

Logic

(Métodos Formais em Engenharia de Software)

Maria João Frade

HASLab - INESC TEC
Departamento de Informática, Universidade do Minho

2013/2014

What is a (formal) logic?

Logic is defined as the study of the principles of reasoning. One of its branches is symbolic logic, that studies formal logic.

- A **formal logic** is a language equipped with rules for deducing the truth of one sentence from that of another.
- A logic consists of
 - ▶ A *logical language* in which (well-formed) sentences are expressed.
 - ▶ A *semantics* that defines the intended interpretation of the symbols and expressions of the logical language.
 - ▶ A *proof system* that is a framework of rules for deriving valid judgments.
- Examples: propositional logic, first-order logic, higher-order logic, modal logics, ...

What is a logical language?

A logical language consists of

- *logical symbols* whose interpretations are fixed
- *non-logical symbols* whose interpretations vary

These symbols are combined together to form *well-formed formulas*.

Logic and computer science

- Logic and computer science share a symbiotic relationship
 - ▶ Logic provides language and methods for the study of theoretical computer science.
 - ▶ Computers provide a concrete setting for the implementation of logic.
- Formal logic makes it possible to calculate consequences at the symbolic level, so computers can be used to automate such symbolic calculations.
- Moreover, logic can be used to model the situations we encounter as computer science professionals, in such a way that we can reason about them formally.

Motivation

- Many applications of formal methods rely on generating formulas of a logical system and investigate about their validity or satisfiability.
- Constraint-satisfaction problems arise in diverse application areas, such as
 - ▶ software and hardware verification
 - ▶ static program analysis
 - ▶ test-case generation
 - ▶ scheduling and planning
 - ▶ ...
- These problems can be encoded by logical formulas. Solvers for such formulations (SAT solvers and SMT solvers) play a crucial role in their resolution.

Motivation

- Increased attention has led to enormous progress in this area in the last decade. Modern SAT procedures can check formulas with hundreds of thousands of variables. Similar progress has been observed for SMT solvers for more commonly occurring theories.
- SMT solvers are the core engine of many tools for program analysis, testing and verification.
- Modern SMT solvers integrate specialized theory solvers with propositional satisfiability search techniques.

Goals of this course

- Review the basic concepts of Propositional Logic and First-Order Logic.
- Address the issues of decidability of logical systems.
- Talk about the algorithms that underlie a large number of automatic proof tools.
- Illustrate the use of automatic theorem provers and proof assistants.

Bibliography

Books

- [Huth&Ryan 2004] *Logic in Computer Science: Modelling and Reasoning About Systems*. Michael Huth & Mark Ryan. Cambridge University Press; 2nd edition (2004).
- [Bradley&Manna 2007] *The Calculus of Computation: Decision Procedures with Applications to Verification*. Aaron R. Bradley & Zohar Manna. Springer (2007).
- [RSD 2011] *Rigorous Software Development: An Introduction to Program Verification*. J.B. Almeida & M.J. Frade & J.S. Pinto & S.M. de Sousa. Springer (2011)
- [Bertot&Casteran 2004] *Interactive Theorem Proving and Program Development Coq'Art: The Calculus of Inductive Constructions*. Yves Bertot & Pierre Casteran. Springer (2004)

Bibliography

Papers

- *Satisfiability Modulo Theories: An Appetizer*. Leonardo de Moura & Nikolaj Bjørner. Invited paper to SBMF 2009, Gramado, Brazil.
- *Satisfiability Modulo Theories: Introduction and Applications*. Leonardo de Moura & Nikolaj Bjørner. Communications of the ACM, September 2011.
- *Automated Deduction for Verification*. Natarajan Shankar, ACM Computing Surveys, Vol. 41, No. 4, Article 20, October 2009.
- *Coq in a Hurry*. Yves Bertot. February 2010.

(Classical) Propositional Logic

Roadmap

Classical Propositional Logic

- syntax; semantics; decision problems SAT and VAL;
- normal forms; equisatisfiability; Tseitin's encoding;
- SAT solving algorithms: basic concepts;
- DPLL framework and its optimisations;
- modeling with PL;
- exercises.

Introduction

- The language of propositional logic is based on **propositions**, or **declarative sentences** which one can, in principle, argue as being “true” or “false”.
 “The capital of Portugal is Braga.”
 “D. Afonso Henriques was the first king of Portugal.”
- Propositional symbols are the **atomic formulas** of the language. More complex sentences are constructed using **logical connectives**.
- In classical propositional logic (PL) each sentence is either true or false.
- In fact, the content of the propositions is not relevant to PL. PL is not the study of truth, but of the relationship between the truth of one statement and that of another.

Syntax

The alphabet of the propositional language is organised into the following categories.

- **Propositional variables:** $P, Q, R, \dots \in \mathcal{V}_{\text{Prop}}$ (a countably infinite set)
- **Logical connectives:** \perp (*false*), \top (*true*), \neg (*not*), \wedge (*and*), \vee (*or*), \rightarrow (*implies*), \leftrightarrow (*equivalent*)
- **Auxiliary symbols:** “(“ and “)”.

The set **Form** of *formulas* of propositional logic is given by the abstract syntax

Form $\ni A, B ::= P \mid \perp \mid \top \mid (\neg A) \mid (A \wedge B) \mid (A \vee B) \mid (A \rightarrow B) \mid (A \leftrightarrow B)$

We let A, B, C, F, G, H, \dots range over **Form**.

Outermost parenthesis are usually dropped. In absence of parentheses, we adopt the following convention about precedence. Ranging from the highest precedence to the lowest, we have respectively: \neg , \wedge , \vee , \rightarrow and \leftrightarrow . All binary connectives are right-associative.

Semantics

The semantics of a logic provides its meaning. **What exactly is meaning?** In propositional logic, meaning is given by the truth values *true* and *false*, where $\text{true} \neq \text{false}$. We will represent true by 1 and false by 0.

An **assignment** is a function $\mathcal{A} : \mathcal{V}_{\text{Prop}} \rightarrow \{0, 1\}$, that assigns to every propositional variable a truth value.

An assignment \mathcal{A} naturally extends to all formulas, $\mathcal{A} : \mathbf{Form} \rightarrow \{0, 1\}$.

