ELSEVIER

Quantifier elimination for the reals with a predicate for the powers of two

Jeremy Avigad*, Yimu Yin

Department of Philosophy, Carnegie Mellon University, Pittsburgh, PA 15213, United States

Abstract

In 1985, van den Dries showed that the theory of the reals with a predicate for the integer powers of two admits quantifier elimination in an expanded language, and is hence decidable. He gave a model-theoretical argument, which provides no apparent bounds on the complexity of a decision procedure. We provide a syntactical argument that yields a procedure that is primitive recursive, although not elementary. In particular, we show that it is possible to eliminate a single block of existential quantifiers in time $2^0_{O(n)}$, where n is the length of the input formula and 2^x_k denotes k-fold iterated exponentiation.
© 2006 Elsevier B.V. All rights reserved.

Keywords: Real closed fields; Quantifier elimination

1. Introduction

Consider the theory of real closed fields, in a language with 0, 1, $+$, $-$, $\times$, $<$. Extend the language with a predicate A which, in the intended interpretation, holds of the powers of two, $2^{\mathbb{Z}}$. Adopting the obvious conventions and abbreviations, extend the theory by adding the following axioms:

- $\forall x \ (A(x) \rightarrow x > 0)$
- $\forall x, y \ (A(x) \rightarrow (A(y) \leftrightarrow A(xy)))$
- $A(2) \wedge \forall x \ (1 < x < 2 \rightarrow \neg A(x))$
- $\forall x \ (x > 0 \rightarrow \exists y \ (A(y) \wedge y \leq x < 2y))$.

The first two imply that the A picks out a multiplicative subgroup of the positive elements. In [5], van den Dries showed that the resulting theory admits quantifier elimination in an expanded language. As a result, it is complete and decidable, and, in particular, axiomatizes the real numbers with a predicate for the powers of two.

The theory we have just described includes not only the theory of real closed fields, but also, via an interpretation of integers as exponents, Presburger arithmetic. Thus, van den Dries's result is particularly interesting in that it subsumes

* Corresponding author.
   E-mail addresses: avigad@cmu.edu (J. Avigad), yimuy@andrew.cmu.edu (Y. Yin).
   URL: http://www.andrew.cmu.edu/user/avigad (J. Avigad).

two of the most important decidability results of the twentieth century. In recent years, this result has been extended in various directions (see, for example, [10,7]).

To establish quantifier elimination, van den Dries gave a model-theoretical argument. The proof does not provide an explicit procedure, nor does it provide a bound on the length of the resulting formula. Here, we present a proof that makes use of nested calls to a quantifier-elimination procedure for real closed fields, yielding a procedure that is primitive recursive but not elementary. In particular, it requires time $2^0_{O(n)}$ to eliminate a single block of existential quantifiers, or even a single existential quantifier, where n is the length of the input formula and 2^0_k denotes a stack of k exponents. Thus, the best bound we can give on the time complexity of the full quantifier-elimination procedure involves $O(n)$ iterates of the stack-of-twos function. We leave it as an open question as to whether one can avoid such nesting and, say, obtain elementary bounds for the elimination of a single existential quantifier.

In Section 2, we describe the extension of the theory above that admits elimination of quantifiers. Our method of eliminating an existential quantifier proceeds in two steps: first, we eliminate that quantifier in favour of a multiple existential quantifiers over powers of two (the number of which is bounded by the length of the original formula); then we successively eliminate each of these. The first step is described in Section 2. In Section 3, we prove a number of lemmas that fill out the relationship between the powers of two and the underlying model of real closed fields in a model of the relevant theory; this contains the bulk of the syntactical and algebraic work. In Section 4, we use these results to carry out the second step. Finally, in Section 5, we show that our procedure satisfies the complexity bounds indicated above.

2. The first step

Expand the language of real closed fields to include a unary function λ and a unary predicate D_n for each natural number $n \geq 1$. Let T be the theory given by the axioms above together with the following:

- $D_n(x) \leftrightarrow \exists y \, (A(y) \wedge y^n = x)$
- $\forall x \, (x \leq 0 \rightarrow \lambda(x) = 0)$
- $\forall x \, (x > 0 \rightarrow A(\lambda(x)) \wedge \lambda(x) \leq x < 2\lambda(x))$.

In the standard interpretation, λ maps negative real numbers to 0 and rounds positive reals down to the nearest power of two, and D_n holds of numbers of the form 2^i where i is an integer divisible by n.[1] Note that A and D_1 are equivalent; we will treat them as the same symbol and use the two notations interchangeably.

Our goal is to prove the following:

Theorem 2.1. *T admits quantifier elimination.*

This is Theorem II of [5]. Henceforth, by 'formula', we mean 'formula in the language of T'. We will use $\vec{x}$ to denote a sequence of variables $x_0, x_1, \ldots, x_{k-1}$, and we will use notation like $A(\vec{x})$ to denote $A(x_0) \wedge A(x_1) \cdots \wedge A(x_{k-1})$.

To eliminate quantifiers from any formula it suffices to be able to eliminate a single existential quantifier, i.e. transform a formula $\exists x \, \varphi$, where φ is quantifier free, to an equivalent quantifier-free formula. Since $\exists x \, (\varphi \vee \psi)$ is equivalent to $\exists x \, \varphi \vee \exists x \, \psi$, we can always factor existential quantifiers through a disjunction. In particular, since any quantifier-free formula can be put in disjunctive normal form, it suffices to eliminate existential quantifiers from conjunctions of atomic formulas and their negations. Also, since $\exists x \, (\varphi \wedge \psi)$ is equivalent to $\exists x \, \varphi \wedge \psi$ when x is not free in ψ, we can factor out any formulas that do not involve x. Furthermore, whenever we can prove $\forall x \, (\theta \vee \eta)$, $\exists x \, \varphi$ is equivalent to $\exists x \, (\varphi \wedge \theta) \vee \exists x \, (\varphi \wedge \eta)$; so we can 'split across cases' as necessary. We will use all of these facts freely below.

In [5], van den Dries established quantifier elimination by demonstrating the following two facts:

(1) Every model of $T^{\forall}$, the universal fragment of T, has a 'T-closure'; in other words, every model $\mathcal{M}$ of $T^{\forall}$ can be extended to a model of T which can be embedded, over $\mathcal{M}$, into any other T-extension of $\mathcal{M}$.

