Deciding Simple Infinity Axiom Sets with One Binary Relation by Means of Superpostulates

Timm Lampert[1(✉)] and Anderson Nakano[2]

[1] Humboldt University Berlin, Unter den Linden 6, 10099 Berlin, Germany
lampertt@staff.hu-berlin.de
[2] Pontifícia Universidade Católica, R. Monte Alegre, 984, São Paulo 05014-901, Brazil
alnakano@pucsp.br

Abstract. Modern logic engines widely fail to decide axiom sets that are satisfiable only in an infinite domain. This paper specifies an algorithm that automatically generates a database of independent infinity axiom sets with fewer than 1000 characters. It starts with complete theories of pure first-order logic with only one binary relation (FOL_R) and generates further infinity axiom sets S of FOL_R with fewer than 1000 characters such that no other infinity axiom set with fewer than 1000 characters exists in the database that implies S. We call the generated infinity axiom sets S "superpostulates". Any formula that is derivable from (satisfiable) superpostulates is also satisfiable. Thus far, we have generated a database with 2346 infinity superpostulates by running our algorithm. This paper ends by identifying three practical uses of the algorithmic generation of such a database: (i) for systematic investigations of infinity axiom sets, (ii) for deciding infinity axiom sets and (iii) for the development of saturation algorithms.

Keywords: First-order logic · Decision problem · Complete theories · Infinity axioms · Reduction classes · Dyadic logic

1 Introduction

Modern logic engines for first-order formulas (FOL), such as Vampire, CVC4, Spass or i-Prover, are very powerful in proving theorems (or refutations) or in deciding formulas with finite models (counter-models). For formulas that have only infinite models, however, these engines widely fail.[1] One exception is

[1] Exceptions include infinity axiom sets that can be decided by saturation, such as the theory of immediate successors as stated in FOL with identity (cf. (14) on p. 15). However, in regard to FOL without identity, infinity axiom sets are only rarely solved, and only a few of them are included in the Thousands of Problems for Theorem Provers (TPTP) library, namely, problems SYO635+1 to SYO638+1. Among these problems, only SYO638+1 can be solved by SPASS due to chaining rules, which directly apply to the transitivity axioms. Furthermore, Infinox decides the finite unsatisfiability of SY0638+1 only by virtue of the plain properties of the relation involved.

N. Peltier and V. Sofronie-Stokkermans (Eds.): IJCAR 2020, LNAI 12166, pp. 201–217, 2020.
https://doi.org/10.1007/978-3-030-51074-9_12

Infinox, which decides the finite unsatisfiability of some relevant TPTP problems due to specific model-theoretic principles (cf. [3]).[2] In contrast to Infinox, we investigate how to *generate* hitherto unknown and intricate infinity axiom sets (*superpostulates*) of FOL_R to make it possible to decide an infinite number of infinity axiom sets related to these superpostulates. No engine is currently able to solve most of the superpostulates we generate.

In this paper, we refer only to FOL_R, i.e., FOL without function symbols, without identity and with only one binary relation. Based on Herbrand's reduction of FOL to dyadic FOL, Kalmar has proven that FOL_R is a reduction class (cf. [7]). Boolos, Jeffrey and Burgess have defined and proven the algorithm for reducing FOL to FOL_R in modern terms (cf. [2], Sect. 21.3). Börger, Grädel and Gurevich have specified prenex normal forms of FOL_R-formulas with prenexes of the form $\forall\exists\forall$ (= $[\forall\exists\forall, (0,1)]$-class) or of the form $\forall\forall\forall\exists$ (= $[\forall^3\exists, (0,1)]$-class) as the two *minimal* classes of FOL_R with infinity axioms (cf. Theorem 6.5.4, p. 309) and have presented the following two examples for the two classes (cf. [1], p. 307):

$$\forall x\exists y\forall z(\neg Rxx \wedge Rxy \wedge (Ryz \rightarrow Rxz)) \tag{1}$$

$$\forall x\forall y\forall z\exists u(\neg Rxx \wedge Rxu \wedge (Rxy \wedge Ryz \rightarrow Rxz)) \tag{2}$$

They also specify another formula with the prenex $\forall\exists\forall$ (cf. [1], p. 33):

$$\forall x\exists u\forall y(\neg Rxx \wedge Rxu \wedge (Ryx \rightarrow Ryu)) \tag{3}$$

(1) and (3) are both derivable from (2). Thus, to decide that the latter is satisfiable, it is sufficient to decide that the former are likewise satisfiable. This method of deciding infinity axioms can be extended by specifying infinity axioms that imply other infinity axioms.

However, prenex normal forms do not provide a syntax that is suitable for this task. Different prenex normal forms can be converted into each other by equivalence transformation. For example, formula (2), which is a member of the minimal class $[\forall^3\exists, (0,1)]$, can easily be converted into a member of the minimal class $[\forall\exists\forall, (0,1)]$, as shown in Table 1.

Table 1. Converting different prenex normal forms into each other

No.	Formula	Strategy
(i)	$\forall x\forall y\forall z\exists u(\neg Rxx \wedge Rxu \wedge (Rxy \wedge Ryz \rightarrow Rxz))$	(2)
(ii)	$\forall x\neg Rxx \wedge \forall x\exists u Rxu \wedge \forall x\forall y(\neg Rxy \vee \forall z(\neg Ryz \vee Rxz))$	Miniscoping
(iii)	$\forall x_2\neg Rx_2x_2 \wedge \forall x_1\exists y_1 Rx_1y_1 \wedge \forall x_3\forall x_4(\neg Rx_3x_4 \vee \forall x_5(\neg Rx_4x_5 \vee Rx_3x_5))$	Renaming
(iv)	$\forall x_1\exists y_1\forall x_2\forall x_3\forall x_4\forall x_5(\neg Rx_2x_2 \wedge Rx_1y_1 \wedge (\neg Rx_3x_4 \vee (\neg Rx_4x_5 \vee Rx_3x_5)))$	Prenexing

[2] A further exception with respect to the engines on the TPTP site is Decider. This engine is able to identify infinity axiom sets of pure FOL without identity due to the implementation of a strong saturation algorithm that is not based on resolution. However, the complexity of the implemented algorithm is exponential, and the Decider engine is not optimized for rapid decision making. Thus, it fails to decide complex formulas within reasonable time and memory limits.

