ISRAEL JOURNAL OF MATHEMATICS 109 (1999), 349–377

TWIN TREES II: LOCAL STRUCTURE AND A UNIVERSAL CONSTRUCTION

BY

M. A. RONAN*

Mathematics Department, University of Illinois at Chicago
851 S. Morgan, mc249, Chicago, IL 60607, USA
e-mail: ronan@uic.edu

AND

J. TITS

Collège de France
Place Marcelin Berthelot, Paris Cedex 05, F-75231, France

ABSTRACT

We investigate the local structure of twin tress, and produce for any semi-homogeneous tree T a "universal twin" that contains, in a natural way, all trees twinned with T. Our methods allow us to show that there are uncountably many isomorphism classes of twinnings involving any semi-homogeneous tree.

1. Introduction

Twin trees were introduced by the authors in [RT] where a further paper on the subject was promised. That paper is still in preparation, but in the meantime we present here a procedure for obtaining all twin trees, showing incidentally that there is a great variety of them, including many that are rigid. In the process we construct, for any semi-homogeneous tree T, a graph T^* that may have some independent interest. To help the reader through some rather technical developments, this introduction will describe, in a somewhat loose and heuristic way, the main ideas and methods.

First we recall that a twin tree is a pair of trees together with a codistance function between vertices of one and vertices of the other. This codistance is a

* Partially supported by the USA National Science Foundation.
Received September 10, 1997 and in revised form July 6, 1998

non-negative integer satisfying the following simple property. If $\mathrm{codist}(x, y) = n$ and if y' is adjacent to y, then $\mathrm{codist}(x, y') = n \pm 1$; furthermore, if $n > 0$, then $+1$ occurs for a unique such neighbour y'.

In [RT] we gave examples of twin trees associated to the group GL_2 over a ring $k[t, t^{-1}]$ of Laurent polynomials. In these examples the valency of each vertex was one greater than the cardinality of k. In particular, if the trees were locally finite then they were homogeneous of valency $1 + q$, where q is a prime power. An elementary result (Prop. 1 in [RT]) shows that if the trees are thick (i.e. all valencies are at least three), then vertices at even distance from each other must have the same valency; one calls such trees **semi-homogeneous**. We now let T denote a thick semi-homogeneous tree.

Our main tool is a certain graph T^* called the **universal twin** of T. It is introduced in section 3 and has the following properties:

(i) there is a codistance function between the vertices of T and those of T^*, analogous to the codistance between vertices of two twinned trees;

(ii) every tree twinned with T embeds in T^* as a full subgraph, in a unique way compatible with the codistances.

The vertices of T divide naturally into two types, vertices at even distance being of the same type. The vertices of T^* divide into the same two types, a vertex of T and a vertex of T^* having the same type if they are at even codistance. Because of the canonical way T^* is defined, any automorphism of T operates also on T^*, and the full automorphism group of T acts transitively on the set of vertices of a given type in T^*. We arbitrarily fix one of the types, and call it the first type, the other being called the second type. For the entire paper we choose once and for all a vertex σ in T^* of the first type; in view of the transitivity this choice is indifferent. All constructions we shall be performing in T^* will occur in the connected component containing σ, and we denote this component by the symbol T^0.

Our version of the construction problem is now to find all subtrees S of T^0 containing σ and twinned with T. The principle is to obtain any such tree S as the union of its balls of radii $1, 2, 3, \ldots$ centred at σ. A ball of radius n in T^0 is called **uniform** if it embeds as the full ball of radius n in a tree twinned with T. Once we have uniform 1-balls we can construct a tree twinned with T because a straightforward induction, given in section 5, does the rest. However, in order to study the local structure of the twinning, uniform 2-balls are needed, and these are used in section 6 onwards.

To illustrate some applications of our construction, we have chosen to exhibit a vast number of isomorphism classes of twinnings, and show the existence of rigid twinnings, i.e. twinnings with no non-trivial automorphisms. The main result of this paper, proved in section 8, is that if α denotes the cardinality of the vertex set of T, then

(1.1) *there exist 2^α isomorphism classes of twin trees in which T is one of the two trees, and among these 2^α are rigid.*

This of course proves the existence, not evident a priori, of twinnings for any thick semi-homogeneous tree. Our method of proving (1.1) involves counting conjugacy classes of uniform balls. More precisely, we define two subgraphs of T^0, or of T, to be σ-conjugate if there is an automorphism of T fixing σ and carrying one to the other. First, let us deal with the possible hope that one could obtain (1.1) by using only uniform 1-balls. We shall prove in section 6 that

(1.2) *any two uniform 1-balls centred at σ are σ-conjugate.*

By contrast we shall prove in section 8 that

(1.3) *there are 2^α σ-conjugacy classes of uniform 2-balls.*

The first part of (1.1) follows readily from (1.3). The second part uses the following theorem, proved in section 6. Let S be a second copy of T, and for each vertex v in S choose a uniform 2-ball B_v in T^0 whose centre has the same type as v. Then there exists a twinning of T and S–hence by (ii) an embedding of S in T^0–such that, for each vertex v in S, its second neighbourhood in S is conjugate to B_v. This fact, along with (1.3), is used to prove the second part of (1.1).

We now describe in a little more detail our method for studying uniform 1- and 2-balls. Heuristically speaking we look at them "from infinity" using ends of T. To explain this, we first recall from [RT] that if S is a tree twinned with T, then some ends of T are naturally identified with some ends of S. The ends of T that one obtains in this way depend of course on S. However, when we regard S as a subtree of T^* then this set of ends depends only on the connected component of T^* containing S. Since we deal with only one connected component T^0 of T^*, namely that containing σ, we always have the same set of ends, and we call this set E.

Let us now get back to uniform n-balls centred at σ. There is a natural map from E to the boundary of such a ball; it sends an element e of E to the boundary vertex x "closest" to it (closest means that x is at maximal codistance from vertices of T as these vertices approach e). The boundary vertices therefore partition E, and this partition determines the ball in question. For a uniform

1-ball the partition is called a σ-**colouring** of E. A uniform 2-ball S_2 centred at σ contains a unique uniform 1-ball S_1 centred at σ, and the pair (S_2, S_1) gives a pair of partitions (P_2, P_1) with P_2 being a refinement of P_1. We call this a σ-**binary colouring**.

The term σ-conjugacy extends naturally to subsets and partitions of E, and assertions (1.2) and (1.3) can be reformulated as:

(1.2′) *all σ-colourings are σ-conjugate,*

and

(1.3′) *there exist 2^α σ-conjugacy classes of σ-binary colourings.*

The first assertion is proved in section 6, and the second in section 8.

The idea of σ-conjugacy is formulated in terms of automorphisms of T, so it is important not to lose track of this tree. However, only a small part of it is really used in our characterization of colourings and binary colourings. We use the set E along with the vertices of T at codistance 0 and 1 from σ (called 0- and 1-**vertices**). More precisely, for σ-colourings we use a tree called T/σ whose vertex set is E along with the set of 0-vertices, and for σ-binary colourings we use a tree called $T/(\sigma - 1)$ whose vertex set is E along with the set of 0- and 1-vertices.

A σ-colouring is a partition of E characterized by the property that each 0-vertex of T/σ be adjacent to exactly one end-vertex from each part of the partition. This condition makes the construction of uniform 1-balls almost a triviality, and it can be shown that the assertion (1.2′) just reflects a simple property of the action on T/σ of automorphisms of T fixing σ. A uniform n-ball is the union of a uniform $(n-1)$-ball and uniform 1-balls centred at the boundary vertices of that uniform $(n-1)$-ball. Moreover, a subtree of T^0 twinned with T is the union of a nested sequence of uniform n-balls for $n = 1, 2, \ldots$. This provides our general construction, described in section 5, of trees twinned with T.

In dealing with uniform 2-balls, the pair (P_2, P_1) of partitions mentioned above is given by a sequence of maps $E \xrightarrow{\chi} J \xrightarrow{\pi} I$. The partition P_1 is provided by the fibres of $\pi\chi$, and its refinement P_2 by the fibres of χ. The fibres of $\pi\chi$ must of course form a σ-colouring, and this, along with a second condition given in section 7, ensures that (χ, π) determines a σ-binary colouring.

In section 8 we use these maps (χ, π) to describe certain vertices of E as "neutral". The set of neutral vertices depends only on the partitions given by χ and $\pi\chi$, and is therefore an invariant of the σ-binary colouring that they define. If two such colourings are σ-conjugate, then so are their neutral sets. We show that given any subset E_0 of E there is a σ-binary colouring for which E_0 is the neutral set. The desired result (1.3′) then follows from the fact that there are 2^α σ-conjugacy classes of subsets of E.

So far, we have used uniform 2-balls as a means to prove the existence of a huge variety of non-isomorphic twinnings (for any semi-homogeneous tree), but they can also be interesting geometric objects in their own right. We illustrate this in section 9 where we define and briefly discuss a very special type of uniform 2-ball that we call "flat". These occur in twin trees associated to important groups, such as $\mathrm{GL}_2(k[t,t^{-1}])$ (already considered in part I of this work), and in some Kac–Moody groups (which will appear in part III). They also connect the general theory of projective planes to that of twin trees.

The idea is as follows. If B is a uniform 2-ball with centre σ, then every vertex v of T at codistance 0 from σ determines a set σ^v of elements in ∂B at codistance 2 from v. We call these sets the **lines** of ∂B; they have the property that each neighbour of σ in B is adjacent to exactly one point of such a line. If τ is a neighbour of σ in B, then the neighbourhood of τ is a subset of $\partial B \cup \{\sigma\}$ that we also call a **line**. We therefore have two types of lines in $\partial B \cup \{\sigma\}$: those containing σ, and those of the form σ^v lying entirely in ∂B.