The truth value of a formula is computed using **truth tables**:

F	A	B	$\neg A$	$A \wedge B$	$A \vee B$	$A \rightarrow B$	$A \leftrightarrow B$	\perp	\top
$\mathcal{A}_1(F)$	0	1	1	0	1	1	0	0	1
$\mathcal{A}_2(F)$	0	0	1	0	0	1	1	0	1
$\mathcal{A}_3(F)$	1	1	0	1	1	1	1	0	1
$\mathcal{A}_4(F)$	1	0	0	0	1	0	0	0	1

Semantics

Let \mathcal{A} be an assignment and let F be a formula.

If $\mathcal{A}(F) = 1$, then we say F **holds** under assignment \mathcal{A} , or \mathcal{A} **models** F .

We write $\mathcal{A} \models F$ iff $\mathcal{A}(F) = 1$, and $\mathcal{A} \not\models F$ iff $\mathcal{A}(F) = 0$.

An alternative (inductive) definition of $\mathcal{A} \models F$ is

$\mathcal{A} \models \top$	
$\mathcal{A} \not\models \perp$	
$\mathcal{A} \models P$	iff $\mathcal{A}(P) = 1$
$\mathcal{A} \models \neg A$	iff $\mathcal{A} \not\models A$
$\mathcal{A} \models A \wedge B$	iff $\mathcal{A} \models A$ and $\mathcal{A} \models B$
$\mathcal{A} \models A \vee B$	iff $\mathcal{A} \models A$ or $\mathcal{A} \models B$
$\mathcal{A} \models A \rightarrow B$	iff $\mathcal{A} \not\models A$ or $\mathcal{A} \models B$
$\mathcal{A} \models A \leftrightarrow B$	iff $\mathcal{A} \models A$ iff $\mathcal{A} \models B$

Validity, satisfiability, and contradiction

A formula F is

valid iff it holds under every assignment. We write $\models F$.
A valid formula is called a **tautology**.

satisfiable iff it holds under some assignment.

unsatisfiable iff it holds under no assignment.
An unsatisfiable formula is called a **contradiction**.

refutable iff it is not valid.

Proposition

F is valid iff $\neg F$ is a contradiction

$(A \wedge (A \rightarrow B)) \rightarrow B$ is valid. $A \rightarrow B$ is satisfiable and refutable.
 $A \wedge \neg A$ is a contradiction.

Consequence and equivalence

- $F \models G$ iff for every assignment \mathcal{A} , if $\mathcal{A} \models F$ then $\mathcal{A} \models G$. We say G is a *consequence* of F .
- $F \equiv G$ iff $F \models G$ and $G \models F$. We say F and G are *equivalent*.
- Let $\Gamma = \{F_1, F_2, F_3, \dots\}$ be a set of formulas.
 $\mathcal{A} \models \Gamma$ iff $\mathcal{A} \models F_i$ for each formula F_i in Γ . We say \mathcal{A} *models* Γ .
 $\Gamma \models G$ iff $\mathcal{A} \models \Gamma$ implies $\mathcal{A} \models G$ for every assignment \mathcal{A} . We say G is a *consequence* of Γ .

Proposition

- $F \models G$ iff $\models F \rightarrow G$
- $\Gamma \models G$ and Γ finite iff $\models \bigwedge \Gamma \rightarrow G$

Consistency

Let $\Gamma = \{F_1, F_2, F_3, \dots\}$ be a set of formulas.

- Γ is *consistent* or *satisfiable* iff there is an assignment that models Γ .
- We say that Γ is *inconsistent* iff it is not consistent and denote this by $\Gamma \models \perp$.

Proposition

- $\{F, \neg F\} \models \perp$
- If $\Gamma \models \perp$ and $\Gamma \subseteq \Gamma'$, then $\Gamma' \models \perp$.
- $\Gamma \models F$ iff $\Gamma, \neg F \models \perp$

Some basic equivalences

$$\begin{aligned} A \vee A &\equiv A \\ A \wedge A &\equiv A \end{aligned}$$

$$\begin{aligned} A \vee B &\equiv B \vee A \\ A \wedge B &\equiv B \wedge A \end{aligned}$$

$$A \wedge (A \vee B) \equiv A$$

$$\begin{aligned} A \wedge (B \vee C) &\equiv (A \wedge B) \vee (A \wedge C) \\ A \vee (B \wedge C) &\equiv (A \vee B) \wedge (A \vee C) \end{aligned}$$

$$\begin{aligned} \neg(A \vee B) &\equiv \neg A \wedge \neg B \\ \neg(A \wedge B) &\equiv \neg A \vee \neg B \end{aligned}$$

$$\begin{aligned} A \wedge \neg A &\equiv \perp \\ A \vee \neg A &\equiv \top \end{aligned}$$

$$\begin{aligned} A \wedge \top &\equiv A \\ A \vee \top &\equiv \top \end{aligned}$$

$$\begin{aligned} A \wedge \perp &\equiv \perp \\ A \vee \perp &\equiv A \end{aligned}$$

$$\neg\neg A \equiv A$$

$$A \rightarrow B \equiv \neg A \vee B$$

Theories

A set of formulas \mathcal{T} is *closed* under logical consequence iff for all formulas F , if $\mathcal{T} \models F$ then $F \in \mathcal{T}$.

\mathcal{T} is a *theory* iff it is closed under logical consequence. The elements of \mathcal{T} are called *theorems*.

Let Γ be a set of formulas.

$\mathcal{T}(\Gamma) = \{F \mid \Gamma \models F\}$ is called the *theory* of Γ .

The formulas of Γ are called *axioms* and the theory $\mathcal{T}(\Gamma)$ is *axiomatizable*.

Substitution

- Formula G is a *subformula* of formula F if it occurs syntactically within F .
- Formula G is a *strict subformula* of F if G is a subformula of F and $G \neq F$

Substitution theorem

Suppose $F \equiv G$. Let H be a formula that contains F as a subformula. Let H' be the formula obtained by replacing some occurrence of F in H with G . Then $H \equiv H'$.

Adquate sets of connectives for PL

There is some *redundancy* among the logical connectives.

