

On Conflict-Driven Reasoning

Maria Paola Bonacina

Dipartimento di Informatica
Università degli Studi di Verona
Strada Le Grazie 15
I-37134 Verona, Italy, EU
mariapaola.bonacina@univr.it

Abstract

Automated formal methods and automated reasoning are interconnected, as formal methods generate reasoning problems and incorporate reasoning techniques. For example, formal methods tools employ reasoning engines to find solutions of sets of constraints, or proofs of conjectures. From a reasoning perspective, the expressivity of the logical language is often directly proportional to the difficulty of the problem. In propositional logic, Conflict-Driven Clause Learning (CDCL) is one of the key features of state-of-the-art satisfiability solvers. The idea is to restrict inferences to those needed to explain conflicts, and use conflicts to prune a backtracking search. A current research direction in automated reasoning is to generalize this notion of *conflict-driven satisfiability* to a paradigm of *conflict-driven reasoning* in first-order theories for satisfiability modulo theories and assignments, and even in full first-order logic for generic automated theorem proving. While this is a promising and exciting lead, it also poses formidable challenges.

1 Introduction

Automated reasoning and automated formal methods, for the specification, analysis, verification, or synthesis of systems, are interconnected, because logic is the calculus of computation, and reasoning about computer systems [19, 52, 5] may be more amenable to automation than other less machine-oriented domains. In automated reasoning, problems are typically presented as *validity* queries. A validity query asks whether a conjecture φ follows from a set H of assumptions, written $H \models \varphi$. Since mechanical methods preferably work refutationally, a validity query is usually reformulated in refutational form, by asking whether $H \cup \{\neg\varphi\}$ is *unsatisfiable*. Assumptions, conjectures, constraints are logical formulæ that express properties of an object of study, such as a system, a program, a data type, a circuit, a protocol, a mathematical structure. As mechanical methods usually adopt clausal form, $H \cup \{\neg\varphi\}$ is transformed into a set S of clauses, where a set is interpreted as a conjunction. Alternatively, one may be interested in knowing whether a constraint φ can be added to a set of constraints H so that $H \cup \{\varphi\}$ is still *satisfiable*. Once $H \cup \{\varphi\}$ has been turned into a set of clauses S , the problem is the same, namely determining whether S has a model or is unsatisfiable.

The answer is either a proof $S \vdash \square$ that S is inconsistent, hence unsatisfiable, where \square is the empty clause, which represents a contradiction, or else a model of S . If the problem was originally formulated as a validity query, a proof means that φ follows from H , while a model represents a counter-example. If the problem was originally formulated as a satisfiability query, a model represents a solution, while a proof means that there is no solution. Depending on the logic, these queries may be *decidable* (validity and satisfiability are both decidable), *semi-decidable* (validity is semi-decidable, satisfiability is not semi-decidable), or *undecidable* (validity and satisfiability are both undecidable).

An automated reasoning method or strategy is typically defined by an *inference system* and a *search plan*. The inference system is a set of inference rules, and the search plan is an algorithm equipped with heuristics to control the application of the inference rules. The application of an inference rule moves the system from one *state* of the *derivation* to the next. When the problem is decidable, an automated reasoning strategy is expected to be a *decision procedure*, that requires it to be *sound*, *complete*, and *terminating*, returning a proof whenever the input is unsatisfiable and a model whenever the input is satisfiable. When the problem is semi-decidable, an automated reasoning strategy is expected to be a *semi-decision procedure*, that requires it to be *sound* and *complete*, returning a proof whenever the input is unsatisfiable. In practice, however, instances of decidable problems may be too difficult for the available computational resources, or complete strategies may be too onerous, so that regardless of decidability, automated reasoning tools may return either a proof, or a model, or a “don’t know” answer. The degree to which “don’t know” answers may be tolerated depends on the application.

Similar to other subfields of artificial intelligence, problems in automated reasoning involve so much knowledge, that it is often too cumbersome or too inefficient to encode all of it in H and φ , hence in the set S of clauses. Therefore, a common paradigm is to reason *modulo* \mathcal{T} , seeking proofs modulo \mathcal{T} and restricting the attention to \mathcal{T} -models. For example, if \mathcal{T} is the theory of equality, we have *equational reasoning*, where the axioms of the theory of equality are built into the inference system. \mathcal{T} may also contain additional axioms stating properties of symbols other than equality, such as *associativity* and *commutativity* of function symbols.

If \mathcal{T} is a theory such that \mathcal{T} -satisfiability is decidable, reasoning modulo \mathcal{T} is known as *satisfiability modulo a theory* (SMT), and the knowledge about \mathcal{T} is built in the algorithm implementing the decision procedure for \mathcal{T} -satisfiability. For example, if \mathcal{T} is the quantifier-free fragment of the theory of equality, a *congruence closure* algorithm decides the \mathcal{T} -satisfiability of a set of equalities and inequalities (see Chapter 9 of [19]). An algorithm that decides the \mathcal{T} -satisfiability of a set of ground literals in the signature of \mathcal{T} , or \mathcal{T} -literals for short, is called a *\mathcal{T} -satisfiability procedure*. An algorithm that decides the \mathcal{T} -satisfiability of a quantifier-free formula in the signature of \mathcal{T} , or quantifier-free \mathcal{T} -formula for short, is called a *\mathcal{T} -decision procedure*. A quantifier-free formula φ is satisfiable if and only if its existential closure $\exists \bar{x}.\varphi$ is satisfiable, where \bar{x} are all the variables in φ . Then, $\exists \bar{x}.\varphi$ is satisfiable if and only if $\hat{\varphi}$ is satisfiable, where $\hat{\varphi}$ is φ with all variables replaced by Skolem constants. Thus, the problem of deciding the \mathcal{T} -satisfiability of a quantifier-free \mathcal{T} -formula is equivalent to that of deciding the \mathcal{T} -satisfiability of a ground \mathcal{T} -formula, or, equivalently, of a set of ground clauses.

A reasoning method is *model-based*, if the state of a derivation contains a representation of a candidate partial model, and inference and search for a model are intertwined, as inferences build and transform the model while the model drives the inferences [9]. In a model-based strategy, a *conflict* arises if one of the clauses of S is false in the current candidate model. The strategy is deemed *conflict-driven*, if it uses inferences to *explain* and *solve* the conflict repairing the model. This paper offers in Section 2 a necessarily incomplete overview of the state of the art in conflict-driven methods, while Section 3 advertises two recent conflict-driven methods: *Semantically-Guided Goal-Sensitive reasoning* (SGGS), for full first-order logic [16, 17, 18], and *Conflict-Driven Satisfiability* (CDSAT), for satisfiability modulo a *generic* combination of theories, and for a new class of problems called *satisfiability modulo assignments* (SMA) [10, 11].

2 Conflict-Driven Methods

The conflict-driven paradigm was pioneered by *Conflict-Driven Clause Learning* (CDCL) for propositional satisfiability [43, 46, 42]. In conflict-driven methods that incorporate a CDCL-

based SAT-solver as a black-box, the conflict-driven reasoning is propositional, even if the method applies to a more general logic. These methods are covered in Section 2.1. Other conflict-driven methods generalize the conflict-driven principle to satisfiability modulo a theory. These methods are treated in Section 2.2.

2.1 Conflict-Driven Propositional Reasoning

This section summarizes methods whose conflict-driven component is restricted to propositional logic. In other words, the candidate model and the conflict-driven inferences are propositional, with an abstraction function mapping first-order atoms to propositional atoms.

2.1.1 The DPLL Procedure

Satisfiability (SAT) is the problem of deciding the satisfiability of a set S of clauses in propositional logic. The DPLL (Davis-Putnam-Logemann-Loveland) procedure for SAT [24, 23, 21, 60] represents a candidate partial model by a *sequence* of literals, called a *trail*, and named M . The trail represents the partial model, also called M , where all literals on the trail are true. If a literal L is in M , its complement $\neg L$ is false in M . If neither L nor $\neg L$ is in M , both literals are *undefined*.

The procedure starts by putting in M all input unit clauses, and *propagating* their consequences in the form of *implications* and *conflicts*, an activity called *Boolean clausal propagation* (BCP). For implications, assume that all literals of a clause $C \in S$ but one, say L , are false in M . Then literal L is an *implied literal*, and is added to M with C as *justification*, because extending M with L is the only way to satisfy C . The discovery of an implied literal can be seen in terms of inferences as the result of a sequence of *unit resolution* steps using the literals in M as unit premises. For conflicts, whenever all literals of a clause $C \in S$ are false in M , a *conflict* emerges with C as the *conflict clause*. The discovery of a conflict can be seen in terms of inferences as the result of a sequence of unit resolution steps yielding the empty clause \square .

When no more propagations are possible and in the absence of a conflict, the procedure *decides* that a literal L is true by adding it to M . A literal added to M by a decision is termed a *decided literal*. A decision is merely a guess to advance the search. This operation is also termed *case analysis* or *splitting*, because for a literal L there are two cases, as L is either true or false. After every decision, the procedure applies BCP to discover more implied literals or a conflict. When a conflict arises, the procedure backtracks chronologically, undoing the latest decision and all the propagations that depend on it. The procedure returns “satisfiable” if $M \models S$, and “unsatisfiable” if there is a conflict and no decision to undo.