(2) If $\mathcal{M}$ is a proper substructure of $\mathcal{N}$ and both are models of T, there is some $b \in \mathcal{N} - \mathcal{M}$ such that $\mathcal{M}(b)$, the model of $T^{\forall}$ generated by $\mathcal{M} \cup \{b\}$, can be embedded into an elementary extension of $\mathcal{M}$.

---

[1] For parsimony, 0 can be defined as $1 - 1$ and $A(x)$ by $x > 0 \wedge \lambda(x) = x$. In the next section, we will see that the division symbol is another inessential addition to the language. But in contrast to q.e. for real closed fields, one can't eliminate $-$ in terms of $+$; for example, the quantifier-free formula $A(x - y)$, if replaced by $\exists z \, (z + y = x \wedge A(z))$, would have no quantifier-free equivalent.

The novelty of this test, as compared to more common ones (see e.g. [12]), lies in the prover's right to choose an appropriate b in the second clause (see also the discussion in [6]). This clause implies that any existential formula with parameters from $\mathcal{M}$ that is true in the T-closure of $\mathcal{M}(b)$ is true in $\mathcal{M}$; the test works because this clause can be iterated in a countable model to obtain a sequence of T-extensions $\mathcal{M} = \mathcal{M}_0 \subseteq \mathcal{M}_1 \subseteq \mathcal{M}_2 \cdots \subseteq \mathcal{N}$ that eventually picks up every element of $\mathcal{N}$, so any existential formula with parameters from $\mathcal{M}$ true in $\mathcal{N}$ is true in $\mathcal{M}$. On the syntactic side, this iteration translates to the simple observation that to eliminate a single existential quantifier from an otherwise quantifier-free formula, it suffices to eliminate additional existential quantifiers from an equivalent existential formula. Thus, our effective proof is based on the following two lemmas:

Lemma 2.2. *Every formula of the form $\exists w\ \psi$, with ψ quantifier free, is equivalent to a disjunction of formulas of the form $\exists \vec{x}\ (A(\vec{x}) \wedge \varphi)$, with φ quantifier free.*

Lemma 2.3. *Every formula of the form $\exists x\ (A(x) \wedge \varphi)$, with φ quantifier-free, is equivalent to a formula that is quantifier free.*

The remainder of this section is devoted to proving the first of these two lemmas. The next lemma explains why the new existentially quantified variables are helpful.

Lemma 2.4. *Every existential formula is equivalent, in T, to an existential formula in which λ does not occur and the predicates D_i are applied only to variables.*

Proof. First, replace $\ldots D_i(t) \ldots$ by $\exists z\ (z = t \wedge \cdots D_i(z) \ldots)$. Then, iteratively simplify terms involving λ, noting that $\psi(\lambda(t))$ is equivalent to

$$(t \leq 0 \wedge \psi(0)) \vee \exists z\ (A(z) \wedge z \leq t < 2z \wedge \psi(z)),$$

and that the existential quantifier can be brought to the front. $\quad\square$

Thus to prove Lemma 2.2, we are reduced to showing that when ψ is quantifier free, λ does not occur in ψ, and the predicates D_i occurring in ψ are applied only to variables, the formula $\exists \vec{x}\ \psi$ is equivalent to one of the form $\exists \vec{x}\ (A(\vec{x}) \wedge \varphi)$, where φ is quantifier free. In general, $\exists x\ \theta(x)$ is equivalent to

$$\exists x > 0\ \theta(x) \vee \theta(0) \vee \exists x > 0\ \theta(-x).$$

Moreover, assuming $x > 0$, any subformula of the form $D_i(-x)$ is equivalent to falsity. So, across a disjunction, we are reduced to proving the claim for formulas of the form $\exists \vec{x} > 0\ \psi(\vec{x})$, where ψ satisfies the criteria above.

In T we can factor out the greatest power of two from any positive x, i.e. we can prove

$$x > 0 \leftrightarrow \exists y\ \exists z\ (A(y) \wedge 1 \leq z < 2 \wedge x = yz).$$

Since we have $1 \leq z < 2 \leftrightarrow (z = 1 \vee 1 < z < 2)$, we can transform our formula into a disjunction of formulas of the form

$$\exists \vec{y}, \vec{z}\ (A(\vec{y}) \wedge 1 < \vec{z} < 2 \wedge \psi)$$

where ψ once again meets the criteria above, except that the predicates D_i are applied to expressions of the form yz. When $1 < z < 2$, each $D_i(yz)$ is false, so we can rewrite the formula above as

$$\exists \vec{y}\ (A(\vec{y}) \wedge \theta \wedge \exists \vec{z}\ \eta)$$

where θ is a conjunction of predicates of the form $D_n(y)$ and negations of such, and $\exists \vec{z}\ \eta$ is in the language of real closed fields. We can therefore replace $\exists \vec{z}\ \eta$ by a quantifier-free formula, using any q.e. procedure for real closed fields.

3. Reasoning about powers of two

Our goal in this section is to establish some general relationships between the powers of two in a model of our theory, T, and the underlying real closed field.

Definition 3.1. Let φ be a quantifier-free formula. We say φ is *simple in x* if the following hold:

(1) every equality or inequality occurring in φ is either of the form $p(x) = 0$ or $q(x) > 0$, where $p(x)$, $q(x)$ are polynomials in x; that is, they are of the form $\sum_{i \le n} s_i x^i$ where each s_i is a term that does not involve x.
(2) For every atomic formula $D_n(t)$ occurring in φ, either t does not contain x or t is of the form $2^r x$ for some integer r such that $0 \le r < n$.

The main goal of this section is to prove the following proposition:

Proposition 3.2. *Let φ be any quantifier-free formula. Then there is a quantifier-free formula φ' such that φ' is simple in x and T proves $A(x) \to (\varphi \leftrightarrow \varphi')$.*

In semantic terms, this says the following: let $\mathcal{N}$ be any model of T, let $\mathcal{M} \subseteq \mathcal{N}$ be a model of $T^{\forall}$, and let x be a power of two in $\mathcal{N}$. Then the structure of $\mathcal{M}(x)$ is completely determined by the structure of $\mathcal{M}$, the structure of $\mathcal{M}(x)$ as an ordered ring, and the divisibility properties of the exponent of x.

First, we need to note some easy facts about λ and the predicates D_i.

