



UNIVERSITY
OF TRENTO

DEPARTMENT OF INFORMATION AND COMMUNICATION TECHNOLOGY

38050 Povo – Trento (Italy), Via Sommarive 14
<http://www.dit.unitn.it>

DELAYED THEORY COMBINATION VS. NELSON-OPPEN FOR
SATISFIABILITY MODULO THEORIES: A COMPARATIVE
ANALYSIS

Roberto Bruttomesso, Alessandro Cimatti, Anders Franzén
Alberto Griggio and Roberto Sebastiani

May 2006

Technical Report # DIT-06-032

Delayed Theory Combination vs. Nelson-Oppen for Satisfiability Modulo Theories: a Comparative Analysis^{*}

Roberto Bruttomesso¹, Alessandro Cimatti¹, Anders Franzen^{1,2},
Alberto Griggio², and Roberto Sebastiani²

¹ ITC-IRST, Povo, Trento, Italy. {bruttomesso, cimatti, franzen}@itc.it

² DIT, Università di Trento, Italy. {griggio, rseba}@dit.unitn.it

Abstract. Many approaches for Satisfiability Modulo Theory ($SMT(\mathcal{T})$) rely on the integration between a SAT solver and a decision procedure for sets of literals in the background theory \mathcal{T} (\mathcal{T} -solver). When \mathcal{T} is the combination $\mathcal{T}_1 \cup \mathcal{T}_2$ of two simpler theories, the approach is typically handled by means of Nelson-Oppen's (NO) theory combination schema in which two specific \mathcal{T} -solvers deduce and exchange (disjunctions of) interface equalities.

In recent papers we have proposed a new approach to $SMT(\mathcal{T}_1 \cup \mathcal{T}_2)$, called *Delayed Theory Combination* (DTC). Here part or all the (possibly very expensive) task of deducing interface equalities is played by the SAT solver itself, at the potential cost of an enlargement of the boolean search space. In principle this enlargement could be up to exponential in the number of interface equalities generated.

In this paper we show that this estimate was too pessimistic. We present a comparative analysis of DTC vs. NO for $SMT(\mathcal{T}_1 \cup \mathcal{T}_2)$, which shows that, using state-of-the-art SAT-solving techniques, the amount of boolean branches performed by DTC can be upper bounded by the number of deductions and boolean branches performed by NO on the same problem. We prove the result for different deduction capabilities of the \mathcal{T} -solvers and for both convex and non-convex theories.

1 Introduction

Satisfiability Modulo a Theory \mathcal{T} ($SMT(\mathcal{T})$) is the problem of checking the satisfiability of a quantifier-free (or ground) first-order formula with respect to a given first-order theory \mathcal{T} . Theories of interest for many applications are, e.g., the theory of difference logic \mathcal{DL} , the theory \mathcal{EUF} of equality and uninterpreted functions, the quantifier-free fragment of Linear Arithmetic over the rationals $\mathcal{LA}(\mathbb{Q})$ and that over the integers $\mathcal{LA}(\mathbb{Z})$. Particularly relevant is the case of $SMT(\mathcal{T}_1 \cup \mathcal{T}_2)$, where the background theory \mathcal{T} is the combination of two (or more) simpler theories \mathcal{T}_1 and \mathcal{T}_2 .¹

^{*} This work has been partly supported by ISAAC, an European sponsored project, contract no. AST3-CT-2003-501848, by ORCHID, a project sponsored by Provincia Autonoma di Trento, and by a grant from Intel Corporation.

¹ For better readability, and as it is common practice in papers dealing with combination of theories, in this paper we always deal with only two theories \mathcal{T}_1 and \mathcal{T}_2 . The discourse generalizes to more than two theories.

A prominent approach to $SMT(\mathcal{T})$ which underlies several systems (e.g., CVCLITE [2], DLSAT [8], DPLL(T)/BarceLogic [10], MATHSAT [4], TSAT++ [1], ICS/YICES [9]), is based on extensions of propositional SAT technology: a SAT engine is modified to enumerate boolean assignments, and integrated with a decision procedure for sets of literals in the theory \mathcal{T} (\mathcal{T} -solver). The above schema is also followed to tackle the $SMT(\mathcal{T}_1 \cup \mathcal{T}_2)$ problem. The approach relies on a decision procedure able to decide the satisfiability of sets of literals in $\mathcal{T}_1 \cup \mathcal{T}_2$, that is typically based on an integration schema like Nelson-Oppen (NO) [11] (or its variant due to Shostak [13]): the \mathcal{T}_i -solvers are combined by means of a structured exchange of (disjunctions of) interface equalities (e_{ij} 's).

Unfortunately from a practical point of view this schema poses some challenges. First, the integration between the two \mathcal{T}_i -solvers is not trivial to implement. Second, the ability of \mathcal{T}_i -solvers of inferring (disjunctions of) interface equalities (hereafter e_{ij} -deduction completeness) required by NO is neither always easy to achieve nor always cheap to perform. (E.g., e_{ij} -deduction is cheap for \mathcal{EUF} but can be very expensive for $\mathcal{LA}(\mathbb{Z})$.) Third, in case of non-convex theories (e.g., $\mathcal{LA}(\mathbb{Z})$), a backtrack search must be used to take care of the disjunctions that need to be managed.

In recent papers [3, 6] we have proposed a novel approach to $SMT(\mathcal{T}_1 \cup \mathcal{T}_2)$, called *Delayed Theory Combination* (DTC). The main idea is to avoid the integration schema between \mathcal{T}_1 and \mathcal{T}_2 , and tighten the connection between each \mathcal{T}_i and the SAT engine. While the truth assignment is being constructed, it is checked for consistency with respect to each theory in isolation. This can be seen as constructing two (possibly inconsistent) partial models for the original formula; the “merging” of the two partial models is enforced, on demand, since the solver is requested to find a complete assignment to the e_{ij} 's.

Compared to the NO schema, this approach has several advantages [3, 6]. First, it is easier to implement and analyze. Second, the approach does not rely on the \mathcal{T}_i -solvers being e_{ij} -deduction complete, although it can fully benefit from this property. Third, the DTC nicely encompasses the case of non-convex theories. On the negative side, in [3, 6] we noticed that these benefits are traded with a potential enlargement of the boolean search space which, in principle, could be up to exponential in the number of interface equalities generated. Thus, despite the positive empirical results presented in [3, 6], the latter fact represented, at least in theory, one possible drawback of DTC.

In this paper we show that this latter point was way too pessimistic. We present a comparative analysis of DTC vs. NO for $SMT(\mathcal{T}_1 \cup \mathcal{T}_2)$, and we introduce some novel theoretical results, for both convex and non-convex theories and for different deduction capabilities of the \mathcal{T} -solvers. These results show that, by exploiting the full power of advanced SAT techniques like backjumping and learning, DTC can be implemented in such a way as to mimic the behavior of NO, so that the amount of boolean branches required by DTC can be upper-bounded by the sum of the number of deductions and branches required by NO in order to perform the same tasks.

From these results we have that DTC generalizes NO, in the sense that:

- under the same hypotheses of e_{ij} -deduction-completeness of the \mathcal{T}_i -solvers required by NO, DTC emulates NO with no extra cost in terms of boolean search;

- in the more general case (\mathcal{T}_i -solvers with partial or no e_{ij} -deduction capabilities) DTC can mimic the behavior of NO, in such a way that all or part of the (possibly very expensive) e_{ij} -deductions are substituted with only few extra boolean branches.

We also notice that the capability of learning conflict clauses containing interface equalities, which is typical of DTC, allows for cutting branches corresponding to repeated deductions in an equivalent NO schema.

The paper is structured as follows. In Section 2 we present some background and introduce the Nelson-Oppen combination schema for $SMT(\mathcal{T}_1 \cup \mathcal{T}_2)$. DTC is then discussed in Section 3. We present our analysis in Sections 5 (where the case of e_{ij} -deduction completeness in the \mathcal{T}_i -solvers of DTC is examined) and 4 (where the \mathcal{T}_i -solvers employed by DTC are assumed to have limited or no deduction capabilities). Finally, in Section 6 we draw some conclusions.

The proofs of the theorems are reported in appendix.

2 SMT for combined theories via Nelson-Oppen's integration

2.1 Basic definitions and properties

Consider a theory \mathcal{T} with equality. \mathcal{T} is *stably-infinite* iff every quantifier-free \mathcal{T} -satisfiable formula is satisfiable in an infinite model of \mathcal{T} . Notice that \mathcal{EUF} , $\mathcal{DL}(\mathbb{Q})$, $\mathcal{DL}(\mathbb{Z})$, $\mathcal{LA}(\mathbb{Q})$, $\mathcal{LA}(\mathbb{Z})$ are stably-infinite, whereas e.g. theories of bit-vectors \mathcal{BV} are typically not. In what follows, we shall assume to deal only with stably-infinite theories with equality and with disjoint signatures.

\mathcal{T} is *convex* iff, for every collection l_1, \dots, l_k, e, e' of literals in \mathcal{T} s.t. e, e' are in the form $(x = y)$, x, y being variables, we have that

$$\{l_1, \dots, l_k\} \models_{\mathcal{T}} (e \vee e') \iff \{l_1, \dots, l_k\} \models_{\mathcal{T}} e \text{ or } \{l_1, \dots, l_k\} \models_{\mathcal{T}} e'.$$

Notice that \mathcal{EUF} , $\mathcal{DL}(\mathbb{Q})$, $\mathcal{LA}(\mathbb{Q})$ are convex, whereas $\mathcal{DL}(\mathbb{Z})$ and $\mathcal{LA}(\mathbb{Z})$ are not.

