Enseign. Math. (2) 72 (2026), 191–213
DOI 10.4171/LEM/1097

Sidon sequences and nonpositive curvature

Sylvain Barré and Mikaël Pichot

Abstract. A sequence $a_0 < a_1 < \cdots < a_n$ of nonnegative integers is called a Sidon sequence if the sums of pairs $a_i + a_j$ are all different. In this paper we construct CAT(0) groups and spaces from Sidon sequences. The arithmetic condition of Sidon is shown to be equivalent to nonpositive curvature, and the number of ways to represent an integer as an alternating sum of triples $a_i - a_j + a_k$ of integers from the Sidon sequence, is shown to determine the structure of the space of embedded flat planes in the associated CAT(0) complex.

Mathematics Subject Classification 2020: 20F65 (primary); 57M07 (secondary).

Keywords: nonpositive curvature, Sidon sequences, ring puzzles.

1. Introduction

A sequence $a_0 < a_1 < \cdots < a_n$ of nonnegative integers is called a Sidon sequence if the sums of pairs $a_i + a_j$ (for $i \leq j$) are pairwise different. Here we assume that $a_0 = 0$. An example is $0, 2, 7, 8, 11$.

Every Sidon sequence $a_0 = 0 < a_1 < \cdots < a_n$ can be extended by letting a_{n+1} be the smallest integer such that $a_0 < \cdots < a_{n+1}$ satisfies the condition of Sidon. Starting at $a_0 = 0$, this defines an infinite Sidon sequence

$$0, 1, 3, 7, 12, 20, 30, 44, 65, 80, 96, \ldots$$

called the Mian–Chowla sequence.

Sidon, in connection with his investigations on Fourier series, was interested in estimates of the number of terms not exceeding N a Sidon sequence can have. In general, there are at most $\ll N^{1/2}$ terms in the interval $[0, N]$, and the Mian–Chowla sequence, for example, which verifies $a_n \leq n^3$, contains $\gg N^{1/3}$ terms in this interval (see [13]). Erdös conjectured that for every $\varepsilon > 0$, there should exist denser infinite Sidon sequences with $\gg N^{1/2-\varepsilon}$ terms. This problem is well studied. We quote here the well-known theorem of Ruzsa [11] which constructs sequences with $\gg N^{\gamma+o(1)}$ terms, where $\gamma = \sqrt{2} - 1$. For finite sequences, a famous theorem of Singer [8, 12] in

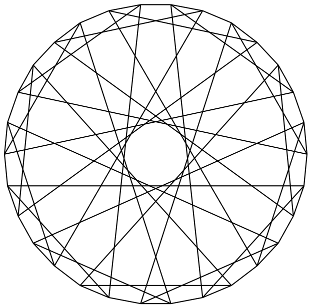

Figure 1
The Mian–Chowla sequence 0, 1, 3, 7 (for $n = 3$) viewed as a nonpositive curvature condition modulo 15 via the link condition of Gromov.

projective geometry provides, for any fixed $\varepsilon > 0$, Sidon sequences with $> N^{1/2}(1-\varepsilon)$ terms for any large N. This uses the fact that the asymptotic ratio of consecutive primes is 1, can be improved by using a sharper estimate for $p - p'$, where p and p' are consecutive primes (see [9, §2.3]). For the purpose of the present paper, we shall require the stronger condition that the sums of pairs $a_i + a_j$ $(i \leq j)$ are all different modulo some integer $N \geq 2$. Such a sequence is called a *Sidon sequence modulo N*. Clearly, every finite Sidon sequence is a Sidon sequence modulo N for every large enough integer N.

The present paper describes a connection between Sidon sequences and nonpositive curvature. The basic idea is to associate with a Sidon sequence $a_0 < a_1 < \cdots < a_n$ modulo N a group G acting geometrically a 2-complex X, whose geometric properties depend on the arithmetic properties of the sequence $a_0 < a_1 < \cdots < a_n$, as follows:

(1) the sum of pairs condition of Sidon ensures that the ambient space X is nonpositively curved;

(2) the alternating sums of triples $a_i - a_j + a_k$ of elements of the Sidon sequence regulates the structure of embedded flat planes in the space X.

We shall refer to [7] for nonpositive curvature and CAT(0) spaces.

In the simplest situation, we shall introduce for every Sidon sequence $a_0 = 0, \ldots, a_n$ modulo N a CAT(0) complex X of dimension 2 called the modular complex of the Sidon sequence (see Section 7.3) whose automorphism group is transitive on the vertex set. The common vertex link of this complex is a graph S, defined in Section 2, which depends on the Sidon sequence in a simple manner as follows: S is the quotient, modulo $2N$, of the graph structure on $\mathbb{Z}$ in which $n+1$ edges are issued from every even vertex with an increment of $2a_r - 1$, for every $r \in [0, \ldots, n]$. It is shown in

Section 2 that the link condition of Gromov (which ensures that X is nonpositively curved) is *equivalent* to the standard arithmetic condition of Sidon.

These complexes appear to be new geometric objects for the most part. For an example, Figure 1 represents the link of the modular complex (modulo 15) associated with a Mian–Chowla sequence of length 4.

We shall describe more examples (old and new) in Section 7 below. In general, in the construction of groups and complexes "of intermediate rank", one is interested in the structure of the space of embedded flat planes, and especially the existence of periodic flats in the nonhyperbolic case (i.e., the flat closing conjecture). Our main theorem is that the structure of embedded flat is entirely determined by a tiling problem associated with alternating sums of triples of the original Sidon sequence. We have given several other construction methods of groups of intermediate rank in the past; in Section 7, it is shown in particular that the odd Moebius–Kantor complex (which has the mesoscopic rank property) can be recovered by the method of Sidon.

Before we state our theorem, we indicate how the structure of embedded flat planes can be related to a Sidon sequence. In some CAT(0) 2-complexes, including the ones defined in the present paper, the space of embedded flat planes can encoded by means of a ring puzzle problem [2]. A ring puzzle is a tessellation of the Euclidean plane obtained by using a (finite) set of shapes which are constrained locally by a set of conditions around every vertex. In this context, the local conditions are appropriately described by a (finite) set of length 2π circles called the rings. Every ring is labelled and specifies unambiguously the shapes that can be used in the neighbourhood of the given vertex. Here we require ring puzzles in which all the shapes are equilateral triangles.

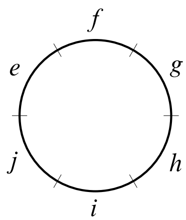

FIGURE 2
Local geometry in a ring puzzle.

We are given a set of $n+1$ equilateral triangles, labelled by integers in $\{0,\ldots,n\}$, together with a set of rings (see Figure 2), where $e, f, g, h, i, j \in \{0,\ldots,n\}$ specify which triangle label can be used around a vertex. The local condition ensures that the labels on the triangles and links are consistent. We shall in fact consider slightly more elaborate triangle ring puzzles in which the tessellation edges, and correspondingly, the link vertices, are labelled. This type of puzzle arises in complexes constructed from Sidon sequences, which are endowed naturally with both an edge and a face labelling. The edge labels are elements in $\mathbb{Z}/3\mathbb{Z}$ satisfying the property that every triangle contains three distinct labels, and every vertex is incident to edges which omit exactly one label. Such a labelling is uniquely determined from the choice of labels on a single triangle, and will be fixed implicitly throughout the paper. We will show that the complexes themselves do not depend (up to isomorphism) on the choices of edge labels made during the constructions, due to the existence of "polarities" (as shown in Section 5) between them.

