Math. Log. Quart. **65,** No. 2, 237–250 (2019) / **DOI** 10.1002/malq.201800068

Parametric Presburger arithmetic: complexity of counting and quantifier elimination

Tristram Bogart[1], **John Goodrick**[1]*, **Danny Nguyen**[2], and **Kevin Woods**[3]

[1] Departamento de Matemáticas, Universidad de Los Andes, Carrera 1 No. 18a 10, 111711 Bogotá, Colombia
[2] Department of Mathematics, University of Michigan, Ann Arbor, 530 Church Street, Ann Arbor, MI 48109-1043, United States of America
[3] Department of Mathematics, Oberlin College, 10 N. Professor St., Oberlin, OH 44074, United States of America

Received 16 October 2018, accepted 16 May 2019
Published online 2 September 2019

We consider an expansion of Presburger arithmetic which allows multiplication by k parameters $t_1, \ldots, t_k$. A formula in this language defines a parametric set $S_{\mathbf{t}} \subseteq \mathbb{Z}^d$ as $\mathbf{t}$ varies in $\mathbb{Z}^k$, and we examine the counting function $|S_{\mathbf{t}}|$ as a function of $\mathbf{t}$. For a single parameter, it is known that $|S_t|$ can be expressed as an eventual quasi-polynomial (there is a period m such that, for sufficiently large t, the function is polynomial on each of the residue classes mod m). We show that such a nice expression is impossible with 2 or more parameters. Indeed (assuming $\mathbf{P} \neq \mathbf{NP}$) we construct a parametric set S_{t_1,t_2} such that $|S_{t_1,t_2}|$ is not even polynomial-time computable on input (t_1, t_2). In contrast, for parametric sets $S_{\mathbf{t}} \subseteq \mathbb{Z}^d$ with arbitrarily many parameters, defined in a similar language without the ordering relation, we show that $|S_{\mathbf{t}}|$ is always polynomial-time computable in the size of $\mathbf{t}$, and in fact can be represented using the gcd and similar functions.

1 Introduction

We study the difficulty of counting points in parametric sets of the form

$$S_{\mathbf{t}} = \{\mathbf{x} \in \mathbb{Z}^d \ : \ Q_1 y_1 \ Q_2 y_2 \ \ldots Q_m y_m \ \ \Theta_{\mathbf{t}}(\mathbf{x}, \mathbf{y})\}. \tag{1}$$

Here $\mathbf{t} = (t_1, \ldots, t_k)$ are the *parameters*, $\mathbf{x} = (x_1, \ldots, x_d)$ are the *free* variables, and $\mathbf{y} = (y_1, \ldots, y_m)$ are the *quantified* variables, all ranging over $\mathbb{Z}$; $Q_i \in \{\forall, \exists\}$ are the quantifiers; and $\Theta_{\mathbf{t}}(\mathbf{x}, \mathbf{y})$ is a Boolean combination, in disjunctive normal form, of linear inequalities in $\mathbf{x}, \mathbf{y}$ with coefficients in $\mathbb{Z}[\mathbf{t}]$. That is,

$$\Theta_{\mathbf{t}}(\mathbf{x}, \mathbf{y}) = \left[A_1(\mathbf{t}) \cdot (\mathbf{x}, \mathbf{y})^T \ \leq \ \overline{b}_1(\mathbf{t})\right] \vee \ldots \vee \left[A_\ell(\mathbf{t}) \cdot (\mathbf{x}, \mathbf{y})^T \leq \overline{b}_\ell(\mathbf{t})\right], \tag{2}$$

where each $A_i(\mathbf{t})$ is a $r_i \times (d+m)$ matrix, each $\overline{b}_i(\mathbf{t})$ is a length r_i column vector, all with entries in $\mathbb{Z}[\mathbf{t}]$, and the concatenation $(\mathbf{x}, \mathbf{y})$ of the $\mathbf{x}$ and $\mathbf{y}$ variables is treated as a row vector.[1] If there are k parameters $t_1, \ldots, t_k$, we say that the family of sets $\{S_{\mathbf{t}} : \mathbf{t} \in \mathbb{Z}^k\}$ is a *k-parametric Presburger family*. A general expression of the type

$$\Phi_{\mathbf{t}}(\mathbf{x}) = Q_1 y_1 \ Q_2 y_2 \ \ldots Q_m y_m \ \ \Theta_{\mathbf{t}}(\mathbf{x}, \mathbf{y})$$

with $\Theta_{\mathbf{t}}(\mathbf{x}, \mathbf{y})$ as in (1) is called a *formula in k-parametric Presburger Arithmetic* (often abbreviated as k-parametric PA). Classic Presburger arithmetic corresponds to $k = 0$.

Question 1.1 Given a k-parametric Presburger family defined by $S_{\mathbf{t}} = \{\mathbf{x} \in \mathbb{Z}^d : \Phi_t(\mathbf{x})\}$, under what conditions on the formula $\Phi_{\mathbf{t}}$ is the counting function $|S_{\mathbf{t}}|$ a "nice" function of $\mathbf{t}$?

Of course, "nice" is a vague qualifier, so let us start with some nice examples. We shall assume that the parameters t_i are nonnegative in the following examples, which simplifies the number of cases:

* Corresponding author; e-mail: jr.goodrick427@uniandes.edu.co

[1] By a simple trick, we do not need to worry about negations $\neg(\lambda_1 x_1 + \cdots + \lambda_{d+m} y_m \leq c)$ of basic inequalities, since these are equivalent to strict inequalities "$\ldots > c$," which in turn are equivalent to non-strict inequalities "$\ldots \geq c + 1$" since we are working over the integers.

Example 1.2 If we define $S_{t_1,t_2} = \{x \in \mathbb{Z} : \ x \geq 0 \ \wedge \ t_1 x \leq t_2\}$, then $|S_{t_1,t_2}| = \lfloor t_2/t_1 \rfloor + 1$.

Example 1.3 The set $S_{t_1,t_2} = \{(x_1, x_2) \in \mathbb{Z}^2 : \ x_1, x_2 \geq 0 \ \wedge \ t_1 x_1 + t_2 x_2 = t_1 t_2\}$ consists of the integer points on a line segment with endpoints $(t_2, 0)$ and $(0, t_1)$, and so $|S_{t_1,t_2}| = \gcd(t_1, t_2) + 1$.

Example 1.4 If $S_{t_1,t_2} = \{(x_1, x_2) \in \mathbb{Z}^2 : \ x_1, x_2 \geq 0 \ \wedge \ x_1 + x_2 = t_1 \ \wedge \ 2x_1 + x_2 \leq t_2\}$, then the equality forces $x_2 = t_1 - x_1$ (which is only valid if $x_1 \leq t_1$) and substituting into the inequality shows that

$$|S_{t_1,t_2}| = |\{x_1 \in \mathbb{Z} : \ 0 \leq x_1 \leq \min(t_1, t_2 - t_1)\}|$$

$$= \begin{cases} t_1 + 1 & \text{if } 2t_1 \leq t_2, \\ t_2 - t_1 + 1 & \text{if } t_1 \leq t_2 < 2t_1, \\ 0 & \text{if } t_2 < t_1. \end{cases}$$

Example 1.5 If $S_t = \{x \in \mathbb{Z} : \ \exists y \in \mathbb{Z}, \ x, y \geq 0 \ \wedge \ 2x + 2y + 2 = t\}$, then

$$|S_t| = \begin{cases} t/2 & \text{if } t \text{ even}, t \geq 2, \\ 0 & \text{if } t \text{ odd}. \end{cases}$$

We are seeing many types of "nice" functions in these examples, and the question is now how to generalize. In fact, Example 1.5 generalizes to any family in 1-parametic Presburger arithmetic [3], as described in the next section.

1.1 1-parametric Presburger arithmetic

In the case of a single parameter t, our perspective means studying families $\{S_t \ : \ t \in \mathbb{Z}\}$ of subsets of $\mathbb{Z}^d$ of the form

$$S_t = \{\mathbf{x} \in \mathbb{Z}^d \ \ : \ \ Q_1 y_1 \ Q_2 y_2 \ \ldots Q_m y_m \ \ \Theta_t(\mathbf{x}, \mathbf{y})\},$$

where $\Theta_t(\mathbf{x}, \mathbf{y})$ is exactly as in (2) except that the entries of the A_is and the $\overline{b}_i$s come from the univariate polynomial ring $\mathbb{Z}[t]$. The study of such *1-parametric PA families* was proposed by Woods in [14]. These families were further analyzed in [3], in which the main result is that they exhibit *quasi-polynomial* behavior:

A function $g : \mathbb{Z} \to \mathbb{Z}$ is a *quasi-polynomial* if there exists a period m and polynomials $f_0, \ldots, f_{m-1} \in \mathbb{Q}[t]$ such that $g(t) = f_i(t)$, for $t \equiv i \bmod m$. A function $g : \mathbb{Z} \to \mathbb{Z}$ is an *eventual quasi-polynomial*, abbreviated *EQP*, if it agrees with a quasi-polynomial for sufficiently large $|t|$. Example 1.5 is a family where $|S_t|$ is an EQP.

Theorem 1.6 (Bogart, Goodrick, & Woods; [3]) *Let $\{S_t : t \in \mathbb{Z}\}$ be a 1-parametric PA family. There exists an EQP $g : \mathbb{Z} \to \mathbb{N}$ such that, if S_t has finite cardinality, then $g(t) = |S_t|$. The set of t such that S_t has finite cardinality is eventually periodic.*

In [3], the parameter t takes values in $\mathbb{N}$ instead of $\mathbb{Z}$. However, one can see that the same proofs and conclusions also hold when t ranges over $\mathbb{Z}$.

