

Detecting and breaking symmetries on specifications

Marco Cadoli and Toni Mancini

Dipartimento di Informatica e Sistemistica
Università di Roma "La Sapienza"
Via Salaria 113, I-00198 Roma, Italy
cadoli|tmancini@dis.uniroma1.it

Abstract. In this paper we address symmetry on combinatorial problems by following the approach of imposing additional symmetry-breaking constraints. Differently from other works in the literature, we attack the problem at the *specification* level. In fact, sometimes symmetries on specifications can be very easily detected, and symmetry-breaking formulae generated. We give formal definitions of symmetries and symmetry-breaking formulae on specifications written in existential second-order logic, clarifying the new definitions on some specifications: the graph 3-coloring, the not-all-equal (3-)Sat and the social golfer problems. Finally, we show the results of a preliminary experimentation of our techniques on some of the examples using state-of-the-art linear and constraint programming solvers.

1 Introduction

In order to speed up the solving process of combinatorial and constraint satisfaction problems, much work has been done on detecting and breaking symmetries, with the aim of greatly reducing the size of the search space. In the literature, three approaches have been followed to deal with symmetries:

1. Imposing additional constraints to the model of the problem, which are satisfied only for one of the symmetrical points in the search space, cf., e.g., [15–17];
2. Introducing additional constraints during the search process, to avoid the traversal of symmetrical points, cf., e.g., [2, 11, 8];
3. Defining a search strategy able to break symmetries as soon as possible (e.g., by first selecting variables involved in the greatest number of local symmetries), cf., e.g., [13].

In this paper we follow the first approach, but, differently from other works in the literature, we aim to attack the problem at the *specification* level. Here are our motivations:

- Many systems and languages for the solution of constraint problems (e.g., AMPL [9], OPL [18], GAMS [4], DLV [6], SMODELS [14], and NP-SPEC [3]) clearly

separate the *specification* of a problem, e.g., graph three-coloring, and its *instance*, e.g., a graph, using a two-level architecture for finding solutions: the specification is *instantiated* against the instance, and then an appropriate solver is invoked.

- In some cases, symmetries on specifications can be very easily detected by humans (cf. Section 3), and convenient symmetry breaking formulae can be added to the specification itself. Since specifications are logical formulae, in principle, a theorem prover can be used for detecting and breaking symmetries in complex cases.
- Detecting and breaking symmetries at the specification level does not rule out the possibility to use symmetry-breaking techniques at the instance level. As an example, since some systems generate a SAT instance, e.g., [3], or an instance of integer linear programming, e.g., [18], it is possible to do symmetry breaking on such instances, using existing techniques, cf. e.g., [5].

In this work we show some preliminary results of our research. The outline of the paper is as follows: in Section 2, after recalling basic definitions of CSPs and symmetries, we lift to the specification level, defining the concepts of symmetry on a specification and of symmetry-breaking formula. In Section 3 we clarify the new definitions by showing some specifications: the graph 3-coloring, the not-all-equal 3-Sat and Sat, the social golfer, and the BIBD generation problems. For all the specifications we detect and break symmetries. Finally, in Section 4 we show the results of a preliminary experimentation of our symmetry-breaking techniques on the first example of Section 3, using state-of-the-art linear and constraint programming solvers CPLEX and SOLVER, invoked through OPLSTUDIO.

2 Basic definitions

In this section we briefly recall the standard definitions of Constraint Satisfaction Problem (CSP) and symmetry. Afterwards, we lift up to the definition of problem specification, which is independent on a particular instance, and give the definition of symmetry in the new context.

2.1 CSPs and their symmetries

Definition 1 (Constraint satisfaction problem (CSP)). *A CSP is a tuple $\langle \mathbf{V}, \mathbf{D}, \mathbf{C} \rangle$, where $\mathbf{V} = \{X_1, \dots, X_n\}$ is a set of variables, $\mathbf{D} = \{D_1, \dots, D_n\}$ is a set of finite domains, one for each variable, $\mathbf{C} = \{C_1, \dots, C_m\}$ is a set of constraints on \mathbf{V} and \mathbf{D} , i.e., a set of relations $C_i(X_{i_1}, \dots, X_{i_k}) \subseteq D_{i_1} \times \dots \times D_{i_k}$, for every $i \in [1..m]$, containing those configurations of assignments allowed by the constraint.*

Given a CSP π , we can transform it in a new CSP π' by exchanging variables and/or domains order. A CSP transformation is defined in the following way:

Definition 2 (CSP transformation). Given a CSP $\pi = \langle \mathbf{V}, \mathbf{D}, \mathbf{C} \rangle$, with $\mathbf{V} = \{X_1, \dots, X_n\}$ and $\mathbf{D} = \{D_1, \dots, D_n\}$, a transformation of π is a set of bijections $\sigma_0, \sigma_1, \dots, \sigma_n$ such that:

$$\sigma_0 : \mathbf{V} \rightarrow \mathbf{V}$$

for every $i \in [1, n]$, $\sigma_i : D_i \rightarrow D_{\text{index}(\sigma_0(X_i))}$,

where $\text{index}()$ is a function that returns the index of the argument variable. It has to be observed that the transformation is defined on \mathbf{V} and \mathbf{D} only (not on \mathbf{C}).

