

Article

Complexity Assessments for Decidable Fragments of Set Theory. IV: A Quadratic Reduction from Constraints over Nested Sets to Boolean Formulae

Domenico Cantone ^{1,*} , Andrea De Domenico ² , Pietro Maugeri ^{1,*}  and Eugenio G. Omodeo ³ ¹ Department of Mathematics and Computer Science, University of Catania, 95125 Catania, Italy² IMDEA Software Institute, 28223 Madrid, Spain; andrea.domenico@imdea.org³ Department of Mathematics, Informatics, and Geosciences, University of Trieste, 34127 Trieste, Italy; eomodeo@units.it

* Correspondence: domenico.cantone@unict.it (D.C.); pietro.maugeri@unict.it (P.M.)

Abstract

As a contribution to automated set-theoretic inferencing, a translation is proposed of conjunctions of literals of the forms $x = y \setminus z$, $x \neq y \setminus z$, and $z = \{x\}$, where x, y, z stand for variables ranging over the von Neumann universe of sets, into quantifier-free Boolean formulae of a rather simple conjunctive normal form. The formulae in the target language involve variables ranging over a Boolean ring of sets, along with a difference operator and relators designating equality, non-disjointness, and inclusion. Moreover, the result of each translation is a conjunction of literals of the forms $x = y \setminus z$ and $x \neq y \setminus z$ and of implications whose antecedents are isolated literals and whose consequents are either inclusions (strict or non-strict) between variables, or equalities between variables. Besides reflecting a simple and natural semantics, which ensures satisfiability preservation, the proposed translation has quadratic algorithmic time complexity and bridges two languages, both of which are known to have an NP-complete satisfiability problem.

Keywords: satisfiability problem; computable set theory; expressibility; proof verification; NP-completeness; quantitative logical inference

1. Introduction

As reported in the survey paper [1], the line of investigation to which this paper belongs began in 1978 (ref. [2]), driven by the goal of developing an automated system for proof-checking and program correctness verification. Since then, identifying decidable fragments of set theory has been regarded as a key step toward achieving that goal. As a matter of fact, in the organization of the proof-checker Referee/ \AE tnaNova, which came into existence some thirty years later (ref. [3]), a satisfiability decision tester for a specific class of unquantified set-theoretic formulae plays a central and pervasive role.

When it comes to implementations, complexity emerges as an inescapable issue; this is why we recently undertook a systematic study on the algorithmic complexities of satisfiability testers (see, e.g., [4,5]). In the same frame of mind, this paper enhances a quadratic-cost method (announced ref. [4] and then presented in D. Cantone, A. De Domenico, P. Maugeri, and E.G. Omodeo; a quadratic reduction of constraints over nested sets to purely Boolean formulae in CNF; in *Proc. 35th Italian Conference on Computational Logic*, volume 2710 of *CEUR Workshop Proceedings*, pages 214–230. 2020). This translates the formulae of an



Academic Editor: Ricardo Lopez-Ruiz

Received: 23 September 2025

Revised: 8 January 2026

Accepted: 16 January 2026

Published: 30 January 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article

distributed under the terms and

conditions of the [Creative Commons](https://creativecommons.org/licenses/by/4.0/)[Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

unquantified language involving set-theoretic variables, Boolean operators, and membership and equality relators into propositional combinations of purely Boolean literals.

The enhancement lies in the availability of a singleton-formation operator ‘{■}’, which increases the expressive power of the source language without affecting the algorithmic time complexity of the translation process; however, the fact that the translation preserves satisfiability as well as unsatisfiability must be proved anew and differently. The material that follows bridges two complexity taxonomies, developed and analyzed in the respective papers [4,5]. These taxonomies, one of which—unlike the other—involves membership, concern fragments of set theory for which satisfiability is decidable.

A fit theoretical framework for the study of our target language is the theory of Boolean rings, a merely equational first-order theory endowed with finitely many axioms—at times, one blends this theory with an arithmetic of cardinals (see, e.g., [6])—or just a surrogate of it, cf. Figure 1 in Section 7. Frameworks for the study of the source language are such all-embracing theories as ZF and NBG (the Zermelo–Fraenkel and von Neumann–Bernays–Gödel theories), within which one can cast the whole corpus of mathematical disciplines. Boolean algebra is decidable in its entirety (cf. Section 3.7 ref. [7]); ZF is essentially undecidable, yet nonetheless an effort to find practical decision algorithms for fragments of it began in 1979. The rationale of this long-standing research is that satisfiability testers embodying some knowledge about ZF can act as key inference mechanisms within a programmed system apt to verifying the correctness of large-scale mathematical proofs as envisaged ref. [3].

(D.1)	$x \setminus (y \setminus y) = x$	Existence of zero
(D.2)	$(x \setminus y) \setminus z = (x \setminus z) \setminus y$	Permutativity
(D.3)	$x \setminus (x \setminus y) = y \setminus (y \setminus x)$	Commutativity (of intersection)
(D.4)	$(x \setminus y) \setminus y = x \setminus y$	Double subtraction

Figure 1. Axioms of the variety of difference algebras.

A priori, one would expect the distance between the performances of decision algorithms for fragments of Boolean algebra, and of the seemingly much more expressive languages whose dictionaries embody nested membership, to be abysmal. Luckily, though, as we will see, this is not the case.

We introduce in Section 2 an interpreted formal language, dubbed \mathbb{BST} , within which one can formulate unquantified Boolean constraints. Despite its syntax being quite minimal— \mathbb{BST} only encompasses conjunctions of primitive literals of two forms, namely, $x = y \setminus z$ and $x \neq y \setminus z$ —, the satisfiability problem for \mathbb{BST} is NP-complete (ref. [4]). By way of abbreviations, a number of additional constraints (e.g., literals of the form $x \neq y \cup z$) can be expressed in \mathbb{BST} .

Alongside \mathbb{BST} , we also consider its natural *Boolean closure* \mathbb{BST}^+ , defined as the language of all propositional combinations—using $\wedge, \vee, \rightarrow, \leftrightarrow,$ and \neg with arbitrary nesting—of atoms of the form $x = y \setminus z$. It is straightforward to verify that the satisfiability problem for \mathbb{BST}^+ can be reduced to that for \mathbb{BST} in nondeterministic polynomial time; therefore, it is NP-complete in turn.

According to our semantics, the domain of discourse to which \mathbb{BST} (and \mathbb{BST}^+) refers is a universe of nested sets; however, as will be seen in Section 5, every satisfiable propositional combination of \mathbb{BST} literals (in particular, every satisfiable \mathbb{BST} constraint) admits a model consisting of sets which are, in a certain sense, ‘flat’. This makes it clear that there is no straightforward way of expressing membership atoms $x \in y$ in \mathbb{BST} . Yet, many natural set-theoretic constructors (most notably the singleton operator $\{\blacksquare\}$) inherently interact with

membership, and understanding whether and how such constructs can be eliminated while preserving satisfiability and complexity is the main motivation of this paper.

To address this research gap, Section 3 introduces a complexity-aware notion of expressibility, which requires not only a semantic reduction but also an explicit bound on the cost of the translation. While related notions already appear ref. [4], the present formulation is tailored to our setting and is a valuable technical tool that enables the complexity-preservation arguments developed here.

In terms of this notion, Section 6 provides an explicit translation of conjunctions of literals of the three forms $x = y \setminus z$, $x \neq y \setminus z$, and $z = \{x\}$ into propositional combinations of \mathbb{BST} literals, i.e., into \mathbb{BST}^+ . As shown in Sections 6.3 and 6.4, the translation preserves satisfiability. It yields a conjunction whose conjuncts are either \mathbb{BST} literals or simple disjunctions. Moreover, the translation can be computed in quadratic time (see Section 6.1), which entails that satisfiability remains NP-complete when the singleton operator $\{\bullet\}$ is added to the constructs of \mathbb{BST} . This NP-completeness result was already known (see, e.g., [8]), but our approach derives it anew via an explicit, complexity-controlled elimination procedure.

The concluding section, Section 7, recapitulates what has been accomplished, indicates what we intend to pursue next, and situates our line of investigation in relation to that of other scholars.

2. The Theories \mathbb{BST} and \mathbb{BST}^+

Boolean set theory (\mathbb{BST}) is the quantifier-free theory consisting of all finite conjunctions of literals of the forms

$$x = y \setminus z, \quad x \neq y \setminus z, \tag{1}$$

where x , y , and z are *set variables*, ranging over the universe of well-founded sets.

Remark 1. *Throughout, we use the term theory in a nonstandard sense, referring to a collection of set-theoretic formulae—possibly involving, along with primitive symbols, some derived constructs—rather than a deductively closed collection of sentences.*

Semantics for the theory \mathbb{BST} is defined in terms of set assignments. Specifically, given a (finite) collection V of set variables, a *set assignment* M over V —the *domain of definition* of M , denoted by $\text{dom}(M)$ —is any map from V into the *von Neumann universe* \mathcal{V} (see below). (Note that our semantics of \mathbb{BST} does not rely on flat sets of urelements (as would be doable). Working with those would call for minor adjustments, unjustified—and perhaps disturbing—in the economy of this paper.) The *support* of M , denoted by $\text{supp}(M)$, is the union of the sets assigned by M to the variables in its domain, namely,

$$\text{supp}(M) := \bigcup_{x \in \text{dom}(M)} Mx.$$

We shall refer to the cardinality of the support of M as the *cardinality of the assignment* M . A set assignment M is said to *satisfy*, or to *model*, a given literal $x = y \setminus z$, with $x, y, z \in \text{dom}(M)$, if $Mx = My \setminus Mz$ holds, where ‘ \setminus ’ denotes standard set difference. Likewise, M is said to satisfy the literal $x \neq y \setminus z$ if $Mx \neq My \setminus Mz$ holds. Finally, M satisfies a \mathbb{BST} -conjunction φ such that $\text{Vars}(\varphi) \subseteq \text{dom}(M)$ (where $\text{Vars}(\varphi)$ denotes the collection of the variables occurring free in φ) if it satisfies all of the conjuncts of φ , in which case we say that M is a *model* of φ and write

$$M \models \varphi$$

(a short for $M \models_{\mathcal{V}} \varphi$). A BST-conjunction φ is said to be *satisfiable* if it has some model; otherwise, it is said to be *unsatisfiable*. It is said to be *valid* if it is true in every model, in which case we write

$$\models \varphi$$

(again, a short for $\models_{\mathcal{V}} \varphi$). Two BST-conjunctions φ and ψ are said to be *equisatisfiable* if either both are satisfiable or both are unsatisfiable—that is, if φ has a model if and only if ψ has a model. Equisatisfiability does not require the two formulae to have the same models, but only that they agree in satisfiability, irrespective of their models. The notions introduced in this paragraph, here tailored to the case of BST-conjunctions, extend naturally to all theories discussed in the sequel.

In ref. [4], it is proved that the satisfiability problem for BST, namely, the problem of establishing algorithmically the satisfiability status of any given BST-conjunction, is NP-complete.

The extension BST^+ .

Recall that BST only allows for *conjunctions* of its primitive literals—in particular, a BST-formula is a finite conjunction of constraints of the form $x = y \setminus z$ and $x \neq y \setminus z$. In contrast, BST^+ is the *Boolean closure* of BST: it consists of all propositional combinations of atoms of the form $x = y \setminus z$, obtained by unrestrained use of the logical connectives \wedge , \vee , \longrightarrow , \longleftarrow , and \neg . Thus, BST^+ strictly extends BST by allowing for arbitrary Boolean structure (e.g., disjunctions, implications, and negations at any depth), whereas BST is restricted to conjunctive constraints.

The satisfiability problem for BST^+ reduces to that for BST in nondeterministic polynomial time (by guessing a satisfying assignment for the Boolean structure and checking the resulting conjunction), and hence the satisfiability problem for BST^+ is NP-complete as well.

The von Neumann Universe

We recall that the von Neumann universe \mathcal{V} of (well-founded) sets, also dubbed *von Neumann cumulative hierarchy*, is built up through a transfinite sequence of steps as the union $\mathcal{V} := \bigcup_{\alpha \in On} \mathcal{V}_{\alpha}$ of the levels $\mathcal{V}_{\alpha} := \bigcup_{\beta < \alpha} \mathcal{P}(\mathcal{V}_{\beta})$, with $\mathcal{P}(\blacksquare)$ denoting the powerset operator and α ranging over the class *On* of all ordinals.

It can easily be seen that, for every ordinal α (and, in particular, for every integer), we have

$$\mathcal{V}_{\alpha+1} = \mathcal{P}(\mathcal{V}_{\alpha}),$$

and consequently

$$|\mathcal{V}_{\alpha+1}| = 2^{|\mathcal{V}_{\alpha}|}, \tag{2}$$

so that $|\mathcal{V}_{\alpha}| < |\mathcal{V}_{\alpha+1}|$.

Based on the level of first appearance in the von Neumann hierarchy, one can define the *rank* of any set s , denoted $\text{rk}(s)$. Specifically, $\text{rk}(s)$ is the ordinal α such that $s \in \mathcal{V}_{\alpha+1} \setminus \mathcal{V}_{\alpha}$. Hence, for every $\alpha \in On$, the set $\mathcal{V}_{\alpha+1} \setminus \mathcal{V}_{\alpha}$, hereinafter denoted $\mathcal{W}_{\alpha}^{\#}$, collects all sets having rank α .

The following lower bound on the number of well-founded sets of any positive integer rank n , to be proved as Proposition A2 in Appendix A, will be useful:

$$|\mathcal{W}_n^{\#}| \geq 2^{n-1}. \tag{3}$$

A particularly important subclass of \mathcal{V} is that of *hereditarily finite sets*, denoted by HF. By definition,

$$\text{HF} := \bigcup_{n \in \mathbb{N}} \mathcal{V}_n,$$

that is, HF consists of all sets of finite rank.

Some handy properties of the rank function that we shall tacitly use are the following, which hold for all sets $s, t \in \mathcal{V}$:

- If $s \in t$, then $\text{rk}(s) < \text{rk}(t)$;
- If $s \subseteq t$, then $\text{rk}(s) \leq \text{rk}(t)$;
- $\text{rk}(s) = \begin{cases} 0 & \text{if } s = \emptyset, \\ \sup_{u \in s} (\text{rk}(u) + 1) & \text{(i.e., } \bigcup_{u \in s} (\text{rk}(u) + 1)) \text{ otherwise.} \end{cases}$

We also recall that well-foundedness, as enforced by the *regularity* (or *foundation*) axiom of set theory, precludes the formation of infinite descending membership chains of the form

$$\cdots \in s_2 \in s_1 \in s_0,$$

and, in particular, membership cycles of the form

$$s_0 \in s_n \in \cdots \in s_2 \in s_1 \in s_0,$$

for any sets s_i .

