

**Critical Reading of “Optimization Methods
for Logical Inference” [1]**

Undergraduate Research Internship
Department of Management Sciences
Fall 2007

Supervisor: Dr. Miguel Anjos
UNIVERSITY OF WATERLOO

Rajesh Kumar Swaminathan

February 4, 2008

Table of Contents

| | | |
|----------|--|-----------|
| 1 | Introduction to SAT | 1 |
| 2 | Special Cases in Propositional Logic | 1 |
| 2.1 | Basic Concepts | 1 |
| 2.1.1 | Unit Resolution | 2 |
| 2.2 | Integer Linear Programming Models | 3 |
| 2.2.1 | Optimization and Inference | 4 |
| 2.2.2 | The Linear Programming Relaxation | 6 |
| 2.3 | Horn Polytopes | 7 |
| 2.3.1 | Horn Resolution | 8 |
| 2.3.2 | The Integer Least Element of a Horn Polytope | 9 |
| 2.3.3 | Dual Integrality of Horn Polytopes | 11 |
| 3 | Summary of Findings | 12 |
| 4 | List of Theorems | 12 |
| | Bibliography | 13 |

1 Introduction to SAT

The boolean satisfiability (SAT) problem consists of answering the following questions:

1. Is a given proposition (or formula) S satisfiable? That is, is it possible to find truth assignments for the variables (or atomic propositions) in S such that S evaluates to *true*? If the answer to this question is no, we call S *unsatisfiable*.
2. Is a given proposition S a tautology? That is, is every possible combination of 2^n truth assignments a model for S ? (A model is any satisfying truth assignment.) This is the same as asking if there does not exist a truth assignment for which S evaluates to *false*. The simplest tautology would be $(x_1 \vee \bar{x}_1)$.
3. Does a proposition S_1 *imply* (\rightarrow) another proposition S_2 defined on the same set of variables? That is, is every model of S_1 a model of S_2 without explicitly solving for all the models of S_1 (if this is even possible)?

One way to answer all three questions is to enumerate all 2^n possible truth assignments, where n is the number of atomic propositions, and to run them through the formula. This number of possible combinations unfortunately grows far too quickly (exponentially) with the number of atomic propositions and becomes unmanageable for even as few as 30 atomic propositions (1 billion possible truth combinations). The problems we are likely targeting have from 1000 to 10,000 atomic propositions in them.

The implication question posed in q. 3 can be rewritten as a satisfiability problem similar to the one posed in q. 1 since asking if $S_1 \rightarrow S_2$ is the same as asking if $(\bar{S}_1 \vee S_2)$ is a tautology. Conversely, every proposition S in a satisfiability problem similar to the one posed in q. 1 can be written as an implication problem like the one posed in q. 3 by simply solving $(T \rightarrow S)$ since this reduces to solving $(F \vee S)$ which is the same as (S) . The disjunction of a proposition with another proposition that is always unsatisfiable is pointless and can be simplified by dropping the unsatisfiable proposition. Thus questions 1 and 3 are equivalent and are therefore equally “hard”.

Question 2 is a little bit “harder” in an algorithmic sense since one would have to explore a significantly larger feasible set that contains every single possible model which can be as large as 2^n , n being the number of atomic propositions, if the proposition is indeed a tautology.

There is a simple equivalence between all three problems. The following statements are equivalent

1. S_1 implies S_2 .
2. $(\bar{S}_1 \vee S_2)$ is a tautology.
3. $(S_1 \wedge \bar{S}_2)$ is unsatisfiable.

since every model of S_1 is also a model of S_2 . We showed above that statements 1 and 3 are the same. But questions 2 and 3 are simply negations of each other and are therefore the same. Thus all three statements above are equivalent and a simple equivalence between all three problems has been demonstrated.

2 Special Cases in Propositional Logic

2.1 Basic Concepts

Conjunction in logic is a compound proposition that is true if and only if all of its component propositions are true. Conjunctions are represented by the symbol “ \wedge ” and correspond to a logical AND. Disjunction in logic is a compound proposition that is true if and only if at least one of its

component propositions is true. Disjunctions are represented by the symbol “ \vee ” and correspond to a logical OR. Negations are represented by a “bar” on top of the formula or literal such as \bar{S} or \bar{x} and correspond to a logical NOT.

We therefore have the following simple rules:

1. $(S_1 \wedge S_2)$ is T if and only if both S_1 and S_2 are T.
2. $(S_1 \vee S_2)$ is F if and only if both S_1 and S_2 are F.
3. S is T if and only if \bar{S} is F.

