Kobayashi length bounds on bordered surfaces and generalized integral points on abelian varieties

Paolo Dolce

Abstract

Let B be a compact Riemann surface and $B_0 \subseteq B$ a hyperbolic surface obtained by removing finitely many disjoint closed disks. Fix a nontrivial loop α in B_0. For $s \geq 0$, let $L(\alpha, s)$ denote the supremum, over all finite subsets $S \subset B_0$ with $\#S \leq s$, of the minimal Kobayashi length of a loop in $B_0 \smallsetminus S$ that is freely homotopic to α in B_0. Phung in [22] proved that $L(\alpha, s)$ grows at most linearly and at least as $\sqrt{s}/\log s$. We sharpen the upper bound to $O(\sqrt{s+1})$, which determines $\lim_{s\to\infty} \log L(\alpha, s)/\log s = 1/2$, answering [22, Question 1.4]. As an application, we improve Phung's counting bound for generalized integral points on abelian varieties over complex function fields. For an abelian variety of dimension n over $\mathbb{C}(B)$ satisfying his geometric hypotheses, Phung proved a bound of order $(s+1)^{2nk}$ modulo the constant trace, where k is the minimal number of generators of $\pi_1(B_0)$. We explain his method in detail and insert our stronger length estimate to obtain $(s+1)^{nk}$, halving the exponent. Counting in the Mordell–Weil lattice of rank r gives the more precise bound $O((s+1)^{r/2})$ and a linear bound in $s+1$ for the canonical heights, with the family and B_0 fixed. For normal-crossings D with Cartier model closure, this proves the fixed-model geometric Lang–Vojta conjecture in the stated abelian setting, without an exceptional set.

Contents

Contents **1**

1 Introduction **1**
1.1 Kobayashi lengths on punctured bordered surfaces 2
1.2 The Lang–Vojta conjecture and generalized integral points 3
1.3 Applications to counting and heights 4

2 Kobayashi length estimates **7**
2.1 Kobayashi pseudo-metric 7
2.2 Growth of loop lengths in terms of the cusps 8
2.3 Reduction to the simple case 14

3 Bounding generalized integral points **15**
3.1 Phung's bound 15
3.2 Improved bound 21
3.3 Counting in the Mordell–Weil lattice and bounding heights 22

References **24**

1 Introduction

The main result of this paper is an upper bound for the Kobayashi length of representatives of a fixed homotopy class on a bordered Riemann surface with finitely many punctures. The bound is uniform over all configurations of the punctures and, together with Phung's lower bound, answers his question about the polynomial growth exponent. We first present this analytic result and its proof strategy, then explain its application to generalized integral points through Phung's implementation of the Parshin cocycle method.

1.1 Kobayashi lengths on punctured bordered surfaces

Let B be a compact connected Riemann surface and let $\overline{\Delta}_1, \dots, \overline{\Delta}_t$ be finitely many pairwise disjoint closed topological disks in B, allowing $t = 0$. Put

$$B_0 = B \setminus \bigsqcup_{i=1}^{t} \overline{\Delta}_i$$

and assume that B_0 is hyperbolic. The disks and the surface are fixed, while the finite puncture set will vary. We specify the distinction between the removed disks and the metric balls used in the estimates.

Fix a smooth Riemannian metric ρ on B, with length ℓ_ρ and distance d_ρ. We use two different notions of disk and keep their notation and roles separate.

- A *closed (topological) disk* is an embedded subset $\overline{\Delta} \subset B$ homeomorphic to the closed unit disk, with piecewise-C^1 boundary. These are the disks $\overline{\Delta}_1, \dots, \overline{\Delta}_t$ removed from B to form B_0. They are not required to be metric balls; this flexibility is needed in §3, where disks are enlarged by collars and by joining auxiliary disks with thin tubes.
- A *metric ball* is defined by the fixed Riemannian distance:

  $$\Delta_\rho(x, r) = \{b \in B : d_\rho(b, x) < r\}, \qquad \overline{\Delta}_\rho(x, r) = \{b \in B : d_\rho(b, x) \le r\}.$$

  These balls are used in the local metric and area estimates of §2.2. For r sufficiently small (in particular, below the injectivity radius at x), the closed metric ball is a closed topological disk with smooth boundary.

Thus $\overline{\Delta}_i$ denotes a removed topological disk, while the subscript ρ and the arguments (x, r) identify a metric ball. When $t \ge 1$, B_0 is the interior of a compact bordered surface and is itself an open Riemann surface without boundary. When $t = 0$, it is the compact surface B.

We study the Kobayashi length of loops on punctured bordered surfaces. For a finite $S \subset B_0$, let ℓ_S be the length induced by the Kobayashi metric of $B_0 \setminus S$. The inclusion into B_0 decreases Kobayashi lengths, so a path avoiding S can only become longer when measured in the punctured surface. Geometrically, the punctures create new cusps. It is natural to ask how short a representative of a fixed homotopy class can remain in the worst configuration of at most s punctures. For a nontrivial loop α in B_0, set

$$L(\alpha, s) = \sup_{\substack{S \subset B_0 \\ \#S \le s}} \inf\big\{\ell_S(\gamma) \colon \gamma \text{ is a loop in } B_0 \setminus S \text{ freely homotopic to } \alpha \text{ in } B_0\big\}.$$

Theorem A. *For every such α there is a constant $C > 0$, depending only on α, B_0 and the fixed background metric, such that*

$$L(\alpha, s) \le C\sqrt{s+1} \qquad (s \in \mathbb{Z}_{\ge 0}). \tag{1}$$

Phung proved a linear upper bound and a lower bound $c\sqrt{s}/\log(s+2)$ [22, Theorems B and C]. Together with Theorem A, the latter implies

$$\lim_{s \to \infty} \frac{\log L(\alpha, s)}{\log s} = \frac{1}{2}, \tag{2}$$

answering [22, Question 1.4].

The main idea of the proof of Theorem A is to average lengths over a family of nearby loops. Suppose first that α has a smooth simple representative. Moving this loop slightly in the normal direction gives a family of parallel loops filling a thin tubular neighborhood. Phung's local estimates control how much puncturing at S increases the Kobayashi metric. Although this increase can be arbitrarily large near the punctures, the region where it exceeds a large value v has area at most Cs/v^2, where $s = \#S$. Its area is also bounded by that of the fixed neighborhood. Integrating the minimum of these two bounds gives an average length of order $\sqrt{s+1}$. Since only finitely many parallel loops meet S, one avoids the punctures and has the required length, without changing its free homotopy class in B_0.

For a general class, we write α as a word in simple generators of $\pi_1(B_0, b_0)$. We arrange their representatives to have the same tangent direction at b_0, so that moving them normally gives loops with a common moving base point. Averaging the sum of their lengths makes all generators short simultaneously, and concatenation gives the bound for α. The common motion of the base point also supplies the single conjugation needed in Phung's counting argument. The compact case is handled by removing an auxiliary disk.

We now recall the geometric Lang–Vojta setting and Phung's generalized integral points, before explaining how these length bounds enter his counting method and yield a linear height bound.

1.2 The Lang–Vojta conjecture and generalized integral points

Let k be an algebraically closed field of characteristic zero and let B be a smooth projective connected curve over k, with function field $K = k(B)$. Let X be a smooth integral projective K-variety and let D be a reduced effective divisor on X. A model $(\mathcal{X}, \mathcal{D})$ consists of a normal k-variety with a proper flat morphism $f\colon \mathcal{X} \to B$, an identification $\mathcal{X}_K \simeq X$, and the reduced horizontal closure $\mathcal{D}$ of D. Such models exist.

We distinguish the divisor on the generic fibre from its closure on the total space. Since X is smooth, every Weil divisor on X, in particular D, is Cartier. The total space $\mathcal{X}$, however, is only assumed normal and may be singular, so its Weil divisor $\mathcal{D}$ need not be Cartier. Whenever we use $\deg_B \sigma_P^*\mathcal{D}$, we assume that $\mathcal{D}$ is Cartier. This is automatic on a regular model, which can be chosen projective in characteristic zero. This assumption concerns the total-space divisor and is unnecessary for the incidence conditions and counting statements below.

Properness and the valuative criterion identify $X(K)$ with the sections of f; write σ_P for the section associated to P. Put

$$\mathrm{Sing}(f) = \{b \in B : \mathcal{X}_b \text{ is singular}\}.$$

For a finite set $S \subset B$, the set of $(S, \mathcal{D})$-integral points is

$$\mathcal{X}(\mathcal{O}_{S,\mathcal{D}}) = \{P \in X(K) : f(\sigma_P(B) \cap \mathcal{D}) \subseteq S\}.$$

In other words, these are precisely the sections whose image does not intersect $\mathcal{D}$ "above" $B \smallsetminus S$. In particular, an integral section is not contained in $\mathcal{D}$. If $\mathcal{D}$ is Cartier, its pullback along such a section is an effective divisor on B, and we have the usual notion of height of a point:

$$h_{\mathcal{D}}(P) := \deg_B \sigma_P^*\mathcal{O}_{\mathcal{X}}(\mathcal{D}) = \deg_B \sigma_P^*\mathcal{D}. \tag{3}$$

For a smooth projective pair (X, D) with normal-crossings boundary divisor D, *log general type* means that $\omega_X(D)$ is big.

Conjecture 1.1 (Geometric Lang–Vojta, fixed-model boundary-degree form)**.** Suppose that the pair (X, D) is of log general type. Fix a model $(\mathcal{X}, \mathcal{D})$ with $\mathcal{D}$ Cartier. There exist a proper closed subset $Z \subset X \setminus D$ and a constant $m > 0$, depending on the fixed pair, model and base curve B, such that for every finite $S \subset B$ and every $P \in \mathcal{X}(\mathcal{O}_{S,\mathcal{D}})$ with $P \notin Z$,

$$\deg_B \sigma_P^*\mathcal{D} \le m \max\{1, 2g(B) - 2 + \#S\}. \tag{4}$$

Roughly speaking, every integral point either lies in a fixed proper exceptional subset or has intersection degree bounded linearly in $\#S$, with the contribution of the base expressed by $2g(B) - 2$. The constants in this formulation are independent of S and P, while the model and B remain fixed. There are stronger uniform versions of the geometric Lang–Vojta conjecture, requiring independence of the base curve under the specified geometric data and base changes. We use only the fixed-model formulation above.

Earlier results include Hindry–Silverman's height bounds for elliptic curves [16] and Buium's *abc* theorem for abelian varieties with zero trace [2]. Here an *abc*-type bound controls heights linearly by the number of distinct boundary intersections, up to an additive constant. For constant abelian varieties and constant divisors, Buium [3] and Noguchi–Winkelmann [19] established bounds for intersection multiplicities under their respective hypotheses; summing these bounds gives an intersection-degree bound linear in the number of distinct intersection points.

Corvaja–Zannier proved cases of the geometric Lang–Vojta conjecture in the split setting, for certain plane complements and ramified covers of $\mathbb{G}_m^2$ [7, 9]. Their related *abcd* theorem gives lower bounds for the number of distinct zeros of $1+u+v$ in terms of the degrees of the rational functions u, v, when these have few distinct zeros and poles [8]. Nonsplit extensions were obtained by Turchet [25] and Capuano–Turchet [4]; see also Guo–Nguyen–Sun–Wang [15] for higher-dimensional toric results. The precise assumptions and the available uniformity differ among these results; for background, see [22, §1] and [4, §1.1]. The results below concern a fixed complex function field and a fixed family. They do not require the boundary to have normal crossings; their geometric hypothesis on the boundary will be stated separately.

For the application we now take $k = \mathbb{C}$ and use the surface notation introduced in §1.1. The notion of $(S, \mathcal{D})$-integral point can be generalized by allowing S to vary while keeping only its cardinality $s = \#S$ bounded, and by requiring the intersection condition to hold only on the complement of finitely many disjoint closed disks. This notion of generalized integral points was introduced in [22] and we recall it below.

Definition 1.2. Choose disjoint closed topological disks $\overline{\Delta}_1, \dots, \overline{\Delta}_t$ whose union contains $\mathrm{Sing}(f)$, with distinct points of $\mathrm{Sing}(f)$ in distinct disks, and put

$$B_0 = B \setminus \bigsqcup_{i=1}^{t} \overline{\Delta}_i.$$

For $s \in \mathbb{Z}_{\geq 0}$, define

$$I(s, B_0) = \{P \in X(K) : \#f(\sigma_P(B_0) \cap \mathcal{D}) \leq s\}.$$

These are the generalized (s, B_0)-integral points of $(\mathcal{X}, \mathcal{D})$.

The definition includes all $(S, \mathcal{D})$-integral points with $\#S \leq s$:

$$\bigcup_{\#S \leq s} \mathcal{X}(\mathcal{O}_{S,\mathcal{D}}) \subseteq I(s, B_0).$$

Thus $I(s, B_0)$ is potentially much larger than the integral points for any single fixed S: it allows the intersection points in B_0 to move, and it puts no restriction at all on intersections above the removed disks. Understanding its growth as a function of s, with B_0 fixed and with the trace quotient taken in the abelian case, is the arithmetic application of the loop-length estimates in this paper.

1.3 Applications to counting and heights

Building on Parshin's foundational work [21], Phung in [22] proved an important quantitative bound for generalized integral points on abelian varieties modulo the constant trace. We work in his setting, making the geometric interpretation of the divisor hypothesis and the abelian-scheme structure of the model explicit.

(P) Let A be an abelian variety of dimension n over $K = \mathbb{C}(B)$, let $D \subset A$ be a nonzero reduced effective divisor, and fix a model $(\mathcal{A}, \mathcal{D})$ with morphism $f\colon \mathcal{A} \to B$. Let σ_O be the section of the identity of A. We assume:

(i) The disks defining B_0 contain $\mathrm{Sing}(f)$, with distinct singular-fibre points in distinct disks; over $B \setminus \mathrm{Sing}(f)$ the model is an abelian scheme with zero section σ_O.

(ii) B_0 is hyperbolic. This is automatic when $t \geq 1$.

(iii) $D_{\overline{K}}$ contains no translate of a positive-dimensional abelian subvariety of $A_{\overline{K}}$.

The abelian-scheme structure in (i) is part of the model data used below. Such models can be obtained by compactifying the abelian scheme over the good-reduction locus.

The ampleness assumption appearing in [22] is automatic under (iii). Indeed, the identity component of the translation stabilizer of $D_{\overline{K}}$ would otherwise give positive-dimensional cosets contained in $D_{\overline{K}}$. Hence that stabilizer is finite and D is ample, by the stabilizer criterion [6, Proposition 3.4.8].

Write

$$T_A = \mathrm{im}\big(\mathrm{Tr}_{K/\mathbb{C}}(A)(\mathbb{C}) \longrightarrow A(K)\big), \qquad M = A(K)/T_A, \qquad r = \mathrm{rk}\, M.$$

The group M is finitely generated by Lang–Néron; see [5]. We write $I(s, B_0)/T_A$ for the image of $I(s, B_0)$ in M; no invariance of $I(s, B_0)$ under trace translations is asserted. Let k denote the minimal number of generators of $\pi_1(B_0, b_0)$. For $t \geq 1$, this group is free and $k = 2g(B) + t - 1$; for $t = 0$, $k = 2g(B)$.

Theorem 1.3 (Phung [22, Theorem A])**.** *In setting (P), there is a constant $m > 0$, depending on the fixed model, divisor and B_0, such that*

$$\#(I(s, B_0)/T_A) \leq m(s+1)^{2nk} \qquad (s \in \mathbb{Z}_{\geq 0}). \tag{5}$$

Note that the bound behaves coherently: adding further removed disks does not decrease k and shrinks B_0, so the set of generalized integral points grows as well.

