22c:295 Seminar in AI — Decision Procedures

Quantifier Elimination

Cesare Tinelli

tinelli@cs.uiowa.edu

The University of Iowa

[a] These notes were developed from a set of lecture notes originally written by Clark Barrett at New York University. These notes are copyrighted material and may not be used in other course settings outside of the University of Iowa in their current form or modified form without the express written permission of the copyright holders.

Outline

- Quantifier Elimination
- Presburger Arithmetic

Sources:
Harrison, John. *Introduction to Logic and Automated Theorem Proving*. Unpublished manuscript. Used by permission.
Enderton, Herbert B. *A Mathematical Introduction to Logic, Second Edition*. Academic Press, 2001.

Quantifi er Elimination

A theory T of signature Σ *admits quantifier elimination* iff for every Σ-formula $\varphi(\vec{x})$ there is a quantifier-free Σ-formula $\varphi'(\vec{x})$ such that

$$T \models \varphi(\vec{x}) \leftrightarrow \varphi'(\vec{x}).$$

Note that if we can

1. compute φ' above and
2. decide the T-satisfiability of qffs,

then we can decide the T-satisfiability of arbitrary formulas.

Notation

$\varphi(\vec{x})$ denotes a formula whose free variables are included in the vector $\vec{x}$.

Quantifi er Elimination

The following theorem reduces the quantifier elimination problem to a particular special case.

Theorem Assume that for every formula $\varphi(\vec{y})$ of the form $\exists x\,(\alpha_0 \wedge \cdots \wedge \alpha_n)$, where each α_i is a literal, there is a quantifier-free formula $\psi(\vec{y})$ such that $T \models (\varphi \leftrightarrow \psi)$. Then T admits quantifier elimination.

Proof By induction on formulas.

(Base) Clearly, every atomic formula is equivalent to a quantifier-free formula with the same free variables: itself.

(Step) Suppose that $\alpha(\vec{y})$ and $\beta(\vec{y})$ are formulas with quantifier-free equivalents $\alpha'(\vec{y})$ and $\beta'(\vec{y})$.

(over)

Quantifi er Elimination

Proof (cont.)

The propositional connective cases are trivial: $T \models \neg\alpha \leftrightarrow \neg\alpha'$, $T \models (\alpha \wedge \beta) \leftrightarrow (\alpha' \wedge \beta')$, etc.

For the quantifier cases, we can rewrite $\forall x.\, \alpha$ as $\neg\exists x.\, \neg\alpha$, so it is sufficient to consider $\exists x.\, \alpha(\vec{y})$.

By induction hypothesis, $\exists x.\, \alpha(\vec{y})$ is T-equivalent to $\exists x.\, \alpha'(\vec{y})$, where $\alpha'(\vec{y})$ is quantifier-free. But now, we can convert $\alpha'(\vec{y})$ to DNF and distribute the existential quantifier over the disjunction to get $(\exists x.\, \gamma_0(\vec{y})) \vee \cdots \vee (\exists x.\, \gamma_n(\vec{y}))$, where each $\gamma_i(\vec{y})$ is a conjunction of literals. But then, by assumption, we can find an equivalent quantifier-free formula for each $\exists x.\, \gamma_i(\vec{y})$, resulting in an equivalent quantifier-free formula for $\exists x.\, \alpha(\vec{y})$. $\square$

Quantifi er Elimination

Example: Dense linear orders without end points

Consider a theory with equality and one predicate symbol $<$, axiomatized as follows:

(Totality)	$\forall x \forall y \, (x \approx y \vee x < y \vee y < x)$
(Transitivity)	$\forall x \forall y \forall z \, (x < y \wedge y < z \rightarrow x < z)$
(Irreflexivity)	$\forall x \, \neg (x < x)$
(Density)	$\forall x \forall y \, (x < y \rightarrow \exists z. \, x < z \wedge z < y)$
(No Left End Point)	$\forall x. \, \exists y \, (y < x)$
(No Right End Point)	$\forall x. \, \exists y \, (x < y)$

Two obvious models of the theory are the reals and the rational numbers with the symbol $<$ interpreted in the usual way.