Consider two theories $\mathcal{T}_1, \mathcal{T}_2$ with equality and disjoint signatures Σ_1, Σ_2 . An atom ψ is *i-pure* if only $=$, variables and symbols from Σ_i occur in ψ . A formula φ is *pure* iff every atom in φ is *i-pure* for some $i \in \{1, 2\}$. Every non-pure $\mathcal{T}_1 \cup \mathcal{T}_2$ formula φ can be converted into an equivalently satisfiable pure formula φ' by recursively labeling terms t with fresh variables v_t , and by adding the atom $(v_t = t)$. E.g.:

$$(f(x+3y) = g(2x-y)) \Rightarrow (f(v_{x+3y}) = g(v_{2x-y})) \wedge (v_{x+3y} = x+3y) \wedge (v_{2x-y} = 2x-y).$$

This process is called *purification*, and is linear in the size of the input formula. Thus, henceforth we assume w.l.o.g. that all input formulas $\varphi \in \mathcal{T}_1 \cup \mathcal{T}_2$ are pure.

If φ is a pure $\mathcal{T}_1 \cup \mathcal{T}_2$ formula, then v is an *interface variable* for φ iff it occurs in both 1-pure and 2-pure atoms of φ . An equality $(v_i = v_j)$ is an *interface equality* for φ iff v_i, v_j are interface variables for φ . We assume a unique representation for $(v_i = v_j)$ and $(v_j = v_i)$. Henceforth we denote the interface equality $(v_i = v_j)$ by " e_{ij} ".

Given a \mathcal{T} -inconsistent set of literals $L = \{l_1, \dots, l_n\}$ in a theory \mathcal{T} , a *conflict set* η is an (\mathcal{T} -)inconsistent subset of L . η is *minimal* if none of its strict subsets is \mathcal{T} -inconsistent. We say that η is *$\neg e_{ij}$ -minimal* iff $\eta \setminus \{\neg e_{ij}\}$ is no more \mathcal{T} -inconsistent, for every $\neg e_{ij} \in \eta$.

```

function Bool+ $\mathcal{T}$  ( $\varphi$ : quantifier-free formula)
1    $\mathcal{A}^p \leftarrow \mathcal{T}2\mathcal{B}(\text{Atoms}(\varphi))$ 
2    $\varphi^p \leftarrow \mathcal{T}2\mathcal{B}(\varphi)$ 
3   while Bool-satisfiable( $\varphi^p$ ) do
4      $\mu^p \leftarrow \text{pick\_total\_assign}(\mathcal{A}^p, \varphi^p)$ 
5      $(\rho, \pi) \leftarrow \mathcal{T}\text{-satisfiable}(\mathcal{B}2\mathcal{T}(\mu^p))$ 
6     if  $\rho = \text{sat}$  then return sat
7      $\varphi^p \leftarrow \varphi^p \wedge \mathcal{T}2\mathcal{B}(\neg\pi)$ 
8   end while
9   return unsat
end function

```

Fig. 1. A simplified view of enumeration-based \mathcal{T} -satisfiability procedure: Bool+ \mathcal{T}

A \mathcal{T} -solver is a procedure that decides the consistency of an assignment μ in \mathcal{T} . An (propositional) assignment μ for a formula φ is a function $\mu : \text{Atoms}(\varphi) \mapsto \{\text{true}, \text{false}\}$. μ can be equivalently represented as a set of literals μ_S , where $\neg A \in \mu_S$ if $\mu(A) = \text{false}$, and $A \in \mu_S$ otherwise. μ can also equivalently be seen as a formula μ_φ , built as the conjunction of the literals in the set μ_S . (In the following, we will denote all such equivalent representations with μ . Moreover, we will denote with $\mu_{\mathcal{T}_i}$ the subassignment of μ containing only i -pure literals.) When a \mathcal{T} -solver detects the inconsistency of μ , it also returns a conflict set η of μ . Finally, we also require every \mathcal{T} -solver involved in either the NO schema or DTC to be *incremental* (it does not need to restart the computation from scratch to decide the satisfiability of μ' if it had already proved that of $\mu \subset \mu'$) and *backtrackable* (it can return to a previous state in an efficient manner) [11].

We say that a \mathcal{T} -solver is $\neg e_{ij}$ -minimal (resp. *minimal*) if the conflict sets it returns are always $\neg e_{ij}$ -minimal (resp. *minimal*). Notice that $\neg e_{ij}$ -minimality is a much weaker requirement than minimality.

2.2 Satisfiability Modulo Theory

Fig. 1 presents Bool+ \mathcal{T} , a (much simplified) decision procedure for $SMT(\mathcal{T})$. The function $\text{Atoms}(\varphi)$ takes a ground formula φ and returns the set of atoms which occur in φ . We use the notation φ^p to denote the *propositional abstraction* of φ , which is formed by the function $\mathcal{T}2\mathcal{B}$ that maps propositional variables to themselves, ground atoms into fresh propositional variables, and is homomorphic w.r.t. boolean operators and set inclusion. The function $\mathcal{B}2\mathcal{T}$ is the inverse of $\mathcal{T}2\mathcal{B}$. We use μ^p to denote a propositional assignment. (If $\mathcal{T}2\mathcal{B}(\mu) \models \mathcal{T}2\mathcal{B}(\varphi)$, then we say that μ *propositionally satisfies* φ .) The idea underlying the algorithm is that the truth assignments for the propositional abstraction of φ are enumerated and checked for satisfiability in \mathcal{T} . The procedure either returns sat if one such model is found, or returns unsat otherwise. The function pick_total_assign returns a total assignment to the propositional variables in φ^p , that is, it assigns a truth value to all variables in \mathcal{A}^p . The function \mathcal{T} -satisfiable(μ) detects if the set of conjuncts μ is \mathcal{T} -satisfiable: if so, it returns (sat, \emptyset); otherwise, it

returns (unsat, π) , where $\pi \subseteq \mu$ is a \mathcal{T} -unsatisfiable set, called a *theory conflict set*. We call the negation of a conflict set, a *conflict clause*.

The algorithm is a coarse abstraction of the ones underlying TSAT++, MATHSAT, DLSAT, DPLL(T)/BarceLogic, CVCLITE, and ICS/YICES. The test for satisfiability and the extraction of the corresponding truth assignment are kept separate in this description only for the sake of simplicity.

In practice, the enumeration of truth assignments is carried out by means of efficient implementations of the DPLL algorithm [15], where a *partial assignment* μ^p is built incrementally, each time selecting an unassigned literal l (*literal selection*), called *decision literal*, according to some heuristic criterion, adding it to μ^p and performing all the other assignments which derive deterministically from this choice (*unit propagation*). When some assignment μ^p falsifies the formula returning a (boolean) conflict set π^p , or when \mathcal{T} -satisfiable($\mathcal{B}2\mathcal{T}(\mu^p)$) fails returning a theory conflict set π , the negation $\neg\pi^p$ of (the boolean abstraction of) the conflict set is passed as a conflict clause to the boolean solver. Then $\neg\pi^p$ is added in conjunction to φ^p either temporarily or permanently (*learning*), and the algorithm backtracks up to the highest point in the search where a literal can be unit-propagated on $\neg\pi^p$ (*backjumping*). Learning also avoids generating the same conflicts in future branches.

An important variant [10] is that of building from $\neg\pi^p$ a “mixed boolean+theory conflict clause”, by recursively removing non-decision literals l from the conflict clause by resolving the latter with the clause C_l which caused the unit-propagation of l ; this is done until the conflict clause contains only decision literals (*last-UIP strategy*) or at most one non-decision literal assigned after the last decision (*first-UIP strategy*).²

Another important improvement is *early pruning (EP)*: before every literal selection, intermediate assignments are checked for \mathcal{T} -satisfiability and, if not \mathcal{T} -satisfiable, they are pruned (since no refinement can be \mathcal{T} -satisfiable). Finally, *theory deduction* can be used to reduce the search space by allowing the \mathcal{T} -solvers to explicitly return truth values for unassigned literals, which can be unit-propagated by the SAT solver. The interested reader is pointed to, e.g., [1, 4, 10, 5] for details and further references.

2.3 Nelson-Oppen’s schema

Given two signature-disjoint stably infinite theories \mathcal{T}_1 and \mathcal{T}_2 , the Nelson-Oppen combination schema [11], in the following referred to as NO, allows for solving the satisfiability problem for $\mathcal{T}_1 \cup \mathcal{T}_2$ (i.e. the problem of checking the $\mathcal{T}_1 \cup \mathcal{T}_2$ -satisfiability of sets of $\Sigma_1 \cup \Sigma_2$ -literals) by using the satisfiability procedures for \mathcal{T}_1 and \mathcal{T}_2 . The procedure is basically a structured interchange of information inferred from either theory and propagated to the other, until convergence is reached. The schema requires the exchange of information, the kind of which depends on the *convexity* of the involved theories. In the case of convex theories, the two solvers communicate to each other single interface equalities. In the case of non-convex theories, the NO schema becomes more complicated, because the two solvers need to exchange *arbitrary disjunctions of interface equalities*, which have to be managed within the decision procedure by means of case

² These are standard techniques implemented in most SAT solvers in order to build the boolean conflict clauses [14].