Here are interesting specializations of Definition 2 (which apply also to forthcoming Definition 3):

Variable transformation: $\sigma_1, \dots, \sigma_n$ are the identity function;

Value transformation: σ_0 is the identity function;

Uniform value transformation: σ_0 is the identity function, $D_1 = D_2 = \dots = D_n$, and $\sigma_1 = \sigma_2 = \dots = \sigma_n$.

Intuitively, a variable transformation exchanges only the order of variables, leaving the domains unchanged. On the other hand, a value transformation does not modify the variable order. In what follows, we focus mainly on uniform value transformations, in which domain values for each variable are swapped *uniformly*.

As widely described in the literature, symmetries of a CSP are transformations that map solutions into solutions:

Definition 3 (CSP symmetry [12]). Given a CSP $\pi = \langle \mathbf{V}, \mathbf{D}, \mathbf{C} \rangle$, a symmetry on π is a CSP transformation on $\langle \mathbf{V}, \mathbf{D} \rangle$ which preserves constraints: an assignment $\tau = \{X_1 = v_1, \dots, X_n = v_n\}$ to variables in \mathbf{V} satisfies every constraint in \mathbf{C} iff $\sigma(\tau) = \{\sigma_0(X_1) = \sigma_{\text{index}(\sigma_0(X_1))}(v_1), \dots, \sigma_0(X_n) = \sigma_{\text{index}(\sigma_0(X_n))}(v_n)\}$ does.

2.2 Specifications and their symmetries

Several languages for specifying problems exist. For the sake of simplicity, in this paper we focus on the most basic one, the existential fragment of second-order logic (ESO). ESO is able to specify all problems in the complexity class NP [7], where the instance is represented as a relational database. An ESO specification is a formula

$$\exists \mathbf{S} \phi(\mathbf{S}, \mathbf{R}), \tag{1}$$

where \mathbf{R} is the relational schema for every input instance, and ϕ is a quantified first-order formula on the relational vocabulary $\mathbf{S} \cup \mathbf{R} \cup \{=\}$ (“=” is always interpreted as identity). Predicates in \mathbf{S} are called *guessed*.

Example 1 (Graph 3-Coloring).

Input: A graph, according to the input relational schema $\mathbf{R} = \{\text{edge}(\cdot, \cdot)\}$.

Question: Is there a way to assign each node of the input graph one out of three colors such that every edge links nodes with different colors?

An ESO specification for the problem is as follows:

$$\begin{aligned}
 & \exists RGB \\
 & \forall X \ R(X) \vee G(X) \vee B(X) \ \wedge & (2) \\
 & \forall X \ R(X) \rightarrow \neg G(X) \ \wedge & (3) \\
 & \forall X \ R(X) \rightarrow \neg B(X) \ \wedge & (4) \\
 & \forall X \ B(X) \rightarrow \neg G(X) \ \wedge & (5) \\
 & \forall XY \ X \neq Y \wedge R(X) \wedge R(Y) \rightarrow \neg edge(X, Y) \ \wedge & (6) \\
 & \forall XY \ X \neq Y \wedge G(X) \wedge G(Y) \rightarrow \neg edge(X, Y) \ \wedge & (7) \\
 & \forall XY \ X \neq Y \wedge B(X) \wedge B(Y) \rightarrow \neg edge(X, Y). & (8)
 \end{aligned}$$

Clauses (2–8) denote the first order part ϕ of the specification. Clause (2) imposes that every node is assigned at least one color; clauses (3–5) impose that every node is assigned at most one color; clauses (6–8) impose every edge to link nodes with different colors. \square

It is worthwhile to note that, when a specification is instantiated against an input database, a CSP, in the sense of Definition 1, is obtained.

In what follows we focus initially on problem specifications which, like the one of Example 1, have only *monadic* guessed predicates. In this way we have a neat conceptual correspondence between the guessed predicates and the *values* of a CSP. In Example 6 we show how our definitions can be used for non-monadic predicates, essentially by unfolding non-unary predicates and exploiting finiteness of the input database.

In this new context we can give the following definition:

Definition 4 (Uniform value transformation (UVT) of a specification). Given a problem specification $\psi \doteq \exists \mathbf{S} \ \phi(\mathbf{S}, \mathbf{R})$, with $\mathbf{S} = \{S_1, \dots, S_n\}$, S_i monadic for every i , and input schema \mathbf{R} , a uniform value transformation (UVT) for ψ is a mapping $\sigma : \mathbf{S} \rightarrow \mathbf{S}$, which is total and onto, i.e., defines a permutation of guessed predicates in \mathbf{S} .

The term “uniform value transformation” in Definition 4 is used because swapping monadic guessed predicates is conceptually the same as uniformly exchanging domain values in a CSP.

From here on, given ϕ and σ as in the above definition, ϕ^σ is defined as $\phi[S_1/\sigma(S_1), \dots, S_n/\sigma(S_n)]$, i.e., ϕ^σ is obtained from ϕ by uniformly substituting every occurrence of each guessed predicate with the one given by the transformation σ . Analogously, ψ^σ is defined as $\exists \mathbf{S} \ \phi^\sigma(\mathbf{S}, \mathbf{R})$.

We now define when a UVT is a symmetry for a given specification.