Lemma 3.3. *For any n, T proves*

$$0 < u < x \le 2^n u \wedge A(x) \to (x = 2\lambda(u) \vee \cdots \vee x = 2^n \lambda(u)).$$

Lemma 3.4. *For any n, T proves*

$$A(x) \to D_n(x) \vee D_n(2x) \vee \cdots \vee D_n(2^{n-1}x).$$

Although we have not included the division symbol in the language of T, we can define the function r/s by making $x/y = z$ equivalent to $x = yz \vee (y = 0 \wedge z = 0)$. In the proof of Proposition 3.2, it will be useful to act as though the division symbol is part of the language. The next few lemmas show that if θ is any quantifier-free formula in the expanded language with division, there is a quantifier-free formula θ' in the language without division such that $T \vdash \theta \leftrightarrow \theta'$.

Lemma 3.5. *From the hypotheses $0 < x$ and $0 < y$, T proves*

$$x\lambda(y) < y\lambda(x) \to \lambda(x/y) = \lambda(x)/2\lambda(y)$$

and

$$x\lambda(y) \ge y\lambda(x) \to \lambda(x/y) = \lambda(x)/\lambda(y).$$

Proof. An easy calculation using the axioms for λ shows that if $x/y < \lambda(x)/\lambda(y)$, then $\lambda(x/y) = \lambda(x)/2\lambda(y)$; and otherwise, $\lambda(x/y) = \lambda(x)/\lambda(y)$. $\square$

Lemma 3.6. *If θ is any quantifier-free formula involving the division symbol, there is a quantifier-free formula θ' in which the division symbol does not occur in the scope of λ, such that $T \vdash \theta \leftrightarrow \theta'$.*

Proof. This can be done by iterating the previous lemma. To measure the nesting of λ's and division symbols, we define the 'λ-depth of the division symbol in t', $\Lambda^{\div}(t)$, recursively, as follows:

(1) $\Lambda^{\div}(t) = 0$ if the division symbol does not occur in the scope of λ in t;
(2) if t is $t_1 + t_2$, $t_1 - t_2$, $t_1 \times t_2$, or t_1/t_2, then $\Lambda^{\div}(t) = \max\{\Lambda^{\div}(t_1), \Lambda^{\div}(t_2)\}$;
(3) assuming the division symbol occurs in t, $\Lambda^{\div}(\lambda(t)) = \Lambda^{\div}(t) + 1$.

The previous lemma shows that, using a case disjunction over the possibilities for the signs of the numerator and denominator, we can eliminate one term t such that the λ-depth of the division symbol in t is maximal, in favour of terms in which the λ-depth of the division symbol is smaller. Lemma 3.6 follows, by a primary induction on this maximal depth, and a secondary induction on the number of terms of this depth. $\square$

Lemma 3.7. $T \vdash A(x) \wedge A(y) \to (D_n(x/y) \leftrightarrow \bigvee_{i<n}(D_n(2^i x) \wedge D_n(2^i y)))$.

Proof. The right-to-left direction is easy: if $z^n = 2^i x$ and $w^n = 2^i y$ then $(z/w)^n = x/y$. Proving the other direction is not much more difficult, using Lemma 3.4. $\square$

Proposition 3.8. *Let θ be any quantifier-free formula involving division. Then there is a quantifier-free formula θ' that does not involve division, such that $T \vdash \theta \leftrightarrow \theta'$.*

Proof. Using Lemma 3.6, we can assume that division does not occur in the scope of any λ in θ. So each atomic formula $D_n(t)$ can be put in the form $D_n(r/s)$, where the division symbol does not occur in r and s. Across a case disjunct, we can assume r and s are positive. Then $D_n(r/s)$ is equivalent to

$$\lambda(r/s) = r/s \wedge D_n(\lambda(r/s)).$$

Using Lemma 3.5, we can replace $\lambda(r/s)$ by either $\lambda(r)/\lambda(s)$ or $\lambda(r)/2\lambda(s)$. Then using Lemma 3.7 we can replace $D_n(\lambda(r)/\lambda(s))$ or $D_n(\lambda(r)/2\lambda(s))$ by a disjunction in which the division symbol does not occur.

Once all divisibility symbols are removed from the λ's and D_n's, we can clear division from the remaining equalities and inequalities by multiplying through. $\quad\square$

It therefore suffices to prove Proposition 3.2 where φ' is a quantifier-free formula in the expanded language with the division symbol. The next few lemmas, then, make use of this expanded language.

Lemma 3.9. *Let $p(x)$ be the term $\sum_{i \leq n} a_i x^i$. Then there is a sequence of quantifier-free formulas $\theta_0, \ldots, \theta_{m-1}$ such that T proves*

$$A(x) \wedge p(x) > 0 \rightarrow \bigvee_{k < m} \theta_k,$$

where each θ_k is of one of the following forms:

- $\lambda(p(x)) = 2^r \lambda(a_i) x^i$ *for some* $-1 \leq r \leq n$,
- $x^e = \frac{2^r \lambda(a_i)}{\lambda(-a_j)}$ *or* $x^e = \frac{2^r \lambda(-a_j)}{\lambda(a_i)}$, *for some* e, i, j, *and* r *such that* $1 \leq e \leq n$, $0 \leq i, j \leq n$, *and* $-(n+1) \leq r \leq (n+1)$.

Proof. Argue in T. Using a disjunction on all possible cases, we can write $p(x)$ as $a_i x^i + a_j x^j + \hat{p}(x)$, where $a_i x^i$ is the largest summand and $a_j x^j$ the least summand. Note that we have $a_i x^i > 0$, $p(x) \leq (n+1) a_i x^i$, and

$$p(x) - a_i x^i = a_j x^j + \hat{p}(x) \geq n a_j x^j.$$

We now distinguish between two cases, depending on whether $p(x)$ is roughly the same size as $a_i x^i$ or sufficiently smaller.

In the first case, suppose we have $p(x) \geq (a_i x^i)/2$. This means we have:

$$(a_i/2)x^i \leq p(x) \leq (n+1) a_i x^i \leq 2^n a_i x^i$$

which yields

$$(\lambda(a_i)/2)x^i \leq \lambda(p(x)) \leq 2^n \lambda(a_i) x^i.$$

This yields a disjunction of clauses of the first type, by Lemma 3.3.

