Logical Decision Procedures in Practice
1: Background & Propositional Logic

John Harrison

Intel Corporation

Marktoberdorf 2005

Wed 10th August 2005 (10:35 – 11:20)

What I will talk about

Aim is to cover some of the most important decidable problems in classical logic, with an emphasis on practical usefulness.

1. Background and propositional logic

2. First-order logic and arithmetical theories

3. Real quantifier elimination

4. Combination and certification of decision procedures

What I won't talk about

- Decision procedures for temporal logic, model checking (well covered in other courses)
- Higher-order logic and interactive theorem proving (my own interest but off the main topic)
- Undecidability and incompleteness (I don't have enough time; there is some material in the notes).
- Decision methods for constructive logic, modal logic, other nonclassical logics (I don't know much anyway)

A practical slant

Our approach to logic will be highly constructive!

Most of what is described is implemented by explicit code that can be obtained here:

`http://www.cl.cam.ac.uk/users/jrh/atp/`

See also my interactive higher-order logic prover HOL Light:

`http://www.cl.cam.ac.uk/users/jrh/hol-light/`

which incorporates many decision procedures in a certified way.

Propositional Logic

We probably all know what propositional logic is.

English	Standard	Boolean	Other
false	$\bot$	0	F
true	$\top$	1	T
not p	$\neg p$	$\overline{p}$	$-p$, $\sim p$
p and q	$p \wedge q$	pq	$p\&q$, $p \cdot q$
p or q	$p \vee q$	$p + q$	$p \mid q$, $p\ or\ q$
p implies q	$p \Rightarrow q$	$p \leq q$	$p \longrightarrow q$, $p \supset q$
p iff q	$p \Leftrightarrow q$	$p = q$	$p \equiv q$, $p \sim q$

In the context of circuits, it's often referred to as 'Boolean algebra', and many designers use the Boolean notation.

Is propositional logic boring?

Traditionally, propositional logic has been regarded as fairly boring.

- There are severe limitations to what can be said with propositional logic.
- Propositional logic is trivially decidable in theory
- The usual methods aren't efficient enough for interesting problems.

But . . .

No!

The last decade has seen a remarkable upsurge of interest in propositional logic.

In fact, it's arguably the hottest topic in automated theorem proving!

Why the resurgence?

- There *are* many interesting problems that can be expressed in propositional logic
- Efficient algorithms *can* often decide large, interesting problems

A practical counterpart to the theoretical reductions in NP-completeness theory.

Logic and circuits

The correspondence between digital logic circuits and propositional logic has been known for a long time.

Digital design	Propositional Logic
circuit	formula
logic gate	propositional connective
input wire	atom
internal wire	subexpression
voltage level	truth value

Many problems in circuit design and verification can be reduced to propositional tautology or satisfiability checking ('SAT').

For example optimization correctess: $\phi \Leftrightarrow \phi'$ is a tautology.

Combinatorial problems

Many other apparently difficult combinatorial problems can be encoded as Boolean satisfiability (SAT), e.g. scheduling, planning, even factorization.

$$
\begin{aligned}
\neg(\quad & (out_0 \Leftrightarrow x_0 \wedge y_0) \wedge \\
& (out_1 \Leftrightarrow (x_0 \wedge y_1 \Leftrightarrow \neg(x_1 \wedge y_0))) \wedge \\
& (v_2^2 \Leftrightarrow (x_0 \wedge y_1) \wedge x_1 \wedge y_0) \wedge \\
& (u_2^0 \Leftrightarrow ((x_1 \wedge y_1) \Leftrightarrow \neg v_2^2)) \wedge \\
& (u_2^1 \Leftrightarrow (x_1 \wedge y_1) \wedge v_2^2) \wedge \\
& (out_2 \Leftrightarrow u_2^0) \wedge (out_3 \Leftrightarrow u_2^1) \wedge \\
& \neg out_0 \wedge out_1 \wedge out_2 \wedge \neg out_3)
\end{aligned}
$$

Read off the factorization $6 = 2 \times 3$ from a refuting assignment.

Efficient methods

The naive truth table method is quite impractical for formulas with more than a dozen primitive propositions.

Practical use of propositional logic mostly relies on one of the following algorithms for deciding tautology or satisfiability:

- Binary decision diagrams (BDDs)
- The Davis-Putnam method (DP, DPLL)
- Stålmarck's method

We'll sketch the basic ideas behind Davis-Putnam and Stålmarck's method.

DP and DPLL

Actually, the original Davis-Putnam procedure is not much used now.

What is usually called the Davis-Putnam method is actually a later refinement due to Davis, Loveland and Logemann (hence DPLL).

We formulate it as a test for *satisfiability*. It has three main components:

- Transformation to conjunctive normal form (CNF)
- Application of simplification rules
- Splitting

Normal forms

In ordinary algebra we can reach a 'sum of products' form of an expression by:

- Eliminating operations other than addition, multiplication and negation, e.g. $x - y \mapsto x + -y$.
- Pushing negations inwards, e.g. $-(-x) \mapsto x$ and $-(x + y) \mapsto -x + -y$.
- Distributing multiplication over addition, e.g. $x(y + z) \mapsto xy + xz$.

