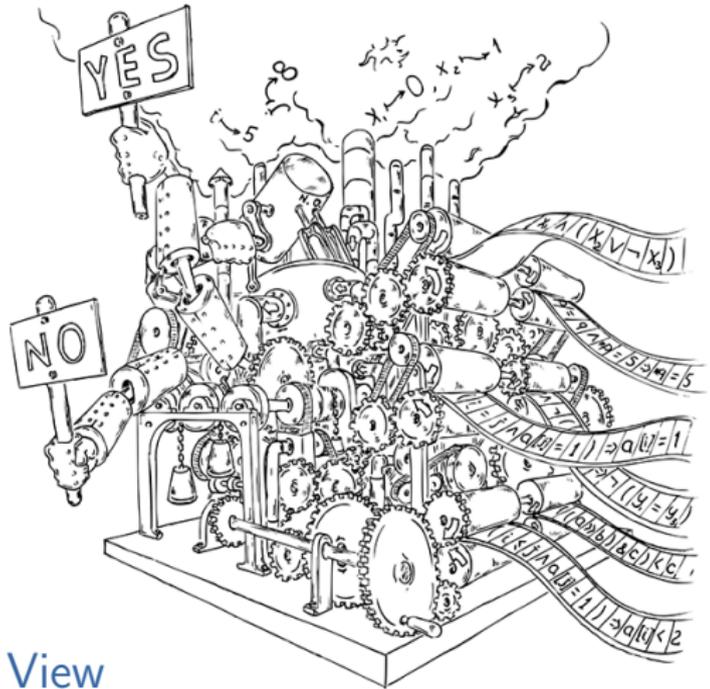


Propositional Encodings

Chapter 11



Decision Procedures An Algorithmic Point of View

- 1 Overview
- 2 Notation
- 3 A Basic Encoding Algorithm
- 4 Integration into DPLL
- 5 Theory Propagation and the DPLL(T) Framework
- 6 Theory Propagation and the DPLL(T) Framework
- 7 Optimizations and Implementation Issues

- Let T be a first-order Σ -theory such that:
 - T is quantifier-free.
 - There exists a decision procedure, denoted DP_T , for the conjunctive fragment of T .

- Example 1:
 - T is equality logic.
 - DP_T is the congruence closure algorithm.

- Example 2:
 - T is disjunctive linear arithmetic.
 - DP_T is the Simplex algorithm.

We will now study a framework that combines

- DP_T , and
- a SAT solver,

in various ways, in order to construct a decision procedure for T .

We will now study a framework that combines

- DP_T , and
- a SAT solver,

in various ways, in order to construct a decision procedure for T .

This method is

- modular,
- efficient,
- competitive (all state-of-the-art SMT solvers work this way).

The two main engines in this framework work in tight collaboration:

- The **SAT solver** chooses those literals that need to be satisfied in order to satisfy the Boolean structure of the formula, and
- The **theory solver** DP_T checks whether this choice is consistent in T .

Let l be a Σ -literal.

- Denote by $e(l)$ the **Boolean encoder of this literal**.

Let t be a Σ -formula,

- Denote by $e(t)$ the Boolean formula resulting from substituting each Σ -literal in t with its Boolean encoder.

Let l be a Σ -literal.

- Denote by $e(l)$ the **Boolean encoder of this literal**.

Let t be a Σ -formula,

- Denote by $e(t)$ the Boolean formula resulting from substituting each Σ -literal in t with its Boolean encoder.

For a Σ -formula t , the resulting Boolean formula $e(t)$ is called the **propositional skeleton of t** .

- **Example 1:** Let $l := x = y$ be a Σ -literal. Then $e(x = y)$, a Boolean variable, is its encoder.

- **Example I:** Let $l := x = y$ be a Σ -literal. Then $e(x = y)$, a Boolean variable, is its encoder.

- **Example II:** Let

$$t := x = y \vee x = z$$

be a Σ -formula. Then

$$e(t) := e(x = y) \vee e(x = z)$$

is its Boolean encoder.

Let T be equality logic. Given an NNF formula

$$\varphi := x = y \wedge ((y = z \wedge x \neq z) \vee x = z), \quad (1)$$

we begin by computing its propositional skeleton:

Let T be equality logic. Given an NNF formula

$$\varphi := x = y \wedge ((y = z \wedge x \neq z) \vee x = z), \quad (1)$$

we begin by computing its propositional skeleton:

$$e(\varphi) := e(x = y) \wedge ((e(y = z) \wedge e(x \neq z)) \vee e(x = z)). \quad (2)$$

Let T be equality logic. Given an NNF formula

$$\varphi := x = y \wedge ((y = z \wedge x \neq z) \vee x = z), \quad (1)$$

we begin by computing its propositional skeleton:

$$e(\varphi) := e(x = y) \wedge ((e(y = z) \wedge e(x \neq z)) \vee e(x = z)). \quad (2)$$

Note that since we are encoding *literals* and not *atoms*, $e(\varphi)$ has no negations and hence is trivially satisfiable.

