Uniform One-Dimensional Fragments with One Equivalence Relation

Emanuel Kieroński[1] and Antti Kuusisto[2]

1   University of Wrocław, Poland, kiero@cs.uni.wroc.pl
2   Stockholm University, Sweden, antti.j.kuusisto@gmail.com

───── **Abstract** ─────

The uniform one-dimensional fragment U_1 of first-order logic was introduced recently as a natural generalization of the two-variable fragment FO^2 to contexts with relation symbols of all arities. It was shown that U_1 has the exponential model property and a NExpTime-complete satisfiability problem. In this paper we investigate two restrictions of U_1 that still contain FO^2. We call these logics RU_1 and SU_1, or the restricted and strongly restricted uniform one-dimensional fragments. We introduce Ehrenfeucht-Fraïssé games for the logics and prove that while SU_1 and RU_1 are expressively equivalent, they are strictly contained in U_1. Furthermore, we consider extensions of the logics SU_1, RU_1 and U_1 with unrestricted use of a single built-in equivalence relation $\sim$. We prove that while all the obtained systems retain the finite model property, their complexities differ. Namely, the satisfiability problem is NExpTime-complete for $SU_1(\sim)$ and 2-NExpTime-complete for both $RU_1(\sim)$ and $U_1(\sim)$. Finally, we show undecidability of some natural extensions of $SU_1(\sim)$.

1 Introduction

Two-variable logic FO^2 was proved decidable in [17], and the satisfiability and finite satisfiability problems of FO^2 were shown NExpTime-complete in [7]. The extension of two-variable logic with counting quantifiers, FOC^2, was proved decidable in [8], [18]. It was subsequently shown to be NExpTime-complete in [20]. Research on extensions and variants of two-variable logic is *currently very active*. Recent research efforts have mainly concerned decidability and complexity issues over restricted classes of structures, and also questions related to different built-in features and operators that increase the expressivity of the base language. Recent articles in the field include for example [3, 11, 21, 23], and several others.

Typical systems of modal logic are contained in two-variable logic, or some variant of it, and hence investigations on two-variable logics have direct implications on various fields of computer science, including verification of software and hardware, distributed systems, knowledge representation and artificial intelligence. However, two-variable logics do not cope well with relations of arities greater than two, and therefore the *scope of related research is significantly restricted*. In database theory contexts, for example, two-variable logics as such are often not directly applicable due to the *severe arity-related limitations*.

Uniform one-dimensional fragment U_1 of first-order logic is a recently introduced formalism that generalizes two-variable logic to contexts with relation symbols of all arities. The fragment was originally defined in [9] and studied further in [10]. The fragment is based on restricting

first-order logic in two ways. Firstly, quantification is restricted to blocks of existential (universal) quantifiers that *leave at most one free variable* in the resulting formula. Secondly, a *uniformity condition* applies to the use of atomic formulas: a Boolean combination of atoms $R(x_1, ..., x_k)$ and $S(y_1, ..., y_n)$, where $k, n \geq 2$, is allowed only if $\{x_1, ..., x_k\} = \{y_1, ..., y_n\}$. Boolean combinations of formulas with at most one free variable can be formed freely, and the use of equality is unrestricted.

It was established in [9] that if either of the two restrictions, one-dimensionality or uniformity, is lifted in a canonical way, the resulting formalism is undecidable. It was also established that already the equality-free fragment of U_1 can define properties not expressible in FOC^2 and also properties not expressible in the recently introduced *guarded negation fragment* [2], which significantly generalizes the *guarded fragment* [1] and *unary negation fragment* [14]. It was later established in [10] that U_1 has the finite model property and that the satisfiability problem of U_1 is NExpTime-complete. Thus the increase in expressivity when going from FO^2 to U_1 comes without cost in complexity. However, it was also proved in [10] that, in contrast to FO^2, adding counting quantifiers to U_1 leads to undecidability.

In this paper we investigate two restrictions of U_1 that still contain FO^2. We call these logics RU_1 and SU_1, or the *restricted* and *strongly restricted* uniform one-dimensional fragments. We begin our study by investigating the expressive power of the logics U_1, RU_1 and SU_1. We first provide Ehrenfeucht-Fraïssé game characterizations for our fragments; the rather simple and natural characterizations provide a nice algebraic perspective on the logics. We then establish that while SU_1 and RU_1 are expressively equivalent, they are strictly contained in U_1. Strictness of the containment follows by use of the EF-game for SU_1.

We then consider extensions of the logics SU_1, RU_1 and U_1 with a single built-in equivalence relation $\sim$ which can be used freely, i.e., the uniformity conditions do not apply to the use of $\sim$. We prove that while all the obtained systems retain the finite model property, their complexities differ. Namely, the satisfiability problem is NExpTime-complete for $SU_1(\sim)$ and 2-NExpTime-complete for both $RU_1(\sim)$ and $U_1(\sim)$. Thus we provide a complete classification of the complexities of the logics $SU_1(\sim)$, $RU_1(\sim)$ and $U_1(\sim)$.

We finish the investigations in this paper by establishing undecidability of some natural extensions of $SU_1(\sim)$. We show undecidability of the extension of SU_1 with *two* equivalence relations as well as the extension with one transitive relation. This contrasts with the case of FO^2 which remains decidable when extended by two equivalence relations [12, 13] or one transitive relation [23]. FO^2 with three equivalence relations is undecidable [12].

Built-in equivalence relations have played a visible role in recent investigations on two-variable logics, see for example [4, 5, 12, 13]. The articles [4, 5] discuss applications of two-variable logics with built-in equivalences in the context of *data words* and XML reasoning. In addition to being relevant in the context of data words, two-variable logics with equivalence relations naturally embed various different kinds of epistemic logics, where equivalence relations naturally correspond to epistemic indistinguishability relations of agents. Furthermore, the idea of adding equivalence relations in order to increase expressivity has been recently investigated in the context of interval temporal logics; see, e.g., [16].

Two-variable logics and guarded fragments are currently the two principal frameworks used for identifying decidable fragments of first-order logic. Originally, the logic U_1 was defined to be a generalization of FO^2, and in this respect U_1 is to FO^2 what the guarded negation fragment is to the guarded fragment—a reasonable generalization. U_1 has the same complexity as FO^2, and—as discussed above—its extension with counting quantifiers as well as its variants without either the uniformity or the one-dimensionality constraint, are undecidable. However, there are of course other decidable generalizations of FO^2, such

as FOC^2 and the novel logics RU$_1$ and SU$_1$. Hence it is important, we believe, to try to better understand the realm of decidable logics above FO^2. The investigations in this article contribute towards that aim. In particular, we observe, e.g., that the generalization SU$_1$ of FO^2 is of *lower complexity* than U$_1$ and RU$_1$ in the presence of a built-in equivalence.

2 Preliminaries

We let $\mathbb{Z}_+$ denote the set of positive integers and $\mathbb{N}$ the natural numbers. If $\overline{a}$ is a finite tuple of elements, we write $b \in \overline{a}$ in order to indicate that b is one of the elements of the tuple. By $(u, ..., u)$ we denote a finite tuple where each position contains the same element u; the arity of the tuple is unimportant or known from the context when this notation is used. We recall that $\bigwedge \emptyset = \top$ and $\bigvee \emptyset = \bot$. The order of priority of logical connectives when brackets are left unwritten is such that first come $\wedge, \vee$, and after that come $\rightarrow, \leftrightarrow$. The length of a formula φ is denoted by $\|\varphi\|$.

Let $\mathcal{V}$ denote a *complete relational vocabulary*, i.e., $\mathcal{V} := \bigcup_{k \in \mathbb{Z}_+} \tau_k$, where τ_k denotes a countably infinite set of k-ary relation symbols. Every vocabulary we consider below is assumed to be a subset of $\mathcal{V}$. In the sections concerning expressivity, we use the symbol σ in order to refer to *finite* vocabularies. In investigations concerning complexities of satisfiability problems, the vocabulary of the set of input formulas is always $\mathcal{V}$ extended with the special built-in symbols such as the equivalence relation symbol $\sim$. In this article a σ-model $\mathfrak{A}$ is a model that interprets at least the relation symbols in the vocabulary σ.

We let VAR $= \{ v_i \,|\, i \in \mathbb{N} \}$ be the set of first-order variables. We mostly use *meta-variables* x, y, z, x_1, x_2, x_3, etc., in order to denote variables in VAR. We let $diff(x_1, ..., x_m)$ denote the conjunction $\bigwedge_{1 \le i < j \le m} x_i \ne x_j$. We define $diff(x) := x = x$.

Let $X = \{x_1, ..., x_m\} \ne \emptyset$ be a finite set of variable symbols. Let R be a k-ary relation symbol. if $\{x_{i_1}, ..., x_{i_k}\} = X$. Equalities $x = y$ are *not* $\{x, y\}$-atoms, since the definition requires a relation symbol to be used. A formula is called an X-*literal* if it is an X-atom or a negated X-atom.

Let τ be a vocabulary. A k-ary τ-*atom* is an atomic τ-formula ψ such that

$$|\{ x \in \text{VAR} \mid \psi \text{ contains an instance of } x \}| = k.$$

For example, if $P \in \tau$ is a unary and $R \in \tau$ a ternary symbol, then $P(x)$, $x = x$, $R(x, x, x)$ are unary τ-atoms, and $R(v_1, v_2, v_2)$, $v_1 = v_2$ are binary τ-atoms.

The set of τ-formulas of the *uniform one-dimensional fragment* U$_1$ is the smallest set $\mathcal{F}$ satisfying the following conditions.
1. Every unary τ-atom is in $\mathcal{F}$. Also $\bot, \top \in \mathcal{F}$.
2. Every identity atom $x = y$ is in $\mathcal{F}$.
3. If $\varphi \in \mathcal{F}$, then $\neg\varphi \in \mathcal{F}$.
4. If $\varphi, \psi \in \mathcal{F}$, then $(\varphi \wedge \psi) \in \mathcal{F}$.
5. Let $Y := \{x_0, ..., x_k\} \subseteq \text{VAR}$ and $X \subseteq Y$. Let φ be a Boolean combination of X-atoms over τ and formulas in $\mathcal{F}$ whose free variables (if any) are in Y. Then
   a. $\exists x_1...\exists x_k \, \varphi \in \mathcal{F}$ and
   b. $\exists x_0...\exists x_k \, \varphi \in \mathcal{F}$.

For example $\exists y \exists z (\neg Sxyz \wedge Py \wedge (Syzx \vee Tzyxz))$ is a U$_1$-formula, while $\exists y \exists z (Rxy \wedge Ryz)$ is not. Now consider the U$_1$-formula $\exists y \exists z (x \ne y \wedge Ryz)$. The free variable x *does not* occur in the set $\{y, z\}$ that corresponds to the set X in clause 5 of the definition of U$_1$. Consider the clause 5.a which states that "$\exists x_1...\exists x_k \varphi \in \mathcal{F}$." Change the clause 5.$a$ to the novel clause "if $x_0 \in X$, then $\exists x_1...\exists x_k \varphi \in \mathcal{F}$." The five clauses with this modified version of clause 5.a

define the set of τ-sentences RU_1. Note that in clause 5, the formula φ does not have to contain any X-atoms, so formulas such as $\exists y \exists z \big(x \neq y \land x \neq z\big)$ are in U_1 and RU_1.

