

A NOTE ON DENSITY OF GEODESICS

J. BECK, W.W.L. CHEN, AND Y. YANG

ABSTRACT. We extend the famous result of Katok and Zemlyakov on the density of half-infinite geodesics on finite flat rational surfaces to half-infinite geodesics on a finite polycube translation 3-manifold. We also extend this original result to establish a weak uniformity statement.

1. A DENSITY RESULT

A finite *polysquare region* P is an arbitrary connected, though not necessarily simply-connected, polygon on the plane which is tiled with closed unit squares, called the *atomic squares* or *faces* of P , and which satisfies the following conditions:

- (i) Any two atomic squares in P either are disjoint, or intersect at a single point, or have a common edge.
- (ii) Any two atomic squares in P are joined by a chain of atomic squares where any two neighbours in the chain have a common edge.

Note that P may have *holes*, and we also allow *whole barriers* which are vertical or horizontal *walls* that consist of one or more boundary edges of atomic squares, as shown in the picture on the left in Figure 1.




Figure 1: bla bla bla

Given such a finite polysquare region P , we can convert it into a finite *polysquare translation surface* $\mathcal{P}$ by identifying pairs of the boundary edges with inward normals in opposite directions, as illustrated in the passage from the picture on the left to the picture on the right in Figure 1. Geodesic flow on such a finite polysquare translation surface $\mathcal{P}$ is then 1-direction geodesic flow.

We say that a geodesic $\mathcal{L}(t)$, $t \geq 0$, on $\mathcal{P}$ is *half-infinite* if it does not hit a singularity of $\mathcal{P}$ and becomes undefined. Furthermore, such a geodesic is *dense* if it gets arbitrarily close to every point of $\mathcal{P}$.

The following result of Katok and Zemlyakov [1] in 1975 includes density of half-infinite geodesics on finite polysquare translation surfaces as a special case.

Theorem A. *With the exception of a countable set of given directions, every half-infinite geodesic on a finite rational surface $\mathcal{P}$ is dense. In the special case when $\mathcal{P}$ is a finite polysquare translation surface, any half-infinite geodesic with irrational slope is dense in $\mathcal{P}$.*

We remark that a finite rational surface is a flat surface with finitely many faces and where every face is a rational polygon, where each angle is a rational multiple of π , and parallel edges with the same length are identified in pairs, while identified edges from the same face have opposite orientation.

In this paper, we consider the 3-dimensional analogue of the special case of this result.

A finite *polycube region* M is an arbitrary connected, though not necessarily simply-connected, polyhedron in 3-space which is tiled with closed unit cubes, called the *atomic cubes* of M , and which satisfies the following conditions:

- (i) Any two atomic cubes in M either are disjoint, or intersect at a single point, or have a common edge, or have a common face.
- (ii) Any two atomic cubes in M are joined by a chain of atomic cubes where any two neighbours in the chain have a common face.

Given such a finite polycube region M , we can convert it into a finite *polycube translation 3-manifold* $\mathcal{M}$ by identifying pairs of the boundary faces with inward normals in opposite directions. Geodesic flow in such a finite polycube translation 3-manifold $\mathcal{M}$ is then 1-direction geodesic flow.

A half-infinite geodesic in a finite polycube translation 3-manifold $\mathcal{M}$ is dense if it gets arbitrarily close to every point of $\mathcal{M}$.

To state our first result, we need a simple definition.

A vector $\mathbf{v}^* = (\alpha_1, \alpha_2, 1) \in \mathbb{R}^3$ is a Kronecker direction if $v_1, v_2, 1$ are linearly independent over $\mathbb{Q}$.

Theorem 1. *Let $\mathcal{M}$ be a polycube translation 3-manifold with s atomic cubes. Then any half-infinite geodesic with a Kronecker direction $\mathbf{v}^* = (\alpha_1, \alpha_2, 1) \in \mathbb{R}^3$ is dense in $\mathcal{M}$.*

Proof. Suppose, on the contrary, that $\mathcal{L}(t)$, $t \geq 0$, is not dense. Then its closure $\overline{\mathcal{L}}$ is a proper subset of $\mathcal{M}$, and the open set $\mathcal{M} \setminus \overline{\mathcal{L}}$ contains an open ball B . We may assume, without loss of generality, that B is contained in a single atomic cube of $\mathcal{M}$. Then the subset

$$\mathcal{I}(B) = \{\mathbf{v} \in \mathcal{M} : \mathbf{v} = \mathbf{v}_0 + t\mathbf{v}^* \text{ for some } \mathbf{v}_0 \in B \text{ and } t \geq 0\} \subset \mathcal{M}$$

is a $\mathbf{v}^*$ -flow-invariant open set.