Let \mathcal{B} be a Boolean formula, initially set to $e(\varphi)$, i.e.,

$$\mathcal{B} := e(\varphi) .$$

Let \mathcal{B} be a Boolean formula, initially set to $e(\varphi)$, i.e.,

$$\mathcal{B} := e(\varphi) .$$

As a second step, we pass \mathcal{B} to a SAT solver.

Let \mathcal{B} be a Boolean formula, initially set to $e(\varphi)$, i.e.,

$$\mathcal{B} := e(\varphi) .$$

As a second step, we pass \mathcal{B} to a SAT solver.

Assume that the SAT solver returns the satisfying assignment

$$\alpha := \{e(x = y) \mapsto \text{TRUE}, e(y = z) \mapsto \text{TRUE}, e(x \neq z) \mapsto \text{TRUE}, \\ e(x = z) \mapsto \text{FALSE}\} .$$

- Denote by $\hat{T}h(\alpha)$ the conjunction of the literals corresponding to this assignment.

$$\hat{T}h(\alpha) := x = y \wedge y = z \wedge x \neq z \wedge \neg(x = z) .$$

- Denote by $\hat{T}h(\alpha)$ the conjunction of the literals corresponding to this assignment.

$$\hat{T}h(\alpha) := x = y \wedge y = z \wedge x \neq z \wedge \neg(x = z) .$$

- The decision procedure DP_T now has to decide whether $\hat{T}h(\alpha)$ is satisfiable.

- Denote by $\hat{T}h(\alpha)$ the conjunction of the literals corresponding to this assignment.

$$\hat{T}h(\alpha) := x = y \wedge y = z \wedge x \neq z \wedge \neg(x = z) .$$

- The decision procedure DP_T now has to decide whether $\hat{T}h(\alpha)$ is satisfiable.

$\hat{T}h(\alpha)$ is not satisfiable, which means that the negation of this formula is a tautology.

Thus \mathcal{B} is conjoined with $e(\neg\hat{T}h(\alpha))$, the Boolean encoding of this tautology:

$$e(\neg\hat{T}h(\alpha)) := (\neg e(x = y) \vee \neg e(y = z) \vee \neg e(x \neq z) \vee e(x = z)) .$$

Thus \mathcal{B} is conjoined with $e(\neg\hat{T}h(\alpha))$, the Boolean encoding of this tautology:

$$e(\neg\hat{T}h(\alpha)) := (\neg e(x = y) \vee \neg e(y = z) \vee \neg e(x \neq z) \vee e(x = z)) .$$

- This clause contradicts the current assignment, and hence **blocks** it from being repeated.
- Such clauses are called **blocking clauses**.

Thus \mathcal{B} is conjoined with $e(\neg\hat{T}h(\alpha))$, the Boolean encoding of this tautology:

$$e(\neg\hat{T}h(\alpha)) := (\neg e(x = y) \vee \neg e(y = z) \vee \neg e(x \neq z) \vee e(x = z)) .$$

- This clause contradicts the current assignment, and hence **blocks** it from being repeated.
- Such clauses are called **blocking clauses**.
- We denote by t the formula – also called the **lemma** – returned by DP_T (in this example $t := \neg\hat{T}h(\alpha)$).

After the blocking clause has been added, the SAT solver is invoked again and suggests another assignment, for example

$$\alpha' := \{e(x = y) \mapsto \text{TRUE}, e(y = z) \mapsto \text{TRUE}, e(x = z) \mapsto \text{TRUE}, \\ e(x \neq z) \mapsto \text{FALSE}\} .$$

After the blocking clause has been added, the SAT solver is invoked again and suggests another assignment, for example

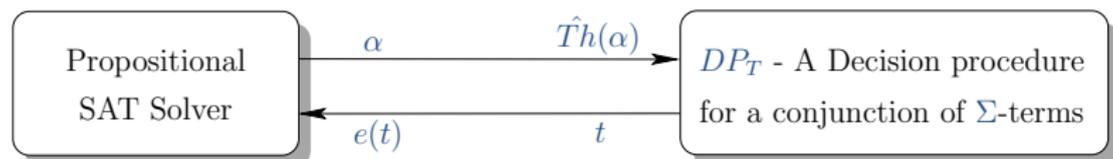
$$\alpha' := \{e(x = y) \mapsto \text{TRUE}, e(y = z) \mapsto \text{TRUE}, e(x = z) \mapsto \text{TRUE}, e(x \neq z) \mapsto \text{FALSE}\} .$$

The corresponding Σ -formula

$$\hat{T}h(\alpha') := x = y \wedge y = z \wedge x = z \wedge \neg(x \neq z) \quad (3)$$

is satisfiable, which proves that φ , the original formula, is satisfiable.