Consider the RU_1-formula $\exists y \exists z \big(Rxy \land y \neq z\big)$. The free variable z is not in the set $\{x, y\}$ which corresponds to the set X in clause 5 of the defintion of U_1. Consider the variant of the clause 5.a stating that "if $x_0 \in X$ and $X = \{x_0, ..., x_k\}$, then $\exists x_1...\exists x_k \varphi \in \mathcal{F}$." Consider also a variant of the rule 5.b which states that "if $X = \{x_0, ..., x_k\}$, then $\exists x_0...\exists x_k \varphi \in \mathcal{F}$." The five clauses with these modified versions of 5.a and 5.b define the set of τ-sentences of SU_1.

The above minor modifications to the syntax of U_1 that lead to RU_1 and SU_1 deal with *the way free and bound variables of formulas interact with relation symbols of higher arities*. The modifications lead to interesting complexity issues, as we will see. Our Ehrenfeucht-Fraïssé characterizations show that the (initially perhaps somewhat complicated) logics correspond to natural algebraic back and forth conditions that extend the well-known two-pebble games for FO^2. It is worth noting that clearly even the weakest of our logics, SU_1, contains FO^2.

We then define extensions of the three logics U_1, RU_1, SU_1 by a single built-in equivalence relation $\sim$. A formula φ is a τ-formula of $\text{U}_1(\sim)$ if and only if it can be obtained from some τ-formula of U_1 by replacing any number of equality symbols $=$ by the equivalence symbol $\sim$. The logics $\text{RU}_1(\sim)$ and $\text{SU}_1(\sim)$ are defined analogously from RU_1 and SU_1.

We define the *quantifier block rank* of a U_1-formula φ, or $qbr(\varphi)$, as follows.
1. $qbr(\varphi) = 0$ iff φ is quantifier-free.
2. $qbr(\varphi \land \psi) = max(qbr(\varphi), qbr(\psi))$; $qbr(\neg\varphi) = qbr(\varphi)$.
3. Assume $\varphi := \exists \overline{x}\, \chi$, where χ does *not* begin with $\exists$. Then $qbr(\varphi) = qbr(\chi) + 1$.

We define the *quantifier width* of a U_1-formula φ, or $qw(\varphi)$, as follows.
1. $qw(\varphi) = 1$ iff φ is atomic and has a free variable. $\top$ and $\bot$ have quantifier width 0.
2. $qw(\varphi \land \psi) = max(qw(\varphi), qw(\psi))$; $qw(\neg\varphi) = qw(\varphi)$.
3. Assume $\varphi := \exists x_1...\exists x_k \chi$, where χ does *not* begin with $\exists$. If φ has a free variable, then $qw(\varphi) = max(1 + k, qw(\chi))$. If φ is a sentence, then $qw(\varphi) = max(k, qw(\chi))$.

The pair $(qbr(\varphi), qw(\varphi))$ is the *rank* of a U_1-formula. Note that SU_1 and RU_1 are fragments of U_1, so various technical definitions, such as the above definition of a notion of rank, automatically concern SU_1 and RU_1 as well.

Let $\overline{x}$ denote a tuple of variables. Let $\chi := \exists \overline{x}\, \varphi$ be a U_1-formula formed by using the formula construction rule 5. Assume φ is quantifier-free. Then we call φ a U_1-*matrix*. If φ does not contain k-ary atoms for any $k \geq 2$, with the possible exception of equality atoms $x = y$, then we define $S_\varphi := \emptyset$. Otherwise we define S_φ to be the set X used in the construction of χ (see rule 5). The set S_φ is the set of *live variables* of φ. Let $\psi(x_0, ..., x_k)$ be a U_1-matrix, where $(x_0, ..., x_k)$ enumerates the variables of ψ. Let $\mathfrak{A}$ be a structure and $a_0, \ldots, a_k \in A$. We let $\text{live}\big(\psi(x_0, ..., x_k)[a_0, \ldots, a_k]\big)$ denote the smallest set $T \subseteq \{a_0, ..., a_k\}$ such that $a_i \in T$ if x_i is a live variable of $\psi(x_0, ..., x_k)$. We may write $\text{live}\big(\psi[a_0, ..., a_k]\big)$ instead of $\text{live}\big(\psi(x_0, ..., x_k)[a_0, \ldots, a_k]\big)$ when no confusion can arise. Notice that since the elements a_i are not required to be distinct, it is possible that $|\text{live}(\psi[a_0, ..., a_k])|$ is smaller than the number of live variables in ψ.

2.1   Normal form and types

We introduce a normal form for our uniform one-dimensional logics which is inspired by the Scott normal form for FO^2 [22]. We say that a $\text{U}_1(\sim)$ ($\text{SU}_1(\sim)$, $\text{RU}_1(\sim)$) formula φ is in *generalized Scott normal form* if φ has the following shape

$$\bigwedge_{1 \leq i \leq m_\exists} \forall x \exists y_1 \ldots y_{k_i} \varphi_i^\exists \land \bigwedge_{1 \leq i \leq m_\forall} \forall x_1 \ldots x_{l_i} \varphi_i^\forall, \tag{1}$$

where $\varphi_i^\exists = \varphi_i^\exists(x, y_1, \ldots, y_{k_i})$ and $\varphi^\forall = \varphi_i^\forall(x_1, \ldots, x_{l_i})$ are quantifier-free. The following proposition is a natural generalisation of Proposition 1 in [10]. It can be proved in the standard fashion, see, e.g., [6].

▶ **Proposition 1.** *For every* $U_1(\sim)$ *(*$SU_1(\sim)$*,* $RU_1(\sim)$*) formula* φ*, one can compute in polynomial time a* $U_1(\sim)$ *(*$SU_1(\sim)$*,* $RU_1(\sim)$*) formula* φ' *in generalized Scott normal form (over a signature extended by some fresh unary symbols) such that* φ *and* φ' *are satisfiable over the same domains. Any model of* φ *can be expanded to a model of* φ' *by defining new unary symbols. Any model of* φ' *restricted to the signature of* φ *is a model of* φ*.*

Let φ be a $U_1(\sim)$ formula in generalized normal form, let $\mathfrak{A} \models \varphi$, $a \in A$, and let $b_1, \ldots, b_{k_i}$ be such that $\mathfrak{A} \models \varphi_i^\exists[a, b_1, \ldots, b_{k_i}]$. We say that $\mathfrak{B} = \mathfrak{A}{\restriction}\{a, b_1, \ldots, b_{k_i}\}$ (i.e., the restriction of $\mathfrak{A}$ to $\{a, b_1, \ldots, b_{k_i}\}$) is a *witness structure* for a and $\varphi_i^\exists$. The substructure of $\mathfrak{B}$ restricted to the elements of $\mathrm{live}(\varphi_i^\exists[a, b_1, \ldots, b_{k_i}])$ is called the *live part* of $\mathfrak{B}$. If the live part of $\mathfrak{B}$ does not contain a, then it is called *free*. Note that $|B|$ may be smaller than $k_i + 1$. Also, a may be a member of the live part of $\mathfrak{B}$ even if the variable x is not live in $\varphi_i^\exists$.

Let σ be a finite vocabulary. Let $\mathfrak{B}$ be a σ-model. Let $k \geq 1$ be an integer and $\bar{b} = (b_1, ..., b_k) \in B^k$ a tuple of *distinct* elements of $\mathfrak{B}$. Let $X = \{x_1, ..., x_k\}$ be a set of k *distinct* variables. Let T be the set of exactly all X-literals $\varphi(x_1, ..., x_k)$ over σ such that $\mathfrak{B} \models \varphi(b_1, ..., b_k)$. The conjunction $\bigwedge T$ is the *diagram type* of $\mathfrak{B}, \bar{b}$ over σ and with respect to the tuple $(x_1, ..., x_k)$. We denote this formula by $\delta_\sigma^{\mathfrak{B}, \bar{b}}(x_1, ..., x_k)$. We assume some standard syntactic form (ordering of conjuncts and bracketing), so that if two formulas $\delta_\sigma^{\mathfrak{A}, \bar{a}}(x_1, ..., x_k)$ and $\delta_\sigma^{\mathfrak{B}, \bar{b}}(x_1, ..., x_k)$ are equivalent, they are one and the same formula.

Let $\tau \subseteq \mathcal{V}$ be a finite vocabulary. A *1-type* over τ is a maximal satisfiable set of literals (atoms and negated atoms) over τ in the variable v_1. The set of all 1-types over τ is denoted by $\boldsymbol{\alpha}[\tau]$, or just by $\boldsymbol{\alpha}$ when τ is clear.

We identify 1-types α and conjunctions $\bigwedge \alpha$. A *k-table* over τ is a maximal satisfiable set of $\{v_1, ..., v_k\}$-atoms and negated $\{v_1, ..., v_k\}$-atoms over τ. Recall that a $\{v_1, ..., v_k\}$-atom must contain exactly the variables in $\{v_1, ..., v_k\}$, and note that a 2-table contains neither equality formulas nor negated equality formulas. We identify k-tables β and conjunctions $\bigwedge \beta$. We note that k-tables and diagram types are closely related notions.

Let $\mathfrak{A}$ be a τ-structure, and let $a \in A$. Let α be a 1-type over τ. We say that a *realizes* α if α is the unique 1-type such that $\mathfrak{A} \models \alpha[a]$. We let $\mathrm{tp}_{\mathfrak{A}}(a)$ denote the 1-type realized by a. Similarly, for *distinct* elements $a_1, \ldots, a_k \in A$, we let $\mathrm{tb}_{\mathfrak{A}}(a_1, \ldots, a_k)$ denote the unique k-table *realized* by the tuple $(a_1, \ldots, a_k)$, i.e., the k-table $\beta(v_1, ..., v_k)$ such that $\mathfrak{A} \models \beta[a_1, \ldots, a_k]$. Note that we have $\mathrm{tp}_{\mathfrak{A}}(a) \equiv \mathrm{tb}_{\mathfrak{A}}(a)$ for every $a \in A$.

Let us further introduce some new helpful terminology. A *multitype* is a function $\boldsymbol{\alpha} \to \mathbb{N}$. We say that a multitype θ is a *k-multitype* if $\sum_{\alpha \in \boldsymbol{\alpha}} \theta(\alpha) = k$. For a given set $\{a_1, \ldots, a_k\}$ of distinct elements from a structure $\mathfrak{A}$, we say that they *realize* a k-multitype θ, if for each $\alpha \in \boldsymbol{\alpha}$, we have that $\theta(\alpha)$ is the number of elements in $\{a_1, \ldots, a_k\}$ of 1-type α. If $\mathfrak{A}$ interprets an equivalence relation $\sim$, then we say that a multitype is realized *in a class D* if it is realized by a subset of elements of the equivalence class D of $\mathfrak{A}$. We say that a multitype is realized *by a class* if it is realized by the set of all elements of this equivalence class.

