

Lecture Notes on Virtual Substitution & Real Equations

André Platzer

Carnegie Mellon University
Lecture 18

1. Introduction

Cyber-physical systems are important technical concepts for building better systems around us. Their safe design requires careful specification and verification, which this course provides using differential dynamic logic and its proof calculus [Pla08, Pla10, Pla12b]. The proof calculus for differential dynamic logic has a number of powerful axioms and proof rules (especially in [Lecture 5](#), [Lecture 6](#), [Lecture 11](#), and [Lecture 12](#)). In theory, the *only* difficult problem in proving hybrid systems safety is finding their invariants or differential invariants [Pla08, Pla12a] ([Lecture 13 on Differential Invariants & Proof Theory](#)). In practice, however, the handling of real arithmetic is another challenge that you have faced in your labs, even though the problem is easier in theory. How arithmetic interfaces with proofs by way of the proof rules \forall, \exists has already been discussed in [Lecture 6 on Truth & Proof](#). But how does the handling of real arithmetic by quantifier elimination really work?

Today's lecture shows one technique for deciding interesting formulas of first-order real arithmetic. Understanding how such techniques for real arithmetic work is interesting for at least two reasons. First of all, it is important to understand why this miracle happens at all that something as complicated and expressive as first-order logic of real arithmetic is decidable. But this lecture is also helpful to get an intuition about how real arithmetic decision procedures work. With such an understanding, you are better prepared to identify the limitations of these techniques, learn when they are likely not to work out in due time, and get a sense of what you can do to help arithmetic prove more complicated properties. For complex proofs, it is often very important to use your insights and intuitions about the system to help the prover along to scale your verification results to more challenging systems in feasible amounts of time. An understanding how arithmetic decision procedures work helps to focus such insights on the parts of

the arithmetic analysis that has a big computational impact. Quite substantial impact has been observed for handling the challenges of real arithmetic [Pla07, dMP13].

There are a number of different approaches to understanding real arithmetic and its decision procedures besides Tarski's original seminal breakthrough [Tar51]. There is an algebraic approach using cylindrical algebraic decompositions [Col75], which leads to practical procedures, but is highly nontrivial. There are simple and elegant model-theoretic approaches using semantic properties of logic and algebra [Rob77], which are easy to understand, but do not lead to any particularly useful algorithms. There is a reasonably simple Cohen-Hörmander algorithm [Coh69, Hör83] that, unfortunately, does not generalize well into a practical algorithm. Other simple but inefficient decision procedures are also described elsewhere [KK71, Eng93]. And there is virtual substitution [Wei97], a syntactical approach that fits well to the understanding of logic that we have developed in this course and leads to highly efficient algorithms (although not in the most general cases). As a good compromise of accessibility and practicality, this lecture focuses on virtual substitution [Wei97].

These lecture notes are loosely based on [Wei97, Pla10, Appendix D]. They add substantial intuition and motivation that is helpful for following the technical development. More information about virtual substitution can be found in the literature [Wei97]. See, e.g., [PQR09, Pas11] for an overview of other techniques for real arithmetic.

The most important learning goals of this lecture are:

Modeling and Control: This lecture has an indirect impact on CPS models and controls by informing the reader about the consequences of the analytic complexity resulting from different arithmetical modeling tradeoffs. There is always more than one way of writing down a model. It becomes easier to find the right tradeoffs for expressing a CPS model with some knowledge of and intuition for the working principles of the workhorse of quantifier elimination that will handle the resulting arithmetic.

Computational Thinking: The primary purpose of today's lecture is to understand how arithmetical reasoning, which is crucial for CPS, can be done rigorously and automatically. Developing an intuition for the working principles of real arithmetic decision procedures can be very helpful for developing strategies to verify CPS models at scale. The lecture also serves the purpose of learning to appreciate the miracle that quantifier elimination in real arithmetic provides by contrasting it with closely related problems that have fundamentally different challenges. We will also see a conceptually very important device in the logical trinity: the flexibility of moving back and forth between syntax and semantics at will. We have seen this principle in action already in the case of differential invariants in [Lecture 10 on Differential Equations & Differential Invariants](#), where we moved back and forth between analytic differentiation $\frac{d}{dt}$ and syntactic derivations $(\cdot)'$ by way of the derivation lemma and the differential substitution lemma as we saw fit. This time, we leverage the same conceptual device for real arithmetic (rather than differential arithmetic) by working with virtual substitutions to bridge the gap between semantic operations that are inexpressible otherwise in first-order logic

of real arithmetic. Virtual substitutions will again allow us to move back and forth at will between syntax and semantics.

CPS Skills: The author is not aware of any impact that this lecture has on CPS skills, because the primary point of this lecture is to get a sense of the pragmatics of CPS analysis.

2. Framing the Miracle

First-order logic is an expressive logic in which many interesting properties and concepts can be expressed, analyzed, and proven. It is certainly significantly more expressive than propositional logic, which is decidable by NP-complete SAT solving.

In classical (uninterpreted) *first-order logic* (FOL), no symbol (except possibly equality) has a special meaning. There are only predicate symbols p, q, r, \dots and function symbols f, g, h, \dots whose meaning is subject to interpretation. And the domain that quantifiers range over is subject to interpretation. In particular, a formula of first-order logic is only valid if it holds true for all interpretations of all predicate and function symbols and all domains.

In contrast, *first-order logic of real arithmetic* ($\text{FOL}_{\mathbb{R}}$ or the theory of real-closed field arithmetic FOL_{RCF} [Pla10, Appendix D]) is interpreted, because its symbols have a special fixed interpretation. The only predicate symbols are $=, \geq, >, \leq, <, \neq$ and they mean exactly equality, greater-or-equals, greater-than, etc., and the only function symbols are $+, -, \cdot$, which mean exactly addition, subtraction, and multiplication of real numbers. Furthermore, the quantifiers quantify over the set \mathbb{R} of all real numbers.¹

The first special interpretation for symbols that comes to mind may not necessarily be the real numbers but maybe the natural numbers \mathbb{N} with $+$ for addition and \cdot for multiplication on natural numbers and where quantifiers range over the natural numbers. That gives the *first-order logic of natural numbers* ($\text{FOL}_{\mathbb{N}}$). Is $\text{FOL}_{\mathbb{N}}$ easier or harder than FOL? How do both compare to $\text{FOL}_{\mathbb{R}}$? What would happen compared to $\text{FOL}_{\mathbb{Q}}$, the first-order logic of rational numbers? $\text{FOL}_{\mathbb{Q}}$ is like $\text{FOL}_{\mathbb{R}}$ and $\text{FOL}_{\mathbb{N}}$, except that the rational numbers \mathbb{Q} are used as the domain of quantification and interpretation of variables, rather than \mathbb{R} and \mathbb{N} , respectively. How do those different flavors of first-order logic compare? How difficult is it to prove validity of logical formulas in each case?

Before you read on, see if you can find the answer for yourself.

¹Respectively over another real-closed field, but that has been shown not to change validity [Tar51].

Uninterpreted first-order logic FOL is semidecidable, because there is a (sound and complete [Göd30]) proof procedure that is algorithmic and able to prove all true sentences of first-order logic [Her30]. The natural numbers are more difficult. Actually much more difficult! By Gödel's incompleteness theorem [Göd31], first-order logic $\text{FOL}_{\mathbb{N}}$ of natural numbers does not have a sound and complete effective axiomatization. $\text{FOL}_{\mathbb{N}}$ is neither semidecidable nor cosemidecidable [Chu36]. There is neither an algorithm that can prove all valid formulas of $\text{FOL}_{\mathbb{N}}$ nor one that can disprove all formulas of $\text{FOL}_{\mathbb{N}}$ that are not valid. One way of realizing the inherent challenge of the logic of natural numbers in retrospect is to use that not all questions about programs can be answered effectively (for example the halting problem of Turing machines is undecidable) [Chu36, Tur37], in fact "none" can [Ric53], and then encode questions about classical programs into the logic of natural numbers.

Yet, a miracle happened. Alfred Tarski proved in 1930 [Tar31, Tar51] that reals are much better behaved and that $\text{FOL}_{\mathbb{R}}$ is decidable, even though this seminal result remained unpublished for many years and only appeared in full in 1951 [Tar51].

The first-order logic $\text{FOL}_{\mathbb{Q}}$ of rational numbers, however, was shown to be undecidable [Rob49], even though rational numbers may appear to be so close to real numbers. Rationals are lacking something important: completeness (in the topological sense). The square root $\sqrt{2}$ of 2 is a perfectly good witness for $\exists x x^2 = 2$ but only a real number, not a rational one.