There are several other forms of quasi-polynomial behavior that 1-parametric PA families exhibit (such as possessing EQP Skolem functions; cf. [3]). Here we focus on the cardinality, $|S_t|$. We hope the reader agrees that EQPs are relatively "nice" functions.

1.2 k-parametric Presburger arithmetic

Let us restate our main definition:

Definition 1.7 A *k-parametric PA family* is a collection $\{S_{\mathbf{t}} : \mathbf{t} = (t_1, \ldots, t_k) \in \mathbb{Z}^k\}$ of subsets of $\mathbb{Z}^d$ of the form

$$S_{\mathbf{t}} = \{\mathbf{x} \in \mathbb{Z}^d \ \ : \ \ Q_1 y_1 \ Q_2 y_2 \ \ldots Q_m y_m \ \ \Theta_{\mathbf{t}}(\mathbf{x}, \mathbf{y})\}, \tag{3}$$

where now $\Theta_{\mathbf{t}}(\mathbf{x}, \mathbf{y})$ is a Boolean combination of linear inequalities with coefficients in $\mathbb{Z}[\mathbf{t}]$.

A *k-parametric PA formula* $\Phi_{\mathbf{t}}$ is an expression "$Q_1 y_1\ Q_2 y_2\ \ldots\ Q_m y_m\ \Theta_{\mathbf{t}}(\mathbf{x}, \mathbf{y})$" as above, or any logically equivalent first-order formula in the language $\mathcal{L} = \{+, 0, 1, \leq, \lambda_p(\mathbf{t}) : p \in \mathbb{Z}[\mathbf{t}]\}$ with a function symbols for $+$, unary function symbols $\lambda_p(\mathbf{t})$ for multiplication by each polynomial $p(\mathbf{t}) \in \mathbb{Z}[\mathbf{t}]$, constant symbols for 0 and 1, and a relation symbol for $\leq$.

Abusing the notation, we also denote the parametric family $\{S_{\mathbf{t}} : \mathbf{t} \in \mathbb{Z}^k\}$ just by $S_{\mathbf{t}}$ when the dimension k is clear.

Examples 1.2, 1.3, & 1.4 show that k-parametric PA families, with $k \geq 2$, *can* have nice counting functions, $|S_{\mathbf{t}}|$. Will they always? We despair of defining "nice" precisely, but we can at least provide a necessary condition: for a fixed family $S_{\mathbf{t}}$, if $|S_{\mathbf{t}}|$ is to qualify as a nice function, there must at least be a polynomial-time algorithm that takes as input $\mathbf{t} \in \mathbb{Z}^k$ and outputs $|S_{\mathbf{t}}|$.

Question 1.8 Given a k-parametric Presburger family defined by $S_{\mathbf{t}} = \{\mathbf{x} \in \mathbb{Z}^d : \Phi_{\mathbf{t}}(\mathbf{x})\}$, under what conditions on the (fixed) formula $\Phi_{\mathbf{t}}$ is the counting function $|S_{\mathbf{t}}|$ polynomial-time computable, taking as input the values of the parameters $\mathbf{t}$?

Note that we define polynomial-time computation in the usual computer-science sense: the number of steps of the algorithm must be polynomial in the *input size* of $\mathbf{t}$ (that is, the number of bits to encode $\mathbf{t}$ into binary), which is $k + \sum_i \log_2 |t_i|$. E.g., the Euclidean algorithm is polynomial-time: it computes $\gcd(t_1, t_2)$ in number of arithmetic operations bounded by a degree 1 polynomial in $2 + \log_2 t_1 + \log_2 t_2$.

The functions $|S_{\mathbf{t}}|$ from Examples 1.2 through 1.5 are all polynomial-time computable. From Theorem 1.6 and the observation that EQPs are polynomial-time computable, we immediately obtain an answer to Question 1.8 in the case of a single parameter t:

Corollary 1.9 *Let S_t be any fixed 1-parametric PA family. Then there are polynomial time algorithms to:*

(i) check if $|S_t| = \infty$,
(ii) compute $|S_t|$ if $|S_t| < \infty$.

The main goal of this paper is to construct a fixed 2-parametric PA family $\{S_{t_1,t_2} : (t_1, t_2) \in \mathbb{Z}^2\}$ for which there is no polynomial-time algorithm computing $|S_{\mathbf{t}}|$ (assuming $\mathbf{P} \neq \mathbf{NP}$). Therefore, while we cannot say with precision what a nice function should be like, we can say that this particular counting function $|S_{\mathbf{t}}|$ is not nice. Furthermore, this implies that certain classes of functions (polynomials, gcds, floor functions, modular reductions, ...) are not expressive enough to capture $|S_{\mathbf{t}}|$, even for a very simple-looking $S_{\mathbf{t}}$. This contrasts with the 1-parameter case, where $|S_t|$ is always an EQP and hence polynomial-time computable.

Definition 1.7 is a generalization of *classical Presburger arithmetic* (PA), in which a formula Φ is given only with explicit integer coefficients and constants (A_i and $\overline{b}_i$) without any parameters $\mathbf{t}$. PA is *decidable*, meaning there is an algorithm to decide the truth of any given well-formed sentence in it. Moreover, PA has full *quantifier elimination* in an expanded language with predicates for divisibility by each fixed integer. This important logical fact permits an algorithm to actually count the cardinality of any set definable by a PA formula Φ with an arbitrary number of quantifiers and inequalities, although with an unpractical triply exponential complexity in the length of Φ (cf. [11]). The complexity of PA is itself a fundamental topic in the study of decidable logical theories and their complexities (cf. [6, 8]).

Returning to k-parametric PA, for a fixed formula $\Phi_{\mathbf{t}}$, given any value $\mathbf{a} \in \mathbb{Z}^k$ for $\mathbf{t}$, we can substitute it into $\Phi_{\mathbf{t}}$ to get a formula $\Phi_{\mathbf{a}}$ in PA. By the above paragraph, the parametric counting problem for (1) is always computable. Moreover, the form of the resulting formula $\Phi_{\mathbf{a}}$, especially its number of quantifiers and inequalities, stays the same for different values $\mathbf{a}$ of $\mathbf{t}$. So we can hope that the complexity of computing $|S_{\mathbf{t}}|$ (for a fixed family $S_{\mathbf{t}}$) is much lower than that of counting solutions to a general PA formula (when the formula is not fixed, but instead given as input to the algorithm). To reiterate, it is critical in our analysis that the formula $\Phi_{\mathbf{t}}$ be fixed throughout, and we look for an efficient algorithm with $\mathbf{t}$ as the only input.

1.3 Summary of results

Our main result is that if $\mathbf{P} \neq \mathbf{NP}$ (technically, we only need the weaker assumption that $\#\mathsf{P} \neq \mathsf{FP}$), then there exists a 2-parametric PA family $S_{\mathbf{t}}$ such that $|S_{\mathbf{t}}|$ is not polynomial-time computable; in fact, such a family exists with limited alternation of quantifiers. First we recall the Σ_n and Π_n hierarchies of first-order formulas based on the number of quantifier alternations.

A k-parametric PA formula $\Phi_{\mathbf{t}}(\mathbf{x})$ is in Σ_1 (Π_1) if it is logically equivalent to one of the form

$$Q_1 y_1 \ Q_2 y_2 \ \ldots Q_m y_m \ \ \Theta_{\mathbf{t}}(\mathbf{x}, \mathbf{y})$$

in which every quantifier Q_i is $\exists$ (every Q_i is $\forall$), and $\Theta_{\mathbf{t}}(\mathbf{x}, \mathbf{y})$ is a Boolean combination of linear inequalities with coefficients in $\mathbb{Z}[\mathbf{t}]$. Inductively, a k-parametric PA formula $\Phi_{\mathbf{t}}(\mathbf{x})$ is in Σ_{n+1} (Π_{n+1}) if it is equivalent to one of the form

$$Q_1 y_1 \ Q_2 y_2 \ \ldots Q_m y_m \ \ \Phi'_{\mathbf{t}}(\mathbf{x}, \mathbf{y})$$

in which every Q_i is $\exists$ ($\forall$) and $\Phi'_{\mathbf{t}}(\mathbf{x}, \mathbf{y})$ is a formula in Π_n (Σ_n).

Theorem 1.10 *Assume* $\mathbf{P} \neq \mathbf{NP}$. *There exists a 2-parametric* Σ_2 *PA family* S_{t_1,t_2} *for which* $|S_{t_1,t_2}|$ *is always finite but cannot be expressed as a polynomial time evaluable function in* t_1 *and* t_2.

Two corollaries are:

Corollary 1.11 *There is a 2-parametric family* S_{t_1,t_2} *such that the set of* $(t_1, t_2) \in \mathbb{Z}^2$ *for which* $|S_{t_1,t_2}|$ *is positive cannot be described using polynomial-time relations in* t_1, t_2.

Corollary 1.12 *Any extension of 2-parametric PA with only polynomial-time computable predicates cannot have full quantifier elimination.*

1.4 Structure of the rest of the paper

We shall present what amount to two different proofs of Theorem 1.10 in the following two sections. In each case, we leverage the main result of Nguyen and Pak [10] which yields a 3-parametric Σ_2 PA formula, and then show how this can be reduced to a 2-parametric Σ_2 PA formula whose points are equally "hard" to count (modulo polynomial-time reductions). The first reduction we present, in § 2, uses a trick due to Glivický and Pudlák [7] to encode multiplication by three different integers using multiplication by only two integers, and this reduction has the advantage of not increasing the number of free variables in the formula. Next, in § 3 we present a more general counting-reduction technique which is less *ad hoc* and reduces any k-parametric PA formula to a 2-parametric PA formula with the same number of quantifier alternations; the idea here is a little more transparent than in § 2, but it has the disadvantage of introducing many more new free and quantified variables to the formula, so we consider that it is interesting to present both reductions.