3 Games for U_1 and SU_1

In this section we provide Ehrenfeucht-Fraïssé game characterizations for U_1 and SU_1. A similar characterization exists for RU_1, but we will not discuss it explicitly.

We will below define the games rigorously, but *roughly*, the game for U_1 involves positions encoded by a bijection between finite subsets of two models. *Spoiler* chooses (at most) one

pair (u, u') of bijectively related points. Then he chooses a finite *blue* set B and a *green* set $G \subseteq B$ from one of the models such that we have $u \in B$ or $u' \in B$, depending on which model B was chosen from. *Duplicator* responds by a *new* bijection from B onto a subset of the other model. Intuitively, the bijection defines counterparts of the blue and green sets in the other model. Information about the relations of the two models in restriction to the sets B, G and their counterparts in the other model, are then compared (in a way specified later).

Let σ be a finite vocabulary. Let $\mathfrak{A}$ and $\mathfrak{D}$ be σ-models. Let $k \in \mathbb{N}$ and $n \in \mathbb{Z}_+$. Let $S \subseteq A$ and $T \subseteq D$ be *finite* (possibly empty) sets such that $|S| = |T|$. Let $f : S \to T$ be a bijection. We next define the game $G_\sigma^{k,n}(\mathfrak{A}, S, f, \mathfrak{D}, T)$ that characterizes expressivity of U_1-formulas of rank (k, n) and over σ.

The game is played between two players, Spoiler and Duplicator. The game begins from the *position* $(\mathfrak{A}, S_k, f_k, \mathfrak{D}, T_k)$, where $S_k = S$, $T_k = T$ and $f_k = f$. If $k = 0$, the (play of the) game ends immediately in the beginning position $(\mathfrak{A}, S, f, \mathfrak{D}, T)$. If $k \neq 0$, the game is played for k *rounds*; the game begins with round k, and each round $j \neq 0$ is followed by round $j - 1$. Round $j \in \{1, ..., k\}$ begins from a position denoted by $(\mathfrak{A}, S_j, f_j, \mathfrak{D}, T_j)$ and ends in a position $(\mathfrak{A}, S_{j-1}, f_{j-1}, \mathfrak{D}, T_{j-1})$, and the game ends in a position $(\mathfrak{A}, S_0, f_0, \mathfrak{D}, T_0)$; for each $j \in \{0, ..., k\}$, we have $S_j \subseteq A$ and $T_j \subseteq D$, while f_j is a bijection from S_j onto T_j. Round $j \in \{1, ..., k\}$ consists of a move by Spoiler and a response by Duplicator. These actions determine how the positions of the game arise and evolve.

In round j, Spoiler first decides whether he wants to make a *local* or a *global* move. Assume first that he decides upon a local move. Local moves are allowed only when S_j and T_j are nonempty. Spoiler chooses one of the pairs $\mathfrak{A}, S_j$ and $\mathfrak{D}, T_j$. Let us assume he chooses $\mathfrak{A}, S_j$. Spoiler then chooses an element $r \in S_j$ and sets $B \subseteq A$ and $G \subseteq B$ such that $|B| \leq n$ and $r \in B$. We call r the *red* element coloured by Spoiler in round j, and we call B and G the sets of *blue* and *green* elements coloured by Spoiler in round j. (Note that green elements are blue as well, and the red element r must be blue and can be green.) Once Spoiler has appointed the element r and the sets G and B, Duplicator chooses an injection $h : B \to D$ such that $h(r) = f_j(r)$. The game continues from the position $(\mathfrak{A}, S_{j-1}, f_{j-1}, \mathfrak{D}, T_{j-1}) := (\mathfrak{A}, B, h, \mathfrak{D}, h(B))$. (We define $h(B) = \{ h(b) \mid b \in B \}$.)

If Spoiler chooses the pair $\mathfrak{D}, T_j$ instead of $\mathfrak{A}, S_j$, the rules of the game are symmetric; Spoiler chooses a red element $r' \in T_j$ and blue and green sets $B' \subseteq D$ and $G' \subseteq B$ such that $|B'| \leq n$ and $r' \in B'$. Duplicator responds by an injection $h : B' \to A$ such that $h(r') = f_j^{-1}(r')$, where f_j^{-1} denotes the inverse function of the bijection f_j. The inverse function of the injection h is the novel bijection f_{j-1} from the novel set $S_{j-1} := h(B')$ onto the blue set B'. Of course $T_{j-1} := B'$.

If Spoiler decides upon a *global* move instead of a local one, he first chooses one of the structures $\mathfrak{A}$ and $\mathfrak{D}$. Let us assume that he chooses $\mathfrak{A}$. Spoiler then chooses a blue set $B \subseteq A$ and a green set $G \subseteq B$ such that $|B| \leq n$. Duplicator responds by an injection $h : B \to D$. The game continues from the position $(\mathfrak{A}, S_{j-1}, f_{j-1}, \mathfrak{D}, T_{j-1}) := (\mathfrak{A}, B, h, \mathfrak{D}, h(B))$. Again if Spoiler chooses the structure $\mathfrak{D}$ instead of $\mathfrak{A}$, he chooses the blue and green sets B' and G' from $\mathfrak{D}$. Duplicator then responds by an injection h from B' into A. The inverse function of h becomes the bijection f_{j-1}. Of course $|B'| \leq n$ and $G' \subseteq B'$.

We then describe the winning conditions of the game. We begin with some auxiliary definitions. Let X be a set, and let $l \in \mathbb{Z}_+$. Let $(u_1, ..., u_l) \in X^l$ be a tuple and $Y \subseteq X$. We say that $(u_1, ..., u_l)$ *spans* the set Y if $\{u_1, ..., u_l\} = Y$. Note that it is possible that $(u_1, ..., u_l)$ spans Y even if $|Y| < l$.

Let $\mathfrak{A}$ and $\mathfrak{D}$ be σ-structures. Let $G \subseteq A$ and $G' \subseteq D$ be finite sets. Let f be a bijection from G onto G'. We say that f *preserves spanning tuples* over σ and write $\mathfrak{A}, G \langle f, \sigma \rangle \mathfrak{D}, G'$,

if for each symbol $R \in \sigma$ and each tuple $\overline{a}$ that spans G, we have $\overline{a} \in R^{\mathfrak{A}} \Leftrightarrow f(\overline{a}) \in R^{\mathfrak{D}}$. (We define $f(\overline{a}) = (f(a_1), ..., f(a_p))$, where $\overline{a} = (a_1, ..., a_p)$.)

Duplicator wins a play of the game $G_\sigma^{k,n}(\mathfrak{A}, S, f, \mathfrak{D}, T)$ iff the conditions below hold.

1. Consider round $j \in \{1, ..., k\}$ of the game. If Spoiler makes his moves in $\mathfrak{A}$, then let $G \subseteq A$ be the green set coloured by Spoiler in round j. If Spoiler makes his moves in $\mathfrak{D}$, let G be the set $h(G')$, where $G' \subseteq D$ is the green set coloured by Spoiler in round j and h is the injection chosen by Duplicator. The restriction of f_{j-1} to G preserves spanning tuples over σ, i.e., $\mathfrak{A}, G \langle f_{j-1} \restriction G, \sigma \rangle \mathfrak{D}, f_{j-1}(G)$.

2. Recall that $(a, ..., a)$ denotes a tuple where each coordinate position contains a. Let $j \in \{0, ..., k\}$. For all $R \in \sigma$ and all $a \in S_j$, we have $(a, ..., a) \in R^{\mathfrak{A}} \Leftrightarrow (f_j(a), ..., f_j(a)) \in R^{\mathfrak{D}}$. In particular, $a \in P^{\mathfrak{A}} \Leftrightarrow f_j(a) \in P^{\mathfrak{D}}$ for each unary symbol $P \in \sigma$ and each $a \in S_j$.

We write $\mathfrak{A}, S \sim_{f,\sigma}^{k,n} \mathfrak{D}, T$ if *Duplicator* has a winning strategy in the game. A strategy of Duplicator is simply a function that takes as an argument a position in the game together with a move of Spoiler in that position; the value of the function with such an input is a specification of the response move of Duplicator. A strategy is a winning strategy if it guarantees a win in every play of the game.

Now consider a variant $\hat{G}_\sigma^{k,n}(\mathfrak{A}, S, f, \mathfrak{D}, T)$ of the game $G_\sigma^{k,n}(\mathfrak{A}, S, f, \mathfrak{D}, T)$ defined by adding to the game $G_\sigma^{k,n}(\mathfrak{A}, S, f, \mathfrak{D}, T)$ the additional rule that the green set chosen by Spoiler must always be either empty or equal to the blue set. In other words, if B and G are the blue and green sets chosen in some round of the game, then we have $G \in \{\emptyset, B\}$. Let $\hat{\sim}_{f,\sigma}^{k,n}$ denote the relation analogous to $\sim_{f,\sigma}^{k,n}$ but concerning the new variant of the game.

Let $\mathfrak{A}$ and $\mathfrak{D}$ be σ-models. Let $S \subseteq A$ and $T \subseteq D$ be equicardinal finite sets, and let $f : S \to T$ be a bijection. We write $\mathfrak{A}, S \equiv_{f,\sigma}^{k,n} \mathfrak{B}, T$ if the equivalence $\mathfrak{A} \models \varphi(\overline{a}) \Leftrightarrow \mathfrak{D} \models \varphi(f(\overline{a}))$ holds for all tuples $\overline{a}$ of elements of S and all U_1-formulas $\varphi(\overline{x})$ of rank (k, n) over σ. We let $\hat{\equiv}_{f,\sigma}^{k,n}$ denote the relation analogous to $\equiv_{f,\sigma}^{k,n}$ but concerning SU_1-formulas instead of U_1-formulas. When σ is clear or irrelevant, we may leave it unwritten.

The following theorem is relatively easy but tedious to prove. A detailed proof will be presented in the full version of the paper.

▶ **Theorem 2.** $\mathfrak{A}, S \equiv_f^{k,n} \mathfrak{D}, T \Leftrightarrow \mathfrak{A}, S \sim_f^{k,n} \mathfrak{D}, T$ *and* $\mathfrak{A}, S \hat{\equiv}_f^{k,n} \mathfrak{D}, T \Leftrightarrow \mathfrak{A}, S \hat{\sim}_f^{k,n} \mathfrak{D}, T$.

4 Comparing the expressive power

In this section we first establish that SU_1 is strictly less expressive than U_1.

Let R be a ternary relation. The U_1-sentence

$$\exists v \forall x \forall y \forall z \big(Rxyz \to (v = x \vee v = y \vee v = z) \big)$$

states that some v belongs to every tuple of R. Let us call this the *covering node property*.

▶ **Theorem 3.** *The covering node property is not expressible in* SU_1.

