15-424: Foundations of Cyber-Physical Systems

Lecture Notes on Virtual Substitution & Real Arithmetic

André Platzer

Carnegie Mellon University
Lecture 19

1 Introduction

Reasoning about cyber-physical systems and hybrid systems requires understanding and handling their real arithmetic, which can be challenging, because cyber-physical systems can have complex behavior. Differential dynamic logic and its proof calculus [Pla08, Pla10, Pla12] reduce the verification of hybrid systems to real arithmetic. How arithmetic interfaces with proofs has already been discussed in Lecture 9 on Proofs & Arithmetic. How real arithmetic with linear and quadratic equations can be handled by virtual substitution has been shown in Lecture 18 on Virtual Substitution & Real Equations. Today's lecture shows how virtual substitution for quantifier elimination in real arithmetic extends to the case of linear and quadratic inequalities.

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 Recall: Square Root $\sqrt{\cdot}$ Substitutions for Quadratics

Recall the way to handle quantifier elimination for linear or quadratic equations from Lecture 18 on Virtual Substitution & Real Equations:

Theorem 1 (Virtual substitution of quadratic equations)**.** *For a quantifier-free formula F, the following equivalence is valid over $\mathbb{R}$:*

$$\begin{aligned} a \neq 0 \vee b \neq 0 \vee c \neq 0 \rightarrow \\ \Big(\exists x\,(ax^2+bx+c=0 \wedge F) \leftrightarrow \\ a=0 \wedge b \neq 0 \wedge F_x^{-c/b} \\ \vee\, a \neq 0 \wedge b^2-4ac \geq 0 \wedge \big(F_x^{(-b+\sqrt{b^2-4ac})/(2a)} \vee F_x^{(-b-\sqrt{b^2-4ac})/(2a)}\big)\Big) \end{aligned} \tag{1}$$

When using virtual substitutions of square roots from Lecture 18, the resulting formula on the right-hand side of the biimplication is quantifier-free and can be chosen for $\mathrm{QE}(\exists x\,(ax^2+bx+c=0 \wedge F))$ as long as it is not the case that $a=b=c=0$. In case $a=b=c=0$, another formula in F needs to be considered for directing quantifier elimination, because the equation $ax^2+bx+c=0$ is noninformative if $a=b=c=0$, e.g. when a,b,c are the zero polynomials or even if they just have a common root.

The formula on the right-hand side of the biimplication in (1) is a formula in the first-order logic of real arithmetic when using the virtual substitution of square root expressions defined in Lecture 18.

3 Infinity ∞ Substitution

Theorem 1 address the case where the quantified variable occurs in a linear or quadratic equation. It might only occur in inequalities, however. Consider a formula of the form

$$\exists x\,(ax^2+bx+c \leq 0 \wedge F) \tag{2}$$

Under the respective conditions on a,b,c from (1), the possible solutions $-c/b, (-b+\sqrt{d})/(2a), (-b-\sqrt{d})/(2a)$ from (1) continue to be options for solutions of (2), because one way of satisfying the weak inequality $ax^2+bx+c \leq 0$ is by satisfying the equation $ax^2+bx+c=0$. So if F is *true* for any of those solutions of the quadratic equation (under the auspices of the additional constraints on a,b,c), then (2) holds as well.

Yet, if those points do not work out, the weak inequality in (2) allows for more possible solutions. For example, if $a=0, b>0$, then sufficiently small values of x would satisfy $0x^2+bx+c \leq 0$. Also, if $a<0$, then sufficiently small values of x would satisfy $ax^2+bx+c \leq 0$, because x^2 grows faster than x and, thus the negative ax^2 ultimately overcomes any contribution of bx and c to the value of ax^2+bx+c. But that would quickly diverge into the principle full substitution principle for the uninsightful case of $T \stackrel{\text{def}}{=} \mathbb{R}$ from Lecture 18.

Now, one possibility of pursuing this line of thought may be to substitute smaller and smaller values for x into (2) and see if that happens to work. There is a much better way though. The only really small value that would have to be substituted into (2) to see if it

happens to work is one that is so negative that it is smaller than all others: $-\infty$, which is the lower limit of all negative real numbers. Alternatively, $-\infty$ can be understood as being "always as negative as needed, i.e. more negative than anything else". Think of $-\infty$ as being built out of elastic rubber so that it always ends up being smaller when being compared to any actual real number.