Basic properties and laws of propositions may be applied without affecting the proposition’s models ($\neg S$ is alternate notation for \bar{S}):

1. Involutionary property of negation: $\neg\neg S = S$
2. De Morgan’s Laws:

$$\neg(S_1 \vee S_2) \Leftrightarrow (\bar{S}_1 \wedge \bar{S}_2)$$

$$\neg(S_1 \wedge S_2) \Leftrightarrow (\bar{S}_1 \vee \bar{S}_2)$$

3. Distributive Law: $S_1 \vee (S_2 \wedge S_3) \Leftrightarrow (S_1 \vee S_2) \wedge (S_1 \vee S_3)$

A satisfiability problem posed as a propositional formula is said to be in *conjunctive normal form* (CNF) if it is a conjunction of one or more *clauses*, each of which is a disjunction of one or more (possibly negated) literals.

Theorem Any formula in propositional logic is equivalent to a CNF formula whose length is linearly related to the length of the original formula.

That there exist rewriting techniques leading to CNF representations that are polynomially bounded in length was first noted by Tseitin as early as 1968.

Consider the worst-case scenario which is rewriting a formula in *disjunctive normal form* (DNF) as CNF. A DNF formula is a disjunction of *terms*, each of which is a conjunction of literals. An example of a DNF formula is:

$$(x_1 \wedge x_2) \vee (x_3 \wedge x_4) \vee (x_5 \wedge x_6) \vee (x_7 \wedge x_8)$$

We introduce new propositions $x_{12}, x_{34}, x_{56}, x_{78}$ that represent each conjunction in parenthesis. For each $x_{2j-1,2j}$ ($j = 1, 2, 3, 4$) we write the clauses $(\bar{x}_{2j-1,2j} \vee x_{2j-1})$ and $(\bar{x}_{2j-1,2j} \vee x_{2j})$

We also need one additional clause to knit the four subformulas together.

$$(x_{12} \vee x_{34} \vee x_{56} \vee x_{78})$$

These clauses put together represent the DNF formula above. It is easy to see that this technique results in a CNF formula with three times the original number of atomic propositions, and twice the original number of clauses plus one additional clause to knit the subformulas together. Both the number of atomic propositions and the number of clauses are therefore linear in growth. \square

2.1.1 Unit Resolution

Consider the following example:

$$(x_1) \wedge (x_2 \vee \bar{x}_3 \vee x_4) \wedge (\bar{x}_1 \vee \bar{x}_4) \wedge (\bar{x}_2 \vee x_3 \vee x_5)$$

It is clear that if this formula is to hold true, one would have to set x_1 to 1 since it stands alone as a unit positive clause. We therefore set x_1 to 1 and remove it from the formula. In addition, we also remove all other clauses where x_1 is present as a positive literal since these clauses will be

automatically satisfied. Furthermore, we get rid of all occurrences of \bar{x}_1 in any of the remaining clauses since this literal will always be false and a disjunction with something that is always false is pointless.

We are now left with

$$(x_2 \vee \bar{x}_3 \vee x_4) \wedge (\bar{x}_4) \wedge (\bar{x}_2 \vee x_3 \vee x_5)$$

We now repeat the process with \bar{x}_4 to obtain the reduced form

$$(x_2 \vee \bar{x}_3) \wedge (\bar{x}_2 \vee x_3 \vee x_5)$$

We now stop as we cannot apply the same process once more due to a lack of unit clauses.

This procedure is referred to as *unit resolution*.

1. If the unit resolution procedure applied to S returns an empty formula then unit resolution has succeeded in finding a satisfying truth assignment for S .
2. On the other hand, if the procedure returns $S = \{\}$, that is the empty clause, then S is unsatisfiable since the only way to obtain an empty clause is to have something of the form $(x_i) \wedge (\bar{x}_i)$ to begin with.

2.2 Integer Linear Programming Models

Integer linear programming deals with linear programming problems where one or more variables are restricted to taking integer values. By introducing suitable bounds on the integer variables (which are expressed as linear inequalities), it is possible to restrict variables to only take values in the nonnegative integers or even just values of 0 or 1. It is the latter “boolean” restriction that captures the semantics of propositional logic since the values of 0 and 1 may be naturally associated with *False* and *True*.

The idea here is to formulate satisfiability of CNF formulas as integer linear programming problems with clauses represented by constraints and atomic propositions represented by 0-1 binary variables.