Given these remarks, the following theorem describes a connection between Sidon sequences, nonpositive curvature, and the structure of flats in the associated complexes, in the case of modular complexes. We refer to Theorem 4.1 for a more general statement.

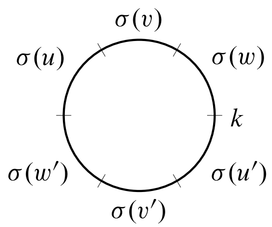

FIGURE 3
Local geometry in a Sidon ring puzzle.

Theorem 1.1. *Let $a_0 = 0 < a_1 < \cdots < a_n$ be a Sidon sequence modulo N and let X be the associated modular complex, as defined in Section 7.3. The following conditions are equivalent for a labelled equilateral triangle tessellation Π of the Euclidean plane:*

(a) *Π embeds in X in a label preserving way;*

(b) *Π is a solution to a triangle ring puzzle problem with rings of the form given in Figure 3 for every $k \in \mathbb{Z}/3\mathbb{Z}$ and every triple (u, v, w) and (u', v', w') of elements of $\{a_0^j, \ldots, a_n^j\}$ with equal alternating sum*

$$u - v + w = u' - v' + w' \bmod N$$

such that $u \neq v \neq w$, $u' \neq v' \neq w'$ and $(u, v, w) \neq (u', v', w')$.

In this theorem, the map $\sigma\colon \{a_0, \ldots, a_n\} \to \{0, \ldots, n\}$ is the increasing bijection, which gives labels to the faces in X by elements in $\{0, \ldots, n\}$. The edge labels in X can remain implicit, since they are uniquely determined by a single triangle of X. They correspond to labels on vertices in the rings, using k and $k + 1$ alternately ($k \in \mathbb{Z}/3\mathbb{Z}$).

We note that this only limits the rings that can be used around a vertex in a puzzle (namely, only rings with vertex labels in $\{k, k+1\}$ can be positioned around a vertex having these edge labels) and does not contribute additional complexity to the solution space of the puzzle problem.

Informally speaking, Theorem 1.1 reduces the description of embedded flat planes in the resulting complexes to the resolution of a paving problem in $\mathbb{R}^2$. The latter problem is purely 2-dimensional in nature and it has a number theoretic component to it, since the rings are obtained from the representations of a modular integer as an alternating sum of triples from the given Sidon sequences.

The construction offers some freedom. Thus, one may use more than one Sidon sequence. It is natural in our setting to formulate a result with three possibly different sequences, all of the same length, say $n+1 \geq 2$, and twist the construction by using permutations in the symmetric group S_{n+1} (twisting the σ's). These generalizations, which may appear to be technical at first sight, enable us to include interesting complexes of intermediate rank including, as shown in Section 7.6, the odd Moebius–Kantor complex, which can be viewed as a particular case of the Sidon construction. This particular construction answers the question of finding an analog of the Pauli construction for odd Moebius–Kantor complexes, and leads to uniqueness of the odd Moebius–Kantor complex up to isomorphism (see Theorem 7.10). We defer the formulation of a general theorem "with parameters" to Section 4.1.

The paper is structured as follows. The proof of Theorem 1.1 (in its general formulation as stated in Theorem 4.1) has been divided into five steps, occupying Section 2 to Section 6. It is shown in Section 2 that the nonpositive curvature condition on links (namely, having girth at least 2π) is equivalent to the arithmetic condition of Sidon. A new class of complexes, called completely homogeneous complexes, is introduced in Section 3. The modular complexes are example of completely homogeneous complexes. In Section 4 we prove a version of Theorem 1.1 with parameters. Some of these parameters are in fact redundant, and the redundancy is removed in Section 5 by showing that the complexes associated with Sidon sequences admit sufficiently many "polarities" between them; in particular, the various constructions in Section 4 are in fact pairwise isomorphic. We prove the equivalence stated in Theorem 1.1 (and its generalizations) in Section 6. The last section of the paper, Section 7, concludes with some examples and applications.

2. Sidon sequences

We refer to [9, 10] for surveys on finite and infinite Sidon sequences; here we begin the proof of Theorem 1.1 (in its general version with parameters, as stated in Theorem 4.1).

The first step in the proof is to relate the arithmetic condition of Sidon to a geometric condition on spaces, i.e., nonpositive curvature via the link condition of Gromov (see [7]).

Let $a_0 = 0 < a_1 < \cdots < a_n$ be a sequence of positive integers. We consider the graph S on $\mathbb{Z}$ in which $n + 1$ edges are issued from every even vertex with an increment of $2a_r - 1$ for every $r \in [0, \ldots, n]$. We write $S = S(a_1, \ldots, a_n)$ to indicate the dependency in the sequence of integers.

Theorem 2.1. *Let $a_0 = 0 < a_1 < \cdots < a_n$ be an increasing sequence of positive integers and let $S = S(a_1, \ldots, a_n)$. Let $N \geq 1$ be an integer. The following are equivalent:*

(1) *the sequence $a_0, a_1, \ldots, a_n$ is a Sidon sequence modulo N;*

(2) *the quotient of S by the action of $2N\mathbb{Z}$ by translations has girth at least 6.*

Proof. Note that the condition of Sidon, that the sums of pairs

$$a_i + a_j \quad (i \leq j)$$

are all different modulo N, is verified if and only if the following alternating sums do not vanish

$$a - b + c - d \neq 0 \bmod N$$

for every quadruple $a \neq b \neq c \neq d \neq a$ chosen from the set $\{a_0, a_1, \ldots, a_n\}$. We write S_N for the quotient graph of S by the translation action of $2N\mathbb{Z}$. By definition, S_N is a bipartite graph.

Suppose that

$$a - b + c - d \neq 0 \bmod N$$

for every $a \neq b \neq c \neq d \neq a$ as above. We claim that S_N has no repeated edges. Indeed, a repeated edge corresponds arithmetically to the existence of distinct integers $a \neq b$ such that $2a - 1 = 2b - 1 \bmod 2N$. Then

$$a - b + a - b = 0 \bmod N.$$

This contradicts the fact that the condition $a \neq b \neq c \neq d \neq a$ is verified when $a \neq b$, $c = a$, $d = b$.

Next we claim that S_N contains no square. Observe that a square in S_N corresponds arithmetically to the choice of four elements a, b, c, d in the Sidon sequence, such that

$$a \neq b \neq c \neq d \neq a,$$

which indicates that consecutive edges in the square are distinct, and

$$(2a-1)-(2b-1)+(2c-1)-(2d-1)=0 \bmod 2N,$$

which indicates that the edges form a square. However, the equality

$$2(a-b+c-d)=0 \bmod 2N$$

again contradicts the Sidon condition modulo N. This establishes that the condition of Sidon translates geometrically into the absence of (possibly degenerate) squares in the graph S_N.

Since the graph S_N is bipartite, it contains no cycle of odd length. This implies that the girth of S_N is ≥ 6. Conversely, suppose that S_N has girth at least 6. Since S contains no square, it follows that

$$(2a-1)-(2b-1)+(2c-1)-(2d-1)\neq 0 \bmod 2N$$

where $a \neq b \neq c \neq d \neq a$. This implies that $a-b+c-d \neq 0 \bmod N$. Therefore, the Sidon condition holds modulo N. ■

Remark 2.2. (a) We refer to [6] for an earlier use of Sidon sequences in graph theory (in a different context).