Likewise, it may well be that a formula of the minimal class $[\forall\exists\forall, (0,1)]$ can be converted into a formula of a decidable class, e.g., with prenexes starting with existential quantifiers followed by universal quantifiers $(=\exists^*\forall^*, (0,1)]$-class). Similarly, the distribution of bound variables in the matrix is not standardized in prenex normal forms. Thus, many equivalent variations are possible. Therefore, the syntax of prenex normal forms is not well suited for investigating the internal properties and relations of infinity axioms.

This is why we do not refer to *infinity axioms* (and, thus, to prenex normal forms) in the following. Instead, we refer to the opposite of prenex normal forms, namely, anti-prenex normal forms, in which the scopes of quantifiers are minimized. We refer to *infinity axiom sets* in terms of sets of *anti-prenex normal forms*. To do so, we define *primary formulas* via negation normal forms (NNFs) as follows:

Definition 1. *A first-order formula ϕ is* primary *(= in* anti-prenex form*) if ϕ is an NNF and either*

1. *does not contain $\wedge$ or $\vee$ or*
2. *contains $\wedge$ or $\vee$ iff either*
   (a) *any conjunction of n conjuncts ($n > 1$) is preceded by a sequence of existential quantifiers of minimal length 1, where all n conjuncts contain each variable of the existential quantifiers in that sequence, or*
   (b) *any disjunction of n disjuncts ($n > 1$) is preceded by a sequence of universal quantifiers of minimal length 1, where all n disjuncts contain each variable of the universal quantifiers in that sequence.*

Thus, lines (ii) and (iii) of Table 1 are conjunctions of three primary formulas, while lines (i) and (iv) do not contain primary formulas. Any first-order formula can be converted into a disjunction of conjunctions of primary formulas (=FOLDNF); cf. [8] for an algorithm for doing so. Since the satisfiability of a disjunction can be decided by deciding the satisfiability of each disjunct, we consider only conjunctions of primary formulas or, analogously, (finite) sets of axioms (= theories) in anti-prenex form.

Definition 2. *An* axiom set *is a finite set of primary formulas.*

Definition 3. *An* infinity axiom set *is a finite set of primary formulas that is satisfiable only within an infinite domain.*

From here on, we will express all axiom sets of FOL_R in standardized notation.

Definition 4. *A set of primary formulas (axioms) of FOL_R is in* standardized notation *if*

1. *the binary predicate used is R and*

2. *for each primary formula (axiom) A,*
   - *the m universal quantifiers of A are binding variables x_1 to x_m from left to right and*
   - *the n existential quantifiers of A are binding variables y_1 to y_n from left to right.*

Example 1. In standardized notation, formula (2) is converted into the following infinity axiom set:

$$\forall x_1 \neg Rx_1x_1, \forall x_1 \exists y_1 Rx_1y_1, \forall x_1 \forall x_2 (\neg Rx_1x_2 \vee \forall x_3 (\neg Rx_2x_3 \vee Rx_1x_3)) \quad (4)$$

This axiom set was the example used by Hilbert and Bernays to motivate proof theory by showing that the consistency of (4) cannot be proven by finite interpretations ("Methode der Aufweisung"; cf. [6], p. 10). The third axiom expresses transitivity.

Formula (3) is equivalent to the following infinity axiom set from [11], p. 183:

$$\forall x_1 \neg Rx_1x_1, \forall x_1 \exists y_1 (Rx_1y_1 \wedge \forall x_2 (\neg Rx_2x_1 \vee Rx_2y_1)) \quad (5)$$

Like (1), (3) and (5) are derivable from (4). The following infinity axiom set (6) is taken from [2], p. 138:

$$\forall x_1 \forall x_2 (\neg Rx_1x_2 \vee \neg Rx_2x_1), \forall x_1 \exists y_1 Rx_1y_1, \forall x_1 \forall x_2 (\neg Rx_1x_2 \vee \forall x_3 (\neg Rx_2x_3 \vee Rx_1x_3)) \quad (6)$$

Like (2), this infinity axiom set is equivalent to (4).

The formulas mentioned up to this point are examples of infinity axiom sets of FOL_R given in the standard literature. They are all satisfiable by interpreting Rxy as $x < y$ over $\mathbb{Q}$. Thus, they are all derivable from the complete dense linear order (DLO) axiom set without endpoints. In the following sections, we will generate a system of DLO variants without identity, to which all of the mentioned standard examples belong. We use the term "*system* of formulas" to refer to a recursive set of formulas generated by rules that can be implemented by a computer program. Our investigation may serve as a case study for investigating systems of infinity axiom sets of FOL_R.

Our goal is the automated generation of a system of consistent infinity axiom sets S of a limited length L such that no other infinity axiom sets $S2$ of length $\leq L$ exist such that $S2$ implies S. We call these infinity axiom sets S "superpostulates". Superpostulates enable the reduction of a large number of infinity axiom sets to a small number of infinity axiom sets of limited length. The term and the general idea of studying systems of superpostulates originated in the work of Sheffer (cf. [12]). Sheffer, however, neither published his work nor developed his logic project to a full extent. His student Langford made use of Sheffer's idea in [9]. Langford intended to prove that the axiom set for linear orderings is a complete superpostulate. Urquhart (cf. [14], p. 44), however, objected to the correctness of his proof and showed how to correct it. Langford proved similar results for the theories of betweenness and cyclic order. All of his examples are equivalent to formulas of decidable classes without existential quantifiers and with identity, which have the finite model property (cf. [1], Theorem 6.5.1, p.

307). To our knowledge, neither Langford nor Sheffer studied superpostulates in terms of infinity axiom sets.

The following sections investigate infinity axiom sets related to the complete DLO axiom set without endpoints. Section 2 introduces further terminology and specifies a superpostulate DLO_R for FOL_R, which is a DLO variant without identity and without endpoints from which all of the mentioned infinity axiom sets follow. After having specified a complete infinity axiom set, Sect. 3 specifies an infinite sequence of infinity axiom sets, each strictly implying the next and all implied by (5). Section 4 then specifies a system of superpostulates based on DLO_R; Sect. 5 refers to a system THS_R based on the complete theory of immediate successors without identity. Finally, we conclude by identifying several practical uses of our method of algorithmic generation of systems of superpostulates in Sect. 6.

2 Superpostulates

We first introduce some terminology as a basis for defining the term *superpostulate*.