3. From Existential Expressibility to $\mathcal{O}(f)$ -Expressibility Across Theories

We first recall the definition of *existential expressibility* (cf. [4] for several applications of this notion).

Definition 1 (Existential expressibility). *A formula $\psi(x)$ is said to be existentially expressible in a theory \mathcal{T} if there exists a \mathcal{T} -formula $\Psi(x, z)$ such that*

$$\models \psi(x) \longleftrightarrow (\exists z) \Psi(x, z),$$

where x and z stand for tuples of set variables.

In spite of the parsimony of $\mathbb{B}\mathcal{S}\mathcal{T}$ as just presented, it turns out (see [4]) that several other Boolean constructs, such as the ones in the list of literals

$$\begin{aligned} x = \emptyset, \quad x \subseteq y, \quad x = y \cap z, \quad x = y \cup z, \quad \text{DISJ}(x, y), \quad x \subsetneq y, \\ x \neq \emptyset, \quad x \not\subseteq y, \quad x \neq y \cap z, \quad x \neq y \cup z, \quad \neg\text{DISJ}(x, y), \end{aligned} \tag{4}$$

can be expressed existentially in $\mathbb{B}\mathcal{S}\mathcal{T}$, where $\text{DISJ}(a, b)$ is a short for $a \cap b = \emptyset$.

Ref. [4], existential expressibility was generalized into $\mathcal{O}(f)$ -expressibility (we use standard asymptotic notations $\mathcal{O}(\cdot)$, $\Omega(\cdot)$, and $\Theta(\cdot)$ throughout the paper; see, e.g., Ch. 3 of [9]), a notion that helped develop a fine-grained complexity taxonomy of the subfragments of $\mathbb{B}\mathcal{S}\mathcal{T}$.

Definition 2 ($\mathcal{O}(f)$ -expressibility). *Let \mathcal{T} be a theory and let $f: \mathbb{N} \rightarrow \mathbb{N}$ be a complexity function. A formula $\psi(x)$ —typically involving a construct one aims to eliminate—is said to be $\mathcal{O}(f)$ -expressible in \mathcal{T} if there exists a transformation*

$$\varphi(y) \mapsto \Psi_\varphi(x, y, z) \tag{5}$$

from \mathcal{T} to \mathcal{T} , where no variable in z occurs in x or y , such that, for every φ , the following conditions hold:

- (a) The transformation (5) can be computed in $\mathcal{O}(f(|\varphi|))$ -time;
- (b) If $\varphi(y) \wedge (\exists z) \Psi_\varphi(x, y, z)$ is satisfiable, then $\varphi(y) \wedge \psi(x)$ is satisfiable;
- (c) $\models (\varphi(y) \wedge \psi(x)) \longrightarrow (\exists z) \Psi_\varphi(x, y, z)$.

Here, we further generalize $\mathcal{O}(f)$ -expressibility in two directions. First, we allow for a collection \mathcal{C} of set-theoretic formulae, rather than a single formula ψ as in [4]. Second, we explicitly distinguish a source theory \mathcal{T}_1 from a target theory \mathcal{T}_2 , whereas [4] works in the single-theory setting, with source and target taken to coincide.

Definition 3 ($\mathcal{O}(f)$ -expressibility across theories). *Let \mathcal{T}_1 and \mathcal{T}_2 be any theories and $f : \mathbb{N} \rightarrow \mathbb{N}$ be a given complexity function. A collection \mathcal{C} of formulae is said to be $\mathcal{O}(f)$ -expressible from \mathcal{T}_1 into \mathcal{T}_2 if there exists a map*

$$\langle \varphi(\mathbf{y}), \psi(\mathbf{x}) \rangle \mapsto \Xi_\varphi^\psi(\mathbf{x}, \mathbf{y}, \mathbf{z}) \tag{6}$$

from $\mathcal{T}_1 \times \mathcal{C}$ into \mathcal{T}_2 , where no variable in \mathbf{z} occurs in either \mathbf{x} or \mathbf{y} , such that the following conditions are satisfied:

- (a) The mapping (6) can be computed in $\mathcal{O}(f(|\varphi \wedge \psi|))$ time;
- (b) If $\varphi(\mathbf{y}) \wedge \Xi_\varphi^\psi(\mathbf{x}, \mathbf{y}, \mathbf{z})$ is satisfiable, so is $\varphi(\mathbf{y}) \wedge \psi(\mathbf{x})$;
- (c) $\models (\varphi(\mathbf{y}) \wedge \psi(\mathbf{x})) \rightarrow (\exists \mathbf{z}) \Xi_\varphi^\psi(\mathbf{x}, \mathbf{y}, \mathbf{z})$.

Observe that the two formulae appearing in condition (b) of Definition 3—namely, $\varphi(\mathbf{y}) \wedge \Xi_\varphi^\psi(\mathbf{x}, \mathbf{y}, \mathbf{z})$ and $\varphi(\mathbf{y}) \wedge \psi(\mathbf{x})$ —are in fact equisatisfiable. This is formalized in the following lemma.

Lemma 1. *Let \mathcal{T}_1 and \mathcal{T}_2 be theories, and let \mathcal{C} be a collection of formulae that is $\mathcal{O}(f)$ -expressible from \mathcal{T}_1 into \mathcal{T}_2 via a given mapping $\langle \varphi(\mathbf{y}), \psi(\mathbf{x}) \rangle \mapsto \Xi_\varphi^\psi(\mathbf{x}, \mathbf{y}, \mathbf{z})$, as defined in Definition 3. Then, for every $\varphi(\mathbf{y}) \in \mathcal{T}_1$ and every $\psi(\mathbf{x}) \in \mathcal{C}$, the formulae*

$$\varphi(\mathbf{y}) \wedge \Xi_\varphi^\psi(\mathbf{x}, \mathbf{y}, \mathbf{z}) \quad \text{and} \quad \varphi(\mathbf{y}) \wedge \psi(\mathbf{x})$$

are equisatisfiable.

Proof. Let $\varphi(\mathbf{y}) \in \mathcal{T}_1$ and $\psi(\mathbf{x}) \in \mathcal{C}$. We prove that the two conjunctions $\varphi(\mathbf{y}) \wedge \psi(\mathbf{x})$ and $\varphi(\mathbf{y}) \wedge \Xi_\varphi^\psi(\mathbf{x}, \mathbf{y}, \mathbf{z})$ are equisatisfiable.

(\Rightarrow) Suppose $\varphi(\mathbf{y}) \wedge \Xi_\varphi^\psi(\mathbf{x}, \mathbf{y}, \mathbf{z})$ is satisfiable. Then, by condition (b) of Definition 3, it follows that $\varphi(\mathbf{y}) \wedge \psi(\mathbf{x})$ is satisfiable.

(\Leftarrow) Conversely, suppose $\varphi(\mathbf{y}) \wedge \psi(\mathbf{x})$ is satisfiable, and let M be a model of this formula. Then, by condition (c) of Definition 3, we have

$$M \models \varphi(\mathbf{y}) \wedge (\exists \mathbf{z}) \Xi_\varphi^\psi(\mathbf{x}, \mathbf{y}, \mathbf{z}).$$

This means that there exists an extension M^* of M to the variables in \mathbf{z} —which, by hypothesis, occur neither in \mathbf{x} nor in \mathbf{y} —such that $M^* \models \Xi_\varphi^\psi(\mathbf{x}, \mathbf{y}, \mathbf{z})$; hence, $M^* \models \varphi(\mathbf{y}) \wedge \Xi_\varphi^\psi(\mathbf{x}, \mathbf{y}, \mathbf{z})$.

From the satisfiability of one formula, we have derived that of the other; hence, the two formulae are equisatisfiable. \square

While Lemma 1 shows that the two formulae $\varphi(\mathbf{y}) \wedge \psi(\mathbf{x})$ and $\varphi(\mathbf{y}) \wedge \Xi_\varphi^\psi(\mathbf{x}, \mathbf{y}, \mathbf{z})$ are equisatisfiable, Definition 3 does not adopt equisatisfiability as a primitive requirement. Instead, it explicitly requires the two separate conditions (b) and (c), which together entail equisatisfiability but provide strictly more structure. This asymmetry is intentional: clause (b) ensures that satisfiability of the translated formula readily implies satisfiability of the original, supporting the soundness of the translation; clause (c) provides a constructive existential guarantee—namely, that any model of $\varphi \wedge \psi$ can be extended to a model of $\varphi \wedge \Xi_\varphi^\psi$, thus ensuring a form of witness-preserving completeness.

To analyze complexity preservation under $\mathcal{O}(f)$ -expressibility, we encounter expressions like $g(n + f(n))$ and, more generally, $g(\mathcal{O}(f(n)))$. It is therefore useful to assume a mild robustness of g under constant scaling, formalized below.

Definition 4 (Scale-invariant function). *A function $g: \mathbb{N} \rightarrow \mathbb{N}$ is said to be scale-invariant (up to constants) if, for every real constant $c > 0$, there exist constants $K_c > 0$ and $N_c \in \mathbb{N}$ such that*

$$g(\lceil cn \rceil) \leq K_c g(n), \quad \text{for all } n \geq N_c.$$

Equivalently, $g(\lceil cn \rceil) = \mathcal{O}(g(n))$ for every fixed $c > 0$.

Remark 2. *Definition 4 is closely related to classical doubling-type assumptions. For instance, in the theory of Orlicz spaces, the Δ_2 -condition for a Young function Φ requires $\Phi(2t) \leq C \Phi(t)$ for all $t \geq 0$ (ref. [10]). Similarly, if g is (eventually) nondecreasing and satisfies a doubling-type bound, namely, there exist constants $K > 0$ and $N \in \mathbb{N}$ such that*

$$g(2n) \leq K g(n) \quad \text{for all } n \geq N,$$

then the estimate extends to any fixed scaling factor $c > 1$ by iteration. Indeed, choose $k \in \mathbb{N}$ with $c \leq 2^k$. For all $n \geq 2^{k-1}N$, we have

$$g(\lceil cn \rceil) \leq g(2^k n) \leq K^k g(n).$$

Finally, our assumption can be viewed as a discrete, one-sided variant of \mathcal{O} -regular variation. In the Karamata–Matuszewska framework, a function f is called \mathcal{O} -regularly varying if, for every $\lambda \geq 1$, the ratio $f(\lambda x)/f(x)$ stays asymptotically bounded away from both 0 and $+\infty$ (via \liminf and \limsup) (cf. [11]). In contrast, we only require the upper control $g(\lceil cn \rceil) = \mathcal{O}(g(n))$, which is sufficient for the complexity-preservation steps where constant-factor rescalings arise.

Example 1. *Let $g(n) = n^2$ and take $c > 0$. Then,*

$$g(\lceil cn \rceil) = \lceil cn \rceil^2 \leq \lceil c \rceil^2 n^2 = \lceil c \rceil^2 g(n),$$

and hence $g(\lceil cn \rceil) = \mathcal{O}(g(n))$.

More generally, the scale-invariance property holds for all standard polynomially bounded complexity functions, such as n^k and $n \log n$, as well as for poly-logarithmic functions. More generally, it holds for functions of at most polynomial growth, and it also follows from a weak subadditivity condition: there exist constants $C_1 \geq 1$ and $C_2 \geq 0$ such that

$$g(m + n) \leq C_1(g(m) + g(n)) + C_2,$$

for all $m, n \in \mathbb{N}$.

In contrast, $g(n) = 2^n$ is not scale-invariant: for $c = 2$ we have $g(2n) = 2^{2n} = (2^n)^2$, and the ratio $g(2n)/g(n) = 2^n$ is unbounded, so $g(2n) \notin \mathcal{O}(g(n))$.

The next lemma provides the quantitative part of the reduction: under $\mathcal{O}(f)$ -expressibility, it transfers the time bound for $\mathcal{T}_1 \wedge \mathcal{T}_2$ to $\mathcal{T}_1 \wedge \mathcal{C}$. The assumptions that g is nondecreasing and scale-invariant control the rescaling of input sizes, while the at least linear growth of g ensures that the preprocessing cost $\mathcal{O}(f(n))$ is absorbed in the final bound.

Lemma 2. *Let \mathcal{T}_1 and \mathcal{T}_2 be two theories (in the sense of Remark 1) and suppose that their conjunctive product*

$$\mathcal{T}_1 \wedge \mathcal{T}_2 := \{\varphi_1 \wedge \varphi_2 : \varphi_1 \in \mathcal{T}_1, \varphi_2 \in \mathcal{T}_2\}$$

admits a decision procedure running in time $\mathcal{O}(g(n))$, for some nondecreasing, scale-invariant function $g: \mathbb{N} \rightarrow \mathbb{N}$ of at least linear growth, i.e., $g(n) = \Omega(n)$.

Suppose moreover that the collection \mathcal{C} of formulae is $\mathcal{O}(f)$ -expressible from \mathcal{T}_1 into \mathcal{T}_2 , for some complexity function $f: \mathbb{N} \rightarrow \mathbb{N}$.

Then, the combined theory $\mathcal{T}_1 \wedge \mathcal{C}$ is decidable in time $\mathcal{O}(g(n + f(n)))$.

Proof. Let $\varphi(\mathbf{y}) \in \mathcal{T}_1$ and $\psi(\mathbf{x}) \in \mathcal{C}$, and set $n := |\varphi \wedge \psi|$. Since \mathcal{C} is $\mathcal{O}(f)$ -expressible from \mathcal{T}_1 into \mathcal{T}_2 , we can compute, in time $\mathcal{O}(f(n))$, a formula $\Xi_\varphi^\psi(\mathbf{x}, \mathbf{y}, \mathbf{z}) \in \mathcal{T}_2$ satisfying conditions (a)–(c) of Definition 3. Moreover, by condition (a) (and the fact that the output of an $\mathcal{O}(f(n))$ -time transformation has size $\mathcal{O}(f(n))$), we have $|\Xi_\varphi^\psi| = \mathcal{O}(f(n))$. Hence, there exist a constant $c > 0$ and a threshold $N_0 \in \mathbb{N}$ such that, for every input pair $(\varphi, \psi) \in \mathcal{T}_1 \times \mathcal{C}$ with $|\varphi \wedge \psi| \geq N_0$, we have

$$|\varphi \wedge \Xi_\varphi^\psi| \leq c(n + f(n)) \quad \text{and thus} \quad |\varphi \wedge \Xi_\varphi^\psi| \leq \lceil c(n + f(n)) \rceil,$$

where $n := |\varphi \wedge \psi|$ as stipulated above.

