

2.6 The DPLL Procedure

Goal:

Given a propositional formula in CNF (or alternatively, a finite set N of clauses), check whether it is satisfiable (and optionally: output *one* solution, if it is satisfiable).

Preliminaries

Recall:

$\mathcal{A} \models N$ if and only if $\mathcal{A} \models C$ for all clauses C in N .

$\mathcal{A} \models C$ if and only if $\mathcal{A} \models L$ for some literal $L \in C$.

Assumptions:

Clauses contain neither duplicated literals nor complementary literals.

The order of literals in a clause is irrelevant.

\Rightarrow Clauses behave like *sets* of literals.

Notation:

We use the notation $C \vee L$ to denote a clause with some literal L and a clause rest C . Here L need *not* be the last literal of the clause and C may be empty.

\overline{L} is the complementary literal of L , i. e., $\overline{P} = \neg P$ and $\overline{\overline{P}} = P$.

Partial Valuations

Since we will construct satisfying valuations incrementally, we consider *partial valuations* (that is, partial mappings $\mathcal{A} : \Pi \rightarrow \{0, 1\}$).

Every partial valuation \mathcal{A} corresponds to a set M of literals that does not contain complementary literals, and vice versa:

$\mathcal{A}(L)$ is true, if $L \in M$.

$\mathcal{A}(L)$ is false, if $\overline{L} \in M$.

$\mathcal{A}(L)$ is undefined, if neither $L \in M$ nor $\overline{L} \in M$.

We will use \mathcal{A} and M interchangeably.

A clause is true under a partial valuation \mathcal{A} (or under a set M of literals) if one of its literals is true; it is false (or “*conflicting*”) if all its literals are false; otherwise it is undefined (or “*unresolved*”).

Unit Clauses

Observation:

Let \mathcal{A} be a partial valuation. If the set N contains a clause C , such that all literals but one in C are false under \mathcal{A} , then the following properties are equivalent:

- there is a valuation that is a model of N and extends \mathcal{A} .
- there is a valuation that is a model of N and extends \mathcal{A} and makes the remaining literal L of C true.

C is called a *unit clause*; L is called a *unit literal*.

Pure Literals

One more observation:

Let \mathcal{A} be a partial valuation and P a variable that is undefined under \mathcal{A} . If P occurs only positively (or only negatively) in the unresolved clauses in N , then the following properties are equivalent:

- there is a valuation that is a model of N and extends \mathcal{A} .
- there is a valuation that is a model of N and extends \mathcal{A} and assigns 1 (0) to P .

P is called a *pure literal*.

The Davis-Putnam-Logemann-Loveland Proc.

```
boolean DPLL(literal set  $M$ , clause set  $N$ ) {
  if (all clauses in  $N$  are true under  $M$ ) return true;
  elif (some clause in  $N$  is false under  $M$ ) return false;
  elif ( $N$  contains unit literal  $P$ ) return DPLL( $M \cup \{P\}$ ,  $N$ );
  elif ( $N$  contains unit literal  $\neg P$ ) return DPLL( $M \cup \{\neg P\}$ ,  $N$ );
  elif ( $N$  contains pure literal  $P$ ) return DPLL( $M \cup \{P\}$ ,  $N$ );
  elif ( $N$  contains pure literal  $\neg P$ ) return DPLL( $M \cup \{\neg P\}$ ,  $N$ );
  else {
    let  $P$  be some undefined variable in  $N$ ;
    if (DPLL( $M \cup \{\neg P\}$ ,  $N$ )) return true;
    else return DPLL( $M \cup \{P\}$ ,  $N$ );
  }
}
```

Initially, DPLL is called with an empty literal set and the clause set N .

2.7 From DPLL to CDCL

In practice, there are several changes to the procedure:

The pure literal check is only done while preprocessing (otherwise is too expensive).

The algorithm is implemented iteratively \Rightarrow the backtrack stack is managed explicitly (it may be possible and useful to backtrack more than one level).