In the second case, we have $p(x) < (a_i x^i)/2$ and $i \neq j$. This means that $a_j x^j$ must be negative and roughly comparable to $a_i x^i$ in absolute value. That is, we have $a_j < 0$ and

$$(a_i/2)x^i < a_i x^i - p(x) \leq -n a_j x^j,$$

and so

$$(a_i/(-a_j))x^{i-j} \leq 2n \leq 2^n.$$

Also, $p(x) > 0$ implies $n a_i x^i \geq -a_j x^j$, which yields

$$0 < 2^{-n} < 1/n \leq (a_i/(-a_j))x^{i-j}.$$

Combining these, we have $2^{-n} < (a_i/(-a_j))x^{i-j} \leq 2^n$. Using Lemma 3.5 and Lemma 3.3 we get a disjunction of clauses of the second type. $\quad\square$

Lemma 3.10. *In Lemma* 3.9, *if the assumption is changed to $A(x) \land p(x) = 0$, then in the conclusion we can assume that each θ_k is of the second form.*

Proof. This is exactly as in the second case of the previous proof. □

Lemma 3.11. *In the conclusion of Lemma* 3.9, *we may demand that each θ_k is of the form $\lambda(p(x)) = sx^i$ for some $0 \le i \le n$ and some term s that does not contain x.*

Proof. The proof is by induction on the degree of x in $p(x)$. The lemma is trivial if the degree of x in $p(x)$ is 0.

Now assume that the degree of x in $p(x)$ is n and the lemma holds whenever the degree is less than n. By Lemma 3.9, T proves a disjunction $\bigvee \sigma_l$, with σ_l of one of those two forms. Each σ_l of the first form there is already as required. For each σ_l of the second form, consider a new term $\hat{p}(x)$, which is obtained by substituting the right-hand side of σ_l for x^e in $p(x)$. Notice that the degree of x in $\hat{p}(x)$ is less than n, and clearly T proves $p(x) = \hat{p}(x) \land \hat{p}(x) > 0$. By the inductive hypothesis we may replace σ_l in $\bigvee \sigma_l$ by a disjunction $\bigvee \theta_k$ which is of the required form. □

As was the case with the division symbol, we will iteratively 'squeeze' x's out from within the λ symbols. Thus we introduce the following definitions:

Definition 3.12. Let t be a term. Define the λ-*depth of x in t*, $\Lambda(x, t)$, recursively, as follows:

(1) $\Lambda(x, t) = 0$ if x is not in the scope of any λ;
(2) if t is $t_1 + t_2$, $t_1 - t_2$, $t_1 \times t_2$, or t_1/t_2, then $\Lambda(x, t) = \max\{\Lambda(x, t_1), \Lambda(x, t_2)\}$;
(3) if t is $\lambda(t_1)$ and t_1 contains x, then $\Lambda(x, t) = \Lambda(x, t_1) + 1$.

Definition 3.13. Let φ be a formula. Define the λ-*depth of x in φ* by

$$\Lambda(x, \varphi) = \max\{\Lambda(x, t) : t \text{ is a term that contains } x \text{ and occurs in } \varphi\}.$$

Lemma 3.14. *Let φ be any quantifier-free formula. Then there is a quantifier-free formula φ' such that $T \vdash A(x) \to (\varphi \leftrightarrow \varphi')$, and $\Lambda(x, \varphi') = 0$.*

Proof. The proof is by induction on the λ-depth of x in φ. The lemma is trivial if $\Lambda(x, \varphi) = 0$.

Assume $\Lambda(x, \varphi) = n > 0$ and the lemma holds for every quantifier-free formula ψ if $\Lambda(x, \psi) < n$. Let $\lambda(p_0), \ldots, \lambda(p_{m-1})$ be all the different terms in φ with $\Lambda(x, p_i) = 0$ for all $i < m$. Across a case disjunction we can assume $p_i > 0$ for all $i < m$, since otherwise we can replace $\lambda(p_i)$ by 0. By Lemma 3.6, we may assume that each p_i is a polynomial in x. By Lemma 3.11, T proves $\varphi \leftrightarrow \bigvee(\tau_l \land \sigma_l)$, where each τ_l is of the form $\bigwedge_{i<m} \lambda(p_i(x)) = s_i x^{j_i}$, and each σ_l is obtained by substituting $s_i x^{j_i}$ for $\lambda(p_i)$ in φ. Clearly T proves

$$A(x) \to (\lambda(p_i(x)) = s_i x^{j_i} \leftrightarrow A(s_i) \land s_i x^{j_i} \le p_i(x) < 2s_i x^{j_i}).$$

Now since $\Lambda(x, \sigma_l) < n$, we may apply the inductive hypothesis to each σ_l and the lemma is proved. □

Lemma 3.15. *Let p be a term such that $\Lambda(x, p) = 0$. Then for any n there is a sequence of terms p_k such that*

- *T proves $A(x) \land p > 0 \to (D_n(p) \leftrightarrow \bigvee(p = p_k \land D_n(p_k)))$,*
- *each p_k is of the form sx^i, where s is a term that does not contain x.*

Proof. Using Lemma 3.7, we can assume that p is a polynomial in x. We can replace $D_n(p)$ by $p = \lambda(p) \land D_n(\lambda(p))$, and then by Lemma 3.11, across a disjunction we may replace $\lambda(p)$ in each disjunct by a term of the form sx^i, where s does not contain x. (Note that here no formulas like the τ_l's in the previous lemma are needed.) □

Lemma 3.16. *Let s be a term that does not contain x. Then for any n, i there is a sequence of formulas θ_k such that T proves*

$$A(x) \to (D_n(sx^i) \leftrightarrow \bigvee \theta_k),$$

and each θ_k is of the form $D_n(2^w s) \land D_n(2^r x)$ for some $0 \le w, r < n$.

Proof. Since for each n, from the assumption $A(x)$, T proves $\bigvee_{j<n} D_n(2^j x)$, it is straightforward to see that $D_n(sx^i)$ is equivalent to a disjunction each of whose disjuncts is of the specified form. □

We are finally ready to prove Proposition 3.2.

