Moving Out the Edges of a Lattice Polygon

Wouter Castryck

Received: 23 January 2011 / Revised: 7 July 2011 / Accepted: 8 August 2011 /

Published online: 4 October 2011

© Springer Science+Business Media, LLC 2011

Abstract We review previous work of (mainly) Koelman, Haase and Schicho, and Poonen and Rodriguez-Villegas on the dual operations of (i) taking the interior hull and (ii) moving out the edges of a two-dimensional lattice polygon. We show how the latter operation naturally gives rise to an algorithm for enumerating lattice polygons by their genus. We then report on an implementation of this algorithm, by means of which we produce the list of all lattice polygons (up to equivalence) whose genus is contained in $\{1, \dots, 30\}$. In particular, we obtain the number of inequivalent lattice polygons for each of these genera. As a byproduct, we prove that the minimal possible genus for a lattice 15-gon is 45.

Keywords Lattice polygon · Genus · Legal loop

1 Introduction

(1.1) A *lattice polygon* is a (nonempty) convex polygon in $\mathbb{R}^2$ with vertices in $\mathbb{Z}^2$. Points of $\mathbb{Z}^2$ are called *lattice points*. The *dimension* of a polygon Δ is the dimension of the smallest affine subspace of $\mathbb{R}^2$ containing Δ . The *genus* of a two-dimensional lattice polygon is the number of lattice points in its topological interior (when equipped with the subspace topology of $\mathbb{R}^2$). The genus of a lower-dimensional lattice polygon is considered 0. A $\mathbb{Z}$ -*affine transformation* of $\mathbb{R}^2$ is a map of the form $\varphi : \mathbb{R}^2 \rightarrow \mathbb{R}^2 : p \mapsto pA + b$ with $A \in \mathrm{GL}_2(\mathbb{Z})$ and $b \in \mathbb{Z}^2$. Two lattice polygons Δ and Δ' are called *equivalent* if and only if there exists a $\mathbb{Z}$ -affine transformation φ such that $\varphi(\Delta) = \Delta'$.

W. Castryck (✉)

Departement Wiskunde, Katholieke Universiteit Leuven, Celestijnenlaan 200B, 3001 Leuven (Heverlee), Belgium

e-mail: wouter.castryck@gmail.com

(1.2) We review a useful tool in the study of the combinatorics of lattice polygons. The rough idea is to gradually peel off the lattice polygon by consecutively considering the convex hull of the interior lattice points. Although this ‘onion skin’ principle dates (at least) back to the work of Rabinowitz [24] and Koelman [15], Haase and Schicho [12] noticed that to each step of the peeling, one can associate a so-called legal loop catching the corresponding global information. This allows one to apply a remarkable theorem due to Poonen and Rodriguez-Villegas [23], in which the number 12 makes an intriguing appearance.

The interior hull of a lattice polygon and the associated legal loop

Whereas Haase and Schicho worked toward a refinement of a theorem by Scott [25], our endpoint is a new and conceptual proof of a more precise conjecture due to Coleman [8].

Theorem 1 (Coleman’s conjecture, 1978) *Let Δ be a lattice n -gon of genus $g \geq 1$. Let R be the number of lattice points on the boundary of Δ . Then $R \leq 2g + 10 - n$.*

The first complete proof was provided in 2006 by Kołodziejczyk and Olszewska [17]. However, already in his 1991 Ph.D. thesis [15, Lemma 4.5.2(2)], Koelman must have been unaware of the existence of this conjecture and was only one sentence left from a proof. His argument heavily relies on another ‘12 theorem’ due to Oda [21, Remark on p. 45], and can in fact be extended to cover a non-trivial part of Poonen and Rodriguez-Villegas’ result, see (2.6). In their turn, Haase and Schicho must have been unaware of this entire story: our proof will merely add a couple of lines to their arguments. We therefore do not claim many credits, but hope that this proof gives an indication of the unacknowledged potential of the machinery. At the same time, we smoothen the theoretical and historical framework.

(1.3) We then reverse the process of gradually peeling off a lattice polygon by instead consecutively moving out its edges, following ideas that were discovered by Koelman [15, Sect. 2.2], Haase and Schicho [12], and Kołodziejczyk and Olszewska [16]. This gives a natural and efficient way of enumerating lattice polygons by their genus, up to equivalence. We will report on an implementation of this procedure using the MAGMA computer algebra system [5], by means of which we produced the list of all equivalence classes of lattice polygons of genus $1 \leq g \leq 30$. This comprises approximately 368 MB of data that we made available for download at <http://wis.kuleuven.be/algebra/castryck/>. As a consequence, we can now answer vir-

tually every reasonable question on lattice polygons of genus $1 \leq g \leq 30$. In particular, we obtain the number of inequivalent lattice polygons for each of these genera. Up to our knowledge, these numbers did not appear in the literature thus far, even for g as small as 3. Among the other consequences, we find the following.

Theorem 2 *The minimal genus of a lattice 15-gon is 45.*

This fills in the smallest open entry of a list whose study was initiated by Arkinstall [1] and that since invoked fair interest. E.g. only recently, it was proven that the minimal genus of a lattice 11-gon is 17 (see [22]). Finally, we introduce the *lifespan* of a lattice polygon, which measures how often its edges can be moved out without tumbling off the lattice. We prove the following fact:

Theorem 3 *A lattice n -gon has finite lifespan as soon as $n \geq 10$.*

For each $3 \leq n \leq 9$, there exists an n -gon having infinite lifespan. Explicit examples are provided in (3.6) below.

2 Legal Loops Associated to a Lattice Polygon

(2.1) In 1976, Scott proved the following theorem [25].

Theorem 4 (Scott, 1976) *Let $\Delta \subset \mathbb{R}^2$ be a lattice polygon having $g \geq 1$ lattice points in its interior. Let R be the number of lattice points on the boundary of Δ . Then $R \leq 2g + 7$.*

Moreover, Scott proved that equality holds if and only if Δ is equivalent to $\text{Conv}\{(0, 0), (3, 0), (0, 3)\}$. Two years later, Coleman conjectured that the refinement mentioned in Theorem 1 of (1.2) should hold. The usage of the word ‘theorem’ is now justified, due to a recent proof by Kołodziejczyk and Olszewska [17]. However, as already mentioned, it should be attributed in part to Koelman [15, Lemma 4.5.2(2)]: see (2.6).

In 2009, Haase and Schicho revisited Scott’s bound and provided an alternative proof. Implicitly, though, they gave a new proof of Coleman’s conjecture, along with a substantial refinement of the statement. An explicit version of this proof will be given in (2.5). Along the way, we give an overview of the machinery used, adapt certain definitions, and prove a number of facts that were merely sketched in the literature before.

(2.2) The following two basic operations on lattice polygons will be crucial throughout.

Definition (Moving out the edges) Let $\Delta \subset \mathbb{R}^2$ be a two-dimensional lattice polygon. Then each of its edges $\tau \subset \Delta$ corresponds to a unique half-plane

$$\mathcal{H}_\tau = \{(x, y) \in \mathbb{R}^2 \mid a_\tau x + b_\tau y \leq c_\tau\}$$

jointly satisfying $\Delta = \bigcap_{\tau} \mathcal{H}_{\tau}$. In this set, $a_{\tau}, b_{\tau}, c_{\tau} \in \mathbb{Z}$ are uniquely determined by the condition $\gcd(a_{\tau}, b_{\tau}) = 1$. Then we define

$$\Delta^{(-1)} := \bigcap_{\tau} \mathcal{H}_{\tau}^{(-1)},$$

where

$$\mathcal{H}_{\tau}^{(-1)} = \{(x, y) \in \mathbb{R}^2 \mid a_{\tau}x + b_{\tau}y \leq c_{\tau} + 1\}.$$

We say that $\Delta^{(-1)}$ is obtained from Δ by *moving out the edges*.

Definition (Interior hull) Let $\Delta \subset \mathbb{R}^2$ be a two-dimensional lattice polygon of genus at least 1. Then we define $\Delta^{(1)}$ as the convex hull of the lattice points in the interior of Δ . We say that $\Delta^{(1)}$ is the *interior hull* of Δ .

