max planck institut
informatik

Decidable Fragments of First-Order Logic modulo Linear Rational Arithmetic

Marco Voigt

July 09/16, 2019

SIC Saarland Informatics Campus

Introduction

Our goal in the next two lectures is to identify decidable FOL(T) fragments beyond Bernays–Schönfinkel with simple bounds.

Since our approach will use model-theoretic arguments, we start with some basics illustrating the model-theoretic way of thinking:

(1) reminder of FOL semantics
(2) finite and infinite models
(3) the finite model property for BS sentences
(4) domain constraints and the Löwenheim–Skolem Theorem
(5) the finite model property for monadic FO sentences

1. Some Basics from Model Theory

Reminder: Semantics of FOL formulas

Let $\Sigma = (\Pi, \Omega)$ be a single-sorted signature, where

- Π is a finite set of predicates
- Ω is a finite set of functions

FO formulas over Σ are interpreted by Σ-structures.

Definition (Σ-algebra / Σ-interpretation / Σ-structure)

A *Σ-structure* $\mathcal{A}$ comprises
(1) a nonempty set $\mathcal{U}^{\mathcal{A}}$, called *universe* or *domain*,
(2) for every $P \in \Pi$ with arity m a set

$$P^{\mathcal{A}} \subseteq (\mathcal{U}^{\mathcal{A}})^m,$$

(3) a for every constant $c \in \Omega$ a domain element $c^{\mathcal{A}} \in \mathcal{U}^{\mathcal{A}}$,
(4) for every $f \in \Omega$ with arity $m \geq 1$ a total function

$$f^{\mathcal{A}} : (\mathcal{U}^{\mathcal{A}})^m \to \mathcal{U}^{\mathcal{A}}.$$

Reminder: Semantics of FOL formulas

Example:

Signature $\Sigma = (\Omega, \Pi)$ with unary $s \in \Omega$ and binary $E \in \Pi$.

Consider the Σ-structure $\mathcal{A}$ with

$\mathcal{U}^{\mathcal{A}}$:

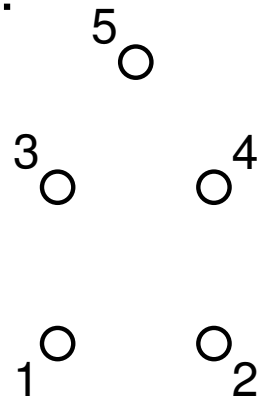

$E^{\mathcal{A}}$:

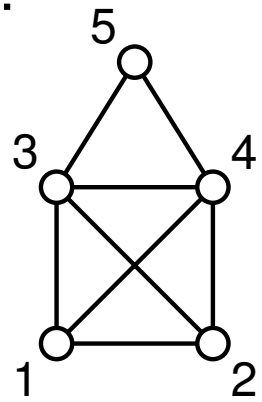

symmetric

$s^{\mathcal{A}}$:

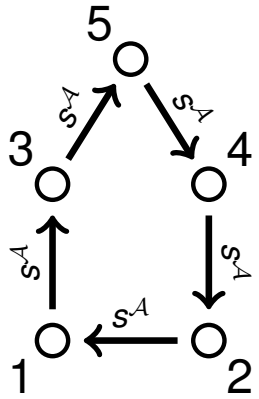

Reminder: Semantics of FOL formulas

Definition (Satisfaction relation)

Given some Σ-structure $\mathcal{A}$ and a variable assignment $\beta : \text{Var} \to \mathcal{U}^{\mathcal{A}}$, we define the *satisfaction relation* $\models$ such that

$$\begin{aligned}
\mathcal{A}, \beta \models s \approx t &\text{ iff } \mathcal{A}(\beta)(s) = \mathcal{A}(\beta)(t) \\
\mathcal{A}, \beta \models P(s_1, \ldots, s_m) &\text{ iff } \big(\mathcal{A}(\beta)(s_1), \ldots, \mathcal{A}(\beta)(s_m)\big) \in P^{\mathcal{A}}
\end{aligned}$$

$$\begin{aligned}
\mathcal{A}, \beta \models \neg\varphi &\text{ iff } \mathcal{A}, \beta \not\models \varphi \\
\mathcal{A}, \beta \models \varphi \wedge \psi &\text{ iff } \mathcal{A}, \beta \models \varphi \text{ and } \mathcal{A}, \beta \models \psi \\
\mathcal{A}, \beta \models \varphi \vee \psi &\text{ iff } \mathcal{A}, \beta \models \varphi \text{ or } \mathcal{A}, \beta \models \psi \\
\mathcal{A}, \beta \models \forall x.\, \varphi &\text{ iff } \mathcal{A}, \beta[x \mapsto a] \models \varphi \text{ for every } a \in \mathcal{U}^{\mathcal{A}} \\
\mathcal{A}, \beta \models \exists x.\, \varphi &\text{ iff } \mathcal{A}, \beta[x \mapsto a] \models \varphi \text{ for some } a \in \mathcal{U}^{\mathcal{A}}
\end{aligned}$$

We write $\varphi(\bar{x})$ to say that all free variables in φ belong to $\bar{x}$. If φ does not contain free variables, we call it a *sentence* and simply write $\mathcal{A} \models \varphi$ or $\mathcal{A} \not\models \varphi$. In case of $\mathcal{A} \models \varphi$ we call $\mathcal{A}$ a *model* of φ.