Some smaller adquate sets of connectives for PL:

$$\{\wedge, \neg\} \quad \perp \equiv P \wedge \neg P, \quad \top \equiv \neg(P \wedge \neg P), \\ A \vee B \equiv \neg(\neg A \wedge \neg B), \quad A \rightarrow B \equiv \neg(A \wedge \neg B)$$

$$\{\vee, \neg\} \quad \top \equiv A \vee \neg A, \quad \perp \equiv \neg(A \vee \neg A), \\ A \wedge B \equiv \neg(\neg A \vee \neg B), \quad A \rightarrow B \equiv A \vee B$$

$$\{\rightarrow, \neg\} \quad \top \equiv A \rightarrow A, \quad \perp \equiv \neg(A \rightarrow A), \\ A \vee B \equiv \neg A \rightarrow B, \quad A \wedge B \equiv \neg(A \rightarrow \neg B)$$

$$\{\rightarrow, \perp\} \quad \neg A \equiv A \rightarrow \perp, \quad \top \equiv A \rightarrow A, \\ A \vee B \equiv (A \rightarrow \perp) \rightarrow B, \quad A \wedge B \equiv (A \rightarrow B \rightarrow \perp) \rightarrow \perp$$

$$A \leftrightarrow B \equiv (A \rightarrow B) \wedge (B \rightarrow A)$$

Decidability

A *decision problem* is any problem that, given certain input, asks a question to be answered with a “yes” or a “no”.

A *solution* to a decision problem is a program that takes problem instances as input and *always* terminates, producing a correct “yes” or “no” output. A decision problem is *decidable* if it has a solution.

Given formulas F and G as input, we may ask:

Decision problems

- Validity problem:* “Is F valid?”
- Satisfiability problem:* “Is F satisfiable?”
- Consequence problem:* “Is G a consequence of F ?”
- Equivalence problem:* “Are F and G equivalent?”

All these problems are *decidable*!

Decidability

Any algorithm that works for one of these problems also works for all of these problems!

F is satisfiable	iff	$\neg F$ is not valid
$F \models G$	iff	$\neg(F \rightarrow G)$ is not satisfiable
$F \equiv G$	iff	$F \models G$ and $G \models F$
F is valid	iff	$F \equiv \top$

Truth-table method

For the satisfiability problem, we first compute a truth table for F and then check to see if its truth value is ever one.

This algorithm certainly works, but is very inefficient.

Its *exponential-time*! $O(2^n)$

If F has n atomic formulas, then the truth table for F has 2^n rows.

Complexity

An algorithm is *polynomial-time* if there exists a polynomial $p(x)$ such that given input of size n , the algorithm halts in fewer than $p(n)$ steps. The class of all decision problems that can be resolved by some polynomial-time algorithm is denoted by **P** (or **PTIME**).

It is not known whether the satisfiability problem (and the other three decision problems) is in **P**.

We do not know of a polynomial-time algorithm for satisfiability.

If it exists, then **P** = **NP** !

The Satisfiability problem for PL (PSAT) is **NP**-complete (it was the first one to be shown NP-complete).

Complexity

- A *deterministic* algorithm is a step-by-step procedure. At any stage of the algorithm, the next step is completely determined.
- In contrast, a *non-deterministic* algorithm may have more than one possible “next step” at a given stage. That is, there may be more than one computation for a given input.

NP (*non-deterministic polynomial-time*) decision problems

Let PROB be an arbitrary decision problem. Given certain input, PROB produces an output of either “yes” or “no”. Let Y be the set of all inputs for which PROB produces the output of “yes” and let N be the analogous set of inputs that produce output “no”.

- If there exists a non-deterministic algorithm which, given input x , can produce the output “yes” in polynomial-time if and only if $x \in Y$, then PROB is in **NP**.
- If there exists a non-deterministic algorithm which, given input x , can produce the output “no” in polynomial-time if and only if $x \in N$, then PROB is in **coNP**.

Complexity

Essentially, a decision problem is in **NP** (**coNP**) if a “yes” (“no”) answer can be obtained in polynomial-time by guessing.

Satisfiability problem is **NP**. Given a formula F compute an assignment \mathcal{A} for F . If $\mathcal{A}(F) = 1$, then F is satisfiable.

Validity problem is **coNP**.

NP-complete

A decision problem Π is **NP-complete** if it is in **NP** and for every problem Π_1 in **NP**, Π_1 is *polynomially reducible* to Π ($\Pi_1 \propto \Pi$).


Cook's theorem (1971)

PSAT is **NP**-complete.

BDDs

- A *Binary Decision Diagram* (BDD) is a data structure that is used to represent a Boolean formulas.
- The Boolean formula is represented as a rooted, directed, acyclic graph, which consists of several decision nodes (labeled by Boolean variables) and terminal nodes (0 and 1).
- Each decision node x has two child nodes called *low child* and *high child*. The edge from node x to a low (resp. high) child represents an assignment of x to 0 (resp. 1).
- A BDD is called *ordered* if different variables appear in the same order on all paths from the root. A BDD is said to be *reduced* if the following rules have been applied to its graph:
 - Merge the terminal nodes into two nodes 0 and 1.
 - Merge isomorphic subtrees.
 - Remove redundant nodes.

From Binary Decision Trees to ROBDDs



ROBDDs [Bryant, 1986]

Reduced Ordered Binary Decision Diagrams (ROBDDs, or **BDDs** for short) are graph-based data structure for manipulating Boolean formulas.

- BDDs are *canonical* representation of Boolean formulas (if two formulas are equivalent, then their BDD representations are isomorphic) assuming the same variable ordering.
- Implications of canonicity:
 - ▶ All tautologies have the same BDD (a single node with “true”).
 - ▶ All contradictions have the same BDD (a single node with “false”).
 - ▶ To check if a formula is satisfiable check if its BDD is isomorphic to false.
- Checking for satisfiability, validity, or contradiction can be done in **constant time** for a given BDD.

However ...

BDDs

- A path from the root node to the 1-terminal (resp. 0-terminal) represents a (possibly partial) variable assignment for which the represented Boolean formula is true (resp. false).
- The process of turning a binary tree into a BDD is not very useful by itself, as one does not want to build the binary decision tree in the first place, owing to its exponential size.
- Instead, one creates the reduced ordered BDD directly: given a formula, its BDD is built recursively from the BDDs of its subexpressions.

BDDs

However,

- Building the BDD for a given formula can take **exponential space and time**.
- The size of the BDD is extremely sensitive to variable ordering.
- Computing a optimal variable ordering for a BDD is an NP-complete problem.
- It is know that the BDD representation of certain Boolean formulas is exponential in size regardless of variable order.
- In practice, the Boolean formula that BDDs can represent are often limited to several hundred variables in size.