The proof of Theorem 1.3 rests on Phung's implementation of the Parshin cocycle argument. Let $G = \pi_1(B_0, b_0)$ and fix generators $\alpha_1, \dots, \alpha_k$. By Ehresmann's theorem the family over B_0 gives a short exact sequence

$$1 \longrightarrow \Gamma \longrightarrow \pi_1(\mathcal{A}_{B_0}, w_0) \longrightarrow G \longrightarrow 1, \qquad \Gamma = H_1(\mathcal{A}_{b_0}, \mathbb{Z}) \simeq \mathbb{Z}^{2n}.$$

Each section σ_P induces a splitting. Its difference from the splitting defined by σ_O is a 1-cocycle $c_P\colon G \to \Gamma$. The cohomology class determines the image of P modulo the trace up to finitely many possibilities. Thus it suffices to count the possible cocycle classes.

Phung does this by choosing loops representing the generators and estimating the corresponding displacements in the universal cover of $\mathcal{A}_{B_0}$. After removing the bad fibres of $\mathcal{D}$, a Brody–Green comparison gives

$$\ell_h(\sigma_P(\gamma_j)) \le c^{-1}\ell_S(\gamma_j),$$

where h is a fixed Hermitian metric and S is the set of intersections of the section with $\mathcal{D}$ inside B_0. Phung's linear bound $\ell_S(\gamma_j) = O(s+1)$ yields a displacement bound $H(s) = O(s+1)$. A lattice ball in $\Gamma \simeq \mathbb{Z}^{2n}$ contains $O((H(s)+1)^{2n})$ points. There are k generators, which explains the exponent $2nk$ in (5).

This counting method, including the passage from short loops to bounded cocycles, is Phung's. We recall it in detail in §3, with the metric, base-point and monodromy verifications needed for the argument. Our improvement enters through a sharper estimate for the loops in the base. Substituting $H(s) = O(\sqrt{s+1})$ in the same lattice count gives the following result.

Theorem B. *In setting (P), there is a constant $m > 0$, depending on the fixed model, divisor and B_0, such that*

$$\#(I(s,B_0)/T_A) \le m(s+1)^{nk} \qquad (s \in \mathbb{Z}_{\ge 0}). \tag{6}$$

More precisely, after changing m if necessary,

$$\#(I(s,B_0)/T_A) \le m(s+1)^{r/2}. \tag{7}$$

The passage from the exponent nk to $r/2$ uses the same radius $O(\sqrt{s+1})$ in a smaller lattice. The first count treats the k cocycle values as a tuple in Γ^k, a lattice of rank $2nk$. For the refinement, the Parshin–Phung homomorphism from $M = A(K)/T_A$ to the cohomology of the fibre lattice has finite kernel, so it becomes injective after tensoring with $\mathbb{R}$. The bounded cocycle representatives therefore bound the coordinates of P in $M \otimes \mathbb{R}$ by $O(\sqrt{s+1})$, by equivalence of norms on fixed finite-dimensional spaces. Counting in the rank-r lattice M/M_{tors} now gives $O((s+1)^{r/2})$; torsion contributes only a fixed finite factor. In particular, $r \le 2nk$. The details are given in §3.3.

The same method gives a linear height estimate, establishing the fixed-model geometric Lang–Vojta conjecture in this setting.

Corollary C. *In setting (P), fix a symmetric ample line bundle L on A and its canonical height $\widehat{h}_L$. There is a constant $C > 0$ such that*

$$\widehat{h}_L(P) \le C(s+1) \qquad (P \in I(s,B_0)).$$

If $\mathcal{D}$ is Cartier, there is also a constant $C' > 0$ such that

$$\deg_B \sigma_P^*\mathcal{D} \le C'(s+1) \qquad (P \in I(s,B_0)).$$

All constants refer to the fixed family, divisor, B_0 and chosen height data; they are independent of s and P.

Phung's finiteness theorem also has a height-theoretic consequence. In the ample-divisor setting, [22, Corollary 1.3] gives a finite bound, for each fixed s, for the heights of generalized integral points. This relates his result to geometric Lang–Vojta in two ways. The intersection condition is weaker, because it is imposed only over B_0, whereas the height conclusion supplies a bound $M(s)$ without a specified linear rate as $s \to \infty$. The issue is therefore the growth of $M(s)$, not merely its dependence on s.

Corollary C supplies that linear rate, with the family and base fixed. Its proof uses a norm bound for the image of the Mordell–Weil group in the cohomology of the fibre lattice; the cardinality estimate alone would not imply it.

More precisely, Corollary C proves Conjecture 1.1 for abelian varieties over $\mathbb{C}(B)$ in setting (P), when D has normal crossings and $\mathcal{D}$ is Cartier, with $Z = \varnothing$. Indeed, fix one admissible B_0 once and for all. For any finite $S \subset B$, every $(S,\mathcal{D})$-integral point belongs to $I(s,B_0)$ with $s = \#S$. Since

$$s+1 \le 4\max\{1, 2g(B) - 2 + s\} \qquad (g(B), s \in \mathbb{Z}_{\ge 0}),$$

the corollary gives

$$\deg_B \sigma_P^*\mathcal{D} \le 4C' \max\{1, 2g(B) - 2 + \#S\},$$

which is the conjectured inequality with $m = 4C'$, independent of S and P. For the numerical inequality, if $s \le 3$ then $s+1 \le 4$, while if $s \ge 4$ then $s+1 \le 4(s-2) \le 4(2g(B)-2+s)$. Thus the cases of genus zero or one and small S are included. The dependence on the auxiliary B_0 is harmless here, since it has been fixed for the given family. The corollary also applies without the normal-crossings assumption and allows unrestricted intersections above the removed disks.

The exponent in (5) can also be compared with the classical theory of ordinary integral points. Put

$$J(s) := \{P \in A(K) : \#f(\sigma_P(B) \cap \mathcal{D}) \le s\} \subseteq I(s, B_0).$$

For nonisotrivial elliptic curves, Hindry–Silverman's height estimates and Shioda's rank bound give polynomial counting bounds for ordinary integral points; see [22, §2.3] and [16, 24]. This is one motivation for improving the exponent for the larger set $I(s, B_0)$. The trace qualification is essential: on a constant elliptic curve with its constant zero divisor there are infinitely many constant integral sections. Our statements consistently count the image in $A(K)/T_A$.

In §3 we explain Phung's six-step proof in detail. Replacing his linear loop bound with Proposition 2.5 changes the lattice radius from $O(s+1)$ to $O(\sqrt{s+1})$ and halves the exponent. Finally, the same bounded cocycle representatives control coordinates in the Mordell–Weil lattice. Counting in its actual rank r proves (7); the quadratic nature of the canonical height proves Corollary C.

Note added. After the completion of this paper, Gao's preprint [13] became available on arXiv. Gao proves Vojta's geometric *abc* conjecture for rational points on arbitrary abelian varieties A over $K = \mathbb{C}(B)$ relative to reduced effective divisors D. In particular, his results imply the geometric Lang–Vojta bound of Conjecture 1.1 in the case $X = A$ over $\mathbb{C}(B)$.

More precisely, fix model data and a finite set $E \subset B$. By [13, Theorem 1.3], for every $0 < \varepsilon < 1$ there exist a constant $c > 0$ and a proper Zariski closed subset $Z_\varepsilon \subset A$ such that

$$N^{(1)}_{D,E}(P) \ge (1-\varepsilon)h_D(P) - c$$

for every $P \in (A \setminus \operatorname{Supp} D)(K)$ with $P \notin Z_\varepsilon$. Here h_D is a Weil height attached to D, and $N^{(1)}_{D,E}$ counts distinct boundary intersections outside E. Thus E is a set on the base curve where intersections are omitted, whereas Z_ε is the exceptional subset in the abelian variety. The constant and exceptional subset may depend on ε, E and the fixed data, but not on P.

In the notation of this paper, for a fixed model $(\mathcal{A}, \mathcal{D})$ with $\mathcal{D}$ effective Cartier, we may choose the associated height and counting function so that

$$h_D(P) = h_{\mathcal{D}}(P) = \deg_B \sigma_P^* \mathcal{D}, \qquad N^{(1)}_{D,E}(P) = \#\big(f(\sigma_P(B) \cap \mathcal{D}) \setminus E\big).$$

Taking $E = \varnothing$ counts intersections over the whole base curve; it does not remove the exceptional subset Z_ε. For every finite $S \subset B$ and every $(S, \mathcal{D})$-integral point $P \notin Z_\varepsilon$, we therefore obtain

$$\deg_B \sigma_P^* \mathcal{D} \le \frac{\#S + c}{1-\varepsilon}.$$

The constant and exceptional subset are independent of S and P. This implies the fixed-model boundary-degree bound in our formulation of geometric Lang–Vojta.

Gao also proves Buium's conjecture: for any effective divisor D, including nonreduced divisors, [13, Proposition 1.4] gives constants $c_1, c_2 > 0$ such that

$$N^{(1)}_{D,E}(P) \ge c_1 h_D(P) - c_2$$

for every $P \in (A \setminus \operatorname{Supp} D)(K)$, without an exceptional subset. Taking $E = \varnothing$ consequently gives a linear boundary-degree bound for all $(S, \mathcal{D})$-integral points, but with an unspecified positive coefficient c_1. Neither result requires trivial trace or `(P)`(iii).

The bounds of the present paper concern the larger sets $I(s, B_0)$ of generalized integral points: only intersections above B_0 are counted, while intersections above the removed disks are unrestricted. They arise from the uniform square-root bound for loop lengths in Theorem A.

Funding. The author is supported by the NSFC RFIS-I 2025, project titled *"Interdisciplinary Perspectives in Diophantine Geometry: Analytic, Dynamical and Arithmetic Explorations"*, and by the Zhejiang Science and Technology Development Fund [2024] No. 28, project titled *"New Perspectives in Diophantine Geometry"*.

2 Kobayashi length estimates

2.1 Kobayashi pseudo-metric

Let X be a connected complex manifold with complex structure J. Throughout, TX denotes the real tangent bundle $T^{\mathbb{R}}X$, equipped with the complex scalar multiplication

$$\lambda v = (\operatorname{Re}\lambda)v + (\operatorname{Im}\lambda)J_x v, \qquad x \in X, \quad v \in T^{\mathbb{R}}_x X, \quad \lambda \in \mathbb{C}.$$

It has complex rank $\dim_{\mathbb{C}} X$ and is canonically identified with the holomorphic tangent bundle $T^{1,0}X$ by

$$v \longmapsto \frac{1}{2}(v - iJ_x v).$$

We use these two descriptions interchangeably. This should be distinguished from the complexification

$$T^{\mathbb{R}}X \otimes_{\mathbb{R}} \mathbb{C} = T^{1,0}X \oplus T^{0,1}X,$$

whose complex rank is $2\dim_{\mathbb{C}} X$.

Definition 2.1. A *Finsler pseudo-metric* on X is a function

$$\begin{aligned} F\colon TX &\to \mathbb{R}_{\geq 0} \\ (x,v) &\mapsto F(x,v) \end{aligned}$$

satisfying the homogeneity condition $F(x,\lambda v) = |\lambda|\, F(x,v)$ for all $\lambda \in \mathbb{C}$. If $F(x,v) > 0$ for all $x \in X$ and $v \in T_{X,x} \setminus \{0\}$, then we say that F is a *Finsler metric*.

Throughout, by a *path* we mean a piecewise C^1 map $\gamma\colon [0,1] \to X$; this entails no loss of generality for the distances defined below. Given a Borel measurable Finsler pseudo-metric F and a path γ, the function $t \mapsto F(\gamma(t), \gamma'(t))$ is measurable, and we define the length of γ as

$$\ell_F(\gamma) := \int_0^1 F(\gamma(t), \gamma'(t))\, dt \ \in [0, +\infty].$$

When F is *upper semicontinuous* we have $\ell_F(\gamma) < +\infty$, since F is then locally bounded above and $\gamma([0,1])$ is compact. Given a set of paths Γ we set

$$\ell_F(\Gamma) := \inf_{\gamma \in \Gamma} \ell_F(\gamma)\,.$$

We endow X with the structure of extended pseudo-metric space by setting

$$d_F(x,y) := \inf_{\gamma} \ell_F(\gamma)\,, \quad \text{where the infimum runs over all paths } \gamma \text{ joining } x \text{ and } y\,.$$

Moreover, if Y, Z are two subsets of X, we put

$$\operatorname{diam}_F(Y) := \sup_{y,y' \in Y} d_F(y,y')\,, \quad d_F(Y,Z) := \inf_{y \in Y,\, z \in Z} d_F(y,z)\,.$$

One can show that a *continuous* Finsler metric induces a metric d_F. Any Hermitian metric h on X induces a Finsler metric $F_h(x,v) = \sqrt{h_x(v,v)}$, hence a distance which we denote by d_h.

An important example of a Finsler pseudo-metric is the *Kobayashi–Royden pseudo-metric*

$$\kappa_X(x,v) := \inf\left\{ \frac{2}{R} \in \mathbb{R}_{>0} \colon \exists\, f \in \operatorname{Hol}(\Delta(0,R), X) \text{ such that } f(0) = x,\ f'(0) = v \right\},$$

where $\operatorname{Hol}(\Delta(0,R), X)$ is the set of holomorphic maps from the one-dimensional complex open disk of radius R centered at 0 to X. Royden [23, Proposition 3] shows that κ_X is upper semicontinuous, hence Borel measurable. Since κ_X is intrinsic to X, the induced topological pseudo-metric is denoted by d_X and called the *Kobayashi pseudo-metric*. We say that X is *Kobayashi hyperbolic* if d_X is a metric.

From the definition one deduces the distance-decreasing property of the Kobayashi–Royden metric: if $\phi\colon X \to Y$ is holomorphic, then

$$\kappa_Y(\phi(x), d_x\phi(v)) \leq \kappa_X(x,v)\,, \quad \forall (x,v) \in TX\,.$$

This in turn implies the distance-decreasing property of the Kobayashi distance

$$d_Y(\phi(x),\phi(x')) \le d_X(x,x')\,, \quad \forall x,x' \in X\,.$$

If X carries a Hermitian metric h whose holomorphic sectional curvature is bounded above by a negative constant, then by the Ahlfors–Schwarz lemma one has $\kappa_X \ge c\,\|\cdot\|_h$ for some constant $c>0$; in particular $d_X \ge c\,d_h$ and X is Kobayashi hyperbolic. In dimensions at least two the converse is false in general: Demailly constructed a smooth projective Kobayashi hyperbolic surface which admits no Hermitian metric of negative holomorphic sectional curvature [10, Theorem 8.2]. On a Riemann surface, however, the converse holds. Suppose that X is Kobayashi hyperbolic. By the uniformization theorem, its universal cover is biholomorphic to $\mathbb{H}$. Let $\pi\colon \mathbb{H} \to X$ be the resulting holomorphic covering. Its deck-transformation group $\Gamma \subset \mathrm{Aut}(\mathbb{H}) = \mathrm{PSL}_2(\mathbb{R})$ acts freely and properly discontinuously. Thus Γ is a torsion-free Fuchsian group, and covering-space theory gives a biholomorphism $X \simeq \Gamma\backslash\mathbb{H}$.

Every element of Γ preserves the curvature -1 Poincaré metric

$$g_{\mathbb{H}} = \frac{dx^2+dy^2}{y^2},$$

which therefore descends to a complete Riemannian metric g_X on X, compatible with its complex structure and of curvature -1. Since holomorphic coverings preserve the Kobayashi–Royden infinitesimal metric, and $\kappa_{\mathbb{H}}(z,v) = \sqrt{g_{\mathbb{H},z}(v,v)}$, we obtain

$$\kappa_X(x,v) = \sqrt{g_{X,x}(v,v)}.$$

Consequently, the distance induced by g_X is the Kobayashi distance d_X. Thus, for Riemann surfaces, Kobayashi hyperbolicity is equivalent to admitting a Hermitian metric whose holomorphic sectional curvature is bounded above by a negative constant.