(b) One can use Theorem 2.1 to introduce a notion of "girth" for a Sidon sequence, namely, an integer $g_N(a_1, \dots, a_n)$ defined to be the girth of the associated graph $S_N(a_1, \dots, a_n)$. Arithmetically, it is easy to characterize the girth as the smallest integer g for which there exists a nontrivial alternating sum of g terms of the Sidon sequence equalling zero modulo $2N$. We may also define

$$g(a_0, \dots, a_n) := g_N(a_0, \dots, a_n),$$

where $N := N_{00}(a_0, \dots, a_n)$ is defined in Section 7.

(c) The group in Theorem 1.1 associated with a Sidon sequence of girth > 6 is hyperbolic in the sense of Gromov.

(d) The definition of S is motivated by the paper of Singer [12], in which he constructs a collineation of order n^2+n+1 in a classical projective plane of order n, and uses this collineation to organize the incidence matrix as a "regular array" as shown in [12, Table 9]. Singer's sequences (as defined in [12, eq. (10)]) have the property that $a_n < n^2+n+1$ and the n^2+n differences $a_i - a_j$ are in fact congruent to the integers $1, 2, 3, \dots, n^2+n$ modulo n^2+n+1 (ibid., p. 381). (See also [8].)

3. Homogeneous labellings in 2-complexes

The second step constructs a simplicial complex of dimension 2, together with a homogeneous labellings (or colourings) of its cells in each dimension, assuming that a similarly homogeneous colouring exists in dimension 1. We shall first introduce some definitions which are useful in our context. Let $p, q \geq 2$ be integers and let L be a graph.

Definition 3.1. We call (p, q)*-labelling* of L an assignment of p labels to the vertex set and q labels to the edge set, in such a way that any two adjacent vertices, and any two incident edges, have distinct labels.

Such a labelling can be viewed as a map $\delta: L \to C$, where C is the set of colours (or labels) such that $\delta(v) \neq \delta(v')$ (resp. $\delta(e) \neq \delta(e')$) if v and v' (resp. e and e') are adjacent.

We shall use graphs endowed with $(2, m)$-labellings (such a graph is necessarily bipartite) which are *homogeneous* in the sense that the label preserving automorphism group

$$\mathrm{Aut}(L, \delta) := \{\theta \in \mathrm{Aut}(L) : \delta \circ \theta = \delta\}$$

is transitive on the sets of vertices and edges with a given label. We call a graph $(2, m)$-homogeneous if it admits such a labelling.

Similarly, let $p, q, r \geq 2$ be integers and let X be a complex of dimension 2.

Definition 3.2. We call (p, q, r)-labelling of X a map $\delta: X \to C$ into a set C such that

(a) $|\delta(X^0)| = p$, $|\delta(X^1)| = q$, $|\delta(X^3)| = r$;

(b) $\delta(v) \neq \delta(v')$ if $v \neq v' \in X^0$ are in the same edge or the same face;

(c) $\delta(e) \neq \delta(e')$ if $e \neq e' \in X^1$ are in the same face;

(d) $\delta(f) \neq \delta(f')$ if $f \neq f' \in X^2$ contain the same edge.

If $\delta: X \to C$ and $\delta': X' \to C$ are (p, q, r)-labellings of 2-complexes, we call a homomorphism $f: X \to X'$ a label preserving homomorphism if $\delta = \delta' \circ f$. We let $\mathrm{Aut}(X, \delta)$ denote the group of label preserving automorphisms of X. We are interested in the following.

Definition 3.3. A 2-complex X is (p, q, r)*-homogeneous* if it admits a (p, q, r)-colouring δ such that $\mathrm{Aut}(X, \delta)$ acts transitively on the sets $X_c^k := \{e \in X^k : \delta(e) = c\}$ of in every dimension $k = 0, 1, 2$, for every $c \in \delta(X^k)$.

The next theorem shows that one can use $(2, m)$-homogeneous labellings in dimension 1 to construct $(3, 3, m)$-homogeneous labellings in dimension 2 for every $m \geq 3$.

It is similar in spirit to the classical development theorems in the theory of complexes of groups [7, Part III.C].

Theorem 3.4. *Let $m \geq 3$. Suppose we are given three graphs L_1, L_2, L_3 of order m and girth 6, endowed with a $(2, m)$-homogeneous labelling on these graphs, with a common set of labels on the edges, and such that the set of labels on the vertices on L_k is $\{1, 2, 3\} \setminus \{k\}$. Then there exists a unique simplicial complex X of dimension 2 admitting a $(3, 3, m)$-labelling with vertex labels the set $\{1, 2, 3\}$, such that L_k is label isomorphic to the link of every vertex of X of label k. Furthermore, this complex is $(3, 3, m)$-homogeneous.*

Proof. We let G_k denote the group of label preserving automorphisms of the graph L_k. By our assumption on L_k, the action of G_k on L_k admits two orbits of vertices and n orbits of edges.

We establish some claims first.

Claim 3.5. *The graph L_k is bipartite.*

Proof. This follows by the fact that the extremities of an edge contain both vertex labels. ∎

Claim 3.6. *The group G_k acts freely on the sets of vertices.*

Proof. The stabilizer of a vertex is trivial since the edges incident to a vertex have to have pairwise distinct labels. ∎

We remark that the fact that L_k is bipartite and that G_k acts transitively on the set of vertices with same label necessarily implies that G_k acts transitively on the set of edges having the same label.

Claim 3.7. *Every element stabilizing an edge of L_k is an involution.*

Proof. If θ stabilizes an edge, then θ^2 fixes the end points of this edge, and therefore $\theta^2 = \mathrm{Id}$ by the previous lemma. ∎

Suppose we are given three graphs as in the assumptions of the theorem and let us construct the 2-complex X.

The first step is the construct a ball B_1 of simplicial radius 1 with links isomorphic to L_1 as a graph. We assign to the faces of B_1 the label of the edges in L_1 they correspond to. We give to the centre of B_1 label 1, and to every vertex y incident to x the label of the vertex in L_1 associated with the edge $[x, y]$ in B_1. We obtain in this way a labelled complex B_1, by associating to every edge in B_1 the unique pair in $\{\{1, 2\}, \{2, 3\}, \{1, 3\}\}$ corresponding to the colour of its extremities.

We proceed by induction. Suppose a labelled simplicial ball B_n is constructed in such a way that it satisfies the assumptions of the theorem in its interior vertices. We write S_n for its boundary and call $\widetilde{B}_n$ the labelled complex obtained by attaching $m-1$ labelled triangles above every edge in S_n.

Lemma 3.8. *The link of a vertex y of S_n in $\widetilde{B}_n$ is a labelled tree T_y of diameter ≤ 5.*

Proof. Since B_n is a simplicial ball, the intersection of T_y with B_n has diameter ≤ 3 (the diameter is exactly 3 at the vertices in $\widetilde{B}_{n-1}$ and it is 2 elsewhere). Adding the faces of $\widetilde{B}_n$, the diameter is now at most 5. Finally, if y has label k, the label of a vertex in T_y is the unique k' such that $\{k, k'\}$ is the label of the corresponding edge in $\widetilde{B}_n$. ∎

Lemma 3.9. *Every labelled metric tree of diameter ≤ 5 having the same vertex and edge labels as L_k embeds in L_k in a label preserving way (for every $k = 1, 2, 3$).*

Proof. Since the girth of L_k is at least 6, it follows by assumption that the balls of diameter 5 in L_k are labelled trees, with m edges adjacent to every vertex. This shows every labelled tree having the same sets of labels embeds in L_k. ∎

Let B_{n+1} denote the simplicial ball obtained from $\widetilde{B}_n$ by fixing a completion $\varphi_y : T_y \to L_k$ for every vertex y of S_n of colour k, adding the labelled cells corresponding to $L_k \setminus \varphi_y(T_y)$.