Fig. 2. Representation of the search tree for the formula of Example 1

splitting and of *backtrack* search. In the latter case, the NO schema performs a number of *branches* to check the consistency of a set of literals which depends on how many disjunctions of equalities are exchanged at each step: if the current set of literals is μ , and one of the \mathcal{T}_i -solver sends the disjunction $\bigvee_{k=1}^n (e_{ij})_k$ to the other, the latter must further investigate up to n branches to check the consistency of each of the $\mu \cup \{(e_{ij})_k\}$ sets separately.

We notice that the ability to carry out deductions of (disjunctions of) e_{ij} 's is crucial for NO: if the current set of literals $\mu_{\mathcal{T}_i}$ in input to each \mathcal{T}_i -solver is \mathcal{T}_i -consistent, then \mathcal{T}_i -solver must be able to derive the (disjunctions of) interface equalities e_{ij} which are entailed by $\mu_{\mathcal{T}_i}$ in \mathcal{T}_i (if any), or say that no (disjunction of) e_{ij} 's is entailed if this is the case. If the \mathcal{T}_i -solver is always capable of doing this, we say that it is *e_{ij} -deduction complete*. In what follows we will assume that all the \mathcal{T}_i -solvers used in a NO schema are *e_{ij} -deduction complete*.

Example 1 (convex case). Consider the following $\mathcal{EUF} \cup \mathcal{LA}(\mathbb{Q})$ formula ϕ (see Fig. 2)

$$\begin{aligned}
\mathcal{EUF} : & (v_3 = h(v_0)) \wedge (v_4 = h(v_1)) \wedge (v_6 = f(v_2)) \wedge (v_7 = f(v_5)) \wedge \\
\mathcal{LA}(\mathbb{Q}) : & (v_0 \geq v_1) \wedge (v_0 \leq v_1) \wedge (v_2 = v_3 - v_4) \wedge (RESET_5 \rightarrow (v_5 = 0)) \wedge \\
\text{Both} : & (\neg RESET_5 \rightarrow (v_5 = v_8)) \wedge \neg(v_6 = v_7).
\end{aligned} \tag{1}$$

$v_0, v_1, v_2, v_3, v_4, v_5$ are interface variables, v_6, v_7, v_8 are not. (Thus, e.g., $(v_0 = v_1)$ is an interface equality, whilst $(v_0 = v_6)$ is not.) $RESET_5$ is a boolean variable.

After the first run of unit propagations, assume DPLL selects the literal $RESET_5$, resulting in the assignment

$$\mu = \{ (v_3 = h(v_0)), (v_4 = h(v_1)), (v_6 = f(v_2)), (v_7 = f(v_5)), (v_0 \geq v_1), (v_0 \leq v_1), (v_2 = v_3 - v_4), \neg(v_6 = v_7), RESET_5, (v_5 = 0) \}, \tag{2}$$

which propositionally satisfies ϕ . Now, the set of literals $\mu_{\mathcal{EUF}} \subset \mu$ is given to the \mathcal{EUF} solver, which reports its consistency and deduces no new interface equality. Then the set $\mu_{\mathcal{LA}(\mathbb{Q})} \subset \mu$ is given to the $\mathcal{LA}(\mathbb{Q})$ solver, which reports consistency and deduces the interface equality $(v_0 = v_1)$, which is passed to the \mathcal{EUF} solver. The new set $\mu_{\mathcal{EUF}} \cup \{(v_0 = v_1)\}$ is still \mathcal{EUF} -consistent, but this time the \mathcal{EUF} solver deduces the equality




Fig. 3. Representation of the search tree for the formula of Example 2

($v_3 = v_4$), which is in turn passed to the $\mathcal{L}\mathcal{A}(\mathbb{Q})$ solver, that now as a consequence of this and the assignment $\mu_{\mathcal{L}\mathcal{A}(\mathbb{Q})}$ deduces ($v_2 = v_5$). The $\mathcal{E}\mathcal{U}\mathcal{F}$ solver is then invoked again to check the $\mathcal{E}\mathcal{U}\mathcal{F}$ -consistency of the assignment $\mu_{\mathcal{E}\mathcal{U}\mathcal{F}} \cup \{(v_0 = v_1), (v_2 = v_5)\}$: since this check fails, the Nelson-Oppen method reports the $\mathcal{E}\mathcal{U}\mathcal{F} \cup \mathcal{L}\mathcal{A}(\mathbb{Q})$ -unsatisfiability of φ under the whole assignment μ . At this point, then, DPLL backtracks and tries assigning false to $RESET_5$, resulting in the new assignment

$$\mu = \{ (v_3 = h(v_0)), (v_4 = h(v_1)), (v_6 = f(v_2)), (v_7 = f(v_5)), (v_0 \geq v_1), (v_0 \leq v_1), (v_2 = v_3 - v_4), \neg(v_6 = v_7), \neg RESET_5, (v_5 = v_8) \},$$

which is found $\mathcal{E}\mathcal{U}\mathcal{F} \cup \mathcal{L}\mathcal{A}(\mathbb{Q})$ -satisfiable (see Fig. 2). \diamond

Example 2 (non-convex case). Consider the following $\mathcal{E}\mathcal{U}\mathcal{F} \cup \mathcal{L}\mathcal{A}(\mathbb{Z})$ formula φ

$$\begin{aligned} \mathcal{E}\mathcal{U}\mathcal{F} : & \neg(f(v_1) = f(v_2)) \wedge \neg(f(v_2) = f(v_4)) \wedge (f(v_3) = v_5) \wedge (f(v_1) = v_6) \wedge \\ \mathcal{L}\mathcal{A}(\mathbb{Z}) : & (v_1 \geq 0) \wedge (v_1 \leq 1) \wedge (v_5 = v_4 - 1) \wedge (v_3 = 0) \wedge (v_4 = 1) \wedge \\ & (v_2 \geq v_6) \wedge (v_2 \leq v_6 + 1). \end{aligned} \quad (3)$$

Here (see Fig. 3) all the variables (v_1, \dots, v_6) are interface ones. φ contains only unit clauses, so after the first run of unit propagations, DPLL generates the assignment μ which is simply the set of literals in φ . The Nelson-Oppen combination schema then runs as follows. First, the sub-assignment $\mu_{\mathcal{E}\mathcal{U}\mathcal{F}}$ is given to the $\mathcal{E}\mathcal{U}\mathcal{F}$ solver, which reports its consistency and deduces no interface equality. Then, the sub-assignment $\mu_{\mathcal{L}\mathcal{A}(\mathbb{Z})}$ is given to the $\mathcal{L}\mathcal{A}(\mathbb{Z})$ solver, which reports its consistency and deduces the disjunction $(v_1 = v_3) \vee (v_1 = v_4)$. Next, there is a case-splitting and the two equalities $(v_1 = v_3)$ and $(v_1 = v_4)$ are passed to the $\mathcal{E}\mathcal{U}\mathcal{F}$ solver. The first branch, corresponding to selecting $(v_1 = v_3)$, is opened: then the set $\mu_{\mathcal{E}\mathcal{U}\mathcal{F}} \cup \{(v_1 = v_3)\}$ is $\mathcal{E}\mathcal{U}\mathcal{F}$ -consistent, and the equality $(v_5 = v_6)$ is deduced. After that, the assignment $\mu_{\mathcal{L}\mathcal{A}(\mathbb{Z})} \cup \{(v_5 = v_6)\}$ is passed to the $\mathcal{L}\mathcal{A}(\mathbb{Z})$ solver, that reports its consistency and deduces another disjunction, $(v_2 = v_3) \vee (v_2 = v_4)$. At this point, another case-splitting is needed in the $\mathcal{E}\mathcal{U}\mathcal{F}$ solver, resulting in the two branches $\mu_{\mathcal{E}\mathcal{U}\mathcal{F}} \cup \{(v_1 = v_3), (v_2 = v_3)\}$

```

function Bool+ $\mathcal{T}_1+\mathcal{T}_2$  ( $\varphi_i$ : quantifier-free formula)
1    $\varphi \leftarrow \text{purify}(\varphi_i)$ 
2    $\mathcal{A}^p \leftarrow \mathcal{T}2\mathcal{B}(\text{Atoms}(\varphi) \cup \text{interface\_equalities}(\varphi))$ 
3    $\varphi^p \leftarrow \mathcal{T}2\mathcal{B}(\varphi)$ 
4   while Bool-satisfiable ( $\varphi^p$ ) do
5      $\mu_1^p \wedge \mu_2^p \wedge \mu_e^p = \mu^p \leftarrow \text{pick\_total\_assign}(\mathcal{A}^p, \varphi^p)$ 
6      $(\rho_1, \pi_1) \leftarrow \mathcal{T}_1\text{-satisfiable}(\mathcal{B}2\mathcal{T}(\mu_1^p \wedge \mu_e^p))$ 
7      $(\rho_2, \pi_2) \leftarrow \mathcal{T}_2\text{-satisfiable}(\mathcal{B}2\mathcal{T}(\mu_2^p \wedge \mu_e^p))$ 
8     if ( $\rho_1 = \text{sat} \wedge \rho_2 = \text{sat}$ ) then return sat else
9       if  $\rho_1 = \text{unsat}$  then  $\varphi^p \leftarrow \varphi^p \wedge \mathcal{T}2\mathcal{B}(\neg\pi_1)$ 
10      if  $\rho_2 = \text{unsat}$  then  $\varphi^p \leftarrow \varphi^p \wedge \mathcal{T}2\mathcal{B}(\neg\pi_2)$ 
11  end while
12  return unsat
end function