We will abuse notation and write $\Delta^{(k)}$ for $\Delta^{(1)(1)\dots(1)}$ (the interior hull taken k times consecutively), given that this is well-defined: note that $\Delta^{(1)}$ need not have interior lattice points. Likewise, we will write $\Delta^{(-k)}$ for $\Delta^{(-1)(-1)\dots(-1)}$ (moving out the edges k times consecutively). Again, this may not be well-defined since $\Delta^{(-1)}$ need not be a lattice polygon: it may take vertices outside $\mathbb{Z}^2$. It is sometimes convenient to write $\Delta^{(0)}$ for Δ .

Examples where $\Delta^{(2)}$ resp. $\Delta^{(-2)}$ would not be well-defined

A crucial property of lattice polygons is the following.

Theorem 5 (Koelman, 1991) *Let $\Delta \subset \mathbb{R}^2$ be a two-dimensional lattice polygon, such that $\Delta^{(1)}$ is again two-dimensional. Then $\Delta^{(1)(-1)}$ is a lattice polygon containing Δ .*

Proof See Koelman [15, Lemma 2.2.13] or Haase–Schicho [12, Lemma 11]. $\square$

Note that this theorem gives a criterion for a two-dimensional lattice polygon Γ to satisfy that $\Gamma^{(-1)}$ is a lattice polygon: this will be the case *if and only if* there exists a lattice polygon Δ such that $\Delta^{(1)} = \Gamma$. Although the notions of moving out the facets and taking the interior hull straightforwardly generalize to higher dimensions, Theorem 5 does not. This is the main reason why we restrict to dimension two in this article.

Definition (Maximal polygon) A lattice polygon Δ for which $\Delta^{(1)}$ is two-dimensional is called *maximal* if $\Delta = \Delta^{(1)(-1)}$.

(2.3) We now review the theory of legal loops, in the sense of [23]. Throughout, we will write o for the origin $(0, 0) \in \mathbb{R}^2$.

Definition (Legal move) A *legal move* is a couple of points (p_1, p_2) with $p_1, p_2 \in \mathbb{Z}^2$ such that $\text{Conv}\{o, p_1, p_2\} \cap \mathbb{Z}^2 = \{o\} \sqcup (\text{Conv}\{p_1, p_2\} \cap \mathbb{Z}^2)$.

Note that $p_1 = p_2$ is a priori allowed. If $p_1 \neq p_2$, then the condition reads that the line connecting p_1 and p_2 lies at integral distance 1 from o , i.e. it has an equation of the form $aX + bY = 1$ for (necessarily coprime) $a, b \in \mathbb{Z}$.

Definition (Legal loop) A *legal loop* is a finite sequence $\mathcal{P} = (p_0, \dots, p_{n-1})$, where $n \geq 1$, such that (p_i, p_{i+1}) is a legal move for all $i = 0, \dots, n-1$. For any primitive vector $p_0 \in \mathbb{Z}^2$, the legal loop (p_0) will be called *trivial*.

In the above, indices should be considered modulo n , i.e. $p_n = p_0$. Such abuse of notation will be repeated throughout.

Definition (Length) The *length* of a legal move $s = (p_1, p_2)$ is defined to be

$$\det \begin{pmatrix} p_1 \\ p_2 \end{pmatrix}$$

and will be denoted by $\ell(s)$. The *length* of a legal loop $\mathcal{P}$ is the sum of the lengths of its legal moves and will be denoted by $\ell(\mathcal{P})$.

A legal loop $\mathcal{P} = (p_0, \dots, p_{n-1})$ gives rise to a closed curve $\gamma(\mathcal{P})$ in $\mathbb{R}^2 \setminus \{o\}$ by ‘connecting the dots’. One way of making this precise is

$$[0, 1] \rightarrow \mathbb{R}^2 \setminus \{o\} : t \mapsto (nt - \lfloor nt \rfloor)p_{\lfloor nt \rfloor + 1} + (1 - nt + \lfloor nt \rfloor)p_{\lfloor nt \rfloor},$$

although we will only be interested in $\gamma(\mathcal{P})$ up to homotopy.

Definition (Winding number) The *winding number* of a legal loop $\mathcal{P}$ is the winding number of $\gamma(\mathcal{P})$ around o in the sense of algebraic topology, i.e. the image of its homotopy class under the unique isomorphism $\pi_1(\mathbb{R}^2 \setminus \{o\}) \rightarrow \mathbb{Z}$ mapping the class of a counterclockwise loop around o to 1.

Definition (Inverse loop) Let $\mathcal{P} = (p_0, p_1, \dots, p_{n-1})$ be a legal loop. Then we define the *inverse loop* $\mathcal{P}^{-1}$ to be $(p_{n-1}, p_{n-2}, \dots, p_0)$.

Taking the inverse of a legal loop alters the sign of both the length and the winding number.

Definition (Equivalence) We equip the set of legal loops with the smallest equivalence relation satisfying

1. (*shifting*) a legal loop $(p_0, p_1, \dots, p_{n-1})$ is equivalent to $(p_1, \dots, p_{n-1}, p_0)$;
2. (*merging and splitting moves*) a legal loop $(p_0, p_1, \dots, p_{n-1})$ is equivalent to the legal loop $(p_0, q, p_1, \dots, p_{n-1})$, where q is any lattice point on a line through p_0 and p_1 at integral distance 1 from o ;
3. (*orientation-preserving lattice equivalence*) a legal loop $(p_0, p_1, \dots, p_{n-1})$ is equivalent to $(p_0 A, p_1 A, \dots, p_{n-1} A)$ for any matrix $A \in \text{SL}_2(\mathbb{Z})$.

One easily verifies that equivalence preserves the length and the homotopy class of the corresponding curve. For that reason, we can unambiguously talk about the length and the winding number of an equivalence class of legal loops $\overline{\mathcal{P}}$. The former will be denoted by $\ell(\overline{\mathcal{P}})$.

A pathological remark is that all trivial legal loops are equivalent. Indeed, if $p_0, p_1 \in \mathbb{Z}^2$ are primitive vectors, one can always find a distinct point $p_2 \in \mathbb{Z}^2$ such that both the line through p_0 and p_2 and the line through p_1 and p_2 lie at integral distance 1 from o . By merging and splitting, $(p_0) \sim (p_0, p_2) \sim (p_2) \sim (p_2, p_1) \sim (p_1)$. The corresponding equivalence class will be called the *trivial class*.

Definition (Dual class) Let $\overline{\mathcal{P}}$ be an equivalence class of legal loops. Take a representant $\mathcal{P} = (p_0, \dots, p_{n-1})$ for which $p_i \neq p_{i+1}$ for all $i = 0, \dots, n-1$. Define

$$q_i = \frac{p_{i+1} - p_i}{\det \begin{pmatrix} p_{i+1} \\ p_i \end{pmatrix}}.$$

Then the dual class $\overline{\mathcal{P}}^\vee$ is defined to be the class of $(q_0, \dots, q_{n-1})$.

The reader can check that this is well-defined. Note that the trivial class is self-dual. In the non-trivial case, one can take a representant having no two consecutive moves along the same line, by means of which one easily verifies that $\overline{\mathcal{P}}^{\vee\vee} = \overline{\mathcal{P}}$.

The above series of figures illustrates the construction of the dual class. The notation $\mathcal{P}^\vee$ should in principle be read as ‘a representant of $\overline{\mathcal{P}}^\vee$ ’, but since $\mathcal{P}$ does not contain any moves of length 0 we can unambiguously write $\mathcal{P}^\vee$. Note that the length of $\mathcal{P}$ is $1 - 3 + 1 + 3 + 1 + 1 = 4$, whereas the length of $\mathcal{P}^\vee$ is $1 + 1 + 3 + 1 - 1 + 3 = 8$.

Theorem 6 (Poonen and Rodriguez-Villegas, 2000) *Let $\mathcal{P}$ be a legal loop of winding number ω . Then $\ell(\overline{\mathcal{P}}) + \ell(\overline{\mathcal{P}}^\vee) = 12\omega$.*

Proof This is hinted at in the paper by Poonen and Rodriguez-Villegas [23], which contains the details of the case where $\mathcal{P}$ is the boundary of a lattice polygon of genus 1, ran through counterclockwise. Since the necessary adaptations for the general case are not entirely trivial, we include the details here.