Claim 3.10. *The direct limit X of the B_n is a $(3, 3, m)$-labelled complex with vertex colour set $\{1, 2, 3\}$, for which L_k is label isomorphic to the link of every vertex of X of colour k.*

Proof. This is clear since the requirements are local and satisfied by B_n viewed as a subcomplex of B_{n+1}. ∎

Let us now turn the uniqueness of X up to isomorphism.

Lemma 3.11. *Suppose φ_1 and φ_2 are two completely coloured isometric embeddings of a metric completely coloured tree of diameter at most 5 (with edges of length 1 in L_k). Then $\varphi_2 \circ \varphi_1^{-1}$ extends in a unique way to a labelled isomorphism of L_k (for every $k = 1, 2, 3$).*

Proof. Since the automorphism group of L_k is transitive on the set of edges having the same label, two balls of diameter 5 centred at edges having the same label in L_k are pairwise label isomorphic. Since the codomains of φ_1 and φ_2 are label isomorphic and included in such balls, the map $\varphi_2 \circ \varphi_1^{-1}$ extends in a unique way. ∎

In order to show that two complexes X and X' as described in the theorem are labelled isomorphic, one then proceeds by induction, by showing that the balls $B_n(x)$ and $B_n(x')$ of radius n, where $x \in X$ and $x' \in X'$ are two vertices of the same colour, are uniquely labelled isomorphic. This results from the previous lemma. In fact, for any pair of vertices (x, y), there exists a unique colour preserving isometric isomorphism of X taking x to y. Similar reasoning starting from an edge or a face implies that X is a $(3, 3m)$-homogeneous complex. ∎

4. Main theorem "with parameters"

The third step states a more general version of Theorem 1.1, that includes "twisting parameters", and establishes (a) ⇒ (b) in this more general context. The twisting parameters are fixed bijections, which are necessary for the group G and the complex X to be well defined a priori. The bijections are denoted σ_j and τ_j where $j \in \{1, 2, 3\}$. The σ_j's allow to twist the face labelling around a vertex in X. The τ_j's can be viewed as signs appearing at the vertices in a fundamental domain. There are three choices for the signs for a total of eight possibles constructions $(\varepsilon_1, \varepsilon_2, \varepsilon_3)$, where $\varepsilon_i \in \{\pm\}$, for every fixed choice of Sidon sequences and bijections σ_i. It is natural in view of the results in the previous section to work with three possibly distinct Sidon sequences.

Theorem 4.1 (Main theorem). *For every integer $n \geq 2$, every triple of increasing sequences*

$$0 \leq a_0^j < a_1^j < \cdots < a_n^j, \quad j = 1, 2, 3,$$

every triple of integers r_j, $j = 1, 2, 3$, such that for every $j \in \{1, 2, 3\}$, the sequence $a_0^j < a_1^j < \cdots < a_n^j$ is a Sidon sequence modulo r_j, and every family of bijections

$$\sigma_j \colon \{a_0^j, \ldots, a_n^j\} \to \{0, \ldots, n\}, \quad j = 1, 2, 3,$$

and

$$\tau_j \colon \{0, 1\} \to \{1, 2, 3\} \setminus \{j\}, \quad j = 1, 2, 3,$$

there exists a countable group G, acting properly discontinuous on a CAT(0) *complex X of dimension* 2 *with compact quotient and three orbits of vertices, a G-invariant labelling $X^2 \to \{0, \ldots, n\}$ of the faces of X, and a G-invariant labelling $X^1 \to \{1, 2, 3\}$ of the edges of X, such that the following are equivalent for a labelled equilateral triangle tessellation Π of the Euclidean plane:*

(a) Π *embeds in X in a label preserving way;*

(b) Π *is solution to a triangle ring puzzle problem with rings of the form given in Figure 4 such that* $\sigma = \sigma_j$, $l = \tau_j(0)$, $k = \tau_j(1)$, $j \in \{1,2,3\}$, *for every triple* (a,b,c) *and* (a',b',c') *of elements of* $\{a_0^j, \ldots, a_n^j\}$ *with equal alternating sum*

$$a - b + c = a' - b' + c' \mod r_j$$

such that $a \neq b \neq c$, $a' \neq b' \neq c'$ *and* $(a,b,c) \neq (a',b',c')$.

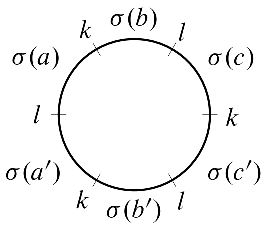

Figure 4
A Sidon ring puzzle problem.

We postpone the proof of (b) ⇒ (a) to Section 6.

Proof. We apply Theorem 3.4, assuming that all triangles are equilateral, and correspondingly that all link edges have length $\pi/3$.

Consider the three graphs $S_{r_j}(a_0^j, a_1^j, \ldots, a_n^j)$ for $j = 1,2,3$. It is easy to verify that the colouring δ_j defined by

(1) $\delta_j(v) = \tau_j(v \bmod 2)$ for every vertex v;

(2) $\delta_j(e) = \sigma_j(a)$ for every edge e labelled by $a \in \{a_0^j, a_1^j \ldots, a_n^j\}$.

is homogeneous. Since the sequences $a_0^j, a_1^j, \ldots, a_n^j$ are Sidon, the graphs S_{r_j} have girth at least 6. This shows that Theorem 3.4 applies. We let X be the corresponding $(3,3,m)$-homogeneous complex, and G be the group of label preserving automorphisms of X. Using the maps σ_j we obtain a G-invariant labelling $X^2 \to \{0,\ldots,n\}$ of the faces of X, and using the maps τ_j we obtain a G-invariant labelling $X^1 \to \{1,2,3\}$ of the edges of X. It follows by Theorem 3.4 that G has three orbits of vertices, and it therefore remains to prove the assertion regarding the structure of the space of embedded flat planes.

Here we show (a) ⇒ (b), i.e., that every tessellation which embeds in X in a label preserving way is the solution of a ring puzzle problem. A ring of X is a cycle of length 2π included in a vertex link of X. Every ring is endowed with the induced vertex and edge labelling from X. Since $S_{r_j}(a_0^j, a_1^j, \ldots, a_n^j)$ is bipartite, the vertex labels of a ring at a vertex of type j alternate k and l where $l = \tau_j(0)$ and $k = \tau_j(1)$.

We must show that the edge labelling has the form stated in the theorem, for some triples (a,b,c) and (a',b',c') of elements of $\{a_0^j,\dots,a_n^j\}$ with equal alternating sum

$$a-b+c=a'-b'+c' \bmod r_j$$

such that $a\neq b\neq c$, $a'\neq b'\neq c'$ and $(a,b,c)\neq(a',b',c')$. Let R be a ring in $S_{r_j}(a_0^j,a_1^j,\dots,a_n^j)$. We fix an arbitrary base vertex $m\in\mathbb{N}$ of R with label l and consider the two triples (a,b,c) and (a',b',c') of elements of $\{a_0^j,\dots,a_n^j\}$ describing the edge labelling of R. Then by the definition of $S_{r_j}(a_0^j,a_1^j,\dots,a_n^j)$ the following integers:

$$m+(2a-1)-(2b-1)+(2c-1),\quad m+(2a'-1)-(2b'-1)+(2c'-1),$$

must coincide modulo $2r_j$. This shows that

$$2a-2b+2c=2a'-2b'+2c' \bmod 2r_j,$$

which implies the desired result. It is clear conversely that if this condition is satisfied, then the two integers:

$$m+(2a-1)-(2b-1)+(2c-1),\quad m+(2a'-1)-(2b'-1)+(2c'-1),$$

coincide modulo $2r_j$ for every $m\in\mathbb{N}$ with label l. The corresponding paths in $S_{r_j}(a_0^j,a_1^j,\dots,a_n^j)$ are distinct assuming that $(a,b,c)\neq(a',b',c')$, and therefore they define a cycle, which is a ring with the correct labelling. ■

As mentioned above, we may encode the bijections τ_j by a sign $\varepsilon_j\in\{\pm\}$ at every vertex, where by convention $\varepsilon_j=+$ if and only if the bijection τ_j is increasing. Given the Sidon sequences and the bijections σ_j, the choice of τ_j defines eight complexes $X_{(\varepsilon_1,\varepsilon_2,\varepsilon_3)}$. In the next section we prove these complexes are in fact pairwise isomorphic.