The geometric structure comprising $\partial B \cup \{\sigma\}$ and its lines is a useful invariant of the 2-ball B. Any two points lie on a common line, and when T is homogeneous it is possible for $\partial B \cup \{\sigma\}$ to be a projective plane; this is the **flat** case. In a given twin tree it may happen that some balls are flat, or that some balls are and some not, or that none are. (A similar phenomenon occurs in the more elementary context of generalized polygons. Here also, one defines 2-balls and, in exactly the same way as above, the system of lines on such 2-balls. When the polygon in question is the generalized triangle associated to a projective plane Π, the 2-balls, along with their lines, are projective planes—hence flat—isomorphic to Π or to its dual. The next simplest case is that of the generalized quadrangle associated to $\mathrm{Sp}_4(k)$, k a field. In this case let B be a 2-ball centred at a vertex σ. Depending on the type of σ, B is either isomorphic to the projective plane over k, or to a quadratic cone along with its generating lines and non-degenerate conic plane sections; in the former case it is flat, and in the latter case it is not flat unless k is perfect of characteristic 2.)

Now suppose T is homogeneous of valency $q+1$, and suppose we are given a projective plane Π of order q and a point ω in Π. Then there exists a flat 2-ball B in T^0 centred at σ such that $\partial B \cup \{\sigma\}$ and its lines is isomorphic to Π, via an isomorphism sending σ to ω (Lemma 9.2). Moreover, that 2-ball is unique up to conjugacy (a fact whose proof we leave to the reader). More generally, assume

T and S are homogeneous trees of valency $q+1$, and distinguish one of the two types of vertices of S. For each vertex of that type choose a punctured projective plane of order q. Then (Proposition 10) *there exists a twinning of S and T, such that the 2-ball of S centred at any vertex of the distinguished type is flat and as a punctured projective plane is isomorphic to the one chosen for that vertex.* If there exist two non-isomorphic projective planes of order q, then when T is homogeneous of valency $q+1$ this provides 2^α non-isomorphic twinnings in which all 2-balls of one type in S are flat, and a fortiori an alternative proof of (1.3).

In rounding off our introduction it is appropriate to mention the work of Fon-Der-Flaass [F], who gives a construction of twin trees by taking a single tree and associating a non-negative integer to each pair of its vertices. His methods are particularly illuminating for trees of low valency.

2. Horospheres and horoballs

Let T be a tree having no vertex of valency 1, and recall the usual "buildings" terminology of **apartments** and **half-apartments** as in [RT; sect. 3]. Given an end e and a vertex x of T, let (xe) denote the half-apartment starting at x and having end e. We say that two vertices x and y are **equidistant** from e if they are equidistant from some (hence any) vertex in the half-apartment $(xe)\cap(ye)$. Being equidistant from e is an equivalence relation on the set of vertices, and the equivalence classes are called the **horospheres of T centred at** e. If C is such a horosphere, we call the union H of the (xe) for $x \in C$ a **horoball centred at** e, and refer to C as its **boundary** ∂H.

Given horoballs H and H' having the same centre e, we let $H-H'$ denote their radial difference, defined as the distance from ∂H to $\partial H'$ taken in the direction towards e. In particular, $H-H'$ is positive when H contains H' and negative otherwise. When $H-H'=r$ we write $H=H'+r$.

TWIN TREES. Let T_- be a tree twinned with T.

(2.1) PROPOSITION 1: *Let v_- be a vertex in T_- and e an end of the twinning. Let H_0 denote the set of vertices x in T for which the codistance from v_- increases monotonically along (xe). Given a non-negative integer n, the subset H_n (resp. C_n) of H_0 consisting of those x with* $\operatorname{codist}(x,v_-)\geq n$ *(resp.* $=n$*) forms a horoball (resp. a horosphere) with centre e.*

Proof: Let $x\in C_n$. If y lies in the horosphere containing x and centred at e, then x and y are at distance d from some common vertex $z\in(xe)\cap(ye)$. The definition of H_0 implies that $\operatorname{codist}(z,v_-)=n+d$, and since the codistance from v_- increases along (ze) it must decrease monotonically from z to y, so $\operatorname{codist}(y, v_-) = n$ and $y \in C_n$. Conversely, if $y \in C_n$ then the codistance increases monotonically from n along both (xe) and (ye), so both x and y are equidistant from a vertex of $(xe) \cap (ye)$, proving that they lie in the same horosphere centred at e. Thus C_n is a horosphere, and H_n, as the union of the C_m for $m \geq n$, is a horoball. ∎

(2.2) Corollary: *If V denotes the set of vertices opposite v_-, then V is a union of horospheres, and the complement of V in T is a union of disjoint horoballs, the centres of which are all the ends of the twinning.*

Proof: If x is any vertex of T at codistance n from v_-, then the codistance from v_- increases monotonically along (xe) for some end e, and if $n \geq 1$ then e is unique. Apply Proposition 1: the first statement uses $n = 0$, and the second statement uses $n \geq 1$; the disjointness of the horoballs follows from the uniqueness of e when $n \geq 1$. The ends of the twinning are precisely the ends that arise in this way—see [RT; sect. 3]. ∎

We call the disjoint horoballs in this Corollary the **components** of v_- in T. The next proposition and its corollary show how they alter as we move from v_- to an adjacent vertex w_-.

(2.3) Proposition 2: *Let x and y be vertices in the same component of v_-, and let w_- be a neighbour of v_- in T_-. If $\operatorname{codist}(w_-, x) = \operatorname{codist}(v_-, x) + 1$ (resp. -1), then $\operatorname{codist}(w_-, y) = \operatorname{codist}(v_-, y) + 1$ (resp. -1).*

Proof: Let e be the centre of the component of v_- containing x and y, and let z be a vertex of $(xe) \cap (ye)$. The codistances from v_- increase monotonically along (xe) and (ye), hence in particular from x to z, and from y to z. The codistances from w_- must do likewise or else they would start decreasing to zero from some point, and this would contradict the fact that the codistance from w_- equals the codistance from v_- plus or minus 1. Therefore, if the codistance from w_- to x is 1 greater (or 1 less) than the codistance from v_- to x, then it must remain so along the path from x to z and from y to z. ∎

(2.4) Corollary: *Let v_- and w_- be adjacent vertices of T_-, and let e be an end of the twinning. Then the components of v_- and w_- centred at e have radial difference 1.*

3. The universal twin T^*

We begin this section by defining two sets **F** and **V** that are in canonical one-to-one correspondence with one another.

- **F** is the set of all functions f: $\mathrm{Vert}(T) \to \{0,1,2,3,\dots\}$ such that for adjacent vertices x and y, $f(x)-f(y) = \pm 1$, and if $f(x) \neq 0$ then $f(x)-f(y)$ takes the value -1 exactly once as y runs over all neighbours of x.
- **V** is the set of all discrete subsets V of $\mathrm{Vert}(T)$ (i.e. any two points of V are non-adjacent) such that the complement of V in T is a union of disjoint horoballs (called the **components** of $T - V$).

The canonical bijection is given by $f \mapsto V$ where $V = f^{-1}(0)$. On the other hand, starting with V one defines f as follows: if x lies in V then $f(x) = 0$; otherwise x lies in a unique component of $T - V$, and $f(x) = 1 + d$ where d is the distance from x to the boundary of this component. We let Σ denote the sets **F** and **V** identified by this bijection, and for $\sigma \in \Sigma$ we let f_σ and V_σ denote the representatives of σ in **F** and **V**, respectively. Notice that Σ is always non-empty because each vertex of T has valency at least 2.

The **components** of σ will mean the components of $T - V_\sigma$. The centres of these components are certain ends of T; we let $E(\sigma)$ denote the set of such ends, and call them the **ends** of σ. If e is such an end then $\sigma(e)$ will denote the component of σ whose centre is e. If T_- is a tree twinned with T, and v_- a vertex of T_-, then the codistance from v_- is an element of **F**, and hence of Σ. Its set of ends $E(v_-)$ does not depend on the choice of v_- in T_-; any two vertices of T_- give the same set of ends (see [RT; sect. 3, remark 1]).

THE UNIVERSAL TWIN T^*. We now define a graph T^* that we call the **universal twin** of T. Its vertices are the elements of Σ, and we define the codistance between vertices of T and T^* by $\mathrm{codist}(\sigma, x) = f_\sigma(x)$. Two vertices σ and σ' of T^* are adjacent if they have the same set of ends (i.e. $E(\sigma) = E(\sigma')$) and if $\sigma(e) = \sigma'(e) \pm 1$ for each $e \in E(\sigma)$. This is the same as requiring that $|f_\sigma(x) - f_{\sigma'}(x)| = 1$ for all vertices x in T. Notice that any two vertices σ and τ in the same connected component of T^* have the same set of ends (i.e. $E(\sigma) = E(\tau)$). Note also that the vertices of T^* can, like those of T_-, be separated into two types depending on the parity of the codistance from a vertex of T.

An automorphism α of T induces an automorphism α^* of T^* as follows. For any $\sigma \in T^*$ one defines $\alpha^*(\sigma)$ to be the function on T given by $\mathrm{codist}(\alpha^*(\sigma), v) = \mathrm{codist}(\sigma, \alpha^{-1}v)$.

(3.1) PROPOSITION 3: *When T is semi-homogeneous, the group of type preserving automorphisms of T is transitive on the set of all vertices in T^* having a given type.*

Proof: Given two vertices σ and σ' of the same type in T^*, it is straightforward to construct an automorphism α of T such that $\text{codist}(\sigma', \alpha(v)) = \text{codist}(\sigma, v)$, and hence $\alpha^*(\sigma) = \sigma'$. We leave the details to the reader.