In logic we can do exactly the same, e.g. $p \Rightarrow q \mapsto \neg p \vee q$, $\neg(p \wedge q) \mapsto \neg p \vee \neg q$ and $p \wedge (q \vee r) \mapsto (p \wedge q) \vee (p \wedge r)$.

The first two steps give 'negation normal form' (NNF).

Following with the last (distribution) step gives 'disjunctive normal form' (DNF), analogous to a sum-of-products.

Conjunctive normal form

Conjunctive normal form (CNF) is the dual of DNF, where we reverse the roles of ‘and’ and ‘or’ in the distribution step to reach a ‘product of sums’:

$$\begin{array}{rcl} p \lor (q \land r) & \mapsto & (p \lor q) \land (p \lor r) \\ (p \land q) \lor r & \mapsto & (p \lor r) \land (q \lor r) \end{array}$$

Reaching such a CNF is the first step of the Davis-Putnam procedure.

Unfortunately the naive distribution algorithm can cause the size of the formula to grow exponentially — not a good start. Consider for example:

$$(p_1 \land p_2 \land \cdots \land p_n) \lor (q_1 \land p_2 \land \cdots \land q_n)$$

Definitional CNF

A cleverer approach is to introduce new variables for subformulas. Although this isn't logically equivalent, it does preserve satisfiability.

$$(p \lor (q \land \neg r)) \land s$$

introduce new variables for subformulas:

$$(p_1 \Leftrightarrow q \land \neg r) \land (p_2 \Leftrightarrow p \lor p_1) \land (p_3 \Leftrightarrow p_2 \land s) \land p_3$$

then transform to (3-)CNF in the usual way:

$$(\neg p_1 \lor q) \land (\neg p_1 \lor \neg r) \land (p_1 \lor \neg q \lor r) \land$$

$$(\neg p_2 \lor p \lor p_1) \land (p_2 \lor \neg p) \land (p_2 \lor \neg p_1) \land$$

$$(\neg p_3 \lor p_2) \land (\neg p_3 \lor s) \land (p_3 \lor \neg p_2 \lor \neg s) \land p_3$$

Clausal form

It's convenient to think of the CNF form as a set of sets:

- Each disjunction $p_1 \vee \cdots \vee p_n$ is thought of as the set $\{p_1, \ldots, p_n\}$, called a *clause*.
- The overall formula, a conjunction of clauses $C_1 \wedge \cdots \wedge C_m$ is thought of as a set $\{C_1, \ldots, C_m\}$.

Since 'and' and 'or' are associative, commutative and idempotent, nothing of logical significance is lost in this interpretation.

Special cases: an empty clause means $\bot$ (and is hence unsatisfiable) and an empty set of clauses means $\top$ (and is hence satisfiable).

Simplification rules

At the core of the Davis-Putnam method are two transformations on the set of clauses:

I The 1-literal rule: if a unit clause p appears, remove $\neg p$ from other clauses and remove all clauses including p.

II The affirmative-negative rule: if p occurs *only* negated, or *only* unnegated, delete all clauses involving p.

These both preserve satisfiability of the set of clause sets.

Splitting

In general, the simplification rules will not lead to a conclusion. We need to perform case splits.

Given a clause set Δ, simply choose a variable p, and consider the two new sets $\Delta \cup \{p\}$ and $\Delta \cup \{\neg p\}$.

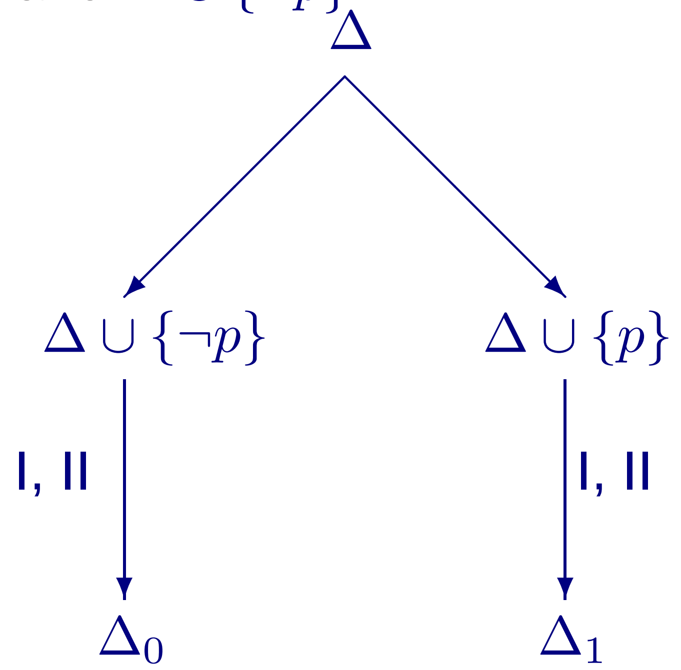

In general, these case-splits need to be nested.

