Università degli studi di Udine

Set-syllogistics meet combinatorics

Original

Availability:
This version is available http://hdl.handle.net/11390/1119226 since 2017-10-19T13:55:02Z

Publisher:

Published
DOI:10.1017/S0960129515000122

Terms of use:
The institutional repository of the University of Udine (http://air.uniud.it) is provided by ARIC services. The aim is to enable open access to all the world.

Publisher copyright

(Article begins on next page)

12 October 2024

Under consideration for publication in Math. Struct. in Comp. Science

Set-syllogistics meet combinatorics†

Eugenio G. Omodeo[1], Alberto Policriti[2], Alexandru I. Tomescu[3]

[1] *Dipartimento di Matematica e Geoscienze, Università di Trieste, Via Valerio 12/1, I-34127 – Trieste, Italy. email:* `eomodeo@units.it`
[2] *Dipartimento di Matematica e Informatica, Università di Udine, Via delle Scienze 206, I-33100 – Udine, Italy. email:* `alberto.policriti@uniud.it`
[3] *Helsinki Institute for Information Technology HIIT, Department of Computer Science, University of Helsinki, P.O. 68 (Gustaf Hällströmin katu 2b), FI-00014 – Helsinki, Finland. email:* `tomescu@cs.helsinki.fi`

Received 26 February 2015

This paper considers $\exists^*\forall^*$ prenex sentences of pure first-order predicate calculus with equality. This is the set of formulas which F.P. Ramsey's treated in a famous article of 1930. We demonstrate that the satisfiability problem and the problem of existence of arbitrarily large models for these formulas can be reduced to the satisfability problem for $\exists^*\forall^*$ prenex sentences of Set Theory (in the relators $\in, =$).
We present two satisfiability-preserving (in a broad sense) translations $\Phi \mapsto \dot{\Phi}$ and $\Phi \mapsto \Phi^{\boldsymbol{\sigma}}$ of $\exists^*\forall^*$ sentences from pure logic to well-founded Set Theory, so that if $\dot{\Phi}$ is satisfiable (in the domain of Set Theory) then so is Φ, and if $\Phi^{\boldsymbol{\sigma}}$ is satisfiable (again, in the domain of Set Theory) then Φ can be satisfied in arbitrarily large finite structures of pure logic. It turns out that $\left|\dot{\Phi}\right| = \mathcal{O}(|\Phi|)$ and $|\Phi^{\boldsymbol{\sigma}}| = \mathcal{O}\left(|\Phi|^2\right)$.
Our main result makes use of the fact that $\exists^*\forall^*$ sentences, even though constituting a decidable fragment of Set Theory, offer ways to describe infinite sets. Such a possibility is exploited to glue together infinitely many models of increasing cardinalities of a given $\exists^*\forall^*$ logical formula, within a single pair of infinite sets.

Keywords: Bernays-Schönfinkel-Ramsey class, spectrum of a first-order prenex sentence, infinite sets, satisfiability decision algorithms, computable Set Theory.

Introduction

Multi-level syllogistic (Ferro *et al.* 1980; Breban *et al.* 1981; Cantone *et al.* 2001; Cantone 2012) is a decision algorithm which determines whether a given formula involving only individual variables, which designate sets, and a restricted collection of set operators, is satisfiable.

By and large, multi-level syllogistic has the ability to check a prenex $\exists^*\forall$-sentence in

† Work partially supported by the INdAM/GNCS 2013 project *"Specifica e verifica di algoritmi tramite strumenti basati sulla teoria degli insiemi"*, by the Academy of Finland under grants 250345 (CoECGR) and 274977, and by FRA-UniTS PUMA.

the relators $\in, =$ for truth over sets. In practice, the decision algorithm does not handle quantifiers explicitly and needs not eliminate the dyadic operators $\cup, \setminus$ or the monadic operator $\{\,_\,\}$ (to mention a few). But the very possibility to do this reduction gives us a clue on the power of the decision method; to the authors, it made plain how to adapt the method to Aczel's non-standard view on sets (Omodeo and Policriti 1995);[†] moreover, by bringing set-theoretic syllogistics closer to the stream of classical research on the decision problem for predicate calculus (Börger *et al.* 1997), it suggested ways to reinforce the known decidability results about those syllogistics.

In recent papers we have moved on to the much larger class, named BSR,[‡] of all $\exists^*\forall^*$-sentences, studied in the framework of Set Theory.

Concerning *pure* logic, namely first-order predicate calculus with equality, the $\exists^*\forall^*$ satisfiability problem was solved long ago by Bernays and Schönfinkel. Frank Plumpton Ramsey, by analyzing the full spectrum of interpretations modeling each sentence in this class (over an arbitrary, uninterpreted signature), got a foundational result in combinatorics (Ramsey 1930).

In pure logic without equality, it is easy to arbitrarily *enlarge* the size of a structure satisfying a given BSR formula Φ. When equality constraints enter into play, they provide means to bound *from above* the cardinality of the underlying domain. The essence of Ramsey's combinatorial analysis was the proof that when an $\exists^*\forall^*$ sentence Φ with equality can be satisfied in a structure whose domain's cardinality is an integer exceeding a specific computable threshold $\mathfrak{r}(\Phi)$, then Φ admits models of every size larger than $\mathfrak{r}(\Phi)$. Consequently, infinity cannot be captured in pure logic by $\exists^*\forall^*$ sentences.

Partly influenced by Ramsey's historical success, we tackled the BSR truth problem in the context of Set Theory. Today that problem has been solved (Omodeo and Policriti 2010; Omodeo and Policriti 2012) for sets in von Neumann's hierarchy of well-founded sets. In this paper we continue to study the connections between BSR formulae in the framework of pure logic and in the one of Set Theory. More specifically, we reduce Ramsey's spectral problem for a BSR logical formula to the solvable satisfaction problem for a set-theoretic BSR formula. As will turn out, the length of the target formula of the reduction will be quadratic in the length of the original formula.

Instrumental to our result is the fact that within Set Theory one can express the existence of infinite sets by way of a prenex $\exists\exists\forall\forall$ sentence, e.g. by the sentence[§]

$$\exists\, x_0\, \exists\, x_1\; \iota\!\iota(x_0, x_1),$$

where

$$\iota\!\iota(x_0, x_1) \;\leftrightarrow_{\text{Def}}\; \left(\begin{array}{c} x_0 \neq x_1 \;\wedge\; x_0 \notin x_1 \;\wedge\; x_1 \notin x_0 \;\wedge\; \bigcup x_0 \subseteq x_1 \;\wedge\; \bigcup x_1 \subseteq x_0 \;\wedge \\ (\forall\, y_0 \in x_0)\,(\forall\, y_1 \in x_1)\,(y_0 \in y_1 \vee y_1 \in y_0) \end{array} \right),$$

whose existential variables admit no simpler model than $\boldsymbol{x}_0 = \omega_1$, $\boldsymbol{x}_1 = \omega_0$, and $\boldsymbol{x}_2 = \emptyset$,

[†] Save for occasional mentions—like here—of Aczel's theory, this paper will take for granted that the membership relation is well-founded all over the universe of sets.

[‡] This is an acronym for *Bernays-Schönfinkel-Ramsey*.