By Lemma 1, the formulae $\varphi \wedge \psi$ and $\varphi \wedge \Xi_\varphi^\psi$ are equisatisfiable. Therefore, to decide satisfiability of $\varphi \wedge \psi$, it suffices to decide satisfiability of $\varphi \wedge \Xi_\varphi^\psi$, which belongs to $\mathcal{T}_1 \wedge \mathcal{T}_2$. By hypothesis, this can be done in time $\mathcal{O}(g(|\varphi \wedge \Xi_\varphi^\psi|))$. Using that g is nondecreasing and scale-invariant, we obtain

$$g(|\varphi \wedge \Xi_\varphi^\psi|) \leq g(\lceil c(n + f(n)) \rceil) = \mathcal{O}(g(n + f(n))).$$

Finally, the overall running time is the sum of the preprocessing time $\mathcal{O}(f(n))$ needed to compute Ξ_φ^ψ and the decision time for $\varphi \wedge \Xi_\varphi^\psi$; hence,

$$\mathcal{O}(f(n)) + \mathcal{O}(g(n + f(n))).$$

Since g has at least linear growth, there exist constants $a > 0$ and $N_1 \in \mathbb{N}$ such that $g(t) \geq at$ for all $t \geq N_1$. For all sufficiently large n , we have $n + f(n) \geq N_1$, and therefore

$$g(n + f(n)) \geq a(n + f(n)) \geq af(n),$$

whence $\mathcal{O}(f(n)) \subseteq \mathcal{O}(g(n + f(n)))$. Consequently, the total running time is $\mathcal{O}(g(n + f(n)))$.

Since this procedure works uniformly for every input $\varphi \wedge \psi \in \mathcal{T}_1 \wedge \mathcal{C}$, the claim follows. \square

4. BST-Replacements and Flat Models

In this section, we develop two central notions that will be used in the rest of the paper. The first is that of *BST-replacement*, which captures how set assignments can be modified in a controlled way while preserving the satisfaction of \mathbb{BST}^+ -formulae.

The second is the notion of a *b-flat model*, namely, a set assignment in which every element occurring in the interpretation of a variable has rank exactly b . This ‘single-layer support’ property makes it possible to reason cleanly about disjointness and membership when dealing with singleton atoms under replacement.

4.1. Replacement Assignments and BST-Replacements

In preparation for the results that follow, we need a method to modify a set assignment for a BST formula without disrupting satisfiability. Such a method is implemented by means of BST-replacements.

Definition 5 (Replacement assignments and BST-replacements). Let M be a set assignment over a collection V of variables, let $W \subseteq V$, and let S, T be nonempty sets. The replacement (assignment) restricted to W of M from S to T (or with respect to the pair (S, T)), denoted $(M|_W)_T^S$, is the set assignment defined for each $x \in V$ by

$$(M|_W)_T^S x := \begin{cases} (Mx \setminus S) \cup T & \text{if } x \in W \wedge \neg \text{DISJ}(S, Mx) \\ Mx & \text{otherwise.} \end{cases} \tag{7}$$

The assignment $(M|_W)_T^S$ is a BST-REPLACEMENT RESTRICTED TO W (or a BST_W -replacement) of M from S to T if in addition the following condition holds:

$$(\forall x \in W) ((\text{DISJ}(S, Mx) \vee S \subseteq Mx) \wedge \text{DISJ}(T \setminus S, Mx)). \tag{8}$$

When $W = V$, we omit W from the notation and write M_T^S , calling it simply a REPLACEMENT of M from S to T . If condition (8) also holds, it is a BST-REPLACEMENT of M from S to T .

Note that the requirement that S and T are nonempty is essential only for T : if $S = \emptyset$, then $(M|_W)_T^S = M$ for any T , so the replacement has no effect and is of no interest.

Example 2 (BST_W -replacement). Let

$$V = \{x, y, z\}, \quad W = \{x, y\}, \quad S = \{2, 3\}, \quad T = \{3, 7\},$$

and define the set assignment M by

$$Mx = \{1, 2, 3\}, \quad My = \{2, 3, 4\}, \quad Mz = \{5\}.$$

Note that $S, T \neq \emptyset$.

Since $x, y \in W$ with $\neg \text{DISJ}(S, Mx)$ and $\neg \text{DISJ}(S, My)$, the restricted replacement $(M|_W)_T^S$ acts on x and y . On the other hand, as $z \notin W$, z is left unchanged:

$$(M|_W)_T^S(x) = (Mx \setminus S) \cup T = (\{1, 2, 3\} \setminus \{2, 3\}) \cup \{3, 7\} = \{1\} \cup \{3, 7\} = \{1, 3, 7\},$$

$$(M|_W)_T^S(y) = (My \setminus S) \cup T = (\{2, 3, 4\} \setminus \{2, 3\}) \cup \{3, 7\} = \{4\} \cup \{3, 7\} = \{3, 4, 7\},$$

$$(M|_W)_T^S(z) = Mz = \{5\}.$$

We verify the BST-condition (8) for all $u \in W$:

$$(\text{DISJ}(S, Mu) \vee S \subseteq Mu) \wedge \text{DISJ}(T \setminus S, Mu).$$

For $u = x$ and $u = y$,

$$S = \{2, 3\} \subseteq Mx = \{1, 2, 3\}, \quad S = \{2, 3\} \subseteq My = \{2, 3, 4\},$$

and

$$T \setminus S = \{7\} \text{ is disjoint from } Mx = \{1, 2, 3\} \text{ and } My = \{2, 3, 4\}.$$

Hence, the condition holds for every $u \in W$. Therefore,

$$(M|_W)_T^S \text{ is a } \text{BST}_W\text{-replacement of } M \text{ from } S \text{ to } T.$$

Remark 3 (Restricted setting convention). All results in this subsection are stated in the unrestricted case ($W = V$), except where explicitly noted otherwise (cf. Lemma 6), but they extend plainly to the restricted setting $(M|_W)_T^S$ with $W \subseteq V$.

To study more closely how BST-replacements affect a set assignment, it is useful to introduce the *partition induced by a set assignment*, which decomposes its support into blocks determined by membership patterns of elements with respect to the assignment’s variable interpretations.

Definition 6. Given a set assignment M over a collection V of variables, the PARTITION INDUCED BY M is defined as

$$\Sigma_M := \left\{ \bigcap_{x \in U} Mx \setminus \bigcup_{y \in V \setminus U} My \mid \emptyset \neq U \subseteq V \right\} \setminus \{\emptyset\}.$$

This notion provides a precise way to track the effect of a replacement, as illustrated in the subsequent lemma.

Lemma 3. Let Φ be a formula of BST^+ , M a set assignment for it, and T a nonempty set such that $\text{DISJ}(T, Mv)$ holds for all $v \in \text{Vars}(\Phi)$. Then, for all $\sigma \in \Sigma_M$

- (a) M_T^σ is a BST-replacement of M ;
- (b) $\Sigma_{M_T^\sigma} = (\Sigma_M \setminus \{\sigma\}) \cup \{T\}$.

Proof. Let $\sigma \in \Sigma_M$ and set $M' := M_T^\sigma$.

(a) By the definition of Σ_M , for every $x \in \text{Vars}(\Phi)$ either $\sigma \subseteq Mx$ or $\text{DISJ}(\sigma, Mx)$. Together with the assumption that $\text{DISJ}(T, Mv)$ holds for all $v \in \text{Vars}(\Phi)$, this is precisely condition (8). Hence, M_T^σ is a BST-replacement of M .

(b) Each block of Σ_M collects the elements that share the same membership pattern with respect to the variables of Φ . The replacement M_T^σ removes σ wherever it occurs and inserts T instead. Since T is disjoint from every Mv , no other block is modified, and the pattern corresponding to σ is now realized by T . Therefore,

$$\Sigma_{M'} = (\Sigma_M \setminus \{\sigma\}) \cup \{T\}. \quad \square$$

We note that the hypothesis on T in Lemma 3 (namely, its disjointness from every Mv) is stronger than what is needed in general. The next result shows that a replacement can always be reversed under the sole assumption that the BST-replacement conditions (8) are satisfied.

Lemma 4. Let M be a set assignment over a collection V of variables, and let S, T be nonempty sets such that M_T^S is a BST-replacement of M from S to T . Then, $(M_T^S)^T_S$ is a BST-replacement of M from T to S , and

$$(M_T^S)^T_S = M.$$

Proof. Set $M' := M_T^S$. We first check that the BST-conditions (8) hold for the pair (T, S) relative to M' and then show that $(M')^T_S = M$.

Fix $x \in V$. From the BST-conditions for (S, T) relative to M , we know that either $\text{DISJ}(S, Mx)$ or $S \subseteq Mx$, and in all cases $\text{DISJ}(T \setminus S, Mx)$.

Case 1: $\text{DISJ}(S, Mx)$. Then, $M'x = Mx$. Furthermore,

$$T \cap Mx = ((T \setminus S) \cup (T \cap S)) \cap Mx = ((T \setminus S) \cap Mx) \cup ((T \cap S) \cap Mx) = \emptyset,$$

since $\text{DISJ}(T \setminus S, Mx)$ and $\text{DISJ}(S, Mx)$. Hence, $\text{DISJ}(T, M'x)$. Also, $\text{DISJ}(S \setminus T, M'x)$ holds because $S \setminus T \subseteq S$ and $\text{DISJ}(S, Mx)$. Thus, the BST-conditions for (T, S) hold at x , and the reverse update does nothing:

$$(M')^T_S x = M'x = Mx.$$

Case 2: $S \subseteq Mx$. Then, $M'x = (Mx \setminus S) \cup T$, so $T \subseteq M'x$. Moreover, $S \setminus T$ is disjoint from $Mx \setminus S$ and from T ; hence, $\text{DISJ}(S \setminus T, M'x)$. Therefore, the BST-conditions for (T, S) hold at x , and

$$(M')_S^T x = (M'x \setminus T) \cup S = (((Mx \setminus S) \cup T) \setminus T) \cup S = (Mx \setminus S) \cup S = Mx.$$

In both cases, the (T, S) -conditions hold relative to M' , and $(M')_S^T x = Mx$. Since x was arbitrary, we conclude that $(M')_S^T$ is a BST-replacement of M' from T to S , and indeed $(M')_S^T = M$. \square

The next lemma shows that satisfiability of a BST^+ -formula is preserved under BST-replacements: if a formula holds in a given assignment, it continues to hold after any valid BST-replacement.

Lemma 5. For every set assignment M , every formula Φ of BST^+ , and every BST-replacement M' of M , we have

$$M \models \Phi \iff M' \models \Phi.$$

Proof. Let (S, T) be any pair of nonempty sets fulfilling the BST-replacement conditions (8) with respect to M . As BST^+ -formulae are Boolean combinations of atoms of the form $x = y \setminus z$, it is enough to show that for every such atom (with variables in the domain of M) we have

$$M \models x = y \setminus z \iff M_T^S \models x = y \setminus z.$$

Suppose $M \models x = y \setminus z$, that is, $Mx = My \setminus Mz$.

We analyze separately the two cases $\text{DISJ}(S, Mx)$ and $\neg \text{DISJ}(S, Mx)$.

Case $\text{DISJ}(S, Mx)$. Then, $Mx = M_T^S x$, and since $S \subseteq My$ implies $S \subseteq Mz$, we have

$$M_T^S x = Mx = My \setminus Mz = (My \setminus S) \setminus (Mz \setminus S) = ((My \setminus S) \cup T) \setminus ((Mz \setminus S) \cup T),$$

using (8) to ensure $\text{DISJ}(T \setminus S, Mz)$. Hence, $M_T^S \models x = y \setminus z$.

Case $\neg \text{DISJ}(S, Mx)$. By (8), we have $S \subseteq Mx \subseteq My$, and hence $\text{DISJ}(S, Mz)$. Using the definition of M_T^S , we compute

$$\begin{aligned} M_T^S x &= (Mx \setminus S) \cup T = ((My \setminus Mz) \setminus S) \cup T \\ &= ((My \setminus S) \setminus Mz) \cup T = ((My \setminus S) \cup T) \setminus Mz = M_T^S y \setminus M_T^S z, \end{aligned}$$

where we used $\text{DISJ}(T \setminus S, Mz)$ and $\text{DISJ}(S, Mz)$ to conclude $\text{DISJ}(T, Mz)$.

Thus, in all cases,

$$M \models x = y \setminus z \implies M_T^S \models x = y \setminus z.$$

For the converse implication, assume $M_T^S \models x = y \setminus z$. By Lemma 4, $(M_T^S)_S^T$ is a BST-replacement of M_T^S and moreover $(M_T^S)_S^T = M$; hence, by the first part,

$$M_T^S \models x = y \setminus z \implies M \models x = y \setminus z.$$

Therefore, for every atomic formula of the form $x = y \setminus z$,

$$M \models x = y \setminus z \iff M_T^S \models x = y \setminus z.$$

It follows that the logical equivalence holds for all BST^+ -formulae Φ . This proves the lemma. \square

For later use, we also state a restricted variant, in which replacements are carried out only with respect to a chosen subset of variables.

Lemma 6. *Let V be a collection of variables and $W \subseteq V$. Let M be a set assignment over V , and let S, T be nonempty sets such that*

$$M' := (M|_W)_T^S$$

is a BST-replacement of M restricted to W . If Φ is a \mathbb{BST}^+ -formula such that for every literal ℓ in Φ either $\text{Vars}(\ell) \subseteq W$ or $\text{Vars}(\ell) \cap W = \emptyset$ holds, then

$$M \models \Phi \iff M' \models \Phi.$$

Proof. As \mathbb{BST}^+ -formulae are Boolean combinations of atoms $x = y \setminus z$, it suffices to check atoms. Fix an atom α .

If $\text{Vars}(\alpha) \cap W = \emptyset$, then none of its variables is modified by the replacement, so $M \models \alpha$ iff $M' \models \alpha$.

If $\text{Vars}(\alpha) \subseteq W$, then satisfaction of α depends only on the restrictions $M|_W$ and $M' = (M|_W)_T^S$. By Lemma 5, $M|_W \models \alpha$ iff $M'|_W \models \alpha$; hence, $M \models \alpha$ iff $M' \models \alpha$.

The claim follows by propositional logic. \square

With the same restriction on variables—namely, every literal lies wholly inside W or wholly outside it—Lemma 6 yields invariance of satisfiability under any finite sequence of BST-replacements restricted to W .

