

SGGS Decision Procedures^{*}

Maria Paola Bonacina and Sarah Winkler

Università degli Studi di Verona, Verona, Italy
{mariapaola.bonacina,sarahmaria.winkler}@univr.it

Abstract. SGGS (Semantically-Guided Goal-Sensitive reasoning) is a conflict-driven first-order theorem-proving method which is refutationally complete and model complete in the limit. These features make it attractive as a basis for decision procedures. In this paper we show that SGGS decides the *stratified fragment* which generalizes EPR, the *PVD fragment*, and a new fragment that we dub *restrained*. The new class has the *small model property*, as the size of SGGS-generated models can be upper-bounded, and is also decided by hyperresolution and ordered resolution. We report on experiments with a termination tool implementing a restrainedness test, and with an SGGS prototype named *Koala*.

1 Introduction

Many applications of automated reasoning require to combine the decidability of satisfiability with an expressive logic. Since first-order theorem proving is only semidecidable, the quest for decidable fragments of first-order logic is key in advancing the field, and many classes of formulæ were shown decidable. Without claiming completeness (see [18,15,22] for surveys), we mention: the *Bernays-Schönfinkel class*, also known as EPR for *effectively propositional* [8,39,37,3,19]; the *Ackermann class* [2,23]; the *monadic class* with and without equality [2,23,5]; the *positive variable dominated* (PVD) fragment [17]; the *two-variable* fragment (FO^2) [21]; the *guarded* fragment [4,20]; the *modal fragment* [31,9], which is included in the EPR [22], FO^2 [32], and guarded [4] fragments; and the *stratified* fragment [1,25,36], which generalizes EPR. However, many theorem proving problems from the practice fall in none of these classes.

Example 1. Problem HWV036-2 from TPTP 7.3.0 [41] specifies a full-adder in 51 clauses, including for instance:

$$\begin{aligned} \neg \text{and}_{\text{ok}}(x) \vee \neg 1(\text{in}_1(x)) \vee \neg 1(\text{in}_2(x)) \vee 1(\text{out}_1(x)), & \quad \neg \text{lor}(x) \vee \text{or}_{\text{ok}}(x) \vee \text{error}(x), \\ \neg \text{halfadd}(x) \vee \text{connection}(\text{in}_1(x), \text{in}_1(\text{or}_1(x))), & \quad \neg \text{fulladd}(x) \vee \text{halfadd}(\text{h}_1(x)). \end{aligned}$$

This set is satisfiable, which means that termination of a theorem prover is *a priori* not guaranteed. However, it is neither EPR ($\exists^* \forall^* \varphi$ formulæ with φ quantifier- and function-free), nor Ackermann ($\exists^* \forall \exists^* \varphi$ formulæ with φ as above) nor FO^2 (only two variables, no functions), nor monadic (only unary predicates, no functions). One can also check that it is neither guarded nor stratified.

^{*} This research was funded in part by grant “Ricerca di base 2017” of the Università degli Studi di Verona. Authors are listed alphabetically.

As refutational completeness guarantees termination on unsatisfiable inputs, if one can prove termination of an inference system on satisfiable inputs in a certain class, any strategy given by that inference system and a fair search plan is a decision procedure for satisfiability in that class. Here we consider *Semantically-Guided Goal-Sensitive reasoning* (SGGS) [13,14], that is a refutationally complete instance-based theorem-proving method especially suitable for decision procedures: SGGS is *model-based* (it searches for a model by building candidates), *semantically guided* (the search is guided by a fixed initial interpretation), *conflict-driven* (it applies inferences such as resolution only to explain conflicts), *proof confluent* (it never needs to undo inferences), and *model complete in the limit* (if the input is satisfiable, the limit of the derivation represents a model), so that model generation is guaranteed if termination is.

This paper shows that SGGS decides the *stratified fragment* [1], which includes EPR and finds application in verification [1,25,36], and the *PVD fragment*. Then, we discover a *new decidable class* named the *restrained fragment*, and show that SGGS, ordered resolution, and hyperresolution all decide it. Since it is possible to compute *bounds on the size of SGGS-generated models*, this new class enjoys the *small model property*. We give a sufficient condition for membership in the restrained fragment that can be tested *automatically* by termination tools for rewriting. The relevance of this new class is evaluated empirically by applying this test to problems in TPTP. For instance, the axiomatization in Ex. 1, as well as all the TPTP problems including it, turn out to be restrained. We also summarize the outcomes of experiments with an SGGS prototype, named *Koala*, built reusing code from Konstantin Korovin’s *iProver* [24,26].

The paper is structured as follows. Since the stratified fragment has sorts, Sect. 2 presents SGGS for a language with sorts. Sect. 3 shows that SGGS decides the stratified fragment. In Sect. 4 we define the restrained fragment, establish the small model property, and prove that SGGS decides both this class and PVD. Ordered resolution and hyperresolution also decide restrained sets (Sect. 5). The experimental results are reported in Sect. 6, and Sect. 7 concludes the paper.

2 Preliminaries: SGGS for Many-Sorted Logic

Let S be a set of clauses in many-sorted logic with non-empty sorts (there is a ground term for every sort). We use \mathbf{a}, \mathbf{b} for constants, P, Q for predicates, f, g for functions, w, x, y, z for variables, t, u for terms, L, M for literals, $at(L)$ for L ’s atom, C, D for clauses, $Var(C)$ for the set of variables in C , α, σ for substitutions, I, J for interpretations, and we extend the at notation to sets of literals, clauses, and sets of clauses. C^+ and C^- are the disjunctions of the positive and negative literals in C , respectively; C is *positive* if $C = C^+$ and *negative* if $C = C^-$.

In SGGS, a clause C may have a constraint A , written $A \triangleright C$. An atomic *SGGS constraint* is *true*, *false*, $t \equiv u$, and $top(t) = f$, where \equiv is syntactic identity, and $top(t)$ is the top symbol of term t . The negation, conjunction, and disjunction of constraints is a constraint. Constraints in *standard form* are *true*, *false*, and conjunctions of distinct atomic constraints $x \neq y$ and $top(x) \neq f$.

Substitutions are *sort-preserving* ($x\sigma$ has the same sort as x) so that instantiation respects sorts. The set $Gr(A \triangleright C)$ of *constrained ground instances* (cgi) of $A \triangleright C$ is the set of ground instances of C that satisfy A . Literals $A \triangleright L$ and $B \triangleright M$ intersect if $at(Gr(A \triangleright L)) \cap at(Gr(B \triangleright M)) \neq \emptyset$, and are *disjoint* otherwise.

Example 2. In a signature with sorts $\{s_1, s_2\}$ and symbols $\mathbf{a} : s_1, \mathbf{b} : s_2, \mathbf{f} : s_1 \rightarrow s_2$, and $\mathbf{P} \subseteq s_2 \times s_2$, the only term of sort s_1 is \mathbf{a} , and \mathbf{b} and $\mathbf{f}(\mathbf{a})$ are the only terms of sort s_2 . Thus, $Gr(\mathbf{P}(x, y)) = \{\mathbf{P}(\mathbf{b}, \mathbf{b}), \mathbf{P}(\mathbf{f}(\mathbf{a}), \mathbf{b}), \mathbf{P}(\mathbf{b}, \mathbf{f}(\mathbf{a})), \mathbf{P}(\mathbf{f}(\mathbf{a}), \mathbf{f}(\mathbf{a}))\}$. For $\mathbf{P}(\mathbf{f}(x), y)$, with $x : s_1$ and $y : s_2$, constraint $top(x) \neq \mathbf{a}$ is unsatisfiable, while $top(y) \neq \mathbf{a}$ is valid. Then, $top(x) \neq \mathbf{a} \triangleright \mathbf{P}(\mathbf{f}(x), y)$ is equivalent to $false \triangleright \mathbf{P}(\mathbf{f}(x), y)$ and has no cgi's, while $top(y) \neq \mathbf{a} \triangleright \mathbf{P}(\mathbf{f}(x), y)$ is equivalent to $true \triangleright \mathbf{P}(\mathbf{f}(x), y)$, or simply $\mathbf{P}(\mathbf{f}(x), y)$, and has all cgi's, namely $\mathbf{P}(\mathbf{f}(\mathbf{a}), \mathbf{b})$ and $\mathbf{P}(\mathbf{f}(\mathbf{a}), \mathbf{f}(\mathbf{a}))$.

SGGS is *semantically guided* by an *initial interpretation* I : unless $I \models S$, SGGs seeks a model of S , by building candidate partial interpretations different from I , and using I as default to complete them. If the empty clause \perp arises in the process, unsatisfiability is reported. If I is the *all-negative interpretation* I^- that makes all negative literals true, SGGs tries to discover which positive literals need to be true to satisfy S , and dually if I is the *all-positive interpretation* I^+ . While I can be any Herbrand interpretation, I^+ and I^- suffice in this paper.