Let $\infty, -\infty$ be *positive and negative infinities*, respectively, i.e. choose extra elements $\infty, -\infty \notin \mathbb{R}$ with $-\infty < r < \infty$ for all $r \in \mathbb{R}$. Formulas of real arithmetic can be substituted with $\pm\infty$ for a variable x if the compactified reals $\mathbb{R} \cup \{\infty, -\infty\}$. Yet, just like with square root expressions, $\pm\infty$ do not actually need to ever occur in the resulting formula, because substitution of infinities can be defined differently. For example, $(x + 5 > 0)^{\infty}_{x}$ simplifies to *false*, while $(x + 5 < 0)^{\infty}_{x}$ simplifies to *true*.

Note 2. *Substitution of the infinity $-\infty$ for x into an atomic formula for a polynomial $p \overset{def}{=} \sum_{i=0}^{n} a_i x^i$ with polynomials a_i that do not contain x is defined by the following equivalences (accordingly for substituting ∞ for x).*

$$(p = 0)^{-\infty}_{x} \equiv \bigwedge_{i=0}^{n} a_i = 0 \tag{3}$$

$$(p \leq 0)^{-\infty}_{x} \equiv (p < 0)^{-\infty}_{x} \vee (p = 0)^{-\infty}_{x} \tag{4}$$

$$(p < 0)^{-\infty}_{x} \equiv @_{-\infty}(p) < 0 \tag{5}$$

$$(p \neq 0)^{-\infty}_{x} \equiv \bigvee_{i=0}^{n} a_i \neq 0 \tag{6}$$

Lines (3) and (6) use that the only equation of real arithmetic that infinities $\pm\infty$ satisfy is the trivial equation $0 = 0$. Line (4) uses the equivalence $p \leq 0 \equiv p < 0 \vee p = 0$. Line (5) uses a simple inductive definition based on the degree, $\deg(p)$, in the variable x of the polynomial p to characterize whether p is ultimately negative at $-\infty$ (or for sufficiently negative numbers):

Note 3. *Let $p \overset{def}{=} \sum_{i=0}^{n} a_i x^i$ with polynomials a_i that do not contain x. Whether p is ultimately negative (written $@_{-\infty}(p) < 0$) at $-\infty$ is easy to characterize:*

$$@_{-\infty}(p) < 0 \overset{def}{\equiv} \begin{cases} p < 0 & \text{if } \deg(p) = 0 \\ (-1)^n a_n < 0 \vee (a_n = 0 \wedge @_{-\infty}(\sum_{i=0}^{n-1} a_i x^i) < 0) & \text{if } \deg(p) > 0 \end{cases}$$

Substitution of ∞ for x into an atomic formula is defined similarly, except that the sign factor $(-1)^n$ disappears. Substitution of ∞ or of $-\infty$ for x into first-order formulas is then defined as in Lecture 18.

Example 2. Using this principle to check under which circumstance the quadratic inequality from (2) evaluates to *true* yields the answer from our earlier ad-hoc analysis of

what happens for sufficiently small values of x:

$$(ax^2 + bx + c \leq 0)^{-\infty}_{x} \equiv (-1)^2 a < 0 \vee a = 0 \wedge ((-1)b < 0 \vee b = 0 \wedge c < 0)$$

In the same way, the virtual substitution can be used to see under which circumstance F would also evaluate to *true* for sufficiently small values of x, exactly when $F^{-\infty}_{x}$ holds. Note that (at least if $a \neq 0$), the virtual substitution of ∞ for x would not make sense to check (2) at, because in that case, the inequality $ax^2 + bx + c \leq 0$ is violated. That would be different for an inequality such as $ax^2 + bx + c \geq 0$.

The crucial thing to note is again that the *virtual substitution* of infinities $\pm\infty$ for x in F giving $F^{\pm\infty}_{x}$ is semantically equivalent to the result $F^{\pm\infty}_{x}$ of the literal substitution replacing x with $\pm\infty$, but operationally different, because the virtual substitution never introduces actual infinities. Because of their semantical equivalence, we use the same notation by abuse of notation.