Consider the multiplicity function $m : [0, 1)^3 \rightarrow \mathbb{N}$, given by

$$m(\mathbf{x}) = |\{\mathbf{p} \in \mathcal{I}(B) : \mathbf{p} \equiv \mathbf{x} \pmod{1}\}|, \quad \mathbf{x} \in [0, 1)^3.$$

In other words, $m(\mathbf{x})$ counts the number of points in $\mathcal{I}(B)$ that have common image $\mathbf{x} \in [0, 1)^3$ under modulo 1 projection from $\mathcal{M}$ to the unit torus $[0, 1)^3$. Since the $\mathbf{v}^*$ -flow in the unit torus $[0, 1)^3$ is ergodic, the function m is constant almost everywhere, and the constant is an integer m_0 satisfying $0 \leq m_0 \leq s$.

As $\mathcal{I}(B)$ is an open set, clearly $m_0 \geq 1$. On the other hand, the complement of $\mathcal{I}(B)$ in $\mathcal{M}$ is a closed set that contains $\overline{\mathcal{L}}$. Clearly the image under modulo 1 projection of $\overline{\mathcal{L}}$ is $[0, 1)^3$, as the image of $\mathcal{L}(t)$, $t \geq 0$, under modulo 1 projection from $\mathcal{M}$ to the unit torus $[0, 1)^3$ is a geodesic in the unit torus $[0, 1)^3$ with a Kronecker direction, and is therefore dense in $[0, 1)^3$. Hence $1 \leq m_0 \leq s - 1$.

Lemma 1.1. *There exist m_0 open balls $B'_1, \dots, B'_{m_0}$ of identical positive radius such that $B'_1 \subset B$ and*

$$B'_i \subset \mathcal{I}(B), \quad i = 1, \dots, m_0,$$

are in distinct atomic cubes of $\mathcal{M}$ and their images under modulo 1 projection from $\mathcal{M}$ to the unit torus $[0, 1)^3$ are identical.

Proof. Let B_0 be the image of B under modulo 1 projection from $\mathcal{M}$ to the unit torus $[0, 1]^3$. Clearly there exists $\mathbf{x}_0 \in B_0$ such that $m(\mathbf{x}_0) = m_0$. This means that there exist $\mathbf{p}_1 \in B \subset \mathcal{I}(B)$ and $\mathbf{p}_2, \dots, \mathbf{p}_{m_0} \in \mathcal{I}(B)$ in $m_0 - 1$ other distinct atomic cubes of $\mathcal{M}$ such that their common image under modulo 1 projection from $\mathcal{M}$ to the unit torus $[0, 1]^3$ is $\mathbf{x}_0$.

Since $\mathcal{I}(B)$ is an open set, it follows that for each $i = 1, \dots, m_0$, there exists a real number $r_i > 0$ such that $B(\mathbf{p}_i; r_i) \subset \mathcal{I}(B)$, where $B(\mathbf{p}_i; r_i)$ denotes the open ball with centre $\mathbf{p}_i$ and radius r_i . Let $r = \min\{r_1, \dots, r_{m_0}\} > 0$. Then the open balls

$$B'_i = B(\mathbf{p}_i; r) \subset \mathcal{I}(B), \quad i = 1, \dots, m_0,$$

have common image $B(\mathbf{x}_0; r) \subset [0, 1]^3$ under modulo 1 projection from $\mathcal{M}$ to the unit torus $[0, 1]^3$. $\square$

Let $B'_{m_0+1}, \dots, B'_s$ denote open balls in the other $s - m_0$ atomic cubes of $\mathcal{M}$ such that $B'_1, \dots, B'_s$ have common image under modulo 1 projection from $\mathcal{M}$ to the unit torus $[0, 1]^3$. Observe that all these open balls have the same relative position within their own atomic cubes.

It is convenient to 2-colour the polycube translation 3-manifold $\mathcal{M}$ as follows. We colour the $\mathbf{v}^*$ -flow-invariant set $\mathcal{I}(B)$ *white* and colour its complement *silver*. Then each of the open balls $B'_1, \dots, B'_{m_0}$ is almost everywhere white, while each of the open balls $B'_{m_0+1}, \dots, B'_s$ is almost everywhere silver.

For technical reasons which will become clear later, we consider a collection of open balls $B_1, \dots, B_s$, where for every $i = 1, \dots, s$, the open ball B_i has the same centre as the open ball B'_i , but has radius equal to half of the radius of B'_i . Of course, each of the open balls $B_1, \dots, B_{m_0}$ is almost everywhere white, while each of the open balls $B_{m_0+1}, \dots, B_s$ is almost everywhere silver.