Industrial strength SAT solvers

For big applications, there are several important modifications to the basic DPLL algorithm:

- Highly efficient data structures
- Good heuristics for picking ‘split’ variables
- Intelligent non-chronological backtracking / conflict clauses

Some well-known provers are GRASP, SATO, Chaff and BerkMin.

These often shine because of careful attention to low-level details like memory hierarchy, not cool algorithmic ideas.

Stålmarck's algorithm

Stålmarck's 'dilemma' rule attempts to avoid nested case splits by feeding back common information from both branches.

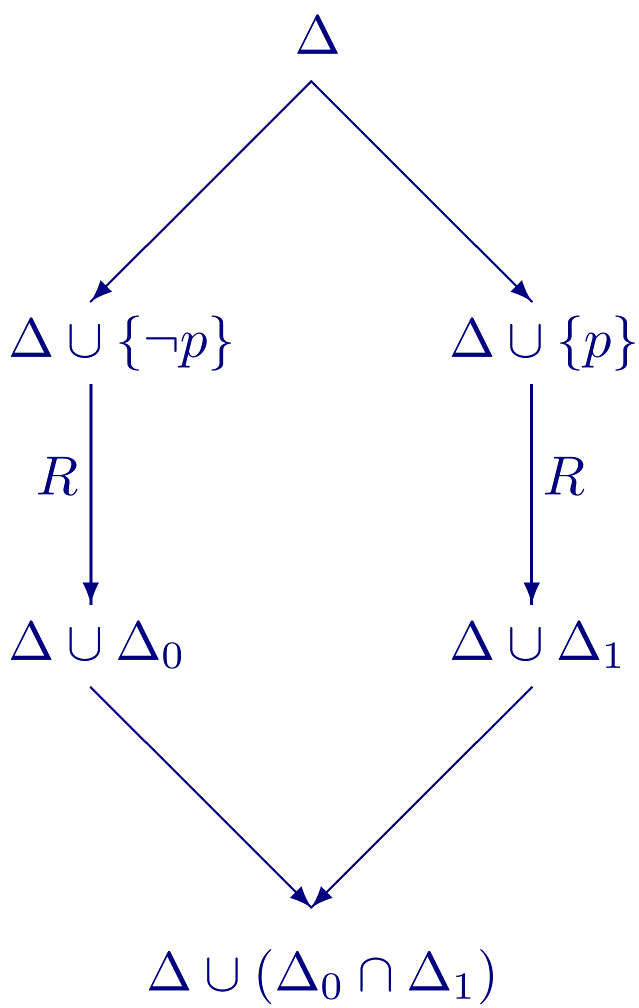

Summary

- Propositional logic is no longer a neglected area of theorem proving
- A wide variety of practical problems can usefully be encoded in SAT
- There is intense interest in efficient algorithms for SAT
- Many of the most successful systems are still based on minor refinements of the ancient Davis-Putnam procedure
- Can we invent a better SAT algorithm?

Wednesday puzzle

A problem in digital circuit design due to E. Snow (Intel). Show how to construct a digital circuit with three inputs:

i_1, i_2, i_3

and three outputs:

o_1, o_2, o_3

satisfying the specification:

$$(o_1 \Leftrightarrow \neg i_1) \wedge (o_2 \Leftrightarrow \neg i_2) \wedge (o_3 \Leftrightarrow \neg i_3)$$

subject to the constraint that you can use *at most two* ‘NOT’ gates (inverters), but any number of ‘AND’ and ‘OR’ gates.

Logical Decision Procedures in Practice
2: First-order logic and arithmetical theories

John Harrison

Intel Corporation

Marktoberdorf 2005

Thu 11th August 2005 (10:35 – 11:20)

Summary

- First order logic
- Naive Herbrand procedures
- Unification
- Decidable classes
- Decidable theories
- Quantifier elimination

First-order logic

Start with a set of *terms* built up from variables and constants using function application:

$$x + 2 \cdot y \quad \equiv \quad +(x, \cdot(2(), y))$$

Create atomic formulas by applying relation symbols to a set of terms

$$x > y \quad \equiv \; > (x, y)$$

Create complex formulas using quantifiers

- $\forall x.\ P[x]$ — for all x, $P[x]$
- $\exists x.\ P[x]$ — there exists an x such that $P[x]$

Quantifier examples

The order of quantifier nesting is important. For example

$\forall x.\ \exists y.\ loves(x, y)$ — everyone loves someone
$\exists x.\ \forall y.\ loves(x, y)$ — somebody loves everyone
$\exists y.\ \forall x.\ loves(x, y)$ — someone is loved by everyone

This says that a function $\mathbb{R} \to \mathbb{R}$ is continuous:

$$\forall \epsilon.\ \epsilon > 0 \Rightarrow \forall x.\ \exists \delta.\ \delta > 0 \wedge \forall x'.\ |x' - x| < \delta \Rightarrow |f(x') - f(x)| < \varepsilon$$

while this one says it is *uniformly* continuous, an important distinction

$$\forall \epsilon.\ \epsilon > 0 \Rightarrow \exists \delta.\ \delta > 0 \wedge \forall x.\ \forall x'.\ |x' - x| < \delta \Rightarrow |f(x') - f(x)| < \varepsilon$$

Skolemization

Skolemization relies on this observation (related to the axiom of choice):

$$(\forall x.\ \exists y.\ P[x, y]) \Leftrightarrow \exists f.\ \forall x.\ P[x, f(x)]$$

For example, a function is surjective (onto) iff it has a right inverse:

$$(\forall x.\ \exists y.\ g(y) = x) \Leftrightarrow (\exists f.\ \forall x.\ g(f(x)) = x$$

Can't quantify over functions in first-order logic.