Nonetheless, BDDs are extensively used in the verification of HW circuits, and optimization and synthesis of logic circuits.

Another usual application of BDDs is symbolic model checking.

SAT solving algorithms

- There are several techniques and algorithms for SAT solving.
- The majority of modern SAT solvers can be classified into two main categories:
 - ▶ SAT solvers based on a *stochastic local search*: the solver guesses a full assignment, and then, if the formula is evaluated to false under this assignment, starts to flip values of variables according to some (greedy) heuristic.
 - ▶ SAT solvers based on the *DPLL framework*: optimizations to the Davis-Putnam-Logemann-Loveland algorithm (DPLL) which corresponds to backtrack search through the space of possible variable assignments.
- DPLL-based SAT solvers, however, are considered better in most cases.
- Usually SAT solvers receive as input a formula in a specific syntactical format. So, one has first to transform the input formula to this specific format preserving satisfiability.

Local search

- Local search is *incomplete*; usually it cannot prove unsatisfiability.
- However, it can be very effective in specific contexts.
- The algorithm:
 - ▶ Start with a (random) assignment,
 - ▶ And repeat a number of times:
 - ★ If not all clauses satisfied, change the value of a variable.
 - ★ If all clauses satisfied, it is done.
 - ▶ Repeat (random) selection of assignment a number of times.
- The algorithm terminates when a satisfying assignment is found or when a time bound is elapsed (inconclusive answer).

Normal forms

SAT solvers usually take input in *conjunctive normal form*.

- A *literal* is a propositional variable or its negation.
 - ▶ A literal is *negative* if it is a negated atom, and *positive* otherwise.
- A formula A is in *negation normal form (NNF)*, if the only connectives used in A are \neg , \wedge and \vee , and negation only appear in literals.
- A *clause* is a disjunction of literals.
- A formula is in *conjunctive normal form (CNF)* if it is a conjunction of clauses, i.e., it has the form

$$\bigwedge_i \left(\bigvee_j l_{ij} \right)$$

where l_{ij} is the j -th literal in the i -th clause.

Normalization

Transforming a formula F to equivalent formula F' in NNF can be computed by repeatedly replace any subformula that is an instance of the left-hand-side of one of the following equivalences by the corresponding right-hand-side

$$\begin{array}{ll} A \rightarrow B \equiv \neg A \vee B & \neg \neg A \equiv A \\ \neg(A \wedge B) \equiv \neg A \vee \neg B & \neg(A \vee B) \equiv \neg A \wedge \neg B \end{array}$$

This algorithm is *linear on the size of the formula*.

Normalization

To transform a formula already in NNF into an equivalent CNF, apply recursively the following equivalences (left-to-right):

$$\begin{aligned} A \vee (B \wedge C) &\equiv (A \vee B) \wedge (A \vee C) & (A \wedge B) \vee C &\equiv (A \vee C) \wedge (B \vee C) \\ A \wedge \perp &\equiv \perp & \perp \wedge A &\equiv \perp & A \wedge \top &\equiv A & \top \wedge A &\equiv A \\ A \vee \perp &\equiv A & \perp \vee A &\equiv A & A \vee \top &\equiv \top & \top \vee A &\equiv \top \end{aligned}$$

This algorithm converts a NNF formula into an equivalent CNF, but its worst case is **exponential on the size of the formula**.

Example

Compute the CNF of $((P \rightarrow Q) \rightarrow P) \rightarrow P$

The first step is to compute its NNF by transforming implications into disjunctions and pushing negations to proposition symbols:

$$\begin{aligned} ((P \rightarrow Q) \rightarrow P) \rightarrow P &\equiv \neg((P \rightarrow Q) \rightarrow P) \vee P \\ &\equiv \neg(\neg(P \rightarrow Q) \vee P) \vee P \\ &\equiv \neg(\neg(\neg P \vee Q) \vee P) \vee P \\ &\equiv \neg((P \wedge \neg Q) \vee P) \vee P \\ &\equiv (\neg(P \wedge \neg Q) \wedge \neg P) \vee P \\ &\equiv ((\neg P \vee Q) \wedge \neg P) \vee P \end{aligned}$$

To reach a CNF, distributivity is then applied to pull the conjunction outside:

$$((\neg P \vee Q) \wedge \neg P) \vee P \equiv (\neg P \vee Q \vee P) \wedge (\neg P \vee P)$$

Worst-case example

Compute the CNF of $(P_1 \wedge Q_1) \vee (P_2 \wedge Q_2) \vee \dots \vee (P_n \wedge Q_n)$

$$\begin{aligned} &(P_1 \wedge Q_1) \vee (P_2 \wedge Q_2) \vee \dots \vee (P_n \wedge Q_n) \\ \equiv &(P_1 \vee (P_2 \wedge Q_2) \vee \dots \vee (P_n \wedge Q_n)) \wedge (Q_1 \vee (P_2 \wedge Q_2) \vee \dots \vee (P_n \wedge Q_n)) \\ \equiv &\dots \\ \equiv &(P_1 \vee \dots \vee P_n) \wedge \\ &(P_1 \vee \dots \vee P_{n-1} \vee Q_n) \wedge \\ &(P_1 \vee \dots \vee P_{n-2} \vee Q_{n-1} \vee P_n) \wedge \\ &(P_1 \vee \dots \vee P_{n-2} \vee Q_{n-1} \vee Q_n) \wedge \\ &\dots \wedge \\ &(Q_1 \vee \dots \vee Q_n) \end{aligned}$$

The original formula has $2n$ literals, while the equivalent CNF has 2^n clauses, each with n literals.

The size of the formula increases exponentially.

Definitional CNF

Equisatisfiability

Two formulas F and F' are *equisatisfiable* when F is satisfiable iff F' is satisfiable.

Any propositional formula can be transformed into a equisatisfiable CNF formula with only **linear** increase in the size of the formula.

The price to be paid is n new Boolean variables, where n is the number of logical connectives in the formula.

This transformation can be done via *Tseitin's encoding* [Tseitin, 1968].

This transformation compute what is called the *definitional CNF* of a formula, because they rely on the introduction of new proposition symbols that act as names for subformulas of the original formula.

Tseitin's encoding

Tseitin transformation

- 1 Introduce a new fresh variable for each compound subformula.
- 2 Assign new variable to each subformula.
- 3 Encode local constraints as CNF.
- 4 Make conjunction of local constraints and the root variable.