SGGS works with a *trail* of clauses $\Gamma = A_1 \triangleright C_1[L_1], \dots, A_n \triangleright C_n[L_n]$, where $C[L]$ means that literal $L \in C$ is *selected* in C . The length of Γ and its prefix of length j are denoted $|\Gamma|$ and $\Gamma|_j$, respectively. An SGGs trail Γ represents a partial interpretation $I^P(\Gamma)$: if Γ is empty, denoted ε , $I^P(\Gamma) = \emptyset$; otherwise, $I^P(\Gamma) = I^P(\Gamma|_{n-1}) \cup pcgi(A_n \triangleright L_n, \Gamma)$, where *pcgi* abbreviates *proper constrained ground instances*. A *pcgi* of $A_n \triangleright C_n[L_n]$ is a cgi $C[L]$ that is not satisfied by $I^P(\Gamma|_{n-1})$ (i.e., $I^P(\Gamma|_{n-1}) \cap C[L] = \emptyset$) and can be satisfied by adding L as $\neg L \notin I^P(\Gamma|_{n-1})$. For the selected literal, $pcgi(A_n \triangleright L_n, \Gamma) = \{L : C[L] \in pcgi(A_n \triangleright C_n[L_n], \Gamma)\}$. $I^P(\Gamma)$ is completed into an interpretation $I[\Gamma]$ by consulting I for the truth value of any literal undefined in $I^P(\Gamma)$.

A literal L is *uniformly false* in an interpretation J , if all $L' \in Gr(L)$ are false in J . Then, L is said to be *I -false* if it is uniformly false in I , and *I -true* if it is true in I . SGGs requires that if a clause in Γ has I -false literals, one is selected, so as to differentiate $I[\Gamma]$ from I . A clause whose literals are all I -true is an *I -all-true clause*, and only in such a clause an I -true literal is selected.

A *conflict clause* is one whose literals are all uniformly false in $I[\Gamma]$. SGGs ensures that every I -all-true clause $C[L]$ in Γ is either a conflict clause or the *justification* of its selected literal L , meaning that all literals of $C[L]$ except L are uniformly false in $I[\Gamma]$, so that L must be true in $I[\Gamma]$ to satisfy $C[L]$. In the latter case $C[L]$ is in the *disjoint prefix* of Γ , denoted $dp(\Gamma)$, which is the longest prefix such that $pcgi(A \triangleright C[L], \Gamma) = Gr(A \triangleright C[L])$ for all its clauses $A \triangleright C[L]$.

An *SGGS-derivation* is a series of trails $\Gamma_0 \vdash \Gamma_1 \vdash \dots \Gamma_j \vdash \dots$, where $\Gamma_0 = \varepsilon$, and $\forall j, j > 0$, an SGGs-inference generates Γ_j from Γ_{j-1} and S . If $\perp \notin \Gamma$ and $I[\Gamma] \not\models S$, SGGs has *two ways to make progress*. If $\Gamma = dp(\Gamma)$, the trail is in order, but $I[\Gamma] \not\models C'$ for some $C' \in Gr(C)$ and $C \in S$. Then, SGGs applies *SGGS-extension* to generate from C and Γ a clause $A \triangleright E$, such that E is an instance

of C and $C' \in Gr(A \triangleright E)$. If $\Gamma \neq dp(\Gamma)$, the trail needs repair: either there is a conflict, or there are intersections between selected literals to be removed by *SGGS-splitting*. The *SGGS-extension rules* specialize the *SGGS-extension scheme* ([14, Def. 12]) of which we give here the instance for I based on sign:

Definition 1. *Given input clause set S and trail Γ , if there is a clause $C \in S$ such that for all its I -true literals L_1, \dots, L_n ($n \geq 0$) there are clauses $B_1 \triangleright D_1[M_1], \dots, B_n \triangleright D_n[M_n]$ in $dp(\Gamma)$, such that literals M_1, \dots, M_n are I -false, and $\forall j, 1 \leq j \leq n, L_j \alpha = \neg M_j \alpha$ with simultaneous most general unifier (mgu) α , then SGGS-extension adds $A \triangleright E = (\bigwedge_{j=1}^n B_j \alpha) \triangleright C \alpha$ to Γ .*

SGGS-splitting decomposes a clause into instances to isolate and remove intersections between literals, and it is the only rule that introduces constraints. Splitting a ground clause is trivial and never done. Let $A \triangleright C[L]$ be a clause where A is satisfiable. Roughly speaking (see [14, Sect. 3.2]), a *partition* of $A \triangleright C[L]$ is a set $\{A_i \triangleright C_i[L_i]\}_{i=1}^n$ such that $Gr(A \triangleright C) = \bigcup_{i=1}^n \{Gr(A_i \triangleright C_i)\}$ and the literals $A_i \triangleright L_i$ are pairwise disjoint. Adding predicate symbol $Q \subseteq s_1 \times s_2$ to Ex. 2, a partition of $[P(f(x), y) \vee Q(x, y)] \vee Q(x, b)$ is $\{[P(f(x), b)] \vee Q(x, b), [P(f(x), f(a))] \vee Q(x, f(a))\}$. Given trail clauses $A \triangleright C[L]$ and $B \triangleright D[M]$, a *splitting* of C by D , denoted $split(C, D)$, is a partition of $A \triangleright C[L]$ such that $at(Gr(A_j \triangleright L_j))$ for some j is the intersection of $A \triangleright L$ and $B \triangleright M$, and all other $A_i \triangleright L_i$ are disjoint from $B \triangleright M$. *SGGS-splitting* replaces $A \triangleright C[L]$ with $split(C, D)$. Not all clauses in $split(C, D)$ need to be kept for completeness (see [14, Sect. 4.2]).

Clause $A_n \triangleright C_n[L_n]$ is *disposable*, if $I^p(\Gamma|_{n-1}) \models A_n \triangleright C_n[L_n]$, and *SGGS-deletion* removes *all* disposable clauses from the trail. The following example shows that SGGS halts, if applied to the EPR set used to show that another semantically-guided method, hyperresolution, cannot decide EPR.

Example 3. The set S consists of four clauses ([18, Ex. 4.8] and [15, Ex. 3.17]):

$$\begin{array}{llll} P(x, x, a) & (i), & P(x, y, w) \vee P(y, z, w) \vee \neg P(x, z, w) & (ii), \\ \neg P(x, x, b) & (iii), & P(x, z, w) \vee \neg P(x, y, w) \vee \neg P(y, z, w) & (iv). \end{array}$$

SGGS with all-negative initial interpretation I^- yields the following derivation:

$$\begin{array}{ll} \Gamma_0: \varepsilon \vdash \Gamma_1: [P(x, x, a)] & \text{extend (i)} \\ \vdash \Gamma_2: [P(x, x, a)], P(x, y, a) \vee [P(y, x, a)] \vee \neg P(x, x, a) & \text{extend (ii)} \\ \vdash \Gamma_3: [P(x, x, a)], P(x, x, a) \vee [P(x, x, a)] \vee \neg P(x, x, a), \\ \quad y \neq x \triangleright P(x, y, a) \vee [P(y, x, a)] \vee \neg P(x, x, a) & \text{split} \\ \vdash \Gamma_4: [P(x, x, a)], \\ \quad y \neq x \triangleright P(x, y, a) \vee [P(y, x, a)] \vee \neg P(x, x, a) & \text{delete} \end{array}$$

Since $I^- \not\models P(x, x, a)$, SGGS-extension puts it on the trail. As $I[\Gamma_1]$ satisfies $P(x, x, a)$, but no other positive literal, $I[\Gamma_1] \not\models (ii)$. Thus, SGGS-extension unifies the third literal of clause (ii) with $[P(x, x, a)]$ on the trail, producing Γ_2 ,

where an I^- -false (i.e., positive) literal is selected in the added clause (choosing the other makes no difference). As the selected literals intersect, the second clause gets split, yielding Γ_3 . The second clause in Γ_3 is disposable and SGGs-deletion removes it. Since $I[\Gamma_4] \models S$, the derivation halts reporting **satisfiable**. In contrast, hyperresolution generates infinitely many clauses of the form $P(x_1, x_2, \mathbf{a}) \vee P(x_2, x_3, \mathbf{a}) \vee \dots \vee P(x_{n-1}, x_n, \mathbf{a}) \vee P(x_n, x_1, \mathbf{a})$.

If a clause $A \triangleright C[L]$ added by SGGs-extension is in conflict with $I[\Gamma]$ and $C[L]$ contains I -false literals, *SGGS-resolution explains the conflict* by resolving $A \triangleright C[L]$ with a justification in $dp(\Gamma)$: the resolvent is also a conflict clause that replaces $A \triangleright C[L]$. As SGGs-extension (see [14, Def. 19]) ensures that all I -false literals of a conflict clause can be resolved away, conflict explanation generates either \perp or an I -all-true conflict clause $B \triangleright D[M]$. The conflict represented by $B \triangleright D[M]$ is *solved* by moving (*SGGS-move*) $B \triangleright D[M]$ to the left of the clause in $dp(\Gamma)$ whose selected literal makes M uniformly false in $I[\Gamma]$: the effect is to *flip* M from being uniformly false to being an implied literal.