Information is reused by conflict analysis and learning.

The branching variable is not chosen randomly.

Under certain circumstances, the procedure is restarted.

Conflict Analysis and Learning

Conflict analysis serves a dual purpose:

Backjumping (non-chronological backtracking): If we detect that the conflict is independent of some earlier branch, we can skip over that backtrack level.

Learning: By deriving a new clause from the conflict that is added to the current set of clauses, we can reuse information that is obtained in one branch in further branches. (Note: This may produce a large number of new clauses; therefore it may become necessary to delete some of them afterwards to save space.)

These ideas are implemented in all modern SAT solvers.

Because of the importance of clause learning the algorithm is now called CDCL: Conflict Driven Clause Learning.

Formalizing DPLL with Refinements

We model the improved DPLL procedure by a transition relation $\Rightarrow_{\text{CDCL}}$ on a set of states.

States:

- *fail*
- $M \parallel N$,

where M is a *list of annotated literals* and N is a set of clauses.

Annotated literal:

- L : deduced literal, due to unit propagation.
- L^d : decision literal (guessed literal).

Unit Propagate:

$$M \parallel N \cup \{C \vee L\} \Rightarrow_{\text{CDCL}} M L \parallel N \cup \{C \vee L\}$$

if C is false under M and L is undefined under M .

Decide:

$$M \parallel N \Rightarrow_{\text{CDCL}} M L^d \parallel N$$

if L is undefined under M and contained in N .

Fail:

$$M \parallel N \cup \{C\} \Rightarrow_{\text{CDCL}} \text{fail}$$

if C is false under M and M contains no decision literals.

Backjump:

$$M' L^d M'' \parallel N \Rightarrow_{\text{CDCL}} M' L' \parallel N$$

if there is some “backjump clause” $C \vee L'$ such that

$$N \models C \vee L',$$

C is false under M' , and

L' is undefined under M' .

We will see later that the Backjump rule is always applicable, if the list of literals M contains at least one decision literal and some clause in N is false under M .

There are many possible backjump clauses. One candidate: $\overline{L_1} \vee \dots \vee \overline{L_n}$, where the L_i are all the decision literals in $M' L^d M''$. (But usually there are better choices.)

Lemma 2.15 *If we reach a state $M \parallel N$ starting from $\varepsilon \parallel N$, then:*

- (1) M does not contain complementary literals.
- (2) Every deduced literal L in M follows from N and decision literals occurring before L in M .

Proof. By induction on the length of the derivation. □

Lemma 2.16 *Every derivation starting from $\varepsilon \parallel N$ terminates.*

Proof. Let $M \parallel N$ and $M' \parallel N'$ be two CDCL states, such that $M = M_0 L_1^d M_1 \dots L_k^d M_k$ and $M' = M'_0 L_1'^d M'_1 \dots L_{k'}^d M'_{k'}$. We define a relation \succ on lists of annotated literals by $M \succ M'$ if and only if

- (i) there is some j such that $0 \leq j \leq \min(k, k')$, $|M_i| = |M'_i|$ for all $0 \leq i < j$, and $|M_j| < |M'_j|$, or
- (ii) $|M_i| = |M'_i|$ for all $0 < i \leq k < k'$ and $|M| < |M'|$.

It is routine to check that \succ is irreflexive and transitive, hence a strict partial ordering, and that for every CDCL step $M \parallel N \Rightarrow_{\text{CDCL}} M' \parallel N'$ we have $M \succ M'$. Moreover, the set of propositional variables in N is finite, and each of these variables can occur at most once in a literal list (positively or negatively, with or without a d-superscript). So there are only finitely many literal lists that can occur in a CDCL derivation. Consequently, if there were an infinite CDCL derivation, there would be some cycle $M \parallel N \Rightarrow_{\text{CDCL}}^+ M \parallel N'$, so by transitivity $M \succ M$, but that would contradict the irreflexivity of \succ .