Definition 5 (Uniform value symmetry (UVS) of a specification). Let $\psi \doteq \exists \mathbf{S} \phi(\mathbf{S}, \mathbf{R})$, be a specification, with $\mathbf{S} = \{S_1, \dots, S_n\}$, S_i monadic for every $i \in [1, n]$, and input schema \mathbf{R} , and let σ be a UVT for ψ . Transformation σ is a uniform value symmetry (UVS) for ψ if the following holds:

$$\forall \mathbf{S} \forall \mathcal{I} \phi(\mathbf{S}, \mathcal{I}) \leftrightarrow \phi^\sigma(\mathbf{S}, \mathcal{I}). \quad (9)$$

Intuitively, formula (9) says that σ is a UVS for ψ if every extension for \mathbf{S} which satisfies ϕ , satisfies also ϕ^σ and vice versa, for every input instance \mathcal{I} , i.e., for every extension of the input schema \mathbf{R} . Note that every CSP obtained by instantiating a specification with σ has at least the corresponding uniform value symmetry. The above definition could in principle be stated with a different pattern of quantifiers, as an example, we could require equivalence only for some instances, instead that for each instance. Generalization of the definition is left for future research.

The following proposition shows that checking whether a UVT is a UVS reduces to checking equivalence of two first-order formulae. As a consequence this task can in principle be performed by a first-order theorem prover.

Proposition 1. Let ψ be a problem specification of the kind (1), with only monadic guessed predicates, and σ a UVT for ψ . Transformation σ is a UVS for ψ if and only if $\phi \equiv \phi^\sigma$.

Proof.

If part. If $\phi \equiv \phi^\sigma$, then formula (9) holds for σ , so σ is a UVS for ψ .

Only if part. If σ is a UVS for ψ , then formula (9) holds. This implies that every model of ϕ must be a model of ϕ^σ and vice versa, therefore $\phi \equiv \phi^\sigma$ must hold.

The following proposition shows that also the converse reduction can be proven, thus implying that checking whether a UVT is a UVS is not decidable.

Proposition 2 (Undecidability of the UVS checking problem). Given a specification ψ and a UVT σ for it, the problem of checking whether σ is a UVS for ψ is undecidable.

Proof. We prove the statement by reducing it to the problem of checking whether an arbitrary closed first-order formula is a tautology.

Let $\psi \doteq \exists \mathbf{S} \phi(\mathbf{S}, \mathbf{R})$ be any fixed specification on input schema \mathbf{R} , and let σ be any fixed UVT which is not a UVS for it, i.e., from Proposition 1:

$$\phi \not\equiv \phi^\sigma. \quad (10)$$

Let now $\gamma(\mathbf{R})$ be a given arbitrary closed first-order formula on the relational vocabulary \mathbf{R} , and consider the following new specification ψ' :

$$\exists \mathbf{S} \phi(\mathbf{S}, \mathbf{R}) \vee \gamma(\mathbf{R}).$$

The UVT σ is also a UVT for ψ' , since ψ' has the same set of guessed predicates as ψ . By Proposition 1, the problem of deciding whether σ is a UVS for ψ' reduces to checking whether:

$$\phi(\mathbf{S}, \mathbf{R}) \vee \gamma(\mathbf{R}) \equiv \phi^\sigma(\mathbf{S}, \mathbf{R}) \vee \gamma^\sigma(\mathbf{R}). \quad (11)$$

Since (10) and $\gamma^\sigma \equiv \gamma$ (γ has no occurrences of predicates in \mathbf{S}), if (11) holds then γ must be a tautology; on the other hand, if γ is a tautology, then (11) trivially holds. So, an arbitrary given first-order closed formula γ is a tautology if and only if a fixed UVT σ is a UVS for the specification $\exists \mathbf{S} \phi(\mathbf{S}, \mathbf{R}) \vee \gamma(\mathbf{R})$. Since the former problem is undecidable [1], the latter is undecidable as well.

Example 2 (Graph 3-Coloring: Example 1, continued). Three UVTs are as follows:

- $\sigma^{R,G}$ s.t. $\sigma^{R,G}(R) = G$, $\sigma^{R,G}(G) = R$, and $\sigma^{R,G}(B) = B$;
- $\sigma^{R,B}$ s.t. $\sigma^{R,B}(R) = B$, $\sigma^{R,B}(G) = G$, and $\sigma^{R,B}(B) = R$;
- $\sigma^{G,B}$ s.t. $\sigma^{G,B}(R) = R$, $\sigma^{G,B}(G) = B$ and $\sigma^{G,B}(B) = G$.

It is easy to observe that, for each of the above transformations, formulae $\phi^{\sigma^{R,G}}$, $\phi^{\sigma^{R,B}}$, and $\phi^{\sigma^{G,B}}$ are all equivalent to ϕ , because clauses of the former ones are syntactically equivalent to clauses of ϕ and vice versa. This implies, by Proposition 1, that they are all UVSs. \square

Example 3 (Graph 3-Coloring with red self-loops). To play the devil's advocate, we consider a modification of the problem of Example 1, and show that only one of the UVTs in Example 2 is indeed a UVS in the new problem. The problem is obtained by adding one more constraint.

Question: Is there a way to assign each node of the input graph one out of three colors such that: (i) every edge links nodes with different colors, and (ii) every self loop insists on a red node?