Fairness of an SGGs-derivation involves several properties: SGGs-deletion and other clause removals are applied eagerly; trivial splitting is avoided; progress is made whenever possible; every SGGs-extension generating a conflict clause is *bundled* with explanation and conflict-solving inferences to eliminate the conflict before new extensions occur; and inferences applying to shorter prefixes of the trail are never neglected in favor of others applying to longer prefixes (see [14, Defs. 32, 37, and 39]). The *limit* of a fair derivation $\Gamma_0 \vdash \Gamma_1 \vdash \dots \Gamma_j \vdash \Gamma_{j+1} \vdash \dots$ is the longest trail Γ_∞ such that $\forall i, i \leq |\Gamma_\infty|$, there is an n_i such that $\forall j, j \geq n_i$, if $|\Gamma_j| \geq i$ then $\Gamma_j|_i$ is equivalent to $\Gamma_\infty|_i$ (see [14, Def. 50]). In words, all prefixes of the trail stabilize eventually. Both derivation and Γ_∞ may be infinite, but if the derivation halts at stage k , $\Gamma_\infty = \Gamma_k$. The following results employ an *SGGS-suitable* (i.e., total and extending the size ordering, hence well-founded) ordering on ground atoms (see [14, Def. 16]) and a *convergence ordering* $>^c$ on SGGs trails (see [14, Def. 46]):

- *Completeness*: For all input clause sets S , initial interpretations I , and fair SGGs-derivations, if S is satisfiable, $I[\Gamma_\infty] \models S$, and if S is unsatisfiable, $\perp \in \Gamma_k$ for some k (see [14, Thm. 9 and 11]).
- *Descending chain theorem*: A fair SGGs-derivation forms a descending chain $\Gamma_0 >^c \Gamma_1 >^c \dots >^c \Gamma_j >^c \Gamma_{j+1} \dots$ (see [14, Thm. 8]).
- *Finiteness of descending chains of length-bounded trails*: A chain $\Gamma_0 >^c \Gamma_1 >^c \dots \Gamma_j >^c \Gamma_{j+1} \dots$ where $\forall j, j \geq 0, |\Gamma_j| \leq n$, for some $n \geq 0$, is finite (see [14, Thm. 6 and Cor. 2]).

The gist of this paper is to find fragments where the length of SGGs trails is bounded so that termination of fair derivations is guaranteed.

3 SGGs Decides the Stratified Fragment

A way to ensure termination is to restrict an inference engine to produce only terms or atoms from a finite set \mathcal{B} , usually called a *basis*. For SGGs, let \mathcal{B} be a finite subset of the Herbrand base \mathcal{A} of the input clause set S .

Definition 2. An SGGS-trail $\Gamma = A_1 \triangleright C_1[L_1], \dots, A_n \triangleright C_n[L_n]$ is in \mathcal{B} if for all i , $1 \leq i \leq n$, $\text{at}(\text{Gr}(A_i \triangleright C_i)) \subseteq \mathcal{B}$.

An SGGS-derivation is in \mathcal{B} if all its trails are.

Lemma 1. If a fair SGGS-derivation $\Gamma_0 \vdash \Gamma_1 \vdash \dots \Gamma_j \vdash \Gamma_{j+1} \vdash \dots$ is in a finite basis \mathcal{B} , then for all j , $j \geq 0$, $|\Gamma_j| \leq |\mathcal{B}|+1$, and if the derivation halts at stage k , $k \geq 0$, then $|\Gamma_k| \leq |\mathcal{B}|$.

Proof. SGGS cannot do worse than generating a ground trail where every atom in \mathcal{B} appears selected with either sign: any trail with non-ground clauses will be shorter, because a non-ground clause covers many (possibly infinitely many) ground instances. By fairness, if the trail contains an intersection given by clauses $C[L]$ and $D[L]$, or $C[L]$ and $D[\neg L]$ with $L \in \mathcal{B}$, the clause on the right is either deleted eagerly by SGGS-deletion, or replaced with a resolvent by SGGS-resolution before SGGS-extension applies. Thus, there can be at most one such intersection, and the first claim follows. The second claim holds, because the intersection is removed by fairness prior to termination. \square

By the descending chain theorem and the finiteness of descending chains of length-bounded trails, the following general result follows:

Theorem 1. A fair SGGS-derivation in a finite basis is finite.

In order to apply this result, we need to find fragments that admit a finite basis. We begin with the *stratified fragment*. A signature is *stratified*, if there is a well-founded ordering $<_s$ on sorts, and for all functions $f: s_1 \times \dots \times s_n \rightarrow s$, it holds that $s <_s s_i$ for all $1 \leq i \leq n$ [1,25,36]. Thus, there are *no cycles over sorts* when applying functions. The signature from Ex. 2 is stratified with ordering $s_1 >_s s_2$. If a sentence over a stratified signature belongs to the $\exists^* \forall^*$ fragment, Skolemization only introduces constants and preserves stratification [25]. If there is only one sort, this fragment reduces to EPR, because stratification over a single sort implies that there are no function symbols. However, also stratified sentences with a prefix other than $\exists^* \forall^*$ can yield stratified clauses [33].

Example 4. Assume a stratified signature with sorts s_1 and s_2 such that $s_1 <_s s_2$, and symbols $f: s_1 \rightarrow s_2$, and $P \subseteq s_2 \times s_2$. The Skolemization of $\forall x \exists y. P(f(x), y)$ preserves stratification, as clause $P(f(x), g(x))$ with Skolem symbol $g: s_1 \rightarrow s_2$ is still stratified. On the other hand, the Skolemization of $\forall x \exists y. P(f(y), x)$ yields $P(f(g(x)), x)$ with Skolem symbol $g: s_2 \rightarrow s_1$, so that stratification is lost.

Given a set S of clauses whose signature is stratified, or stratified clause set for short, the Herbrand universe \mathcal{H} and the Herbrand base \mathcal{A} are *finite*, because stratification prevents building terms of unbounded depth [1]. Therefore, it suffices to pick \mathcal{A} itself as the finite basis for Thm. 1.

Theorem 2. Any fair SGGS-derivation from a stratified clause set S halts, is a refutation if S is unsatisfiable, and constructs a model of S if S is satisfiable.

However, SGGs-derivations can get exponentially long.

Example 5. Consider the following clause set S_k describing a k -digits binary counter [38, Def. 2.4.10]. Let Q be a predicate symbol of arity k , and for all i , $1 \leq i \leq k$, let $\bar{0}_i$, $\bar{1}_i$, and \bar{x}_i be i -tuples of 0's, 1's, and distinct variables x_1, \dots, x_i , respectively. S_k consists of the $k + 2$ clauses, for $1 \leq m \leq k$,

$$C_0: Q(\bar{0}_k) \quad C_m: \neg Q(\bar{x}_m, 0, \bar{1}_{k-m-1}) \vee Q(\bar{x}_m, 1, \bar{0}_{k-m-1}) \quad C_{k+1}: \neg Q(\bar{1}_k)$$

so that it is in EPR. Guided by I^- , SGGs generates a derivation

$$\begin{array}{ll} \Gamma_0: \varepsilon \vdash \Gamma_1: [Q(\bar{0}_k)] & \text{extend } (C_0) \\ \vdash \Gamma_2: \dots, \neg Q(\bar{0}_k) \vee [Q(\bar{0}_{k-1}, 1)] & \text{extend } (C_{k-1}) \\ \vdash \Gamma_3: \dots, \neg Q(\bar{0}_{k-1}, 1) \vee [Q(\bar{0}_{k-2}, 1, 0)] & \text{extend } (C_{k-2}) \\ \vdash \Gamma_4: \dots, \neg Q(\bar{0}_{k-2}, 1, 0) \vee [Q(\bar{0}_{k-2}, 1, 1)] & \text{extend } (C_{k-1}) \\ \dots & \dots \\ \vdash \Gamma_{2^k-1}: \dots, \neg Q(\bar{1}_{k-2}, 0, 1) \vee [Q(\bar{1}_{k-1}, 0)] & \text{extend } (C_{k-2}) \\ \vdash \Gamma_{2^k}: \dots, \neg Q(\bar{1}_{k-1}, 0) \vee [Q(\bar{1}_k)] & \text{extend } (C_{k-1}) \\ \vdash \Gamma_{2^k+1}: \dots, \neg Q(\bar{1}_{k-1}, 0) \vee [Q(\bar{1}_k)], [\neg Q(\bar{1}_k)] & \text{extend } (C_{k+1}) \end{array}$$

that simulates binary counting by adding a clause in each of these $2^k + 1$ steps till a conflict emerges. Then it takes another 2^{k+1} steps to detect unsatisfiability:

$$\begin{array}{ll} \vdash \Gamma_{2^k+2}: \dots, [\neg Q(\bar{1}_k)], \neg Q(\bar{1}_{k-1}, 0) \vee [Q(\bar{1}_k)] & \text{move} \\ \vdash \Gamma_{2^k+3}: \dots, [\neg Q(\bar{1}_k)], [\neg Q(\bar{1}_{k-1}, 0)] & \text{resolve} \\ \vdash \Gamma_{2^k+2}: \dots, [\neg Q(\bar{1}_{k-1}, 0)], \neg Q(\bar{1}_{k-2}, 0, 1) \vee [Q(\bar{1}_{k-1}, 0)], [\neg Q(\bar{1}_k)] & \text{move} \\ \vdash \Gamma_{2^k+3}: \dots, [\neg Q(\bar{1}_{k-1}, 0)], [\neg Q(\bar{1}_{k-2}, 0, 1)], [\neg Q(\bar{1}_k)] & \text{resolve} \\ \dots & \dots \\ \vdash \Gamma_{2^k+2}: [\neg Q(\bar{0}_k)], [Q(\bar{0}_k)], \dots & \text{move} \\ \vdash \Gamma_{2^k+2+1}: \perp, \dots & \text{resolve} \end{array}$$