Proof. We begin by defining two models $\mathfrak{M}$ and $\mathfrak{N}$ with a *ternary* relation R. Intuitively, both of these models represent a hypergraph where each edge has exactly three elements. We define the model $\mathfrak{M} = (M, R^{\mathfrak{M}})$ such that $M = \{0, 1, 2, 3, 4, 5, 6\}$ and for each $(u, v, w) \in M^3$, we have $(u, v, w) \in R^{\mathfrak{M}}$ iff $\{u, v, w\} \in \{ \{0, 1, 2\}, \{0, 3, 4\}, \{0, 5, 6\} \}$. We define the model $\mathfrak{N} = (N, R^{\mathfrak{N}})$ such that $N = \{a, b, c, d, e, f, g\}$, and for each $(u, v, w) \in N^3$, we have $(u, v, w) \in R^{\mathfrak{N}}$ iff $\{u, v, w\} \in \{ \{a, b, c\}, \{c, d, e\}, \{e, f, g\} \}$. We note that while $\mathfrak{M}$ satisfies the covering node property, $\mathfrak{N}$ does not.

We then fix some terminology for later use. Let $\mathfrak{A} = (A, R^{\mathfrak{A}})$ be a model that represents a hypergraph where each edge has exactly three elements, meaning that for each $(u, v, w) \in A^3$, if $(u, v, w) \in R^{\mathfrak{A}}$, then every permutation of the tuple (u, v, w) is also in $R^{\mathfrak{A}}$ and $|\{u, v, w\}| = 3$. Let $t \in A$. We say that t is *incident to an edge* if we have $(t, t', t'') \in R^{\mathfrak{A}}$ for some elements $t', t'' \in A$. We say that t is *incident to a gap* if there exist elements $t', t'' \in A$ such that $(t, t', t'') \notin R^{\mathfrak{A}}$ and $|\{t, t', t''\}| = 3$. A subset S of A is an *edge* iff $S = \{s, s', s''\}$ for some elements s, s', s'' such that $(s, s', s'') \in R^{\mathfrak{A}}$.

Fix an arbitrary pair (k, n); we will show that $\mathfrak{M}, \emptyset \triangleq_{\emptyset}^{k,n} \mathfrak{N}, \emptyset$ by using the game for SU_1.

Assume first a position $(\mathfrak{M}, S, f, \mathfrak{N}, T)$ has been reached in the game. We show how Duplicator plays in that position.

Assume Spoiler chooses a blue set B and a green set $G \in \{B, \emptyset\}$. If Spoiler is making a local move, he also chooses a red element r which is in $S \cap B$ if Spoiler moves in $\mathfrak{A}$ and in $T \cap B$ otherwise. If $|B| \neq 3$, Duplicator responds by choosing an *arbitrary* injection h that maps the elements of B into the other model, and if Spoiler has made a local move, then also $h(r) = f(r)$ or $h(r) = f^{-1}(r)$ holds. Duplicator can always do this since $|M| = |N|$.

If $|B| = 3$, then the move of Duplicator depends on whether B is an edge. We assume that $G = B$. (In the case where $G = \emptyset$, Duplicator acts precisely the same way as in the case $G = B$.) Duplicator must choose an injection h' such that $h'(B)$ is an edge iff also B is an edge. Furthermore, if Spoiler has made a local move and thus appointed a red element r', which is necessarily in $S \cap B$ or in $T \cap B$ (depending on which model Spoiler moves in), the injection h' must map r' to $f(r')$ or $f^{-1}(r')$. Duplicator can always choose such an injection h' since *in both models, each element is incident to an edge as well as a gap*. ◀

On the other hand, it turns out that RU_1 and SU_1 have the same expressive power.

▶ **Theorem 4.** RU_1 *and* SU_1 *are expressively equivalent.*

Proof. Let σ be the vocabulary of a formula to be translated. It is easy to show that σ-formulas of RU_1 can be represented in a normal form where each formula $\exists \overline{x} \, \varphi$ is such that φ has the following shape

$$\delta(x_1, ..., x_q) \wedge \mathit{diff}(x_1, ..., x_r) \wedge \bigwedge_{i \in \{1, ..., r\}} \tau_i(x_i),$$

where $\delta(x_1, ..., x_q)$ is a diagram type, $q \leq r$, and the formulas $\tau_i(x_i)$ are so-called *types of rank (k, n)*. Types of rank (k, n) have the property that for each i and $j \neq i$, either $\tau_i(y)$ and $\tau_j(y)$ are equivalent, or the conjunction $\tau_i(y) \wedge \tau_j(y)$ is not satisfiable. Types of rank (k, n) have various analogous incarnations in various different contexts of finite model theory; see for example [15] for *rank-k types* for FO and also similar types for finite variable logics.

We define a translation t from such normal form formulas into SU_1 such that $t(\varphi) = \varphi$ for atoms and $t(\varphi \wedge \psi) = t(\varphi) \wedge t(\psi)$ and $t(\neg \varphi) = \neg t(\varphi)$. Formulas $\exists \overline{x} \, \varphi$ are trickier to translate. Let $\chi(x_1) := \exists x_2 ... \exists x_r \psi$, where ψ is the formula $\delta(x_1, ..., x_q) \wedge \mathit{diff}(x_1, ..., x_r) \wedge \bigwedge_{i \in \{1, ..., r\}} \tau_i(x_i)$. Let us translate $\chi(x_1)$ into SU_1. Define $\chi'(x_1)$ to be the formula

$$\exists x_2 ... \exists x_q \big(\delta(x_1, ..., x_q) \ \wedge \ \mathit{diff}(x_1, ..., x_q) \ \wedge \bigwedge_{i \in \{1, ..., q\}} \tau_i(x_i) \big)$$
$$\wedge \ \exists x_2 ... \exists x_r \big(\mathit{diff}(x_1, ..., x_r) \wedge \bigwedge_{i \in \{1, ..., r\}} \tau_i(x_i) \big).$$

(Notice carefully how the indices q and r are now placed.) Recalling the properties of the formulas $\tau_i(x_i)$ discussed above, it is easy to see that $\chi'(x_1)$ is equivalent to $\chi(x_1)$. The

translation $t(\chi(x_1))$ is obtained from $\chi'(x_1)$ by replacing the conjuncts $\tau_i(x_i)$ in the above conjunctions by $t(\tau_i(x_i))$.

Now consider a formula $\gamma := \exists \overline{x}\, \psi$ *without* free variables. Remove one variable y from $\overline{x}$ and let $\eta(y)$ denote the obtained formula; the variable y is assumed to be part of the diagram type in ψ. Now obtain from $\eta(y)$ the formula $t(\eta(y))$ in the way $t(\chi(x_1))$ was obtained from $\chi(x_1)$ above. We define $t(\gamma) := \exists y\, t(\eta(y))$. ◀

5 Built-in equivalence relations and complexity

It is not difficult to show (e.g., using the games we introduced in this paper) that uniform one-dimensional logics cannot express that a binary relation is an equivalence. In this section we consider the logics $U_1(\sim)$, $RU_1(\sim)$, $SU_1(\sim)$ that have free use of the equivalence relation $\sim$. (An alternative, but less interesting and less expressive variant would allow only uniform use of $\sim$, i.e., the use of $\sim$ as if it was an ordinary binary relation symbol.)

Even in the simplest of the three logics, $SU_1(\sim)$, one can express pretty complex properties such as, e.g., the existence of at most two equivalence classes with more than two elements:

$$\forall x_1...x_9 \big(\mathit{diff}(x_1,\ldots,x_9) \wedge (x_1 \sim x_2 \sim x_3) \wedge (x_4 \sim x_5 \sim x_6) \wedge x_1 \not\sim x_4$$
$$\rightarrow x_7 \sim x_1 \vee x_7 \sim x_4 \vee x_7 \not\sim x_8 \vee x_7 \not\sim x_9 \big).$$

Unrestricted use of $\sim$ allows also for a non-trivial interaction of $\sim$ with relations of arity greater than 2. One can express, e.g, that if, say, four elements are connected by a four-ary predicate R, then they are members of at least two equivalence classes.

We first observe that in $RU_1(\sim)$, models of doubly exponential size can be enforced, and use this to show a 2-NExpTime-lower bound for the satisfiability problem. Then we show that all our logics have the exponential classes property: if a formula is satisfiable, then it has a model in which all equivalence classes are bounded exponentially. We further use this result to show that $SU_1(\sim)$ has the exponential model property, and that both $RU_1(\sim)$ and $U_1(\sim)$ have the doubly exponential model property. This leads to tight complexity bounds for each logic.

5.1 Lower bound for $RU_1(\sim)$

In this section we show that the satisfiability and finite satisfiability problems for $RU_1(\sim)$ (and thus also for $U_1(\sim)$) are 2-NExpTime-hard. In particular this demonstrates that in $RU_1(\sim)$ one can construct satisfiable formulas whose models are of at least doubly exponential size with respect to their length.

We employ a reduction from a variant of the tiling problem. Let $\mathfrak{G}_m$ denote the standard $m \times m$ grid, $\mathfrak{G}_m = ([0, m-1]^2, H, V)$ with the horizontal and vertical successor relations H and V. A *tiling system* is a quadruple $\mathcal{T} = \langle C, c_0, \mathit{Hor}, \mathit{Ver} \rangle$, where C is a non-empty, finite set of *colours*, c_0 is an element of C, and Hor, Ver are binary relations on C called the *horizontal* and *vertical* constraints, respectively. A *tiling* for $\mathcal{T}$ of a grid $\mathfrak{G}_m$ is a function $f : G_m \rightarrow C$ such that $f(0,0) = c_0$, and for all $(d, d') \in H$, the pair $\langle f(d), f(d') \rangle$ is in Hor, and for all $(d, d') \in V$, the pair $\langle f(d), f(d') \rangle$ is in Ver. The *doubly exponential tiling problem* consists in checking for a given $n \in \mathbb{N}$ written in unary, and a tiling system $\mathcal{T}$, if $\mathcal{T}$ has a tiling of the grid $\mathfrak{G}_m$, where $m = 2^{2^n}$. It is well known that the doubly exponential tiling problem is 2-NExpTime-complete (see, e.g., [19], p. 501).

▶ **Theorem 5.** *The satisfiability and the finite satisfiability problems for $RU_1(\sim)$ are hard for 2-NExpTime.*

The proof is similar in spirit to the proof of the 2-NExpTime-lower bound for the two-variable fragment with two equivalence relations given in [11]. The crux is a succinct axiomatization of a grid structure of doubly exponential size.

