

Satisfiability Modulo Theory Chains with DPLL(T)

Scott Cotton and Oded Maler

Verimag, Centre Équation
2, Avenue de Vignate
38610 Gières, France
{Scott.Cotton,Oded.Maler}@imag.fr

Abstract. We extend the DPLL(T) framework for satisfiability modulo theories to address richer theories by means of increased flexibility in the interaction between the propositional and theory-specific solvers. We decompose a rich theory into a chain of increasingly more complex subtheories, and define a corresponding propagation strategy which favors the simpler subtheories using two mechanisms. First, subtheory propagation is prioritized so that more expensive propagation is avoided whenever possible. Second, constraints are filtered along the path from simpler to more complex propagation, thus easing the task of propagation for each subtheory. We present this strategy formally in a refined abstract DPLL(T) system and provide a concrete algorithmic skeleton with a proof of correctness.

1 Introduction

The problem of determining the satisfiability of a Boolean combination of constraints is a fundamental and computationally-intractable problem with applications in many fields including model checking, theorem proving, scheduling, circuit analysis and planning. The field of *satisfiability modulo theories* (SMT) studies effective mechanisms for solving this problem in its full generality, that is, in a manner independent of the form of the constraints. Thus far, the field has focused on modular integration of standard decision procedures for propositional logic together with pluggable decision procedures for the theory that corresponds to the specific type of constraints.

Amongst a plethora of recent research on satisfiability modulo theories [20, 8, 21, 18, 16, 17, 1, 12, 11, 3], the DPLL(T) framework has proven successful [8, 16]. Within this framework, a major issue in the interaction between the DPLL procedure and the theory-specific solver is that of *theory propagation*, that is, at what frequency and to what extent are the theory consequences of newly-assigned constraints computed. Since reasoning inside a theory is typically more expensive than purely-Boolean reasoning, different choices of theory propagation strategies may significantly affect performance.

In this paper we present a prioritized strategy for theory propagation of richer theories which can be decomposed into chains of increasingly complex sub-theories. The strategy consists primarily of an abstract *constraint filter* which is described in terms of the roles that constraints play with respect to the subtheories in the chain. We also provide a simple, prioritized implementation of this filter given a chain of sub-solvers for each theory, and prove its correctness. This rest of this

paper is organized as follows. In section 2, we review the DPLL(T) framework and discuss theory propagation strategies. Our new methodology is presented in section 3 as a refinement of an abstract DPLL(T) system. In section 4, we present a simple implementation of the abstract constraint filter, followed by a proof of its correctness in section 5. Conclusions and future work are presented briefly in section 6.

2 Background

2.1 DPLL(T)

DPLL(T) is a framework for determining the satisfiability of a Boolean combination of constraints, where the constraints belong to some specified theory. In particular it makes use of two solvers, a slightly modified Davis-Putnam-Loveland-Logeman (DPLL) style Boolean satisfiability solver [7, 6], and a solver for a given theory T . As the modified DPLL solver is parameterized on the theory solver, we refer to it as DPLL(X); the theory solver is called Solver $_T$.

The DPLL satisfiability checking procedure is concerned with propositional logic. Its input formula φ is assumed to be in conjunctive normal form. This is to say that φ is a conjunction of clauses, where each clause is a disjunct of literals and each literal is either a propositional variable or its negation. Given such a formula, a DPLL solver will search the space of truth assignments to the variables by incrementally building up partial assignments and backtracking whenever a partial assignment falsifies a clause. Assignments are extended both automatically by *unit propagation* and by guessing truth values. Unit propagation occurs by keeping track of the effect of partial assignments on the clauses in the input formula. For example, the clause $x \vee \neg y$ is solved under an assignment including $y \mapsto \mathbf{false}$ and is reduced to the clause x under an assignment which includes $y \mapsto \mathbf{true}$ but which contains no mapping for x . Whenever a clause is reduced to contain a single literal, *i.e.* it is in the form v or the form $\neg v$, the DPLL procedure deduces by unit propagation the assignment to v which solves the clause. If a satisfying assignment is found, the formula is satisfiable. If the procedure exhausts the assignment space without finding a satisfying assignment, the formula is not satisfiable.

Modern DPLL solvers are highly optimized, employing various dynamic variable ordering schemes [13, 9, 19], non-chronologic backtracking [10], efficient clause status tracking [13], and dynamic learning of clauses [2]. These procedures have made advances by leaps and bounds in the last years and are routinely able to solve real-world problems consisting of tens of thousands of variables and clauses within a few seconds. As DPLL(X) is a slightly modified version of DPLL, these optimizations may readily be employed in the DPLL(T) framework.

DPLL(X) interacts with a theory-specific solver Solver $_T$, which is a satisfiability solver for *conjunctions* of constraints in the language of a given theory T . The DPLL(X) procedure interprets every constraint c as a propositional variable, and interacts with Solver $_T$ in the following way.

SetTrue. DPLL(X) calls Solver $_T$ after assigning a truth value of some constraint c . Solver $_T$ then performs a consistency check on the current assignment. If the assignment is consistent, Solver $_T$ may optionally return a list $I(c)$ of constraints which are implied by the theory under the current assignment. Otherwise, it

indicates that the current assignment is unsatisfiable. In this case, Solver_T may moreover identify a small subset of the assignment which is unsatisfiable. The $\text{DPLL}(X)$ solver is expected to backtrack in response.

Explain. $\text{DPLL}(X)$ may ask Solver_T to provide a list of reasons R_c for any implied constraint c previously returned by SetTrue . This occurs when learning new clauses. The reasons R_c for c must not include c and it must be the case that c follows from the reasons under the theory, *i.e.* that $R_c \models_T c$. Moreover, every $c' \in R_c$ must have been assigned prior to c .