```

Fig. 4. A simplified view of the Delayed Theory Combination procedure for $SMT(\mathcal{T}_1 \cup \mathcal{T}_2)$

and $\mu_{\mathcal{E}\mathcal{U}\mathcal{F}} \cup \{(v_1 = v_3), (v_2 = v_4)\}$. Both of them are found inconsistent, so the whole branch previously opened by the selection of $(v_1 = v_3)$ is found inconsistent; at this point, the other case of the branch (i.e. the equality $(v_1 = v_4)$) is selected, and since the assignment $\mu_{\mathcal{E}\mathcal{U}\mathcal{F}} \cup \{(v_1 = v_4)\}$ is $\mathcal{E}\mathcal{U}\mathcal{F}$ -consistent and no new interface equality is deduced, the Nelson-Oppen method reports the $\mathcal{E}\mathcal{U}\mathcal{F} \cup \mathcal{L}\mathcal{A}(\mathbb{Z})$ -satisfiability of φ under the whole assignment μ . \diamond

3 SMT for combined theories via Delayed Theory Combination

In the Delayed Theory Combination (DTC) schema [3, 6], the $SMT(\mathcal{T}_1 \cup \mathcal{T}_2)$ problem is tackled in a different way: instead of using the Nelson-Oppen schema to combine two decision procedures for \mathcal{T}_1 and \mathcal{T}_2 by exchange of (disjunctions of) interface equalities, each of the two \mathcal{T}_i solvers works in isolation, without direct exchange of information. Their mutual consistency is ensured by augmenting the input problem with all interface equalities e_{ij} , even if these do not occur in the original problem. The enumeration of assignments includes not only the atoms in the formula, but also the interface equalities e_{ij} . Both theory solvers receive, from the boolean level, the same truth assignment μ_e for e_{ij} : under such conditions, the two “partial” models found by each decision procedure can be merged into a model for the input formula.

A simplified view of the algorithm is presented in Fig. 4. Initially (lines 1–3), the formula is purified, the *new* e_{ij} ’s are created and added to the set of propositional symbols \mathcal{A}^p , and the propositional abstraction φ^p of φ is created. Then, the main loop is entered (lines 4–11): while φ^p is propositionally satisfiable (line 4), a satisfying truth assignment μ^p is selected (line 5). It is important to stress that truth values are associated not only to atoms in φ , but also to the e_{ij} atoms, even though they do not occur in φ . μ^p is then (implicitly) separated into $\mu_1^p \wedge \mu_e^p \wedge \mu_2^p$, where $\mathcal{B}2\mathcal{T}(\mu_i^p)$ is a set of i -pure literals and $\mathcal{B}2\mathcal{T}(\mu_e^p)$ is a set of e_{ij} -literals. The relevant parts of μ^p are checked for consistency against each theory (lines 6–7); \mathcal{T}_i -satisfiable(μ) returns a pair (ρ_i, π_i) , where

ρ_i is unsat iff μ is unsatisfiable in \mathcal{T}_i , and sat otherwise. If both calls to \mathcal{T}_i -satisfiable return sat, then the formula is satisfiable. Otherwise, when ρ_i is unsat, then π_i is a theory conflict set, i.e. $\pi_i \subseteq \mu$ and π_i is \mathcal{T}_i -unsatisfiable. Then, φ^p is strengthened to exclude truth assignments which may fail in the same way (line 9–10), and the loop is resumed. Unsatisfiability is returned (line 12) when the loop is exited without having found a model.

In practical implementations of DTC, the search for a satisfactory assignment is based on a modern DPLL engine, performing literal selection, unit-propagation, back-jumping and learning, early pruning, and theory deduction, as explained in §2.2. In particular, DTC can be enhanced by e_{ij} -deduction, in which e_{ij} 's can be deduced by the \mathcal{T}_i -solvers and hence unit-propagated. (See Step 4. of Strategy1 below). We refer the reader to [3, 6] for a more detailed discussion.

Notation-wise, we call “new” e_{ij} 's all the interface equalities e_{ij} 's which do not occur in any clause of the input formula φ (including all the clauses learned). Moreover, we often write sets of literals $\{l_1, \dots, l_n\}$ as conjunctions $l_1 \wedge \dots \wedge l_n$, and we often write clauses $(\bigvee_i l_i) \vee (\bigvee_j l_j)$ as implications: $(\bigwedge_i \neg l_i) \rightarrow (\bigvee_j l_j)$ or $(\bigwedge_i \neg l_i \wedge \bigwedge_j \neg l_j) \rightarrow \perp$.

Hereafter, for the sake of proving the theoretical results in §4 and §5, we assume that DTC implements the following strategy.

Strategy 1 (NO emulation)

1. All the conflict clauses derived by theory conflicts are learned, either temporarily or permanently.³
2. Each conflict clause in 1. is a mixed boolean+theory conflict clause which is built from the theory conflict set by means of the last-UIP strategy described in §2.2.⁴
3. The literal selection heuristic and the \mathcal{T}_i -solvers calls are such that:
 - (i) new e_{ij} 's are selected only after all the other literals have been assigned,
 - (ii) Early pruning (EP) is applied before every selection of a new e_{ij} ,⁵
 - (iii) the new e_{ij} 's selected are always assigned false,
 - (iv) each \mathcal{T}_i -solver is invoked only if at least one literal (which has not been deduced singularly by \mathcal{T}_i -solver itself) has been added to its input since the last call.⁶
4. At every early-pruning call on a branch (namely μ) which is found both \mathcal{T}_1 - and \mathcal{T}_2 -consistent, if one \mathcal{T}_i -solver performs the e_{ij} -deduction $\mu^* \models_{\mathcal{T}_i} \bigvee_{j=1}^k e_j$, s.t. $\mu^* \subseteq \mu_{\mathcal{T}_i}$, each e_j being an unassigned interface equality on variables in μ , then:
 - (i) the clause $\mathcal{T}2\mathcal{B}(\mu^* \rightarrow \bigvee_{j=1}^k e_j)$ is learned immediately;

³ That is, if one \mathcal{T}_i -solver returns a conflict set π , then the conflict clauses $\mathcal{T}2\mathcal{B}(-\pi)$ is always added to φ^p , either temporarily or permanently.

⁴ That is, each conflict clause contains all and only (the negation of) the decision literals which forced the unit-propagation or the e_{ij} -deduction of those in the theory conflict.

⁵ That is, before selecting and adding a new (negated) e_{ij} to μ , the \mathcal{T}_i -satisfiability of μ is checked for both \mathcal{T}_i 's by calling the \mathcal{T}_i -solver's. If μ is found \mathcal{T}_i -inconsistent for some \mathcal{T}_i , then the procedure backtracks.

⁶ This avoids invoking a \mathcal{T}_i -solver twice in sequence on the same input. The restriction “which ... by \mathcal{T}_i -solver itself” means that, if \mathcal{T}_i -solver (μ) returns “Sat” and deduces e_{ij} , then \mathcal{T}_i -solver is not invoked on $\mu \cup \{e_{ij}\}$.

- (ii) if $k = 1$, then e_k is added to the current assignment and unit-propagated immediately;
 - (iii) if $k > 1$, then $\neg e_1, \dots, \neg e_k$ are put on the top of the literal selection list, so that to be the next $\neg e_{ij}$'s selected by the literal selection heuristic.
5. **[If and only if both \mathcal{T}_i -solvers are e_{ij} -deduction complete]**
 If a total assignment μ which propositionally satisfies φ is found \mathcal{T}_i -satisfiable for both \mathcal{T}_i 's, and neither \mathcal{T}_i -solver performs any e_{ij} -deduction from μ , then DTC stops returning "Sat".⁷

4 DTC with non e_{ij} -deduction-complete \mathcal{T}_i -solvers vs. NO

In this section, we assume that both the \mathcal{T}_i -solvers employed by DTC are $\neg e_{ij}$ -minimal and have limited or no e_{ij} -deduction capabilities. Under these assumptions, we have the following result.

Theorem 1. *Let \mathcal{T}_1 and \mathcal{T}_2 be two stably-infinite (possibly non-convex) theories. Let both \mathcal{T}_i -solvers be $\neg e_{ij}$ -minimal, and possibly have some e_{ij} -deduction capabilities; let φ be a pure $\mathcal{T}_1 \cup \mathcal{T}_2$ formula and let μ be a total assignment propositionally satisfying φ . Let DTC with Strategy 1 prove the $\mathcal{T}_1 \cup \mathcal{T}_2$ -consistency (resp. $\mathcal{T}_1 \cup \mathcal{T}_2$ -inconsistency) of μ , returning a conflict set η in the case of inconsistency. Let dtc_br and dtc_ded be the number of boolean branches and of e_{ij} -deductions performed in the DTC proof. Then we have:*

$$dtc_br + dtc_ded \leq no_br + no_ded, \quad (4)$$

no_ded and no_br being respectively the number of deductions and of branches performed by a corresponding NO proof of the $\mathcal{T}_1 \cup \mathcal{T}_2$ -consistency (resp. $\mathcal{T}_1 \cup \mathcal{T}_2$ -inconsistency) of μ .