Definition 5. *An axiom A is* independent *of a set of axioms S if A is not implied by S and $S \wedge A$ is consistent.*

Definition 6. *An axiom set A is* complete *in FOL_R if for any FOL_R-formula B, either B or $\neg B$ follows from A.*

Remark 1. Throughout this paper, we consider FOL_R. This means that an axiom set A is complete iff no FOL_R-formula B is independent of A.

We measure the length of an axiom set by the number of characters it contains in standardized notation.

Definition 7. *An axiom set A of length $\leq L$ is* L-complete *in FOL_R if for any FOL_R-formula B of length $\leq L$, either B or $\neg B$ follows from A.*

Remark 2. Definition 7 implies that no formula B of length $\leq L$ exists such that A of length $\leq L$ is strictly implied by B if A is L-complete.

Definition 8. *An axiom set A is* minimal *if A does not contain any redundant part that can be eliminated to result in a logically equivalent axiom set.*

Definition 9. *A* superpostulate *(SP) is a minimal and consistent set of independent axioms that is L-complete.*

It is desirable to define superpostulates as complete (and not merely L-complete) theories. However, we abstain from doing so because we intend to define a *practical* algorithm for generating infinity axiom sets. The standard of preserving completeness, however, faces both practical and theoretical problems, as we will show in Sect. 4. Thus, we merely presume that the generation of a system of superpostulates is based on complete theories and aims for completeness relative to a length L of up to 1000 characters (cf. p. 11).

Definition 10. *An* infinity superpostulate *is a satisfiable superpostulate that has only infinite models.*

Any semi-decider is able to prove that a formula is implied by or inconsistent with another formula or finite axiom set. Since these are the only two options in the case of *complete* superpostulates, it is decidable whether a given formula is implied by a complete superpostulate. Furthermore, all (finite) infinity axiom sets that are implied by a set of infinity superpostulates are decidable.

From the following axiom set DLO_R, all infinity axiom sets of the standard literature mentioned in Sect. 1 follow:

$$\begin{gathered}\forall x_1 \neg Rx_1x_1, \forall x_1 \exists y_1 Rx_1y_1, \forall x_1 \exists y_1 Ry_1x_1, \\ \forall x_1 \forall x_2 (\neg Rx_1x_2 \vee \forall x_3 (\neg Rx_2x_3 \vee Rx_1x_3)), \\ \forall x_1 \forall x_2 (Rx_1x_2 \vee \forall x_3 (Rx_2x_3 \vee \neg Rx_1x_3)), \\ \forall x_1 \forall x_2 (\neg Rx_1x_2 \vee \exists y_1 (Rx_1y_1 \wedge Ry_1x_2))\end{gathered} \tag{7}$$

Axiom 1 (irreflexivity), Axioms 2 and 3 (no right endpoint and no left endpoint), Axiom 4 (transitivity) and Axiom 6 (density) are identical to axioms of DLO without endpoints. However, DLO also contains trichotomy, i.e., the axiom $\forall x \forall y (Rxy \vee Ryx \vee x = y)$, which involves identity. The trichotomy axiom of DLO is replaced by Axiom 5 (transitivity of $\neg R$) in DLO_R.

Formula (7) contains all of the axioms of the infinity axiom set given in formula (4). Additionally, Axioms 3, 5 and 6 are added to those in (4). These axioms are not essential for satisfiability in an infinite domain.

Definition 11. *An axiom A of a superpostulate SP is* essential *if SP without A is no longer an infinity axiom.*

The set of the essential axioms plus the negation of an inessential axiom is also an infinity axiom set, but it is not necessarily a complete one.

In the following, we prove that (7) is a complete superpostulate by proving the relevant properties.

Theorem 1. *(7) is an infinity axiom set.*

Proof. If (7) is satisfiable, then it is satisfiable only in an infinite domain; this is because it implies (4), which is known to be an infinity axiom set. Thus, it remains to be shown that (7) is satisfiable. This is done by providing the following model over $\mathbb{Q}$ for (7): $\Im(R) = \{(x, y) \mid x < y\}$. It can be proven that this interpretation is indeed a model by paraphrasing each axiom. We use the common forms of the transitivity axioms (Axiom 4 and Axiom 5) and the density axiom (Axiom 6) for convenience:

Axiom 1	$\forall x \neg Rxx$
$\Im$(Axiom 1)	All rational numbers (r.n.) are no less than ($<$) themselves (irreflexivity of $<$)
Axiom 2	$\forall x_1 \exists y_1 Rx_1y_1$
$\Im$(Axiom 2)	For all r.n. x_1, some r.n. y_1 exists such that $x_1 < y_1$ (no right endpoint)
Axiom 3	$\forall x_1 \exists y_1 Ry_1x_1$
$\Im$(Axiom 3)	For all r.n. x_1, some r.n. y_1 exists such that $y_1 < x_1$ (no left endpoint)
Axiom 4	$\forall x_1 \forall x_2 \forall x_3 (Rx_1x_2 \wedge Rx_2x_3 \rightarrow Rx_1x_3)$
$\Im$(Axiom 4)	For all r.n. x_1, x_2, and x_3, if $x_1 < x_2$ and $x_2 < x_3$, then $x_1 < x_3$ (transitivity of $<$)
Axiom 5	$\forall x_1 \forall x_2 \forall x_3 (\neg Rx_1x_2 \wedge \neg Rx_2x_3 \rightarrow \neg Rx_1x_3)$
$\Im$(Axiom 5)	For all r.n. x_1, x_2, and x_3, if $x_1 \not< x_2$ and $x_2 \not< x_3$, then $x_1 \not< x_3$ (transitivity of $\not<$)
Axiom 6	$\forall x_1 \forall x_2 (Rx_1x_2 \rightarrow \exists y_1 (Rx_1y_1 \wedge Ry_1x_2))$
$\Im$(Axiom 6)	For all r.n. x_1 and x_2, if $x_1 < x_2$, then there exists an r.n. y_1 such that $x_1 < y_1$ and $y_1 < x_2$ (density)

Therefore, (7) is satisfiable; this is so only in an infinite domain, which means that (7) is an infinity axiom set. □

Theorem 2. *(7) is* minimal.

Proof. This is ensured by systematically producing the following:

1. all possible axiom sets A' from (7) with one axiom eliminated (case 1a) or with the conjunct in Axiom 6 eliminated (case 1b),
2. all possible axiom sets A'' with one disjunct in Axioms 4 to 6 eliminated (case 2), and
3. all possible axiom sets A''' obtained by (α) reducing the universal quantifiers in Axioms 4 to 6 and (β) replacing Axiom 2 and Axiom 3 with $\forall x_1 Rx_1x_1$ or $\exists y_1 Ry_1y_1$ (case 3).