Similar to positive resolution¹ [38, Thm. 2.4.12] or SCL [19], SGGs behaves exponentially, whereas resolution offers a refutation in $2k+1$ steps [38, Thm. 2.4.11], which shows that in EPR ground resolution (same as positive resolution for S_k) can do exponentially worse than resolution [34]. Encoding S_k in propositional logic requires exponentially many clauses, as each clause C_m , $1 \leq m \leq k$, has 2^m ground instances that need to be modeled as distinct propositional clauses. This indicates that generating instances is not good for this example.

4 SGGs Decides the Restrained Fragment

We begin with the notion of *ground-preserving* clause, which is convenient for sign-based semantic guidance.

¹ Every resolution step has a positive parent.

Definition 3. A clause C is positively ground-preserving if $\text{Var}(C) \subseteq \text{Var}(C^-)$, and negatively ground-preserving if $\text{Var}(C) \subseteq \text{Var}(C^+)$. A set of clauses is positively/negatively ground-preserving if all its clauses are, and ground-preserving if it is positively or negatively ground-preserving.

For example, $\neg P(x, y, z) \vee Q(y) \vee Q(f(z))$ and $\neg Q(x) \vee \neg Q(y)$ are positively ground-preserving, while the clauses in Ex. 5 are both positively and negatively ground-preserving. We say that I is *suitable* for a ground-preserving set S if either I is I^- and S is positively ground-preserving, or I is I^+ and S is negatively ground-preserving. If S is positively ground-preserving, its positive clauses are ground, positive hyperresolution only generates ground clauses ([12], Lem. 3), and similarly for the negative variant. We show that this also holds for SGGs.

Lemma 2. *If the input clause set S is ground-preserving and the initial interpretation is suitable for S , any fair SGGs-derivation from S is ground.*

Proof. We consider S positively ground-preserving and I^- (for the dual case one exchanges the signs). The proof is by induction on the length n of the derivation. The base case ($n = 0$) is vacuously true. The induction hypothesis is that the claim holds for a derivation of length n producing trail Γ . Let $\Gamma \vdash \Gamma'$ be the $(n+1)$ -th step. Since Γ is ground, $\Gamma \vdash \Gamma'$ cannot be a splitting step, because any splitting of a ground clause yields the clause itself, and fairness excludes such trivial splittings ([14, Defs. 32, 47, and 49]). If $\Gamma \vdash \Gamma'$ is an SGGs-resolution step, it is a ground resolution step, and also Γ' is ground. If $\Gamma \vdash \Gamma'$ is an SGGs-extension step, it adds an instance $C\alpha$ of a clause $C \in S$, where α is the simultaneous mgu of all I^- -true (i.e., negative) literals L_1, \dots, L_n in C with as many I^- -false (i.e., positive) selected literals M_1, \dots, M_n in Γ (see Def. 1). Since Γ is ground by induction hypothesis, the clauses containing M_1, \dots, M_n are ground and do not have constraints. Thus, $L_1\alpha, \dots, L_n\alpha$ are also ground. The I^- -false literals of $C\alpha$ are ground, because C is positively ground-preserving (i.e., $\text{Var}(C) \subseteq \text{Var}(C^-)$), so that all its variables get grounded by α . Hence $C\alpha$ and Γ' are ground. \square

The next example illustrates Lem. 2 and gives the intuition for restrainedness.

Example 6. Assume a positively ground-preserving set S which includes:

$$P(s^{10}(0), s^9(0)) \quad (i), \quad \neg P(s(s(x)), y) \vee P(x, s(y)) \quad (ii), \quad \neg P(s(0), 0) \quad (iii),$$

and I^- is the initial interpretation. SGGs starts with an extension that puts the positive clause $P(10, 9)$ on the trail, where we write n for $s^n(0)$. Subsequent extensions unify the negative literal in clause (ii) with some positive ground literal on the trail, so that new literals in added clauses are positive:

$$\begin{aligned} \Gamma_0: \varepsilon &\vdash \Gamma_1: [P(10, 9)] \\ &\vdash \Gamma_2: [P(10, 9), \neg P(10, 9) \vee [P(8, 10)]] \\ &\vdash \Gamma_3: [P(10, 9), \neg P(10, 9) \vee [P(8, 10)], \neg P(8, 10) \vee [P(6, 11)]]. \end{aligned}$$

The positive literals have decreasing number of symbols, matching the fact that $P(s(x), y) \succ P(x, s(y))$ in (ii) for \succ any lexicographic path ordering (LPO).

This suggests to strengthen ground-preservingness with an ordering (unrelated to the SGGs-suitable ordering of Sect. 2) to get a finite basis.

Definition 4. A quasi-ordering \succeq on terms and atoms is *restraining*, if (i) it is stable under substitution, (ii) the strict ordering $\succ = \succeq \setminus \preceq$ is well-founded, and (iii) the equivalence $\approx = \succeq \cap \preceq$ has finite equivalence classes.

Note that Condition (i) implies that \succ and \approx are stable under substitution. From now on, \succeq is a restraining quasi-ordering. Let \mathcal{A}_S be the set of ground atoms occurring in S , and \mathcal{A}_S^{\preceq} the subset of the Herbrand base \mathcal{A} of ground atoms upper-bounded by \mathcal{A}_S , so $\mathcal{A}_S^{\preceq} = \{L : L \in \mathcal{A}, \exists M \in \mathcal{A}_S \text{ with } M \succeq L\}$. By Conditions (ii) and (iii) in Def. 4, \mathcal{A}_S^{\preceq} is finite and thus can serve as basis.

Definition 5. A clause C is (strictly) *positively restrained* if it is *positively ground preserving*, and for all non-ground literals $L \in C^+$ there is a literal $M \in C^-$ such that $at(M) \succeq at(L)$ ($at(M) \succ at(L)$). A set of clauses is *positively restrained* if all its clauses are.

Negatively restrained clauses and clause sets are defined similarly, and a set of clauses is *restrained* if it is positively or negatively restrained. The set in Ex. 6 is strictly positively restrained. We see next the role of the *quasi-ordering*.

Example 7. Problem PLA030-1 in TPTP is neither stratified, nor monadic, nor guarded. Its clause $\text{differ}(x, y) \vee \neg \text{differ}(y, x)$ cannot be shown *strictly* restrained. Let \succ_{acrpo} be an AC-compatible [40] recursive path ordering with differ as an AC-symbol, meaning associative-commutative. The quasi-ordering \succeq_{acrpo} , built from \succ_{acrpo} and the AC-equivalence \approx_{AC} that has finite equivalence classes, satisfies $\text{differ}(x, y) \succeq_{\text{acrpo}} \text{differ}(y, x)$, and shows that PLA030-1 negatively restrained.

Restrainedness is undecidable in general, but decidable for fixed, suitable orderings. If S is restrained, a fair SGGs-derivation will be in \mathcal{A}_S^{\preceq} .

Lemma 3. If the input clause set S is restrained and the initial interpretation is suitable for S , any fair SGGs-derivation from S is in \mathcal{A}_S^{\preceq} .

Proof. We consider S positively ground-preserving and I^- (for the dual case one exchanges the signs). Since the set is restrained hence ground-preserving, the derivation is ground by Lem. 2 (†). The proof is by induction on the length n of the derivation, and it follows the same pattern as that of Lem. 2. Let $\Gamma \vdash \Gamma'$ be the $(n+1)$ -th step. By induction hypothesis, Γ is in \mathcal{A}_S^{\preceq} . If $\Gamma \vdash \Gamma'$ is an SGGs-resolution step, it is a ground resolution step which does not generate new atoms, and also Γ' is in \mathcal{A}_S^{\preceq} . If $\Gamma \vdash \Gamma'$ is an SGGs-extension step, it adds an instance $C\alpha$ of a clause $C \in S$, where α is the simultaneous mgu of all I^- -true (i.e., negative) literals $\neg L_1, \dots, \neg L_n$ in C with as many I^- -false (i.e., positive) selected literals M_1, \dots, M_n in Γ . The literals M_1, \dots, M_n are ground by (†), and by induction hypothesis they are in \mathcal{A}_S^{\preceq} . We have to show $at(C\alpha) \subseteq \mathcal{A}_S^{\preceq}$.