4 Infinitesimal ε Substitutions

Theorem 1 address the case where the quantified variable occurs in a linear or quadratic equation and the virtual substitution in Sect. 3 adds the case of sufficiently small values for x. Consider a formula of the form

$$\exists x\, (ax^2 + bx + c < 0 \wedge F) \tag{7}$$

In this case, the roots from Theorem 1 will not help, because they satisfy the equation $ax^2 + bx + c = 0$ but not the strict inequality $ax^2 + bx + c < 0$. The virtual substitution of $-\infty$ for x from Sect. 3 still makes sense to consider, because that one might satisfy F and $ax^2 + bx + c < 0$. If $-\infty$ does not work, however, the solution of (7) could be near one of the roots of $ax^2 + bx + c = 0$, just slightly off so that $ax^2 + bx + c < 0$ is satisfied. How far off? Well, saying that exactly by any real number is again difficult, because any particular real number might already have been too large in absolute value, depending on the constraints in the remainder of F. Again, this calls for quantities that are always as small as we need them to be.

Sect. 3 used a negative quantity that is so small that it is smaller than all negative numbers and hence infinitely small (but infinitely large in absolute value). Analyzing (7) needs positive quantities that are infinitely small and hence also infinitely small in absolute value. Infinitesimals are positive quantities that are always smaller than all positive real numbers, i.e. "always as small as needed". Think of them as built out of elastic rubber so that they always shrink as needed when compared with any actual positive real number so that the infinitesimals end up being smaller than positive reals. Another way of looking at infinitesimals is that they are the multiplicative inverses of $\pm\infty$.

A positive infinitesimal $\infty > \varepsilon > 0$ is positive and an extended real that is infinitesimal, i.e. positive but smaller than all positive real numbers ($\varepsilon < r$ for all $r \in \mathbb{R}$ with $r > 0$). For all polynomials $p \in \mathbb{R}[x] \setminus \{0\}$, $\zeta \in \mathbb{R}$, the Taylor expansion of p around ζ evaluated at $\zeta + \varepsilon$ can be used to show:

1. $p(\zeta + \varepsilon) \neq 0$
   that is, infinitesimal are positive and always so small that they never yield roots of any equation, except the trivial zero polynomial. Whenever it looks like there might be a root, the infinitesimal just became a bit smaller. And nonzero univariate polynomials only have finitely many roots, so the infinitesimals will take care to avoid them.

2. $p(\zeta) \neq 0 \Rightarrow p(\zeta)p(\zeta + \varepsilon) > 0$,
   that is, p has constant sign on infinitesimal neighborhoods of nonroots ζ. If the neighborhood around ζ is small enough (and for an infinitesimal it will be), then the polynomial will not yet have changed sign then.

3. $0 = p(\zeta) = p'(\zeta) = p''(\zeta) = \cdots = p^{(k-1)}(\zeta) \neq p^{(k)}(\zeta) \Rightarrow p^{(k)}(\zeta)p(\zeta + \varepsilon) > 0$,
   that is the first nonzero derivative of p at ζ determines the sign of p in an infinitesimal neighborhood of ζ.

Note 4. *Substitution of an infinitesimal expression $e + \varepsilon$ with a square root expression e and a positive infinitesimal ε for x into a polynomial $p = \sum_{i=0}^{n} a_i x^i$ with polynomials a_i that do not contain x is defined by the following equivalences.*

$$(p = 0)_x^{e+\varepsilon} \equiv \bigwedge_{i=0}^{n} a_i = 0 \tag{8}$$

$$(p \leq 0)_x^{e+\varepsilon} \equiv (p < 0)_x^{e+\varepsilon} \vee (p = 0)_x^{e+\varepsilon} \tag{9}$$