BackTrack. $\text{DPLL}(X)$ calls Solver_T every time backtracking occurs, indicating those constraints which are removed from the partial truth assignment.

In hopes of making this paper both self-contained and readably short, the high level description given here of DPLL , and $\text{DPPL}(T)$ is prosaic and informal. The reader is referred to [10, 13] for a general overview of modern DPLL SAT solving and to [8, 20] for a general overview of $\text{DPLL}(T)$ and SMT .

2.2 A $\text{DPLL}(T)$ Transition System

We describe our extension of $\text{DPLL}(T)$ in terms of the abstract $\text{DPLL}(T)$ transition system due to Nieuwenhuis *et al.* [17]. The system is both more faithful to real implementations than the prosaic description above and at the same time facilitates formal reasoning about a wide class of $\text{DPLL}(T)$ implementations. We describe this system briefly below and extend it in section 3, to better accommodate flexible theory propagation.

To begin at the beginning, a *theory* is a deductively closed set of valid formulas, each of which has no free variables. The formulas are expressed in a language consisting of variables, logical constants ($\wedge, \neg, \forall, \dots$), and a set of function and relation symbols called a *signature*. Each theory has an associated class of *models* each of which is an interpretation of all the symbols in the signature over some set, called the *domain* of the model. One usually reasons about properties of models by means of the axioms under which the theory is deductively closed. For example, a theory of equality T_{eq} may be specified by a relation symbol “ R ”, and a set of axioms expressing that “ R ” is reflexive, symmetric, and transitive. Any model $M = \langle D, R^D \rangle$ where R^D is in fact a reflexive, symmetric and transitive relation over D is a T_{eq} -model. In particular, a model M is called a T -model, also written $M \models T$, whenever for every formula $\varphi \in T$, M interprets φ in such a way that φ is true. The satisfaction relation \models_T is then used to indicate whether T -models satisfy arbitrary formulas in the language of T . In particular, we write $M \models_T \varphi$ if $M \models T$ and $M \models \varphi$, *i.e.* if M is a T -model and there is an assignment mapping the free variables of φ to the domain of M such that φ is true. In this case, we say that φ is *satisfiable in T* . Similarly, the expression $\Gamma \models_T \varphi$, when Γ is a set of formulas in the language of T , denotes that for any T -model M such that $M \models \Gamma$, $M \models \varphi$.

A *literal* is either a *propositional literal* or a *constraint*. A propositional literal is either a Boolean variable v or its negation $\neg v$. A *constraint* is an atom, possibly negated, in the language of the given theory T . For example, $x = 2$ is a constraint in the language of arithmetic. A *clause* C is a non-trivial disjunction $l_0 \vee l_1 \vee \dots \vee l_k$ of literals. For any literal $l \in C$, we assume its negation is not a member of C .

With these definitions at hand, we briefly present the DPLL(T) framework as an abstract *state transition system* \mathbb{D}_T with conditional transition rules. Each state in \mathbb{D}_T consists of a partial truth assignment A in the form of an annotated set of literals, coupled with a set of clauses F . Every literal in a truth assignment either appears in F , or is in the form $\neg c$ for some constraint in F . There is a single initial state which consists of the empty truth assignment coupled with the initial set of clauses H . As transitions are taken, the partial assignment is built up, and the set of clauses may change, subject to the conditions in the transition rules. Every literal in a given partial assignment is labelled as either *guessed* or *not guessed*, depending on what transition rule is used to add it to the truth assignment. The transition rules are given in Figure 1.

Rule	Transition	Conditions
U-prop	$A \circ F, C \vee l \rightarrow A, l \circ F, C \vee l$	$A \models_T \neg C$
T-prop	$A \circ F \rightarrow A, l \circ F$	$A \models_T l$
Guess	$A \circ F \rightarrow A, l_g \circ F$	-
Back	$A, l_g, A' \circ F, C \rightarrow A, l' \circ F, C$	$A, l_g, A' \models_T \neg C$ There exists a clause $C' \vee l'$ s.t. $A \models_T \neg C'$ $F, C \models_T C' \vee l'$
Learn	$A \circ F \rightarrow A \circ F, C$	$F \models_T C$
Forget	$A \circ F, C \rightarrow A \circ F$	$F \models_T C$
Fail	$A \circ F, C \rightarrow \text{fail}$	$A \models_T \neg C$ There are no guessed literals in A
Restart	$A \circ F \rightarrow \emptyset \circ F$	-

Fig. 1. Transition relation \rightarrow for DPLL(T) system \mathbb{D}_T , originating from [17]. We write $A \circ F$ for a state, indicating the assignment A and a set of clauses C . A clause $C \vee l$ is taken to be a possibly empty clause C , which does not include l or $\neg l$, disjoined with the literal l . We write l_g for guessed literals and simply l otherwise. We assume all literals or their negations occur in the original input formula and that any transition involving an assignment assigns only undefined literals.

The system \mathbb{D}_T is, in general, non-deterministic because several transition rules may be enabled in a given state. We call the mechanism by which a solver chooses among transitions a *strategy*. One easy way to define a strategy is to define *priorities* for to the various transition rules. In general, some priorities are more or less universal. For example, the fail transition is almost always implemented as taking priority over all other rules, and the Guess transition is almost never implemented when the U-prop rule is applicable. On the other hand, the priority of the T-prop rule with respect to the U-prop and Guess rules is more complex and varied.

2.3 Theory Propagation Strategies

Theory propagation, *i.e.* application of the T-prop rule, is generally expensive in practice. However, it also has the benefit of reducing the guessing space of the whole

DPLL(T) procedure. Every literal implied by theory propagation is a propositional variable whose truth value DPLL(X) need not guess; moreover every such literal may enable further unit propagation which, in turn, reduces the guessing space further and may also lead to more theory propagation. As increased guessing leads in the worst case to exponentially longer runs and hence even more theory specific processing, there is a balance to be struck between applying theory propagation and avoiding it. That balance is unfortunately both theory- and problem-specific.