We next spread the open balls $B_1, \dots, B_s$ with the $\mathbf{v}^*$ -flow. Suppose that $B_\dagger$ is the common image of $B_1, \dots, B_s$ under modulo 1 projection from $\mathcal{M}$ to the unit torus $[0, 1]^3$. Then the projected $\mathbf{v}^*$ -flow spreads $B_\dagger$ over the unit torus $[0, 1]^3$ splitting free such that every point in $[0, 1]^3$ is covered. This can be achieved in a sufficiently large finite time, in view of the Kronecker density theorem.

Consider now what happens to the open balls $B_1, \dots, B_s$ under the $\mathbf{v}^*$ -flow. Since the system is in general non-integrable, there is splitting. We make use of this splitting, and we first explain this using a very simple example.

Consider the polysquare translation surface with 4 atomic squares in the picture on the left in Figure 2. Here the vertical edges on the bottom row of atomic squares are barriers. We now consider the polycube translation 3-manifold $\mathcal{M} = \mathcal{P} \times [0, 1]$ which is the cartesian product of the surface $\mathcal{P}$ with the unit torus $[0, 1]$. Here the X -faces on the bottom row of atomic cubes are barriers, as shown in the picture on the right in Figure 2.




Figure 2: a very simple polycube translation 3-manifold $\mathcal{M} = \mathcal{P} \times [0, 1]$

Figure 3, where the integrable direction corresponding to the unit torus $[0, 1]$ is suppressed, shows 4 open balls $B_1, \dots, B_4$ in the same relative positions within their

own atomic cubes. The lightgray colour shows part of their passages under the $\mathbf{v}^*$ -flow when they encounter singularities caused by the barriers on the bottom X -faces. Let us see what happens to the open ball B_1 . Clearly it is cut into 2 parts by a plane containing the vector $\mathbf{v}^*$ and the top edge of barrier which is the middle bottom X -face. The part labelled $B_1(+)$ continues unhindered and eventually reaches $B_1^*(+)$. The part labelled $B_1(-)$ hits the middle bottom X -face, then bounces to the left bottom X -face and eventually reaches $B_1^*(-)$. However, $B_1^*(+)$ and $B_1^*(-)$ are now in different atomic cubes. Instead, $B_1^*(+)$ now forms an open ball with $B_2^*(-)$. In summary, each of $B_1, \dots, B_4$ is cut into 2 parts in an identical way, and each individual part may after splitting pair up with the complementary part of a different open ball to form an open ball.




Figure 3: explaining the splitting method on $\mathcal{M} = \mathcal{P} \times [0, 1)$

After this example, we return to the general problem concerning the open balls $B_1, \dots, B_s$. Each of these open balls is monochromatic. After encountering the splittings, we have s new open balls, some of which may not be monochromatic any more, as the 2 individual parts of a new open ball may come from distinctly coloured open balls.

Recall that the projected $\mathbf{v}^*$ -flow spreads $B_{\dagger}$ over the unit torus $[0, 1)^3$ splitting free such that every point in $[0, 1)^3$ is covered, and this can be achieved in finite time. Let us move the open balls $B_1, \dots, B_s$ by the $\mathbf{v}^*$ -flow for this finite time, and call this the finite $\mathbf{v}^*$ -flow process. Splitting will occur a finite number of times during this finite process. There are two possibilities:

Case 1. At some stage of the finite $\mathbf{v}^*$ -flow process, there is a colour split open ball. Thus a positive proportion of this ball is white, and a positive proportion of this ball is silver.

Case 2. After each splitting, all the new open balls remain monochromatic. Thus there is no colour split balls at any stage of the finite $\mathbf{v}^*$ -flow process.

We show that both possibilities lead to a contradiction.

Case 1: colour split. We first show that colour splitting is not reversible, in the sense that any further splitting cannot lead to a monochromatic open ball.

Lemma 1.2. *If there is a colour split open ball at any stage of the finite $\mathbf{v}^*$ -flow process, then there is a colour split open ball at the end of the finite $\mathbf{v}^*$ -flow process.*

Proof. It suffices to consider the effect of a splitting edge in the y -direction, as the planes that cut S into 2 parts corresponding to splitting edges in the x -, y - and z -directions are perpendicular to the directions given by the vector products

$$\mathbf{v}^* \times (1, 0, 0), \quad \mathbf{v}^* \times (0, 1, 0), \quad \mathbf{v}^* \times (0, 0, 1),$$

respectively, and so are not parallel.