Theorem 1 states that, if the \mathcal{T}_i -solvers are both $\neg e_{ij}$ -minimal, then there is a strategy for DTC which emulates some NO proof (even though the \mathcal{T}_i -solvers have limited or no e_{ij} -deduction capabilities!) at the cost of (at most) one extra boolean branch for every e_{ij} -deduction performed by NO. Therefore the (possibly very expensive) e_{ij} -deduction steps of the NO schema can be avoided at the cost of one extra boolean branch each.

More generally, we notice that one key idea in the proof of Theorem 1 is that, when the DPLL engine fails and generates a conflict set π , it backjumps up to the second-most-recently-assigned $\neg e_{ij}$ in π , if any [7]. (See, e.g., the case of C_{23} in Figure 6.) Therefore, in a more general case than that of Theorem 1 (no $\neg e_{ij}$ -minimality), *the more redundant $\neg e_{ij}$'s the \mathcal{T}_i -solvers are able to remove from the conflict set returned, the more boolean branches are skipped by backjumping.*

We consider some relevant subcases. Assume both \mathcal{T}_i -solvers have no deduction capabilities. If so, $dtc_ded = 0$. If $\mathcal{T}_1, \mathcal{T}_2$ are both convex theories, then NO proves the $\mathcal{T}_1 \cup \mathcal{T}_2$ -(un)satisfiability of μ by a chain of e_{ij} -deductions in one branch (i.e., $no_br = 1$). From Theorem 1 we have $dtc_br \leq no_ded + 1$.

⁷ Step 5. is identical to the $\mathcal{T}_1 \cup \mathcal{T}_2$ -satisfiability termination condition of NO.




Fig. 5. DTC execution of Example 3 on $\mathcal{LAr}(\mathbb{Q}) \cup \mathcal{EUF}$, with no e_{ij} -deduction. The clauses C_{ij} 's are the same as those of Fig. 7.

Example 3 (no e_{ij} -deduction, convex case). Consider the $\mathcal{EUF} \cup \mathcal{LAr}(\mathbb{Q})$ formula φ (1) and the assignment μ (2) of Example 1. Look at Fig. 5. Both $\mu_{\mathcal{LAr}(\mathbb{Q})}$ and $\mu_{\mathcal{EUF}}$ are found consistent in the respective theories by the respective solvers.

Then DTC starts selecting the new negated e_{ij} 's, and proceeds without causing conflicts, until it selects $\neg(v_0 = v_1)$, which causes a $\mathcal{LAr}(\mathbb{Q})$ conflict. As $\mathcal{LAr}(\mathbb{Q})$ is convex and $\mathcal{LAr}(\mathbb{Q})$ -Solver is $\neg e_{ij}$ -minimal, it returns a conflict set in the form $\mu'_{\mathcal{LAr}(\mathbb{Q})} \cup \{\neg(v_0 = v_1)\}$ s.t. $\{(v_0 \geq v_1), (v_0 \leq v_1)\} \subseteq \mu'_{\mathcal{LAr}(\mathbb{Q})} \subseteq \mu_{\mathcal{LAr}(\mathbb{Q})}$. Thus DTC learns the corresponding clause C_{01} and backjumps up to μ (or even higher), hence unit propagating $(v_0 = v_1)$.

What happens next depends on whether the learned clause C_{01} contains the redundant $\mathcal{LAr}(\mathbb{Q})$ atom $(v_5 = 0)$ or not. Here (Fig. 5) we consider the “worst” case, that is, when such atom is present in C_{01} . This means that DTC backjumps after the unit-propagation of $(v_5 = 0)$ ⁸. Then $(v_0 = v_1)$ is unit-propagated and new unassigned $\neg e_{ij}$'s are selected again, until $\neg(v_3 = v_4)$ generates another conflict represented by clause C_{34} , which causes backjumping and unit-propagating $(v_3 = v_4)$. The same is repeated for $(v_2 = v_5)$. Then $\mu \cup \{(v_0 = v_1), (v_3 = v_4), (v_2 = v_5)\}$ is found \mathcal{EUF} -inconsistent s.t. the conflict is represented by the clause C_{67} , and the whole procedure backtracks, causing the unit-propagation of $\neg RESET_5$ and $(v_5 = v_8)$ as in Example 5.

We notice that, so far, the whole process mimics the NO deduction process of the first branch in Example 1 (see Fig. 2), requiring a number of extra boolean branches equal to the number of deductions performed by the corresponding NO process ($d_{tc_br} = 4$, $d_{tc_ded} = 0$, $no_br = 1$ and $no_ded = 3$.)

⁸ More precisely, by Step 2. of Strategy 1, DTC eliminates $(v_5 = 0)$ from the conflict clause C_{01} by resolving the latter with the clause $RESET_5 \rightarrow (v_5 = 0)$ in φ , and thus by introducing $RESET_5$ in the conflict clause. Then it backjumps at the point before it started selecting $\neg e_{ij}$'s.




Fig. 6. DTC execution of Example 4 on $\mathcal{LA}(\mathbb{Z}) \cup \mathcal{EUF}$, with no e_{ij} -deduction. The clauses C_{ij} 's are the same as those of Fig. 8.

In the other branch of the boolean search (after the unit propagation of $\neg RESET_5$ and $(v_5 = v_8)$), since we are supposing that $(v_5 = 0)$ is present in C_{01} , then this clause can't be used to unit-propagate $(v_0 = v_1)$, and DTC starts selecting $\neg e_{ij}$'s until $\neg(v_0 = v_1)$ causes again a $\mathcal{LA}(\mathbb{Q})$ -conflict. At this point the new clause C'_{01} (similar to C_{01} but possibly involving $(v_5 = v_8)$ instead of $(v_5 = 0)$) is learned and (after backjumping) $(v_0 = v_1)$ is unit-propagated on it. However, as a consequence of this unit-propagation, $(v_3 = v_4)$ is unit-propagated on C_{34} ⁹, and after this step DTC starts selecting unassigned $\neg e_{ij}$'s again, this time without generating conflicts, so that to conclude that the formula is $\mathcal{T}_1 \cup \mathcal{T}_2$ -satisfiable. (Comparing with the right branch in Fig. 2, we have $d_{tc_br} = 2$, $d_{tc_ded} = 0$, $no_br = 1$ and $no_ded = 2$.) \diamond

Notice that the three leftmost diagonal branches in Fig. 5 obtain the same effect as the e_{ij} -deduction steps in Fig. 7 (and in Fig. 2).

As before, the fact that the DPLL solver is aware a priori of the e_{ij} 's allows it for learning clauses containing e_{ij} 's, which can be used in subsequent branches in order to prune the search and to avoid redoing the same deductions from scratch.

We consider now the case where some \mathcal{T}_i 's are non-convex.

Example 4 (no e_{ij} -deduction, non-convex case). Consider the $\mathcal{EUF} \cup \mathcal{LA}(\mathbb{Z})$ formula φ (3) and the assignment μ of Example 2. Look at Fig. 6. Both $\mu_{\mathcal{LA}(\mathbb{Z})}$ and $\mu_{\mathcal{EUF}}$ are found consistent in the respective theories by the respective solvers.

Then DTC starts selecting new $\neg e_{ij}$'s, and proceeds without causing conflicts, until it selects $\neg(v_1 = v_4)$ and $\neg(v_1 = v_3)$, which cause a $\mathcal{LA}(\mathbb{Z})$ conflict. The branch is in the form $\mu \cup \bigcup_j \neg e_j$, so that, the $\neg e_{ij}$ -minimal conflict set η_{13} returned is in the form $\mu'_{\mathcal{LA}(\mathbb{Z})} \cup \{\neg(v_1 = v_3), \neg(v_1 = v_4)\}$. Thus DTC learns the corresponding clause

⁹ Here we are sure that $(v_5 = 0)$ does not appear on C_{34} , since it was learned as a consequence of a \mathcal{EUF} conflict.

C_{13} (see Fig 6) and backjumps up to the highest point which allows for unit-propagating ($v_1 = v_3$) on C_{13} , and performs such unit propagation. Then DTC starts and proceeds selecting new $\neg e_{ij}$'s without causing conflicts, until it selects $\neg(v_5 = v_6)$, which causes a \mathcal{EUF} conflict represented by the clause C_{56} . As \mathcal{EUF} is convex, $\neg(v_5 = v_6)$ is the only $\neg e_{ij}$ occurring in the conflict set, so that DTC backtracks over the last chain of $\neg e_{ij}$'s and unit-propagates ($v_5 = v_6$).