- For the negative literals $\neg L_1\alpha, \dots, \neg L_n\alpha$ we have $L_i\alpha = M_i\alpha = M_i \in \mathcal{A}_S^\prec$.
- Let L be a literal in C^+ . If L is ground, then $L\alpha = L \in \mathcal{A}_S \subseteq \mathcal{A}_S^\prec$. If L is not ground, by positive restrainedness there exists a $\neg L_i, 1 \leq i \leq n$, such that $L_i \succeq L$. By stability, $L_i\alpha \succeq L\alpha$. Since for all $i, 1 \leq i \leq n$, $M_i \in \mathcal{A}_S^\prec$ and $M_i = M_i\alpha = L_i\alpha \succeq L\alpha$, we have $L\alpha \in \mathcal{A}_S^\prec$. \square

Therefore, as \mathcal{A}_S^\prec is finite, Thm. 1 applies.

Theorem 3. *Any fair SGGS-derivation from a restrained clause set S with a suitable initial interpretation for S halts, is a refutation if S is unsatisfiable, and constructs a model of S if S is satisfiable.*

Restrainedness also makes it possible to derive an upper bound on the cardinality of a single-sorted model, defined as the cardinality of its domain. Let \mathcal{H}_S^\prec be the set $\mathcal{H}_S^\prec = \{t : t \text{ is a strict subterm of } L \text{ for some } L \in \mathcal{A}_S^\prec\}$.

Theorem 4. *A restrained satisfiable clause set S has a model of cardinality at most $|\mathcal{H}_S^\prec| + 1$ that can be extracted from the limit of any fair SGGS-derivation from S with suitable initial interpretation I .*

Proof. By Thm. 3 the derivation halts with some trail Γ , and by Lems. 2 and 3, Γ contains only ground clauses whose atoms are in \mathcal{A}_S^\prec . Since SGGS is model complete, $I[\Gamma] \models S$. Consider the following interpretation J with domain $\mathcal{H}_S^\prec \uplus \{u\}$, where u is a new constant symbol: For every constant symbol c we set $c^J = c$ if $c \in \mathcal{H}_S^\prec$, and $c^J = u$ otherwise. For every n -ary ($n \geq 1$) function symbol f , we set $f^J(t_1, \dots, t_n) = f(t_1^J, \dots, t_n^J)$ if $f(t_1, \dots, t_n) \in \mathcal{H}_S^\prec$, and $f^J(t_1, \dots, t_n) = u$ otherwise. For every predicate symbol P , $(t_1, \dots, t_n) \in P^J$ if $I[\Gamma] \models P(t_1, \dots, t_n)$. Note that J is well-defined because if $f(t_1, \dots, t_n) \in \mathcal{H}_S^\prec$ then t_1, \dots, t_n are also, hence all terms are interpreted in $\mathcal{H}_S^\prec \uplus \{u\}$. As J agrees with $I[\Gamma]$ on all atoms, $J \models S$, and it has cardinality $|\mathcal{H}_S^\prec| + 1$ by construction. \square

Therefore the restrained fragment also enjoys the *small model property*.

Example 8. The satisfiable clause set S (PUZ054-1 in TPTP) extending Ex. 6:

$$\begin{aligned} & \mathsf{P}(\mathsf{s}^{10}(0), \mathsf{s}^9(0)), \quad \neg \mathsf{P}(\mathsf{s}(\mathsf{s}(x)), y) \vee \mathsf{P}(x, \mathsf{s}(y)), \quad \neg \mathsf{P}(x, \mathsf{s}(\mathsf{s}(y))) \vee \mathsf{P}(x, \mathsf{s}(y)), \\ & \neg \mathsf{P}(\mathsf{s}(0), 0), \quad \neg \mathsf{P}(\mathsf{s}(x), \mathsf{s}(y)) \vee \mathsf{P}(\mathsf{s}(x), y), \end{aligned}$$

is neither in EPR, nor in FO^2 , nor in the monadic class. However, it can be shown strictly positively restrained by a Knuth-Bendix ordering (KBO) \succ with empty precedence and weights $w(\mathsf{P}) = 0$ and $w_0 = w(\mathsf{s}) = w(0) = 1$, where w_0 is the weight of variables. The largest atom in $\mathcal{A}_S = \{\mathsf{P}(\mathsf{s}^{10}(0), \mathsf{s}^9(0)), \mathsf{P}(\mathsf{s}(0), 0)\}$ has weight $w(\mathsf{P}(\mathsf{s}^{10}(0), \mathsf{s}^9(0))) = 21$. No atom L of the form $\mathsf{P}(\mathsf{s}^n(0), \mathsf{s}^m(0))$ in \mathcal{A}_S^\prec can have a subterm $\mathsf{s}^k(0)$ with $k > 19$, because otherwise $w(L) > w(\mathsf{P}(\mathsf{s}^{10}(0), \mathsf{s}^9(0)))$. Therefore, we have $\mathcal{H}_S^\prec = \{\mathsf{s}^i(0) : 0 \leq i \leq 19\}$ and there is a model of cardinality at most 21 by Thm. 4.

We next consider PVD [17]. Let $\text{depth}(C)$ be the maximum depth of an atom in clause C , and $\text{depth}_x(C)$ the maximum occurrence depth in C of $x \in \text{Var}(C)$.

Definition 6. A clause set S is in PVD if every clause $C \in S$ is positively ground-preserving and $\forall x \in \text{Var}(C^+)$ it holds that $\text{depth}_x(C^+) \leq \text{depth}_x(C^-)$.

Let \mathcal{A}_d be the subset of the Herbrand base \mathcal{A} containing all ground atoms whose depth does not exceed d .

Lemma 4. If the input clause set S is in PVD and has maximal depth d then any fair SGGs-derivation from S using I^- is in \mathcal{A}_d .

Proof. Since S is in PVD and hence ground-preserving, by Lem. 2 the derivation is ground (\dagger). The proof is by induction on the length n of the derivation, and follows that of Lem. 3 except for the extension case in the inductive step. Let $\Gamma \vdash \Gamma'$ be an SGGs-extension step as described in the proof of Lem. 3. The positive literals M_1, \dots, M_n are ground by (\dagger) and by induction hypothesis in \mathcal{A}_d . Since $L_i\alpha = M_i\alpha = M_i$ for all i , $1 \leq i \leq n$, the atoms $L_i\alpha$ in $C^-\alpha$ are ground and in \mathcal{A}_d . Let L be a literal in C^+ . Since S is in PVD, every variable in L occurs at a lower or equal depth in some L_i ($1 \leq i \leq n$). Therefore, $L\alpha$ is ground and its depth cannot exceed that of $L_i\alpha$, so that $L\alpha \in \mathcal{A}_d$. \square

Since \mathcal{A}_d is a finite basis, termination follows from Thm. 1.

Theorem 5. Any fair SGGs-derivation using I^- from a PVD set S halts, is a refutation if S is unsatisfiable, and constructs a model of S if S is satisfiable.

5 Ordered Resolution Decides the Restrained Fragment

In this section we work with the positively restrained fragment; the case for the negatively restrained one is symmetric. Let $>$ be a stable and well-founded ordering on literals, such that positive literals are maximal only in positive clauses, so that ordered resolution is positive ordered resolution. In this way, the ordering embeds the suitable sign-based semantic guidance for the positively restrained fragment. The ordering $>$ could be the extension of the restraining ordering \succ (see Def. 5) to literals, but does not have to be. We use the following notations [18]: $\text{Res}_>(S)$ denotes the set of ordered resolvents generated from parents in S ; $R^0_>(S) = S$, $R^{k+1}_>(S) = R^k_>(S) \cup \text{Res}_>(R^k_>(S))$, and $R^*_>(S) = \bigcup_{k \geq 0} R^k_>(S)$. We begin with an auxiliary lemma.

Lemma 5. If S is positively restrained, then for all $C \in R^*_>(S)$, for all $L \in C^+$ either (i) $L \in \mathcal{A}_S^<$, or (ii) $\text{at}(M) \succeq \text{at}(L)$ for some $M \in C^-$.