Indeed, any assignment that satisfies $\hat{T}h(\alpha')$ also satisfies φ .



The information flow between the two components of the decision procedure.

There are many improvements to this basic procedure.

There are many improvements to this basic procedure.

One such improvement is:

“Invoke the decision procedure DP_T after **partial assignments**, rather than waiting for a full assignment.”

There are many improvements to this basic procedure.

One such improvement is:

“Invoke the decision procedure DP_T after **partial assignments**, rather than waiting for a full assignment.”

- A contradicting partial assignment leads to a more powerful lemma t , as it blocks all assignments that extend it.

There are many improvements to this basic procedure.

One such improvement is:

“Invoke the decision procedure DP_T after **partial assignments**, rather than waiting for a full assignment.”

- A contradicting partial assignment leads to a more powerful lemma t , as it blocks all assignments that extend it.
- **Theory propagation**: When the partial assignment is not contradictory, it can be used to derive implications that are propagated back to the SAT solver.

Continuing the example above, consider the partial assignment

$$\alpha := \{e(x = y) \mapsto \text{TRUE}, e(y = z) \mapsto \text{TRUE}\}, \quad (4)$$

Continuing the example above, consider the partial assignment

$$\alpha := \{e(x = y) \mapsto \text{TRUE}, e(y = z) \mapsto \text{TRUE}\}, \quad (4)$$

and the corresponding formula that is transferred to DP_T ,

$$\hat{T}h(\alpha) := x = y \wedge y = z. \quad (5)$$

Continuing the example above, consider the partial assignment

$$\alpha := \{e(x = y) \mapsto \text{TRUE}, e(y = z) \mapsto \text{TRUE}\}, \quad (4)$$

and the corresponding formula that is transferred to DP_T ,

$$\hat{T}h(\alpha) := x = y \wedge y = z. \quad (5)$$

DP_T concludes that $x = z$ is implied, and hence inform the SAT solver that $e(x = z) \mapsto \text{TRUE}$ and $e(x \neq z) \mapsto \text{FALSE}$ are implied by the current partial assignment α .

We will now formalize three versions of the algorithm:

- 1 Simple
- 2 Incremental
- 3 DPLL(T)

- $lit(\varphi)$ – the set of literals in a given NNF formula φ .
- $lit_i(\varphi)$ – the i -th distinct literal in φ
(assuming some predefined order on the literals).

- $lit(\varphi)$ – the set of literals in a given NNF formula φ .
- $lit_i(\varphi)$ – the i -th distinct literal in φ
(assuming some predefined order on the literals).
- α – For a given encoding $e(\varphi)$, denotes an assignment (either full or partial), to the encoders in $e(\varphi)$.

- $Th(lit_i, \alpha)$ – For an encoder $e(lit_i)$ that is assigned a truth value by α , denotes the corresponding literal:

$$Th(lit_i, \alpha) \doteq \begin{cases} lit_i & \alpha(lit_i) = \text{TRUE} \\ \neg lit_i & \alpha(lit_i) = \text{FALSE} . \end{cases} \quad (6)$$

- $Th(lit_i, \alpha)$ – For an encoder $e(lit_i)$ that is assigned a truth value by α , denotes the corresponding literal:

$$Th(lit_i, \alpha) \doteq \begin{cases} lit_i & \alpha(lit_i) = \text{TRUE} \\ \neg lit_i & \alpha(lit_i) = \text{FALSE} . \end{cases} \quad (6)$$

- $Th(\alpha) \doteq \{Th(lit_i, \alpha) \mid e(lit_i) \text{ is assigned by } \alpha\}$

- $Th(lit_i, \alpha)$ – For an encoder $e(lit_i)$ that is assigned a truth value by α , denotes the corresponding literal:

$$Th(lit_i, \alpha) \doteq \begin{cases} lit_i & \alpha(lit_i) = \text{TRUE} \\ \neg lit_i & \alpha(lit_i) = \text{FALSE} . \end{cases} \quad (6)$$

- $Th(\alpha) \doteq \{Th(lit_i, \alpha) \mid e(lit_i) \text{ is assigned by } \alpha\}$
- $\hat{Th}(\alpha)$ – a conjunction over the elements in $Th(\alpha)$.