Corollary 1. *Let V be a collection of variables and let $W \subseteq V$. Let $M^{(0)}, \dots, M^{(m)}$ ($m \geq 1$) be set assignments over V such that, for each $i = 1, \dots, m$,*

$$M^{(i)} = (M^{(i-1)})_{T_i}^{S_i}$$

is a BST-replacement of $M^{(i-1)}$ restricted to W , for some nonempty S_i, T_i . Let Φ be a \mathbb{BST}^+ -formula such that for every literal ℓ occurring in Φ we have either $\text{Vars}(\ell) \subseteq W$ or $\text{DISJ}(\text{Vars}(\ell), W)$. Then, for all $0 \leq h, j \leq m$,

$$M^{(h)} \models \Phi \iff M^{(j)} \models \Phi.$$

4.2. \flat -Flat Models

Definition 7. *For every ordinal $\flat \geq 1$, a set assignment M over a collection V of variables is said to be \flat -flat if all sets in the realm $\cup\{Mv \mid v \in V\}$ of M have rank \flat .*

No membership atom $x \in y$ is satisfied by any \flat -flat set assignment:

Lemma 7. *Let M be a \flat -flat set assignment over a collection V of variables. Then, $Mx \notin My$ for any $x, y \in V$.*

Proof. By the \flat -flatness of M , for all $x \in V$, either $\text{rk}(Mx) = 0$ (when $Mx = \emptyset$), or $\text{rk}(Mx) = \flat + 1$ (when $Mx \neq \emptyset$). In either case, $\text{rk}(Mx) \neq \flat$, since \flat -flatness presupposes $\flat \geq 1$. Thus, $Mx \notin \cup\{My \mid y \in V\}$ for all $x \in V$. \square

A satisfiable formula of \mathbb{BST}^+ always admits a \flat -flat model, for sufficiently large \flat . This is proved in the next lemma.

Lemma 8. *Let Φ be a satisfiable \mathbb{BST}^+ -formula with n distinct variables. Then, for every ordinal $\flat \geq n + 1$, Φ admits a \flat -flat set model of cardinality at most $2^n - 1$.*

Proof. Let Φ be satisfiable, M a model of Φ , and $\Sigma_M = \{\sigma_0, \dots, \sigma_{m-1}\}$ its induced partition. Fix an ordinal $b \geq n + 1$, where $n = |\text{Vars}(\Phi)|$.

We associate to each block $\sigma_i \in \Sigma_M$ a distinct set s_i of rank b , as follows:

1. For every block σ_i that already contains elements of rank b , choose one such element as s_i .
2. For the remaining blocks, choose s_i among elements of rank b (possibly belonging to other blocks), ensuring that the chosen s_i are all distinct and different from those fixed in step 1

This is possible because the number of blocks is at most exponential in $n = |\text{Vars}(\Phi)|$, while the collection $\mathcal{W}_b^\#$ of sets of rank exactly b is large enough to supply that many distinct elements. Indeed, since each block of Σ_M is nonempty and determined by a membership pattern over the n variables, it follows that $|\Sigma_M| \leq 2^n - 1$. On the other hand, the functions $\kappa \mapsto 2^\kappa - \kappa$ and $\kappa \mapsto |\mathcal{V}_\kappa|$ are strictly increasing for $\kappa \geq 1$, so from (2) and (3) it follows that, for every integer $\ell \geq n + 1$,

$$\begin{aligned} |\mathcal{W}_\ell^\#| &= |\mathcal{V}_{\ell+1}| - |\mathcal{V}_\ell| = 2^{|\mathcal{V}_\ell|} - |\mathcal{V}_\ell| \\ &\geq 2^{|\mathcal{V}_{n+1}|} - |\mathcal{V}_{n+1}| = |\mathcal{V}_{n+2}| - |\mathcal{V}_{n+1}| = |\mathcal{W}_{n+1}^\#| \geq 2^n. \end{aligned}$$

Therefore, whether b is finite or infinite, $\mathcal{W}_b^\#$ contains at least 2^n elements of rank b , enough to assign distinct representatives to all the blocks of Σ_M .

The argument that follows requires that the distinct blocks $\sigma_0, \dots, \sigma_{m-1}$ of Σ_M be ordered so that all blocks containing elements of rank b come first. Based on this ordering, we recursively construct a sequence of assignments M_0, \dots, M_m as follows:

$$\begin{cases} M_0 := M, \\ M_{i+1} := (M_i)_{\{s_i\}}^{\sigma_i}, \quad i = 0, \dots, m - 1. \end{cases}$$

For the blocks considered in step 1, we have $\{s_i\} \subseteq \sigma_i$; hence, $(M_i)_{\{s_i\}}^{\sigma_i}$ is trivially a BST-replacement of M_i from σ_i to $\{s_i\}$. For the blocks in step 2, the choice of s_i guarantees that $\{s_i\}$ is disjoint from every block not yet replaced, so the BST-conditions also hold.

Therefore, each M_{i+1} is a BST-replacement of M_i , and since M_0 is a model of Φ , induction and Lemma 5 yield that every M_{i+1} is as well. In particular, M_m is a model of Φ with support $\{s_0, \dots, s_{m-1}\}$, consisting solely of sets of rank b . Therefore, M_m provides the desired b -flat model of Φ , with cardinality at most $2^n - 1$. \square

Remark 4. Building on an argument developed in [12] (see also [13]), one can show that every satisfiable formula Φ of BST^+ admits a model of cardinality less than the number of its distinct variables. Accordingly, in the proof of Lemma 8, the initial model M may be taken with support of size at most $n - 1$, where $n = |\text{Vars}(\Phi)|$. The remainder of the argument then produces a b -flat model of cardinality at most $n - 1$.

This observation leads to the following sharper result.

Corollary 2. Let Φ be a satisfiable BST^+ -formula with n distinct variables. Then, for every ordinal $b \geq n + 1$, Φ admits a b -flat set model of cardinality at most $n - 1$. \square

5. Existential Inexpressibility of $z = \{x\}$ in BST^+

We investigate whether atoms of the form $z = \{x\}$ can be existentially expressed in BST^+ . If this were possible, then membership atoms $x \in y$ would also be expressible: indeed, the presence of ‘ $\{\blacksquare\}$ ’ allows one to derive ‘ \in ’ as a definable construct.

Lemma 9. $\models x \in y \iff (\exists z)(z = \{x\} \wedge z \subseteq y)$, if z is distinct from the variables x, y .

Proof. If $M \models z = \{x\} \wedge z \subseteq y$ holds, then $M \models x \in y$, since $Mx \in Mz \subseteq My$.

Conversely, if $M \models x \in y$ holds, extend M by putting $Mz = \{Mx\}$; then, $Mx \in My$ yields $M \models z = \{x\} \wedge z \subseteq y$, so that $M \models (\exists z)(z = \{x\} \wedge z \subseteq y)$. \square

In what follows we show that membership atoms $x \in y$ are *not* existentially expressible in \mathbb{BST}^+ . Therefore, by Lemma 9, atoms of the form $z = \{x\}$ are not existentially expressible either. Specifically, we prove that every satisfiable formula Ψ of \mathbb{BST}^+ admits a *flat* model M , namely, a model whose support consists solely of elements all having the same positive rank. As a consequence, $Mx \notin My$ for any $x, y \in \text{Vars}(\Psi)$, and thus $M \not\models x \in y$.

We are now ready to prove that membership atoms $x \in y$ —and hence, by Lemma 9, atoms of the form $z = \{x\}$ —are not existentially expressible in \mathbb{BST}^+ .

Theorem 1. *The atom $x \in y$ is not existentially expressible in \mathbb{BST}^+ .*

Proof. By way of contradiction, assume that $x \in y$ is existentially expressible by a formula $\Psi(x, y, z)$ of \mathbb{BST}^+ involving only atoms of the form $x' = y' \setminus z'$, i.e.,

$$\models x \in y \iff (\exists z) \Psi(x, y, z). \tag{9}$$

Since $x \in y$ is trivially satisfiable, so are—by (9)— $(\exists z) \Psi(x, y, z)$ and $\Psi(x, y, z)$. Thus, by Lemma 8, $\Psi(x, y, z)$ is modeled by a *b-flat* set assignment M^* ; hence, by Lemma 7, $M^*x \notin M^*y$ holds. It follows that $M^* \models x \notin y \wedge (\exists z) \Psi(x, y, z)$ holds, and therefore $M^* \not\models (\exists z) \Psi(x, y, z) \rightarrow x \in y$, which contradicts (9). \square

6. $\mathcal{O}(n^2)$ -Expressibility in \mathbb{BST}^+ of Singleton-Atom Conjunctions

In accordance with Definition 3, we shall prove that any conjunction $\psi(x)$ of atoms of the form $x = \{y\}$ is $\mathcal{O}(n^2)$ -expressible from \mathbb{BST} into \mathbb{BST}^+ by exhibiting a map

$$\langle \varphi(y), \psi(x) \rangle \mapsto \Xi_\varphi^\psi(x, y, z),$$

computable in quadratic time, such that conditions (a)–(c) of Definition 3 hold. Here, $\varphi(y)$ ranges over \mathbb{BST} -conjunctions, and the variables in z are distinct from those in x and in y .

Thus, let $\varphi(y)$ and $\psi(x)$ be of the said form. For each variable $x \in \text{Vars}(\varphi \wedge \psi)$ (note that $\text{Vars}(\varphi \wedge \psi) = \text{Vars}(\varphi) \cup \text{Vars}(\psi)$), introduce a fresh auxiliary variable \tilde{x} , chosen distinct from all others. These variables will be used to mimic $x \in y$ via the relation $\tilde{x} \subsetneq \tilde{y}$. Also, let z be an additional fresh auxiliary variable, to be interpreted as the empty set. Then, we put

$$\begin{aligned} \Xi_\varphi^\psi := & z = z \setminus z \wedge \bigwedge_{x=\{y\} \text{ in } \psi} x \not\subseteq y \wedge \\ & \bigwedge_{\substack{x=\{y\} \text{ in } \psi \\ x'=\{y'\} \text{ in } \psi}} (y = y' \iff x = x') \wedge \bigwedge_{\substack{x=\{y\} \text{ in } \psi \\ v \in \text{Vars}(\varphi \wedge \psi)}} (\neg \text{DISJ}(x, v) \rightarrow x \subseteq v) \wedge \\ & \bigwedge_{\substack{x=\{y\} \text{ in } \psi \\ v \in \text{Vars}(\varphi \wedge \psi)}} (\neg \text{DISJ}(x, v) \rightarrow \tilde{y} \subsetneq \tilde{v}) \wedge \bigwedge_{x, y \in \text{Vars}(\varphi \wedge \psi)} (x = y \rightarrow \tilde{x} = \tilde{y}) \end{aligned} \tag{10}$$

(thus, the list z of variables in Definition 3 consists of the collection \tilde{x} of all auxiliary set variables \tilde{x} together with z).

Remark 5. In the formula Ξ_φ^ψ , the following shorthands are used:

$$\begin{aligned} \text{DISJ}(x, v) &\stackrel{\text{def}}{\iff} x = x \setminus v, \\ x = v &\stackrel{\text{def}}{\iff} x = v \setminus z, \\ x \subseteq v &\stackrel{\text{def}}{\iff} x \setminus v = z, \\ x \subsetneq v &\stackrel{\text{def}}{\iff} x \setminus v = z \wedge v \setminus x \neq z. \end{aligned}$$

Here, z is the auxiliary variable intended to denote the empty set. Indeed, for every model M of Ξ_φ^ψ , the conjunct $z = z \setminus z$ ensures $Mz = \emptyset$, from which the following equivalences follow:

$$\begin{aligned} M \models \text{DISJ}(x, v) &\iff \text{DISJ}(Mx, Mv); \\ M \models x = v &\iff Mx = Mv; \\ M \models x \subseteq v &\iff Mx \subseteq Mv; \\ M \models x \subsetneq v &\iff Mx \subsetneq Mv. \end{aligned} \tag{11}$$

For instance, the second equivalence can be justified as follows:

$$M \models x = v \iff M \models x = v \setminus z \iff Mx = Mv \setminus Mz = Mv.$$

These shorthands are introduced purely for readability: their syntactic form mirrors the underlying semantics, making formulae such as Ξ_φ^ψ easier to parse at a glance without repeatedly unfolding their full definitions.

By the preceding remark, Ξ_φ^ψ is a formula of BST^+ ; indeed, it is just a conjunction of a rather simple form. As a preliminary step, we record a basic semantic property that will be used repeatedly.

Lemma 10. For every model M of Ξ_φ^ψ , if $x = \{y\}$ occurs in ψ , then

$$Mx \neq \emptyset \quad \text{and} \quad \text{DISJ}(Mx, My).$$

Proof. Let $M \models \Xi_\varphi^\psi$, and let $x = \{y\}$ be a literal in ψ . Since Ξ_φ^ψ includes the conjuncts

$$z = z \setminus z, \quad x \not\subseteq y, \quad \text{and} \quad (\neg \text{DISJ}(x, y) \rightarrow x \subseteq y),$$

the equivalences (11) in Remark 5 imply

$$Mx \not\subseteq My \quad \text{and} \quad \neg \text{DISJ}(Mx, My) \implies Mx \subseteq My.$$

From this, it follows that $Mx \neq \emptyset$ and $\text{DISJ}(Mx, My)$, as claimed. \square

6.1. Design and Analysis of the Translation Algorithm

We first present an explicit procedure that constructs Ξ_φ^ψ from the input conjunction $\varphi \wedge \psi$; see Algorithm 1. We then analyze its running time, thereby establishing point (a) of Definition 3.

We use the auxiliary program variables V and Ξ in the pseudocode below to store, respectively, the list of variables encountered in $\varphi \wedge \psi$ and the list of generated conjuncts.

Algorithm 1 Construction of Ξ_φ^ψ

- 1: Initialize V as an empty list of set variables;
 - 2: Initialize Ξ as an empty list of conjuncts;
 - 3: **for** each set variable x that appears in φ **do**
 - 4: add x to V ;
 - 5: **for** each conjunct $x = \{y\}$ in ψ **do**
 - 6: add x and y to V ;
 - 7: add $x \not\subseteq y$ to Ξ ;
 - 8: **for** each conjunct $x = \{y\}$ in ψ **do**
 - 9: **for** each $v \in V$ **do**
 - 10: add $(\neg\text{DISJ}(x, v) \rightarrow x \subseteq v) \wedge (\neg\text{DISJ}(x, v) \rightarrow \tilde{y} \subsetneq \tilde{v})$ to Ξ ;
 - 11: **for** each pair $x = \{y\}, x' = \{y'\}$ of distinct conjuncts in ψ **do**
 - 12: add $(y = y' \leftrightarrow x = x')$ to Ξ ;
 - 13: **for** all $x, y \in V$ **do**
 - 14: add $(x = y \rightarrow \tilde{x} = \tilde{y})$ to Ξ ;
 - 15: add $z = z \setminus z$ to Ξ ;
 - 16: **return** Ξ .
-

The running time of Algorithm 1 is bounded as follows.

