

Sums of Squares for Real-Closed Fields

John Harrison

Intel Corporation

CMU Seminar, Pittsburgh

Mon 19th March 2007

Summary

- The theory of reals and its universal fragment
- Nonnegativity via sum-of-squares
- Semidefinite programming
- The real Positivstellensatz
- Experiences
- The univariate case

The theory of reals

Consider the theory of reals (i.e. formulas true in $\mathbb{R}$) based on the following language:

- All rational constants p/q
- Operators of negation (' $-$ '), addition (' $+$ '), subtraction (' $-$ ') and multiplication (' $\cdot$ ')
- Relations ' $=$ ', ' $<$ ', ' $\leq$ ', ' $>$ ', ' $\geq$ '

An interesting theory that can express many nontrivial (indeed open) problems:

Kissing problem: how many disjoint n -dimensional spheres can be packed into space so that they touch a given unit sphere?

Axiomatizing the theory of reals (1)

$$1 \neq 0$$

$$\forall x y. x + y = y + x$$

$$\forall x y z. x + (y + z) = (x + y) + z$$

$$\forall x. 0 + x = x$$

$$\forall x. (-x) + x = 0$$

$$\forall x y. xy = yx$$

$$\forall x y z. x(yz) = (xy)z$$

$$\forall x. 1x = x$$

$$\forall x. x \neq 0 \Rightarrow x^{-1}x = 1$$

$$\forall x y z. x(y + z) = xy + xz$$

Axiomatizing the theory of reals (2)

Axioms for an *ordered* field:

$$\forall x y. x = y \vee x < y \vee y < x$$

$$\forall x y z. x < y \wedge y < z \Rightarrow x < z$$

$$\forall x. x \not< x$$

$$\forall y z. y < z \Rightarrow \forall x. x + y < x + z$$

$$\forall x y. 0 < x \wedge 0 < y \Rightarrow 0 < xy$$

and the *higher-order* axiom of completeness:

$$\forall S. (\exists x. x \in S) \wedge (\exists M. \forall x \in S. x \leq M)$$

$$\Rightarrow \exists m. (\forall x \in S. x \leq m) \wedge \forall m'. (\forall x \in S. x \leq m') \Rightarrow m \leq m'$$

These axioms are categorical, i.e. determine $\mathbb{R}$ up to isomorphism.

Real-closed fields

The theory of real-closed fields takes instead of completeness just the existence of square roots:

$$\forall x. x \geq 0 \Rightarrow \exists y. x = y^2$$

and that all polynomials of odd degree have a root (one of these for each odd n):

$$\forall a_0, \dots, a_n. a_n \neq 0 \Rightarrow \exists x. a_n x^n + a_{n-1} x^{n-1} + \dots + a_1 x + a_0 = 0$$

This theory is *not* categorical: other models include the computable real numbers. However, it *is* complete, i.e. determines all *first-order* consequences.

Completeness and decidability

Tarski proved in the 1930s that the theory of real-closed fields is complete and decidable, and even exhibited a quantifier elimination procedure for it.

This was only published in 1948 (by RAND!)

$$\mathbb{R} \models (\exists x. ax^2 + bx + c = 0) \Leftrightarrow a \neq 0 \wedge b^2 \geq 4ac \vee a = 0 \wedge (b \neq 0 \vee c = 0)$$

Collins's CAD algorithm is much more efficient and the first decision method actually to be implemented.

Some good implementations like qepcad and REDLOG, but theoretical and practical complexity issues limit its application.

Cohen-Hörmander algorithm is significantly simpler and has been implemented in Coq and HOL to generate proofs, but even slower.

The universal fragment

Many interesting problems fall into the purely universal fragment:

- Everyday trivialities like $\forall x y. x \geq 0 \wedge y \geq 0 \Rightarrow xy \geq 0$
- Polynomial bound problems like $\forall x \in [0, 1]. |p(x)| \leq k$ (used for some of my verifications).
- Most classical geometrical theorems

NB: geometry theorems with no use of ordering often turn out to be true over $\mathbb{C}$, which makes things easier.