In § 4 we consider a variant of Question 1.8 in which there is no order relation in our language; that is, we can only express linear equations but not linear inequalities. Quantifier-free formulas in this language define finite unions of *lattice translates*. This setting was studied in detail from a model-theoretic perspective by van den Dries and Holly [13], and we apply their results to show that, in contrast to Theorem 1.10, the counting functions in the unordered setting can be computed in polynomial time, regardless of the number of parameters and of quantifier alternations. Indeed, these functions can be expressed using gcd and related functions.

Finally, in § 5 we discuss the optimality of Theorem 1.10 by explaining what happens when we weaken or modify some of the hypotheses.

2 Proof of Theorem 1.10 and its corollaries

In what follows, it will be convenient to allow k-parametric PA formulas in which the quantifiers are not necessarily outside the scope of all Boolean operations, but these are always logically equivalent to expressions as in (3); e.g.,

$$\exists y_1[\Theta_t(\mathbf{x}, y_1)] \wedge \exists y_1[\Theta'_{\mathbf{t}}(\mathbf{x}, y_1)]$$

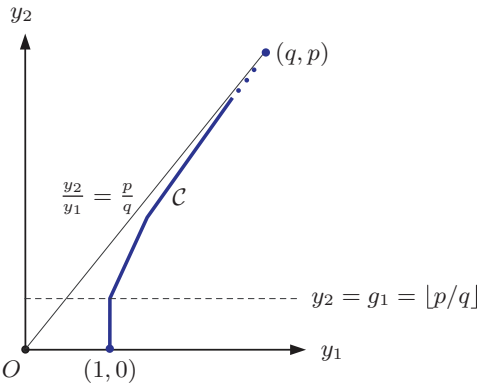

Fig. 1 The (bold) sail $\mathcal{C}$ below the line $y_2/y_1 = p/q$.

is equivalent to

$$\exists y_1\, \exists y_2\, [\Theta_{\mathbf{t}}(\mathbf{x}, y_1) \wedge \Theta'_{\mathbf{t}}(\mathbf{x}, y_2)].$$

In [10], certain subclasses of classical PA formulas, called *short PA formulas*, were investigated. The PA formulas in each such subclass are allowed to have only a bounded number of variables, quantifiers and inequalities (atomic formulas). The main problem was to classify the complexity (of counting and decision) for those short PA subclasses. It was proved that a simple subclass with only 5 variables, 2 quantifier alternations and 10 inequalities is **NP**-complete to decide, and also #P-complete to count. Combined with the positive results in [1,2], this settled the last open subcase of classical PA complexity problems. The main reduction in [10] started with the following **NP**-complete problem:[2]

$$\begin{array}{l}\text{Given an interval}^2 [\mu, \nu] \subset \mathbb{Z} \text{ and } n \text{ arithmetic progressions } \mathrm{AP}_i = \\ \mathrm{AP}(g_i, h_i, e_i) := \{g_i, g_i + e_i, \ldots, g_i + h_i e_i\}, \text{with} 1 \le \mu \le \nu, g_i, h_i, e_i \in \mathbb{Z}, \\ h_i \ge 1, \text{decide if there exists some} z \in [\mu, \nu] \setminus \bigcup_{i=1}^{n} \mathrm{AP}_i.\end{array} \qquad \text{(AP-COVER)}$$

In other words, the problem asks whether there is some element in the interval $[\mu, \nu]$ not covered by the given arithmetic progressions. The problem is clearly invariant under a translation of both $[\mu, \nu]$ and the AP_i's, so we can assume $\mu = 1$. Also without affecting the complexity, we can assume that $g_1 = \nu$, $h_1 = 1$, $e_1 = 0$, i.e., $\mathrm{AP}_1 = \{\nu\}$. The main argument in [10] uses continued fractions to construct an integer M and a rational number p/q such that the *best approximations* of p/q, in the terminology of continued fractions, encode $\bigcup_{i=1}^{n} \mathrm{AP}_i$ modulo M. The main point is that p/q should satisfy $\lfloor p/q \rfloor = g_1 = \nu$, so that $[\mu, \nu] = [1, p/q]$, and the formula

$$\Phi_{p,q,M}(z) = 1 \le z \le p/q \ \wedge\ \exists \mathbf{y}\ \ y_2 \equiv z (\mathrm{mod}\ M) \wedge \lfloor p/q \rfloor \le y_2 < p \ \wedge\ q y_2 < p y_1 \wedge \forall \mathbf{x}\ \ \neg \left\{ \begin{array}{rcl} p y_1 - q y_2 & \ge & p x_1 - q x_2 \ \ge\ 0 \\ y_2 & > & x_2 \ > \ 0 \end{array} \right\} \qquad (4)$$

satisfies the property

$$\{z \in \mathbb{Z} : \Phi_{p,q,M}(z)\} = [\mu, \nu] \cap (\bigcup_{i=1}^{n} \mathrm{AP}_i). \qquad (5)$$

Thus, the original AP-COVER/instance is *not* satisfied if and only if $|S_{p,q,M}| = |[\mu, \nu]| = \lfloor p/q \rfloor$. We emphasize that p, q, M can be computed in polynomial time from μ, ν, g_i, h_i, e_i. The meaning behind this formula can be explained as follows.

In Figure 1, the line $y_2/y_1 = p/q$ divides the positive orthant into two parts. The integer hull of the points strictly below this line and above the horizontal axis form a polyhedron, whose boundary is the (bold) convex polygonal curve $\mathcal{C}$, starting at (1,0) and ending at (q, p). Denote by $\mathcal{C}_i$ the i-th edge of $\mathcal{C}$ above the (dotted) horizontal line $y_2 = g_1 = \lfloor p/q \rfloor$. Then for every $1 \le i \le n$ we have $\mathrm{AP}_i = \{y_2 \bmod M \ :\ (y_1, y_2) \in \mathcal{C}_i\}$, and thus $\bigcup_{i=1}^{n} \mathrm{AP}_i = \{y_2 \bmod M \ :\ (y_1, y_2) \in \mathcal{C},\ y_2 \ge g_1\}$.

[2] All intervals in the paper are over $\mathbb{Z}$, so $[a, b]$ with $a, b \in \mathbb{R}$ should be understood as $[a, b] \cap \mathbb{Z}$.

In (4), we express $z \in [\mu, \nu] \cap (\bigcup_{i=1}^{n} \mathrm{AP}_i)$ as $z \equiv y_2 \pmod{M}$ for some (y_1, y_2) with $\lfloor p/q \rfloor \leq y_2 < p$ and $(y_1, y_2) \in \mathcal{C}$.[3] By a basic property of continued fractions (cf., e.g, [9]), the condition $(y_1, y_2) \in \mathcal{C}$ is equivalent to saying that $qy_2 < py_1$, and there is no other integer point (x_1, x_2) with $y_2 > x_2 > 0$ such that x_2/x_1 approximates p/q better than y_2/y_1. This last condition is expressed by the $\forall \mathbf{x} \ldots$ clause in $\Phi_{p,q,M}$.

A hardness result for 3-parameter PA immediately follows.

Proposition 2.1 *Assume* $\mathbf{P} \neq \mathbf{NP}$. *There exists a 3-parametric* Σ_2 *PA family* $S_{p,q,M}$ *such that* $|S_{p,q,M}|$ *is always finite but cannot be expressed as a polynomial-time evaluable function in* p, q, *and* M.

Proof. We can clear the integer denominators in (4) by cross multiplications. The condition $y_2 \equiv z \pmod{M}$ can be expressed with existential quantifiers. Thus we obtain a 3-parametric Σ_2 PA formula $\Phi_{p,q,M}$, which defines a family $S_{p,q,M}$. The set of satisfying values z is finite by $1 \leq z \leq p/q$. Now assume $|S_{p,q,M}|$ is a polynomial-time evaluable function $f(p, q, M)$. Then given any AP-COVER instance, we can compute p, q, M in polynomial time from the AP_i's, and then evaluate $f(p, q, M)$ in polynomial time to check whether $f(p, q, M) = \lfloor p/q \rfloor$. This contradicts $\mathbf{P} \neq \mathbf{NP}$. □

It remains to reduce the three parameters p, q, M to two. To do this, we shall adapt a trick of Glivický and Pudlák [7]. Their context is slightly different from ours in that they use nonstandard integers rather than parameters that range over $\mathbb{Z}$, and that their results involve computability rather than complexity. However their key idea and its proof apply in our context. The two parameters that will be involved are

$$t_1 = pM, \quad t_2 = pqM^2 + M. \tag{6}$$

For convenience, we shall assume for the rest of this section that all the parameters in our formulas (t_1, t_2, p, q, and M) only take nonnegative integer values. Although in other parts of this paper the parameters are assumed to range over $\mathbb{Z}$, this restriction does not affect the hardness results we are proving here.

Proposition 2.2 (Glivický & Pudlák; [7, §3.2]) *For* $0 \leq j < p$, *the three multiplications* $j \mapsto pMj$, $j \mapsto qMj$, $j \mapsto Mj$ *can be defined by using just two multiplications* $j \mapsto t_1 j$ *and* $j \mapsto t_2 j$.