The only external fact we need concerns the set $\widetilde{\mathrm{SL}_2(\mathbb{Z})}$, an element of which is a pair $(M, [\gamma])$. Here $M = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbb{Z})$ and γ is a homotopy class of paths in $\mathbb{R}^2 \setminus \{o\}$ from $(0, 1)$ to (c, d) . If $(c, d) \neq (0, -1)$, one can consider a straight-line path γ from $(0, 1)$ to (c, d) . In that case, we simply write $\tilde{M}$ instead of $(M, [\gamma])$. The set $\widetilde{\mathrm{SL}_2(\mathbb{Z})}$ is turned into a group by the rule

$$(M_1, [\gamma_1]) \cdot (M_2, [\gamma_2]) = (M_1 M_2, [\gamma_2 * \gamma_1^{M_2}]),$$

where $*$ is the concatenation of paths sharing an endpoint, and where $\gamma_1^{M_2}$ is the path obtained by composing γ_1 with $(a, b) \mapsto (a, b)M_2$. A prominent role is played by the element $(\mathbb{I}, \text{loop})$, where $\mathbb{I}$ is the identity matrix and ‘loop’ is the homotopy class of a counterclockwise loop around the origin. Namely, the property of $\widetilde{\mathrm{SL}_2(\mathbb{Z})}$ that we need is the existence of a group homomorphism $\Phi : \widetilde{\mathrm{SL}_2(\mathbb{Z})} \rightarrow \mathbb{Z}$ under which $(\mathbb{I}, \text{loop})$ is mapped to 12. This can be achieved in various ways; see [23, Sect. 8.4] for a fancy proof in which the ‘12’ appears as the weight of the modular discriminant.

By merging, splitting and switching to $\mathcal{P}^{-1}$ if necessary, we may assume that $\mathcal{P} = (p_0, \dots, p_{n-1})$ consists of moves of length 1 or -1 only, and that the move (p_0, p_1) has length 1. Then an orientation-preserving lattice transformation brings us to the case where $p_0 = (1, 0)$ and $p_1 = (0, 1)$. For $i = 0, \dots, n-1$, let s_i denote the legal move from p_i to p_{i+1} , and let $M_i \in \mathrm{SL}_2(\mathbb{Z})$ be defined inductively by

$$M_i \cdot M_{i-1} \cdots M_0 = \begin{pmatrix} \ell(s_{i+1}) \cdot p_{i+1} \\ p_{i+2} \end{pmatrix}.$$

In particular $M_{n-1} \cdots M_0 = \mathbb{I}$, but note that one even has

$$\tilde{M}_{n-1} \cdots \tilde{M}_0 = (\mathbb{I}, [\gamma(\mathcal{P})]).$$

It follows that

$$\sum_{i=0}^{n-1} \Phi(\tilde{M}_i) = 12\omega.$$

Now there are two types of M_i :

$$\text{either } M_i = \begin{pmatrix} 0 & 1 \\ -1 & d_i \end{pmatrix} \quad \text{or} \quad M_i = \begin{pmatrix} 0 & -1 \\ 1 & d_i \end{pmatrix},$$

depending on whether $\ell(s_{i+1}) = 1$ or $\ell(s_{i+1}) = -1$. Using that $\Phi((\mathbb{I}, \text{loop})) = 12$, one accordingly finds that

$$\Phi(\tilde{M}_i) = 3 - d_i \quad \text{resp.} \quad \Phi(\tilde{M}_i) = d_i - 3.$$

Thus we conclude

$$\sum_{i=0}^{n-1} \ell(s_{i+1}) \cdot (3 - d_i) = 12\omega. \quad (1)$$

On the other hand, if we define q_i as in the above definition, and if we let $s_i^\vee$ be the legal move from q_i to q_{i+1} , then d_i contains information about $\ell(s_i) + \ell(s_i^\vee)$. Namely,

1. if $\ell(s_i) = \ell(s_{i+1}) = 1$, then $\ell(s_i) + \ell(s_i^\vee) = 3 - d_i$;
2. if $\ell(s_i) = \ell(s_{i+1}) = -1$, then $\ell(s_i) + \ell(s_i^\vee) = d_i - 3$;
3. if $\ell(s_i) = 1$ and $\ell(s_{i+1}) = -1$, then $\ell(s_i) + \ell(s_i^\vee) = d_i - 1$;
4. if $\ell(s_i) = -1$ and $\ell(s_{i+1}) = 1$, then $\ell(s_i) + \ell(s_i^\vee) = 1 - d_i$.

Since a closed loop must switch as many times from being positively oriented to being negatively oriented as conversely, we find that

$$\sum_{i=0}^{n-1} \ell(s_i) + \ell(s_i^\vee) = \sum_{i=0}^{n-1} \ell(s_{i+1})(3 - d_i).$$

Together with (1), this concludes the proof. $\square$

Remark In case $\mathcal{P}$ is the boundary of a genus 1 polygon, Theorem 6 can be easily proven by exhaustively verifying it for the 16 representants of Theorem 10(b). In a quest for explaining the ‘12’, Poonen and Rodriguez-Villegas gave three alternative proofs. One of these is, essentially, the proof produced above. Of the four approaches, it seems best-suited for addressing the general case. A fifth proof manages to deal with the intermediate case where $\mathcal{P}$ has winding number 1 and consists of positively oriented moves only, and is implicitly contained in the work of Koelman [15]; this is briefly elaborated in (2.6) below.

Remark The notion of a legal loop can be generalized by dropping the condition that the endpoints should coincide. Such curves have been studied by Karpenkov in the context of lattice trigonometry, where they are called *o-broken lines* [14, Definition 3.1]. In the same trigonometric philosophy, Poonen and Rodriguez-Villegas suggested a connection between Theorem 6 and the Gauss-Bonnet theorem. See [23, Sect. 10].

(2.4) Let Δ be a lattice polygon such that $\Delta^{(1)}$ is two-dimensional. To $\Delta^{(1)}$, we can associate in a natural way two legal loops, up to shifting.

Definition (Edge-moving loop) Let $p_0, p_1, \dots, p_{n-1}$ be the vertices of $\Delta^{(1)}$, enumerated counterclockwise. Let τ_i and τ'_i be the edges adjacent to p_i , so that p_i is the top of the cone $\mathcal{H}_{\tau_i} \cap \mathcal{H}_{\tau'_i}$. Let $p_i^{(-1)}$ be the top of the cone $\mathcal{H}_{\tau_i}^{(-1)} \cap \mathcal{H}_{\tau'_i}^{(-1)}$. The *edge-moving loop* $\mathcal{P}(\Delta^{(1)})$ is defined to be $(p_0^{(-1)} - p_0, \dots, p_{n-1}^{(-1)} - p_{n-1})$. This is well-defined up to shifting.

Definition (Normal fan loop) Let $t_0, \dots, t_{n-1}$ be the primitive generators of the rays of the normal fan of $\Delta^{(1)}$, enumerated counterclockwise. Along with o , two consecutive primitive generators always span a triangle having no interior lattice points. Therefore, $\mathcal{N}(\Delta^{(1)}) := (t_0, \dots, t_{n-1})$ is a legal loop, well-defined up to shifting. It will be called the *normal fan loop* of $\Delta^{(1)}$.

Remark If, conversely, the primitive generators of the normal fan of a two-dimensional lattice polygon Γ span a legal loop, then this does not guarantee that $\Gamma = \Delta^{(1)}$ for some lattice polygon Δ . However, it does guarantee that $k\Gamma = \Delta^{(1)}$ for some lattice polygon Δ and some Minkowski multiple $k\Gamma$. In other words, there are two distinct reasons why a lattice polygon Δ can fail to be interior to a larger lattice polygon: either its normal fan is not a legal loop (i.e. the fan is not *Gorenstein* using the terminology of (4.2)), or the polygon is just too small.

The following lemma gives relationships between $\mathcal{P}(\Delta^{(1)})$ and $\mathcal{N}(\Delta^{(1)})$. We call a legal loop $\mathcal{P} = (p_0, p_1, \dots, p_{n-1})$ *convex* if every move has positive length and each p_i lies on an edge of $\text{Conv}\{p_0, p_1, \dots, p_{n-1}\}$. We call $\mathcal{P}$ *strictly convex* if moreover each p_i appears as a vertex. Note that convexity nor strict convexity are properties of the equivalence class: they are not invariant under merging and splitting. Observation (d) below is crucial and is essentially due to Haase and Schicho.

Lemma 1

(a) $\mathcal{N}(\Delta^{(1)})$ has moves of strictly positive length only.

(b) The following are equivalent.