2.2 Growth of loop lengths in terms of the cusps

We fix a compact Riemann surface B of genus g endowed with a Riemannian metric ρ. Let $\overline{\Delta}_1,\dots,\overline{\Delta}_t$ be disjoint disks on B and define $B_0 := B \smallsetminus \sqcup_{i=1}^t \overline{\Delta}_i$. We initially assume that $t \ge 1$, so that B_0 is hyperbolic; the case $t=0$ is treated at the end of §2.3. Let $S \subseteq B_0$ be a finite set (possibly empty), put $s := \#S$, and let κ_S be the Kobayashi–Royden metric of the hyperbolic surface $B_0 \smallsetminus S$, with induced length and distance ℓ_S, d_S. We define the *distortion function*

$$\lambda_S(b) := \sup_{v \in T_bB \smallsetminus \{0\}} \frac{\kappa_S(b,v)}{\sqrt{\rho_b(v,v)}}\,, \qquad b \in B_0 \smallsetminus S\,, \tag{8}$$

where $T_bB = T_b^{\mathbb{R}}B$ is the *real* tangent plane (real dimension 2), on which both $\kappa_S(b,\cdot)$ and $\sqrt{\rho_b(\cdot,\cdot)}$ are defined. The ratio is homogeneous of degree 0, so the supremum may be taken over the compact ρ-unit circle. On this circle the denominator equals 1, and the numerator $\kappa_S(b,\cdot)$ is upper semicontinuous. Thus $\lambda_S(b)$ is finite and λ_S is Borel on $B_0 \smallsetminus S$. In a holomorphic coordinate $z = x+iy$, identifying $a\partial_x + b'\partial_y \leftrightarrow (a+ib')\partial_z$, $\mathbb{C}$-homogeneity gives $\kappa_S(b, a\partial_x + b'\partial_y) = \kappa_S(b,\partial_z)\sqrt{a^2+b'^2}$, while ρ_b has eigenvalues $\mu_1(b) \le \mu_2(b)$ in the frame (∂_x,∂_y). Maximizing the ratio over the Euclidean unit circle thus minimizes $\rho_b(v,v)$, giving

$$\lambda_S(b) = \frac{\kappa_S(b,\partial_z)}{\sqrt{\mu_1(b)}}\,.$$

If ρ is conformal ($\mu_1=\mu_2$, or equivalently ρ is hermitian) the ratio of Equation (8) is independent of v (so that the supremum on v is superfluous) and λ_S is smooth; we do not assume this.

Standing assumption for §2.2 (lifted in §2.3). Throughout this subsection, we fix a loop α in B_0 that doesn't represent a trivial class in $\pi_1(B_0,b_0)$ *and that admits a smooth simple representative* $\gamma \Subset B_0$ freely homotopic to α. Here and throughout, a *smooth simple loop* is an injective smooth immersion $\gamma\colon S^1 \to B_0$; since S^1 is compact, γ is then a smooth embedding, and its image is a smoothly embedded circle. In particular γ' never vanishes, so γ may be parametrized by ρ-arc length, with $\gamma'(0) = \gamma'(\ell_\rho(\gamma))$. Such loops exist: the basis loops of $\pi_1(B_0,b_0)$ described in §2.3 are smooth simple, and the general case (no simplicity hypothesis) is reduced to this one in §2.3.

Parametrize γ by its ρ-arc-length and let $\eta(\tau)$ be the unit normal vector of γ at τ; here and throughout we fix the convention that $\eta(\tau)$ is the rotation of the unit tangent vector $\gamma'(\tau)$ by $+\pi/2$, with respect to

ρ and the orientation of the Riemann surface B. Since γ is smooth and simple and compactly contained in B_0, for $\delta \in \mathbb{R}_{>0}$ sufficiently small the exponential map from the normal bundle is a diffeomorphism, yielding:

$$\begin{aligned} \mathbb{R}/\ell_\rho(\gamma)\mathbb{Z} \times [-\delta, \delta] &\to \{b \in B_0 \colon d_\rho(b, \gamma) \le \delta\} =: T_\delta(\gamma) \Subset B_0 \\ (\tau, u) &\mapsto \exp_{\gamma(\tau)}(u\eta(\tau)) \end{aligned}$$

The coordinates (τ, u) are called *Fermi coordinates*, and $T_\delta(\gamma)$ is the *Fermi strip* with respect to γ. For any $u \in [-\delta, \delta]$ we define the *parallel curve* $\gamma_u \colon [0, \ell_\rho(\gamma)] \to B_0$, $\gamma_u(\tau) := \exp_{\gamma(\tau)}(u\eta(\tau))$.

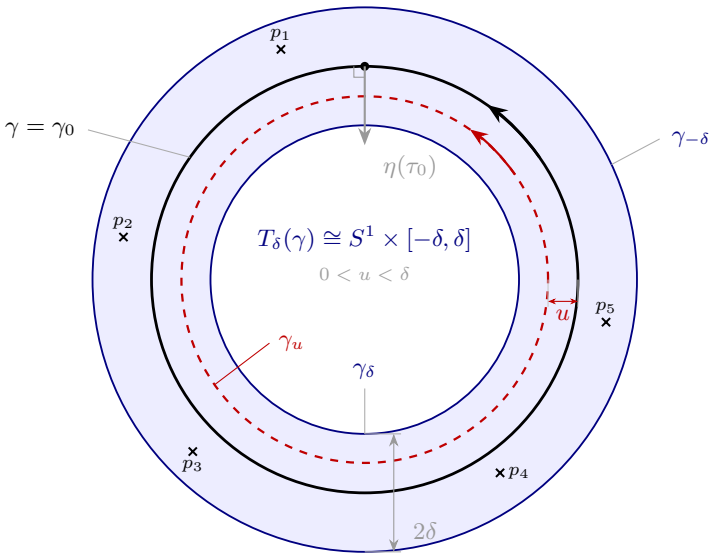


Figure 1: The Fermi strip drawn schematically as an annulus. The central loop γ and the parallel curve γ_u are closed; positive u points along the chosen normal. The punctures $p_1, \dots, p_5 \in S$ are avoided by the displayed γ_u.

Since γ is a hypersurface in B, its Fermi coordinates are semigeodesic [17, Example 6.43 and Corollary 6.42(a)]. This means that the u-coordinate curves are unit-speed geodesics orthogonal to the curves $u =$ constant. Consequently,

$$\rho_{uu} = 1, \qquad \rho_{\tau u} = 0, \qquad \rho = \rho_{\tau\tau}(\tau, u)\, d\tau^2 + du^2.$$

The area density is the square root of the determinant of the metric matrix. Since

$$\det \begin{pmatrix} \rho_{\tau\tau} & 0 \\ 0 & 1 \end{pmatrix} = \rho_{\tau\tau},$$

we obtain

$$dA_\rho = \sqrt{\rho_{\tau\tau}(\tau, u)}\, d\tau\, du. \tag{9}$$

Along a fixed parallel curve γ_u, we have $du = 0$, so its length element is $ds_\rho = \sqrt{\rho_{\tau\tau}(\tau, u)}\, d\tau$. Thus, for every nonnegative Borel function f on $T_\delta(\gamma)$,

$$\begin{aligned} \int_{T_\delta(\gamma)} f\, dA_\rho &= \int_{-\delta}^{\delta} \left(\int_0^{\ell_\rho(\gamma)} f(\gamma_u(\tau)) \sqrt{\rho_{\tau\tau}(\tau, u)}\, d\tau \right) du \\ &= \int_{-\delta}^{\delta} \left(\int_{\gamma_u} f\, ds_\rho \right) du. \end{aligned} \tag{10}$$

The first equality is the change-of-variables formula for $(\tau, u) \mapsto \gamma_u(\tau)$; the second uses the length element along γ_u.

Lemma 2.2. *Let $T \Subset B_0$ be a fixed compact set. There exists a constant $A_T > 0$, independent of the finite set $S \subset B_0$, such that*

$$\int_T \lambda_S\, dA_\rho \le A_T \sqrt{\#S + 1}.$$

The values of λ_S at points of S may be defined arbitrarily for this integral.

Proof. Put $s = \#S$ and $a = \mathrm{Area}_\rho(T)$, and set $\lambda_S = 0$ at the points of S. If $a = 0$ there is nothing to prove. We treat $s \geq 1$ first and address $s = 0$ at the end. We divide the argument into three steps, keeping track of why the constants are independent of the positions of the punctures.

Step 1: A local pointwise bound. There are constants $c_1, c_2, r_1, r_2 > 0$, depending only on B and ρ, such that

$$\kappa_{\Delta_\rho(b,r)}(b, v) \leq \frac{c_1}{r}|v|_\rho \quad (0 < r < r_1), \qquad \mathrm{Area}_\rho(\Delta_\rho(b, r)) \leq c_2 r^2 \quad (0 < r \leq r_2).$$

These are the local metric and area estimates of [22, Lemmas 3.3 and 3.5]. We recall why their constants can be chosen uniformly, including when ρ is not conformal. Choose finitely many holomorphic coordinate charts $z_i \colon U_i \to \mathbb{C}$, extending to neighborhoods of $\overline{U_i}$, and open sets $V_i \Subset U_i$ covering B. Smoothness and positive definiteness of ρ give constants $0 < \mu \leq \nu$ such that

$$\mu\,|dz_i(v)| \leq |v|_\rho \leq \nu\,|dz_i(v)| \qquad \text{on every } U_i.$$

Compactness of the finitely many $\overline{V_i}$ allows us to choose $r_0 > 0$ such that, for $b \in V_i$ and $0 < r < r_0$, the metric ball $\Delta_\rho(b, r)$ stays in U_i, and the Euclidean disk of radius r/μ centered at $z_i(b)$ stays in $z_i(U_i)$. In these coordinates,

$$\Delta(z_i(b), r/\nu) \subseteq z_i(\Delta_\rho(b, r)) \subseteq \Delta(z_i(b), r/\mu).$$

The first inclusion follows by estimating the ρ-length of a straight coordinate segment. For the second, a minimizing ρ-geodesic from b to a point of the metric ball stays in that ball, hence in U_i, and its Euclidean coordinate length is at most its ρ-length divided by μ. The Kobayashi metric at the centre of a Euclidean disk of radius R is $2|\cdot|/R$, with our normalization. Holomorphic contraction for the first inclusion therefore gives

$$\kappa_{\Delta_\rho(b,r)}(b, v) \leq \frac{2\nu}{r}|dz_i(v)| \leq \frac{2\nu}{\mu r}|v|_\rho.$$

Also $dA_\rho \leq \nu^2\, dx\, dy$ in the chart, so the second inclusion gives

$$\mathrm{Area}_\rho(\Delta_\rho(b, r)) \leq \pi \frac{\nu^2}{\mu^2} r^2.$$

Thus one may take $c_1 = 2\nu/\mu$, $c_2 = \pi\nu^2/\mu^2$, $r_1 = r_0$ and $r_2 = r_0/2$. These choices use only the fixed surface and background metric, and impose no separation condition on the punctures. Now choose

$$d_* := \tfrac{1}{2}\min\{d_\rho(T, \partial B_0), r_1, r_2\} > 0.$$

For $b \in T \setminus S$, set

$$r := \min\{d_*, d_\rho(b, S)\}.$$

This choice ensures that the open metric ball $\Delta_\rho(b, r)$ stays inside B_0 and avoids S. Indeed, $r \leq d_* < d_\rho(T, \partial B_0)$ prevents the ball from reaching ∂B_0, while $r \leq d_\rho(b, S)$ excludes every puncture. A puncture at distance exactly r is not in the open ball. Moreover, $0 < r \leq d_* \leq r_1/2 < r_1$, so the local metric estimate applies. The holomorphic inclusion $\Delta_\rho(b, r) \hookrightarrow B_0 \setminus S$ and that estimate give, for every $v \in T_b B$,

$$\kappa_S(b, v) \leq \kappa_{\Delta_\rho(b,r)}(b, v) \leq \frac{c_1}{r}|v|_\rho.$$

Dividing by $|v|_\rho$ for $v \neq 0$ and taking the supremum yields

$$\lambda_S(b) \leq \frac{c_1}{\min\{d_*, d_\rho(b, S)\}}. \tag{11}$$

For $S = \varnothing$, we use $d_\rho(b, S) = +\infty$; then $r = d_*$ and the same argument gives $\lambda_\varnothing(b) \leq c_1/d_*$ on T.

Step 2: The area of the high-distortion region. For $u > c_1/d_*$ and $b \in T \setminus S$, (11) gives

$$\lambda_S(b) > u \implies \min\{d_*, d_\rho(b, S)\} < c_1/u < d_* \implies d_\rho(b, S) < c_1/u.$$

Since $\lambda_S = 0$ on S, it follows that

$$E_u := \{b \in T : \lambda_S(b) > u\} \subseteq T \cap \bigcup_{p \in S} \Delta_\rho(p, c_1/u).$$

The area estimate from Step 1,

$$\operatorname{Area}_\rho \Delta_\rho(p,r) \le c_2 r^2 \qquad (p \in B,\ 0 < r \le r_2),$$

applies with $r = c_1/u < d_* \le r_2/2$. Since $E_u \subseteq T$, we have $\operatorname{Area}_\rho(E_u) \le a$. On the other hand, the covering by metric balls and subadditivity give

$$\operatorname{Area}_\rho(E_u) \le \sum_{p\in S} \operatorname{Area}_\rho \Delta_\rho(p, c_1/u) \le c_2 c_1^2 s/u^2,$$

where $\#S = s$; the balls need not be disjoint. Combining these two upper bounds yields

$$\operatorname{Area}_\rho(E_u) \le \min\{a, c_2 c_1^2 s/u^2\}. \tag{12}$$

Step 3: Integrating the distribution function. For $s \ge 1$, set $u_* = c_1\sqrt{c_2 s/a}$. By the layer-cake formula [18, Theorem 1.13], the bound $\operatorname{Area}_\rho\{b \in T : \lambda_S(b) > u\} \le a$, and (12),

$$\begin{aligned}
\int_T \lambda_S \, dA_\rho &= \int_0^\infty \operatorname{Area}_\rho(E_u)\, du \\
&\le \int_0^{c_1/d_*} a\, du + \int_{c_1/d_*}^\infty \min\{a, c_2 c_1^2 s/u^2\}\, du \\
&\le \frac{c_1 a}{d_*} + \int_0^\infty \min\{a, c_2 c_1^2 s/u^2\}\, du \\
&= \frac{c_1 a}{d_*} + \int_0^{u_*} a\, du + c_2 c_1^2 s \int_{u_*}^\infty u^{-2}\, du \\
&= \frac{c_1 a}{d_*} + a u_* + \frac{c_2 c_1^2 s}{u_*} \\
&= \frac{c_1 a}{d_*} + 2c_1\sqrt{c_2 a s} \le A_T\sqrt{s+1},
\end{aligned}$$

where

$$A_T = \frac{c_1 a}{d_*} + 2c_1\sqrt{c_2 a}$$

is independent of S. For $s = 0$, Step 1 gives $\int_T \lambda_\varnothing \, dA_\rho \le c_1 a/d_* \le A_T$. □

Definition 2.3. For a nontrivial loop α in B_0, let Γ_S^α be the set of loops in $B_0 \setminus S$ freely homotopic to α in B_0, and set

$$L(\alpha, s) = \sup_{\substack{S \subset B_0 \\ \#S \le s}} \inf_{\gamma \in \Gamma_S^\alpha} \ell_S(\gamma).$$

For the bordered surfaces considered here, nontrivial free homotopy classes have positive length in B_0; see [1, Proposition 2.24]. Distance-decreasing therefore gives a positive lower bound for the infimum defining $L(\alpha, s)$.