Proof. Given φ, first use Lemma 3.14 to eliminate x from the scope of any λ. Then use Lemma 3.15 to ensure the atomic formulas involving D_n are in the form $D_n(sx^i)$, where s does not involve x. (This will require splitting across cases depending on whether $p > 0$ or $p \leq 0$; in the latter case, $D_n(p)$ is equivalent to $\bot$.) Finally, use Lemma 3.16 to ensure that all the atomic formulas involving D_n are in the required form. □

We close with some consideration about the predicates D_n which are analogous to considerations that arise in the context of quantifier elimination for Presburger arithmetic. Remember that when n is a positive integer and s is a non-negative integer, $D_n(2^s x)$ asserts, in the intended interpretation, that x is equal to 2^t for some integer t, and n divides $s + t$; in other words, the exponent of x is congruent to $-s$ modulo n. Let θ be any boolean combination of predicates of the form $D_n(2^s x)$, and let M be the least common multiple of these various n. Then in T one can show that there is an x satisfying θ if and only if for any w satisfying $A(w)$ we have

$$\theta(w) \vee \theta(2w) \vee \theta(4w) \vee \cdots \vee \theta(2^{M-1}w),$$

and, in particular, if and only if

$$\theta(1) \vee \theta(2) \vee \theta(4) \vee \cdots \vee \theta(2^{M-1}).$$

Moreover, T can decide the truth or falsity of this last sentence. So we have:

Lemma 3.17. *With θ and M as above, either T proves $\forall x \ \neg\theta$, or it proves*

$$\forall u \ (0 < u \rightarrow \exists x \ (u \leq x < 2^M u \wedge \theta)).$$

4. Eliminating a quantifier over powers of two

We are now ready to prove Lemma 2.3, which asserts that every formula of the form $\exists x \ (A(x) \wedge \varphi)$, with φ quantifier free, is equivalent to a formula that is quantifier-free. By Proposition 3.2, we can assume that φ is simple, which is to say, x does not occur in the scope of any λ and all divisibility assertions involving x are of the form $D_n(2^r x)$. Put φ in disjunctive normal form, replace negated equalities $s \neq t$ by $s < t \vee t < s$, and replace negated inequalities $s \not< t$ by $t < s \vee t = s$. Rewrite equalities and inequalities so that they are of the form $p(x) = 0$ and $q(x) > 0$, where $p(x)$ and $q(x)$ are polynomials in x. Factoring existential quantifiers through disjunctions and getting rid of atomic formulas that do not depend on x, we are reduced to eliminating quantifiers of the form $\exists x \ (A(x) \wedge \varphi)$ where φ is a conjunction of formulas of the following types:

- $p(x) = 0$, where p is a polynomial,
- $q(x) > 0$, where q is a polynomial,
- $D_n(2^r x)$, where $0 \leq r < n$, or
- $\neg D_n(2^r x)$, where $0 \leq r < n$.

Splitting across a disjunction, we can assume that in a conjunct of the form $p(x) = 0$, not all the coefficients are zero. By Lemma 3.10, we can assume that one of the conjuncts is of the form $x^e = s$, where x does not occur in s. In that case, each conjunct $D_n(2^r x)$ is equivalent to $D_{ne}(2^{re} x^e)$ and hence $D_{ne}(2^{re} s)$ (and $A(x)$, in particular, is equivalent to $D_e(s)$). But now x no longer occurs in these formulas, and so they can be brought outside the scope of the existential quantifier. The resulting existential formula is then essentially in the language of real closed fields. By this last phrase we mean that it is of the form $\exists x \ \alpha(x, t_0, \ldots, t_{k-1})$, where $\alpha(x, y_0, \ldots, y_{k-1})$ is in the language of real closed fields. Treating the terms $t_0, \ldots, t_{k-1}$ in the expanded language as parameters, we can therefore replace it by an equivalent quantifier-free formula using any q.e. procedure for real closed fields.

We are thus reduced to eliminating an existential quantifier of the form

$$\exists x \ \left(\bigwedge q_i(x) > 0 \wedge \theta(x) \right) \tag{1}$$

where θ is a conjunction of formulas of the form $D_n(2^r x)$ and negations of such that includes at least the formula $A(x)$. By Lemma 3.17, either T proves that θ is false for every x, or there is a natural number M such that T proves that for any $u > 0$, that θ is satisfied by some x in the interval $[u, 2^M u]$. In the first case, T proves that formula (1) is false. So we only have to worry about the second case. Fix such an M for the remainder of the discussion.

Arguing in T, suppose formula (1) holds. There are two possibilities: either there is a 'large' interval on which $\bigwedge q_i(x) > 0$, that is, an interval of the form $[u, 2^M u]$; or there is an x satisfying $A(x) \wedge \bigwedge q_i(x) > 0 \wedge \theta$, but it is trapped between a u and a v with $q_i(u) = 0$ for some i, $q_j(v) = 0$ for some j, and $v < 2^M u$. Thus formula (1) is equivalent to a disjunction of the formula

$$\exists u > 0 \, \forall x \, \left(u \leq x \leq 2^M u \to \bigwedge q_i(x) > 0 \right)$$

and the formulas

$$\exists u > 0 \, \left(q_j(u) = 0 \wedge \exists x \, \left(u < x \leq 2^M u \wedge \bigwedge q_i(x) > 0 \wedge \theta(x) \right) \right)$$

for the various j. To see this, note that if formula (1) holds, then by the previous discussion one of these formulas holds; and conversely, each of these formulas implies (1).

The first of these formulas is essentially in the language of real closed fields, so these quantifiers can be eliminated. The second formula is equivalent to:

$$\exists u_1, u_2 \, \left(A(u_1) \wedge 1 \leq u_2 < 2 \wedge q_j(u_1 u_2) = 0 \wedge \exists x \, \left(u_1 < x \leq 2^M u_1 \wedge \bigwedge q_i(x) > 0 \wedge \theta(x) \right) \right).$$

In this case, we can replace the inner existential quantifier over x by a disjunction, so that the entire formula is equivalent to a disjunction of formulas of the form

$$\exists u_1, u_2 \, \left(A(u_1) \wedge 1 \leq u_2 < 2 \wedge q_j(u_1 u_2) = 0 \wedge \bigwedge \hat{q}_i(u_1) > 0 \wedge \hat{\theta}(u_1) \right),$$

where each $\hat{q}_i(u_1)$ is $q_i(2^r u_1)$ for some r such that $1 \leq r \leq M$, and similarly for $\hat{\theta}(u_1)$. In particular, $\hat{\theta}(u_1)$ is a conjunction of formulas of the form $D_i(2^r u_1)$, and their negations.