Quantifi er Elimination

Theorem The theory of dense linear orders admits elimination of quantifiers.

Proof Consider a formula $\varphi = \exists x\,(\beta_0 \wedge \cdots \wedge \beta_l)$, where each β_i is a literal. Note that atomic formulas have only two forms: $u \approx v$ or $u < v$.

First, by the Totality axiom, we can eliminate negative literals from φ as follows . We replace literals of the form $\neg(u < v)$ by $v \approx u \vee v < u$ and literals of the form $\neg(u \approx v)$ by $u < v \vee v < u$.

By converting into DNF and distributing $\exists$ over $\vee$, we obtain a disjunction of formulas ψ of the form $\psi = \exists x\,(\alpha_0 \wedge \cdots \wedge \alpha_p)$, where each α_i is an atomic formula.

If α_i does not contain x, we can move it outside the scope of the quantifier in ψ.

Quantifier Elimination

Proof (continued)

If some α_i is an equation $x \approx x$, we can simply remove it. If it is an equation $x \approx v$ (or $v \approx x$) with v distinct from x, we can replace x by v everywhere in the scope of the quantifier and then eliminate both α_i and the quantifier. This is justified by the logical equivalence: $\exists\, x\, (x \approx v \wedge \varphi[x, v]) \leftrightarrow \varphi[v, v]$.

The remaining literals must have the form $x < x$, $u < x$ or $x < v$ with u, v distinct from x. If $x < x$ occurs, we can replace the whole formula by *false* by the Irreflexivity axiom. Otherwise, we can rewrite the formula as

$$\exists\, x\, (\bigwedge_i u_i < x \wedge \bigwedge_j x < v_j).$$

It is not difficult to show (**exercise**) that the latter formula is equivalent in the theory to $\bigwedge_{i,j} u_i < v_j$. □

Presburger Arithmetic

Let $\mathcal{Z}$ be the integers, let $\Sigma_{\mathcal{Z}}$ be the signature $(0, 1, +, -, <)$, and let $\mathcal{A}_{\mathcal{Z}}$ be the standard model of the integers with domain $\mathcal{Z}$.

We define the theory of *Presburger arithmetic*, $T_{\mathcal{Z}}$, to be the set $\mathit{Th}(\mathcal{A}_{\mathcal{Z}})$ of all $\Sigma_{\mathcal{Z}}$-sentences true in $\mathcal{A}_{\mathcal{Z}}$.

This theory does *not* admit quantifier elimination. For example, the formula

$$\varphi(y) = \exists\, x.\, y \approx x + x$$

(expressing that y is even) has no quantifier-free equivalent.

Nevertheless, as we will show, $T_{\mathcal{Z}}$ is decidable using an algorithm based on quantifier elimination for an extension of $T_{\mathcal{Z}}$ to a larger signature.

Presburger Arithmetic

Notational Conventions

For convenience, we will use

- the predicates $t_1 \leq t_2,\ t_1 > t_2,\ t_1 \geq t_2$ to respectively abbreviate the formulas $t_1 < t_2 + 1,\ t_2 < t_1,\ t_2 < t_1 + 1$.
- each numeral n to abbreviate the sum of n 1's;
- expressions of the form $n \cdot t$ to abbreviate $\underbrace{t + \cdots + t}_{n \text{ times}}$.

Note The new symbols above are just syntactic sugar. They do not really extend the signature $\Sigma_{\mathcal{Z}}$.

Quantifi er elimination for Presburger Arithmetic

Let's now instead extend $\Sigma_{\mathcal{Z}}$ to $\Sigma^*_{\mathcal{Z}}$ by adding an infinite number of unary predicate symbols D_k, one for each integer $k \geq 2$.

Let $T^*_{\mathcal{Z}}$, be the set all $\Sigma^*_{\mathcal{Z}}$-sentences true in the expansion $\mathcal{A}^*_{\mathcal{Z}}$ of $\mathcal{A}_{\mathcal{Z}}$ that interprets $D_k(x)$ as "x is divisible by k".

The theory $T^*_{\mathcal{Z}}$ *does* admit quantifier elimination.