As a result of this quandary, the role of theory propagation in SMT solvers is extremely varied. Some solvers, such as MathSat [11] don't make use of expensive theory propagation at all but rather simply perform theory consistency checks on the partial assignments. Others, such as MxSolver [15] and BarceLogicTools [16] (for difference logic) apply theory propagation whenever the T-prop rule is enabled to avoid expensive guessing. Between these two extremes, some solvers such as Jat [3] apply theory propagation lazily and incompletely and others, such as BarceLogicTools (for the logic of uninterpreted functions) apply an incomplete propagation procedure exhaustively. However, each of these strategies lack flexibility. The lazy strategy of [11] disallows theory propagation altogether. The eager strategy of [16] requires special treatment of theory consequences by the Boolean solver and requires that theory consequences be complete. The lazy round-robin strategy of Jat makes incremental theory propagation difficult at best.

In all these cases, the theory propagation is monolithic in the sense that there is no decomposition of theory propagation into distinct subprocedures or subtheories. In addition, the relative priorities of unit propagation and theory propagation are hard-coded and in no way configurable. Finally, with the notable exception of the greedy strategy, the strategies make no distinction between the role of constraints with respect to theory propagation. In particular, T -consequences are subject to the same processing as any other constraint, although they are guaranteed to be T -consistent and, under reasonable assumptions, not to lead to more T -consequences.

2.4 Overview of a New Methodology

In this paper we present a methodology for theory propagation as a refinement of the system \mathbb{D}_T , providing a framework for reasoning about a broad class of propagation strategies. We describe an implementation scheme for this methodology. The following concepts underly this work:

Decomposition. A theory T is decomposed into a chain of theories $T_1 \subset T_2 \subset \dots \subset T_k$ with $T = T_k$. Each step in the chain properly extends the set of valid formulas and each T_i is a deductively closed theory. Decomposition provides a chain of refinements for a given formula in the language of T . As an example, the theory of the real field may be decomposed into a chain consisting of pure equality, difference logic, linear arithmetic, quadratic arithmetic, and arbitrary order arithmetic.

Prioritization. The decomposition above is prioritized: propagation for each theory T_i always takes priority over propagation for any T_j with $i < j$. We present a transition system \mathbb{D}_{T° , a variant of \mathbb{D}_T , which provides for modelling prioritization more flexibly, including for the case of a trivial decomposition. The

system \mathbb{D}_{T° can be mapped back to the original \mathbb{D}_T , thus preserving useful properties of \mathbb{D}_T such as termination and correctness.

Identification of constraint roles. We refine the states of \mathbb{D}_{T° by decorating constraints with labels that reflect their role with respect to theory propagation in each T_i . We use the constraint roles to build a *constraint filter* along the theory chain. This filter prevents theory-deduced constraints from undergoing a lot of redundant processing to which they are subject in the standard DPLL(T) framework, even in the case of a trivial decomposition.

3 Abstract Methodology

3.1 Preliminaries

Definition 1 (T-Decomposition). *A decomposition T_k° of a theory T is a set of theories $\{T_i \mid 1 \leq i \leq k\}$ such that each $T_i \subset T_{i+1}$ and $T = T_k$. In addition, for each theory T_i with $i \in [1..k]$, we associate a language \mathcal{L}_i , a satisfaction relation \models_i and a decision procedure \vdash_i . We require that each \vdash_i be able to answer queries of the form $\Gamma \models_i c$ for a finite set $\Gamma \cup \{c\}$ of constraints in \mathcal{L}_i .*

The decomposition can be seen as a sequence of refinements. Just as DPLL(X) views constraints propositionally while Solver_T views them in refined form as theory constraints, a decomposition allows one to view a constraint as a propositional literal, as a constraint in a simple theory, as well as a constraint in a more complex theory. The following lemma and corollary are central to our handling of decompositions.

Lemma 1 (Refinement). *Given two theories T, T' , if $T \subset T'$ then any T' -model is a T -model*

Proof. Any T' -model M satisfies every formula in T' , and $T \subset T'$. Hence M satisfies every formula in T and so by definition is a T -model. \square

Corollary 1 (Consequence lifting). *Given two theories $T_i, T_j \in T_k^\circ$ with $i < j$ and a set of formula $\Gamma \cup \{\varphi\}$, if $\Gamma \models_i \varphi$, then $\Gamma \models_j \varphi$.*

Proof. Suppose that any T_i -model of Γ is a model of φ . We need to show that any T_j model M of Γ is also a model of φ . By lemma 1, since M is a T_j -model, M is also T_i -model, and so $M \models_i \Gamma$. By the hypothesis, M is a model of φ . \square

Definition 2 (The set χ). *Given a set of clauses F , we define the set χ as the smallest set containing every constraint c and its negation $\neg c$ such that c appears in some clause in F .*

The consequence finders ρ_x^i defined below provide a means to model actual deduction procedures associated with a given theory T_i . A sound and *complete* consequence finder may be implemented using the decision procedure \vdash_i , simply by querying for every $c \in \chi$ whether $\Gamma \models_i c$, but we allow also incomplete finders to reflect more faithfully the working of real solvers.

Definition 3 (Consequence finders). Given a theory decomposition T_k° and the set χ , a consequence finder ρ_x^i for $T_i \in T_k^\circ$ is an arbitrary function $\rho_x^i : 2^X \rightarrow 2^X$ satisfying the following three conditions.