[§] One can eliminate '=' from this formula, at the price of introducing one more existential quantifier.

$$\iota\iota(x_0, x_1) \leftrightarrow_{\mathrm{Def}} \quad \exists\, x_2\, \forall\, y_0\, \forall\, y_1 \left(\begin{array}{c} (x_2 \in x_0 \leftrightarrow x_2 \notin x_1) \ \wedge\ x_0 \notin x_1 \ \wedge\ x_1 \notin x_0 \ \wedge \\ \big((y_0 \in y_1 \wedge y_1 \in x_0) \rightarrow y_0 \in x_1\big) \ \wedge \\ \big((y_0 \in y_1 \wedge y_1 \in x_1) \rightarrow y_0 \in x_0\big) \ \wedge \\ \big((y_0 \in x_0 \wedge y_1 \in x_1) \rightarrow (y_0 \in y_1 \vee y_1 \in y_0\big) \end{array} \right)$$

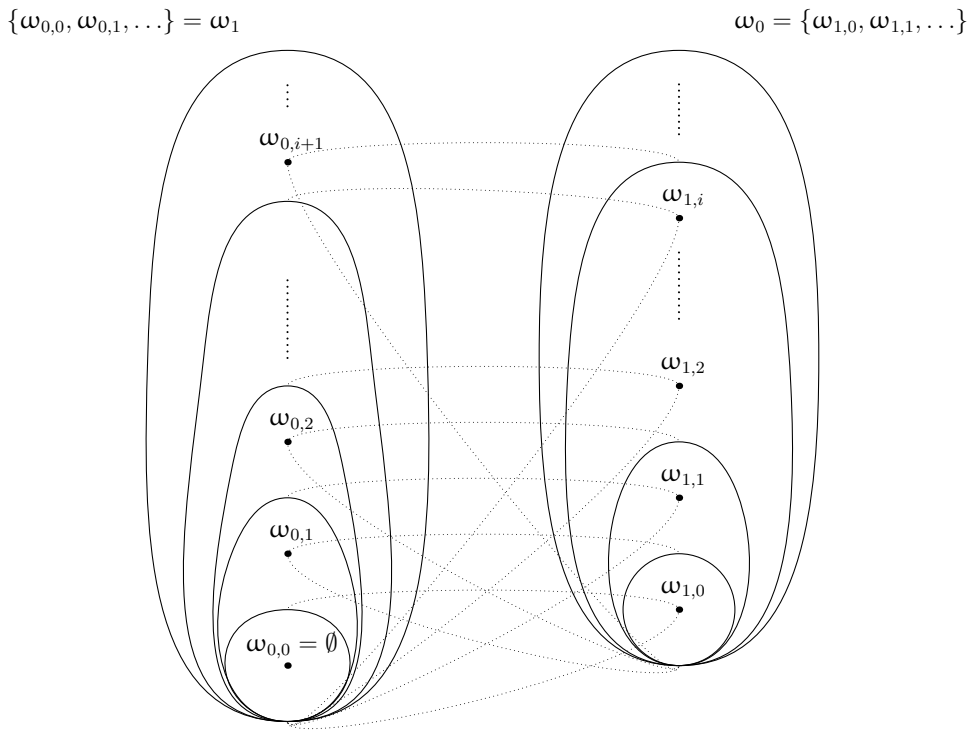

Fig. 1. In the upper part $\iota\iota(x_0, x_1)$ is reformulated without derived symbols: along with $\bigcup$ and $\subseteq$, even $=$ has been eliminated, causing a third existentially quantified variable to appear. The lower part shows a double-stranded infinity such that $\iota\iota(\omega_1, \omega_0)$ holds.

where ω_0 and ω_1 are as shown in Fig. 1. More generally, for each $n > 1$, one can state the existence of n infinite sets by means of an $\underbrace{\exists \cdots \exists}_{n \text{ times}} \exists \underbrace{\forall \cdots \forall}_{n \text{ times}}$ sentence, e.g. in the way shown by the template in Fig. 2.

On the basis of this remark and by paralleling the techniques involved in our decision method with Ramsey's combinatorics, in (Omodeo *et al.* 2012) we have begun to study the possibility of analyzing the spectrum of any $\exists^*\forall^*$ logical sentence by translating it into a set-theoretic BSR formula, so that the infinitude of the spectrum of the former can be revealed simply through the satisfaction of the latter. This paper concretizes

$$\exists\, x_1 \cdots \exists\, x_n \exists\, x_{n+1} \Big(\quad x_{n+1} \in x_1 \;\wedge \\ \bigwedge_{i=1}^{n} \big(x_i \notin x_{(i \mod n)+1}\big) \;\wedge\; \bigwedge_{i=1}^{n} \big(\bigcup x_{(i \mod n)+1} \subseteq x_i\big) \;\wedge \\ (\forall y_1 \in x_1) \cdots (\forall y_n \in x_n) \big(\bigvee_{i=1}^{n} y_i \in y_{(i \mod n)+1}\big) \Big)$$

Fig. 2. BSR formula whose satisfaction calls for infinite $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n$ $(n > 1)$.

that plan; as a consequence, it makes previous results on syllogistics, i.e. on decidable fragments of Set Theory, exploitable not only as an aid to correct reasoning but also to offer a combinatorial means to collectively specify all possible ways of satisfying a given logical sentence.

1. Testing set-theoretic BSR sentences for truth

Testing set-theoretic BSR sentences is not an easy task and we can only give, in this section, a very sketchy account of the result in (Omodeo and Policriti 2010) and (Omodeo and Policriti 2012), to which the reader is referred for a complete account.

The task consists in establishing whether a given formula

$$\forall\, y_1 \cdots \forall\, y_m\, \varphi(x_1, \ldots, x_n, y_1, \ldots, y_m),$$

in the relators $\in, =$ can, or cannot, be made true by an assignment of sets to its existential variables x_i. In the affirmative case our algorithm also produces a (finite representation of a) *model*, i.e., a satisfying assignment. In this sense, it does not act as a simple-minded satisfiability tester, but as a *satisfaction algorithm* which constructs a model whenever possible.

Within Set Theory one can express the existence of infinite sets by way of a prenex $\exists\exists\exists\forall\forall$ sentence (as recalled above), but *not* by way of an $\exists^*\forall$ sentence (Parlamento and Policriti 1988; Omodeo *et al.* 2012). In raising the skills of a decision method from the $\exists^*\forall$- to the $\exists^*\forall^*$-class, one encounters here a major challenge; also, each universal quantifier can add intricacy to the interplay among the infinite sets in a satisfying assignment.

Addressing the decision problem for the entire $\exists^*\forall^*$ class in a single shot, offers a pleasant initial facilitation: thanks to *extensionality* (according to which, distinct sets cannot have exactly the same elements), one can get rid of the equality symbol. In practice, one replaces the given sentence $\exists\, x_1 \cdots \exists\, x_n \forall\, y_1 \cdots \forall\, y_m \varphi$ by a finite collection Ψ of $\forall^*$-formulae so that $\forall\, y_1 \cdots \forall\, y_m \varphi$ can be satisfied through an assignment $x_i \mapsto \boldsymbol{x}_i$ of sets to its existential variables if and only if at least one formula ψ in Ψ can be satisfied *injectively*, i.e. by means of an assignment whose images are pairwise different sets. One can manage that each ψ in Ψ be devoid of the symbol $=$, usually at the price of introducing new existential variables.

Another essential preparation of the formulae to be tested for injective satisfiability consists in bounding the universal variables: specifically, on the grounds of a reduction carried out in (Omodeo and Policriti 2010, pp. 468–470), one can assume the following

RESTRICTED format for formulae of the BSR class:

$$\Phi \quad = \quad \bigwedge_{i=1}^{\kappa} (\forall y_1 \in z_1) \cdots (\forall y_{m_i} \in z_{m_i})\ \phi_i(x_1, \ldots, x_n, y_1, \ldots, y_{m_i})\ ,$$

where $z_h \in \{x_1, \ldots, x_n, y_1, \ldots, y_{h-1}\}$ for $h \in \{1, \ldots, m_i\}$, and *equality does not appear* in any of the unquantified matrices $\phi_i(x_1, \ldots, x_n, y_1, \ldots, y_{m_i})$.