But we get an *equisatisfiable* formula if we just introduce a new function symbol.

$$\begin{array}{l} \forall x_1, \ldots, x_n.\ \exists y.\ P[x_1, \ldots, x_n, y] \\ \to \forall x_1, \ldots, x_n.\ P[x_1, \ldots, x_n, f(x_1, \ldots, x_n)] \end{array}$$

Now we just need a satisfiability test for universal formulas.

First-order automation

The underlying domains can be arbitrary, so we can't do an exhaustive analysis, but must be slightly subtler.

We can reduce the problem to propositional logic using the so-called *Herbrand theorem*:

> Let $\forall x_1, \ldots, x_n.\ P[x_1, \ldots, x_n]$ be a first order formula with only the indicated universal quantifiers (i.e. the body $P[x_1, \ldots, x_n]$ is quantifier-free). Then the formula is satisfiable iff the infinite set of 'ground instances' $P[t_1^i, \ldots, t_n^i]$ that arise by replacing the variables by arbitrary variable-free terms made up from functions and constants in the original formula is *propositionally* satisfiable.

Still only gives a *semidecision* procedure, a kind of proof search.

Example

Suppose we want to prove the 'drinker's principle'

$$\exists x.\, \forall y.\, D(x) \Rightarrow D(y)$$

Negate the formula, and prove negation unsatisfiable:

$$\neg(\exists x.\, \forall y.\, D(x) \Rightarrow D(y))$$

Convert to prenex normal form: $\forall x.\, \exists y.\, D(x) \wedge \neg D(y)$

Skolemize: $\forall x.\, D(x) \wedge \neg D(f(x))$

Enumerate set of ground instances, first $D(c) \wedge \neg D(f(c))$ is not unsatisfiable, but the next is:

$$(D(c) \wedge \neg D(f(c))) \wedge (D(f(c)) \wedge \neg D(f(f(c)))$$

Unification

The first automated theorem provers actually used that approach.

It was to test the propositional formulas resulting from the set of ground-instances that the Davis-Putnam method was developed.

However, more efficient than enumerating ground instances is to use *unification* to choose instantiations intelligently.

Many theorem-proving algorithms based on unification exist:

- Tableaux
- Resolution
- Model elimination
- Connection method
- ...

Decidable problems

Although first order validity is undecidable, there are special cases where it is decidable, e.g.

- AE formulas: no function symbols, universal quantifiers before existentials in prenex form (so finite Herbrand base).
- Monadic formulas: no function symbols, only unary predicates

These are not particularly useful in practice, though they can be used to automate syllogistic reasoning.

If all M are P, and all S are M, then all S are P

can be expressed as the monadic formula:

$$(\forall x.\ M(x) \Rightarrow P(x)) \wedge (\forall x.\ S(x) \Rightarrow M(x)) \Rightarrow (\forall x.\ S(x) \Rightarrow P(x))$$

The theory of equality

A simple but useful decidable theory is the universal theory of equality with function symbols, e.g.

$$\forall x.\ f(f(f(x)) = x \land f(f(f(f(f(x))))) = x \Rightarrow f(x) = x$$

after negating and Skolemizing we need to test a ground formula for satisfiability:

$$f(f(f(c)) = c \land f(f(f(f(f(c))))) = c \land \neg(f(c) = c)$$

Two well-known algorithms:

- Put the formula in DNF and test each disjunct using one of the classic 'congruence closure' algorithms.
- Reduce to SAT by introducing a propositional variable for each equation between subterms and adding constraints.

Decidable theories

More useful in practical applications are cases not of *pure* validity, but validity in special (classes of) models, or consequence from useful axioms, e.g.

- Does a formula hold over all rings (Boolean rings, non-nilpotent rings, integral domains, fields, algebraically closed fields, . . .)
- Does a formula hold in the natural numbers or the integers?
- Does a formula hold over the real numbers?
- Does a formula hold in all real-closed fields?
- . . .

Because arithmetic comes up in practice all the time, there's particular interest in theories of arithmetic.

Theories

These can all be subsumed under the notion of a *theory*, a set of formulas T closed under logical validity. A theory T is:

- *Consistent* if we never have $p \in T$ and $(\neg p) \in T$.
- *Complete* if for closed p we have $p \in T$ or $(\neg p) \in T$.
- *Decidable* if there's an algorithm to tell us whether a given closed p is in T

Note that a complete theory generated by an r.e. axiom set is also decidable.