1. *Soundness.* For every $i \in [1..k]$, $\rho_x^i(\Gamma) \subseteq \{c \in \mathcal{L}_i \mid \Gamma \cap \mathcal{L}_i \models_i c\}$.
2. *Monotonicity.* For every $i \in [1..k]$, and $\Gamma \subseteq \Gamma' \subseteq \chi$, $\rho_x^i(\Gamma) \subseteq \rho_x^i(\Gamma')$.
3. *Chain closure.* For every $i, j \in [1..k]$ with $j \leq i$, $\rho_x^i(\rho_x^j(\Gamma)) \subseteq \rho_x^i(\Gamma)$.

Monotonicity simply states that as more constraints are handed to a consequence finder, more consequences are found. Chain closure which, by Corollary 1, holds for any sound and complete consequence finder, is not strictly necessary in our methodology but helps simplifying the results and their presentation.

3.2 Modelling Decompositions with Priorities

Given a T -decomposition T_k° , consider replacing the T-prop rule in \mathbb{D}_T with one rule T-prop $_i$ for each $T_i \in T_k^\circ$:

$$A \circ F \rightarrow A, l \circ F \text{ if } A \models_i l$$

All the runs of the resulting transition system can easily be mapped back to runs in \mathbb{D}_T simply by renaming applications of each T-prop $_i$ to the original T-prop rule. This follows directly from Corollary 1 as any T_i -propagation is a valid T -propagation. We refer to the modified transition system as \mathbb{D}_{T° and view it as a model of DPLL(T) executions in which we can differentiate between the effects of different T_i consequences in a given T -decomposition.

However, we would also like to reason about priorities given to rules and hence need to take a few more things into consideration. In particular, for one propagation rule to take precedence over another, it is necessary to know when propagation by one rule is “finished” so that a lower priority rule may be applied. It is also necessary to know what constitutes the smallest possible unit of propagation for a given rule, so that propagation by a lower priority rule may defer to a higher priority rule, but only after having completed a conveniently small chunk of work.

It is reasonable to define a propagation by a rule as finished when that rule is no longer enabled. This in turn may be expressed by a closure condition associated with each propagation rule. In particular, we call a rule T-prop $_i$ *closed* in a state $A \circ F$ whenever $\rho_x^i(A) \subseteq A$, *i.e.* whenever it is not enabled. Similarly, we consider the U-prop rule *closed* in a state $A \circ F$ if there are no unit clauses in F under the assignment A .

While closure is easily defined in the system \mathbb{D}_{T° , a flexible notion of a “conveniently small chunk of work” is not. Often, propagation procedures will find a set of consequences in reaction to an extension of a set of antecedents. For example, if $S = \{w < x, x < y\}$ is a set of antecedents and $S' = S \cup \{y < z\}$, then a consequence finder may find that $w < y$ is a consequence of S and then that $\{w < z, x < z\}$ are consequences of S' . In the latter case, the procedure has found more than one consequence in response to antecedent extension. But a transition $A \circ F \rightarrow A, l \circ F$ in the system \mathbb{D}_{T° would require that only one consequence $l \in \{w < z, x < z\}$ become part of the assignment at a time. If some other propagation rule has higher priority, then the consequence finder should remember unassigned consequences

while other higher priority rules are brought to closure. This may be inconvenient since it is natural to implement consequence finders simply as a function of a set of antecedents. To allow for prioritized theory propagation based on antecedent extension rather than based on consequence extension, we define alternate T -prop $_i$ rules in below.

$$A \circ F \rightarrow A, B \circ F \text{ if } \begin{cases} \exists A' \subseteq A . B = \rho_x^i(A') \setminus A \\ B \neq \emptyset \end{cases} \quad (1)$$

With closure and antecedent based propagation rules, we can easily express realistic rule preferences for \mathbb{D}_{T° . If one rule R has a higher priority than R' , then we add the closure condition of R to the set of preconditions of R' . We always consider that the Guess rule has lowest priority, and hence that it is conditioned on the closure of both theory and unit propagation.

Rule prioritization is effective in the DPLL(T) framework for two reasons. Consider the case that one rule R is much more expensive than a rule R' . If a literal is impliable under both R and R' , then it need not be implied under R when R' has priority. In this case, one may “save” an application of the rule R by applying the rule R' first. Secondly, and more profoundly, the DPLL(T) framework is a *backtracking search* framework. Consequently, there are many contexts in which applications of the cheaper R' may lead to a backtrack. In all these contexts, if R' has higher priority than R , then *no* applications of the more expensive rule occur.

3.3 Extended Constraint States

In every state of \mathbb{D}_{T° , there is an annotated partial truth assignment A which indicates whether or not constraints are assigned and whether or not assigned constraints are guessed. Below, we define *extended states* which refine the original states by indicating, in addition, the role of constraints with respect to each subtheory in a T -decomposition. This additional information is used by Solver $_T$ and is invisible to DPLL(X), hence all the good properties of \mathbb{D}_{T° are retained.

To motivate this refinement we illustrate an inherent inefficiency in standard implementations of the DPLL(T) algorithm, manifested already in the case of a monolithic theory T . Suppose Solver $_T$ is invoked by DPLL(X) to find consequences of some newly-assigned constraint c , and it finds a consequence c' . It then hands c' to DPLL(X) which puts it in its implication queue. When the time comes for c' to be processed, DPLL(X) will hand c' back to Solver $_T$, where it will undergo a redundant consistency check, and moreover it will participate, together with its antecedents, in each and every further theory propagation although, being a consequent, it will *not* contribute *any new consequences*. This can needlessly increase the amount of work done by the consistency checker and the consequence finder. By identifying the role of constraints with respect to theory propagation, *i.e.* which constraints are consequences and which are antecedents with respect to each subtheory, we can avoid this processing, even in the case of a trivial decomposition.