Think of $q_j(u_1 u_2)$ as a polynomial in u_1 with coefficients of the form $s u_2^n$, where s does not involve u_1 or u_2. By Lemma 3.10, across a disjunction we may add a clause of the form $u_1^e = 2^r \lambda(s u_2^n)/\lambda(t u_2^m)$. Splitting on cases of the form $2^l \leq u_2^h < 2^{l+1}$ we can simplify each of these to an expression of the form $u_1^e = 2^k \lambda(s)/\lambda(t)$ for some integer k. By Lemma 3.17, $A(u_1) \wedge \hat{\theta}(u_1)$ is equivalent to a formula $\bar{\theta}$ which now involves neither u_1 nor u_2, and hence can be brought outside the existential quantifier. We are thus reduced to eliminating quantifiers from a formula of the form:

$$\exists u_1, u_2 \, \left(1 \leq u_2 < 2 \wedge u_1^e = 2^k \lambda(s)/\lambda(t) \wedge 2^l \leq u_2^h < 2^{l+1} \wedge q_j(u_1 u_2) = 0 \wedge \bigwedge \hat{q}_i(u_1) = 0 \right).$$

We can eliminate these quantifiers using a q.e. procedure for real closed fields. This completes the proof of Lemma 2.3, and hence the proof of our main theorem, Theorem 2.1.

Note that there is nothing special about the number 2 in our quantifier elimination procedure: inspection of the proofs shows that the arguments go through unchanged for any real algebraic number $\alpha > 1$. There are various ways to represent the real algebraical numbers; for example, we can represent α by providing a polynomial, $p(x)$, of which it is a root, together by a pair of rational numbers u and v isolating α from the other roots of p. In that case, we simply replace 2 by a new constant, c, in the axioms, and then add the following:

- $p(c) = 0$
- $u < c < v$.

As noted in [7], this implies that the resulting theory is decidable. To see this, it suffices to see that any quantifier-free sentence φ is decidable. But we can do this using the decision procedure for real closed fields to iteratively compute the values of $\lambda(t)$ for any t involving the field operations and c, and then to determine the truth of terms of atomic formulas $D_n(t)$. (For explicit algorithms for computing with real algebraical numbers, see [2,11].)

5. Complexity analysis

In this section we establish an upper bound on the complexity of our elimination procedure.

For the theory of real closed fields, the best known upper bound for a quantifier-elimination procedure, in terms of the length of the input formula, is $2^{2^{O(n)}}$. This is originally due to Collins [4], and, independently, Monk and Solovay. There are more precise bounds that depend on various parameters, such as the number of quantifier alternations and the degrees of the polynomials in the formula; see, for example, [1,2]. In particular, a block of existential quantifiers can be eliminated in time $2^{O(n)}$. The best lower bound for the full quantifier-elimination procedure is $2^{O(n)}$, by Fischer and Rabin [9], and applies even to just the additive fragment. The best upper bound for Presburger arithmetic is $2_3^{O(n)}$ (see [8,13]) and is essentially sharp (see [14]).

Our bounds are far worse. Consider what our procedure does when given a formula with a single block of existential quantifiers:

(1) First, replace this by a disjunction of formulas of the form
$$\exists \vec{y} \ (A(\vec{y}) \wedge \exists \vec{z} \ (1 < \vec{z} < 2 \wedge \psi))$$
where ψ is in the language of real closed fields.

(2) Then, use an elimination procedure for real closed fields to eliminate the quantifiers $\exists \vec{z}$.

(3) Successively eliminate the innermost quantifier over a power of two, as follows:

   (a) Call the relevant formula $\exists x \ (A(x) \wedge \varphi)$. Apply Proposition 3.2, to reduce φ to a formula that is simple in x.

   (b) Put the new φ in disjunctive normal form, split across a disjunction, and remove atomic formulas that do not involve x, so that each formula is of the form
   $$\exists x \ (A(x) \wedge \bigwedge p_i(x) = 0 \wedge \bigwedge q_j(x) = 0 \wedge \theta)$$
   where θ is a conjunction of formulas of the form $D_n(2^r x)$ and negations of such, and in each disjunction where a disjunct of the form $p(x) = 0$ occurs, we can assume p is not identically 0.

   (c) In each disjunct where a conjunct of the form $p(x) = 0$ occurs, apply Lemma 3.10, factor out the divisibility predicates, D_n, and call a quantifier-elimination procedure for real closed fields.

   (d) In the remaining disjuncts, again, split across a disjunct; in one case, we call a quantifier-elimination procedure for real closed fields right away; in another, we expand a bounded existential quantifier into a disjunction, and then call the elimination procedure for real closed fields.

Note that each iteration of the inner loop (3) requires at least one call to a quantifier-elimination procedure for real closed fields. Each of these calls can be carried out in time, say, $2^{2^{O(n)}}$, where n is the length of the relevant formula. But then the next iteration of the loop will involve calls to the q.e. procedure for real closed fields on a formula that is potentially much longer. Thus, part (3) of the procedure requires an exponential stack of Cm twos, for some constant C, where m is the number of existential quantifiers over powers of two that need to be eliminated.

In this section, we will confirm that such an upper bound can be obtained. To that end, it is sufficient to show that each pass of the inner loop is elementary, which is to say, it can be computed in time bounded by some fixed stack of exponents to the base 2. Note that after the first step, the number of quantifiers over powers of two is bounded by the length of the original formula (in fact, it is bounded by the number of A's and λ's in the original formula). Thus our procedure for eliminating a block of existential quantifiers runs in time $2_{O(n)}^0$, where n is the length of the original formula.

We have been unable to eliminate this nesting of calls to a procedure for real closed fields. Efficient procedures for this latter theory avoid putting formulas in disjunctive normal form; for example, Collins's cylindrical algebraic decomposition procedure obtains a description of cells, depending on the coefficients, on which a set of polynomials have constant sign. In our setting, suppose we are given a formula $\exists \vec{x} \ (A(\vec{x}) \wedge \eta \wedge \theta)$, where η contains only equalities and inequalities between polynomials, and θ consists of divisibility conditions D_n on the exponents of the x's. One might start by applying Collins's procedure to the polynomials occurring in η. Then, given a description of the various cells (depending on the other parameters in the formula), one needs to determine which cells contain points with coordinates that are powers of two, with exponents satisfying the requisite divisibility conditions. For one-dimensional cells, our procedure relies on a simple disjunction: if the cell is large enough, one is guaranteed a solution, and otherwise one need only test a finite number of cases. For multidimensional cells, however, the situation is more complex, and we do not see how one can proceed except along the lines we have described above. It is thus

an interesting question of whether it is possible to obtain elementary bounds on a procedure for eliminating a single block of quantifiers. Given our failure to do so, we have not taken great pains to bound the number of exponents in the time bound on the inner loop, which would merely improve the constant bound implicit in the $O(n)$.

