Colors Make Theories Hard

Roberto Sebastiani$^{(\boxtimes)}$

DISI, University of Trento, Trento, Italy
`roberto.sebastiani@unitn.it`

Abstract. The satisfiability problem for conjunctions of quantifier-free literals in first-order theories $\mathcal{T}$ of interest–"$\mathcal{T}$-*solving*" for short–has been deeply investigated for more than three decades from both theoretical and practical perspectives, and it is currently a core issue of state-of-the-art SMT solving. Given some theory $\mathcal{T}$ of interest, a key theoretical problem is to establish the computational *(in)tractability* of $\mathcal{T}$-solving, or to identify intractable fragments of $\mathcal{T}$.

In this paper we investigate this problem from a general perspective, and we present a simple and general criterion for establishing the NP-hardness of $\mathcal{T}$-solving, which is based on the novel concept of "*colorer*" for a theory $\mathcal{T}$.

As a proof of concept, we show the effectiveness and simplicity of this novel criterion by easily producing very simple proofs of the NP-hardness for many theories of interest for SMT, or of some of their fragments.

1 Introduction

Since the pioneering works of the late 70's and early 80's by Nelson, Oppen, Shostak and others [16,17,19–21,25,26], the satisfiability problem for conjunctions of quantifier-free literals in first-order theories $\mathcal{T}$ of interest–hereafter "$\mathcal{T}$-*solving*" for short–has been deeply investigated from both theoretical and practical perspectives, and it is currently a core issue of state-of-the-art SMT solving.

Given some theory $\mathcal{T}$ of interest, or some fragment thereof, a key theoretical problem is that of establishing the computational *(in)tractability* of $\mathcal{T}$-solving, or to identify (in)tractable fragments of $\mathcal{T}$. Although in the pool of theories of interest $\mathcal{T}$-solving presents many levels of intractability, the main divide is between polynomiality and NP-hardness. Despite a wide literature studying the complexity of single theories or of families of theories (e.g. [5,7,8,10,11,13–15, 17,19–21]) and some more general work on complexity of $\mathcal{T}$-solving [3,20,21], we are not aware of any previous work explicitly addressing NP-hardness of $\mathcal{T}$-solving for a generic theory $\mathcal{T}$.

In this paper we try to fill this gap, and we present a simple and general criterion for establishing the NP-hardness of $\mathcal{T}$-solving for theories with equality– and in some cases also for theories without equality–which is based on the novel concept of "*colorer*" for a theory $\mathcal{T}$, inducing the notion of "*colorable*" theory.

Our work started from the heuristic observation that the *graph k-colorability problem*, which is NP-complete for $k \geq 3$, fits very naturally as a candidate problem to be polynomially encoded into $\mathcal{T}$-solving for theories with equality. (We believe, more naturally than the very frequently-used 3-SAT problem.) In fact, we notice that the set of the arcs in a graph and the coloring of the vertexes can be encoded respectively into a conjunction of disequalities between "vertex" variables and into a conjunction of equalities between "vertex" and "color" variables, *both of which are theory-independent.* Therefore, in designing a reduction from k-colorability to $\mathcal{T}$-solving, the only facts one needs formalizing by $\mathcal{T}$-specific literals is a coherent definition of k distinct "colors" and the fact that a generic vertex can be "colored" with and only with k colors.

Following this line of thought, in this paper we present a general framework for producing reductions from graph k-colorability with $k \geq 3$ to $\mathcal{T}$-solving for generic theories $\mathcal{T}$ with equality. This framework decouples the $\mathcal{T}$-specific part of a reduction from its $\mathcal{T}$-independent part: the former is formalized into the definition of a $\mathcal{T}$-specific object, called "*k-colorer*", the latter is formalized and proven once forall in this paper. Thus, the task of proving the NP-hardness of a theory $\mathcal{T}$ via reduction from k-colorability reduces to that of finding a k-colorer for $\mathcal{T}$.

To this extent, we also provide some general criteria for producing k-colorers, with hints and tips to achieve this simplified task. As a proof of concept, we show the effectiveness and simplicity of this novel approach by easily producing k-colorers with $k \geq 3$ for many theories of interest for SMT, or for some of their fragments.

We notice that this technique can be used not only to investigate the intractability of major theories, but also to investigate that of *fragments* of such theories, so that to pinpoint the subsets of constructs (i.e. functions and predicates in the signature) which cause a theory to be intractable. We stress the fact that the problem of identifying such intractable fragments is not only of theoretical interest, but also of practical importance in the development of SMT solvers, in order to drive the activation of ad-hoc techniques–including e.g. *weakened early pruning*, *layering*, *splitting-on-demand* [1,4]–which partition the search load among distinct specialized $\mathcal{T}$-solvers and between the $\mathcal{T}$-solvers and the underlining SAT solver [2,23].

Note. An extended version of this paper with more details is publicly available [24].

Content. The rest of the paper is organized as follows: Sect. 2 provides the necessary background knowledge and terminology for logic and graph coloring; Sect. 3 introduces our main definitions of k-colorer and k-colorability and presents our main results; Sect. 4 explains how to produce k-colorers for given theories, providing a list of examples; Sect. 5 provides some discussion about k-colorability vs. non-convexity; Sect. 6 extends the framework to theories without equality; Sect. 7 discusses ongoing and future developments.

2 Background and Terminology

Logic. We assume the reader is familiar with the standard syntax and semantics of first-order logic. (We report a full description in [24].) We add some terminology.

Given a signature Σ, we call Σ-*theory* $\mathcal{T}$ a class of Σ-models. Given a theory $\mathcal{T}$, we call $\mathcal{T}$-*interpretation* an extension of some Σ-model $\mathcal{M}$ in $\mathcal{T}$ which maps free variables into elements of the domain of $\mathcal{M}$. (The map is denoted by $\langle . \rangle^{\mathcal{I}}$.) A Σ-formula φ–possibly with free variables–is $\mathcal{T}$-*satisfiable* if $\mathcal{I} \models \varphi$ for some $\mathcal{T}$-interpretation $\mathcal{I}$. (Hereafter we will use the symbol "$\models_{\mathcal{T}}$" to denote the $\mathcal{T}$-satisfiability relation; we will also drop the prefix "Σ-" when the signature is implicit by context.) We say that a set/conjunction of formulas Ψ $\mathcal{T}$-*entails* another formula φ, written $\Psi \models_{\mathcal{T}} \varphi$, if every $\mathcal{T}$-interpretation $\mathcal{T}$-satisfying Ψ also $\mathcal{T}$-satisfies φ. We say that φ is $\mathcal{T}$ -*valid*, written $\models_{\mathcal{T}} \varphi$, if $\emptyset \models_{\mathcal{T}} \varphi$. We call a *cube* any finite quantifier-free conjunction of literals. For short, we call "$\mathcal{T}$-*solving*" the problem of deciding the $\mathcal{T}$-satisfiability of a cube.

Finally, a theory $\mathcal{T}$ is *convex* if for all cubes μ and all sets E of equalities between variables, $\mu \models_{\mathcal{T}} \bigvee_{e \in E} e$ iff $\mu \models_{\mathcal{T}} e$ for some $e \in E$.

Remark 1. In SMT and other contexts it is often convenient to use formulas with *uninterpreted symbols* (see e.g. [2]). Notice, however, that the presence of uninterpreted function or predicate symbols of arity >0 may cause the complexity of $\mathcal{T}$-solving scale up (see e.g. the example in [21]). Thus, when not explicitly specified otherwise, we implicitly assume that a theory $\mathcal{T}$ does *not* admit such symbols.   ⬦

We are often interested in fragments of a theory obtained by restricting its signature. Let Σ, Σ' be two signatures s.t. $\Sigma' \subseteq \Sigma$; we say that a Σ'-model $\mathcal{M}'$ is a *restriction to Σ'* of a Σ-model $\mathcal{M}$ iff $\mathcal{M}'$ and $\mathcal{M}$ agree on all the symbols in Σ', and that a Σ'-theory $\mathcal{T}'$ is the *signature-restriction fragment* of a Σ-theory $\mathcal{T}$ wrt. Σ' iff $\mathcal{T}'$ is the set of the restrictions to Σ' of the Σ-models in $\mathcal{T}$.

Graph Coloring. We recall a few notions from [9].

Definition 1 (k-Colorability of a graph (see [9])). *Let $\mathcal{G} \stackrel{def}{=} \langle \mathcal{V}, \mathcal{E} \rangle$ be an un-directed graph, where $\mathcal{V} \stackrel{def}{=} \{V_1, ..., V_n\}$ is the set of vertexes and $\mathcal{E} \stackrel{def}{=} \{E_1, ..., E_m\}$ is the set of edges in the form $\langle V_i, V_{i'} \rangle$ for some i, i'. Let $\mathcal{C} \stackrel{def}{=} \{C_1, ..., C_k\}$ be a set of distinct values, namely "colors", for $k>0$. Then $\mathcal{G}$ is k-**Colorable** if and only if there exists a total map $color : \mathcal{V} \longmapsto \mathcal{C}$ s.t. $color(V_i) \neq color(V_{i'})$ for every $\langle V_i, V_{i'} \rangle \in \mathcal{E}$. The problem of deciding if $\mathcal{G}$ is k-colorable is called the k-colorability problem for $\mathcal{G}$.*

Lemma 1 (see [9]). *The k-colorability problem for un-directed graphs is NP-complete for $k \geq 3$, it is in P for $k<3$.*

Figure 1 (top) shows two small graph 3-colorability problems.

$$\mathsf{Enc}_{[\mathcal{G}_1 \Rightarrow \mathcal{LA}(\mathbb{Z})]} \overset{\text{def}}{=} \left(\begin{array}{c} (c_1 = 1) \wedge (c_2 = 2) \wedge (c_3 = 3) \wedge \bigwedge_{i=1}^{4}((v_i \geq 1) \wedge (v_i \leq 3)) \wedge \\ \neg(v_1 = v_2) \wedge \neg(v_1 = v_3) \wedge \neg(v_1 = v_4) \wedge \neg(v_2 = v_3) \wedge \neg(v_2 = v_4) \end{array} \right)$$

$$\mathsf{Enc}_{[\mathcal{G}_2 \Rightarrow \mathcal{LA}(\mathbb{Z})]} \overset{\text{def}}{=} \mathsf{Enc}_{[\mathcal{G}_1 \Rightarrow \mathcal{LA}(\mathbb{Z})]} \wedge \neg(v_3 = v_4)$$

Fig. 1. Top Left: a small 3-colorable graph ($\mathcal{G}_1$), with $C_1 = blue$, $C_2 = red$, $C_3 = green$. Top Right: the same graph augmented with the vertex $\langle V_3, V_4 \rangle$ ($\mathcal{G}_2$) is no more 3-colorable. Bottom: example of encodings of the 3-colorability of $\mathcal{G}_1$ and $\mathcal{G}_2$ into $\mathcal{LA}(\mathbb{Z})$ -solving. (Color figure online)

3 k-Colorers and k-Colorable Theories with Equality

Hereafter we focus w.l.o.g. on theories $\mathcal{T}$ of domain size ≥ 2, i.e., s.t. $\neg(v_1 = v_2)$ is $\mathcal{T}$-consistent. In fact, if not so, then it is easy to see that $\mathcal{T}$-solving is in P (see [24]).

