Master Sciences Informatiques

Solvers Principles and Architectures (SPA)

Final Exam, Fall 2019

Khalil Ghorbal

1 SAT/SMT Solvers

A. Combining Theories

A typical SMT solver has several back-ends (one for each supported theory) as well as an underlying SAT solver to handle the Boolean structure of the problem to solve. The solver relies crucially on combining theories, that is the ability to propagate learned results from one back-end to another to either prove falsity or find a model.

1. Cite one generic method that allows combining several theories. What restriction does it have?
2. Cite other generic means to combine theories that overcome such restriction? (hint: remember the seminar!)

Consider the formula: $\varphi := x^2 + y^2 \leq 1 \wedge 2x + 2y \geq 3$, and suppose you only have the polyheral abstract domain (the theory of linear inequalities) at your disposal.

3. How would you proceed to prove that φ is UNSAT over the reals? (hint: think of an appropriate logical cut.)

3′. What open goal you will still have to discharge?

4. Is it always possible to find such cuts? Can you slightly edit φ to make it a counter-example?

B. Bonus

The satisfiability of the formula φ, and in general any logical combination of polynomial equations and inequalities, is decidable. Indeed, φ is SAT if and only if the sentence $\exists x.\exists y. \quad \varphi$ is true.

5. Perform the Cylindrical Algebraic Decomposition (CAD) of φ and prove it is UNSAT.
6. What are the advantages of combining linear theories (using eventually over-approximations for non-linear expressions) compared to CAD? Explain.

Answer.

1. The Nelson-Oppen method [Nelson and Oppen, 1979] combines decision procedures of individual theories to construct a decision procedure for a combination of theories. The most fundamental restriction being that the combined theories have to be essentially disjoint sharing information solely with the equality symbol, the only allowed common symbol.
2. During the seminar, we saw how the abstract interpretation framework is used to share a richer set of information beyond equality via using several different (numerical) abstract domains.
3. $x^2 + y^2 \leq 1$ implies $x \leq 1 \wedge y \leq 1$ implies $x + y < \frac{3}{2}$.

3′. Showing that the disk is included in $x \leq 1 \wedge y \leq 1$ is beyond the capabilities of the polyhedral abstract domain as it requires handling a non-linear inequality, namely $x^2 + y^2 \leq 1$.

4. One can move the slope arbitrarily close to the circle making the previous cut non-sufficient to separate the disk for the half-space. One can also construct examples beyond the expressiveness

of the underlying domains requiring to either increase the expressiveness of the domain or to lose termination (when a cut exists at infinity). For instance consider

$$\varphi' := x^2 + y^2 \leq 1 \wedge x^2 + y > \frac{5}{4} \ .$$

A line that separates the hyperbola from the circle cannot be below the tangent to the circle at that point. But then since the parabola touches the circle at the same point, the line intersects with the hyperbola preventing separation (see Figure 1).

5. For the disk, the CAD is as follows:

$$(x = -1 \wedge y = 0) \vee \left(-1 < x < 1 \wedge -\sqrt{1 - x^2} \leq y \leq \sqrt{1 - x^2}\right) \vee (x = 1 \wedge y = 0)$$

For the half-space $x + y \geq \frac{3}{2}$, there is only one cell, namely

$$y \geq \frac{1}{2}(3 - 2x)$$

The intersection of this one cell with the decomposition of the disk leads to False: no intersection is possible.

6. The number of cells grows exponentially in CAD. There is therefore a trade-off between decidability and efficiency. Over-approximations are much more efficient in practice, but any procedure using over-approximation is necessarily non-complete.

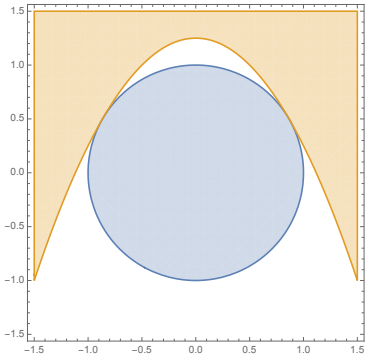

Figure 1: Plot of the formula φ'.

2 Convex Optimization

Dantzig's Lexicographic Rule

The simplex algorithm iterates over the bases while decreasing the reduced cost.