- $\mathcal{P}(\Delta^{(1)})$ has moves of positive length only,
- $\mathcal{P}(\Delta^{(1)})$ is convex,
- $\mathcal{P}(\Delta^{(1)})$ is strictly convex,
- $\mathcal{N}(\Delta^{(1)})$ is convex.

(c) The following are equivalent.

- $\mathcal{P}(\Delta^{(1)})$ has moves of strictly positive length only,
- $\mathcal{N}(\Delta^{(1)})$ is strictly convex.

(d) $\overline{\mathcal{P}(\Delta^{(1)})}^\vee = \overline{\mathcal{N}(\Delta^{(1)})}$

Proof We only include the details for (d). Instead of $\mathcal{N}(\Delta^{(1)})$, we will consider the legal loop $\mathcal{T}(\Delta^{(1)})$ obtained by considering the consecutive counterclockwise direction vectors of the edges of $\Delta^{(1)}$. Since $\mathcal{T}(\Delta^{(1)})$ is obtained from $\mathcal{N}(\Delta^{(1)})$ by applying a 90° counterclockwise rotation, both legal loops are clearly equivalent.

If $\mathcal{P}(\Delta^{(1)})$ contains no legal moves of length 0, it follows by construction that $\overline{\mathcal{P}(\Delta^{(1)})}^\vee = \overline{\mathcal{T}(\Delta^{(1)})}$. In general, the situation is more subtle, and it is convenient to start from $\mathcal{T}(\Delta^{(1)})$ instead. The latter has no moves of length 0, so it makes sense to talk of the dual $\mathcal{T}(\Delta^{(1)})^\vee$ of $\mathcal{T}(\Delta^{(1)})$, rather than of its class.

An example where $\mathcal{P}(\Delta^{(1)})$ has several moves of length 0

After an orientation-preserving lattice transformation and a translation over an integral vector if necessary, we may assume that

$$p_0 = (a, 0), \quad p_1 = (0, 0), \quad p_2 = (p, q),$$

for integers $a > 0$ and coprime $p, q < 0$ such that $q \leq p < 0$. Because moving out the edges should result in a lattice polygon again, we must have $p = -1$. Then $p_1^{(-1)} - p_1 = (0, 1)$. On the other hand, the corresponding move of $\mathcal{T}(\Delta^{(1)})$ is from $(-1, 0)$ to $(-1, q)$. This gives a vertex $(0, -1)$ on the dual $\mathcal{T}(\Delta^{(1)})^\vee$. One concludes that $\mathcal{T}(\Delta^{(1)})$ equals $\mathcal{P}(\Delta^{(1)})$ modulo a 180° rotation (and modulo shifting). $\square$

Lemma 2 *Let Δ be a lattice polygon such that $\Delta^{(1)}$ is two-dimensional. Suppose that Δ is maximal. Let R be the number of lattice points on the boundary of Δ and let $R^{(1)}$ be the number of lattice points on the boundary of $\Delta^{(1)}$. Then $\ell(\mathcal{P}(\Delta^{(1)})) = R - R^{(1)}$.*

Proof Using a normalization as above, one easily verifies that the length of $(p_i^{(-1)} - p_i, p_{i+1}^{(-1)} - p_{i+1})$ equals the difference between the number of lattice points on the face (edge or vertex) of $\Delta = \Delta^{(1)(-1)}$ connecting $p_i^{(-1)}$ and $p_{i+1}^{(-1)}$ and the number of lattice points on the edge of $\Delta^{(1)}$ connecting p_i and p_{i+1} . $\square$

Remark As pointed out in [27], there is a natural way of associating a legal loop to any lattice polygon Δ for which $\Delta^{(1)}$ is two-dimensional, in such a way that its length still measures $R - R^{(1)}$: for each vertex $p_i^{(-1)}$ of $\Delta^{(1)(-1)}$, let a_i and b_i be the nearest-by lattice points on the adjacent edges of $\Delta^{(1)(-1)}$ that are contained in Δ (considered counterclockwise). Then in the definition of $\mathcal{P}(\Delta^{(1)})$, one should replace $p_i^{(-1)} - p_i$ by $a_i - p_i, b_i - p_i$.

(2.5) We are now ready to prove Coleman's conjecture.

Proof of Theorem 1 Using the classification given in Theorem 10 below, the statement is easily verified in case $\Delta^{(1)}$ is not two-dimensional. So suppose to the contrary that $\Delta^{(1)}$ is two-dimensional. Then a second observation is that it suffices to give a proof for the case where Δ is maximal, i.e. $\Delta = \Delta^{(1)(-1)}$. Indeed, if not, Δ is obtained from $\Delta^{(1)(-1)}$ by repeatedly clipping off a vertex. At each step, the number

of lattice points on the boundary is reduced by one, whereas the number of vertices increases by at most one. Hence the validity of Coleman's conjecture for Δ follows from its validity for $\Delta^{(1)(-1)}$.

Now let $n^{(1)}$ be the number of vertices of $\Delta^{(1)}$ and let $R^{(1)}$ be the number of lattice points on its boundary. From the definition of moving out the edges, we see that $n \leq n^{(1)}$. From Lemmata 1 and 2, it follows that

$$R - R^{(1)} = \ell(\mathcal{P}(\Delta^{(1)})) = 12 - \ell(\mathcal{N}(\Delta^{(1)})) \leq 12 - n^{(1)} \leq 12 - n.$$

The statement then follows from $R^{(1)} \leq g$ and $g \geq 2$. □

Note that the proof yields the much stronger statement that

$$R \leq R^{(1)} + 12 - n \tag{2}$$

as soon as $\Delta^{(1)}$ is two-dimensional (regardless of whether Δ is maximal or not).

(2.6) Building on work of Oda [21, Remark on p. 45], Koelman proved a statement which immediately implies Coleman's conjecture. Let Δ be a lattice polygon with two-dimensional interior $\Delta^{(1)}$. Let η be number of rays of the smooth completion of the normal fan of Δ . Let R resp. $R^{(1)}$ be the number of lattice points on the boundary of Δ resp. $\Delta^{(1)}$. Then [15, Lemma 4.5.2(2)] states

$$R^{(1)} = R + \eta - 12. \tag{3}$$

Since $\eta \geq n$, with n the number of vertices of Δ , Coleman's conjecture follows.

Equality (3) even implies the '12 theorem' for legal loops of winding number 1, all of whose segments are positively oriented (and for their duals, of course). Indeed, let $\mathcal{P} = (p_0, \dots, p_{n-1})$ be such a legal loop, then the p_i can be thought of as the generators of the rays of a fan. This fan can always be realized as the normal fan of a certain two-dimensional lattice polygon. By the legal-loop-properties of the fan, a sufficiently large Minkowski multiple Δ of this lattice polygon will be such that $\Delta^{(-1)}$ takes vertices in $\mathbb{Z}^2$. Applying (3) to $\Delta^{(-1)}$ then implies the theorem, by noting that $\eta = \ell(\mathcal{T}(\Delta))$. Digging into Oda's work, one sees that this proof is somehow related to Poonen and Rodriguez-Villegas' second proof [23, Sect. 6] and the exercises in Fulton's book [11, Sect. 2.5] to which they refer. Using work of Hille and Skarke [13], it should be possible to generalize the above to arbitrary winding numbers.

(2.7) We end this section by briefly commenting on Haase and Schicho's 'onion skin theorem' [12, Theorem 8]. Using $n \geq 3$, inequality (3) yields $R \leq R^{(1)} + 9$, which is the above-mentioned refinement of Scott's bound that Haase and Schicho obtained. In this case, one additionally checks that equality holds if and only if Δ is equivalent to $d\Sigma$ for some integer $d \geq 4$. Here $\Sigma = \text{Conv}\{(0, 0), (1, 0), (0, 1)\}$ is the standard 2-simplex. By recursively applying $R \leq R^{(1)} + 9$, while gradually 'peeling off' the lattice polygon, one obtains an inequality relating R to the genus g of Δ and to the 'number of onion skins'. This led Haase and Schicho to introducing the notion of

level. Let $n \geq 0$ be the maximal integer for which $\Delta^{(n)}$ is defined. The level of Δ is (i) equal to n if $\Delta^{(n)}$ is a point or a line segment, (ii) equal to $n + 1/3$ if $\Delta^{(n)}$ is equivalent to Σ , (iii) equal to $n + 2/3$ if $\Delta^{(n)}$ is equivalent to 2Σ , and (iv) is equal to $n + 1/2$ if $\Delta^{(n)}$ is any other two-dimensional lattice polygon of genus 0. Then the onion skin theorem reads as follows.