In general, integer linear programming (ILP) problems may be solved by solving the linear relaxation and then applying a branch-and-bound procedure. But since all of our variables are not only integer but also binary, we may alternatively apply special-case algorithms such as Balas’ additive algorithm to solve the binary ILP problem.

Consider, for example, the single clause $x_2 \vee \bar{x}_3 \vee x_4$. This clause may be represented by the inequality $x_2 + (1 - x_3) + x_4 \geq 1$ with x_2, x_3 and x_4 restricted to boolean values of 0 and 1. Thus the SAT problem

$$(x_1) \wedge (x_2 \vee \bar{x}_3 \vee x_4) \wedge (\bar{x}_1 \vee \bar{x}_4) \wedge (\bar{x}_2 \vee x_3 \vee x_5) \tag{1}$$

may be represented by the following set of constraints, one for each clause:

$$\begin{aligned} x_1 &\geq 1 \\ x_2 + (1 - x_3) + x_4 &\geq 1 \\ (1 - x_1) + (1 - x_4) &\geq 1 \\ (1 - x_2) + x_3 + x_5 &\geq 1 \\ x_1, \dots, x_5 &= 0 \text{ or } 1 \end{aligned}$$

Moving all constants to the right hand side, we get the following *clausal* form:

$$\begin{aligned}
 x_1 &\geq 1 \\
 x_2 - x_3 + x_4 &\geq 0 \\
 -x_1 - x_4 &\geq -1 \\
 -x_2 + x_3 + x_5 &\geq 0 \\
 x_1, \dots, x_5 &= 0 \text{ or } 1
 \end{aligned} \tag{2}$$

In general, satisfiability in propositional logic is equivalent to the solvability of the linear system $Ax \geq b, x \in \{0, 1\}^n$ where the inequalities $A_i x \geq b_i$ are clausal.

A few quick notes and observations:

1. A is an $m \times n$ matrix of 0s and ± 1 s, where m is the number of clauses and n is the total number of atomic propositions in the formula. A_{ij} is $+1$ if x_j is positive in clause i , -1 if negated, and 0 otherwise.
2. $b_i = (1 - \# \text{ of } -1\text{s in row } i \text{ of matrix } A)$. b can therefore be calculated given a matrix A of 0s and ± 1 s.
3. The geometric interpretation of a SAT problem reduces to looking for an extreme point of the unit hypercube in \mathfrak{R}^n that is contained in all the half-spaces defined by the clausal inequalities.

We may verify the involutory property of the negation operator “ \neg ” and De Morgan’s laws stated previously by looking at their corresponding linear integer constraint representations:

1. $(\neg \neg x_1) \Leftrightarrow 1 - (1 - x_1) \geq 1 \Leftrightarrow x_1 \geq 1 \Leftrightarrow (x_1)$
2. $\neg(x_1 \vee x_2) \Leftrightarrow 1 - (x_1 + x_2) \geq 1 \Leftrightarrow x_1 + x_2 \leq 0$ which is equivalent to the set of constraints

$$\begin{aligned}
 x_1 &\leq 0 \\
 x_2 &\leq 0
 \end{aligned}$$

since if any one of x_1 or x_2 (say x_1) were allowed to be strictly positive, it would take the value of 1 (recall the additional constraint $x_i \in \{0, 1\}$) and so x_2 would need to be at most -1 to satisfy $x_1 + x_2 \leq 0$ which is of course not possible. We reach a contradiction and hence conclude that both x_1 and x_2 need to be at most 0 if their sum is to be at most 0.

One may now proceed to verify the other half of De Morgan’s laws and the distributive property in a similar fashion.

2.2.1 Optimization and Inference

We have already seen that the intersection of clausal half-spaces defines a convex polyhedron. If the additional box constraints $0 \leq x_j \leq 1$ are added, we obtain a *bounded polyhedron* also known as a *polytope*. Satisfiability now essentially implies finding a feasible integer point inside this polytope.

We may check to see if an integer linear programming problem is feasible or not by using a 2-phase method, that is, by adding an *artificial variable* and then trying to optimize the artificial variable to 0. Thus the satisfiability of (1) may be tested by checking the feasibility of (2) which is in turn done by solving the following optimization problem:

$$\begin{array}{ll}
\text{MIN} & x_0 \\
\text{s.t.} & x_0 + x_1 \geq 1 \\
& x_0 + x_2 - x_3 + x_4 \geq 0 \\
& x_0 - x_1 - x_4 \geq -1 \\
& x_0 - x_2 + x_3 + x_5 \geq 0
\end{array}$$

$$x_j \in \{0, 1\}, j = 0, 1 \dots 5$$

The above optimization problem is always feasible with $x_0 = 1$ and all other $x_j = 0$ and an optimal value of $x_0 = 0$. Thus, the original formula is satisfiable if and only if the optimization problem above is solved with x_0 at 0. When $x_0 = 0$, the x_0 's in the constraints drop out and we are left with a feasible solution to the original set of constraints put down in (2).