The first-order logic $\text{FOL}_{\mathbb{C}}$ of complex numbers, though, is again perfectly decidable [Tar51, CC56].

Note 1 (The miracle of reals. Overview of validity problems of first-order logics).

<i>Logic</i>	<i>Validity</i>
FOL	<i>semidecidable</i>
$\text{FOL}_{\mathbb{N}}$	<i>not semidecidable nor cosemidecidable</i>
$\text{FOL}_{\mathbb{Q}}$	<i>not semidecidable nor cosemidecidable</i>
$\text{FOL}_{\mathbb{R}}$	<i>decidable</i>
$\text{FOL}_{\mathbb{C}}$	<i>decidable</i>

3. Quantifier Elimination

Alfred Tarski's seminal insight for deciding real arithmetic is based on quantifier elimination, i.e. the successive elimination of quantifiers from formulas so that the remaining formula is equivalent but structurally significantly easier, because it has less quantifiers. Why does eliminating quantifiers help? When evaluating a logical formula for whether it is true or false in a given state (i.e. an assignment of real numbers to all its free variables), arithmetic comparisons and polynomial terms are easy, because all we need to do is plug the numbers in and compute according to their semantics (recall [Lecture 2 on Differential Equations & Domains](#)). For example, for a state ν with $\nu(x) = 2$, we can

easily evaluate the logical formula

$$x^2 > 2 \wedge 2x < 3 \vee x^3 < x^2$$

to *false* by following the semantics, which ultimately just plugs in 2 for x :

$$\llbracket x^2 > 2 \wedge 2x < 3 \vee x^3 < x^2 \rrbracket_\nu = 2^2 > 2 \wedge 2 \cdot 2 < 3 \vee 2^3 < 2^2 = \textit{false}$$

Similarly, in a state ω with $\omega(x) = -1$, the same formula evaluates to *true*:

$$\llbracket x^2 > 2 \wedge 2x < 3 \vee x^3 < x^2 \rrbracket_\omega = (-1)^2 > 2 \wedge 2 \cdot (-1) < 3 \vee (-1)^3 < (-1)^2 = \textit{true}$$

But quantifiers are a difficult matter, because they require us to check for all possible values of a variable (in the case $\forall x F$) or to find exactly the right value for a variable that makes the formula true (in the case of $\exists x F$). The easiest formulas to evaluate are the ones that have no free variables (because then their value does not depend on the state ν) and that also have no quantifiers (because then there are no choices for the values of the quantified variables during the evaluation). Quantifier elimination can take a logical formula that is closed, i.e. has no free variables, and equivalently remove its quantifiers, so that it becomes easy to evaluate the formula to *true* or *false*. Quantifier elimination also works for formulas that still have free variables. Then it will eliminate all quantifiers in the formula but the original free variables will remain in the resulting formula, unless it simplifies in the quantifier elimination process.

Definition 1 (Quantifier elimination). A first-order theory admits *quantifier elimination* if, with each formula ϕ , a quantifier-free formula $\text{QE}(\phi)$ can be associated effectively that is equivalent, i.e. $\phi \leftrightarrow \text{QE}(\phi)$ is valid (in that theory).

That is, a first-order theory that admits quantifier elimination if there is a computer program that outputs a quantifier-free formula $\text{QE}(\phi)$ for any input formula ϕ in that theory such that the input and output are equivalent ($\phi \leftrightarrow \text{QE}(\phi)$ is valid) and such that the output $\text{QE}(\phi)$ is quantifier-free.

Theorem 2 (Tarski [Tar51]). *The first-order logic of real arithmetic admits quantifier elimination and is, thus, decidable.*

The operation QE is further assumed to evaluate ground formulas (i.e., without variables), yielding a decision procedure for closed formulas of this theory (i.e., formulas without free variables). For a closed formula ϕ , all it takes is to compute its quantifier-free equivalent $\text{QE}(\phi)$ by quantifier elimination. The closed formula ϕ is closed, so has no free variables or other free symbols, and neither will $\text{QE}(\phi)$. Hence, ϕ as well as its equivalent $\text{QE}(\phi)$ are either equivalent to *true* or to *false*. Yet, $\text{QE}(\phi)$ is quantifier-free, so which one it is can be found out simply by evaluating the (variable-free) concrete arithmetic in $\text{QE}(\phi)$ as in the above examples.

Example 3. Quantifier elimination uses the special structure of real arithmetic to express quantified arithmetic formulas equivalently without quantifiers and without using more free variables. For instance, QE yields the following equivalence:

$$\text{QE}(\exists x (2x^2 + c \leq 5)) \equiv c \leq 5.$$

In particular, the formula $\exists x (2x^2 + c \leq 5)$ is not valid, but only true if $c \leq 5$ holds, as has been so aptly described by the outcome of the above quantifier elimination result.

Example 4. Quantifier elimination can be used to find out whether a first-order formula of real arithmetic is valid. Take $\exists x (2x^2 + c \leq 5)$, for example. A formula is valid iff its universal closure is, i.e. the formula obtained by universally quantifying all free variables. After all, valid means that a formula is true for all interpretations. Hence, consider the universal closure $\forall c \exists x (2x^2 + c \leq 5)$, which is a closed formula, because it has no free variables. Quantifier elimination could, for example, lead to

$$\text{QE}(\forall c \exists x (2x^2 + c \leq 5)) \equiv \text{QE}(\forall c \text{QE}(\exists x (2x^2 + c \leq 5))) \equiv \text{QE}(\forall c (c \leq 5)) \equiv -100 \leq 5 \wedge 5 \leq 5 \wedge 100 \leq 5$$

The resulting formula still has no free variables but is now quantifier-free, so it can simply be evaluated arithmetically. Since the conjunct $100 \leq 5$ evaluates to *false*, the universal closure $\forall c \exists x (2x^2 + c \leq 5)$ is equivalent to *false* and, hence, the original formula $\exists x (2x^2 + c \leq 5)$ is not valid (although still satisfiable for $c = 1$).

Geometrically, quantifier elimination corresponds to projection, see Fig. 1.




Figure 1: The geometric counterpart of quantifier elimination for $\exists y$ is projection onto the x axis

Note that, when using QE, we usually assume it would already evaluate ground arithmetic, so that the only two possible outcomes of applying QE to a closed formula are *true* and *false*.

Alfred Tarski's result that quantifier elimination over the reals is possible and that real arithmetic is decidable was groundbreaking. The only issue is that the complexity of Tarski's decision procedure is non-elementary, i.e. cannot be bounded by any tower of exponentials $2^{2^{\dots^n}}$, which made it quite impractical. Still, it was a seminal breakthrough because it showed reals to be decidable at all. It was not until another seminal result in 1949 by Julia Robinson, who proved the rationals to be undecidable [Rob49]. It took many further advances [Sei54, Coh69, KK71, Hör83, Eng93] and a major breakthrough by George Collins in 1975 [Col75] until more practical procedures had been found [Col75, CH91, Wei97]. The virtual substitution technique shown in this lecture has been implemented in Redlog [DS97], which has an interface for KeYmaera [PQ08]. There is also a recent approach of combining ideas from SMT solving with nonlinear real arithmetic [JdM12] implemented in the SMT solver Z3, which has an interface for KeYmaera.