Proof of Theorem A for simple loops. Choose a smooth simple representative $\gamma \Subset B_0$ of α and a Fermi strip $T = T_\delta(\gamma) \Subset B_0$, with coordinates defined for $|u| \le \delta$. Set $I = (-\delta, \delta)$ and

$$E = \{u \in I : \gamma_u \cap S \ne \varnothing\}.$$

Injectivity of the strip coordinates gives $\#E \le \#S$. Recall that in Fermi coordinates

$$\gamma_u'(\tau) = \partial_\tau|_{(\tau,u)}, \qquad \rho_{\tau\tau}(\tau, u) = \rho_{\gamma_u(\tau)}(\gamma_u'(\tau), \gamma_u'(\tau)).$$

For $u \in I \setminus E$, the definition of λ_S as a directional supremum therefore gives

$$\begin{aligned}
\kappa_S(\gamma_u(\tau), \gamma_u'(\tau)) &\le \lambda_S(\gamma_u(\tau))\sqrt{\rho_{\gamma_u(\tau)}(\gamma_u'(\tau), \gamma_u'(\tau))} \\
&= \lambda_S(\gamma_u(\tau))\sqrt{\rho_{\tau\tau}(\tau, u)}.
\end{aligned}$$

By the definition of Kobayashi length,

$$\begin{aligned}\ell_S(\gamma_u) &= \int_0^{\ell_\rho(\gamma)} \kappa_S(\gamma_u(\tau), \gamma_u'(\tau))\, d\tau \\ &\le \int_0^{\ell_\rho(\gamma)} \lambda_S(\gamma_u(\tau))\sqrt{\rho_{\tau\tau}(\tau, u)}\, d\tau.\end{aligned}$$

The length function is continuous on $I \setminus E$, since locally the parallel curves vary smoothly in a compact subset of $B_0 \setminus S$. Extend it by zero on the finite set E; the resulting function is measurable. Integrating the preceding inequality in u and using (10) with $f = \lambda_S$ extended by zero on S, we obtain the equality below; the final inequality follows from Lemma 2.2:

$$\begin{aligned}\int_I \ell_S(\gamma_u)\, du &\le \int_I \int_0^{\ell_\rho(\gamma)} \lambda_S(\gamma_u(\tau))\sqrt{\rho_{\tau\tau}(\tau, u)}\, d\tau\, du \\ &= \int_T \lambda_S\, dA_\rho \le A_T\sqrt{\#S + 1}.\end{aligned} \tag{13}$$

Set $M = A_T\sqrt{\#S+1}/(2\delta)$. We claim that $\ell_S(\gamma_{u_0}) \le M$ for some $u_0 \in I \setminus E$. Otherwise, the measurable sets

$$F_m = \{u \in I \setminus E : \ell_S(\gamma_u) \ge M + 1/m\}, \qquad m \ge 1,$$

satisfy $\bigcup_{m\ge1} F_m = I \setminus E$. Since $|I \setminus E| = 2\delta$, some F_m has positive Lebesgue measure. Hence

$$\int_I \ell_S(\gamma_u)\, du \ge M|I \setminus E| + \frac{|F_m|}{m} > 2\delta M = A_T\sqrt{\#S+1},$$

contrary to (13).

The family γ_{tu_0}, $0 \le t \le 1$, gives a free homotopy in B_0 from γ to γ_{u_0}. This homotopy may cross S, whereas γ_{u_0} avoids S; thus $\gamma_{u_0} \in \Gamma_S^\alpha$. Consequently, for every S with $\#S \le s$,

$$\inf_{\eta\in\Gamma_S^\alpha} \ell_S(\eta) \le \frac{A_T}{2\delta}\sqrt{\#S+1} \le \frac{A_T}{2\delta}\sqrt{s+1}.$$

The constants are independent of the cardinality and positions of S. Taking the supremum over all such S, including $S = \varnothing$, gives $L(\alpha, s) \le (A_T/2\delta)\sqrt{s+1}$. Moreover, all selected loops lie in the fixed compact strip T, and

$$L_0 := \max_{|u|\le\delta} \ell_\rho(\gamma_u) < \infty$$

by smooth dependence on u and compactness of $[-\delta, \delta]$. □

For the application to generalized integral points, the Parshin cocycle argument requires loops representing the generators of $\pi_1(B_0, b_0)$ that share a *common base point* and are conjugated to the generators by a *single* path. Theorem A, applied independently to each generator, produces loops with the square-root Kobayashi length bound but with potentially different base points. The next proposition shows that a common base point can be arranged without degrading the asymptotic bound.

We first record the local tangent adjustment that makes this possible.

Lemma 2.4. *Let M be a smooth orientable Riemannian surface without boundary, let $p \in M$, and let $w \in T_pM$ be a unit tangent vector. Then for every smooth simple loop γ at p with $\gamma \Subset M$, there exists a smooth simple loop $\tilde\gamma$ at p with $\tilde\gamma \Subset M$, homotopic to γ relative to p, and satisfying*

$$\frac{\tilde\gamma'(0)}{|\tilde\gamma'(0)|} = w\,.$$

Proof. Let $w_0 := \gamma'(0)/|\gamma'(0)|$ be the unit tangent vector of γ at p. If $w_0 = w$ there is nothing to prove, so assume $w_0 \ne w$. Let $\theta \in]0, 2\pi[$ be the angle from w_0 to w and for each $t \in [0,1]$ let $R_t \in \mathrm{SO}(2)$ denote the rotation of T_pM by angle $t\theta$, so that $R_0 = \mathrm{Id}$ and $R_1(w_0) = w$.

Choose $\epsilon > 0$ small enough that the exponential map $\exp_p$ is a diffeomorphism from $B(0,\epsilon) \subset T_pM$ onto an open neighborhood U of p in M with $U \Subset M$. Choose a smooth cutoff function $\chi\colon [0,\infty) \to [0,1]$ with $\chi \equiv 1$ on $[0, \epsilon/3]$ and $\chi \equiv 0$ on $[2\epsilon/3, \infty)$. Define the map $\phi\colon M \to M$ by

$$\phi(x) \;:=\; \begin{cases} \exp_p\big(R_{\chi(|y|)}\, y\big), & \text{if } x = \exp_p(y) \in U, \\ x, & \text{if } x \notin U, \end{cases}$$

where $|y|$ denotes the norm in T_pM. Since $\chi \equiv 1$ near $|y| = 0$, the map ϕ equals $\exp_p \circ R_1 \circ \exp_p^{-1}$ in a neighborhood of p, and since $\chi \equiv 0$ for $|y| \geq 2\epsilon/3$, ϕ is the identity outside U. In geodesic polar coordinates (r, ψ) centered at p (via $\exp_p$), the map reads $\phi(r, \psi) = (r,\ \psi + \chi(r)\theta)$: it preserves the radius r and rotates each geodesic circle $\{r = \text{const}\}$ by the angle $\chi(r)\theta$. On each circle this is a bijection, and distinct circles map to distinct circles, so ϕ is a bijection of U fixing p; together with $\phi = \mathrm{id}$ outside U this makes ϕ a bijection of M. It is smooth with smooth inverse $\phi^{-1}(r, \psi) = (r, \psi - \chi(r)\theta)$ (both expressions are smooth across $r = 0$ because $\chi \equiv 1$ there, where ϕ is the fixed rotation R_1), so ϕ is a diffeomorphism of M. Moreover:

(i) $\phi(p) = p$ and $d_p\phi = R_1$, so $(\phi \circ \gamma)'(0) = R_1(\gamma'(0)) = |\gamma'(0)|\, w$, which has unit tangent direction w;

(ii) ϕ is the identity outside U, so $\phi \circ \gamma$ coincides with γ outside U and in particular $\phi \circ \gamma \Subset M$;

(iii) ϕ is a diffeomorphism, so $\phi \circ \gamma$ is simple;

(iv) the family ϕ_t (defined by replacing χ with $t\chi$ for $t \in [0, 1]$) satisfies $\phi_0 = \mathrm{Id}$ and $\phi_1 = \phi$, and $\phi_t(p) = p$ for all t. Hence $t \mapsto \phi_t \circ \gamma$ is a homotopy of loops based at p from γ to $\tilde{\gamma}$, so $[\tilde{\gamma}] = [\gamma]$ in $\pi_1(M, p)$.

Setting $\tilde{\gamma} := \phi \circ \gamma$ completes the proof. □

Proposition 2.5. *Assume $t \geq 1$ and fix a free basis $\alpha_1, \ldots, \alpha_k$ of $\pi_1(B_0, b_0)$ whose elements have smooth simple representatives, with $k \geq 1$. There are constants $C, \delta', L_0 > 0$ and a compact set $\mathcal{K}_0 \Subset B_0$ such that for every finite $S \subset B_0$ there are a point $q \in B_0 \setminus S$, a path σ from b_0 to q in $\mathcal{K}_0$, and loops $\widehat{\gamma}_j \subset \mathcal{K}_0 \setminus S$ based at q, satisfying*

$$[\sigma^{-1} \circ \widehat{\gamma}_j \circ \sigma] = \alpha_j, \qquad \ell_S(\widehat{\gamma}_j) \leq C\sqrt{\#S + 1},$$

and

$$\ell_\rho(\sigma) \leq \delta', \qquad \ell_\rho(\widehat{\gamma}_j) \leq L_0.$$

The constants and $\mathcal{K}_0$ are independent of S.

Proof. Fix a unit tangent vector $w \in T_{b_0}B_0$ and let η_0 be the rotation of w by $+\pi/2$ with respect to ρ and the orientation of B. By Lemma 2.4, for each $j = 1, \ldots, k$ we can choose a smooth simple loop $\gamma^{(j)}$ at b_0, compactly contained in B_0, homotopic to α_j relative to b_0, and whose unit tangent vector at b_0 equals w. Since all loops share the tangent direction w at b_0, and since the unit normal field along every loop is the $+\pi/2$-rotation of the unit tangent (the convention fixed in §2.2), their unit normal vectors at b_0 all equal η_0. For each j, construct the Fermi strip $T_\delta(\gamma^{(j)})$ as in the proof of Theorem A, with a common width $\delta > 0$ small enough for all k strips. The parallel curve $\gamma_u^{(j)}$ has base point

$$q(u) \ :=\ \exp_{b_0}(u\,\eta_0),$$

which is independent of j.

The strips need not be disjoint from one another; only the coordinate map on each individual strip must be injective. For each j, let

$$E_j = \{u \in (-\delta, \delta) : \gamma_u^{(j)} \cap S \neq \varnothing\}, \qquad E = \bigcup_{j=1}^{k} E_j.$$

Each E_j has at most $\#S$ elements, so E is finite. Define $g_j(u) = \ell_S(\gamma_u^{(j)})$ outside E and set $g_j = 0$ on E. Formula (13), applied separately to the finitely many strips $T_j = T_\delta(\gamma^{(j)})$, gives

$$\frac{1}{2\delta}\int_{-\delta}^{\delta} \sum_{j=1}^{k} g_j(u)\, du \leq \frac{\sum_{j=1}^{k} A_{T_j}}{2\delta}\sqrt{\#S + 1} =: M_S.$$

By the mean-value argument above, there is $u_0 \notin E$ with $\sum_j g_j(u_0) \leq M_S$. Since every summand is nonnegative, $g_j(u_0) \leq M_S$ for every j. This is why we average the sum: selecting an averaging parameter separately for each loop would not provide a common base point.

Set $\widehat{\gamma}_j = \gamma_{u_0}^{(j)}$ and $q = q(u_0)$. Since the loops all avoid S and pass through q, we also have $q \notin S$. The path

$$\sigma(t) = \exp_{b_0}(t u_0 \eta_0), \qquad 0 \leq t \leq 1,$$

is defined also for $u_0 = 0$, when it is constant, and satisfies $\ell_\rho(\sigma) = |u_0| \le \delta$. For every j, the homotopy $\gamma^{(j)}_{tu_0}$ moves its base point along this same path. Consequently

$$[\sigma^{-1} \circ \hat{\gamma}_j \circ \sigma] = [\gamma^{(j)}] = \alpha_j \quad \text{in } \pi_1(B_0, b_0).$$

Only the loops, not the conjugating path σ, are required to avoid S. All these homotopies take place in B_0.

Finally, take

$$\mathcal{K}_0 = \bigcup_{j=1}^{k} T_j, \qquad L_0 = \max_j \sup_{|u| \le \delta} \ell_\rho(\gamma^{(j)}_u), \qquad \delta' = \delta.$$

The finite union $\mathcal{K}_0$ is compactly contained in B_0, contains every selected loop and the normal segment σ, and L_0 is finite by smoothness and compactness. Taking $C = (\sum_j A_{T_j})/(2\delta)$ proves all the assertions, with constants independent of S. □

2.3 Reduction to the simple case

In this subsection we lift the standing simplicity assumption of §2.2 and prove Theorem A for every nontrivial class.

For $t \ge 1$, B_0 has the homotopy type of a compact orientable surface of genus g with t boundary circles, so its fundamental group is free of rank $k = 2g + t - 1$ [26, Theorem 3.2.8]. The disk-and-bands model gives a basis of embedded piecewise smooth loops. Smoothing each loop separately, with the base point fixed, gives smooth simple representatives of the same based classes. We do not require the smoothed representatives to remain pairwise disjoint away from the base point; this is unnecessary for the argument. For background on curves in surfaces see [12, Chapter 1]. Up to this smoothing, this is the *simple base* of [22, §3.1], which we call a *simple basis*.

Our path convention is that $c_2 \circ c_1$ first traverses c_1 and then c_2, whereas the product $[c_1][c_2]$ in the fundamental group is represented by $c_2 \circ c_1$.

Proof of Theorem A, general case. First suppose $t \ge 1$. Fix a nontrivial loop α in B_0. Choose once and for all a representative of its free homotopy class based at b_0, and denote its based class by α as well. Fix a simple basis $\alpha_1, \dots, \alpha_k$ of $\pi_1(B_0, b_0)$ as above. Write α as a reduced word in this basis:

$$\alpha = \alpha_{i_1}^{\epsilon_1} \cdots \alpha_{i_m}^{\epsilon_m} \quad \text{in } \pi_1(B_0, b_0), \qquad \epsilon_j \in \{\pm 1\}, \quad i_j \in \{1, \dots, k\}.$$

The word length $m \ge 1$ is determined by the chosen based class and the basis, which are fixed independently of S.

Fix $S \subset B_0$ with $\#S = s \ge 0$. By Proposition 2.5 applied to the simple basis $\alpha_1, \dots, \alpha_k$, there exist

- a point $q \in B_0 \setminus S$,
- a path $\sigma\colon [0,1] \to B_0$ from b_0 to q with $\ell_\rho(\sigma) \le \delta'$,
- loops $\hat{\gamma}_1, \dots, \hat{\gamma}_k \subset B_0 \setminus S$ based at q,

such that $[\sigma^{-1} \circ \hat{\gamma}_j \circ \sigma] = \alpha_j$ in $\pi_1(B_0, b_0)$ and

$$\ell_S(\hat{\gamma}_j) \le C_0 \sqrt{s+1} \quad \text{for all } j = 1, \dots, k,$$

where $C_0 > 0$ depends only on $\alpha_1, \dots, \alpha_k$, B_0, and ρ.