Quantifier elimination

Often, a quantified formula is T-equivalent to a quantifier-free one:

- $\mathbb{C} \models (\exists x.\ x^2 + 1 = 0) \Leftrightarrow \top$
- $\mathbb{R} \models (\exists x. ax^2 + bx + c = 0) \Leftrightarrow a \neq 0 \land b^2 \geq 4ac \lor a = 0 \land (b \neq 0 \lor c = 0)$
- $\mathbb{Q} \models (\forall x.\ x < a \Rightarrow x < b) \Leftrightarrow a \leq b$
- $\mathbb{Z} \models (\exists k\ x\ y.\ ax = (5k + 2)y + 1) \Leftrightarrow \neg(a = 0)$

We say a theory T admits *quantifier elimination* if *every* formula has this property.

Assuming we can decide variable-free formulas, quantifier elimination implies completeness.

And then an *algorithm* for quantifier elimination gives a decision method.

Important arithmetical examples

- Presburger arithmetic: arithmetic equations and inequalities with addition but *not multiplication*, interpreted over $\mathbb{Z}$ or $\mathbb{N}$.
- Tarski arithmetic: arithmetic equations and inequalities with addition and multiplication, interpreted over $\mathbb{R}$ (or any real-closed field)
- General algebra: arithmetic equations with addition and multiplication interpreted over $\mathbb{C}$ (or other algebraically closed field).

However, arithmetic with multiplication over $\mathbb{Z}$ is not even semidecidable, by Gödel's theorem.

Nor is arithmetic over $\mathbb{Q}$ (Julia Robinson), nor just solvability of equations over $\mathbb{Z}$ (Matiyasevich). Equations over $\mathbb{Q}$ unknown.

Summary

- Can't solve first-order logic by naive method, but Herbrand's theorem gives a proof search procedure
- Unification is normally a big improvement on straightforward search through the Herbrand base
- A few fragments of first-order logic are decidable, but few are very useful.
- We are often more interested in arithmetic theories than pure logic
- Quantifier elimination usually gives a nice decision method and more

Thursday puzzle

Here is a 'non-obvious' fact of first-order logic due to Łoś:

$$(\forall x\ y\ z.\ P(x,y) \land P(y,z) \Rightarrow P(x,z)) \land$$
$$(\forall x\ y\ z.\ Q(x,y) \land Q(y,z) \Rightarrow Q(x,z)) \land$$
$$(\forall x\ y.\ Q(x,y) \Rightarrow Q(y,x)) \land$$
$$(\forall x\ y.\ P(x,y) \lor Q(x,y))$$
$$\Rightarrow (\forall x\ y.\ P(x,y)) \lor (\forall x\ y.\ Q(x,y))$$

It's very easy for most automated theorem provers, and if intelligently prenexed, falls in the decidable AE subset.

Can you find a short intuitive 'human' proof?

Logical Decision Procedures in Practice
3: Real quantifier elimination

John Harrison

Intel Corporation

Marktoberdorf 2005

Fri 12th August 2005 (10:35 – 11:20)

Summary

- What we'll prove
- History
- Sign matrices
- The key recursion
- Parametrization
- Real-closed fields

What we'll prove

Take a first-order language:

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

We'll prove that every formula in the language has a quantifier-free equivalent, and will give a systematic algorithm for finding it.

Applications

In principle, this method can be used to solve many non-trivial problems.

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

Pretty much *any* geometrical assertion can be expressed in this theory.

If theorem holds for *complex* values of the coordinates, and then simpler methods are available (Gröbner bases, Wu-Ritt triangulation. . .).

History

- 1930: Tarski discovers quantifier elimination procedure for this theory.
- 1948: Tarski's algorithm published by RAND
- 1954: Seidenberg publishes simpler algorithm
- 1975: Collins develops and *implements* cylindrical algebraic decomposition (CAD) algorithm
- 1983: Hörmander publishes very simple algorithm based on ideas by Cohen.
- 1990: Vorobjov improves complexity bound to doubly exponential in number of quantifier *alternations*.

We'll present the Cohen-Hörmander algorithm.

Current implementations

There are quite a few simple versions of real quantifier elimination, even in computer algebra systems like Mathematica.

Among the more heavyweight implementations are:

- `qepcad` — `http://www.cs.usna.edu/~qepcad/B/QEPCAD.html`
- `REDLOG` — `http://www.fmi.uni-passau.de/~redlog/`

One quantifier at a time

For a general quantifier elimination procedure, we just need one for a formula

$$\exists x.\ P[a_1, \ldots, a_n, x]$$

where $P[a_1, \ldots, a_n, x]$ involves no other quantifiers but may involve other variables.

Then we can apply the procedure successively inside to outside, dealing with universal quantifiers via $(\forall x.\ P[x]) \Leftrightarrow (\neg \exists x.\ \neg P[x])$.

Forget parametrization for now

First we'll ignore the fact that the polynomials contain variables other than the one being eliminated.

This keeps the technicalities a bit simpler and shows the main ideas clearly.