2.1.2 The CDCL Procedure

The Conflict-Driven Clause Learning (CDCL) procedure [43, 46, 42] inherits from the DPLL procedure the representation of the candidate model, BCP, and decisions. It also maintains the initialization of the trail with input unit clauses, said to be stored at level 0. Then, every decision opens a subsequent *decision level* in the trail: the decision level numbered n contains the n -th decided literal in the current trail and all implied literals discovered by BCP as a consequence. The CDCL procedure behaves in a markedly different manner when a conflict arises. Suppose that C is a conflict clause and contains a literal L , such that $\neg L$ is in M with justification D . Then propositional resolution is applied to resolve C and D upon L and $\neg L$. This inference is said to *explain* the conflict, as L is false because $\neg L$ is true, and $\neg L$ is true, because D is in S and all other literals of D appear negated in M . The generated resolvent is

still a conflict clause, since all other literals in C and D are false in M . A resolvent is a logical consequence of S and can be added to S as a *learned clause* or *lemma*. Such a step is called *learning*. In practice, S may be huge and the procedure learns one clause per conflict. How many resolutions to do and which resolvent to learn is a heuristic choice.

The *first unique implication point* (1UIP) heuristic prescribes to perform resolution until an *asserting* conflict clause C is generated. Assume that n is the number of the current decision level. A conflict clause $C = L_1 \vee \dots \vee L_i \vee \dots \vee L_m$ is asserting, or is an *assertion clause*, if for only one of its literals, say L_i , the complement appears in decision level n of the trail M . For all other literals L_j in C , with $1 \leq j \leq i - 1$ or $i + 1 \leq j \leq m$, the complement appears in a decision level smaller than n . As a special case, $\neg L_j$ appears in decision level 0, if $\neg L_j$ is a unit clause in the input set S . The 1UIP heuristic lets the procedure *learn* clause C and *backjump* to the smallest decision level where L_i is undefined and all other literals of C are false. Note that L_i is undefined in every level smaller than n . This smallest decision level is guaranteed to exist, because C is a conflict clause, and therefore for all its literals the complement is on the trail at some level. If this smallest decision level is $n - 1$, backjumping reduces to backtracking. If this smallest decision level is 0, the procedure backjumps to a state where only input unit clauses are on the trail. After backjumping, the procedure adds L_i to the trail, so that C is satisfied and the conflict is solved. The CDCL procedure returns “satisfiable” if $M \models S$, and “unsatisfiable” if there is a conflict at level 0.

2.1.3 The DPLL(\mathcal{T}) Framework

Satisfiability modulo theory (SMT) is the problem of deciding the satisfiability of a set S of *ground* clauses modulo a theory \mathcal{T} . The DPLL(\mathcal{T}) paradigm for SMT obtains a \mathcal{T} -decision procedure by integrating a CDCL-based SAT-solver and a theory solver, or \mathcal{T} -solver for short, implementing a \mathcal{T} -satisfiability procedure [49]. Since the SAT-solver accepts only propositional clauses, first-order ground atoms are abstracted to propositional variables, sometimes called *proxy variables*.

The interface between SAT-solver and \mathcal{T} -solver consists essentially of two rules. The \mathcal{T} -*conflict* rule detects that a set of literals L_1, \dots, L_r in M is inconsistent in \mathcal{T} . The \mathcal{T} -*propagation* rule discovers that a set of literals L_1, \dots, L_r in M derives in \mathcal{T} a literal L , and adds L to M with the \mathcal{T} -*lemma* $\neg L_1 \vee \dots \vee \neg L_r \vee L$ as justification. Thus, the \mathcal{T} -solver is a black-box for the SAT-solver and vice versa. However, the relationship among them is asymmetric, as only the CDCL-based SAT-solver operates on the trail, while the \mathcal{T} -solver acts as a satellite that submits \mathcal{T} -lemmas and signal \mathcal{T} -conflicts to the SAT-solver.

DPLL(\mathcal{T}) features *no creation of new atoms*, meaning atoms that do not appear in S . Indeed, the \mathcal{T} -propagation rule requires that the atom of L occurs in the existing set of clauses, and clauses learned by CDCL are propositional resolvents made of input atoms.

2.1.4 Combination of Theories by Equality Sharing

If \mathcal{T} is a combination of theories $\mathcal{T}_1, \dots, \mathcal{T}_n$, the \mathcal{T}_k -solvers, $1 \leq k \leq n$, need to agree on the interpretation of whatever is shared among the theories. If the theories are *disjoint*, meaning that they do not share function or predicate symbols other than equality, the theory solvers need to agree on the cardinalities of the domains for shared sorts and on an *arrangement* of shared constant symbols, that tells which are equal and which are not.

The *equality sharing* method is the standard approach to this combination problem (see [48, 47] and Chapter 10 of [19]). It requires the theories to be *stably infinite*, so that the common cardinality of the shared domains can be implicitly assumed to be countably infinite.

An arrangement is computed by having each \mathcal{T}_k -solver propagate any disjunction of equalities $c_1 \simeq d_1 \vee \dots \vee c_n \simeq d_n$ between shared constants that is entailed in \mathcal{T}_k by the \mathcal{T}_k -subproblem. Thus, every \mathcal{T}_k -solver is a black-box for the others. The case analysis for the propagated disjunctions, as well as for any other disjunction generated by a \mathcal{T}_k -solver, is entrusted to the SAT-solver.

If \mathcal{T} is a combination of theories $\mathcal{T}_1, \dots, \mathcal{T}_n$, the \mathcal{T} -solver integrated in $\text{DPLL}(\mathcal{T})$ is a combination of the \mathcal{T}_k -solvers by equality sharing, and the notion that a disjunction $c_1 \simeq d_1 \vee \dots \vee c_n \simeq d_n$ is handled by the SAT-solver is termed *splitting on demand* [4, 41]. The $\text{DPLL}(\mathcal{T})$ framework is extended to allow the generation of a finite number of “new” atoms, namely the proxy variables for the equalities $c_1 \simeq d_1, \dots, c_n \simeq d_n$.

2.1.5 Model-Based Theory Combination

Another way to implement equality sharing is *model-based theory combination* (MBTC) [55, 26]. It assumes that the \mathcal{T}_k -solvers build partial \mathcal{T}_k -models. Then, each \mathcal{T}_k -solver propagates, by adding it to M , any equality $s \simeq t$ between ground terms that is true in the current candidate \mathcal{T}_k -model, rather than entailed (disjunctions of) equalities between shared constants. Such an equality $s \simeq t$ may cause a conflict, precisely because it is not necessarily a logical consequence in \mathcal{T}_k of the \mathcal{T}_k -subproblem. If this happens, the backjumping mechanism of the CDCL-based SAT-solver will retract it. MBTC *does not generate new atoms either*, because the propagation of an equality $s \simeq t$ is allowed only if s and t appear in the existing set of clauses. MBTC applies mostly to fragments of arithmetic, where domain of interpretation and interpretation of theory symbols are fixed by an intended model (e.g., the integers), and algorithms that can update the candidate partial model after a conflict are known [31, 26].

MBTC is an approach to the implementation of equality sharing in the context of an SMT-solver built on top of a CDCL-based SAT-solver. Thus, conflicts are still handled and reasoned about in propositional logic. However, with its notion of allowing the propagation of equalities that are true in a current candidate partial theory model, but not necessarily in all models, MBTC contributed to prepare the ground for conflict-driven theory reasoning.

2.1.6 The $\text{DPLL}(\Gamma+\mathcal{T})$ Framework

MBTC is applied also in the $\text{DPLL}(\Gamma+\mathcal{T})$ framework that integrates an *ordering-based inference system* Γ for first-order logic with equality in $\text{DPLL}(\mathcal{T})$ [25, 13, 14]. Both ordering-based inference system Γ and theory or combination of theories \mathcal{T} are regarded as two parameters of the framework.

An ordering-based inference system assumes a *well-founded* ordering on terms, literals, and clauses, and comprises *expansion* inference rules, such as ordered resolution, ordered paramodulation, and superposition, and *contraction* inference rules, such as subsumption and simplification. The well-founded ordering is used to define the contraction rules and to restrict the expansion rules. Equipped with a fair search plan, such an inference system provides (1) a semi-decision procedure for validity in first-order logic with equality, and (2) \mathcal{T} -satisfiability procedures for the quantifier-free fragments of the theory of equality and of several theories of data structures [2, 3, 7, 8], including arrays with or without extensionality.

$\text{DPLL}(\Gamma+\mathcal{T})$ is designed to determine the \mathcal{T} -satisfiability of sets of clauses in the form $S = P \uplus R$, where \mathcal{T} is a combination of theories $\mathcal{T}_1, \dots, \mathcal{T}_n$, P is a set of *ground* clauses with occurrences of \mathcal{T} -symbols, and R is a set of *non-ground* clauses where \mathcal{T} -symbols do *not* occur. Variables in non-ground clauses are implicitly universally quantified. The set R may be the axiomatization of a theory for which a built-in satisfiability procedure is not available. This

kind of problem is more general than the standard SMT problem of deciding the \mathcal{T} -satisfiability of a set of ground clauses. The idea is to use the generic inference system Γ to reason about the axiomatized theory, precisely because Γ offers complete quantifier reasoning, since it is refutationally complete for first-order logic with equality.