We now make the following hypothesis which is satisfied for all thick twin trees (see [RT]; Prop. 1).

STANDING HYPOTHESIS: *We assume all vertices in T of the same type have the same valency; i.e., T is semi-homogeneous.*

(3.2) PROPOSITION 4: *If T_- is a tree twinned with T, then the identity map on T extends uniquely to a codistance preserving map from (T, T_-) to (T, T^*); its restriction to T_- maps this tree isomorphically onto a full subgraph of T^*.*

Proof: This is straightforward; every vertex of T_- is represented by a unique element of Σ, and if σ and σ' represent neighbouring vertices of T_-, then by the Corollary to Proposition 2, σ and σ' are neighbours in T^*. Furthermore, if v_1 and v_2 represent distinct vertices v_1 and v_2 of T_-, then $\sigma_1 \neq \sigma_2$ because there are vertices of T opposite one but not the other, for example in any twin apartment containing v_1 and v_2. This shows that the canonical image of T_- in T^* is a full subgraph of T^*. ∎

THE SET E. We intend to obtain twin trees by constructing sub-trees of T^*, and in view of (3.1) it makes no difference which connected component of T^* we work in. Let us therefore pick, once and for all, a given component T^0. Any two vertices σ and σ' of T^0 have the same set of ends $E(\sigma) = E(\sigma')$, and we let E denote this set. All vertices σ of T^* that we deal with from now on will lie in T^0.

Remark: In Propositions 1 and 2 and their corollaries, T_- can be replaced by T^0 modulo suitable adjustments. One has:

PROPOSITION 1′: *Let σ be a vertex of T^0, and e an end of σ. Given a non-negative integer n, the set of vertices x in T for which* $\text{codist}(\sigma, x) \geq n$ *(resp. $= n$) and such that the codistance from σ increases monotonically along (xe) forms a horoball (resp. a horosphere) with centre e.*

PROPOSITION 2′: *Let x and y be vertices in the same component of σ, and let τ be a neighbour of σ in T^0. If* $\text{codist}(\tau, x) = \text{codist}(\sigma, x) + 1$ *(resp. -1), then* $\text{codist}(\tau, y) = \text{codist}(\sigma, y) + 1$ *(resp. -1).*

COROLLARY: *Let σ and τ be adjacent vertices of T^0, and let e be an end of σ (and τ). Then the components of σ and τ centred at e have radial difference 1.*

UNIFORM SETS OF NEIGHBOURS IN T^0. Let σ be a vertex of T^0. We wish to distinguish subsets of the neighbourhood of σ in T^0—called "uniform" subsets—that arise in sub-trees of T^0 containing σ and twinned with T. Let v be a vertex of T in V_σ (i.e. $\text{codist}(\sigma, v) = 0$), and notice that the codistance between a neighbour of σ and a neighbour of v is necessarily 0 or 2. Using $T(v)$ to denote the neighbourhood of v in T, we call a set U of neighbours of σ in T^0 **uniform** if for each $v \in V_\sigma$, the relationship of being at codistance 2 sets up a bijection between U and $T(v)$.

(3.3) THEOREM 1: *Any tree twinned with T embeds as a full subgraph of T^0, preserving its codistance with T. Furthermore, a connected subgraph T of T^0 is a tree twinned with T if and only if, for each vertex σ in S, the neighbourhood $S(\sigma)$ is a uniform set of neighbours in T^0.*

Proof: The first statement is Proposition 4. Now suppose S is a connected subgraph of T^0. If S is a tree twinned with T, then by the definition of "uniform" every neighbourhood $S(\sigma)$ is a uniform set of neighbours. Conversely, suppose that for each σ in S the neighbourhood $S(\sigma)$ is uniform. Then the lemma below (applied in this case with $\partial S = \emptyset$) shows that S is a tree twinned with T. ∎

For any connected graph X let ∂X (the **boundary** of X) denote the set of vertices of valency at most 1 in X. The next lemma, used in the theorem above, will also be applied later when we construct a tree twinned with T.

(3.4) LEMMA: *Let S be a connected subgraph of T^0 such that the neighbourhood in S of each vertex in $S - \partial S$ is a uniform set of neighbours in T^0. Let σ be a vertex of S, and x a vertex of T with* $\text{codist}(\sigma, x) = n \geq 1$. *Then*

(i) *for all neighbours x' of x in T,* $\text{codist}(\sigma, x') = n \pm 1$ *with $+1$ occurring for a unique such neighbour;*

(ii) *for all neighbours σ' of σ in S,* $\text{codist}(\sigma', x) = n \pm 1$ *and if $\sigma \notin \partial S$ then $+1$ occurs for a unique such neighbour;*

(iii) *S is a tree.*

Proof: (i) is immediate from the definition of f_σ and the fact that $\text{codist}(\sigma, x') = f_\sigma(x')$.

(ii) We have $f_\sigma(x) = \text{codist}(\sigma, x) \geq 1$. If $f_\sigma(x) = 1$, then $x \in T(v)$ for some v opposite σ, and by uniformity there is exactly one σ' in $S(\sigma)$ for which $f_{\sigma'}(x) = 2$, as required. If $f_\sigma(x) = n > 1$, let e be the end for which $x \in \sigma(e)$

and take any y in $\partial\sigma(e)$. Since $f_\sigma(y) = 1$ there is, as above, a unique σ' in $S(\sigma)$ for which $f_{\sigma'}(y) = 2$, and by Proposition 2′ applied to vertices of $S(\sigma)$ we have $f_{\sigma'}(x) = n + 1$ for exactly this one σ' in $S(\sigma)$, as required.

(iii) We postpone the proof of 3.4 (iii) until after the next lemma, and first proceed with a definition.

UNIFORM PATHS. We use the word **path** to mean a sequence of vertices $v_0, v_1, \ldots$ with v_{i-1} adjacent to v_i and distinct from v_{i+1}, for each $i = 1, 2, \ldots$. Define a path L in T^0 to be **uniform** if the neighbourhood of each vertex σ in L lies in a uniform set of neighbours of σ. The lemma below implies that a uniform path is a geodesic between its end points. We remark without proof that there exist paths of length 2 that are not uniform and paths of length 3 that are not geodesics.

(3.5) LEMMA: *If $\sigma_0, \ldots, \sigma_n$ is a uniform path in T^0, then there exists a vertex x in T such that* $\operatorname{codist}(\sigma_i, x) = i$ *for all i. In particular, the distance between σ_0 and σ_n in T^0 is n.*

Proof: Pick an end e such that $\sigma_n(e) = \sigma_{n-1}(e) + 1$, and a vertex $x \in \sigma_n(e)$ with $\operatorname{codist}(\sigma_n, x) = n$. Then $\operatorname{codist}(\sigma_{n-1}, x) = n - 1$. Since $\{\sigma_n, \sigma_{n-2}\}$ lies in a uniform subset of the neighbours of σ_{n-1}, part (ii) of (3.4) implies that $\operatorname{codist}(\sigma_{n-2}, x) = n - 2$, and an obvious induction implies that $\operatorname{codist}(\sigma_{n-r}, x) = n - r$ for $r = 0, \ldots, n$. ∎

Proof of 3.4 (iii): That S is a tree is clear from Lemma 3.5, since each path in S is uniform. ∎

We end this section with a lemma that we shall use when counting isomorphism classes of twin trees. As above, C denotes a horosphere in T, and E denotes the set of ends of σ; we let $\operatorname{Vert} T$ denote the set of vertices of T.

(3.6) LEMMA: *If T is semi-homogeneous and not thin, then the sets* $\operatorname{Vert} T$, C, *and E all have the same cardinality.*

Proof: The conditions on T imply that all three sets are of infinite cardinality. Let α, γ, and ε denote the cardinalities of $\operatorname{Vert} T$, C, and E, respectively. Let e denote the centre of C. Since all $\aleph_0$ horospheres centred at e have the same cardinality γ, and since $\operatorname{Vert} T$ is the disjoint union of these horospheres, we have $\gamma\aleph_0 = \alpha$. This shows that $\alpha = \gamma$.

Now for any vertex c in C, there exist ends f in E such that $\{c\} = C \cap (ef)$, showing $\gamma \le \varepsilon$. On the other hand, there exists a disjoint set of horoballs whose centres form the set E, showing $\varepsilon \le \alpha$. Therefore $\alpha = \gamma = \varepsilon$. ∎

4. Uniform sets of neighbours in T^0 and the tree T/σ

To study uniform sets of neighbours further, we first describe the neighbourhood $T^0(\sigma)$ of a vertex σ in T^0, by introducing a new tree associated to σ.

THE TREE T/σ. Define T/σ to be the tree obtained from T by collapsing each component $\sigma(e)$ of σ onto a single vertex that we identify with its centre e (it is a tree because the components of σ are connected, so any path in T/σ lifts to a path in T). With this identification, the set of vertices of T/σ is $V_\sigma \cup E$ (we recall that E is the set of ends of σ). Given v in V_σ, the sets of neighbours of v in T and in T/σ are in canonical one-to-one correspondence, because each neighbour of v in T lies in a component of σ and no two neighbours lie in the same component. If $e \in E$, then the neighbours of e in T/σ are the vertices of the horosphere bounding $\sigma(e)+1$.

Given a neighbour τ of σ in T^0, Proposition 2′ tells us that as we move from σ to τ the codistance goes up by 1, or down by 1, throughout each component of σ. In other words, the function $f_\tau - f_\sigma$ remains constant on each of these components. Let $E_{\sigma\tau}$ denote the set of centres of those components for which $f_\tau - f_\sigma$ takes the value $+1$. In T/σ every v in V_σ is adjacent to exactly one vertex of $E_{\sigma\tau}$, namely the centre of the τ-component containing v.