1. Explain why the termination is ensured when the decrease is strict at each iteration.

The decrease can however be non-strict leading to degenerate cases for which termination is no longer guaranteed. For instance, consider the following pyramid:

$$P := \{(x, y, z) \in \mathbb{R}^3 \mid y + z \leq 1 \wedge z - y \leq 1 \wedge x + z \leq 1 \wedge z - x \leq 1 \wedge z \geq 0\},$$

and the following optimization problem (see Figure 2)

$$\begin{aligned} \min \quad & x - y - 3z \\ \text{s.t.} \quad & (x, y, z) \in P \end{aligned}$$

The point $s := (0, 0, 1)$ is a (geometric) vertex that corresponds to the top of the pyramid.

2. Rewrite the problem to apply the simplex algorithm.
3. Perform one step of the algorithm assuming the initial point is s (you can use the tableau presentation). What decision the algorithm could make, why?

In general, consider the linear optimization problem ($A \in \mathbb{R}^{m\times n}$):

$$\begin{aligned} \min \quad & c \cdot x \\ \text{s.t.} \quad & Ax \geq b \end{aligned} \qquad (\mathcal{P})$$

and its perturbed version:

$$\begin{aligned} \min \quad & c \cdot x \\ \text{s.t.} \quad & Ax \geq \tilde{b} \qquad \text{where } \tilde{b}_i = b_i - \epsilon^i \end{aligned}$$

for some positive $\epsilon \ll 1$.

4. Does the perturbed problem have degenerate bases? Explain why.

4′. Give the geometric intuition for the pyramid P.

Dantzig's idea is to encode a real number r as a polynomial in ϵ of degree at most m:

$$r + r_1\epsilon + r_2\epsilon^2 + \cdots + r_m\epsilon^m,$$

represented as a row: $(r, r_1, \ldots, r_m)$. The usual order over the reals is replaced by the following *lexicographic* order [1]

$$(r, r_1, \ldots, r_m) \geq_{\texttt{lex}} (s, s_1, \ldots, s_m) \iff r + r_1\epsilon + \cdots + r_m\epsilon^m \geq s + s_1\epsilon + \cdots + s_m\epsilon^m, \quad \forall \epsilon.\, 0 < \epsilon \ll 1$$

Now the feasible set $\{x \in \mathbb{R}^n \mid Ax \geq \tilde{b}\}$ can be encoded as:

$$\left\{ y \in \mathbb{R}^{n\times(1+m)} \mid Ay \geq_{\texttt{lex}} \begin{pmatrix} b & -I_m \end{pmatrix} \right\},$$

where $\geq_{\texttt{lex}}$ over $\mathbb{R}^{m\times(1+m)}$ is interpreted row by row.

Consider the following optimization problem ($y \in \mathbb{R}^{n\times(1+m)}$):

$$\begin{aligned} \min \quad & \begin{pmatrix} c & 0 \end{pmatrix} \cdot y \\ \text{s.t.} \quad & Ay \geq_{\texttt{lex}} \begin{pmatrix} b & -I_m \end{pmatrix} \end{aligned} \qquad (\tilde{\mathcal{P}})$$

5. Define the bases (algebraic vertices) for $(\tilde{\mathcal{P}})$. We will call them `lex`-bases.

6. Prove that `lex`-bases of $(\tilde{\mathcal{P}})$ form a subset of the bases of $(\mathcal{P})$. (Observe that the matrix A is left unchanged in $(\tilde{\mathcal{P}})$.)

Answer.

1. The number of (geometric) vertices defining the feasible set is finite. If the decrease is strict at each vertex ($\rho > 0$) the termination is guaranteed: either the minimum is finite and therefore will be necessarily reached or the problem is unbounded (which will be detected via the unboundedness criteria).

2. The feasible set can be rewritten as $As = b$ for some matrix A, vector b and $s \in \mathbb{R}^9_+$. First, we encode the original variables x, y and z as follows: $x := s_1 - s_2$, $y := s_3 - s_4$, and $z = s_5$ to get

$$x \in P \iff \begin{pmatrix} 0 & 0 & 1 & -1 & 1 \\ 0 & 0 & -1 & 1 & 1 \\ 1 & -1 & 0 & 0 & 1 \\ -1 & 1 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} s_1 \\ s_2 \\ s_3 \\ s_4 \\ s_5 \end{pmatrix} \leq \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} = b \wedge s_i \geq 0$$

   We then add 4 extra (slack) variables $(s_6, \ldots, s_9)$ to saturate the inequalities. The matrix A is the concatenation of the above matrix and the identity matrix of dimension 4.