Let

$$lit_1 = (x = y), \quad lit_2 = (y = z), \quad lit_3 = (z = w), \quad (7)$$

Let

$$lit_1 = (x = y), \quad lit_2 = (y = z), \quad lit_3 = (z = w), \quad (7)$$

and let α be a partial assignment such that

$$\alpha := \{e(lit_1) \mapsto \text{FALSE}, e(lit_2) \mapsto \text{TRUE}\} .$$

Let

$$lit_1 = (x = y), lit_2 = (y = z), lit_3 = (z = w), \quad (7)$$

and let α be a partial assignment such that

$$\alpha := \{e(lit_1) \mapsto \text{FALSE}, e(lit_2) \mapsto \text{TRUE}\}.$$

Then

$$Th(lit_1, \alpha) := \neg(x = y), Th(lit_2, \alpha) := (y = z),$$

Let

$$lit_1 = (x = y), lit_2 = (y = z), lit_3 = (z = w), \quad (7)$$

and let α be a partial assignment such that

$$\alpha := \{e(lit_1) \mapsto \text{FALSE}, e(lit_2) \mapsto \text{TRUE}\}.$$

Then

$$Th(lit_1, \alpha) := \neg(x = y), Th(lit_2, \alpha) := (y = z),$$

and

$$Th(\alpha) := \{\neg(x = y), (y = z)\}.$$

Let

$$lit_1 = (x = y), lit_2 = (y = z), lit_3 = (z = w), \quad (7)$$

and let α be a partial assignment such that

$$\alpha := \{e(lit_1) \mapsto \text{FALSE}, e(lit_2) \mapsto \text{TRUE}\}.$$

Then

$$Th(lit_1, \alpha) := \neg(x = y), Th(lit_2, \alpha) := (y = z),$$

and

$$Th(\alpha) := \{\neg(x = y), (y = z)\}.$$

Conjoining these terms gives us

$$\hat{Th}(\alpha) := \neg(x = y) \wedge (y = z).$$

- T – a Σ -theory.

- T – a Σ -theory.
- DP_T a decision procedure for the conjunctive fragment of T .

- T – a Σ -theory.
- DP_T a decision procedure for the conjunctive fragment of T .
- Let DEDUCTION be a procedure based on DP_T , which receives a conjunction of Σ -literals as input, and
 - decides whether it is satisfiable, and,
 - if the answer is negative, returns constraints over these literals.

1. A Basic Algorithm

```
1: function LAZY-BASIC( $\varphi$ )
2:    $\mathcal{B} := e(\varphi)$ ;
3:   while (TRUE) do
4:      $\langle \alpha, res \rangle := \text{SAT-SOLVER}(\mathcal{B})$ ;
5:     if  $res = \text{"Unsatisfiable"}$  then return "Unsatisfiable";
6:     else
7:        $\langle t, res \rangle := \text{DEDUCTION}(\hat{T}h(\alpha))$ ;
8:       if  $res = \text{"Satisfiable"}$  then return "Satisfiable";
9:        $\mathcal{B} := \mathcal{B} \wedge e(t)$ ;
```

Consider the following three requirements from the clause t that is returned by DEDUCTION:

- 1 The formula t is T -valid, i.e., t is a tautology in T . For example, if T is the theory of equality, then $x = y \wedge y = z \longrightarrow x = z$ is T -valid.

Consider the following three requirements from the clause t that is returned by DEDUCTION:

- 1 The formula t is T -valid, i.e., t is a tautology in T . For example, if T is the theory of equality, then $x = y \wedge y = z \longrightarrow x = z$ is T -valid.
- 2 The atoms in t are restricted to those appearing in φ .

Consider the following three requirements from the clause t that is returned by DEDUCTION:

- 1 The formula t is T -valid, i.e., t is a tautology in T . For example, if T is the theory of equality, then $x = y \wedge y = z \longrightarrow x = z$ is T -valid.
- 2 The atoms in t are restricted to those appearing in φ .
- 3 The encoding of t contradicts α , i.e., $e(t)$ is a blocking clause.

Consider the following three requirements from the clause t that is returned by DEDUCTION:

- 1 The formula t is T -valid, i.e., t is a tautology in T . For example, if T is the theory of equality, then $x = y \wedge y = z \longrightarrow x = z$ is T -valid.
- 2 The atoms in t are restricted to those appearing in φ .
- 3 The encoding of t contradicts α , i.e., $e(t)$ is a blocking clause.

The first requirement is sufficient for guaranteeing soundness.

Consider the following three requirements from the clause t that is returned by DEDUCTION:

- 1 The formula t is T -valid, i.e., t is a tautology in T . For example, if T is the theory of equality, then $x = y \wedge y = z \longrightarrow x = z$ is T -valid.
- 2 The atoms in t are restricted to those appearing in φ .
- 3 The encoding of t contradicts α , i.e., $e(t)$ is a blocking clause.

The first requirement is sufficient for guaranteeing soundness.

The second and third requirements are sufficient for guaranteeing termination.

Two of the requirements can be weakened:

Two of the requirements can be weakened:

- **Requirement 1:** the clause t can be any formula that is implied by φ , and not just a T -valid formula.

Two of the requirements can be weakened:

- **Requirement 1:** the clause t can be any formula that is implied by φ , and not just a T -valid formula.
- **Requirement 2:** the clause t may refer to atoms that do not appear in φ , as long as the number of such new atoms is finite.