Note The $T^*_{\mathcal{Z}}$-validity of ground $\Sigma^*_{\mathcal{Z}}$-formulas is easily decidable: one simply has to evaluate the formula over the integers.

Since quantifier elimination maps every closed $\Sigma^*_{\mathcal{Z}}$-formula to a $T^*_{\mathcal{Z}}$-equivalent ground $\Sigma^*_{\mathcal{Z}}$-formula, the *whole* $T^*_{\mathcal{Z}}$ is then decidable.

Since $T_{\mathcal{Z}} \subseteq T^*_{\mathcal{Z}}$, it follows that $T_{\mathcal{Z}}$ is decidable as well.

Canonical Forms

Given a total ordering $\prec$ on variables, each $\Sigma^*_{\mathcal{Z}}$-term can be effectively translated into the $T^*_{\mathcal{Z}}$-equivalent form:

$$c_1 \cdot x_1 + \cdots + c_n \cdot x_n + k$$

where each c_i is a numeral or a negated numeral othen than 0, k is a numeral or a negated numeral, and $x_1 \prec x_2 \prec \cdots \prec x_n$.

If σ is the function that maps each term to its normal form above, the following holds:

$$T^*_{\mathcal{Z}} \models t_1 \approx t_2 \text{ iff } \sigma(t_1) = \sigma(t_2)$$

Thus, we call σ a *canonizer*, and the above form a *canonical* form.

Note The function σ is idempotent: $\sigma(\sigma(t)) = \sigma(t)$ for all terms t.

Cooper's QE Algorithm

Presburger's original quantifier elimination algorithm for $T^*_{\mathcal{Z}}$ (1930) follows the classic pattern of dealing with an existential quantifier applied to a conjunction of literals.

In 1972, Cooper developed an improved version based on eliminating the quantifiers from formulas of the form $\exists x\ \varphi$ where φ is an arbitrary quantifier-free formula.

Because Cooper's algorithm avoids the transformation to DNF, it can be substantially more efficient.

This is the algorithm we will present.

Note In the rest this discussion, when we speak about formula equivalence or validity we mean equivalence and validity in $T^*_{\mathcal{Z}}$, which, notice, is the same as equivalence and validity in $\mathcal{A}^*_{\mathcal{Z}}$.

Preprocessing Step

For convenience, before eliminating quantifiers from a given formula, we apply (in order) the following transformations to it.

- All logical connectives other then $\neg$, $\wedge$ and $\vee$ are replaced by their definition in terms of $\neg$, $\wedge$ and $\vee$.
- The defined predicates $\leq, >, \geq$ are replaced by their definitions (e.g., $s \leq t$ is replaced by $s < t + 1$).
- All negated inequalities of the form $\neg(s < t)$ are replaced by the inequality $t < s + 1$.
- All equations and inequations are rewritten to have zero on the left-hand side (i.e., $s \approx t$ is replaced by $0 \approx t - s$, and $s < t$ by $0 < t - s$).
- All arguments of predicates are replaced by their canonical form.

Presburger Arithmetic

Consider a formula $\exists\, x.\, \varphi$ where φ is quantifier-free.

After applying the preprocessing transformations and putting the formula in negated normal form, we can assume that φ is composed of conjunctions and disjunctions of literals of the form:

$$0 \approx t,\ \neg(0 \approx t),\ 0 < t,\ D_k(t),\ \neg D_k(t)$$

where t is a term in canonical form.

Let's call such a φ a *restricted* formula.

Presburger Arithmetic

We have $\exists x.\, \varphi$ where φ is a restricted formula.

The next step is to turn φ into a formula where all the occurrences of x have the same coefficient.

We do this by computing the least common multiple (LCM) l of all of the coefficients of x in φ.

- For equations and disequations, $0 \approx t$ and $\neg(0 \approx t)$, we simply multiply t by the factor which results in x having the coefficient l.
- Similarly, for divisibility predicates $D_k(t)$, we multiply both t and k by the factor that results in x having the coefficient l.
- For inequalities $0 < t$, we multiply by the absolute value of the required factor so that the coefficient of x becomes $\pm l$.