DPLL($\Gamma + \mathcal{T}$) integrates Γ into DPLL(\mathcal{T}) by letting it use R -literals in M , including those propagated by MBTC, as premises of Γ -inferences. Since these literals may be withdrawn upon backjumping, they are memorized in clauses as *hypotheses*, and DPLL($\Gamma + \mathcal{T}$) works with *hypothetical clauses*. Conclusions of Γ -inferences inherit the hypotheses of their parents. When backjumping removes literals from M , the hypothetical clauses that depend on them are also removed.

Integrating an ordering-based inference system with a solver that performs a backtracking search presents both difficulties and opportunities. A difficulty is that one needs to prevent the unsound situation where a clause C is deleted by subsumption or simplification with a clause D and then D is removed upon backjumping. DPLL($\Gamma + \mathcal{T}$) solves this problem by adapting the contraction inference rules of Γ for hypothetical clauses in such a way that C is deleted, if D cannot be removed by backjumping before C , and only *disabled* otherwise. While deletion is final, a disabled clause C will be re-enabled, if D is removed by backjumping.

A distinctive opportunity is the possibility of allowing *speculative inferences*: the user can tentatively add to the set of clauses an arbitrary clause. The system will search for a model that satisfies *both* the input set S and the speculatively added clauses. DPLL($\Gamma + \mathcal{T}$) keeps track of the speculative addition of a clause C by placing a special propositional variable $[C]$ on the trail M . Clause C is added to the set of clauses as a hypothetical clause with hypothesis $[C]$. If an inconsistency results, $[C]$ will be retracted upon backjumping, and the clause will be removed from the set as a consequence. In this way, the speculative inferences are *reversible*.

The crux is to add clauses that may induce termination on satisfiable inputs, such as equations that limit term depth by rewriting: if S is satisfiable, it may happen that the search for a model of S does not terminate, but the search for a model of S that also satisfies the speculatively added clauses terminates. DPLL($\Gamma + \mathcal{T}$) is (1) a semi-decision procedure for validity of generic problems in the form $S = P \uplus R$, and (2) a \mathcal{T} -decision procedure with *speculative inferences* for problems $S = P \uplus R$ that satisfy additional hypotheses. For example, DPLL($\Gamma + \mathcal{T}$) offers \mathcal{T} -decision procedures with speculative inferences for several *axiomatizations of type systems* [14].

A feature of DPLL($\Gamma + \mathcal{T}$) is that it applies each reasoner to handle the part of the problem that it is best for: DPLL(\mathcal{T}) deals with ground clauses, while Γ sees non-ground R -clauses and ground unit R -clauses in M . The two engines communicate through M , making DPLL($\Gamma + \mathcal{T}$) *model-based*. However, the conflict-driven part is propositional as in DPLL(\mathcal{T}). We consider next methods that lift conflict-driven reasoning to the theory level.

2.2 Conflict-Driven Theory Reasoning

In conflict-driven theory reasoning, the mechanisms to *explain* a conflict, *learn* a lemma, and *solve* the conflict, work within the \mathcal{T} -solver itself, and not only at the propositional level in the SAT-solver. In other words, the \mathcal{T} -solver implements a *conflict-driven \mathcal{T} -satisfiability procedure*. Such procedures exist for linear real arithmetic [45, 39, 22], linear integer arithmetic [57, 55, 37], non-linear arithmetic [38], and floating-point binary arithmetic [32].

Some progress has been made towards a conflict-driven \mathcal{T} -satisfiability procedure for the theory of *arrays with extensionality* [20], by developing the notion of *lemmas on demand* [30]. The idea of *lemmas on demand* is that a theory solver should generate only theory lemmas that

explain why some contents of the trail M is inconsistent with respect to the theory. In other words, theory propagation should be model-based and conflict-driven. In propositional logic, lemmas on demand is the same as CDCL, with propositional resolvents as lemmas.

Although there are decision procedures for the theory of *arrays with extensionality* [54], SMT-solvers often reason about it by reading the theory axioms as part of the input set S , and heuristically instantiating the universally quantified variables in the theory axioms. For this theory, the difference of approach between SMT-solvers and superposition-based theorem provers that instantiate by unification the universally quantified variables in the theory axioms [3] is less dramatic than commonly thought.

Of greater relevance to this analysis is the difference between generating potentially all lemmas, as in a saturation process, and generating lemmas in a conflict-driven manner. The decision procedure with lemmas on demand features rules that propagate *read* operations over arrays, and generate lemmas of the form $\neg L_1 \vee \dots \vee \neg L_r \vee L$, where L_1, \dots, L_r are true and L is false in the current candidate model M , whereas L should be true according to the axioms of the theory [20]. The lemma reveals that M is not a theory model and tells why. Often lemmas are instances of axioms, so that lemmas on demand can be regarded as model-based conflict-driven axiom instantiation.

2.2.1 The MCSAT Framework

The next problem is how to get a *conflict-driven \mathcal{T} -decision procedure*. Conflict-driven \mathcal{T} -satisfiability procedures [45, 39, 22, 37, 38, 32] are not compatible in general with $\text{DPLL}(\mathcal{T})$, and therefore one cannot get a conflict-driven \mathcal{T} -decision procedure by plugging a conflict-driven \mathcal{T} -satisfiability procedure into $\text{DPLL}(\mathcal{T})$. A reason of incompatibility is that $\text{DPLL}(\mathcal{T})$ does not allow the creation of new atoms, whereas a conflict-driven \mathcal{T} -satisfiability procedure may need to generate a clause that contains *new* atoms in order to *explain* a conflict [27]. Another reason is that the trail in $\text{DPLL}(\mathcal{T})$ is defined to contain only propositional literals, whereas conflict-driven \mathcal{T} -satisfiability procedures need to store on the trail also assignments to first-order variables. Addressing such issues was the motivation for the design of MCSAT, that stands for *Model-Constructing SATisfiability* [27]. MCSAT is a paradigm for *conflict-driven \mathcal{T} -decision procedures* for satisfiability modulo a single generic theory \mathcal{T} [27]. It has been instantiated to the combined theories of equality and linear real arithmetic [36], to non-linear integer arithmetic [35], and to the theory of bit-vectors [58].

MCSAT merges the propositional model of CDCL with the theory models of MBTC, by allowing the trail M to contain both literals and assignments of domain values to free first-order variables. For example, the trail may contain a literal L , meaning the assignment $L \leftarrow \text{true}$, and assignments such as $x \leftarrow 3$. Therefore, the trail is viewed as carrying an *assignment* that partially represents a candidate first-order model. Furthermore, MCSAT generalizes CDCL to any theory that can be equipped with clausal inference rules to *explain* theory conflicts. Thus, the existence of a *conflict-explanation inference mechanism* emerges as the key ingredient for a conflict-driven procedure. The possibility of learning a clause generated by the conflict-explanation inference, and using it to amend the candidate partial model follows.

The conflict-explanation inferences generate clauses that may contain *new* ground atoms in the signature of the theory, beyond what is allowed by $\text{DPLL}(\mathcal{T})$ even with splitting on demand. Assignments to first-order variables and new atoms are involved in decisions, propagations, conflict detections, and explanations, on a par with Boolean assignments and input atoms. This means that the conflict-driven \mathcal{T} -satisfiability procedure is not integrated as a black-box satellite as in $\text{DPLL}(\mathcal{T})$, but cooperates with the SAT-solver on the same level. The CDCL procedure itself is a conflict-driven \mathcal{T} -satisfiability procedure where \mathcal{T} is propositional logic.

For termination, MCSAT requires that new atoms come from a *finite basis*. A procedure that applies systematically the inference rules to enumerate all atoms in the finite basis would be too inefficient. The key point is that the inference rules are applied only to explain conflicts and amend the current partial assignment, so that the generation of new atoms is model-based and conflict-driven. In this sense, MCSAT is a faithful lifting of CDCL to SMT, with first-order inferences for theory explanation, beyond explanation by propositional resolution.

3 General Conflict-Driven Methods

While satisfiability in propositional logic is decidable, in first-order logic validity is semi-decidable and satisfiability is not even semi-decidable. Nonetheless, theorem-proving approaches often are conceived and understood first for propositional logic and then generalized to full first-order logic. Section 3.1 presents the main features of a method named SGGS that lifts the CDCL procedure to first-order logic [16, 17, 18]. An alternative approach is *conflict resolution*, which focuses on lifting to first-order logic the conflict-driven generation of resolvents [53, 34]. Other approaches lift the CDCL procedure to the Bernays-Schönfinkel fragment, which allows only formulæ of the form $\exists^*\forall^*\varphi$, where φ contains neither quantifiers nor occurrences of function symbols [50, 1]. Section 3.2 returns to SMT with a summary of an inference system named CDSAT, that generalizes MCSAT to *generic* combinations of theories [10, 11]. Section 3.3 discusses how the SMT problem itself can be generalized to the SMA problem, for *satisfiability modulo assignments*.