Theorem 7 (Haase and Schicho, 2009) *Let Δ be a convex lattice polygon of level $\ell \geq 1$ and genus g , containing R lattice points on the boundary. Then $(2\ell - 1)R \leq 2g + 9\ell^2 - 2$.*

However, although the proof is beautiful, the resulting statement is not as deep as one might hope. The reason is that applying $R \leq R^{(1)} + 9$ at each step is too rough; it would be more powerful to include the number of vertices in the argument, although we did not find an elegant way of doing so.

An alternative, more classically flavored measure for the number of onion skins is the *lattice width* of Δ , which is the minimal integer $s \geq 0$ for which there is a $\mathbb{Z}$ -affine transformation mapping Δ into the strip $\{(x, y) \in \mathbb{R}^2 \mid 0 \leq y \leq s\}$; it is denoted $\text{lw}(\Delta)$. Indeed, one can prove (see [6, Theorem 4] or [19, Theorem 13]) that for every lattice polygon Δ of genus at least 1 one has $\text{lw}(\Delta) = \text{lw}(\Delta^{(1)}) + 2$, unless Δ is equivalent to $d\Sigma$ for some integer $d \geq 3$, in which case $\text{lw}(\Delta) = \text{lw}(\Delta^{(1)}) + 3 = d$. Then by redoing the onion skin argument, using the lattice width rather than the level, one obtains a closely related statement:

Theorem 8 *Let Δ be a two-dimensional lattice polygon of genus g , containing R lattice points on its boundary. Then $(\text{lw}(\Delta) - 1) \cdot R \leq 2g + 2 \cdot (\text{lw}(\Delta)^2 - 1)$.*

Proof If Δ is a lattice polygon which has genus 0, or for which $\Delta^{(1)}$ is not two-dimensional, then the inequality can be verified by hand (using, e.g. Theorem 10 below). Therefore, suppose that the lemma holds for all lattice widths up to $k - 1$, $k \geq 3$. Let Δ be a lattice polygon with $\text{lw}(\Delta) = k$. If Δ is equivalent to $k\Sigma$, then the inequality holds by explicit verification (use $\text{lw}(\Delta) = k$, $g = (k - 1)(k - 2)/2$, $R = 3k$). If not, then the result easily follows by induction, using $R \leq R^{(1)} + 8$. $\square$

As said, this is not a deep statement. Indeed, imagine Δ being caught in a horizontal strip $\{(x, y) \in \mathbb{R}^2 \mid 0 \leq y \leq \text{lw}(\Delta)\}$. Assume for ease of exposition that both $\Delta \cap \{y = 0\}$ and $\Delta \cap \{y = \text{lw}(\Delta)\}$ are line segments of length $a \geq 3$ and $b \geq 3$, respectively. By focusing on the interior lattice points of these, one sees that $R \leq (a - 1) + (b - 1) + 2 \cdot (\text{lw}(\Delta) + 1)$. Applying Pick's theorem to the trapezoid spanned by these interior lattice points then yields

$$R \leq \frac{2}{\text{lw}(\Delta) - 1} \cdot g + 2 \cdot (\text{lw}(\Delta) + 1)$$

which is indeed a rephrasing of Theorem 8.

3 Consecutively Moving Out the Edges

(3.1) Although it is not entirely clear whom it should be attributed to, the following fact is well-known.

Theorem 9 *Let g be a positive integer. If $g \geq 1$ then there exists only a finite number of equivalence classes of lattice polygons of genus g .*

The classical argument works by noting that a lattice polygon of genus $g \geq 1$ satisfies $\text{Vol}(\Delta) \leq 2g + \frac{5}{2}$ (using Scott's bound from Theorem 4 along with Pick's theorem) and that a lattice polygon with volume V can always be caught in a lattice square of side length $4V$ —see Lagarias and Ziegler for an account that deals with arbitrary dimension [18]. This yields an algorithm for enumerating all equivalence classes of lattice polygons of genus g : consider all lattice polygons that are contained in

$$[0, 8g + 10] \times [0, 8g + 10]$$

and filter out unique representants of each conjugacy class. However, this is too slow to be of any practical use.

(3.2) We present an alternative proof of Theorem 9 that leads to a more efficient algorithm. The idea is to proceed by induction on g , based on Theorem 5. We call a lattice polygon *elliptic* if it contains a unique lattice point in its interior. A lattice polygon of genus $g \geq 2$ is called *hyperelliptic* if its interior lattice points are contained in a line.

Proof of Theorem 9 Suppose that the theorem holds for $0, 1, \dots, g - 1$, where $g \geq 1$. We partition the set of (equivalence classes of) lattice polygons of genus g as

- (i) {elliptic or hyperelliptic lattice polygons}
- (ii) $\sqcup$ {lattice polygons Δ for which $\Delta^{(1)}$ is two-dimensional of genus 0}
- (iii) $\sqcup$ {lattice polygons Δ for which $\Delta^{(1)}$ has genus ≥ 1 }

and prove the finiteness of each subset. For set (i), this follows from Theorem 10(b–c) below. For set (ii), Theorem 10(a) shows that there is only a finite number of possibilities for $\Delta^{(1)}$; for each $\Delta^{(1)}$, Theorem 5 states that $\Delta \subset \Delta^{(1)(-1)}$, hence there is only a finite number of possibilities for Δ . Finally, for set (iii), the induction hypothesis shows that there is only a finite number of possibilities for $\Delta^{(1)}$, and again Theorem 5 shows that there is only a finite number of possibilities for Δ . $\square$

Theorem 10 (Koelman, 1991)

- (a) Every two-dimensional lattice polygon of genus 0 having $R \geq 3$ lattice points on the boundary is equivalent to exactly one of the following $\lfloor R/2 \rfloor$ polygons:

for $i \in \{0, \dots, \lfloor R/2 \rfloor - 1\}$, except if $R = 6$, where in addition one has the possibility

- (b) Every lattice polygon of genus 1 is equivalent to exactly one of the following 16 polygons:

- (c) Every hyperelliptic lattice polygon of genus $g \geq 2$ is equivalent to exactly one of the following $\frac{1}{6}(g+3)(2g^2+15g+16)$ polygons:

$$g \leq i \leq 2g$$

$$0 \leq i \leq g \text{ \& } 0 \leq j \leq i$$

$$g < i \leq 2g+1 \text{ \& } 0 \leq j \leq 2g-i+1$$

$$0 \leq k \leq g+1 \text{ \& } 0 \leq i \leq g+1-k \text{ \& } 0 \leq j \leq i$$

$$0 \leq k \leq g+1 \text{ \& } g+1-k < i \leq 2g+2-2k \text{ \& } 0 \leq j \leq 2g-i-2k+2.$$

Proof A complete proof can be found in Chap. 4 of Koelman's Ph.D. thesis [15], but since there are no surprising ingredients, the proof could also be left as a patience-involving exercise. Note that (a) and (b) have been (re)discovered multiple times before and since. $\square$

For our alternative proof of Theorem 9, a strongly simplified version of Theorem 10 only involving finiteness statements would have been sufficient.

(3.3) Our proof of Theorem 9 results in the following algorithm for enumerating all lattice polygons of genus at most g up to equivalence.

INPUT: an integer $g \geq 1$.

OUTPUT: a list $[L_1, L_2, \dots, L_g]$ where each L_i is a list containing a unique representant of each equivalence class of lattice polygons of genus i .

In fact, the algorithm produces a list $\mathfrak{L} = [\mathfrak{L}_1, \mathfrak{L}_2, \dots, \mathfrak{L}_g]$, where each $\mathfrak{L}_i$ is a list

$$[\ell_{i,3}, \ell_{i,4}, \dots, \ell_{i,2i+7}],$$

and each $\ell_{i,R}$ is a list containing a unique representant of each equivalence class of lattice polygons of genus i having R lattice points on the boundary. The lists L_i are then obtained by concatenating the $\ell_{i,R}$'s for $R = 3, \dots, 2i + 7$. Remark that the number of lattice points on the boundary of a genus $i \geq 1$ lattice polygon is indeed at most $2i + 7$ by Theorem 4.