Proof. The proof is by induction on the stage k of the construction of $R^*_>(S)$. For $k=0$, the clauses in $R^0_>(S) = S$ satisfy the claim by the definitions of restrained-ness, \mathcal{A}_S , and $\mathcal{A}_S^<$. The induction hypothesis is that all clauses in $R^k_>(S)$ satisfy the claim. For the inductive step, let $C\sigma \vee D\sigma$ be a resolvent in $\text{Res}_>(R^k_>(S))$ generated from parents $\neg L \vee C$ and $L' \vee D$, where $\neg L\sigma$ and $L'\sigma$ are $>$ -maximal literals, and $L\sigma = L'\sigma$ for mgu σ . The clause $L' \vee D$ must be positive, otherwise $L'\sigma$ cannot be $>$ -maximal. Since $L' \vee D \in R^k_>(S)$, by induction hypothesis

$at(L' \vee D) \subseteq \mathcal{A}_{\mathcal{S}}^{\prec}(\dagger)$, which means $L' \vee D$ is ground, $(L' \vee D)\sigma = L' \vee D$, and all atoms in $D\sigma$ are in $\mathcal{A}_{\mathcal{S}}^{\prec}$. For the positive literals in $C\sigma$, let $M\sigma$ be one of them, so $M \in C^+$. Since $\neg L \vee C$ is in $R_{>}^k(S)$, by induction hypothesis, either (i) $M \in \mathcal{A}_{\mathcal{S}}^{\prec}$, or (ii) $M' \succeq M$ for some negative literal $\neg M'$ in $\neg L \vee C$. In case (i), M is ground, $M\sigma = M$, and $M\sigma \in \mathcal{A}_{\mathcal{S}}^{\prec}$. In case (ii), if $\neg M'$ occurs in C then $\neg M'\sigma \in C\sigma$, and $M'\sigma \succeq M\sigma$ holds by stability, so that the claim holds. Otherwise, $\neg M'$ is the resolved-upon literal $\neg L$ with $L\sigma = L'\sigma$. Thus, $L = M' \succeq M$, which implies $L\sigma \succeq M\sigma$ by stability, and since L' is ground, $L'\sigma = L'$. By (\dagger) , $L' \in \mathcal{A}_{\mathcal{S}}^{\prec}$. Since $L' = L'\sigma = L\sigma \succeq M\sigma$, we have $M\sigma \in \mathcal{A}_{\mathcal{S}}^{\prec}$ by definition of $\mathcal{A}_{\mathcal{S}}^{\prec}$. \square

Theorem 6. *Any fair ordered resolution derivation using $>$ from a restrained clause set S terminates, and is a refutation if S is unsatisfiable.*

Proof. Let S be positively restrained. We show that $R_{>}^*(S)$ is finite. The claim then follows from refutational completeness of ordered resolution. We first show the following (\dagger) : for every ordered resolvent $C\sigma \vee D\sigma$ in $R_{>}^*(S)$ from parents $\neg L \vee C$ and $L' \vee D$ and mgu σ , $L' \vee D$ is ground and positive, and $C\sigma \vee D\sigma$ has strictly fewer negative literals than $\neg L \vee C$. Indeed, by the definition of ordered resolution, $\neg L\sigma$ and $L'\sigma$ are $>$ -maximal in $(\neg L \vee C)\sigma$ and $(L' \vee D)\sigma$, respectively. The clause $L' \vee D$ must be positive, otherwise $L'\sigma$ cannot be $>$ -maximal. By Lem. 5, every positive literal in a positive clause is in $\mathcal{A}_{\mathcal{S}}^{\prec}$, and therefore it is ground. Thus, in the resolvent $C\sigma \vee D\sigma$ the literals in $D\sigma$ are positive, hence $C\sigma \vee D\sigma$ has fewer negative literals than $\neg L \vee C$.

Now suppose that $R_{>}^*(S)$ is infinite. Observation (\dagger) reveals that the number of negative literals decreases with every resolution step. Hence, if $R_{>}^*(S)$ is infinite, it must contain infinitely many positive clauses. By Lem. 5, these positive clauses are ground, and all their atoms are in the finite basis $\mathcal{A}_{\mathcal{S}}^{\prec}$. As repeated literals in ground clauses disappear by merging, the number of ground clauses that can be built from $\mathcal{A}_{\mathcal{S}}^{\prec}$ is finite. This contradicts $R_{>}^*(S)$ being infinite. \square

This result extends to other positive resolution strategies.

Corollary 1. *Hyperresolution and $>$ -ordered resolution with negative selection decide the positively restrained fragment.*

The next example shows that SGGS can be exponentially more efficient than these saturation-based resolution strategies because it is model-based.

Example 9. Consider the following parametric clause set S_n consisting of $n + 1$ clauses, using $i+1$ -ary predicates P_i and constants c_i , for all i , $0 \leq i \leq n$:

$$P_0(c_0) \vee P_0(c_1) \vee \cdots \vee P_0(c_n) \quad (C_0),$$

$$\neg P_0(x_1) \vee P_1(x_1, c_0) \vee P_1(x_1, c_1) \vee \cdots \vee P_1(x_1, c_n) \quad (C_1),$$

$$\neg P_1(x_1, x_2) \vee P_2(x_1, x_2, c_0) \vee \cdots \vee P_2(x_1, x_2, c_n) \quad (C_2),$$

...

$$\neg P_{n-1}(x_1, \dots, x_n) \vee P_n(x_1, \dots, x_n, c_0) \vee \cdots \vee P_n(x_1, \dots, x_n, c_n) \quad (C_n).$$

The set S_n is positively restrained by an LPO with precedence $P_0 > \dots > P_n > c_i$ for all i , $0 \leq i \leq n$. SGGS with I^- simply selects one positive literal per clause and detects satisfiability after $n + 1$ SGGS-extension steps, producing for instance the model where $P_0(c_0), P_1(c_0, c_0), \dots, P_n(c_0, \dots, c_0)$ are true and all other atoms are false. A saturation by any of the above positive resolution strategies produces exponentially many clauses, because for all i , $0 \leq i \leq n$, all n positive literals in C_i unify with the negative literal in C_{i+1} , generating n^{i+1} positive clauses, so that the clause count is given by $\sum_{k=1}^n n^k$.

6 Experiments

We begin by reducing positive restrainedness (the negative case is similar) to termination of rewrite systems, so that termination tools can yield a partial practical test. Since restrainedness employs a quasi-ordering, we consider unsorted rewrite systems \mathcal{R} and \mathcal{E} , and *rewriting modulo* $\rightarrow_{\mathcal{R}/\mathcal{E}}$ defined by $\leftrightarrow_{\mathcal{E}}^* \circ \rightarrow_{\mathcal{R}} \circ \leftrightarrow_{\mathcal{E}}^*$.

Definition 7. *Given a clause set S , a pair of rewrite systems $(\mathcal{R}_S, \mathcal{E}_S)$ is positively restraining for S if for all clauses $C \in S$ and all non-ground literals $L \in C^+$, there is a rule $at(M) \rightarrow at(L)$ in $\mathcal{R}_S \cup \mathcal{E}_S$ for some literal $M \in C^-$.*

Let $(\mathcal{R}_S, \mathcal{E}_S)$ be positively restraining for S . For instance, for Ex. 8, \mathcal{R}_S will consist of $P(s(s(x)), y) \rightarrow P(x, s(y))$, $P(x, s(s(y))) \rightarrow P(x, s(y))$, and $P(s(x), s(y)) \rightarrow P(s(x), y)$, while $\mathcal{E}_S = \emptyset$. A common situation is that \mathcal{E}_S has permutative rules, as in Ex. 7 where $\text{differ}(x, y) \rightarrow \text{differ}(y, x)$ will be in \mathcal{E}_S .

Lemma 6. (1) *If $\rightarrow_{\mathcal{R}_S}$ is terminating and $\mathcal{E}_S = \emptyset$, clause set S is strictly positively restrained.* (2) *If $\rightarrow_{\mathcal{R}_S/\mathcal{E}_S}$ is terminating, $\text{Var}(t) = \text{Var}(u)$ for all $t \rightarrow u$ in \mathcal{E}_S , and $\leftrightarrow_{\mathcal{E}}^*$ has finite equivalence classes, S is positively restrained.*

Proof. (1) Since $\rightarrow_{\mathcal{R}_S}$ is terminating, for all $t \rightarrow u$ in \mathcal{R}_S , $\text{Var}(u) \subseteq \text{Var}(t)$, so that S is positively ground-preserving. S is strictly positively restrained by the quasi-ordering $\rightarrow_{\mathcal{R}_S}^*$: indeed, $\rightarrow_{\mathcal{R}_S}^*$ is stable, $\rightarrow_{\mathcal{R}_S}^+$ is well-founded, and the equivalence classes of $\rightarrow_{\mathcal{R}_S}^* \cap \mathcal{R}_S^* \leftarrow$, which is identity, are finite. (2) Since $\rightarrow_{\mathcal{R}_S/\mathcal{E}_S}$ is terminating, for all $t \rightarrow u$ in \mathcal{R}_S , $\text{Var}(u) \subseteq \text{Var}(t)$, and for all $t \rightarrow u$ in \mathcal{E}_S , $\text{Var}(u) = \text{Var}(t)$ by hypothesis, so that S is positively ground-preserving. S is positively restrained by $\rightarrow_{\mathcal{R}_S/\mathcal{E}_S}^*$: Indeed, $\rightarrow_{\mathcal{R}_S/\mathcal{E}_S}^*$ is stable, $\rightarrow_{\mathcal{R}_S/\mathcal{E}_S}^+$ is well-founded, and the equivalence classes of $\leftrightarrow_{\mathcal{E}}^*$ are finite by hypothesis. \square