Now let $\mathbf{X}(\sigma)$ denote the set of all subsets $X \subset E$ having the property that every v in V_σ is adjacent to exactly one member of X. We have just shown that $E_{\sigma\tau} \in \mathbf{X}(\sigma)$ for each τ in $T^0(\sigma)$, so we have a map φ: $T^0(\sigma) \to \mathbf{X}(\sigma)$.

(4.1) LEMMA: *The map φ is a bijection from $T^0(\sigma)$ to $\mathbf{X}(\sigma)$.*

Proof: We construct an inverse for φ. Let X be any member of $\mathbf{X}(\sigma)$. Define a function f on the vertices of T as follows. If $x \in V_\sigma$, then $f(x) = 1$, otherwise x is in a σ-component, and $f(x) = f_\sigma(x) \pm 1$ with $+1$ if the centre of this component lies in X, and -1 if not. Then f belongs to the set $\mathbf{F}$ of section 3 and therefore defines a neighbour $\tau \in T^0(\sigma)$. It is straightforward to see that $\mathbf{X} \mapsto f$ is an inverse for φ.

σ-COLOURINGS. We now wish to show that a set of neighbours τ of σ in T^0 is uniform precisely when the corresponding sets $\varphi(\tau)$ partition E. Call a partition of E a **σ-colouring of E** when each part lies in $\mathbf{X}(\sigma)$, i.e. a partition in which each vertex v in V_σ is adjacent in T/σ to exactly one vertex from each part. Given a set I whose cardinality equals the number of parts (the same as the valency of v in T), a bijection from I to the set of parts, we call an indexed σ-colouring, **indexed** by I. We refer to I as the set of **colours**, so each e in E acquires some colour i in I. Indexed colourings will be used in section 6.

(4.2) PROPOSITION 5: *The map φ above gives a bijection between the set of all uniform sets $U \subset T^0(\sigma)$ and the set of all σ-colourings of E.*

Proof: We show first that any uniform set U gives a σ-colouring of E. Given any v in V_σ and any e adjacent to v in T/σ, let w be the vertex of $\sigma(e)$ adjacent to v in T. Since U is uniform it contains a unique vertex τ at codistance 2 from w, in which case $e \in \varphi(\tau)$. Moreover, each τ in U is at codistance 2 from some neighbour w of v in T, and hence $\varphi(\tau)$ contains some e adjacent to v in T/τ. This shows that as τ ranges over U the $\varphi(\tau)$ form a σ-colouring of E, as required.

Conversely, by Lemma 4.1 a σ-colouring gives a set U of vertices of $T^0(\sigma)$, and reversing the argument above shows that U is uniform. ∎

CONSTRUCTING σ-COLOURINGS. In the next section we construct a sub-graph of T^0 in which every vertex neighbourhood is uniform. To do this we shall need to know that any neighbour lies in a uniform set of neighbours. This is a corollary of Proposition 6, for which we need a new definition. A **partial σ-colouring of E** will mean a collection of disjoint sets X in $\mathbf{X}(\sigma)$ such that for any vertex v in V_σ the set of its neighbours in T/σ not lying in any of these sets X has a cardinality independent of v. This cardinality condition means that the number of "uncoloured" vertices adjacent to v is the same for all v. It is automatically satisfied if the number of "coloured" vertices adjacent to v (i.e. the number of sets X) is finite. We also call a set $U_1 \subset T^0(\sigma)$ **partially uniform** if $\varphi(U_1)$ gives a partial σ-colouring of E.

Examples: The following subsets of $T^0(\sigma)$ are partially uniform: any subset of a uniform set, the empty subset of $T^0(\sigma)$ and the singleton set $\{\tau\}$ for any neighbour τ of σ in T^0.

(4.3) PROPOSITION 6: *Every partial σ-colouring of E can be extended to a σ-colouring.*

Proof: Let B denote the set of vertices in E which are "uncoloured" (i.e., do not lie in one of the sets of the partial colouring). For v in V_σ the neighbours of v which lie in B form a set whose cardinality is independent of v; let C be a set of that cardinality. To extend the given partial colouring to a full colouring, we need only exhibit a function $\lambda: B \to C$ such that, for every v in V_σ, the restriction of λ to $B \cap (T/\sigma)(v)$ is a bijection onto C. To do that, choose a vertex v_0 in V_σ and construct, by induction on n, the restriction λ_n of λ to the intersection of B with the ball of radius $2n+1$ in the tree T/σ centred at v_0. A bijection $\lambda_0: B \cap (T/\sigma)(v_0) \to C$ exists because of the assumption made on the cardinality of C. As for the inductive step ($n > 0$), one supposes λ_{n-1} already constructed and, for every v in V_σ at distance $2n$ from v_0, one takes the restriction of λ_n to $B \cap (T/\sigma)(v)$ to be any bijection onto C subject only to the condition that if the unique neighbour $e \in (T/\sigma)(v)$ at distance $2n-1$ from v_0 belongs to B, then $\lambda_n(e) = \lambda_{n-1}(e)$. By construction, $\lambda^{-1}(c) \subset \mathbf{X}(\sigma)$ for each c in C, and the sets $\lambda^{-1}(c)$ extend the partial colouring to a full σ-colouring. ∎

(4.4) COROLLARY: *Every partially uniform set U_0 can be extended to a uniform set.*

Remark: For U_0 empty, one can view the foregoing proof as an explicit construction of a uniform set.

5. The construction

To construct a tree twinned with T we create a sub-graph S of T^0 in which every vertex neighbourhood is uniform. By Theorem 1, S is twinned with T and every tree twinned with T arises in this way.

To construct S, we start with a single vertex σ of T^0 and work outwards. For this purpose a **uniform n-ball** of T^0 with **centre** σ will mean a sub-tree of T^0 having the following three properties:

(i) all its vertices are at distance $\leq n$ from σ;
(ii) all vertices on the boundary are at distance n from σ;
(iii) the neighbourhood of each vertex not on the boundary is uniform.

In particular, a uniform 1-ball is simply σ together with a uniform set of neighbours. The following characterization of uniform n-balls will be useful in section 6.

(5.1) LEMMA: *Let σ be a vertex of T^0, and let S be a subgraph of T^0 having the following two properties:*

(i) the distance in S from σ to any other vertex of S is at most n;
(ii) if v is a vertex of S whose distance from σ in S is less than n, then v and its neighbours in S form a uniform 1-ball.

Then S is a full subgraph of T^0 and is a uniform n-ball.

Proof: Let v and w be vertices of S that are adjacent in T^0. We intend to show that v and w are adjacent in S, and if v is at distance n from σ in S then w is unique. Without loss of generality, the distance in S from σ to w is less than or equal to the distance in S from σ to v. By hypothesis (i) there are paths γ and δ of length at most n from σ to v and to w, respectively. Each vertex on either of these paths, with the possible exception of the end vertices v and w, is at distance $< n$ from σ, and therefore by hypothesis (ii) these paths are uniform.

By the assumption on distances from σ to v and w (namely $\operatorname{dist}(\sigma, w) \leq \operatorname{dist}(\sigma, v)$) v cannot lie on δ. Let x be the last vertex of γ that is also on δ. If $x = w$ then v and w are adjacent in S. If $x \neq w$ then there exists a uniform path in S from v via x to w. By (3.5) this implies that the distance from v to w in T^0 is at least two, contradicting the adjacency of v and w. This leads to the following two conclusions:

(a) vertices v and w of S that are adjacent in T^0 are also adjacent in S;

(b) if v is at distance n from σ and adjacent to w, then w is the penultimate vertex on a path of length n from σ to v.

Conclusion (a) says that S is a full subgraph of T^0. Conclusion (b) implies that each vertex of S at distance n from σ is a boundary vertex of S; this, along with hypothesis (ii), implies that S is a uniform n-ball. ∎

Now given a uniform n-ball S_n, for $n \geq 1$, each vertex τ on its boundary has a single neighbour in S_n; by (4.4), we may extend this to a uniform set $U(\tau)$ of neighbours. Let S_{n+1} denote the union of S_n and all $U(\tau)$, as τ runs over ∂S_n. By (3.5) each vertex of $S_{n+1} - S_n$ is at distance $n+1$ from the centre of S_n and not adjacent to any other vertex of S_{n+1}. Therefore $\partial S_{n+1} = S_{n+1} - S_n$, and S_{n+1} itself is a uniform $(n+1)$-ball. By Theorem 1 and induction on n, we have:

(5.2) THEOREM 2: *Every uniform n-ball in T^0 lies in a tree twinned with T, and in any tree twinned with T the vertices at distance $\leq n$ from a given vertex form a uniform n-ball. In particular, every semi-homogeneous tree admits a twinning.*

If we are interested only in isomorphism classes of twin trees, then (3.1) shows that the vertex at which we start our construction is irrelevant.

PARTIAL TWINS. In the construction above we started with a single vertex and worked outwards. More generally we could have started with a non-empty subtree S_0 of T^0, as in Lemma 3.4, where the neighbours of any vertex not on the boundary form a uniform set of neighbours; such a graph we call a **partial twin** for T. For each vertex on the boundary, (4.4) allows us to extend its singleton neighbour in S_0 to a uniform set of neighbours; let S_1 denote the union of S_0 and these uniform sets of neighbours. By the argument above, using (3.5), $\partial S_1 = S_1 - S_0$, and hence S_1 is a partial twin. Continuing inductively we obtain a sequence of trees $S_1 \subset \cdots \subset S_n \subset \cdots$ whose union S is a tree twinned with T. The following proposition is now a straightforward consequence of Theorem 1.

(5.3) PROPOSITION 7: *Any partial twin S_0 lies in a tree twinned with T.*

6. Uniform 1-balls and $(1,1)$-balls

In this section, and for the rest of the paper, we shall assume thickness (i.e. the existence of at least three edges per vertex). This is no essential loss (see [J]).

