider rectangles having the form $Q = I_1 \times I_2 \times I_3$, where I_2 has length at most π. Let $\theta(Q)$ be the midpoint of I_2. Let $\theta'(Q)$ be the angle such that $\theta(Q) - \theta'(Q) \equiv \pi$, mod 2π. Let Q' be the 2-dimensional rectangle $I_1 \times \theta'(Q) \times I_3$. By construction, $Z(Q)$ and $Z(Q')$ are disjoint, and every θ-curve of $Z(Q)$ is contained in a circle which intersects $Z(Q')$.

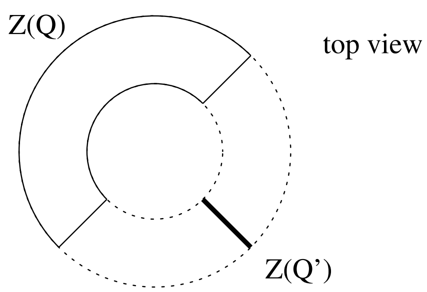

Figure 3.4.

We define a *standard patch cluster* to be triple of the form (A, A', p), where

$$\boldsymbol{B}(A) = Z(Q); \quad \boldsymbol{B}(A') = Z(Q'); \quad \boldsymbol{B}(p) = \infty.$$

Here $\boldsymbol{B}$ is a Heisenberg stereographic projection. We transfer the notion of θ-curves, t-curves, and r-curves, in the obvious way, to the patches A and A'. By construction, each θ-curve of A is contained on a circle which intersects A'. Each r-curve of A is contained on a $\boldsymbol{R}$-circles which contains p. Each t-curve of A is contained on an $\boldsymbol{C}$-circles which contains p. The same statements are true for A'.

The patch A has 12 distinguished edges. These edges have the form $\boldsymbol{B}^{-1} \circ Z(Q_1)$. Here Q_1 is the 1-skeleton of Q. Also, A has 8 distinguished vertices. Likewise A' has 4 distinguished edges and 4 distinghished vertices.

3.6 The Guarantee Algorithm

Let (A, A', p) be a standard patch cluster, as above. In the description of the algorithm to follow, we will assume that the computer halts if it fails to verify any given step. If the computer gets to the end of the list of things to check, without failure, it means that the foliated patch A is guaranteed. In what follows, for any $\gamma \in \mathcal{F}$, the continuation $\hat{\gamma}$ is a closed curve.

Let B be the smallest ball centered at $c(A)$ and containing all vertices of A. Likewise define B'. Recall that K and Δ are our curvature and diameter guarantees. Let ρ and ρ' be the radii of B and B'. Let c and C' be the centers of B and B'.

1. Each r-edge γ of A' defines a canonical element $(\hat{\gamma}, p)$ of $\mathcal{F}_{\boldsymbol{R}}(1,0)$. Use $K_{\boldsymbol{R}}$ and $\Delta_{\boldsymbol{R}}$ on this element to verify that γ is straight enough for A'.

2. Each t-curve α of A' determines the element $(\hat{\alpha}, p)$ of $\mathcal{F}_{\boldsymbol{C}}(1,0)$. From the previous step, each endpoint of α is within ρ' of c'. Therefore, α has diameter at least $\|c' - p\| - \rho'$. Use $K_{\boldsymbol{C}}$ to check that $\rho(B)\kappa(\alpha) < 1$. That is, check that
$$\frac{\rho'}{\|c' - p\| - \rho'} < 1.$$
This shows that $\rho'\kappa(\alpha) < 1$ for any such choice of α.

3. Let α be one of the t-edges of A'. By Step 1, the emdpoints of α are within $2\rho(B')$ of each other. As in Step 2, each endpoint is at least $\|c(A') - p\| - \rho(B')$ from p. Use $\Delta_{\boldsymbol{C}}$ to check that α is straight enough for B'. One only has to do this for a single choice of α.

4. Let A_1 and A_2 be the two tr-faces of A. Repeat Steps 1-3 for both A_1 and A_2, using the ball B in place of B'.

5. At this point we know that $A' \subset B'$, and $A_1, A_2 \subset B$. Let α be any θ-curve of A. Let $q_\alpha = \alpha \cap A'$. We know that q_α is within $\rho(B')$ of $c(A')$. We also know that both endpoints of α are within $\rho(B)$ of $c(A)$. Hence, q_α is at least $\|c(A') - c(A)\| - \rho(B) - \rho(B')$ from either endpoint of α. Repeat the check in Step 2, with B replacing B', using the bound just derived. This checks that $\kappa(\alpha)\rho(B) < 1$ for all choices of α.

6. Choose a single θ-curve α of A, and repeat Step 3, with B replacing B', and with the bound from Step 5 replacing the bound from Step 2.

3.7 Parabolic Hybrid Cones

Our foliated patch cluster is designed to deal with a certain kind of surface, which we call a *hybrid cone*. There are several kinds of hybrid cones. We will treat the simplest one, the *parabolic hybrid cone*.

We call $H_0 = \{0\} \times \boldsymbol{R}$ the *center* of Heisenberg space. We say that a *flag* is a pair (E, p), where E is a chain and $p \in E$ is a point.