We must focus on models of a special, *irredundant* nature which can be captured by a finite (di)graph structure $\mathcal{G}$ on the one hand and can also suggest, on the other hand, how to compute a bound on the size of $\mathcal{G}$. Arcs will represent the inverse $\ni$ of membership restricted to the sets associated with the nodes.

Let $x_i \overset{\mathcal{M}}{\longmapsto} \boldsymbol{x}_i$ be an injective model of Φ and consider the transitive membership closure $\mathsf{TrCl}(\mathcal{F})$ of the family $\mathcal{F}$ of sets onto which the x_i's are mapped by $\mathcal{M}$. Redundancy might derive from the presence of overly complex infinite sets in $\mathsf{TrCl}(\mathcal{F})$. As proved in (Omodeo and Policriti 2010), the only unescapable kinds of infinitude can be described by means of formulae falling under the template of Fig. 2. These infinite sets are internally organized in regular structures: in a faithful graph representation of $\mathsf{TrCl}(\mathcal{F})$, each one of these structures would form a peculiar ascending membership spiral. In $\mathcal{G}$ these situations will be encoded by finite cycles. $\mathsf{TrCl}(\mathcal{F})$ will consist of nodes appearing in the said spirals, and of additional nodes forming the so-called *core* of $\mathcal{M}$, which includes the $\boldsymbol{x}_i$'s.

In (Omodeo and Policriti 2012) we tackled the problem of setting a bound on the size of the core, and to compute it on the basis of how many existential/universal variables appear in Φ. Thanks to this computable bound, the semi-decision algorithm proposed in (Omodeo and Policriti 2010) evolved into a decision algorithm.

To pinpoint additional restrictions on the nature of a model $\mathcal{M}$ worth of consideration, we can insist that $\mathsf{TrCl}(\mathcal{F})$ owns no more elements per rank¶ than the number n of x_i's. To these restrictions (and a few more), which appeared already in (Omodeo and Policriti 2010), we added an important one in (Omodeo and Policriti 2012): *the core has least possible cardinality.* Altogether, the irredundancy assumptions enable us to get a bound on the cardinality of $\mathcal{G}$. In particular, the bound on the size of the core is obtained very much in the spirit of the original Ramsey's result. Two steps are necessary: an equivalence relation of finite index on tuples of sets in the core (actually, on their membership graphs) and an application of the pigeonhole principle to a "striped" version of the core. The first step allows one to classify the elements of $\mathsf{TrCl}(\mathcal{F})$ into finitely many *types*, in such a way that different elements of the same type can be interchangeably used to construct a model, as far as the satisfaction of the given BSR formula is concerned. Then, after having subdivided the core into "stripes", one uses the pigeonhole principle to contract $\mathcal{M}$ into another satisfying assignment *if any of its stripes repeats.* Such a contraction, if doable, would lead to a smaller core, which is absurd.

2. Expressiveness of the BSR set-theoretic class

The BSR set-theoretic class turns out to be much more expressive than the corresponding class of formulae interpreted in merely logical terms. The observation, already made, that

¶ A recursive formulation of the *rank* function from sets to ordinals is: $\mathsf{rk}(X) = \sup\{\, \mathsf{rk}(y) + 1 : y \in X \,\}$.

$$\Psi_1 \equiv \exists\, x_1\, \exists\, x_2\, \exists\, x_3\, \forall\, y \big(x_1 \in x_3 \wedge x_2 \notin x_3 \wedge (\, y \in x_1 \leftrightarrow y \in x_2\,)\big),$$

$$\Psi_2 \equiv \forall\, y_1 \cdots \forall\, y_n \left(\bigvee_{i=0}^{n-1} \bigvee_{j=i+1}^{n} \bigwedge_{k=1}^{n} (\, y_k \in x_i \leftrightarrow y_k \in x_j\,) \right),$$

$$\Psi_3 \equiv \exists\, y_0 \cdots \exists\, y_n \Big(\bigwedge_{i=0}^{n-1} \bigwedge_{j=i+1}^{n} \bigvee_{k=0}^{n} (\, y_k \in x_i \leftrightarrow y_k \notin x_j\,) \wedge \bigvee_{k=0}^{n} \bigwedge_{i=0}^{n} (\, y_k \notin y_i \wedge (\, i \neq k \rightarrow y_i \in x_i\,)\,) \Big).$$

Fig. 3. In Set Theory, Ψ_1 is a false $\exists^*\forall$ sentence, Ψ_2 is an injectively unsatisfiable $\forall^*$ scheme, and Ψ_3 is the negation of an injectively unsatisfiable $\forall^*$ scheme.

infinity can be captured by a BSR formula in the set-theoretic framework but not in the purely logical one, gives evidence of the higher expressiveness of the former language. At a more elementary level, this can be seen from the formulae Ψ_1, Ψ_2, and Ψ_3 displayed in Fig. 3: their status, which is indicated in the caption of that figure, depends either on extensionality alone or (as for the third of them and richer variants of it, cf. (Omodeo and Policriti 1995)) on very little more.

To be more specific about the expressive power of the BSR set-theoretic class, we will now discuss a satisfiability-preserving translation of BSR sentences from an uninterpreted, purely logical context into one referring to a model $\mathcal{U} = (\mathcal{U}, \in)$ of the standard Zermelo-Fraenkel theory of sets. We make the simplifying assumption that the language $\mathcal{L}$ of pure logic has only one dyadic relator ϱ and equality: $\mathcal{L} \equiv \mathcal{L}_\varrho$. To see that this assumption is, in fact, inessential, it is sufficient to observe that any occurrence of an n-ary relational symbol $R(z_1, \ldots, z_n)$ other than ϱ can be replaced by the following conjunction of n atomic formulae

$$\varrho(z_1, x_1^R) \wedge \ldots \wedge \varrho(z_n, x_n^R),$$

where $x_1^R, \ldots, x_n^R$ are (fresh) existentially quantified variables, to be used to eliminate R only. Roughly speaking, $\varrho(\cdot, x_j^R)$ captures the j-th *projection* $R_j = \{z_j \mid R(z_1, \ldots, z_n)\}$ of R.

Let us stress again that in set theory—as opposed to the case of logic—and in connection to the satisfiability problem at hand, it is immaterial whether or not we regard equality as a *primitive* relator in the signature of the language. Anyhow, since we know that we can eliminate '=' from set-theoretic BSR sentences without leaving the BSR class, we feel free to use it when this can improve readability.

We want to convert any given BSR sentence

$$\Phi \quad \equiv \quad \exists\, x_1 \cdots \exists\, x_n\, \forall\, y_1 \cdots \forall\, y_m\, \varphi(x_1, \ldots, x_n, y_1, \ldots, y_m)\,,$$

where φ is an unquantified matrix in the language $\mathcal{L}_\varrho$, into a BSR sentence $\dot{\Phi}$ in the language $\mathcal{L}_\in$ *interpreted* in $\mathcal{U}$, much as we did in (Omodeo *et al.* 2012, Sec. 4), whose target language referred to Aczel's non-well-founded sets. In that paper, taking advantage of the non-well-foundedness of membership, we could simply translate Φ into

$$(\exists\, d)\ (\exists\, x_1 \in d) \cdots (\exists\, x_n \in d)\ (\forall\, y_1 \in d) \cdots (\forall\, y_m \in d)\ \varphi_\in^\varrho(x_1, \ldots, x_n, y_1, \ldots, y_m)\,,$$

where $\varphi^{\varrho}_{\in}$ results from φ through uniform replacement of ϱ by the membership sign. Here we prefer to replace ϱ by the converse, $\ni$, of membership. Moreover, we must proceed in a slightly more roundabout fashion, because ϱ can be cyclic while $\in$, by axiomatic assumption, cannot. We overcome this problem by representing each logical variable z in *split* form, by means of a source-target pair, $z_{\mathrm{s}}, z_{\mathrm{t}}$, of set-variables. This transformation reflects a common way of proceeding in graph theory, for example to reduce cycle cover problems to matching problems in bipartite graphs (cf. (Plummer and Lovász 1986)).

