Lecture Notes on Propositional Encodings

Ruben Martins

Carnegie Mellon University
Lecture 13
Thursday, March 14, 2024

1 Introduction

In the last lecture, we talked about resolution and the input of SAT solvers. As we have seen, SAT solvers take as input a formula in Conjunctive Normal Form (CNF). To tranform a formula to CNF, we can use De Morgan's laws and distribituve law but this may lead to an exponential blow-up as we will see in the next section. In this lecture, we will learn how to transform any formula into CNF without having an exponential explosion on the size of the formula.

Learning goals. After this lecture, you should be able to:

- Use the **Tseitin encoding** to convert any propositional formula into an **equisatisfiable** CNF formula with a linear increase in the size of formula.

- Represent finite domais using a **unary** or **binary** representation.

- Encode linear such as **at-most-one constraints** in CNF.

- Encode problems into CNF and use a SAT solver to solve them.

2 Tseitin Encoding

Given a propositional formula, one can use De Morgan's laws and distributive law to convert it to CNF. However, in some cases, converting a formula to CNF can have an exponential explosion on the size of the formula.

Suppose we have the following formula φ,

$$\varphi = (x_1 \wedge y_1) \vee (x_2 \wedge y_2) \vee \ldots \vee (x_n \wedge y_n)$$

and want to convert φ to CNF. If we apply De Morgan's laws and distribute law then we will obtain a formula φ' such that:

$$\varphi' = (x_1 \vee x_2 \vee \ldots \vee x_n) \wedge (y_1 \vee x_2 \vee \ldots \vee x_n) \wedge (y_1 \vee y_2 \vee \ldots \vee y_n)$$

.

Note that φ' has an exponential number of clauses, namely 2^n clauses. Can we avoid this exponential blowup on the size of the formula? Yes, with the Tseitin encoding we can transform any propositional formula into an *equisatisfiable* CNF formula.

Definition 1 (Equisatisfiable). Two formulas φ and ϕ are *equisatisfiable* if φ is satisfiable iff ϕ is satisfiable.

Note that equisatisfiability is weaker than equivalence but useful if all we want to do is to determine the satisfiability of a formula.

The key idea behind the *Tseitin Encoding* is to introduce fresh variables to encode subformulas and to encode the meaning of these fresh variables with clauses. This procedure avoids duplicating whole subformulas and can transform a propositional formula into CNF with a linear increase in the size of the formula.

Example 2. Consider the formula $\phi = (x \wedge \neg y) \vee (z \vee (x \wedge \neg w))$. This formula can be viewed as a tree as depicted in Figure 1. The terminal nodes denote the atoms of the formula and the intermediate nodes denote fresh variables that encode each subformula.


Figure 1: Tree representation of a propositional formula

For each fresh variable f, a, b, c, we introduce clauses that represent their equivalence with the respective subformula. In particular, we add the following clauses:

- $f \leftrightarrow (a \vee b) \equiv (\neg f \vee a \vee b) \wedge (\neg a \vee f) \wedge (\neg b \vee f)$

- $a \leftrightarrow (x \wedge \neg y) \equiv (\neg a \vee x) \wedge (\neg a \vee \neg y) \wedge (\neg x \vee y \vee a)$

- $b \leftrightarrow (z \vee c) \equiv (\neg b \vee z \vee c) \wedge (\neg z \vee b) \wedge (\neg c \vee b)$

- $c \leftrightarrow (x \wedge \neg w) \equiv (\neg a \vee x) \wedge (\neg a \vee \neg w) \wedge (\neg x \vee w \vee a)$

Since we want the formula to hold, we additionally need to add the unit clause (f). Note that by adding this unit clause, unit propagation would simplify the first three clauses to $(a \vee b)$.

Let's take a closer look at the previous formula $\varphi = (x_1 \wedge y_1) \vee (x_2 \wedge y_2) \vee \ldots \vee (x_n \wedge y_n)$. Recall that this formula would require an exponential number of clauses if we would use De Morgan's laws and distribute law. If instead, we use the Tseitin Encoding we can have an equisatisfiable formula φ'' in CNF composed by the following clauses:

- $w_1 \leftrightarrow (x_1 \wedge y_1) \equiv (\neg w_1 \vee x_1) \wedge (\neg w_1 \vee y_1) \wedge (w_1 \vee \neg x_1 \vee \neg y_1)$

- $\ldots$

- $w_n \leftrightarrow (x_n \wedge y_n) \equiv (\neg w_n \vee x_n) \wedge (\neg w_n \vee y_n) \wedge (w_n \vee \neg x_n \vee \neg y_n)$

- $(w_1 \vee w_2 \vee \ldots \vee w_n)$

This would result in a formula φ'' with $3n + 1$ clauses and with n auxiliary variables.