Definition 2 (k-Colorer, k-Colored Theory). *Let $\mathcal{T}$ be some theory with equality and k be some integer value s.t. $k \geq 2$. Let v_i be a variable, called* **vertex variable***, (implicitly) denoting the i-th vertex in an un-directed graph; let $\underline{\mathbf{c}} \overset{\text{def}}{=} \{c_1, .., c_k\}$ be a set of variables, called* **color variables***, denoting the set of colors; let $\underline{\mathbf{y}}_i \overset{\text{def}}{=} \{y_{i1}, ..., y_{il}\}$ denote a possibly-empty set of variables, which is indexed with the same index i of the vertex variable v_i. Let $\mathsf{AllDifferent}_k(\underline{\mathbf{c}}) \overset{\text{def}}{=} \bigwedge_{j=1}^{k} \bigwedge_{j'=j+1}^{k} \neg(c_j = c_{j'})$.*
*   We call k-***colorer** *for $\mathcal{T}$, namely $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i)$, a finite quantifier-free conjunction of $\mathcal{T}$-literals (cube) over v_i, $\underline{\mathbf{c}}$ and $\underline{\mathbf{y}}_i$ which verify the following properties:*

$$\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i) \models_{\mathcal{T}} \mathsf{AllDifferent}_k(\underline{\mathbf{c}}), \tag{1}$$

$$\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i) \models_{\mathcal{T}} \bigvee_{j=1}^{k}(v_i = c_j), \tag{2}$$

there exist k $\mathcal{T}$-interpretations $\{\mathcal{I}_{i,1}, ..., \mathcal{I}_{i,k}\}$ s.t. $\qquad\qquad$ (3)
$\qquad$ *for every $j \in [1..k]$, $\langle c_j \rangle^{\mathcal{I}_{i,1}} = \langle c_j \rangle^{\mathcal{I}_{i,2}} = ... = \langle c_j \rangle^{\mathcal{I}_{i,k}}$, and*
$\qquad$ *for every $j \in [1..k]$, $\mathcal{I}_{i,j} \models_{\mathcal{T}} \mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i) \wedge (v_i = c_j)$.*

*We say that $\mathcal{T}$ is k-***colorable** *if and only if it has a k-colorer.*

$\underline{\mathbf{y}}_i$ is a (possibly-empty) set of auxiliary variables, one distinct set for each vertex variable v_i, which sometimes may be needed to express (1), (2) and (3) (see Examples 7 and 9), or to make $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i)$ more readable by renaming

internal terms (see Example 9). If $\underline{\mathbf{y}}_i = \emptyset$, we may write "$\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}})$" instead of "$\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\emptyset)$".[1]

$\{\mathcal{I}_{i,1}, ..., \mathcal{I}_{i,k}\}$ denotes a set of $\mathcal{T}$-interpretations each satisfying $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i)$ s.t. all the $\mathcal{T}$-interpretations in $\{\mathcal{I}_{i,1}, ..., \mathcal{I}_{i,k}\}$ agree on the values assigned to the color variables in $\{c_1, ..., c_k\}$ and s.t. each $\mathcal{I}_{i,j}$ assigns to the vertex variable v_i the same value assigned to the jth color variable c_j. The condition $\langle c_j \rangle^{\mathcal{I}_{i,1}} = ... = \langle c_j \rangle^{\mathcal{I}_{i,k}}$ of (3) expresses the fact that, when passing from the scenario $\mathcal{I}_{i,j}$ in which v_i is assigned the color c_j–expressed by the equality $(v_i = c_j)$ in (3)–to the scenario $\mathcal{I}_{i,j'}$ in which v_i is assigned the color $c_{j'}$–expressed by the equality $(v_i = c_{j'})$– it is the value of the vertex variable v_i who must change, not those of the color variables $c_1, ..., c_k$.

Intuitively, $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i)$ expresses the following facts: (1) that $c_1, ..., c_k$ represent the names of distinct "color" values, (2) that each vertex represented by the variable v_i can be tagged ("colored") only with one of such color names c_j, (3) that the values associated to the color names are not affected by the choice of the color name c_j tagged to v_i–represented by the index j in $\mathcal{I}_{i,j}$–and that each tagging choice is admissible.

There may be many distinct k-colorers for a theory $\mathcal{T}$, as shown in Example 1.

Example 1 ($\mathcal{LA}(\mathbb{Z})$). We consider the theory of linear arithmetic over the integers ($\mathcal{LA}(\mathbb{Z})$), assuming the standard model of integers, so that the symbols $+, -, \leq, \geq$ and the interpreted constants $0, 1, ...$ are interpreted in the standard way by all $\mathcal{LA}(\mathbb{Z})$-interpretations. $\mathcal{LA}(\mathbb{Z})$ is 3-colorable, since we can define, e.g., $k \stackrel{\text{def}}{=} 3$, $\mathbf{y}_i \stackrel{\text{def}}{=} \emptyset$, and

$$\mathsf{Colorer}_3(v_i, c_1, c_2, c_3) \stackrel{\text{def}}{=} (c_1 = 1) \wedge (c_2 = 2) \wedge (c_3 = 3) \wedge (v \geq 1) \wedge (v \leq 3). \quad (4)$$

It is straightforward to see that $\mathsf{Colorer}_3(v_i, c_1, c_2, c_3)$ verifies (1), (2) and (3), with $\mathcal{I}_{i,j} \stackrel{\text{def}}{=} \{c_1 \to 1, c_2 \to 2, c_3 \to 3, v_i \to j\}$ for every $j \in [1..3]$. Notice that in this case $\mathbf{y}_i = \emptyset$, i.e. $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i)$ requires no auxiliary variables. Notice also that $\mathsf{AllDifferent}_k(\underline{\mathbf{c}})$ is implied by the usage of the interpreted constants $1, 2, 3$.

An alternative 3-colorer which does not explicitly assign fixed values to the c_j's is:

$$\mathsf{Colorer}_3(v_i, c_1, c_2, c_3) \stackrel{\text{def}}{=} \begin{pmatrix} \mathsf{AllDifferent}_3(\underline{\mathbf{c}}) \wedge \bigwedge_{j=1}^{3}((c_j \geq 1) \wedge (c_j \leq 3)) \wedge \\ (v \geq 1) \wedge (v \leq 3) \end{pmatrix}, (5)$$

which verifies (1), (2) and (3), e.g., with the same $\mathcal{I}_{i,j}$'s as above. Consider instead:

$$\mathsf{Colorer}_3(v_i, c_1, c_2, c_3) \stackrel{\text{def}}{=} \begin{pmatrix} \mathsf{AllDifferent}_3(\underline{\mathbf{c}}) \wedge \bigwedge_{j=1}^{3}((c_j \geq 1) \wedge (c_j \leq 3)) \wedge \\ (v_i = 1) \end{pmatrix}. (6)$$

This is not a 3-colorer, because it does not verify (3): there is no pair of $\mathcal{LA}(\mathbb{Z})$-interpretations $\mathcal{I}_{i,1}$ and $\mathcal{I}_{i,2}$ s.t. $\mathcal{I}_{i,1} \models_{\mathcal{LA}(\mathbb{Z})} \mathsf{Colorer}_3(v_i, c_1, c_2, c_3) \wedge (v_i = c_1)$ and $\mathcal{I}_{i,2} \models_{\mathcal{LA}(\mathbb{Z})} \mathsf{Colorer}_3(v_i, c_1, c_2, c_3) \wedge (v_i = c_2)$ which agree on the values of c_1, c_2, c_3. ◇

---

[1] The symbol "|" is used to separate color and node variables from auxiliary ones.

Remark 2. The choice of using *variables* $c_1, ..., c_k$ to represent colors is due to the fact that some theories do not provide k distinct interpreted constant symbols in their signature (see Example 9). If this is not the case, then $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i)$ can be built to force $c_1, ..., c_k$ to assume fixed values expressed by interpreted constant symbols, like $1, 2, 3$ in (4), so that the condition $\langle c_j \rangle^{\mathcal{I}_{i,1}} = ... = \langle c_j \rangle^{\mathcal{I}_{i,k}}$ of (3) is verified a priori.

The following properties of k-colorable theories follow straightforwardly.

Property 1. Let $\mathcal{T}$ be a k-colorable theory for some $k \geq 2$. Then we have that:

(a) $\exists \underline{\mathbf{c}}.\mathsf{AllDifferent}_k(\underline{\mathbf{c}})$ is $\mathcal{T}$-valid;
(b) $\mathcal{T}$ is non-convex.

Proof. Consider the definition of $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i)$ in Definition 2.

(a) By (3) $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i)$ is $\mathcal{T}$-satisfiable; thus by (1) $\mathsf{AllDifferent}_k(\underline{\mathbf{c}})$ is $\mathcal{T}$-satisfiable, so that $\models_{\mathcal{T}} \exists \underline{\mathbf{c}}.\mathsf{AllDifferent}_k(\underline{\mathbf{c}})$;
(b) By (2), $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i) \models_{\mathcal{T}} \bigvee_{j=1}^{k}(v_i = c_j)$. By (3), for every $j_1 \in [1..k]$ there exists an interpretation $\mathcal{I}_{i,j_1}$ s.t. $\mathcal{I}_{i,j_1} \models_{\mathcal{T}} \mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i) \wedge (v_i = c_{j_1})$. Then, by (1), for every $j_2 \in [1..k]$ s.t. $j_2 \neq j_1$ we have that $\mathcal{I}_{i,j_1} \models_{\mathcal{T}} \mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i) \wedge \neg(v_i = c_{j_2})$. Thus for every $j \in [1..k]$ $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i) \not\models (v_i = c_j)$. Therefore $\mathcal{T}$ is non-convex. $\square$

Property 2. If $\mathcal{T}'$ is a k-colorable theory with equality for some $k \geq 2$, and $\mathcal{T}'$ is a signature-restriction fragment of another theory $\mathcal{T}$, then $\mathcal{T}$ is k-colorable.