In case 2 and case 3 (β), the resulting sets A'' and A''' have been proven to be inconsistent by Vampire in any arbitrary case; in case 3 (α), the resulting axioms in A''' have been proven to be redundant by Vampire. In case 1a, the results of eliminating Axiom 1 or Axiom 4 have been proven to be finite axiom sets by Vampire. Therefore, the resulting axiom sets are not equivalent. Axioms 2, 3, 5 and 6 are not redundant (and, thus, cannot be eliminated) because they can be proven to be independent axioms by replacing them with their negations, which, in turn, results in infinity axiom sets. This can be seen by referring to the following interpretations:

Axiom 1	Axiom 2	Axiom 3	Axiom 4	Axiom 5	Axiom 6	$\Im$
+	¬	+	+	+	+	$x < y$ over $\mathbb{Q}^+$
+	+	¬	+	+	+	$x < y$ over $\mathbb{Q}^-$
+	+	+	+	¬	+	$x < y - 1$ over $\mathbb{Q}$
+	+	+	+	+	¬	$x < y$ over $\mathbb{Z}$

The conjunct in Axiom 6 is not redundant because eliminating it results in a redundant axiom, whereas Axiom 6 itself is not redundant. □

Theorem 3. *(7) is complete.*

Proof. We extend (7) by the following defining axiom of $x_1 = x_2$:

$$\forall x_1 \forall x_2 (x_1 = x_2 \leftrightarrow \forall x_3 (Rx_1x_3 \leftrightarrow Rx_2x_3)) \tag{8}$$

The union of (7) and (8) is equivalent to DLO without endpoints (as proven by Vampire). Thus, DLO without endpoints is a conservative extension of (7) and, by conservativeness, (7) is complete if DLO is complete.[3] Since DLO without endpoints is known to be complete, (7) is also complete. □

We say that an infinity axiom set B is *weaker* than A if B is strictly implied by A, that is, $A \vdash B$, but $B \vdash A$ does not hold. Before we consider how to generate further superpostulates, let us first consider the opposite: infinity axiom sets that are as weak as possible.

3 Weak Infinity Axiom Sets

The set given in formula (5) is the weakest infinity axiom set mentioned in the standard literature. The following infinite axiom sets, however, strictly follow from (5):

$$\forall x_1 \exists y_1 (Rx_1y_1 \wedge \forall x_2 (\neg Rx_2x_1 \vee Rx_2y_1) \wedge \neg Ry_1y_1) \tag{9}$$

$$\forall x_1 \exists y_1 (Rx_1y_1 \wedge \neg Ry_1x_1 \wedge \forall x_2 (\neg Rx_2x_1 \vee Rx_2y_1)) \tag{10}$$

Formula (9) strictly follows from (5), and (10), in turn, strictly follows from (9). We first prove that (10) is an infinity axiom. We then show that it can still be systematically weakened without losing the property of being an infinity axiom set.

Theorem 4. *(10) is an infinity axiom.*

Proof. We first show that there exists a denumerable model of (10) and then show that (10) has no finite model.

We assume the natural numbers as the domain and interpret Rxy as $x < y$. Then, it is true that for all natural numbers x_1, there exists a number y_1 such

[3] For extensions by definitions, cf., e.g., [13], section 4.6.

that $x_1 < y_1$ and $y_1 \not< x_1$, and that for all x_2, if $x_2 < x_1$, then $x_2 < y_1$. Thus, a denumerable model of (10) exists.

We prove that (10) has no finite model by contradiction. Suppose that there is a finite model M of (10), where the elements of its domain are listed as $m_1, \ldots, m_k$. Let $n_1 = m_1$. In the following, we consider how to satisfy the three conjuncts in the scope of $\forall x_1 \exists y_1$. By the first conjunct, there must be some n in the domain such that (n_1, n) is in $\Im(R)$. Let n_2 be the first element for which this is true. Thus, we have $(n_1, n_2) \in \Im(R)$. By the second conjunct, we have that $(n_2, n_1) \notin \Im(R)$. It follows that $n_1 \neq n_2$ is in $\Im(R)$. Again, by the first conjunct, there must be some n such that (n_2, n). Let n_3 be the first element for which this is the case; thus, we have $(n_2, n_3) \in \Im(R)$ and, by the second conjunct, $(n_3, n_2) \notin \Im(R)$, and therefore, $n_2 \neq n_3$. By the third conjunct, we have either $(n_1, n_2) \notin \Im(R)$ or $(n_1, n_3) \in \Im(R)$. Since we already have $(n_1, n_2) \in \Im(R)$, the second of these statements, i.e., $(n_1, n_3) \in \Im(R)$, must be true. Then, by the second conjunct, we have $(n_3, n_1) \notin \Im(R)$. Thus, we also have $n_1 \neq n_3$ (in addition to $n_2 \neq n_3$) since we have $(n_3, n_1) \notin \Im(R)$ and $(n_1, n_3) \in \Im(R)$. Continuing in this manner, we obtain n_4, which is different from all of n_1, n_2, and n_3; then, we obtain n_5, which is different from n_1 to n_4; and so on. However, by the time we reach n_{k+1}, we will have exceeded the number of elements of the domain. Thus, our supposition that the domain of M is finite leads to a contradiction. Therefore, (10) has no finite model. □

Formula (10) is a simple and weak infinity axiom that no TPTP engine is able to decide.[4]

Theorem 5. *From (10), an infinite number of strictly weaker infinity axiom sets follow.*

Proof. We consider the iterative derivation of strictly weaker axioms from (10) by applying the rule $A \vdash B \vee A$ ($\vee$I), as shown in Table 2.