Two of the requirements can be weakened:

- **Requirement 1:** the clause t can be any formula that is implied by φ , and not just a T -valid formula.
- **Requirement 2:** the clause t may refer to atoms that do not appear in φ , as long as the number of such new atoms is finite.
 - For example, in equality logic, we may allow t to refer to all atoms of the form $x_i = x_j$ where x_i, x_j are variables in $\text{var}(\varphi)$, even if only some of these equality predicates appear in φ .

2. We can do better...

- Let \mathcal{B}^i be the formula \mathcal{B} in the i -th iteration of the loop.

2. We can do better...

- Let \mathcal{B}^i be the formula \mathcal{B} in the i -th iteration of the loop.
- The constraint \mathcal{B}^{i+1} is strictly stronger than \mathcal{B}^i for all $i \geq 1$, because clauses are added but not removed between iterations.

2. We can do better...

- Let \mathcal{B}^i be the formula \mathcal{B} in the i -th iteration of the loop.
- The constraint \mathcal{B}^{i+1} is strictly stronger than \mathcal{B}^i for all $i \geq 1$, because clauses are added but not removed between iterations.
- As a result, any conflict clause that is learned while solving \mathcal{B}^i can be reused when solving \mathcal{B}^j for $i < j$.

2. We can do better...

- Let \mathcal{B}^i be the formula \mathcal{B} in the i -th iteration of the loop.
- The constraint \mathcal{B}^{i+1} is strictly stronger than \mathcal{B}^i for all $i \geq 1$, because clauses are added but not removed between iterations.
- As a result, any conflict clause that is learned while solving \mathcal{B}^i can be reused when solving \mathcal{B}^j for $i < j$.
- This is a special case of **incremental satisfiability**.

2. We can do better...

- Hence, invoking an incremental SAT solver in line 4 can increase the efficiency of the algorithm.

2. We can do better...

- Hence, invoking an incremental SAT solver in line 4 can increase the efficiency of the algorithm.
- A better option is to **integrate** DEDUCTION into the DPLL-SAT algorithm, as shown in the following algorithm.

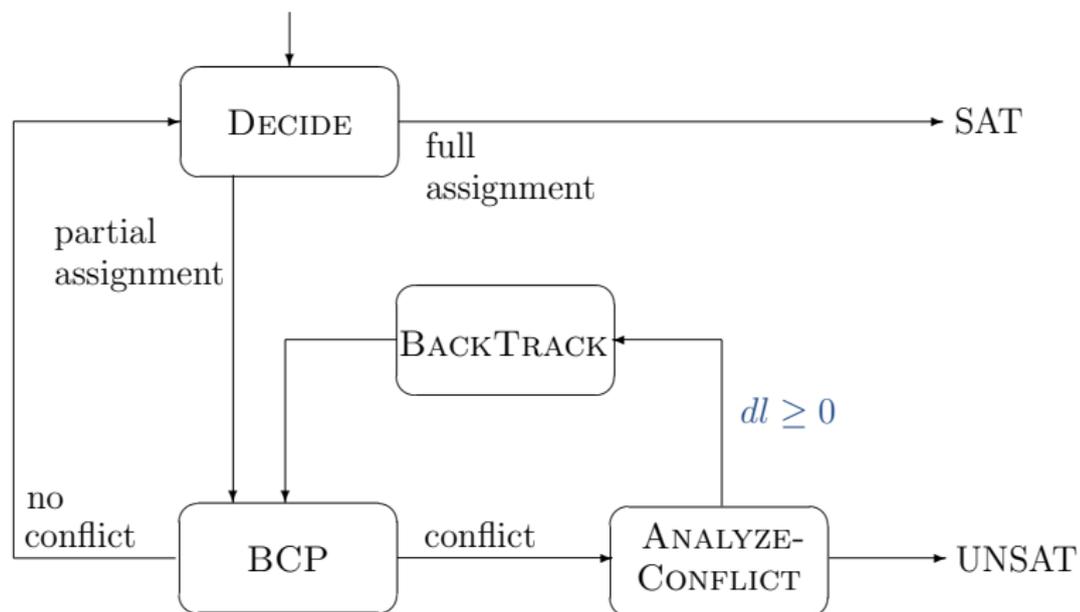
2. We can do better...

- Hence, invoking an incremental SAT solver in line 4 can increase the efficiency of the algorithm.
- A better option is to **integrate** DEDUCTION into the DPLL-SAT algorithm, as shown in the following algorithm.
- This algorithm uses a procedure `ADDCLAUSES`, which adds new clauses to the current set of clauses at run time.

2. We can do better...

- Hence, invoking an incremental SAT solver in line 4 can increase the efficiency of the algorithm.
- A better option is to **integrate** DEDUCTION into the DPLL-SAT algorithm, as shown in the following algorithm.
- This algorithm uses a procedure ADDCLAUSES, which adds new clauses to the current set of clauses at run time.
- Before seeing this algorithm let us first recall DPLL...