Lemma 2.17 *Suppose that we reach a state $M \parallel N$ starting from $\varepsilon \parallel N$ such that some clause $D \in N$ is false under M . Then:*

- (1) *If M does not contain any decision literal, then “Fail” is applicable.*
- (2) *Otherwise, “Backjump” is applicable.*

Proof. (1) Obvious.

(2) Let L_1, \dots, L_n be the decision literals occurring in M (in this order). Since $M \models \neg D$, we obtain, by Lemma 2.15, $N \cup \{L_1, \dots, L_n\} \models \neg D$. Since $D \in N$, this is a contradiction, so $N \cup \{L_1, \dots, L_n\}$ is unsatisfiable. Consequently, $N \models \overline{L_1} \vee \dots \vee \overline{L_n}$. Now let $C = \overline{L_1} \vee \dots \vee \overline{L_{n-1}}$, $L' = \overline{L_n}$, $L = L_n$, and let M' be the list of all literals of M occurring before L_n , then the condition of “Backjump” is satisfied. \square

Theorem 2.18 *Suppose that we reach a final state starting from $\varepsilon \parallel N$.*

- (1) *If the final state is $M \parallel N$, then N is satisfiable and M is a model of N .*
- (2) *If the final state is fail, then N is unsatisfiable.*

Proof. (1) Observe that the “Decide” rule is applicable as long as literals in N are undefined under M . Hence, in a final state, all literals must be defined. Furthermore, in a final state, no clause in N can be false under M , otherwise “Fail” or “Backjump” would be applicable. Hence M is a model of every clause in N .

(2) If we reach *fail*, then in the previous step we must have reached a state $M \parallel N$ such that some $C \in N$ is false under M and M contains no decision literals. By part (2) of Lemma 2.15, every literal in M follows from N . On the other hand, $C \in N$, so N must be unsatisfiable. \square

Getting Better Backjump Clauses

Suppose that we have reached a state $M \parallel N$ such that some clause $C \in N$ (or following from N) is false under M .

Consequently, every literal of C is the complement of some literal in M .

- (1) If every literal in C is the complement of a decision literal of M , then C is a backjump clause.
- (2) Otherwise, $C = C' \vee \bar{L}$, such that L is a deduced literal.

For every deduced literal L , there is a clause $D \vee L$, such that $N \models D \vee L$ and D is false under M .

Then $N \models D \vee C'$ and $D \vee C'$ is also false under M . ($D \vee C'$ is a *resolvent* of $C' \vee \bar{L}$ and $D \vee L$.)

By repeating this process, we will eventually obtain a clause that satisfies the requirements of a backjump clause.

Usually, one resolves the literals in the reverse order in which they were added to M and stops as soon as one obtains a clause in which all but one literal are complements of literals occurring in M before the last decision literal.

\Rightarrow 1UIP (first unique implication point) strategy.

Learning Clauses

Backjump clauses are good candidates for learning.

To model learning, the CDCL system is extended by the following two rules:

Learn:

$$\begin{aligned} M \parallel N &\Rightarrow_{\text{CDCL}} M \parallel N \cup \{C\} \\ &\text{if } N \models C. \end{aligned}$$

Forget:

$$\begin{aligned} M \parallel N \cup \{C\} &\Rightarrow_{\text{CDCL}} M \parallel N \\ &\text{if } N \models C. \end{aligned}$$

If we ensure that no clause is learned infinitely often, then termination is guaranteed.

The other properties of the basic CDCL system hold also for the extended system.

Restart

Runtimes of CDCL-style procedures depend extremely on the choice of branching variables.

If no solution is found within a certain time limit, it can be useful to *restart* from scratch with an adapted variable selection heuristics. Learned clauses, however, are kept.

In addition, it is useful to restart after a unit clause has been learned.