Lemma 11. *The formula Ξ_φ^ψ can be constructed from φ and ψ in time $\mathcal{O}(|\varphi \wedge \psi|^2)$.*

Proof. In order to prove the lemma, it is convenient to spell out in Algorithm 1 the procedure that constructs Ξ_φ^ψ from the conjunction $\varphi \wedge \psi$. We maintain the program variable V as a list of variables and Ξ as a list of conjuncts, so that appending a variable or a conjunct takes constant time. Moreover, each conjunct appended to Ξ has constant size.

The **for**-loop at lines 3–4 can be performed in $\Theta(|\varphi|)$ time, where $|\varphi|$ denotes the total length of the conjunction φ . Similarly, the **for**-loop at lines 5–7 can be performed in $\Theta(|\psi|)$ time.

Letting m be the length of the list V after the execution of the **for**-loops 3–4 and 5–7, we have $m \leq |\varphi| + |\psi|$.

Upper bound (running time). The only nontrivial operations performed by Algorithm 1 are iterating through the **for**-loops and appending variables/conjuncts to the lists V and Ξ . By assumption, each append takes constant time, and each conjunct appended to Ξ has constant size. Hence, the running time is asymptotically proportional to the total number of loop iterations.

As noted, lines 5–7 perform $\Theta(|\psi|)$ iterations. Lines 8–10 perform $\Theta(|\psi| \cdot m)$ iterations (for each conjunct in ψ , the inner loop ranges over all v in V). Lines 11–12 perform $\Theta(|\psi|^2)$ iterations, since they range over all pairs of distinct conjuncts in ψ . Finally, lines 13–14 perform $\Theta(m^2)$ iterations, since they range over all pairs (x, y) with x and y in V , whereas line 15 clearly takes constant time.

Therefore, the overall running time is

$$\Theta(|\varphi| + |\psi| + |\psi| \cdot m + |\psi|^2 + m^2).$$

Since $m \leq |\varphi| + |\psi| \leq |\varphi \wedge \psi|$ and $|\psi| \leq |\varphi \wedge \psi|$, we conclude that Algorithm 1 runs in time

$$\mathcal{O}(|\varphi \wedge \psi|^2),$$

as claimed. \square

Before proving the semantic conditions (b) and (c) of Definition 3 for the translation map (10), we briefly illustrate Ξ_φ^ψ on two concrete examples. We then return to the main line of the argument and show that

- If $\varphi(\mathbf{y}) \wedge \Xi_\varphi^\psi(\mathbf{x}, \mathbf{y}, \mathbf{z})$ is satisfiable, then so is $\varphi(\mathbf{y}) \wedge \psi(\mathbf{x})$;
- Every model of $\varphi(\mathbf{y}) \wedge \psi(\mathbf{x})$ can be extended into a model of $\Xi_\varphi^\psi(\mathbf{x}, \mathbf{y}, \mathbf{z})$.

As above, this is where the list \mathbf{z} of variables consists of the collection $\tilde{\mathbf{x}}$ of all auxiliary set variables \tilde{x} together with \mathbf{z} . This will prove that conditions (b) and (c) of Definition 3 hold and hence that every singleton-atom conjunction is $\mathcal{O}(n^2)$ -expressible from BST into BST⁺.

6.2. Illustrating the Translation Ξ_φ^ψ on Examples

We start with an unsatisfiable input, showing how Ξ_φ^ψ preserves unsatisfiability, and then consider a satisfiable instance illustrating the model extension mechanism used for condition (c).

Example 3. *The conjunction $\varphi \wedge \psi$, where $\varphi := y = x \setminus z$ and $\psi := x = \{y\} \wedge y = \{z\}$, is not satisfiable; in fact, any set assignment M satisfying this formula is such that $M_y = M_x \setminus M_z$, $M_x = \{M_y\}$, and $M_y = \{M_z\}$; thus, $M_y = M_z$, and hence $M_z = \{M_z\}$, a contradiction.*

Our translation Ξ_φ^ψ comprises

$$y \neq \emptyset \wedge x \not\subseteq y \wedge (\neg \text{DISJ}(x, y) \longrightarrow x \subseteq y).$$

Any model M for $\varphi \wedge \Xi_\varphi^\psi$ is such that $M_y \subseteq M_x$ (since $M_y = M_x \setminus M_z$), $M_y \neq \emptyset$, and $M_x \not\subseteq M_y$ hold, so that $\neg \text{DISJ}(M_x, M_y)$; but then $M_x \subseteq M_y$ must hold, which conflicts with $M_x \not\subseteq M_y$.

The next example complements the previous one: it is satisfiable and highlights the extension of a model of $\varphi \wedge \psi$ to the auxiliary variables so as to satisfy Ξ_φ^ψ (cf. Section 6.4).

As we will prove in Section 6.4, satisfiability carries over from $\varphi \wedge \psi$ to $\varphi \wedge \Xi_\varphi^\psi$. Here is an interesting example of this:

Example 4. *Consider the conjunction $\varphi \wedge \psi$, where*

$$\varphi := x = y \setminus y' \wedge x = z \setminus z', \quad \text{and} \quad \psi := z = \{y\}.$$

This formula is satisfiable: for instance, it is satisfied by every set assignment M over $\text{Vars}(\varphi \wedge \psi)$ of the form

$$M_x := \emptyset, \quad M_y := s, \quad M_z := \{s\}, \quad M_{y'} := s \cup s', \quad M_{z'} := \{s\} \cup s'', \quad (12)$$

where s , s' , and s'' are any well-founded sets.

It can easily be checked that, for every set assignment M of the form (12), the extension M^+ of M over the auxiliary variables \tilde{v} , for $v \in \text{Vars}(\varphi \wedge \psi)$ and where

$$M^+ \tilde{v} := \begin{cases} s \cup \{s\} & \text{if } s \in Mv \\ \emptyset & \text{otherwise,} \end{cases}$$

satisfies Ξ_φ^ψ , so that $M \models (\exists \tilde{x}, \tilde{y}, \tilde{z}, \tilde{y}', \tilde{z}') \Xi_\varphi^\psi$ holds. Thus, to show that condition (c) of Definition 3 is satisfied, namely, that $\models (\varphi \wedge \psi) \longrightarrow (\exists \tilde{x}, \tilde{y}, \tilde{z}, \tilde{y}', \tilde{z}') \Xi_\varphi^\psi$ holds, it is enough to check that the conjunction $\varphi \wedge \psi$ is satisfied by set assignments of the form (12) only. Let then \bar{M} be any model for $\varphi \wedge \psi$. Since $\bar{M}_x \subseteq \bar{M}_z$ and $\bar{M}_z = \{\bar{M}_y\}$, either $\bar{M}_x = \emptyset$ or $\bar{M}_x = \{\bar{M}_y\}$ holds. The latter case can be readily ruled out, for in view of $\bar{M}_x \subseteq \bar{M}_y$ it would follow $\bar{M}_y \in \bar{M}_y$, which

is untenable in the realm of well-founded sets. Thus, $\overline{M}x = \emptyset$ must hold. Letting $s := \overline{M}y$ —so that $\overline{M}z = \{s\}$ —, since $\overline{M}y \subseteq \overline{M}y'$ and $\overline{M}z \subseteq \overline{M}z'$, we have $\overline{M}y' = s \cup s'$ and $\overline{M}z' = \{s\} \cup s''$, for some sets s' and s'' , and therefore \overline{M} has the form (12).

6.3. Satisfiability Preservation

We first establish that if $\varphi \wedge \Xi_\varphi^\psi$ is satisfiable, then so is $\varphi \wedge \psi$, thereby fulfilling condition (b) of Definition 3.

Let M be a model of $\varphi \wedge \Xi_\varphi^\psi$. In order to convert it into a model of $\varphi \wedge \psi$, we can assume that M is b -flat over $\text{Vars}(\varphi \wedge \Xi_\varphi^\psi)$ for some integer

$$b > \max(|\text{Vars}(\varphi \wedge \Xi_\varphi^\psi)|, |\psi|).$$

Indeed, on the one hand, Lemma 8 enables us to do so; on the other hand, Lemma 7 tells us that such an M does not model any of the atoms $x = \{y\}$ in ψ .

Set $M^{(0)} := M$, and let m be the number of singleton conjuncts in ψ . For $i = 1, \dots, m$, we select an atom $x_i = \{y_i\}$ in ψ and lift $M^{(i-1)}$ to a model $M^{(i)}$ of φ such that every conjunct of ψ already satisfied in $M^{(i-1)}$ remains satisfied in $M^{(i)}$, and the selected conjunct is also satisfied in $M^{(i)}$. At the end of this iterative process, all conjuncts of ψ will be satisfied. More precisely, if ψ_i denotes the conjunction of the atoms selected up to step i , then $\psi_1 \wedge \dots \wedge \psi_m$ coincides with ψ (possibly reordered and with repetitions).

In our set up, each ψ_i must hence comprise a conjunct $x_i = \{y_i\}$ from ψ not appearing in any ψ_j with $j < i$. The selection of $x_i = \{y_i\}$ determines the transformation from $M^{(i-1)}$ to $M^{(i)}$ and is made as follows: among the conjuncts of ψ that are not already present in $\psi_1 \wedge \dots \wedge \psi_{i-1}$, select one that is *minimal* with respect to the ordering \prec_ψ^M on the conjuncts of ψ , defined below:

Definition 8. The relation \prec_ψ^M is the minimal transitively closed relation over the conjuncts of ψ such that, for all singleton atoms $x = \{y\}$ and $x' = \{y'\}$ in ψ , one has

$$(x = \{y\}) \prec_\psi^M (x' = \{y'\}) \quad \text{whenever} \quad \neg \text{DISJ}(Mx, My').$$

This definition enforces the following:

Lemma 12. The relation \prec_ψ^M is a strict partial order.

Proof. By definition \prec_ψ^M is transitive, so it remains only to prove that it is acyclic.

Suppose, by way of contradiction, that \prec_ψ^M contains some cycle. Then, there are atoms $\ell_0, \ell_1, \dots, \ell_h$ of ψ with

$$\ell_0 \prec^M \ell_1 \prec^M \dots \prec^M \ell_{m-1} \prec^M \ell_h = \ell_0,$$

where each ℓ_i is of the form $x_i = \{y_i\}$ for $i = 0, 1, \dots, h$. By the definition of \prec_ψ^M , we have

$$\neg \text{DISJ}(Mx_0, My_1), \dots, \neg \text{DISJ}(Mx_{h-1}, My_0).$$

Since $M \models \Xi_\varphi^\psi$, we know in particular that

$$M \models \bigwedge_{\substack{x = \{y\} \text{ in } \psi \\ v \in \text{Vars}(\varphi \wedge \psi)}} (\neg \text{DISJ}(x, v) \longrightarrow \tilde{y} \subsetneq \tilde{v}),$$

and so we obtain the strict chain

$$M\tilde{y}_0 \subsetneq M\tilde{y}_1 \subsetneq \cdots \subsetneq M\tilde{y}_{h-1} \subsetneq M\tilde{y}_0,$$

which yields the contradiction $M\tilde{y}_0 \subsetneq M\tilde{y}_0$.

Therefore, \prec_{ψ}^M has no cycles and is thus acyclic. Together with its transitivity by definition, this shows that \prec_{ψ}^M is a strict partial order, as claimed. \square

From the previous lemma, we may arrange the literals in ψ as

$$\ell_1, \dots, \ell_m,$$

where each ℓ_i has the form $x_i = \{y_i\}$, in such a way that the sequence complies with the strict partial order \prec_{ψ}^M ; that is, for all $1 \leq i < j \leq m$, we have $\ell_j \not\prec_{\psi}^M \ell_i$.

Such an ordering will guide the iterative process: it is not arbitrary but rather ensures that the satisfaction of atoms established in earlier steps is preserved in all subsequent steps.

The iterative process proceeds as follows: at each step i , an atom $\ell_i := x_i = \{y_i\}$ is selected from ψ , as specified above, so that ℓ_i is minimal, in ψ with all atoms selected at previous steps removed, with respect to \prec_{ψ}^M ; that is, for all $i, j \in \{1, \dots, m\}$,

$$\text{if } i < j \text{ then } \ell_j \not\prec_{\psi}^M \ell_i, \tag{13}$$

Then, letting

$$X_i := M^{(i-1)}x_i \quad \text{and} \quad Y_i := M^{(i-1)}y_i, \tag{14}$$

we define

$$M^{(i)} := (M_{|W}^{(i-1)})_{\{Y_i\}}^{X_i}, \tag{15}$$

where

$$W := \text{Vars}(\phi \wedge \psi).$$

That is, $M^{(i)}$ is obtained from $M^{(i-1)}$ by replacing X_i with $\{Y_i\}$ while restricting to the variables in $\text{Vars}(\phi \wedge \psi)$.

Remark 6. Throughout the remainder of this subsection, we may occasionally omit to mention the explicit restriction to $W = \text{Vars}(\phi \wedge \psi)$ and, for instance, write $(M^{(i-1)})_{\{Y_i\}}^{X_i}$ in place of $(M_{|W}^{(i-1)})_{\{Y_i\}}^{X_i}$.

The following lemma establishes a useful dichotomy on the possible ranks of the values assigned to variables during the replacement process.

Lemma 13. For each $i = 1, \dots, m$, we have

$$\text{either } \text{rk}(M^{(i-1)}y_i) \leq i - 1 < m \leq b \quad \text{or} \quad \text{rk}(M^{(i-1)}y_i) > b.$$

Proof. Preliminarily, it is immediate to check that for every $i = 1, \dots, m$ we have

$$\text{supp}(M^{(i)}) \subseteq \text{supp}(M^{(i-1)}) \cup \{M^{(i-1)}y_i\}. \tag{16}$$

We proceed by induction on i .

Base case ($i = 1$). Before any replacement we are at $M^{(0)}$, which by construction is b -flat; hence, every member of every $M^{(0)}v$ has rank b . Thus, $\text{rk}(M^{(0)}y_1)$ is either 0 or $b + 1$, so the stated dichotomy holds.

Inductive step. Fix $i > 1$ and assume the statement holds for all indices $1, \dots, i - 1$. By iterating (16) we obtain

$$\text{supp}(M^{(i-1)}) \subseteq \text{supp}(M^{(0)}) \cup \{M^{(t-1)}y_t \mid 1 \leq t \leq i-1\}. \tag{17}$$

Let $s \in M^{(i-1)}y_i$ be arbitrary. Then, $s \in \text{supp}(M^{(i-1)})$, so by (17) either

- (a) $s \in \text{supp}(M^{(0)})$, in which case $\text{rk}(s) = b$ (since $M^{(0)}$ is b -flat);
- (b) $s = M^{(t-1)}y_t$ for some $t \in \{1, \dots, i-1\}$. By the inductive hypothesis at index t , either $\text{rk}(s) = \text{rk}(M^{(t-1)}y_t) \leq t-1 \leq i-2$, or $\text{rk}(s) > b$.