For the experiments, given a clause set S , a script named **StoR** generates rewrite systems \mathcal{R}_S and \mathcal{E}_S that can be fed to a termination tool. The source of clause sets is TPTP 7.2.0 and the termination tool is **T₁T₂** [27]. All problems in the FOF category are transformed into conjunctive normal form, excluding those with equality (whether sets with equality can be restrained is a topic for future work). Besides 1,539 inputs where either **StoR** or **T₁T₂** timed out, out of the remaining 3,462 problems, **T₁T₂** found 313 restrained ones. For those still undetermined, we tested whether it is sufficient to flip the sign of all literals with a certain predicate to get a restrained problem, which succeeded in 36 cases, for

a total of 349 restrained problems. Of these, 277 are positively restrained, 181 negatively restrained, and 109 are both; 74 are ground, 232 are PVD, 277 are stratified, 252 are EPR, 169 are monadic, 204 are FO^2 , 209 are guarded, but 43 problems do not fall in any of these classes, and therefore, to the best of our knowledge, they are proved decidable here for the first time. The average TPTP rating of the 349 problems is 0.1, and that of the 43 problems is 0.015, where 0.1 means that the problem can be solved by most provers. However, the group of 349 includes hard problems such as instances of the binary counter problem in Ex. 5 (MSC015-1.n), and Rubik’s cube problems (e.g., PUZ052-1). For example, MSC015-1.030 is restrained and has rating 1.00, that is, no theorem prover could solve it so far within the timeout allowed in the CASC competition.

Example 10. Problem HWV036-2 (cf. Ex. 1) is a set of axioms with no ground atoms, so that $\mathcal{H}_S^>$ is empty, and the model constructed according to Thm. 4 is trivial. Several other problems combine this set with ground clauses to prove theorems from those axioms. For example, HWV008-2.002 adds 23 ground clauses. As we found a terminating positively restraining rewrite system for HWV008-2.002, this problem, as well as HWV036-2, is strictly restrained.

An SGGS prototype named *Koala* was built reusing code for basic data structures, term indexing, and type inference from *iProver* [24,26]. In *Koala*, the SGGS trail is represented as a list of constrained clauses, with constraints maintained in *standardized form* (see Sect. 2 and [14, Sect. 7]), and selected literals stored in a *discrimination tree*, since SGGS-extension requires to find selected literals that unify with literals in an input clause (cf. Def. 1). *Koala* takes sorts into account when checking satisfiability of constraints (e.g., Ex. 2), and implements a *fair search plan* which ensures that all derivations are fair (see Sect. 2). The *SGGS-suitable ordering* is a KBO with built-in precedence, $w_0 = 1$, and weight 1 for all function and predicate symbols in order to extend the size ordering. *Koala* also sorts by this ordering the clauses in a splitting, according to the SGGS notion of *preferred clause* in a splitting [14, Def. 22]).

Koala picks either I^- or I^+ as initial interpretation, based on whether the problem is positively or negatively ground-preserving, which overapproximates positively or negatively restrained. *Koala* implements the above mentioned sign-flipping test to obtain more ground-preserving sets. In order to recognize stratified input problems, one can compute the sort dependency graph and check for acyclicity [25], or let *Koala* apply type inference. For all experiments with *Koala*, the time-out was 300 sec of wall clock time. Out of the above 349 restrained problems, *Koala* shows that 50 are satisfiable and 283 unsatisfiable. Of 351 PVD problems, not including the 232 ones in the set of 349 restrained problems, *Koala* discovers 232 unsatisfiable and 76 satisfiable instances. Of the 1,246 stratified problems found in the FOF category by the above acyclicity test, *Koala* solves 643 unsatisfiable and 277 satisfiable instances. The list of the 349 restrained sets with their properties and restraining rewrite systems, as well as detailed results of the experiments with *Koala* are available.²

² <http://profs.scienze.univr.it/winkler/sggsgdp/>

7 Discussion

SGGS [13,14] is an attractive candidate for decision procedures, because it is conflict-driven and it builds models. In this paper, we showed that if the generated clauses are in a *finite basis*, SGGS is guaranteed to terminate, and we instantiated this result to yield SGGS decision procedures for the *stratified fragment*, *PVD*, and the newly introduced *restrained fragment*, all without equality.

While also Inst-Gen [24] and the model evolution calculus (MEC) [7] decide the stratified fragment [25], this is not the case for the restrained fragment: for instance, if Inst-Gen uses an unfortunate literal selection in Ex. 8, it does not halt. SGGS avoids this phenomenon thanks to semantic guidance. Since MEC starts with I^+ as candidate model, it may not terminate on satisfiable negatively restrained sets such as Ex. 7. Indeed, E-Darwin [6] does not halt on this example that Koala solves in a few seconds.

However, it is generally difficult to tame a refutationally complete first-order inference system to yield decision procedures. For instance, SGGS does not halt given the following set of clauses that belongs to several decidable fragments.³

Example 11. The following set is in the Ackermann, monadic, and FO^2 classes:

$$P(0) \quad (i) \quad P(x) \vee P(f(x)) \quad (ii) \quad \neg P(x) \vee \neg P(f(x)) \quad (iii),$$

where membership in the Ackermann class stems from the Skolemization of $\exists v \forall x \exists y. P(v) \wedge (P(x) \vee P(y)) \wedge (\neg P(x) \vee \neg P(y))$. It has the finite model property, as witnessed by model \mathcal{I} with domain $\{0, 1\}$, $f^{\mathcal{I}}(x) = 1 - x$, $0^{\mathcal{I}} = 0$, and $P^{\mathcal{I}} = \{0\}$. With I^- , SGGS performs an infinite series of non-conflicting extensions:

$$\begin{array}{ll} \varepsilon \vdash [P(0)] & \text{extend (i)} \\ \vdash [P(0)], \neg P(0) \vee [\neg P(f(0))] & \text{extend (iii)} \\ \vdash [P(0)], \neg P(0) \vee [\neg P(f(0))], P(f(0)) \vee [P(f(f(0)))] & \text{extend (ii)} \\ \vdash \dots & \end{array}$$

It can be shown that SGGS does not terminate with I^+ either. Inst-Gen may terminate if appropriate candidate models are constructed, but not in general. Ordered resolution terminates if the ordering satisfies $\neg P(f(x)) > \neg P(x)$.

The stratified fragment was presented as the first of three fragments of increasing expressivity, named St_0 , St_1 , and St_2 , and suitable to capture Alloy specifications [1]. If the generic predicate symbols P of Ex. 4 is replaced with equality for sort s_2 , the second sentence of Ex. 4 becomes $\forall x \exists y. f(y) \simeq x$, which states that x is in the image of f : this formula separates St_0 and St_1 , since it is allowed in St_1 , but not in St_0 . The LISBQ logic for the verification of unbounded data structures [29] can be translated into St_0 with equality [1].

The *restrained fragment* is defined starting from a notion of *ground-preservingness*, which is convenient for positive or negative strategies. Negative ground-preservingness was used for Horn theories with equality [28] and [10, Sect. 5.2].

³ Renate A. Schmidt and Marco Voigt suggested this example to the first author.

Positive ground-preservingness, also named *range restrictedness* [15], is implicit in PVD and recent extensions [30]. Unlike these settings, the restrained fragment does not limit literal depth. Positive ground-preservingness was used to show that $\text{DPLL}(\Gamma+\mathcal{T})$, where Γ is an inference system including hyperresolution, superposition with negative selection, and simplification, decides *essentially finite* (only one monadic function f with finite range) and positively ground-preserving axiomatizations, provided that *speculative axioms* $f^j(x) \simeq f^k(x)$ ($j > k$) are tried for increasing values of j and k ([12], Thm. 7 and Lem. 3). Simplification applies the speculative axioms to limit the depth of generated terms. Without this feature, it is not surprising that SGGS does not halt.