- This transformation produces a formula that is **equisatisfiable**: the result is satisfiable iff and only the original formula is satisfiable.
- One can get a satisfying assignment for original formula by projecting the satisfying assignment onto the original variables.

There are various optimizations that can be performed in order to reduce the size of the resulting formula and the number of additional variables.

Navigation icons

Tseitin's encoding: an example

Encode $P \rightarrow Q \wedge R$

$$\begin{array}{c} A_1 \\ \overbrace{P \rightarrow Q \wedge R} \\ A_2 \end{array}$$

- 2 We need to satisfy A_1 together with the following equivalences

$$A_1 \leftrightarrow (P \rightarrow A_2)$$


$$A_2 \leftrightarrow (Q \wedge R)$$

- 3 These equivalences can be rewritten in CNF as $(A_1 \vee P) \wedge (A_1 \vee \neg A_2) \wedge (\neg A_1 \vee \neg P \vee A_2)$ and $(\neg A_2 \vee Q) \wedge (\neg A_2 \vee R) \wedge (A_2 \vee \neg Q \vee \neg R)$, respectively.
- 4 The CNF which is equisatisfiable with $P \rightarrow (Q \wedge R)$ is

$$\begin{aligned} A_1 \wedge & (A_1 \vee P) \wedge (A_1 \vee \neg A_2) \wedge (\neg A_1 \vee \neg P \vee A_2) \\ & \wedge (\neg A_2 \vee Q) \wedge (\neg A_2 \vee R) \wedge (A_2 \vee \neg Q \vee \neg R) \end{aligned}$$

Navigation icons

Tseitin's encoding: a circuit to CNF



$$\begin{aligned} & o \wedge \\ & (x \leftrightarrow a \wedge c) \wedge \\ & (y \leftrightarrow b \vee x) \wedge \\ & (u \leftrightarrow a \vee b) \wedge \\ & (v \leftrightarrow b \vee c) \wedge \\ & (w \leftrightarrow u \wedge v) \wedge \\ & (o \leftrightarrow y \oplus w) \end{aligned}$$

$$o \wedge (x \rightarrow a) \wedge (x \rightarrow c) \wedge (x \leftarrow a \wedge c) \wedge \dots$$

$$o \wedge (\bar{x} \vee a) \wedge (\bar{x} \vee c) \wedge (x \vee \bar{a} \vee \bar{c}) \wedge \dots$$

Navigation icons

CNFs validity

- The strict shape of CNFs make them particularly suited for checking validity problems.
 - A CNF is a tautology iff all of its clauses are tautologies.
 - A clause C is a tautology precisely when there exists a proposition symbol P such that both P and $\neg P$ are in C (such clauses are said to be **closed**).
 - So, a CNF is a tautology iff all of its clauses are closed.
- However, the applicability of this simple criterion for validity is compromised by the potential exponential growth in the CNF transformation.

- This limitation is overcome considering instead SAT, with satisfiability preserving CNFs (definitional CNF). Recall that

$$F \text{ is valid} \quad \text{iff} \quad \neg F \text{ is unsatisfiable}$$

Navigation icons

CNFs satisfiability

- A CNF is *satisfied* by an assignment if all its clauses are satisfied. And a *clause is satisfied* if at least one of its literals is satisfied.
- The idea is to incrementally construct an assignment compatible with a CNF.
 - ▶ An *assignment* of a formula F is a function mapping F 's variables to 1 or 0. We say it is
 - ★ *full* if all of F 's variables are assigned,
 - ★ and *partial* otherwise.
- Most current state-of-the-art SAT solvers are based on the *Davis-Putnam-Logemann-Loveland (DPLL)* framework: in this framework the tool can be thought of as traversing and backtracking on a binary tree, in which
 - ▶ internal nodes represent *partial assignments*
 - ▶ and leaves represent *full assignments*

Basic concepts

Pure literals

- A literal is *pure* if only occurs as a positive literal or as a negative literal in a CNF formula.
- *Pure literal rule*
Clauses containing pure literals can be removed from the formula (i.e. just assign pure literals to the values that satisfy the clauses).
- This technique was extensively used until the mid 90s, but nowadays seldom used.

Example

Let F be $(Q \vee P) \wedge (R \vee Q \vee \neg P) \wedge (\neg R \vee P) \wedge (P \vee \neg X)$.

- Q and $\neg X$ are pure literals in F .
- The resulting (equisatisfiable) formula is $(\neg R \vee P)$.

Basic concepts

Unit propagation (also called Boolean Constraint Propagation (BCP))

- A clause is *unit* if all literals but one are assigned value 0, and the remaining literal is unassigned.
- *Unit clause rule*
Given a unit clause, its only unassigned literal *must* be assigned value 1 for the clause to be satisfied.
- *Unit propagation* is the iterated application of the unit clause rule.
- This technique is extensively used.

- Consider the partial assignment $P = 0, Q = 1$. Under this assignment $(P \vee \neg R \vee \neg Q)$ is a unit clause. R must be assigned the value 0.
- Consider the partial assignment $R = 1, Q = 1$.
 - ▶ By unit propagation we can conclude that $(P \vee \neg R \vee \neg Q) \wedge (\neg P \vee \neg Q \vee X) \wedge (\neg P \vee \neg R \vee X)$ is satisfiable.
 - ▶ What about $(P \vee \neg R \vee \neg Q) \wedge (\neg P \vee \neg Q \vee X) \wedge (\neg P \vee \neg R \vee \neg X)$?

Basic concepts

Resolution

- *Resolution rule*
If a formula F contains clauses $(A \vee P)$ and $(\neg P \vee B)$, then one can infer $(A \vee B)$. The formula $(A \vee B)$ is called the *resolvent*.
- The resolvent can be added as a conjunction to F to produce an equivalent formula still in CNF.
- If even \perp is deduced via resolution, F must be unsatisfiable.
- A CNF formula that does not contain \perp and to which no more resolutions can be applied represents all possible satisfying interpretations.

What happens if we apply resolution between $(\neg Q \vee R \vee P)$ and $(Q \vee \neg R \vee X)$?