STANDING HYPOTHESIS: *We assume that all trees T are semi-homogeneous and thick.*

Define two subgraphs of T^0 to be **conjugate** if there is an automorphism α° of T^0, induced by an automorphism α of T, sending one to the other. The first proposition in this section is that any two uniform 1-balls centred at σ are conjugate. This is the same as showing that given two σ-colourings of E, there is an automorphism of T fixing σ and sending one colouring to the other. To do this we show that T can be recovered from T/σ together with an ultra-metric on the neighbourhood of each end-vertex, induced by the metric on T. The automorphisms of T/σ preserving these ultra-metrics are precisely those induced by automorphisms of T.

T/σ AND THE ULTRA-METRIC. Recall that T/σ has two types of vertices: the set E of **end-vertices** and the set V_σ of vertices in T opposite σ, which we call **0-vertices**. Given an end-vertex e, its neighbourhood in T/σ is a set C of 0-vertices. This set C is a horosphere in T with centre e, and the metric on T (namely the distance between vertices in T) when restricted to C is an **ultra-metric** in the sense that for any three vertices x, y, and z in C, one has $\mathrm{dist}(x,y) \le \max\{\mathrm{dist}(x,z), \mathrm{dist}(y,z)\}$.

Let H denote the horoball whose boundary is C, and let h be any vertex of H. If $\mathrm{codist}(\sigma,h)=r$, then h is at distance r from C, and since we assume thickness, the set of vertices in C at distance r from h forms a ball of diameter $2r$ (using the ultra-metric on C induced by T). Conversely, every ball of diameter $2r$ in C arises in this way from a unique vertex of H, so the vertices of H correspond to balls of even diameter in C. Two such vertices are adjacent when one ball contains the other and their diameters differ by 2. This implies that the datum T together with V_σ is equivalent to the datum T/σ together with the ultra-metric on the neighborhood of each end-vertex. In particular, each automorphism of T/σ preserving these ultra-metrics gives a unique automorphism of T preserving V_σ, and vice-versa; we call such an automorphism a **T-automorphism of T/σ**. We have proved:

(6.1) LEMMA: *The datum T/σ together with the ultra-metric on the neighbourhood of each end-vertex is equivalent to the datum (T, V_σ). Furthermore, there is a canonical isomorphism between the group of T-automorphisms of T/σ, and the group of automorphisms of T preserving σ, characterized by the property that corresponding automorphisms induce the same permutation on V_σ.*

(6.2) PROPOSITION 8: *All uniform 1-balls of T^0 centred at σ are conjugate, and the automorphisms of T stabilizing a uniform 1-ball induce the full symmetric group on its set of boundary vertices.*

Proof: Let X and X' be two indexed σ-colourings of E, both indexed by the same set I. It suffices to show that there is an automorphism of T sending X to X', and fixing I pointwise. By the lemma above it is enough to define this as an automorphism φ of T/σ preserving the ultra-metrics.

Let e be an end having colour i for X, and e' an end having the same colour i for X'. Set $\varphi(e) = e'$. Extend φ to a bijection from the neighbourhood N of e in T/σ to the neighbourhood N' of e' preserving the ultra-metric; this is possible because the semi-homogeneity of T ensures that all these neighbourhoods are isometric. This defines $\varphi(v)$ for each vertex v in N, and there is a unique way of extending φ to the neighbourhood of v preserving colours. Working outwards from e in the tree T/σ, an obvious induction completes the proof. ∎

UNIFORM $(1,1)$-BALLS. Given two adjacent vertices σ and τ of T^0, a **uniform $(1,1)$-ball** with **foci** σ and τ will mean the union of two uniform 1-balls, one centred at σ and containing τ, the other centred at τ and containing σ. The two foci partition the boundary into two **components**, one being the 1-ball centred at σ and punctured by τ, the other being the 1-ball centred at τ and punctured by σ.

In order to study these objects we first partition E into the two subsets $E_{\sigma\tau} = \{e \in E \mid \tau(e) - \sigma(e) = 1\}$ and $E_{\tau\sigma} = \{e \in E \mid \sigma(e) - \tau(e) = 1\}$. These subsets were defined in section 4; the first determines τ as a neighbour of σ, and the second determines σ as a neighbour of τ. Now define a tree $T/(\sigma,\tau)$ whose vertex set is $V_\sigma \cup V_\tau \cup E$. A vertex of V_σ and a vertex of V_τ are joined in $T/(\sigma,\tau)$ if and only if they are joined in T. A vertex of V_σ is joined to no vertex of $E_{\sigma\tau}$, but is joined to the unique vertex of $E_{\sigma\tau}$ that is adjacent to it in T/σ. If e is an end-vertex in $E_{\sigma\tau}$, then its neighbourhood in $T/(\sigma,\tau)$ is the same as its neighbourhood in T/σ, namely the horosphere bounding $e(\sigma)$. Similarly with the roles of σ and τ interchanged.

The discussion earlier implies that the datum T together with V_σ and V_τ is equivalent to the datum $T/(\sigma,\tau)$ together with the ultra-metric on the neighborhood of each end-vertex. In particular, each automorphism of $T/(\sigma,\tau)$ preserving these ultra-metrics determines, and is determined by, a unique automorphism of T preserving V_σ and V_τ. Let us call this a **T-automorphism of $T/(\sigma,\tau)$**. By the discussion preceding (6.1) we have the following lemma.

(6.3) LEMMA: *The datum $T/(\sigma,\tau)$ together with the ultra-metric on the neighbourhood of each end-vertex is equivalent to the datum (T, V_σ, V_τ). Furthermore, there is a canonical isomorphism between the group of T-automorphisms of $T/(\sigma,\tau)$, and the group of automorphisms of T preserving σ and τ, characterized by the property that corresponding automorphisms induce the same permutation on $V_\sigma \cup V_\tau$.*

The point of this lemma is to deal with uniform $(1,1)$-balls. By section 4 a uniform $(1,1)$-ball is equivalent to the following datum: a σ-colouring having $E_{\sigma\tau}$ as one of its parts, and a τ-colouring having $E_{\tau\sigma}$ as one of its parts. Using these colourings we prove an analogue of Proposition 8, stated below as (6.4). The proof is almost identical to that of Proposition 8, but uses the tree $T/(\sigma,\tau)$ in place of T/σ.

(6.4) LEMMA: *All uniform $(1,1)$-balls of T^0 are conjugate. Moreover, the group of automorphisms of T stabilizing a uniform $(1,1)$-ball induces on its boundary the direct product of the full symmetric groups of its two components.*

Proof: Let X and X' be two indexed σ-colourings of E, both indexed by the same set I_1, and extending the partial σ-colouring determined by $E_{\sigma\tau}$, and let i_τ in I_1 denote the "colour" assigned to the elements of $E_{\sigma\tau}$. Thus every e in $E_{\tau\sigma}$ is assigned a "colour" in $I_1 - \{i_\tau\}$. Similarly, let Y and Y' be two indexed τ-colourings of E, both indexed by the same set I_2, and extending the partial τ-colouring in which each element of $E_{\tau\sigma}$ is indexed by i_σ in I_2. It suffices to show that there is an automorphism of T sending X to X', and Y to Y', and fixing I_1 and I_2 pointwise. By (6.3) it is enough to define this as an automorphism φ of $T/(\sigma,\tau)$ preserving the ultra-metrics on the neighbourhoods of the end-vertices.

Let e be an end-vertex in $E_{\tau\sigma}$ to which the σ-colouring X assigns some colour i in $I_1 - \{i_\tau\}$, and let e' be an end to which X' assigns the same colour i. Set $\varphi(e) = e'$. Extend φ to a bijection from the neighbourhood N of e in $T/(\sigma,\tau)$, to the neighbourhood N' of e' preserving the ultra-metric; this is possible because the semi-homogeneity of T ensures that all these neighbourhoods are isometric. Then $\varphi(w)$ is defined for each vertex w in N (N is a subset of V_τ). Each such vertex w is adjacent to vertices in V_σ, and each of these vertices v in turn is adjacent to a unique end-vertex e_v in $E_{\sigma\tau}$. There is a unique way of extending

φ to these end-vertices e_v, and thus to the vertices v in V_σ neighbouring w, so as to preserve colours.

Having defined $\varphi(v)$ and $\varphi(e_v)$, one then defines φ on the neighbourhood of e_v in $T/(\sigma,\tau)$ so that it is sent to the neighbourhood of $\varphi(e_v)$ in $T/(\sigma,\tau)$, preserving the ultra-metric and sending v to $\varphi(v)$. Working outwards from e in the tree $T/(\sigma,\tau)$ in this way, an induction, left to the reader, completes the proof. ∎

(6.5) LEMMA: *Let B be a uniform 2-ball, let R be a uniform $(1,1)$-ball, and let τ denote the focus of R having the same type as the centre of B. Then there exists a uniform 2-ball conjugate to B, centred at τ, and containing R.*

Proof: Let σ denote the centre of B, and let σ_1 be a vertex of B adjacent to σ. Let P denote the uniform $(1,1)$-ball in B having foci at σ and σ_1. By (6.4) there is an automorphism α° of T^0 conjugating P to R with $\alpha^\circ(\sigma)=\tau$. Then $\alpha^\circ(B)$ has the required properties. ∎

(6.6) THEOREM 3: *Let S be a tree isomorphic to T. For each vertex v in S let $S_2(v)$ denote the 2-ball of S centred at v, and choose a uniform 2-ball B_v in T^0 whose centre has the same type as v. Then there is an embedding α of S as a full uniform sub-tree in T^0 (and hence a twinning of T and S) such that $\alpha(S_2(v))$ is conjugate to B_v for each v.*

Proof: Let s be a given vertex in S, and for any integer $n\geq 0$, let S_n denote the ball of radius n in S having centre s. If v is a boundary vertex of S_m then $S_2(v)$ is a subset of S_{m+2}, so to establish the proposition, it will suffice to construct a sequence of type-preserving graph embeddings $\alpha_m\colon S_m\to T^0$ having the following properties:

(i) $\alpha_m(S_m)$ is a uniform m-ball;
(ii) the restriction of α_{m+1} to S_m is α_m;
(iii) for each vertex v on the boundary of S_m, $\alpha_{m+2}(S_2(v))$ is conjugate to B_v.