In ESO, one more clause (which forces the nodes with self loops to be colored in red) must be added.

$$\begin{aligned} & \exists RGB \\ & \dots \\ & \forall X \text{ edge}(X, X) \rightarrow R(X). \end{aligned} \quad (12)$$

- $\sigma^{G,B}$: it is a UVS, because the argument of previous example applies.
- $\sigma^{R,G}$: in this case, $\phi^{\sigma^{R,G}}$ is not equivalent to ϕ : as an example, for the input instance $\text{edge} = \{(v, v)\}$, the color assignment $\bar{R}, \bar{G}, \bar{B}$ such that $\bar{R} = \{v\}, \bar{G} = \bar{B} = \emptyset$ is a model for the original problem, i.e., $\bar{R}, \bar{G}, \bar{B} \models \phi(R, G, B, \text{edge})$. It is however easy to observe that $\bar{R}, \bar{G}, \bar{B} \not\models \phi^{\sigma^{R,G}}(R, G, B, \text{edge})$, because $\phi^{\sigma^{R,G}}$ is verified only by color assignments for which $\bar{G}(v)$ holds. This implies, by Proposition 1, that σ is not a UVS.

- $\sigma^{R,B}$: it is not a UVS, because the same argument of the previous point applies. \square

Once symmetries of a specification have been detected, additional constraints can be added in order to *break* them, i.e., to wipe out from the solution space (some of) the symmetrical points. These kind of constraints are called *symmetry-breaking formulae*, and are defined as follows.

Definition 6 (Symmetry-breaking formula). *Let $\psi \doteq \exists \mathbf{S} \phi(\mathbf{S}, \mathbf{R})$, be a specification, with $\mathbf{S} = \{S_1, \dots, S_n\}$, S_i monadic for every $i \in [1, n]$, and input schema \mathbf{R} , and let σ be a UVS for ψ . A symmetry-breaking formula for ψ with respect to symmetry σ is a closed (except for \mathbf{S}) formula $\beta(\mathbf{S})$ such that the following two conditions hold:*

1. Transformation σ is no longer a symmetry for $\phi \wedge \beta$:

$$(\phi \wedge \beta(\mathbf{S})) \not\equiv (\phi \wedge \beta(\mathbf{S}))^\sigma;$$

2. Every model of ϕ can be obtained by those of $\phi \wedge \beta(\mathbf{S})$ and their symmetric with respect to σ :

$$\phi \equiv (\phi \wedge \beta(\mathbf{S})) \vee (\phi \wedge \beta(\mathbf{S}))^\sigma.$$

If $\beta(\mathbf{S})$ matches the above definition, then we are entitled to solve the problem $\exists \mathbf{S} \phi(\mathbf{S}, \mathbf{R}) \wedge \beta(\mathbf{S})$ instead of the original one $\exists \mathbf{S} \phi(\mathbf{S}, \mathbf{R})$. In fact, point 1 in the above definition states that formula $\beta(\mathbf{S})$ actually breaks σ , since, by Proposition 1, σ is not a symmetry of the rewritten problem. Point 2 guarantees that all solutions are preserved in the rewritten problem, up to symmetric ones, which can be readily obtained by applying σ .

It is important to observe that breaking a symmetry is sound, i.e., it preserves at least one solution, as shown by the following proposition:

Proposition 3 (Symmetry-breaking formulae preserve satisfiability).

Let $\psi \doteq \exists \mathbf{S} \phi(\mathbf{S}, \mathbf{R})$ be a problem specification with $\mathbf{S} = \{S_1, \dots, S_n\}$, S_i monadic for every $i \in [1, n]$, and input schema \mathbf{R} , and let σ be a UVS for ψ . Furthermore, let $\beta(\mathbf{S})$ be a symmetry-breaking formula for ψ with respect to σ . For each input instance \mathcal{I} , i.e., for each extension of the input schema \mathbf{R} , if $\exists \mathbf{S} \phi(\mathbf{S}, \mathcal{I})$ has solutions, then also $\exists \mathbf{S} \phi(\mathbf{S}, \mathcal{I}) \wedge \beta(\mathbf{S})$ has solutions.

Proof. First of all, we observe that for each input instance \mathcal{I} , if problem $\exists \mathbf{S} \phi(\mathbf{S}, \mathcal{I})$ has solutions, then, for every UVT σ , also $\exists \mathbf{S} (\phi(\mathbf{S}, \mathcal{I}))^\sigma$ has solutions. In fact, supposing $\overline{\mathbf{S}} = \{\overline{S}_1, \dots, \overline{S}_n\}$ a solution of the former, the extension $\sigma(\overline{\mathbf{S}}) = \{\sigma(\overline{S}_1), \dots, \sigma(\overline{S}_n)\}$ is a solution of the latter.