Table 2. Infinity axioms implied by (10)

No.	Axiom	Rule
(i)	$\forall x_1 \exists y_1 (Rx_1y_1 \wedge \neg Ry_1x_1 \wedge \forall x_2 (\neg Rx_2x_1 \vee Rx_2y_1))$	(10)
(ii)	$\forall x_1 \exists y_1 (Rx_1y_1 \wedge \neg Ry_1x_1 \wedge \forall x_2 (\forall x_3 \neg Rx_3x_2 \vee \neg Rx_2x_1 \vee Rx_2y_1))$	$\vee$I
(iii)	$\forall x_1 \exists y_1 (Rx_1y_1 \wedge \neg Ry_1x_1 \wedge \forall x_2 (\forall x_3 (\forall x_4 \neg Rx_4x_3 \vee \neg Rx_3x_2) \vee \neg Rx_2x_1 \vee Rx_2y_1))$	$\vee$I
(iv)	$\forall x_1 \exists y_1 (Rx_1y_1 \wedge \neg Ry_1x_1 \wedge \forall x_2 (\forall x_3 (\forall x_4 (\forall x_5 \neg Rx_5x_4 \vee \neg Rx_4x_3) \vee \neg Rx_3x_2) \vee \neg Rx_2x_1 \vee Rx_2y_1))$	$\vee$I
⋮	⋮	⋮

Lines (ii), (iii), etc., follow from (10) with the application of nothing but $\vee$I, which is a valid derivation rule. Each result is, in turn, a primary formula. To verify that iteratively applying $\vee$I generates a sequence of weaker *infinity* axiom

[4] This can be checked on the site http://tptp.org/cgi-bin/SystemOnTPTP.

sets, one must determine how the proof of Theorem 4 still applies. Formula (10) allows one to directly conclude that $n_1 \neq n_3$ once $n_1 \neq n_2$ and $n_2 \neq n_3$ are established. Applying $\vee$I once (cf. line (ii)), however, allows one to conclude only that either $(n_1, n_3) \in \Im(R)$ or $(n_4, n_2) \notin \Im(R)$ (where n_4 is the element introduced by $\forall x_3$, which cannot be presumed to be either identical to or different from n_1, n_2, or n_3). Due to the first two conjuncts, $n_1 \neq n_3$ follows from $(n_1, n_3) \in \Im(R)$, and $n_2 \neq n_4$ follows from $(n_4, n_2) \notin \Im(R)$. Considering further objects of the domain introduces further objects that must be different from all objects in one of the established alternative sets of objects. With each further application of $\vee$I, further alternative sets arise. Thus, each axiom resulting from $\vee$I establishes a wider range of objects that cannot be identical. Therefore, each application allows for less powerful inferences concerning the non-identity of the elements in pairs satisfying $\Im(R)$, and thus, each resulting axiom is strictly weaker. Nevertheless, the resulting axioms still necessitate that at least one of the newly considered objects (e.g., n_3 and n_4 in step 2 of the argument after $n_1 \neq n_2$ is established in step 1) in each step of the argument must be different from at least one of the already considered objects (e.g., n_1 and n_2 in step 1). Therefore, the resulting axioms still cannot be satisfied under the assumption of only a finite number of objects. Consequently, we generate an endless sequence of increasingly weaker infinity axioms. □

There are many further infinite sequences of non-equivalent infinity axiom sets that are implied by (7). They all can be proven to be satisfiable by proving that they follow from (7).

4 A System of Superpostulates: The DLO$_R$-System

We are searching for a system of infinity superpostulates that is automatically generated by transformation rules from a complete theory such as (7). The rules specified in this section are a starting point for the general project of generating systems of infinity superpostulates from a given complete axiom set. We do not claim that our third rule preserves completeness.

Definition 12. *A DLO$_R$-system is a system of infinity superpostulates generated from (7).*

Definition 13. *An infinity axiom set is a DLO$_R$-variant if it is derivable from a superpostulate of a DLO$_R$-system.*

The following two rules yield structurally similar infinity axiom sets of FOL$_R$. They trivially preserve the property of being a complete infinity axiom set in the case of FOL$_R$:

Rule 1: Exchange all negated literals for non-negated literals and all non-negated literals for negated literals.

Rule 2: Exchange the variables in positions 1 and 2 in all atomic propositional functions.

These rules do not necessarily yield new (non-equivalent) infinity axiom sets. Rule 1, for example, converts Axiom 4 of superpostulate (7) into Axiom 5 and vice versa. Rule 2 does not change Axiom 1, does convert Axiom 2 into Axiom 3 and vice versa, and yields equivalent axioms in the cases of Axioms 4 and 5. In fact, applying these rules to our superpostulate (7) results in only one further complete superpostulate:

$$\begin{array}{c}\forall x_1 R x_1 x_1, \forall x_1 \exists y_1 \neg R x_1 y_1, \forall x_1 \exists y_1 \neg R y_1 x_1, \\ \forall x_1 \forall x_2 (R x_1 x_2 \vee \forall x_3 (R x_2 x_3 \vee \neg R x_1 x_3)), \\ \forall x_1 \forall x_2 (\neg R x_1 x_2 \vee \forall x_3 (\neg R x_2 x_3 \vee R x_1 x_3)), \\ \forall x_1 \forall x_2 (R x_1 x_2 \vee \exists y_1 (\neg R x_1 y_1 \wedge \neg R y_1 x_2))\end{array} \tag{11}$$

The crucial problem in considering systems of superpostulates is specifying rules beyond Rules 1 and 2. Replacing Axiom 2 or Axiom 3 with its negation results in a complete variant of DLO_R with endpoints.[5] Thus, one might consider a rule that replaces axioms with their negations. However, the case in which completeness is preserved by simply replacing an axiom with its negation is an exceptional one. This is true only if one cannot replace an axiom A of a complete theory S with a strictly weaker axiom that is still independent of S without A. For example, there is no primary formula X that is strictly implied by $\forall x_1 \exists y_1 R x_1 y_1$ (=Axiom 2) and is not implied by Axiom 1 and Axioms 3–6 of (7). The same does not hold, e.g., for Axiom 1. Therefore, simply replacing Axiom 1 of (7) with its negation does not preserve completeness since $\exists y_1 \neg R y_1 y_1$ could be added as an independent axiom.

To preserve completeness when replacing an axiom A of an axiom set S with its negation, one must generate an axiom set S' that is *minimally weaker* than S and does not imply A.

Definition 14. *An axiom set* $S2$ *is* minimally weaker *than an axiom set* $S1$ *if* $S2$ *is strictly weaker than* $S1$ *and no intermediate axiom set* $S3$ *exists such that* $S3$ *is strictly weaker than* $S1$ *and* $S2$ *is strictly weaker than* $S3$.