4. Homomorphic Normalization for Quantifier Elimination

The first insight for defining quantifier elimination is to understand that the quantifier elimination operation commutes with almost all logical connectives, so that QE only needs to be defined for existential quantifiers. Consequently, as soon as we understand how to eliminate existential quantifiers, universal quantifiers can be eliminated as well just by double negation.

$$\text{QE}(A \wedge B) \equiv \text{QE}(A) \wedge \text{QE}(B)$$

$$\text{QE}(A \vee B) \equiv \text{QE}(A) \vee \text{QE}(B)$$

$$\text{QE}(\neg A) \equiv \neg \text{QE}(A)$$

$$\text{QE}(\forall x A) \equiv \text{QE}(\neg \exists x \neg A)$$

These transformations isolate existential quantifiers for quantifier elimination. In particular, it is sufficient if quantifier elimination focuses on existentially quantified variables. When using the QE operation inside out, i.e. when using it repeatedly to eliminate the inner-most quantifier to a quantifier-free equivalent and then again eliminating the inner-most quantifier, the quantifier elimination is solved if only we manage to solve it for $\exists x A$ with a quantifier-free formula A . If A is not quantifier-free, its quantifiers can be eliminated from inside out:

$$\text{QE}(\exists x A) \equiv \text{QE}(\exists x \text{QE}(A)) \quad \text{if } A \text{ not quantifier-free}$$

It is possible, although not necessary and not even necessarily helpful, to simplify the form of A as well. The following transformations transform the kernel of a quantifier into negation normal form using deMorgan's equivalences.

$$\begin{aligned} \text{QE}(\exists x (A \vee B)) &\equiv \text{QE}(\exists x A) \vee \text{QE}(\exists x B) \\ \text{QE}(\exists x \neg(A \wedge B)) &\equiv \text{QE}(\exists x (\neg A \vee \neg B)) \\ \text{QE}(\exists x \neg(A \vee B)) &\equiv \text{QE}(\exists x (\neg A \wedge \neg B)) \\ \text{QE}(\exists x \neg\neg A) &\equiv \text{QE}(\exists x A) \end{aligned}$$

Distributivity can be used to simplify the form of the quantifier-free *kernel* A to disjunctive normal form and split existential quantifiers over disjuncts:

$$\begin{aligned} \text{QE}(\exists x (A \wedge (B \vee C))) &\equiv \text{QE}(\exists x ((A \wedge B) \vee (A \wedge C))) \\ \text{QE}(\exists x ((A \vee B) \wedge C)) &\equiv \text{QE}(\exists x ((A \wedge C) \vee (B \wedge C))) \\ \text{QE}(\exists x (A \vee B)) &\equiv \text{QE}((\exists x A) \vee (\exists x B)) \end{aligned}$$

The only remaining case to address is the case $\text{QE}(\exists x (A \wedge B))$ where $A \wedge B$ is a purely conjunctive formula (yet it can actually have any number of conjuncts, not just two). Using the following normalizing equivalences,

$$\begin{aligned} p = q &\equiv p - q = 0 \\ p \geq q &\equiv p - q \geq 0 \\ p > q &\equiv p - q > 0 \\ p \neq q &\equiv p - q \neq 0 \\ p \leq q &\equiv q - p \geq 0 \\ p < q &\equiv q - p > 0 \\ \neg(p \geq q) &\equiv p < q \\ \neg(p > q) &\equiv p \leq q \\ \neg(p = q) &\equiv p \neq q \\ \neg(p \neq q) &\equiv p = q \end{aligned}$$

it is further possible to normalize all atomic formulas equivalently to one of the forms $p = 0, p > 0, p \geq 0, p \neq 0$. Since $p \neq 0 \equiv p > 0 \vee p < 0$, disequations \neq are unnecessary *in theory* as well (although they are quite useful in practice).

5. Substitution Base

Virtual substitution is a quantifier elimination technique that is based on substituting extended terms into formulas virtually, i.e. without the extended terms² actually occurring in the resulting constraints.

²Being an *extended real term* really means it is not a real term, but somehow closely related. We will see more concrete extended real terms and how to get rid of them again later.

Note 4. *Virtual substitution essentially leads to an equivalence of the form*

$$\exists x F \leftrightarrow \bigvee_{t \in T} A_t \wedge F_x^t \quad (1)$$

for a suitable finite set T of extended terms that depends on the formula F and that gets substituted into F virtually, i.e. in a way that results in standard real arithmetic terms, not extended terms.

Such an equivalence is how quantifier elimination can work. Certainly if the right-hand side of (1) is true, then t is a witness for $\exists x F$. The key to establishing an equivalence of the form (1) is to ensure that if F has a solution at all (in the sense of $\exists x F$ being true), then F must already hold for one of the cases in T . That is, T must cover all representative cases. There might be many more solutions, but if there is one at all, one of the possibilities in T must be a solution as well. If we were to choose all real numbers $T \stackrel{\text{def}}{=} \mathbb{R}$, then (1) would be trivially valid, but then the right-hand side is not a formula because it is uncountably infinitely long, which is even worse than the quantified form on the left-hand side. But if a finite set T is sufficient for the equivalence (1) and the extra formulas A_t are quantifier-free, then the right-hand side of (1) is structurally simpler than the left-hand side, even if it may be (sometimes significantly) less compact.

The various ways of virtually substituting various forms of extended reals e into logical formulas equivalently without having to mention the actual extended reals is the secret of virtual substitution. The first step is to see that it is enough to define substitutions only on atomic formulas of the form $p = 0, p < 0, p \leq 0$ (or, just as well, on $p = 0, p > 0, p \geq 0$). If σ denotes such an extended substitution of θ for x , then σ lifts to arbitrary first-order formulas homomorphically³ as follows

$$\begin{aligned} \sigma(A \wedge B) &\equiv \sigma A \wedge \sigma B \\ \sigma(A \vee B) &\equiv \sigma A \vee \sigma B \\ \sigma(\neg A) &\equiv \neg \sigma A \\ \sigma(\forall y A) &\equiv \forall y \sigma A && \text{if } x \neq y \text{ and } y \notin \theta \\ \sigma(\exists y A) &\equiv \exists y \sigma A && \text{if } x \neq y \text{ and } y \notin \theta \\ \sigma(p = q) &\equiv \sigma(p - q = 0) \\ \sigma(p < q) &\equiv \sigma(p - q < 0) \\ \sigma(p \leq q) &\equiv \sigma(p - q \leq 0) \\ \sigma(p > q) &\equiv \sigma(q - p < 0) \\ \sigma(p \geq q) &\equiv \sigma(q - p \leq 0) \\ \sigma(p \neq q) &\equiv \sigma(\neg(p - q = 0)) \end{aligned}$$

This lifting applies the substitution σ to all subformulas, with minor twists on quantifiers for admissibility and normalization of atomic formulas into the canonical forms

³With a caveat on admissibility for quantifiers to avoid capture of variables.

$p = 0, p < 0, p \leq 0$ for which σ has been assumed to already have been defined.

From now on, all that remains to be done for defining a substitution or virtual substitution is to define it on atomic formulas of the remaining forms $p = 0, p < 0, p \leq 0$ and the above construction will take care of substituting in any first-order formulas. Of course, the above construction is only helpful for normalizing atomic formulas that are not already of one of those forms, so the term q above can be assumed not to be the term 0.

6. Term Substitutions

Consider a formula of the form

$$\exists x (bx + c = 0 \wedge F) \quad (x \notin b, c) \quad (2)$$

where x does not occur in the terms b, c . Let's consider how a first mathematical solution to this formula might look like. The only solution that the conjunct $bx + c = 0$ has is $x = -c/b$. Hence, the left conjunct in (2) only holds for $x = -c/b$, so formula (2) can only be true if F also holds for that single solution $-c/b$ in place of x . That is, formula (2) holds only if $F_x^{-c/b}$ does. Hence, (2) is equivalent to the formula $F_x^{-c/b}$, which is quantifier-free.

So, how can we eliminate the quantifier in (2) equivalently?

Before you read on, see if you can find the answer for yourself.

Most certainly, $F_x^{-c/b}$ is quantifier-free. But it is not exactly always equivalent to (2) and, thus, does not necessarily qualify as its quantifier eliminate form. Oh no! What we wrote down is a good intuitive start, but does not make any sense at all if $b = 0$, for then $-c/b$ would have been a rather ill-devised division by zero. Performing such divisions by zero sounds like a fairly shaky start for an equivalence transformation such as quantifier elimination. And certainly sounds like a shaky start for anything that is supposed to ultimately turn into a proof.

Let's start over. The first conjunct in (2) has the only solution $x = -c/b$ if $b \neq 0$. In that case, indeed, (2) is equivalent to $F_x^{-c/b}$, because the only way for (2) to be true then is exactly when the second conjunct F holds for the solution of the first conjunct, i.e. when $F_x^{-c/b}$ holds. But there is, in general, no way of knowing whether evaluation could yield $b \neq 0$ or not, because b might be a complicated polynomial term that is only zero under some interpretations, not under all. Certainly if b is the zero polynomial, we know for sure. Or if b is a polynomial that is never zero, such as a sum of squares plus a positive constant. In general, if $b = 0$, then, the first conjunct in (2) has all numbers for x as solutions if $c = 0$ and, otherwise, has no solution at all if $c \neq 0$. In the latter case, $b = 0, c \neq 0$, (2) is false, because its first conjunct is already false. In the former case, $b = c = 0$, however, the first conjunct $bx + c = 0$ is trivial and does not impose any constraints on x , nor does it help for finding out a quantifier-free equivalent of (2). In that case $b = c = 0$, the trivial constraint will be dropped and the remaining formula will be considered recursively instead.