Suppose that a point of an open ball S hits some splitting edge of the polycube translation 3-manifold $\mathcal{M}$ in the y -direction at a point (m_1, y_1, n_1) , where for the sake of convenience, we take $m_1, n_1 \in \mathbb{Z}$. Then this point belongs to the plane that cuts S into 2 parts. As S continues in the direction $\mathbf{v}^* = (\alpha_1, \alpha_2, 1)$ of the flow, this point takes on values of the form

$$(m_1, y_1, n_1) + t(\alpha_1, \alpha_2, 1), \quad t \geq 0. \quad (1.1)$$

To establish the lemma, it suffices to show that the half-infinite geodesic (1.1) never hits a splitting edge in the y -direction again. Suppose, on the contrary, that it does. Then there exist $t \in \mathbb{R}$, $m_2, n_2 \in \mathbb{Z}$ and $y_2 \in \mathbb{R}$ such that

$$(m_1, y_1, n_1) + t(\alpha_1, \alpha_2, 1) = (m_2, y_2, n_2), \quad \text{so that} \quad \alpha_1 = \frac{m_2 - m_1}{n_2 - n_1},$$

contradicting the assumption that $\mathbf{v}^* = (\alpha_1, \alpha_2, 1)$ is a Kronecker direction. $\square$

Let $B'_\dagger$ be the common image of $B'_1, \dots, B'_s$ under modulo 1 projection from $\mathcal{M}$ to the unit torus $[0, 1)^3$. In view of the Kronecker density theorem, the centre of the open ball $B_\dagger$ will move under the projected $\mathbf{v}^*$ -flow in the unit torus $[0, 1)^3$ to a point close to the centre of the open ball $B'_\dagger$. This means that the open ball $B_\dagger$ will move under the projected $\mathbf{v}^*$ -flow in the unit torus $[0, 1)^3$ to inside the bigger open ball $B'_\dagger$.

The implication of this is that the colour split open ball will eventually move under the $\mathbf{v}^*$ -flow in $\mathcal{M}$ to within one of the monochromatic open balls $B'_1, \dots, B'_s$, and this is clearly absurd.

Case 2: no colour split. In this case, the $\mathbf{v}^*$ -flow always exhibits m_0 white open balls and $s - m_0$ silver open balls. The $\mathbf{v}^*$ -flow spreads the open balls $B_1, \dots, B_s$ and eventually covers the whole polycube translation 3-manifold $\mathcal{M}$. Within any atomic cube of $\mathcal{M}$, consider any two overlapping open balls that are the $\mathbf{v}^*$ -flow images of two of $B_1, \dots, B_s$. Since both are monochromatic, they must have the same colour. Thus this colour-coincidence spreads within the atomic cube, and so it follows that every atomic cube of $\mathcal{M}$ is monochromatic. Thus there are m_0 white atomic cubes and $s - m_0$ silver atomic cubes. However, since $\mathcal{M}$ is face-connected, there must be 2 neighbouring atomic cubes where one is white and the other is silver. Then the $\mathbf{v}^*$ -flow transports a set of positive volume from an atomic cube of one colour to an atomic cube of another colour, contradicting that the colour is $\mathbf{v}^*$ -flow invariant.

This completes the proof of the theorem. $\square$

2. WHEN FACES HAVE GATES AND BARRIERS

The idea of the proof of Theorem 1 can be extended to prove the following more general result.

Let $\mathcal{M}$ be modified from a finite polycube translation 3-manifold as follows. Each square face is in part a barrier, coloured red (shaded), and in part a gate, coloured green (white). The latter permits travel between the two cubes that share the common square face. We then modify boundary identification accordingly to ensure that $\mathcal{M}$ is boundary free and remains connected.

Theorem 2. *Let $\mathcal{M}$ be any polycube 3-manifold with s atomic cubes and barriers, where each square face has a 2-colouring such that each of the red and green parts is the union of finitely many polygons and $\mathcal{M}$ is connected. Then for almost every direction $\mathbf{v}^* \in \mathbb{R}^3$, any half-infinite geodesic with direction $\mathbf{v}^*$ is dense in $\mathcal{M}$.*

Proof. We need to modify the proof of Theorem 1 appropriately at a number of points. These concern those planes that cut the open balls B into 2 parts. Apart from arising from the edges of $\mathcal{M}$, these now also arise from the edges of the polygons on the faces of $\mathcal{M}$. Let us first deal with the latter, and let us restrict ourselves first to those that are on the X -faces of $\mathcal{M}$.

The collection of polygons on one X -face may differ significantly from that on another X -face. To handle this, we first observe that not all edges of the atomic cubes of $\mathcal{M}$ are singularities, but we still consider *splitting* on them. If no splitting actually occurs, then each part of an open ball after splitting reunites with its complement.

Let Γ_i denote the collection of colour split boundaries arising from the green and red polygons on the right X -face of the i -th atomic cube of $\mathcal{M}$, and let Γ'_i denote the image of Γ_i under modulo 1 projection from $\mathcal{M}$ to the unit torus $[0, 1)^3$. Then Γ'_i lies on the right X -face of the unit torus $[0, 1)^3$. Let

$$\Gamma = \bigcup_{i=1}^s \Gamma'_i,$$

and now replace Γ_i on the right X -face of the i -th atomic cube by Γ_i^* , where the image of Γ_i^* under modulo 1 projection from $\mathcal{M}$ to the unit torus $[0, 1)^3$ is Γ . Thus all the right X -faces of the atomic cubes of $\mathcal{M}$ have essentially the same collection of *splitting edges* arising from the green and red polygons. We refer to these as the *X -face splitting edges*.