Proof. If $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i)$ is a k-colorer for $\mathcal{T}'$, then by definition of signature-restriction fragment it is also a k-colorer for $\mathcal{T}$. $\square$

Lemma 2. *Let k be an integer value s.t. $k \geq 3$. Let $\mathcal{G}$ and $\mathcal{C}$ be respectively an un-directed graph with n vertexes $V_1, ..., V_n$ and a set of k distinct colors $C_1, ..., C_k$, like in Definition 1. Let $\mathcal{T}$ be a k-colorable theory with equality. We consider the following conjunctions of $\mathcal{T}$-literals:*

$$\mathsf{Colorable}(v_1, ..., v_n, \underline{\mathbf{c}}|\underline{\mathbf{y}}_1, ..., \underline{\mathbf{y}}_n) \overset{def}{=} \bigwedge_{V_i \in \mathcal{V}} \mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i) \tag{7}$$

$$\mathsf{Graph}_{[\mathcal{G}]}(v_1, ..., v_n) \overset{def}{=} \bigwedge_{\langle V_{i_1}, V_{i_2} \rangle \in \mathcal{E}} \neg(v_{i_1} = v_{i_2}) \tag{8}$$

$$\mathsf{Enc}_{[\mathcal{G} \Rightarrow \mathcal{T}]}(v_1, ..., v_n, \underline{\mathbf{c}}|\underline{\mathbf{y}}_1, ..., \underline{\mathbf{y}}_n) \overset{def}{=} \mathsf{Colorable}(v_1, ..., v_n, \underline{\mathbf{c}}|\underline{\mathbf{y}}_1, ..., \underline{\mathbf{y}}_n) \wedge \tag{9}$$
$$\mathsf{Graph}_{[\mathcal{G}]}(v_1, ..., v_n),$$

where $v_1, ..., v_n$, $c_1, ..., c_k$ and $y_{11}, ..., y_{1l}, ...y_{i1}, ..., y_{il}, ..., y_{n1}, ..., y_{nl}$ are free variables,[2] and all the k-colorers $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i)$ in (7) are identical modulo the renaming of the variables v_i and $\underline{\mathbf{y}}_i$, but not of the color variables $\underline{\mathbf{c}}$.
Then $\mathcal{G}$ is k-colorable iff $\mathsf{Enc}_{[\mathcal{G} \Rightarrow \mathcal{T}]}(v_1, ..., v_n, \underline{\mathbf{c}}|\underline{\mathbf{y}}_1, ..., \underline{\mathbf{y}}_n)$ is $\mathcal{T}$-satisfiable.

---

[2] Notice that each c_j is implicitly associated with the color $C_j \in \mathcal{C}$ for every $j \in [1..k]$ and each v_i and $\underline{\mathbf{y}}_i$ is implicitly associated to the vertex $V_i \in \mathcal{V}$ for every $i \in [1..n]$.

Proof.

If: Suppose $\mathsf{Enc}_{[\mathcal{G} \Rightarrow \mathcal{T}]}(v_1, ..., v_n, \underline{\mathbf{c}}|\underline{\mathbf{y}}_1, ..., \underline{\mathbf{y}}_n)$ is $\mathcal{T}$-satisfiable, that is, there exist an interpretation $\mathcal{I}$ in $\mathcal{T}$ s.t. $\mathcal{I} \models_{\mathcal{T}} \mathsf{Colorable}(v_1, ..., v_n, \underline{\mathbf{c}}|\underline{\mathbf{y}}_1, ..., \underline{\mathbf{y}}_n)$ and $\mathcal{I} \models_{\mathcal{T}}$ $\mathsf{Graph}_{[\mathcal{G}]}(v_1, ..., v_n)$. Thus:
  (i) By (7) and (1), $\langle c_{j_1} \rangle^{\mathcal{I}} \neq \langle c_{j_2} \rangle^{\mathcal{I}}$ for every $j_1, j_2 \in [1, ..., k]$ s.t. $j_1 \neq j_2$.
  (ii) By (7), (2) and (1), for every $i \in [1...n]$ there exists some $j \in [1...k]$ s.t. $\langle v_i \rangle^{\mathcal{I}} = \langle c_j \rangle^{\mathcal{I}}$ and s.t. $\langle v_i \rangle^{\mathcal{I}} \neq \langle c_{j'} \rangle^{\mathcal{I}}$ for every $j' \neq j$.
  (iii) By (8), $\langle v_{i_1} \rangle^{\mathcal{I}} \neq \langle v_{i_2} \rangle^{\mathcal{I}}$ for every $\langle V_{i_1}, V_{i_2} \rangle \in \mathcal{E}$.
  Then by (i) and (ii) we can build a map $color : \mathcal{V} \longmapsto \mathcal{C}$ s.t., for every $V_i \in \mathcal{V}$, $color(V_i) = C_j$ iff $\langle v_i \rangle^{\mathcal{I}} = \langle c_j \rangle^{\mathcal{I}}$. By (iii) we have that $color(V_{i_1}) \neq color(V_{i_2})$ for every $\langle V_{i_1}, V_{i_2} \rangle \in \mathcal{E}$. Thus $\mathcal{G}$ is k-colorable.
Only if: Suppose $\mathcal{G}$ is k-colorable, that is, there exist a map $color : \mathcal{V} \longmapsto \mathcal{C}$ s.t. $color(V_{i_1}) \neq color(V_{i_2})$ for every $\langle V_{i_1}, V_{i_2} \rangle \in \mathcal{E}$.
  Consider $i = 1$, and let $\{\mathcal{I}_{1,1}, ..., \mathcal{I}_{1,k}\}$ be the set of $\mathcal{T}$-interpretations for $\mathsf{Colorer}_k(v_1, \underline{\mathbf{c}}|\underline{\mathbf{y}}_1)$ as in (3), so that:
  (a) for every $j \in [1..k]$, $\mathcal{I}_{1,j} \models_{\mathcal{T}} \mathsf{Colorer}_k(v_1, \underline{\mathbf{c}}|\underline{\mathbf{y}}_1) \wedge (v_1 = c_j)$,
  (b) for every $j \in [1..k]$, $\langle c_j \rangle^{\mathcal{I}_{1,1}} = ... = \langle c_j \rangle^{\mathcal{I}_{1,k}}$.
  For every $i \in [1..n]$ we consider $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i)$ and we build a replica $\{\mathcal{I}_{i,1}, ..., \mathcal{I}_{i,k}\}$ of the set of $\mathcal{T}$-interpretations $\{\mathcal{I}_{1,1}, ..., \mathcal{I}_{1,k}\}$ in such a way that:
  (i) $\langle v_i \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} \langle v_1 \rangle^{\mathcal{I}_{1,j}} = \langle c_j \rangle^{\mathcal{I}_{1,j}}$ (each $\mathcal{I}_{i,j}$ maps its vertex variable v_i into the same color as $\mathcal{I}_{1,j}$ maps its vertex variable v_1);
  (ii) $\langle c_j \rangle^{\mathcal{I}_{i,1}} \stackrel{\text{def}}{=} \langle c_j \rangle^{\mathcal{I}_{1,1}}, ..., \langle c_j \rangle^{\mathcal{I}_{i,k}} \stackrel{\text{def}}{=} \langle c_j \rangle^{\mathcal{I}_{1,k}}$, so that, by ($a$), $\langle c_j \rangle^{\mathcal{I}_{i,1}} = ... = \langle c_j \rangle^{\mathcal{I}_{i,k}} = \langle c_j \rangle^{\mathcal{I}_{1,1}} = ... = \langle c_j \rangle^{\mathcal{I}_{1,k}}$ (all $\mathcal{I}_{i,j}$ agree on the values of the color variables, for every $i \in [1..n]$ and $j \in [1..k]$);
  (iii) $\langle y_{i1} \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} \langle y_{11} \rangle^{\mathcal{I}_{1,j}}, ..., \langle y_{il} \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} \langle y_{1l} \rangle^{\mathcal{I}_{1,j}}$ (each $\mathcal{I}_{i,j}$ maps its auxiliary variables $\underline{\mathbf{y}}_i$ into the same domain values as $\mathcal{I}_{1,j}$ maps $\underline{\mathbf{y}}_1$).
  Consequently, by (3), for every $v_i \in \{v_1, ..., v_n\}$, $\{\mathcal{I}_{i,1}, ..., \mathcal{I}_{i,k}\}$ are s.t.
  (a) for every $j \in [1..k]$, $\mathcal{I}_{i,j} \models_{\mathcal{T}} \mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i) \wedge (v_i = c_j)$,
  (b) for every $j \in [1..k]$, $\langle c_j \rangle^{\mathcal{I}_{i,1}} = ... = \langle c_j \rangle^{\mathcal{I}_{i,k}}$.
  For every $i \in [1...n]$, let $j_i \in [1..k]$ be the index s.t. $C_{j_i} = color(V_i)$, and we pick the $\mathcal{T}$-interpretation $\mathcal{I}_{i,j_i}$. Thus, since all the $\mathcal{I}_{i,j_i}$s agree on the common variables $\underline{\mathbf{c}}$, we can merge them and create a global $\mathcal{T}$-interpretation $\mathcal{I}$ as follows:
  (i) $\langle v_i \rangle^{\mathcal{I}} \stackrel{\text{def}}{=} \langle v_i \rangle^{\mathcal{I}_{i,j_i}} = \langle c_{j_i} \rangle^{\mathcal{I}_{i,j_i}} = \langle c_{j_i} \rangle^{\mathcal{I}}$, for every $i \in [1..n]$;
  (ii) $\langle c_j \rangle^{\mathcal{I}} \stackrel{\text{def}}{=} \langle c_j \rangle^{\mathcal{I}_{i,j_i}}$, for every $j \in [1..k]$;
  (iii) $\langle y_{ir} \rangle^{\mathcal{I}} \stackrel{\text{def}}{=} \langle y_{ir} \rangle^{\mathcal{I}_{i,j_i}}$, for every $i \in [1..n]$ and for every $r \in [1..l]$.
  By construction, for every $i \in 1..n$, $\mathcal{I}$ agrees with $\mathcal{I}_{i,j_i}$ on $\underline{\mathbf{c}}$, v_i, and $\underline{\mathbf{y}}_i$, so that, by point (a), $\mathcal{I} \models_{\mathcal{T}} (\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i) \wedge (v_i = c_{j_i}))$.
  Thus $\mathcal{I} \models_{\mathcal{T}} \mathsf{Colorable}(v_1, ..., v_n, \underline{\mathbf{c}}|\underline{\mathbf{y}}_1, ..., \underline{\mathbf{y}}_n)$.
  Since the values $\langle c_1 \rangle^{\mathcal{I}}, ..., \langle c_k \rangle^{\mathcal{I}}$ are all distinct, we can build a bijection linking each domain value $\langle c_j \rangle^{\mathcal{I}}$ to the color C_j, for every $j \in [1..k]$. Hence $\langle c_j \rangle^{\mathcal{I}} = \langle c_{j'} \rangle^{\mathcal{I}}$ iff $C_j = C_{j'}$. For every $\langle V_i, V_{i'} \rangle \in \mathcal{E}$, $color(V_i) \neq color(V_{i'})$, that is, $C_{j_i} \neq C_{j_{i'}}$. Therefore $\langle c_{j_i} \rangle^{\mathcal{I}} \neq \langle c_{j_{i'}} \rangle^{\mathcal{I}}$, and $\langle v_i \rangle^{\mathcal{I}} = \langle c_{j_i} \rangle^{\mathcal{I}} \neq \langle c_{j_{i'}} \rangle^{\mathcal{I}} = \langle v_{i'} \rangle^{\mathcal{I}}$.
  Consequently $\mathcal{I} \models_{\mathcal{T}} \mathsf{Graph}_{[\mathcal{G}]}(v_1, ..., v_n)$.
  Thus $\mathsf{Enc}_{[\mathcal{G} \Rightarrow \mathcal{T}]}(v_1, ..., v_n, \underline{\mathbf{c}}|\underline{\mathbf{y}}_1, ..., \underline{\mathbf{y}}_n)$ is $\mathcal{T}$-satisfiable. $\qquad \square$