Now, suppose that, for a given input instance \mathcal{I} , $\exists \mathbf{S} \phi(\mathbf{S}, \mathcal{I})$ has solutions; this implies that $\phi(\mathbf{S}, \mathcal{I})$ is satisfiable. From point 2 of Definition 6 it follows that also $(\phi(\mathbf{S}, \mathcal{I}) \wedge \beta(\mathbf{S})) \vee (\phi(\mathbf{S}, \mathcal{I}) \wedge \beta(\mathbf{S}))^\sigma$ is satisfiable, and, since σ is a UVT for the problem $\exists \mathbf{S} \phi(\mathbf{S}, \mathcal{I}) \wedge \beta(\mathbf{S})$, from the observation above, it follows that both $\phi(\mathbf{S}, \mathcal{I}) \wedge \beta(\mathbf{S})$ and $(\phi(\mathbf{S}, \mathcal{I}) \wedge \beta(\mathbf{S}))^\sigma$ are satisfiable. This implies that the rewritten problem $\exists \mathbf{S} \phi(\mathbf{S}, \mathcal{I}) \wedge \beta(\mathbf{S})$ has solutions.

Intuitively, the “quality” of a symmetry-breaking formula is higher when $(\phi \wedge \beta(\mathbf{S}))$ and $(\phi \wedge \beta(\mathbf{S}))^\sigma$ have few common models, i.e., the less models in the conjunction $(\phi \wedge \beta(\mathbf{S})) \wedge (\phi \wedge \beta(\mathbf{S}))^\sigma$, the better. In next section we see several examples of UVs and different breaking formulas for them.

The above definition deals with breaking a single symmetry, and it can be generalized with breaking *multiple* symmetries simultaneously. The definition, omitted for lack of space, deals with a set $\Sigma = \{\sigma_1, \dots, \sigma_n\}$ of UVs. The first condition must be repeated for all symmetries in Σ , and the the second one becomes $\phi \equiv (\phi \wedge \beta(\mathbf{S})) \vee (\phi \wedge \beta(\mathbf{S}))^{\sigma_1} \vee \dots \vee (\phi \wedge \beta(\mathbf{S}))^{\sigma_n}$.

3 Examples of symmetry-breaking formulas

In this section we discuss some examples of well-known combinatorial problems which have many symmetries, show some of them by using Proposition 1, and present breaking formulae respecting conditions of Definition 6.

Example 4 (Graph 3-Coloring: Example 2, continued).

In the following we present different symmetry-breaking formulae for this problem, referring to $\sigma^{R,G}$, $\sigma^{R,B}$, and $\sigma^{G,B}$.

1. Let us consider $\sigma^{G,B}$: a symmetry-breaking formula is:

$$\beta^{G,B}(R, G, B) \doteq |G| \leq |B|, \quad (13)$$

i.e., the number of green nodes is less than or equal to the number of blue ones (this constraint can be written in ESO using standard techniques).

Formula (13) matches Definition 6. In fact we have that:

- Symmetry $\sigma^{G,B}$ is broken, since:

$$(\phi \wedge \beta^{G,B}(R, G, B)) \not\equiv (\phi \wedge \beta^{G,B}(R, G, B))^{\sigma^{G,B}};$$

- All solutions –up to symmetric ones– are preserved, since:

$$\phi \equiv (\phi \wedge \beta^{G,B}(R, G, B)) \vee (\phi \wedge \beta^{G,B}(R, G, B))^{\sigma^{G,B}}.$$

We note that $|G| < |B|$ is not a symmetry-breaking formula, since point 2 of Definition 6 is not satisfied: in fact adding such a constraint is obviously unsound, since some satisfiable instances may become unsatisfiable.

2. Let us now consider the two symmetries $\sigma^{R,G}$ and $\sigma^{R,B}$. A multiple symmetry-breaking formula for them is the following one:

$$\beta^R(R, G, B) \doteq |R| \leq |G| \wedge |R| \leq |B|, \quad (14)$$

which uses a partial order among the sets of nodes with the same color, R being the minimal element. For lack of space, we don't formally prove here that β^R is a multiple symmetry-breaking formula.

3. The next formula breaks all three symmetries:

$$\beta(R, G, B) \doteq |R| \leq |G| \wedge |R| \leq |B| \wedge |G| \leq |B|, \quad (15)$$

which uses a total order among colors, i.e., $|R| \leq |G| \leq |B|$. It is interesting to note that $\beta(R, G, B)$ can be obtained by composing the above formulae. Such a composition can be extended to the k -coloring problem for handling any number of colors. \square

Example 5 (The not-all-equal-3-Sat problem [10]).

Input: A propositional formula in 3-CNF, encoded by the following relations:

- $clause(\cdot, \cdot, \cdot)$; a tuple $\langle l_1, l_2, l_3 \rangle$ is in $clause$ iff the formula contains clause $l_1 \vee l_2 \vee l_3$;
- $lit^+(\cdot, \cdot)$; a tuple $\langle l, v \rangle$ is in lit^+ iff l is the positive literal relative to the propositional variable v , i.e., v itself;
- $lit^-(\cdot, \cdot)$; a tuple $\langle l, v \rangle$ is in lit^- iff l is the negative literal relative to the propositional variable v , i.e., $\neg v$;
- $var(\cdot)$, containing the set of propositional variables occurring in the formula.

Question: Is there a way to assign a truth value to all the variables such that (i) the input formula is satisfied, and (ii) every clause contains at least one literal which truth value is false?