3.1 A Taste of SGGS

SGGS, or *Semantically-Guided Goal-Sensitive* reasoning, brings the conflict-driven style to first-order logic [16, 17, 18]. It is *simultaneously* first-order, *model-based*, *semantically-guided*, *goal-sensitive*, and *proof confluent*, a rare combination of features.

In first-order logic variables in clauses are implicitly universally quantified, atoms have infinitely many ground instances, and there are infinitely many interpretations, so that guessing truth values of atoms is too uninformed. SGGS adopts an *initial interpretation* I for *semantic guidance*, and employs a new kind of structures, called *SGGS clause sequences*, to represent first-order models [17]. An SGGS clause sequence is a sequence of possibly constrained clauses with *selected literals*. A sequence Γ represents an interpretation $I[\Gamma]$, that is I modified to satisfy the selected literals in Γ . Thus, the SGGS clause sequence plays the role of the trail in CDCL, and *literal selection* is the first-order analogue of propositional decision.

Example 3.1. Assume that S includes the clauses $on(a, b)$, $on(b, c)$, $green(a)$, and $\neg green(c)$. If I is the all-negative interpretation, that makes all negative literals true, the SGGS-derivation starts with the SGGS clause sequence $\Gamma = on(a, b), on(b, c), green(a)$. In a unit clause its only literal is obviously selected. $I[\Gamma]$ is the interpretation that makes all positive literals false except $on(a, b)$, $on(b, c)$, and $green(a)$. If I is the all-positive interpretation, that makes all positive literals true, the SGGS-derivation starts with $\Gamma = \neg green(c)$, and $I[\Gamma]$ is the interpretation that makes all negative literals false except $\neg green(c)$. CDCL would put all input unit clauses on the trail. SGGS assumes a guiding interpretation and modifies it lazily, because dealing with first-order models is much heavier than dealing with propositional models.

SGGS generalizes BCP to *first-order clausal propagation*. BCP is based on the symmetry of truth values in propositional logic: if L is true, $\neg L$ is false, and if L is false, $\neg L$ is true. Since variables in first-order literals are implicitly universally quantified, if L is true, $\neg L$ is false, but

if L is false, we only know that at least one ground instance of $\neg L$ is true. To address this discrepancy, SGGS introduces *uniform falsity*: a first-order literal is *uniformly false*, if *all* its ground instances are false. This stronger notion of falsity restores the symmetry: if L is true, $\neg L$ is uniformly false, and if L is uniformly false, $\neg L$ is true.

Every literal in an SGGS clause sequence Γ must be either *I-true* (true in I) or *I-false* (uniformly false in I), so that it represents the truth value in I of all its ground instances. Every clause C in Γ must have a *selected literal* L : the clause with L selected is written $C[L]$. *I-false* literals are preferred for selection. An *I-true* literal is selected only in a clause whose literals are all *I-true*; such a clause is termed *I-all-true*. SGGS aims at building a model of S : if $I \models S$, the search halts immediately; if $I \not\models S$, SGGS seeks to build an $I[\Gamma]$ that differs from I where needed to satisfy S , hence the preference for *I-false* literals.

Example 3.2. $S = \{R(x, f(x)), \neg R(x, x), \neg R(x, y) \vee R(y, x)\}$ presents an irreflexive and symmetric reachability relation R such that every state x has a successor $f(x)$. If I is all-negative, SGGS builds the sequence $\Gamma = [R(x, f(x))], \neg R(x, f(x)) \vee [R(f(x), x)]$: in the second clause, which is binary, the positive literal is preferred for selection, denoted by the square brackets, because it is *I-false*. Then SGGS halts as $I[\Gamma] \models S$.

A first-order clause is a *conflict clause* if all its literals are uniformly false in $I[\Gamma]$. A literal L is uniformly false in $I[\Gamma]$, if all its ground instances appear negated among those that a preceding selected literal M makes true in $I[\Gamma]$. In this sense, L *depends* on M .

Example 3.3. Given $S = \{P(x), \neg P(x) \vee R(a, x), \neg P(x) \vee \neg R(x, b)\}$ and I all-negative, SGGS builds the sequence $\Gamma = [P(x)], \neg P(x) \vee [R(a, x)], \neg P(a) \vee [\neg R(a, b)]$: literals $\neg P(x)$ and $\neg P(a)$ are uniformly false in $I[\Gamma]$, because $P(x)$ is selected; literal $\neg R(a, b)$ is uniformly false in $I[\Gamma]$, because $R(a, x)$ is selected; and the last clause in Γ is in conflict with $I[\Gamma]$. Note that this clause is *I-all-true*.

A first-order literal L is *implied*, with clause C as *justification*, if L is the only literal of C that is not uniformly false in $I[\Gamma]$. SGGS ensures that every *I-all-true* clause in Γ is either a conflict clause or the justification of its selected literal. To this end, SGGS uses *assignment functions* to keep track of the dependence of *I-true* literals on *I-false* selected literals: an *I-all-true* clause whose literals are all assigned to *I-false* selected literals is a conflict clause; an *I-all-true* clause whose literals, except the selected one, are assigned, is a justification.

Example 3.4. Continuing Example 3.3, literals $\neg P(x)$ and $\neg P(a)$ are assigned to $[P(x)]$; literal $\neg R(a, b)$ is assigned to $[R(a, x)]$; and the last clause in Γ is in conflict with $I[\Gamma]$ as all its literals are assigned.

All SGGS clause sequences in the above examples are generated by applications of the *SGGS-extension* inference rule, that adds to the sequence an instance E of a clause $C \in S$ and selects one of its literals [18]. The instance E is built in order to capture the ground instances of C such that $I[\Gamma] \not\models C$, so that the resulting sequence (e.g., ΓE) will satisfy them.

Example 3.5. In Example 3.2, clause $\neg R(x, f(x)) \vee [R(f(x), x)]$ is an instance of input clause $\neg R(x, y) \vee R(y, x)$. SGGS generates it by unifying literal $\neg R(x, y)$ in this input clause with the selected literal $[R(x, f(x))]$ which is already on the trail. Recall that every first-order clause has its own variables. Let us rename the variables of the input clause as $\neg R(u, v) \vee R(v, u)$. Then the applied most general unifier (mgu) is $\alpha = \{u \leftarrow x, v \leftarrow f(x)\}$. The meaning is as follows. Initially, because I is all-negative, the second and the third clause in S are satisfied by I , but the first one is not. Thus, SGGS generates $\Gamma = [R(x, f(x))]$ by an SGGS-extension

with empty mgu. At this point, $I[\Gamma]$ satisfies the first and the second clause, but not the third one. Which ground instances of the third clause have been lost? Precisely those where $\neg R(u, v)$ unifies with $[R(x, f(x))]$. Thus, SGGS extends the model to recapture these instances by adding $\neg R(x, f(x)) \vee [R(f(x), x)]$.

However, it is not always the case that an SGGS-extension adds a clause E whose selected literal extends $I[\Gamma]$, because E may be a conflict clause. In such a case, SGGS *explains* the conflict by a restricted form of first-order resolution, called *SGGS-resolution*. SGGS-resolution resolves an I -false literal L in E with the implied I -true literal M , whose selection in Γ makes L uniformly false in $I[\Gamma]$. Thus, SGGS-resolution resolves the conflict clause E with the I -all-true clause D that is the justification of M in Γ . The resolvent is still in conflict. This series of *explanation inferences* by SGGS-resolution terminates when either the empty clause \square or an I -all-true conflict clause is generated.

The generation of \square signals that the input set S is unsatisfiable. Otherwise, SGGS applies an inference rule called *SGGS-move*, that *moves* the I -all-true conflict clause, say $E[L]$, to the left of the clause $D[M]$ whose selected I -false literal M makes E 's I -true selected literal L uniformly false in $I[\Gamma]$. The effect of this SGGS-move inference is to *learn* $E[L]$ and *solve* the conflict by *flipping* the truth value of *all* ground instances of L . At this point, $D[M]$ is in conflict, so that SGGS-resolution intervenes to resolve $E[L]$ and $D[M]$ upon L and M . Prior to the move, SGGS may *partition* $D[M]$ by $E[L]$ as in the following example.