Example 2. Figure 1 shows a simple example of encoding a graph 3-colorability problem into $\mathcal{LA}(\mathbb{Z})$-solving, using the k-colorer (4) of Example 1. (Notice that the literals which do not contain v_i and $\mathbf{y}_i$ can be moved out of the conjunction $\bigwedge_{V_i \in \mathcal{V}} ...$ in (7).) The first formula is $\mathcal{LA}(\mathbb{Z})$-satisfied, e.g., by an interpretation $\mathcal{I}$ s.t. $\langle c_j \rangle^{\mathcal{I}} \stackrel{\text{def}}{=} j$ for every $j \in [1..3]$, $\langle v_1 \rangle^{\mathcal{I}} \stackrel{\text{def}}{=} 1$, $\langle v_2 \rangle^{\mathcal{I}} \stackrel{\text{def}}{=} 2$, $\langle v_3 \rangle^{\mathcal{I}} \stackrel{\text{def}}{=} 3$ and $\langle v_4 \rangle^{\mathcal{I}} \stackrel{\text{def}}{=} 3$, which mimics the coloring in Fig. 1 (left). The second formula is $\mathcal{LA}(\mathbb{Z})$-unsatisfiable, as expected.   ◇

Lemma 3. *Let* k, n, $\mathcal{G}$, $\mathcal{C}$, $\mathcal{T}$ *and* $\mathsf{Enc}_{[\mathcal{G} \Rightarrow \mathcal{T}]}(v_1, ..., v_n, \underline{\mathbf{c}}|\underline{\mathbf{y}}_1, ..., \underline{\mathbf{y}}_n)$ *be as in Lemma 2. Then* $||\mathsf{Enc}_{[\mathcal{G} \Rightarrow \mathcal{T}]}(v_1, ..., v_n, \underline{\mathbf{c}}|\underline{\mathbf{y}}_1, ..., \underline{\mathbf{y}}_n)||$ *is polynomial in* $||\mathcal{G}|| \stackrel{\text{def}}{=} ||\mathcal{V}|| + ||\mathcal{E}||$.[3]

Proof. By Definition 2 we have that $||\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i)||$ is constant wrt. $||\mathcal{V}||$ or $||\mathcal{E}||$. From (7), (8) and (9), $||\mathsf{Enc}_{[\mathcal{G} \Rightarrow \mathcal{T}]}(v_1, ..., v_n, \underline{\mathbf{c}}|\underline{\mathbf{y}}_1, ..., \underline{\mathbf{y}}_n)||$ is $O(||\mathcal{V}|| + ||\mathcal{E}||)$. □

Combining Lemmas 1, 2 and 3 we have directly the following main result.

Theorem 1. *If a theory with equality* $\mathcal{T}$ *is* k-colorable for some $k \geq 3$, then the problem of deciding the $\mathcal{T}$-satisfiability of a quantifier-free conjunction of $\mathcal{T}$-literals is NP-hard.*

Notice that the key source of hardness is condition (2) in Definition 2: intuitively, a k-colorable theory is expressive enough to represent with a quantifier-free conjunction of $\mathcal{T}$-literals–without disjunctions!–the fact that one variable must assume a value among a choice of $k \geq 3$ possible candidates–in addition to the fact that a list of pairs of variables cannot pairwise assume the same value. This source of non-deterministic choices has a high computational cost, as stated in Theorem 1.

4   Proving k-Colorabilty

Theorem 1 suggests a general technique for proving the NP-hardness of a theory $\mathcal{T}$: pick some $k \geq 3$ and then try to build a k-colorer $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i)$. Also, when $\mathcal{T}$ is known to be NP-hard, one may want to identify smaller –and possibly minimal– signature-restriction fragments $\mathcal{T}'$ which are k-colorable for some k, by identifying increasingly-smaller subsets of the signature of $\mathcal{T}$ which are needed to define a k-colorer.

We introduce some sufficient criteria for a theory to be k-colorable with some $k \geq 3$. As a proof of concept, we use these criteria to prove the k-colorability with some $k \geq 3$, and hence the NP-hardness, of some theories $\mathcal{T}$ of practical interest, and of some of their signature-restriction fragments.

We remark that the ultimate goal here is not to provide fully-detailed proofs of NP-hardness–all the main theories presented here are already well-known to

---

[3] Notice that k is fixed a priori and as such it is a *constant value* for the input graph k-colorability problem: e.g., depending on $\mathcal{T}$, we are speaking of reducing graph 3-colorability–or 4-colorability, or even 2^{64}-colorability–to $\mathcal{T}$-solving.

be NP-hard, although to the best of our knowledge the complexity of not all of their fragments has been investigated explicitly–rather to present proof of concept of the convenience and effectiveness of our proposed colorability-based technique, using various theories/fragments as examples. To this extent, for the sake of simplicity and space needs, and when this does not affect comprehension, sometimes we skip some formal details of the syntax and semantics of the theories under analysis, referring the reader to the proper literature. Rather, we dedicate a few lines to give some hints and tips on how to apply our colorability-based technique in potentially-typical scenarios.

4.1   Exploiting Interpreted Constants, Closed Terms and Provably-Distinct Terms

Proposition 1. *Let $\mathcal{T}$ be a theory which admits at least $k \geq 3$ terms $t_1(\underline{\mathbf{x}}_i), ..., t_k(\underline{\mathbf{x}}_i)$, where $\underline{\mathbf{x}}_i$ are the set of variables which are free in t_j (if any), let $\underline{\mathbf{y}}_i$ being a possibly-empty set of auxiliary variables, and let*

$$\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{x}}_i, \underline{\mathbf{y}}_i) \stackrel{def}{=} \bigwedge\nolimits_{j=1}^{k}(c_j = t_j(\underline{\mathbf{x}}_i)) \wedge \Psi(v_i|\underline{\mathbf{x}}_i, \underline{\mathbf{y}}_i) \tag{10}$$

be a quantifier-free conjunction of literals s.t.

$$\models_{\mathcal{T}} \forall \underline{\mathbf{x}}_i. \ \mathsf{AllDifferent}_k(\{t_1(\underline{\mathbf{x}}_i), ..., t_k(\underline{\mathbf{x}}_i)\}) \tag{11}$$

$$\Psi(v_i|\underline{\mathbf{x}}_i, \underline{\mathbf{y}}_i) \models_{\mathcal{T}} \bigvee\nolimits_{j=1}^{k}(v_i = t_j(\underline{\mathbf{x}}_i)) \tag{12}$$

$$\text{there exist } k \ \mathcal{T}\text{-interpretations } \{\mathcal{I}_{i,1}, ..., \mathcal{I}_{i,k}\} \text{ s.t.} \tag{13}$$
$$\text{for every } j \in [1..k], \ \langle c_j \rangle^{\mathcal{I}_{i,1}} = \langle c_j \rangle^{\mathcal{I}_{i,2}} = ... = \langle c_j \rangle^{\mathcal{I}_{i,k}}, \text{ and}$$
$$\text{for every } j \in [1..k], \ \mathcal{I}_{i,j} \models_{\mathcal{T}} \mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{x}}_i, \underline{\mathbf{y}}_i) \wedge (v_i = t_j(\underline{\mathbf{x}}_i)).$$

Importantly, if $t_1, .., t_k$ are closed terms, then (13) reduces to he following:

$$\text{there exist } k \ \mathcal{T}\text{-interpretations } \{\mathcal{I}_{i,1}, ..., \mathcal{I}_{i,k}\} \text{ s.t.} \tag{14}$$
$$\text{for every } j \in [1..k], \ \mathcal{I}_{i,j} \models_{\mathcal{T}} \mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i) \wedge (v_i = t_j).$$

Then $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{x}}_i, \underline{\mathbf{y}}_i)$ is a k-colorer for $\mathcal{T}$.

Proof. By (11), $\bigwedge_{j=1}^{k}(c_j = t_j(\underline{\mathbf{x}}_i)) \models_{\mathcal{T}} \mathsf{AllDifferent}_k(\underline{\mathbf{c}})$, s.t. (1) holds. By (10) and (12), $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{x}}_i, \underline{\mathbf{y}}_i)$ verifies (2). By (10) and (13) we have that (3) holds. □

Theories of Arithmetic. We use Proposition 1–where $t_1, ..., t_k$ are numerical constants–to prove the k-colorability of (various signature-restriction fragments of) the theories of arithmetic.