To this end, for every constraint c we keep track of the role it plays in T_i propagation, namely, consequence, antecedent or no role at all. Whenever a constraint c is identified as a consequence in theory T_i , we consider it to be a consequence in

T_j for all $j \geq i$. Likewise, whenever a constraint c is known to be an antecedent in a subtheory T_i , it is considered an antecedent in subtheories T_j with $j \leq i$. All this information is compactly represented by associating with each constraint c a pair of labels $\langle l_c, u_c \rangle$ with $l_c, u_c \in [0..k + 1]$, such that c has participated as antecedent in all $\{T_i \mid i < l_c\}$, and as a consequence in all $\{T_i \mid i > u_c\}$. The role of c in $\{T_i \mid i \in [l_c, u_c]\}$ is not yet determined. We will maintain that whenever $u_c < l_c$, $u_c = l_c - 1$. This way, for every subtheory, a constraint may participate either as an antecedent or as a consequence but not both.

An extended state is then simply a state $A \circ F$ together with a labelling $\langle l_c, u_c \rangle$ of every constraint $c \in \chi$, satisfying certain properties detailed below. To express these properties, we make use of the following definition.

Definition 4 (T_i -basis and T_i -cobasis). *In any given state in the system \mathbb{D}_{T° , we refer to the T_i -basis \mathbf{b}_i as the set $\{c \in \chi \mid l_c > i\}$. Similarly, we write the T_i -cobasis $\bar{\mathbf{b}}_i$, denoting the set $\{c \in \chi \mid u_c < i\}$.*

Below we give several extended state properties which we require to hold throughout the transition system \mathbb{D}_{T° with extended states. Taken together, these properties facilitate safe optimization of DPLL(T) for richer theories by identifying redundant constraints for each T_i , safely allowing a great deal of flexibility in when consequence finding and consistency checking occurs, for which constraints they occur, and guaranteeing that all possible consequence finding eventually occurs.

1. *Basis-cobasis assignment.* In any state $A \circ F$, $\bigcup_{i \in [1..k]} \mathbf{b}_i \subseteq A$, and $A \cap \chi \subseteq \bigcup_{i \in [0..k]} \mathbf{b}_i \cup \bar{\mathbf{b}}_i$. This property states the relationship between assignments and constraint roles. It simply requires that we know the role of all assigned constraints with respect to some subtheories, while allowing the latitude that the role of some unassigned constraints may also be known.
2. *Basis consistency.* In any state, for every $i \in [1..k]$, $\mathbf{b}_i \cap \mathcal{L}_i$ is T_i -consistent. This is a weak consistency condition, allowing the system to be T -inconsistent at times while requiring piecewise T_i -consistency. If some consistency checks are expensive, they can be temporarily avoided, allowing the possibility of arriving at an inconsistency by less expensive means.
3. *Full consistency.* In any state $A \circ F$ which precedes a Guess or is a final state, for every $i \in [1..k]$, $A \cap \mathcal{L}_i$ is T_i -consistent. This is a periodic strong consistency condition which ensures that the system is consistent, but only in the states in which full T -consistency is required.
4. *Cobasis irrelevancy.* In any state $A \circ F$, for every $i \in [1..k]$ and $j \geq i$, $A \setminus \bar{\mathbf{b}}_i \models_j A$. This is a safety condition which allows the system to correctly filter out any constraint in $\bar{\mathbf{b}}_i$ from all T_i specific processing.
5. *Basis-cobasis closure.* In any state, for every $i \in [1..k]$, $\rho_\chi^i(\mathbf{b}_i) \subseteq \mathbf{b}_i \cup \bar{\mathbf{b}}_i$. This property requires that all consequence finding of each basis has occurred. As a T_i -local property, it ensures that all T_i consequences of \mathbf{b}_i have been identified.
6. *Full closure.* In any state $A \circ F$ which precedes a Guess or is a final state, for every $i \in [1..k]$, $\rho_\chi^i(A) \subseteq A$. This property states that all possible theory propagation must have occurred before the system solves a problem or Guesses.

4 Concrete Methodology

We present a relatively simple skeleton implementation of $\text{DPLL}(X)$ and Solver_T which realizes the system \mathbb{D}_{T° with extended states, satisfying all extended state properties and using the theory propagation rules based on antecedent extension as presented in equation 1. In addition, this implementation enforces prioritization of propagation rules which respect the decomposition T_k° . We assume that consequence finding in theory T_i is less expensive than consequence finding in T_{i+1} . We also assume that consequence finding in T_i is comensurate with consistency checks in T_{i+1} , coupling consistency checks in T_{i+1} with propagation in T_i . We present these implementations in terms of changes to the original $\text{DPLL}(T)$ system described in section 2.1.

4.1 $\text{DPLL}(X)$

Recall that the $\text{DPLL}(X)$ engine presented in section 2.1 interacts with Solver_T via the three methods `SetTrue`, `Backtrack`, and `Explain`. In this framework, Solver_T is unable to distinguish between states of the $\text{DPLL}(X)$ engine. In particular, Solver_T has no way of knowing whether or not the $\text{DPLL}(X)$ engine will guess if Solver_T does not apply the T-prop rule.

We remedy this situation by adding a method `TheoryProp` to Solver_T . This method is called by $\text{DPLL}(X)$ as a last-resort form of constraint propagation. In particular, as in a standard DPLL engine, applicable instances of the propagation rules are kept in an implication queue, and processed in a first-in-first-out manner. Processing consists of checking for new instances of the U-prop rule and calling `SetTrue` for each assigned constraint. If the queue becomes empty and no inconsistencies are found, our $\text{DPLL}(X)$ engine calls `TheoryProp`, which in turn may re-fill the implication queue with some instances of the T-prop rule. If `TheoryProp` does not re-fill the implication queue, $\text{DPLL}(X)$ will guess because there are no more instances of propagation rules. Also, as in `SetTrue`, the method `TheoryProp` may indicate that an inconsistency has occurred. In this case, it also identifies an inconsistent subset of the assignment. Unlike `SetTrue`, the method `TheoryProp` is not associated with any particular constraint. The presence of `TheoryProp` simply enables Solver_T to know in which states the U-prop rule is not enabled, so that theory propagation can occur at a lower priority than unit propagation.