The above optimization problem is called a phase 1 construction which takes the general form

$$\begin{array}{ll}
\text{MIN} & x_0 \\
\text{s.t.} & x_0 e + Ax \geq b \\
& x_j \in \{0, 1\}, j = 0, 1 \dots n
\end{array}$$

where $Ax \geq b$ represents the original clausal inequalities and e is a column of ones.

We therefore now have a way of checking to see if a given proposition is satisfiable or not. How about the inference problems in the form of implications? That is, does a formula S_1 imply another formula S_2 ($S_1 \rightarrow S_2$)? We have already seen that the problem ($S_1 \rightarrow S_2$) is identical to asking if $(\bar{S}_1 \vee S_2)$ is a tautology.

To begin with, assume S_1 is CNF and S_2 is given by a single clause C . The optimization model is given by:

$$\begin{array}{ll}
\text{MIN} & cx \\
\text{s.t.} & Ax \geq b \\
& x \in \{0, 1\}^n
\end{array}$$

where c is the incidence vector of clause C and $Ax \geq b$ are the clausal inequalities representing S_1 . The incidence vector c is constructed by assigning a value of +1 to c_i if the literal x_i is positive in C , -1 if negated, and 0 otherwise.

If this optimization yields a minimum value of $1 - n(C)$ (1 minus the number of negative literals in C) or larger, S_1 implies S_2 . Otherwise the implication does not hold.

Why is this the case?

Suppose that there exists a model that satisfies both S_1 and S_2 , then we are guaranteed a feasible solution for the set of constraints

$$\begin{array}{ll}
Ax & \geq b \\
cx & \geq 1 - n(C) \\
x & \in \{0, 1\}^n
\end{array}$$

Now we take the left-hand-side of the second constraint and move it to the objective function and then try to minimize it. Since $Ax \geq b$ still remains as a constraint (note that we do not

associate any artificial variables with this constraint), we are ensured that we are only working with models of S_1 (since $Ax \geq b$ are the clausal inequalities representing S_1).

The question is whether all these models of S_1 are also models of S_2 whose inequality is represented by $cx \geq 1 - n(C)$. If all models of S_1 are also models of S_2 , cx would always remain greater than $1 - n(C)$ (meaning that $cx \geq 1 - n(C)$ always remains feasible when $Ax \geq b$ is) for if there was even one model of S_1 that was not a model of S_2 then the minimization would pick this up to yield an objective value cx that is strictly less than $1 - n(C)$ which causes the inequality $cx \geq 1 - n(C)$ representing S_2 to be violated.

Thus S_1 implies S_2 if and only if the optimal minimum value is greater than or equal to $1 - n(C)$.

Now, what if S_2 is given by more than just a clause? The idea is here is to show that $(S_1 \wedge \bar{S}_2)$ is unsatisfiable as discussed in Section 1.

Consider the situation

(i) S_1 , CNF, with clausal inequalities $Ax \geq b$

(ii) \bar{S}_2 , CNF, with clausal inequalities $Bx \geq d$

To test if S_2 is logically implied by S_1 we solve

$$\begin{array}{ll} \text{MIN} & x_0 \\ \text{s.t.} & x_0e + Bx \geq d \quad (\text{corresponds to } \bar{S}_2) \\ & Ax \geq b \quad (\text{corresponds to } S_1) \\ & x \in \{0, 1\}^n \end{array}$$

If this problem is optimized at 0, then both S_1 and \bar{S}_2 are true which means there exists a model of S_1 that is not a model of S_2 indicating that S_1 *does not* imply S_2 . Thus S_1 implies S_2 if and only if the minimum value of x_0 is 1, for if x_0 were allowed to be 0, the minimization would pick it up. The fact that the minimization does not pick it up indicates that there is no model of S_1 that is not also a model of S_2 .

Once again, we do not associate an artificial x_0 with the constraint $Ax \geq b$ since we assume that S_1 is satisfiable. If S_1 had no models to begin with, i.e. if it were unsatisfiable, the implication $S_1 \rightarrow S_2$ always holds.