3 Finite Domains

Many real-world problems require the encoding of finite domains to propositional logic. In this section, we will present two different ways of encoding integer domains in propositional logic by using *unary* and *binary* representations of these finite domains. The intuition behind these representations is that an *unary* representation considers a Boolean variable for each possible value, while a *binary* representation considers the binary representation of an integer.

Example 3. Suppose we want to encode the domain of an integer variable $\mathcal{X} = \{1, 2, 3\}$.

Unary representation

Consider the auxiliary variables x_1, x_2, x_3. We want to encode the meaning that x_i is *true* iff $X = i$. To encode this property we need to encode that:

1. At least one of these variables must occur:
   $(x_1 \vee x_2 \vee x_3)$

2. At most one of these variables must occur:
   $(\neg x_1 \vee \neg x_2) \wedge (\neg x_1 \vee \neg x_3) \wedge (\neg x_2 \vee \neg x_3)$

Binary representation

Consider the binary representation of integers and the auxiliary variables b_1, b_0. We want to encode the following property:

- If $X = 1$ then $b_0 = 0 \wedge b_1 = 0$

- If $X = 2$ then $b_0 = 1 \wedge b_1 = 0$

- If $X = 3$ then $b_0 = 0 \wedge b_1 = 1$

In this case, the meaning of each variable can be used to implicitly encode the possible values of X. The only information we need to encode is possible integer values that are *not part of the domain* of X. In this case, $X = 4$ is not part of the domain but can be encoded using these two variables, therefore we need to disallow this value from occurring by adding the clause $(\neg b_0 \vee \neg b_1)$.

Unary vs. Binary representations

The main advantage of the binary representation is that only requires a logarithmic number of auxiliary variables to encode the finite domain. In contrast, we need a linear number of auxiliary variables for the unary encoding. However, when encoding problems using a binary encoding, it can be cumbersome to express constraints that relate to different numbers since each number is represented by a conjunction of variables instead of a single variable. Moreover, unit propagation is able to infer more information when using a unary encoding than when using binary encoding. In practice, the size of the domain is usually the decider between choosing one or other encoding. For small domains, unary encoding is usually preferred while for large domains the binary encoding is usually the best choice.

4 Linear Constraints

Linear constraints are common when encoding many real-world properties and impose a limit on a subset of variables that can be assigned to true. The most common linear constraint is the *at-most-one* constraint.

Definition 4 (At-most-one constraints). Let $x_1 + \ldots + x_n \leq 1$ be a linear constraint over Boolean variables $x_1, \ldots, x_n$. This constraint is called *at-most-one* and denotes that at most one variable x_i can be assigned to *true*.

A naive encoding for at-most-one constraints is the one that we implicitly used in the previous sections. For each pairwise combination x_i and x_j with $1 \leq i < j \leq n$, we add the clause $(\neg x_i \vee \neg x_j)$ that encodes that if x_i is assigned to *true* then x_j is assigned to *false*. Similarly, it also encodes that if x_j is assigned to *true* then x_i is assigned to *false*. This encoding is simple but requires a quadratic number of clauses. If n is very large, this encoding will lead to a very large CNF formula and should be avoided.

How can we encode linear constraints such as the at-most-one constraint using a linear number of clauses? When encoding a problem into SAT, there is usually a trade-off between reducing the number of clauses and increasing the number of variables. In this case, it is possible to use auxiliary variables in order to use only a linear number of clauses.

Sequential encoding

Consider the constraint $x_1 + x_2 + x_3 \leq 1$. Consider additionally, the auxiliary variables S_1, S_2 with the following meaning: S_i is assigned to *true* iff the sum up to x_i is exactly one. Using these auxiliary variables we can now encode this kind of constraint as follows:

- If S_i is set to *true* then S_{i+1} is also set to *true*:
  $(\neg S_1 \vee S_2)$

- If x_i is set to *true* then S_i is also set to *true*:
  $(\neg x_1 \vee S_1) \wedge (\neg x_2 \vee S_2)$

- If S_i is set to *true* then x_{i+1} is set to *false* since the sum is already 1:
  $(\neg S_1 \vee \neg x_2) \wedge (\neg S_2 \vee \neg x_3)$

In general, this encoding will require $3n - 4$ clauses which are much fewer clauses than the naive encoding which requires a quadratic number of clauses. Note that this reduction is achieved at the cost of $n - 1$ auxiliary variables. This reasoning can be further generalized to *at-most-k* constraints and many CNF encodings exist for these kinds of constraints. For the generalization of the sequential encoding to at-most-k constraints, we refer the interested student to the literature [Sin05].