$$(p < 0)_x^{e+\varepsilon} \equiv (@_-(p) < 0)_x^{e} \tag{10}$$

$$(p \neq 0)_x^{e+\varepsilon} \equiv \bigvee_{i=0}^{n} a_i \neq 0 \tag{11}$$

Lines (8) and (11) use that infinitesimals offsets satisfy no equation except the trivial equation 0=0 (case 1). Line (9) again uses the equivalence $p \leq 0 \equiv p < 0 \vee p = 0$. Line (10) checks that the sign of p at e is negative (which will make p inherit the same negative sign at $e + \varepsilon$ by case 2) or will immediately become negative right away using a recursive formulation of immediately becoming negative that uses higher derivatives (which determine the sign by case 3). The lifting to arbitrary quantifier-free formulas of real arithmetic is again by substitution into all atomic subformulas and equivalences such as $(p > q) \equiv (p - q > 0)$ as defined in Lecture 18. Note that, for the case $(p < 0)_x^{e+\varepsilon}$, the (non-infinitesimal) square root expression e gets virtually substituted in for x into a formula $@_-(p) < 0$, which characterizes whether p becomes negative immediately at or after x (which will be substituted by e).

Note 5. *Whether p is immediately negative, i.e. negative itself or with a derivative p' that makes it negative on an infinitesimal interval, can be characterized recursively:*

$$@_-(p) < 0 \overset{def}{\equiv} \begin{cases} p < 0 & \text{if } deg(p) = 0 \\ p < 0 \vee (p = 0 \wedge @_-(p')) & \text{if } deg(p) > 0 \end{cases}$$

Example 3. Using this principle to check under which circumstance the quadratic strict inequality from (7) evaluates to *true* at $(-b + \sqrt{b^2 - 4ac})/(2a) + \varepsilon$, i.e. right after its root $(-b + \sqrt{b^2 - 4ac})/(2a)$, leads to the following computation.

$$@_-(ax^2+bx+c) \equiv ax^2+bx+c < 0 \vee ax^2+bc+c = 0 \wedge (2ax+b < 0 \vee 2ax+b = 0 \wedge 2a < 0)$$

Hence,

$$\begin{aligned}
&(ax^2 + bx + c < 0)_x^{(-b+\sqrt{b^2-4ac})/(2a)+\varepsilon} \equiv (@_-(ax^2 + bx + c))_x^{(-b+\sqrt{b^2-4ac})/(2a)+\varepsilon} \equiv \\
&(ax^2 + bx + c < 0 \vee ax^2 + bc + c = 0 \wedge (2ax + b < 0 \vee 2ax + b = 0 \wedge 2a < 0))_x^{(-b+\sqrt{b^2-4ac})/(2a)} \\
&\equiv 0{\cdot}1 < 0 \vee 0 = 0 \wedge (\underbrace{(0 < 0 \vee 4a^2 \leq 0 \wedge (0 < 0 \vee -4a^2(b^2 - 4ac) < 0))}_{2ax+b<0_x^{(-b+\sqrt{b^2-4ac})/(2a)}} \vee \underbrace{0 = 0}_{2ax+b=0_x^{\cdots}} \wedge \underbrace{2a1 < 0}_{2a<0}) \\
&\equiv 4a^2 \leq 0 \wedge -4a^2(b^2 - 4ac) < 0 \vee 2a < 0
\end{aligned}$$

because the square root virtual substitution gives $(ax^2 + bx + c)_x^{(-b+\sqrt{b^2-4ac})/(2a)} = 0$ by construction (compare example from Lecture 18). The virtual substitution into the polynomial $2ax + b$ computes as follows:

$$\begin{aligned}
(2ax + b)_x^{(-b\pm\sqrt{b^2-4ac})/(2a)} &\equiv 2a \cdot (-b \pm \sqrt{b^2 - 4ac})/(2a) + b \\
&\equiv (-2ab + \pm 2a\sqrt{b^2 - 4ac})/(2a) + b \\
&\equiv (\cancel{-2ab} + \cancel{2ab} + \pm 2a\sqrt{b^2 - 4ac})/(2a)
\end{aligned}$$

The resulting formula can be further simplified internally to

$$(ax^2 + bx + c < 0)_x^{(-b+\sqrt{b^2-4ac})/(2a)+\varepsilon} \equiv 4a^2 \leq 0 \wedge -4a^2(b^2 - 4ac) < 0 \vee 2a < 0 \equiv 2a < 0$$

because the first conjunct $4a^2 \leq 0 \equiv a = 0$ and, with $a = 0$, the second conjunct simplifies to $-4a^2(b^2 - 4ac)_a^0 = -0(b^2) < 0$, which is impossible in the reals. This answer makes sense. Because, indeed, exactly if $2a < 0$ will a quadratic polynomial still evaluate to $ax^2 + bx + c < 0$ right after its second root $(-b + \sqrt{b^2 - 4ac})/(2a)$.