Define the concatenation

$$\hat{\gamma}_\alpha := \hat{\gamma}_{i_m}^{\epsilon_m} \circ \hat{\gamma}_{i_{m-1}}^{\epsilon_{m-1}} \circ \cdots \circ \hat{\gamma}_{i_1}^{\epsilon_1},$$

a piecewise smooth loop based at q. We verify that $\hat{\gamma}_\alpha \in \Gamma_S^\alpha$ and bound its Kobayashi length.

Membership in Γ_S^α. Each $\hat{\gamma}_{i_j}$ is contained in $B_0 \setminus S$, hence so is the concatenation $\hat{\gamma}_\alpha$. To see that $\hat{\gamma}_\alpha$ is freely homotopic to α in B_0, insert the null-homotopic backtracking loop $\sigma \circ \sigma^{-1}$ between consecutive factors. The change-of-base-point map is a group homomorphism, so

$$\begin{aligned}[\sigma^{-1} \circ \hat{\gamma}_\alpha \circ \sigma] &= [\sigma^{-1} \circ \hat{\gamma}_{i_1}^{\epsilon_1} \circ \sigma] \cdots [\sigma^{-1} \circ \hat{\gamma}_{i_m}^{\epsilon_m} \circ \sigma] \\ &= \alpha_{i_1}^{\epsilon_1} \cdots \alpha_{i_m}^{\epsilon_m} = \alpha \quad \text{in } \pi_1(B_0, b_0).\end{aligned}$$

Thus $\hat{\gamma}_\alpha$, viewed as a loop based at q, is conjugate (via the path σ) to a loop based at b_0 representing α; equivalently, the free homotopy class of $\hat{\gamma}_\alpha$ in B_0 coincides with the free homotopy class of α.

Kobayashi length bound. Since κ_S is a Finsler pseudo-metric, its homogeneity gives $\kappa_S(b,-v)=\kappa_S(b,v)$, hence $\ell_S(\hat{\gamma}_{i_j}^{-1})=\ell_S(\hat{\gamma}_{i_j})$. By additivity of ℓ_S over concatenation,

$$\ell_S(\hat{\gamma}_\alpha)=\sum_{j=1}^{m}\ell_S\big(\hat{\gamma}_{i_j}^{\epsilon_j}\big)=\sum_{j=1}^{m}\ell_S(\hat{\gamma}_{i_j})\le m\cdot C_0\sqrt{s+1}.$$

Conclusion. Setting $C:=m\cdot C_0$, which depends only on the chosen based class α (through its word length m in the fixed basis), the simple basis $\alpha_1,\dots,\alpha_k$, B_0, and ρ, we obtain

$$\ell_S(\Gamma_S^\alpha)\le\ell_S(\hat{\gamma}_\alpha)\le C\sqrt{s+1}.$$

As $S\subset B_0$ with $\#S=s\ge 0$ was arbitrary, and since sets with $\#S<s$ satisfy the same bound a fortiori (the right-hand side being increasing in s), taking the supremum over all S with $\#S\le s$ yields

$$L(\alpha,s)\le C\sqrt{s+1}\quad\text{for all } s\ge 0.$$

The case $S=\varnothing$ is already included in Proposition 2.5.

If $t=0$, choose a fixed small closed disk $\overline{\Delta}\subset B$ disjoint from a piecewise smooth representative of α, independently of S, and put $U=B\setminus\overline{\Delta}$. The class remains nontrivial in U, since a null-homotopy there would also be one in B. Apply the bordered result to U and $S\cap U$. The resulting loop avoids all of S, and the inclusion $U\setminus(S\cap U)\hookrightarrow B\setminus S$ decreases Kobayashi lengths. Its free homotopy class in B is that of α. This proves the assertion for compact hyperbolic B as well. □

Corollary 2.6. *For every nontrivial α as above,*

$$\lim_{s\to\infty}\frac{\log L(\alpha,s)}{\log s}=\frac{1}{2}.$$

Proof. Combine Theorem A with [22, Theorem C]:

$$\frac{c\sqrt{s}}{\log(s+2)}\le L(\alpha,s)\le C\sqrt{s+1}.$$

Taking logarithms and dividing by $\log s$ proves the assertion. □

Remark 2.7. The quoted lower bound and our upper bound differ by a factor of order $\log(s+2)$. These estimates determine the polynomial growth exponent, but do not determine a leading constant or imply $L(\alpha,s)\asymp\sqrt{s}$.

3 Bounding generalized integral points

3.1 Phung's bound

In this section we revisit the proof of Theorem 1.3. We follow Phung's strategy [22] closely, and we recall the structure of the argument in enough detail to identify where our improvement enters (§3.2). Our account differs from *loc. cit.* in the following respects; each point is taken up again, with full notation, at the corresponding place of the proof below.

- Phung bounds the Kobayashi distance of $(\mathcal{A}\smallsetminus\mathcal{D})|_{B_0}$ from below by a fixed hermitian metric by invoking [22, Theorem 7.5], a criterion of Brody–Green type. We give a short self-contained proof of the infinitesimal comparison actually needed here, based on Brody's reparametrization lemma and Hurwitz's theorem [20, Lemmas 7.2.1 and 7.2.11]. In particular, we check the hyperbolicity of the two possible strata containing a limiting entire curve.
- At one point the argument requires a path of uniformly bounded length joining two points of the same fibre *within that fibre*. In *loc. cit.* this length is bounded by the diameter of the closure of $\mathcal{A}_{B_0}$ in $\mathcal{A}$. This is not sufficient: two points of a fibre may be close to each other in $\mathcal{A}$ even though every path joining them inside the fibre is long. We bound the length instead by the diameter of the fibre itself, measured with respect to the metric induced on the fibre, and we check that these diameters are uniformly bounded.

- Phung converts bounds on the lengths of certain loops in $\mathcal{A}_{B_0}$ into bounds on the displacements, along the fibre direction of the universal cover, of the corresponding deck transformations. This conversion needs justification: the metric pulled back to the universal cover need not be a product metric, so the projection onto the fibre direction does not decrease lengths in general. We show that all the lifted paths in question stay in a fixed region, independent of s and of the section, on which the differential of this projection has a uniform bound, giving the required estimate for lengths of projected paths.
- Several routine reductions are carried out in detail: the case where $\pi_1(B_0)$ is trivial, the reduction to $t \geq 1$, the enlargement of the disks $\overline{\Delta}_i$ absorbing the (countable) set of points of B whose fibre in $\mathcal{D}$ fails to be hyperbolic, and the verification that the two conjugating paths naturally appearing in the construction differ by a loop whose homotopy class lies in the fibre subgroup, so that the associated cocycles are cohomologous.

The statement and the strategy of Theorem 1.3 are entirely Phung's; the purpose of the details above is to make the argument self-contained in the form needed for §3.2.

Proof of Theorem 1.3. Fix a hermitian metric h on the smooth complex manifold $\mathcal{A}_U$, where $U := B \smallsetminus \mathrm{Sing}(f)$. Let k be the minimal number of generators of $\pi_1(B_0, b_0)$, as in the introduction.
The rank-zero case. If $k = 0$ then $\pi_1(B_0, b_0)$ is trivial, so $H^1(\pi_1(B_0, b_0), \Gamma) = 0$ for the fibre lattice $\Gamma = H_1(A_{b_0}, \mathbb{Z}) \cong \mathbb{Z}^{2n}$. By the Parshin–Phung injectivity ([21, §3, Proposition 1], [22, Proposition 7.2]; cf. Proposition 3.1 below) the map $A(K)/T_A \to H^1(\pi_1(B_0, b_0), \Gamma) = 0$ is the zero homomorphism; hence every two points of $A(K)$ are congruent modulo $T_A + A(K)_{\mathrm{tors}}$, so $A(K)/T_A$ is finite of order $t_{\mathcal{A}} := \#\big(A(K)/T_A\big)_{\mathrm{tors}} < \infty$ (Lang–Néron). Consequently $\#\big(I(s, B_0)/T_A\big) \leq t_{\mathcal{A}}$ for every s, and the bounds of Theorem 1.3 and Theorem B hold trivially (take $m \geq t_{\mathcal{A}}$). We therefore assume $k \geq 1$ from now on.
Reduction to $t \geq 1$. As in the proof of [22, Theorem A], we may moreover assume $t \geq 1$: if $t = 0$ (so that $\mathrm{Sing}(f) = \varnothing$ and $B_0 = B$ is a compact hyperbolic surface of genus $g \geq 2$), remove one closed disk $\overline{\Delta}_1 \subset B$. This shrinks B_0, so $I(s, B_0)$ only grows, and it leaves the rank unchanged, $\mathrm{rk}\,\pi_1(B \smallsetminus \overline{\Delta}_1) = 2g + 1 - 1 = 2g = \mathrm{rk}\,\pi_1(B)$; hence the bound for the new B_0 implies the bound for the old one. With $t \geq 1$ the surface B_0 has the homotopy type of a compact surface of genus g with $t \geq 1$ boundary circles, so $\pi_1(B_0, b_0)$ is a *free* group of rank $k = 2g + t - 1$; this freeness will be used in Steps 3 and 6, and it places us within the standing hypotheses of §2.2.
Reduction to $s \geq 1$. Since $I(0, B_0) \subseteq I(1, B_0)$, the bound for $s = 1$ implies the bound for $s = 0$ after enlarging m. We therefore assume $s \geq 1$ from now on.

Step 1: Reduction to the hyperbolic locus. Consider the non-hyperbolic locus

$$V := \{b \in B \colon (\mathcal{D}_b)_{\mathrm{red}} \text{ is not hyperbolic}\}$$

It is well known that V is a closed subset of B in the analytic (Euclidean) topology: the hyperbolic locus of a proper holomorphic map between complex spaces is open ([20, Theorem 7.2.15], applied to $f|_{\mathcal{D}} \colon \mathcal{D} \to B$). Moreover, by [22, Lemma B.2] applied to the abelian scheme $\mathcal{A}|_{B \smallsetminus \mathrm{Sing}(f)}$, via the inclusion $V \smallsetminus \mathrm{Sing}(f) \subseteq Z\big(\mathcal{A}|_{B \smallsetminus \mathrm{Sing}(f)}, \mathcal{D}|_{B \smallsetminus \mathrm{Sing}(f)}\big)$ furnished by [22, Theorem 7.6], and using crucially the assumption (P)(iii) that $D_{\overline{K}}$ contains no translate of a positive-dimensional abelian subvariety, the set $V \smallsetminus \mathrm{Sing}(f)$ — and hence V, as $\mathrm{Sing}(f)$ is finite — is *at most countable*. The geometric form of (P)(iii) is used here: a dominating component of the scheme parametrizing positive-dimensional translates would produce such a translate after a finite extension of K.

We may therefore enlarge the disks $\overline{\Delta}_1, \dots, \overline{\Delta}_t$ so as to remove $V \cup \mathrm{Sing}(f)$ from the base. First, a collar of each $\partial\overline{\Delta}_i$ is foliated by parallel Jordan curves, and only countably many of these meet the countable set $V \cup \mathrm{Sing}(f)$; we may thus enlarge each $\overline{\Delta}_i$ inside its collar to a closed disk $\overline{\Delta}_i^+$ whose boundary avoids $V \cup \mathrm{Sing}(f)$, the $\overline{\Delta}_i^+$ remaining pairwise disjoint. In particular $\mathrm{Sing}(f) \subset \bigsqcup_i \mathrm{int}\,\overline{\Delta}_i^+$. The set $W := (V \cup \mathrm{Sing}(f)) \smallsetminus \bigsqcup_i \mathrm{int}\,\overline{\Delta}_i^+$ is closed and disjoint from the compact set $\bigsqcup_i \overline{\Delta}_i^+$. If W is empty, set $\overline{\Delta}_i' = \overline{\Delta}_i^+$ and omit the following attachments. Otherwise its ρ-distance d_0 from that compact set is positive. Applying [22, Lemma 8.1] to W with disks of radius at most ε, for some $0 < \varepsilon < d_0/2$, yields finitely many pairwise disjoint closed cover disks whose union contains W in its interior and each of which is disjoint from every $\overline{\Delta}_j^+$. The boundaries can be chosen disjoint from W, as in the proof of that lemma, so that the interiors cover W. Annex the cover disks successively to the first enlarged original disk: join the current disk to the next cover disk V' by an embedded arc avoiding the other disks, and take a small

regular neighborhood of their union. Such an arc exists because the complement of finitely many disjoint disks is connected. Equivalently, attach V' by a thin tube β meeting the two disks only in short boundary arcs at its ends. At each stage the resulting region is again a closed disk. Denote the final enlarged disks by $\overline{\Delta}'_i$. Each $\overline{\Delta}'_i$ is simply connected with piecewise-C^1 boundary, hence again a closed disk (in general not a metric ball), and the $\overline{\Delta}'_i$ are pairwise disjoint, with distinct points of $\mathrm{Sing}(f)$ still lying in distinct disks. By construction $V \cup \mathrm{Sing}(f) \subseteq \bigsqcup_{i=1}^{t} \operatorname{int} \overline{\Delta}'_i$. Writing $B'_0 := B \smallsetminus \bigsqcup_{i=1}^{t} \overline{\Delta}'_i \subseteq B_0$, the inclusion $B'_0 \hookrightarrow B_0$ is a homotopy equivalence: the region between each original disk and the corresponding enlarged disk is an annulus, which retracts to its outer boundary on the compact surfaces with boundary; collars identify each such compact surface up to homotopy with its interior. In particular the induced map on fundamental groups is an isomorphism, so $\operatorname{rk} \pi_1(B'_0) = \operatorname{rk} \pi_1(B_0) = k$ is unchanged. Since $B'_0 \subseteq B_0$ we have $I(s, B_0) \subseteq I(s, B'_0)$ for every s, so it suffices to prove the bound for B'_0. Replacing B_0 by B'_0 and choosing b_0 in the new surface, we may thus assume from now on that $V \cup \mathrm{Sing}(f) \subseteq \bigsqcup_{i=1}^{t} \operatorname{int} \overline{\Delta}_i$; equivalently, the compact set $\overline{B_0}$ is disjoint from $V \cup \mathrm{Sing}(f)$. In particular $(\mathcal{D}_b)_{\mathrm{red}}$ is Kobayashi hyperbolic for every $b \in \overline{B_0}$; hence $\mathcal{D}_b$ contains no translate of a positive-dimensional subtorus of $\mathcal{A}_b$, since such a translate would be a non-hyperbolic complex subspace of the hyperbolic space $\mathcal{D}_b$. By a result of Green ([20, Theorem 7.3.8(ii)], applied to the underlying reduced divisor of $\mathcal{D}_b$; see also [22, Theorem 7.6]) it follows that $\mathcal{A}_b \setminus \mathcal{D}_b$ is hyperbolically embedded into $\mathcal{A}_b$, in particular Kobayashi hyperbolic, for any $b \in \overline{B_0}$. Finally, fix a compact connected smooth subsurface $N \Subset B_0$ with $b_0 \in N$ (it will be enlarged below finitely many times, by fixed compact sets independent of s and P). The conjugating paths $c_{b_0,b}$ are chosen as in the finite reference-loop construction of [22, Lemma 4.1]: choose a fixed path from b_0 to each reference base point, then move to b within the corresponding simply connected neighborhood. These choices give a collection $\mathcal{C}_{b_0}$ of paths in a fixed compact set, with uniformly bounded ρ-length; the same path is used for all generators. Enlarge N to contain them and choose $\delta_N > 0$ so that

$$\ell_\rho(c_{b_0,b}) \leq \delta_N < +\infty\,.$$

Step 2: Bounding loop lengths in the base. Let $P \in I(s, B_0)$ and set $S := f(\sigma_P(B_0) \cap \mathcal{D}) \cap B_0$, so that $\#S \leq s$ and $\sigma_P(B_0 \smallsetminus S) \subseteq (\mathcal{A} \smallsetminus \mathcal{D})|_{B_0 \smallsetminus S}$; we may assume $S \neq \emptyset$ (otherwise replace S by an arbitrary singleton of B_0: this only increases the Kobayashi lengths below, and $\#S \leq s$ still holds since $s \geq 1$). Let $\alpha_1, \ldots, \alpha_k$ be a simple basis of $\pi_1(B_0, b_0)$. By [22, Theorem 5.1], applied with $\mathcal{C}_{b_0}$ as the fixed collection of paths of [22, Definition 3.2] (with the paths chosen in Step 1; the base points b produced there lie in $B_\varepsilon \subseteq N$, cf. below), there exist $b \in B_0$ and piecewise smooth loops $\gamma_j \in \Gamma_S^{\alpha_j}$ such that $\alpha_j = c_{b_0,b}^{-1}\, \gamma_j\, c_{b_0,b}$ in $\pi_1(B_0, b_0)$ where $c_{b_0,b} \in \mathcal{C}_{b_0}$ and the following bound holds:

$$\ell_S(\gamma_j) \leq L(s+1) \quad \forall j = 1, \ldots, k \tag{14}$$

for $L \in \mathbb{R}_{>0}$ independent of s, S, b. Two further facts, immediate from the proof of [22, Theorem 5.1], will be needed in Step 5: the loops γ_j lie in a fixed compact set $\mathcal{K}_1 \Subset B_0$ and satisfy $\ell_\rho(\gamma_j) \leq L_1$, with $\mathcal{K}_1$ and L_1 independent of s, S, P. Indeed, by [22, Lemmas 3.7 and 4.1] the γ_j are obtained by modifying, within ρ-distance $2a$, finitely many fixed loops of bounded ρ-length, while keeping ρ-distance $\geq 2a$ from ∂B_0; the length bound follows from [22, (5.2)–(5.3)]. Moreover, the base points b lie in the fixed compact set $B_\varepsilon \Subset B_0$ of *loc. cit.* We enlarge N once and for all so that $\mathcal{K}_1 \cup B_\varepsilon \subseteq N$.