Only adding the negation of an axiom A of S to S', which is minimally weaker than S, preserves completeness. However, the generation of a minimally weaker axiom set S' from a given axiom set S gives rise to intricate practical and theoretical problems. We cannot even prove whether a minimal weaker axiom set S' exists for any axiom set S. The main theoretical and practical problem is that one must specify some upper bound for an axiom set that preserves completeness if one axiom is replaced with its negation. We leave it as an open

[5] The same is not true in the case of negating Axiom 2 *and* Axiom 3. DLO with both endpoints is complete only if one adds the axiom $\exists y_1 \exists y_2 y_1 \neq y_2$. Without this further axiom, the remaining axioms have a trivial finite model in a domain with only one object. We abstain from generating a DLO_R-variant with left and right endpoints.

question whether this problem is solvable and, if so, how it can be solved. Instead, we confine our DLO_R-system to superpostulates of a certain length.

To do so, we adopt the following two restrictions:

Restriction 1: We set an upper bound for any considered axiom set with a length L of 1000 characters.

Restriction 2: We make use of an incomplete calculus that does not significantly increase the length of formulas when generating implied formulas.

These restrictions are justified by the *practical* aim of generating a system of infinity superpostulates.

Due to our restrictions, the considered superpostulates are simple enough to make use of the following assumption:

Assumption 1. If one of the following cases holds, this is proven by the standard casc-mode of Vampire within a time limit of 300 s:

1. a considered part of a considered axiom set is redundant,
2. a considered minimal axiom set is inconsistent, or
3. the considered axiom set is implied by another considered axiom set or implies either (10) or (12).

In fact, Vampire is able to almost immediately identify these cases for the simple axiom sets that we consider here in nearly all cases. By Assumption 1, we *assume* that Vampire can solve the semi-decidable problems 1 to 3 for axiom sets with fewer than 1000 characters to obtain a positive result within 300 s. In other words, our algorithm runs Vampire with a time limit of 300 s, and if the problem is not solved within this time limit, then we assume a negative solution. The reliability of Assumption 1 is based on extensive experience.

To generate strictly weaker axioms A' from an axiom A by applying an inference rule, we make use of an incomplete calculus that does not include rules for the introduction of disjuncts ($\vee I$) and conjuncts ($\wedge I$), which would increase the length of axioms. For this reason, we cannot ensure the generation of *minimally* weaker axiom sets and, thus, do not necessarily preserve completeness. Instead, we consider only rules concerning quantifiers: one rule for changing the order of quantifiers (QEx), two for universal quantifier elimination ($\forall$E1 and $\forall$E2) and two for existential quantifier introduction ($\exists$I1 and $\exists$I2); cf. Table 3. The abbreviated notations for the rules specify only the relevant syntactic changes. For example, $\exists\mu\forall\nu \vdash \forall\nu\exists\mu$ means that the order of two quantifiers is changed in a primary formula. We tacitly presume that a sequence of universal (existential) quantifiers is orderless; each of the quantifiers of such a sequence can be considered as either the leftmost or rightmost quantifier of the sequence. μ/ν means that μ is replaced with ν. In the case of $\forall$E1, ν may also be replaced with the variable of a new existential quantifier preceding the resulting axiom. $\forall\nu\varphi(\nu,\nu)$ indicates that ν occurs more than once in the scope of $\forall\nu$. $\varphi(\mu/\nu,\mu)$ means that at least one occurrence of μ is replaced with ν. We tacitly presume that new variables are used if a new quantifier is introduced.

Table 3. Quantifier rules

$$\text{QEx: } \exists\mu\forall\nu \vdash \forall\nu\exists\mu$$
$$\forall\text{E1: } \exists\mu\forall\nu\varphi(\mu,\nu) \vdash \exists\mu\varphi(\mu,\nu/\mu) \qquad \exists\text{I1: } \forall\nu\varphi(\nu,\nu) \vdash \forall\nu\exists\mu\varphi(\nu,\nu/\mu)$$
$$\exists\text{I2: } \exists\mu\varphi(\mu,\mu) \vdash \exists\mu\exists\nu\varphi(\mu/\nu,\mu) \quad \forall\text{E2: } \forall\mu\forall\nu\varphi(\mu,\nu,\nu) \vdash \forall\mu\varphi(\mu,\nu/\mu)$$

We establish the following rules for the application of the quantifier rules. $\forall$E1 can be applied such that $\exists\mu$ is taken from an axiom that differs from A after having variables renamed and PN laws applied to the effect that $\forall\nu$ is directly to the right of $\exists\mu$. Before $\forall$E1 can be applied to quantifiers of one and the same axiom, PN laws must be applied to the effect that existential quantifiers that are separated by $\wedge$ from other existential quantifiers are pulled outwards. The same holds for universal quantifiers that are separated by $\vee$ from other universal quantifiers. Thus, the numbers of quantifiers in a sequence of existential quantifiers and a sequence of universal quantifiers from which $\exists\mu$ and $\forall\nu$ are to be chosen are increased. $\forall$E2 and QEx can also be applied to universal quantifiers that are separated by $\vee$; QEx additionally can be applied to existential quantifiers that are separated by $\wedge$ from the universal quantifier. In these cases, PN laws are applied to exchange the order of the universal quantifiers and $\wedge$ or $\vee$. The results of the application of these rules are converted into a minimized FOLDNF in standardized notation (cf. p. 3). *One* application of a quantifier rule results in conversion from a set of primary formulas to a set of primary formulas. Thus, each application comprises the application of equivalence rules to prepare for the application of QEx, $\forall$E1 or $\forall$E2 and to convert the result into an FOLDNF. In the case in which the resulting FOLDNF is, in fact, a disjunction, we consider only disjuncts as axiom sets that are strictly implied by the axiom set prior to the application of a quantifier rule. Furthermore, we consider each conjunct of a disjunct of an FOLDNF as a separate axiom.

Given an axiom set S that includes axiom A, we consider the totality of all possible applications of the 5 quantifier rules to A. In the case of $\forall$E1, this may involve different orders of the prior application of PN laws in addition to the elimination of the universal quantified variable by variables bound by different existential quantifiers. QEx can be applied to all $\exists\mu$ and all $\forall\nu$ such that $\exists\mu$ occurs in a sequence of existential quantifiers directly to the left of a sequence of universal quantifiers containing $\forall\nu$. $\exists$I1 and $\exists$I2 may replace variables in different positions. The totality of possible applications of QEx, $\forall$E1, $\forall$E2, $\exists$I1, and $\exists$I2 to A realizes all different applications of the five rules.