Historical perspective

- The DP algorithm [Davis&Putnam, 1960]
 - ▶ Based on the resolution rule.
 - ▶ Eliminate one variable at each step, using resolution.
 - ▶ Applied the pure literal rule and unit propagation.
- The DPLL algorithm [Davis&Putnam&Logemann&Loveland, 1962]
 - ▶ Based on backtrack search.
 - ▶ Progresses by making a decision about a variable and its value.
 - ▶ Propagates implications of this decision that are easy to detect.
 - ▶ Backtracks in case a conflict is detected in the form of a falsified clause.
- In the last two decades, several enhancements have been introduced to improve the efficiency of DPLL-based SAT solving.

DP algorithm

- Iteratively apply the following steps:
 - ▶ Select variable X .
 - ▶ Apply resolution rule between every pair of clauses of the form $(X \vee A)$ and $(\neg X \vee B)$.
 - ▶ Remove all clauses containing either X or $\neg X$.
 - ▶ Apply the pure literal rule and unit propagation.
- Terminate when either the empty clause or the empty formula is derived.
- The algorithm is complete but **inefficient**.

DPLL algorithm

- Traditionally the DPLL algorithm is presented as a recursive procedure.
- The procedure DPLL is called with the CNF and a partial assignment.
- We will represent a CNF by a set of sets of literals.
- We will represent the partial assignment by a set of literals (P denote that P is set to 1, and $\neg P$ that P is set to 0).
- The algorithm:
 - ▶ Progresses by making a decision about a variable and its value.
 - ▶ Propagates implications of this decision that are easy to detect, simplifying the clauses.
 - ▶ Backtracks in case a conflict is detected in the form of a falsified clause.

CNFs (as sets of sets of literals)

- Recall that CNFs are formulas with the following shape (each l_{ij} denotes a literal):

$$(l_{11} \vee l_{12} \vee \dots \vee l_{1k}) \wedge \dots \wedge (l_{n1} \vee l_{n2} \vee \dots \vee l_{nj})$$

- Associativity, commutativity and idempotence of both disjunction and conjunction allow us to treat each CNF as **a set of sets of literals** S

$$S = \{\{l_{11}, l_{12}, \dots, l_{1k}\}, \dots, \{l_{n1}, l_{n2}, \dots, l_{nj}\}\}$$

- An empty inner set will be identified with \perp , and an empty outer set with \top . Therefore,
 - ▶ if $\{\} \in S$, then S is equivalent to \perp ;
 - ▶ if $S = \{\}$, then S is \top .

Simplification of a clause under an assignment

If we fix the assignment of a particular proposition symbol, we are able to simplify the corresponding CNF accordingly.

The *opposite* of a literal l , written $\neg l$, is defined by

$$\neg l = \begin{cases} \neg P & , \text{if } l = P \\ P & , \text{if } l = \neg P \end{cases}$$

When we set a literal l to be true,

- any clause that has the literal l is now guaranteed to be satisfied, so we throw it away for the next part of the search.
- any clause that had the literal $\neg l$, on the other hand, must rely on one of the other literals in the clause, hence we throw out the literal $\neg l$ before going forward.

Simplification of S assuming l holds

$$S|_l = \{c \setminus \{\neg l\} \mid c \in S \text{ and } l \notin c\}$$

Simplification of a clause under an assignment

If a CNF S contains a clause that consists of a single literal (called *unit clause*), we know for certain that the literal must be set to true and S can be simplified.

One should apply this rule while it is possible and worthwhile.

```
UNIT_PROPAGATE( $S, \mathcal{A}$ ) {
  while  $\{\} \notin S$  and  $S$  has a unit clause  $l$  do {
     $S \leftarrow S|_l$ ;
     $\mathcal{A} \leftarrow \mathcal{A} \cup \{l\}$ 
  }
}
```

DPLL algorithm

DPLL is called with a CNF S and a partial assignment \mathcal{A} (initially \emptyset).

```
DPLL( $S, \mathcal{A}$ ) {
  UNIT_PROPAGATE( $S, \mathcal{A}$ );
  if  $S = \{\}$  then return SAT;
  else if  $\{\} \in S$  then return UNSAT;
  else {  $l \leftarrow$  a literal of  $S$ ;
    if DPLL( $S|_l, \mathcal{A} \cup \{l\}$ ) = SAT then return SAT;
    else return DPLL( $S|_{\neg l}, \mathcal{A} \cup \{\neg l\}$ )
  }
}
```

- DPLL complete algorithm for SAT.
- Unsatisfiability of the complete formula can only be detected after exhaustive search.

DPLL algorithm

Is $(\neg P \vee Q) \wedge (\neg P \vee R) \wedge (Q \vee R) \wedge (\neg Q \vee \neg R) \wedge (P \vee \neg R \vee Q)$ satisfiable?

	S	\mathcal{A}
DPLL	$\{\{\neg P, Q\}, \{\neg P, R\}, \{Q, R\}, \{\neg Q, \neg R\}, \{P, \neg R, Q\}\}$	\emptyset
UNIT_PROPAGATE	$\{\{\neg P, Q\}, \{\neg P, R\}, \{Q, R\}, \{\neg Q, \neg R\}, \{P, \neg R, Q\}\}$	\emptyset
choose $l = P$		
DPLL	$S _l = \{\{Q\}, \{R\}, \{Q, R\}, \{\neg Q, \neg R\}\}$	$\{P\}$
UNIT_PROPAGATE	$\{\{Q\}, \{R\}, \{Q, R\}, \{\neg Q, \neg R\}\}$	$\{P\}$
$\neg l = \neg P$		
DPLL	$S _{\neg l} = \{\{Q, R\}, \{\neg Q, \neg R\}, \{\neg R, Q\}\}$	$\{\neg P\}$
UNIT_PROPAGATE	$\{\{Q, R\}, \{\neg Q, \neg R\}, \{\neg R, Q\}\}$	$\{\neg P\}$
choose $l = Q$		
DPLL	$S _l = \{\{\neg R\}\}$	$\{\neg P, Q\}$
UNIT_PROPAGATE	$\{\{\neg R\}\}$	$\{\neg P, Q\}$
	$\{\}$	$\{\neg P, Q, \neg R\}$

DPLL recursive algorithm

```
DPLL(formula, assignment) {
  if (deduce(formula, assignment) == SAT)
    return SAT;
  else if (deduce(formula, assignment) == CONF)
    return CONF;
  else {
    v = new_variable(formula, assignment);
    a = new_assignment(formula, assignment, v, 0);
    if (DPLL(formula, a) == SAT)
      return SAT;
    else {
      a = new_assignment(formula, assignment, v, 1);
      return DPLL(formula, a);
    }
  }
}
```

DPLL-based iterative algorithm [Marques-Silva&Sakallah,1996]

At each step:

- **Decide** on the assignment of a variable (which is called the *decision variable*, and it will have a *decision level* associated with it).
- **Deduce** the consequences of the decision made. (Variables assigned will have the same decision level as the decision variable.)
 - ▶ If all the clauses are satisfied, then the instance is satisfiable.
 - ▶ If there exists a conflicting clause, then **analyze** the conflict and determine the decision level to backtrack. (The solver may perform some analysis and record some information from the current conflict in order to prune the search space for the future.)
 - ★ Decision level ≤ 0 indicates that the formula is unsatisfiable.
 - ▶ Otherwise, proceed with another decision.