The generalization to the parametrized case will then be very easy:

- Replace polynomial division by pseudo-division
- Perform case-splits to determine signs of coefficients

Sign matrices

Take a set of univariate polynomials $p_1(x)$, ..., $p_n(x)$.

A *sign matrix* for those polynomials is a division of the real line into alternating points and intervals:

$$(-\infty, x_1), x_1, (x_1, x_2), x_2, \ldots, x_{m-1}, (x_{m-1}, x_m), x_m, (x_m, +\infty)$$

and a matrix giving the sign of each polynomial on each interval:

- Positive ($+$)
- Negative ($-$)
- Zero (0)

Sign matrix example

The polynomials $p_1(x) = x^2 - 3x + 2$ and $p_2(x) = 2x - 3$ have the following sign matrix:

Point/Interval	p_1	p_2
$(-\infty, x_1)$	$+$	$-$
x_1	0	$-$
(x_1, x_2)	$-$	$-$
x_2	$-$	0
(x_2, x_3)	$-$	$+$
x_3	0	$+$
$(x_3, +\infty)$	$+$	$+$

Using the sign matrix

Using the sign matrix for all polynomials appearing in $P[x]$ we can answer any quantifier elimination problem: $\exists x.\ P[x]$

- Look to see if any row of the matrix satisfies the formula (hence dealing with existential)
- For each row, just see if the corresponding set of signs satisfies the formula.

We have replaced the quantifier elimination problem with sign matrix determination

Finding the sign matrix

For constant polynomials, the sign matrix is trivial (2 has sign '$+$' etc.)

To find a sign matrix for $p, p_1, \ldots, p_n$ it suffices to find one for $p', p_1, \ldots, p_n, r_0, r_1, \ldots, r_n$, where

- $p_0 \equiv p'$ is the derivative of p
- $r_i = \text{rem}(p, p_i)$

(Remaindering means we have some q_i so $p = q_i \cdot p_i + r_i$.)

Taking p to be the polynomial of highest degree we get a simple recursive algorithm for sign matrix determination.

Details of recursive step

So, suppose we have a sign matrix for $p', p_1, \ldots, p_n, r_0, r_1, \ldots, r_n$.

We need to construct a sign matrix for $p, p_1, \ldots, p_n$.

- May need to add more points and hence intervals for roots of p
- Need to determine signs of $p_1, \ldots, p_n$ at the new points and intervals
- Need the sign of p itself everywhere.

Step 1

Split the given sign matrix into two parts, but keep all the points for now:

- M for $p', p_1, \ldots, p_n$
- M' for $r_0, r_1, \ldots, r_n$

We can infer the sign of p at all the 'significant' *points* of M as follows:

$$p = q_i p_i + r_i$$

and for each of our points, one of the p_i is zero, so $p = r_i$ there and we can read off p's sign from r_i's.

Step 2

Now we're done with M' and we can throw it away.

We also 'condense' M by eliminating points that are not roots of one of the $p', p_1, \ldots, p_n$.

Note that the sign of any of these polynomials is stable on the condensed intervals, since they have no roots there.

- We know the sign of p at all the points of this matrix.
- However, p itself may have additional roots, and we don't know anything about the intervals yet.

Step 3

There can be at most one root of p in each of the existing intervals, because otherwise p' would have a root there.

We can tell whether there is a root by checking the signs of p (determined in Step 1) at the two endpoints of the interval.

Insert a new point precisely if p has strictly opposite signs at the two endpoints (simple variant for the two end intervals).

None of the other polynomials change sign over the original interval, so just copy the values to the point and subintervals.

Throw away p' and we're done!

Multivariate generalization

In the multivariate context, we can't simply divide polynomials. Instead of

$$p = p_i \cdot q_i + r_i$$

we get

$$a^k p = p_i \cdot q_i + r_i$$

where a is the leading coefficient of p_i.

The same logic works, but we need case splits to fix the sign of a.

Real-closed fields

With more effort, all the 'analytical' facts can be deduced from the axioms for *real-closed fields*.

- Usual ordered field axioms
- Existence of square roots: $\forall x.\ x \geq 0 \Rightarrow \exists y.\ x = y^2$
- Solvability of odd-degree equations:
  $\forall a_0, \ldots, a_n.\ a_n \neq 0 \Rightarrow \exists x.\ a_n x^n + a_{n-1}x^{n-1} + \cdots + a_1 x + a_0 = 0$

Examples include computable reals and algebraic reals. So this already gives a complete theory, without a stronger completeness axiom.

Summary

- Real quantifier elimination one of the most significant logical decidability results known.
- Original result due to Tarski, for general real closed fields.
- A half-century of research has resulted in simpler and more efficient algorithms (not always at the same time).
- The Cohen-Hörmander algorithm is remarkably simple (relatively speaking).
- The complexity, both theoretical and practical, is still bad, so there's limited success on non-trivial problems.

Friday puzzle

A famous example in real quantifier elimination is the *Kahan ellipse problem*.

$$\exists x\ y.\ a^2(x-c)^2 + b^2(y-d)^2 - 1 = 0 \land x^2 + y^2 > 1$$

This is asking for conditions on the parameter of an ellipse for it to lie inside or outside the unit circle.