Our method of generating a DLO_R-system is based on all possible *non-redundant* applications of the quantifier rules starting from (7).

Definition 15. *An application of a quantifier rule to an axiom A of an axiom set S is* non-redundant *if it results in a strictly weaker axiom set S' and the axiom A' that replaces A in S is not implied by S' without A'.*

If, however, S' implies S or if A' is not independent in S', then the application of a quantifier rule is redundant.

To ensure that all superpostulates are, in fact, *infinity* axiom sets, we consider only superpostulates that imply the very weak infinity axiom set given in (10) or its counterpart given in (12), as generated by Rule 2:

$$\forall x_1 \exists y_1 (Ry_1x_1 \wedge \neg Rx_1y_1 \wedge \forall x_2 (\neg Rx_1x_2 \vee Ry_1x_2)) \tag{12}$$

DLO_R without endpoints, with a left endpoint, and with a right endpoint all imply either (10) or (12).

By Assumption 1 and our quantifier rules, it is possible to apply the following algorithm to generate a set S^* of further superpostulates from a superpostulate S.

Algorithm 1. Generate S^* from S, which includes A, as follows:

1. Set $S^* = \{\}$.
2. Specify the set A' of all non-redundant applications of the quantifier rules to A.
3. Generate the powerset A'' of A'.
4. Traverse the members $A_1, A_2, \ldots, A_k$ of A'', beginning with the largest set A'. Replace A in S with the members of A_i and denote the result by S'; if $A_i = \{\}$, then $S' = S$ without A.
   (a) Delete redundant axioms and redundant subexpressions in S'.
   (b) If $S', \neg A$ is fewer than 1000 characters and consistent, then delete all proper subsets of A_i in A''.
   (c) If $\neg A$ is additionally independent of S', then denote the conversion of the resulting axiom set $S' \cup \{\neg A\}$ set into standard notation by S''.
   (d) If S'' implies (10) or (12), then append S'' to S^*.

Remark 3. For simplicity, our implemented algorithm in fact considers only A' (instead of its powerset A'') and merely deletes redundant axioms (and no other subexpressions) in S'.

Example 2. Applying Algorithm 1 to Axiom 1 of (7) results in a set S^* with the following superpostulate as its only member:

$$\begin{gathered} \exists y_1 Ry_1y_1, \forall x_1 \exists y_1 \neg Rx_1y_1, \forall x_1 \exists y_1 \neg Ry_1x_1, \\ \forall x_1 \exists y_1 (Rx_1y_1 \wedge \neg Ry_1y_1), \forall x_1 \exists y_1 (Ry_1x_1 \wedge \neg Ry_1y_1), \\ \text{Axioms 4, 5, and 6 of (7)} \end{gathered} \tag{13}$$

Note that the application of ∀E1 to replace x_1 in Axiom 1 of (7) with y_1 in Axiom 6 is redundant. Set (13) can, in turn, serve as a starting point for the application of Algorithm 1. Applying Algorithm 1 to Axiom 1 of (13) results in (7) once again (consider the implied minimization strategies!).

Let us use S^{**} to denote the result of applying Algorithm 1 to all of the axioms of S; S^{**} is the union of all results S^*. Furthermore, let us use S_T to denote the totality of the generated superpostulates. Finally, we use S^{***} to denote the subset of S^{**} that contains only superpostulates that are not implied

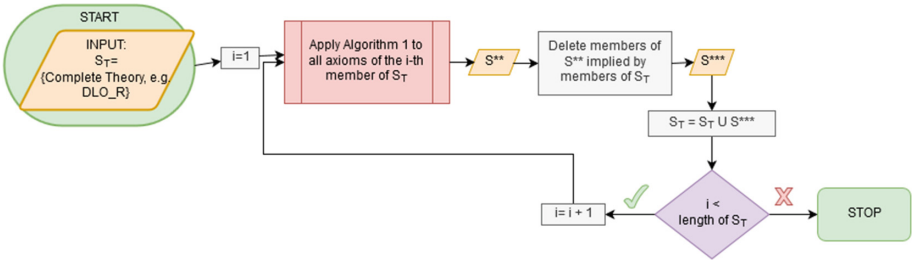

Fig. 1. Flowchart of Rule 3

by superpostulates in S_T. Figure 1 provides a flowchart of our general algorithm (Rule 3).

Based on Algorithm 1, we generate a DLO_R-system by means of the following rule:

Rule 3: Starting from (7), generate infinity superpostulates by iteratively applying Algorithm 1 to all axioms of all generated superpostulates until no further superpostulates are generated. Restriction: Delete superpostulates if they are implied by previously generated superpostulates.

The restriction ensures that no other infinity axiom set of S_T strictly implies S; this is the best we can do to generate L-complete infinity axiom sets up to a length of 1000 characters. Furthermore, the restriction avoids repetitions and loops (cf. the application of Algorithm 1 to Axiom 1 of (13) in Example 2).

Remark 4. With our implementation of Rule 3, we have generated 732 superpostulates in approximately 1000 h of run time. We have generated a database of 2196 infinity superpostulates in TPTP format by applying Rules 1 to 3 starting with DLO_R. None of these superpostulates can be currently solved by any of the TPTP-engines. Additionally, we have generated a database of 60 weak infinity DLO_R-variants. 8 of them can be solved by SPASS within 60 s.

5 The THS_R-System

A prominent alternative complete theory to start with is the theory of immediate successors (THS), which can be defined as follows in familiar notation:

$$\begin{gathered} \forall x \neg Rxx, \\ \forall x \exists y \forall z (Rxz \leftrightarrow z = y), \\ \exists x \forall y (\forall z \neg Rzy \leftrightarrow y = x), \\ \forall x (\exists y Ryx \rightarrow \exists y \forall z (Rzx \rightarrow z = y)) \end{gathered} \tag{14}$$