Note 5. In the non-degenerate case $b \neq 0$ with $x \notin b, c$, (2) can be rephrased into a quantifier-free equivalent over \mathbb{R} as follows:

$$b \neq 0 \rightarrow (\exists x (bx + c = 0 \wedge F) \leftrightarrow b \neq 0 \wedge F_x^{-c/b}) \quad (3)$$


All it takes is, thus, the ability to substitute the term $-c/b$ for x in the formula F . The division $-c/b$ that will occur in $F_x^{-c/b}$ for ordinary term substitutions can cause technical annoyances but at least it is well-defined, because $b \neq 0$ holds in that context. Instead of pursuing the looming question how exactly this substitution in $F_x^{-c/b}$ works, we make the question more general by moving the quadratic case already.

7. Square Root $\sqrt{\cdot}$ Substitutions for Quadratics

Consider a formula of the form

$$\exists x (ax^2 + bx + c = 0 \wedge F) \quad (x \notin a, b, c) \quad (4)$$

where x does not occur in the terms a, b, c . The generic solution of its first conjunct is $x = (-b \pm \sqrt{b^2 - 4ac})/(2a)$, but that, of course, again depends on whether a could evaluate to zero, in which case linear solutions may be possible and the division by $2a$ is most certainly not well-defined; see Fig. 2. Whether a could be zero may again

Figure 2: Roots of quadratic functions p

sometimes be hard to say when a is a polynomial term that has roots, but does not always evaluate to 0 either (which only the zero polynomial would). So let's be more careful this time to find an equivalent formulation right away for all possible cases of a, b, c . The cases to consider are where the first conjunct is either a constant equation (in which case the equation imposes no interesting constraint on x) or a linear equation (in which case $x = -c/b$ is the solution Sect. 6) or a proper quadratic equation with $a \neq 0$ (in which case $x = (-b \pm \sqrt{b^2 - 4ac})/(2a)$ is the solution). The trivial equation $0 = 0$ when $a = b = c = 0$ is again useless, so another part of F would have to be considered in that case, and the equation $c = 0$ for $a = b = 0, c \neq 0$ is again *false*.

When $ax^2 + bx = 0$ is either a proper linear or a proper quadratic equation, its respective solutions single out the only points that can solve (4), so the only points in which it remains to be checked whether the second conjunct F also holds.

Theorem 5 (Virtual substitution of quadratic equations). *For a quantifier-free formula F with $x \notin a, b, c$, the following equivalence is valid over \mathbb{R} :*

$$\begin{aligned}
 a \neq 0 \vee b \neq 0 \vee c \neq 0 \rightarrow \\
 & \left(\exists x (ax^2 + bx + c = 0 \wedge F) \leftrightarrow \right. \\
 & \quad a = 0 \wedge b \neq 0 \wedge F_x^{-c/b} \\
 & \quad \left. \vee a \neq 0 \wedge b^2 - 4ac \geq 0 \wedge \left(F_x^{(-b + \sqrt{b^2 - 4ac})/(2a)} \vee F_x^{(-b - \sqrt{b^2 - 4ac})/(2a)} \right) \right) \quad (5)
 \end{aligned}$$

Hold on, we fortunately noticed just in time for writing down the formula (5) that $(-b + \sqrt{b^2 - 4ac})/(2a)$ only ever makes actual sense in the reals if $b^2 - 4ac \geq 0$, because the square root is otherwise imaginary, which is hard to find in $\text{FOL}_{\mathbb{R}}$.

The resulting formula on the right-hand side of the bimplication is quantifier-free and, thus, sounds like it could be chosen for $\text{QE}(\exists x (ax^2 + bx + c = 0 \wedge F))$ as long as it is not the case that $a = b = c = 0$.

Note 7. *The important thing to notice, though, is that $(-b \pm \sqrt{b^2 - 4ac})/(2a)$ is not a polynomial term, nor even a rational term, because it involves a square root $\sqrt{\cdot}$. Hence, (5) is not generally a formula of first-order real arithmetic! Unless we do something about its square roots and divisions.*

Recall from [Lecture 2 on Differential Equations & Domains](#) that the terms of $\text{FOL}_{\mathbb{R}}$ are polynomials with rational coefficients in \mathbb{Q} . So $4x^2 + \frac{1}{7}x - 1.41$ is a polynomial term of $\text{FOL}_{\mathbb{R}}$. But $4x^2 + \frac{1}{y}x - 1.41$ is not, because of the division by variable y , which should make us panic about y possibly being zero in any case. And $4x^2 + \frac{1}{7}x - \sqrt{2}$ is not either, because of the square root $\sqrt{2}$.

Note 8 (Semantic domains versus syntactic expressions). *While the domains that the quantifiers \forall and \exists of first-order logic $\text{FOL}_{\mathbb{R}}$ of real arithmetic quantify over includes reals like $\sqrt{2}$, the terms and logical formulas themselves are syntactically restricted to be built from polynomials with rational coefficients. Square roots (and all higher roots) are already part of the semantic domain \mathbb{R} , but not allowed in the syntax of $\text{FOL}_{\mathbb{R}}$.*

Of course, it is still easy to write down a formula such as $\exists x x^2 = 2$ which indirectly makes sure that x will have to assume the value $\sqrt{2}$, but that formula mentions a quantifier again.

Square roots are really not part of real arithmetic. But they can be defined, still, by appropriate quadratures. For example, the positive root $x = \sqrt{y}$ can be defined as $x^2 = y \wedge y \geq 0$. Let's find out how square roots such as $(-b \pm \sqrt{b^2 - 4ac})/(2a)$ can be substituted into first-order formulas systematically without the need for involving any square roots in the resulting formula.

A square root expression is an expression of the form

$$(a + b\sqrt{c})/d$$

with polynomials $a, b, c, d \in \mathbb{Q}[x_1, \dots, x_n]$ of rational coefficients in the variables x_1, \dots, x_n and, for well-definedness, $d \neq 0 \wedge c \geq 0$. Square root expressions with the same \sqrt{c} can be added and multiplied symbolically by considering them as algebraic objects.⁴

$$\begin{aligned} (a + b\sqrt{c})/d + (a' + b'\sqrt{c})/d' &= ((ad' + da') + (bd' + db')\sqrt{c})/(dd') \\ ((a + b\sqrt{c})/d) \cdot ((a' + b'\sqrt{c})/d') &= ((aa' + bb'c) + (ab' + ba')\sqrt{c})/(dd') \end{aligned} \quad (6)$$

Another way of saying that is that square root expressions with the same \sqrt{c} provide an addition and a multiplication operation that leads to square root expressions. Substituting $(a + b\sqrt{c})/d$ for a variable x in a polynomial term p , thus, leads to a square root

⁴Despite the poor notation, please don't mistake the primes for derivatives here. The name a' is not the derivative of a here but just meant as a name for a polynomial term that happens to go by the misleading name a' .

expression $p_x^{(a+b\sqrt{c})/d} = (\tilde{a} + \tilde{b}\sqrt{c})/\tilde{d}$ with the same \sqrt{c} , because the arithmetic resulting from evaluating the polynomial only requires addition and multiplication using (6).⁵

Note 9. *Subsequent symbolic addition and multiplication makes it possible to substitute a square root expression in for a variable in a polynomial to form. Yet, the result $p_x^{(a+b\sqrt{c})/d}$ is still a square root expression, which still cannot be written down directly in first-order real arithmetic. Yet, as soon as a square root expression appears in an atomic formula of first-order real arithmetic, that square root can be rephrased equivalently to disappear.*

The substitution of a square root expression $(a' + b'\sqrt{c})/d'$ into a polynomial p for x to form $p_x^{(a'+b'\sqrt{c})/d'}$ by polynomial evaluation leads to a square root expression, say the square root expression $p_x^{(a'+b'\sqrt{c})/d'} = (a + b\sqrt{c})/d$.

The next step is to handle the comparison of the resulting square root expression to 0 in atomic formulas $p \sim 0$ for some $\sim \in \{=, \leq, <\}$. That works by characterizing it using the square root expression $p_x^{(a'+b'\sqrt{c})/d'}$:

$$(p \sim 0)_{\tilde{x}}^{(a'+b'\sqrt{c})/d'} \equiv (p_x^{(a'+b'\sqrt{c})/d'} \sim 0)$$

Suppose the square root expression $p_x^{(a'+b'\sqrt{c})/d'}$ from the polynomial evaluation is $(a + b\sqrt{c})/d$. All that remains to be done is to rewrite $(a + b\sqrt{c})/d \sim 0$ equivalently in $\text{FOL}_{\mathbb{R}}$.