Proof. By definition, we have $t_1 j = pMj$ for all j, so it remains to define the multiplications by qMj and Mj for $0 \leq j < p$. By the division algorithm, for every $j \geq 0$ we can uniquely write $(pqM^2 + M)j = (pM)r + s$, where $0 \leq r$ and $0 \leq s < pM$. If $0 \leq j < p$, then $s = Mj \pmod{pM} = Mj$ and we can then solve to obtain $r = qMj$. Thus for $0 \leq j < p$, the formula

$$t_2 j = t_1 r + s \ \wedge \ 0 \leq r \ \wedge \ 0 \leq s < t_1 \tag{$\mathrm{Div}_{t_1,t_2}(j, r, s)$}$$

is satisfied by the triple (j, qMj, Mj). Furthermore, for such j this formula cannot be satisfied by any other values of the second and third arguments. □

We now prove some additional capabilities of the parameters $t_1 = pM, t_2 = pqM^2 + M$ that will be required in order to transform the entire formula (4) into a formula in t_1 and t_2 alone.

Lemma 2.3 *The congruence relation modulo* M *is definable using just the multiplications by* t_1 *and* t_2.

Proof. Let $\text{Cong-M}_{t_1,t_2}(b, c, w_1, w_2)$ be the formula

$$b - c - t_1 w_1 - t_2 w_2 = 0.$$

Since $\gcd(t_1, t_2) = M$, the condition $b \equiv c \pmod{M}$ is expressed as $\exists w_1 \exists w_2 \, \text{Cong-M}_{t_1,t_2}(b, c, w_1, w_2)$. □

Lemma 2.4 *The constant* p *is definable using just the multiplications by* t_1 *and* t_2.

Proof. Since $t_2/t_1 = qM + 1/p$, p is the smallest positive integer v such that $t_1 | t_2 v$. Since $t_2 p/t_1 = t_2/M = pqM + 1$, we can express that a pair of variables u, v satisfy $(u, v) = (pqM + 1, p)$ by the formula

$$u > 0 \ \wedge \ t_2 v = t_1 u \ \wedge \ \forall v', u' \ \ 0 < v' < v \rightarrow t_2 v' \neq t_1 u'$$

which we denote by $\text{Equal-p}_{t_1,t_2}(v, u)$. □

[3] The curve $\mathcal{C}$ includes (p, q) in [10], but not here. This small difference is not very significant as one can easily check.

Lemma 2.5 *Suppose p, q, and M are positive integers such that $p/q \notin \mathbb{Z}$. If $t_1 = pM$ and $t_2 = pqM^2 + M$ then $\lfloor t_1^2/t_2 \rfloor = \lfloor p/q \rfloor$.*

Proof. First, we have $t_1^2/t_2 = p^2M^2/(pqM^2 + M) = p/(q + 1/pM) < p/q$, so $\lfloor t_1^2/t_2 \rfloor \leq \lfloor p/q \rfloor$. On the other hand, since $p/q \notin \mathbb{Z}$ we have $p \geq \lfloor p/q \rfloor q + 1 > \lfloor p/q \rfloor q + \lfloor p/q \rfloor / pM = \lfloor p/q \rfloor (q + 1/pM)$. This means $t_1^2/t_2 = p/(q + 1/pM) > \lfloor p/q \rfloor$, and thus $\lfloor t_1^2/t_2 \rfloor = \lfloor p/q \rfloor$. □

Proof of Theorem 1.10. In order to apply Proposition 2.2, we must first multiply by M every inequality in (4) that involves multiplication by p or q. This works because multiplications by p, q, and M appear separately in (4). After doing so and clearing some denominators, we obtain the equivalent formula

$$\Phi'_{p,q,M}(z) = \exists y_1, y_2 :$$

$$0 < z \leq p/q \tag{7}$$

$$\wedge y_2 \equiv z \pmod{pM} \tag{8}$$

$$\wedge p/q < y_2 + 1 \leq p \tag{9}$$

$$\wedge qMy_2 < pMy_1 \tag{10}$$

$$\wedge \forall x_1, x_2 \quad \neg \left\{ \begin{matrix} pMy_1 - qMy_2 \geq pMx_1 - qMx_2 \geq 0 \\ y_2 > x_2 > 0 \end{matrix} \right\} \tag{11}$$

Here (9) is equivalent to $\lfloor p/q \rfloor \leq y_2 < p$ in (4) because $y_2 \in \mathbb{Z}$. Now consider the formula

$$\Psi_{t_1,t_2}(z) = \exists y_1, y_2, w_1, w_2, u, v, r, s :$$

$$0 < t_2 z \leq t_1^2 \tag{7'}$$

$$\wedge \text{Cong-M}_{t_1,t_2}(y_2, z, w_1, w_2) \tag{8'}$$

$$\wedge \text{Equal-p}_{t_1,t_2}(u, v) \wedge t_1^2 < t_2(y_2 + 1) \leq t_2 v \tag{9'}$$

$$\wedge \text{Div}_{t_1,t_2}(y_2, r, s) \wedge r < t_1 y_1 \tag{10'}$$

$$\wedge \forall x_1, x_2 \quad (0 < x_2 < y_2 \wedge \text{Div}_{t_1,t_2}(x_2, r', s')) \rightarrow \neg(0 \leq t_1 x_1 - r' \leq t_1 y_1 - r). \tag{11'}$$

It only remains to show that $\Phi'_{p,q,M}(z)$ and $\Psi_{t_1,t_2}(z)$ are equivalent. We have the following:

"(7) $\Leftrightarrow$ (7′)" follows by rounding down both equations to the nearest integer and applying Lemma 2.5. "(8) $\Leftrightarrow$ (8′)" is Lemma 2.3. In order to prove "(9) $\Leftrightarrow$ (9′)", we can again apply Lemma 2.5 to replace p/q in (9) by t_1^2/t_2, since every other quantity in (9) is an integer. By Lemma 2.4, the formula Equal-$\text{p}_{t_1,t_2}(v, u)$ fixes the value of v to be p, so we can now replace p by v to obtain 9′.

We now show that (9) implies "(10) $\Leftrightarrow$ (10′)": By (9), we have $0 \leq y_2 < p$, so by Proposition 2.2, the condition $\text{Div}_{t_1,t_2}(y_2, r, s)$ fixes the value of r to be qMy_2. Here we modify (10) by replacing qMy_2 by r and pMy_1 by $t_1 y_1$ to obtain (10′).

Finally, we use (10) to show "(11) $\Leftrightarrow$ (11′)": Using (10′) which we have already shown to be equivalent to (10), we can replace qMy_2 by r. Using the definition of t_1, we can also replace pMy_1 by ty_1 and pMx_1 by $t_1 x_1$. So (11) is equivalent to

$$\forall x_1, x_2 \quad \neg \left\{ \begin{matrix} ty_1 - r \geq t_1 x_1 - qMx_2 \geq 0 \\ y_2 > x_2 > 0 \end{matrix} \right\},$$

or in another form

$$\forall x_1, x_2 \quad 0 < x_2 < y_2 \rightarrow \neg[ty_1 - r \geq t_1 x_1 - qMx_2 \geq 0].$$

Since the hypothesis $x_2 < y_2$ along with $y_2 < p$ from (9) implies $x_2 < p$, we can (by Proposition 2.2) insert the condition $\text{Div}_{t_1,t_2}(x_2, r', s')$ into the hypothesis to fix r' equal to qMx_2. Accordingly substituting in r' for qMx_2, we obtain (11′).

So $\Phi_{p,q,M}$, $\Phi'_{p,q,M}$ and Ψ_{t_1,t_2} are all equivalent. This finishes the proof of Theorem 1.10. □

Proof of Corollaries 1.12 & 1.11. The formula $\Psi'_{t_1,t_2}(z) := (0 < z \le t_1^2/t_2) \wedge \neg\Psi_{t_1,t_2}(z)$ is satisfied only by those $z \in [\mu, \nu]\setminus\bigcup_{i=1}^{n} \mathrm{AP}_i$ (cf. (5)). This formula defines a 2-parametric family S_{t_1,t_2}. So the condition $|S_{t_1,t_2}| > 0$, which is equivalent to AP-Cover, cannot be expressed using polynomial-time relations in t_1 and t_2. Similarly, any expansion of parametric PA with polynomial-time predicates cannot have full quantifier elimination. For otherwise we can apply it to the sentence $\exists z\ \Psi'_{t_1,t_2}(z)$ and get an equivalent Boolean combination of polynomial-time relations in t_1, t_2.

3 Counting-universality of 2-parametric Presburger formulas

Consider a k-parametric PA formula:

$$\Phi_{\mathbf{u}}(\mathbf{x}) = Q_1 y_1\ Q_2 y_2\ \dots Q_m y_m\ \ \Theta_{\mathbf{u}}(\mathbf{x}, \mathbf{y}). \tag{12}$$

Here $\mathbf{u} \in \mathbb{Z}^k$ are the k scalar parameters, $\mathbf{x} \in \mathbb{Z}^d$ are the free variables, $\mathbf{y} = (y_1, \dots, y_m) \in \mathbb{Z}^m$ are the quantified variables, $Q_1, \dots Q_m \in \{\forall, \exists\}$ are the quantifiers, and $\Theta_{\mathbf{u}}(\mathbf{x}, \mathbf{y})$ is a Boolean combination of linear inequalities in $\mathbf{x}, \mathbf{y}$ with coefficients and constants from $\mathbb{Z}[\mathbf{u}]$. This formula defines a parametric family $S_{\mathbf{u}}$.