The restart rule is typically applied after a certain number of clauses have been learned or a unit is derived:

Restart:

$$M \parallel N \Rightarrow_{\text{CDCL}} \varepsilon \parallel N$$

If Restart is only applied finitely often, termination is guaranteed.

2.8 Implementing CDCL

The formalization of CDCL that we have seen so far leaves many aspects unspecified.

To get a fast solver, we must use good heuristics, for instance to choose the next undefined variable, and we must implement basic operations efficiently.

Variable Order Heuristic

Choosing the right undefined variable to branch is important for efficiency, but the branching heuristics may be expensive itself.

State of the art: Use branching heuristics that need not be recomputed too frequently.

In general: Choose variables that occur frequently; after a restart prefer variables from recent conflicts.

The VSIDS (Variable State Independent Decaying Sum) heuristic:

- We associate a positive *score* to every propositional variable P_i . At the start, k_i is the number of occurrences of P_i in N .
- The variable order is then the descending ordering of the P_i according to the k_i .

The scores k_i are adjusted during a CDCL run.

- Every time a learned clause is computed after a conflict, the propositional variables in the learned clause obtain a bonus b , i.e., $k_i := k_i + b$.
- Periodically, the scores are leveled: $k_i := k_i/l$ for some l .

- After each restart, the variable order is recomputed, using the new scores.

The purpose of these mechanisms is to keep the search focused. The parameter b directs the search around the conflict,

Further refinements:

- Add the bonus to all literals in the clauses that occur in the resolution steps to generate a backjump clause.
- If the score of a variable reaches a certain limit, all scores are rescaled by a constant.
- Occasionally (with low probability) choose a variable at random, otherwise choose the undefined variable with the highest score.

Implementing Unit Propagation Efficiently

For applying the unit rule, we need to know the number of literals in a clause that are not false.

Maintaining this number is expensive, however.

Better approach: “*Two watched literals*”:

In each clause, select two (currently undefined) “watched” literals.

For each variable P , keep a list of all clauses in which P is watched and a list of all clauses in which $\neg P$ is watched.

If an undefined variable is set to 0 (or to 1), check all clauses in which P (or $\neg P$) is watched and watch another literal (that is true or undefined) in this clause if possible.

Watched literal information need not be restored upon backtracking.

Preprocessing

Some operations are only needed once at the beginning of the CDCL run.

- (i) Deletion of tautologies
- (ii) Deletion of duplicated literals

Some operations are useful, but so expensive that they are performed only initially and after restarts (before computation of the variable order heuristics):

(iii) Subsumption

$$N \cup \{C\} \cup \{D\} \Rightarrow N \cup \{C\}$$

if $C \subseteq D$ considering C, D as multisets of literals.

(iv) Purity deletion

Delete all clauses containing a literal L where \bar{L} does not occur in the clause set.

(v) Merging replacement resolution

$$N \cup \{C \vee L\} \cup \{D \vee \bar{L}\} \Rightarrow N \cup \{C \vee L\} \cup \{D\}$$

if $C \subseteq D$ considering C, D as multisets of literals.

(vi) Literal elimination

Compute all possible resolution steps

$$\frac{C \vee L \quad D \vee \bar{L}}{C \vee D}$$

on a literal L with premises in N ; add the conclusions to N and then throw away all clauses containing L or \bar{L} ; repeat this as long as $|N|$ does not grow.

Literature

Lintao Zhang and Sharad Malik: The Quest for Efficient Boolean Satisfiability Solvers; Proc. CADE-18, LNAI 2392, pp. 295–312, Springer, 2002.

Robert Nieuwenhuis, Albert Oliveras, Cesare Tinelli: Solving SAT and SAT Modulo Theories; From an abstract Davis-Putnam-Logemann-Loveland procedure to DPLL(T), pp. 937–977, Journal of the ACM, 53(6), 2006.