Example 3 ($\mathcal{A}^{\{\geq,=\}}(\mathbb{Z})$, $\mathcal{LA}(\mathbb{Z})$, $\mathcal{NLA}(\mathbb{Z})$). Let $\mathcal{A}^{\{\geq,=\}}(\mathbb{Z})$ be the basic theory of integers under successor [20,21], that is, whose atoms are in the form $(s_1 \odot s_2)$, where $\odot \in \{\geq,=\}$ and s_1, s_2 are variables or positive numerical constants.

Then $\mathcal{A}^{\{\geq,=\}}(\mathbb{Z})$ is 3-colorable, because we can define a 3-colorer like that of (4) in Example 1. (Notice that this is an instance of Proposition 1.) $\mathcal{A}^{\{\geq,=\}}(\mathbb{Z})$ is a signature-restriction fragment of $\mathcal{LA}(\mathbb{Z})$ and $\mathcal{NLA}(\mathbb{Z})$ (see e.g. [24]), which are then 3-colorable by Proposition 2. Therefore, $\mathcal{T}$-solving for all these theories is NP-hard by Theorem 1.[4]   ◇

Notice that conjunctions of only *positive* equalities and inequalities in the form $(s_1 \odot s_2)$, without negated literals, are instead well-known to be solvable in polynomial time (see e.g. [2,18]). Notice also that, on the rational domain, the corresponding theories $\mathcal{A}^{\{\geq,=\}}(\mathbb{Q})$ and $\mathcal{LA}(\mathbb{Q})$ are convex and hence they are not colorable by Property 1. In fact, $\mathcal{T}$-solving for such theories is notoriously in P [10].

Example 4 ($\mathcal{NLA}(\mathbb{R})^{\backslash\{\geq,>\}}$, $\mathcal{NLA}(\mathbb{R})$). We consider $\mathcal{NLA}(\mathbb{R})^{\backslash\{\geq,>\}}$, the signature-restriction fragment of the non-linear arithmetic over the reals ($\mathcal{NLA}(\mathbb{R})$) without inequality symbols $\{\geq,\leq\}$. As an instance of Proposition 1, we show that $\mathcal{NLA}(\mathbb{R})^{\backslash\{\geq,>\}}$ is 3-colorable, because we can define, e.g., $k \stackrel{\text{def}}{=} 3$, $\underline{\mathbf{y}} \stackrel{\text{def}}{=} \emptyset$, and

$$\mathsf{Colorer}_3(v_i, c_1, c_2, c_3) \stackrel{\text{def}}{=} \left(\begin{array}{c} (c_1 = -1) \wedge (c_2 = 0) \wedge (c_3 = 1) \wedge \\ (v_i \cdot (v_i - 1) \cdot (v_i + 1) = 0) \end{array} \right).$$

By Proposition 1, it is straightforward to see that $\mathsf{Colorer}_3(v_i, c_1, c_2, c_3)$ verifies (1), (2) and (3), with $\langle c_1 \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} -1, \langle c_2 \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} 0, \langle c_3 \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} 1$, and $\langle v_i \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} \langle c_j \rangle^{\mathcal{I}_{i,j}}$ s.t. $j \in [1..3]$. Then by Proposition 2 the full $\mathcal{NLA}(\mathbb{R})$ is 3-colorable, so that $\mathcal{T}$-solving for both theories is NP-hard by Theorem 1.   ◇

4.2   Exploiting Finite Domains of Fixed Size

Proposition 2. *Let $\mathcal{T}$ be some theory with finite domain of fixed size $k \geq 3$. Then $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}) \stackrel{\text{def}}{=} \mathsf{AllDifferent}_k(\underline{\mathbf{c}})$ is a k-colorer for $\mathcal{T}$.*

Proof. Let $\underline{\mathbf{c}} \stackrel{\text{def}}{=} \{c_1, ..., c_k\}$. Since the domain of $\mathcal{T}$ has fixed size $k \geq 3$, we have:

$$\mathsf{AllDifferent}_k(\underline{\mathbf{c}}) \not\models_{\mathcal{T}} \bot \tag{15}$$

$$\mathsf{AllDifferent}_{k+1}(\underline{\mathbf{c}} \cup \{v_i\}) \models_{\mathcal{T}} \bot. \tag{16}$$

$\mathsf{AllDifferent}_k(\underline{\mathbf{c}})$ entails itself, so that (1) holds. $\mathsf{AllDifferent}_k(\underline{\mathbf{c}}) \wedge \bigwedge_{j=1}^{k} \neg(v_i = c_j)$ is the same as $\mathsf{AllDifferent}_{k+1}(\underline{\mathbf{c}} \cup \{v_i\})$ which is $\mathcal{T}$-unsatisfiable by (16), so that $\mathsf{AllDifferent}_k(\underline{\mathbf{c}}) \models_{\mathcal{T}} \bigvee_{j=1}^{k}(v_i = c_j)$. Hence (2) holds. By (15) there exists some $\mathcal{T}$-interpretation $\mathcal{I}$ s.t. $\mathcal{I} \models_{\mathcal{T}} \mathsf{AllDifferent}_k(\underline{\mathbf{c}})$. For every $j \in [1..k]$ we build an extension $\mathcal{I}_{i,j}$ of $\mathcal{I}$ with the same domain s.t. $\langle c_1 \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} \langle c_1 \rangle^{\mathcal{I}}, ..., \langle c_k \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} \langle c_k \rangle^{\mathcal{I}}$, and $\langle v_i \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} \langle c_j \rangle^{\mathcal{I}}$. Hence (3) holds.   □

---

[4] Notice that $\mathcal{NLA}(\mathbb{Z})$-solving is undecidable.

Theories of Fixed-Width Bit-Vectors and Floating-Point Arithmetic.
We prove the k-colorability of (the signature-restriction fragments of) the theories of Fixed-width Bit-vectors and Floating-point Arithmetic by instantiating Proposition 2.

Example 5 ($\mathcal{BV}_w$, $w>1$). Let $\mathcal{BV}_w^{\{=\}}$ be the simplest possible signature-restriction fragment of the fixed-width bit-vectors theory with equality $=$ and width $w>1$, with no interpreted constant, function or predicate symbol in its signature. Then by Proposition 2, $\mathcal{BV}_w^{\{=\}}$ is k-colorable, where $k = 2^w$. Hence, by Property 2 all theories $\mathcal{BV}_w^*$ obtained by augmenting the signature of $\mathcal{BV}_w^{\{=\}}$ with various combinations of interpreted constants (e.g. $\mathsf{bv_w\_0...00}$, $\mathsf{bv_w\_0...01}$,...), functions (e.g. $\mathsf{bv_w\_and}$, $\mathsf{bv_w\_or}$,...) and predicates (e.g. $\mathsf{bv_w\_{\geq}}$,...)–are k-colorable with $k = 2^w$. Hence, when $w>1$, by Theorem 1, $\mathcal{T}$-solving is NP-hard for all such theories.   ◇

[7] shows that the $\mathcal{T}$-satisfiability of quantifier-free conjunctions of atoms for the fragment of $\mathcal{BV}$ involving only concatenation and partition of words is in P. Notice however that neither Example 5 contradicts the results in [7], nor Example 5 plus [7] build a proof of $P = NP$, because the polynomial procedure in [7] does not admit *negative equalities* $\neg(v_i = v_i')$ in the conjunction.

Example 6 ($\mathcal{FPA}_{e,s}$). Let $\mathcal{FPA}_{e,s}$ be the theory of floating-point arithmetic s.t. $e \geq 1$ and $s \geq 1$ are the number of available bits for the exponent and the significant respectively [22]. (E.g., $\mathcal{FPA}_{11,53}$ represents the binary64 format of IEEE 754-2008 [22].) As with Example 5, let $\mathcal{FPA}_{e,s}^=$ be the simplest possible signature-restriction fragment of $\mathcal{FPA}_{e,s}$ with equality $=$,[5] with no interpreted constant, function and predicate symbol in its signature. Then by Proposition 2, $\mathcal{FPA}_{e,s}^=$ is k-colorable, where $k = 2^{e+s}$. Hence, by Property 2, all theories $\mathcal{FPA}_{e,s}^*$ obtained by augmenting the signature of $\mathcal{FPA}_{e,s}^=$ with various combinations of interpreted constants, functions or predicates are k-colorable with $k \geq 4$, so that $\mathcal{T}$-solving is NP-hard.   ◇

4.3   Dealing with Collection Datatypes

A class of theories of big interest in SMT-based formal verification are these describing collection datatypes (see e.g. [6,12])–e.g., lists, arrays, sets, etc. In general these are "families" of theories, each being a combination of a "basic" theory (e.g., the basic theory of lists) with one or more theories describing the elements or the indexes of the datatype. In what follows we consider the basic theories, where elements are represented by generic variables representing values in some infinite domain.

One potential problem if finding k-colorers for most of these "basic" theories is that neither we have interpreted constants in the domain of the elements, so that we cannot apply Proposition 1 as we did with arithmetical theories, nor

---

[5] Here "$=$" is the equality symbol and it is not the $\mathcal{FPA}_{e,s}$-specific symbol "$==$", see [22].

we have any information on the size of the domain of the elements, so that we cannot apply Proposition 2.

We analyze different potential scenarios. One first scenario is where we have at least one "structural" interpreted constant–e.g., that representing the empty collection–plus some function symbols, which we can use to build $k \geq 3$ closed terms $t_1, ..., t_k$ and then use the schema of Proposition 1 to build a k-colorer.

Theories of Lists. The above scenario is illustrated in the next example.