Let $U_0, \ldots, U_{n-1}$ be unary predicates. By taking the predicates U_i to indicate the values of binary digits, we may take each element in any structure interpreting these predicates to have a 'local coordinate' in the range $[0, 2^n - 1]$; a point u of a model encodes the binary string s such that the ith bit of s is 1 iff $U_i(u)$ holds. It helps to think that an element's local coordinate fixes its position inside its equivalence class. We employ the abbreviation $\lambda^=(x, y)$ in order to state that x and y (which may be from different classes) have the same local coordinates; $\lambda^<(x, y)$ to state that the local coordinate of y is greater than the local coordinate of x; and $\lambda^{+1}(x, y)$ to state that the local coordinate of y is one greater than the local coordinate of x (addition modulo 2^n). All these abbreviations can be defined in the standard way using quantifier-free formulas of length polynomial in n. The formula

$$\forall x \exists y \big(x \sim y \wedge \lambda^{+1}(x, y) \big) \tag{2}$$

then ensures that each class contains a collection of 2^n elements, distinguished by local coordinates in the range $[0, 2^n - 1]$.

We now endow each class with a pair of 'global coordinates', corresponding to the grid coordinates in the range $[0, 2^{2^n} - 1]$. Let X and Y be unary predicates. The conjunct

$$\forall xy \big(x \sim y \wedge \lambda^=(x, y) \rightarrow ((X(x) \leftrightarrow X(y)) \wedge (Y(x) \leftrightarrow Y(y))) \big) \tag{3}$$

ensures that elements of the same class with the same local coordinates agree on the satisfaction of X and Y. For simplicity we allow ourselves to speak of *the* element of some class with a given local coordinate, since all such elements will turn out to have identical properties. If D is a class, we take the global X-coordinate of D to be the number in the range $[0, 2^{2^n} - 1]$ whose jth bit ($0 \leq j \leq 2^n - 1$) is 1 iff the element of D whose local coordinate is j satisfies the predicate X. Likewise, we define the global Y-coordinate of D using the predicate Y.

Now we enforce that for a class with global coordinates (p, q), there exists a class with coordinates $(p + 1, q)$ (if $p < 2^{2^n} - 1$) and a class with coordinates $(p, q + 1)$ (if $q < 2^{2^n} - 1$).

We take the predicate X^1 to mark in each class the least significant position satisfying X, and we define X^0 symmetrically:

$$\forall x \big(X^1(x) \leftrightarrow (X(x) \wedge \forall y(x \sim y \wedge \lambda^<(y, x) \rightarrow \neg X(y))) \big), \tag{4}$$

$$\forall x \big(X^0(x) \leftrightarrow (\neg X(x) \wedge \forall y(x \sim y \wedge \lambda^<(y, x) \rightarrow X(y))) \big). \tag{5}$$

Consider now the following formulas.

$$\forall x \big(X^0(x) \rightarrow \exists y(X^1(y) \wedge \lambda^=(x, y) \wedge H(x, y)) \big), \tag{6}$$

$$\forall xyx'y' \big(x \sim y \wedge x' \sim y' \wedge \lambda^=(x, x') \wedge H(y, y') \rightarrow$$
$$((Y(x) \leftrightarrow Y(x')) \wedge (\lambda^<(y, x) \rightarrow (X(x) \leftrightarrow X(x')))) \big) \tag{7}$$

They link via H the element marked by X^0 from one class to the element marked by X^1 from another class with the X-coordinate greater by one and with the same Y-coordinate.

Let (8)–(11) be formulas analogous to (4)-(7) using predicates Y^1, Y^0 and linking via the binary predicate V the element marked by Y^0 from one class to the element marked by Y^1 from another class with the Y-coordinate greater by one and with the same X-coordinate. The following formula which states that there exists a class with global coordinates $(0, 0)$,

$$\exists x \forall y(x \sim y \rightarrow \neg X(y) \wedge \neg Y(y)), \tag{12}$$

guarantees that a class with any pair of coordinates from the range $[0, 2^{2^n} - 1]$, exists. To finish our axiomatization of the grid, it remains to enforce that any pair of global coordinates appears at most once, or, in other words, that any pair of distinct classes have different global coordinates. This is done by means of additional binary predicates $R_0, \ldots, R_{n-1}$ that connect elements from two different classes. The binary predicates define a bit string that indicates the local coordinate on which the bits of X- or Y-coordinates of these classes differ.

$$\forall xyx'y'\big(x \sim y \wedge x' \sim y' \wedge \neg x \sim x' \wedge \bigwedge_i (R_i(x,x') \leftrightarrow U_i(y)) \wedge \lambda^=(y,y')$$
$$\rightarrow ((X(y) \leftrightarrow \neg X(y')) \vee (Y(y) \leftrightarrow \neg Y(y')))) \quad (13)$$

Consider the conjunction of (2)-(13). It should be clear that each of its models contains, for any $0 \le p, q < 2^{2^n}$, precisely one class with global coordinates (p, q).

Having established a grid of doubly exponential size, the encoding of any instance of the doubly exponential tiling problem on some tiling system (C, c_0, Hor, Ver) is routine. We simply use the following formulas.

$$\forall x \Big(\bigvee_{c \in C} P_c(x) \wedge \bigwedge_{c \ne d} \neg(P_c(x) \wedge P_d(x)) \Big), \quad (14)$$

$$\bigwedge_{c \in C} \forall xy \big(x \sim y \wedge P_c(x) \rightarrow P_c(y)\big), \quad (15)$$

$$\bigwedge_{\langle c,d \rangle \notin Hor} \forall xy \big(H(x,y) \wedge P_c(x) \rightarrow \neg P_d(y)\big), \quad (16)$$

$$\bigwedge_{\langle c,d \rangle \notin Ver} \forall xy \big(V(x,y) \wedge P_c(x) \rightarrow \neg P_d(y)\big), \quad (17)$$

$$\forall x \left(\forall y(x \sim y \rightarrow (\neg X(y) \wedge \neg Y(y))) \rightarrow P_{c_0}(x) \right). \quad (18)$$

Notice that (18) states that the grid point with coordinates (0,0) is coloured with c_0.

Let Ω be the conjunction of (2)-(18). From any model of Ω, we can read off a $\mathcal{T}$-tiling of size 2^{2^n} for example by inspecting the colours assigned to the elements with local coordinate 0 in each of the $2^{2 \cdot 2^n}$ classes. On the other hand, given any tiling for $\mathcal{T}$, we can construct a finite model of Ω in the obvious way. Thus we see that: (i) if Ω is satisfiable, then $(\mathcal{T}, n)$ has a tiling; (ii) if $(\mathcal{T}, n)$ has a tiling, then Ω is finitely satisfiable. This proves the theorem.

Note that formulas (7) and (13) are in the restricted uniform but not in strongly restricted uniform fragment. The use of $\mathrm{RU}_1(\sim)$ formulas is indeed crucial, since, as we will show later, $\mathrm{SU}_1(\sim)$ has the exponential model property.

5.2   Exponential classes property

In this section we show the following *exponential classes property* of our logics, which then will be used as an important tool in our decidability proofs.

▶ **Lemma 6.** *Let φ be a satisfiable formula in any of the logics $\mathrm{SU}_1(\sim)$, $\mathrm{RU}_1(\sim)$, $\mathrm{U}_1(\sim)$. Then φ has a model in which each equivalence class is bounded exponentially in $\|\varphi\|$.*

Here we show this property for $\mathrm{RU}_1(\sim)$ (which obviously covers also the case of $\mathrm{SU}_1(\sim)$). An extension of the proof covering the case of $\mathrm{U}_1(\sim)$ will be presented in the full version of the paper. The approach we employ is based up to an extent on the approach used in [12] to establish the *small substructures property* for FO^2 (which was further used in that paper to show the exponential classes property for $\mathrm{FO}^2(\sim)$). However, due to considering a richer logic, our proof is technically much more involved.

By Proposition 1, we can restrict attention to normal form formulas. Let us fix a normal form $\mathrm{RU}_1(\sim)$-formula φ of the form given in Equation (1) and a model $\mathfrak{A} \models \varphi$. Let n be the

width of φ, i.e., $n = \max(\{k_i + 1\}_{1 \le i \le m_\exists} \cup \{l_i\}_{1 \le i \le m_\forall})$. Recall that $m_\exists$ is the number of $\forall\exists^*$-conjuncts of φ. To simplify notation, we denote $m := m_\exists$.

In a single step of our construction we consider an equivalence class D in $\mathfrak{A}$, a 1-type α, and the fragment $D_\alpha \subseteq D$ of D consisting of all realizations of α. If $|D_\alpha| \le n$, then we do nothing. Otherwise, we replace D_α by a new fragment bounded polynomially in $\|\varphi\|$, obtaining a new model model $\mathfrak{A}' \models \varphi$. The universe of $\mathfrak{A}'$ consists of $A \setminus D_\alpha$ and a new set D'_α of realizations of α; $\mathfrak{A}'{\restriction}(A \setminus D'_\alpha) = \mathfrak{A}{\restriction}(A \setminus D_\alpha)$; and D'_α is formed out of three new disjoint sets D^0_α, D^1_α, D^2_α such that $D^i_\alpha = \{a^i_1, \ldots, a^i_{mn}\}$ for $i = 1, 2$, and $D^0_\alpha = \{a^0_1, \ldots, a^0_{(m+1)^3 n^2}\}$. For a set $B \subseteq A$, we call $B \setminus D_\alpha$ its *external fragment* and $B \cap D_\alpha$ its *internal* fragment. We also speak about external and internal fragments of witness structures and use analogous terminology for $\mathfrak{A}'$. The construction of $\mathfrak{A}'$ is divided into several stages.

Labelling of subsets. We take a set of labels $L_1, \ldots, L_m$. For each subset B of $A \setminus D_\alpha$ such that $1 \le |B| < n$, for each $a \in D_\alpha$, for each $1 \le i \le m$: if B forms the external fragment of the live part of a witness structure for a and $\varphi^\exists_i$ in $\mathfrak{A}$, then label B with L_i. Note that some subsets B may be labelled by several L_is, and some may have no labels.

Let L^* be a fresh label. For every $b \in A \setminus D_\alpha$ and every $1 \le i \le m$, choose a witness structure for b and $\varphi^\exists_i$ in $\mathfrak{A}$. If the internal fragment of the live part of this witness structure is not empty, then label the set of elements of this live part with L^*. Later on, we will take care of replicating such a witness structure for b in $\mathfrak{A}'$.

Special subsets. We collect some subsets of $A \setminus D_\alpha$ into a *set Θ of special subsets.* This set will be sufficiently rich to provide the external fragments of the live parts of witness structures for any element in D'_α. For each label L_i, $1 \le i \le m$, if there are at most m subsets of $A \setminus D_\alpha$ labelled by L_i, then make all of them members of Θ; call such a label *rare*. Otherwise, choose $m + 1$ such subsets and make them members of Θ. Note that $|\Theta| \le (m + 1)m$, and thus it is bounded polynomially in $\|\varphi\|$.