To conclude, we can always easily rewrite inference problems in propositional logic as integer linear programming (ILP) problems. But in general, ILP problems are just as hard (if not harder in particular instances) to solve as the satisfiability problem itself. It seems then that we have taken a hard problem—the one of satisfiability—and made it even harder by converting it into an ILP problem. However in practice, special mathematical structure makes many ILP problems easy to solve and this happens to be true of many inference problems in propositional logic. We now proceed to investigate special mathematical structure within an ILP representation of SAT by looking at its corresponding linear programming (LP) relaxation.

2.2.2 The Linear Programming Relaxation

The linear programming relaxation consists of taking the ILP problem formulated above and relaxing the integer (binary) constraints $x_j \in \{0, 1\}$ to the weaker condition $0 \leq x_j \leq 1$.

The idea is to analyze the properties of this linear programming relaxation to obtain ideas on computational strategies for solving the integer model. It turns out that the linear programming relaxation retains sufficient structure to be a useful representation of the original inference problem. The relaxation provides a means of understanding special structures in propositions that permit efficient solvability of inference.

It is interesting to observe that the actions performed by the unit resolution procedure discussed in Section 2.1.1 earlier are implicitly encoded into the linear inequalities of the linear programming relaxation. For example, a unit clause (x_j) is given by $x_j \geq 1$. However, there is an explicit upper bound $x_j \leq 1$ on x_j . Thus x_j is fixed to have a value of exactly 1 (as fixed by unit resolution). Similarly, a unit clause with a negated literal (\bar{x}_k) is given by the linear inequality $-x_k \geq 0$, but there exists a lower bound on x_k given by $x_k \geq 0$ which fixes the variable x_k to have a value 0. Also if at any stage of unit resolution two conflicting unit clauses (x_j) and (\bar{x}_j) are obtained, the corresponding implied inequalities are $x_j \geq 1$ and $-x_j \geq 0$, which are in conflict. The linear program is inconsistent and therefore infeasible in such a situation.

Lemma The linear programming relaxation of a unit clause free CNF formula is always feasible with the trivial fractional solution $x_j = \frac{1}{2}$ for all j .

It is easy to convince ourselves that the lemma is indeed true. If the clause contained (two or more) positive-only literals, the left hand side of the corresponding constraint would always add up to 1 or more. For every additional negated literal that is added, the left hand side decreases by a $\frac{1}{2}$, but the right hand side decreases by 1, a larger quantity, thereby preserving the inequality. If the clause contained one positive and one negated literal, the left hand side would sum up to 0 and so would the right hand side.

The lemma implies the following theorem:

Theorem A proposition S has a unit refutation (unsatisfiability of S proved by unit resolution) if and only if the linear programming relaxation of S is infeasible.

For satisfiable propositions, the power of the linear programming relaxation exceeds that of unit resolution. This is because for satisfiable propositions we may get lucky while solving the linear programming relaxation and hit on an integer solution and thus prove satisfiability. The statement implies that there might exist a special class of satisfiable propositions for which unit resolution fails to come up with a solution (perhaps because we're left with a unit clause free formula at some point), but a feasible solution for the linear programming relaxation happens to be integer.

One such class of propositions are called *balanced propositions*. These are propositions whose relaxations are integral polytopes. Using a simplex method to solve the relaxation will guarantee us a vertex feasible solution that is integer and hence prove satisfiability. However, if the proposition itself has no unit clauses, unit resolution would not achieve anything.

For refutable propositions on the other hand, the power of unit resolution and the relaxation are identical. That is, if unit resolution can refute a proposition, then so can solving the linear programming relaxation and vice versa. Inversely, if unit resolution fails to refute a refutable proposition, then the linear programming relaxation also has no hope of refuting it and vice versa.

2.3 Horn Polytopes

We now proceed to analyze special classes of propositions by looking at properties of their linear relaxations and attempting to derive insight into solving their corresponding ILP representations efficiently. One such special class of propositions are Horn propositions.

A *Horn* rule has either no atoms or a single atom in its consequent. A Horn clause must therefore contain *at most* one positive literal since consequents of a rule correspond to positive literals inside a clause. A Horn clause system is a system in which all clauses are Horn.

2.3.1 Horn Resolution

Horn systems are highly structured propositions where satisfiability can be solved in linear time (in the number of literals) using a restricted form of unit resolution. The restricted form of unit resolution resolves only unit *positive* clauses since a Horn clause system with no unit positive clauses is trivially satisfiable by assigning 0 to the rest of the literals. This is because each clause must then have at least one negated literal, for none of the clauses may contain a single positive literal (it would then be unit positive) or two or more positive literals (it would then be non-Horn). Thus assigning 0 to all literals will ensure all clauses are satisfied simultaneously.