Presburger Arithmetic

We have $\exists x.\, \varphi$ where φ is a restricted formula and the coefficients of x in φ are $\pm l$ for some positive integer l.

We replace it with the equivalent formula $\exists x.\, D_l(x) \wedge \psi$, where ψ is obtained from φ by replacing every occurence of $l \cdot x$ by x.

Note that $\varphi' = D_l(x) \wedge \varphi$ is a restricted formula and all coefficients of x in φ' are ± 1.

Observation Given a valuation of the free variables of φ', either

1. φ' holds for arbitrarily small values of x,
   i.e., $\forall y.\, \exists x.\, x < y \wedge \varphi'$ (with y fresh) is valid, or
2. there is a minimal value for x such that φ' holds,
   i.e., $\exists x.\, \varphi' \wedge \forall y.\, (y < x \rightarrow \neg \varphi'[x \mapsto y])$ (with y fresh) is valid.

We consider these two cases separately and then take the disjunction of the result.

Presburger Arithmetic

Before considering the two separate cases, we need some preliminary results.

For an atomic formula α appearing in φ', we define $\alpha_{-\infty}$ as follows:

α	$\alpha_{-\infty}$
$0 \approx t$ with $1 \cdot x$ in t	*false*
$0 < t$ with $1 \cdot x$ in t	*false*
$0 < t$ with $-1 \cdot x$ in t	*true*
other atomic formula α	α

Now, we define $\varphi'_{-\infty}$ to be the result of replacing each atomic formula α in φ' with $\alpha_{-\infty}$.

Presburger Arithmetic

Lemma For sufficiently small values of x, $\varphi'[x]$ and $\varphi'_{-\infty}[x]$ are equivalent.

Proof By induction on the structure of φ'. **Exercise** $\square$

Lemma $\forall\, y.\, \exists\, x.\, x < y \wedge \varphi'[x]$ is equivalent to $\exists\, x.\, \varphi'_{-\infty}[x]$.

Proof $(\Rightarrow)$ Assume that $\forall\, y.\, \exists\, x.\, x < y \wedge \varphi'[x]$ holds. Then $\varphi'[x]$ holds for arbitrarily small values of x. In particular, $\varphi'[x]$ holds for a sufficiently small value of x that makes $\varphi'[x]$ and $\varphi'_{-\infty}[x]$ equivalent by the previous lemma. If follows that $\exists\, x.\, \varphi'_{-\infty}[x]$ holds.

(over)

Presburger Arithmetic

Proof (cont.) $(\Leftarrow)$ Suppose $\exists\, x.\, \varphi'_{-\infty}[x]$ holds, which means that $\varphi'_{-\infty}[x]$ holds for some value i of x.
We first show that $\varphi'_{-\infty}[x]$ holds for infinitely many values of x smaller than i.
To do that, observe that the truth of predicates of the form $D_k(t)$ in which x occurs positively (i.e., with coefficient 1) is unchanged if we subract from the value of x a multiple of k. Let m be the LCM of all k's such that $D_k(t)$ is an atomic formula in $\varphi'_{-\infty}$.
Then, since these are the *only* atomic formulas involving x in $\varphi'_{-\infty}$, the truth of $\varphi'_{-\infty}$ is unchanged if we subtract from the value of x a multiple of m.
Given that $\varphi'_{-\infty}[x]$ holds for some value i of x, it follows that it also holds for the value $i - n \cdot m$ of x for all $n \geq 0$.

It follows from the previous lemma (**How?**) that $\forall\, y.\, \exists\, x.\, x < y \wedge \varphi'[x]$ also holds. $\square$

Presburger Arithmetic

So far, we have that $\forall\, y.\, \exists\, x.\, x < y \wedge \varphi'[x]$ and $\exists\, x.\, \varphi'_{-\infty}[x]$ are equivalent.

Corollary Let m be the LCM of all k such that $D_k(t)$ is a subformula of $\varphi'_{-\infty}$ containing x. Then $\forall\, y.\, \exists\, x.\, x < y \wedge \varphi'[x]$ is equivalent to $\bigvee_{i=1}^{m} \varphi'_{-\infty}[x \mapsto i]$.