Lemma 3.3 *Suppose $X \in S^3 - E$. There is a unique $\boldsymbol{R}$-circle $\gamma = \gamma(E, p; X)$ such that $X \in \gamma$ and $p \in \gamma$ and $\gamma \cap (E - p) \neq \emptyset$.*

Proof: We normalize by a Heisenberg stereographic projection so that $E = H_0$, the center of $\mathcal{H}$, and $p = \infty$. In this case, there is a unique level Heisenberg $\boldsymbol{R}$-circle containing X. This is $\gamma(H_0, \infty; X)$. ♠

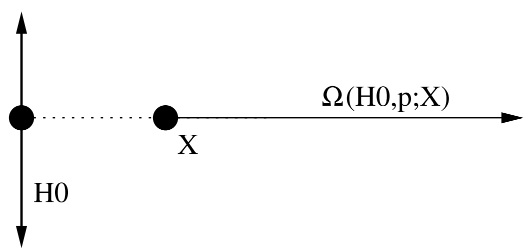

Figure 3.1.

Let $\Omega(E, p; X)$ be the portion of γ which connects p to X but which avoids $E - p$. Given a set $S \subset S^3 - E$, we define

$$\Omega(E, p; S) = \bigcup_{X \in S} \Omega(E, p; X).$$

We call Ω the *parabolic hybrid cone*. When the context is clear, we will call Ω a *hybrid cone*, as we did in [**S1**]. Our construction is natural. The $PU(2,1)$-image of a hybrid cone is again a hybrid cone.

We say that $\Omega(E, p; S)$ is in *standard position* if, as in Lemma 3.3, it is normalized so that E is the center of $\mathcal{H}$ and $p = \infty$.

One special case we have found particularly useful is when S is a $\boldsymbol{C}$-arc, contained on a $\boldsymbol{C}$-circle which links E. In this case, we call $\Omega(E, p; S)$ a *hybrid sector*.

3.8 Constructing Patch Maps

Let $T_\lambda : \mathcal{H} \to \mathcal{H}$ be the map $T_\lambda(z,t) = (\lambda z, t)$. Note that T_λ is not conjugate to an element of $PU(2,1)$, via Heisenberg stereographic projection.

Suppose $\gamma : [0,1] \to \mathcal{H}$ is a continuous curve. We say that a *dominating map* is a map Q, which assigns to each triple $(I, \lambda_1, \lambda_2)$, a rectangle $Q(I, \lambda_1, \lambda_2) \subset \Sigma$ such that

$$T_{\lambda_1}(\gamma(I)) \cup T_{\lambda_2}(\gamma(I)) \subset Z(Q(I, \lambda_1, \lambda_2)).$$

We do not insist that $Q(I, \lambda)$ is always the optimal—that is, smallest—possible rectangle which has this property. In case γ is a Heisenberg $\boldsymbol{C}$-arc—or more generally an arc of an ellipse—the optimal dominating map is quite easy to work out, and may be implemented as a computer program. We observe that $Q(I, \lambda_1, \lambda_2)$ also contains the level $\boldsymbol{R}$-arcs joining $T_{\lambda_1}(\gamma(I))$ to to $T_{\lambda_2}(\gamma(I))$.

Suppose $\Omega = \Omega(H_0, \infty; \gamma)$ is in standard position. We say that a *standard parametrization* of Ω is a map $\Phi : [0,1] \times [1, \infty) \to \Omega - \infty$ such that $\Phi(*, 1)$ is a parametrization of γ, and

$$\Phi(t, \lambda) = T_\lambda(\Phi(t, 1)).$$

Such a map is uniquely determined, and also easy to construct, given a parametrization of γ.

We will explain how to construct a patch map for the restriction of Φ to a subrectangle $R_N = [0,1] \times [0, N]$. If N is sufficiently large, then the only part of Ω left uncovered is a small neighborhood of ∞ which can be estimated separately. Given a subrectangle $R \subset R_N$, we extract the triple $(I, \lambda_1, \lambda_2)$ such that $R = I_1 \times [\lambda_1, \lambda_2]$. We define

$$\Psi(R) = (Z(Q), Z(Q'), \infty); \qquad Q = Q(I, \lambda_1, \lambda_2).$$

By construction, $Z(Q)$ contains $\Phi(R)$. In other words, $\Phi(R) \subset (\Psi(R))_1$. Hence, Ψ is a patch map dominating Φ.

If $\Omega(E, p; S)$ is not in standard position, we move Ω into standard position, where it is easy to construct parametrizations and patch maps. Then, we pull these objects back to Ω.

4 References

[**E**] D.B.A. Epstein, *Complex Hyperbolic Geometry*, London Mathematical Society Lecture Notes, **111**, 1987

[**G**], W. Goldman, *Complex Hyperbolic Geometry*, Oxford University Press, 1999.

[**GP**] W. Goldman and J. Parker, *Complex hyperbolic ideal triangle groups*, J reine agnew. Math **425**, 1992.

[**S1**] Rich Schwartz, *Ideal Triangle Groups, Dented Tori, and Numerical Analysis*, Annals of Mathematics (to appear)

[**S2**] Rich Schwartz, *Degenerating the Complex Hyperbolic Ideal Triangle Groups*, preprint, 1999.