It only contains two quantifiers. Nevertheless this is a significant challenge for most quantifier elimination algorithms. Can you find a nice quantifier-free equivalent?

Logical Decision Procedures in Practice
4: Combination and certification of decision procedures

John Harrison

Intel Corporation

Marktoberdorf 2005

Sat 13th August 2005 (10:35 – 11:20)

Summary

- Need to combine multiple decision procedures
- Basics of Nelson-Oppen method
- Proof-producing decision procedures
- Separate certification
- LCF-style implementation and reflection

Need for combinations

In applications we often need to combine decision methods from different domains.

$x - 1 < n \wedge \neg(x < n) \Rightarrow a[x] = a[n]$

An arithmetic decision procedure could easily prove

$x - 1 < n \wedge \neg(x < n) \Rightarrow x = n$

but could not make the additional final step, even though it looks trivial.

Most combinations are undecidable

Adding almost any additions, especially uninterpreted, to the usual decidable arithmetic theories destroys decidability.

Some exceptions like BAPA ('Boolean algebra + Presburger arithmetic').

This formula over the reals constrains P to define the integers:

$$(\forall n.\, P(n+1) \Leftrightarrow P(n)) \land (\forall n.\, 0 \leq n \land n < 1 \Rightarrow (P(n) \Leftrightarrow n = 0))$$

and this one in Presburger arithmetic defines squaring:

$$(\forall n.\, f(-n) = f(n)) \land (f(0) = 0) \land$$
$$(\forall n.\, 0 \leq n \Rightarrow f(n+1) = f(n) + n + n + 1)$$

and so we can define multiplication.

Quantifier-free theories

However, if we stick to so-called ‘quantifier-free’ theories, i.e. deciding universal formulas, things are better.

Two well-known methods for combining such decision procedures:

- Nelson-Oppen
- Shostak

Nelson-Oppen is more general and conceptually simpler.

Shostak seems more efficient where it does work, and only recently has it really been understood.

Nelson-Oppen basics

Key idea is to combine theories $T_1, \ldots, T_n$ with *disjoint signatures*. For instance

- T_1: numerical constants, arithmetic operations
- T_2: list operations like cons, head and tail.
- T_3: other uninterpreted function symbols.

The only common function or relation symbol is '$=$'.

This means that we only need to share formulas built from equations among the component decision procedure, thanks to the *Craig interpolation theorem*.

The interpolation theorem

Several slightly different forms; we'll use this one (by compactness, generalizes to theories):

> If $\models \phi_1 \wedge \phi_2 \Rightarrow \bot$ then there is an 'interpolant' ψ, whose only free variables and function and predicate symbols are those occurring in *both* ϕ_1 and ϕ_2, such that $\models \phi_1 \Rightarrow \psi$ and $\models \phi_2 \Rightarrow \neg\psi$.

This is used to assure us that the Nelson-Oppen method is complete, though we don't need to produce general interpolants in the method.

In fact, interpolants can be found quite easily from proofs, including Herbrand-type proofs produced by resolution etc.

Nelson-Oppen I

Proof by example: refute the following formula in a mixture of Presburger arithmetic and uninterpreted functions:

$$f(v-1) - 1 = v + 1 \land f(u) + 1 = u - 1 \land u + 1 = v$$

First step is to *homogenize*, i.e. get rid of atomic formulas involving a mix of signatures:

$$u + 1 = v \land v_1 + 1 = u - 1 \land v_2 - 1 = v + 1 \land v_2 = f(v_3) \land v_1 = f(u) \land v_3 = v - 1$$

so now we can split the conjuncts according to signature:

$$(u + 1 = v \land v_1 + 1 = u - 1 \land v_2 - 1 = v + 1 \land v_3 = v - 1) \land$$
$$(v_2 = f(v_3) \land v_1 = f(u))$$

Nelson-Oppen II

If the entire formula is contradictory, then there's an interpolant ψ such that in Presburger arithmetic:

$$\mathbb{Z} \models u + 1 = v \land v_1 + 1 = u - 1 \land v_2 - 1 = v + 1 \land v_3 = v - 1 \Rightarrow \psi$$

and in pure logic:

$$\models v_2 = f(v_3) \land v_1 = f(u) \land \psi \Rightarrow \bot$$

We can assume it only involves variables and equality, by the interpolant property and disjointness of signatures.

Subject to a technical condition about finite models, the pure equality theory admits quantifier elimination.

So we can assume ψ is a propositional combination of equations between variables.

Nelson-Oppen III

In our running example, $u = v_3 \wedge \neg(v_1 = v_2)$ is one suitable interpolant, so

$$\mathbb{Z} \models u + 1 = v \wedge v_1 + 1 = u - 1 \wedge v_2 - 1 = v + 1 \wedge v_3 = v - 1 \Rightarrow u = v_3 \wedge \neg(v_1 = v_2)$$

in Presburger arithmetic, and in pure logic:

$$\models v_2 = f(v_3) \wedge v_1 = f(u) \Rightarrow u = v_3 \wedge \neg(v_1 = v_2) \Rightarrow \bot$$

The component decision procedures can deal with those, and the result is proved.