The restricted form of unit resolution can therefore completely solve the satisfiability problem on a Horn system:

1. Look for only unit positive clauses and apply unit resolution on them.
2. If no unit positive clauses are found, assign 0 to all literals in the formula, declare the proposition as satisfiable and stop.
3. If an empty clause is obtained, the proposition is unsatisfiable. Stop.
4. If an empty formula is obtained, unit resolution has succeeded in finding a satisfying truth assignment. Stop.
5. Repeat step 1 using the resulting simplified formula.

In doing so, we take advantage of a basic property that Horn systems are closed under the deletion of literals and clauses and therefore applying unit resolution on a Horn system results in a new system that is also Horn.

Consider the Horn clause system $(x_3) \wedge (x_1 \vee \bar{x}_2) \wedge (\bar{x}_1 \vee \bar{x}_3)$. Applying unit resolution once yields $(x_1 \vee \bar{x}_2) \wedge (\bar{x}_1)$ with x_3 set to 1. In the usual case, we might re-apply unit resolution on the unit clause (\bar{x}_1) , but because this system is Horn and there are no other unit positive clauses, we instead assign 0 to both x_1 and x_2 and declare the proposition as satisfiable.

Theorem A *satisfiable* Horn proposition has a *unique minimal model*. A unique minimal model for a satisfiable proposition S is achieved by:

- (i) Setting all atoms to T *only* for those atoms that *must* be true in *all* models of S .
- (ii) Setting everything else to F.

The restricted form of unit resolution satisfies (3i). The minimal model tries to minimize the number of T's assigned to the variables, or equivalently, tries to maximize the number of F's assigned.

The minimal model is unique because of the following argument. Assume there exists two minimal models T_1 and T_2 different from each other. If an atom set to T in T_1 is set to F in T_2 and T_1 is a model, then T_2 cannot be a model because (3i) sets *only* those atoms that must be true in all models to T since an atom is set to T only when it shows up as a unit positive clause (positive singleton).

Thus if there exists two minimal models T_1 and T_2 , they must be identical to each other. \square

Even if there exists more than one unit positive clause to resolve on, the procedure will yield the same minimal model regardless of which clause unit resolution picks to resolve on.

The above theorem also follows from a very strong *closure* property satisfied by the set of models of a Horn proposition.

Lemma If T_1 and T_2 are models for a Horn proposition S , then so is $T_1 \wedge T_2$. (An atom is set to true in $T_1 \wedge T_2$ if and only if it is true in both T_1 and T_2 .)

Proof: Assume T_1 and T_2 are models for a Horn proposition S , but $T_1 \wedge T_2$ isn't.

1. If a clause C is negative (meaning no positive literals), then $T_1 \wedge T_2$ must set all literals in C to T since setting even one literal to F will satisfy the clause since all literals are negated. But the only way all literals can be set to T is if both T_1 and T_2 set all corresponding literals to be T which causes both T_1 and T_2 to falsify C since all literals are negated. Therefore neither T_1 or T_2 can be models and this is a contradiction.
2. If C has a single positive literal x_k then at least one of T_1 or T_2 (say T_1) must set x_k to F, for otherwise $T_1 \wedge T_2$ would satisfy C . But since T_1 is a model, there must be at least one negated literal \bar{x}_j in C (no further positive literals are allowed since C is Horn) set to F, but this would automatically make the corresponding negated literal in $T_1 \wedge T_2$ F causing C to be satisfied. Again a contradiction.

This concludes that models of Horn propositions are closed under the “ \wedge ” operation. \square

Definition Two propositions S_1 and S_2 are said to be *equivalent* if they are built on the same ground set of atoms and if they have the same set of models.

It is easy to see that two propositions are equivalent if and only if S_1 implies S_2 and S_2 implies S_1 . We shall use the notation $S_1 \Leftrightarrow S_2$ for equivalent propositions S_1 and S_2 .

Definition A proposition S is called *H-equivalent* (or Horn equivalent) if it is equivalent to some Horn proposition.

So if S is Horn equivalent to some Horn proposition H , the above lemma suggests that the set of models of H and hence of S is closed under the \wedge operation. Now suppose we have a non-Horn proposition S that is closed under the \wedge operation, does it suggest that there exists some Horn proposition H to which S is *H-equivalent*? It turns out that this is indeed the case because of the following reasoning.