Assume $d \neq 0 \wedge c \geq 0$ for well-definedness. For square-root-free expressions ($b = 0$) with just divisions, i.e. those of the form $(a + 0\sqrt{c})/d$, the following equivalences hold:

$$\begin{aligned} (a + 0\sqrt{c})/d = 0 &\equiv a = 0 \\ (a + 0\sqrt{c})/d \leq 0 &\equiv ad \leq 0 \\ (a + 0\sqrt{c})/d < 0 &\equiv ad < 0 \end{aligned}$$

Assume $d \neq 0 \wedge c \geq 0$ for well-definedness. For square root expressions $(a + b\sqrt{c})/d$ with arbitrary b , the following equivalences hold:

$$\begin{aligned} (a + b\sqrt{c})/d = 0 &\equiv ab \leq 0 \wedge a^2 - b^2c = 0 \\ (a + b\sqrt{c})/d \leq 0 &\equiv ad \leq 0 \wedge a^2 - b^2c \geq 0 \vee bd \leq 0 \wedge a^2 - b^2c \leq 0 \\ (a + b\sqrt{c})/d < 0 &\equiv ad < 0 \wedge a^2 - b^2c > 0 \vee bd \leq 0 \wedge (ad < 0 \vee a^2 - b^2c < 0) \end{aligned}$$

The first line characterizes that $= 0$ holds iff a, b have different signs (possibly 0) and their squares cancel, because $a^2 = b^2c$. The second line characterizes that ≤ 0 holds iff $a^2 \geq b^2c$ so that a will dominate, which has a different sign than d , or if $a^2 \leq b^2c$ so that $b\sqrt{c}$ will dominate, which has a different sign than d (possibly 0). The squares $a^2 - b^2c = a^2 - b^2\sqrt{c}^2$ is the square of the absolute value of the involved terms, which uniquely identifies the truth-values along with the accompanying sign conditions. The

⁵In practice, the polynomial addition and multiplication operations for a polynomial η are performed by Horner's scheme for dense polynomials η and by repeated squaring for sparse polynomials η .

third line characterizes that < 0 holds iff a strictly dominates, because $a^2 > b^2c$ and the dominant a, d have different signs or if b, d have different signs and either a, d have different signs as well (so a, b have the same sign or 0 and different than d) or b strictly dominates because $a^2 < b^2c$. The last case involves a little extra care for the required sign conditions to avoid the $= 0$ case.

This defines the substitution of a square root $(a + b\sqrt{c})/d$ for x into atomic formulas and can be lifted to all first-order logic formulas as explained in Sect. 5. The important thing to observe is that the result of this substitution does not introduce square root expressions nor divisions even though the square root expression $(a + b\sqrt{c})/d$ had the square root \sqrt{c} and the division $/d$. Substitution of a square root $(a + b\sqrt{c})/d$ for x into a (quantifier-free) first-order formula F then works as usual by substitution in all atomic formulas (as defined in Sect. 5). The result of such a *virtual* substitution is denoted by $F_{\hat{x}}^{(a+b\sqrt{c})/d}$.

It is crucial to note that the *virtual substitution* of square root expression $(a + b\sqrt{c})/d$ for x in F giving $F_{\hat{x}}^{(a+b\sqrt{c})/d}$ is semantically equivalent to the result $F_x^{(a+b\sqrt{c})/d}$ of the literal substitution replacing x with $(a + b\sqrt{c})/d$, but operationally quite different, because the virtual substitution never introduces square roots or divisions. Because of their semantical equivalence, we use the same notation by abuse of notation.

Lemma 6 (Virtual substitution lemma for square roots). *The result $F_{\hat{x}}^{(a+b\sqrt{c})/d}$ of the virtual substitution is semantically equivalent to the the result $F_x^{(a+b\sqrt{c})/d}$ of the literal substitution, but better behaved, because it stays within $\text{FOL}_{\mathbb{R}}$ proper. Essentially, the following equivalence of virtual substitution and literal substitution for square root expressions is valid:*

$$F_x^{(a+b\sqrt{c})/d} \leftrightarrow F_{\hat{x}}^{(a+b\sqrt{c})/d}$$

Keep in mind, though, that the result $F_{\hat{x}}^{(a+b\sqrt{c})/d}$ of virtual substitution is a proper formula of $\text{FOL}_{\mathbb{R}}$, while the literal substitution $F_x^{(a+b\sqrt{c})/d}$ could actually only even be considered a formula in an extended logic that allows for a syntactic representation of divisions and square root expressions within a context in which they are meaningful (no divisions by zero, no imaginary roots).

Using Lemma 6, Theorem 5 continues to hold when using the so-defined *square root virtual substitutions* $F_{\hat{x}}^{(-b \pm \sqrt{b^2 - 4ac})/(2a)}$ that turn (5) into a valid formula of first-order real arithmetic, without scary square root expressions. In particular, since the fraction $-c/b$ also is a (somewhat impoverished) square root expression $(-c + 0\sqrt{0})/b$, the $\text{FOL}_{\mathbb{R}}$ formula $F_{\hat{x}}^{-c/b}$ in (5) can be formed and rephrased equivalently using the square root virtual substitution as well. Hence, the quantifier-free right-hand side of (5) neither introduces square roots nor divisions but happily remains a proper formula in $\text{FOL}_{\mathbb{R}}$.

With this virtual substitution, the right-hand side of the biimplication (5) can be chosen as $\text{QE}(\exists x (ax^2 + bx + c = 0 \wedge F))$ if it is not the case that $a = b = c = 0$.

When using square root substitutions, divisions could, thus, also have been avoided

in the quantifier elimination (3) for the linear case. Thus, the right-hand side of (3) can be chosen as $\text{QE}(\exists x (bx + c = 0 \wedge F))$ if it is not the case that $b = c = 0$.

8. Optimizations

Before going any further, it is helpful to notice that virtual substitutions admit a number of useful optimizations that make it more practical. When substituting a square root expression $(a + b\sqrt{c})/d$ for a variable x in a polynomial p , the resulting square root expression $p_{\tilde{x}}^{(a+b\sqrt{c})/d} = (\tilde{a} + \tilde{b}\sqrt{c})/\tilde{d}$ will end up occurring with a higher power of the form $\tilde{d} = d^k$ where k is the degree of p in variable x . This is easy to see just by inspecting the definitions of addition and multiplication from (6). Such larger powers of d can be avoided using the equivalences $(pq^3 \sim 0) \equiv (pq \sim 0)$ and, if $q \neq 0$, using also $(pq^2 \sim 0) \equiv (p \sim 0)$ for arithmetic relations $\sim \in \{=, >, \geq, \neq, <, \leq\}$. Since $d \neq 0$ needs to be assumed for well-definedness of a square root expression $(a + b\sqrt{c})/d$, the degree of d in the result $F_{\tilde{x}}^{(a+b\sqrt{c})/d}$ of the virtual substitution can, thus, be lowered to either 0 or 1 depending on whether it ultimately occurs as an even or as an odd power (Exercise 7). If d occurs as an odd power, its occurrence can be lowered to degree 1. If d occurs as an even power, its occurrence can be reduced to degree 0, which makes it disappear entirely.

The significance of lowering degrees does not just come from the conceptual and computational impact that large degrees have on the problem of quantifier elimination, but, for the case of virtual substitution, also from the fact that virtual substitution only works for certain bounded but common degrees.

Example 7. Using this principle to check under which circumstance the quadratic equality from (4) evaluates to *true* requires a nontrivial number of algebraic and logical computations to handle the virtual substitution of the respective roots of $ax^2 + bx + c = 0$ into F .