Example 7 ($\mathcal{L}^+$). Let $\mathcal{L}$ be the simplest theory of lists of generic elements, with the signature $\Sigma \stackrel{\text{def}}{=} \{\mathsf{nil}, \mathsf{car}(\cdot), \mathsf{cdr}(\cdot), \mathsf{cons}(\cdot, \cdot)\}$ and described by the axioms:

$$\forall xy.(\mathsf{car}(\mathsf{cons}(x, y)) = x)), \ \forall xy.(\mathsf{cdr}(\mathsf{cons}(x, y)) = y)), \tag{17}$$
$$\forall xy.(\neg(\mathsf{cons}(x, y) = \mathsf{nil})), \ \forall x.(\neg(x = \mathsf{nil}) \rightarrow (\mathsf{cons}(\mathsf{car}(x), \mathsf{cdr}(x)) = x)), \tag{18}$$

and let $\mathcal{L}^+$ be $\mathcal{L}$ enriched by the axioms

$$(\mathsf{car}(\mathsf{nil}) = \mathsf{nil}), \ \ (\mathsf{cdr}(\mathsf{nil}) = \mathsf{nil}). \tag{19}$$

$\mathcal{L}^+$-solving is NP-complete whilst $\mathcal{L}$-solving is in P [17]. A more general theory of lists, which has $\mathcal{L}^+$ as a signature-restriction fragment, is described in [6,12]. Following Proposition 1, we prove that $\mathcal{L}^+$ is 4-colorable, by setting $k \stackrel{\text{def}}{=} 4$, $\underline{\mathbf{y}} \stackrel{\text{def}}{=} \{x_1, x_2, y_1, y_2\}$,

$$\mathsf{Colorer}_4(v_i, c_{11}, c_{21}, c_{12}, c_{22} | x_1, x_2, y_1, y_2) \stackrel{\text{def}}{=} \tag{20}$$
$$\left(\begin{array}{l} (c_{11} = \mathsf{cons}(\mathsf{nil}, \mathsf{nil})) \wedge (c_{21} = \mathsf{cons}(\mathsf{cons}(\mathsf{nil}, \mathsf{nil}), \mathsf{nil})) \wedge \\ (c_{12} = \mathsf{cons}(\mathsf{nil}, \mathsf{cons}(\mathsf{nil}, \mathsf{nil}))) \wedge (c_{22} = \mathsf{cons}(\mathsf{cons}(\mathsf{nil}, \mathsf{nil}), \mathsf{cons}(\mathsf{nil}, \mathsf{nil}))) \wedge \\ \bigwedge_{i=1}^{2} ((\mathsf{car}(x_i) = \mathsf{car}(y_i)) \wedge (\mathsf{cdr}(x_i) = \mathsf{cdr}(y_i)) \wedge \neg(x_i = y_i)) \wedge \\ (v_i = \mathsf{cons}(x_1, x_2)). \end{array} \right)$$

To prove (11) we notice that we can deduce $\neg(\mathsf{cons}(\mathsf{nil}, \mathsf{nil}) = \mathsf{nil})$ from (18), so that, by construction, all the c_i's are pairwise different. Let $\Psi(v_i \underline{\mathbf{y}}_i)$ be the formula given by the last two rows in (20), so that (20) matches the definition in Proposition 1. Then we derive (12) from the following observation [17], with $i \in \{1, 2\}$:

$$((\mathsf{car}(x_i) = \mathsf{car}(y_i)) \wedge (\mathsf{cdr}(x_i) = \mathsf{cdr}(y_i)) \wedge \neg(x_i = y_i)) \tag{21}$$
$$\models_{\mathcal{L}^+} (x_i = \mathsf{nil}) \vee (x_i = \mathsf{cons}(\mathsf{nil}, \mathsf{nil})),$$

which derives from the fact that (18) and (19) imply that either $(x_i = \mathsf{nil})$ or $(y_i = \mathsf{nil})$ must hold. Therefore $v_i \stackrel{\text{def}}{=} \mathsf{cons}(x_1, x_2)$ can consistently assume one and only one of the values $c_{11}, ..., c_{22}$ in the first two rows in (20).

To prove (14), since the c_is are closed, we deterministically define each $\mathcal{I}_{i,j}$'s using the standard interpretation of $\mathsf{nil}, \mathsf{cons}, \mathsf{car}$, and cdr: $\langle c_{11} \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} (\mathsf{NIL.NIL})$, $\langle c_{21} \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} ((\mathsf{NIL.NIL}).\mathsf{NIL}), ... \langle v_i \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} \langle c_j \rangle^{\mathcal{I}_{i,j}}$, checking that, for every $j \in [1..k]$,

$$\mathcal{I}_{i,j} \models_{\mathcal{L}^+} \mathsf{Colorer}_4(v_i, c_{11}, c_{21}, c_{12}, c_{22} | x_1, x_2, y_1, y_2) \wedge (v_i = c_j).$$

Thus $\mathcal{L}^+$-solving is NP-hard by Theorem 1, so that also the more general theory described in [6,12] is NP-hard.   ◇

Remark 3. The k-colorer (20) was produced along the following heuristic process.

1. Look for an entailment in the form: $\mu_1(x_1, \underline{\mathbf{y}}_1) \models_{\mathcal{T}} (x_1 = t_1) \vee (x_1 = t_2)$, s.t. t_1, t_2 are closed terms representing distinct values in the domain (21).
2. Define $(v_i = \mathsf{cons}(x_1, x_2))$ and $(c_{r_1 r_2} = \mathsf{cons}(t_{r_1}, t_{r_2}))$, s.t. $r_1, r_2 \in \{1, 2\}$
3. Define the k-colorer as

$$\bigwedge\nolimits_{i \in \{1,2\}} \mu_i(x_i, \underline{\mathbf{y}}_i) \wedge \bigwedge\nolimits_{r_1, r_2 \in \{1,2\}} (c_{r_1 r_2} = \mathsf{cons}(t_{r_1}, t_{r_2})) \wedge (v_i = \mathsf{cons}(x_1, x_2)).$$

4. Check (11), (12), (14).

Notice that the only non-obvious step is 1, the other come out nearly deterministically.

Theories of Finite Sets. Another scenario is where we cannot use interpreted constants to build closed terms, but we can build k non-closed terms $t_1(\mathbf{x}_i), ..., t_k(\mathbf{x}_i)$ which match the requirements of Proposition 1 anyway, which allows to build a k-colorer. This scenario is illustrated in the next example.

Example 8. Let $\mathcal{S}$ be the theory of finite sets as defined, e.g., in [6,12].[6] Let $\mathcal{S}^{\{\subseteq, \{\}\}}$ be the signature-restriction fragment of the $\mathcal{S}$ which considers only the subset and the enumerator operators $\{\subseteq, \{\}\}$. We show that $\mathcal{S}^{\{\subseteq, \{\}\}}$ is 4-colorable by Proposition 1.

In fact, consider the following set of literals:

$$\mathsf{Colorer}_4(v_i, \underline{\mathbf{c}} | y_1, y_2) \stackrel{\text{def}}{=} \begin{pmatrix} (c_1 = \{y_1, y_2\}) \wedge (c_2 = \{y_1\}) \wedge \\ (c_3 = \{y_2\}) \qquad \wedge (c_4 = \{\}) \qquad \wedge \\ \neg(y_1 = y_2) \qquad \wedge (v_i \subseteq c_1) \end{pmatrix}. \tag{22}$$

(22) is a 4-colorer. It is easy to see from the semantics of $\{\subseteq, \{\}\}$ that (11) and (12) hold. Let Y_1, Y_2 s.t. $Y_1 \neq Y_2$ be two domain elements so that we can set $\langle y_r \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} Y_r$ for every $r \in [1..2]$ and $j \in [1..k]$. Then, for every $j \in [1..k]$, we define $\mathcal{I}_{i,j}$ s.t. $\langle c_1 \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} \{Y_1, Y_2\}$, $\langle c_2 \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} \{Y_1\}$, $\langle c_3 \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} \{Y_2\}$, $\langle c_4 \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} \{\}$, $\langle v_i \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} \langle c_j \rangle^{\mathcal{I}_{i,j}}$. Then $\mathcal{I}_{i,1}, ..., \mathcal{I}_{i,k}$ verify (13).   ◇

In this case the k-colorer (22) was really immediate to build, upon the observation that the operator $\subseteq$ can produce 4 distinct subsets of a 2-element set.

---

[6] $\mathcal{S}$ includes the operators $\{\{...\}), (\cdot \subseteq \cdot), (\cdot \cup \cdot), (\cdot \cap \cdot), (\cdot \setminus \cdot), (\cdot \mathcal{P} \cdot), | \cdot |, (\cdot \in \cdot)\}$, following their standard semantics. We refer the reader to [6,12] for a precise description of the theory.

Theories of Arrays. In the following case we cannot apply Proposition 1, so that we apply Definition 2 directly.

Example 9. ($\mathcal{AR}$). Let $\mathcal{AR}$ be the theory of arrays of generic elements and indexes, with the signature $\Sigma \stackrel{\text{def}}{=} \{\cdot[\cdot], \cdot\langle\cdot \leftarrow \cdot\rangle\}$ [7] and described by the axioms:

$$\forall Aijv. \ ((i = j) \rightarrow (A\langle i \leftarrow v\rangle[j] = v), \tag{23}$$

$$\forall Aijv. \ (\neg(i = j) \rightarrow (A\langle i \leftarrow v\rangle[j] = A[j]), \tag{24}$$

$$\forall AB. \ ((\forall i. \ A[i] = B[i]) \rightarrow (A = B)). \tag{25}$$

$\mathcal{AR}$ is 3-colorable, because we can define, e.g., $k \stackrel{\text{def}}{=} 3$, $\underline{\mathbf{y}} \stackrel{\text{def}}{=} \{A_1, ..., A_4, i_1, ..., i_3\}$ and

$$\mathsf{Colorer}_3(v_i, c_1, c_2, c_3 | A_1, ..., A_4, i_1, ..., i_3) \stackrel{\text{def}}{=} \begin{pmatrix} \mathsf{AllDifferent}_3(\underline{\mathbf{c}}) & \wedge \\ \neg(i_2 = i_3) & \wedge \\ (A_2 = A_1\langle i_1 \leftarrow c_1\rangle) \wedge \\ (A_3 = A_2\langle i_2 \leftarrow c_2\rangle) \wedge \\ (A_4 = A_3\langle i_3 \leftarrow c_3\rangle) \wedge \\ (v_i = A_4[i_1]) \end{pmatrix} \tag{26}$$

so that obviously (1) holds, and also (2) holds, because $\mathsf{Colorer}_3(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}})$ entails $(v_i = c_1)$ when $\langle i_1\rangle^{\mathcal{I}} \neq \langle i_3\rangle^{\mathcal{I}}$ and $\langle i_1\rangle^{\mathcal{I}} \neq \langle i_2\rangle^{\mathcal{I}}$, entails $(v_i = c_2)$ when $\langle i_1\rangle^{\mathcal{I}} = \langle i_2\rangle^{\mathcal{I}}$, and entails $(v_i = c_3)$ when $\langle i_1\rangle^{\mathcal{I}} = \langle i_3\rangle^{\mathcal{I}}$. Also (3) holds: given three distinct domain values C_1, C_2, C_3, the $\mathcal{T}$-interpretations $\mathcal{I}_{i,j}$ can be built straightforwardly as follows:

	c_1	c_2	c_3	v_i	i_1	i_2	i_3	A_4	...
$\mathcal{I}_{i,1}$	C_1	C_2	C_3	C_1	1	2	3	$[C_1, C_2, C_3, ...]$	
$\mathcal{I}_{i,2}$	C_1	C_2	C_3	C_2	2	2	3	$[**, C_2, C_3, ...]$	
$\mathcal{I}_{i,3}$	C_1	C_2	C_3	C_3	3	2	3	$[**, C_2, C_3, ...]$	

$\diamond$

Notice that in Example 9, $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i)$ uses the auxiliary variables $A_1, ..., A_4$ representing arrays and $i_1, ..., i_3$ representing indexes. The A_2, A_3, A_4, however, are not strictly necessary and can be eliminated by inlining. Notice also that $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i)$ includes explicitly $\mathsf{AllDifferent}_3(\underline{\mathbf{c}})$ because no interpreted constants come into play.