Again, DTC selects a chain of new $\neg e_{ij}$'s without causing conflicts, until it selects $\neg(v_2 = v_4)$ and $\neg(v_2 = v_3)$, which cause a $\mathcal{LA}(\mathbb{Z})$ conflict represented by clause C_{23} . As before, it backjumps to the highest point where it can unit-propagate ($v_2 = v_3$). Performing the latter unit propagation causes a \mathcal{EUF} conflict, learning the clause C_{24} . By applying Step 2. of Strategy 1, resolving on literals ($v_2 = v_3$), ($v_5 = v_6$), ($v_1 = v_3$) the conflicting clause C_{24} with the clauses C_{23} , C_{56} and C_{13} , (which caused the unit-propagation of ($v_2 = v_3$), ($v_5 = v_6$) and ($v_1 = v_3$) respectively), DTC obtains a clause $C'_{24} : (\mu'_{\mathcal{LA}(\mathbb{Z})} \wedge \mu'_{\mathcal{EUF}} \wedge \mu''_{\mathcal{LA}(\mathbb{Z})} \wedge \mu''_{\mathcal{EUF}}) \rightarrow ((v_1 = v_4) \vee (v_2 = v_4))$, which allows it for backjumping over all the remaining $\neg e_{ij}$'s of the current chain and unit-propagating ($v_2 = v_4$).

The latter causes a new \mathcal{EUF} conflict represented by the clause C_{14} . By Step 2. of Strategy 1, C_{14} is resolved with the clauses C'_{24} , C_{13} (which caused the propagation of ($v_2 = v_4$) and ($v_1 = v_3$) respectively), obtaining the clause $C'_{14} : (\mu'_{\mathcal{LA}(\mathbb{Z})} \wedge \mu''_{\mathcal{LA}(\mathbb{Z})} \wedge \mu'_{\mathcal{EUF}} \wedge \mu''_{\mathcal{EUF}} \wedge \mu'''_{\mathcal{EUF}}) \rightarrow (v_1 = v_4)$, which allows for backjumping up to μ and unit-propagating ($v_1 = v_4$).

Finally, DTC starts and proceeds selecting $\neg e_{ij}$'s (possibly unit-propagating some value due to the clauses learned) without generating conflicts, so that to conclude that the formula is $\mathcal{T}_1 \cup \mathcal{T}_2$ -satisfiable.

Comparing with Fig. 3, $dtc_br = 6$, $dtc_ded = 0$, $no_ded = 3$ and $no_br = 3$. \diamond

Notice that the three leftmost diagonal branches in Fig. 6 obtain the same effect as the e_{ij} -deduction steps in Fig. 8 (and in Fig. 3).

5 DTC with e_{ij} -deduction-complete \mathcal{T}_i -solvers vs. NO

In this section, we assume that both the \mathcal{T}_i -solvers employed by DTC are e_{ij} -deduction complete. Under these assumptions, we have the following result.

Theorem 2. *Let \mathcal{T}_1 and \mathcal{T}_2 be two stably-infinite (possibly non-convex) theories and let both \mathcal{T}_i -solvers be e_{ij} -deduction complete; let φ be a pure $\mathcal{T}_1 \cup \mathcal{T}_2$ formula and let μ be a total assignment propositionally satisfying φ . Let DTC with Strategy 1 prove the $\mathcal{T}_1 \cup \mathcal{T}_2$ -consistency (resp. $\mathcal{T}_1 \cup \mathcal{T}_2$ -inconsistency) of μ , returning a conflict set η in the case of inconsistency. Let dtc_br be the number of boolean branches required in the DTC proof. Then we have:*

$$dtc_br \leq no_br \tag{5}$$

no_br being the number of branches performed by a corresponding NO proof of the $\mathcal{T}_1 \cup \mathcal{T}_2$ -consistency (resp. $\mathcal{T}_1 \cup \mathcal{T}_2$ -inconsistency) of μ .




Fig. 7. DTC execution of Example 5 on $\mathcal{LA}(\mathbb{Q}) \cup \mathcal{EUF}$, with e_{ij} -deduction-complete \mathcal{T}_i -solvers.

Theorem 2 states that, under the same hypotheses of e_{ij} -deduction as NO, DTC emulates NO with no extra cost in terms of boolean search.

Example 5 (convex case). Consider again the $\mathcal{EUF} \cup \mathcal{LA}(\mathbb{Q})$ formula ϕ of Example 1. Figure 7 illustrates a DTC execution when both \mathcal{T}_i -solvers are e_{ij} -deduction complete.

On the left branch (when $RESET_5$ is selected), after the unit-propagation of $(v_5 = 0)$, the $\mathcal{LA}(\mathbb{Q})$ solver deduces $(v_0 = v_1)$, and thus by Step 4. (i) of Strategy 1, the clause C_{01} is learned and $(v_0 = v_1)$ is unit-propagated. As a consequence of this, the \mathcal{EUF} solver can deduce $(v_3 = v_4)$, resulting in the learning of C_{34} and the unit-propagation of $(v_3 = v_4)$, which in turn causes the $\mathcal{LA}(\mathbb{Q})$ -deduction of $(v_2 = v_5)$, with the resulting learning of C_{25} and unit-propagation of the deduced equality.

At this point, $\mu''_{\mathcal{EUF}} \cup \{(v_2 = v_5)\}$ ¹⁰ is found \mathcal{EUF} -inconsistent, so that the \mathcal{EUF} -solver returns (the negation of) the clause C_{67} , which is resolved backward with the clauses C_{25} , C_{34} , C_{01} , $\neg(v_6 = v_7)$, and $(RESET_5 \rightarrow (v_5 = 0))$ as explained in Step 2. of Strategy 1, obtaining a mixed theory+boolean conflict clause C'_{67} in the form $(\mu^* \wedge RESET_5) \rightarrow \perp$ s.t. μ^* contains no interface equality. C'_{67} forces DTC to backjump up to the last branching point. Then the execution of the right branch begins with the unit-propagation of $\neg RESET_5$ on C'_{67} and hence of $(v_5 = v_8)$ on $\neg RESET_5 \rightarrow (v_5 = v_8)$, which produces an assignment propositionally satisfying ϕ . The theory solvers are invoked, and the $\mathcal{LA}(\mathbb{Q})$ solver deduces again $(v_0 = v_1)$, learning a clause C'_{01} which is similar to C_{01} except for the fact that it may contain the redundant literal $(v_5 = v_8)$ instead of $(v_5 = 0)$.¹¹ Then $(v_3 = v_4)$ is unit-propagated on C_{34} . At this point, since both theory solvers cannot deduce any new e_{ij} , by Step 5. of Strategy 1 DTC concludes that ϕ is $\mathcal{EUF} \cup \mathcal{LA}(\mathbb{Q})$ -satisfiable. \diamond

Notice that the left branch of the DTC search tree of Figure 7 mimics directly that of the NO execution of Figure 2. The main difference relies on the fact that, unlike with

¹⁰ Hereafter, $\mu'_T, \mu''_T, \mu'''_T$ will denote generic subsets of μ_T , $T \in \{\mathcal{EUF}, \mathcal{LA}(\mathbb{Q}), \mathcal{LA}(\mathbb{Z})\}$.

¹¹ Here we assume the “worst” case in which $\mu'_{\mathcal{LA}(\mathbb{Q})}$ in C_{01} contains the (redundant) literal $(v_5 = 0)$. If this is not the case, then $(v_0 = v_1)$ is directly unit-propagated on C_{01} , without calling the theory solvers.




Fig. 8. DTC execution of Ex 6 on $\mathcal{LA}(\mathbb{Z}) \cup \mathcal{EUF}$, with e_{ij} -deduction-complete \mathcal{T}_i -solvers.

NO, the deduced e_{ij} 's are not exchanged directly by the \mathcal{T}_i -solvers, but rather they are added to the current assignment μ and unit-propagated.

In the right branch, instead, all values are assigned directly by unit-propagation. This fact illustrates one further potential advantage of DTC with respect to NO: the fact that new e_{ij} 's are known a priori to the DPLL engine allows their inclusion in the learned clauses derived by theory conflicts. Thanks to unit-propagation, this makes it possible to assign truth values to them *directly at the boolean level*, without performing the (potentially costly) invocation of the \mathcal{T}_i -solvers. In the traditional NO schema, this fact does not come naturally, because the boolean solver knows nothing about the e_{ij} 's.

We consider now the case where some \mathcal{T}_i 's are non-convex.

Example 6 (non-convex case). Consider the $\mathcal{EUF} \cup \mathcal{LA}(\mathbb{Z})$ formula φ and assignment μ of Example 2. Figure 8 illustrates a DTC execution when both \mathcal{T}_i -solvers are e_{ij} -deduction complete.

The first invocation of the $\mathcal{LA}(\mathbb{Z})$ solver results in deducing of the disjunction $(v_1 = v_4) \vee (v_1 = v_3)$ and learning of the corresponding clause C_{13} . By Step 4.(iii) of Strategy 1, then, $(v_1 = v_4)$ and $(v_1 = v_3)$ are put on the top of the literal selection list. As a consequence, DTC selects $\neg(v_1 = v_4)$, and thanks to C_{13} it immediately unit-propagates $(v_1 = v_3)$. At this point the \mathcal{EUF} solver can deduce $(v_5 = v_6)$, so that the clause C_{56} is learned and the deduced equality is unit-propagated immediately. When $\mu_{\mathcal{LA}(\mathbb{Z})} \cup \{(v_5 = v_6)\}$ is passed to the $\mathcal{LA}(\mathbb{Z})$ solver, this deduces the disjunction $(v_2 = v_4) \vee (v_2 = v_3)$, learning C_{23} . Selecting $\neg(v_2 = v_4)$ results in the unit-propagation of $(v_2 = v_3)$, which in turn causes a \mathcal{EUF} conflict. After the \mathcal{EUF} -solver returns (the negation of) C_{24} , DTC backjumps up to a point where $(v_2 = v_4)$ can be unit-propagated. This results again in an \mathcal{EUF} -conflict, so that the \mathcal{EUF} -solver returns (the negation of) C_{14} , which causes another backjumping up to where $(v_1 = v_4)$ can be unit-propagated. Then, after another invocation to the theory solvers, DTC stops, declaring φ to be $\mathcal{EUF} \cup \mathcal{LA}(\mathbb{Z})$ -satisfiable. \diamond

As with the convex example, notice that the DTC search tree of Figure 8 mimics directly that of the NO execution of Figure 3 (both *dtc.br* and *no.br* are equal to 3.)