A specification for this problem is as follows (T and F list the set of variables whose truth value is true and false, respectively):

$$\exists T F \quad \forall X \quad var(X) \leftrightarrow T(X) \vee F(X) \wedge \quad (16)$$

$$\forall X \quad \neg(T(X) \wedge F(X)) \wedge \quad (17)$$

$$\begin{aligned} &\forall L_1, L_2, L_3 \quad clause(L_1, L_2, L_3) \rightarrow \\ &\quad [(\forall V_1 (lit^+(L_1, V_1) \rightarrow T(V_1)) \wedge (lit^-(L_1, V_1) \rightarrow F(V_1))) \vee \\ &\quad (\forall V_2 (lit^+(L_2, V_2) \rightarrow T(V_2)) \wedge (lit^-(L_2, V_2) \rightarrow F(V_2))) \vee \\ &\quad (\forall V_3 (lit^+(L_3, V_3) \rightarrow T(V_3)) \wedge (lit^-(L_3, V_3) \rightarrow F(V_3)))] \wedge \end{aligned} \quad (18)$$

$$\begin{aligned} &\forall L_1, L_2, L_3 \quad clause(L_1, L_2, L_3) \rightarrow \\ &\quad [(\forall V_1 (lit^+(L_1, V_1) \rightarrow F(V_1)) \wedge (lit^-(L_1, V_1) \rightarrow T(V_1))) \vee \\ &\quad (\forall V_2 (lit^+(L_2, V_2) \rightarrow F(V_2)) \wedge (lit^-(L_2, V_2) \rightarrow T(V_2))) \vee \\ &\quad (\forall V_3 (lit^+(L_3, V_3) \rightarrow F(V_3)) \wedge (lit^-(L_3, V_3) \rightarrow T(V_3)))] . \end{aligned} \quad (19)$$

Constraints (16–17) force every variable to be assigned exactly one truth value; moreover, (18) forces the assignment to be a model of the formula, while (19) leaves in every clause at least one literal whose truth value is false.

UVT σ , such that $\sigma(T) = F$ and $\sigma(F) = T$ is a UVS for this problem, since ϕ^σ is clearly equivalent to ϕ .

A symmetry-breaking formula for σ is the following one:

$$\beta(T, F) \doteq |T| \leq |F| \quad (20)$$

which forces the truth assignment to assign true at most the number of variables assigned to false. The above formula satisfies both conditions of Definition 6.

A generalization of the above problem is *not-all-equal-Sat*, in which the input formula in conjunctive form can have arbitrarily long clauses. The input database of the new problem has the following schema:

- *inclause*(\cdot, \cdot); tuple $\langle l, c \rangle$ is in *inclause* iff literal l is in clause c ;
- *lit*⁺(\cdot, \cdot); a tuple $\langle l, v \rangle$ is in *lit*⁺ iff l is the positive literal relative to the propositional variable v , i.e., v itself;
- *lit*⁻(\cdot, \cdot); a tuple $\langle l, v \rangle$ is in *lit*⁻ iff l is the negative literal relative to the propositional variable v , i.e., $\neg v$;
- *var*(\cdot), containing the set of propositional variables occurring in the formula;
- *clause*(\cdot), containing the set of clauses of the formula;

and a specification is as follows:

$$\exists T F \quad \forall X \text{ var}(X) \leftrightarrow T(X) \vee F(X) \wedge \quad (21)$$

$$\forall X \neg (T(X) \wedge F(X)) \wedge \quad (22)$$

$$\forall C \text{ clause}(C) \rightarrow$$

$$\left[\exists L \text{ inclosure}(L, C) \wedge \quad (23)$$

$$\forall V (\text{lit}^+(L, V) \rightarrow T(V)) \wedge (\text{lit}^-(L, V) \rightarrow F(V)) \right] \wedge$$

$$\forall C \text{ clause}(C) \rightarrow$$

$$\left[\exists L \text{ inclosure}(L, C) \wedge \quad (24)$$

$$\forall V (\text{lit}^+(L, V) \rightarrow F(V)) \wedge (\text{lit}^-(L, V) \rightarrow T(V)) \right].$$

Similarly to the not-all-equal-3-Sat problem, UVT σ , such that $\sigma(T) = F$ and $\sigma(F) = T$ is a UVS for this problem, and (20) is a symmetry-breaking formula for it. \square

Example 6 (The social golfer problem (Problem 10 at www.csplib.org)).

Input: Positive integers ng , nw and ngr , denoting respectively the number of golfers, weeks, and groups.

Question: Is there a way to arrange, for w weeks, ng golfers in ngr groups of the same size, such that any two golfers play in the same group at most once?

A specification for this problem (supposing the ratio ng/ngr , i.e., the group size, integral) is the following one (*ASSIGN*(G, W, GR) states that golfer G plays in group GR on week W):

$\exists ASSIGN$

$$\forall G, W \exists GR \text{ ASSIGN}(G, W, GR) \wedge \quad (25)$$

$$\forall G, W, GR, GR' \text{ ASSIGN}(G, W, GR) \wedge \text{ASSIGN}(G, W, GR') \rightarrow \quad (26)$$

$$GR = GR' \wedge$$

$$\forall G, G', W, W' (G \neq G' \wedge W \neq W') \rightarrow$$

$$\{\forall GR, GR' (GR \neq GR' \wedge \text{ASSIGN}(G, W, GR) \wedge \quad (27)$$

$$\text{ASSIGN}(G', W, GR)) \rightarrow$$

$$\neg[\text{ASSIGN}(G, W', GR') \wedge \text{ASSIGN}(G', W', GR')]\} \wedge$$

$$\forall GR, GR', W, W' |\{G : \text{ASSIGN}(G, W, GR)\}| = \quad (28)$$

$$|\{G : \text{ASSIGN}(G, W', GR')\}|.$$

Constraints (25–26) force *ASSIGN* to be a total function assigning a group to each golfer on each week; moreover, (27) is the *meet only once* constraint, while (28) forces groups to be of the same size.