Example 12. Consider a positively ground-preserving variant of Ex. 11:

$$\text{P}(0) \quad (i) \quad \neg\text{P}(x) \vee \text{P}(f(f(x))) \quad (ii) \quad \neg\text{P}(x) \vee \neg\text{P}(f(x)) \quad (iii)$$

The finite model property implies that it is essentially finite. SGGS terminates with neither I^- nor I^+ as initial interpretation. With I^- it generates:

$$\begin{array}{ll} \varepsilon \vdash [\text{P}(0)] & \text{extend (i)} \\ \vdash [\text{P}(0)], \neg\text{P}(0) \vee [\text{P}(f^2(0))] & \text{extend (ii)} \\ \vdash [\text{P}(0)], \neg\text{P}(0) \vee [\text{P}(f^2(0))], \neg\text{P}(f^2(0)) \vee [\text{P}(f^4(0))] & \text{extend (ii)} \\ \vdash \dots & \end{array}$$

Positive hyperresolution generates the series $\{\text{P}(f^{2k}(0))\}_{k \geq 0}$, which is essentially the same behavior as SGGS. $\text{DPLL}(\Gamma+\mathcal{T})$ tries $f(x) \simeq x$, detects a conflict, backtracks, tries $f^2(x) \simeq x$, and halts reporting satisfiability.

Ensuring termination by restricting new (i.e., non-input) terms to come from a finite basis is common for conflict-driven decision procedures [16,11]. A major direction for future work towards decision procedures for richer languages is the integration of SGGS with CDSAT [11], in order to endow SGGS with equality and CDSAT with quantifiers. Speculative inferences and initial interpretations not based on sign are additional leads. For the restrained fragment, one may consider its relations with other decidable fragments and its relevance to applications beyond TPTP. While positive clauses in a restrained set are ground, one may study the decidability of sets where positive clauses are not necessarily ground, but admit a *restraining* rewrite system (Def. 7) such that *narrowing* halts. Techniques to detect the termination of narrowing are known [35].

Although Koala is only a prototype, the experiments show potential and allow us to identify critical issues for the performance of an SGGS prover. For example, instance generation by SGGS-extension is a bottleneck for problems with many input clauses, and forms of *caching* should be considered to avoid repeating computations. Further development of Koala and more experiments may contribute to discover classes of first-order problems where the conflict-driven style of SGGS reasoning is especially rewarding.

References

1. A. Abadi, A. Rabinovich, and M. Sagiv. Decidable fragments of many-sorted logic. *J. Symb. Comput.*, 45(2):153–172, 2010.
2. W. Ackermann. *Solvable Cases of the Decision Problem*. North Holland, 1954.
3. G. Alagi and C. Weidenbach. NRCL – a model building approach to the Bernays-Schönfinkel fragment. In *Proc. FroCoS-10*, volume 9322 of *LNAI*, pages 69–84. Springer, 2015.
4. H. Andréka, J. van Benthem, and I. Nemeti. Modal logics and bounded fragments of predicate logic. *J. Phil. Logic*, 27(3):217–274, 1998.
5. L. Bachmair, H. Ganzinger, and U. Waldmann. Superposition with simplification as a decision procedure for the monadic class with equality. In *Proc. Computational Logic and Proof Theory (KGC 1993)*, volume 713 of *LNCS*, pages 83–96, 1993.
6. P. Baumgartner, A. Fuchs, and C. Tinelli. Implementing the model evolution calculus. *Int. J. Artif. Intell. T.*, 15(1):21–52, 2006.
7. P. Baumgartner and C. Tinelli. The model evolution calculus as a first-order DPLL method. *Artificial Intelligence*, 172(4-5):591–632, 2008.
8. P. Bernays and M. Schönfinkel. Zum Entscheidungsproblem der mathematischen Logik. *Math. Annalen*, 99:342–372, 1928.
9. P. Blackburn, M. Rijke, and Y. Venema. *Modal Logic*. Cambridge Univ. Press, 2001.
10. M. P. Bonacina and N. Dershowitz. Canonical ground Horn theories. In *Programming Logics: Essays in Memory of H. Ganzinger*, volume 7797 of *LNAI*, pages 35–71. Springer, 2013.
11. M. P. Bonacina, S. Graham-Lengrand, and N. Shankar. Conflict-driven satisfiability for theory combination: transition system and completeness. *J. Automat. Reason.*, 64(3):579–609, 2020.
12. M. P. Bonacina, C. A. Lynch, and L. de Moura. On deciding satisfiability by theorem proving with speculative inferences. *J. Automat. Reason.*, 47(2):161–189, 2011.
13. M. P. Bonacina and D. A. Plaisted. Semantically-guided goal-sensitive reasoning: model representation. *J. Automat. Reason.*, 56(2):113–141, 2016.
14. M. P. Bonacina and D. A. Plaisted. Semantically-guided goal-sensitive reasoning: inference system and completeness. *J. Automat. Reason.*, 59(2):165–218, 2017.
15. R. Caferra, A. Leitsch, and N. Peltier. *Automated Model Building*. Kluwer, 2004.
16. L. de Moura and D. Jovanović. A model-constructing satisfiability calculus. In *Proc. VMCAI-14*, volume 7737 of *LNCS*, pages 1–12. Springer, 2013.
17. C. Fermüller and A. Leitsch. Model building by resolution. In *Proc. CSL-6*, volume 702 of *LNCS*, pages 134–148, 1993.
18. C. Fermüller, A. Leitsch, U. Hustadt, and T. Tammet. Resolution decision procedures. In *Handbook of Automated Reasoning*, pages 1791–1849. Elsevier and MIT Press, 2001.
19. A. Fiori and C. Weidenbach. SCL clause learning from simple models. In *Proc. CADE-27*, volume 11716 of *LNAI*, pages 233–249, 2019.
20. H. Ganzinger and H. de Nivelle. A superposition decision procedure for the guarded fragment with equality. In *Proc. LICS-14*. IEEE Computer Society Press, 1999.
21. E. Grädel, P. Kolaitis, and M. Vardi. On the decision problem for two-variable first-order logic. *Bull. Symb. Logic*, 3:53–69, 1997.
22. U. Hustadt, R. A. Schmidt, and L. Georgieva. A survey of decidable first-order fragments and description logics. *J. Rel. Meth. in Comp. Sci.*, 1:251–276, 2004.

23. W. H. Joyner Jr. Resolution strategies as decision procedures. *J. ACM*, 23(3):398–417, 1976.
24. K. Korovin. Inst-Gen: a modular approach to instantiation-based automated reasoning. In *Programming Logics: Essays in Memory of H. Ganzinger*, volume 7797 of *LNAI*, pages 239–270. Springer, 2013.
25. K. Korovin. Non-cyclic sorts for first-order satisfiability. In *Proc. FroCoS-9*, volume 8152 of *LNAI*, pages 214–228, 2013.
26. K. Korovin. iProver - an instantiation-based theorem prover for first-order logic (v2.8). <http://www.cs.man.ac.uk/~korovink/iprover/index.html>, 2018.
27. M. Korp, C. Sternagel, H. Zankl, and A. Middeldorp. Tyrolean Termination Tool 2. In *Proc. RTA-20*, volume 5595 of *LNCS*, pages 295–304, 2009.
28. E. Kounalis and M. Rusinowitch. On word problems in Horn theories. *J. Symb. Comput.*, 11(1–2):113–128, 1991.
29. S. Lahiri and S. Qadeer. Back to the future: revisiting precise program verification using SMT solvers. In *Proc. POPL-31*, pages 171–182. ACM, 2004.
30. M. Lamotte-Schubert and C. Weidenbach. BDI: a new decidable clause class. *J. Logic and Comput.*, 27(2):441–468, 2017.
31. C. I. Lewis and C. H. Langford. *Symbolic Logic*. Dover Publications Inc., 1932.
32. C. Lutz, U. Sattler, and F. Wolter. Modal logic and the two-variable fragment. In *Proc. CSL-15*, volume 2142 of *LNCS*, pages 247–261, 2001.
33. K. McMillan. Developing distributed protocols with Ivy. Slides from <http://vmcaischool19.tecnico.ulisboa.pt/>, 2019.
34. J. Navarro and A. Voronkov. Proof systems for effectively propositional logic. In *Proc. IJCAR-4*, volume 5195 of *LNAI*, pages 426–440, 2008.
35. N. Nishida and G. Vidal. Termination of narrowing via termination of rewriting. *Appl. Algebr. Eng. Comm.*, 21(3):177–225, 2010.
36. O. Padon, K. McMillan, A. Panda, M. Sagiv, and S. Shoham. Ivy: Safety verification by interactive generalization. *SIGPLAN Not.*, 51(6):614–630, 2016.
37. R. Piskac, L. de Moura, and N. Bjørner. Deciding effectively propositional logic using DPLL and substitution sets. *J. Automat. Reason.*, 44(4):401–424, 2010.
38. D. A. Plaisted and Y. Zhu. *The Efficiency of Theorem Proving Strategies*. Friedr. Vieweg & Sohns, 1997.
39. F. Ramsey. On a problem in formal logic. *Proc. London Math. Soc.*, 30:264–286, 1930.
40. A. Rubio. A fully syntactic AC-RPO. *Inform. Comput.*, 178(2):515–533, 2002.
41. G. Sutcliffe. The TPTP Problem Library and Associated Infrastructure. From CNF to TH0, TPTP v6.4.0. *J. Automat. Reason.*, 59(4):483–502, 2017.