6 Conclusions

Theorem 2 shows that, under the same hypotheses of e_{ij} -deduction-completeness as NO, DTC can emulate NO, with no extra boolean search. Theorem 1 shows that, under the hypothesis of $\neg e_{ij}$ -minimality, even \mathcal{T}_i -solvers with limited or no e_{ij} -deduction capabilities allow DTC to emulate NO, at the cost of (at most) one extra boolean branch for every (possibly very expensive) e_{ij} -deduction performed by NO. Both results also highlight the fact that DTC naturally allows for learning clauses containing e_{ij} 's, which can be used in subsequent branches to prune search and avoid redoing the same search/deductions from scratch.

We remark that Strategy 1 has been conceived only for mimicking NO, and by no means it is assumed to be the most efficient strategy for DTC. (E.g., Step 3.(ii) can be substituted with a weakened version of EP [4], and more efficient literal selection strategies might be preferable to Step 3.(i) and (iii).) Some alternatives are currently under investigation, and their theoretical properties and practical performance are subject for future work.

As far as the $\neg e_{ij}$ -minimality hypothesis is concerned, we notice that, at least for theories like \mathcal{EUF} and $\mathcal{LA}(\mathbb{Q})$, there are known decision procedures that fulfill this requirement (see [12] and [4] respectively.) For other theories, the problem of $\neg e_{ij}$ -minimization opens a novel research branch.¹² However, we remark that DTC works also when the \mathcal{T}_i -solvers are not $\neg e_{ij}$ -minimal, at the cost of (at most) one extra branch to explore for each redundant $\neg e_{ij}$ returned in a conflict set.

It is also important to notice that, in general, only a fraction of the assignments μ enumerated turn out to be \mathcal{T}_i -satisfiable for both \mathcal{T}_i 's, so that to require the boolean search on the e_{ij} 's. Thus, for all the other branches, DTC may save the effort of many failed attempts of deducing implied e_{ij} 's.

On the whole, the results presented in this paper show that DTC allows for trading boolean search for e_{ij} -deduction. Thus everyone can choose and implement the most suitable \mathcal{T}_i -solvers without being forced by the e_{ij} -deduction-completeness strait-jacket: for theories for which efficient e_{ij} -deduction complete procedures are available (e.g., \mathcal{EUF} [12]), DTC allows for exploiting the full power of e_{ij} -deduction; for harder theories (e.g., $\mathcal{LA}(\mathbb{Z})$), the research task changes from that of finding e_{ij} -deduction complete \mathcal{T} -solvers to that of finding $\neg e_{ij}$ -minimal or nearly- $\neg e_{ij}$ -minimal ones.

References

1. A. Armando, C. Castellini, E. Giunchiglia, and M. Maratea. A SAT-based Decision Procedure for the Boolean Combination of Difference Constraints. In *Proc. SAT'04*, 2004.
2. C.L. Barrett and S. Berezin. CVC Lite: A New Implementation of the Cooperating Validity Checker. In *Proc. CAV'04*, volume 3114 of *LNCS*. Springer, 2004.

¹² Bottom line, one can always make μ $\neg e_{ij}$ -minimal by dropping the remaining $\neg e_{ij}$'s one by one, each time checking $\mu \setminus \{\neg e_{ij}\}$. Notice that, in general, with $\neg e_{ij}$ -minimization the search for the candidate $\neg e_{ij}$'s to drop is restricted to only those occurring in μ , whilst with e_{ij} -deduction the search for the candidate e_{ij} 's to deduce extends to all the unassigned e_{ij} 's.

3. M. Bozzano, R. Bruttomesso, A. Cimatti, T. Junttila, P. van Rossum, S. Ranise, and R. Sebastiani. Efficient Satisfiability Modulo Theories via Delayed Theory Combination. In *Proc. Int. Conf. on Computer-Aided Verification, CAV 2005.*, volume 3576 of *LNCS*. Springer, 2005.
4. M. Bozzano, R. Bruttomesso, A. Cimatti, T. Junttila, P. van Rossum, S. Schulz, and R. Sebastiani. An incremental and Layered Procedure for the Satisfiability of Linear Arithmetic Logic. In *Proc. TACAS'05*, volume 3440 of *LNCS*. Springer, 2005.
5. M. Bozzano, R. Bruttomesso, A. Cimatti, T. Junttila, P. van Rossum, S. Schulz, and R. Sebastiani. MathSAT: A Tight Integration of SAT and Mathematical Decision Procedure. *Journal of Automated Reasoning*, 35(1-3), October 2005.
6. M. Bozzano, R. Bruttomesso, A. Cimatti, T. Junttila, P. van Rossum, S. Ranise, and R. Sebastiani. Efficient Theory Combination via Boolean Search. *Information and Computation*, 204(10), 2006.
7. R. Bruttomesso, A. Cimatti, A. Franzén, A. Griggio, and R. Sebastiani. Delayed Theory Combination vs. Nelson-Oppen for Satisfiability Modulo Theories: a Comparative Analysis. Technical Report DIT-06-032, DIT, University of Trento, 2006. Available at http://dit.unitn.it/~rseba/papers/lpar06_dtc_extended.pdf.
8. S. Cotton, E. Asarin, O. Maler, and P. Niebert. Some Progress in Satisfiability Checking for Difference Logic. In *Proc. FORMATS-FTRTFT 2004*, 2004.
9. J.-C. Filliâtre, S. Owre, H. Rueß, and N. Shankar. ICS: Integrated Canonizer and Solver. In *Proc. CAV'01*, volume 2102 of *LNCS*, pages 246–249, 2001.
10. H. Ganzinger, G. Hagen, R. Nieuwenhuis, A. Oliveras, and C. Tinelli. DPLL(T): Fast decision procedures. In *Proc. CAV'04*, volume 3114 of *LNCS*, pages 175–188. Springer, 2004.
11. G. Nelson and D.C. Oppen. Simplification by Cooperating Decision Procedures. *ACM Transactions on Programming Languages and Systems*, 1(2):245–257, 1979.
12. R. Nieuwenhuis and A. Oliveras. Congruence Closure with Integer Offsets. In *Proc. 10th LPAR*, number 2850 in *LNAI*, pages 77–89. Springer, 2003.
13. R.E. Shostak. Deciding Combinations of Theories. *Journal of the ACM*, 31:1–12, 1984.
14. L. Zhang, C. F. Madigan, M. H. Moskewicz, and S. Malik. Efficient conflict driven learning in a boolean satisfiability solver. In *Proc. ICCAD '01*. IEEE Press, 2001.
15. L. Zhang and S. Malik. The quest for efficient boolean satisfiability solvers. In *Proc. CAV'02*, number 2404 in *LNCS*, pages 17–36. Springer, 2002.




Fig. 9. Graphical representations of base cases 1. and 2. and Step cases (a), (b), (c).

A Appendix: The proofs of the theorems

Proof of Theorem 1

Proof. By induction on the structure of the DTC boolean search tree.

Base: We have two basic cases (Figure 9, left).

1. Let μ be \mathcal{T}_i -unsatisfiable for some \mathcal{T}_i . The \mathcal{T}_i -solver detects this fact returning an $\neg e_{ij}$ -minimal conflict set η . Thus $dtc_br = 1$ and $dtc_ded = 0$. Similarly, in every NO refutation \mathcal{T}_i -solver detects the \mathcal{T}_i -unsatisfiability of μ . Thus $no_ded = 0$ and $no_br = 1$, so that (4) holds.
2. Let μ be \mathcal{T}_i -satisfiable for both \mathcal{T}_i 's, and no e_{ij} -deduction is performed by the \mathcal{T}_i -solvers. DTC selects a chain of new negated e_{ij} 's, invoking an early-pruning check before each new selection which cause no e_{ij} -deduction, until no new e_{ij} 's are available, from which it concludes that μ is $\mathcal{T}_1 \cup \mathcal{T}_2$ -satisfiable. Here $dtc_br = 1$ and $dtc_ded = 0$. In every corresponding NO proof, both \mathcal{T}_i -solvers return "Sat" without performing e_{ij} -deductions (i.e., $no_br = 1, no_ded = 0$). Therefore (4) holds.