1. $\mathfrak{L} := [[], [], \dots, []]$ (g entries, indexed by $1, \dots, g$);
2. for $i = 1, \dots, g$ do;
3. $\mathfrak{L}[i] := [[], [], \dots, []]$ ($2i + 5$ entries, indexed by $3, \dots, 2i + 7$)
4. for $R = 3, \dots, 2i + 7$ do
5. $\mathfrak{L}[i][R] :=$ list of (hyper)elliptic polygons of genus i
6. with R boundary points (using Theorem 10(b–c));
7. end for;
8. for all two-dimensional Δ of genus 0 having i boundary points
9. (using Theorem 10(a)) do
10. $L :=$ [unique representants of all polygons with interior Δ];
11. for $\Gamma \in L$ do
12. $\mathfrak{L}[i][\#\partial\Gamma \cap \mathbb{Z}^2]$ cat: $= [\Gamma]$;
13. end for;
14. end for;
15. for all (j, R) such that $j + R = i$ and $3 \leq R \leq 2j + 7$ do
16. for all $\Delta \in \mathfrak{L}[j][R]$ do
17. $L :=$ [unique representants of all polygons with interior Δ];
18. for $\Gamma \in L$ do
19. $\mathfrak{L}[i][\#\partial\Gamma \cap \mathbb{Z}^2]$ cat: $= [\Gamma]$;
20. end for;
21. end for
22. end for;
23. end for;

The three disjoint for-loops correspond, respectively, to the cases (i), (ii), (iii) of our proof of Theorem 9. The operation ‘cat:=’ abbreviates ‘concatenate with’. Listing all polygons with interior Δ (see steps 10. and 17.) is done using Theorem 5: one checks whether $\Delta^{(-1)}$ is a lattice polygon. If not, then the resulting set L is empty. If yes, then L consists of all lattice polygons Γ that can be obtained from $\Delta^{(-1)}$ by taking away boundary points without affecting the interior. Each time a lattice polygon is to be added, one checks whether or not it is equivalent to a polygon that is already contained in the list. In our implementation below, checking the equivalence of two lattice polygons Γ and Γ' is done very naively: we simply try for all triples of consecutive vertices $v_1, v_2, v_3 \in \Gamma$ (ordered counterclockwise) and $v'_1, v'_2, v'_3 \in \Gamma'$ (ordered either clockwise or counterclockwise) whether there is a $\mathbb{Z}$ -affine transformation φ taking v_i to v'_i (for $i = 1, 2, 3$); if yes, it is necessarily unique. We then check whether $\varphi(\Gamma) = \Gamma'$. The algorithm can be sped up by instead keeping track of the automorphisms of Δ (i.e., $\mathbb{Z}$ -affine transformations φ for which $\varphi(\Delta) = \Delta$): indeed, if there is a $\mathbb{Z}$ -affine transformation taking Γ to Γ' , it must be an automorphism of $\Delta = \Gamma^{(1)}$. For almost all polygons, this automorphism group will consist of the identity map only, resulting in a substantial speed-up.

While much more efficient than the naive method suggested in (3.1), note that the problem is in itself exponential and that one cannot expect being able to push the computation very far: if $N(g)$ denotes the number of equivalence classes of lattice polygons of genus g , then one can show that $\log N(g)$ grows like $\sqrt[3]{g}$, although it is unknown whether $\lim_{g \rightarrow \infty} \log N(g) / \sqrt[3]{g}$ exists. See Bárány’s survey paper [2] for some discussions on this matter.

We finally remark that Koelman already briefly described and implemented a similar algorithm for enumerating lattice polygons by their total number of lattice points, rather than their genus—see [15, Sect. 4.4].

(3.4) We have implemented the above algorithm in MAGMA [5], along with several basic functions for dealing with lattice polygons, such as functions for computing the genus, the number of boundary points, the interior hull, the polygon obtained by moving out the edges, ... We have also implemented an algorithm due to Feschet [10, Sect. 3] for computing the lattice width of a lattice polygon Δ . The code can be found at <http://wis.kuleuven.be/algebra/castryck/>. The intention is to make the code cleaner and more efficient in the future. We have executed our current implementation on the input $g = 30$. This took roughly one month of computation, although it is likely that keeping track of the automorphism groups, as explained at the end of (3.3), would have shortened this span considerably. The resulting output is stored in a file of approximately 368 MB, which has also been made available for download. We include some summarizing data here—see Table 1, but of course our output can be used to answer virtually every reasonable question on lattice polygons of genus ≤ 30 . Many of these questions have been asked explicitly in the literature before, see e.g. [16, 24]. It is somewhat remarkable that the merit of exhaustive computation in tackling these questions has not been fully acknowledged thus far, with the exception of the Ph.D. thesis of Koelman and some preliminary attempts by Rabinowitz.

(3.5) We now focus on one particular problem: for every positive integer n , what is the minimal genus $g(n)$ of a lattice n -gon? Arkinstall [1], Rabinowitz [24], Simpson [26],

Table 1 A superscript ⁽¹⁾ denotes that the corresponding invariant was obtained by restricting the count to those polygons that are *interior* to another polygon (which can be easily checked using Theorem 5). Then for each integer $1 \leq g \leq 30$, the table shows the number of equivalence classes (N) of polygons of genus g , the maximal resp. minimal number of vertices ($n_{\max}$ resp. $n_{\min}$), the maximal lattice width ($lw_{\max}$) and the minimal number of lattice points on the boundary ($R_{\min}$) that are possible for that genus. Note that $n_{\min}$ and $R_{\min}$ without superscript are always equal to 3, hence not included in the table. The maximal number of points on the boundary (in both settings) is $2g + 6$ except if $g = 1$ (where it is 9) by Theorem 4, so again we did not include this

g	N	$N^{(1)}$	$n_{\max}$	$n_{\max}^{(1)}$	$n_{\min}^{(1)}$	$lw_{\max}$	$lw_{\max}^{(1)}$	$R_{\min}^{(1)}$
1	16	16	6	6	3	3	3	3
2	45	22	6	6	4	2	2	5
3	120	63	6	6	3	4	4	5
4	211	78	8	8	3	4	4	6
5	403	122	7	7	3	4	4	7
6	714	192	8	8	3	5	5	6
7	1023	239	9	9	3	4	4	7
8	1830	316	8	8	4	4	4	8
9	2700	508	8	8	3	5	5	8
10	3659	509	10	10	3	6	6	8
11	6125	700	9	9	4	5	5	8
12	8101	1044	9	9	4	6	6	8
13	11027	1113	10	10	3	6	6	9
14	17280	1429	10	10	4	6	6	9
15	21499	2052	10	10	3	7	7	9
16	28689	1962	10	10	3	6	6	9
17	43012	2651	11	11	4	6	6	9
18	52736	3543	10	10	4	7	7	10
19	68557	3638	12	12	3	8	8	9
20	97733	4594	12	11	4	7	7	9
21	117776	5996	12	12	3	8	8	10
22	152344	6364	11	11	4	8	8	10
23	209409	7922	11	11	4	8	7	10
24	248983	9693	12	12	4	8	8	10
25	319957	10208	12	12	3	8	8	10
26	420714	12727	12	12	4	8	8	11
27	497676	15431	13	12	4	8	8	9
28	641229	15918	12	12	3	9	9	10
29	813814	20354	12	12	4	8	8	11
30	957001	23874	13	12	4	9	9	11

and Olszewska [22] elaborated this for various small values of n , leaving $n = 15$ as the smallest open entry. Up to $n = 13$, their results are immediately confirmed by our computation. E.g. using Table 1, one can check that $g(11) = 17$, a case which has provoked particular interest in the past and was settled in 2006 only [22]. In fact, our output shows that there are three inequivalent 11-gons realizing the bound:

Note that the second polygon is obtained from the third by clipping off the right-most of the lower-most vertices. This is a general phenomenon: every lattice n -gon having minimal genus can be transformed to a lattice n -gon having minimal genus and all of whose boundary lattice points are vertices. This implies that the minimal genus $g(n)$ and the minimal volume $V(n)$ of a lattice n -gon are related through Pick's theorem by $V(n) = g(n) + n/2 - 1$. These are the contents of [26, Theorem 1]. In our case, it yields that every lattice 11-gon has an area of at least $43/2$. This bound is achieved by (and only by) the first two of the above polygons.

Using a refined search, we managed to settle the case $n = 15$.