Universality of real-closed fields

By the Artin-Schreier theory *every ordered field can be embedded in a real-closed field*.

This means that a *universal* formula holds in all real-closed fields iff it holds in all ordered fields, or even in all ordered integral domains.

So we will never need to use anything beyond the axioms for an ordered integral domain!

(Compare the case of fields in general: a universal formula holds in $\mathbb{C}$ iff it holds in all fields/integral domains of characteristic 0.)

Positivity

Consider first an even more special case of proving *positive semidefiniteness*:

$$\forall x_1, \dots, x_n. p(x_1, \dots, x_n) \geq 0$$

Not as limited as it may appear: can express polynomial bounds by change of variables like $x \mapsto \frac{y^2}{1+y^2}$

Illustrates the core techniques of SOS and SDP methods while avoiding some technicalities.

Sum-of-squares proofs

A *sufficient* condition for

$$\forall x_1, \dots, x_n. p(x_1, \dots, x_n) \geq 0$$

is the expressibility of p as a sum of squares (SOS)

$$p(x_1, \dots, x_n) = s_1(x_1, \dots, x_n)^2 + \dots + s_k(x_1, \dots, x_n)^2$$

In general it is *not* a necessary condition; a concrete counterexample is the *Motzkin form* $1 + x^4y^2 + x^2y^4 - 3x^2y^2$.

The solution to Hilbert's 17th problem shows that a polynomial is PSD iff it is a sum of squares of *rational* functions.

Sufficiency of sum-of-squares

PSD and SOS are equivalent in several special cases, the most important being

- Univariate polynomials of any degree
- Quadratic forms (all terms have degree exactly 2) in any number of variables ('complete the square')

Moreover, one can base complete approaches on various "Positivstellensatz" results that also depend essentially on sums of squares.

Example (problem)

Consider the following (Zeng et al, JSC vol 37, 2004, p83-99).

$$\forall w \ x \ y \ z. \ w^6 + 2z^2w^3 + x^4 + y^4 + z^4 + 2x^2w + 2x^2z + \\ 3x^2 + w^2 + 2zw + z^2 + 2z + 2w + 1 \geq 0$$

Constraint problems of this sort are in general quite hard to solve.

Example (solution)

We can express the polynomial as a SOS:

$$\begin{aligned} &w^6 + 2z^2w^3 + x^4 + y^4 + z^4 + 2x^2w + 2x^2z + \\ &3x^2 + w^2 + 2zw + z^2 + 2z + 2w + 1 = \\ &(y^2)^2 + (x^2 + w + z + 1)^2 + x^2 + (w^3 + z^2)^2 \end{aligned}$$

Note how nice this is for LCF-style proving: the SOS decomposition can be checked without any tricky decision procedures.

But how do we find the SOS decomposition? By semidefinite programming (SDP)!

Reduction to quadratic form

By introducing new variables for monomials, we can express a polynomial as a quadratic form subject to linear constraints. Example:

$$2x^4 + 2x^3y - x^2y^2 + 5y^4$$

We consider all monomials (only need homogenous ones since original is a form): $z_1 = x^2$, $z_2 = y^2$, $z_3 = xy$ and write the polynomial as a quadratic form. In matrix notation:

$$\begin{bmatrix} z_1 \\ z_2 \\ z_3 \end{bmatrix}^T \begin{bmatrix} q_{11} & q_{12} & q_{13} \\ q_{12} & q_{22} & q_{23} \\ q_{13} & q_{23} & q_{33} \end{bmatrix} \begin{bmatrix} z_1 \\ z_2 \\ z_3 \end{bmatrix}$$

Linear parametrization

By comparing coefficients we get linear constraints; in this case we end up with only one parameter.

$$q_{11} = 5$$

$$q_{22} = 5$$

$$q_{33} + 2q_{12} = -1$$

$$2q_{13} = 2$$

$$2q_{23} = 0$$

In general we'll get more, but the key point is that the parametrization is linear.