Just out of curiosity: What would happen if we tried to apply the same virtual substitution coming from this equation to $ax^2 + bx + c = 0$ itself instead of to F ? Imagine, for example, that $ax^2 + bx + c = 0$ shows up a second time in F . Let's only consider the case of quadratic solutions, i.e. where $a \neq 0$. And let's only consider the root $(-b + \sqrt{b^2 - 4ac})/(2a)$. The other cases are left as an exercise. First virtually substitute $(-b + \sqrt{b^2 - 4ac})/(2a)$ into the polynomial $ax^2 + bx + c$ leading to symbolic square root

expression arithmetic:

$$\begin{aligned}
& (ax^2 + bx + c)_{\hat{x}}^{(-b+\sqrt{b^2-4ac})/(2a)} \\
&= a((-b + \sqrt{b^2 - 4ac})/(2a))^2 + b((-b + \sqrt{b^2 - 4ac})/(2a)) + c \\
&= a((b^2 + b^2 - 4ac + (-b - b)\sqrt{b^2 - 4ac})/(4a^2)) + (-b^2 + b\sqrt{b^2 - 4ac})/(2a) + c \\
&= (ab^2 + ab^2 - 4a^2c + (-ab - ab)\sqrt{b^2 - 4ac})/(4a^2) + (-b^2 + 2ac + b\sqrt{b^2 - 4ac})/(2a) + c \\
&= ((ab^2 + ab^2 - 4a^2c)2a + (-b^2 + 2ac)4a^2 + ((-ab - ab)2a + b4a^2)\sqrt{b^2 - 4ac})/(4a^2) \\
&= (\cancel{2a^2b^2} + \cancel{2a^2b^2} - \cancel{8a^3c} + \cancel{-4a^2b^2} + \cancel{8a^3c} + (\cancel{-2a^2b} - \cancel{2a^2b} + \cancel{4a^2b})\sqrt{b^2 - 4ac})/(4a^2) \\
&= (0 + 0\sqrt{0})/1 = 0
\end{aligned}$$

So $(ax^2 + bx + c)_{\hat{x}}^{(-b+\sqrt{b^2-4ac})/(2a)}$ is the zero square root expression? That is actually exactly as expected by construction, because $(-b \pm \sqrt{b^2 - 4ac})/(2a)$ is supposed to be the root of $ax^2 + bx + c$ in the case where $a \neq 0 \wedge b^2 - 4ac \geq 0$. In particular, if $ax^2 + bx + c$ occurs again in F as either an equation or inequality, its virtual substitute in the various cases just ends up being:

$$\begin{aligned}
(ax^2 + bx + c = 0)_{\hat{x}}^{(-b+\sqrt{b^2-4ac})/(2a)} &\equiv ((0 + 0\sqrt{0})/1 = 0) \equiv (0 \cdot 1 = 0) \equiv \text{true} \\
(ax^2 + bx + c \leq 0)_{\hat{x}}^{(-b+\sqrt{b^2-4ac})/(2a)} &\equiv ((0 + 0\sqrt{0})/1 \leq 0) \equiv (0 \cdot 1 \leq 0) \equiv \text{true} \\
(ax^2 + bx + c < 0)_{\hat{x}}^{(-b+\sqrt{b^2-4ac})/(2a)} &\equiv ((0 + 0\sqrt{0})/1 < 0) \equiv (0 \cdot 1 < 0) \equiv \text{false} \\
(ax^2 + bx + c \neq 0)_{\hat{x}}^{(-b+\sqrt{b^2-4ac})/(2a)} &\equiv ((0 + 0\sqrt{0})/1 \neq 0) \equiv (0 \cdot 1 \neq 0) \equiv \text{false}
\end{aligned}$$

And that makes sense as well. After all, the roots of $ax^2 + bx + c = 0$ satisfy the weak inequality $ax^2 + bx + c \leq 0$ but not the strict inequality $ax^2 + bx + c < 0$. In particular, Theorem 5 could substitute the roots of $ax^2 + bx + c = 0$ also into the full formula $ax^2 + bx + c = 0 \wedge F$ under the quantifier, but the formula resulting from the left conjunct $ax^2 + bx + c = 0$ will always simplify to *true* so that only the virtual substitution into F will remain, where actual logic with real arithmetic happens.

The above computations are all that is needed for Theorem 5 to show the following quantifier elimination equivalences:

$$\begin{aligned}
a \neq 0 \rightarrow (\exists x (ax^2 + bx + c = 0 \wedge ax^2 + bx + c = 0)) &\leftrightarrow b^2 - 4ac \geq 0 \wedge \text{true} \\
a \neq 0 \rightarrow (\exists x (ax^2 + bx + c = 0 \wedge ax^2 + bx + c \leq 0)) &\leftrightarrow b^2 - 4ac \geq 0 \wedge \text{true}
\end{aligned}$$

With analog computations for the case $(-b - \sqrt{b^2 - 4ac})/(2a)$, this also justifies:

$$\begin{aligned}
a \neq 0 \rightarrow (\exists x (ax^2 + bx + c = 0 \wedge ax^2 + bx + c < 0)) &\leftrightarrow b^2 - 4ac \geq 0 \wedge \text{false} \\
a \neq 0 \rightarrow (\exists x (ax^2 + bx + c = 0 \wedge ax^2 + bx + c \neq 0)) &\leftrightarrow b^2 - 4ac \geq 0 \wedge \text{false}
\end{aligned}$$

Consequently, in a context where $a \neq 0$ is known, for example because it is a term such as 5 or $y^2 + 1$, Theorem 5 and simplification yields the following quantifier elimination results:

$$\text{QE}(\exists x (ax^2 + bx + c = 0 \wedge ax^2 + bx + c = 0)) \equiv b^2 - 4ac \geq 0$$

$$\text{QE}(\exists x (ax^2 + bx + c = 0 \wedge ax^2 + bx + c \leq 0)) \equiv b^2 - 4ac \geq 0$$

$$\text{QE}(\exists x (ax^2 + bx + c = 0 \wedge ax^2 + bx + c < 0)) \equiv \text{false}$$

$$\text{QE}(\exists x (ax^2 + bx + c = 0 \wedge ax^2 + bx + c \neq 0)) \equiv \text{false}$$

In a context where $a \neq 0$ is not known, more cases become possible and the disjunctive structure in Theorem 5 remains, leading to a case distinction on whether $a = 0$ or $a \neq 0$.

Example 8 (Nonnegative roots of quadratic polynomials). Consider the formula

$$\exists x (ax^2 + bx + c = 0 \wedge x \geq 0) \quad (7)$$

for the purpose of eliminating quantifiers using Theorem 5. For simplicity, again assume $a \neq 0$ is known, e.g., because $a = 5$. Since $a \neq 0$, Theorem 5 will only consider the two square root expressions $(-b + \sqrt{b^2 - 4ac})/(2a)$ and $(-b - \sqrt{b^2 - 4ac})/(2a)$ and no linear roots. The first thing that happens during the virtual substitution of those roots into the remaining formula $F \equiv (x \geq 0)$ is that the construction in Sect. 5 will flip $x \geq 0$ around to a base case $-x \leq 0$. On that base case, the substitution of the square root expression $(-b + \sqrt{b^2 - 4ac})/(2a)$ into the polynomial $-x$ leads to the following square root computations following (6):

$$-(-b + \sqrt{b^2 - 4ac})/(2a) = ((-1 + 0\sqrt{b^2 - 4ac})/1) \cdot ((-b + \sqrt{b^2 - 4ac})/(2a)) = (b - \sqrt{b^2 - 4ac})/(2a)$$

Observe how the unary minus operator expands to multiplication by -1, whose representation as a square root expression with square root $\sqrt{b^2 - 4ac}$ is $(-1 + 0\sqrt{b^2 - 4ac})/1$. The virtual square root substitution of this square root expression, thus, yields

$$\begin{aligned} & (-x \leq 0)_{\hat{x}}^{(b - \sqrt{b^2 - 4ac})/(2a)} \\ & \equiv b2a \leq 0 \wedge b^{\cancel{2}} - (-1)^2(b^{\cancel{2}} - 4ac) \geq 0 \vee -1 \cdot 2a \leq 0 \wedge b^{\cancel{2}} - (-1)^2(b^{\cancel{2}} - 4ac) \leq 0 \\ & \equiv 2ba \leq 0 \wedge 4ac \geq 0 \vee -2a \leq 0 \wedge 4ac \leq 0 \end{aligned}$$

For the second square root expression $(-b - \sqrt{b^2 - 4ac})/(2a)$, the corresponding polynomial evaluation leads to

$$-(-b - \sqrt{b^2 - 4ac})/(2a) = ((-1 + 0\sqrt{b^2 - 4ac})/1) \cdot ((-b - \sqrt{b^2 - 4ac})/(2a)) = (b + \sqrt{b^2 - 4ac})/(2a)$$

The virtual square root substitution of this square root expression, thus, yields

$$\begin{aligned} & (-x \leq 0)_{\hat{x}}^{(b + \sqrt{b^2 - 4ac})/(2a)} \\ & \equiv b2a \leq 0 \wedge b^{\cancel{2}} - 1^2(b^{\cancel{2}} - 4ac) \geq 0 \vee 1 \cdot 2a \leq 0 \wedge b^{\cancel{2}} - 1^2(b^{\cancel{2}} - 4ac) \leq 0 \\ & \equiv 2ba \leq 0 \wedge 4ac \geq 0 \vee 2a \leq 0 \wedge 4ac \leq 0 \end{aligned}$$