5. Existence of polarities

The fourth step concerns the automorphism group of $X:=X_{(\varepsilon_1,\varepsilon_2,\varepsilon_3)}$. We prove the existence of sufficiently many "polarities". This implies that

$$X_{(\varepsilon_1,\varepsilon_2,\varepsilon_3)}\simeq X_{(\varepsilon_1',\varepsilon_2',\varepsilon_3')}$$

for every $\varepsilon_1,\varepsilon_2,\varepsilon_3,\varepsilon_1',\varepsilon_2',\varepsilon_3'\in\{\pm\}$.

Definition 5.1. Let S be a bipartite graph. A *polarity* of S is an automorphism of S of order 2 which permutes the vertex types of S.

If e is an edge of S, we call *polarity at e* a polarity which permutes the extremities of e.

In the next two lemmas, we let $a_0 < \cdots < a_n$ be a Sidon sequence modulo N and consider the quotient graph S_N of $S = S(a_0, \dots, a_n)$ by the action of $2N\mathbb{Z}$ by translations.

We assume S_N is endowed with a fixed edge labelling

$$\sigma\colon \{a_0, \dots, a_n\} \to \{0, \dots, n\}$$

(as in Theorem 4.1).

Lemma 5.2. *S_N admits an edge label preserving polarity at every edge.*

Proof. For every $0 \leq p \leq n$ consider the map $\varphi_p\colon S \to S$ taking k to $-k + 2a_p - 1$. For every $0 \leq q \leq n$ and l even, we have

$$\varphi_p\big([l, l + 2a_q - 1]\big) = \big[-l + 2a_p - 1, -l + 2a_p - 1 - (2a_q - 1)\big].$$

This shows φ_p takes edges to edges and defines an automorphism of S. This automorphism factorizes to an automorphism of S_N, which is a label preserving polarity that exchanges the end vertices of edge $[0, 2a_p - 1]$. By transitivity of the group of edge label preserving automorphisms, such a polarity exists at every edge of S_N. ∎

Lemma 5.3. *Let T, T' be isomorphic subtrees of S_N containing at least a tripod, and let $\varphi_0\colon T \to T'$ be a isomorphism preserving the edge labels. Then φ_0 extends in a unique way to an automorphism of S_N preserving the edge labels.*

Proof. If T and T' are adjacent tripod it suffices to choose the polarity at the common edge. If T and T' are arbitrary tripods, one can use transitivity of the group of automorphisms on vertices of the same type. If T and T' are arbitrary trees, one may fix a tripod $T_0 \subset T$ can consider the unique edge label preserving automorphism of S_N whose restriction to T_0 is φ_0. It takes T to T'. Uniqueness is clear since we have required the automorphisms to preserve be the edge labelling. ∎

Theorem 5.4. *There exists an isomorphism $X_{(\varepsilon_1,\varepsilon_2,\varepsilon_3)} \simeq X_{(\varepsilon'_1,\varepsilon'_2,\varepsilon'_3)}$ which preserve the labels on faces, for every $\varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon'_1, \varepsilon'_2, \varepsilon'_3 \in \{\pm\}$.*

Proof. We write $X = X_{(\varepsilon_1,\varepsilon_2,\varepsilon_3)}$ and $X' = X_{(\varepsilon'_1,\varepsilon'_2,\varepsilon'_3)}$. By symmetry, it is enough to consider the case $\varepsilon_1 \neq \varepsilon'_1$, $\varepsilon_2 = \varepsilon'_2$, $\varepsilon_3 = \varepsilon'_3$. We let x and x' respectively denote a vertex of type 1 in X and X'. By Lemma 5.2, there exists an edge label preserving isomorphism of the corresponding links. This isomorphism induces an isomorphism $\varphi_1\colon B_1(x) \to B_1(x')$ between the ball of radius 1, which preserves the labels on the faces. Then, it follows by Lemma 5.3 that φ_1 extends in a unique way to label preserving isomorphisms φ_n between successive balls $B_n(x)$ and $B_n(x')$. The direct limit of these maps is an isomorphism which preserve the labels of the faces. ■

Thus, X is well defined up to isomorphism regardless of the choice of τ_j. One may for example fix these bijections to be

$$\tau_1(0) = 2, \quad \tau_1(1) = 3, \quad \tau_2(0) = 1, \quad \tau_2(1) = 3, \quad \tau_3(0) = 1, \quad \tau_3(1) = 2.$$

Remark 5.5. Assuming that the Sidon sequences (a_j^j) and the integers r_j are fixed, the choice of the σ_j's generates a priori $(n+1)!^3$ distinct groups and spaces associated with these data. The associated 2-dimensional paving problems depends on the choice of these elements a priori, and so does the space of embedded flats in X. It is straightforward to provide examples showing that different choices of bijections may lead to nonisomorphic complexes and groups, for example by using constructions analogous to that of the odd Moebius–Kantor complex described in Section 7.2.

6. Every ring puzzle embeds in X

The fifth and last step is to prove that every solution to the given ring puzzle problem defines a flat in X, showing (b) $\Rightarrow$ (a) in Theorem 1.1, which is the last assertion that remains to be established. In fact, we prove this equivalence in the more general setting of Theorem 4.1.

Theorem 6.1. *Suppose* Π *is a labelled flat obtained as a solution of the puzzle problem described in Theorem* 4.1. *Then* Π *embeds in* X *in a label preserving way.*

In fact, every labelled preserving embedding of a 1-ball of Π into X extends uniquely to an embedding of Π into X.

Proof. Let Π be a labelled flat obtained as a solution of the puzzle problem described in Theorem 4.1, and let B_1 be a 1-ball of Π. Consider a labelled preserving embedding φ_1 of B_1 into X. We show that φ_1 extends uniquely to an embedding of Π into X.

The proof is by induction using the following lemma.

Lemma 6.2. *Let* R *be a ring of type* $j = 1, 2, 3$ *(as described in Theorem* 4.1*) and let* $P \subset R$ *be a segment in* R. *Then every label preserving embedding* $\psi\colon P \to L_j$ *extends to a label preserving embedding* $\psi'\colon R \to L_j$.