Similarly there are *Y -face splitting edges* and *Z -face splitting edges*.

We have already commented at the beginning of the proof of Lemma 1.2 that the planes that cut S into 2 parts corresponding to splitting edges in the x -, y - and z -directions are not parallel. Consider now an X -face splitting edge E . The plane corresponding E that cuts S into 2 parts contains both $\mathbf{v}^*$ and E . We need to be careful in case this plane coincides with one of the following:

- (1) the plane corresponding to an edge of an atomic cube in direction z ;
- (2) the plane corresponding to an edge of an atomic cube in direction y ;
- (3) the plane corresponding to an edge of an atomic cube in direction x ;
- (4) the plane corresponding to a Y -face splitting edge E_y ;
- (5) the plane corresponding to a Z -face splitting edge E_z ;
- (6) the plane corresponding to a parallel X -face splitting edge E_x ; or
- (7) the plane corresponding to a non-parallel X -face splitting edge.

In case (1), the common plane must contain $\mathbf{v}^*$, E and the edge of the atomic cube in direction z . This can only happen if E is in the direction z and the $\mathbf{v}^*$ -flow takes E to this edge. It follows that the vector $\mathbf{v}^*$ can take only one of countably many different directions in $\mathbb{R}^3$. A similar conclusion can be drawn in case (2).

In case (3), the common plane must contain $\mathbf{v}^*$, E and the edge of the atomic cube in direction x . Note that E is not parallel to the edge of the atomic cube in direction x . This case can therefore only happen if the vector $\mathbf{v}^*$ lies on the plane containing E and the edge of the atomic cube in direction x . This gives rise to a set of directions in $\mathbb{R}^3$ of measure 0.

In case (4), the common plane must contain $\mathbf{v}^*$, E and E_y . Note that E is not parallel to E_y . This case can therefore only happen if the vector $\mathbf{v}^*$ lies on the plane containing E and E_y . This gives rise to a set of directions in $\mathbb{R}^3$ of measure 0. A similar conclusion can be drawn in case (5).

In case (6), the common plane must contain $\mathbf{v}^*$, E and E_x . Since E and E_x are parallel, this can only happen if the $\mathbf{v}^*$ -flow takes E to E_x . It follows that the vector $\mathbf{v}^*$ can take only one of countably many different directions in $\mathbb{R}^3$.

In case (7), the plane containing the two X -face splitting edges has constant x -coordinate, and so can only contain vectors of the form $(0, \alpha_1, \alpha_2) \in \mathbb{R}^3$. These give rise to a set of directions in $\mathbb{R}^3$ of measure 0.

Note that there are only finitely many X -, Y - or Z -face splitting edges. Hence the above give rise to a set of exceptional directions in $\mathbb{R}^3$ of measure 0. In order to complete the proof of Theorem 2, it suffices to establish the following analogue of Lemma 1.2.

Lemma 2.1. *Let E be a Z -face splitting edge, and let $\mathcal{E}(E)$ denote the set of exceptional directions $\mathbf{v}^* = (\alpha_1, \alpha_2, 1) \in \mathbb{R}^3$ such that the $\mathbf{v}^*$ -flow takes some point of E back to E or to some Z -face splitting edge analogous to E . Then the set $\mathcal{E}(E)$ has measure 0.*

Proof. The edge E can be described by a linear equation $c_1x + c_2y = c_3$, where the coefficients $c_1, c_2, c_3 \in \mathbb{R}$ are fixed. Consider a point (x_1, y_1, m_1) on E , where $x_1, y_1 \in \mathbb{R}$ and for the sake of convenience, we take $m_1 \in \mathbb{Z}$. Then

$$c_1x_1 + c_2y_1 = c_3.$$

The effect of the $\mathbf{v}^*$ -flow on this point is described by the half-infinite geodesic

$$(x_1, y_1, m_1) + t(\alpha_1, \alpha_2, 1), \quad t \geq 0.$$

Suppose that this half-infinite geodesic takes the point (x_1, y_1, m_1) back to E or to some Z -face splitting edge analogous to E . Then there exist $t \in \mathbb{R}$, $m_2, n_2, q_2 \in \mathbb{Z}$ and $x_2, y_2 \in \mathbb{R}$ such that

$$(x_1, y_1, m_1) + t(\alpha_1, \alpha_2, 1) = (x_2, y_2, m_2) \quad \text{and} \quad c_1(x_2 - n_2) + c_2(y_2 - q_2) = c_3,$$

from which we deduce that

$$c_1(m_2 - m_1)\alpha_1 + c_2(m_2 - m_1)\alpha_2 = c_1n_2 + c_2q_2, \quad (2.1)$$

where c_1 and c_2 are fixed and m_1 is determined by E . For any fixed choice of the integers (m_2, n_2, q_2) , the equation (2.1) shows that the exceptional points (α_1, α_2) must lie on a line. Since there are only countably many choices for the integers (m_2, n_2, q_2) , it follows that the set $\mathcal{E}(E)$ can be described as a countable union of lines, and so has measure 0. $\square$