Different DPLL-based modern solvers differ mainly in the detailed implementation of each of these functions.

DPLL-based iterative algorithm

```
while(1) {
  decide_next_branch(); //branching heuristics
  while (true) {
    status = deduce(); //deduction mechanism
    if (status == CONF) {
      bl = analyze_conflict(); //conflict analysis
      if (bl == 0) return UNSAT;
      else backtrack(bl);
    }
    else if (status == SATISFIABLE) return SAT;
    else break;
  }
}
```

DPLL framework: heuristics & optimizations

Many different techniques are applied to achieve efficiency in DPLL-based SAT solvers.

- **Look-ahead**: exploit information about the remaining search space.
 - ▶ unit propagation
 - ▶ pure literal rule
 - ▶ decision (splitting) heuristics
- **Look-back**: exploit information about search which has already taken place.
 - ▶ non-chronological backtracking (a.k.a. *backjumping*)
 - ▶ clause learning
- **Other techniques**:
 - ▶ preprocessing (detection of subsumed clauses, simplification, ...)
 - ▶ (random) restart (restarting the solver when it seems to be in a hopeless branch of the search tree)
- ...

Decision heuristics

Probably the most important element in SAT solving is the strategy by which the literals are chosen. This strategy is called the *decision heuristic* of the SAT solver.

- *MOMS heuristics*
Pick the literal occurring most often in the minimal size clauses.
- *Jeroslow-Wang*
Selects literals that appear frequently in short clauses.
- *DLIS: Dynamic Large Individual Sum*
Selects the literal that appears most frequently in unresolved clauses. (Introduced in GRASP)
- *VSIDS: Variable State Independent Decaying Sum*
Similar to DLIS. (Introduced in Chaff)
- *Berkmin method*
- ...

Conflict analysis and learning

- *Non-chronological backtracking*: does not necessarily flip the last assignment and can backtrack to a earlier decision level.
- The process of adding conflict clauses is generally referred to as *learning*.
- The conflict clauses record the reasons deduced from the conflict to avoid making the same mistake in the future search. For that *implication graphs* are used.
- *Conflict-driven backtracking* uses the conflict clauses learned to determine the actual reasons for the conflict and the decision level to backtrack in order to prevent the repetition of the same conflict.

Conflict-Driven Clause Learning (CDCL) solvers

- DPLL framework.
- New clauses are *learned* from conflicts.
- Structure (*implication graphs*) of conflicts exploited.
- Backtracking can be *non-chronological*.
- Efficient *data structures* (compact and reduced maintenance overhead).
- Backtrack search is periodically *restarted*.
- Can deal with hundreds of thousand variables and tens of million clauses!

Modern SAT solvers

- In the last two decades, satisfiability procedures have undergone dramatic improvements in efficiency and expressiveness. Breakthrough systems like *GRASP* (1996), *SATO* (1997), *Chaff* (2001) and *MiniSAT* (2003) have introduced several enhancements to the efficiency of DPLL-based SAT solving.
- Modern SAT solvers can check *formulas with hundreds of thousands variables and millions of clauses* in a reasonable amount of time.
- New SAT solvers are introduced every year.
 - ▶ The satisfiability library *SATLIB*¹ is an online resource that proposes, as a standard, a unified notation and a collection of benchmarks for performance evaluation and comparison of tools.
 - ▶ Such a uniform test-bed has been serving as a framework for regular tool competitions organised in the context of the regular SAT conferences.²

¹<http://www.satlib.org>

²<http://www.satcompetition.org>

DIMACS CNF format

- DIMACS CNF format is a standard format for CNF used by most SAT solvers.
- Plain text file with following structure:

```
c <comments>
...
p cnf <num.of variables> <num.of clauses>
<clause> 0
<clause> 0
...
```
- Every number 1, 2, . . . corresponds to a variable (variable names have to be mapped to a variable).
- A negative number denote the negation of the corresponding variable.
- Every clause is a list of numbers, separated by spaces. (One or more lines per clause).

DIMACS CNF format

Example

$$A_1 \wedge (A_1 \vee P) \wedge (\neg A_1 \vee \neg P \vee A_2) \wedge (A_1 \vee \neg A_2)$$

- We have 3 variables and 4 clauses.
- CNF file:

```
p cnf 3 4
1 0
1 3 0
-1 -3 2 0
1 -2 0
```

Applications of SAT

- A large number of problems can be described in terms of satisfiability, including graph problems, planning, games, scheduling, software and hardware verification, extended static checking, optimization, test-case generation, among others.
- These problems can be encoded by propositional formulas and solved using SAT solvers.

problem $\mathcal{P} \rightsquigarrow$ formula $F \longrightarrow$ CNF converter \longrightarrow SAT solver

SAT solver output: If F is satisfiable: **sat** + model
 If F is unsatisfiable: **unsat** + proof

The satisfying assignments (models) of F are the solutions of \mathcal{P} .

- SAT solvers are core engines for other solvers (like SMT solvers).
- SAT solver may be integrated into theorem provers.

Modeling with PL

When can the meeting take place?

- Maria cannot meet on Wednesday.
- Peter can only meet either on Monday, Wednesday or Thursday.
- Anne cannot meet on Friday.
- Mike cannot meet neither on Tuesday nor on Thursday.

Encode into the following proposition:

$$\neg \text{Wed} \wedge (\text{Mon} \vee \text{Wed} \vee \text{Thu}) \wedge \neg \text{Fri} \wedge (\neg \text{Tue} \wedge \neg \text{Thu})$$

Modeling with PL

Graph coloring

Can one assign one of K colors to each of the vertices of graph $G = (V, E)$ such that adjacent vertices are assigned different colors?