Armin Biere, Marijn Heule, Hans van Maaren, Toby Walsh (eds.): Handbook of Satisfiability; IOS Press, 2009

Daniel Le Berre's slides at VTSA'09: <http://www.mpi-inf.mpg.de/vtsa09/>.

2.9 Other Calculi

OBDDs (Ordered Binary Decision Diagrams):

Minimized graph representation of decision trees, based on a fixed ordering on propositional variables,

⇒ canonical representation of formulas.

see Chapter 6.1/6.2 of Michael Huth and Mark Ryan: *Logic in Computer Science: Modelling and Reasoning about Systems*, Cambridge Univ. Press, 2000.

FRAIGs (Fully Reduced And-Inverter Graphs)

Minimized graph representation of boolean circuits.

⇒ semi-canonical representation of formulas.

Implementation needs DPLL (and OBDDs) as subroutines.

Ordered resolution

Tableau calculus

Hilbert calculus

Sequent calculus

Natural deduction

see next chapter

3 First-Order Logic

First-order logic

- formalizes fundamental mathematical concepts
- is expressive (Turing-complete)
- is not too expressive (e. g. not axiomatizable: natural numbers, uncountable sets)
- has a rich structure of decidable fragments
- has a rich model and proof theory

First-order logic is also called (first-order) *predicate logic*.

3.1 Syntax

Syntax:

- non-logical symbols (domain-specific)
⇒ terms, atomic formulas
- logical connectives (domain-independent)
⇒ Boolean combinations, quantifiers

Signatures

A signature $\Sigma = (\Omega, \Pi)$ fixes an alphabet of non-logical symbols, where

- Ω is a set of *function symbols* f with *arity* $n \geq 0$, written $\text{arity}(f) = n$,
- Π is a set of *predicate symbols* P with *arity* $m \geq 0$, written $\text{arity}(P) = m$.

Function symbols are also called *operator symbols*.

If $n = 0$ then f is also called a *constant (symbol)*.

If $m = 0$ then P is also called a *propositional variable*.

We will usually use

b, c, d for constant symbols,

f, g, h for non-constant function symbols,

P, Q, R, S for predicate symbols.

Convention: We will usually write $f/n \in \Omega$ instead of $f \in \Omega$, $\text{arity}(f) = n$ (analogously for predicate symbols).

Refined concept for practical applications:

many-sorted signatures (corresponds to simple type systems in programming languages);
no big change from a logical point of view.

Variables

Predicate logic admits the formulation of abstract, schematic assertions. (Object) variables are the technical tool for schematization.

We assume that X is a given countably infinite set of symbols which we use to denote *variables*.

Terms

Terms over Σ and X (Σ -terms) are formed according to these syntactic rules:

$$s, t, u, v ::= x \quad , x \in X \quad \text{(variable)}$$

$$\quad \quad \quad | \quad f(s_1, \dots, s_n) \quad , f/n \in \Omega \quad \text{(functional term)}$$

By $T_\Sigma(X)$ we denote the set of Σ -terms (over X). A term not containing any variable is called a *ground term*. By T_Σ we denote the set of Σ -ground terms.

In other words, terms are formal expressions with well-balanced brackets which we may also view as marked, ordered trees. The markings are function symbols or variables. The nodes correspond to the *subterms* of the term. A node v that is marked with a function symbol f of arity n has exactly n subtrees representing the n immediate subterms of v .

Atoms

Atoms (also called atomic formulas) over Σ are formed according to this syntax:

$$A, B ::= P(s_1, \dots, s_m) \quad , P/m \in \Pi \quad \text{(non-equational atom)}$$

$$\quad \quad \quad \left[\quad | \quad (s \approx t) \quad \quad \quad \text{(equation)} \quad \right]$$

Whenever we admit equations as atomic formulas we are in the realm of *first-order logic with equality*. Admitting equality does not really increase the expressiveness of first-order logic, (cf. exercises). But deductive systems where equality is treated specifically are much more efficient.