2. A Reminder: DPLL



2. Pseudo-code for DPLL

```
1: function DPLL
2:   if BCP() = "conflict" then return "Unsatisfiable";
3:   while (TRUE) do
4:     if  $\neg$ DECIDE() then return "Satisfiable";
5:     else
6:       while (BCP() = "conflict") do
7:         backtrack-level := ANALYZE-CONFLICT();
8:         if backtrack-level < 0 then return
9:           "Unsatisfiable";
10:        else BackTrack(backtrack-level);
```

2. Integration into DPLL

```
1: function LAZY-DPLL
2:   ADDCLAUSES(cnf( $e(\varphi)$ ));
3:   if BCP() = "conflict" then return "Unsatisfiable";
4:   while (TRUE) do
5:     if  $\neg$ DECIDE() then ▷ Full assignment
6:        $\langle t, res \rangle :=$  DEDUCTION( $\hat{T}h(\alpha)$ );
7:       if res = "Satisfiable" then return "Satisfiable";
8:       ADDCLAUSES( $e(t)$ );
9:       while (BCP() = "conflict") do
10:        backtrack-level := ANALYZE-CONFLICT();
11:        if backtrack-level < 0 then return "Unsatisfiable";
12:        else BackTrack(backtrack-level);
13:     else
14:       while (BCP() = "conflict") do
15:        backtrack-level := ANALYZE-CONFLICT();
16:        if backtrack-level < 0 then return "Unsatisfiable";
17:        else BackTrack(backtrack-level);
```

3. DPLL(T)

- Consider a formula φ that contains an integer variable x_1 and, among others, the literals $x_1 \geq 10$ and $x_1 < 0$.

3. DPLL(T)

- Consider a formula φ that contains an integer variable x_1 and, among others, the literals $x_1 \geq 10$ and $x_1 < 0$.
- Assume that the DECIDE procedure assigns $e(x_1 \geq 10) \mapsto \text{TRUE}$ and $e(x_1 < 0) \mapsto \text{TRUE}$.

3. DPLL(T)

- Consider a formula φ that contains an integer variable x_1 and, among others, the literals $x_1 \geq 10$ and $x_1 < 0$.
- Assume that the DECIDE procedure assigns $e(x_1 \geq 10) \mapsto \text{TRUE}$ and $e(x_1 < 0) \mapsto \text{TRUE}$.
- Inevitably, any call to DEDUCTION results in a **contradiction** between these two facts, independently of any other decisions that are made.

- However, the algorithms we saw so far do not call DEDUCTION until a **full satisfying assignment** is found.
 - Thus, the time taken to complete the assignment is wasted.

3. DPLL(T)

- However, the algorithms we saw so far do not call DEDUCTION until a **full satisfying assignment** is found.
 - Thus, the time taken to complete the assignment is wasted.
- Further, the refutation of this full assignment may be due to other reasons (i.e., a proof that a different subset of the assignment is contradictory).
 - Hence, additional assignments that include the same wrong assignment to $e(x_1 \geq 10)$ and $e(x_1 < 0)$ are not ruled out.

Early call to DEDUCTION can serve two purposes:

Early call to DEDUCTION can serve two purposes:

- 1 Contradictory partial assignments are ruled out early.

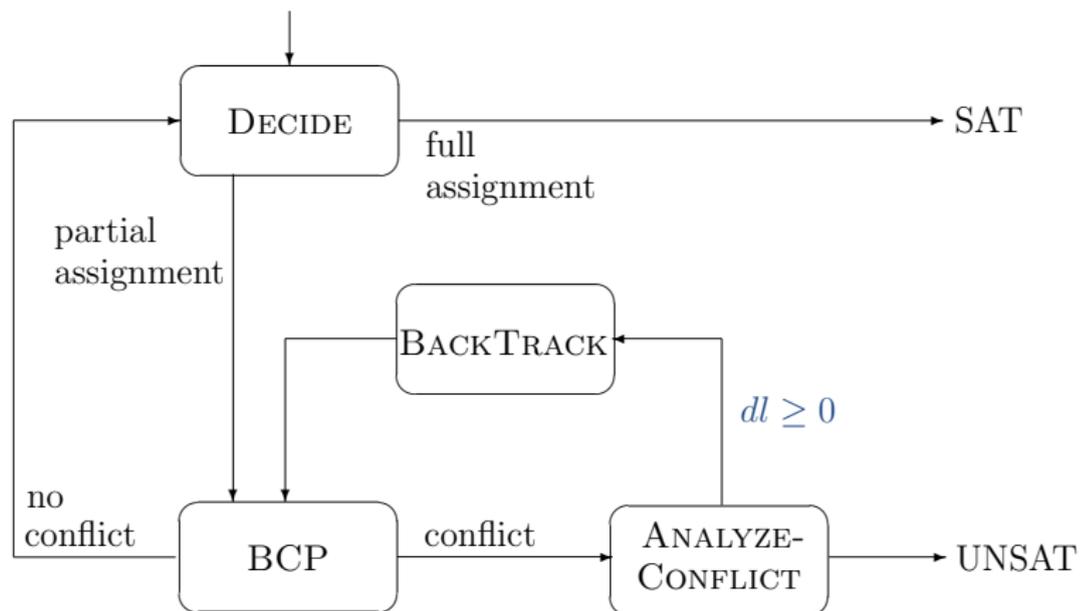
3. DPLL(T)