If case (a) occurs for some member s of $M^{(i-1)}y_i$, then $\text{rk}(M^{(i-1)}y_i) \geq b+1 > b$. If case (b) occurs with $\text{rk}(s) > b$ for some member s , again $\text{rk}(M^{(i-1)}y_i) > b$. Otherwise, all members of $M^{(i-1)}y_i$ (if any) have rank at most $i-2$, whence

$$\text{rk}(M^{(i-1)}y_i) \leq (i-2) + 1 = i-1.$$

Finally, by our choice of b we have $m \leq b$, so in particular $i-1 < m \leq b$. This yields exactly the stated dichotomy for i .

This completes the proof. \square

Our goal is to prove that the final assignment $M^{(m)}$ is a model of $\varphi \wedge \psi$. To this end, we first establish a stability lemma, from which it follows that after step i the sets assigned to the variables in the atom $x_i = \{y_i\}$ remain unchanged in all later assignments $M^{(t)}$, for $t > i$.

Lemma 14. *For every $i = 1, \dots, m$, letting $x_i = \{y_i\}$ be the atom of ψ selected at step i , we have the following:*

- (a) $M^{(i)}$ is a BST-replacement of $M^{(i-1)}$ restricted to $\text{Vars}(\varphi \wedge \psi)$ such that $M^{(i)} \models \varphi \wedge \Xi_\varphi^\psi$;
- (b) for each j with $1 \leq j < i$,
 - (b₁) $M^{(j)}y_j = M^{(i)}y_j$;
 - (b₂) $M^{(j)}x_j = M^{(i)}x_j$;
- (c) $M^{(i)}y_i = M^{(i-1)}y_i$;
- (d) $M^{(i)}x_i = \{M^{(i)}y_i\}$.

Proof Sketch. The proof proceeds by induction on i , relying on the preservation properties of BST-replacements and the constraints encoded in Ξ_φ^ψ . The detailed argument, which involves a careful case analysis for clauses (a)–(d), is rather technical. For readability, we defer the full proof to Appendix B. \square

The replacement construction of Lemma 14 immediately entails a monotonicity condition on the cardinalities of the assigned sets, stated next.

Corollary 3. *For every $i = 1, \dots, m$ and every $v \in \text{Vars}(\varphi \wedge \Xi_\varphi^\psi)$, we have*

$$|M^{(i)}v| \leq |M^{(i-1)}v|.$$

Consequently, for every $v \in \text{Vars}(\varphi \wedge \Xi_\varphi^\psi)$, we have

$$|M^{(m)}v| \leq |M^{(0)}v|.$$

Proof. Fix $i \in \{1, \dots, m\}$ and $v \in \text{Vars}(\varphi \wedge \Xi_\varphi^\psi)$. If $M^{(i)}v = M^{(i-1)}v$, the claim is immediate. Otherwise, since $M^{(i)}$ is obtained from $M^{(i-1)}$ by a BST-replacement (with respect to the pair

$(M^{(i-1)}x_i, \{M^{(i-1)}y_i\})$, the only way $M^{(i-1)}v$ can change is when $\emptyset \neq M^{(i-1)}x_i \subseteq M^{(i-1)}v$, in which case

$$M^{(i)}v = (M^{(i-1)}v \setminus M^{(i-1)}x_i) \cup \{M^{(i-1)}y_i\}.$$

Plainly, $|M^{(i-1)}x_i| \geq |\{M^{(i-1)}y_i\}|$; hence, $|M^{(i)}v| \leq |M^{(i-1)}v|$. Applying this inequality successively for $i = 1, 2, \dots, m$ yields

$$|M^{(m)}v| \leq |M^{(m-1)}v| \leq \dots \leq |M^{(0)}v|,$$

as required. \square

We are now ready to prove that $M^{(m)}$ satisfies $\varphi \wedge \psi$. By Lemma 14(a), we have

$$M^{(m)} \models \varphi. \tag{18}$$

Moreover, Lemma 14(c) guarantees that, for every i with $1 \leq i \leq m$,

$$M^{(i)}x_i = \{M^{(i)}y_i\}. \tag{19}$$

Thus, to establish $M^{(m)} \models \psi$, it suffices to prove that, for all i ,

$$M^{(m)}x_i = \{M^{(m)}y_i\}. \tag{20}$$

If $i = m$, then (19) immediately gives (20). If $i < m$, clauses (b₁) and (b₂) of Lemma 14 yield $M^{(i)}y_i = M^{(m)}y_i$ and $M^{(i)}x_i = M^{(m)}x_i$, respectively. Combining these with (19), we again obtain (20). Hence, $M^{(m)} \models \psi$. From this and (18), it follows that $M^{(m)} \models \varphi \wedge \psi$, thus completing the proof of satisfiability preservation.

6.4. Model Extension

We next prove that every model of $\varphi \wedge \psi$ can be extended into a model of $\varphi \wedge \Xi_\varphi^\psi$, which corresponds to condition (c) of Definition 3.

Let $\varphi \wedge \psi$ be satisfiable, and let M be a model of it. Define \prec as the minimal transitive relation on $\text{Vars}(\varphi \wedge \psi)$ such that

$$y \prec x \text{ whenever } My \in Mx, \quad \text{for all } x, y \in \text{Vars}(\varphi \wedge \psi).$$

By set well-foundedness, \prec is irreflexive; together with transitivity, this makes \prec a strict ordering on $\text{Vars}(\varphi \wedge \psi)$.

Extend M to the auxiliary variables \tilde{v} , for each $v \in \text{Vars}(\varphi \wedge \psi)$, by setting for each of them:

$$M\tilde{v} := \{Mx \mid x \prec v\}.$$

Moreover, extend M to the auxiliary variable z by stipulating

$$Mz := \emptyset,$$

so that in particular

$$M \models z = z \setminus z. \tag{21}$$

We prove that $M \models \Xi_\varphi^\psi$ holds.

Preliminarily, note that Mx is a singleton for every singleton atom $x = \{y\}$ in ψ . It then follows directly that

$$M \models \bigwedge_{\substack{x=\{y\} \text{ in } \psi \\ v \in \text{Vars}(\varphi \wedge \psi)}} (\neg \text{DISJ}(x, v) \longrightarrow x \subseteq v). \tag{22}$$

Moreover, for each singleton atom $x = \{y\}$ in ψ , the acyclicity of the membership relation yields $Mx \not\subseteq My$, and so $M \models x \not\subseteq y$. Therefore,

$$M \models \bigwedge_{x=\{y\} \in \psi} x \not\subseteq y. \tag{23}$$

In addition, for each pair of singleton atoms $x = \{y\}$ and $x' = \{y'\}$ in ψ , we plainly have

$$M \models y = y' \iff x = x',$$

whence

$$M \models \bigwedge_{\substack{x=\{y\} \in \psi \\ x'=\{y'\} \in \psi}} (y = y' \iff x = x'). \tag{24}$$

Now, consider any implication of the form

$$\neg \text{DISJ}(x, v) \longrightarrow \tilde{y} \subseteq \tilde{v},$$

with $x = \{y\}$ in ψ and $v \in \text{Vars}(\varphi \wedge \psi)$, and assume that $M \models \neg \text{DISJ}(x, v)$. Since $Mx = \{My\}$, this yields $Mx \subseteq Mv$; hence, $My \in Mv$. Thus, $y \prec v$. By the transitivity and strictness of \prec , we obtain

$$M\tilde{y} = \{My' \mid y' \prec y\} \subseteq \{My' \mid y' \prec v\} = M\tilde{v},$$

namely, $M \models \tilde{y} \subseteq \tilde{v}$.

Hence, by the arbitrariness of $x = \{y\}$ in ψ and $v \in \text{Vars}(\varphi \wedge \psi)$, it follows that

$$M \models \bigwedge_{\substack{x=\{y\} \in \psi \\ v \in \text{Vars}(\varphi \wedge \psi)}} (\neg \text{DISJ}(x, v) \longrightarrow \tilde{y} \subseteq \tilde{v}). \tag{25}$$

Finally, consider the implications

$$x = y \longrightarrow \tilde{x} = \tilde{y},$$

with $x, y \in \text{Vars}(\varphi \wedge \psi)$. Assume $Mx = My$. Suppose $v \prec x$. By definition of \prec , there exists a finite chain of variables v_1, \dots, v_n such that

$$v = v_1, \quad Mv_i \in Mv_{i+1} \text{ for each } i, \quad v_n = x.$$

Hence,

$$Mv = Mv_1 \in \dots \in Mv_n = Mx = My,$$

which yields $v \prec y$. By a symmetric argument, $v \prec x$ follows from $v \prec y$. Therefore,

$$\{v \mid v \prec x\} = \{v \mid v \prec y\},$$

and thus $M\tilde{x} = M\tilde{y}$. Hence,

$$M \models \bigwedge_{x, y \in \text{Vars}(\varphi \wedge \psi)} (x = y \longrightarrow \tilde{x} = \tilde{y}). \tag{26}$$

Collecting (21)–(25), and the above argument, we see that all conjuncts of Ξ_φ^ψ are satisfied by M . Therefore,

$$M \models \Xi_\varphi^\psi,$$

which establishes the model extension property.

We have so extended a generic M such that $M \models \varphi \wedge \psi$ into a model of $\varphi \wedge \Xi_\varphi^\psi$; therefore, we get

$$\models (\varphi(\mathbf{y}) \wedge \psi(\mathbf{x})) \longrightarrow (\exists \mathbf{z}) \Xi_\varphi^\psi(\mathbf{x}, \mathbf{y}, \mathbf{z}),$$

where the list \mathbf{z} consists of the collection $\tilde{\mathbf{x}}$ of all auxiliary set variables \tilde{x} together with \mathbf{z} .

Finally, in view of Definition 3, combining the satisfiability preservation established in Section 6.3 with the model extension property of Section 6.4, and relying on Lemma 11, we obtain our main expressibility result:

Theorem 2. *Singleton-atom conjunctions are $\mathcal{O}(n^2)$ -expressible from \mathbb{BST} into \mathbb{BST}^+ .*

7. Conclusions: Related and Planned Work

The main contributions of this paper are as follows:

- The introduction of the notion of $\mathcal{O}(f)$ -expressibility across theories, refining the existing notion of existential expressibility, which, while useful, is too coarse for our purposes.
- The proof that atoms of the form $z = \{x\}$ are not existentially expressible in \mathbb{BST}^+ .
- The proof that, by contrast, any conjunction of such atoms is $\mathcal{O}(n^2)$ -expressible from \mathbb{BST} into \mathbb{BST}^+ , using a construction we call the *nested-to-flat translation*.

As noted in the Introduction, the authors’ interest in satisfiability mechanisms for restricted fragments of set theory arises primarily from the design and experimental use of the *ÆtnaNova* proof verifier.

As discussed ref. [3], ELEM is the core inference mechanism in *ÆtnaNova*, where it often operates implicitly alongside other methods. It is based on an enhanced form of multilevel syllogistic (ref. [1]), a decision procedure for checking the satisfiability of certain unquantified set-theoretic formulae. This procedure allows *ÆtnaNova* to establish that a statement follows from a given proof context by showing that its negation yields an unsatisfiable conjunction with earlier statements.

When parts of a proof involve constructs that fall outside the scope of ELEM’s built-in syllogistic, a preprocessing step replaces them with fresh variables, ensuring uniform treatment of identical structures. However, proof steps are rarely uniquely determined by prior lines and hints alone. This is because ELEM often generates a range of easy consequences from a given context, leaving the user free to choose which of these consequences to assert as the next proof step.

Since inference mechanisms akin to ELEM are likely to play a central role in proof technology, it is important to develop translation methods of low algorithmic cost—such as the nested-to-flat translation discussed above—that make their integration possible.

7.1. Envisaged Enhancements to the Nested-to-Flat Translation

We expect that the satisfiability-preserving translation treated in Section 6 can be tuned to theories richer than \mathbb{BST} , such as the following. Let \mathbb{BSTS} denote the collection of conjunctions of literals of the forms

$$x = y \setminus z, \quad x \neq y \setminus z, \quad x = \{y\}$$

and let \mathbb{BSTS}^+ be the Boolean closure of atoms of the forms

$$x = y \setminus z, \quad x = \{y\}.$$

Then, it seems plausible that conjunctions of atoms of the form

$$\text{FINITE}(x), \quad \text{DENUMERABLE}(x),$$

together with their negations—where $\text{FINITE}(x)$ expresses that the set assigned to x is finite and $\text{DENUMERABLE}(x)$ expresses that the set assigned to x is denumerable (i.e., finite or countably infinite)—are likewise $\mathcal{O}(n^2)$ -expressible from \mathbb{BSTS} to \mathbb{BSTS}^+ . Exploring this broader scenario is part of our ongoing research agenda.

7.2. *Towards Integrating Set-Theoretic, Boolean, and Numerical Constraint Reasoning*

Besides its independent interest, the nested-to-flat translation discussed so far can be seen as a preparatory step toward the quantitative approach to logical inference (cf. [14,15]), as specialized to the field of computable set theory (cf. Ch. 11 ref. [16]).

Our translation, in fact, lays the groundwork for reducing satisfiability problems about sets to satisfiability problems about nonnegative integers—or, if we move from the theory of hereditarily finite sets to a more general setting, to the language of the additive theory of cardinals (which is decidable; ref. [17]).

One such reduction, where the source language embodies a cardinality operator $|x|$ indicating how many elements belong to the set x , is presented in Sec. 11.1 ref. [16]. Let us briefly outline how it proceeds. We are given a conjunction ψ of literals of the forms

$$x = y \setminus z, \quad x = \{y\}, \quad h = |x|, \quad h = i + j, \quad h = 1.$$

Suppose we want to test ψ for satisfiability over the hereditarily finite sets. Here, x, y, z stand for set variables while i, j, h stand for numeric variables drawn from a disjoint infinite set of symbols and are intended to range over the nonnegative integers. Let \mathcal{X} be the collection of all set variables occurring in ψ ; let, moreover, \mathcal{Q} be the collection of all nonempty subsets Q of \mathcal{X} such that the assignment

$$B_Q(v) := \begin{cases} \{\emptyset\} & \text{if } v \in Q, \\ \emptyset & \text{if } v \in \mathcal{X} \setminus Q, \end{cases}$$

satisfies all equalities of the form $x = y \setminus z$ occurring in ψ . Associate two numeric variables ν_Q, ρ_Q with each Q in \mathcal{Q} . We have an algorithm that, given ψ , constructs a system \mathcal{C} of purely arithmetic constraints such that ψ is satisfiable over sets and natural numbers if and only if \mathcal{C} has a solution over the natural numbers. Those literals of the forms $h = i + j$ and $h = 1$ that were in ψ from the outset are retained in \mathcal{C} ; in addition to them, \mathcal{C} encompasses conditions specifying the intended meanings of ν_Q, ρ_Q relative to a (hypothetical) model $x \mapsto \xi_x$ of ψ , namely,

- Each ν_Q represents the cardinality of the set

$$\widehat{Q} := \left(\bigcap_{v \in Q} \xi_v \right) \setminus \left(\bigcup_{w \in \mathcal{X} \setminus Q} \xi_w \right);$$
- Each ρ_Q represents the position occupied by \widehat{Q} in a total ordering that extends the rank comparison relation.