Witnesses for Θ. Now, for each subset $B \in \Theta$, we replicate in $\mathfrak{A}'$ those witness structures from $\mathfrak{A}$ whose live parts are labelled by L^* and their external fragments equal B. Assume $B = \{b_1, \ldots, b_k\}$. For each $\{a_1, \ldots, a_l\} \subseteq D_\alpha$, $l \ge 1$, such that $\{b_1, \ldots, b_k, a_1, \ldots, a_l\}$ is labelled by L^*, take fresh elements $a'_1, \ldots, a'_l$ from D^0_α and set $\text{tb}_{\mathfrak{A}'}(b_1, \ldots, b_k, a'_1, \ldots, a'_l) := \text{tb}_{\mathfrak{A}}(b_1, \ldots, b_k, a_1, \ldots, a_l)$. We simultanously begin defining a *pattern function* $f : D'_\alpha \to D_\alpha$ by setting $f(a'_i) := a_i$ for $1 \le i \le l$. Let us estimate the number of elements in D^0_α required for this step: There are at most $(m + 1)m$ subsets B in Θ, each of them of size smaller than n. In Step *Labelling of subsets*, each element of such B could produce at most m witness structures labelled with L^*, and each such structure has less than n elements in D_α. Thus we need at most $(m + 1)m(n - 1)m(n - 1)$ elements, and we indeed have that many, as we declared D^0_α to have $(m + 1)^3 n^2$ elements.

For all elements a' of D^0_α not used in the above step, as well as for elements a' from $D^1_\alpha \cup D^2_\alpha$, choose an arbitrary element $a \in D_\alpha$ and set $f(a') := a$.

Witnesses for elements of D^0_α. Let $a' \in D^0_\alpha$. If there is a subset in Θ such that a' was used to replicate a witness structure labelled by L^* in stage *Witnesses for Θ*, then call this set $B^*_{a'}$ (by our construction, there is at most one such set). We then continue without assuming that a' was necessarily used for replicating a set labelled L^*. Let $a = f(a')$ be the *pattern element* for a'. For each $1 \le j \le m$ such that L_j is rare, find a witness structure $\mathfrak{B}_j$ for a and $\varphi^\exists_j$ in $\mathfrak{A}$. Assume that $B_j = \{a, a_1 \ldots, a_k, a_{k+1}, \ldots, a_s, b_1, \ldots, b_l, b_{l+1}, \ldots, b_t\}$,

with $a_i \in D_\alpha$ for $1 \leq i \leq s$ and $b_i \in A \setminus D_\alpha$ for $1 \leq i \leq t$, and that the live part of $\mathfrak{B}_j$ is $\bar{B}_j = \{a, a_1 \ldots, a_k, b_1, \ldots, b_l\}$. It may happen that $l = 0$, which means that the live part of the witness structure is contained in D_α. Otherwise, the set $\{b_1, \ldots, b_l\}$ is labelled by L_j (and possibly some other labels) and is a member of Θ. We set $\mathrm{tb}_{\mathfrak{A}'}(a, a^1_{(j-1)n+1}, \ldots, a^1_{(j-1)n+k}, b_1, \ldots, b_l) := \mathrm{tb}_{\mathfrak{A}}(a, a_1, \ldots, a_k, b_1, \ldots, b_l)$. Note that this guarantees that the structure defined on the set $\{a', a^1_{(j-1)n+1}, \ldots, a^1_{(j-1)n+k}, a^1_{(j-1)n+k+1}, \ldots, a^1_{(j-1)n+s}, b_1, \ldots, b_l, b_{l+1}, \ldots, b_t\}$ will be a witness structure for a' and $\varphi^\exists_j$ in $\mathfrak{A}'$. Let us here comment one subtlety: if $k = 0$, then it is possible that $\mathrm{tb}_{\mathfrak{A}'}(a, b_1, \ldots, b_l)$ was defined before (either in this stage for some different j, or if $B^*_{a'} = \{b_1, \ldots, b_l\}$, in the step *Witnesses for Θ*). Note, however, that there is no danger of conflict here, since this earlier definition must agree with the new one.

For each $1 \leq j \leq m$ such that L_j is not rare, let $\{b_1, \ldots, b_l\}$ be a subset from Θ labelled by L_j, different from $B^*_{a'}$, and not yet used for a' for any other j (note that such a subset exists, since there are $m + 1$ subsets labelled by L_j in Θ, and at most $m - 1$ of them can be used as the external parts of witness structures for a' for other js). Let $a \in D_\alpha$, and let $\mathfrak{B}_j$ be a witness structure for a and $\varphi^\exists_j$ in $\mathfrak{A}$ in which the external fragment of the live part is formed by $\{b_1, \ldots, b_l\}$. Such a and $\mathfrak{B}_j$ exist due to the construction of Θ. Assume that $B_j = \{a, a_1 \ldots, a_k, a_{k+1}, \ldots, a_s, b_1, \ldots, b_l, b_{l+1}, \ldots, b_t\}$, with $a_i \in D_\alpha$ for $1 \leq i \leq s$, $b_i \in A \setminus D_\alpha$ for $1 \leq i \leq t$. Assume the live part of B_j is $\bar{B}_j = \{a, a_1 \ldots, a_k, b_1, \ldots, b_l\}$. We set $\mathrm{tb}_{\mathfrak{A}'}(a', a^1_{(j-1)n+1}, \ldots, a^1_{(j-1)n+k}, b_1, \ldots, b_l) := \mathrm{tb}_{\mathfrak{A}}(a, a_1, \ldots, a_k, b_1, \ldots, b_l)$. This guarantees that the structure defined on the set $\{a', a^1_{(j-1)n+1}, \ldots, a^1_{(j-1)n+k}, a^1_{(j-1)n+k+1}, \ldots, a^1_{(j-1)n+s}, b_1, \ldots, b_l, b_{l+1}, \ldots, b_t\}$ will be a witness structure for a' and $\varphi^\exists_j$ in $\mathfrak{A}'$.

Witnesses for D^1_α and D^2_α. Witness structures for $a' \in D^1_\alpha \cup D^2_\alpha$ are provided in a similar way to that described for elements of D^0_α, but if $a' \in D^1_\alpha$ ($a' \in D^2_\alpha$), then we take elements a'_i from D^2_α (D^0_α). This cyclic scheme guarantees that the procedure avoids conflicts.

Witnesses for subsets of $A \setminus D'_\alpha$ not belonging to Θ. Let $B = \{b_1, \ldots, b_l\}$ be a subset of $A \setminus D_\alpha$, $l \leq n$, not belonging to Θ. Note that no table with external part B has yet been defined. The number of subsets labelled by L^* whose external part equals B is bounded above by mn. Since the internal part of each of them has less than n elements, we can replicate them without conflicts using mn^2 elements of D^0_α.

Completion. Let $\{a'_1, \ldots, a'_k, b_1, \ldots, b_l\} \subseteq A'$ be such that $a'_i \in D'_\alpha$ for $1 \leq i \leq k$, $b_i \in A \setminus D_\alpha$ for $1 \leq i \leq l$, $k + l \leq n$, and such that $\mathrm{tb}_{\mathfrak{A}'}(a'_1, \ldots, a'_k, b_1, \ldots, b_l)$ has not been defined yet. Take any pairwise distinct elements $a_1, \ldots, a_k$ from D_α (such elements exists, since $k \leq n$ and $|D_\alpha| \geq n$) and set $\mathrm{tb}_{\mathfrak{A}'}(a'_1, \ldots, a'_k, b_1, \ldots, b_l) := \mathrm{tb}_{\mathfrak{A}}(a_1, \ldots, a_k, b_1, \ldots, b_l)$.

This finishes the construction for replacing D_α by a small set D'_α. We now argue that the obtained model $\mathfrak{A}'$ is indeed a model of φ.

▶ **Claim 7.** $\mathfrak{A}' \models \varphi$.

Proof. Let us see first that all elements have the required witness structures. Let $b \in A'$ and $1 \leq i \leq m$. If $b \in D'_\alpha$, then an appropriate witness structure for b and $\varphi^\exists_j$ was constructed in the step *Witnesses for D^0_α* or *Witnesses for D^1_α and D^2_α*. Assume that $b \in A' \setminus D'_\alpha (= A \setminus D_\alpha)$. Either b has a witness structure for $\varphi^\exists_i$ in $\mathfrak{A} \restriction A \setminus D_\alpha$ and this structure is inherited into $\mathfrak{A}'$, or, in the step *Labelling* we labelled with L^* at least one witness structure $\mathfrak{B}$ for b and $\varphi^\exists_j$ in $\mathfrak{A}$. If the external fragment of the live part $\bar{B}$ of this structure is in Θ, then the live part of the corresponding witness structure in $\mathfrak{A}'$ was defined in step *Witnesses for Θ*. Otherwise it was

defined in step *Witness for subsets of $A \setminus D'_\alpha$ not belonging to* Θ. The necessary members of the witness structure which are not live can easily be found.

Consider now a conjunct of φ of the form $\forall x_1, \ldots, x_u \varphi_j^\forall$ and a tuple of not necessarily distinct elements $d'_1, \ldots, d'_u \in A'$. We want to see that $\mathfrak{A}' \models \varphi_j^\forall[d'_1, \ldots, d'_u]$. Let us enumerate the elements of $\{d'_1, \ldots, d'_u\}$ by $a'_1, \ldots, a'_k, a'_{k+1}, \ldots, a'_s, b_1, \ldots, b_l, b_{l+1}, \ldots, b_t$, where $a'_i \in D'_\alpha$ for $1 \leq i \leq s$, $b_i \in A \setminus D'_\alpha$ for $1 \leq i \leq t$, and $\mathrm{live}(\varphi_j^\forall[d'_1, \ldots, d'_u]) = \{a'_1, \ldots, a'_k, b_1, \ldots, b_l\}$. By our construction, $\mathrm{tb}_{\mathfrak{A}'}(a'_1, \ldots, a'_k, b_1, \ldots, b_l) = \mathrm{tb}_{\mathfrak{A}}(a_1, \ldots, a_k, b_1, \ldots, b_l)$ for some $a_1, \ldots, a_k \in D_\alpha$ (if $k = 0$, then simply $\mathrm{tb}_{\mathfrak{A}'}(b_1, \ldots, b_l) = \mathrm{tb}_{\mathfrak{A}}(b_1, \ldots, b_l)$; otherwise the table was set either in one of *Witnesses for ...* -stages or in the *Completion* stage). Take now any pairwise distinct elements $a_{k+1}, \ldots, a_s \in D_\alpha$ different from $a_1, \ldots, a_k$ (this is possible since $s \leq n$ and there are at least n elements in D_α) and observe that the equivalence relations among $a_1, \ldots, a_s, b_1, \ldots, b_t$ are isomorphic to those among $a'_1, \ldots, a'_s, b_1, \ldots, b_t$. This guarantees that $\mathfrak{A}' \models \varphi_j^\forall[d'_1, \ldots, d'_u]$. ◀

Now, to find a small replacement of a whole class, we apply the described construction iteratively to all D_α, where α is a 1-type realized in this class. Let $D_1, D_2, \ldots$ be a (possibly infinite) sequence of the classes in a $\mathfrak{A}$ (we can assume that $\mathfrak{A}$ is countable due to the Löwenheim-Skolem property), let $\mathfrak{A}_0 = \mathfrak{A}$, and let $\mathfrak{A}_{j+1}$ be the structure $\mathfrak{A}_j$ modified by replacing class D_{j+1} by the small replacement class D'_{j+1} as described above. The obtained natural limit structure is the desired model with exponentially bounded classes.