The direct limit of the maps α_m is an embedding of S in T^0 having the desired properties. We shall obtain the maps α_m by induction on m. The choice of α_0 is trivial, and we now assume that α_n is already constructed, and satisfies: (i) for $m=n$; (ii) for $m=n-1$; and (iii) for $m=n-2$.

To define α_{n+1} we first observe that for every boundary vertex x of S_{n-1}, $\alpha_n(S_2(x)\cap S_n)$ is a uniform $(1,1)$-ball having $\alpha_n(x)$ as one of its foci (the other focus is $\alpha_n(v)$ where v is the unique vertex of S_{n-2} adjacent to x). For each such x, (6.5) allows us to choose a uniform 2-ball A_x centred at $\alpha_n(x)$, conjugate to B_x and containing $\alpha_n(S_2(x)\cap S_n)$, along with a graph isomorphism α_x from $S_2(x)$ to A_x agreeing with α_n on $S_2(x)\cap S_n$. For every boundary vertex y of S_{n+1}, there is a unique x as above at distance 2 from y and we set $\alpha_{n+1}(y) = \alpha_x(y)$. Having thus defined α_{n+1} on the boundary of S_{n+1}, we extend it to all of S_{n+1} by setting $\alpha_{n+1}| \, S_n = \alpha_n$, so property (ii) is satisfied for $m = n$. Since α_x agrees with α_n on $S_2(x) \cap S_n$, property (iii) is satisfied for $m = n-1$. Moreover, for any vertex z on the boundary of S_n its neighbourhood in S is sent by α_{n+1} to a uniform neighbourhood (in some uniform 2-ball) in T^0. This and condition (i) for $m = n$ shows that the neighbourhood of each interior vertex of $\alpha_{n+1}(S_{n+1})$ is uniform. Therefore, by (5.1) condition (i) is satisfied for $m = n+1$. This completes the inductive step, hence the proof of the proposition.

7. Uniform 2-balls

In section 4 we showed that uniform 1-balls centred at σ are equivalent to certain partitions of E called σ-colourings, the parts of the partition being in natural correspondence with the boundary vertices of the 1-ball. In this section we show that uniform 2-balls B centred at σ are equivalent to certain pairs of partitions of E, one of which is a refinement of the other. The parts of the coarser partition correspond to the neighbours of σ in B, and those of the finer partition to the boundary vertices of B. The main goal of the section is a way of constructing such pairs of partitions. If P is one of the parts of a σ-colouring we let τ_P denote the corresponding neighbour of σ in T^0. We now define a σ-**binary colouring** to be a pair of partitions of E satisfying the following conditions:

(i) The first partition is a σ-colouring.
(ii) The second partition is a refinement of the first.
(iii) For each part P of the first partition, the parts of P in the second partition, along with $E - P$, form a τ_P-colouring of E.

By definition a uniform 2-ball centred at σ is a uniform 1-ball S_1 with, for each vertex τ of ∂S_1, a uniform 1-ball $S(\tau)$ centred at τ and containing σ. The following lemma is therefore immediate from (4.2).

(7.1) LEMMA: *There is a one-to-one correspondence between σ-binary colourings of E, and uniform 2-balls B centred at σ. The parts of the first partition correspond to the neighbours of σ in B, and the parts of the second partition to the boundary vertices of B.*

In order to construct σ-binary colourings we avoid dealing with the vertices τ_P of (iii) by introducing a new tree $T/(\sigma - 1)$.

THE TREE $T/(\sigma-1)$. Define $T/(\sigma - 1)$ to be the tree obtained from T by collapsing each horoball $\sigma(e) - 1$ to a point that we identify with e. It has three kinds of vertices: those at codistance 0 or 1 from σ, called 0- or 1-**vertices** accordingly, and all e in E, called **end-vertices**. The neighbourhood in $T/(\sigma-1)$ of a 0-vertex is the same as its neighbourhood in T and consists of 1-vertices. The neighbourhood of a 1-vertex consists of 0-vertices and a single end-vertex, and the neighbourhood of an end-vertex e is the horosphere bounding $\sigma(e)$. Given a 1-vertex y, let $\varepsilon(y)$ denote the unique end vertex adjacent to y, and given a 0-vertex v, let E_v denote the set of $\varepsilon(y)$ as y ranges over the neighbours of v–this is the same as the set of end-vertices adjacent to v in T/σ.

To construct a pair of partitions, one of which is a refinement of the other, we use a pair of surjective maps $E \xrightarrow{\chi} J \xrightarrow{\pi} I$. The first partition is given by the fibres of $\pi\chi$ and its refinement by the fibres of χ. Note that the cardinality of I equals the valency of a 0-vertex, and the cardinality of each fibre of π is one less than the valency of a 1-vertex. For example, if T is homogeneous of valency $q+1$ then $|I| = q+1$, $|J| = q(q+1)$ and π is a q-to-1 map. Conditions for this set-up to give a σ-binary colouring will use the following terminology.

In the context of the map $J \xrightarrow{\pi} I$ a subset of J will be called a **line** if it contains exactly one element in each fibre of π. The essential example of this is the set $\chi(E_v)$ for any 0-vertex v in $T/(\sigma-1)$; it will be called the χ-**line** (or simply the **line**) associated to v. It is the same as the set of all $\chi\varepsilon(y)$ as y ranges over the 1-vertices neighbouring v. Our conditions are now as follows:

(S1) *for each 0-vertex v, the map $\pi\chi$ restricted to E_v is a bijection from E_v to* I (this is equivalent to saying that the partition of E given by the fibres of $\pi\chi$ is a σ-colouring);

(S2) *for each 1-vertex y, every element of J whose image by π is different from that of $\chi\varepsilon(y)$ belongs to one and only one line $\chi(E_v)$ where v is a 0-vertex neighbouring y.*

For a 1-vertex y, let $L(y)$ be the set of lines $\chi(E_v)$ as v ranges over the 0-vertices adjacent to y, as in (S2). These lines all have the point $j = \chi\varepsilon(y)$ in common and, as v varies, the sets $\chi(E_v) - \{j\}$ partition the set $J - \pi^{-1}(\pi(j))$. Such a set of lines we call a **pencil**, centred at j, or if we wish to be more precise about χ and y, a χ-pencil, denoted by $L\chi(y)$.

Given $E \xrightarrow{\chi} J \xrightarrow{\pi} I$ satisfying (S1) and (S2), consider the partitions $\{X_i|\, i \in I\}$ and $\{X_j|\, j \in J\}$ of E defined by:

$$X_i = \{e \in E|\, \pi\chi(e) = i\} \quad \text{and} \quad X_j = \{e \in E|\, \chi(e) = j\}.$$

(7.2) PROPOSITION 9: *If $E \xrightarrow{\chi} J \xrightarrow{\pi} I$ satisfies* (S1) *and* (S2), *then the pair of partitions above forms a σ-binary colouring of E.*

Proof: As remarked above, (S1) implies that $\{X_i | \ i \in I\}$ is a σ-colouring. Let τ_i denote the neighbour of σ corresponding to X_i. We must show that the partition $\{E - X_i, X_j | \ j \in \pi^{-1}(i)\}$ is a τ_i-colouring of E. To do this, let y be a 0-vertex of T/τ_i and note that y is a 1-vertex of $T/(\sigma - 1)$. The neighbourhood of y in T/τ_i contains a unique end-vertex e for which $\sigma(e) - \tau_i(e) = 1$; this e coincides with $\varepsilon(y)$ and lies in $E - X_i$. All other neighbours e of y in T/τ_i satisfy $\tau_i(e) - \sigma(e) = 1$ and hence lie in X_i. If v is the neighbour of y in the half-apartment (ye), then v is a 0-vertex of $T/(\sigma - 1)$ and $e \in E_v \cap X_i$. Conversely, if $e \in E_v \cap X_i$ for a 0-vertex v neighbouring y, then $v \in \tau_i(e)$, so e is a neighbour of y in T/τ_i. Therefore, the neighbours of y in T/τ_i consist of $\varepsilon(y)$ and all e in $E_v \cap X_i$ as v ranges over the 0-vertices of $T/(\sigma - 1)$ neighbouring y. By (S2) the $\chi(e)$ range over all j in $\pi^{-1}(i)$, so $\{E - X_i, X_j | \ j \in \pi^{-1}(i)\}$ is a τ_i-colouring of E. ∎

(7.3) COROLLARY: *Any pair of maps (χ, π) satisfying (S1) and (S2) determines a uniform 2-ball centred at σ. Moreover, if (χ, π) and (χ', π') determine conjugate 2-balls B and B', then there is an automorphism φ of $T/(\sigma - 1)$ sending the pair of partitions determined by $(\pi'\chi', \chi')$ to the pair determined by $(\pi\chi, \chi)$.*

Proof: The first statement is immediate from (7.1) and (7.2). The second statement follows from the definitions of conjugacy and $T/(\sigma - 1)$.

8. Non-isomorphic twin trees

In this section we look in more detail at the construction of uniform 2-balls. Our main purpose is to prove the following theorem, where the **cardinality** of a tree means the cardinality of its set of vertices.