The k-colorer (26) was produced straightforwardly by noticing that the combination of (23) and (24) produces a case-split in the form "*if $i = j$ then* $(A\langle i \leftarrow v\rangle[j] = v)$ *else* $(A\langle i \leftarrow v\rangle[j] = A[j])$", which could be reiterated so that to produce a 3-branch decision tree, producing 3 different expressions for the term $A[i_1]$. This could be rewritten into k-colorer by means of some term renaming.

---

[7] We use the following notation: "$A[i]$" (aka "$\mathsf{read}(A, i)$") is the value returned by reading the i-th element of the array A, whilst "$A\langle i \leftarrow v_i\rangle$" (aka "$\mathsf{write}(A, i, v)$") is the array resulting from assigning the value v to the i-th element of array A.

5 k-Colorability vs. Non-Convexity

Although related by Property 1, k-colorability and non-convexity are distinct properties. First, we recall that the non-convexity of a theory $\mathcal{T}$ does not imply the NP-hardness of $\mathcal{T}$-solving. (In [24] we report a simple example.) Second, by Property 1, having domain size ≥ 3 is a strict requirement for proving NP-hardness via colorability, whereas there exist non-convex theories with domain size 2 whose $\mathcal{T}$-solving is NP-Hard. (E.g., the theory $\mathcal{BV}_1$ of bit vectors with fixed width 1, see [24].)

In what follows we introduce a theory with domain size ≥ 3 whose $\mathcal{T}$-solving is NP-hard, which is non-convex and which is not k-colorable for any $k \geq 3$. This shows that not every theory with domain size ≥ 3 can be proven NP-hard by k-colorability. The same example shows also that k-colorability is strictly stronger than non-convexity, even when the theory has domain size ≥ 3.

Example 10. Consider the theory $\mathcal{T}$ with equality whose signature consists in the interpreted constant symbols $\{0, 1, 2, ...\}$ with the standard meaning plus the function symbols $\{\mathsf{and}(\cdot, \cdot), \mathsf{not}(\cdot)\}$ which are interpreted as follows:

$$\langle \mathsf{and}(x, y) \rangle^{\mathcal{I}} \stackrel{\text{def}}{=} \begin{cases} 1 \text{ if } \langle x \rangle^{\mathcal{I}} > 0 \text{ and } \langle y \rangle^{\mathcal{I}} > 0 \\ 0 \text{ otherwise} \end{cases} , \quad \langle \mathsf{not}(x) \rangle^{\mathcal{I}} \stackrel{\text{def}}{=} \begin{cases} 0 \text{ if } \langle x \rangle^{\mathcal{I}} > 0 \\ 1 \text{ otherwise.} \end{cases} \quad (27)$$

(Importantly, the $\geq, >, \leq, <$ predicates are not part of the signature.) $\mathcal{T}$-satisfiability is NP-complete since you can polynomially reduce SAT to it and you can always have a polynomial-size witness for every $\mathcal{T}$-satisfied formula.

Also, as with $\mathcal{BV}_1$, $\mathcal{T}$ is non-convex, because we have:

$$(x_0 = 0) \wedge (\mathsf{and}(x_1, x_2) = 0) \models_{\mathcal{T}} (((x_0 = x_1) \vee (x_0 = x_2)) \qquad (28)$$
$$(x_0 = 0) \wedge (\mathsf{and}(x_1, x_2) = 0) \not\models_{\mathcal{T}} (x_0 = x_i) \ \ i \in \{1, 2\}. \qquad (29)$$

We show that $\mathcal{T}$ is not k-colorable for any $k \geq 3$. We notice that every literal l including v_i must be in one of the following forms (modulo the symmetry of $=$ and and): $(v_i = t)$, $(v_i = \mathsf{not}(t))$, $(v_i = \mathsf{and}(t_1, t_2))$, $(t = t^*(v_i, ...))$, and their negations, where t, t_1, t_2 are generic terms in $\mathcal{T}$ and $t^*(v_i, ...)$ is any term in $\mathcal{T}$ containing v_i. Looking at the above literal forms, we notice that the presence of the subterms $\mathsf{not}(v_i)$ and $\mathsf{and}(v_i, t_2)$ in a term entails either $\langle v_i \rangle^{\mathcal{I}} > \langle 0 \rangle^{\mathcal{I}}$, or $\langle v_i \rangle^{\mathcal{I}} = \langle 0 \rangle^{\mathcal{I}}$ or $\langle v_i \rangle^{\mathcal{I}} \geq \langle 0 \rangle^{\mathcal{I}}$, so that one single literal l can express only the following facts about one variable v_i:[8]

(i) for every $\mathcal{T}$-interpretation $\mathcal{I}$ s.t. $\mathcal{I} \models_{\mathcal{T}} l$, $\langle v_i \rangle^{\mathcal{I}} = \langle n \rangle^{\mathcal{I}}$ for some $n \in \{0, 1, 2, 3, ...\}$;

(ii) for every $\mathcal{T}$-interpretation $\mathcal{I}$ s.t. $\mathcal{I} \models_{\mathcal{T}} l$, $\langle v_i \rangle^{\mathcal{I}} \neq \langle n \rangle^{\mathcal{I}}$ for some $n \in \{0, 1, 2, 3, ...\}$;

---

[8] Whereas (i) and (ii) can be also written as $l \models_{\mathcal{T}} (v_i = n)$ and $l \models_{\mathcal{T}} (v_i \neq n)$, (iii) and (iv) cannot be rewritten as $l \models_{\mathcal{T}} (v_i \geq 0)$ and $l \models_{\mathcal{T}} (v_i > 0)$ because $\geq$ and $>$ are not part of the signature.

(iii) for every $\mathcal{T}$-interpretation $\mathcal{I}$ s.t. $\mathcal{I} \models_{\mathcal{T}} l$, $\langle v_i \rangle^{\mathcal{I}} \geq \langle 0 \rangle^{\mathcal{I}}$ (equivalent to true);
(iv) for every $\mathcal{T}$-interpretation $\mathcal{I}$ s.t. $\mathcal{I} \models_{\mathcal{T}} l$, $\langle v_i \rangle^{\mathcal{I}} > \langle 0 \rangle^{\mathcal{I}}$ (equivalent to $\langle v_i \rangle^{\mathcal{I}} \neq 0$);
 (v) for every $\mathcal{T}$-interpretation $\mathcal{I}$ s.t. $\mathcal{I} \models_{\mathcal{T}} l$, $\langle v_i \rangle^{\mathcal{I}} = \langle v_i \rangle^{\mathcal{I}}$ (equivalent to true);
(vi) for every $\mathcal{T}$-interpretation $\mathcal{I}$ s.t. $\mathcal{I} \models_{\mathcal{T}} l$, $\langle v_i \rangle^{\mathcal{I}} \neq \langle v_i \rangle^{\mathcal{I}}$ (equivalent to false).

Thus, for $k \geq 3$, no finite conjunction of $\mathcal{T}$-literals $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i)$ complying with (1) and (3) can also comply with (2). $\diamond$

6   Colorable Theories without Equality

In previous sections we have restricted our interest to theories *with equality*. In this section we extend the technique by dropping this restriction. The following definition extends Definition 2 to the case of general theories.

Definition 3 (*k*-Colorer, *k*-Colored Theory). *Let $\mathcal{T}$ be some theory and k be some integer value s.t. $k \geq 2$. Let v_i be a variable, called* **vertex variable***, (implicitly) denoting the i-th vertex in an un-directed graph; let $\underline{\mathbf{c}} \overset{def}{=} \{c_1, .., c_k\}$ be a set of variables, called* **color variables***, denoting the set of colors; let $\underline{\mathbf{y}}_i \overset{def}{=} \{y_{i1}, ..., y_{il}\}$ denote a possibly-empty set of variables, which is indexed with the same index i of the vertex variable v_i. We call k-**colorer** for $\mathcal{T}$, namely $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i)$, a finite quantifier-free conjunction of $\mathcal{T}$-literals (cube) over v_i, $\underline{\mathbf{c}}$ and $\underline{\mathbf{y}}_i$ which verify the following properties:*

– *For every $\mathcal{T}$-intepretation $\mathcal{I}$, if $\mathcal{I} \models_{\mathcal{T}} \mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i)$, then:*

$$\text{for every } j, j' \in [1..k] \text{ s.t. } j \neq j', \qquad \langle c_j \rangle^{\mathcal{I}} \neq \langle c_{j'} \rangle^{\mathcal{I}}, \tag{30}$$

$$\text{for some } j \in [1..k], \qquad \langle v \rangle^{\mathcal{I}} = \langle c_j \rangle^{\mathcal{I}}, \tag{31}$$

– *There exist k $\mathcal{T}$-interpretations $\{\mathcal{I}_{i,1}, ..., \mathcal{I}_{i,k}\}$ s.t.*

$$\text{for every } j \in [1..k], \qquad \langle c_j \rangle^{\mathcal{I}_{i,1}} = \langle c_j \rangle^{\mathcal{I}_{i,2}} = ... = \langle c_j \rangle^{\mathcal{I}_{i,k}}, \text{ and} \tag{32}$$

$$\text{for every } j \in [1..k], \qquad \begin{cases} \langle v \rangle^{\mathcal{I}_{i,j}} = \langle c_j \rangle^{\mathcal{I}_{i,j}} \text{ and} \\ \mathcal{I}_{i,j} \models_{\mathcal{T}} \mathsf{Colorer}_k(v_i, \underline{\mathbf{c}}|\underline{\mathbf{y}}_i). \end{cases}$$

*We say that $\mathcal{T}$ is k-**colorable** iff it has a k-colorer.*

Notice that if $\mathcal{T}$ is a theory with equality, then Definitions 2 and 3 are equivalent.