If S is non-Horn, it must contain a clause C of the form $x_k \vee x_l \vee D$, where D is a disjunction of zero or more literals. If every model of S satisfies $x_k \vee D$ we can delete x_l from C and obtain an equivalent proposition. Likewise if $x_l \vee D$ is satisfied by all models, we delete x_k from C . We continue this process, gradually deleting the excess positive literals from clauses in S , until we obtain a Horn proposition equivalent to S or we are unable to apply the deletion criteria. If we are unable to apply the deletion criteria it is because we have two models T_1 and T_2 and a clause $x_k \vee x_l \vee D$ such that

$$\begin{aligned} T_1(x_k) &= True, T_1(x_l) = T_1(D) = False \\ T_2(x_l) &= True, T_2(x_k) = T_2(D) = False \end{aligned}$$

But then $T_1 \wedge T_2$ cannot satisfy C , which contradicts our assumption that the models of S are closed under the \wedge operation. Therefore the deletion process does not get stuck and S is reduced to an equivalent Horn formula.

We thus have the following result:

Theorem The set of models of a proposition is closed under the “ \wedge ” operation if and only if the proposition is *H-equivalent*.

2.3.2 The Integer Least Element of a Horn Polytope

The rich syntactic and semantic structure of Horn propositions is revealed as special integrality properties of the LP relaxation. This helps shape characteristics of the polytopes formed from the linear programming relaxation of Horn propositions, that is, Horn polytopes.

Definition A least element of a polyhedron P is a point $x_{min} \in P$, all of whose individual components are no larger than the corresponding components of any x in P . In mathematical terms, x_{min} is a least element if $x_{min} \leq x \forall x \in P$.

Of course, not every convex polyhedron has a least element. In the two-dimensional case, the least element, assuming one exists, will be somewhere at the bottom-left of P . If the bottom-most left-most point is not a vertex of the polyhedron, it is clear that it cannot possibly be a least element. It is also clear that if a least element exists, then it must be unique since if there exists a second least element different from the first, it would automatically imply one of the two least elements is not least anymore.

Theorem A convex polyhedron defined by a system of linear inequalities

$$P = \{ x \mid Ax \geq b, x \geq 0 \}$$

has an *integral* least element if the following conditions are met:

1. b must be integral
2. b must be such that P is non-empty
3. each row of A must have at most one positive component
4. all positive components of A must be equal to 1

If $b \leq 0$ it is evident that P contains a least element $x_{min} = 0$. If $b > 0$, [1] (p. 35) proposes a technique to obtain the least element of P using a lower bound escalation scheme which performs a sequence of translations until the least element is resolved.

Since each inequality associated with a Horn clause will have at most one positive coefficient on the left-hand side and further all positive left-hand-side coefficients will be equal to 1, we are assured the existence of an integral least element for all Horn polytopes generated by satisfiable Horn propositions. Further, this least element is the incidence vector of the unique minimum model of the proposition found by Horn resolution since both methods try to maximize the number of F's (0s) assigned to the variables.

It is important to note that although polytopes generated by satisfiable Horn propositions are guaranteed to have an integral least element, the polytope need not itself be integral.

Thus, one can find the integral least element by optimization since we are guaranteed that for satisfiable Horn propositions, a least element exists and this least element needs to be a vertex of the polytope from the way a least element was defined. Hence a simplex algorithm that hops from vertex to vertex will eventually hit upon this integral least element which corresponds to the unique minimal model, thus proving feasibility of the ILP formulation, and hence satisfiability.

Summarizing this idea, for any vector $c \in \mathfrak{R}$, all of whose components are positive (the simplest would be a vector of all ones), the linear program

$$\min \{ c^T x \mid x \in \text{Horn Polytope} \}$$

is optimized uniquely by the integral least element. Of course, this optimization model is nowhere as efficient as Horn resolution in proving satisfiability. However, the *dual* of the optimization model above may provide additional mathematical insight to give us some idea for polytopes generated from *unsatisfiable* Horn propositions.

2.3.3 Dual Integrality of Horn Polytopes

The dual of the integer linear representation of a satisfiability problem has an interesting interpretation discovered by Jeroslow and Wang. When the clauses are unsatisfiable, the values of the dual variables are proportional to the number of times the corresponding clauses serve as premises in a refutation.