Secondly, we modify a standard DPLL procedure to communicate the event that a constraint is removed from the implication queue as a result of a backtrack. A standard DPLL empties the implication queue when a conflict is found, just prior to backtracking. We assume this behavior and add a call to a new method `Dequeue(c)` for each constraint c that is removed from the implication queue upon backtracking. The $\text{DPLL}(X)$ procedure makes no further accommodations for our methodology and in fact, may be readily used for lazy theory propagation [3] or standard $\text{DPLL}(T)$ -style theory propagation.

4.2 A Generic Solver_{T°

We present a generic implementation of Solver_{T° which is provided with a T -decomposition T_k° . Solver_{T° allocates an extended state constraint label $\langle l_c, u_c \rangle$ for

each constraint $c \in \chi$. Initially, Solver_{T° identifies $\rho_x^i(\emptyset)$ for each $i \in [1..k]$, and constraints c are labelled as follows. All l_c are assigned 0. If c is a member of some $\rho_x^i(\emptyset)$, then u_c is assigned $j-1$ where j is the least j such that $c \in \rho_x^j(\emptyset)$. Otherwise, u_c is assigned k . All constraints with $u_c < k$ are added to the $\text{DPLL}(X)$ implication queue.

After initialization all constraint relabelling occurs via the methods **SetTrue**, **Backtrack**, **Dequeue**, and **TheoryProp**¹. The method **SetTrue** occurs as a constraint c becomes assigned. The procedure first checks whether or not $u_c < 1$. If $u_c < 1$ or $c \notin \mathcal{L}_1$, no consistency check is performed. Otherwise the T_1 -consistency of $\mathbf{b}_0 \cap \mathcal{L}_1 \cup \{c\}$ is tested. If $u_c < 1$ or the consistency check succeeds, l_c is assigned the value 1. The methods **Backtrack** and **Dequeue** relabel every constraint c which become unassigned or leave the implication queue as a result of backtracking. Each such constraint c is labelled $\langle 0, k \rangle$. The method **TheoryProp** performs all relabelling which establishes the basis and cobasis of each T_i , and is detailed in figure 2.

```

R ← ∅
while {c ∈ R | l_c = 0} = ∅ and ∃l_c ∈ [1..k] . u_c ≥ l_c do
  select some c with the minimal l_c ∈ [1..k] such that u_c ≥ l_c
  i ← l_c
  if c ∈ L_{i+1} ∧ b_i ∩ L_{i+1} ⊨_{i+1} ¬c then
    return ⟨⊥, Explain(¬c) ∪ {c}⟩
  l_c ← i + 1
  R ← ρ_x^i(b_i) ∩ {c ∈ χ | l_c ∈ {0, i} ∧ u_c ≥ l_c}
  for c ∈ R do
    u_c ← i - 1
  return ⟨⊤, {c ∈ R | l_c = 0}⟩

```

Fig. 2. **TheoryProp** returns a pair $\langle ok, \Gamma \rangle$. If ok is true, then Γ is a set of T -consequences of the current assignment, otherwise Γ is a small inconsistent subset of the current assignment. **TheoryProp** relabels constraints c one at a time. Each such constraint satisfies $l_c \leq u_c$ and minimizes l_c , prior to the relabelling. Let $i = l_c$ before c is relabelled. If $c \in \mathcal{L}_{i+1}$ then T_{i+1} -consistency of $\mathbf{b}_i \cap \mathcal{L}_{i+1} \cup \{c\}$ is tested. If the test succeeds, then l_c is assigned $i+1$ and then the consequence finder ρ_x^i is invoked. For every consequence c , u_c is assigned $i-1$. If new consequences c with $l_c = 0$ are found, they are unassigned and the procedure returns them to the $\text{DPLL}(X)$ engine to add to its implication queue.

The implementation above exhibits some “instant” optimizations which are not part of a standard $\text{DPLL}(T)$ system. First, no constraints c with $u_c < k$ are subject to processing in any T_i with $i > u_c$. These constraints do not trigger T_i -consistency checks, do not trigger T_i consequence finding, and do not form part of the T_i -basis \mathbf{b}_i . Thus in comparison to a standard $\text{DPLL}(T)$ system, fewer consistency checks occur, fewer calls to theory propagation occur, and theory propagation is

¹ The method **Explain** does not perform any relabelling. Nonetheless, it is worth mentioning that **Explain** is always called with a constraint c which is implied by some $\mathbf{b}_i \cap \mathcal{L}_i$. This set may be used as a set of reasons for c in case a particular implementation is unable to provide a smaller set of reasons.

itself a function of a smaller set. This optimization holds for single theories as well and its correctness follows from the correctness of extended state properties, which we address in the next section. In practice, the dramatic effect of these optimizations for the case of the (undecomposed) theory of difference constraints was demonstrated empirically by the *Jat* solver [4].

5 Correctness of Extended State Properties

We now give some proof sketches showing how the required extended state properties of section 3.3 are realized by the implementation. Full proofs are available in the technical report version of this paper [5]. We begin by simply stating a few of the simpler results which are crucial for proving the higher level properties.