Definition 4. *We say that a theory $\mathcal{T}$ **emulates equality** [resp. **disequality**] if and only if there exists a finite quantifier-free conjunction of $\mathcal{T}$-literals $\mathsf{Eq}(x_1, x_2)$ [resp. $\mathsf{Neq}(x_1, x_2)$] such that, for every $\mathcal{T}$-interpretation $\mathcal{I}$, $\mathcal{I} \models_{\mathcal{T}} \mathsf{Eq}(x_1, x_2)$ [resp. $\mathcal{I} \models_{\mathcal{T}} \mathsf{Neq}(x_1, x_2)$] if and only if $\langle x_1 \rangle^{\mathcal{I}} = \langle x_2 \rangle^{\mathcal{I}}$ [resp. $\langle x_1 \rangle^{\mathcal{I}} \neq \langle x_2 \rangle^{\mathcal{I}}$].*

Obviously every theory $\mathcal{T}$ with equality emulates both equality and disequality, with $\mathsf{Eq}(x_1, x_2) \overset{def}{=} (x_1 = x_2)$ and $\mathsf{Neq}(x_1, x_2) \overset{def}{=} \neg(x_1 = x_2)$.

Theorem 2. *If a theory $\mathcal{T}$ is k-colorable for some $k \geq 3$ and $\mathcal{T}$ emulates equality and inequality, then the problem of deciding the $\mathcal{T}$-satisfiability of a finite conjunction of quantifier-free $\mathcal{T}$-literals is $\mathcal{T}$-satisfiable is NP-hard.*

Proof. Identical to that of Theorem 1, referring to Definition 3 instead of Definition 2 and substituting every positive equality in the form $(x_1 = x_2)$ with $\mathsf{Eq}(x_1, x_2)$ and every negative equality in the form $\neg(x_1 = x_2)$ with $\mathsf{Neq}(x_1, x_2)$. □

Example 11. Let $\mathcal{NLA}(\mathbb{R})^{\backslash\{=\}}$ be the signature-restriction fragment of $\mathcal{NLA}(\mathbb{R})$ without equality. We notice that $\mathcal{NLA}(\mathbb{R})^{\backslash\{=\}}$ emulates both equality and inequality:

$$\mathsf{Eq}(x_1, x_2) \stackrel{\text{def}}{=} (x_1 \geq x_2) \wedge (x_2 \geq x_1) \tag{33}$$

$$\mathsf{Neq}(x_1, x_2) \stackrel{\text{def}}{=} ((x_1 - x_2) * (x_1 - x_2) > 0). \tag{34}$$

$\mathcal{T}$ is 3-colorable because, like in Example 4, we can define, e.g., $k \stackrel{\text{def}}{=} 3$, $\underline{\mathbf{y}} \stackrel{\text{def}}{=} \emptyset$, and

$$\mathsf{Colorer}_3(v_i, c_1, c_2, c_3) \stackrel{\text{def}}{=} \mathsf{Eq}(c_1, -1) \wedge \mathsf{Eq}(c_2, 0) \wedge \mathsf{Eq}(c_3, 1) \wedge \mathsf{Eq}(v_1 * (v_2 - 1) * (v_1 + 1), 0).$$

Like in Example 4, it is straighforward to see that $\mathsf{Colorer}_3(v, c_1, c_2, c_3)$ verifies (30), (31) and (32), with $\langle c_1 \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} -1, \langle c_2 \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} 0, \langle c_3 \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} 1$, and $\langle v_i \rangle^{\mathcal{I}_{i,j}} \stackrel{\text{def}}{=} \langle c_j \rangle^{\mathcal{I}_{i,j}}$ for every $j \in [1..3]$. Thus $\mathcal{NLA}(\mathbb{R})^{\backslash\{=\}}$-solving is NP-hard by Theorem 2.   ◇

7   Open Issues, Ongoing and Future Work

We believe that our framework can be generalized along the following directions, which we are currently working on: (i) adopt some more general notion of fragment, so that to extend the range of applicability of Property 2; (ii) extend the applicability of our technique for the case of theories without equality by providing a more general definition of $\mathsf{Eq}(.,.)$ and $\mathsf{Neq}(.,.)$ enriched with auxiliary variables –or uninterpreted function/predicate symbols– adapting Theorem 2 accordingly; (iii) extend $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}} | \underline{\mathbf{y}}_i)$ so that to use also uninterpreted function/predicate symbols as auxiliary symbols $\underline{\mathbf{y}}_i$; (iv) to overcome the restriction of domain size ≥ 3, extend $\mathsf{Colorer}_k(v_i, \underline{\mathbf{c}} | \underline{\mathbf{y}}_i)$ to use pairs of variables $\underline{\mathbf{v}}_i \ \underline{\mathbf{c}}_1, .., \underline{\mathbf{c}}_k$ instead of single variables to encode vertexes and colors, including ad hoc $\mathsf{Neq}(.,.)$ functions.

The above work should be run in parallel and interleaved with an extensive exploration of the pool of available NP-hard theories, proving the k-colorability of as many theories/fragments as possible. To this extent, we would like to investigate the boundary of k-colorability, looking for theories of domain size ≥ 3 which are not k-colorable.

References

1. Barrett, C.W., Nieuwenhuis, R., Oliveras, A., Tinelli, C.: Splitting on demand in SAT modulo theories. In: Hermann, M., Voronkov, A. (eds.) LPAR 2006. LNCS (LNAI), vol. 4246, pp. 512–526. Springer, Heidelberg (2006)
2. Barrett, C., Sebastiani, R., Seshia, S.A., Tinelli, C.: Satisfiability Modulo Theories, chapter 26, pp. 825–885. IOS Press, Amsterdam (2009)

3. Berman, L.: The complexity of logical theories. Theor. Comput. Sci. **11**, 71 (1980)
4. Bozzano, M., Bruttomesso, R., Cimatti, A., Junttila, T.A., van Rossum, P., Schulz, S., Sebastiani, R.: Mathsat: tight integration of sat and mathematical decision procedures. J. Autom. Reasoning **35**(1–3), 265–293 (2005)
5. Bradley, A., Manna, Z.: The Calculus of Computation: Decision Procedures with Applications to Verification. Springer, Heidelberg (2010)
6. Bruun, H., Damm, F., Dawes, J., Hansen, B., Larsen, P., Parkin, G., Plat, N., Toetenel, H.: A formal definition of vdm-sl. Technical report, University of Leicester (1998)
7. Cyrluk, D., Möller, M.O., Rueß, H.: An efficient decision procedure for the theory of fixed-sized bit-vectors. In: Grumberg, O. (ed.) Computer Aided Verification. LNCS, vol. 1254, pp. 60–71. Springer, Heidelberg (1997)
8. Fröhlich, A., Kovásznai, G., Biere, A.: More on the complexity of quantifier-free fixed-size bit-vector logics with binary encoding. In: Bulatov, A.A., Shur, A.M. (eds.) CSR 2013. LNCS, vol. 7913, pp. 378–390. Springer, Heidelberg (2013)
9. Garey, M.R., Johnson, D.S.: Computers and Intractability. Freeman and Company, New York (1979)
10. Karmakar, N.: A new polynomial-time algorithm for linear programming. Combinatorica **4**(4), 373 (1984)
11. Kovásznai, G., Fröhlich, A., Biere, A.: On the complexity of fixed-size bit-vector logics with binary encoded bit-width. In: Proceedings of the SMT (2012)
12. Kroening, D., Rümmer, P., Weissenbacher, G.: A proposal for a theory of finite sets, lists, and maps for the SMT-LIB standard. In: Proceedings of the SMT Workshop (2007)
13. Kroening, D., Strichman, O.: Decision Procedures. Springer, Heidelberg (2008)
14. Kuncak, V., Rinard, M.: Towards efficient satisfiability checking for Boolean algebra with Presburger arithmetic. In: Pfenning, F. (ed.) CADE 2007. LNCS (LNAI), vol. 4603, pp. 215–230. Springer, Heidelberg (2007)
15. Lahiri, S.K., Musuvathi, M.: An efficient decision procedure for UTVPI constraints. In: Gramlich, B. (ed.) FroCos 2005. LNCS (LNAI), vol. 3717, pp. 168–183. Springer, Heidelberg (2005)
16. Nelson, G., Oppen, D.: Simplification by cooperating decision procedures. ACM Trans. Program. Lang. Syst. **1**(2), 245–257 (1979)
17. Nelson, G., Oppen, D.: Fast decision procedures based on congruence closure. J. ACM **27**(2), 356–364 (1980)
18. Nieuwenhuis, R., Oliveras, A.: DPLL(T) with exhaustive theory propagation and its application to difference logic. In: Etessami, K., Rajamani, S.K. (eds.) CAV 2005. LNCS, vol. 3576, pp. 321–334. Springer, Heidelberg (2005)
19. Oppen, D.: Reasoning about recursively defined data structures. J. ACM **27**(3), 403–411 (1980)
20. Oppen, D.C.: Complexity, convexity and combinations of theories. Theor. Comput. Sci. **12**, 291–302 (1980)
21. Pratt, V.R.: Two Easy Theories Whose Combination is Hard. Technical report, M.I.T. (1977)
22. Rümmer, P., Wahl, T.: An SMT-LIB theory of binary floating-point arithmetic. In: Proceedings of the SMT (2010)
23. Sebastiani, R.: Lazy satisfiability modulo theories. J. Satisfiability Boolean Model. Comput. (JSAT) **3**(3–4), 141–224 (2007)

24. Sebastiani, R.: Colors Make Theories hard. Technical Report DISI-16-001, DISI, University of Trento, Extended version. http://disi.unitn.it/rseba/ijcar16extended. pdf
25. Shostak, R.: An algorithm for reasoning about equality. In: Proceedings of the 7th International Joint Conference on Artificial Intelligence, pp. 526–527 (1977)
26. Shostak, R.: A pratical decision procedure for arithmetic with function symbols. J. ACM **26**(2), 351–360 (1979)