- Create $|V| \times K$ variables: $x_{ij} = 1$ iff vertex i is assigned color j ; 0 otherwise.
- For each edge (u, v) , require different assigned colors to u and v :
for each $1 \leq j \leq K$, $(x_{uj} \rightarrow \neg x_{vj})$

- Each vertex is assigned exactly one color.

► At least one color to each vertex:

$$\text{for each } 1 \leq i \leq |V|, \bigvee_{j=1}^K x_{ij}$$

► At most one color to each vertex:

$$\text{for each } 1 \leq i \leq |V|, \bigwedge_{a=1}^{K-1} (x_{ia} \rightarrow \bigwedge_{b=a+1}^K \neg x_{ib})$$

Modeling with PL

At least, at most, exactly one...

How to represent in CNF the following constraints

- At least one: $\sum_{j=1}^N x_j \geq 1$?

Standard solution:

$$\bigvee_{j=1}^N x_j$$

- At most one: $\sum_{j=1}^N x_j \leq 1$?

Naive solution:

$$\bigwedge_{a=1}^{N-1} \bigwedge_{b=a+1}^N (\neg x_a \vee \neg x_b)$$

More compact solutions are possible.

- Exactly one: $\sum_{j=1}^N x_j = 1$?

Standard solution: at least 1 and at most 1 constraints.

Modeling with PL

Placement of guests

We have three chairs in a row and we need to place Anne, Susan and Peter.

- Anne does not want to sit near Peter.
- Anne does not want to sit in the left chair.
- Susan does not want to sit to the right of Peter.

Can we satisfy these constraints?

- Denote: Anne = 1, Susan = 2, Peter = 3
left chair = 1, middle chair = 2, right chair = 3
- Introduce a propositional variable for each pair (*person, place*)
- $x_{ij} = 1$ iff person i is sited in place j ; 0 otherwise

Modeling with PL

Placement of guests (cont.)

- Anne does not want to sit near Peter.
 $((x_{11} \vee x_{13}) \rightarrow \neg x_{32}) \wedge (x_{12} \rightarrow (\neg x_{31} \wedge \neg x_{33}))$
- Anne does not want to sit in the left chair. $\neg x_{11}$
- Susan does not want to sit to the right of Peter.
 $(x_{31} \rightarrow \neg x_{22}) \wedge (x_{32} \rightarrow \neg x_{23})$

- Each person is placed.

$$\bigwedge_{i=1}^3 \bigvee_{j=1}^3 x_{ij}$$

- No more than one person per chair.

$$\bigwedge_{i=1}^3 \bigwedge_{a=1}^2 \bigwedge_{b=a+1}^3 (\neg x_{ia} \vee \neg x_{ib})$$

Modeling with PL

Equivalence checking of if-then-else chains

Original C code

```
if(!a && !b) h();  
else if(!a) g();  
else f();
```

Optimized C code

```
if(a) f();  
else if(b) g();  
else h();
```

Are these two programs equivalent?

- 1 Model the variables a and b and the procedures that are called using the Boolean variables a , b , f , g , and h .
- 2 Compile if-then-else chains into Boolean formulae
 $\text{compile}(\text{if } x \text{ then } y \text{ else } z) \equiv (x \wedge y) \vee (\neg x \wedge z)$
- 3 Check the validity of the following formula
 $\text{compile}(\text{original}) \leftrightarrow \text{compile}(\text{optimized})$
Reformulate it as a SAT problem: Is the Boolean formula
 $\neg (\text{compile}(\text{original}) \leftrightarrow \text{compile}(\text{optimized}))$
satisfiable?

Proof system

- So far we have taken the “*semantic*” approach to logic, with the aim of characterising the semantic concept of model, from which validity, satisfiability and semantic entailment were derived.
- However, this is not the only possible point of view.
- Instead of adopting the view based on the notion of truth, we can think of logic as a codification of reasoning. This alternative approach to logic, called “*deductive*”, focuses directly on the deduction relation that is induced on formulas, i.e., on what formulas are logical consequences of other formulas.
- We will explore this perspective later in this course.

Exercises

- Solve the following logic puzzle.
 - ▶ If the unicorn is mythical, then it is immortal.
 - ▶ If the unicorn is not mythical, then it is a mortal mammal.
 - ▶ If the unicorn is either immortal or a mammal, then it is horned.
 - ▶ The unicorn is magical if it is the horned.Is the unicorn mythical? Is it magical? Is it horned?
- Encode into SAT a N-queen puzzle.
Place N queens on a $N \times N$ board, such that no two queens attack each other.

Exercises

- Encode into SAT a Sudoku puzzle.
 - ▶ 9×9 square divided into 9 sub-squares
 - ▶ General rules:
 - ★ Values 1-9, one value per cell
 - ★ No duplicates in rows
 - ★ No duplicates in columns
 - ★ No duplicates in sub-squares
 - ▶ A particular instance of the Sudoku puzzle has some known initial values.
- Use a SAT solver to show that the following two if-then-else expressions are equivalent.
 $!(a||b) ? h : !(a==b) ? f : g$
 $!(!a||!b) ? g : (!a\&\&!b) ? h : f$

Exercises

- Convert into an equivalent CNF the following formulas.
 - ▶ $A \vee (A \rightarrow B) \rightarrow A \vee \neg B$
 - ▶ $(A \rightarrow B \vee C) \wedge \neg(A \wedge \neg B \rightarrow C)$
 - ▶ $(\neg A \rightarrow \neg B) \rightarrow (\neg A \rightarrow B) \rightarrow A$
- Convert $P \wedge Q \vee (R \wedge P)$ into a equisatisfiable formula in CNF by using the Tseitin transformation.
- Run by hand the DPLL procedure to decide about the satisfiability of the formulas above.

Exercises

- Pick up a SAT solver.
 - Play with simple examples.
 - Use the SAT solver to test if each of the following formulas is satisfiable, valid, refutable or a contradiction.
 - ▶ $A \vee (A \rightarrow B) \rightarrow A \vee \neg B$
 - ▶ $(A \rightarrow B \vee C) \wedge \neg(A \wedge \neg B \rightarrow C)$
 - ▶ $(\neg A \rightarrow \neg B) \rightarrow (\neg A \rightarrow B) \rightarrow A$
- Note that CNF equivalents of these formulas where already calculated.
- Search the web for “SAT benchmarks” and experiment.