The reduction just outlined can certainly be refined beyond the treatment ref. [16], as we will strive to do in the future. We also expect that it can be boosted with the treatment of explicit rank-related constructs, such as the comparison relator $\text{rk}(x) < \text{rk}(y)$.

Some reductions of the set-satisfiability problem to integer programming can be found ref. [18], whose line of research aimed at integrating linear programming problems and set constraint manipulation methods in a single logic programming language, as explained refs. [19,20] (the endeavor of integrating cardinality constraints into constraint logic program-

ming with sets has been carried out with a different approach, as reported ref. [21]). A technique for reducing the problem of multilevel syllogistic (cf. [3]) to propositional consistency testing was described ref. [18] (an account of it can also be found in Sec. 11.3 ref. [16]).

In Sec. 4.5 ref. [22], the authors show that satisfiability of nonrecursive Tarskian set constraints is decidable in nondeterministic double-exponential time by reducing the problem to a class of Diophantine constraints called *prequadratic*. They prove that satisfiability of prequadratic Diophantine constraints is decidable in nondeterministic exponential time and conjecture that it is in NP. If this conjecture holds, satisfiability of nonrecursive Tarskian constraints would be decidable in nondeterministic single-exponential time.

Along a related line of research, ref. [23] reduces quantifier-free constraints on sets involving cardinalities along with direct and inverse images of functions on sets, to systems of numerical constraints in linear integer arithmetic or of the form $x \leq y^d$, where d is a positive integer.

7.3. Difference Algebras

We take this opportunity to mention an issue concerning an alternative semantics for \mathbb{BST} , which, although only loosely related to the central theme of this paper, is nevertheless of independent interest. In his bachelor’s thesis *Unificazione semantica in strutture booleane* (‘Semantic unification in Boolean structures’), defended at the University of Trieste in 2020, Mattia Furlan isolated the valid formulae involving Boolean difference that are displayed in Figure 1.

Let us adopt the universal closures of these formulae as the axioms of a theory in quantificational first-order logic with equality. These axioms characterize an algebraic variety, whose instances we provisionally call *difference algebras*.

A natural open question is whether every difference algebra $\mathbb{D} = (\mathcal{D}, \setminus_{\mathcal{D}})$ is isomorphic to an algebra of the form $\mathbb{S} = (\mathcal{S}, \setminus)$, in which the operator ‘ \setminus ’ is interpreted as ordinary set-theoretic difference. In this case, \mathcal{S} must be a family of sets closed under difference, and hence under intersection, since

$$X \cap Y = X \setminus (X \setminus Y)$$

holds for all sets X, Y .

One might hope to settle this question by appealing to Stone’s celebrated representation theorem, which states that every Boolean algebra is isomorphic to a field of sets. However, we see no direct way to apply this theorem, since there exist difference algebras \mathbb{D} whose carrier \mathcal{D} is not closed under symmetric difference, viewed as an operation $\langle Y, Z \rangle \mapsto Y \triangle_{\mathcal{D}} Z$ satisfying, for all X, Y, Z in \mathcal{D} ,

$$X = Y \triangle_{\mathcal{D}} Z \iff (X \setminus_{\mathcal{D}} (Y \setminus_{\mathcal{D}} Z) = Z \setminus_{\mathcal{D}} Y \wedge Y \setminus_{\mathcal{D}} Z = X \setminus_{\mathcal{D}} Z).$$

Moreover, it is unclear how to embed an arbitrary difference algebra into one that forms a proper Boolean ring by virtue of satisfying this closure property.

Author Contributions: Conceptualization, D.C., A.D.D., P.M. and E.G.O.; Methodology, D.C., A.D.D., P.M. and E.G.O.; Writing – original draft, D.C., A.D.D., P.M. and E.G.O.; Writing – review & editing, D.C. and E.G.O. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data supporting this theoretical study are contained within the article.

Acknowledgments: This work benefited from interactions with the Research Program PIA_no di inCEntivi per la Ricerca di Ateneo 2024–2026—Linea di Intervento I “Progetti di ricerca collabora-

tiva”—University of Catania—Project “Semantic Web of EveryThing through Ontological Protocols” (SWETOP). We also acknowledge INdAM/GNCS.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. A Lower Bound on the Number of Sets of a Fixed Integer Rank

Here, we figure out inequalities preparatory to the proof of Proposition A2 below.

Proposition A1. *For every positive integer n , we have*

$$2^{2^n} - 2^n \geq 2(2^n - n).$$

Proof. We prove the proposition by establishing the equivalent inequality

$$2^{2^n} \geq 3 \cdot 2^n - 2n, \tag{A1}$$

which follows by simply adding 2^n to both sides of the original inequality.

We proceed by induction on $n \geq 1$.

Base case:

For $n = 1$, we have

$$2^{2^1} = 4 \quad \text{and} \quad 3 \cdot 2^1 - 2 \cdot 1 = 6 - 2 = 4.$$

Thus, the base case holds with equality.

Inductive step:

Fix $n \geq 1$ and assume that

$$2^{2^n} \geq 3 \cdot 2^n - 2n. \tag{A2}$$

We aim to prove that

$$2^{2^{n+1}} \geq 3 \cdot 2^{n+1} - 2(n + 1). \tag{A3}$$

Note that $2^{2^{n+1}} = (2^{2^n})^2$. Since $n \geq 1$, we have $2^{2^n} \geq 2^2 = 4$, and thus

$$(2^{2^n})^2 \geq 4 \cdot 2^{2^n}.$$

Using the inductive hypothesis (A2), we obtain

$$2^{2^{n+1}} = (2^{2^n})^2 \geq 4 \cdot 2^{2^n} \geq 4(3 \cdot 2^n - 2n) = 12 \cdot 2^n - 8n.$$

Now, observe that

$$3 \cdot 2^{n+1} - 2(n + 1) = 6 \cdot 2^n - 2n - 2,$$

so it suffices to verify the inequality

$$12 \cdot 2^n - 8n \geq 6 \cdot 2^n - 2n - 2,$$

which reduces to

$$6 \cdot 2^n - 6n + 2 \geq 0.$$

This inequality clearly holds, since $n \geq 1$ implies $2^n - n \geq 0$, and the leftover $+2$ ensures strict positivity.

Thus, the inductive step is verified. \square

Next, we come to a proposition which lies in the background of this paper:

Proposition A2. For every positive integer n , the number of well-founded sets of rank equal to n is greater than or equal to 2^{n-1} , namely,

$$|\mathcal{W}_n^\#| \geq 2^{n-1}.$$

Proof. We proceed by induction on $n \geq 1$.

Base case:

For $n = 1$, we have $|\mathcal{W}_1^\#| = 1 = 2^{1-1}$.

Inductive step:

Let $n \geq 1$ be such that $|\mathcal{W}_n^\#| \geq 2^{n-1}$. We aim to show that

$$|\mathcal{W}_{n+1}^\#| \geq 2^n. \tag{A4}$$

From the equality $|\mathcal{W}_n^\#| = |\mathcal{V}_{n+1}| - |\mathcal{V}_n|$ and the inductive hypothesis, it follows that

$$|\mathcal{V}_{n+1}| = |\mathcal{W}_n^\#| + |\mathcal{V}_n| \geq 2^{n-1} + |\mathcal{V}_n|.$$

Hence, by (2),

$$\begin{aligned} |\mathcal{W}_{n+1}^\#| &= |\mathcal{V}_{n+2}| - |\mathcal{V}_{n+1}| \\ &= 2^{|\mathcal{V}_{n+1}|} - |\mathcal{V}_{n+1}| \\ &= 2^{2^{|\mathcal{V}_n|}} - 2^{|\mathcal{V}_n|} \\ &\geq 2(2^{|\mathcal{V}_n|} - |\mathcal{V}_n|) && \text{(by Proposition A1, since } |\mathcal{V}_n| \geq 1) \\ &= 2(|\mathcal{V}_{n+1}| - |\mathcal{V}_n|) \\ &= 2|\mathcal{W}_n^\#| \\ &\geq 2 \cdot 2^{n-1} && \text{(by the induction hypothesis)} \\ &= 2^n. \quad \square \end{aligned}$$

Appendix B. Proof of Lemma 14

For convenience, we restate Lemma 14 before giving its proof.

Lemma 14. For every $i = 1, \dots, m$, letting $x_i = \{y_i\}$ be the atom of ψ selected at step i , we have

- (a) $M^{(i)}$ is a BST-replacement of $M^{(i-1)}$ such that $M^{(i)} \models \varphi \wedge \Xi_\varphi^\psi$;
- (b) For each j with $1 \leq j < i$
 - (b₁) $M^{(j)}y_j = M^{(i)}y_j$;
 - (b₂) $M^{(j)}x_j = M^{(i)}x_j$;
- (c) $M^{(i)}y_i = M^{(i-1)}y_i$;
- (d) $M^{(i)}x_i = \{M^{(i)}y_i\}$.

Proof. We prove the lemma by strong induction on $i = 1, \dots, m$.

BASE CASE ($i = 1$). Since $M^{(0)} \models \varphi \wedge \Xi_\varphi^\psi$, we have in particular

$$M^{(0)} \models \bigwedge_{x \in \text{Vars}(\varphi \wedge \psi)} (\neg \text{DISJ}(x_1, x) \longrightarrow x_1 \subseteq x). \tag{A5}$$

Moreover, by Lemma 10, $M^{(0)}x_1 \neq \emptyset$. In addition, since $\text{rk}(M^{(0)}y_1)$ is either 0 or $b + 1$, the b -flatness of $M^{(0)}$ ensures that $\text{DISJ}(M^{(0)}y_1, Mv)$, for all $v \in \text{Vars}(\varphi \wedge \psi)$. Hence, the BST-conditions for the pair $(M^{(0)}x_1, \{M^{(0)}y_1\})$ are satisfied, and therefore $M^{(1)}$ is a BST-replacement of $M^{(0)}$ from $M^{(0)}x_1$ to $\{M^{(0)}y_1\}$ (restricted to $\text{Vars}(\varphi \wedge \psi)$). Since every literal ℓ in $\varphi \wedge \Xi_\varphi^\psi$ satisfies either $\text{Vars}(\ell) \subseteq \text{Vars}(\varphi \wedge \psi)$ or $\text{Vars}(\ell) \cap \text{Vars}(\varphi \wedge \psi) = \emptyset$, Lemma 6 ensures that $M^{(1)} \models \varphi \wedge \Xi_\varphi^\psi$. Therefore, clause (a) holds in the base case.

Clause (b) is vacuous when $i = 1$.

Finally, concerning clause (c), from (A5) and Lemma 10 we obtain $\text{DISJ}(M^{(0)}x_1, M^{(0)}y_1)$. Therefore, by the definition of $M^{(1)}$, it follows that $M^{(1)}y_1 = M^{(0)}y_1$.

INDUCTIVE STEP ($1 < i \leq m$). Assume (a)–(d) hold for all $1 \leq t < i$. \square

Proof of (a). We show that $M^{(i)}$ is a BST-replacement of $M^{(i-1)}$ by verifying that the BST-conditions (8) hold for the pair $(M^{(i-1)}x_i, \{M^{(i-1)}y_i\})$ relative to $M^{(i-1)}$.

We first show that $M^{(i-1)}x_i \neq \emptyset$. By the inductive hypothesis (a) at step $i - 1$, we have $M^{(i-1)} \models \varphi \wedge \Xi_\varphi^\psi$. Therefore, by Lemma 10, we conclude that $M^{(i-1)}x_i \neq \emptyset$. Moreover, $\{M^{(i-1)}y_i\} \neq \emptyset$ holds trivially. Thus, to complete the proof that $M^{(i)}$ is a BST-replacement of $M^{(i-1)}$, it remains to check that the following conditions hold for all $v \in \text{Vars}(\varphi \wedge \psi)$:

- (A) $\text{DISJ}(M^{(i-1)}x_i, M^{(i-1)}v) \vee M^{(i-1)}x_i \subseteq M^{(i-1)}v$;
- (B) $\text{DISJ}(\{M^{(i-1)}y_i\} \setminus M^{(i-1)}x_i, M^{(i-1)}v)$.

Concerning (A), recalling that $M^{(i-1)} \models \varphi \wedge \Xi_\varphi^\psi$, in particular we have

$$M^{(i-1)} \models \bigwedge_{v \in \text{Vars}(\varphi \wedge \psi)} (\neg \text{DISJ}(x_i, v) \longrightarrow x_i \subseteq v),$$

as such conjunction is explicitly contained in Ξ_φ^ψ . Thus,

$$\neg \text{DISJ}(M^{(i-1)}x_i, M^{(i-1)}v) \implies M^{(i-1)}x_i \subseteq M^{(i-1)}v, \quad \text{for all } v \in \text{Vars}(\varphi \wedge \psi),$$

which establishes (A). \square

We now turn to clause (B). Its verification requires a case analysis, carried out for an arbitrary $v \in \text{Vars}(\varphi \wedge \psi)$.

If $\text{DISJ}(\{M^{(i-1)}y_i\}, M^{(i-1)}v)$ holds, then *a fortiori* so does $\text{DISJ}(\{M^{(i-1)}y_i\} \setminus M^{(i-1)}x_i, M^{(i-1)}v)$.

The complementary case $\neg \text{DISJ}(\{M^{(i-1)}y_i\}, M^{(i-1)}v)$ requires a much more detailed argument, which we provide below.

Thus, let $\{M^{(i-1)}y_i\} \cap M^{(i-1)}v \neq \emptyset$, so that $M^{(i-1)}y_i \in M^{(i-1)}v$.