Example 3.6. Continuing Example 3.4, we can see why $\neg R(a, b)$ is selected in conflict clause $\neg P(a) \vee [\neg R(a, b)]$: in an I -all-true conflict clause, SGGS prescribes to select the literal that is assigned *rightmost*, so that when the clause moves left to solve the conflict, the only literal in the clause that will be unassigned is the selected one, and the clause changes status from conflict clause to learned justification of an implied literal. The move consists of moving $\neg P(a) \vee [\neg R(a, b)]$ to the left of $\neg P(x) \vee [R(a, x)]$. However, SGGS does not do that, because changing the truth value of all ground instances of $[R(a, x)]$ in order to satisfy $[\neg R(a, b)]$ is too much. The philosophy of SGGS is to be *conflict-driven* and change $I[\Gamma]$ only as far as it is needed to solve the conflict. SGGS *partitions* $\neg P(x) \vee [R(a, x)]$ by $\neg P(a) \vee [\neg R(a, b)]$ producing $\Gamma = [P(x), x \neq b \triangleright \neg P(x) \vee [R(a, x)], \neg P(b) \vee [R(a, b)], \neg P(a) \vee [\neg R(a, b)]]$. Next, SGGS-move yields $\Gamma = [P(x), x \neq b \triangleright \neg P(x) \vee [R(a, x)], \neg P(a) \vee [\neg R(a, b)], \neg P(b) \vee [R(a, b)]]$. SGGS-resolution resolves $\neg P(a) \vee [\neg R(a, b)]$ and $\neg P(b) \vee [R(a, b)]$ to generate $\Gamma = [P(x), x \neq b \triangleright \neg P(x) \vee [R(a, x)], \neg P(a) \vee [\neg R(a, b)], \neg P(b) \vee [\neg P(a)]]$, where the resolvent $\neg P(b) \vee [\neg P(a)]$ is another I -all-true conflict clause. The selection of $\neg P(a)$ is arbitrary, since both $\neg P(b)$ and $\neg P(a)$ are assigned to $[P(x)]$.

As shown in the above example, in SGGS-resolution, the resolvent *replaces* the parent that is not I -all-true. In other words, the resolvent replaces the conflict clause, not the justification, as in CDCL. All clauses that have literals assigned to the deleted resolution parent are also deleted. Partitioning a clause $D[M]$ by a clause $E[L]$ replaces $D[M]$ by a *partition*, $D_1[M_1], \dots, D_n[M_n]$, that is, a set of clauses that together represent the same ground instances as $D[M]$, but have *disjoint* selected literals. Furthermore, the set of ground instances of $atom(L)$ is equal to the set of ground instances of $atom(M_j)$ for some j , $1 \leq j \leq n$, where $atom(L)$ denotes the atom of literal L . In other words, partitioning $D[M]$ by $E[L]$ splinters $D[M]$ in such a way to expose the non-empty intersection between the ground instances of L and those of M , where intersection ignores sign. Partitioning introduces *constraints*, that are a kind of *Herbrand constraints* [15, 18].

Example 3.7. Continuing Example 3.6, clause $\neg P(b) \vee [\neg P(a)]$ partitions clause $[P(x)]$, yielding $\Gamma = x \neq a \triangleright [P(x), [P(a)], \neg P(a) \vee [\neg R(a, b)], \neg P(b) \vee [\neg P(a)]]$, where $x \neq$

$b \triangleright \neg P(x) \vee [R(a, x)]$ has been deleted: SGGs allows us to delete a clause that has a literal (here $\neg P(x)$) assigned to a clause (here $[P(x)]$) that gets partitioned. The alternative is to recursively partition $x \neq b \triangleright \neg P(x) \vee [R(a, x)]$ into $\neg P(a) \vee [R(a, a)]$ and $x \neq b, x \neq a \triangleright \neg P(x) \vee [R(a, x)]$, and assign $\neg P(a)$ to $[P(a)]$ and $x \neq b, x \neq a \triangleright \neg P(x)$ to $x \neq a \triangleright [P(x)]$. Note that $x \neq b \triangleright \neg P(x) \vee [R(a, x)]$ cannot simply remain in Γ , because $x \neq b \triangleright \neg P(x)$ has nowhere to be assigned after $[P(x)]$ has been partitioned. Resuming from the above Γ , an SGGs-move step generates $\Gamma = x \neq a \triangleright [P(x)], \neg P(b) \vee [\neg P(a)], [P(a)], \neg P(a) \vee [\neg R(a, b)]$. Then SGGs-resolution resolves $\neg P(b) \vee [\neg P(a)]$ and $[P(a)]$ to produce $\Gamma = x \neq a \triangleright [P(x)], \neg P(b) \vee [\neg P(a)], [\neg P(b)]$, where the resolvent $\neg P(b)$ replaces the non-*I*-all-true parent $P(a)$. Clause $\neg P(a) \vee [\neg R(a, b)]$ is deleted too, because its literal $\neg P(a)$ was assigned to the deleted resolution parent $P(a)$.

Another reason for deleting $\neg P(a) \vee [\neg R(a, b)]$ in the above example is that it is *disposable*: in SGGs a clause C in $\Gamma C \Gamma'$ is *disposable*, if it is satisfied by $I[\Gamma]$. SGGs features an inference rule, called *SGGS-deletion*, that deletes all disposable clauses in the given SGGs clause sequence. A typical SGGs search plan applies SGGs-deletion eagerly.

Example 3.8. Continuing Example 3.7 from $\Gamma = x \neq a \triangleright [P(x)], \neg P(b) \vee [\neg P(a)], [\neg P(b)]$, clause $[\neg P(b)]$ is in conflict, and should move left of the clause that makes it false, namely $x \neq a \triangleright [P(x)]$. As before, SGGs partitions before moving: clause $[\neg P(b)]$ partitions clause $x \neq a \triangleright [P(x)]$, generating $\Gamma = x \neq a, x \neq b \triangleright [P(x)], [P(b)], \neg P(b) \vee [\neg P(a)], [\neg P(b)]$. By an application of SGGs-move we get $\Gamma = x \neq a, x \neq b \triangleright [P(x)], [\neg P(b)], [P(b)], \neg P(b) \vee [\neg P(a)]$. Then SGGs-resolution resolves $[\neg P(b)]$ and $[P(b)]$ yielding $\Gamma = x \neq a, x \neq b \triangleright [P(x)], [\neg P(b)], \square, \neg P(b) \vee [\neg P(a)]$, where the generation of \square terminates the derivation, as unsatisfiability has been discovered.

Fairness of an SGGs-derivation ensures that inferences that are infinitely often possible are not neglected. It also ensures that every conflict is solved *before* further SGGs-extensions, which is another similarity with CDCL, where the procedure does not venture new decisions when an extant conflict needs to be solved.

SGGs is *refutationally complete*: if the input set of clauses S is unsatisfiable, any fair SGGs-derivation from S is a refutation. SGGs is also *model complete* in the limit: if S is satisfiable, the *limiting sequence* of any fair SGGs-derivation from S represents a model of S , where both limiting sequence and derivation may be infinite, because first-order unsatisfiability is only semi-decidable. The limiting sequence is the *limit* of the derivation: the concept of limit of a derivation is commonly used in first-order theorem proving to give meaning to possibly infinite derivations, and the notion of limiting sequence defines this concept for SGGs [18].

SGGs is flexible with respect to *goal-sensitivity*. Assume that S was obtained by transforming $H \cup \{\neg\varphi\}$, from a problem $H \models \varphi$, into clausal form. Then SGGs is goal-sensitive, if the initial interpretation I satisfies the clauses issued from the reduction of H to clausal form, but not those issued from the reduction of $\neg\varphi$ to clausal form. SGGs is *proof confluent*, because it gets out of conflict by moving a clause in Γ , without undoing inference steps by backtracking or backjumping. This suggests that a backtracking search may not be an essential ingredient of conflict-driven reasoning.

3.2 A Taste of the CDSAT Framework

CDSAT, for *Conflict-Driven SATisfiability*, extends MCSAT to *generic* combinations of disjoint theories $\mathcal{T}_1, \dots, \mathcal{T}_n$ [10, 11]. This extension is crucial for the development of the conflict-driven paradigm, because problems from applications seldom involve only one theory. Furthermore, it is a major extension, because the MCSAT calculus [27] is not a combination calculus.

Conflict-driven \mathcal{T}_k -satisfiability procedures cannot be combined as black-boxes as done in equality sharing, if they are to retain their conflict-driven character. Neither can they be combined in a hierarchic framework such as DPLL(\mathcal{T}), where only the CDCL-based SAT-solver operates on the trail, and the \mathcal{T}_k -solvers are satellites that submit \mathcal{T}_k -lemmas and signal \mathcal{T}_k -conflicts to the SAT-solver. All conflict-driven \mathcal{T}_k -satisfiability procedures need to access the trail on a par with the SAT-solver. One needs to clarify how multiple conflict-driven \mathcal{T}_k -satisfiability procedures can cooperate, and which requirements the theories and their solvers should fulfill, in order to ensure the *soundness*, *completeness*, and *termination* of the combined system. Furthermore, a combination may include a theory \mathcal{T}_k for which a conflict-driven \mathcal{T}_k -satisfiability procedure is not available, so that one also needs to address the issue of a “mixed” combination, involving some conflict-driven \mathcal{T}_k -satisfiability procedures and some black-box \mathcal{T}_k -satisfiability procedures. The CDSAT framework for *conflict-driven theory combination* solves all these problems.