Step: If none of the previous cases holds, then DTC selects a chain of new negated e_{ij} 's, invoking an early-pruning check before each new selection, until either (Figure 9, right):

- (a) one early-pruning call to one \mathcal{T}_i -solver returns Unsat. In this case, let B denote the current branch. The \mathcal{T}_i -solver returns also a $\neg e_{ij}$ -minimal conflict set η , corresponding to the conflicting clause

$$\eta^* \rightarrow \bigvee_{j=1}^k e_j, \quad (6)$$

where $\neg e_1, \dots, \neg e_k$ are all the $\neg e_{ij}$'s occurring in η , and $\eta^* := \eta \setminus \{\neg e_1, \dots, \neg e_k\}$. (In the corresponding NO refutation, this corresponds to the e_{ij} -deduction $\mu \models_{\mathcal{T}_i} \bigvee_{j=1}^k e_j$.) DTC learns the conflicting clause (6) and backjumps, popping up the literals from the branch (if $k > 1$) up to $\neg e_{k-1}$ or (if $k = 1$) up to the highest point in μ where e_k can be unit-propagated on (6), and hence unit propagating e_k . (In the NO refutation, this corresponds to selecting the branch $\mu \cup \{e_k\}$.)




Fig. 10. Graphical representation of the recursive behaviour. Case 1 (left) and 2 (right).

- (b) one positive¹³ e_{ij} , namely e_k , is unit-propagated due to some previously-learned conflict clause, which we can think w.l.o.g. in the form (6), s.t. all $\neg e_1, \dots, \neg e_{j-1}$ and all the literals in η^* are in the current branch, and η^* contains no negated e_{ij} 's. (Notice that, in the NO refutation, this might require a novel e_{ij} -deduction $\mu \models_{\mathcal{T}_i} \bigvee_{j=1}^k e_j$.¹⁴)
- (c) One \mathcal{T}_i -solver performs a e_{ij} -deduction, namely $\eta^* \models_{\mathcal{T}_i} \bigvee_{j=1}^k e_j$, s.t. η^* is part of the current branch and the e_j 's are not. By Step 4. of Strategy 1, the clause (6) is learned immediately, DTC selects in order $\neg e_1, \dots, \neg e_{k-1}$, and hence it unit-propagates e_k on clause (6). (In the corresponding NO refutation, this corresponds to the e_{ij} -deduction $\eta^* \models_{\mathcal{T}_i} \bigvee_{j=1}^k e_j$ and to the selection of the branch e_k .)

Let $B_{e_k} := \mu_k \cup \{e_k\}$ be the current branch, s.t. $\mu_k \supseteq \mu$ and $\neg e_j \in \mu_k$ for every $j < k$. Now DTC checks recursively the $\mathcal{T}_1 \cup \mathcal{T}_2$ -satisfiability of B_{e_k} .

If B_{e_k} is recursively found $\mathcal{T}_1 \cup \mathcal{T}_2$ -satisfiable, then DTC concludes that μ is $\mathcal{T}_1 \cup \mathcal{T}_2$ -satisfiable.

Otherwise, B_{e_k} is recursively found $\mathcal{T}_1 \cup \mathcal{T}_2$ -unsatisfiable. By inductive hypothesis, the DTC sub-proof requires dtc_br_k branches, whilst the NO subproof requires no_ded_k deductions and no_br_k branches, s.t. $dtc_br_k \leq no_br_k + no_ded_k$. Let η_k be the conflict set returned and let N_k be the number of negated equalities in η_k . We distinguish two subcases (Figure 10, left and right):

1. $e_k \in \eta_k$. Thus $\eta_k = \eta_k^* \cup \{\neg e'_j\}_{j=1}^{N_k} \cup \{e_k\}$, s.t. η_k^* does not contain $\neg e_{ij}$'s. By Step 2. of Strategy 1, DTC eliminates e_k from the conflicting clause $\neg \eta_k$ by resolving it

¹³ The case where one *negative* e_{ij} is unit-propagated due to some previously-learned conflict clause C does not affect the overall discussion, because it will be eliminated from every conflict clause by means of of Step 2. in Strategy 1. No literal other than (negated) e_{ij} 's can be unit propagated, because μ assigns a truth value to all the atomic expressions in φ .

¹⁴ E.g., the \mathcal{EUF} deduction of $(v_3 = v_4)$ in the right branch of Figure 2 corresponds to a simple unit-propagation on clause C_{34} in Figure 5.

with (6), obtaining the new conflict clause:

$$(\eta^* \cup \eta_k^*) \rightarrow \left(\bigvee_{j=1}^{k-1} e_j \vee \bigvee_{j=1}^{N_k} e'_j \right). \quad (7)$$

Let $\neg e$ be the most recently assigned $\neg e_{ij}$ in $\{\neg e_j\}_{j=1}^{k-1} \cup \{\neg e'_j\}_{j=1}^{N_k}$. Then DTC backjumps up to the highest point in B_{e_k} where e can be unit-propagated on (7), and hence unit-propagates e . (Notice that (7) dominates (6) in driving the backjumping mechanism because all the $\neg e'_j$'s occur higher in B than $\neg e_k$.)

2. $e_k \notin \eta_k$. Thus $\eta_k = \eta_k^* \cup \{\neg e'_j\}_{j=1}^{N_k}$, s.t. η_k^* does not contain $\neg e_{ij}$'s, corresponding to the clause:

$$\eta_k^* \rightarrow \bigvee_{j=1}^{N_k} e'_j. \quad (8)$$

Let $\neg e$ be the most recently assigned $\neg e_{ij}$ in $\{\neg e'_j\}_{j=1}^{N_k}$. Then DTC backjumps up to the highest point in B_{e_k} where e can be unit-propagated on (8), and hence unit-propagates e . (Notice that also (8) dominates (6) in driving the backjumping mechanism because all the $\neg e'_j$'s occur higher in B than $\neg e_k$.)

Then, DTC proceeds, each time checking recursively the $\mathcal{T}_1 \cup \mathcal{T}_2$ -satisfiability on one open branch $B_{e_j} := \mu_j \cup \{e_j\}$, each e_j corresponding to either one of the original negated e_{ij} 's $\neg e_1, \dots, \neg e_k$ in B , or to one of the negated e_{ij} 's occurring in the conflict sets reported by the recursive sub-proofs. This is done until either a subbranch is recursively found to be $\mathcal{T}_1 \cup \mathcal{T}_2$ -satisfiable, or the current dominating conflict clause forces DTC backjumping up to a point within μ , so that μ can be declared $\mathcal{T}_1 \cup \mathcal{T}_2$ -unsatisfiable, and the negation of the dominating clause is the conflict set returned.

Let N be the number of sub-proofs performed. By inductive hypothesis, the j -th DTC sub-proof requires dtc_br_j branches and dtc_ded_j e_{ij} -deductions, whilst the corresponding NO sub-proof requires no_ded_j e_{ij} -deductions and no_br_j branches, s.t. $dtc_br_j + dtc_ded_j \leq no_br_j + no_ded_j$. In the cases (a)-(c) above we have respectively:

Case (a): $dtc_br := 1 + \sum_{j=1}^N dtc_br_j$, $dtc_ded := \sum_{j=1}^N dtc_ded_j$, $no_br := \sum_{j=1}^N no_br_j$, and $no_ded := 1 + \sum_{j=1}^N no_ded_j$.

Case (b): $no_br := \sum_{j=1}^N no_br_j$, $dtc_ded := \sum_{j=1}^N dtc_ded_j$, $dtc_br := \sum_{j=1}^N dtc_br_j$, and either $no_ded := \sum_{j=1}^N no_ded_j$ or $no_ded := 1 + \sum_{j=1}^N no_ded_j$.

Case (c): $dtc_br := \sum_{j=1}^N dtc_br_j$, $dtc_ded := 1 + \sum_{j=1}^N dtc_ded_j$, $no_br := \sum_{j=1}^N no_br_j$, and $no_ded := 1 + \sum_{j=1}^N no_ded_j$.

In all cases, (4) holds. \square

Proof of Theorem 2

Proof. By induction on the structure of the DTC boolean search tree. .

Base: Let μ be \mathcal{T}_i -unsatisfiable for some \mathcal{T}_i 's. In this case, the proof is as for Theorem 1.

Let μ be \mathcal{T}_i -satisfiable for both \mathcal{T}_i 's, and neither \mathcal{T}_i -solver can perform any e_{ij} -deduction from μ . Then by step 5. of Strategy 1, DTC can conclude that μ is $\mathcal{T}_1 \cup \mathcal{T}_2$ -satisfiable. The same would do every NO tool. Thus $dtc_br = no_br = 1$, so that (5) holds.

Step: If some previously-learned clauses forces some unit-propagation, the proof is as for Theorem 1.

If none of the two base conditions hold, and no unit propagation can be performed on μ , then μ is \mathcal{T}_i -satisfiable for both \mathcal{T}_i 's and one \mathcal{T}_i -solver performs a deduction, namely $\mu \models_{\mathcal{T}_i} \bigvee_{j=1}^k e_j$. By Step 4. of Strategy 1, the clause (6) is learned immediately, and hence DTC selects in order $\neg e_1, \dots, \neg e_{k-1}$, and hence it unit-propagates e_k on clause (6). (In the corresponding NO refutation, this corresponds to the e_{ij} -deduction $\mu \models_{\mathcal{T}_i} \bigvee_{j=1}^k e_j$ and to the selection of the branch e_k .)

Henceforth, the proof is as for Theorem 1, except for the fact that, by induction hypothesis, $dtc_br_k \leq no_br_k$, and $dtc_br := \sum_{j=1}^N dtc_br_j$, $no_br := \sum_{j=1}^N no_br_j$ for every j . Thus (5) holds. \square