Formulas such as this one ($2a > 0$) are the result of a quantifier elimination procedure. If the formula after quantifier elimination is either *true* or *false*, then you know for sure that the formula is valid (*true*) or unsatisfiable (*false*), respectively. If the result of quantifier elimination is *true*, for example, KeYmaera can close proof branches (marked by

proof rule $\mathbb{R}$ in our sequent proofs). Yet, quantifier elimination can also return other formulas, such as $2a > 0$, which are equivalent to the formula where quantifier elimination has been applied. In particular, they identify exactly under which circumstance that respective quantified formula is true. This can be very useful for identifying the missing assumptions to make a proof work and the corresponding statement true.

The crucial thing to note is again that the *virtual substitution* of infinitesimal expressions $e + \varepsilon$ for x in F giving $F_x^{e+\varepsilon}$ is semantically equivalent to the result $F_x^{e+\varepsilon}$ of the literal substitution replacing x with $e + \varepsilon$, but operationally different, because the virtual substitution never introduces actual infinitesimals. Because of their semantical equivalence, we use the same notation by abuse of notation.

Computationally more efficient substitutions of infinitesimals have been reported elsewhere [BD07].

5 Quantifier Elimination by Virtual Substitution

The following quantifier elimination technique works for formulas with a quantified variable that occurs at most quadratically.

Theorem 4 (Virtual substitution of quadratic constraints [Wei97]). *Let F be a quantifier-free formula in which all atomic formulas are of the form $ax^2 + bx + c \sim 0$ for x-free polynomials a, b, c and $\sim \in \{=, \leq, <, \neq\}$, with corresponding discriminant $d \overset{def}{=} b^2 - 4ac$. Then $\exists x\, F$ is equivalent over $\mathbb{R}$ to the following quantifier-free formula:*

$$F_x^{-\infty} \vee \bigvee_{(ax^2+bx+c\{\substack{=\\\leq}\}0)\in F} \big(a = 0 \wedge b \neq 0 \wedge F_x^{-c/b} \vee a \neq 0 \wedge d \geq 0 \wedge (F_x^{(-b+\sqrt{d})/(2a)} \vee F_x^{(-b-\sqrt{d})/(2a)})\big)$$
$$\vee \bigvee_{(ax^2+bx+c\{\substack{\neq\\<}\}0)\in F} \big(a = 0 \wedge b \neq 0 \wedge F_x^{-c/b+\varepsilon} \vee a \neq 0 \wedge d \geq 0 \wedge (F_x^{(-b+\sqrt{d})/(2a)+\varepsilon} \vee F_x^{(-b-\sqrt{d})/(2a)+\varepsilon})\big)$$

Proof. The proof first considers the literal substitution of square root expressions, infinities, and infinitesimals and then, as a second step, uses that the virtual substitutions that avoid square root expressions, infinities, and infinitesimals are equivalent.

The implication from the quantifier-free formula on the right-hand side (denoted G) to $\exists x\, F$ is obvious, because each disjunct of the quantifier-free formula has a conjunct of the form F_x^t for some (extended) term t even if it may be a square root expression or infinity or term involving infinitesimals.

The converse implication from $\exists x\, F$ to the quantifier-free formula depends on showing that the quantifier-free formula covers all possible representative cases and that the accompanying constraints on a, b, c, d are actually necessary.