5 Encoding Graph Coloring as a SAT problem

Suppose that we want to encode the graph coloring problem to SAT, i.e. we want to ask the question, given a graph if there exists a k-coloring such that no two nodes that are connected have the same color. Consider for this example the graph in Figure 2.

When encoding a problem to SAT, we start by defining the meaning of the *variables* that we will use in our formula. In this case, we can use an unary encoding and consider 3 variables per color for each node. Let's denote A^y, A^b, A^r Boolean variables that are true if A is colored yellow (y), blue (b), or red (r), respectively. Similarly, we can define variables $B^y, B^b, B^r, C^y, C^b, C^r, D^y, D^b, D^r, E^y, E^b, E^r$, for the remaining nodes. Given these variables, we can now encode the problem by adding the following clauses:

- If two nodes are connected then they do not have the same color:
  $(\neg A^y \vee \neg E^y) \wedge (\neg A^b \vee \neg E^b) \wedge (\neg A^r \vee \neg E^r)$
  $(\neg A^y \vee \neg C^y) \wedge (\neg A^b \vee \neg C^b) \wedge (\neg A^r \vee \neg C^r)$
  $(\neg C^y \vee \neg B^y) \wedge (\neg C^b \vee \neg B^b) \wedge (\neg C^r \vee \neg B^r)$

Figure 2: 3-coloring of a graph.

$(\neg C^y \vee \neg D^y) \wedge (\neg C^b \vee \neg D^b) \wedge (\neg C^r \vee \neg D^r)$
$(\neg B^y \vee \neg D^y) \wedge (\neg B^b \vee \neg D^b) \wedge (\neg B^r \vee \neg D^r)$
$(\neg B^y \vee \neg E^y) \wedge (\neg B^b \vee \neg E^b) \wedge (\neg B^r \vee \neg E^r)$
$(\neg D^y \vee \neg E^y) \wedge (\neg D^b \vee \neg E^b) \wedge (\neg D^r \vee \neg E^r)$

- Each node has at-least-one color:
  $(A^y \vee A^b \vee A^r)$
  $(B^y \vee B^b \vee B^r)$
  $(C^y \vee C^b \vee C^r)$
  $(D^y \vee D^b \vee D^r)$
  $(E^y \vee E^b \vee E^r)$

- Each node has at-most-one color:
  $(\neg A^y \vee \neg A^b) \wedge (\neg A^y \vee \neg A^r) \wedge (\neg A^r \vee \neg A^b)$
  $(\neg B^y \vee \neg B^b) \wedge (\neg B^y \vee \neg B^r) \wedge (\neg B^r \vee \neg B^b)$
  $(\neg C^y \vee \neg C^b) \wedge (\neg C^y \vee \neg C^r) \wedge (\neg C^r \vee \neg C^b)$
  $(\neg D^y \vee \neg D^b) \wedge (\neg D^y \vee \neg D^r) \wedge (\neg D^r \vee \neg D^b)$
  $(\neg E^y \vee \neg E^b) \wedge (\neg E^y \vee \neg E^r) \wedge (\neg E^r \vee \neg E^b)$

A SAT solver can solve this formula and return the interpretation $I = \{A^r, B^b, C^y, D^r, E^y\}$ (for simplicity omit the variables assigned to false from the interpretation). If we decode this interpretation to the original problem, we obtain the coloring presented in Figure 2.

When encoding problems into SAT, it is useful to use a framework like PySAT with APIs that allows us to create a CNF formula, solve it with a SAT solver, and then interpret the solution and map it back to our original problem. An example of using PySAT to solve this problem was presented during the lecture and is available in color.py.

References

[Sin05] Carsten Sinz. Towards an optimal CNF encoding of boolean cardinality constraints. In *International Conference on Principles and Practice of Constraint Programming*, volume 3709, pages 827–831. Springer, 2005.