There exists $c \in \mathbb{R}_{>0}$ such that, at the infinitesimal level,

$$\kappa_{(\mathcal{A} \smallsetminus \mathcal{D})|_{B_0}}(x, v) \geq c\, |v|_h \qquad \forall\, (x, v)\,. \tag{15}$$

This is the infinitesimal comparison corresponding to [22, Theorem 7.5] (stated there at the level of distances), a criterion of Brody–Green type; see also [14]. We give the short argument directly for our family, making explicit why a limiting entire curve must lie in one of the two hyperbolic strata of a fibre. If no such c exists then, by the definition of the Kobayashi–Royden metric (§2.1), there are holomorphic disks $g_\nu \colon \Delta(0,1) \to (\mathcal{A} \smallsetminus \mathcal{D})|_{B_0}$ with $|g'_\nu(0)|_h \to \infty$. Since $(\mathcal{A} \smallsetminus \mathcal{D})|_{B_0}$ is relatively compact in the complex manifold $\mathcal{A}|_{B \smallsetminus \mathrm{Sing}(f)}$ (its closure lies over $\overline{B_0}$, which is disjoint from $\mathrm{Sing}(f)$ by the reduction above), Brody's reparametrization lemma [20, Lemma 7.2.1] produces a non-constant entire curve $F \colon \mathbb{C} \to \mathcal{A}_U$ which is a locally uniform limit of holomorphic maps with values in $(\mathcal{A} \smallsetminus \mathcal{D})|_{B_0}$. In particular $f \circ F(\mathbb{C}) \subseteq \overline{B_0}$ omits the nonempty open set $\operatorname{int} \overline{\Delta}_1$, hence omits three points of B; so $f \circ F$ maps into the complement of three points of B, which is hyperbolic by uniformization, and $f \circ F$ is therefore constant. Hence $F(\mathbb{C}) \subseteq \mathcal{A}_b$ for some $b \in \overline{B_0}$, and $b \notin V \cup \mathrm{Sing}(f)$ by the reduction above. By Hurwitz's theorem in the form [20, Lemma 7.2.11], applied on an exhaustion of $\mathbb{C}$ by disks inside the complex manifold $\mathcal{A}|_{B \smallsetminus \mathrm{Sing}(f)}$, where

$\mathcal{D}$ restricts to a reduced divisor, either $F(\mathbb{C}) \subseteq \mathcal{D}_b$ or $F(\mathbb{C}) \subseteq \mathcal{A}_b \smallsetminus \mathcal{D}_b$; both are Kobayashi hyperbolic since $b \notin V$ (Step 1), so F is constant — a contradiction.

The divisor is locally Cartier on the smooth total space $\mathcal{A}_U$, so this use of Hurwitz is legitimate even if the compact model is singular.

Now notice that

$$\sigma_P(\gamma_j) \subset \sigma_P(B_0 \smallsetminus S) \subseteq (\mathcal{A} \smallsetminus \mathcal{D})|_{B_0 \smallsetminus S}$$

So by (14), (15), and the infinitesimal contraction property of the Kobayashi–Royden metric, integrated along γ_j, we deduce

$$\ell_h(\sigma_P(\gamma_j)) \leq c^{-1} L(s+1)\,. \tag{16}$$

Step 3: The Parshin cocycle. Since $f\colon \mathcal{A}_{B_0} \to B_0$ is a proper smooth submersion, Ehresmann's theorem gives a fiber bundle in the differentiable category. Let σ_O be the zero section and fix $w_0 := \sigma_O(b_0) \in \mathcal{A}_{b_0}$. Since B_0 is a $K(\pi,1)$-space (i.e. $\pi_i(B_0) = 0$ for $i \geq 2$, which holds because its universal cover is the disc $\mathbb{D}$), the long exact sequence of homotopy groups for the fibration $\mathcal{A}_{b_0} \to \mathcal{A}_{B_0} \to B_0$ yields the short exact sequence

$$1 \longrightarrow \pi_1(\mathcal{A}_{b_0}, w_0) \longrightarrow \pi_1(\mathcal{A}_{B_0}, w_0) \xrightarrow{f_*} \pi_1(B_0, b_0) \longrightarrow 1\,. \tag{17}$$

As $\mathcal{A}_{b_0}$ is a complex torus of dimension n, we have $\pi_1(\mathcal{A}_{b_0}, w_0) = H_1(\mathcal{A}_{b_0}, \mathbb{Z}) =: \Gamma \cong \mathbb{Z}^{2n}$. The zero section σ_O splits (17) via $i_O(\alpha) := [\sigma_O(\gamma)]$ where $\alpha = [\gamma] \in G := \pi_1(B_0, b_0)$. Each rational point $P \in A(K)$ with section $\sigma_P\colon B_0 \to \mathcal{A}_{B_0}$ gives a second splitting

$$i_P(\alpha) := [\eta^{-1}_{w_0, \sigma_P(b_0)} \circ \sigma_P(\gamma) \circ \eta_{w_0, \sigma_P(b_0)}]\,,$$

where $\eta_{w_0,\sigma_P(b_0)}$ is a path in the fiber $\mathcal{A}_{b_0}$ from w_0 to $\sigma_P(b_0)$. Since both i_P and i_O are sections of f_*, for every $\alpha \in G$:

$$f_*\big(i_P(\alpha) \cdot i_O(\alpha)^{-1}\big) = \alpha \cdot \alpha^{-1} = 1\,,$$

so $i_P(\alpha) \cdot i_O(\alpha)^{-1} \in \ker(f_*) = \Gamma$. We define the *Parshin cocycle*

$$c_P\colon G \to \Gamma\,, \quad c_P(\alpha) := i_P(\alpha) \cdot i_O(\alpha)^{-1}\,.$$

One verifies that c_P is a 1-cocycle of G with coefficients in the G-module Γ, where the G-action on Γ is by conjugation

$$\begin{aligned} \varphi\colon G &\to \mathrm{Aut}(\Gamma) \\ \alpha &\mapsto \varphi_\alpha\colon \gamma \mapsto i_O(\alpha)\,\gamma\, i_O(\alpha)^{-1} \end{aligned}$$

Note that the action φ doesn't depend on the choice of the splitting i_O, in fact for another splitting i_P, we have $i_P(\alpha) = c_P(\alpha) \cdot i_O(\alpha)$ with $c_P(\alpha) \in \Gamma$, and

$$i_P(\alpha)\,\gamma\, i_P(\alpha)^{-1} = c_P(\alpha) \cdot \varphi_\alpha(\gamma) \cdot c_P(\alpha)^{-1} = \varphi_\alpha(\gamma)\,,$$

where the last equality uses the commutativity of Γ. The cohomology class $[c_P] \in H^1(G, \Gamma)$ is independent of the choice of path $\eta_{w_0,\sigma_P(b_0)}$ (a different choice changes c_P by a coboundary). The following result, which goes back to Parshin [21] and is proved in [22, Proposition 7.2], will be crucial:

Proposition 3.1 (Parshin-Phung)**.** *The homomorphism*

$$\begin{aligned} \Psi\colon A(K) &\to H^1(G, \Gamma) \\ P &\mapsto [c_P] \end{aligned}$$

factors through $A(K)/T_A$ and has the following injectivity property: if $\Psi(P) = \Psi(Q)$, then $P - Q \in T_A + A(K)_{\mathrm{tors}}$. In particular, each element $[c] \in H^1(G,\Gamma)$ in the image of Ψ arises from at most $t_{\mathcal{A}} := \#\big(A(K)/T_A\big)_{\mathrm{tors}} < \infty$ rational points modulo the trace (finite by the Lang–Néron theorem).

Remark 3.2. The homomorphism Ψ is the connecting map of the relative exponential sequence [22, Equation (7.3)], so it is additive. Since G is free of finite rank, $H^1(G,\Gamma)$ is the cokernel of $\Gamma \to \Gamma^k$, $\lambda \mapsto (\varphi_{\alpha_j}(\lambda) - \lambda)_j$, and is finitely generated. The divisible group T_A therefore maps to zero. This cokernel need not be torsion-free, contrary to the assertion used in the proof of [22, Proposition 7.2].

For completeness, that assertion can be avoided as follows. If $\Psi(P) = 0$, then for every $m \geq 2$ its reduction modulo m is the zero Kummer class of $[m]^{-1}(\sigma_P)$ over B_0, using $\Gamma/m\Gamma \simeq \mathcal{A}[m]_{b_0}$ as monodromy modules; this follows by dividing local logarithms by m, as in [22, Proposition 7.1]. Thus this finite étale

torsor has an analytic section over B_0. The map $\pi_1(B_0) \to \pi_1(U)$ is surjective by van Kampen, since each removed disk contains at most one point of $B \setminus U$. Hence its monodromy fixes a sheet over U, giving a degree-one component and an algebraic section over U. Consequently $P \in mA(K)$ for every m, and its image in the free abelian group M/M_{tors} is zero. This gives the stated finite-kernel conclusion without a torsion-freeness assumption on $H^1(G,\Gamma)$. Here torsion classes in M lift to torsion points modulo T_A, because T_A is divisible.

Remark 3.3. It would be interesting to understand the relationship between the Parshin cocycle and the "analytic cocycle" induced by the logarithm of the section σ_P, as defined in [11, Equation (13)].

To extract a quantitative bound from the cocycle c_P, we work in the semidirect product coordinates provided by i_O. The splitting i_O exhibits

$$\pi_1(\mathcal{A}_{B_0}, w_0) \;\cong\; \Gamma \rtimes_{\varphi} G\,. \tag{18}$$

In these coordinates, $i_O(\alpha_j) = (0, \alpha_j)$ and $i_P(\alpha_j) = (\beta_j, \alpha_j)$ for some $\beta_j \in \Gamma$, so the Parshin cocycle reads $c_P(\alpha_j) = \beta_j$. Since G is free, the k-tuple $(\beta_1, \dots, \beta_k) \in \Gamma^k$ determines c_P completely. After the simultaneous change of splitting in Step 4, the displacement bound in Step 5 will give a norm bound on each β_j.

Step 4: Bounding loop lengths in $\mathcal{A}$. Let us now go back to the special loops (based in b) $\gamma_j \in \Gamma_S^{\alpha_j}$ for $j = 1, \dots, k$ considered in Step 2. Consider the following loop based at w_0:

$$\sigma_P(\gamma_j)^{\#} := (v_b \circ \sigma_O(c_{b_0,b}))^{-1} \circ \sigma_P(\gamma_j) \circ (v_b \circ \sigma_O(c_{b_0,b}))$$

where $v_b \subset \mathcal{A}_b$ is an intrinsic h-geodesic from $\sigma_O(b)$ to $\sigma_P(b)$. Since $\sigma_P(\gamma_j)$ is homotopic relative to its base point to $\sigma_P(c_{b_0,b}) \circ \sigma_P(\alpha_j) \circ \sigma_P(c_{b_0,b})^{-1}$ (with a loop representative of α_j understood), $\sigma_P(\gamma_j)^{\#}$ is freely homotopic to $\sigma_P(\alpha_j)$ and each of its components satisfies the following uniform bounds:

- the loop $\sigma_P(\gamma_j)$ satisfies:
$$\ell_h(\sigma_P(\gamma_j)) \le c^{-1} L(s+1)$$
by Equation (16);
- the fiber path $v_b \subset \mathcal{A}_b$ satisfies:
$$\ell_h(v_b) \le \delta_0 := \sup_{b \in \overline{B_0}} \mathrm{diam}_h(\mathcal{A}_b) < +\infty\,;$$
(here $\mathrm{diam}_h(\mathcal{A}_b)$ is the *intrinsic* h-diameter of the fibre, which bounds the length of a fibre geodesic; the supremum is finite because $\overline{B_0} \cap \mathrm{Sing}(f) = \varnothing$ by Step 1. On finitely many Ehresmann trivializations covering this compact base, the vertical metric is uniformly comparable with a fixed metric on the reference compact fibre, so its intrinsic diameter is uniformly bounded)
- since $N \Subset U$ is compact $\|d\sigma_O\|_\infty := \sup_{b \in N} \|d_b \sigma_O\|_{\rho,h} < +\infty$, so
$$\ell_h(\sigma_O(c_{b_0,b})) \le \|d\sigma_O\|_\infty \, \ell_\rho(c_{b_0,b}) \le \|d\sigma_O\|_\infty \, \delta_N =: \delta_0'$$

In other words for any element of the simple basis α_j we can construct a loop $\sigma_P(\gamma_j)^{\#}$ based at w_0 that is freely homotopic to $\sigma_P(\alpha_j)$ and satisfies

$$\ell_h\left(\sigma_P(\gamma_j)^{\#}\right) \le H(s) := c^{-1} L(s+1) + 2(\delta_0 + \delta_0')$$

where the constants L, δ_0, δ_0' don't depend on S and P.