Early call to DEDUCTION can serve two purposes:

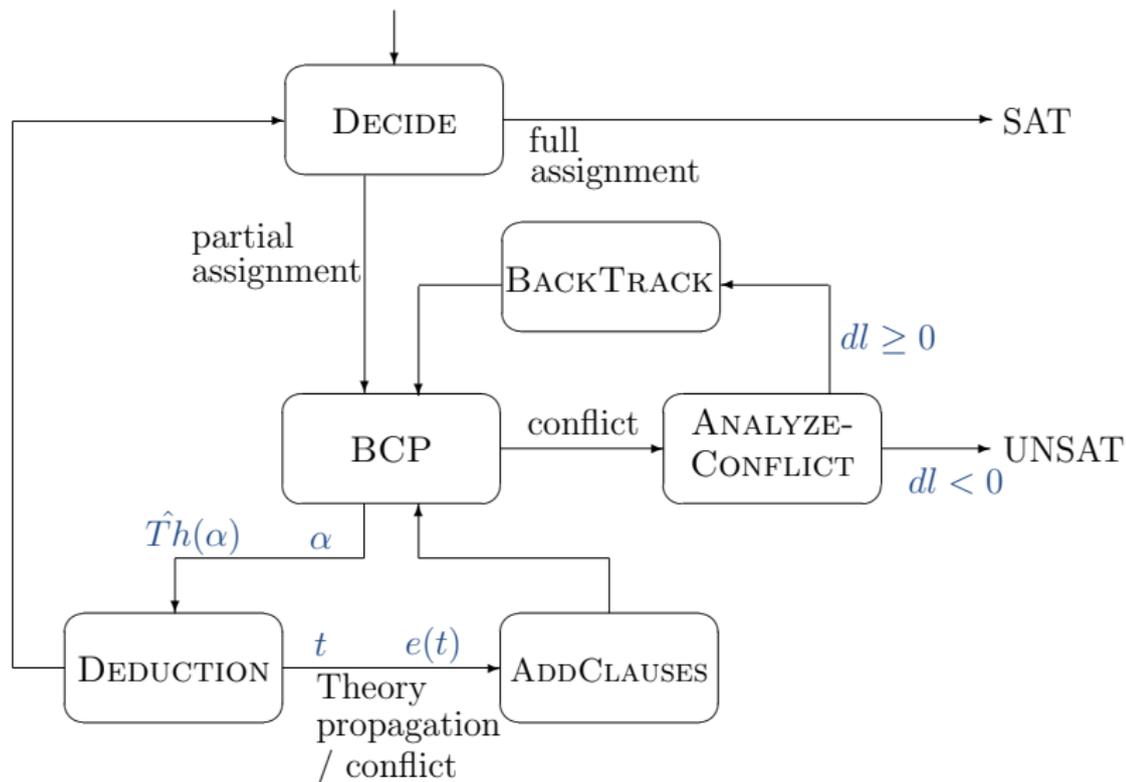
- 1 Contradictory partial assignments are ruled out early.
- 2 Allows **theory propagation**.
 - Continuing our example, once $e(x_1 \geq 10)$ has been assigned TRUE, we can infer that $e(x_1 < 0)$ must be FALSE and avoid the conflict altogether.

This brings us to the next version of the algorithm, called DPLL(T).

3. Reminder: DPLL



3. ... and now DPLL(T)



```

1: function DPLL( $T$ )
2:   ADDCLAUSES( $cnf(e(\varphi))$ );
3:   if BCP() = "conflict" then return "Unsatisfiable";
4:   while (TRUE) do
5:     if  $\neg$ DECIDE() then return "Satisfiable";           ▷ Full
assignment
6:     repeat
7:       while (BCP() = "conflict") do
8:          $backtrack-level :=$  ANALYZE-CONFLICT();
9:         if  $backtrack-level < 0$  then return
"Unsatisfiable";
10:        else BackTrack( $backtrack-level$ );
11:         $\langle t, res \rangle :=$  DEDUCTION( $\hat{T}h(\alpha)$ );
12:        ADDCLAUSES( $e(t)$ );
13:    until  $t \equiv$  TRUE

```

If $\hat{T}h(\alpha)$ is satisfiable, we require t to fulfill one of the following two conditions in order to guarantee termination:

If $\hat{T}h(\alpha)$ is satisfiable, we require t to fulfill one of the following two conditions in order to guarantee termination:

- 1 The clause $e(t)$ is an **asserting clause** under α . This implies that the addition of $e(t)$ to \mathcal{B} and a call to BCP leads to an assignment to the encoder of some literal.