By Lemma 13, $\text{rk}(M^{(i-1)}y_i) \neq b$, so $M^{(i-1)}y_i \notin M^{(0)}v$, i.e., $\text{DISJ}(\{M^{(i-1)}y_i\}, M^{(0)}v)$. Therefore, there must exist an index k with $1 \leq k \leq i - 1$ such that

$$M^{(i-1)}y_i \notin M^{(k-1)}v \quad \text{and} \quad M^{(i-1)}y_i \in M^{(k)}v.$$

Consequently,

$$M^{(k)}v = (M^{(k-1)}v \setminus M^{(k-1)}x_k) \cup \{M^{(k-1)}y_k\},$$

hence,

$$M^{(k-1)}y_k = M^{(i-1)}y_i. \tag{A6}$$

Since by inductive hypothesis (a) at step k the assignment $M^{(k)}$ is a BST-replacement of $M^{(k-1)}$, it must be the case that $M^{(k-1)}x_k \neq \emptyset$, that is, $\neg \text{DISJ}(M^{(k-1)}x_k, M^{(k-1)}x_k)$. Thus, we have

$$M^{(k)}x_k = (M^{(k-1)}x_k \setminus M^{(k-1)}x_k) \cup \{M^{(k-1)}y_k\} = \{M^{(k-1)}y_k\} = \{M^{(k)}y_k\}, \tag{A7}$$

where the final equality follows by the inductive hypothesis (c).

Moreover, whether $k = i - 1$ or $k < i - 1$, it follows in both cases that

$$M^{(k)}y_k = M^{(i-1)}y_k \tag{A8}$$

(the former trivially; the latter by the inductive hypothesis (b₁)).

Thus, we have

$$\begin{aligned} M^{(i-1)}y_i &= M^{(k-1)}y_k && \text{(by (A6))} \\ &= M^{(k)}y_k && \text{(by inductive hypothesis (c))} \\ &= M^{(i-1)}y_k && \text{(by (A8)).} \end{aligned} \tag{A9}$$

Recalling again that $M^{(i-1)} \models \Xi_\varphi^\psi$, we obtain

$$M^{(i-1)} \models y_k = y_i \iff x_k = x_i,$$

since Ξ_φ^ψ contains the conjunct $y_k = y_i \iff x_k = x_i$. In combination with (A9) and the equivalences (11) in Remark 5, this yields

$$M^{(i-1)}x_k = M^{(i-1)}x_i. \tag{A10}$$

Whether $k = i - 1$ (immediate) or $k < i - 1$ (by the inductive hypothesis (b₂)), we obtain

$$M^{(i-1)}x_k = M^{(k)}x_k. \tag{A11}$$

Thus,

$$\begin{aligned} M^{(i-1)}x_i &= M^{(i-1)}x_k && \text{(by (A10))} \\ &= M^{(k)}x_k && \text{(by (A11))} \\ &= \{M^{(k)}y_k\} && \text{(by (A7))} \\ &= \{M^{(i-1)}y_i\} && \text{(by (A8)).} \end{aligned}$$

Hence, $\{M^{(i-1)}y_i\} = M^{(i-1)}x_i$. Therefore, $\{M^{(i-1)}y_i\} \setminus M^{(i-1)}x_i = \emptyset$, and so the disjointness condition holds trivially. This completes the verification of (B) also in the case $\neg \text{DISJ}(\{M^{(i-1)}y_i\}, M^{(i-1)}v)$. Since v was arbitrary, property (B) holds for all variables in $\text{Vars}(\varphi \wedge \psi)$. Together with (A), it follows that $M^{(i)}$ is a BST-replacement of $M^{(i-1)}$. Finally, recalling that, by inductive hypothesis (a) at step $i - 1$, we have $M^{(i-1)} \models \varphi \wedge \Xi_\varphi^\psi$, Lemma 6 yields $M^{(i)} \models \varphi \wedge \Xi_\varphi^\psi$. Thus, clause (a) is established at step i .

Proof of (b). Concerning (b₁), assume by way of contradiction that $M^{(j)}y_j \neq M^{(i)}y_j$ for some j with $1 \leq j < i$, and let k be an index such that $j < k \leq i$ and

$$M^{(k-1)}y_j = M^{(j)}y_j \quad \text{while} \quad M^{(k)}y_j \neq M^{(j)}y_j.$$

Such an index k must exist; otherwise, $M^{(h)}y_j = M^{(j)}y_j$ for all h with $j \leq h \leq i$, which would give $M^{(i)}y_j = M^{(j)}y_j$ and contradict the assumption. Since $M^{(k-1)}y_j \neq M^{(k)}y_j$,

it follows that $M^{(k-1)}x_k \subseteq M^{(k-1)}y_j$ and $M^{(k)}y_j = (M^{(k-1)}y_j \setminus M^{(k-1)}x_k) \cup \{M^{(k-1)}y_k\}$. Moreover, as $M^{(k-1)} \models x_k \subseteq y_j$, the inductive hypothesis (a) guarantees that this inclusion already holds in $M^{(0)}$. Namely, $M^{(0)}x_k \subseteq M^{(0)}y_j$. Finally, by Lemma 10, we also know that $M^{(0)}x_k \neq \emptyset$. Thus, $\ell_k \prec^{M^{(0)}} \ell_j$, implying $k < j$, a contradiction to $j < k$. Therefore, $M^{(j)}y_j = M^{(i)}y_j$ for all $1 \leq j < i$. \square

Regarding (b₂), assume by way of contradiction that $M^{(j)}x_j \neq M^{(i)}x_j$ for some j with $1 \leq j < i$. Then, there must exist a k with $j < k \leq i$ and such that

$$M^{(k-1)}x_j = M^{(j)}x_j \quad \text{while} \quad M^{(k)}x_j \neq M^{(j)}x_j, \tag{A12}$$

otherwise $M^{(h)}x_j = M^{(j)}x_j$ for all h with $j \leq h \leq i$, yielding $M^{(i)}x_j = M^{(j)}x_j$ and contradicting the assumption.

By inductive hypothesis (d) at step j , we have

$$M^{(j)}x_j = \{M^{(j)}y_j\}. \tag{A13}$$

Moreover, since $j \leq k - 1 < i$, the inductive hypothesis (b₂) gives

$$M^{(k-1)}x_j = M^{(j)}x_j. \tag{A14}$$

At stage k , $M^{(k)}$ is a BST-replacement of $M^{(k-1)}$ (by the inductive hypothesis (a) if $k < i$, or by clause (a) at step i if $k = i$). Thus, the pair $(M^{(k-1)}x_k, \{M^{(k-1)}y_k\})$ satisfies the BST-replacement conditions relative to $M^{(k-1)}$, and combined with (A12) this yields

$$\emptyset \neq M^{(k-1)}x_k \subseteq M^{(k-1)}x_j. \tag{A15}$$

Since $M^{(j)}x_j$ is a singleton, it follows from (A14) and (A15) that

$$M^{(k-1)}x_k = M^{(k-1)}x_j. \tag{A16}$$

Therefore, $M^{(k-1)} \models x_k = x_j$.

Because $M^{(0)} \models \Xi_\varphi^\psi$, the inductive hypothesis (a) ensures $M^{(k-1)} \models \Xi_\varphi^\psi$. As $y_k = y_j \iff x_k = x_j$ is a conjunct of Ξ_φ^ψ , we deduce

$$M^{(k-1)} \models y_k = y_j \iff x_k = x_j,$$

hence $M^{(k-1)} \models y_k = y_j$. By the equivalences (11) in Remark 5, this gives

$$M^{(k-1)}y_k = M^{(k-1)}y_j. \tag{A17}$$

Finally, since $j < k \leq i$, the inductive hypothesis (b₁) entails

$$M^{(j)}y_j = M^{(k-1)}y_j. \tag{A18}$$

Putting everything together, we obtain

$$\begin{aligned} M^{(k)}x_j &= (M^{(k-1)}x_j \setminus M^{(k-1)}x_k) \cup \{M^{(k-1)}y_k\} && \text{(by (A15) and (A16))} \\ &= \{M^{(k-1)}y_k\} \\ &= \{M^{(j)}y_j\} && \text{(by (A17) and (A18))} \\ &= M^{(j)}x_j && \text{(by (A13)).} \end{aligned}$$

This contradicts (A12), and therefore clause (b₂) is established at step i .

Proof of (c). We need to show that $M^{(i)}y_i = M^{(i-1)}y_i$. To this end, recall first that $M^{(0)} \models \Xi_\varphi^\psi$. As already noted, we have $M^{(i-1)} \models \Xi_\varphi^\psi$. Lemma 10 then gives $\text{DISJ}(M^{(i-1)}x_i, M^{(i-1)}y_i)$, and by the definition (7) of $M^{(i)}y_i$ we conclude that $M^{(i)}y_i = M^{(i-1)}y_i$. This completes the verification of clause (c) at step i . \square

Proof of (d). Since, by (a), $M^{(i)}$ is a BST-replacement of $M^{(i-1)}$ from $M^{(i-1)}x_i$ to $\{M^{(i-1)}y_i\}$, we must have $M^{(i-1)}x_i \neq \emptyset$, equivalently $\neg\text{DISJ}(M^{(i-1)}x_i, M^{(i-1)}x_i)$. By the definition of BST-replacement, this yields

$$\begin{aligned} M^{(i)}x_i &= (M^{(i-1)}x_i \setminus M^{(i-1)}x_i) \cup \{M^{(i-1)}y_i\} \\ &= \{M^{(i-1)}y_i\} = \{M^{(i)}y_i\}, \end{aligned}$$

where the last equality follows from clause (c). This establishes (d). Since the inductive step is now complete, and the base case was already established, the lemma follows. \square

References

1. Cantone, D.; Omodeo, E.G. Onset and today’s perspectives of multilevel syllogistic. In *From Computational Logic to Computational Biology: Essays Dedicated to Alfredo Ferro to Celebrate His Scientific Career*; Cantone, D., Pulvirenti, A., Eds.; Springer Nature: Cham, Switzerland, 2024; Volume 14070 of LNCS, pp. 9–55. [\[CrossRef\]](#)
2. Schwartz, J.T. *Instantiation and Decision Procedures for Certain Classes of Quantified Set-Theoretic Formulae*; Technical Report 78-10; Institute for Computer Applications in Science and Engineering, NASA Langley Research Center: Hampton, Virginia, 1978.
3. Schwartz, J.T.; Cantone, D.; Omodeo, E.G. *Computational Logic and Set Theory—Applying Formalized Logic to Analysis*; Springer: Berlin/Heidelberg, Germany, 2011.
4. Cantone, D.; Domenico, A.D.; Maugeri, P.; Omodeo, E.G. Complexity assessments for decidable fragments of set theory. I: A taxonomy for the Boolean case. *Fundam. Informaticae* **2021**, *181*, 37–69. [\[CrossRef\]](#)
5. Cantone, D.; Maugeri, P.; Omodeo, E.G. Complexity assessments for decidable fragments of set theory. II: A taxonomy for ‘small’ languages involving membership. *Theor. Comput. Sci.* **2020**, *848*, 28–46. [\[CrossRef\]](#)
6. Kuncak, V.; Nguyen, H.H.; Rinard, M. Deciding Boolean algebra with Presburger arithmetic. *J. Autom. Reason.* **2006**, *36*, 213–239. [\[CrossRef\]](#)
7. Rabin, M.O. Decidable theories. In *Handbook of Mathematical Logic, Volume 90 of Studies in Logic and the Foundations of Mathematics*; Barwise, J., Ed.; Elsevier (North-Holland Publishing Co.): Amsterdam, The Netherlands, 1977; pp. 595–629.
8. Cantone, D.; Omodeo, E.G.; Policriti, A. The automation of syllogistic. II: Optimization and complexity issues. *J. Autom. Reason.* **1990**, *6*, 173–188. [\[CrossRef\]](#)
9. Cormen, T.H.; Leiserson, C.E.; Rivest, R.L.; Stein, C. *Introduction to Algorithms*, 3rd ed.; MIT Press: Cambridge, MA, USA, 2009.
10. Angrisani, F.; Ascione, G.; Manzo, G. Orlicz spaces with a O- Type Structure. *Ric. Mat.* **2019**, *68*, 841–857. [\[CrossRef\]](#)
11. Jiménez-Garrido, J.; Sanz, J.; Schindl, G. Indices of \mathcal{O} -regular variation for weight functions and weight sequences. *Rev. Real Acad. Cienc. Exactas Físicas Nat. Ser. A Matemáticas* **2018**, *113*, 3659–3697. [\[CrossRef\]](#)
12. Cantone, D.; Ferro, A. Techniques of computable set theory with applications to proof verification. *Commun. Pure Appl. Math.* **1995**, *48*, 901–945. [\[CrossRef\]](#)
13. Parlamento, F.; Policriti, A.; Rao, K.P.S.B. Witnessing Differences without Redundancies. *Proc. Am. Math. Soc.* **1997**, *125*, 587–594. [\[CrossRef\]](#)
14. Hooker, J.N. A quantitative approach to logical inference. *Decis. Support Syst.* **1988**, *4*, 45–69. [\[CrossRef\]](#)
15. Hooker, J.N.; Fedjiki, C. Branch-and-cut solution of inference problems in propositional logic. *Ann. Math. Artif. Intell.* **1990**, *1*, 123–139. [\[CrossRef\]](#)
16. Cantone, D.; Omodeo, E.G.; Policriti, A. *Set Theory for Computing. From Decision Procedures to Declarative Programming with Sets*; Monographs in Computer Science; Springer: Berlin/Heidelberg, Germany, 2001.
17. Tarski, A. *Ordinal Algebras*; Studies in Logic and the Foundations of Mathematics; North-Holland: Amsterdam, The Netherlands, 1956;
18. Hibti, M. *Décidabilité et Complexité de Systèmes de Contraintes Ensemblistes*. PhD Thesis, Université de Franche-Comté, Besançon, France, 1995.
19. Hibti, M.; Lombardi, H.; Legeard, B. Deciding in HFS-theory via linear integer programming. In *Logic Programming and Automated Reasoning, Proceedings of the 4th Int’l Conference, LPAR’93, St. Petersburg, Russia, 13–20 July 1993*; Voronkov, A., Ed.; Springer: Berlin/Heidelberg, Germany, 1993; Volume 698, pp. 170–181.

20. Hibti, M.; Legeard, B.; Lombardi, H. Une procédure de décision pour un problème de satisfiabilité dans un univers ensembliste héréditairement fini. *RAIRO Theor. Inform. Appl.* **1997**, *31*, 205–236. [[CrossRef](#)]
21. Cristiá, M.; Rossi, G. Integrating cardinality constraints into constraint logic programming with sets. *Theory Pract. Log. Program.* **2021**, *23*, 468–502. [[CrossRef](#)]
22. Givan, R.; McAllester, D.A.; Witty, C.; Kozen, D. Tarskian Set Constraints. *Inf. Comput.* **2002**, *174*, 105–131. [[CrossRef](#)]
23. Raya, R.; Hamza, J.; Kunčák, V. On the Complexity of Convex and Reverse Convex Prequadratic Constraints. In *LPAR 2023: Proceedings of the 24th International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Manizales, Colombia, 4–9 June 2023*; Piskac, R., Voronkov, A., Eds.; EasyChair: Stockport, UK, 2023; Volume 94, pp. 350–368. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.