Theorem 2.1. To each BSR sentence Φ in $\mathcal{L}_{\varrho}$ there corresponds a BSR sentence $\dot{\Phi}$ in $\mathcal{L}_{\in}$ such that

Φ is satisfiable if and only if $\dot{\Phi}$ is satisfiable in well-founded Set Theory.

Proof. For any given BSR sentence

$$\Phi \equiv \exists\, x_1 \cdots \exists\, x_n\, \forall\, y_1 \cdots \forall\, y_m\, \varphi(x_1, \ldots, x_n, y_1, \ldots, y_m)\ ,$$

where φ is an unquantified matrix in the language $\mathcal{L}_{\varrho}$, put

$$\begin{aligned}\dot{\Phi} \equiv &(\exists d)\ (\exists x_{\mathrm{s},1}, x_{\mathrm{t},1}, \ldots, x_{\mathrm{s},n}, x_{\mathrm{t},n} \in d)\ (\forall y_{\mathrm{s},1}, y_{\mathrm{t},1}, \ldots, y_{\mathrm{s},m}, y_{\mathrm{t},m} \in d)\\ &\dot{\varphi}(x_{\mathrm{s},1}, x_{\mathrm{t},1}, \ldots, x_{\mathrm{s},n}, x_{\mathrm{t},n}, y_{\mathrm{s},1}, y_{\mathrm{t},1}, \ldots, y_{\mathrm{s},m}, y_{\mathrm{t},m})\ ,\end{aligned}$$

where $\dot{\varphi}$ results from φ through replacement of each literal of the form $z_i\, \varrho\, w_j$, with $z, w \in \{x, y\}$, by

$$z_{\mathrm{s},i} \ni w_{\mathrm{t},j}\ ,$$

and of each literal of the form $z_i = w_j$, with $z, w \in \{x, y\}$, by the conjunction

$$z_{\mathrm{s},i} = w_{\mathrm{s},j} \wedge z_{\mathrm{t},i} = w_{\mathrm{t},j}\ .$$

To prove one of the implications in our claim, assume first that Φ is satisfiable, that $\langle D, R\rangle$ is a finite structure satisfying it, and think of $\langle D, R\rangle$ as a directed graph which could undergo the following cycle-untying transformation: replacement of each node v by distinct nodes v_{s} and v_{t}, and of each arc $\langle u, w\rangle \in R$ by an arc leaving u_{s} and entering w_{t}. Bearing this transformation in mind, consider functions f, g from $D_{\mathrm{s,t}} = \{v_{\mathrm{s}}, v_{\mathrm{t}} : v \in D\}$ into sets subject to the following constraints:

— $f(u_{\mathrm{s}}) = \{\, f(w_{\mathrm{t}}) : \langle u, w\rangle \in R\} \cup \{g(u_{\mathrm{s}})\}$;
— $f(u_{\mathrm{t}}) = \{g(u_{\mathrm{t}})\}$;
— the function g is injective and $|g(v)| \neq 1$ for every v.

Once fixed the function g, the function f is determined uniquely in view of the acyclicity of the graph $\langle D_{\mathrm{s,t}}, \{\langle u_{\mathrm{s}}, w_{\mathrm{t}}\rangle : \langle u, w\rangle \in R\}\rangle$.

The functions f and g associate two sets with each node in $D_{\mathrm{s,t}}$, mimicking R by the acyclic relation $\ni$ even in case R has cycles. The function f is injective on $\{v_{\mathrm{t}} : v \in D\}$, by the injectivity of g. Moreover, for every u and w, $f(w_{\mathrm{t}}) \neq g(u_{\mathrm{s}})$, since $|f(w_{\mathrm{t}})| = 1$ while $|g(u_{\mathrm{s}})| \neq 1$. Therefore, f is injective on the whole $D_{\mathrm{s,t}}$, since if $u_{\mathrm{s}} \neq u'_{\mathrm{s}}$ then $g(u_{\mathrm{s}}) \in f(u_{\mathrm{s}}) \setminus f(u'_{\mathrm{s}})$ (and, symmetrically, $g(u'_{\mathrm{s}}) \in f(u'_{\mathrm{s}}) \setminus f(u_{\mathrm{s}})$).

Based on the injectivity—just proved—of the Mostowski-like collapsing function f, equality as well as membership literals are properly modelled: in fact, if one interprets

d as $\{f(v) : v \in D_{\mathrm{s,t}}\}$ and $x_{\mathrm{s},i}, x_{\mathrm{t},i}$ as $f(v_{\mathrm{s},i}), f(v_{\mathrm{t},i})$, where v_i is the node assigned to x_i by the satisfying assignment for Φ, then the resulting set-assignment will satisfy $\dot{\Phi}$.

Conversely, assuming $\dot{\Phi}$ is satisfied by a set-theoretic interpretation, define $\langle D, R\rangle$ to be the graph with nodes $D = \{v_1, \ldots, v_n\}$ such that $v_i = v_j$ holds precisely when the interpretations $\boldsymbol{x}_{\mathrm{t},i}, \boldsymbol{x}_{\mathrm{s},i}$ of $x_{\mathrm{t},i}$ and $x_{\mathrm{s},i}$ equal the corresponding interpretations, $\boldsymbol{x}_{\mathrm{t},j}, \boldsymbol{x}_{\mathrm{s},j}$, of $x_{\mathrm{t},j}$ and $x_{\mathrm{s},j}$, and with arcs $R = \{\langle v_i, v_j\rangle : i, j = 1, \ldots, n \mid \boldsymbol{x}_{\mathrm{t},j} \in \boldsymbol{x}_{\mathrm{s},i}\}$. On these grounds we have that

$$\begin{aligned} v_i = v_j \text{ if and only if } \boldsymbol{x}_{\mathrm{s},i} = \boldsymbol{x}_{\mathrm{s},j} \wedge \boldsymbol{x}_{\mathrm{t},i} = \boldsymbol{x}_{\mathrm{t},j}, \\ \langle v_i, v_j\rangle \in R \text{ if and only if } \boldsymbol{x}_{\mathrm{t},j} \in \boldsymbol{x}_{\mathrm{s},i}, \end{aligned}$$

from which it plainly follows that Φ is satisfiable in $\langle D, R\rangle$. □

3. Specifying the infinite spectrum of a BSR sentence

The BSR class, even in pure logic, has an—admittedly limited—self-referential ability: we can easily write a formula that can force every structure satisfying it, to have at least a certain amount of elements. As a consequence, Ramsey's celebrated combinatorial theorem enables one to capture, via BSR sentences, interesting properties of the *collection* of models of a BSR sentence Φ of first-order predicate calculus. Very straightforwardly, if Φ belongs to $\mathcal{L}_\varrho$, we can state that Φ owns models whose domains of support are arbitrarily large by constructing another BSR sentence, $\widehat{\Phi}$, of $\mathcal{L}_\varrho$ which is satisfiable if and only if Φ is as wanted. We can, in fact, obtain $\widehat{\Phi}$ by introducing $\mathfrak{r}(\Phi)$ new existential variables, where $\mathfrak{r}(\Phi)$ is Ramsey's threshold number mentioned earlier. But, notice, with this approach the size of $\widehat{\Phi}$ will be very big, because $\mathfrak{r}(\Phi)$ is known to grow exponentially relative to the size of Φ (see (Radziszowski 2014) for an updated survey). Proceeding less naïvely, we will now specify this same property, that a given Φ has an infinite spectrum, by means of a sentence $\Phi^{\boldsymbol{\sigma}}$ of $\mathcal{L}_\in$. In proving the correctness of our translation $\Phi \mapsto \Phi^{\boldsymbol{\sigma}}$, we will rely on Ramsey's combinatorial theorem; nevertheless, the size of $\Phi^{\boldsymbol{\sigma}}$ will depend quadratically on the size of Φ: an improvement which adds evidence of the greater expressive power of the BSR set-theoretic class with respect to the BSR class of pure logic.

