15-424: Foundations of Cyber-Physical Systems

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 15). In theory, the *only* difficult problem in proving hybrid systems safety is finding their invariants or differential invariants [Pla08, Pla12a] (Lecture 14 on Differential Invariants & Proof Theory). In practice, however, the handling of real arithmetic is another challenge that you have faced, even though the problem is easier in theory. How arithmetic interfaces with proofs has already been discussed in Lecture 9 on Proofs & Arithmetic. 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 that something as complicated and expressive as first-order logic of real arithmetic is decidable. But it 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 more.

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.

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 $p, q, r, \ldots$ and function symbols $f, g, h, \ldots$ 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* ($\mathrm{FOL}_{\mathbb{R}}$ or the theory of real-closed field arithmetic $\mathrm{FOL}_{\mathrm{RCF}}$ [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 by 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* ($\mathrm{FOL}_{\mathbb{N}}$). Is $\mathrm{FOL}_{\mathbb{N}}$ easier or harder than FOL? How do both compare to $\mathrm{FOL}_{\mathbb{R}}$? What would happen compared to $\mathrm{FOL}_{\mathbb{Q}}$, the first-order logic of rational numbers? $\mathrm{FOL}_{\mathbb{Q}}$ is like $\mathrm{FOL}_{\mathbb{R}}$ and $\mathrm{FOL}_{\mathbb{N}}$, 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.

[1] Respectively 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 much more difficult. By Gödel's incompleteness theorem, 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 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].

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).

Note 1 (Overview of validity problems of first-order logics)**.**

Logic	*Validity*
FOL	*semidecidable*
$\text{FOL}_{\mathbb{N}}$	*not semidecidable nor cosemidecidable*
$\text{FOL}_{\mathbb{Q}}$	*not semidecidable nor cosemidecidable*
$\text{FOL}_{\mathbb{R}}$	*decidable*
$\text{FOL}_{\mathbb{C}}$	*decidable*

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. 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). For example, for a state ν with $\nu(x) = 2$, we can easily evaluate the logical formula

$$x^2 > 2 \land 2x < 3 \lor x^3 < x^2$$

to true just by plugging in 2 for x. But quantifiers are difficult, 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) 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 $\mathrm{QE}(\phi)$ can be associated effectively that is equivalent, i.e. $\phi \leftrightarrow \mathrm{QE}(\phi)$ is valid (in that theory).

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 $\mathrm{QE}(\phi)$ by quantifier elimination. The closed formula ϕ is closed, so has no free variables or other free symbols, and neither will $\mathrm{QE}(\phi)$. Hence, ϕ as well as its equivalent $\mathrm{QE}(\phi)$ are either equivalent to *true* or to *false*. Yet, $\mathrm{QE}(\phi)$ is quantifier-free, so which one it is can be found out simply by evaluating the (variable-free) concrete arithmetic in $\mathrm{QE}(\phi)$.

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:

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

In particular, the formula $\exists x\,(2x^2 + c \leq 5)$ is not valid, but only if $c \leq 5$, 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 intepretations. Hence, consider the universal closure $\forall c\,\exists x\,(2x^2 + c \leq 5)$, which is a closed formula. Quantifier elimination might, for example, lead to

$$\mathrm{QE}(\forall c\,\exists x\,(2x^2{+}c \leq 5)) \equiv \mathrm{QE}(\forall c\;\mathrm{QE}(\exists x\,(2x^2{+}c \leq 5))) \equiv \mathrm{QE}(\forall c\,(c \leq 5)) \equiv -100 \leq 5 \land 5 \leq 5 \land 100 \leq 5$$

The resulting formula is 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 $c = 1$).

The complexity of Alfred Tarski's decision procedure is non-elementary, i.e. cannot be bounded by any tower of exponentials $2^{2^{2^{\cdots n}}}$. 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].

4 Homomorphic Normalization

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. Especially, as soon as we understand how to eliminate existential quantifiers, universal quantifiers can be eliminated as well just by double negation.

$$\begin{aligned}
\mathrm{QE}(A \land B) &\equiv \mathrm{QE}(A) \land \mathrm{QE}(B) \\
\mathrm{QE}(A \lor B) &\equiv \mathrm{QE}(A) \lor \mathrm{QE}(B) \\
\mathrm{QE}(\neg A) &\equiv \neg\, \mathrm{QE}(A) \\
\mathrm{QE}(\forall x\, A) &\equiv \mathrm{QE}(\neg \exists x\, \neg A)
\end{aligned}$$

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:

$$\mathrm{QE}(\exists x\, A) \equiv \mathrm{QE}(\exists x\; \mathrm{QE}(A)) \qquad\qquad \text{if } A \text{ not quantifier-free}$$