The CDSAT approach considers the CDCL procedure as one of n conflict-driven \mathcal{T} -satisfiability procedures to be combined. Accordingly, CDSAT regards atoms, literals, clauses and even formulæ as *terms* of sort **prop**, from proposition. The notion of *assignment* is generalized to allow assignments to *terms*, including non-variable terms. This applies to both terms of sort **prop**, that is, Boolean terms, and first-order terms of any other sort. Since the theories have different signatures and the signatures are mixed in an assignment, CDSAT defines the *theory view* of an assignment for each theory. Input problems are also read as assignments: if the input problem calls for determining the satisfiability of a set of clauses $S = \{l_1, \dots, l_m\}$, CDSAT starts with the initial assignment $\{l_1 \leftarrow \text{true}, \dots, l_m \leftarrow \text{true}\}$.

Domain values such as 3 are not necessarily in the signatures of the theories. Therefore, CDSAT assumes *theory extensions* that add to the signatures as many constants as needed to name the domain values (e.g., all the integers), including truth values. These extensions are required to be *conservative*, meaning that reasoning in the extension does not change the problem: an extension \mathcal{T}_k^+ of theory \mathcal{T}_k is *conservative*, if any \mathcal{T}_k^+ -unsatisfiable set S of clauses is also \mathcal{T}_k -unsatisfiable. If CDSAT discovers \mathcal{T}_k^+ -unsatisfiability, the problem is \mathcal{T}_k -unsatisfiable; if the problem is \mathcal{T}_k -satisfiable, there is a \mathcal{T}_k^+ -model that CDSAT can build.

A conflict-driven \mathcal{T}_k -satisfiability procedure can be viewed as an inference system and a conflict-driven search plan. Since the conflict-driven search is performed centrally by CDSAT for all theories, every theory needs to provide only an inference system. Thus, a key abstraction in CDSAT is to view a combination of theories $\mathcal{T}_1, \dots, \mathcal{T}_n$ as a combination of *inference systems* $\mathcal{I}_1, \dots, \mathcal{I}_n$, called *theory modules*. As CDSAT works with assignments, so do the theory modules: an inference deduces a Boolean assignment from a set of assignments of any sorts. A black-box \mathcal{T}_k -satisfiability procedure is included as a theory module whose only inference rule consists of invoking the procedure to detect that a set of assignments is inconsistent in the theory. Therefore, it is precisely the abstraction of combining inference systems rather than procedures that allows CDSAT to handle both conflict-driven and black-box \mathcal{T}_k -satisfiability procedures, thereby generalizing not only MCSAT, but also equality sharing. Theory modules for propositional logic, and the quantifier-free fragments of the theory of equality, linear rational arithmetic, and the theory of arrays with extensionality are provided as examples [10, 11].

The trail is defined as a sequence of singleton assignments, so that it can be seen as an assignment by forgetting the order. Every assignment A in the trail is either a *decision* or a *justified assignment*. A decision is placed on the trail by an application of the *Decide* rule of the CDSAT inference system; it represents a guess. A justified assignment A is associated with a *justification*, given by a set J of assignments that appear on the trail before A . A justified assignment A may be due to an application of the *Deduce* rule of the CDSAT inference system,

that places A on the trail with justification J , if a theory module \mathcal{I}_k infers A from J . Initial assignments are justified assignments with empty justification. Other justified assignments are due to the conflict-solving rules of the CDSAT inference system [10, 11].

In CDSAT a *conflict* is a set of assignments. CDSAT develops further the intuition, already in MCSAT and SGGS, that the essence of a conflict-driven approach is the *explanation* of conflicts. It is for the purpose of explanation that every assignment A in the trail that is not a decision has a justification J . If A becomes part of a conflict, it can be explained away by replacing it with J . Propositional resolution, as in a CDCL explanation, is a special instance of this explanation mechanism performed by the *Resolve* rule of the CDSAT inference system.

The CDSAT inference system is parametrized by a *global basis*, which is the source of new terms that theory modules can employ in their inferences. CDSAT is *sound* and *complete* for combinations of disjoint theories, assuming that at least one of the theories has information about the cardinalities of the domains to interpret the shared sorts. Assuming that all theories are stably infinite is a special way of having this information. Finiteness of the global basis ensures termination.

Clearly, there is no reason to restrict CDSAT to inputs of the form $\{l_1 \leftarrow true, \dots, l_m \leftarrow true\}$. CDSAT accepts input problems containing both Boolean and first-order assignments. For example, one may need to decide the satisfiability of a quantifier-free formula φ in a combination of theories, given an assignment to some of the free variables in φ , whether propositional or first-order. Therefore, CDSAT addresses a more general problem than SMT, that we call SMA for *satisfiability modulo assignments*. For SMA problems, the input format presupposes the theory extensions.

3.3 A New Class of Problems: Satisfiability Modulo Assignment

During the search, a conflict-driven reasoner maintains a partial candidate model represented by an assignment. This suggests the more general problem of *satisfiability modulo assignments* (SMA), defined as the problem of deciding the satisfiability of a set S of clauses modulo a theory \mathcal{T} with respect to an initial assignment J to some of the terms in S , including *both* propositional and first-order terms. If J is empty, SMA reduces to SMT; if both J and \mathcal{T} are empty, SMA reduces to SAT, while an intermediate state of a SAT or SMT search is an SMA instance. In CDSAT, there is no distinction between S and J , that are united to form the input assignment.

The answer to an SMA problem is either “satisfiable” with a model of S extending J , or “unsatisfiable” with a set of clauses E that follows from S and is false in J . The set E is an *explanation*, because it explains why S is unsatisfiable under J . The concept of *explanation* generalizes those of *unsatisfiable core* and *interpolant*. In SAT, an *unsatisfiable core* of S is a set of clauses that follows from S and is unsatisfiable. An unsatisfiable core explains why S is unsatisfiable, and the smaller it is with respect to the subset ordering \subseteq , the more precise it is regarded. If J is also written as a set of clauses, a (*reverse*) *interpolant* of S and J is a formula that follows from S and is inconsistent with J (see [12] for a survey of interpolation systems for ground proofs). MCSAT uses interpolants in arithmetic as explanations [27].

SMA arises in several contexts, such as *enumeration* of models, *parallelization*, and *optimization*. The models of a SAT or SMT problem can be enumerated by solving a series of SMA problems where each initial assignment J excludes the models already found. Approaches to parallel SAT by distributed search (e.g. Section 4.1 in [6]) solve a SAT problem with input set S , by solving in parallel multiple instances of SMA with input set S and initial assignments J each containing a distinct *guiding path* [59] or *cube* [33]. An optimization problem can be

approached by solving a series of SMA problems where each initial assignment J contains information generated by the previous runs, in such a way that the series converges towards an optimal solution. For example, this concept appeared in the presentation of [29] about adapting to optimization the satisfiability procedure of [38, 28] for the theory of algebraic reals.

4 Discussion

The big picture sees various approaches to extend conflict-driven reasoning to the first-order level. From the SMT side, the process started with generalizations of the Conflict-Driven Clause Learning (CDCL) procedure from propositional logic to several fragments of arithmetic [45, 39, 22, 37, 38, 32]. These methods offer *conflict-driven \mathcal{T} -satisfiability procedures*. By being generic with respect to the theory, *Model-Constructing Satisfiability* (MCSAT) encompasses these predecessors, and by integrating theory reasoning and propositional reasoning in all aspects of deduction and search, it provides a paradigm for *conflict-driven \mathcal{T} -decision procedures* [27, 36, 35, 58]. In turn, *Conflict-Driven Satisfiability* (CDSAT) generalizes MCSAT to generic combinations of theories and satisfiability modulo assignments (SMA) problems, where a partial assignment may also be part of the input problem [10, 11].

From the theorem-proving side, *Semantically-Guided Goal-Sensitive* (SGGS) reasoning [16, 17, 18] is a method that lifts conflict-driven reasoning to full first-order logic. A comparison between SGGS and ordering-based theorem provers (e.g., [56, 44, 40, 51]) is premature, because SGGS still needs to be implemented and extended to first-order logic with equality. The point of SGGS is not to reprove the theorems that other approaches have already conquered, but rather to explore new domains or hard problems, where its conflict-driven character may be rewarding. The identification of such classes of problems is also an objective. Similarities between SGGS and CDSAT include the notion of mapping a literal, in SGGS, or an assignment, in CDSAT, to the literals, or assignments, respectively, that it depends on, and the notion that a model be part of the input problem, as SGGS assumes an initial interpretation for semantic guidance, while CDSAT accepts SMA problems. The future may witness further convergence.

Acknowledgments This paper was written while the author was visiting as a visiting professor the School of Computer Science of the University of Manchester, in Manchester, England, UK, and as an international observer the Computer Science Laboratory of SRI International, in Menlo Park, California, USA: support from both institutions is greatly appreciated. The research and the visits were funded in part by grants “CooperInt 2016” and “Ricerca di base 2015” both from the Università degli Studi di Verona, in Verona, Italy, EU.