Consequently, since $a \neq 0$, Theorem 5 implies the quantifier elimination equivalence:

$$a \neq 0 \rightarrow (\exists x (ax^2 + bx + c = 0 \wedge x \geq 0)) \\ \leftrightarrow b^2 - 4ac \geq 0 \wedge (2ba \leq 0 \wedge 4ac \geq 0 \vee -2a \leq 0 \wedge 4ac \leq 0 \vee 2ba \leq 0 \wedge 4ac \geq 0 \vee 2a \leq 0 \wedge 4ac \leq 0))$$

Consequently, in a context where $a \neq 0$ is known, 5 yields the following quantifier elimination results:

$$\text{QE}(\exists x (ax^2 + bx + c = 0 \wedge x \geq 0)) \\ \equiv b^2 - 4ac \geq 0 \wedge (2ba \leq 0 \wedge 4ac \geq 0 \vee -2a \leq 0 \wedge 4ac \leq 0 \vee \del{2ba \leq 0 \wedge 4ac \geq 0} \vee 2a \leq 0 \wedge 4ac \leq 0) \\ \equiv b^2 - 4ac \geq 0 \wedge (ba \leq 0 \wedge ac \geq 0 \vee a \geq 0 \wedge ac \leq 0 \vee a \leq 0 \wedge ac \leq 0)$$

The sign conditions that this formula expresses make sense when you consider that the original quantified formula (7) expresses that the quadratic equation has a nonnegative root.

9. Summary

This lecture showed part of the miracle of quantifier elimination and quantifier elimination is possible in first-order real arithmetic. Today's technique works for formulas that normalize into an appropriate form as long as the technique can latch on to a linear or quadratic equation for all quantified variables. Note that there can be higher-degree or inequality occurrences of the variables as well within the formula F of Theorem 5, but there has to be at least one linear or quadratic equation. Commuting the formula so that it has the required form is easily done if such an equation is anywhere at all. What is to be done if there is no quadratic equation but only other quadratic inequalities is the topic of the next lecture.

It is also foreseeable that the virtual substitution approach will ultimately run into difficulties for pure high-degree polynomials, because those generally have no radicals to solve the equations. That is where other more algebraic quantifier elimination techniques come into play that are beyond the scope of this lecture.

Virtual substitution of square root expressions uses simple symbolic computations:

$$(\alpha + \beta\sqrt{\gamma})/\delta + (\alpha' + \beta'\sqrt{\gamma})/\delta' = ((\alpha\delta' + \delta\alpha') + (\beta\delta' + \delta\beta')\sqrt{\gamma})/(\delta\delta') \\ ((\alpha + \beta\sqrt{\gamma})/\delta) \cdot ((\alpha' + \beta'\sqrt{\gamma})/\delta') = ((\alpha\alpha' + \beta\beta'\gamma) + (\alpha\beta' + \beta\alpha')\sqrt{\gamma})/(\delta\delta')$$

The following expansions were the core of eliminating square root expressions by virtual substitutions. For square root expressions $(\alpha + \beta\sqrt{\gamma})/\delta$ with $\delta \neq 0 \wedge \gamma \geq 0$ for well-definedness, the following equivalences hold:

$$(\alpha + \beta\sqrt{\gamma})/\delta = 0 \equiv \alpha\beta \leq 0 \wedge \alpha^2 - \beta^2\gamma = 0 \\ (\alpha + \beta\sqrt{\gamma})/\delta \leq 0 \equiv \alpha\delta \leq 0 \wedge \alpha^2 - \beta^2\gamma \geq 0 \vee \beta\delta \leq 0 \wedge \alpha^2 - \beta^2\gamma \leq 0 \\ (\alpha + \beta\sqrt{\gamma})/\delta < 0 \equiv \alpha\delta < 0 \wedge \alpha^2 - \beta^2\gamma > 0 \vee \beta\delta \leq 0 \wedge (\alpha\delta < 0 \vee \alpha^2 - \beta^2\gamma < 0)$$

A. Real Algebraic Geometry

This course follows a logical view on cyber-physical systems. It is helpful to develop an intuition to what geometric objects the various logical concepts correspond. The part that is most interesting in this context is real algebraic geometry [BCR98] as it relates to real arithmetic [BPR06]. General algebraic geometry is also very elegant and beautiful, especially over algebraically closed fields [Har95, CLO92].

The geometric counterpart of polynomial equations are real affine algebraic varieties. Every set F of polynomials defines a geometric object, its variety, i.e. the set of points on which all those polynomials are zero.

Definition 9 (Real Affine Algebraic Variety). $V \subseteq \mathbb{R}^n$ is an *affine variety* iff, for some set $F \subseteq \mathbb{R}[X_1, \dots, X_n]$ of polynomials over \mathbb{R} :

$$V = V(F) := \{x \in \mathbb{R}^n : f(x) = 0 \text{ for all } f \in F\}$$

i.e., affine varieties are subsets of \mathbb{R}^n that are definable by a set of polynomial equations.

The converse construction is that of the vanishing ideal, which describes the set of all polynomials that are zero on a given set V .

Definition 10 (Vanishing Ideal). $I \subseteq \mathbb{R}[X_1, \dots, X_n]$ is the *vanishing ideal* of $V \subseteq \mathbb{R}^n$:

$$I(V) := \{f \in \mathbb{R}[X_1, \dots, X_n] : f(x) = 0 \text{ for all } f \in V\}$$

i.e., all polynomials that are zero on all of V .

Affine varieties and vanishing ideals are related by

$$\begin{aligned} S \subseteq V(I(S)) & & \text{for any set } S \subseteq \mathbb{R}^n \\ V = V(I(V)) & & \text{if } V \text{ an affine variety} \\ F \subseteq G \Rightarrow V(F) \supseteq V(G) \end{aligned}$$

Affine varieties and vanishing ideals are intimately related by Hilbert's Nullstellensatz over algebraically closed fields such as \mathbb{C} and by Stengle's Nullstellensatz over real-closed fields such as \mathbb{R} .

The affine varieties corresponding to a number of interesting polynomials are illustrated in Fig. 3.

Exercises

Exercise 1. Example 7 showed that $ax^2 + bx + c = 0$ simplifies to *true* for the virtual substitution of the root $(-b + \sqrt{b^2 - 4ac})/(2a)$. Show that the same thing happens for the root $(-b - \sqrt{b^2 - 4ac})/(2a)$ and the root $(-c + 0\sqrt{0})/b$.




Figure 3: Polynomial equations describe (real) affine (algebraic) varieties

Exercise 2. Example 7 argued that the simplification of $ax^2 + bx + c = 0$ to *true* for the virtual substitution of the root $(-b + \sqrt{b^2 - 4ac})/(2a)$ is to be expected, because the real number to which $(-b + \sqrt{b^2 - 4ac})/(2a)$ evaluates is a root of $ax^2 + bx + c = 0$ in the case where $a \neq 0 \wedge b^2 - 4ac \geq 0$. Yet, what happens in the case where the extra assumption $a \neq 0 \wedge b^2 - 4ac \geq 0$ does not hold? What is the value of the virtual substitution in that case? Is that a problem? Discuss carefully!

Exercise 3. Use Theorem 5 to eliminate quantifiers in the following formula, assuming $a \neq 0$ is known:

$$\exists x (ax^2 + bx + c = 0 \wedge x < 1)$$

Exercise 4. Use Theorem 5 to eliminate quantifiers in the following formula, assuming $a \neq 0$ is known:

$$\exists x (ax^2 + bx + c = 0 \wedge x^3 + x \leq 0)$$

Exercise 5. How does Example 8 change when removing the assumption that $a \neq 0$?

Exercise 6. Would first-order logic of real arithmetic miss the presence of π ? That is, if we delete π from the domain and make all quantifiers range only over $\mathbb{R} \setminus \{\pi\}$, would there be any formula that notices by having a different truth-value? If we delete $\sqrt[3]{5}$ from the domain, would $\text{FOL}_{\mathbb{R}}$ notice?