In order to highlight UVs according to Definition 5, we need to substitute the ternary guessed predicate *ASSIGN* by means of $nw \times ngr$ monadic predicates $ASSIGN_{W,GR}$ (each one listing golfers playing in group *GR* on week *W*). The above specification must be unfolded accordingly. In this way, UVTs $\sigma_W^{GR,GR'}$, swapping $ASSIGN_{W,GR}$ and $ASSIGN_{W,GR'}$, i.e., given a week *W*, and two groups *GR* and *GR'*, assign to group *GR'* on week *W* all golfers assigned to group *GR* on week *W*, and vice versa, are symmetries for the social golfer problem. Intuitively, UVTs $\sigma_W^{GR,GR'}$ are UVs for the unfolded specification because group renamings have no effect.

After fixing arbitrarily linear orders ' $<$ ' on golfers and groups, each symmetry $\sigma_W^{GR,GR'}$ (with $GR \neq GR'$) can be broken by, e.g., the following formula:

$$\beta_W^{GR,GR'}(ASSIGN_{W,GR}, ASSIGN_{W,GR'}) \doteq$$

$$\forall G \text{ least}(G, ASSIGN_{W,GR} \cup ASSIGN_{W,GR'}) \rightarrow \quad (29)$$

$$\rightarrow \begin{cases} ASSIGN_{W,GR}(G) & \text{if } GR < GR' \\ ASSIGN_{W,GR'}(G) & \text{otherwise,} \end{cases}$$

which forces the least golfer in $ASSIGN_{W,GR} \cup ASSIGN_{W,GR'}$ to play in the least group between *GR* and *GR'*.

Similarly to Example 4, we can simultaneously break several of the above symmetries, by adding the following constraints, one for each week *W* and each pair *GR, GR'* of groups such that *GR'* is the successor of *GR* in the given linear order ' $<$ ':

$$\forall G \text{ least}(G, ASSIGN_{W,GR}) < \text{least}(G, ASSIGN_{W,GR'}). \quad (30)$$

The formulae say that, for each week, the least golfer in group GR precedes the least golfer in group GR' . An implication of the above formulae is that the first golfer always plays in the first group.

As another example, the following family of transformations (which we call $\sigma^{W,W'}$) are also symmetries for the original social golfer specification: given two weeks W and W' , swap uniformly all groups in week W with the ones in week W' , and vice versa.

Each symmetry $\sigma^{W,W'}$ can be broken: in the following, we give only an intuitive description of a possible symmetry-breaking formula $\beta^{W,W'}$, once total orders ' $<$ ' on weeks and on weekly assignments, i.e., on sets $\{ASSIGN_{W,GR_1}, \dots, ASSIGN_{W,GR_{n_{gr}}}\}$ have been fixed: given two assignments of golfers to groups $\{ASSIGN_{W,GR_1}, \dots, ASSIGN_{W,GR_{n_{gr}}}\}$ and $\{ASSIGN_{W',GR_1}, \dots, ASSIGN_{W',GR_{n_{gr}}}\}$, on weeks W and W' , they are swapped iff the first is greater than the second one.

By breaking several of the aforementioned symmetries, we can add the following constraints, one for each pair W, W' of weeks such that W' is the successor of W in the given linear order ' $<$ ':

$$\{ASSIGN_{W,GR_1}, \dots, ASSIGN_{W,GR_{n_{gr}}}\} < \{ASSIGN_{W',GR_1}, \dots, ASSIGN_{W',GR_{n_{gr}}}\}, \quad (31)$$

saying that weekly assignments are in increasing order wrt weeks.

It is worth noting that, by combining the symmetry-breaking formulae for the two families of symmetries described above, we are able to get some of the symmetry-breaking constraints described in [17], like the one which arbitrarily fixes the assignment in the first week. Moreover, we can further fix one golfer (the one having the least order number) to be assigned always to the same group (the one having the least order number). \square

Example 7 (The Balanced Incomplete Block Design (BIBD) generation problem (Problem 28 at www.csplib.org)).

Input: Positive integers v, b, r, k, λ .

Question: Is there a way to arrange v distinct objects into b blocks such that each block contains exactly k distinct objects, each object occurs in exactly r different blocks, and every two distinct objects occur together in exactly λ blocks?