We say that a k_1-parametric family $S_{\mathbf{u}}$ *counting-reduces* to an k_2-parametric family $S'_{\mathbf{t}}$ if there exists $f = (f_1, \dots, f_{k_2}) : \mathbb{Z}^{k_1} \to \mathbb{Z}^{k_2}$ with $f_i \in \mathbb{Z}[\mathbf{u}]$ such that for every $\mathbf{u} \in \mathbb{Z}^{k_1}$ we have that $|S_{\mathbf{u}}| = \infty$ implies $|S'_{f(\mathbf{u})}| = \infty$ and $|S_{\mathbf{u}}| < \infty$ implies $|S_{\mathbf{u}}| = |S'_{f(\mathbf{u})}|$.

Theorem 3.1 *Every k-parametric PA family $S_{\mathbf{u}}$ counting-reduces to another 2-parametric PA family $F_{s,t}$ with the same number of alternations. In other words, 2-parametric PA families are counting-universal.*

First we prove the following lemma.

Lemma 3.2 *For every formula $\Phi_{\mathbf{u}}$ of the form* (12)*, there exist $\mu, \mu', \nu_1, \dots, \nu_m \in \mathbb{Z}[\mathbf{u}]$ such that for every value $\mathbf{u} \in \mathbb{Z}^k$ we have:*

(i) $|S_{\mathbf{u}}| = \infty$ *if and only if*

$$\exists \mathbf{x}\ \left[\mu(\mathbf{u}) \le \|\mathbf{x}\|_\infty \le \mu'(\mathbf{u})\ \wedge\ Q_1(|y_1| \le \nu_1(\mathbf{u}))\ \dots\ Q_m(|y_m| \le \nu_m(\mathbf{u}))\ \ \Theta_{\mathbf{u}}(\mathbf{x}, \mathbf{y})\right]$$

(ii) If $|S_{\mathbf{u}}| < \infty$ *then for every* $\mathbf{x} \in \mathbb{Z}^d$*, we have*

$$S_{\mathbf{u}}(\mathbf{x}) = \mathit{true} \iff \|\mathbf{x}\|_\infty \le \mu(\mathbf{u})\ \wedge\ Q_1(|y_1| \le \nu_1(\mathbf{u}))\ \dots\ Q_m(|y_m| \le \nu_m(\mathbf{u}))\ \ \Theta_{\mathbf{u}}(\mathbf{x}, \mathbf{y}).$$

Here $\|\cdot\|_\infty$ is the ℓ_∞–norm. So $\mu(\mathbf{u}) \le \|\mathbf{x}\|_\infty$ stands for $\bigvee_{i=1}^{d}(x_i \le -\mu(\mathbf{u})\ \vee\ \mu(\mathbf{u}) \le x_i)$ and $\|\mathbf{x}\|_\infty \le \mu'(\mathbf{u})$ stands for $\bigwedge_{i=1}^{d}(-\mu'(\mathbf{u}) \le x_i \le \mu'(\mathbf{u}))$. Each restricted quantifier $Q_i(|y_i| \le \nu_i(\mathbf{u}))$ means exits/for all y_i in the interval $[-\nu_i(\mathbf{u}), \nu_i(\mathbf{u})]$.[4]

Proof. Consider a usual, non-parametric PA formula:

$$\Phi(\mathbf{x})\ =\ Q_1 y_1\ Q_2 y_2\ \dots Q_m y_m\ \ \Theta(\mathbf{x}, \mathbf{y}), \quad \mathbf{x} \in \mathbb{Z}^n,$$

which defines some set $S \subseteq \mathbb{Z}^n$. Recall Cooper's quantifier elimination procedure for Presburger arithmetic (cf. [11]). Applying it to $\Phi(\mathbf{x})$, we obtain an *equivalent* quantifier free formula $\Phi'(\mathbf{x})$, which may contain some extra divisibility predicates. By [11, Theorem 2], after eliminating all m quantifiers from Φ, we obtain the following bounds:

$$c' \ \le\ c^{4^m}, \quad s' \ \le\ s^{(4c)^{4^m}}, \quad a' \ \le\ a^{4^m} s^{(4c)^{4^m}},$$

where c is the number of distinct integers that appeared as coefficients or divisors in Φ, s is the largest absolute value of all integers that appeared in Φ (coefficients + divisors + constants), a is the total number of atomic formulas in Φ (inequalities + divisibilities), and c', s', a' are the corresponding quantities for Φ'. Now assume c, m and n are fixed. Then we have

$$c' \ \le\ \text{const}, s' \ \le\ s^{\text{const}}, \text{and } a' \ \le\ a^{\text{const}} s^{\text{const}},$$

[4] Here we understand that μ, μ', ν_i have positive values for all $\mathbf{u} \in \mathbb{Z}^k$.

where $\text{const} = \text{const}(c, m)$ is fixed. So in this case Φ' has at most a fixed number of coefficients and divisors.

Denote by D the common multiple of all divisors in Φ'. We have $D \le s^{\text{const}}$. Let $\mathcal{L} = \langle De_1, \ldots, De_n \rangle$ be the lattice of $\mathbb{Z}^n$ consisting of $\mathbf{x} \in \mathbb{Z}^n$ whose coordinates are all divisible by D. Fix some particular coset $\mathcal{C}$ of $\mathcal{L}$ and restrict $\mathbf{x}$ to $\mathcal{C}$. Then in $\Phi'(\mathbf{x})$, all divisor predicates have fixed values (either true or false) as $\mathbf{x}$ varies over $\mathcal{C}$. So over $\mathcal{C}$, the formula $\Phi'(\mathbf{x})$ is just a Boolean combination of linear inequalities in $\mathbf{x}$, which represents a *disjoint union* of some rational polyhedra in $\mathbb{R}^n$. Each such polyhedron P can be described by a system of *fixed* length, because there are only at most c' different coefficients for the $\mathbf{x}$ variables. The integers in the system are also bounded by s^{const}. We consider $P \cap \mathcal{C}$. By the fundamental theorem of Integer Programming[5] (cf. [12, Theorems 16.4 and 7.1]), we have:

$$P \cap \mathcal{C} \;=\; \text{conv}(\overline{v}_1, \ldots, \overline{v}_p) \;+\; \mathbb{Z}_+\langle \overline{w}_1, \ldots, \overline{w}_q \rangle$$

for some $\overline{v}_i, \overline{w}_j \in \mathbb{Z}^n$ with $\|\overline{v}_i\|_\infty, \|\overline{w}_j\|_\infty < s^{\text{const}'}$. Here $\text{const}' = \text{const}'(c, m, n)$ is fixed. From this, it is easy to see that there is $\text{const}'' = \text{const}''(c, m, n)$ such that for every polyhedron P in the disjoint union, we have:

$$|P \cap \mathcal{C}| = \infty \quad\Longleftrightarrow\quad \text{there is } \mathbf{x} \in P \cap \mathcal{C} \text{ with } s^{\text{const}''} < \|\mathbf{x}\|_\infty < s^{2\text{const}''},$$

$$|P \cap \mathcal{C}| < \infty \quad\Longrightarrow\quad P \cap \mathcal{C} \subseteq [-s^{\text{const}''}, s^{\text{const}''}]^n.$$

Since this holds for every coset $\mathcal{C}$ of $\mathcal{L}$, we conclude that there is $\text{const}_0 = \text{const}_0(c, m, n)$ such that:

$$|S| = \infty \quad\Longleftrightarrow\quad \exists \mathbf{x} \text{ with } s^{\text{const}_0} < \|\mathbf{x}\|_\infty < s^{2\text{const}_0} \text{ and } \Phi'(\mathbf{x}) = \text{true} \tag{13}$$

$$|S| < \infty \quad\Longrightarrow\quad \forall \mathbf{x}\ \big(\Phi'(\mathbf{x}) = \text{true} \;\to\; \|\mathbf{x}\|_\infty \le s^{\text{const}_0}\big). \tag{14}$$

This gives us a bound for $\mathbf{x}$. Now for every $\mathbf{x}$ with $\|\mathbf{x}\|_\infty \le s^{\text{const}_0}$, by the same argument, it is enough to decide the (substituted) sentence $\Phi(\mathbf{x})$ over those y_1 with $|y_1| \le s^{\text{const}_1}$. In other words, for every such value for $\mathbf{x}$, we may replace $Q_1 y_1$ by $Q_1(|y_1| \le s^{\text{const}_1})$ in $\Phi(\mathbf{x})$ to obtain a new formula $\Phi_1(\mathbf{x})$, which is equivalent to the original formula $\Phi(\mathbf{x})$. Working inwards, we can likewise bound $|y_2|$ by s^{const_2}, $|y_3|$ by s^{const_3}, etc. Therefore, in case $|S| < \infty$, the whole formula Φ is equivalent to one with bounded quantifiers on all y_i. Also by (13), we have $|S| = \infty$ if and only if some $s^{\text{const}_0} < \|\mathbf{x}\|_\infty < s^{2\text{const}_0}$ satisfies it. For $\mathbf{x}$ in this range, we can again bound y_1, y_2, etc., accordingly by some other powers of s. Note that we can bound each y_i by a common larger power of s for both cases (13) and (14).

In a k-parametric PA formula $\Phi_{\mathbf{u}}(\mathbf{x})$, we consider m, n and c to be fixed. Since all coefficients and constants of $\Phi_{\mathbf{u}}$ are in $\mathbb{Z}[\mathbf{u}]$, we can bound s by some polynomial in $\mathbf{u}$. Thus, every s^{const} is also bounded by some polynomial in $\mathbf{u}$. This proves Lemma 3.2. □

In the above application of Cooper's elimination, if only m, n are fixed but not c, then we no longer have the bound $s' \le s^{\text{const}}$. Instead, we would have $c', \log s' \le \text{poly}(c, \log s)$. A bound of this type is important for showing that the decision problem for classical PA with a bounded number of variables falls within the Polynomial Hierarchy (cf., e.g., [8]). However, it would not be strong enough for our argument, which crucially needs $\log s' = O(\log s)$.