References

- [1] Gábor Alagi and Christoph Weidenbach. NRCL – a model building approach to the Bernays-Schönfinkel fragment. In Carsten Lutz and Silvio Ranise, editors, *Proceedings of the Tenth International Symposium on Frontiers of Combining Systems (FroCoS)*, volume 9322 of *Lecture Notes in Artificial Intelligence*, pages 69–84. Springer, 2015.
- [2] Alessandro Armando, Maria Paola Bonacina, Silvio Ranise, and Stephan Schulz. On a rewriting approach to satisfiability procedures: extension, combination of theories and an experimental appraisal. In Bernhard Gramlich, editor, *Proceedings of the Fifth International Workshop on Frontiers of Combining Systems (FroCoS)*, volume 3717 of *Lecture Notes in Artificial Intelligence*, pages 65–80. Springer, 2005.

- [3] Alessandro Armando, Maria Paola Bonacina, Silvio Ranise, and Stephan Schulz. New results on rewrite-based satisfiability procedures. *ACM Transactions on Computational Logic*, 10(1):129–179, 2009.
- [4] Clark Barrett, Robert Nieuwenhuis, Albert Oliveras, and Cesare Tinelli. Splitting on demand in SAT modulo theories. In Miki Hermann and Andrei Voronkov, editors, *Proceedings of the Thirteenth International Conference on Logic, Programming and Automated Reasoning (LPAR)*, volume 4246 of *Lecture Notes in Artificial Intelligence*, pages 512–526. Springer, 2006.
- [5] Maria Paola Bonacina. On theorem proving for program checking – Historical perspective and recent developments. In Maribel Fernández, editor, *Proceedings of the Twelfth International Symposium on Principles and Practice of Declarative Programming (PPDP)*, pages 1–11. ACM, 2010.
- [6] Maria Paola Bonacina. Parallel theorem proving. In Youssef Hamadi and Lakhdar Sais, editors, *Handbook of Parallel Constraint Reasoning*, volume in press of *Lecture Notes in Computer Science*, pages 177–233. Springer, Berlin, Germany, EU, 2018.
- [7] Maria Paola Bonacina and Mnacho Echenim. Rewrite-based satisfiability procedures for recursive data structures. In Byron Cook and Roberto Sebastiani, editors, *Proceedings of the Fourth Workshop on Pragmatics of Decision Procedures in Automated Reasoning (PDPAR) at the Fourth Federated Logic Conference (FLoC), August 2006*, volume 174(8) of *Electronic Notes in Theoretical Computer Science*, pages 55–70. Elsevier, Amsterdam, The Netherlands, EU, 2007.
- [8] Maria Paola Bonacina and Mnacho Echenim. On variable-inactivity and polynomial \mathcal{T} -satisfiability procedures. *Journal of Logic and Computation*, 18(1):77–96, 2008.
- [9] Maria Paola Bonacina, Ulrich Furbach, and Viorica Sofronie-Stokkermans. On first-order model-based reasoning. In Narciso Martí-Oliet, Peter Olveczky, and Carolyn Talcott, editors, *Logic, Rewriting, and Concurrency: Essays Dedicated to José Meseguer*, volume 9200 of *Lecture Notes in Computer Science*, pages 181–204. Springer, Berlin, Germany, EU, 2015.
- [10] Maria Paola Bonacina, Stéphane Graham-Lengrand, and Natarajan Shankar. A model-constructing framework for theory combination. Technical Report 99/2016, Dipartimento di Informatica, Università degli Studi di Verona, Verona, Italy, EU, November 2016. Also Technical Report of SRI International and INRIA - CNRS - École Polytechnique; revised November 2017.
- [11] Maria Paola Bonacina, Stéphane Graham-Lengrand, and Natarajan Shankar. Satisfiability modulo theories and assignments. In Leonardo de Moura, editor, *Proceedings of the Twenty-Sixth Conference on Automated Deduction (CADE)*, volume 10395 of *Lecture Notes in Artificial Intelligence*, pages 42–59. Springer, 2017.
- [12] Maria Paola Bonacina and Moa Johansson. Interpolation systems for ground proofs in automated deduction: a survey. *Journal of Automated Reasoning*, 54(4):353–390, 2015.
- [13] Maria Paola Bonacina, Christopher A. Lynch, and Leonardo de Moura. On deciding satisfiability by $DPLL(\Gamma + \mathcal{T})$ and unsound theorem proving. In Renate Schmidt, editor, *Proceedings of the Twenty-second International Conference on Automated Deduction (CADE)*, volume 5663 of *Lecture Notes in Artificial Intelligence*, pages 35–50. Springer, 2009.
- [14] Maria Paola Bonacina, Christopher A. Lynch, and Leonardo de Moura. On deciding satisfiability by theorem proving with speculative inferences. *Journal of Automated Reasoning*, 47(2):161–189, 2011.
- [15] Maria Paola Bonacina and David A. Plaisted. Constraint manipulation in SGGs. In Temur Kutsia and Christophe Ringeissen, editors, *Proceedings of the Twenty-Eighth Workshop on Unification (UNIF) at the Sixth Federated Logic Conference (FLoC)*, Technical Reports of the Research Institute for Symbolic Computation, pages 47–54. Johannes Kepler Universität Linz, 2014.
- [16] Maria Paola Bonacina and David A. Plaisted. SGGs theorem proving: an exposition. In Stephan Schulz, Leonardo De Moura, and Boris Konev, editors, *Proceedings of the Fourth Workshop on Practical Aspects in Automated Reasoning (PAAR) at the Sixth Federated Logic Conference (FLoC), July 2014*, volume 31 of *EPiC Series in Computing*, pages 25–38. EasyChair, 2015.
- [17] Maria Paola Bonacina and David A. Plaisted. Semantically-guided goal-sensitive reasoning: model

- representation. *Journal of Automated Reasoning*, 56(2):113–141, 2016.
- [18] Maria Paola Bonacina and David A. Plaisted. Semantically-guided goal-sensitive reasoning: inference system and completeness. *Journal of Automated Reasoning*, 59(2):165–218, 2017.
- [19] Aaron R. Bradley and Zohar Manna. *The Calculus of Computation - Decision Procedures with Applications to Verification*. Springer, Berlin, Germany, EU, 2007.
- [20] Robert Brummayer and Armin Biere. Lemmas on demand for the extensional theory of arrays. *Journal on Satisfiability, Boolean Modeling and Computation*, 6:165–201, 2009.
- [21] Chin-Liang Chang and Richard Char-Tung Lee. *Symbolic Logic and Mechanical Theorem Proving*. Academic Press, Cambridge, England, UK, 1973.
- [22] Scott Cotton. Natural domain SMT: A preliminary assessment. In Krishnendu Chatterjee and Thomas A. Henzinger, editors, *Proceedings of the Eighth International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS)*, volume 6246 of *Lecture Notes in Computer Science*, pages 77–91. Springer, 2010.
- [23] Martin Davis, George Logemann, and Donald Loveland. A machine program for theorem-proving. *Communications of the ACM*, 5(7):394–397, 1962.
- [24] Martin Davis and Hilary Putnam. A computing procedure for quantification theory. *Journal of the ACM*, 7:201–215, 1960.
- [25] Leonardo de Moura and Nikolaj Bjørner. Engineering DPLL(T) + saturation. In Alessandro Armando, Peter Baumgartner, and Gilles Dowek, editors, *Proceedings of the Fourth International Conference on Automated Reasoning (IJCAR)*, volume 5195 of *Lecture Notes in Artificial Intelligence*, pages 475–490. Springer, 2008.
- [26] Leonardo de Moura and Nikolaj Bjørner. Model-based theory combination. In Sava Krstić and Albert Oliveras, editors, *Proceedings of the Fifth International Workshop on Satisfiability Modulo Theories (SMT 2007)*, volume 198(2) of *Electronic Notes in Theoretical Computer Science*, pages 37–49. Elsevier, Amsterdam, The Netherlands, EU, 2008.
- [27] Leonardo de Moura and Dejan Jovanović. A model-constructing satisfiability calculus. In Roberto Giacobazzi, Josh Berdine, and Isabella Mastroeni, editors, *Proceedings of the Fourteenth International Conference on Verification, Model Checking and Abstract Interpretation (VMCAI)*, volume 7737 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2013.
- [28] Leonardo de Moura and Grant Olney Passmore. Computation over real closed infinitesimal and transcendental extensions of the rationals. In Maria Paola Bonacina, editor, *Proceedings of the Twenty-Fourth Conference on Automated Deduction (CADE)*, volume 7898 of *Lecture Notes in Artificial Intelligence*, pages 177–191. Springer, 2013.
- [29] Leonardo de Moura and Grant Olney Passmore. Exact global optimization on demand (presentation only). In Silvio Ghilardi, Viorica Sofronie-Stokkermans, and Ashish Tiwari, editors, *Notes of the Third Workshop on Automated Deduction: Decidability, Complexity, Tractability (ADDCT)*, pages 50–50, 2013. Available at <https://userpages.uni-koblenz.de/~sofronie/addct-2013/>, last seen on May 9, 2017.
- [30] Leonardo de Moura and Harald Rueß. Lemmas on demand for satisfiability solvers. In *Proceedings of the Fifth International Symposium on the Theory and Application of Satisfiability Testing (SAT)*, pages 244–251, 2002.
- [31] Bruno Dutertre and Leonardo de Moura. A fast linear arithmetic solver for DPLL(T). In Tom Ball and R. B. Jones, editors, *Proceedings of the Eighteenth International Conference on Computer Aided Verification (CAV)*, volume 4144 of *Lecture Notes in Computer Science*, pages 81–94. Springer, 2006.
- [32] Leopold Haller, Alberto Griggio, Martin Brain, and Daniel Kroening. Deciding floating-point logic with systematic abstraction. In Gianpiero Cabodi and Satnam Singh, editors, *Proceedings of the Twelfth International Conference on Formal Methods in Computer Aided Design (FMCAD)*. ACM and IEEE, 2012.
- [33] Marijn Heule, Oliver Kullmann, Siert Wieringa, and Armin Biere. Cube and conquer: guiding