The specification proposed above of double-stranded infinity—$\iota\iota(d_0, d_1)$, see Fig. 1—will play a major role in our translation $\Phi \mapsto \Phi^{\boldsymbol{\sigma}}$. Let us recall here some properties enjoyed by any pair $\boldsymbol{d}_0, \boldsymbol{d}_1$ of sets that satisfy $\iota\iota(\boldsymbol{d}_0, \boldsymbol{d}_1)$, which we need for Theorem 3.1:

— $\mathsf{rk}(\boldsymbol{d}_0) = \mathsf{rk}(\boldsymbol{d}_1)$ and this rank is a limit ordinal;
— $\boldsymbol{d}_0 \cap \boldsymbol{d}_1 = \emptyset$;
— $(\bigcup \boldsymbol{d}_0) \cap \boldsymbol{d}_0 = \emptyset$;
— for every $\boldsymbol{z} \in \boldsymbol{d}_1$, the set $\boldsymbol{d}_0 \setminus \boldsymbol{z}$ is the infinite set consisting of all elements of $\boldsymbol{d}_0$ whose rank exceeds $\mathsf{rk}(\boldsymbol{z})$.

The above results are easily seen to follow from the definition of $\iota\iota(\boldsymbol{d}_0, \boldsymbol{d}_1)$. Proofs can also be found in (Omodeo *et al.* 2012).

A forthcoming theorem is the main result in this paper and is proved using a set-theoretic encoding of infinitely many (finite) structures within $\boldsymbol{d}_0 \cup \boldsymbol{d}_1$. The encoding is to be designed building on the idea of splitting graph nodes into source-target pairs, as done for the proof of Theorem 2.1. However, the main problem now is not so much the encoding of a possibly cyclic binary relation via well-founded membership, as is the issue that infinitely many arcs must be encoded. Moreover, this must be done by means of elements of increasing ranks that satisfy the constraints imposed by $\iota\iota(\boldsymbol{d}_0, \boldsymbol{d}_1)$.

In preparation for the announced main theorem, let us digress momentarily to observe a useful combinatorial fact (relying on Ramsey's celebrated theorem) about infinite sequences of finite digraphs.

Definition 3.1. Relative to a digraph G with nodes $1, \ldots, n, n+1, \ldots, n+\ell$, call n-TYPE of each node $w \in \{n+1, \ldots, n+\ell\}$ the digraph resulting from G when its node w gets replaced by 0 and then all nodes other than $0, 1, \ldots, n$ are withdrawn, i.e., they are removed from the graph along with their incident arcs.

For each n-type τ of (a node of) a digraph G as above, indicate by $G \restriction \tau$ the subgraph that results from G when every node w of type other than τ gets withdrawn.

Taking advantage of our simplifying assumption—one dyadic relation only—, we can tailor Ramsey's original notion of *homogeneity* to our context. We will say that a digraph G as above (hence endowed with at least n nodes) is n-HOMOGENEOUS when its nodes other than $1, \ldots, n$ have the same type and they form either an *independent set* or a *clique* in G (i.e., either $G \setminus \{1, \ldots, n\}$ has no arcs or there are arcs in both directions between any two elements of $G \setminus \{1, \ldots, n\}$).

Lemma 3.1. Let $G_1, G_2, G_3, \ldots$ be an infinite sequence of digraphs such that each G_j has nodes $1, \ldots, n+\ell_j$ and $0 < \ell_1 < \ell_2 < \cdots$.

Then, for a suitable n-type θ, there is an infinite subsequence $G_{i_1}, G_{i_2}, G_{i_3}, \ldots$ of the given one such that each graph $G_{i_j} \restriction \theta$ has an n-homogeneous subgraph Γ_j endowed with $n+j$ nodes which include the nodes $1, \ldots, n$.

Proof. For $j = 1, 2, \ldots$, let $\tau_{j,1}, \ldots, \tau_{j,\ell_j}$ be the types of $n+1, \ldots, n+\ell_j$ in G_j. Altogether, the number of distinct n-tpes is bounded by the finite amount $2^{(n+1)^2}$; hence, in order that the sizes of the G_j's can increase indefinitely, there must exist an n-type θ such that, indicating by $t_j(\theta)$ the number of times θ occurs within each sequence $\tau_{j,1}, \ldots, \tau_{j,\ell_j}$, the set $\{t_1(\theta), t_2(\theta), t_3(\theta), \ldots\}$ has no maximum. In fact, arguing by contradiction and indicating by t_θ the maximum corresponding to each θ, we would have $\ell_k \leqslant \sum_\theta t_\theta$ for any k, contradicting $0 < \ell_1 < \ell_2 < \cdots$.

This plainly implies that we can extract an infinite subsequence $G_{i'_1}, G_{i'_2}, G_{i'_3}, \ldots$ of $G_1, G_2, G_3, \ldots$ so that the graphs $G_{i'_1} \restriction \theta, G_{i'_2} \restriction \theta, G_{i'_3} \restriction \theta, \ldots$ have increasing sizes. By the Finite Ramsey Theorem (specifically, Theorem C of (Ramsey 1930)), as applied to the case of binary relations, we can extract from the sequence of the $G_{i'_1} \restriction \theta$'s a subsequence $G_{i''_1}, G_{i''_2}, G_{i''_3}, \ldots$ so that every $G_{i''_j} \setminus \{1, \ldots, n\}$ contains either a clique, or

an independent set, of size greater than or equal to j. For $j \geqslant 1$, one now easily gets Γ_j as a subgraph of $G_{i''_j} \setminus \{1, \ldots, n\}$. $\square$

The above result can be seen as a recasting of Ramsey's result on the existence of (arbitrarily) large homogeneous sets, to the case of graphs with n distinguished nodes. These special nodes always produce the same "scenario" when combined with an additional node: the type θ, intuitively to be chosen in accordance with the input formula.

Theorem 3.1. To each BSR sentence Φ in $\mathcal{L}_\varrho$ there corresponds a BSR sentence $\Phi^{\boldsymbol{\sigma}}$ in $\mathcal{L}_\in$, of size $|\Phi^{\boldsymbol{\sigma}}| = \mathcal{O}\left(|\Phi|^2\right)$, such that Φ is satisfiable by arbitrarily large models if and only if $\Phi^{\boldsymbol{\sigma}}$ is satisfiable in well-founded Set Theory.

Proof. Consider a BSR sentence

$$\Phi \quad \equiv \quad \exists x_1, \ldots, x_n \forall y_1, \ldots, y_m \; \varphi(x_1, \ldots, x_n, y_1, \ldots, y_m)\,,$$

where φ is an unquantified matrix in the language $\mathcal{L}_\varrho$.