Theorem 1 (Basis-cobasis assignment). *In any state $A \circ F$, and for any $i \in [1..k]$ the T_i basis is assigned ($\mathbf{b}_i \subseteq A$) and the assignment is contained in the union of the T_i -basis and T_i -cobasis ($A \cap \chi \subseteq \bigcup_{i \in [0..k]} (\mathbf{b}_i \cup \bar{\mathbf{b}}_i)$).*

Lemma 2. *In every state $A \circ F$ which precedes a *Guess* or is final, for every $c \in \chi$ either $l_c = u_c + 1$ and c is assigned, or $l_c = 0$, $u_c = k$ and c is unassigned.*

For the remaining properties, we give proofs and proof sketches:

Theorem 2 (Basis-cobasis closure). *In any state, every $\rho_x^i(\mathbf{b}_i) \subseteq \mathbf{b}_i \cup \bar{\mathbf{b}}_i$*

Proof. (sketch) Induction on constraint relabelling. □

Theorem 3 (Cobasis irrelevancy). *In any state $A \circ F$ and for every $i, j \in [1..k]$ with $j \geq i$, $A \setminus \bar{\mathbf{b}}_i \models_j A$.*

Proof. Every time a constraint c becomes a member of some $\bar{\mathbf{b}}_i$, $c \in \rho_x^h(\mathbf{b}_h)$ for some $h \leq i$. Since each ρ_x^h is sound, we have that $\mathbf{b}_h \models_h \bar{\mathbf{b}}_i$. By theorem 1, $\mathbf{b}_h \subseteq A$ and hence $A \models_h \bar{\mathbf{b}}_i$. By consequence lifting (corollary 1), then $A \models_j \bar{\mathbf{b}}_i$ for any $j \geq h$. Since $h \leq i$, we have $A \models_j \bar{\mathbf{b}}_i$ for any $j \geq i$ as well. □

Theorem 4 (Basis consistency). *In any state every set $\mathbf{b}_i \cap \mathcal{L}_i$ is T_i -consistent.*

Proof. (sketch) We assume each T_i is consistent, and prove a stronger proposition: each $\mathbf{b}_i \cap \mathcal{L}_{i+1}$ is T_{i+1} -consistent. The proof proceeds by induction on constraint relabelling, where it suffices to consider only relabellings which extend some \mathbf{b}_i , observing that consistency checks occur as needed. □

Theorem 5 (Full consistency). *In any state $A \circ F$ which precedes a *guess* or is final, every $A \cap \mathcal{L}_i$ is T_i -consistent.*

Proof. In any state which precedes a guess or is final, every assigned constraint c satisfies $c \in \mathbf{b}_i \cup \bar{\mathbf{b}}_i$ for every $i \in [1..k]$, since every assigned constraint has $l_c = u_c + 1$ (lemma 2). It follows that $A \cap \chi \subseteq \mathbf{b}_i \cup \bar{\mathbf{b}}_i$ for every $i \in [1..k]$. This together with theorem 1 and lemma 2 gives us $A \cap \chi = \mathbf{b}_i \cup \bar{\mathbf{b}}_i$, and it follows that $A \cap \mathcal{L}_i = (\mathbf{b}_i \cup \bar{\mathbf{b}}_i) \cap \mathcal{L}_i$. By the soundness of each consequence finder, $\mathbf{b}_i \cap \mathcal{L}_i \models_i (\mathbf{b}_i \cup \bar{\mathbf{b}}_i) \cap \mathcal{L}_i$ and so $\mathbf{b}_i \cap \mathcal{L}_i \models_i A \cap \mathcal{L}_i$. Then we have that $A \cap \mathcal{L}_i \models_i \perp \implies \mathbf{b}_i \cap \mathcal{L}_i \models_i \perp$ by transitivity of \models_i . Contrapositing this statement gives $\mathbf{b}_i \cap \mathcal{L}_i \not\models_i \perp \implies A \cap \mathcal{L}_i \not\models_i \perp$. By theorem 4, each $\mathbf{b}_i \cap \mathcal{L}_i$ is T_i -consistent, hence $A \cap \mathcal{L}_i$ is T_i -consistent. □

Corollary 2. *Full consistency occurs even if consistency checks only occur with respect to T_k .*

Theorem 6 (Full closure). *In any state which precedes a Guess or is final every $\rho_x^i(A) \subseteq A$.*

Proof. By theorem 2, the implementation ensures that at every state and for every $i \in [1..k]$, $\rho_x^i(\mathbf{b}_i) \subseteq \mathbf{b}_i \cup \bar{\mathbf{b}}_i$. By lemma 2, every constraint c with $u_c < k$ is assigned, hence every $\bar{\mathbf{b}}_i \subseteq A$. It follows immediately that $\rho_x^i(\mathbf{b}_i) \subseteq \mathbf{b}_i \cup A$. By theorem 1, in addition each $\mathbf{b}_i \subseteq A$, and it follows immediately that each $\rho_x^i(\mathbf{b}_i) \subseteq A$.

By chain-closure and monotonicity of each ρ_x^i , it follows that $\rho_x^i(\mathbf{b}_i) = \rho_x^i(\mathbf{b}_i \cup \bar{\mathbf{b}}_i)$. Consequently, $\rho_x^i(\mathbf{b}_i \cup \bar{\mathbf{b}}_i) \subseteq A$. By lemma 2, every assigned constraint c has $l_c = u_c + 1$. Consequently, for every $i, j \in [1..k]$, $\mathbf{b}_i \cup \bar{\mathbf{b}}_i = \mathbf{b}_j \cup \bar{\mathbf{b}}_j = A \cap \chi$. Hence it follows that $\rho_x^i(A \cap \chi) \subseteq A$. Since each $l \in A \setminus \chi$ has no interpretation in any T_i , $\rho_x^i(\mathbf{b}_i \cup (A \setminus \chi)) = \rho_x^i(\mathbf{b}_i)$. We conclude each $\rho_x^i(A) \subseteq A$. \square