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

$$\begin{aligned}
\mathrm{QE}(\exists x\, (A \lor B)) &\equiv \mathrm{QE}(\exists x\, A) \lor \mathrm{QE}(\exists x\, B) \\
\mathrm{QE}(\exists x\, \neg(A \land B)) &\equiv \mathrm{QE}(\exists x\, (\neg A \lor \neg B)) \\
\mathrm{QE}(\exists x\, \neg(A \lor B)) &\equiv \mathrm{QE}(\exists x\, (\neg A \land \neg B)) \\
\mathrm{QE}(\exists x\, \neg\neg A) &\equiv \mathrm{QE}(\exists x\, A)
\end{aligned}$$

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

$$\begin{aligned}\mathrm{QE}(\exists x\,(A \land (B \lor C))) &\equiv \mathrm{QE}(\exists x\,((A \land B) \lor (A \land C)))\\ \mathrm{QE}(\exists x\,((A \lor B) \land C)) &\equiv \mathrm{QE}(\exists x\,((A \land C) \lor (B \land C)))\\ \mathrm{QE}(\exists x\,(A \lor B)) &\equiv \mathrm{QE}((\exists x\,A) \lor (\exists x\,B))\end{aligned}$$

The remaining case to address is the case $\mathrm{QE}(\exists x\,(A \land B))$ where $A \land B$ is a purely conjunctive formula (yet it can 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 $p \neq 0 \equiv p > 0 \lor p < 0$, disequations $\neq$ are unnecessary *in theory* as well (although they are 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 terms[2] actually occurring in the resulting constraints.

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

$$\exists x\,F \leftrightarrow \bigvee_{t \in T} A_t \land F_x^t \tag{1}$$

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.*

[2]Being 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.

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 (in the sense of $\exists x\, F$ being true), then F must hold for one of the cases in T. That is, T must cover all representative cases. 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 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 homomorphically[3] as follows

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

This lifting applies the substitution σ to all subformulas, with minor twists on quantifiers for admissibility and normalization of atomic formulas into the forms $p = 0, p < 0, p \leq 0$ for which σ has been assumed to already have been defined.

6 Term Substitutions

Consider a formula of the form

$$\exists x\, (bx + c = 0 \land F) \tag{2}$$

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

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

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 like a shaky start for anything that is supposed to 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 $b \neq 0$, (2) can be rephrased into a quantifier-free equivalent over $\mathbb{R}$ as follows:*

$$b \neq 0 \rightarrow \big(\exists x\,(bx + c = 0 \land F) \leftrightarrow b \neq 0 \land F_x^{-c/b}\big) \tag{3}$$

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.

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

Consider a formula of the form

$$\exists x\,(ax^2 + bx + c = 0 \land F) \tag{4}$$

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. Whether a could be zero may again 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 is 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, the following equivalence is valid over $\mathbb{R}$:*

$$
\begin{aligned}
& a \neq 0 \lor b \neq 0 \lor c \neq 0 \rightarrow \\
& \qquad \Big(\exists x\, (ax^2 + bx + c = 0 \land F) \leftrightarrow \\
& \qquad\quad a = 0 \land b \neq 0 \land F_x^{-c/b} \\
& \qquad \lor a \neq 0 \land b^2 - 4ac \geq 0 \land \big(F_x^{(-b+\sqrt{b^2-4ac})/(2a)} \lor F_x^{(-b-\sqrt{b^2-4ac})/(2a)} \big) \Big)
\end{aligned} \tag{5}
$$

The resulting formula on the right-hand side of the biimplication is quantifier-free and, thus, sounds like it could be chosen for $\mathrm{QE}(\exists x\, (ax^2 + bx + c = 0 \land F))$ 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 exactly a polynomial term, not even a rational term, because it involves a square root $\sqrt{\cdot}$. Hence, (5) is not generally a formula of first-order real arithmetic.*

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 $x^2 = y \land y \geq 0$. 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 square roots in the resulting formula.

A *square root expression* is an expression of the form

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

with polynomials $a, b, c, d \in \mathbb{Q}[x_1, \ldots, x_n]$ of rational coefficients in the variables $x_1, \ldots, x_n$ and, for well-definedness, $d \neq 0$. Square roots with the same $\sqrt{d}$ can be added and multiplied as expected:

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

Substituting $(a+b\sqrt{c})/d$ for a variable x in a polynomial term p, thus, leads to a square root $p_x^{(a+b\sqrt{c})/d} = (\tilde{a}+\tilde{b}\sqrt{c})/\tilde{d}$ with the same $\sqrt{c}$, because the arithmetic resulting from evaluating the polynomial only requires addition and multiplication.