To prove our claim, let us assume that Φ is satisfied by arbitrarily large finite structures, that we can represent as a sequence $G_i = \langle D_i, R_i \rangle$, with $i \in \mathbb{N}$, of directed graphs: each G_i has n distinguished nodes $v_1^i, \ldots, v_n^i \in D_i$ used to interpret $x_1, \ldots, x_n$. We can assume that the D_i's have strictly increasing cardinalities (if not, we can achieve this by sieving out a subsequence of the G_i's before moving on).

We will amalgamate all G_i's together inside $\boldsymbol{d}_0 \cup \boldsymbol{d}_1$, where $\boldsymbol{d}_0, \boldsymbol{d}_1$ are infinite sets satisfying the formula $\iota\iota(d_0, d_1)$ seen in the Introduction (cf. Fig. 1).

The embedding of each G_i in $\boldsymbol{d}_0 \cup \boldsymbol{d}_1$ is a modification of the one employed in Theorem 2.1 and can be described as follows: for every node v_k^i we introduce a set $x_{\mathrm{s},k} \in \boldsymbol{d}_0$, acting as its representative when v_k^i is considered as source; moreover, we introduce n nodes $x_{\mathrm{t},k,1}, \ldots, x_{\mathrm{t},k,n} \in \boldsymbol{d}_1$ acting as potential targets (for $x_{\mathrm{s},1}, \ldots, x_{\mathrm{s},n}$, respectively) when v_k^i is playing the role of target. The matrix $\varphi^{\boldsymbol{\sigma}}$ will be designed so as to impose the constraints needed to tie $\in$ with R_i, while $\iota\iota(\boldsymbol{d}_0, \boldsymbol{d}_1)$ will ensure that sufficiently many targets—respecting the corresponding membership conditions—can always be found in $\boldsymbol{d}_1$.

All the sub-formulae to be used must be intended (and verified) to be shortcuts for set-theoretic pure BSR-formulae.

For any $i \in \mathbb{N}$, consider now a generic element $w^i \in D_i \setminus \{v_1^i, \ldots, v_n^i\}$ and consider the subgraph $G_i(w^i, v_1^i, \ldots, v_n^i)$ of G_i induced by $w^i, v_1^i, \ldots, v_n^i$. For any $i, j \in \mathbb{N}$ we say that $G_i(w^i, v_1^i, \ldots, v_n^i)$ is *isomorphic* to $G_j(w^j, v_1^j, \ldots, v_n^j)$ if the correspondence sending $v_1^i, \ldots, v_n^i$ to $v_1^j, \ldots, v_n^j$ and w^i to w^j, respectively, is an isomorphism with respect to the arc relation. In formulae:

$$G_i(w^i, v_1^i, \ldots, v_n^i) \cong G_j(w^j, v_1^j, \ldots, v_n^j).$$

We will assume that the sequence of the G_i's enjoys the following three properties: for all $i, j \in \mathbb{N}$,

i) given any $w^i \in D_i \setminus \{v_1^i, \ldots, v_n^i\}$ and any $w^j \in D_j \setminus \{v_1^j, \ldots, v_n^j\}$,

$$G_i(w^i, v_1^i, \ldots, v_n^i) \cong G_j(w^j, v_1^j, \ldots, v_n^j);$$

ii) if $i < j$, then $|D_i| < |D_j|$;
iii) $D_i \setminus \{v_1^i, \ldots, v_n^i\}$ is either an independent set in G_i, or a clique in G_i: i.e., between any two elements of $D_i \setminus \{v_1^i, \ldots, v_n^i\}$ either G_i has no arc or it has arcs in both directions.

Should these conditions not be met, Lemma 3.1 tells us how we can enforce them by replacing the G_i's by a suitably related sequence of Γ_i's.

We are now in the position to define $\Phi^{\boldsymbol{\sigma}}$ and to prove our main claim. Let $\Phi^{\boldsymbol{\sigma}}$ be the following formula:

$$(\exists d_0, d_1)(\exists x_{\mathrm{s},1}, \ldots, x_{\mathrm{s},n+1} \in d_0)(\exists X_{\mathrm{t},1}, \ldots, X_{\mathrm{t},n+1} \subseteq d_0 \cup d_1)(\exists Y_{\mathrm{s}})(\exists \ell \in d_1)\Big(\iota\iota(d_0, d_1)$$
$$\wedge \bigwedge_{k=1}^{n+1} (X_{\mathrm{t},k} = \{x_{\mathrm{t},k,1}, \ldots, x_{\mathrm{t},k,n+1}\} \wedge X_{\mathrm{t},k} \cap d_1 \subseteq x_{\mathrm{s},k}) \wedge Y_{\mathrm{s}} = (d_0 \setminus \ell) \cup \{x_{\mathrm{s},1}, \ldots, x_{\mathrm{s},n+1}\}$$
$$\wedge (\forall y_{\mathrm{s},1}, \ldots, y_{\mathrm{s},m} \in Y_{\mathrm{s}})\varphi^{\boldsymbol{\sigma}}(x_{\mathrm{s},1}, \ldots, x_{\mathrm{s},n+1}, x_{\mathrm{t},1,1}, \ldots, x_{\mathrm{t},n+1,n+1}, y_{\mathrm{s},1}, \ldots, y_{\mathrm{s},m})\Big)$$

where $\varphi^{\boldsymbol{\sigma}}$ is obtained from φ by replacing every literal of the form $z_h \varrho w_j$, with $z, w \in \{x, y\}$, by $z_{\mathrm{s},h} \ni w_{\mathrm{t},h,j}$ with $w_{\mathrm{t},h,j} \equiv x_{\mathrm{t},h,n+1}$ when $w_j \equiv y_j$, and every literal of the form $z_h = w_j$ by $z_{\mathrm{s},h} = w_{\mathrm{s},j}$. It is plain that this can be formulated in $\mathcal{L}_\in$.

We begin by proving that if Φ is satisfiable by models of arbitrarily large cardinalities, then $\Phi^{\boldsymbol{\sigma}}$ is satisfiable in well-founded Set Theory.

Under our hypothesis, as observed above, we can assume we have a sequence of models G_i, for $i \in \mathbb{N}$ such that *i), ii),* and *iii)* hold.

We claim that we can determine:

a) $n+1$ elements $\alpha_1, \ldots, \alpha_{n+1} \in \boldsymbol{d}_0$,
b) $(n+1)^2$ elements $\beta_{k,1}, \ldots, \beta_{k,n+1} \in \boldsymbol{d}_0 \cup \boldsymbol{d}_1$, with $k = 1, \ldots, n+1$,

so that, for any G_i, we have:

1) there is an arc from v_j^i to v_k^i if and only if $\alpha_j \ni \beta_{k,j}$,
2) there is an arc from w^i to v_k^i if and only if $\alpha_{n+1} \ni \beta_{k,n+1}$,
3) there is an arc from v_j^i to w^i if and only if $\alpha_j \ni \beta_{j,n+1}$, and
4) there is an arc from w^i to w^i if and only if $\alpha_{n+1} \ni \beta_{n+1,n+1}$.

The α's satisfying a) are used to interpret $x_{\mathrm{s},1}, \ldots, x_{\mathrm{s},n+1}$, respectively, while the β's satisfying b) are used to interpret $x_{\mathrm{t},k,1}, \ldots, x_{\mathrm{t},k,n+1}$, respectively.

See Figure 4, where we depicted a scenario in which the various choices described in a) and b) have been made on *stripes* to be seen as associated with $\alpha_1, \ldots, \alpha_n$, followed by a stripe associated with $\alpha_{n+1} \in Y_{\mathrm{s}}$. The elements above α_{n+1} in $\boldsymbol{d}_0$ are meant to constitute, along with $v_1, \ldots, v_n$, the infinite interpretation of Y_{s}.