Corollary 3. *If ρ_x^k is complete then the procedure never guesses an inconsistent constraint.*

6 Conclusions

We presented a method for extending the DPLL(T) procedure to richer theories by means of a prioritized constraint filter. Classical approaches to theory *combination* such as Nelson-Oppen's method [14] may also be used to enrich theories. However, these approaches almost invariably require the signatures of the theories to be disjoint. This restriction inherently precludes the exploitation of simple sub-theories and practically excludes disjoint theories intended to operate over the same domain (such as the additive and multiplicative groups in arithmetic). By contrast, while our methodology doesn't address signature-disjoint theories, it does exploit the simplicity of sub-theories, and allows multiple theories which share the same domain. We hence believe that the DPLL(T) methodology presented in this paper complements theory combination methods.

Nonetheless, several questions remain open with regard DPLL(T) theory chains in general. Perhaps foremost amongst these are empirical questions. We have only begun to test theory chains in DPLL(T). Thus far, experiments suggest that prioritization and filtering are effective mechanisms which reduce calls to the most expensive procedures and also reduces their complexity in general. However, our experiments are thus far limited to trivial decompositions and so our evidence is far from conclusive. Other questions of interest are whether or not formula transforms can be used to boost the effectiveness of theory chains. For example, the Nelson-Oppen procedure performs a *purification* transform on formulas which is satisfiability preserving and may produce a formula in which every constraint has only one relation or function symbol. Perhaps this procedure would also artificially enlarge the set of constraints belonging to less complex theories in a chain, thus making the filter more effective. Finally, it would be interesting to address the question of composing Nelson-Oppen with theory chains to arrive at theory trees. We would like to study these issues in future work.

References

1. A. Armando, C. Castellini, E. Giunchiglia, M. Idini, and M. Maratea. Tsat++: an open platform for satisfiability modulo theories. In *PDPAR'04*, 2004.
2. R. J. Jr. Bayardo and R. C. Schrag. Using CSP look-back techniques to solve real-world SAT instances. In *AAAI'97*, pages 203–208, 1997.
3. S. Cotton. Satisfiability checking with difference constraints. Master's thesis, Max Planck Institute, 2005.
4. S. Cotton and O. Maler. Fast and flexible difference logic propagation for DPLL(T). In *Submitted for publication.*, 2006.
5. S. Cotton and O. Maler. Satisfiability modulo theory chains with DPLL(T). In *Verimag Technical Report TR-2006-04*, 2006.
6. M. Davis, G. Logemann, and D. Loveland. A machine program for theorem proving. In *Communications of the ACM*, volume 5(7), pages pages 394–397, 1962.
7. M. Davis and H. Putnam. A computing procedure for quantification theory. *Journal of the ACM*, 7(1):201–215, 1960.
8. H. Ganzinger, G. Hagen, R. Nieuwenhuis, A. Oliveras, and C. Tinelli. DPLL(T): Fast decision procedures. In *CAV'04*, pages 175–188, 2004.
9. E. Goldberg and Y. Novikov. Berkmin: A fast and robust sat solver. In *DATE'02*, pages 142–149, 2002.
10. J. P. Marquez-Silva and K. A. Sakallah. Grasp – a new search algorithm for satisfiability. In *CAV'96*, pages 220–227, November 1996.
11. M.Bozzano, R.Bruttomesso, A.Cimatti, T.Junttila, P.v.Rossum, S.Schulz, and R.Sebastiani. Mathsat: Tight integration of SAT and mathematical decision procedures. *Journal of Automated Reasoning*, Special Issue on SAT, 2005.
12. M.Bozzano, R.Bruttomesso, A.Cimatti, T.Junttila, S.Ranise, P.v.Rossum, and R.Sebastiani. Efficient satisfiability modulo theories via delayed theory combination. In *CAV'05*, volume LNCS 3576, pages pages 335–349, 2005.
13. M. W. Moskewicz, C. F. Madigan, Y. Zhao, L. Zhang, and S. Malik. Chaff: Engineering an Efficient SAT Solver. In *DAC'01*, 2001.
14. G. Nelson and D. C. Oppen. Simplification by cooperating decision procedures. In *ACM Trans. on Programming Languages and Systems*, volume 1(2), pages 245–257, Oct. 1979.
15. P. Niebert, M. Mahfoudh, E. Asarin, M. Bozga, O. Maler, and N. Jain. Verification of timed automata via satisfiability checking. In *LNCS*, volume Volume 2469, pages 225 – 243, Jan 2002.
16. R. Nieuwenhuis and A. Oliveras. DPLL(T) with Exhaustive Theory Propagation and its Application to Difference Logic. In Kousha Etessami and Sriram K. Rajamani, editors, *CAV'05*, volume 3576 of *Lecture Notes in Computer Science*, pages 321–334. Springer, 2005.
17. R. Nieuwenhuis, A. Oliveras, and C. Tinelli. Abstract DPLL and abstract DPLL modulo theories. In F. Baader and A. Voronkov, editors, *LPAR'04*, volume 3452 of *LNCS*, pages 36–50, 2005.
18. R. Sebastiani. On efficiently integrating boolean and theory-specific solving procedures.
19. J. P. M. Silva. The impact of branching heuristics in propositional satisfiability algorithms. In *Proc. of the 9th Portuguese Conference on Artificial Intelligence*, pages 62–74, 1999.
20. C. Tinelli. A DPLL-based calculus for ground satisfiability modulo theories. In *Proc.of the 8th European Conf. on Logics in Artificial Intelligence*, volume 2424 of *LNCS*, 2002.
21. O. Tveretina. DPLL-based procedure for equality logic with uninterpreted functions.