- CDCL SAT solvers by lookaheads. In K. Eder, J. Lourenço, and O. Shehory, editors, *Proceedings of the Seventh Haifa Verification Conference (HVC)*, volume 7261 of *Lecture Notes in Computer Science*, pages 50–65. Springer, 2012.
- [34] Daniyar Itegulov, John Slaney, and Bruno Woltzenlogel Paleo. Scavenger 0.1: a theorem prover based on conflict resolution. In Leonardo de Moura, editor, *Proceedings of the Twenty-Sixth Conference on Automated Deduction (CADE)*, volume 10395 of *Lecture Notes in Artificial Intelligence*. Springer, 2017.
- [35] Dejan Jovanović. Solving nonlinear integer arithmetic with MCSAT. In Ahmed Bouajjani and David Monniaux, editors, *Proceedings of the Eighteenth International Conference on Verification, Model Checking and Abstract Interpretation (VMCAI)*, volume 10145 of *Lecture Notes in Computer Science*, pages 330–346. Springer, 2017.
- [36] Dejan Jovanović, Clark Barrett, and Leonardo de Moura. The design and implementation of the model-constructing satisfiability calculus. In Barbara Jobstman and Sandip Ray, editors, *Proceedings of the Thirteenth Conference on Formal Methods in Computer Aided Design (FMCAD)*. ACM and IEEE, 2013.
- [37] Dejan Jovanović and Leonardo de Moura. Cutting to the chase: solving linear integer arithmetic. In Nikolaj Bjørner and Viorica Sofronie-Stokkermans, editors, *Proceedings of the Twenty-Third International Conference on Automated Deduction (CADE)*, volume 6803 of *Lecture Notes in Artificial Intelligence*, pages 338–353. Springer, 2011.
- [38] Dejan Jovanović and Leonardo de Moura. Solving non-linear arithmetic. In Bernhard Gramlich, Dale Miller, and Ulrike Sattler, editors, *Proceedings of the Sixth International Joint Conference on Automated Reasoning (IJCAR)*, volume 7364 of *Lecture Notes in Artificial Intelligence*, pages 339–354. Springer, 2012.
- [39] Konstantin Korovin, Nestan Tsiskaridze, and Andrei Voronkov. Conflict resolution. In Ian P. Gent, editor, *Proceedings of the Fifteenth International Conference on Principles and Practice of Constraint Programming (CP)*, volume 5732 of *Lecture Notes in Computer Science*, pages 509–523. Springer, 2009.
- [40] Laura Kovács and Andrei Voronkov. First order theorem proving and Vampire. In Natasha Sharygina and Helmut Veith, editors, *Proceedings of the Twenty-Fifth International Conference on Computer-Aided Verification (CAV)*, volume 8044 of *Lecture Notes in Computer Science*, pages 1–35. Springer, 2013.
- [41] Sava Krstić and Amit Goel. Architecting solvers for SAT modulo theories: Nelson-Oppen with DPPL. In Frank Wolter, editor, *Proceedings of the Sixth International Symposium on Frontiers of Combining Systems (FroCoS)*, volume 4720 of *Lecture Notes in Artificial Intelligence*, pages 1–27. Springer, 2007.
- [42] João P. Marques Silva, Inês Lynce, and Sharad Malik. Conflict-driven clause learning SAT solvers. In Armin Biere, Marjin Heule, Hans Van Maaren, and Toby Walsh, editors, *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*, pages 131–153. IOS Press, Amsterdam, The Netherlands, EU, 2009.
- [43] João P. Marques Silva and Karem A. Sakallah. GRASP: A search algorithm for propositional satisfiability. *IEEE Transactions on Computers*, 48(5):506–521, 1999.
- [44] William W. McCune. Prover9 and Mace4, 2005–2010. <http://www.cs.unm.edu/~mccune/prover9/>, last seen on May 10, 2017.
- [45] Kenneth L. McMillan, A. Kuehlmann, and Mooly Sagiv. Generalizing DPPL to richer logics. In Ahmed Bouajjani and Oded Maler, editors, *Proceedings of the Twenty-First International Conference on Computer Aided Verification (CAV)*, volume 5643 of *Lecture Notes in Computer Science*, pages 462–476. Springer, 2009.
- [46] Matthew W. Moskewicz, Conor F. Madigan, Ying Zhao, Lintao Zhang, and Sharad Malik. Chaff: Engineering an efficient SAT solver. In David Blaauw and Luciano Lavagno, editors, *Proceedings of the Thirty-Ninth Design Automation Conference (DAC)*, pages 530–535. ACM and IEEE, 2001.

- [47] Greg Nelson. Combining satisfiability procedures by equality sharing. In Woodrow W. Bledsoe and Donald W. Loveland, editors, *Automatic Theorem Proving: After 25 Years*, pages 201–211. American Mathematical Society, Providence, Rhode Island, USA, 1983.
- [48] Greg Nelson and Derek C. Oppen. Simplification by cooperating decision procedures. *ACM Transactions on Programming Languages and Systems*, 1(2):245–257, 1979.
- [49] Robert Nieuwenhuis, Albert Oliveras, and Cesare Tinelli. Solving SAT and SAT modulo theories: from an abstract Davis-Putnam-Logemann-Loveland procedure to DPLL(T). *Journal of the ACM*, 53(6):937–977, 2006.
- [50] Ruzica Piskac, Leonardo de Moura, and Nikolaj Bjørner. Deciding effectively propositional logic using DPLL and substitution sets. *Journal of Automated Reasoning*, 44(4):401–424, 2010.
- [51] Stephan Schulz. System description: E 1.8. In Ken McMillan, Aart Middeldorp, and Andrei Voronkov, editors, *Proceedings of the Nineteenth International Conference on Logic, Programming and Automated Reasoning (LPAR)*, volume 8312 of *Lecture Notes in Artificial Intelligence*, pages 735–743. Springer, 2013.
- [52] Natarajan Shankar. Automated deduction for verification. *ACM Computing Surveys*, 41(4):40–96, 2009.
- [53] John Slaney and Bruno Woltzenlogel Paleo. Conflict resolution: a first-order resolution calculus with decision literals and conflict-driven clause learning. *Journal of Automated Reasoning*, in press:1–24, 2017. Published online on 24 February 2017 with DOI 10.1007/s10817-017-9408-6.
- [54] Aaron Stump, Clark W. Barrett, David L. Dill, and Jeremy Levitt. A decision procedure for an extensional theory of arrays. In Joseph Halpern, editor, *Proceedings of the Sixteenth IEEE Symposium on Logic in Computer Science (LICS)*. IEEE Computer Society Press, 2001.
- [55] Chao Wang, Franjo Ivančić, Malay Ganai, and Aarti Gupta. Deciding separation logic formulae by SAT and incremental negative cycle elimination. In Geoff Sutcliffe and Andrei Voronkov, editors, *Proceedings of the Twelfth International Conference on Logic, Programming and Automated Reasoning (LPAR)*, volume 3835 of *Lecture Notes in Artificial Intelligence*, pages 322–336. Springer, 2005.
- [56] Christoph Weidenbach, Dylana Dimova, Arnaud Fietzke, Rohit Kumar, Martin Suda, and Patrick Wischniewski. SPASS version 3.5. In Renate Schmidt, editor, *Proceedings of the Twenty-Second International Conference on Automated Deduction (CADE)*, volume 5663 of *Lecture Notes in Artificial Intelligence*, pages 140–145. Springer, 2009.
- [57] Steven A. Wolfman and Daniel S. Weld. The LPSAT engine and its application to resource planning. In Thomas Dean, editor, *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, volume 1, pages 310–316. Morgan Kaufmann Publishers, 1999.
- [58] Aleksandar Zeljić, Christoph M. Wintersteiger, and Philipp Rümmer. Deciding bit-vector formulas with mcSAT. In Nadia Creignou and Daniel Le Berre, editors, *Proceedings of the Nineteenth International Conference on Theory and Applications of Satisfiability Testing (SAT)*, volume 9710 of *Lecture Notes in Computer Science*, pages 249–266. Springer, 2016.
- [59] Hantao Zhang, Maria Paola Bonacina, and Jieh Hsiang. PSATO: a distributed propositional prover and its application to quasigroup problems. *Journal of Symbolic Computation*, 21(4–6):543–560, 1996.
- [60] Hantao Zhang and Mark E. Stickel. Implementing the Davis-Putnam method. *Journal of Automated Reasoning*, 24(1/2):277–296, 2000.