Exercise 7. Consider the process of substituting a square root expression $(a + b\sqrt{c})/d$ for a variable x in a polynomial p . Let k be the degree of p in variable x , so that d occurs as d^k with power k in the result $p_{\frac{(a+b\sqrt{c})}{d}}^x = (\tilde{a} + \tilde{b}\sqrt{c})/\tilde{d}$. Let $\delta = 1$ when k is odd and $\delta = 0$ when k is even. Show that the following optimization can be used for the virtual substitution. Assume $d \neq 0 \wedge c \geq 0$ for well-definedness. For square-root-free

expressions ($b = 0$) with just divisions, i.e. those of the form $(a + 0\sqrt{c})/d$, the following equivalences hold:

$$\begin{aligned}(a + 0\sqrt{c})/d = 0 &\equiv a = 0 \\(a + 0\sqrt{c})/d \leq 0 &\equiv ad^\delta \leq 0 \\(a + 0\sqrt{c})/d < 0 &\equiv ad^\delta < 0 \\(a + 0\sqrt{c})/d \neq 0 &\equiv a \neq 0\end{aligned}$$

Assume $d \neq 0 \wedge c \geq 0$ for well-definedness. For square root expressions $(a + b\sqrt{c})/d$ with arbitrary b , the following equivalences hold:

$$\begin{aligned}(a + b\sqrt{c})/d = 0 &\equiv ab \leq 0 \wedge a^2 - b^2c = 0 \\(a + b\sqrt{c})/d \leq 0 &\equiv ad^\delta \leq 0 \wedge a^2 - b^2c \geq 0 \vee bd^\delta \leq 0 \wedge a^2 - b^2c \leq 0 \\(a + b\sqrt{c})/d < 0 &\equiv ad^\delta < 0 \wedge a^2 - b^2c > 0 \vee bd^\delta \leq 0 \wedge (ad^\delta < 0 \vee a^2 - b^2c < 0) \\(a + b\sqrt{c})/d \neq 0 &\equiv ab > 0 \vee a^2 - b^2c \neq 0\end{aligned}$$

References

- [BCR98] Jacek Bochnak, Michel Coste, and Marie-Francoise Roy. *Real Algebraic Geometry*, volume 36 of *Ergebnisse der Mathematik und ihrer Grenzgebiete*. Springer, 1998.
- [BPR06] Saugata Basu, Richard Pollack, and Marie-Françoise Roy. *Algorithms in Real Algebraic Geometry*. Springer, 2nd edition, 2006. [doi:10.1007/3-540-33099-2](https://doi.org/10.1007/3-540-33099-2).
- [CC56] Claude Chevalley and Henri Cartan. Schémas normaux; morphismes; ensembles constructibles. In *Séminaire Henri Cartan*, volume 8, pages 1–10. Numdam, 2955-1956.
- [CH91] George E. Collins and Hoon Hong. Partial cylindrical algebraic decomposition for quantifier elimination. *J. Symb. Comput.*, 12(3):299–328, 1991.
- [Chu36] Alonzo Church. A note on the Entscheidungsproblem. *J. Symb. Log.*, 1(1):40–41, 1936.
- [CLO92] David A. Cox, John Little, and Donal O’Shea. *Ideals, Varieties and Algorithms: An Introduction to Computational Algebraic Geometry and Commutative Algebra*. Undergraduate Texts in Mathematics. Springer, New York, 1992.
- [Coh69] Paul J. Cohen. Decision procedures for real and p -adic fields. *Communications in Pure and Applied Mathematics*, 22:131–151, 1969.
- [Col75] George E. Collins. Hauptvortrag: Quantifier elimination for real closed fields by cylindrical algebraic decomposition. In H. Barkhage, editor, *Automata Theory and Formal Languages*, volume 33 of *LNCS*, pages 134–183. Springer, 1975.

- [dMP13] Leonardo Mendonça de Moura and Grant Olney Passmore. The strategy challenge in SMT solving. In Maria Paola Bonacina and Mark E. Stickel, editors, *Automated Reasoning and Mathematics - Essays in Memory of William W. McCune*, volume 7788 of *Lecture Notes in Computer Science*, pages 15–44. Springer, 2013.
- [DS97] Andreas Dolzmann and Thomas Sturm. Redlog: Computer algebra meets computer logic. *ACM SIGSAM Bull.*, 31:2–9, 1997.
- [Eng93] E. Engeler. *Foundations of Mathematics: Questions of Analysis, Geometry and Algorithmics*. Springer, 1993.
- [Göd30] Kurt Gödel. Die Vollständigkeit der Axiome des logischen Funktionenkalküls. *Mon. hefte Math. Phys.*, 37:349–360, 1930.
- [Göd31] Kurt Gödel. Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Mon. hefte Math. Phys.*, 38:173–198, 1931.
- [Har95] Joe Harris. *Algebraic Geometry: A First Course*. Graduate Texts in Mathematics. Springer, 1995.
- [Her30] Jacques Herbrand. Recherches sur la théorie de la démonstration. *Travaux de la Société des Sciences et des Lettres de Varsovie, Class III, Sciences Mathématiques et Physiques*, 33:33–160, 1930.
- [Hör83] L. Hörmander. *The Analysis of Linear Partial Differential Operators II*, volume 257 of *Grundlehren der mathematischen Wissenschaften*. Springer, 1983.
- [JdM12] Dejan Jovanovic and Leonardo Mendonça de Moura. Solving non-linear arithmetic. In Bernhard Gramlich, Dale Miller, and Uli Sattler, editors, *Automated Reasoning - 6th International Joint Conference, IJCAR 2012, Manchester, UK, June 26-29, 2012. Proceedings*, volume 7364 of *LNCS*, pages 339–354. Springer, 2012.
- [KK71] Georg Kreisel and Jean-Louis Krivine. *Elements of mathematical logic: Model Theory*. North-Holland, 2 edition, 1971.
- [LIC12] *Proceedings of the 27th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2012, Dubrovnik, Croatia, June 25–28, 2012*. IEEE, 2012.
- [Pas11] Grant Olney Passmore. *Combined Decision Procedures for Nonlinear Arithmetics, Real and Complex*. PhD thesis, School of Informatics, University of Edinburgh, 2011.
- [Pla07] André Platzer. Combining deduction and algebraic constraints for hybrid system analysis. In Bernhard Beckert, editor, *VERIFY'07 at CADE, Bremen, Germany*, volume 259 of *CEUR Workshop Proceedings*, pages 164–178. CEUR-WS.org, 2007.
- [Pla08] André Platzer. Differential dynamic logic for hybrid systems. *J. Autom. Reas.*, 41(2):143–189, 2008. doi:10.1007/s10817-008-9103-8.
- [Pla10] André Platzer. *Logical Analysis of Hybrid Systems: Proving Theorems for Complex Dynamics*. Springer, Heidelberg, 2010. doi:10.1007/978-3-642-14509-4.
- [Pla12a] André Platzer. The complete proof theory of hybrid systems. In LICS [LIC12], pages 541–550. doi:10.1109/LICS.2012.64.

- [Pla12b] André Platzer. Logics of dynamical systems. In LICS [LIC12], pages 13–24. doi:10.1109/LICS.2012.13.
- [PQ08] André Platzer and Jan-David Quesel. KeYmaera: A hybrid theorem prover for hybrid systems. In Alessandro Armando, Peter Baumgartner, and Gilles Dowek, editors, *IJCAR*, volume 5195 of *LNCS*, pages 171–178. Springer, 2008. doi:10.1007/978-3-540-71070-7_15.
- [PQR09] André Platzer, Jan-David Quesel, and Philipp Rümmer. Real world verification. In Renate A. Schmidt, editor, *CADE*, volume 5663 of *LNCS*, pages 485–501. Springer, 2009. doi:10.1007/978-3-642-02959-2_35.
- [Ric53] H. Gordon Rice. Classes of recursively enumerable sets and their decision problems. *Trans. AMS*, 89:25–59, 1953.
- [Rob49] Julia Robinson. Definability and decision problems in arithmetic. *J. Symb. Log.*, 14(2):98–114, 1949.
- [Rob77] Abraham Robinson. *Complete Theories*. Studies in logic and the foundations of mathematics. North-Holland, 2 edition, 1977.
- [Sei54] Abraham Seidenberg. A new decision method for elementary algebra. *Annals of Mathematics*, 60:365–374, 1954.
- [Tar31] Alfred Tarski. Sur les ensembles définissables de nombres réels I. *Fundam. Math.*, 17:210–239, 1931.
- [Tar51] Alfred Tarski. *A Decision Method for Elementary Algebra and Geometry*. University of California Press, Berkeley, 2nd edition, 1951.
- [Tur37] Alan M. Turing. Computability and lambda-definability. *J. Symb. Log.*, 2(4):153–163, 1937.
- [Wei97] Volker Weispfenning. Quantifier elimination for real algebra — the quadratic case and beyond. *Appl. Algebra Eng. Commun. Comput.*, 8(2):85–101, 1997.