(8.1) THEOREM 4: *If T is a thick, semi-homogeneous tree of cardinality α, there are 2^α conjugacy classes of uniform 2-balls of either type in T^0.*

(8.2) COROLLARY: *A thick, semi-homogeneous tree of cardinality α admits 2^α isomorphism classes of twinnings, and among these 2^α have trivial automorphism group.*

This corollary is an immediate consequence of Theorem 4, by applying (6.6). The question of twin trees having large automorphism groups will be dealt with in a later paper.

As before we fix the vertex σ of T^0, and use the tree $T/(\sigma - 1)$. The connected components of $T/(\sigma - 1) - E$ we call $(\sigma, 1)$-**components**. If Y is one of them, we write $E(Y)$ for the set of end-vertices joined to Y, i.e. all $\varepsilon(y)$ as y ranges over the 1-vertices of Y.

Given surjective maps $E(Y) \xrightarrow{\chi} J \xrightarrow{\pi} I$ we refer to conditions (S1) and (S2) of section 7 with the understanding that the 0- and 1-vertices in these conditions are restricted to vertices of Y. Let us now fix I, J and a surjective map $\pi: J \to I$ having the appropriate cardinalities: namely card I equals the valency of a 0-vertex, and the cardinality of each fibre of π is one less than the cardinality of a 1-vertex. We shall call a map χ from E to J, or from $E(Y)$ to J, **admissible** if the pair (χ, π) satisfies (S1) and (S2).

(8.3) LEMMA: *A map χ: $E \to J$ is admissible if and only if its restriction to $E(Y)$ is admissible, for all $(\sigma, 1)$-components Y.*

Proof: This follows from the definition, and the fact that if v is any 0-vertex of Y then E_v is a subset of $E(Y)$. ∎

Let X denote the tree obtained from $T/(\sigma - 1)$ by collapsing each $(\sigma, 1)$-component to a point; it has two subsets of vertices: E, and the set of $(\sigma, 1)$-components. The neighbourhood of a $(\sigma, 1)$-component Y is $E(Y)$, and an end-vertex e has one neighbour for each 1-vertex on the horosphere $\partial\sigma(e)$.

(8.4) LEMMA: *The tree X is homogeneous of valency α, where α is the cardinality of* Vert T.

Proof: Let e be an end vertex, and Y a $(\sigma, 1)$-component. The valency of e in X is the cardinality of $\partial\sigma(e)$, which by (3.6) is α. The valency of Y in X is the cardinality of the set of 1-vertices in Y. This in turn equals the cardinality of Vert T, because if T has valencies (a, b) then Y has valencies $(a, b - 1)$, and we have assumed $a, b > 2$. ∎

A uniform 2-ball is obtained from an admissible map χ, and we shall prove Theorem 4 by associating to χ a rather crude invariant that is a subset of E, and hence of Vert X. We call a point e of E **neutral** for χ if, as y runs over the set of 1-vertices in $T/(\sigma - 1)$ neighbouring e, the pencils $L_\chi(y)$ (defined in section 7 following condition (S2)) depend only on e and not on y. Let $E_0(\chi)$ denote the set of points e in E that are neutral for χ.

(8.5) LEMMA: *Let χ and χ' be admissible maps from E to J, and let B and B' be the associated 2-balls given by (7.3). If B and B' are conjugate, then there is an automorphism of X sending $E_0(\chi)$ to $E_0(\chi')$.*

Proof: When B and B' are conjugate, (7.3) gives an automorphism of $T/(\sigma - 1)$, and hence of X, having the desired property. ∎

A crucial fact in the proof of Theorem 4 is the following lemma whose proof will be given later.

(8.6) LEMMA: *Given any subset E_0 of E, there is an admissible map χ with $E_0(\chi) = E_0$.*

To show that there are 2^α conjugacy classes of uniform 2-balls we use (8.6) and (8.5) and do some counting. For us, a **binary labelling** of a set V will be an assignment of either $+$ or $-$ to each element of V. If V is the set of vertices of a tree A, then we say that two labellings of V are **isomorphic in** A if there exists an automorphism of A sending one of them to the other. Now, coming back to E, we associate to any subset E_0 of E the binary labelling that assigns $+$ to each element of E_0 and $-$ to each element of $E - E_0$.

(8.7) LEMMA: *With the notation above, where A is a homogeneous tree of valency α, an infinite cardinal, and V the set of vertices of one type in A, there are 2^α binary labellings of V that are non-isomorphic in A.*

Proof: Define a **pointed labelling** to mean a labelling together with a chosen vertex, which for our purposes will be a vertex *not* in V. An isomorphism of two pointed labellings has the obvious meaning, namely an isomorphism of the two labellings induced by an automorphism of A sending the chosen vertex of one to that of the other.

STEP 1: *Let $\beta \leq \alpha$ be an infinite cardinal. If there are β non-isomorphic pointed labellings, then there are 2^β.* Fix one vertex $x \notin V$, and take a set B of neighbours of x having cardinality β. Let $\{L_b\}_{b\in B}$ denote a set of non-isomorphic pointed labellings, different from the labelling with $+$ everywhere (here we use the fact that β is infinite). Given $b \in B$ and any neighbour $y \neq x$ of b, let A_y be the tree spanned by the set of vertices $\{a \in A|\ \mathrm{dist}(a,y) < \mathrm{dist}(a,b)\}$; it is isomorphic to A. Now to each subset S of B assign a pointed labelling M_S having the following properties: the distinguished point is x; the points of S carry the label $-$; for b and y as above, the labelling of A_y induced by M_S is isomorphic to L_b with distinguished point y; finally, all vertices which belong neither to S nor to any A_y carry the label $+$. Now, the conditions imposed on $\{L_b\}_{b\in B}$ imply that, if S and S' are two subsets of B, then any isomorphism from M_S to $M_{S'}$ fixes each point of B, and since S is precisely the subset of B whose elements are labelled by $-$ in M_S we see that all M_S are pairwise non-isomorphic in A. This gives 2^β pointed labellings with x as chosen point.

STEP 2: *There are 2^α non-isomorphic pointed labellings.* Let β be the number of isomorphism classes of pointed labellings; β is clearly infinite. Since $\beta < 2^\beta$, Step 1 implies that $\alpha < \beta$. Therefore there exist α non-isomorphic pointed labellings, and hence by Step 1 again, 2^α.

STEP 3: *There are 2^α non-isomorphic labellings of A.* This is immediate from Step 2 since the number of vertices is α, and $2^\alpha/\alpha = 2^\alpha$. ∎

The proof of Theorem 4: This theorem follows from (8.4)–(8.7), and it only remains to prove (8.6). The proof of (8.6) uses a further result (8.8) that involves the construction of an admissible map $\chi\colon E(Y) \to J$ for a given $(\sigma, 1)$-component Y. That construction is presented later in order to avoid disrupting the main flow of the argument proving Theorem 4.

Suppose, however, that we are given an admissible map $\chi\colon E(Y) \to J$ and some e in $E(Y)$. There is a unique 1-vertex y of Y joined to e, and if v is a 0-vertex joined to y then E_v is a subset of $E(Y)$. Therefore $L_\chi(y)$ is defined, and we call it the **χ-pencil of Y at** e.

Proof of (8.6)—assuming (8.8): In the tree X we start with a single end-vertex e_0, and let X_n denote the vertices of X at distance n from e_0. All end-vertices are at even distance from e_0 in X, and all $(\sigma, 1)$-components Y are at odd distance from e_0. We define $\chi(e_0) \in J$ arbitrarily, and work inductively outwards from e_0. If χ is defined on $X_{\leq 2n}$ in such a way that $\chi|E(Y)$ is admissible for all Y in $X_{\leq 2n-1}$, then given e in $X_{\leq 2n-2}$ we have χ-pencils at Y defined for all Y adjacent to e. Our induction hypothesis is now as follows. Suppose χ is defined on X_{2n} in such a way that:

(i) $\chi|E(Y)$ is admissible for all Y in $X_{\leq 2n-1}$;
(ii) if $e \in X_{\leq 2n-2}$ then as Y ranges over the vertices adjacent to e, the χ-pencils at Y agree if $e \in E_0$ and disagree if $e \notin E_0$.

Given e at distance $2n$ from e_0, we intend to show that when we extend χ from X_{2n} to X_{2n+2} we are free to choose whether or not e will be a neutral vertex for χ. Let Y_0 denote the unique vertex of X adjacent to e and at distance $2n-1$ from e_0. By induction the χ-pencil of Y_0 at e is already defined; call it L_0. If $Y \neq Y_0$ is any other vertex adjacent to e, then $E(Y) \cap X_{\leq 2n} = \{e\}$, and by (8.8) there exists an admissible map $\chi_Y\colon E(Y) \to J$ such that $\chi_Y(e) = \chi(e)$, and such that the χ_Y-pencil of Y at e is any desired pencil centred at $\chi(e)$. There are at least two pencils centred at a given element of J (when T is homogeneous of valency three—the smallest case we consider—there are exactly two), so we may choose the χ_Y-pencil at e to be different from L_0 if we wish. When $e \in E_0$ we choose it to be L_0, and when $e \notin E_0$ we choose it to be different from L_0 for at least one Y adjacent to e. Extend χ by setting $\chi|E(Y) = \chi_Y$. Doing this for all Y adjacent to e, and for all e at distance $2n$ from e_0, gives an extension of χ to X_{2n+2} satisfying the induction hypothesis. By (8.3) we obtain an admissible map $\chi\colon E \to J$ for which the neutral vertices are precisely those in E_0. ∎

(8.8) LEMMA: *Let Y be a $(\sigma,1)$-component, let $e_0 \in E(Y)$, and $j_0 \in J$. Given a pencil L_0 centred at j_0, there exists an admissible map χ: $E(Y) \to J$ such that L_0 is the χ-pencil of Y based at e_0.*

Proof: We prove this lemma below, along with Lemma (8.9), using a construction of admissible maps on $E(Y)$.