From Lemma 3.2, it is easy to see that $S_{\mathbf{u}}$ counting-reduces to the family $\widetilde{S}_{\mathbf{u}}$ defined by the following formula $\widetilde{\Phi}_{\mathbf{u}}(\mathbf{x}, \widetilde{x})$:

$$\begin{aligned}\widetilde{\Phi}_{\mathbf{u}}(\mathbf{x}, \widetilde{x}) = \big[\widetilde{x} \ge 0 \;\wedge\; Q_1(|y_1| \le \nu_1(\mathbf{u})) \;\ldots\; Q_m(|y_m| \le \nu_m(\mathbf{u}))\\ \mu(\mathbf{u}) \le \|\mathbf{x}\|_\infty \le \mu'(\mathbf{u}) \;\wedge\; \Theta_{\mathbf{u}}(\mathbf{x}, \mathbf{y})\big] \;\vee\\ \big[\widetilde{x} = 0 \;\wedge\; Q_1(|y_1| \le \nu_1(\mathbf{u})) \;\ldots\; Q_m(|y_m| \le \nu_m(\mathbf{u}))\\ \|\mathbf{x}\|_\infty \le \mu(\mathbf{u}) \;\wedge\; \Theta_{\mathbf{u}}(\mathbf{x}, \mathbf{y})\big].\end{aligned}$$

Here the bounds on $\|\mathbf{x}\|_\infty$ are moved to after the quantifiers on y_i without changing the meaning. The dummy variable $\widetilde{x}$ is used to make sure that $|\widetilde{S}_{\mathbf{u}}| = \infty$ in the first case.

Proof of Theorem 3.1. We show that $\widetilde{S}_{\mathbf{u}}$ counting-reduces to a 2-parameter family $F_{s,t}$, defined by a new formula $\Psi_{s,t}$. First, we list all the different scalar terms that appear in $\widetilde{\Phi}_{\mathbf{u}}$, either as coefficients or constants

[5] We are rescaling $\mathcal{L}$ to $\mathbb{Z}$ before applying this bound.

(including all μ, μ', ν_i), as $\delta_0(\mathbf{u}), \ldots, \delta_r(\mathbf{u})$. Now suppose we need to multiply some $z \in \mathbb{N}$ by $\delta_0(\mathbf{u}), \ldots, \delta_r(\mathbf{u})$ and also know that

$$-t/2 < \delta_0(\mathbf{u})z,\ \ldots,\ \delta_r(\mathbf{u})z < t/2 \tag{15}$$

for some $t \in \mathbb{Z}$. The following base-t concatenation, which is similar to (6), can be used. Essentially, we encode the "multi"-product $(\delta_0(\mathbf{u})z, \ldots, \delta_r(\mathbf{u})z)$ as a single product:

$$\delta_0(\mathbf{u})z \ +\ t\,\delta_1(\mathbf{u})z + \cdots + t^r\,\delta_r(\mathbf{u})z \ =\ (\delta_0(\mathbf{u}) + t\,\delta_1(\mathbf{u}) + \cdots + t^r\delta_r(\mathbf{u}))\,z.$$

In other words, if $s = \delta_0(\mathbf{u}) + t\,\delta_1(\mathbf{u}) + \cdots + t^r\delta_r(\mathbf{u})$ and:

$$s\,z \ =\ z_0 + t\,z_1 + \cdots + t^r z_r \ \ \wedge\ \ t/2 < z_0, \ldots, z_r < -t/2, \qquad (\mathrm{Div}_{s,t}(z, z_0, \ldots, z_r))$$

then we must have $z_0 = \delta_0(\mathbf{u})z,\ \ldots,\ z_r = \delta_r(\mathbf{u})z$. Indeed, by subtracting we get $z_0 - \delta_0(\mathbf{u})z \equiv 0 \pmod t$, which implies $z_0 = \delta_0(\mathbf{u})z$ because $-t/2 < z_0,\ \delta_0(\mathbf{u})z < t/2$. The same argument applies to other z_i.

Observe that in $\widetilde{\Phi}_{\mathbf{u}}$, all variables $\mathbf{x}$ and $\mathbf{y}$ are bounded by polynomials in $\mathbf{u}$. Hence, we can pick $\eta(\mathbf{u}) \in \mathbb{Z}[\mathbf{u}]$ so that for every value $\mathbf{u} \in \mathbb{Z}^k$, the condition (15) is always satisfied when $t = \eta(\mathbf{u})$ and z is either the constant 1 or any of the possible values of the $\mathbf{x}, \mathbf{y}$ variables. Our reduction map $f : \mathbb{Z}^k \to \mathbb{Z}^2$ can now be defined by letting $t = \eta(\mathbf{u})$ and $s = \delta_0(\mathbf{u}) + t\,\delta_1(\mathbf{u}) + \cdots + t^r\delta_r(\mathbf{u})$. Now we can define $\Psi_{s,t}(\mathbf{x}, \widetilde{x})$ from $\widetilde{\Phi}_{\mathbf{u}}(\mathbf{x}, \widetilde{x})$. We need $(m + d + 1)(r + 1)$ extra variables $\mathbf{w} = (w_{ij})_{1\le i\le d,\, 0\le j\le r}$, $\mathbf{w}' = (w'_{ij})_{1\le i\le m,\, 0\le j\le r}$, and $\mathbf{v} = (v_j)_{0\le j\le r}$. Assuming the last quantifier Q_m in $\widetilde{\Phi}_{\mathbf{u}}$ is $\exists$, we insert

$$\mathrm{Div}_{s,t}(x_i, w_{i0}, \ldots, w_{ir}) \ \wedge\ \bigwedge_{i=1}^{m} \mathrm{Div}_{s,t}(y_i, w'_{i0}, \ldots, w'_{ir}) \ \wedge\ \mathrm{Div}_{s,t}(1, v_0, \ldots, v_r)] \tag{$\star$}$$

right before $\Theta_{\mathbf{u}}(\mathbf{x}, \mathbf{y})$, i.e., replace $\Theta_{\mathbf{u}}(\mathbf{x}, \mathbf{y})$ by $(\star) \wedge \Theta_{\mathbf{u}}(\mathbf{x}, \mathbf{y})$. Then in $\widetilde{\Phi}_{\mathbf{u}}$ we replace every term $\delta_j(\mathbf{u})\,x_i$ by w_{ij}, every term $\delta_j(\mathbf{u})y_i$ by w'_{ij} and every term $\delta_j(\mathbf{u})$ by v_j. Now $\widetilde{\Phi}_{\mathbf{u}}$ becomes the desired $\Psi_{s,t}$. In case $Q_m = \forall$, we insert:

$$\forall\,\mathbf{w}, \mathbf{w}', \mathbf{v}\ \bigvee_{i=1}^{d} \neg\mathrm{Div}_{s,t}(x_i, w_{i0}, \ldots, w_{ir}) \vee \bigvee_{i=1}^{m} \neg\mathrm{Div}_{s,t}(y_i, w'_{i0}, \ldots, w'_{ir}) \ \vee\ \neg\mathrm{Div}_{s,t}(1, v_0, \ldots, v_r) \tag{$\star\star$}$$

right before $\Theta_{\mathbf{u}}(\mathbf{x}, \mathbf{y})$, i.e., replace $\Theta_{\mathbf{u}}(\mathbf{x}, \mathbf{y})$ by $(\star\star) \vee \Theta_{\mathbf{u}}(\mathbf{x}, \mathbf{y})$. Again, replace every term $\delta_j(\mathbf{u})\,x_i$ by w_{ij}, every term $\delta_j(\mathbf{u})y_i$ by w'_{ij} and every term $\delta_j(\mathbf{u})$ by v_j. This gives $\Psi_{s,t}$.

Note that $\Psi_{s,t}$ still has the form $[\ldots] \vee [\ldots]$ with each disjunct containing m alternations $Q_1 \ldots Q_m$. This formula is equivalent to a formula in prenex normal form with m quantifier alternations, so we are done. □

In case $S_{\mathbf{u}}$ is defined by a quantifier-free formula, i.e., $m = 0$, we only need to insert $(\star)$, without the $\exists$ quantifiers, before $\Theta_{\mathbf{u}}(\mathbf{x}, \mathbf{y})$. This is because $\mathrm{Div}_{s,t}(z, z_0, \ldots, z_r)$ uniquely determines $z_0, \ldots, z_r$ in z. So in this case $S_{\mathbf{u}}$ also counting-reduces to a quantifier-free $F_{s,t}$, although the latter has many more free variables. Thus, the study of integer point counting functions on k-parametric polyhedra reduces to the case of 2-parametric polyhedra in higher dimensions.