Proof We know $\forall\, y.\, \exists\, x.\, x < y \wedge \varphi'[x]$ is equivalent to $\exists\, x.\, \varphi'_{-\infty}[x]$. Now, since $\varphi'_{-\infty}[x]$ is invariant modulo m, there is an x such that $\varphi'_{-\infty}[x]$ holds iff at least one of the formulas $\varphi'_{-\infty}[x \mapsto i]$ holds where i ranges over any set of m consecutive integers. In particular, we can choose $i \in [1..m]$. $\square$

In conclusion, we have that if $\forall\, y.\, \exists\, x.\, x < y \wedge \varphi'[x]$ holds, our first case, then we can equivalently replace it by $\bigvee_{i=1}^{m} \varphi'_{-\infty}[i]$.

Presburger Arithmetic

Now consider the formula

$$\exists\, x.\, \varphi'[x] \wedge \forall\, y.\, y < x \rightarrow \neg\varphi'[x \mapsto y] \qquad (*)$$

where φ' is a restricted formula and all the coefficients of x in φ' are ± 1, *our second case*, and assume it holds for some valuation of its free variables.

As before, let m be the LCM of all k such that $D_k(t)$ is a subformula of φ' containing x.

Given that $(*)$ holds, there is some minimal integer i such that $\varphi'[x \mapsto i]$ holds, but $\varphi'[x \mapsto i - m]$ does not.

Since the divisibility predicates are invariant under m, the change in the value of φ' when evaluated at $i - m$ must then be caused by one of the other literals of φ'.

Presburger Arithmetic

For each literal $L[x]$ of φ' that contains x and is not a divisibility predicate, we associate a *boundary point* b:

an integer such that $L[x \mapsto b]$ does not hold but $L[x \mapsto b+1]$ does.

literal type	boundary point
$0 \approx x + t$	the value of $-(t+1)$
$\neg(0 \approx x + t)$	the value of $-t$
$0 < x + t$	the value of $-t$
$0 < -x + t$	none

We call the collection of such boundary points the B-*set* for φ'.

Presburger Arithmetic

Theorem Let m be the LCM of all k such that $D_k(t)$ is a subformula of φ' containing x, and let B be the B-set of φ'. For all integers i, if $\varphi'[x \mapsto i]$ holds but $\varphi'[x \mapsto i - m]$ does not, then $i = b + j$, where $b \in B$ and $j \in [1..m]$.

Proof By structural induction on φ'. Let i be such that $\varphi'[x \mapsto i]$ holds but $\varphi'[x \mapsto i - m]$ does not.

Base case

Suppose φ' is (of the form) $0 \approx x + t$. Then i must be equal to the value of $-t$. Since a boundary point for φ' is equal to the value of $-(t + 1)$, it follows that there is a $b \in B$ such that $i = b + 1$.

Suppose φ' is $\neg(0 \approx x + t)$. A boundary point b for φ' equals the value of $-t$. Note that $\varphi'[x \mapsto i - m]$ can only be false when $i = b + m$.

(over)

Presburger Arithmetic

Base case (continued)

Suppose φ' is $0 < x + t$. Then a boundary point b for φ' equals the value of $-t$. By assumption we have that both $-t + 1 \leq i$ and $i \leq -t + m$ hold. It follows that $i = b + j$ for some $j \in [1..m]$.

This concludes the base case, because if φ' is a literal of other forms, it is not possible that $\varphi'[x \mapsto i]$ holds while $\varphi'[x \mapsto i - m]$ does not.

Inductive Step

Exercise

□

Presburger Arithmetic

We can now state the main quantifier elimination result, obtained by combining the results of the two cases.

Theorem Suppose φ' is a restricted formula and all coefficients of x in φ' are ± 1. Let m be the LCM of all k such that $D_k(t)$ is a subformula of φ' containing x, and that B is the B-set of φ'. Then,

$$(\exists\, x.\, \varphi'[x]) \leftrightarrow \bigvee_{j=1}^{m} \left(\varphi'_{-\infty}[j] \vee \bigvee_{b \in B} (\varphi'[x \mapsto b + j]) \right).$$

Proof By the previous results. (*How?*) □