It is enough to prove this for the case where all variables in F except x have concrete numeric real values, because the equivalence holds iff it holds in all states ν. By a fundamental property of real arithmetic called o-minimality, the set $\mathcal{S}(F)$ of all real values for x that satisfy F forms a finite union of (pairwise disjoint) intervals, because the polynomials in F only change signs at their roots, of which there only are finitely many now that the polynomials have become univariate, i.e. with the only variable x, since all free variables are evaluated to concrete real numbers in ν. Without loss of generality (by merging overlapping or adjacent intervals), we assume all those intervals to be maximal, i.e. no bigger interval would satisfy F. So F actually changes its truth-value at the lower and upper endpoints of these intervals (unless the interval is unbounded).

The endpoints of these intervals can be seen to be of the form $-c/b, (-b+\sqrt{d})/(2a), (-b-\sqrt{d})/(2a), \infty, -\infty$ for any of the polynomials in F, because those polynomials are at most quadratic and all roots of those polynomials are contained in that set. Hence, as usual, $-c/b \in \mathcal{S}(F)$ implies $a = 0, b \neq 0$, because that is the only case where $-c/b$ satisfies F, which has only at most quadratic polynomials, while $(-b+\sqrt{d})/(2a) \in \mathcal{S}(F)$ as well as $(-b-\sqrt{d})/(2a) \in \mathcal{S}(F)$ both imply that $a \neq 0$ and discriminant $d \geq 0$. So the side conditions for the roots considered in the quantifier-free formula are necessary for quadratic polynomials.

Now consider one interval $I \subseteq \mathcal{S}(F)$ (if there is none, $\exists x\, F$ is false). If I has no lower bound, then $F_x^{-\infty}$ is true by construction (by Sect. 3, the virtual substitution $F_x^{-\infty}$ is equivalent to the literal substitution $F_x^{-\infty}$ in $\pm\infty$-extended real arithmetic). Otherwise, let $\alpha \in \mathbb{R}$ be the lower bound of I. If $\alpha \in I$ (I is closed at the lower bound), then α is of the form $-c/b, (-b+\sqrt{d})/(2a), (-b-\sqrt{d})/(2a)$ for some equation $(ax^2+bx+c=0) \in F$ or weak inequality $(ax^2+bx+c \leq 0) \in F$. Since the respective extra conditions on a, b, c, d hold, the quantifier-free formula evaluates to true. If, otherwise, $\alpha \notin I$ (I is open at the lower bound α), then α is of the form $-c/b, (-b+\sqrt{d})/(2a), (-b-\sqrt{d})/(2a)$ for some disequation $(ax^2+bx+c \neq 0) \in F$ or strict inequality $(ax^2+bx+c < 0) \in F$ and the interval I cannot be a single point. Thus, one of the infinitesimal increments $-c/b+\varepsilon, (-b+\sqrt{d})/(2a)+\varepsilon, (-b-\sqrt{d})/(2a)+\varepsilon$ is in $I \subseteq \mathcal{S}(F)$. Since the respective conditions a, b, c, d hold, the quantifier-free formula is again true. Hence, in either case, the quantifier-free formula is equivalent to $\exists x\, F$ in state ν. Since the state ν giving concrete real numbers to all free variables of $\exists x\, F$ was arbitrary, the same equivalence holds for all ν, which means that the quantifier-free formula (call it G) is equivalent to $\exists x\, F$. That is $G \leftrightarrow \exists x\, F$ is valid, i.e. $\vDash G \leftrightarrow \exists x\, F$. □

Optimizations are possible [Wei97] if there is only one quadratic occurrence of x, and that occurrence is not in an equation. If that occurrence is an equation, Theorem 1 already showed what to do. If there is only one occurrence of a quadratic inequality, the following variation works.

Note 7 ([Wei97]). *Let* $\left(Ax^2+Bx+C\left\{\begin{smallmatrix}\leq\\<\\\neq\end{smallmatrix}\right\}0\right) \in F$ *be the only quadratic occurrence of* x*. In that case,* $\exists x\, F$ *is equivalent over* $\mathbb{R}$ *to the following quantifier-free formula:*

$$\begin{aligned}
&A=0\wedge B\neq 0\wedge F_x^{-C/B} \vee A\neq 0\wedge F_x^{-B/(2A)}\\
\vee\ & F_x^{\infty} \vee F_x^{-\infty}\\
\vee\ & \bigvee_{\left(0x^2+bx+c\left\{\begin{smallmatrix}=\\\leq\end{smallmatrix}\right\}0\right)\in F} \left(b\neq 0\wedge F_x^{-c/b}\right)\\
\vee\ & \bigvee_{\left(0x^2+bx+c\left\{\begin{smallmatrix}\neq\\<\end{smallmatrix}\right\}0\right)\in F} \left(b\neq 0\wedge (F_x^{-c/b+\varepsilon}\vee F_x^{-c/b-\varepsilon})\right)
\end{aligned}$$