The loop $\sigma_P(\gamma_j)^{\#}$ is based at w_0, so it defines an element $[\sigma_P(\gamma_j)^{\#}] \in \pi_1(\mathcal{A}_{B_0}, w_0)$. This element need not equal $i_P(\alpha_j) = (\beta_j, \alpha_j)$, because $\sigma_P(\gamma_j)^{\#}$ and the loop defining $i_P(\alpha_j)$ use different conjugating paths from w_0 to the fibre above b: the former uses $v_b \circ \sigma_O(c_{b_0,b})$, while the latter uses $\sigma_P(c_{b_0,b}) \circ \eta_{w_0,\sigma_P(b_0)}$. Both are paths from w_0 to $\sigma_P(b)$, so their concatenation

$$\omega \;:=\; \left(v_b \circ \sigma_O(c_{b_0,b})\right)^{-1} \circ\; \sigma_P(c_{b_0,b}) \circ\; \eta_{w_0,\sigma_P(b_0)}$$

is a loop at w_0 with $[\omega] \in \Gamma$: indeed $f_*[\omega] = 1$, since $f \circ v_b$ and $f \circ \eta_{w_0,\sigma_P(b_0)}$ are constant while $f \circ \sigma_O(c_{b_0,b}) = f \circ \sigma_P(c_{b_0,b}) = c_{b_0,b}$. A direct verification gives

$$[\sigma_P(\gamma_j)^{\#}] \;=\; [\omega]^{-1} \cdot i_P(\alpha_j) \cdot [\omega] \qquad \text{in } \pi_1(\mathcal{A}_{B_0}, w_0)\,.$$

The same ω is used for all j. Its homotopy class lies in the fibre subgroup, although the displayed loop need not lie inside a fibre. In other words, $[\sigma_P(\gamma_j)^{\#}]$ and $i_P(\alpha_j)$ are *conjugate* in $\pi_1(\mathcal{A}_{B_0}, w_0)$. The map $i'_P\colon G \to \pi_1(\mathcal{A}_{B_0}, w_0)$ defined by $i'_P(\alpha_j) := [\sigma_P(\gamma_j)^{\#}]$ is therefore a section of (17) (conjugation by $[\omega] \in \Gamma = \ker f_*$ preserves the section property), Γ-conjugate to i_P, so the associated cocycles are cohomologous: $[c'_P] = [c_P] \in H^1(G, \Gamma)$. By Proposition 3.1, we may freely replace i_P by i'_P for counting purposes. We do so and, to lighten notation, simply write β_j for the components of the new cocycle.

Step 5: Displacement bound via the universal cover. We now explain how the h-length bound on $\sigma_P(\gamma_j)^{\#}$ constrains the lattice element $\beta_j \in \Gamma$. Let $\pi\colon \widetilde{\mathcal{A}}_0 \to \mathcal{A}_{B_0}$ be the universal cover. Since $\mathcal{A}_{B_0} \to B_0$ is a fibre bundle (by Ehresmann's theorem) with fibre $\mathcal{A}_{b_0}$ (a complex torus of dimension n, so $\widetilde{\mathcal{A}_{b_0}} \cong \mathbb{R}^{2n}$) and base B_0 (a hyperbolic Riemann surface, so $\widetilde{B_0} \cong \mathbb{D}$, the unit disk), pulling back to $\mathbb{D}$ trivialises the bundle and gives $\widetilde{\mathcal{A}}_0 \cong \mathbb{R}^{2n} \times \mathbb{D}$. We use the trivialisation induced by the relative exponential (cf. [22, (7.3)]): the lattice local system $(R^1 f_* \mathbb{Z})^\vee$ is trivial over the contractible $\mathbb{D}$, and with this choice the subgroup $\Gamma \subset \pi_1(\mathcal{A}_{B_0}, w_0)$ acts on $\mathbb{R}^{2n} \times \mathbb{D}$ by $(\tilde{x}, \tilde{y}) \mapsto (\tilde{x} + \gamma, \tilde{y})$.

Fix the point $\tilde{w} = (\tilde{x}_0, \tilde{y}_0) \in \pi^{-1}(w_0)$ with $\tilde{x}_0 = 0$: this choice is possible because $w_0 = \sigma_O(b_0)$ lies on the zero section, whose lift in the chosen trivialisation is $\{0\} \times \mathbb{D}$, and $\tilde{y}_0$ is any point of $\mathbb{D}$ lying over b_0 for the universal covering $v\colon \mathbb{D} \to B_0$. Pull back the hermitian metric h to a Riemannian metric $\tilde{h}$ on $\widetilde{\mathcal{A}}_0$. The group of deck transformations of π is $\pi_1(\mathcal{A}_{B_0}, w_0) \cong \Gamma \rtimes_\varphi G$ (with the composition convention for π_1 already implicit in our loop concatenations). The action of an element $(\beta, \alpha) = \beta \cdot i_O(\alpha) \in \Gamma \rtimes_\varphi G$ on $\mathbb{R}^{2n} \times \mathbb{D}$ is given by (homotopy lifting property, cf. the proof of [22, Theorem A]):

$$(\beta, \alpha) \cdot (\tilde{x}, \tilde{y}) \;=\; \big(\varphi_\alpha(\tilde{x}) + \beta,\;\; \alpha \cdot \tilde{y}\big)\,, \tag{19}$$

where $\alpha \cdot \tilde{y}$ is the deck transformation of $v\colon \mathbb{D} \to B_0$ by $\alpha \in G = \pi_1(B_0, b_0)$ (acting on the base), and φ_α acts on the fibre factor $\mathbb{R}^{2n}$ as the linear extension of the monodromy: the fibre component "twists" as one transports along the base loop α, and the semidirect product structure of $\Gamma \rtimes_\varphi G$ encodes precisely this twisting. In particular the fibre component of the action is a translation only when $\varphi_\alpha = \mathrm{id}$; at the chosen point $\tilde{w} = (\tilde{x}_0, \tilde{y}_0)$, however, thanks to the normalisation $\tilde{x}_0 = 0$, (19) reads

$$(\beta, \alpha) \cdot (\tilde{x}_0, \tilde{y}_0) \;=\; \big(\beta \cdot \tilde{x}_0,\;\; \alpha \cdot \tilde{y}_0\big)\,, \tag{20}$$

where $\beta \cdot \tilde{x}_0$ is the deck transformation of $u\colon \mathbb{R}^{2n} \to \mathcal{A}_{b_0}$ by the element $\beta \in \Gamma$ (acting on the fibre).

The loop $\sigma_P(\gamma_j)^{\#}$ is based at w_0 and has homotopy class (β_j, α_j). Its lift to $\widetilde{\mathcal{A}}_0$ from $\tilde{w}$ is a path from $\tilde{w}$ to $(\beta_j, \alpha_j) \cdot \tilde{w}$. Since π is a local isometry for the metrics $\tilde{h}$ and h, the $\tilde{h}$-length of the lift equals $\ell_h(\sigma_P(\gamma_j)^{\#}) \le H(s)$, so:

$$d_{\tilde{h}}\big(\tilde{w},\; (\beta_j, \alpha_j) \cdot \tilde{w}\big) \;\le\; H(s) \;:=\; c^{-1} L(s+1) + 2(\delta_0 + \delta'_0) \;=\; O(s)\,, \tag{21}$$

where L, δ_0, δ'_0 are independent of S and P. Let d_j be the Γ-invariant geodesic metric on the fibre factor $\mathbb{R}^{2n}$ obtained by restricting $\tilde{h}$ to the slice $\mathbb{R}^{2n} \times \{\alpha_j \cdot \tilde{y}_0\}$; its quotient by the translation action of $\Gamma \cong \mathbb{Z}^{2n}$ is a compact real torus. Denote the Riemannian metric tensor on this slice by h_j, so that d_j is its length distance. Let $\mathrm{pr}_1\colon \mathbb{R}^{2n} \times \mathbb{D} \to \mathbb{R}^{2n}$ be the projection; by (20), $\mathrm{pr}_1(\tilde{w}) = \tilde{x}_0$ and $\mathrm{pr}_1((\beta_j, \alpha_j) \cdot \tilde{w}) = \beta_j \cdot \tilde{x}_0$.

The $\mathbb{D}$-component of the lift $\tilde{c}$ of $\sigma_P(\gamma_j)^{\#}$ from $\tilde{w}$ is the lift, from $\tilde{y}_0$, of the base loop $f\big(\sigma_P(\gamma_j)^{\#}\big) = c_{b_0,b}^{-1} \circ \gamma_j \circ c_{b_0,b}$, which by Step 2 has image in N and ρ-length at most $L_2 := L_1 + 2\delta_N$, uniformly in s, P. Hence this component is a path of $\tilde{\rho}$-length $\le L_2$ starting at $\tilde{y}_0$ and contained in $v^{-1}(N)$, so it remains in the closed ball

$$\Omega := \overline{B}_{v^{-1}(N)}(\tilde{y}_0, L_2)$$

of radius L_2 around $\tilde{y}_0$ in the connected component of $v^{-1}(N)$ containing $\tilde{y}_0$, for the induced length metric. This component, with the lifted metric $\tilde{\rho}$, is a covering of the compact length space N with the lifted length metric, hence complete and locally compact, hence a proper metric space by the Hopf–Rinow theorem for length spaces; therefore Ω is compact. Note that Ω is independent of s and P, and $\tilde{c} \subset \mathbb{R}^{2n} \times \Omega$.

The ratio $|\mathrm{pr}_{1*}\xi|_{h_j} / |\xi|_{\tilde{h}}$, for nonzero ambient tangent vectors ξ based at points of $\mathbb{R}^{2n} \times \Omega$, is invariant under Γ: every $\gamma \in \Gamma$ is a $\tilde{h}$-isometry and satisfies $\mathrm{pr}_1 \circ \gamma = \tau_\gamma \circ \mathrm{pr}_1$, with τ_γ the d_j-isometric translation

by γ. (The analogous equivariance fails for a general deck element (β, α), whose fibre component has linear part φ_α by (19); this is why we restrict to the Γ-invariant region $\mathbb{R}^{2n} \times \Omega$.) The ratio therefore descends to a continuous function on the unit sphere bundle over the compact set $(\mathbb{R}^{2n}/\Gamma) \times \Omega$, hence it is bounded by some $C_{h,j} \geq 1$; we set $C_h := \max_{1 \leq j \leq k} C_{h,j}$, which depends only on $\mathcal{A}$, h, Ω and the basis, uniformly in s, P. Integrating this differential bound along the actual lifted path $\tilde{c} \subset \mathbb{R}^{2n} \times \Omega$ shows that $\mathrm{pr}_1 \circ \tilde{c}$ has d_j-length $\leq C_h\, \ell_{\tilde{h}}(\tilde{c})$, and

$$d_j\big(\beta_j \cdot \tilde{x}_0,\ \tilde{x}_0\big) \ \leq\ C_h\, \ell_{\tilde{h}}(\tilde{c}) \ =\ C_h\, \ell_h\big(\sigma_P(\gamma_j)^{\#}\big) \ \leq\ C_h\, H(s)\,. \tag{22}$$

Remark 3.4. The intermediate bound (21) controls the displacement in the *full* universal cover $\widetilde{\mathcal{A}}_0 \cong \mathbb{R}^{2n} \times \mathbb{D}$, involving both fibre and base directions (and the metric $\tilde{h}$ need not be a product metric). The projection step (22) extracts a purely fibre-theoretic bound, which is what we need for the lattice count in the next step. Only lengths of paths staying in $\mathbb{R}^{2n} \times \Omega$ are compared; no claim about global Lipschitz continuity for ambient distances is needed.

Step 6: Counting. For each $j = 1, \ldots, k$, we count the elements $\beta \in \Gamma$ compatible with (22):

$$N_j(s) \ :=\ \#\Big\{\beta \in \Gamma :\ d_j\big(\beta \cdot \tilde{x}_0,\ \tilde{x}_0\big) \leq C_h H(s)\Big\}\,. \tag{23}$$

Via the universal covering $u\colon \mathbb{R}^{2n} \to \mathcal{A}_{b_0}$, the quotient torus $\mathbb{R}^{2n}/\Gamma$ is a compact geodesic Riemannian manifold with the metric induced by h_j, and $\Gamma \cong \mathbb{Z}^{2n}$ acts on $(\mathbb{R}^{2n}, d_j)$ by deck transformations. The set in (23) is in natural bijection with the set of Γ-translates of $\tilde{x}_0$ lying in the d_j-ball of radius $C_h H(s)$ around $\tilde{x}_0$.

By the fundamental lemma of the geometry of groups (cf. [22, Proposition A.8 and Lemma A.9]), since $\Gamma \cong \mathbb{Z}^{2n}$ is a lattice of rank $2n$ acting cocompactly on $(\mathbb{R}^{2n}, d_j)$, there exists $m_j > 0$ (depending on $\mathcal{A}$, h, α_j, but not on s) such that

$$N_j(s) \ \leq\ m_j\,(C_h H(s) + 1)^{2n}\,. \tag{24}$$

The exponent $2n$ reflects the rank of $\Gamma \cong \mathbb{Z}^{2n}$: in a ball of radius R in the universal cover of the $2n$-dimensional real torus $\mathcal{A}_{b_0}$, the number of Γ-translates of $\tilde{x}_0$ grows as R^{2n}.

Since G is free on $\alpha_1, \ldots, \alpha_k$, a cocycle representative is completely determined by the k-tuple $(\beta_1, \ldots, \beta_k) \in \Gamma^k$. Each class arising here has such a bounded representative; distinct tuples can represent the same class, which only decreases the number of classes to be counted. Bounding each component independently via (24) with $R = C_h H(s)$:

$$\#\{\text{possible cohomology classes } [c_P]\} \ \leq\ \prod_{j=1}^{k} m_j\,(C_h H(s) + 1)^{2n} \ =\ m_0\,(C_h H(s) + 1)^{2nk}\,, \tag{25}$$

where $m_0 := \prod_{j=1}^{k} m_j$.

By Proposition 3.1, each class $[c_P] \in H^1(G, \Gamma)$ arises from at most $t_{\mathcal{A}} := \#\big(A(K)/T_A\big)_{\mathrm{tors}} < \infty$ rational points modulo the trace (finite by the Lang–Néron theorem). Therefore:

$$\#\left(I(s, B_0)\big/_{T_A}\right) \ \leq\ t_{\mathcal{A}} \cdot m_0\,(C_h H(s) + 1)^{2nk} \ \leq\ m\,(s+1)^{2nk}\,,$$

where $m := t_{\mathcal{A}} \cdot m_0 \cdot \big(C_h(c^{-1}L + 2(\delta_0 + \delta_0')) + 1\big)^{2nk}$ depends only on B_0, $\mathcal{A}$, D, and h. □

3.2 Improved bound

We now prove a sharper version of Theorem 1.3. The general strategy is unchanged: we follow the same six-step scheme (reduction to the hyperbolic locus, loop construction, Parshin cocycle, lifting to $\mathcal{A}$, displacement in the universal cover, lattice counting). The only modification occurs in Steps 2 and 4, where we replace Phung's linear bound $\ell_S(\gamma_j) = O(s)$ with the sublinear bound $O(\sqrt{s+1})$ provided by Proposition 2.5. This propagates through the counting argument and halves the exponent.