In Y_{s}, in fact, all the domains of the G_i's are "glued" together: a constraint reflecting the satisfiability by structures of arbitrarily large sizes.

To see that our claim holds, it is sufficient to recall that each of $\boldsymbol{d}_0$ and $\boldsymbol{d}_1$ has infinitely many elements and that for any pair of elements $a \in \boldsymbol{d}_0, b \in \boldsymbol{d}_1$, either $a \in b$ or $b \in a$ holds.

At this point we can complete our set-theoretic interpretation as follows:

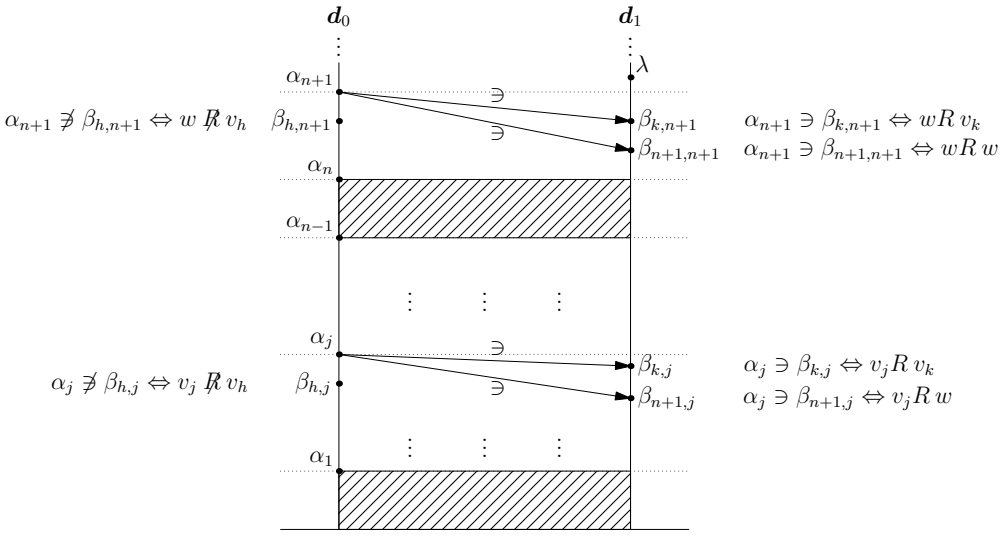

Fig. 4. A possible scenario for the choice of $\alpha_1, \ldots, \alpha_n$ (corresponding to $v_1, \ldots, v_n$ and hence to $x_1, \ldots, x_n$). This illustrates, among other things, the encodings of: presence of an arc between v_j and v_k and between v_j and w; absence of the arc between v_j and v_h. The elements have been chosen in *stripes*, the last stripe being associated with α_{n+1} (which corresponds to w and hence to a generic universal variable).

— interpret ℓ as the element $\lambda \in \boldsymbol{d}_1$ (hence a subset of $\boldsymbol{d}_0$) consisting of elements of rank smaller than the rank of α_{n+1},
— interpret Y_{s} as $(\boldsymbol{d}_0 \setminus \lambda) \cup \{\alpha_1, \ldots, \alpha_{n+1}\}$.

The fact that if $\Phi^{\boldsymbol{\sigma}}$ is satisfiable in well-founded Set Theory then Φ is satisfiable by models of arbitrarily large cardinalities, easily follows from the fact that under our hypothesis Φ is, in fact, satisfied by an infinite model. $\square$

Remark 3.1. In order to establish whether a BSR logical formula Φ admits models of arbitrarily large cardinalities, one can now either search for a model of size $\mathfrak{r}(\Phi)$ (i.e. the original bound established by Ramsey) or test $\Phi^{\boldsymbol{\sigma}}$ for set-theoretic satisfiability. The first method must explore a search space of size exponential in $\mathfrak{r}(\Phi)$, while the second must search for a set-theoretic model of size $\mathcal{O}\left(|\Phi|^2\right)$. Even though neither of the two is—in any practical sense—efficient, the second one is computationally more promising. By inspection of the formula one sees that testing $\Phi^{\boldsymbol{\sigma}}$ for set-theoretic satisfiability does not really require the elaborate machinery developed in (Omodeo and Policriti 2012). In fact, on the one hand the only infinite sets needed to satisfy $\iota\iota$ can be fixed beforehand as ω_0

and ω_1, in their (finite, see (Omodeo and Policriti 2010)) graph-theoretic representation. On the other hand, the search of the (finite) sets to be used to interpret the remaining variables of Φ^{σ}, can be carried out among a bounded collection of subsets and elements of $\omega_0 \cup \omega_1$. Moreover, this search can be performed in a bottom-up fashion, starting from most simple models.

Remark 3.2. As recalled in what precedes, cf. (Omodeo and Policriti 2010), a technique is known for eliminating equality from set-theoretic BSR sentences without leaving the BSR class. Hence, the translation $\Phi \mapsto \Phi^{\sigma}$ could be turned into one producing an equality-free result. However, the authors have never addressed the issue of whether this can be performed in a goal-driven fashion with a reasonable algorithmic cost.

Conclusions

The inception of research on decision algorithms for fragments of Set Theory, many years ago, was motivated by the expectation that such algorithms would play a significant role in the technology of proof assistants. Such expectation has concretized, to a significant extent, in a recent proof-checker: Ref (Schwartz *et al.* 2011; Omodeo and Tomescu 2014). Ref's core inferential mechanism implements, in fact, an enhanced variant of the *multi-level syllogistic* mentioned at the beginning of this paper. The said mechanism intervenes a few times, e.g., during Ref's validation of the two tiny proofs shown in Fig. 5: in either proof, it is used once to check that the statement which starts the argument by contradiction is *equivalent* to the instantiated negation of the claim; then, less shallowly, to establish a *conflict* between that statement and the definition of inj_{Θ}.

The set-theoretic BSR class does not seem easily amenable, in its entirety, to similar direct exploitations, but its decidability is beginning to reveal deep links with combinatorics.

The result discussed in this paper shows that the finite/co-finite spectrum of any given formula Ψ in the BSR class of pure logic, can be expressed with a set-theoretic formula in the same class whose size is proportional to $|\Psi|^2$. After a recasting of the combinatorics in set-theoretic terms, the result essentially exploits—in the proof of Theorem 3.1—only the combinatorial theorem (Ramsey 1930) for complete *graphs* with just *two* colors for arcs. As a matter of fact, this (apparent) simplification of the underlying combinatorics is a consequence of the initial assumption stating that we can deal with uninterpreted *binary* relations only. The remaining technical part of the argument preparing for the set-theoretic embedding—again in the proof of Theorem 3.1—, simply reduces to the use of the infinite case of the pigeonhole principle.

Notice that, for the above mentioned embedding, we did not give a result for non-well-founded Set Theory. We expect that an analogous result can easily be obtained, by exploiting basically the same construction coupled with an infinity-encoding formula, adapted to the non-well-founded case (e.g. the formula $\bar{\imath\imath}$ in (Omodeo *et al.* 2012), originally introduced in (Parlamento and Policriti 1988)—see also (Omodeo *et al.* 2009)). The true limitation, in the non-well-founded case, lies in the lack of a decidability result for the BRS class: an open problem that we rate as challenging.