If $\hat{T}h(\alpha)$ is satisfiable, we require t to fulfill one of the following two conditions in order to guarantee termination:

- 1 The clause $e(t)$ is an **asserting clause** under α . This implies that the addition of $e(t)$ to \mathcal{B} and a call to BCP leads to an assignment to the encoder of some literal.
- 2 When DEDUCTION cannot find an asserting clause t as defined above, t and $e(t)$ are equivalent to TRUE.

The second case occurs, for example, when all the Boolean variables are already assigned, and thus the formula is found to be satisfiable.

3. Theory Propagation

Various ways to perform theory propagation:

- After every decision / after every assignment

Various ways to perform theory propagation:

- After every decision / after every assignment
- Partial / **Exhaustive** theory propagation – propagate all that is implied by the current partial assignment.

Various ways to perform theory propagation:

- After every decision / after every assignment
- Partial / **Exhaustive** theory propagation – propagate all that is implied by the current partial assignment.
- Refer only to existing predicates / add auxiliary ones.

Exhaustive theory propagation after each assignment: what does this mean ?

3. Theory Propagation

Various ways to perform theory propagation:

- After every decision / after every assignment
- Partial / **Exhaustive** theory propagation – propagate all that is implied by the current partial assignment.
- Refer only to existing predicates / add auxiliary ones.

Exhaustive theory propagation after each assignment: what does this mean ?

That's right, no possible conflicts on the theory side.

3. Theory Propagation

How to check whether a predicate p is implied by $\hat{T}h(\alpha)$?

- **Plunging** – is $\hat{T}h(\alpha) \wedge \neg p$ satisfiable ?

3. Theory Propagation

How to check whether a predicate p is implied by $\hat{Th}(\alpha)$?

- **Plunging** – is $\hat{Th}(\alpha) \wedge \neg p$ satisfiable ?
- Theory-specific propagation. For example, in equality logic build the equality graph corresponding to $Th(\alpha)$. Infer equalities/disequalities from the graph.

3. Theory Propagation

How to check whether a predicate p is implied by $\hat{Th}(\alpha)$?

- **Plunging** – is $\hat{Th}(\alpha) \wedge \neg p$ satisfiable ?
- Theory-specific propagation. For example, in equality logic build the equality graph corresponding to $Th(\alpha)$. Infer equalities/disequalities from the graph.
- Note that theory propagation matters for efficiency, not correctness.

3. Theory Propagation

How to check whether a predicate p is implied by $\hat{T}h(\alpha)$?

- **Plunging** – is $\hat{T}h(\alpha) \wedge \neg p$ satisfiable ?
- Theory-specific propagation. For example, in equality logic build the equality graph corresponding to $Th(\alpha)$. Infer equalities/disequalities from the graph.
- Note that theory propagation matters for efficiency, not correctness.
- How much propagation is cost-effective is a subject for research, and depends on T .

3. Theory Propagation – How?

- Normally theory propagation is done by transferring clauses to the the DPLL solver.
- It turns out to be inefficient – few (less than 0.5%) are actually used.
- Instead – add implied literals directly to the implication stack.
 - This causes a **problem** in `ANALYZE-CONFLICT()` – can you see what problem ?

3. Theory Propagation – How?

- The problem: `ANALYZE-CONFLICT()` requires an **antecedent clause** for each implication, in order to compute the conflict clause and backtrack level.

3. Theory Propagation – How?

- The problem: `ANALYZE-CONFLICT()` requires an **antecedent clause** for each implication, in order to compute the conflict clause and backtrack level.
- Theory propagation without clauses breaks this mechanism – there are implications without antecedents.

3. Theory Propagation – How?

- The problem: $\text{ANALYZE-CONFLICT}()$ requires an **antecedent clause** for each implication, in order to compute the conflict clause and backtrack level.
- Theory propagation without clauses breaks this mechanism – there are implications without antecedents.
- Solution – DP_T should be able to **explain** an implication post-mortem, in the form of a clause.

3. Strong Lemmas

- When $\hat{T}h(\alpha)$ is unsatisfiable, the lemma (clause returned by DEDUCTION) rules out α .

3. Strong Lemmas

- When $\hat{T}h(\alpha)$ is unsatisfiable, the lemma (clause returned by DEDUCTION) rules out α .
- Ideally, it should be **generalized** as much as possible.

3. Strong Lemmas

- When $\hat{T}h(\alpha)$ is unsatisfiable, the lemma (clause returned by DEDUCTION) rules out α .
- Ideally, it should be **generalized** as much as possible.
- Solution: analyze the **reason** for unsatisfiability.
Build lemma accordingly.

3. Strong Lemmas – An Example