5.3 Exponential model property for $\mathrm{SU}_1(\sim)$

Recall that in Section 5.1 we proved a 2-NExpTime lower bound for $\mathrm{RU}_1(\sim)$. Here we show that $\mathrm{SU}_1(\sim)$ is easier. To understand why the complexity drop is possible, consider the conjunct (6) of the formula Ω from Section 5.1. When we look for a witness structure for an element a satisfying X^0 and this conjunct, we have to find an appropriate element b (i.e., an element with the same local coordinate as a, satisfying X^1, connected to a by H), but additionally, due to the conjunct (7), we must take into account the 1-types of elements from the classes of a and b that *do not belong to the witness structure*. The restrictions of $\mathrm{SU}_1(\sim)$ would not enable this. Indeed we can now prove the following theorem.

▶ **Theorem 8.** *The satisfiability problem for* $\mathrm{SU}_1(\sim)$ *is* NExpTime-*complete.*

To prove this theorem, we establish the exponential model propery. Thus checking if a given formula is satisfiable can be done by nondeterministically guessing a structure of exponentially bounded size and verifying that it is indeed a model. Such a model checking task (for normal form formulas) can be done in polynomial time in a straighforward way. The matching lower bound follows from the NExpTime-hardness of FO^2.

▶ **Lemma 9.** *Every satisfiable* $\mathrm{SU}_1(\sim)$ *formula* φ *has a finite model of size bounded exponentially in* $\|\varphi\|$.

We next prove this lemma. For the rest of this section, fix a normal form formula φ of $\mathrm{SU}_1(\sim)$ and a model $\mathfrak{A} \models \varphi$. Due to Lemma 6, we may assume that the equivalence classes of $\mathfrak{A}$ are bounded exponentially in $\|\varphi\|$. As previously, let n be the width of φ and m the number of $\forall \exists^*$-conjuncts of φ. We construct a small model $\mathfrak{A}' \models \varphi$ in several stages.

Court. If a k-multitype, for $1 \leq k \leq n$, is realized in less than n classes of $\mathfrak{A}$, then call all these classes *royal*. If a k-multitype, for $1 \leq k \leq n$, is realized only in royal classes, then call this multitype *royal*. Note that it is possible that a multitype realized in more than n

classes is royal. Let K be the union of all royal classes of $\mathfrak{A}$. For each $a \in K$ and for each conjunct $\varphi_i^\exists$ of φ, find a witness structure $\mathfrak{C}_{a,i}$ for a and $\varphi_i^\exists$ in $\mathfrak{A}$. Let C be the union of K and all the classes containing some element from some $C_{a,i}$. Note that the size of C is bounded exponentially in $\|\varphi\|$. $\mathfrak{C}$ is called the *court* of $\mathfrak{A}$.

Universe. For all $1 \leq k \leq n$, for each non-royal k-multitype θ realized in a class of $\mathfrak{A}$, *appoint one* such a class D_θ. We build a new model $\mathfrak{A}' \models \varphi$ whose universe is $C \cup \bigcup D_{\theta,u,w}^s$, where each $D_{\theta,u,w}^s$ is a fresh set and the union is taken over all non-royal k-multitypes θ realized in a class of $\mathfrak{A}$ ($1 \leq k \leq n$), $0 \leq s \leq 2$, $1 \leq u \leq n$, $1 \leq w \leq m$. For $i = 0, 1, 2$, let D^i be the union of all $D_{\theta,u,w}^s$ with $s = i$. We make $\mathfrak{A}' {\restriction} C$ isomorphic to $\mathfrak{A} {\restriction} C$, and $\mathfrak{A}' {\restriction} (K \cup D_{\theta,u,w}^s)$ isomorphic to $\mathfrak{A} {\restriction} (K \cup D_\theta)$, for all relevant θ, s, u, w. For each $D_{\theta,u,w}^s$, let $g_{\theta,u,w}^s : D_{\theta,u,w}^s \to D_\theta$ be an isomorphism. Also, for a class D in $\mathfrak{C}$, let $g_D : D \to D$ be the identity function. Define $g : A' \to A$ to be $g := \bigcup_{D \in C/\sim} g_D \cup \bigcup g_{\theta,u,w}^s$, where the second union is taken over all relevant θ, s, u, w. We call g the *pattern function*. It will return, for each element of $\mathfrak{A}'$, a 'similar' element in $\mathfrak{A}$. At this stage the structure of $\mathfrak{A}'$ is defined on $\mathfrak{C}$, on each equivalence class, and on each union of a non-royal class with K. The size of $\mathfrak{A}'$ is exponentially bounded in $\|\varphi\|$, as required.

Witnesses. Let $a' \in A' \setminus K$. Let a be the *pattern element* for a', $a := g(a')$. For each $1 \leq j \leq m$, find a witness structure $\mathfrak{B}_{a,j}$ for a and $\varphi_j^\exists$ in $\mathfrak{A}$. We want to define a similar witness structure for a' in $\mathfrak{A}'$. We consider explicitly the case in which $a' \in D^0$. Assume that the class of a' is $D_{\theta,u,w}^0$. Then $a \in D_\theta$. Let $T_0, T_1, \ldots, T_s, T_{s+1}, \ldots, T_t$ be the division of $\mathfrak{B}_{a,j}$ into classes such that $a \in T_0 \subseteq D_\theta$, the sets $T_1, \ldots, T_s$ (possibly $s = 0$) are fragments of non-royal classes of $\mathfrak{A}$ and $T_{s+1}, \ldots, T_t$ (possibly $t - s = 0$) are fragments of royal classes of $\mathfrak{A}$. Let θ_i be the multitype of T_i, for $1 \leq i \leq t$. We define a function $h : B_{a,j} \to A'$ whose image is supposed to form a witness structure for a' in $\mathfrak{A}'$. For $i = 1, \ldots, s$, let T_i' be a set of multitype θ_i from $D_{\theta_i,i,j}^1$ (recall that $i \leq s \leq n$), and let $h_i : T_i \to T_i'$ be a bijection preserving 1-types. We set

$$
h(b) := \begin{cases} b' \text{ such that } g(b') = b & \text{if } b \in T_0; \\ h_i(b) & \text{if } b \in T_i \text{ for } 1 \leq i \leq s; \\ b & \text{if } b \in T_i \text{ for } i > s. \end{cases}
$$

Let $b_1, \ldots, b_k$ be an enumeration of the elements of $B_{a,j}$. If $s = 0$, i.e., all elements of $B_{a,j} \setminus \{a\}$ are in royal classes, then $\mathfrak{A}' {\restriction} \{h(b_1), \ldots, h(b_k)\}$ already forms a witness structure for a and $\varphi_j^\exists$. Otherwise we set $\mathrm{tb}_{\mathfrak{A}'}(h(b_1), \ldots, h(b_k)) := \mathrm{tb}_{\mathfrak{A}}(b_1, \ldots, b_k)$.

If $a' \in D^1$, then we proceed similarly, but use elements of D^2 instead of elements of D^1; if $a' \in D^2$ or $a' \in K \setminus C$, then we use elements of D^0 instead of elements of D^1. This circular witnessing scheme, inspired by the one from [7], together with the strategy of using an appropriate number of copies of classes, guarantees that for each subset $B \subseteq A'$, the table for some enumeration of its elements is defined at most once.

Completion. Let $a_1', \ldots, a_k'$, $1 \leq k \leq n$, be a tuple of elements from $\mathfrak{A}'$ whose table is not yet defined. Let $T_1', \ldots, T_s', T_{s+1}', \ldots, T_t'$ be a partition of $\{a_1', \ldots, a_k'\}$ into classes such that the sets $T_1', \ldots, T_s'$ are fragments of non-royal classes of $\mathfrak{A}'$ (this time $s > 0$, since otherwise all elements of the tuple would belong to K, and thus their table would have been already defined) and $T_{s+1}', \ldots, T_t'$ (possibly $t - s = 0$) are fragments of royal classes of $\mathfrak{A}'$. Assume that the multitype of T_i' is θ_i for $1 \leq i \leq t$. Since $s \leq n$, and as the multitypes of

$T'_1, \ldots, T'_s$ are non-royal, we can find in distinct classes of $\mathfrak{A}$ subsets $T_1, \ldots, T_s$ such that the multitype of T_i in $\mathfrak{A}$ equals the multitype of T'_i in $\mathfrak{A}'$. For $1 \leq i \leq t$, let $h'_i : T'_i \to T_i$ be a bijection preserving 1-types (for $i > s$ it can be just the identity). Let $h' = \bigcup_{1 \leq i \leq t} h'_i$. Set $\text{tb}_{\mathfrak{A}'}(a'_1, \ldots, a'_k) := \text{tb}_{\mathfrak{A}}(h'(a'_1), \ldots, h'(a'_k))$. This finishes the construction of $\mathfrak{A}'$.

We provided witness structures for $\forall\exists^*$-conjuncts of φ for elements from K in step *Court*, and for elements from $A' \setminus K$ in step *Witnesses*. All universal conjuncts of φ are satisfied since for any tuple of elements from A', its table is defined precisely as a table in $\mathfrak{A}$ for some elements with the same 1-types and isomorphic equivalence connections. Thus

▶ **Claim 10.** $\mathfrak{A}' \models \varphi$.

5.4 Doubly exponential model property for $\text{RU}_1(\sim)$ and $\text{U}_1(\sim)$

The following result completes our discourse on complexity.

▶ **Theorem 11.** *The satisfiability problems for* $\text{RU}_1(\sim)$ *and* $\text{U}_1(\sim)$ *are* 2-NExpTime-*complete.*

The lower bound for $\text{RU}_1(\sim)$ (and thus $\text{U}_1(\sim)$) was shown in Theorem 5. The matching upper bound follows from the following small model property.

▶ **Lemma 12.** *Every satisfiable formula φ of* $\text{RU}_1(\sim)$ *or* $\text{U}_1(\sim)$ *has a finite model of size bounded doubly exponentially in* $\|\varphi\|$.

A proof of this lemma will appear in the full version of this paper. Here we only remark that instead of working with multitypes of small subsets of classes, as we did in Section 5.3 in the case of $\text{SU}_1(\sim)$, this time we consider multitypes of whole classes. Due to the exponential bound on the size of classes, the number of possible multitypes of classes in a model is bounded doubly exponentially in $\|\varphi\|$. The proof consists of reproducing an appropriate number of realizations of every multitype in the new small model. The basic structure of our small model construction is similar to that in Section 5.3.

5.5 Limits of decidability

We consider two natural generalisations of the weakest of our logics $\text{SU}_1(\sim)$ and show their undecidability. Let us denote by $\text{SU}_1(\sim_1, \sim_2)$ the extension of $\text{SU}_1(\sim)$ in which there are two distinguished binary predicates which must be interpreted as equivalences (and can be used freely), rather than just one. Let $\text{SU}_1(tr)$ be the extension of SU_1 in which a distinguished binary symbol tr must be interpreted as an arbitrary transitive relation (and can be used freely), rather than as an equivalence relation. Note that $\text{SU}_1(tr)$ is an extension of $\text{SU}_1(\sim)$ since reflexivity and symmetry of tr can be enforced in SU_1 in a straighforward way.