Proof of Lemma 6.2. By assumption, there exists an edge label preserving embedding $\psi_0\colon R \to L_j$. We let e denote the initial edge in P. Since L_j is completely transitive, there exists an edge label preserving automorphism $\theta\colon L_j \to L_j$ taking $\psi_0(e)$ to $\psi(e)$. Up to composing with a polarity of L_j fixing $\psi(e)$, we may assume that θ takes $\psi_0(f)$

to $\psi(f)$, where f is the edge adjacent to e in P (assuming that P has length at least 2). It follows by the fact that θ is label preserving that θ takes $\psi_0(P)$ to $\psi(P)$. Then $\psi' := \theta \circ \psi_0$ is an edge label preserving embedding extending ψ. Since ψ is vertex label preserving, so is ψ'. ∎

Suppose φ_n is an embedding the n-ball of Π concentric to B_1 into X and let us show φ_n can be extended to an embedding of the $(n+1)$-ball B_{n+1}. For every vertex x of $S_n = \partial B_n$, we let R_x denote the ring at x in Π and $P_x \subset R_x$ denote the path in R_x associated to B_n. Then φ_n induces a labelled preserving embedding $\psi_x \colon P_x \to L_{\varphi_n(x)}$. By the previous lemma, ψ_x extends uniquely to a label preserving embedding $\psi'_x \colon R_x \to L_{\varphi_n(x)}$. We let B_{n+1} be the union of B_n and $\bigcup_{x\in S_n}[\psi'_x(R_x)]$, where $[\psi'_x(R_x)]$ is the 1-disk in X corresponding to $\psi'_x(R_x)$. Since the maps ψ'_x are label preserving, this is a disk of radius $n+1$ in X which extends B_n. Thus, φ_n extends in a unique way. ∎

7. Examples and applications

7.1. Mian–Chowla complexes. Let $(a_n)_{n\geq 0}$ denote the Mian–Chowla sequence (obtained from $a_0 = 0$ by the greedy algorithm).

Theorem 4.1 enables us to associated groups and complexes to the truncated Mian–Chowla sequence $a_0, \ldots, a_n$ and any integer N such that $a_0, \ldots, a_n$ is a Sidon sequence modulo N. In general, this can be done for every N large enough.

Proposition 7.1. *Let $a_0, \ldots, a_n$ be a Sidon sequence. Then $N_0 = 2a_n + 1$ is the smallest integer with the property that $a_0, \ldots, a_n$ Sidon sequence modulo N for every $N \geq N_0$.*

Proof. It is clear that $N_0 \geq 2a_n + 1$, since $a_n + a_n = a_0 + a_0 \mod 2a_n$. Conversely, if $a_1, \ldots, a_n$ is not Sidon modulo N, then there exist $a_i, a_j, a_p, a_q \in \{a_1, \ldots, a_n\}$ such that

$$a_i + a_j \neq a_p + a_q$$

and

$$N \text{ divides } a_i + a_j - a_p - a_q.$$

Since $0 \leq a_i + a_j, a_p + a_q \leq 2a_n$, this fails if $N \geq 2a_n + 1$. ∎

While the proposition shows the general bound $2a_n + 1$ is sharp, we note that there exist Sidon sequences which are Sidon sequence modulo N for some values of N which are smaller than N_0. Consider for example the Sidon sequence $0, 2, 7$. Then $N_0 = 15$, and it is easily verified that $0, 2, 7$ is a Sidon sequence modulo N, for several values of $N < N_0$ (e.g., $N = 8$).

Question 7.2. Given a Sidon sequence $a_0, \dots, a_n$, what is the value of

$$N_{00}(a_0, \dots, a_n) := \min\{N : a_0, \dots, a_n \text{ is a Sidon sequence modulo } N\}?$$

In some cases $N_{00} = N_0$. Consider for example the Sidon sequence $0, 1, 3$. Then $N_0 = 7$, and it is easily verified that the Mian–Chowla sequence $0, 1, 3$ fails to be a Sidon sequence modulo N, for every $N \leq 6$. On the other hand, it is not difficult to check, for example, that the Mian–Chowla sequence $0, 1, 3, 7, 12, 20$ is a Sidon sequence modulo 35.

For every $n \geq 2$, we call Mian–Chowla complex the CAT(0) 2-complex X_n associated by Theorem 4.1 with the following data:

(1) $a_0^j, \dots, a_n^j$ is the truncated Mian–Chowla sequence;

(2) $\sigma_j : \{a_0^j, \dots, a_n^j\} \to \{0, \dots, n\}$ is the increasing bijection;

(3) $N_j = N_{00}(a_0^j, \dots, a_n^j)$;

for every $j = 1, 2, 3$.

Due to the symmetry in these data, we have the following proposition.

Proposition 7.3. *The automorphism groups of the Mian–Chowla complexes are vertex transitive.*

It would be interesting to study the geometric structure of the Mian–Chowla complexes X_n, and solve the associated ring puzzle problems. We observe that the Mian–Chowla complex X_2 is in fact a Bruhat–Tits building.

7.2. Moebius–Kantor complexes. We say that a CAT(0) 2-complex X is a *Moebius–Kantor complex* if its faces are equilateral triangles and its vertex links are isomorphic to the Moebius–Kantor graph (namely, the unique bipartite cubic symmetric graph on 16 vertices).

Consider the Sidon sequence $a_0 = 0, a_1 = 1, a_2 = 3$ (of length 3) modulo $N = 8$, and the bijections

$$\sigma_i : \{0, 1, 3\} \to \{0, 1, 2\}, \quad i = 1, 2, 3$$

given by $\sigma_1(0) := 0, \sigma_1(1) := 1, \sigma_1(3) := 2$, and $\sigma_i := (0\,1\,2)^i \sigma_1$ for $i = 2, 3$. Applying Theorem 4.1, we find a group G and a CAT(0) complex X with the indicated properties.

We claim the following result.

Proposition 7.4. *The* CAT(0) *complex X is a Moebius–Kantor complex.*

Proof. By Theorem 4.1, the space X is a CAT(0) space with equilateral triangle faces. The links in X are determined by the Sidon sequence. By definition, they are isomorphic to the graph with vertex set $\mathbb{Z}/16\mathbb{Z}$ and edge set $[n, n+1]$ for every n and $[n, n+5]$ for every n even. It is easy to verify that this graph is the Moebius–Kantor graph. ∎

Thus, Theorem 4.1 provides a new construction method for Moebius–Kantor complexes. Here we shall describe here some properties of X, and in particular motivate our choice of bijections.

We call *root* of X an isometric embedding α of a path of length π in a link of X, such that $\alpha(0)$ is a vertex. Thus, the image of every root consists of three edges, which we shall endow with the induced labelling in $\{0, 1, 2\}$. We say that α is a root of rank 2 if there exist precisely two roots distinct from α with the same end points. This definition is a particular case of the notion of rank of a root in a CAT(0) 2-complex—and we shall refer the interested reader to [1, §4] for this generalization.

We write S_i for the link in X at a vertex of type i. For the complex X, the rank of a root in S_i is a function of its labels; the following statement is the most relevant for our purpose.

Lemma 7.5. *Let α be a root in S_i, $i = 1, 2, 3$, with consecutive labels b, a, and b, where $a \neq b$. Then α is a root of rank 2 if and only if $a = \sigma_i(0)$.*

Proof. Suppose an edge $e = [r, r+1]$ has label $a = \sigma_i(0)$ and let $b \neq a$. Then, by our definition of S_i, r is odd. Furthermore, since they have the same label b, then the two edges adjacent to e in α have the same increment, which is either 1 or 5. It is easily seen that α is of rank 2 in both cases.

Suppose now an edge $e = [r, r+1]$ has label $a = \sigma_i(1)$ or $a = \sigma(3)$ and let $b \neq a$. By the same argument, the two edges adjacent to e in α have the same increment. It is easily seen that α is not of rank 2 in both cases. ∎

Let t be a triangle in X. For every side of t, choose a triangle in X adjacent to t. This defines three roots in the links of the vertices of t. We say that t is *odd* if the number of such roots of rank 2 is odd (this is well defined by [3, §2]). We say that X is *odd* if every triangle is odd.