For the discussion which follows, we define the *length* of a formula in the language of T to be the number of symbols in a reasonable formulation of the first-order language, with the following exception: we count the length of each symbol D_n as n, rather than, say, one plus the binary logarithm of n. This choice is a pragmatic one in that it simplifies the analysis, and our results below then imply the corresponding results for the alternative definition of length. A more refined analysis might take both the length of the formula and a bound on the n's occurring in atomic formulas $D_n(t)$, but that does not seem to help much.

It seems that the most delicate part of our task is showing that one can remove the division symbols, and 'squeeze' variables ranging over powers of two out of the λ symbols that are repeatedly introduced after the first step of the procedure, as required in step (3a). A priori, the procedures described in Section 3 look as though they may be non-elementary. The next few lemmas show that this is not the case, by keeping careful track of the terms and formulas that need to be dealt with in the disjunctions.

Lemma 5.1. *Let t be a term with length l. Then there is a sequence of terms $\langle t_k : k < 2^l \rangle$ such that*

- $T \vdash \bigvee_{k<2^l} t = t_k$,
- *each t_k is of the form r/s, where r and s are division-free terms, and*
- *each t_k has length at most 2^l.*

Proof. This can be proved by a straightforward induction on terms. Suppose t is of the form $t_1 + t_2$, where the length of t_1 is l_1 and the length of t_2 is l_2. By the inductive hypothesis, t is equal to one of at most $2^{l_1} 2^{l_2} \le 2^l$ terms of the form $r_1/s_1 + r_2/s_2$, where $r_1, s_1, r_2,$ and s_2 are division-free, the length of r_1/s_1 is at most 2^{l_1}, and the length of r_2/s_2 is at most 2^{l_2}. But then the length of $(r_1 s_2 + r_2 s_1)/s_1 s_2$ is at most $2(2^{l_1} + 2^{l_2}) < 2^l$, as required.

If t is of the form $\lambda(t_1)$, the claim follows from the inductive hypothesis, using Lemma 3.6. The other cases are similar. □

Lemma 5.2. *Let φ be a quantifier-free formula with length l. Then there is a quantifier-free division-free formula φ' with length $2^{O(l)}$ such that $T \vdash \varphi \leftrightarrow \varphi'$.*

Proof. Enumerate all the different terms $t_0, \ldots, t_{m-1}$ in φ such that, for each $i < m$, s_i is not a proper subterm of any term in φ. Using the above lemma we can have a sequence of quantifier-free formulas φ_j for $j < 2^l$ each of which is obtained by replacing each t_i with an appropriate term and therefore has length less than 2^l. Notice that for each φ_j, as indicated in Lemma 3.6, there are some division-free atomic formulas that T used to derive the equalities in question. Clearly for each φ_j there are less than l such atomic formulas, each of which has length less than $2^{O(l)}$. Let σ_j be the conjunction of them all. Let φ' be the formula $\bigvee_{j<2^l} (\varphi_j \wedge \sigma_j)$. The length of φ' is again bounded by $2^{O(l)}$, and clearly $T \vdash \varphi \leftrightarrow \varphi'$.

Finally, we need to clear denominators from atomic formulas of the form $r/s < t/u$ and $r/s = t/u$, and deal with atomic formulas of the form $D_n(r/s)$. The first two require a disjunction over cases, depending on whether denominators are positive, negative, or zero. The third set of atomic formulas is handled as described in the proofs of Lemmas 3.7 and 3.8. But each atomic formula occurring in a disjunct occurs to an atomic formula in the original formula, φ, and there are at most l of these. It is not hard to verify that the corresponding increase in length can be absorbed into the bound $2^{O(l)}$. □

Lemma 5.3. *Let $\lambda(t)$ be a term, where the length of t is l and x does not occur in the scope of any division symbol in t. Then there is a sequence of terms $\langle t_k : k < 2^{8l^2 \log l} \rangle$ such that*

- $T \vdash A(x) \wedge t > 0 \rightarrow \bigvee_{k<2^{8l^2 \log l}} (\lambda(t) = t_k)$,
- *each t_k is of the form sx^i, where s is a term that does not contain x and $i < l$,*
- *each t_k has length at most $2^{2^{4l}}$.*

Proof. For any polynomial p in x, clearly the number of possible values of $\lambda(p)$ of the form sx^i, as in Lemma 3.11, depends on the degree n of x in p. So let $f(n)$ denote the number of possible values of $\lambda(p)$. Observe that the value of $\lambda(p)$ is determined in the first case of Lemma 3.9, and when $e = 1$ in the second case. An calculation shows that

there are no more than $(n+1)(n+2)$ possibilities in the first case, no more than $2n(2n+2)$ possibilities in the second case when $e = 1$, and no more than $(n+1)(n-1)2(n+2)$ possibilities for all the remaining values of e. Hence we have the following equation:

$$f(n) \leq (n+1)(n+2) + 2n(2n+2) + (n+1)(n-1)2(n+2)f(n-1).$$

This can be simplified as $f(n) < 10(n+2)^3 f(n-1)$. So we have $f(n) < 2^{8n \log(n+2)}$. Let the length of p be l. Since $n+2 < l$, we have $f(n) < 2^{8l \log l} < 2^{8l^2 \log l}$.

Now the proof proceeds by induction on the λ-depth of x in t. If $\Lambda(x, t) = 0$, then t is a polynomial in x. So we apply the above analysis to t and obtain no more than $2^{8l \log l}$ possible values of $\lambda(t)$ which are all of the form sx^i for some $i < l$. To compute the length of s, only note that each step of the iteration produces a polynomial whose length is no more than the square of the length of the previous polynomial. So we conclude that the length of s is no more than $l^{2^l} < 2^{2^{4l}}$.