Proof of Theorem 2 We first make some general observations. Let Δ be a lattice n -gon such that $\Delta^{(1)}$ is two-dimensional. Write $\Delta^{\max} = \Delta^{(1)(-1)}$, let $n^{(1)}$ resp. $n^{\max}$ be the number of vertices of $\Delta^{(1)}$ and $\Delta^{\max}$, and let R , $R^{(1)}$ resp. $R^{\max}$ be the number of lattice points on the boundaries of Δ , $\Delta^{(1)}$ and $\Delta^{\max}$. One obviously has

$$n^{\max} \leq n^{(1)}.$$

We also have

$$R^{\max} \geq n^{\max} + 2(n - n^{\max}),$$

which follows because Δ is obtained from $\Delta^{\max}$ by taking away a number of boundary points (indeed, $\Delta \subset \Delta^{\max}$ by Theorem 5), and that each introduction of a new vertex requires the existence of two lattice points on the boundary of $\Delta^{\max}$ that are not vertices. From the proof of Theorem 1 in (2.5) we see that

$$R^{\max} \leq R^{(1)} + 12 - n^{(1)}.$$

Then combining the three inequalities yields

$$R^{(1)} \geq 2n - 12,$$

while it is also clear that

$$n^{(1)} \geq \lceil n/2 \rceil.$$

Now let Δ be a lattice 15-gon of genus $g = g(15)$. Note that by Theorem 10(c), Δ is non-hyperelliptic, thus the above applies. In particular, we have

$$n^{(1)} \geq 8 \quad \text{and} \quad R^{(1)} \geq 18. \quad (4)$$

Now since

Table 2 Known values for $g(n)$; values marked with an asterisk are new contributions

n	3	4	5	6	7	8	9	10	11	12
$g(n)$	0	0	1	1	4	4	7	10	17	19
n	13	14	15	16	17	18	19	20	21	22
$g(n)$	27	34	45*	52	[66, 72*]	79	[96, 105*]	112	[133, 154]	154

is a lattice 15-gon of genus 45, it follows that $g \leq 45$. We also know that $g \geq 43$ by [26, Corollary 11(c)]. So suppose that $g \in \{43, 44\}$. Then with $g^{(1)}$ the genus of $\Delta^{(1)}$ we have $g^{(1)} \geq 1$ because of Theorem 10(a) (note that $n^{(1)} \geq 8$), and $g^{(1)} \leq 26$ because of (4). In particular, $\Delta^{(1)}$ must be contained in our list produced in (3.4).

The remainder of the proof is computational. Out of our list, we have selected those lattice polygons Γ that are of the form $\Delta^{(1)}$, i.e. for which $\Gamma^{(-1)}$ takes vertices in the lattice (following Theorem 5), that have at least eight vertices, that contain at least 18 lattice points on the boundary, that have genus at most 26, and for which the sum of the latter two invariants is contained in $\{43, 44\}$. This resulted in 1929 polygons. For each such polygon Γ , we enumerated all polygons Δ for which $\Delta^{(1)} = \Gamma$, in a similar way as described in (3.3), and checked whether any of these has 15 vertices. In each case, the answer was no. $\square$

Finally, the picture below proves that $g(17) \leq 72$ (Simpson's previous upper bound was 79) and that $g(19) \leq 105$ (versus 112). Our guess is that these bounds are not yet optimal.

A summarizing update of the currently known values of $g(n)$ can be found in Table 2. We conclude by remarking that the asymptotic behavior of $g(n)$ is well-understood. It is known that for all $n \geq 3$

$$\frac{1}{8\pi^2} < \frac{V(n)}{n^3} \leq \frac{1}{54}(1 + o(1))$$

and that

$$\lim_{n \rightarrow \infty} \frac{g(n)}{n^3} = \lim_{n \rightarrow \infty} \frac{V(n)}{n^3}$$

exists and lies close to (but is not equal to) $1/54$. See [3] and the references therein.

(3.6) A notion dual to the level, as introduced by Haase and Schicho and reviewed in (2.7), is the *lifespan* of a two-dimensional lattice polygon $\Delta \subset \mathbb{R}^2$, which is defined to be the maximal $k \in \mathbb{Z}_{\geq 0}$ for which $\Delta^{(-k)}$ is well-defined and takes vertices in $\mathbb{Z}^2$, provided such a k exists. If no such k exists, the lifespan is said to be *infinite*. Theorem 3 claims that an infinite lifespan can only occur if the number of vertices n is at most 9.

Proof of Theorem 3 Suppose that Δ has infinite lifespan. It suffices to prove that all moves of the edge-moving loop $\mathcal{P}(\Delta)$ have positive length. Indeed, by Lemma 1 this implies that the normal fan loop $\mathcal{N}(\Delta)$ is convex. Then the convex hull of its primitive generators (which are in 1-to-1 correspondence with the edges of Δ) must be contained in the list of Theorem 10(b). In particular, the maximal number of primitive generators is 9, hence so is the maximal number of edges (equalling the number of vertices).

So suppose by contradiction that there are two consecutive vertices p_i and p_{i+1} of Δ such that the move $(p_i^{(-1)} - p_i, p_{i+1}^{(-1)} - p_{i+1})$ has negative length. As explained in Lemma 2, this means that the edge of $\Delta^{(-1)}$ connecting to $p_i^{(-1)}$ and $p_{i+1}^{(-1)}$ must have become shorter, unless it has even disappeared, i.e. $p_{i+1}^{(-1)} = p_i^{(-1)}$. If no edge disappears, then it is easy to see that $\mathcal{P}(\Delta) = \mathcal{P}(\Delta^{(-1)})$. By repeating the argument, one eventually must have $p_{i+1}^{(-k)} = p_i^{(-k)}$ for some $k \in \mathbb{Z}_{\geq 1}$ (where $(-k)$ abbreviates $(-1)(-1)\dots(-1)$). We claim that $\Delta^{(-k-1)}$ takes at least one vertex outside $\mathbb{Z}^2$.

To see this, choose i such that $p_{i+1}^{(-k)} = p_i^{(-k)}$, but $p_{i+2}^{(-k)} \neq p_{i+1}^{(-k)}$. Modulo a $\mathbb{Z}$ -affine transformation we may assume that $p_{i+1}^{(-k+1)} = (0, 0)$, that $p_i^{(-k+1)} = (a, 0)$ for some integer $a > 0$, and that $p_{i+1}^{(-k)} = p_i^{(-k)} = (0, 1)$.

Let τ_{left} be the edge of $\Delta^{(-k)}$ that is left adjacent to $p_{i+1}^{(-k)} = p_i^{(-k)}$, and let $L_{\tau_{\text{left}}}$ be its supporting line. Move it out (in the sense of (2.2)) to obtain a line $L_{\tau_{\text{left}}}^{(-1)}$. Because of our choice of i , $L_{\tau_{\text{left}}}^{(-1)}$ contains the point $(0, 2)$. Now similarly define $L_{\tau_{\text{right}}}$ and $L_{\tau_{\text{right}}}^{(-1)}$. The slope σ of $L_{\tau_{\text{right}}}$ (hence of $L_{\tau_{\text{right}}}^{(-1)}$) satisfies $0 > \sigma > -1$. Since $(1, 1)$ is a lattice point that is not contained in $\Delta^{(-k)}$, the lines $L_{\tau_{\text{left}}}^{(-1)}$ and $L_{\tau_{\text{right}}}^{(-1)}$ must intersect in a point that lies strictly between $y = 1$ and $y = 2$. $\square$

Note that the above proof actually gives a criterion for a two-dimensional lattice polygon Δ to have infinite lifespan. This will be the case *if and only if* $\Delta^{(-1)}$ is a lattice polygon and $\mathcal{N}(\Delta)$ is convex. Examples of n -gons (for $n = 3, \dots, 9$) having

infinite lifespan are given in the picture below.

4 Concluding Comments

(4.1) Coleman’s conjecture (or at least parts of the proof given in (2.5)) can be extended to certain *non-convex lattice polygons*. By a non-convex lattice polygon we mean a closed region in $\mathbb{R}^2$ that can be bounded by a closed non-self-intersecting curve that is piece-wise linear, with the endpoints of the linear parts contained in $\mathbb{Z}^2$. For such a non-convex lattice polygon Δ , it makes sense to define $\Delta^{(-1)}$ by extending the corresponding notion of (2.2). Suppose it takes vertices in $\mathbb{Z}^2$ and let R be the number of lattice points on its boundary. Then with $R^{(1)}$ the number of lattice points on the boundary of Δ , we will again have $R = R^{(1)} + 12 - \ell(\mathcal{T}(\Delta))$, with $\mathcal{T}(\Delta)$ the legal loop spanned by the direction vectors of the piece-wise linear boundary components of Δ .