This completes the proof of Theorem 2. $\square$

Remark. We can compare Lemmas 1.2 and 2.1. In the former, we see from the proof that the exceptional set of directions are the non-Kronecker directions, whereas in the latter, this is not always the case.

Suppose that in the statement of Theorem 2, we add an extra condition and require the edges of the finitely many polygons to have *rational* slopes. Then in the proof of Lemma 2.1, the coefficients c_1 and c_2 of the line $c_1x + c_2y = c_3$ can be taken to be rational. Then the equation (2.1) becomes an equation in the variables α_1 and α_2 with rational coefficients, so that any solution leads to a non-Kronecker direction.

This observation suggests that face splitting edges with rational slopes behave in a similar fashion to edges of the atomic cubes, whereas face splitting edges with irrational slopes are somewhat more delicate.

3. WEAK UNIFORMITY WITH BOUNDED RETURNS

For Theorem A in its full generality, we can go beyond the class of finite flat rational surfaces, and consider the larger class of finite *translation surfaces*, where the restriction on the angles is dropped. A finite translation surface $\mathcal{P}$ is a flat surface with finitely many faces and where every face is a polygon, and parallel

edges with the same length are identified in pairs, while identified edges from the same face have opposite orientation.

In this generality, the directions in $\mathcal{P}$ such that no geodesic on $\mathcal{P}$ can contain more than one vertex of $\mathcal{P}$ defines the *good* directions and plays the perfect analogue to irrational slopes in the special case when $\mathcal{P}$ is a polysquare translation surface. It remains to show that the set of *bad* directions is countable. A geodesic on the unit torus $[0, 1)^2$ can be extended to a straight line on the plane. Here every geodesic on $\mathcal{P}$ containing at least two vertices of $\mathcal{P}$ can be extended to a straight line on a countable covering of the plane. Each plane on this countable covering contains a countable number of extended vertices. Since two points determine a straight line, the set of bad directions is countable.

For a finite translation surface $\mathcal{P}$, we can identify the bad directions in terms of *saddle connections*. A finite geodesic segment on a finite translation surface $\mathcal{P}$ is called a saddle connection if both endpoints of the geodesic segment are vertices of $\mathcal{P}$ and there are no vertices of $\mathcal{P}$ in between. The set of saddle connections of a finite translation surface is countable. Furthermore, the collection of saddle connections with a given fixed slope is finite. Then the exceptional directions in Theorem A are precisely those directions of the saddle connections of $\mathcal{P}$.

Let $\mathcal{P}$ be a finite translation surface. Density of a half-infinite geodesic on $\mathcal{P}$, as given by the full generality of Theorem A, means that the orbit of the geodesic visits every non-empty open set in $\mathcal{P}$. Since the orbit is infinite, this trivially means that it returns to this open set infinitely many times.

A less trivial consequence of density of a half-infinite geodesic is a new concept which we call *weak uniformity with bounded returns*. This means that for every non-empty open set $G \subset \mathcal{P}$, there is a constant $\mathfrak{F} = \mathfrak{F}(G; \mathcal{P}; \alpha) > 0$, where α is the slope of the geodesic, and a finite threshold $c_5 = c_5(G; \mathcal{P}; \alpha) > 0$ such that for every geodesic segment L with slope α and length $|L| \geq c_5$, the inequality

$$\lambda_1(G \cap L) \geq \mathfrak{F}|L|$$

holds, where λ_1 denotes 1-dimensional Lebesgue measure, so that L visits G in total length at least $\mathfrak{F}|L|$. Intuitively, $\mathfrak{F} > 0$ gives a lower bound to what may be called the *visiting frequency* of G relative to geodesic segments of slope α .

Every open set in a finite rational translation surface $\mathcal{P}$ is Lebesgue measurable. However, some open sets $G \subset \mathcal{P}$ are not Jordan measurable, and the characteristic function χ_G has no well defined 2-dimensional Riemann integral. The traditional Weyl type uniformity requires that the asymptotic visiting frequency of G is precisely equal to the relative volume of G , assuming that G is Jordan measurable. This motivates us to describe this new concept as a weak form of uniformity.

We establish the following result.