Now suppose the lemma holds for each term s with $\Lambda(x, s) < d$, and suppose $\Lambda(x, t) = d$. Enumerate all the different terms $\lambda(s_0), \ldots, \lambda(s_{m-1})$ in t such that $\lambda(s_i)$ is not in the scope of any λ for each $i < m$. Clearly $\Lambda(x, s_i) < n$ for each $i < m$. So by the inductive hypothesis there are less than $2^{8l_i^2 \log l_i}$ possible values for each $\lambda(s_i)$, where l_i is the length of s_i. Since $\sum_{i<m} l_i < l - 1$, there are no more than $2^{8(l-1)^2 \log l}$ possible values for t. Enumerate these possibilities as $\langle t_k : k < 2^{8(l-1)^2 \log l} \rangle$. In each t_k, $\lambda(s_i)$ is replaced by a term of the form sx^j with $j < l_i$. So t_k is a polynomial in x whose degree in x is less than $l - 2$. So there are $2^{8(l-1)^2 \log l} \cdot 2^{8l \log l} \leq 2^{8l^2 \log l}$ possible values for $\lambda(t)$. The length of each t_k is bounded by $2^{2^{4(l-1)}}$, so the length of each possible value of $\lambda(t)$ is bounded by $2^{2^{4(l-1)}} \cdot l \cdot l^{2^l} < 2^{2^{4l}}$. $\square$

Lemma 5.4. *Let φ be a quantifier-free formula with length l. Assume x does not occur in the scope of any division symbol in φ. Then there is a quantifier-free formula φ' with length at most $2^{2^{O(l)}}$ such that φ' is simple in x and $T \vdash A(x) \to (\varphi \leftrightarrow \varphi')$.*

Proof. First we claim there is a quantifier-free formula φ^* with length at most $2^{2^{O(l)}}$ such that

- $T \vdash A(x) \to (\varphi \leftrightarrow \varphi^*)$,
- x does not occur in the scope of any division in φ^*,
- $\Lambda(x, \varphi^*) = 0$.

The proof is essentially the same as the proof of Lemma 5.2, using Lemma 5.3 instead of Lemma 5.1.

Next we need to deal with atomic formulas of the form $D_n(p)$ in φ^*, as shown in Lemma 3.15. So p is a polynomial in x whose degree in x is less than l. So there are at most $2^{2^{O(l)}}$ possible values for $\lambda(p)$, the length of each of which is bounded by $2^{2^{O(l)}}$. So each $D_n(p)$ can be replaced by a disjunction whose length is less than $2^{2^{O(l)}}$. So the bound does not change.

The increase in length in transforming φ^* to a formula that is simple in x, as described in the proof of Lemma 3.16, can be absorbed in the bound $2^{2^{O(l)}}$. $\square$

Lemma 5.5. *Let φ be a quantifier-free formula with length l. Then there is a quantifier-free formula φ' with length at most $2_3^{O(l)}$ such that φ' is simple in x and T proves $A(x) \to (\varphi \leftrightarrow \varphi')$.*

Proof. Immediate by Lemmas 5.2 and 5.4. $\square$

Lemma 5.6. *Each iteration of step 3 can be performed by an elementary function.*

Proof. It is straightforward to verify that the procedure implicit in Lemma 5.5 runs in time polynomial in its output. As a result, step 3(a) is elementary. Step 3(b) is also clearly elementary. In fact, even though putting a formula in disjunctive normal form can result in exponentially many disjuncts, since each disjunct only involves atomic formulas from the original formula, the length of each disjunct is bounded in the length of the original formula.

After step 3(a), the main increase therefore comes from the handling of the cases in (c) and (d), each of which is easily seen to be elementary. Case (c) involves a call to a quantifier-elimination procedure for real closed fields, with a $\forall\exists$ formula; case (d) involves calls to such a procedure, on existential formulas, across a number of disjuncts that is exponential in the length of the original formula. $\square$

Theorem 5.7. *There is a procedure for eliminating a single block of existential quantifiers in theory T in time $2^0_{O(l)}$, where l is the length of the original formula.*

Proof. Steps 1 and 2 are clearly elementary, after which the procedure performs an elementary operation for each quantifier over a power of two. As noted above, the number of such quantifiers can even be bounded by the number of predicates D_n and λ's in the original formula. $\quad\square$

Corollary 5.8. *There is a procedure for eliminating quantifiers in theory T that runs in time bounded by $O(l)$ iterations of the stack-of-twos function, where l is the length of the original formula.*

Proof. Put the formula in prenex form, and iteratively apply the previous theorem to eliminate each block of quantifiers. $\quad\square$

Acknowledgements

References

[1] S. Basu, New results on quantifier elimination over real closed fields and applications to constraint databases, Journal of the ACM 46 (1999) 537–555.

[2] S. Basu, R. Pollack, M.-F. Roy, Algorithms in Real Algebraic Geometry, in: Algorithms and Computation in Mathematics, vol. 10, Springer, Berlin, 2003.

[3] B.F. Caviness, J.R. Johnson (Eds.), Quantifier Elimination and Cylindrical Algebraic Decomposition, in: Texts and Monographs in Symbolic Computation, Springer, Vienna, 1998.

[4] G.E. Collins, Quantifier elimination for real closed fields by cylindrical algebraic decomposition, in: Automata Theory and Formal Languages (Second GI Conf., Kaiserslautern, 1975), in: Lecture Notes in Computer Science, vol. 33, Springer, 1975, pp. 134–183. Reprinted in [3].

[5] L. van den Dries, The field of reals with a predicate for the powers of two, Manuscripta Mathematica 54 (1985) 187–195.

[6] L. van den Dries, Alfred Tarski's elimination theory for real closed fields, Journal of Symbolic Logic 53 (1988) 7–19.

[7] L. van den Dries, A. Günaydin, The field of real and complex numbers with a small multiplicative group, Proceedings of the London Mathematical Society 93 (2006) 43–81.

[8] J. Ferrante, C.W. Rackoff, The Computational Complexity of Logical Theories, in: Lecture Notes in Mathematics, vol. 718, Springer, Berlin, 1979.

[9] M.J. Fischer, M.O. Rabin, Super-exponential complexity of Presburger arithmetic, in: Complexity of Computation (Proc. SIAM-AMS Sympos., New York, 1973), Amer. Math. Soc., Providence, R.I., 1974, pp. 27–41. Reprinted in [3].

[10] H. Friedman, C. Miller, Expansions of o-minimal structures by fast sequences, Journal of Symbolic Logic 70 (2005) 410–418.

[11] B. Mishra, Algorithmic Algebra, Springer, Berlin, 1993.

[12] J.R. Shoenfield, Mathematical logic, Association for Symbolic Logic, Urbana, IL, 2001. Reprint of the 1973 second printing.

[13] V. Weispfenning, The complexity of almost linear Diophantine problems, Journal of Symbolic Computation 10 (1990) 395–403.

[14] V. Weispfenning, Complexity and uniformity of elimination in Presburger arithmetic, in: Proceedings of the 1997 International Symposium on Symbolic and Algebraic Computation, Kihei, HI, ACM, New York, 1997, pp. 48–53. Electronic.