(4.2) Much (if not all) of the foregoing can be related to toric geometry. It lies beyond the scope of this article to go into much detail here, but we briefly mention a few facts. We fully rely on the according references for the background. For a two-dimensional lattice polygon Δ , we denote the according toric surface over $\mathbb{C}$ by $X(\Delta)$, which we assume to be naturally embedded in $\mathbb{P}_{\mathbb{C}}^{\#(\Delta \cap \mathbb{Z}^2) - 1}$.

- (i) If the primitive generators of the normal fan of Δ span a legal loop (see the corresponding remark in (2.4)), then by definition this is a *Gorenstein fan*, which is equivalent to saying that $X(\Delta)$ has only Gorenstein singularities [4, Proposition 2.7]. In particular, if Δ is the interior hull of another lattice polygon, then $X(\Delta)$ has only Gorenstein singularities. The converse is not true.
- (ii) If this legal loop is moreover convex, then $X(\Delta)$ is *weak Fano*, meaning that the anticanonical bundle $-K_{X(\Delta)}$ is nef and big. If it is strictly convex, then $X(\Delta)$ is *Fano*, meaning that $-K_{X(\Delta)}$ is ample. See [20, Sect. 2.3]. In particular, if a lattice polygon Δ has infinite lifespan, then $X(\Delta)$ is Gorenstein and weak Fano. In this case, the converse holds as well.
- (iii) Since convex legal loops (of winding number 1) have length at most 9, the above implies that in the weak Fano case, Δ can have no more than 9 edges and vertices. This also follows from a well-known degree bound for weak Fano surfaces X (namely, $(-K_X)^2 \leq 9$).
- (iv) If $\Delta^{(1)}$ is well-defined and two-dimensional, then $X(\Delta^{(1)})$ is the so-called *adjoint* of $X(\Delta)$. That is, $X(\Delta^{(1)})$ is obtained from $X(\Delta)$ by taking its image

under the map corresponding to $\mathcal{O}_{X(\Delta)}(1) + K_{X(\Delta)}$. See [9, 12] for more details.

- (v) Conversely, if Δ is two-dimensional and $\Delta^{(-1)}$ has the same number of edges as Δ , then $X(\Delta^{(-1)}) \cong X(\Delta)$. The former is then embedded by the ample line bundle $\mathcal{O}_{X(\Delta)}(1) - K_{X(\Delta)}$. Similarly, for $k \geq 0$, if $\Delta^{(-k)}$ has the same number of edges as Δ , then $X(\Delta^{(-k)})$ corresponds to the ample line bundle $\mathcal{O}_{X(\Delta)}(1) - kK_{X(\Delta)}$. If this works for arbitrary k , $-K_{X(\Delta)}$ must be nef (and big, which is automatic), i.e. $X(\Delta)$ is weak Fano. Along with (iii), this gives some geometric insight in Theorem 3.
- (vi) If $\Delta^{(1)}$ is well-defined and two-dimensional, then the dimension of the automorphism group $\text{Aut}(X(\Delta))$ is determined by the number of lattice points that lie in the interior of a positively oriented move of $\mathcal{P}(\Delta^{(1)})$. See [7, Lemma 10.5] and [15, phrase following (2.99)]. This was used in [7] to determine the dimension of the moduli space of generic hyperplane sections of $X(\Delta)$.
- (vii) The genus of a two-dimensional lattice polygon Δ is equal to the genus of a generic hyperplane section of $X(\Delta)$. Such a generic hyperplane section will be (hyper)elliptic if and only if Δ is (hyper)elliptic. See [15, Sect. 3.2] or [7, Lemma 5.1].

Acknowledgements I am grateful to Sofie Verbeke, the supervision of whom's master thesis [27] I benefited from, and to Hendrik Hubrechts and Jan Schepers for some helpful discussions. I would also like to thank the anonymous referees for their helpful comments and thorough reading. Finally, I thank F.W.O. Vlaanderen for its financial support, and the Massachusetts Institute of Technology for its hospitality.

References

1. Arkinstall, J.: Minimal requirements for Minkowski's theorem in the plane I. *Bull. Aust. Math. Soc.* **22**, 259–274 (1980)
2. Bárány, I.: Extremal problems for convex lattice polytopes: a survey. In: *Surveys on Discrete and Computational Geometry. Contemporary Mathematics*, vol. 453, pp. 87–103. AMS, Providence (2008)
3. Bárány, I., Tokushige, N.: The minimum area of convex lattice n -gons. *Combinatorica* **24**(2), 171–185 (2004)
4. Batyrev, V., Borisov, L.: Dual cones and mirror symmetry for generalized Calabi–Yau manifolds. *Mirror symmetry II. AMS/IP Stud. Adv. Math.* **1**, 71–86 (1997)
5. Bosma, W., Cannon, J., Playoust, C.: The Magma algebra system I: The user language. *J. Symb. Comput.* **24**(3–4), 235–265 (1997)
6. Castryck, W., Cools, F.: Lattice polygons and curve gonality. *J. Algebr. Comb.* doi:[10.1007/s10801-011-0304-6](https://doi.org/10.1007/s10801-011-0304-6)
7. Castryck, W., Voight, J.: On nondegeneracy of curves. *Algebra Number Theory* **3**(3), 255–281 (2009)
8. Coleman, D.: Stretch: a geoboard game. *Math. Mag.* **51**, 49–54 (1978)
9. Di Rocco, S., Haase, C., Nill, B., Paffenholz, A.: Polyhedral adjunction theory. Preprint
10. Feschet, F.: The exact lattice width of planar sets and minimal arithmetical thickness. In: *Combinatorial Image Analysis. Lecture Notes in Computer Science*, vol. 4040, pp. 25–33 (2006)
11. Fulton, W.: *Introduction to Toric Varieties*. *Annals of Mathematics Studies*, vol. 131. Princeton University Press, Princeton (1993)
12. Haase, C., Schicho, J.: Lattice polygons and the number $2i + 7$. *Am. Math. Mon.* **116**(2), 151–165 (2009)
13. Hille, L., Skarke, H.: Reflexive polytopes in dimension 2 and certain relations in $\text{SL}_2(\mathbb{Z})$. *J. Algebra Appl.* **1**(2), 159–173 (2002)
14. Karpenkov, O.: Elementary notions of lattice trigonometry. *Math. Scand.* **102**(2), 161–205 (2008)

15. Koelman, R.: The number of moduli of families of curves on toric surfaces. Ph.D. thesis, Katholieke Universiteit Nijmegen (1991)
16. Kołodziejczyk, K., Olszewska, D.: On some conjectures by Rabinowitz. *Ars Comb.* **79**, 171–188 (2006)
17. Kołodziejczyk, K., Olszewska, D.: A proof of Coleman’s conjecture. *Discrete Math.* **307**, 1865–1872 (2007)
18. Lagarias, J., Ziegler, G.: Bounds for lattice polytopes containing a fixed number of interior points in a sublattice. *Can. J. Math.* **43**(5), 1022–1035 (1991)
19. Lubbes, N., Schicho, J.: Lattice polygons and families of curves on rational surfaces. *J. Algebr. Comb.* **34**(2), 213–236 (2011)
20. Nill, B.: Gorenstein toric Fano varieties. Ph.D. thesis, Universität Tübingen (2005)
21. Oda, T.: *Convex Bodies and Algebraic Geometry*. Springer, Berlin (1988)
22. Olszewska, D.: On the first unknown value of the function $g(v)$. *Electron. Notes Discrete Math.* **24**, 181–185 (2006)
23. Poonen, B., Rodriguez-Villegas, F.: Lattice polygons and the number 12. *Am. Math. Mon.* **107**, 238–250 (2000)
24. Rabinowitz, S.: On the number of lattice points inside a convex lattice n -gon. *Congr. Numer.* **73**, 99–124 (1990)
25. Scott, P.: On convex lattice polygons. *Bull. Aust. Math. Soc.* **15**, 395–399 (1976)
26. Simpson, R.: Convex lattice polygons of minimum area. *Bull. Aust. Math. Soc.* **42**, 353–367 (1990)
27. Verbeke, S.: Roosterveelhoeken en het getal 12 (Lattice polygons and the number 12, in Dutch). Master thesis, K.U. Leuven (2010)