Nelson-Oppen IV

Could enumerate all significanctly different potential interpolants.

Better: case-split the original problem over all possible equivalence relations between the variables (5 in our example).

$$T_1, \ldots, T_n \models \phi_1 \wedge \cdots \wedge \phi_n \wedge ar(P) \Rightarrow \bot$$

So by interpolation there's a C with

$$T_1 \models \phi_1 \wedge ar(P) \Rightarrow C$$

$$T_2, \ldots, T_n \models \phi_2 \wedge \cdots \wedge \phi_n \wedge ar(P) \Rightarrow \neg C$$

Since $ar(P) \Rightarrow C$ or $ar(P) \Rightarrow \neg C$, we must have one theory with $T_i \models \phi_i \wedge ar(P) \Rightarrow \bot$.

Nelson-Oppen V

Still, there are quite a lot of possible equivalence relations (bell$(5) = 52$), leading to large case-splits.

An alternative formulation is to repeatedly let each theory deduce new disjunctions of equations, and case-split over them.

$$T_i \models \phi_i \Rightarrow x_1 = y_1 \vee \cdots \vee x_n = y_n$$

This allows two imporant optimizations:

- If theories are *convex*, need only consider pure equations, no disjunctions.
- Component procedures can actually produce equational consequences rather than waiting passively for formulas to test.

Shostak's method

Can be seen as an optimization of Nelson-Oppen method for common special cases. Instead of just a decision method each component theory has a

- Canonizer — puts a term in a T-canonical form
- Solver — solves systems of equations

Shostak's original procedure worked well, but the theory was flawed on many levels. In general his procedure was incomplete and potentially nonterminating.

It's only recently that a full understanding has (apparently) been reached.

See ICS (`http://www.icansolve.com`) for one implementation.

Certification of decision procedures

We might want a decision procedure to produce a ‘proof’ or ‘certificate’

- Doubts over the correctness of the core decision method
- Desire to use the proof in other contexts

This arises in at least two real cases:

- Fully expansive (e.g. ‘LCF-style’) theorem proving.
- Proof-carrying code

Certifiable and non-certifiable

The most desirable situation is that a decision procedure should produce a short certificate that can be checked easily.

Factorization and primality is a good example:

- Certificate that a number is not prime: the factors! (Others are also possible.)
- Certificate that a number is prime: Pratt, Pocklington, Pomerance, . . .

This means that primality checking is in NP $\cap$ co-NP (we now know it's in P).

Certifying universal formulas over $\mathbb{C}$

Use the (weak) *Hilbert Nullstellensatz*:

The polynomial equations $p_1(x_1, \ldots, x_n) = 0$, ..., $p_k(x_1, \ldots, x_n) = 0$ in an algebraically closed field have *no* common solution iff there are polynomials $q_1(x_1, \ldots, x_n)$, ..., $q_k(x_1, \ldots, x_n)$ such that the following polynomial identity holds:

$$q_1(x_1, \ldots, x_n) \cdot p_1(x_1, \ldots, x_n) + \cdots + q_k(x_1, \ldots, x_n) \cdot p_k(x_1, \ldots, x_n) = 1$$

All we need to certify the result is the cofactors $q_i(x_1, \ldots, x_n)$, which we can find by an instrumented Gröbner basis algorithm.

The checking process involves just algebraic normalization (maybe still not totally trivial...)

Certifying universal formulas over $\mathbb{R}$

There is a similar but more complicated Nullstellensatz (and Positivstellensatz) over $\mathbb{R}$.

The general form is similar, but it's more complicated because of all the different orderings.

It inherently involves sums of squares (SOS), and the certificates can be found efficiently using semidefinite programming (Parillo . . .)

Example: easy to check

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

via the following SOS certificate:

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

Less favourable cases

Unfortunately not all decision procedures seem to admit a nice separation of proof from checking.

Then if a proof is required, there seems no significantly easier way than generating proofs along each step of the algorithm.

Example: Cohen-Hörmander algorithm implemented in HOL Light by McLaughlin (CADE 2005).

Works well, useful for small problems, but about $1000\times$ slowdown relative to non-proof-producing implementation.

Summary

- There is a need for combinations of decision methods
- For general quantifier prefixes, relatively few useful results.
- Nelson-Oppen and Shostak give useful methods for universal formulas.
- We sometimes also want decision procedures to produce proofs
- Some procedures admit efficient separation of search and checking, others do not.
- Interesting research topic: new ways of compactly certifying decision methods.

Saturday puzzle (first version)

Give a polynomial-time algorithm for the Boolean satisfiability problem (SAT).

Hence deduce that $P = NP$.

Saturday puzzle (second version)

Give an algorithm A that accepts Boolean formulas and returns 'true' or 'false', with the following characteristics:

- Terminates on all inputs
- Correctly tests whether any formula is satisfiable.
- *If* $P = NP$ *then* there is a polynomial $p(n)$ so that the runtime of A on satisfiable formulas of size n is $\leq p(n)$.