3. Suppose the algorithm starts with $\mathfrak{B} = \{5, 7, 8, 9\}$ and $\mathfrak{N} = \{1, 2, 3, 4, 6\}$. This is one possible choice that corresponds to the top of the pyramid with $A_{\mathfrak{B}}$ of rank $4 = m$. The reduced cost is $(1, -1, 2, -2, 3)^t$. Two coordinates are negative. If one moves with respect to $(0, 1, 0, 0, 0)^t$, the fourth component of $\delta_{\mathfrak{B}}$ is negative. If the algorithm doesn't check for the actual value of the decrease ρ (which is zero here), it will update the base to $\{2, 5, 7, 8\}$ (switching 2 and 9).

[1] Check quickly that it is actually an *order*!

Likewise, if one moves with respect to $(0,0,0,1,0)^t$, the decrease is not strict and the base will be updated by removing one of the indices 7, 8 or 9 from $\mathfrak{B}$ and adding the index 4 instead.

4. Degenerated cases occur when the rank of $A_{\mathfrak{B}}$ is (strictly) less than m the number of inequalities defining the feasible set. When this happens for the perturbed problem, it means that there exist some λ_i, not all zero, such that
$$\sum_i \lambda_i(b_i - \epsilon^i) = 0$$
but then if this equation holds for infinitely many small positive ϵ, then $\lambda_i = 0$ for all i (a univariate polynomial over the reals has finitely many roots) contradicting the assumption: there must exists a positive small epsilon for which there are no degenerate cases.

- Consider the following sets of equations corresponding to the bases, $\{1,5,9\}$ and $\{2,5,9\}$ respectively:
$$\begin{pmatrix} 0 & 1 & 0 \\ 0 & 1 & 0 \\ 1 & 1 & 0 \\ -1 & 1 & 1 \end{pmatrix} \begin{pmatrix} s_1 \\ s_5 \\ s_9 \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} \qquad \begin{pmatrix} 0 & 1 & 0 \\ 0 & 1 & 0 \\ -1 & 1 & 0 \\ 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} s_2 \\ s_5 \\ s_9 \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}$$
In both cases the rank of the matrices is $3 < 4 (= m)$ (an obvious redundant rows is highlighted in red), and the solutions for the base elements are $s_1 = s_2 = s_9 = 0$ and $s_5 = 1$. Both correspond to $(x,y,z) = (0,0,1)$, the top of the pyramid. The geometric intuition is that the top of the pyramid is 'cut' in the perturbed problem: each base corresponds now to a different vertex and $A_{\mathfrak{B}}$ has to be of rank m.

5. A base is, as we have seen, a set of indices such that the sub-matrix $A_{\mathfrak{B}}$ of $\begin{pmatrix} A & I_m \end{pmatrix}$ has rank m (the number of rows of A). The concatenation of I_m to A is meant to saturate the inequalities $\geq_{\texttt{lex}}$ by adding m rows to the matrix y. We will denote the resulted matrix as y'.

6. Suppose that $\mathfrak{B}$ is a base for $(\tilde{\mathcal{P}})$ then
$$y_{\mathfrak{B}} = A_{\mathfrak{B}}^{-1} \begin{pmatrix} b & -I_m \end{pmatrix} = \begin{pmatrix} A_{\mathfrak{B}}^{-1} b & -A_{\mathfrak{B}}^{-1} \end{pmatrix}$$
In particular, the first column of $y_{\mathfrak{B}}$ is equal to $A_{\mathfrak{B}}^{-1} b$ and defines a geometric vertex (if $A_{\mathfrak{B}}^{-1} b$ is an element of the non-negative orthant).

Essentially, Dantzig's idea to avoid cycling is to 'tag' (or label) all bases (using the perturbation) that have the same geometric vertex. By enumerating the lex-bases, we cover all the bases of the original problem without cycling.

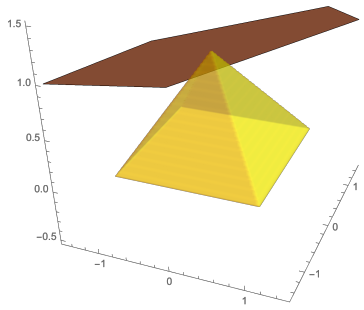

Figure 2: Pyramid P together with the hyperplane defining the objective function at the extreme point.

References

[Nelson and Oppen, 1979] Nelson, G. and Oppen, D. C. (1979). Simplification by cooperating decision procedures. *ACM Trans. Program. Lang. Syst.*, 1(2):245–257.