THEORY $\mathsf{an_injection}(\mathsf{v}_0, \mathsf{d}_0)$
$\mathsf{v}_0 \not\subseteq \mathsf{d}_0$
END $\mathsf{an_injection}$

The following definition requires that inj_Θ send to $\emptyset$:
- every set lying inside d_0,
- an arbitrary but fixed element a of the set-difference $\mathsf{v}_0 \backslash \mathsf{d}_0$.

Moreover, inj_Θ shall send each $\mathsf{w} \in \mathsf{v}_0 \backslash \mathsf{d}_0 \backslash \{\mathsf{a}\}$ to $\{\{\mathsf{v}_0\} \cup (\mathsf{v}_0 \backslash \{\mathsf{w}\})\}$, and each set lying outside v_0 to $\{\mathsf{v}_0 \cup \{\mathsf{v}_0\}\}$.

DEF inj: $\mathsf{inj}_\Theta(\mathsf{W}) =_{\mathrm{Def}}$ **if** $\mathsf{W} \in \mathsf{d}_0 \cup \{\mathbf{arb}(\mathsf{v}_0 \backslash \mathsf{d}_0)\}$ **then** $\emptyset$ **else** $\{\{\mathsf{v}_0\} \cup (\mathsf{v}_0 \backslash \{\mathsf{W}\})\}$ **fi**

THEOREM $\mathsf{an_injection}_0$: [The restriction of inj_Θ to $\mathsf{v}_0 \setminus \mathsf{d}_0$ is 1–1]

$$\mathsf{X} \in \mathsf{v}_0 \backslash \mathsf{d}_0 \ \&\ \mathsf{Y} \notin \mathsf{d}_0 \ \&\ \mathsf{inj}_\Theta(\mathsf{X}) = \mathsf{inj}_\Theta(\mathsf{Y}) \rightarrow \mathsf{X} = \mathsf{Y}.$$

PROOF:
$\mathsf{Suppose_not}(\mathsf{x}_0, \mathsf{y}_0) \Rightarrow \quad \mathsf{x}_0 \in \mathsf{v}_0 \backslash \mathsf{d}_0 \backslash \{\mathsf{y}_0\} \ \&\ \mathsf{y}_0 \notin \mathsf{d}_0 \ \&\ \mathsf{inj}_\Theta(\mathsf{x}_0) = \mathsf{inj}_\Theta(\mathsf{y}_0)$
$\mathsf{Use_def}(\mathsf{inj}_\Theta) \Rightarrow \quad \mathsf{false}; \qquad \mathsf{Discharge} \Rightarrow \quad$ QED

THEOREM $\mathsf{an_injection}_1$: [No membership between inj_Θ images of operands outside d_0]

$$\{\mathsf{X}, \mathsf{Y}\} \cap \mathsf{d}_0 = \emptyset \ \&\ \mathsf{X} \in \mathsf{v}_0 \rightarrow \mathsf{inj}_\Theta(\mathsf{Y}) \notin \mathsf{inj}_\Theta(\mathsf{X}).$$

PROOF:
$\mathsf{Suppose_not}(\mathsf{x}_0, \mathsf{y}_0) \Rightarrow \quad \mathsf{x}_0 \in \mathsf{v}_0 \backslash \mathsf{d}_0 \ \&\ \mathsf{y}_0 \notin \mathsf{d}_0 \ \&\ \mathsf{inj}_\Theta(\mathsf{y}_0) \in \mathsf{inj}_\Theta(\mathsf{x}_0)$
$\mathsf{Use_def}(\mathsf{inj}_\Theta) \Rightarrow \quad \mathsf{false}; \qquad \mathsf{Discharge} \Rightarrow \quad$ QED

Fig. 5. Multi-level syllogistic invisibly at work in a Ref's proof scenario.

As a final consideration on decidability, we observe that the BSR class lies very close to the edge of undecidability. To make the BSR class undecidable, in fact, it would suffice to enhance the unquantified part of the language with the ability to state that a set has *exactly* two elements, cf. (Parlamento and Policriti 1988). On the other hand, within the BSR class treated in this paper it is easy to express the fact that a set is not empty and has *at most* two elements.

Acknowledgements

The authors would like to thank the anonymous referees for their very helpful suggestions.

References

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

M. Breban, A. Ferro, E.G. Omodeo, and J.T. Schwartz. Decision Procedures for Elementary Sublanguages of Set Theory. II: Formulas involving Restricted Quantifiers, together with Ordinal, Integer, Map, and Domain Notions. *Comm. Pure Appl. Math.*, 34:177–195, 1981.

D. Cantone. Decision procedures for elementary sublanguages of set theory. XVII. Commonly occurring decidable extensions of multi-level syllogistic. In Martin Davis and Ed Schonberg, editors, *From Linear Operators to Computational Biology: Essays in Memory of Jacob T. Schwartz*, pages 131–167. Springer, 2012.

D. Cantone, E. G. Omodeo, and A. Policriti. *Set Theory for Computing. From Decision Procedures to Declarative Programming with Sets.* Monographs in Computer Science. Springer-Verlag, 2001.

A. Ferro, E.G. Omodeo, and J.T. Schwartz. Decision procedures for elementary sublanguages of set theory. I. Multi-level syllogistic and some extensions. *Comm. Pure Applied Math.*, 33(5):599–608, 1980.

E. G. Omodeo and A. I. Tomescu. Set graphs. III. Proof Pearl: Claw-free graphs mirrored into transitive hereditarily finite sets. *J. Autom. Reason.*, 52(1):1–29, 2014. DOI: 10.1007/s10817-012-9272-3.

E.G. Omodeo, A.Policriti, and A.I. Tomescu. Infinity, in short. *J. Log. Comput.*, 22(6):1391–1403, 2012.

E.G. Omodeo, A.Policriti, and A.I. Tomescu. Statements of ill-founded infinity in Set Theory, In Patrick Cégielski ed., *Studies in Weak Arithmetics.* No. 196 in Lecture Notes, Center for the Study of Language and Information, Stanford University, pages 173–199, 2009.

E.G. Omodeo and A. Policriti. Solvable set/hyperset contexts: I. Some decision procedures for the pure, finite case. *Comm. Pure Appl. Math.*, 48(9-10):1123–1155, 1995. Special issue in honor of Jacob T. Schwartz.

E.G. Omodeo and A. Policriti. The Bernays-Schönfinkel-Ramsey class for Set Theory: Semidecidability. *J. Symb. Logic*, 75(2):459–480, 2010.

E.G. Omodeo and A. Policriti. The Bernays-Schönfinkel-Ramsey class for set theory: Decidability. *J. Symb. Logic*, 77(3):896–918, 2012.

F. Parlamento and A. Policriti. The Logically Simplest Form of the Infinity Axiom. *Proceedings of the American Mathematical Society*, 103(1):274–276, May 1988.

F. Parlamento and A. Policriti. Expressing innity without foundation. *The Journal of Symbolic Logic*, 56:1230–1235, 1991.

F. Parlamento and A. Policriti. Undecidability results for restricted universally quantified formulae of set theory. *Communications on Pure and Applied Mathematics*, 46(1):57–73, January 1993.

D. Plummer, and L. Lovász. Matching Theory. *North-Holland Mathematics Studies*, Elsevier Science, 1986.

S.P. Radziszowski. Small Ramsey Numbers. *The Electronic Journal of Combinatorics*, 2014. Submitted: June 11, 1994; Revision #14: January 12, 2014.

F.P. Ramsey. On a problem of formal logic. *Proceedings of the London Mathematical Society*, 30:264–286, 1930. Read on December 13, 1928.

J.T. Schwartz, D. Cantone, and E.G. Omodeo. *Computational Logic and Set Theory —Applying Formalized Logic to Analysis.* Springer, 2011.