Further optimizations are possible if some signs of a, b are known, because several cases in the quantifier-free expansion then become impossible and can be simplified to *true* or *false* immediately. This helps simplify the formula in Theorem 4, because one of the cases $a = 0$ versus $a \neq 0$ might drop. But it also reduces the number of disjuncts in $F_x^{-\infty}$, see Example 3, and in the virtual substitutions of square roots (Lecture 18) and of infinitesimals (Sect. 4).

Theorem 4 also applies for polynomials of higher degrees in x if all those factor to polynomials of at most quadratic degree in x [Wei97]. Degree reduction is also possible by renaming based on the greatest common divisor of all powers of x that occur in F. If a quantified variable x occurs only with degrees that are multiples of an odd number d then virtual substitution can use $\exists x\, F(x^d) \equiv \exists y\, F(y)$. If x only occurs with degrees that are multiples of an even number d then $\exists x\, F(x^d) \equiv \exists y\, (y \geq 0 \wedge F(y))$. The cases with infinitesimals $+\varepsilon$ are only needed if x occurs in strict inequalities. The cases $F_x^{(-b+\pm\sqrt{d})/(2a)}$ are only needed if x occurs in equations or weak inequalities.

6 Summary

Virtual substitution is one technique for eliminating quantifiers in real arithmetic. It works for linear and quadratic constraints and can be extended to some cubic cases [Wei94]. Virtual substitution can be applied repeatedly from inside out to eliminate quantifiers. In each case, however, virtual substitution requires the eliminated variable to occur with small enough degrees only. Even if that was the case initially, it may stop to be the case after eliminating the innermost quantifier, because the degrees of the formulas resulting from virtual substitution may increase. In that case, degree optimizations and simplifications may sometimes work. If not, then other quantifier elimination techniques need to be used, which are based on semialgebraic geometry or model theory. Virtual substitution alone always works for mixed quadratic-linear formulas, i.e. those in which all quantified variables occur linearly except for one variable that occurs quadratically. In practice, however, many other cases turn out to work well with virtual substitution.

Exercises

Exercise 1. Consider

$$\exists x\,(ax^2+bx+c\leq 0\wedge F) \tag{12}$$

The virtual substitution of the roots of $ax^2+bx+c=0$ according to Sect. 2 as well as of $-\infty$ according to Sect. 3 will lead to

$$F_x^{-\infty}\vee a=0\wedge b\neq 0\wedge F_x^{-c/b}\vee a\neq 0\wedge b^2-4ac\geq 0\wedge\big(F_x^{(-b+\sqrt{b^2-4ac})/(2a)}\vee F_x^{(-b-\sqrt{b^2-4ac})/(2a)}\big)$$

But when F is $-ax^2+bx+e<0$, then none of those cases necessarily works. Does that mean the result of virtual substitution is not equivalent to (12)? Where is the catch in this argument?

References

[BD07] Christopher W. Brown and James H. Davenport. The complexity of quantifier elimination and cylindrical algebraic decomposition. In Dongming Wang, editor, *ISSAC*, pages 54–60. ACM, 2007.

[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.

[Pla12] André Platzer. Logics of dynamical systems. In *LICS*, pages 13–24. IEEE, 2012. doi:10.1109/LICS.2012.13.

[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.

[Wei94] Volker Weispfenning. Quantifier elimination for real algebra — the cubic case. In *ISSAC*, pages 258–263, 1994.

[Wei97] Volker Weispfenning. Quantifier elimination for real algebra — the quadratic case and beyond. *Appl. Algebra Eng. Commun. Comput.*, 8(2):85–101, 1997.