4 Counting points in parametric unordered Presburger families in polynomial time

In this section, we consider the reduct of multi-parametric Presburger arithmetic to the language without ordering, so that basic quantifier-free formulas are equivalent to Boolean combinations of equations of the form $f_1(\mathbf{t})x_1 + \cdots + f_n(\mathbf{t}) = g(\mathbf{t})$, where $\mathbf{t} = (t_1, \ldots, t_k)$ is a tuple of parameters and $f_1, \ldots, f_m, g \in \mathbb{Z}[\mathbf{t}]$. As always, we are allowed to quantify over the variables x_i but not over the parameters $\mathbf{t}$. Note that if there is no parameter $\mathbf{t}$, this would correspond to studying the first-order logic of the additive group $(\mathbb{Z}; +)$. More precisely:

A *k-parametric unordered PA family* is a collection $\{S_{\mathbf{t}} : \mathbf{t} = (t_1, \ldots, t_k) \in \mathbb{Z}^k\}$ of subsets of $\mathbb{Z}^d$ which can be defined by an equation of the form

$$S_{\mathbf{t}} = \{\mathbf{x} \in \mathbb{Z}^d \quad : \quad Q_1 y_1 \; Q_2 y_2 \; \ldots Q_m y_m \;\; \Theta_{\mathbf{t}}(\mathbf{x}, \mathbf{y})\},$$

where the $Q_i \in \{\forall, \exists\}$ are quantifiers for variables y_i ranging over $\mathbb{Z}$ and $\Theta_{\mathbf{t}}(\mathbf{x}, \mathbf{y})$ is a Boolean combination of linear equations with coefficients in $\mathbb{Z}[\mathbf{t}]$. E.g., $(x_1 = 0) \wedge \exists x_2 \exists x_3 (x_2 t_1 + x_3 t_2 = 1)$ defines a 2-parametric unordered PA family $\{S_{\mathbf{t}} \subseteq \mathbb{Z} : \mathbf{t} \in \mathbb{Z}^2\}$ such that $S_{\mathbf{t}} = \{0\}$ if $\gcd(t_1, t_2) = 1$ and $S_{\mathbf{t}} = \varnothing$ otherwise.

Theorem 4.1 *Suppose that $S_{\mathbf{t}} \subseteq \mathbb{Z}_d$ is a k-parametric unordered PA family. Then*

1. *there is a polynomial-time algorithm to decide whether $S_{\mathbf{t}}$ is nonempty;*
2. *there is a polynomial-time algorithm on input $\mathbf{t}$ which decides whether or not $S_{\mathbf{t}}$ is finite or infinite; and*
3. *there is a polynomial-time evaluable function $g : \mathbb{Z}^k \to \mathbb{N}$ such that whenever $S_{\mathbf{t}}$ is finite, $g(\mathbf{t}) = |S_{\mathbf{t}}|$.*

In fact, the proof of Theorem 4.1 will show that the decision algorithms for (1) and (2) rely upon only a few basic, concrete number-theoretic operations on $\mathbf{t}$, such as gcd and a couple of related functions.

To prove Theorem 4.1, we need to recall some notation from [13]. To eliminate quantifiers, they work in a two-sorted language L_2 in which variables x_i and parameters in $\mathbf{t}$ are assigned to objects of distinct domains, called the *group sort* and the *ring sort*, respectively. For our purposes, the group sort and the ring sort are two disjoint copies of $\mathbb{Z}$. The variables x_i and y_i will always range over values in the group sort, and the parameters t_i will always range over values in the scalar sort. In other words, we can think of the parameters $t_1, \ldots, t_k$ as "typed variables" ranging over a domain of possible parameter values in the scalar sort (a copy of $\mathbb{Z}$), and $x_1, x_2, \ldots$ as variables of a distinct type ranging over values in the group sort (which is a different copy of $\mathbb{Z}$), and the parameters t_i act upon the group sort by scalar multiplication.

The language L_2 consists of the following nonlogical symbols (in addition to equality): within the scalar sort, constant symbols for 0 and 1, a unary operation $-$ for negation, ring operations $+$ and $\cdot$, and four additional binary operations g, α, β, and γ (whose interpretation is explained below); within the group sort, a constant symbol for 0, a unary operation $-$ for negation, and a symbol $+$ for addition; a binary operation $\cdot$ such that $s \cdot x$ is a value in the group sort whenever s is a value in the scalar sort and x is a value in the group sort, denoting multiplication by s in the usual sense; and a binary relation symbol $|$ to be interpreted such that whenever s is in the scalar sort and x is in the group sort,

$$s|x \Leftrightarrow \exists y (s \cdot y = x).$$

The binary operations g, α, β, and γ between values in the scalar sort are interpreted so that $g(r, s) = gcd(r, s)$ and the following axioms hold for all values r, s in the scalar sort:

$$r = \gamma(r, s) \cdot g(r, s),$$
$$1 = \alpha(r, s) \cdot \gamma(r, s) + \beta(r, s) \cdot \gamma(s, r).$$

We shall use the following fact, proved in [13]:

Theorem 4.2 *Any formula $\varphi_{\mathbf{t}}(\overline{x})$ in k-parametric unordered Presburger arithmetic is logically equivalent to a quantifier-free L_2-formula $\psi(\overline{x}, \mathbf{t})$: that is, with the natural interpretations of the symbols from L_2 given above,*

$$\models \forall \overline{x} \in \mathbb{Z}^d \; \forall \mathbf{t} \in \mathbb{Z}^k \; (\varphi_{\mathbf{t}}(\overline{x}) \leftrightarrow \psi(\overline{x}, \mathbf{t})),$$

where $\psi(\overline{x}, \mathbf{t})$ is a Boolean combination of equations $s_1(\overline{x}, \mathbf{t}) = s_2(\overline{x}, \mathbf{t})$ and divisibility relations $s_3(\mathbf{t})|s_1(\overline{x}, \mathbf{t})$, where $s_1(\overline{x}, \mathbf{t})$, $s_2(\overline{x}, \mathbf{t})$, and $s_3(\mathbf{t})$ are L_2-terms, i.e., expressions built up using only the operations in L_2 and the displayed parameters and variables.

Proof of Theorem 4.1 Say $\varphi_{\mathbf{t}}(\overline{x})$ defines a k-parametric unordered PA family in $\mathbb{Z}^d$.

Note that (1) follows almost immediately from quantifier elimination: by Theorem 4.2, the formula $\exists \overline{x} \varphi_t(\overline{x})$ is equivalent to a quantifier-free L_2-formula $\psi(\mathbf{t})$ in only the scalar sort of $\mathbf{t}$, which is a Boolean combination of equations and divisibility relations $|$ in the k parameters using ring operations and the functions g, α, β, and γ, but all of these operations are polynomial-time computable.

For (2), let us assume (by Theorem 4.2) that $\varphi_{\mathbf{t}}(\overline{x})$ is a quantifier-free L_2-formula, and that $\varphi_t(\overline{x})$ is in disjunctive normal form:

$$\varphi_t(\overline{x}) = \bigvee_{i=1}^{m} \vartheta_i(\overline{x}, \mathbf{t}),$$

where each $\vartheta_i(\overline{x}, \mathbf{t})$ is a conjunction of *literals*.[6]

Claim 4.3 *For any fixed value of* $\mathbf{t} \in \mathbb{Z}^k$ *and of* $i \in \{1, \ldots, m\}$, *if* $S_i := \{\overline{x} \in \mathbb{Z}^d : \models \vartheta_i(\overline{x}, \mathbf{t})\}$, *then* $|S_i|$ *is either 0, 1, or* ∞.

Proof. By rearranging terms, we may assume that all atomic L_2-formulas in $\vartheta_i(\overline{x}, \mathbf{t})$ have the form

$$r \mid s(\overline{x}, \mathbf{t}) \tag{A}$$

or

$$s(\overline{x}, \mathbf{t}) = 0, \tag{B}$$

where $s(\overline{x}, \mathbf{t}) = r_0 + \sum_{i=1}^{d} r_i \cdot x_i$ and $r_0, r_1, \ldots, r_n$, and r are terms in the scalar sort. The terms r and r_i may involve the parameters $\mathbf{t}$ and the operations g, α, β, γ, but the details of this are irrelevant since $\mathbf{t}$ has a fixed value.

Write $\vartheta_i(\overline{x}, \mathbf{t}) = \vartheta_A(\overline{x}, t) \wedge \vartheta_B(\overline{x}, \mathbf{t})$ where $\vartheta_A(\overline{x}, \mathbf{t})$ is the conjunctions of all literals of type (A) and $\vartheta_B(\overline{x}, t)$ is the conjunction of all literals of type (B).

First we consider the atomic formulas of type (A). Each one defines some coset of a finite-index subgroup of $\mathbb{Z}^d$, and so the negation of such a formula defines a finite union of cosets of finite-index subgroups. Since the intersection of finitely many finite-index subgroups is of finite index, there is a single subgroup $H \leq \mathbb{Z}^d$ such that $[\mathbb{Z}^d : H] < \infty$ and $\vartheta_A(\overline{x}, \mathbf{t})$ defines a Boolean combination of cosets of H.

Now consider the atomic formulas of type (B). We decompose $\vartheta_B(\overline{x}, \mathbf{t})$ further as $\vartheta_B(\overline{x}, \mathbf{t}) = \vartheta_B^+(\overline{x}, \mathbf{t}) \wedge \vartheta_B^-(\overline{x}, \mathbf{t})$ where $\vartheta_B^+(\overline{x}, \mathbf{t})$ is the conjunction of all positive (non-negated) atomic formulas of type (B) and $\vartheta_B^-(\overline{x}, \mathbf{t})$ is the conjunction of all negative literals of type (B). Note that the set of solutions to $\vartheta_B^+(\overline{x}, \mathbf{t})$ is of the form $(\vec{v} + S) \cap \mathbb{Z}^d$ where S is a vector subspace of $\mathbb{R}^d$ and $\vec{v} \in \mathbb{Z}^d$.