Proposition 7.6. *The complex X is odd.*

Proof. Consider a triangle t in X with label $a \in \{0, 1, 2\}$. Let $b \in \{0, 1, 2\}$, $b \neq a$. Adjacent to t are three triangles with label b. These triangles form, together with t, a larger triangle which we call T. Suppose $a = \sigma_1(k)$ for some $k \in \{0, 1, 3\}$. Then $\sigma_2(k) = (0\ 1\ 2)a, \sigma_3(k) = (0\ 1\ 2)^2 a$; therefore, $\sigma_2((0\ 1\ 3)k) = a, \sigma_3((0\ 1\ 3)^2 k) = a$. This implies that $a = \sigma_i(0)$ for a unique $i = 1, 2, 3$ which in turns shows T contains a single roots of rank 2. This proves that t is odd. ∎

7.3. Modular complexes. Let $a_0 = 0 < a_1 < \cdots < a_n$ be a sequence and N_1, N_2, and N_3 be integers such that $a_0 = 0 < a_1 < \cdots < a_n$ satisfies the condition of Sidon modulo N_i for $i = 1, 2, 3$. Let $G(a_0, a_1, \ldots, a_n : N_1, N_2, N_3)$ and, respectively,

$X(a_0, a_1, \ldots, a_n : N_1, N_2, N_3)$ denote be the group and complex obtained by applying Theorem 4.1 with respect to these data and the increasing bijection

$$\sigma\colon \{a_0, \ldots, a_n\} \to \{0, \ldots, n\}.$$

The notation $X_N(a_0, a_1, \ldots, a_n)$ used in the introduction refers to the case $N = N_1 = N_2 = N_3$, but Theorem 4.1 allows for more freedom in the constructions.

To illustrate this, we explain how Theorem 4.1 provides an alternative approach to [4, §13] for a construction "mixing" the Moebius–Kantor local geometry to that of $\widetilde{A}_2$ buildings in a same complex. Such a complex was said to be "of strict type $A_{\mathrm{MK}} + \widetilde{A}_2$". It was obtained in [4] by a surgery construction, relying on the classification of collars between two "partial complexes" of both types. Here we can obtain similar results in as a direct consequence of Theorem 4.1. For instance, using the terminology of [4], we claim the following.

Proposition 7.7 (Compare [4, Prop. 13.1]). *The modular complex $X(0, 1, 3 : 7, 7, 8)$ is of strict type $A_{\mathrm{MK}} + \widetilde{A}_2$.*

The proof follows that of Proposition 7.4. Similarly, the complex $X(0, 1, 3 : 7, 8, 8)$ is a complex of strict type $A_{\mathrm{MK}} + \widetilde{A}_2$, and it is not isomorphic to $X(0, 1, 3 : 7, 7, 8)$. These "modular complexes" seem particularly interesting when N_i is small relative to a_n. Furthermore, one can use bijections other than σ to twist the construction of modular complexes (for example, the complex X described in Section 7.2 can be viewed as a "twisted modular complex"). We define "the" modular complex to be as untwisted as possible:

Definition 7.8. Let $a_0 = 0 < a_1 < \cdots < a_n$ be a Sidon sequence. We let

$$X(a_0, \ldots, a_n) := X(a_0, a_1, \ldots, a_n : N_{00}, N_{00}, N_{00}),$$

where $N_{00} := N_{00}(a_0, \ldots, a_n)$, and call this complex *the modular complex* associated with $a_0 = 0 < a_1 < \cdots < a_n$.

By definition, the Mian–Chowla complexes are modular complexes in this sense. Due to symmetry in the data, follows by Theorem 4.1 that the automorphism group of the modular complex $X(a_0, \ldots, a_n)$ is transitive on the vertex set. (This would not be true of generalized modular complex, for example, $X(0, 1, 3 : 7, 7, 8)$ is not vertex transitive.)

Remark 7.9. These examples can be further generalized. If A is a finite abelian group, one says a set $\{a_0, a_1, \ldots, a_n\}$ in A is Sidon if the number of pairs of elements in $\{a_0, a_1, \ldots, a_n\}$ with a given sum is at most two. It is not difficult to extend our results to such Sidon sets, giving further natural generalizations of our Sidon complexes.

It would be of considerable interest to study the asymptotic properties of the complexes $X(a_0, \dots, a_n)$, for a fixed infinite Sidon sequence $a_0, a_1, a_2, \dots$, including for example, the Mian–Chowla sequence or the Rusza sequence. We shall not pursue this direction on the present occasion, and conclude this paper with an application to the study of Moebius–Kantor complexes.

7.4. Uniqueness of the odd Moebius–Kantor complex. In this section we prove that the twisted modular complex constructed in Section 7.2 is the unique odd Moebius–Kantor complex up to isomorphism; furthermore, we establish a unique extension theorem for automorphisms. This is done by "mapping" the data associated with the Sidon sequence to an arbitrary odd complex, and may compared to [5], in which a similar result is proved in the even case: the even Moebius–Kantor complex is unique up to isomorphism. This was established in [5] by "mapping" Pauli matrices from the (even) Pauli complex to an arbitrary even complex.

Theorem 7.10. *Let X and X' be odd Moebius–Kantor complexes and let $x \in X$ and $x' \in X'$ be vertices. Let $B_1(x)$ and $B_1(x')$ denote the ball of radius 1 with centre x and x', respectively, and let $\varphi_1 : B_1(x) \to B_1(x')$ be an isomorphism. Then there exists a unique isomorphism $\varphi : X \to X'$ which coincides with φ_1 on $B_1(x)$.*

Proof. We may assume that X is the complex constructed in Section 7.2; furthermore, by symmetry, we may assume that x is a vertex of type 1 in this complex (associated with the map σ_1). Let us first extend φ_1 to the ball $B_2(x)$ of radius 2. We begin with the following lemma.

Lemma 7.11. *Suppose S and S' are Moebius–Kantor graphs, $T \subset S$ and $T' \subset S'$ are tripods, and $\psi_0 : T \to T'$ is an isomorphism. Then there exists a unique isomorphism $\psi : S \to S'$ which coincides with ψ_0 on T.*

Proof. Existence follows because the Moebius-Kantor graph is 2-arc-transitive. If e and f are consecutive edges in T, then there exists a graph isomorphism $S \to S'$ taking respectively e and f to $\psi_0(e)$ and $\psi_0(f)$. This isomorphism is unique since the stabilizer of a tripod is trivial. ∎

For every vertex y in the sphere $\partial B_1(x)$ of radius 1 centred at x, we let ψ_y denote the unique isomorphism between the ball $B_1(y)$ and the ball $B_1(\varphi_0(y))$ induced by the previous lemma, which extends φ_0 on $B_1(x)$.

Lemma 7.12. *The maps ψ_y are consistent.*

Proof. We must show that for every edge $[y, z]$ in $\partial B_1(x)$, the maps ψ_y and ψ_z coincide on set of triangles adjacent to $[y, z]$. Since x is of type 1, we may assume that y is of type 2 and z of type 3. Let $[t_1, y, z]$ be a triangle distinct from $[x, y, z]$, and consider the two triangles $[t_2, x, y]$ and $[t_3, x, z]$ whose labels coincide with that of $[t_1, y, z]$. We write α_x, α_y and α_z for the three roots, respectively at x, y and z, associated with this configuration. There are two cases.