Theorem 3. *Suppose that α is not the slope of any saddle connection of a finite translation surface $\mathcal{P}$. Then every non-empty open set $G \subset \mathcal{P}$ has positive visiting frequency relative to geodesic segments of slope α .*

Before we start the proof, let us draw some simple and useful consequences from Theorem A.

Let α be fixed and not equal to the slope of any saddle connection of $\mathcal{P}$. Consider the direction $(1, \alpha)$. A point $Q \in \mathcal{P}$ is said to be a pathological starting point in this direction if a geodesic in this direction and starting from Q hits a singular vertex of $\mathcal{P}$ and becomes undefined. Otherwise $Q \in \mathcal{P}$ is said to be a non-pathological starting point in this direction.

Let $Q \in \mathcal{P}$ be a non-pathological starting point of a half-infinite geodesic $\mathcal{L}(t)$, $t \geq 0$, in the direction $(1, \alpha)$, so that $\mathcal{L}(0) = Q$. Then Theorem A implies that for

every non-empty open set $G \subset \mathcal{P}$, there is a finite threshold $T = T(G; Q)$ such that

$$\{\mathcal{L}(t) : 0 \leq t \leq T(G; Q)\} \cap G \neq \emptyset,$$

so that a geodesic segment of length $T = T(G; Q)$ in the direction $(1, \alpha)$ and starting from the point Q visits the set G .

Proof of Theorem 3. Suppose that α is fixed and not equal to the slope of any saddle connection of $\mathcal{P}$. Let $G_1 \subset \mathcal{P}$ be an open disk of radius $r > 0$, and consider the set

$$\{T(G_1; Q) : Q \in \mathcal{P} \text{ not pathological in direction } (1, \alpha)\}. \quad (3.1)$$

We first show that this set is bounded.

Suppose, on the contrary, that the set (3.1) is not bounded. Then there exists a sequence of points $Q_i \in \mathcal{P}$, $i = 1, 2, 3, \dots$, such that $T(G_1; Q_i) \rightarrow \infty$ as $i \rightarrow \infty$. Since $\mathcal{P}$ is compact, there exists a subsequence Q_{i_j} , $j = 1, 2, 3, \dots$, such that $Q_{i_j} \rightarrow Q_\infty$ as $j \rightarrow \infty$ for some point $Q_\infty \in \mathcal{P}$ and

$$T(G_1; Q_{i_j}) \rightarrow \infty \quad \text{as } j \rightarrow \infty. \quad (3.2)$$

The limit point $Q_\infty \in \mathcal{P}$ may be non-pathological or pathological.

Suppose that $Q_\infty \in \mathcal{P}$ is pathological. This means that either Q_∞ is a singular vertex of $\mathcal{P}$ and so there is no unique geodesic in the direction $(1, \alpha)$ that starts from Q_∞ , or the geodesic in the direction $(1, \alpha)$ and starting from Q_∞ hits a singular vertex of $\mathcal{P}$ and becomes undefined thereafter. However, there are only finitely many ways to extend the geodesic in the direction $(1, \alpha)$ and starting from Q_∞ beyond the singular vertex.

Let $G_2 \subset \mathcal{P}$ be an open disk of radius $r/2$ and with the same centre as G_1 .

Suppose that $\mathcal{L}^{(\kappa)}$ is one such extension of the geodesic from Q_∞ . Let $Q^{(\kappa)} \in \mathcal{P}$ be a point beyond the singular vertex on $\mathcal{L}^{(\kappa)}$. Since α is not the slope of any saddle connection of $\mathcal{P}$, it follows that $Q^{(\kappa)}$ is non-pathological. Let $d = d(Q_\infty, Q^{(\kappa)})$ be the length of the extended geodesic segment of $\mathcal{L}^{(\kappa)}$ from Q_∞ to $Q^{(\kappa)}$. Then an initial geodesic segment of $\mathcal{L}^{(\kappa)}$ of length $d + T(G_2; Q^{(\kappa)})$ visits G_2 . Furthermore, apart from very short segments near the singular vertex it intersects, this finite geodesic segment of $\mathcal{L}^{(\kappa)}$ from Q_∞ to G_2 has a positive distance $\delta^{(\kappa)}$ from other singular vertices of $\mathcal{P}$. Let $\delta_\kappa = \min\{\delta^{(\kappa)}, r/2\}$.