Semidefinite programming

For quadratic forms, being PSD is equivalent to being SOS.

Finding a parametrization making a matrix PSD, subject to (and optimizing) linear constraints is a standard problem called *semidefinite programming*.

The problem is polynomial-time solvable using interior-point algorithms.

There are many efficient tools to solve the problem effectively in practice. I mostly use CSDP.

The usual Nullstellensatz

Over algebraically closed fields like $\mathbb{C}$ we have a nice simple equivalence.

The polynomial equations $p_1(\overline{x}) = 0, \dots, p_k(\overline{x}) = 0$ in an algebraically closed field have *no* common solution iff there are polynomials $q_1(\overline{x}), \dots, q_k(\overline{x})$ such that the following polynomial identity holds:

$$q_1(\overline{x}) \cdot p_1(\overline{x}) + \dots + q_k(\overline{x}) \cdot p_k(\overline{x}) = 1$$

Thus we can reduce equation-solving to ideal membership and solve it efficiently using Gröbner bases.

The real Nullstellensatz

In the analogous Nullstellensatz result over $\mathbb{R}$, sums of squares play a central role:

The polynomial equations $p_1(\bar{x}) = 0, \dots, p_k(\bar{x}) = 0$ in a real closed field have *no* common solution iff there are polynomials $q_1(\bar{x}), \dots, q_k(\bar{x}), s_1(\bar{x}), \dots, s_m(\bar{x})$ such that

$$q_1(\bar{x}) \cdot p_1(\bar{x}) + \dots + q_k(\bar{x}) \cdot p_k(\bar{x}) + s_1(\bar{x})^2 + \dots + s_m(\bar{x})^2 = -1$$

SDP can also solve this more general problem, either by linear parametrization of possible $q_i(\bar{x})$ or combining with Gröbner bases.

Real Positivstellensatz

There are still more general “Positivstellensatz” results about the inconsistency of a set of equations, negated equations, strict and non-strict inequalities.

Because there are so many different kinds of hypothesis, the exact statement looks a bit daunting.

But here’s a simple example: prove

$$\forall a \ b \ c \ x. ax^2 + bx + c = 0 \Rightarrow b^2 - 4ac \geq 0$$

via the following SOS certificate:

$$b^2 - 4ac = (2ax + b)^2 - 4a(ax^2 + bx + c)$$

Experience and problems

This approach is often much more efficient than competing techniques such as general quantifier elimination.

Lends itself very well to a separation of proof search and LCF-style checking, so fits very well with HOL Light.

Available with HOL Light since 2.0 in `Examples/sos.ml`, and seems quite useful.

Still some awkward numerical problems where the PSD is tight (can become zero) and the rounding to rationals causes loss of PSD-ness.

The univariate case

Alternative based on the simple observation that every nonnegative univariate polynomial is a sum of squares of *real* polynomials.

All roots, real or complex, must occur in conjugate pairs. Thus the polynomial is a product of factors

$$(x - [a_k + ib_k])(x - [a_k - ib_k])$$

and so is of the form

$$(q(x) + ir(x))(q(x) - ir(x)) = q(x)^2 + r(x)^2$$

To get an exact rational decomposition, we need a more intricate algorithm, but this is the basic idea.

Experience of univariate case

Numerical problems can be particularly annoying with some polynomial bound problems in real applications where the coefficients are non-trivial (60-200 bits).

For example, proving $\forall x. |x| \leq k \Rightarrow |f(x) - p(x)| < \epsilon$ where p is a short approximation to a longer polynomial f .

The direct approach is often better than SDP-based methods, for numerical reasons, in such examples.

General conclusion

There's often a lack of communication between researchers in theorem proving and in related fields.

Sometimes we can get important ideas from other theorem provers (e.g. declarative proof from Mizar).

We can learn lots of useful ideas from computer algebra, and even exploit the systems themselves (e.g. Analytica).

As this work shows, there is also a lot of interesting stuff out there in the optimization field that we may be able to exploit.

But a high-precision SDP solver would be desirable!