Any satisfiability problem can be written as the following 0–1 problem:

$$\begin{array}{ll} \text{MIN} & x_0 \\ \text{s.t.} & x_0 e + Ax \geq a \\ & x_j \in \{0, 1\}, j = 0, 1, \dots, n \end{array}$$

The linear relaxation of this problem is:

$$\begin{array}{ll} \text{MIN} & x_0 \\ \text{s.t.} & x_0 e + Ax \geq a \quad (u) \\ & -x \geq -e \quad (v) \\ & x \geq 0 \end{array}$$

and its dual is:

$$\begin{array}{ll} \text{MAX} & ua - v \\ \text{s.t.} & ue \leq 1 \\ & uA - v \leq 0 \\ & -u \leq 0 \\ & -v \leq 0 \end{array}$$

The dual solution implicitly indicates how many times each clause is used in a resolution proof of unsatisfiability.

Theorem Let $Ax \geq a$ represent an unsatisfiable set of Horn clauses. Then if (u, v) is any optimal extreme point solution of the dual, there is an integer N and a refutation proof of unsatisfiability such that Nu_i is the number of times that each clause i is used to obtain the empty clause in the proof.

For example, consider the following *unsatisfiable* Horn clauses:

$$\begin{array}{l} x_1 \\ \bar{x}_1 \vee x_2 \\ \bar{x}_2 \end{array}$$

We add an artificial variable x_0 to check satisfiability:

$$\begin{array}{ll} \text{MIN} & x_0 \\ \text{s.t.} & x_0 + x_1 \geq 1 \quad (u_1) \\ & x_0 - x_1 + x_2 \geq 0 \quad (u_2) \\ & x_0 - x_2 \geq 0 \quad (u_3) \end{array}$$

The optimal solution to the primal is $\bar{x} = (x_0, x_1, x_2) = (\frac{1}{3}, \frac{2}{3}, \frac{1}{3})$ while the corresponding dual solution is $\bar{u} = (u_1, u_2, u_3) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ which is non-integral. The theorem states that for some N , Nu gives the number of times each clause is used to obtain the empty clause. The refutation is achieved by first resolving the first two clauses to obtain x_2 , and then resolving x_2 with the third clause to obtain the empty clause. So each clause is used once and $N = 3$.

Unfortunately, the dual multipliers u_i do not in general encode the *structure* of a refutation proof and therefore do not represent a complete resolution proof.

3 Summary of Findings

1. We show that the two questions of whether a formula is satisfiable and whether one formula implies another are one and the same. Each form may be converted to the other and therefore both problems are equally hard.
2. In Section 2.2.1 we explain *why* the proposed optimization model for the inference problem of the form $(S_1 \rightarrow S_2)$ works. We begin by assuming that S_1 is CNF and S_2 is given by single clause C . We then extend the argument to the case where S_2 is given by more than just a clause, but is still CNF.
3. For satisfiability problems on Horn systems, we may be able to speed up Horn resolution a little by skipping resolution on a unit positive clause by making the substitution $x_j = \bar{x}_j$ since Horn resolution only resolves unit *positive* clauses. We must of course remember to switch the corresponding value for the propositional atom for each of the modified formula's models. However this substitution is permissible only if *every* other clause containing \bar{x}_j does not already contain a positive clause so as to remain Horn after the substitution.

4 List of Theorems

1. **Theorem** Any formula in propositional logic is equivalent to a CNF formula whose length is linearly related to the length of the original formula.
2. **Theorem** A proposition S has a unit refutation (unsatisfiability of S proved by unit resolution) if and only if the linear programming relaxation of S is infeasible.
3. **Theorem** A *satisfiable* Horn proposition has a *unique minimal model*. A unique minimal model for a satisfiable proposition S is achieved by:
 - (i) Setting all atoms to T *only* for those atoms that *must* be true in *all* models of S .
 - (ii) Setting everything else to F.
4. **Theorem** The set of models of a proposition is closed under the “ \wedge ” operation if and only if the proposition is *H-equivalent*.
5. **Theorem** A convex polyhedron defined by a system of linear inequalities

$$P = \{ x \mid Ax \geq b, x \geq 0 \}$$

has an *integral* least element if the following conditions are met:

- (a) b must be integral
 - (b) b must be such that P is non-empty
 - (c) each row of A must have at most one positive component
 - (d) all positive components of A must be equal to 1
6. **Theorem** Let $Ax \geq a$ represent an unsatisfiable set of Horn clauses. Then if (u, v) is any optimal extreme point solution of the dual, there is an integer N and a refutation proof of unsatisfiability such that Nu_i is the number of times that each clause i is used to obtain the empty clause in the proof.

Bibliography

- [1] Vijay Chandru, John N. Hooker, "Optimization Methods for Logical Inference", John Wiley & Sons, Inc., 1999.