A SPECIAL CASE. Suppose T is a homogeneous tree of valency $q+1$, where q is the order of some projective plane. Let Π be a projective plane of order q (in other words with $q+1$ points per line), and let ω be a point of Π. The set $\Pi - \{\omega\}$ has cardinality $q(q+1)$; we may identify it with J in such a way that the lines of Π through ω, punctured by ω, become the fibres of the map $J \xrightarrow{\pi} I$. Such an identification we call a (Π,ω)-**structure** for J.

For $j \in J$, let $p(j)$ denote the corresponding point of Π, and for $p \in \Pi$, let $j(p)$ denote the corresponding element of J. For any $j_0 \in J$, the set of lines through $p(j_0)$, excluding the line $p(j_0)\omega$, or rather their images under j, form a pencil centred at j_0, in the sense defined in section 7 following (S2). We call this the Π-**pencil** centred at j, and denote it by the symbol Π_j.

(8.9) LEMMA: *Suppose T is homogeneous, and suppose we are given a (Π,ω)-structure for J. Let Y be a $(\sigma,1)$-component, let $e_0 \in E(Y)$, and $j_0 \in J$. Then there is an admissible map χ: $E(Y) \to J$ having the following properties:*

(i) *$\chi(e_0) = j_0$;*

(ii) *for each $e \in E(Y)$ the χ-pencil of Y based at e is $\Pi_{\chi(e)}$.*

Proof of (8.8) and (8.9): Consider the sub-tree of $T/(\sigma-1)$ spanned by $Y \cup E(Y)$. Given $e_0 \in E(Y)$, the other $e \in E(Y)$ are at even distance from e_0; let $E(Y)_n$ denote those at distance $2n$. Starting at e_0, we define $\chi(e_0) = j_0$, and work outwards by induction on n. Let y_0 be the unique 1-vertex of Y adjacent to e_0. There are no end-vertices at distance 2 from e_0, so the first step of our induction starts with $E(Y)_2$. This set is the disjoint union of all $E_v - \{e_0\}$ as v ranges over the 0-vertices adjacent to y_0. We may therefore define χ on $E(Y)_2$ such that, as v ranges over these 0-vertices, the lines $\chi(E_v)$ range over the lines of a pencil. This pencil must be centred at j_0, but is otherwise arbitrary. For (8.8) choose it to be L_0. For (8.9) choose it to be Π_{j_0}.

Now for $n \geq 2$ assume that χ has been specified on $E(Y)_{\leq n}$ in such a way that (S1) and (S2) are satisfied. For (8.9) we further assume that the lines and pencils satisfying (S1) and (S2) are members of Π. Each e in $E(Y)_{n+1}$ lies in a unique E_v for some 0-vertex v at distance $2n$ from e_0. This vertex is adjacent to a 1-vertex y at distance $2n-1$ from e_0, and y in turn is adjacent to a 0-vertex

v_0 at distance $2n-2$ from e_0. Since E_{v_0} is a subset of $E(Y)_n$, our induction hypothesis specifies χ on E_{v_0}; moreover $\chi(E_{v_0})$ is a line, and for (8.9) it is a line of Π. Let V_y denote the set of 0-vertices v adjacent to y and different from v_0. As v ranges over V_y the sets $E_v - \{\varepsilon(y)\}$ are disjoint. We may therefore define χ on each such E_v so that $\chi(E_v)$ is a line and so that these lines, along with $\chi(E_{v_0})$, form a pencil centred at $j = \chi\varepsilon(y)$. For (8.8) the choice of pencil is indifferent; for (8.9) we choose it to be Π_j; this is possible because our induction hypothesis guarantees that $\chi(E_{v_0})$ is a line of Π. Doing this for all vertices y at distance $2n-1$ from e_0 defines χ on $E(Y)_{n+1}$ such that (S1) and (S2) are satisfied on $E(Y)_{\leq n+1}$. This completes the inductive step. ∎

9. Uniform 2-balls and projective planes

In this section B will denote a uniform 2-ball with centre σ. Any vertex v of T at codistance 0 from σ determines a set of vertices in ∂B at codistance 2 from v; we call this set σ^v and refer to it as a **line** of ∂B. If τ is a neighbour of σ in B, then the neighbours of τ form a subset of $\partial B \cup \{\sigma\}$ that we also call a **line**. Such lines meet the lines of ∂B in exactly one point. To summarize, we have two types of lines in $\partial B \cup \{\sigma\}$: those containing σ, and those of the form σ^v lying entirely in ∂B.

The lines of ∂B are *mutatis mutandis* the χ-lines of section 7. To explain this recall that uniform 2-balls may be regarded as σ-binary colourings using a pair of maps $E \xrightarrow{\chi} J \xrightarrow{\pi} I$. Each vertex of ∂B corresponds naturally to an element of J, and each neighbour of σ in B to an element of I. Under this correspondence the line σ^v is none other than the χ-line $\chi(E_v)$. The term **pencil** of χ-lines in section 7 applies equally well with χ-lines replaced by lines of ∂B.

(9.1) LEMMA: *Any two vertices of ∂B lie on a common line.*

Proof: If the two vertices have a neighbour in common (necessarily a neighbour of σ) this is immediate. If not, then let j_1 and j_2 be the boundary vertices concerned, and let σ^u be a line containing j_1. The points of σ^u are in natural bijective correspondence with the neighbours of u in the tree T. Let w be the neighbour of u in T corresponding to j_1. By the results of section 7 on χ-lines, applied *mutatis mutandis* to lines of ∂B, the lines σ^v, as v ranges over the neighbours of w at codistance 0 from σ, form a pencil. Therefore one of them contains j_2. ∎

Our construction of uniform 2-balls in section 8 can lead to some rather tangled systems of lines σ^v as v ranges over the vertices of T opposite σ. Two boundary vertices not having a common neighbour in B could be contained in many different lines, but we now consider the extreme case where they lie in a *unique* line. If, in this case, T is homogeneous of finite valency $q+1$, then $\partial B \cup \{\sigma\}$ with its lines is a projective plane. We now define B to be **flat** whenever $\partial B \cup \{\sigma\}$, with its lines, is a projective plane (finite or infinite).

In order to construct flat 2-balls we use (8.9), the essential features of that result being as follows. When T is a homogeneous tree of valency $q+1$, the set I has cardinality $q+1$, and J has cardinality $q(q+1)$. If q is the order of a projective plane Π and ω is a point of Π, then in section 8 we defined a (Π,ω)-*structure* for J. This was an identification of $\Pi - \{\omega\}$ with J in such a way that the lines of Π through ω, punctured by ω, become the fibres of the map $J \xrightarrow{\pi} I$. By (8.9) there exists an admissible map χ such that every $\chi(E_v)$ is a line of Π, and every line of Π not containing ω arises in this way. If B denotes the resulting 2-ball, then under the natural identification of ∂B with J the lines of ∂B are the lines of Π not containing ω. We say that B has a (Π,ω)-**structure.** To summarize:

(9.2) LEMMA: *Let T be homogeneous of valency $q+1$ (finite or infinite), let Π be a projective plane of order q, and let ω be a point of Π. Then there exists a uniform 2-ball having a (Π,ω)-structure.*

Proof: This is immediate from (8.9) along with (8.6) using E_0 as the null-set.

Remark on uniqueness: Although we do not prove it here, the 2-ball in (9.2) is uniquely determined up to conjugacy. This can be shown by using (6.1) and verifying that for any $(\sigma,1)$-component Y the restriction of χ to $E(Y)$ is uniquely determined up to an automorphism of Y.

The following theorem is immediate from (9.2) and (6.6).

(9.3) PROPOSITION 10: *Let S be a tree isomorphic to T of valency $q+1$, and for each vertex s in S let $(\Pi,\omega)_s$ be a punctured projective plane as above. Then there is a twinning of T and S such that for each s in S the uniform 2-ball centred at s in S admits a $(\Pi,\omega)_s$ structure.*

This proposition says nothing about uniqueness, even if all the $(\Pi,\omega)_s$ are identical; in fact, there exist non-isomorphic twin trees whose 2-balls all have the same flat structure. For instance, in the example from our earlier paper [RT]—a twin tree for the group GL_2 over a ring of Laurent series $k[t,t^{-1}]$—all 2-balls are flat, the projective planes in this case being isomorphic to the Desarguesian plane over the field k. Non-isomorphic examples with the same flat 2-balls arise from various Kac–Moody groups over k. We shall prove this fact in a sequel to the present paper.

Kac–Moody groups (of rank 2) also provide other interesting illustrations of the concepts introduced here. Such a group operates on a twin tree, and for each tree of the pair its action is transitive on the set of vertices of a given type. Thus in each tree there are two isomorphism classes of 2-balls (actually, there is also an automorphism of the whole situation permuting the two trees). Now, given an integer $m > 1$, there exists a Kac–Moody group such that the 2-balls of one class are flat (they are projective planes over the ground field), whereas for a 2-ball B of the second class any set of m points in ∂B, no two of which are adjacent to a common point of B, determine a unique line.

References

[F] D. G. Fon-Der-Flaass, *A combinatorial construction for twin trees,* European Journal of Combinatorics **17** (1996), 177–189.

[J] P. Johnson, *Non-thick twin trees,* The Bulletin of the London Mathematical Society, to appear.

[RT] M. A. Ronan and J. Tits, *Twin trees I,* Inventiones Mathematicae **116** (1994), 463–479.