Suppose first that $[x, y, z]$ is labelled by $a = \sigma_1(0)$. Since $a \neq \sigma_2(0)$ and $a \neq \sigma_3(0)$, both roots α_y and α_z fail to be of rank 2. Since φ_0, ψ_y and ψ_z are isomorphism, $\varphi_0(\alpha_x)$ is of rank 2 in X', while $\psi_y(\alpha_y)$ and $\psi_z(\alpha_z)$ are not. Since the triangle $\varphi_0([x, y, z])$ is odd, the two triangles $\psi_y([t_1, y, z])$ and $\psi_z([t_1, y, z])$ must coincide.

Suppose next that $[x, y, z]$ is labelled by $a = \sigma_2(0)$ or $a = \sigma_3(0)$. The two cases are symmetric and we assume $a = \sigma_2(0)$ to fix the ideas. Then α_y is of rank 2, while α_x and α_z are not, and the same must be true of their images. Again, since the triangle $\varphi_0([x, y, z])$ is odd, the two triangles $\psi_y([t_1, y, z])$ and $\psi_z([t_1, y, z])$ coincide. ∎

By Lemma 7.12, the map

$$\varphi_2 := \varphi_1 \vee \bigvee_{y \in \delta B_1(x)} \psi_y$$

is well defined. It induces an isomorphism between $B_2(x)$ and $B_2(x')$ which extends φ_1 by definition. Furthermore, this extension is unique by Lemma 7.11.

Let $n \geq 2$. Let $\varphi_n \colon B_n(x) \to B_n(x')$ be an isomorphism, and fix a vertex y in the sphere $\partial B_n(x)$ of radius n centred at x. If there does not exist a triangle $[y, y_1, y_2]$ in B_n such that $[y, y_1, y_2] \cap \partial B_n = \{y\}$, we let ψ_y denote the unique isomorphism between the ball $B_1(y)$ and the ball $B_1(\varphi_n(y))$ induced by the Lemma 7.11, which extends φ_n on $B_n(x)$.

Suppose there exists a triangle $[y, y_1, y_2]$ in B_n such that $[y, y_1, y_2] \cap \partial B_n = \{y\}$. Lemma 7.11 provides two maps ψ_y^1 and ψ_y^2 between the ball $B_1(y)$ and the ball $B_1(\varphi_n(y))$ induced by the Lemma 7.11, which extends the restriction φ_n to the set of triangles adjacent to $[y, y_1]$ and $[y, y_2]$, respectively. We show the following result.

Lemma 7.13. *It holds that $\psi_y^1 = \psi_y^2$; furthermore, they extend the restriction of φ_n to $B_n \cap B_1(y)$.*

Proof. Let a denote the label of $[y, y_1, y_2]$. Consider triangles $[t_0, y_1, y_2]$, $[t_1, y, y_1]$ and $[t_2, y, y_2]$ with the same label $b \neq a$, and the corresponding roots α_y, α_{y_1} and α_{y_2}, respectively. We assume that y is of type 1. The other cases are similar by symmetry.

Suppose $a = \sigma_1(0)$. Then α_y is a root of rank 2, while α_{y_1} and α_{y_2} are not, and ψ_y^1 takes $[t_2, y, y_2]$ to the unique triangle in X' such that $\psi_y^1(\alpha_y)$ is a root of rank 2 in the link of $\varphi_n(y)$. Since the triangle $\varphi_n([y, y_1, y_2])$ is odd in X', this shows that φ_n and ψ_y^1 coincide on $B_n \cap B_1(y)$. By symmetry, φ_n and ψ_y^2 coincide on B_n. Since ψ_y^1 and ψ_y^2 coincide on (at least) a tripod, they must coincide everywhere by Lemma 7.11. The two other cases, namely, $a = \sigma_2(0)$ and $a = \sigma_3(0)$, are similar. ∎

In the case that there exists a triangle $[y, y_1, y_2]$ in B_n such that $[y, y_1, y_2] \cap \partial B_n = \{y\}$, we let $\psi_y := \psi_y^1 = \psi_y^2$; this now defines ψ_y for all $y \in \delta B_n(x)$. A direct generalization of Lemma 7.12 to larger balls show that the maps ψ_y are consistent.

It follows that

$$\varphi_{n+1} := \varphi_n \vee \bigvee_{y \in \delta B_n(x)} \psi_y$$

is well defined, and induces an isomorphism between $B_{n+1}(x)$ and $B_{n+1}(x')$ which extends φ_n by definition. This extension is unique by Lemma 7.11. Thus, $\varphi := \varinjlim \varphi_n$ is an isomorphism from X to X' which extends φ_1 uniquely. ∎

Acknowledgements. We are grateful to the referees for interesting comments on this paper.

Funding. The authors are partially funded by NSERC Discovery Fund 234313.

References

[1] S. Barré and M. Pichot, Removing chambers in Bruhat–Tits buildings. *Israel J. Math.* **202** (2014), no. 1, 117–160. Zbl 1323.51002 MR 3265316

[2] — Aut(F_2) puzzles. *Geom. Dedicata* **199** (2019), 225–246. Zbl 1414.52013 MR 3928799

[3] — On the isomorphism types of Moebius–Kantor complexes. *J. Geom.* **110** (2019), no. 3, article no. 51. Zbl 1444.57011 MR 4021136

[4] — Surgery on discrete groups. [v1] 2016, [v3] 2021, arXiv:1611.08544

[5] — Root distributions in Moebius–Kantor complexes. *Aus. J. Comb.* **89** (2024), no. 2, 196–214.

[6] B. Bollobás and O. Pikhurko, Integer sets with prescribed pairwise differences being distinct. *European J. Combin.* **26** (2005), no. 5, 607–616. Zbl 1062.05126 MR 2127684

[7] M. R. Bridson and A. Haefliger, *Metric spaces of non-positive curvature*. Grundlehren Math. Wiss. 319, Springer, Berlin, 1999. Zbl 0988.53001 MR 1744486

[8] P. Erdös, On a problem of Sidon in additive number theory and on some related problems. Addendum. *J. London Math. Soc.* **19** (1944), no. 76(4), 208. Zbl 0061.07302 MR 0014111

[9] H. Halberstam and K. F. Roth, *Sequences*. Second edn., Reprint, Springer, New York-Berlin, 2011. Zbl 0498.10001 MR 0687978

[10] K. O'Bryant, A complete annotated bibliography of work related to Sidon sequences. *Electron. J. Combin.* **DS11** (2004), Dynamic Surveys. Zbl 1142.11312 MR 4336213

[11] I. Z. Ruzsa, An infinite Sidon sequence. *J. Number Theory* **68** (1998), no. 1, 63–71. Zbl 0927.11005 MR 1492889

[12] J. Singer, A theorem in finite projective geometry and some applications to number theory. *Trans. Amer. Math. Soc.* **43** (1938), no. 3, 377–385. Zbl 0019.00502 MR 1501951

[13] A. Stöhr, Gelöste und ungelöste Fragen über Basen der natürlichen Zahlenreihe. II. *J. Reine Angew. Math.* **194** (1955), 111–140. Zbl 0066.03101 MR 0075228

(Reçu le 9 juillet 2024)

Sylvain Barré, UMR 6205, LMBA, Université de Bretagne-Sud, BP 573, 56017 Vannes, France; *e-mail:* sylvain.barre@univ-ubs.fr

Mikaël Pichot, Department of Mathematics and Statistics, McGill University, 805 Sherbrooke Street West, Montréal, QC H3A 0B9, Canada; *e-mail:* pichot@math.mcgill.ca