Proof of (6). The proof follows the same six-step scheme as the proof of Theorem 1.3; in particular we may again assume $k \geq 1$ and $s \geq 1$, by the same reductions as there. The main modification is in Step 2: we replace the appeal to [22, Theorem 5.1], which provides loops of Kobayashi length $O(s)$, with Proposition 2.5, applied to the set $S := f(\sigma_P(B_0) \cap \mathcal{D}) \cap B_0$ (as in Step 2, we may assume $S \neq \emptyset$; note $\#S \leq s$, and since the bound below is increasing in the cardinality we may state it with s in place of $\#S$), which provides a point $b \in B_0 \setminus S$, a path σ from b_0 to b in B_0 with $\ell_\rho(\sigma) \leq \delta'$, and loops $\gamma_1, \dots, \gamma_k \subset B_0 \setminus S$ based at b with $[\sigma^{-1} \circ \gamma_j \circ \sigma] = \alpha_j$ in $\pi_1(B_0, b_0)$ and

$$\ell_S(\gamma_j) \;\leq\; C\sqrt{s+1} \qquad \text{for all } j = 1, \dots, k\,.$$

Since the loops share a common base point b and the conjugation uses the single path σ, Steps 3–5 apply verbatim with the conjugating path $v_b \circ \sigma_O(\sigma)$ in $\mathcal{A}_{B_0}$ (in place of $v_b \circ \sigma_O(c_{b_0,b})$ in Theorem 1.3), whose h-length is bounded by $\delta_0 + \|d\sigma_O\|_\infty \delta'$, independently of s. The compactness input of Step 5 is again available: by the last assertion of Proposition 2.5, the base loop $\sigma^{-1} \circ \gamma_j \circ \sigma$ has image in the fixed compact set $\mathcal{K}_0 \Subset B_0$ and ρ-length at most $L_0 + 2\delta'$, uniformly in s, P; enlarging N so that $\mathcal{K}_0 \subseteq N$ (and updating the fixed bound $\|d\sigma_O\|_\infty$ on N if necessary), the compact set Ω of Step 5 (with $L_2 := L_0 + 2\delta'$) is independent of s, P. The displacement bound becomes

$$d_j\big(\beta_j \cdot \tilde{x}_0,\ \tilde{x}_0\big) \;\leq\; C_h\, H(s)\,, \qquad H(s) \;:=\; c^{-1}C\sqrt{s+1} + 2(\delta_0 + \|d\sigma_O\|_\infty \delta')\,,$$

with $C_h \geq 1$ the constant of Step 5. The lattice counting in Step 6 then gives

$$\#\left(I(s, B_0)\big/_{T_A}\right) \;\leq\; t_{\mathcal{A}} m_0 \,(C_h H(s) + 1)^{2nk}\,.$$

It remains to extract the exponent. Since $H(s) = O(\sqrt{s+1})$, there is a constant $C' > 0$, independent of s, such that

$$(C_h H(s) + 1)^{2nk} \leq C'(s+1)^{nk}.$$

Absorbing all constants into m, we obtain

$$\#\left(I(s, B_0)\big/_{T_A}\right) \leq m(s+1)^{nk} \qquad \text{for all } s \geq 0.$$

The case $s = 0$ follows from $I(0, B_0) \subseteq I(1, B_0)$ after enlarging m, as in the proof of Phung's bound. There is no logarithmic or ε loss. This proves (6). □

3.3 Counting in the Mordell–Weil lattice and bounding heights

We retain the fixed reduced base and simple generators used above, with the improved loop bound of §3.2. All choices in this subsection are made once for this base and family, independently of s and P. The reductions that shrink B_0 cause no difficulty: the original set $I(s, B_0)$ is contained in the set for the reduced base, so every bound proved there applies to the original points as well.

The additional input retained from the six-step argument is a bound on a representative of the cohomology class of each point. We first explain how this controls the point in the Mordell–Weil lattice; this will give both the rank refinement and the height estimate.

Put $V_\Gamma = \Gamma \otimes \mathbb{R}$ and first assume $k \geq 1$. Since G is free on $\alpha_1, \dots, \alpha_k$, arbitrary values on these generators extend uniquely to a cocycle, so evaluation identifies

$$Z^1(G, V_\Gamma) \simeq V_\Gamma^k.$$

Under this identification the coboundaries form the linear subspace

$$W = \left\{ \big(\varphi_{\alpha_j}(v) - v\big)_{j=1}^k : v \in V_\Gamma \right\}.$$

Fix a Euclidean norm $|\cdot|$ on $V_\Gamma \simeq \mathbb{R}^{2n}$, and use the product norm $|(u_j)_j| = (\sum_j |u_j|^2)^{1/2}$ on V_Γ^k. Give

$$H^1(G, V_\Gamma) = V_\Gamma^k / W$$

the quotient norm $\|[u]\|_H = \inf_{w \in W} |u + w|$. This is a genuine norm because W is a closed linear subspace of a finite-dimensional vector space. In particular, the norm of a class is at most the norm of any representative. Moreover,

$$H^1(G, V_\Gamma) \simeq H^1(G, \Gamma) \otimes \mathbb{R}\,:$$

both sides are the cokernel of the map $V_\Gamma \to V_\Gamma^k$ displayed above, obtained by tensoring the integral two-term cochain complex with $\mathbb{R}$. Possible torsion in integral cohomology disappears in this operation; no torsion-freeness assertion is needed.

Proof of (7). The induced map

$$\Psi_{\mathbb{R}}\colon M\otimes\mathbb{R}\longrightarrow H^1(G,\Gamma)\otimes\mathbb{R}$$

is injective by Proposition 3.1: its integral kernel is finite, and tensoring with $\mathbb{R}$ kills that kernel and preserves exactness. Choose a basis $e_1,\dots,e_r$ of the free group M/M_{tors}, and use the Euclidean coordinate norm $\|\sum_i a_ie_i\| = (\sum_i a_i^2)^{1/2}$ on $M\otimes\mathbb{R}$. If $r>0$, the continuous function $x\mapsto\|\Psi_{\mathbb{R}}(x)\|_H$ has a positive minimum on the unit sphere, by injectivity and compactness. Its reciprocal gives a constant $C_2>0$ such that

$$\|x\|\le C_2\|\Psi_{\mathbb{R}}(x)\|_H\qquad(x\in M\otimes\mathbb{R}).$$

For $r=0$ the source is zero and the same inequality is automatic.

For the finitely many slice metrics h_j in Step 5, periodicity under Γ and compactness of the quotient give a constant C_7 such that $|\xi|\le C_7|\xi|_{h_j}$ for every tangent vector and every j. Integrating along any fibre path from 0 to β_j, and then taking the infimum of its h_j-length, gives

$$|\beta_j|\le C_7 d_j(\beta_j\cdot\tilde{x}_0,\tilde{x}_0)\le C_7C_hH(s).$$

Here $\tilde{x}_0=0$, as fixed in Step 5, and C_h is obtained as in that step for the fixed compact set Ω used in §3.2; it is independent of s and P. For $P\in I(s,B_0)$, the tuple $(\beta_1,\dots,\beta_k)$ is the representative c'_P of Steps 4–5. The single conjugation used there ensures that $[c'_P]=\Psi(P)$; thus its real class is precisely $\Psi_{\mathbb{R}}(\overline{P}\otimes 1)$, where $\overline{P}$ denotes the image in M. Consequently

$$\|\Psi_{\mathbb{R}}(\overline{P}\otimes1)\|_H\le\left(\sum_{j=1}^{k}|\beta_j|^2\right)^{1/2}\le\sqrt{k}\,C_7C_hH(s).$$

Combining the two norm comparisons, with $C_3=C_2\sqrt{k}\,C_7C_h$, yields

$$\|\overline{P}\otimes1\|\le C_3H(s)\le C_4\sqrt{s+1},\tag{26}$$

where the fixed additive term in $H(s)$ is absorbed using $\sqrt{s+1}\ge1$.

The image of M/M_{tors} is a lattice of rank r in $M\otimes\mathbb{R}$. In the chosen basis its coordinates are integers. Equation (26) bounds each coordinate in absolute value by $C_4\sqrt{s+1}$, so the number of possible free classes is at most

$$\bigl(2\lfloor C_4\sqrt{s+1}\rfloor+1\bigr)^r=O((s+1)^{r/2}).$$

Each free class has exactly $\#M_{\mathrm{tors}}$ lifts to M; counting only those represented by points of $I(s,B_0)$ can decrease this number. Multiplying by this fixed finite factor proves (7), without assuming that $I(s,B_0)$ is invariant under trace translations. For $r=0$ the displayed count is 1. For $k=0$, the rank-zero reduction above gives M finite and the same conclusion directly. The case $s=0$ follows from $I(0,B_0)\subseteq I(1,B_0)$, after enlarging the constants as before, including C_4 in (26). Finally, injectivity gives

$$r\le\dim_{\mathbb{R}}H^1(G,V_\Gamma)\le\dim_{\mathbb{R}}V_\Gamma^k=2nk,$$

which explains the relation between the exponents $r/2$ and nk. □

Proof of Corollary C. The canonical height of a symmetric ample line bundle over a complex function field is invariant under translation by the constant trace and vanishes on torsion; it descends to a quadratic form on $M\otimes\mathbb{R}$; see [5, §9, especially Lemma 9.13]. In the basis chosen above, this form is given by a fixed symmetric real matrix Q: if $\overline{P}$ has free coordinates $(a_1,\dots,a_r)$, then $\widehat{h}_L(P)=a^{\mathsf{T}}Qa$. Taking an upper bound for this quadratic form on the Euclidean unit sphere gives $C_5>0$ with

$$\widehat{h}_L(P)\le C_5\|\overline{P}\otimes1\|^2.$$

Equation (26) therefore proves

$$\widehat{h}_L(P)\le C_5C_4^2(s+1)\qquad(P\in I(s,B_0)).$$

Thus we may take $C=C_5C_4^2$. This uses the bound on the coordinates of each point, rather than the cardinality estimate: a bound on the number of lattice points by itself would give no bound on their

distance from the origin. If $r = 0$ (in particular, if $k = 0$), the height is identically zero and the conclusion holds directly with any $C > 0$.

Suppose now that $\mathcal{D}$ is Cartier. The function

$$h_{\mathcal{D}}(P) = \deg_B \sigma_P^* \mathcal{O}_{\mathcal{A}}(\mathcal{D})$$

is a Weil height attached to $\mathcal{O}_A(D)$ for this fixed model. Choose an integer $a > 0$ such that $E := L^{\otimes a} \otimes \mathcal{O}_A(-D)$ is ample. An ample Weil height h_E over K is bounded below. The additivity of Weil heights and the bounded difference between a Weil height h_L for L and $\widehat{h}_L$ give

$$h_E(P) = a h_L(P) - h_{\mathcal{D}}(P) + O(1), \qquad h_L(P) = \widehat{h}_L(P) + O(1),$$

where the bounded terms are uniform in P. Rearranging, we obtain

$$h_{\mathcal{D}}(P) \le a\widehat{h}_L(P) + C_6$$

with C_6 depending on the fixed height data and model; compare the argument of [22, Corollary 1.3]. For $P \in I(s, B_0)$ the section is not contained in $\mathcal{D}$, and this line-bundle degree equals the degree of the pulled-back effective divisor. Combining the inequalities and using $s + 1 \ge 1$, we conclude that

$$\deg_B \sigma_P^* \mathcal{D} \le aC(s+1) + C_6 \le \big(aC + \max\{C_6, 0\}\big)(s+1).$$

The same comparison covers the rank-zero case. □

Remark 3.5. The quotient norms and the inverse bound for $\Psi_{\mathbb{R}}$ are chosen for the fixed local system and base. Their constants need not be uniform when the family, model or base changes. Thus the linear height bound above is a fixed-data result under (P)(iii), even though it applies to intersections allowed to be unrestricted over the removed disks.

References

[1] D. Borthwick, *Spectral theory of infinite-area hyperbolic surfaces*, second edition, Progress in Mathematics, vol. 318, Birkhäuser/Springer, 2016.

[2] A. Buium, The *abc* theorem for abelian varieties, *Internat. Math. Res. Notices* **1994** (1994), no. 5, 219–233.

[3] A. Buium, Intersection multiplicities on abelian varieties, *Math. Ann.* **310** (1998), 653–659.

[4] L. Capuano and A. Turchet, Lang–Vojta conjecture over function fields for surfaces dominating $\mathbb{G}_m^2$, *Eur. J. Math.* **8** (2022), 573–610.

[5] B. Conrad, Chow's K/k-image and K/k-trace, and the Lang–Néron theorem, *Enseign. Math. (2)* **52** (2006), no. 1–2, 37–108.

[6] B. Conrad, *Abelian varieties*, lecture notes by T. Feng, Stanford University, Spring 2015, https://math.berkeley.edu/~fengt/249C.pdf.

[7] P. Corvaja and U. Zannier, Some cases of Vojta's conjecture on integral points over function fields, *J. Algebraic Geom.* **17** (2008), no. 2, 295–333; addendum, *Asian J. Math.* **14** (2010), no. 4, 581–584.

[8] P. Corvaja and U. Zannier, An *abcd* theorem over function fields and applications, *Bull. Soc. Math. France* **139** (2011), no. 4, 437–454.

[9] P. Corvaja and U. Zannier, Algebraic hyperbolicity of ramified covers of $\mathbb{G}_m^2$ (and integral points on affine subsets of $\mathbb{P}^2$), *J. Differential Geom.* **93** (2013), no. 3, 355–377.

[10] J.-P. Demailly, Algebraic criteria for Kobayashi hyperbolic projective varieties and jet differentials, in *Algebraic geometry—Santa Cruz 1995*, Proceedings of Symposia in Pure Mathematics, vol. 62, Part 2, American Mathematical Society, Providence, RI, 1997, pp. 285–360.

[11] P. Dolce and F. Tropeano, Relative monodromy of ramified sections on abelian schemes, *Internat. Math. Res. Notices* **2026** (2026), no. 4, article rnag025.

[12] B. Farb and D. Margalit, *A primer on mapping class groups*, Princeton Mathematical Series, vol. 49, Princeton University Press, Princeton, NJ, 2012.

[13] G. Gao, GCD and ABC inequalities for abelian varieties over function fields, preprint (2026), arXiv:2610.20379, https://arxiv.org/abs/2610.20379.

[14] M. L. Green, Holomorphic maps to complex tori, *Amer. J. Math.* **100** (1978), no. 3, 615–620.

[15] J. Guo, K. D. Nguyen, C.-L. Sun, and J. T.-Y. Wang, Vojta's *abc* conjecture for algebraic tori and applications over function fields, *Adv. Math.* **476** (2025), article 110358.

[16] M. Hindry and J. H. Silverman, The canonical height and integral points on elliptic curves, *Invent. Math.* **93** (1988), no. 2, 419–450.

[17] J. M. Lee, *Introduction to Riemannian manifolds*, Graduate Texts in Mathematics, vol. 176, Springer, Cham, 2018.

[18] E. H. Lieb and M. Loss, *Analysis*, second edition, Graduate Studies in Mathematics, vol. 14, American Mathematical Society, Providence, RI, 2001.

[19] J. Noguchi and J. Winkelmann, Bounds for curves in abelian varieties, *J. Reine Angew. Math.* **572** (2004), 27–47.

[20] J. Noguchi and J. Winkelmann, *Nevanlinna theory in several complex variables and Diophantine approximation*, Grundlehren der mathematischen Wissenschaften, vol. 350, Springer, 2014.

[21] A. N. Parshin, Finiteness theorems and hyperbolic manifolds, in *The Grothendieck Festschrift, Vol. III*, Progress in Mathematics, vol. 88, Birkhäuser Boston, Boston, MA, 1990, pp. 163–178.

[22] X. K. Phung, Generalized integral points on abelian varieties and the geometric Lang–Vojta conjecture, *J. Algebraic Geom.* **34** (2025), no. 3, 407–446.

[23] H. L. Royden, Remarks on the Kobayashi metric, in *Several complex variables, II (Proc. Internat. Conf., Univ. Maryland, College Park, Md., 1970)*, Lecture Notes in Mathematics, vol. 185, Springer, Berlin–New York, 1971, pp. 125–137.

[24] T. Shioda, On elliptic modular surfaces, *J. Math. Soc. Japan* **24** (1972), 20–59.

[25] A. Turchet, Fibered threefolds and Lang–Vojta's conjecture over function fields, *Trans. Amer. Math. Soc.* **369** (2017), no. 12, 8537–8558.

[26] H. Zieschang, E. Vogt, and H.-D. Coldewey, *Surfaces and planar discontinuous groups*, translated from the German by J. Stillwell, Lecture Notes in Mathematics, vol. 835, Springer, Berlin, 1980.

P. Dolce, Institute for Theoretical Sciences, Westlake University, China
E-mail address: dolce@westlake.edu.cn