Consider next a half-infinite geodesic $\mathcal{L}_j$ in the direction $(1, \alpha)$ and starting from a point Q_{i_j} in the δ_κ -neighbourhood of Q_∞ . If $\mathcal{L}_j$ does not split from $\mathcal{L}^{(\kappa)}$ up to $Q^{(\kappa)}$, then it does not split from $\mathcal{L}^{(\kappa)}$ before it reaches G_1 . More precisely, if an extended geodesic segment of length $d + T(G_2; Q^{(\kappa)})$ starting from Q_∞ reaches a point $Q_\infty^* \in G_2$, then a geodesic segment of length $d + T(G_2; Q^{(\kappa)})$ starting from Q_{i_j} reaches a point in the δ_κ -neighbourhood of Q_∞^* . It is clear that any point in this δ_κ -neighbourhood of Q_∞^* is a distance less than $\delta_\kappa + r/2 \leq r$ from the centre of G_2 , which is also the centre of G_1 , and so is contained in G_1 . It follows that

$$T(G_1; Q_{i_j}) \leq d + T(G_2; Q^{(\kappa)}). \quad (3.3)$$

There are only finitely many possible extensions $\mathcal{L}^{(\kappa)}$, $\kappa \in \mathcal{K}$, of the geodesic from Q_∞ . Let

$$\delta^* = \min_{\kappa \in \mathcal{K}} \delta_\kappa > 0.$$

Then there are clearly infinitely many values of $j = 1, 2, 3, \dots$ such that Q_{i_j} is in the δ^* -neighbourhood of Q_∞ . It follows that there exists $\kappa \in \mathcal{K}$ such that the inequality (3.3) holds for infinitely many values of j . This clearly contradicts (3.2).

For the case when Q_∞ is not pathological, there is a half-infinite geodesic in the direction $(1, \alpha)$ that starts from Q_∞ . This essentially means that there is only one value of κ , and we simply take $Q^{(\kappa)} = Q_\infty$ and $d = 0$.

This completes the deduction that the set (3.1) is bounded. It is also clear that a similar argument shows that both sets

$$\{T(G_2; Q) : Q \in \mathcal{P} \text{ not pathological in direction } (1, \alpha)\}$$

and

$$\{T(G_2; Q) : Q \in \mathcal{P} \text{ not pathological in direction } (-1, -\alpha)\}$$

are bounded. Let $T^*(G_2)$ be a finite upper bound of the two sets.

A point $Q \in \mathcal{P}$ cannot be pathological in both directions $(1, \alpha)$ and $(-1, -\alpha)$, unless Q is a singular vertex of $\mathcal{P}$. It then follows that every geodesic segment of slope α and length $T^*(G_2)$ that does not contain a singular vertex of $\mathcal{P}$ visits G_2 , and so visits G_1 in total length at least $r/2$.

Next, let $G \subset \mathcal{P}$ be a non-empty open set. Then there exists an open disk $G_1 \subset G$ with some radius $r > 0$. Let L be a geodesic segment on $\mathcal{P}$ with length

$$|L| \geq 2T^*(G_2). \quad (3.4)$$

If L contains a singular vertex of $\mathcal{P}$, then this can only be one of the two endpoints. It follows that L contains at least

$$\left\lceil \frac{|L|}{T^*(G_2)} \right\rceil - 1 = k \geq 1$$

distinct geodesic segments $L_1, \dots, L_k$, each of length $T^*(G_2)$ and not containing a singular vertex of $\mathcal{P}$. Each of $L_1, \dots, L_k$ visits G_1 in total length at least $r/2$, so L visits G in total length at least

$$\frac{rk}{2} = \frac{r}{2} \left(\left\lceil \frac{|L|}{T^*(G_2)} \right\rceil - 1 \right) \geq \frac{r}{4} \left\lceil \frac{|L|}{T^*(G_2)} \right\rceil \geq \mathfrak{F}|L|, \quad (3.5)$$

where

$$\mathfrak{F} = \mathfrak{F}(G; \mathcal{P}; \alpha) = \frac{r}{8T^*(G_2)}. \quad (3.6)$$

The theorem follows on combining (3.4)–(3.6). $\square$

Using similar ideas, we can establish the following result.

Theorem 4. *Let $\mathcal{M}$ be a finite polycube translation 3-manifold, and let $\mathbf{v}$ be a Kronecker direction. Then every non-empty open set $G \subset \mathcal{M}$ has positive visiting frequency.*

REFERENCES

- [1] A. Katok, A. Zemlyakov. Topological transitivity of billiards in polygons. *Math. Notes* **18** (1975), 760–764.

DEPARTMENT OF MATHEMATICS, HILL CENTER FOR THE MATHEMATICAL SCIENCES, RUTGERS UNIVERSITY, PISCATAWAY NJ 08854, USA

Email address: jbeck@math.rutgers.edu

SCHOOL OF MATHEMATICAL AND PHYSICAL SCIENCES, FACULTY OF SCIENCE AND ENGINEERING, MACQUARIE UNIVERSITY, SYDNEY NSW 2109, AUSTRALIA

Email address: william.chen@mq.edu.au

SCHOOL OF SCIENCE, BEIJING UNIVERSITY OF POSTS AND TELECOMMUNICATIONS, BEIJING 100876, CHINA

Email address: yangyx@bupt.edu.cn