A specification for this problem is quite similar to the one of Example 6, and also its symmetries are similar. We don't give here more details about this problem. Just as an example, it is always possible to swap the order of two blocks. Symmetry-breaking formulae can be built using the same techniques of Example 6. \square

4 Experiments

In this section we present a preliminary experimentation of the previously described techniques for the graph coloring problem shown in Section 3. Actually,

Table 1. Solving times (seconds) for k -coloring

Instance	Colors	Solvable?	CPLEX			SOLVER		
			W/o s.b.	W s.b.	% saving	W/o s.b.	W s.b.	% saving
anna	11	Y	3.2	1.3	59.4	0.8	–	$-\infty$
DSJC125.9	21	N	1388.7	1206.2	13.1	–	–	–
DSJR500.1	12	Y	22.5	32.1	-42.7	0.7	0.6	14.3
fpsol2.i.2	21	N	139.8	102.2	26.9	–	–	–
le450_15a	13	N	8.6	4.2	51.2	–	–	–
le450_25a	21	N	85.4	23.3	72.7	–	–	–
queen9_9	10	Y	–	–	–	122.9	205.6	-67.3
queen12_12	15	Y	453.6	357.1	21.3	8.7	9.4	-8.0

(‘–’ means that the solver did not terminate in one hour)

rather than 3-coloring, we focused on the k -coloring problem. We used state-of-the-art linear and constraint programming solvers, and wrote both a linear and a non linear specification, solved by using Ilog’s CPLEX and SOLVER, respectively, cf. www.ilog.com.

The symmetry breaking formulae we added are the generalization of (14) to the case of k colors. In fact, using the generalization of (15) resulted in poorer performances: intuitively, solutions of coloring problems are rarely “color unbalanced”, therefore adding too many constraints may not be effective. As shown in Table 1 for some of the graphs in the DIMACS benchmark repository (<ftp://dimacs.rutgers.edu/pub/challenge>), adding our symmetry-breaking formulae seems to be effective especially when CPLEX is used on negative instances. SOLVER seems to be negatively affected, although in some cases is outperformed by CPLEX. Both SOLVER and CPLEX have built-in features for symmetry cuts; all experiments have been done both enabling and disabling such a feature (Table 1 shows times with no symmetry cuts). In general, effectiveness of our technique is independent on using that feature; the built-in feature results anyway in further speed-up when used along with our formulas (thus confirming the intuition of Section 1).

Acknowledgements This research has been supported by MIUR (Italian Ministry for Instruction, University, and Research) under the FIRB project ASTRO (Automazione dell’Ingegneria del Software basata su Conoscenza).

The authors are grateful to Igor Markov for interesting discussions on symmetries and for pointing out references, and to the anonymous Referees, whose comments and suggestions led to substantial improvements in the quality of this paper.

References

1. E. Börger, E. Gräedel, and Y. Gurevich. *The Classical Decision Problem*. Perspectives in Mathematical Logic. Springer, 1997.

2. C. A. Brown, L. Finkelstein, and P. W. Purdom. Backtrack searching in the presence of symmetry. In T. Mora, editor, *Proc. of 6th Intl. Conf. on Applied Algebra, Algebraic Algorithms and Error Correcting codes*, pages 99–110. Springer, 1988.
3. M. Cadoli and A. Schaerf. Compiling problem specifications into SAT. In *Proceedings of the European Symposium On Programming (ESOP 2001)*, volume 2028 of *LNCS*, pages 387–401. Springer, 2001.
4. E. Castillo, A. J. Conejo, P. Pedregal, and N. Alguacil Ricardo Garca. *Building and Solving Mathematical Programming Models in Engineering and Science*. Wiley, 2001.
5. J. M. Crawford, M. L. Ginsberg, E. M. Luks, and A. Roy. Symmetry-breaking predicates for search problems. In *Proc. of KR'96*, pages 148–159, 1996.
6. T. Eiter, N. Leone, C. Mateis, G. Pfeifer, and F. Scarcello. The KR system dlv: Progress report, Comparisons and Benchmarks. In *Proc. of KR'98*, pages 406–417, 1998.
7. R. Fagin. Generalized First-Order Spectra and Polynomial-Time Recognizable Sets. In R. M. Karp, editor, *Complexity of Computation*, pages 43–74. AMS, 1974.
8. F. Focacci and M. Milano. Global cut framework for removing symmetries. In *Proc. of CP 2001*, volume 2239 of *LNCS*, page 77 ff. Springer, 2001.
9. R. Fourer, D. M. Gay, and B. W. Kernigham. *AMPL: A Modeling Language for Mathematical Programming*. International Thomson Publishing, 1993.
10. M. R. Garey and D. S. Johnson. *Computers and Intractability—A guide to NP-completeness*. W.H. Freeman and Company, San Francisco, 1979.
11. I. P. Gent and B. Smith. Symmetry breaking during search in constraint programming. In *Proc. of ECAI 2000*, pages 599–603, 2000.
12. P. Meseguer and C. Torras. Solving strategies for highly symmetric CSPs. In *Proc. of IJCAI'99*, pages 400–405, 1999.
13. P. Meseguer and C. Torras. Exploiting symmetries within constraint satisfaction search. *Journal of Artificial Intelligence*, 129:133–163, 2001.
14. I. Niemelä. Logic programs with stable model semantics as a constraint programming paradigm. *Ann. of Mathematics and Artificial Intelligence*, 25(3,4):241–273, 1999.
15. E. Rothberg. Using cuts to remove symmetry. In *Proc. of ISMP'00*, 2000.
16. H. D. Sherali and J. Cole Smith. Improving Discrete Model Representations Via Symmetry Considerations. *Management Science*, 47:1396–1407, 2001.
17. B. Smith. Reducing symmetry in a combinatorial design problem. In *Proc. of CPAIOR'01*, pages 351–360, 2001.
18. P. Van Hentenryck. *The OPL Optimization Programming Language*. The MIT Press, 1999.