▶ **Theorem 13.** *The satisfiability and finite satisfiability problems for* $\text{SU}_1(\sim_1, \sim_2)$ *and* $\text{SU}_1(tr)$ *are undecidable.*

Proof. Recall the tiling systems from Section 5.1. We define the standard infinite grid as $\mathfrak{G}_{\mathbb{N}} = (\mathbb{N} \times \mathbb{N}, H, V)$, $H = \{((p,q),(p',q)) : p' - p = 1\}$, $V = \{((p,q),(p,q')) : q' - q = 1\}$. The standard grid $\mathfrak{G}_m^*$ on a finite torus is defined as $\mathfrak{G}_m$ from Section 5.1, with additional horizontal H-edges from the last to the first column, and additional vertical V-edges from the last to the first row. It is well known that the problem of checking if a tiling system $\mathcal{T}$ tiles the standard infinite grid $\mathfrak{G}_{\mathbb{N}}$, and the problem of checking if it tiles a toroidal grid $\mathfrak{G}_m^*$
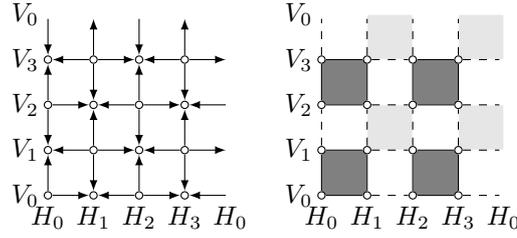
Figure 1 Grid structures $\hat{\mathfrak{G}}_{\mathbb{N}}$ for $\mathrm{SU}_1(tr)$ and $\check{\mathfrak{G}}_{\mathbb{N}}$ for $\mathrm{SU}_1(\sim_1, \sim_2)$. Arrows indicate tr connections, solid edges represent $\sim_1$, dashed edges represent $\sim_2$, equivalence classes of $\sim_1$ and $\sim_2$ are indicated by, respectively, dark and light shadings.

for some $m \in \mathbb{N}$, are undecidable. We can encode these problems in $\mathrm{SU}_1(tr)$ and $\mathrm{SU}_1(\sim_1, \sim_2)$ quite easily. We concentrate on the proof for $\mathrm{SU}_1(tr)$. The proof for $\mathrm{SU}_1(\sim_1, \sim_2)$ is similar.

For a given tiling system $\mathcal{T}$, we construct an $\mathrm{SU}_1(tr)$ formula Θ. Our intended grid expansion $\hat{\mathfrak{G}}_{\mathbb{N}}$ is illustrated in Fig. 1. It interprets auxiliary unary symbols H_i, V_i, for $0 \leq i, j \leq 3$, and, obviously, the transitive symbol tr. It is crucial that $\hat{\mathfrak{G}}_{\mathbb{N}}$ avoids binary connections between points which are distant from each other.

We capture properties of horizontally and vertically neighbouring elements by formulas $\lambda_H(x, y)$ and $\lambda_V(x, y)$,

$$\lambda_H(x, y) \equiv \bigvee_{0 \leq i, j \leq 3} \lambda_H^{i,j}(x, y), \tag{19}$$

where $\lambda_H^{i,j} \equiv H_i(x) \wedge H_{i+1}(y) \wedge V_j(x) \wedge V_j(y) \wedge tr^{i+j}(x, y)$; here $i + 1$ is taken modulo 4, and $tr^k(x, y)$ denotes $tr(x, y)$ for even k, and $tr(y, x)$ for odd k. $\lambda_V(x, y)$ is defined analogously. Grid coordinate points are appropriately completed:

$$\forall x (\exists y \lambda_H(x, y) \wedge \exists y \lambda_V(x, y)), \tag{20}$$

$$\forall xyzt(\lambda_H(y, z) \wedge \lambda_V(y, x) \wedge \lambda_V(z, t) \rightarrow \lambda_H(x, t)). \tag{21}$$

Finally, we encode an instance of the tiling problem $\mathcal{T} = (C, c_0, Hor, Ver)$, similarly to the way we did it in Section 5.1.

$$\exists x (H_0(x) \wedge V_0(x) \wedge P_{c_0}(x)), \tag{22}$$

$$\forall x \Big(\bigvee_{c \in C} P_c(x) \wedge \bigwedge_{c \neq d} \neg(P_c(x) \wedge P_d(x)) \Big), \tag{23}$$

$$\bigwedge_{\langle c,d \rangle \notin Hor} \forall xy \big(\lambda_H(x, y) \wedge P_c(x) \rightarrow \neg P_d(y) \big), \tag{24}$$

$$\bigwedge_{\langle c,d \rangle \notin Ver} \forall xy \big(\lambda_V(x, y) \wedge P_c(x) \rightarrow \neg P_d(y) \big). \tag{25}$$

Let Θ be the conjunction of (20)-(25). We claim that Θ is satisfiable iff $\mathcal{T}$ tiles $\mathfrak{G}_{\mathbb{N}}$, and that Θ is finitely satisfiable iff $\mathcal{T}$ tiles $\mathfrak{G}_m^*$ for some $m \in \mathbb{N}$. We sketch the argument for the first part of the claim. Assume that $\mathcal{T}$ tiles $\mathfrak{G}_{\mathbb{N}}$. Take a tiling $f : G_{\mathbb{N}} \rightarrow C$, and consider the expansion of $\hat{\mathfrak{G}}_{\mathbb{N}}$ which satisfies $P_{f(i,j)}[i, j]$ for every $i, j \in \mathbb{N}$. It is readily verified that it is a model of Θ (here it is important that in our arrangement of the tr-arrows, the transitivity of tr does not enforce connections between distant points). In the opposite direction, if Θ has a model $\mathfrak{M}$, then using (20)-(22) we can define a homomorphism $F : \mathfrak{G}_{\mathbb{N}} \rightarrow \mathfrak{M}$ mapping $\langle 0, 0 \rangle$ to an element that satisfies P_{c_0}. Further, using (22)-(25), we can define a tiling f of $\mathfrak{G}_{\mathbb{N}}$ by

setting $f(i,j) = c$ for the unique c such that $\mathfrak{M} \models P_c[F(i,j)]$. We skip here the (routine) argument for the case of finite satisfiability. This finishes the proof for $SU_1(tr)$.

The case of $SU_1(\sim_1, \sim_2)$ is treated analogously. The only changes are that we use $\breve{\mathfrak{G}}_{\mathbb{N}}$ from Fig. 1 instead of $\hat{\mathfrak{G}}_{\mathbb{N}}$, and modify appropriately the definitions of $\lambda_H(x,y)$ and $\lambda_V(x,y)$. ◄

The above undecidability results contrast with the fact that FO^2 remains decidable when extended by two equivalence relations [12, 13] or one transitive relation [23]. It should be emphasised, however, that our undecidability proofs exploit the free, non-uniform use of the special relation symbols (as in conjunct (21)), rather than transitivity of the relations corresponding to the symbols. (Actually, the presented arguments work in a natural way if we do not require tr to be interpreted as a transitive relation.) It is likely that the decidability can be regained if we require the special symbols to obey the regular uniformity constraints. We leave this, however, for future work.

─── **References** ───

1   H. Andréka, J. van Benthem, and I. Németi. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27:217–274, 1998.

2   V. Bárány, B. ten Cate, and L. Segoufin. Guarded negation. In *ICALP 2011, Proceedings Part II*, pages 356–367, 2011.

3   S. Benaim, M. Benedikt, W. Charatonik, E. nski, R. Lenhardt, F. Mazowiecki, and J. Worrell. Complexity of two-variable logic on finite trees. In *ICALP 2013, Proceedings Part II*, pages 74–88, 2013.

4   M. Bojanczyk, C. David, A. Muscholl, T. Schwentick, and L. Segoufin. Two-variable logic on data words. *ACM Transactions on Computational Logic*, 12(4):27, 2011.

5   M. Bojanczyk, A. Muscholl, T. Schwentick, and L. Segoufin. Two-variable logic on data trees and XML reasoning. *Journal of the ACM*, 56(3), 2009.

6   H.-D. Ebbinghaus and J. Flum. *Finite Model Theory.* Springer, 1885.

7   E. Grädel, P. Kolaitis, and M. Vardi. On the decision problem for two-variable first-order logic. *Bulletin of Symbolic Logic*, 3(1):53–69, 1997.

8   E. Grädel, M. Otto, and E. Rosen. Two-variable logic with counting is decidable. In *LICS 1997, Proceedings*, pages 306–317, 1997.

9   L. Hella and A. Kuusisto. One-dimensional fragment of first-order logic. In *AIML 2014, Proceedings*, pages 274–293, 2014.

10   E. Kieronski and A. Kuusisto. Complexity and expressivity of uniform one-dimensional fragment with equality. In *MFCS 2014, Proceedings Part I*, pages 365–376, 2014.

11   E. Kieronski, J. Michaliszyn, I. Pratt-Hartmann, and L. Tendera. Two-variable first-order logic with equivalence closure. *SIAM Journal on Computing*, 43(3):1012–1063, 2014.

12   E. Kieronski and M. Otto. Small substructures and decidability issues for first-order logic with two variables. *Journal of Symbolic Logic*, 77:729–765, 2012.

13   E. Kieronski and L. Tendera. On finite satisfiability of two-variable first-order logic with equivalence relations. In *LICS 2009, Proceedings*, pages 123–132, 2009.

14   B. ten Cate L. Segoufin. Unary negation. *Logical Methods in Computer Science*, 9(3), 2013.

15   L. Libkin. *Elements of Finite Model Theory.* Springer, 2004.

16   A. Montanari and P. Sala. Adding an equivalence relation to the interval logic $AB\bar{B}$: complexity and expressiveness. In *LICS 2013, Proceedings*, pages 193–202, 2013.

17 M. Mortimer. On languages with two variables. *Zeitschr. f. math. Logik u. Grundlagen d. Math.*, 21:135–140, 1975.

18 L. Pacholski, W. Szwast, and L. Tendera. Complexity of two-variable logic with counting. In *LICS 1997, Proceedings*, pages 318–327, 1997.

19 C.H. Papadimitriou. *Computational Complexity.* Addison Wesley Longman, 1994.

20 I. Pratt-Hartmann. Complexity of the two-variable fragment with counting quantifiers. *Journal of Logic, Language and Information*, 14(3):369–395, 2005.

21 I. Pratt-Hartmann. Logics with counting and equivalence. In *CSL-LICS 2014, Proceedings*, page 76, 2014.

22 D. Scott. A decision method for validity of sentences in two variables. *Journal of Symbolic Logic*, 27:477, 1962.

23 W. Szwast and L. Tendera. FO^2 with one transitive relation is decidable. In *STACS 2013, Proceedings*, pages 317–328, 2013.