The following reduction THS_R of THS to FOL_R can be proven to be complete by adding the defining axiom $\forall x \forall y (x = y \leftrightarrow \forall z (Rxz \leftrightarrow Ryz) \wedge \forall z (Rzx \leftrightarrow Rzy))$ and proving the equivalence to (14):

$$\begin{gathered}\forall x_1 \neg Rx_1x_1, \exists y_1 \forall x_1 \neg Rx_1y_1, \forall x_1 \exists y_1 Rx_1y_1, \\ \forall x_1 \forall x_2 (\forall x_3 (\neg Rx_1x_3 \vee \neg Rx_2x_3) \vee \forall x_4 (\neg Rx_4x_1 \vee Rx_4x_2)), \\ \forall x_1 \forall x_2 (\forall x_3 (\neg Rx_3x_1 \vee \neg Rx_3x_2) \vee \forall x_4 (\neg Rx_1x_4 \vee Rx_2x_4)), \\ \forall x_1 \forall x_2 ((\forall x_3 (\neg Rx_1x_3 \vee Rx_2x_3) \vee \exists y_1 Ry_1x_2) \vee \exists y_2 Ry_2x_1)\end{gathered} \tag{15}$$

Vampire is able to solve (14) by saturation in 300 s, but it cannot solve (15) or variants thereof. Axioms 1, 2, and 3 of THS_R are part of DLO_R with a left endpoint, and Axiom 6 follows from Axiom 5 of DLO_R. Axioms 4 and 5 of THS_R contradict Axioms 1, 2, and 6 of DLO_R. Axioms 4, 5, and 6 of DLO_R contradict Axioms 1, 2, and 4 of THS_R.

Remark 5. Our algorithm has generated a database of 156 superpostulates from (15) and terminated after approximately 200 h of run time. 32 of them are solved by CVC4. None of the 156 superpostulates implies any of the superpostulates of the DLO_R-system or any of the weak DLO_R-variants.

6 Conclusion

From a theoretical point of view, it would be desirable to generate a database of complete infinity axiom sets. However, this would involve the problem of generating minimally weaker infinity axiom sets that may be of arbitrary length. Here, our intent is to provide a *pragmatic* algorithm that, in fact, generates significant infinity axiom sets. Thus, we restrict the generated infinity axiom sets to a certain length. This restriction does not imply that the generated sets are not complete theories; rather, it means only that we are unable to *guarantee* that they are. Our algorithm starts with a complete theory of FOL_R, such as DLO_R or THS_R, and automatically generates further infinity axiom sets S *with fewer than 1000 characters* such that, roughly speaking, no other infinity axiom sets *with fewer than 1000 characters* exist that strictly imply infinity axiom sets S. We say "roughly speaking" here because in addition to limiting the set length, we use an incomplete calculus without inference rules that would significantly increase complexity, and we test the restriction only on the basis of already generated entries in our database. We believe that the theoretical question concerning an (unrestricted) algorithm that does, in fact, preserve completeness is an important one, but we do not aim to answer it in this paper. Instead, our final objective is a practical algorithm, and we maintain that we have achieved this aim. Thus far, we have generated more than 2250 superpostulates in 10 weeks of run time. All of these superpostulates are hard problems.

Note that it is not trivial to automatically generate infinity axiom sets. Infinity axiom sets (not to say *systems* of infinity axiom sets) (i) are rarely encountered, (ii) can seldom be solved by logic engines, (iii) are hard problems even when they are short in length, and (iv) have not been systematically studied; moreover, (v) only a few of them exist in the TPTP library, and no systematic archive of infinity axiom sets is available. Therefore, we believe that our database is of practical importance for (i) systematically studying infinity axiom sets, (ii)

deciding the satisfiability of additional infinity axiom sets that follow from our database, and (iii) improving saturation algorithms by enabling the decidability of infinity axiom sets. In regard to (i), we are investigating the question of whether starting our algorithm with different complete FOL_R theories will yield identical superpostulates. MacPherson conjectured that any ω-categorical axiom set has the strict order property (cf. [10] and [4] for detailed discussions). Proofs of the property of being ω-categorical widely depend on identity. In light of MacPherson's conjecture, it would be interesting to study the limits of the DLO_R-system. We have also reduced Goldbach's conjecture to FOL_R and intend to study its internal relations with other infinity axiom sets in our still-growing database. In regard to (ii), we have proven in Sect. 3 that an infinite number of additional infinity axiom sets follow from our superpostulates. A database of superpostulates will make it possible to decide the satisfiability of infinity axiom sets that are implied by one of the entries of the database. In regard to (iii), we are currently working on saturation algorithms that can cope with infinity axiom sets. We know from our work how important it is to test such algorithms on a variety of intricate examples during development.

Acknowledgement. We thank the three anonymous referees for their comments and Geoff Suttcliffe for checking our database with CVC4, E, Infinox, SPASS and Vampire.

References

1. Börger, E., Grädel, E., Gurevich, Y.: The classical decision problem. J. Logic Lang. Inform. **8**, 478–481 (2001). https://doi.org/10.1023/A:1008334715902
2. Boolos, G.S., Burgess, J.P., Jeffrey, R.C.: Computability and Logic, 4th edn. Cambridge University Press, Cambridge (2003)
3. Claessen, K., Lilleström, A.: Automated inference of finite unsatisfiability. J. Autom. Reason. **47**, 111–132 (2011)
4. Djordjevic, M.: On first-order sentences without finite models. J. Symb. Log. **69**(2), 329–339 (2004)
5. Dreben, B., Goldfarb, W.D.: The Decision Problem. Solvable Classes of Quantificational Formulas. Addison-Wesley, London (1979)
6. Hilbert, D., Bernays, P.: Grundlagen der Mathematik, vol. 1. Springer, Berlin (1970)
7. Kalmár, L.: Zurückfürung des Entscheidungsproblems auf den Fall von Formeln mit einer einzigen, binären Funktions variablen. Compos. Math. **4**, 137–144 (1936)
8. Lampert, T.: Minimizing disjunctive normal forms of pure first-order logic. Log. J. IGPL **25**(3), 325–347 (2017)
9. Langford, C.: Analytic completeness of sets of postulates. Proc. London Math. Soc. **25**, 115–142 (1926)
10. Macpherson, D.: Finite axiomatizability and theories with trivial algebraic closure. Notre Dame J. Form. Log. **32**(2), 188–192 (1991)
11. Quine, W.V.O.: Methods of Logic, 4th edn. Harvard UP, Cambridge (1982)
12. Sheffer, H.M.: The general theory of notational relativity. Typewritten manuscript of 61 pp. Harvard Widener Library (1921)
13. Shoenfield, J.R.: Mathematical Logic. Addison-Wesley, Reading (1967)
14. Urquhart, A.: Henry M. Sheffer and notational relativity. Hist. Philos. Log. **33**(1), 33–47 (2012)