Finally, suppose that there are at least two distinct elements $\overline{x}_1, \overline{x}_2 \in \mathbb{Z}^d$ in S_i, and to finish the proof of the Claim we shall show that S_i has infinitely many elements. In particular, both $\overline{x}_1$ and $\overline{x}_2$ are solutions to $\vartheta_A(\overline{x}, \mathbf{t})$, so there are cosets C_1, C_2 of H such that $\overline{x}_1 \in C_1$, $\overline{x}_2 \in C_2$, and any element $\overline{x} \in C_1 \cup C_2$ satisfies $\vartheta_A(\overline{x}, \mathbf{t})$. Let $L \subseteq \mathbb{R}^d$ be the line passing through $\overline{x}_1$ and $\overline{x}_2$, and observe that since $\overline{x}_1$ and $\overline{x}_2$ satisfy $\vartheta_B^+(\overline{x}, \mathbf{t})$ (which defines the intersection of an affine subspace with $\mathbb{Z}^d$), any other element of $L \cap \mathbb{Z}^d$ will also satisfy $\vartheta_B^+(\overline{x}, \mathbf{t})$.

For any $j \in \mathbb{Z}$, let $\overline{x}(j) := \overline{x}_1 + j \cdot (\overline{x}_2 - \overline{x}_1)$ and $X := \{j \in \mathbb{Z} : \overline{x}(j) \text{ satisfies } \vartheta_i(\overline{x}, \mathbf{t})\}$. Since H is a finite-index subgroup of $\mathbb{Z}^d$, adding successive copies of the element $(\overline{x}_2 - \overline{x}_1)$ to $\overline{x}_1$ causes the $\overline{x}(j)$ to cycle through cosets of H, and the set of j for which $\vartheta_A(\overline{x}(j), \mathbf{t})$ is true is infinite (and periodic). As observed in the previous paragraph, *every* $\overline{x}(j)$ lies on the line L, and hence $\vartheta_B^+(\overline{x}(j), \mathbf{t})$ is always true, and we need only worry about the truth of $\vartheta_B^-(\overline{x}(j), \mathbf{t})$. Now $\vartheta_B^-(\overline{x}(j), \mathbf{t})$ is true whenever $\overline{x}(j)$ *avoids* every one of a finite number of affine subspaces $A_1, \ldots, A_\ell$ of $\mathbb{R}^d$, but given that L is a line which contains some points satisfying the formula $\vartheta_B^-(\overline{x}, \mathbf{t})$, each A_i can only intersect L in at most one point. Therefore X is infinite, as we wanted. □

The Claim shows that we can define the set of values of the parameter $\mathbf{t}$ for which any given $\vartheta_i(\overline{x}, \mathbf{t})$ has infinitely many solutions (for $\overline{x}$) by the formula

$$\exists \overline{x}_1 \exists \overline{x}_2 \, (\overline{x}_1 \neq \overline{x}_2 \wedge \vartheta_i(\overline{x}_1, \mathbf{t}) \wedge \vartheta_i(\overline{x}_2, \mathbf{t})),$$

and as before this is equivalent to a quantifier-free L_2-formula $\psi_i(\mathbf{t})$ whose truth can be decided by a polynomial-time algorithm in $\mathbf{t}$. Finally, our original formula $\bigvee_{i=1}^{m} \vartheta_i(\overline{x}, \mathbf{t})$ has infinitely many solutions just in case any one of the formulas $\vartheta_i(\overline{x}, \mathbf{t})$ does, establishing (2).

By the argument above, for any k-parametric unordered PA family $S_{\mathbf{t}}$, there is a finite partition $\mathbb{Z}^k = X_1 \cup \ldots \cup X_\ell$ which is definable by quantifier-free L_2-formulas in $\mathbf{t}$ and such that $|S_{\mathbf{t}}|$ is constant as $\mathbf{t}$ varies over any of the sets X_i. Since deciding whether $\mathbf{t} \in X_i$ is polynomial-time decidable, this establishes (3). □

[6] A literal is an *atomic* L_2-formula, i.e., one containing no logical operations $\wedge$, $\vee$ or $\neg$, or the negation of an atomic formula.

5 Summary of complexity results

To conclude, we summarize the complexity results which suggest that Theorem 1.10 may be the best we could hope for: weakening or changing various assumptions results in problems which can be resolved in polynomial time, or else (with unrestricted multiplication) have no algorithmic solutions at all.

Recall that Theorem 1.10 states that, if $\mathbf{P} \neq \mathbf{NP}$, then there is a Σ_2 PA family $S_{\mathbf{t}}$ with two parameters $\mathbf{t} = (t_1, t_2)$ such that $|S_{\mathbf{t}}|$ cannot be computed in polynomial time given $\mathbf{t}$ as input.

However:

(i) If we allow only a single parameter $t \in \mathbb{N}$ (or $\mathbf{t} \in \mathbb{Z}$), then for any PA family S_t, we can compute $|S_t|$ in polynomial time, even if S_t has complexity Σ_2 or higher, by Corollary 1.9.

(ii) If $S_{\mathbf{t}}$ is a k-parametric PA family defined by a formula of complexity Π_1 or Σ_1, then [2] implies that there is a polynomial time algorithm to evaluate $|S_{\mathbf{t}}|$, for any finite number k of parameters. If $S_{\mathbf{t}}$ is defined by a quantifier-free formula, then a polynomial-time algorithm was earlier given in [1].

(iii) If $S_{\mathbf{t}}$ is any k-parametric PA family defined by a formula with no inequalities (only equations), as in § 4, then $|S_{\mathbf{t}}|$ can be evaluated in polynomial time, regardless of the number of quantifier alternations in the defining formula or the number of parameters.

(iv) In k-parametric PA formulas, we allow a restricted version of multiplication: the non-quantified parameters in $\mathbf{t}$ can be multiplied by terms containing the variables $\mathbf{x}$ and $\mathbf{y}$, but no multiplication between the $\mathbf{x}$ and $\mathbf{y}$ variables is allowed. Permitting unrestricted multiplication amongst the $\mathbf{x}$ and $\mathbf{y}$ variables in a parametric PA formula would obviously be bad, since the full first-order theory of $(\mathbb{N}, +, \cdot)$ is undecidable (by theorems of Church and Turing—cf., e.g., [4]). In fact, the Matiyasevich-Robinson-Davis-Putnam theorem [5] states that there is a *single* multivariate polynomial $p(t, x_1, \ldots, x_d)$ such that if $\Phi_t(x_1, \ldots, x_d)$ is the formula expressing $p(t, x_1, \ldots, x_d) = 0$, then the set of $t \in \mathbb{N}$ for which $\Phi_t(x_1, \ldots, x_d)$ defines a nonempty subset of $\mathbb{Z}^d$ is not computable (much less in polynomial time). Note that here we have only a single parameter t, no quantifiers in the formula Φ_t, and mere equations rather than inequalities.

(v) On the other hand, if we allow *no multiplication*, even by parameters (cf. Example 1.4), then $|S_{\mathbf{t}}|$ will be computable in polynomial time; in fact, it has a nice form as a piecewise-defined quasi-polynomial [15].

Acknowledgements We thank Igor Pak for interesting conversations and helpful remarks. This work was started when the first and third authors were participating in the MSRI program *Geometric and Topological Combinatorics*; we thank MSRI for their hospitality. The third author was partially supported by the UCLA Dissertation Year Fellowship. The first author would also like to thank San Francisco State University and the second author would like to thank the City University of New York for hosting them as visiting researchers.

References

[1] A. Barvinok, A polynomial time algorithm for counting integral points in polyhedra when the dimension is fixed, Math. Oper. Res. **19**(4), 769–779 (1994).
[2] A. Barvinok and K. Woods, Short rational generating functions for lattice point problems, J. Amer. Math. Soc. **16**(4), 957–979 (2003).
[3] T. Bogart, J. Goodrick, and K. Woods, Parametric Presburger arithmetic: logic, combinatorics, and quasi-polynomial behavior, Discrete Anal. **2017**(4) (2017).
[4] A. Church, An unsolvable problem of elementary number theory, Amer. J. Math. **58**, 345–63 (1936).
[5] M. Davis, Hilbert's tenth problem is unsolvable, Amer. Math. Mon. **80**(3), 233–269 (1973).
[6] M. J. Fischer and M. O. Rabin, Super-exponential complexity of Presburger arithmetic, in: Complexity of Computation, Proceedings of a Symposium Held in New York City, April 18–19, 1973, edited by R. M. Karp, SIAM-AMS Proceedings Vol. 7 (American Mathematical Society, 1974), pp. 27–41.
[7] P. Glivický and P. Pudlák, A wild model of linear arithmetic and discretely ordered modules, Math. Log. Q. **63**(6), 501–508 (2017).
[8] E. Grädel, Subclasses of Presburger arithmetic and the polynomial-time hierarchy, Theoret. Comput. Sci. **56**(3), 289–301 (1988).
[9] O. Karpenkov, Geometry of Continued Fractions, Algorithms and Computation in Mathematics Vol. 26 (Springer, 2013).
[10] D. Nguyen and I. Pak, Short Presburger arithmetic is hard, in: Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science, held in Berkeley, California, October 15–17, 2017, edited by C. Ulmans (IEEE Computer Society, 2017), pp. 37–48.

[11] D. C. Oppen, A $2^{2^{2^{pn}}}$ upper bound on the complexity of Presburger arithmetic, J. Comput. System Sci. **16**(3), 323–332 (1978).
[12] A. Schrijver, Theory of Linear and Integer Programming, Wiley-Interscience Series in Discrete Math. (Wiley, 1986).
[13] L. van den Dries and J. Holly, Quantifier elimination for modules with scalar variables, Ann. Pure Appl. Log. **57**, 161–179 (1992).
[14] K. Woods, The unreasonable ubiquitousness of quasi-polynomials, Electron. J. Comb. **21**(1), P1.44 (2014).
[15] K. Woods, Presburger arithmetic, rational generating functions, and quasi-polynomials, J. Symb. Log. **80**(2), 433–449 (2015).