Note 8. *This explains how a square root expression can be substituted in for a variable in a polynomial. Yet, the result is still a square root expression, which cannot be written down directly in first-order real arithmetic. Yet, as soon as a square root expression, say $(a+b\sqrt{c})/d$, appears in an atomic formula of first-order real arithmetic, the square root can be rephrased equivalently to disappear.*

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

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

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

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

This defines the substitution of a square root $(a+b\sqrt{c})/d$ for x into atomic formulas when normalizing atomic formulas appropriately[4]. 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). Denote the result of such a substitution by $F_{\bar{x}}^{(a+b\sqrt{c})/d}$.

It is crucial to note that the *virtual substitution* of square root expression $(a+b\sqrt{c})/d$ for x in F giving $F_{\bar{x}}^{(a+b\sqrt{c})/d}$ 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 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.

Theorem 5 continues to hold when using the so-defined square root substitutions $F_{\bar{x}}^{(-b\pm\sqrt{b^2-4ac})/(2a)}$ that make (5) a valid formula of first-order real arithmetic, without

[4]E.g. $f > g \equiv f - g > 0$ and $f \leq g \equiv g \geq f$

square root expressions. In particular, since the fraction $-c/b$ also is a (somewhat impoverished) square root expression $(-c+0\sqrt{0})/b$, $F_x^{-c/b}$ in (5) can be formed using the square root substitution, so the quantifier-free right-hand side of (5) neither introduces square roots nor divisions.

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

Before going any further, it is helpful to notice that virtual substitutions admit a number of useful optimizations that make it more practical. For example, when substituting a square root expression $(a+b\sqrt{c})/d$ for a variable x in a polynomial p, the resulting square root expression $p_x^{(a+b\sqrt{c})/d} = (\tilde{a}+\tilde{b}\sqrt{c})/\tilde{d}$ has a higher power $\tilde{d}=d^k$ where k is the degree of p in variable x, just by inspecting the above definitions of addition and multiplication. Such larger powers of d can be avoided. Note the equivalences $(pq^3 \sim 0) \equiv (pq \sim 0)$ and, if $q \neq 0$, even $(pq^2 \sim 0 \equiv (p \sim 0)$ for arithmetic relations $\sim \in \{=,>,\geq,\neq,<,\leq\}$. Since $d \neq 0$ for well-definedness, the degree of d in the result $F_x^{(a+b\sqrt{c})/d}$ of the virtual substitution can be lowered to 0 or 1 depending on whether it occurs as an even or odd power.

Example 6. Using this principle to check under which circumstance the quadratic equality from (4) evaluates to $true$ requires a nontrivial number of computations to handle the virtual substitution of the respective roots of $ax^2+bx+c=0$ into F. What would happen if we tried to apply the same virtual substitution coming from this equation to $ax^2+bx+c=0$ itself? Imagine, for example, that $ax^2+bx+c=0$ shows up again 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:

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

So $(ax^2+bx+c)_x^{(-b+\sqrt{b^2-4ac})/(2a)}$ 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 $a\neq 0 \land b^2-4ac \geq 0$. In particular, if ax^2+bx+c occurs again in F as either an equation or inequality, its virtual substitute in the various cases is

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

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 $ax^2+bx+c=0 \land F$ under the quantifier, but the formula resulting from the left conjunct $ax^2+bc+c=0$ will always simplify to *true* so that only the virtual substitution into F will remain.

Exercises

Exercise 1. Example 6 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 $(-c+0\sqrt{0})/b$.

Exercise 2. Example 6 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 $(-b+\sqrt{b^2-4ac})/(2a)$ is a root of $ax^2+bx+c=0$ in the case where $a\neq 0 \land b^2-4ac\geq 0$. Yet, what happens in the case where the extra assumption $a\neq 0\land b^2-4ac\geq 0$ does not hold? What is the value of the virtual substitution in that case? Is that a problem? Discuss carefully!

References

[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.

[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.

[DBL12] *Proceedings of the 27th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2012, Dubrovnik, Croatia, June 25–28, 2012*. IEEE, 2012.

[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.

[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.

[KK71] Georg Kreisel and Jean-Louis Krivine. *Elements of mathematical logic: Model Theory*. North-Holland, 2 edition, 1971.

[Pas11] Grant Olney Passmore. *Combined Decision Procedures for Nonlinear Arithmetics, Real and Complex*. PhD thesis, School of Informatics, University of Edinburgh, 2011.

[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* [DBL12], pages 541–550. doi:10.1109/LICS.2012.64.

[Pla12b] André Platzer. Logics of dynamical systems. In *LICS* [DBL12], 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.

[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.