Reminder: Semantics of FOL formulas

Example:

Signature $\Sigma = (\Omega, \Pi)$ with unary $s \in \Omega$ and binary $E \in \Pi$.

Consider the Σ-structure $\mathcal{A}$ with

$\mathcal{U}^{\mathcal{A}}$:

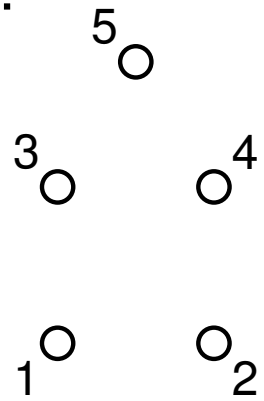

$E^{\mathcal{A}}$:

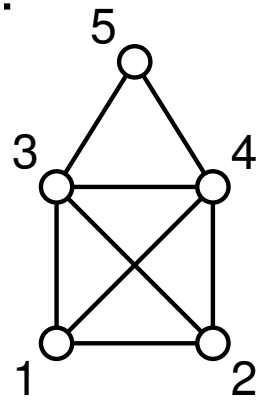

symmetric

$s^{\mathcal{A}}$:

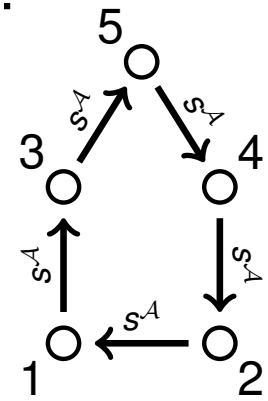

We observe $\mathcal{A} \models \forall x \exists y.\, E(x, y)$ $\quad \mathcal{A} \models \forall xy.\, E(x, y) \rightarrow E(y, x)$

$\mathcal{A} \not\models \exists z \forall x.\, s(x) \not\approx z$ $\quad \mathcal{A} \models \forall x.\, s(x) \not\approx x \wedge s(s(x)) \not\approx x$

Finite and infinite models

Proposition

There are satisfiable FO sentences that do not have finite models.

Proof: Consider the following *infinity axioms* [BGG97], Section 6.5

$$\underbrace{\big(\forall x.\, \neg P(x,x)\big)}_{\text{irreflexivity}} \wedge \underbrace{\big(\forall xyz.\, P(x,y) \wedge P(y,z) \rightarrow P(x,z)\big)}_{\text{transitivity}} \wedge \underbrace{\big(\forall x \exists y.\, P(x,y)\big)}_{\substack{\text{existence of} \\ P\text{-successors}}}$$

$$\big(\forall x.\, \neg P(x,x)\big) \wedge \big(\forall x \exists y.\, P(x,y) \wedge \forall z.\, P(y,z) \rightarrow P(x,z)\big)$$

$$\big(\exists v \forall x.\, f(x) \not\approx v\big) \;\wedge\; \big(\forall xy.\, f(x) \approx f(y) \rightarrow x \approx y\big)$$

Each of these three sentences is satisfiable over infinite structures only.

Finite and infinite models

Definition (Finite model property)

Let $\mathcal{C}$ be any class of FO sentences. We say that $\mathcal{C}$ enjoys the *finite model property* if every satisfiable sentence in $\mathcal{C}$ has a model $\mathcal{A}$ with a finite domain $\mathcal{U}^{\mathcal{A}}$.

We shall see that every fragment of FOL enjoying the finite model property has a decidable satisfiability problem.

Two exemplary fragments:

Bernays–Schönfinkel (BS): $\exists^*\forall^*$ prenex sentences without $\approx$ and without non-constant functions

monadic FO (MFO): all predicates are unary, no $\approx$, neither functions nor constants

We disallow equality only for simplicity and convenience. Moreover, constant symbols in MFO would not do any harm.

Finite and infinite models

Lemma 1.1 (Prop. 6.0.4 in [BGG97])

Let φ be an FO sentence in prenex form with k universal quantifiers and length n. Let m be some positive integer. Whether φ has a model with m domain elements can be decided nondeterministically in time poly$(m^k n)$.

Proof: Assume w.l.o.g. that φ is fully Skolemized, i.e. it is of the form $\forall \bar{x}.\, \psi(\bar{x})$ where ψ is quantifier free and $\bar{x}$ has length k. Consider the following nondeterministic procedure.

(1) Construct $\mathcal{A}$ with the domain $\mathcal{U}^{\mathcal{A}} := \{1, \ldots, m\}$. For every k-tuple $a_1, \ldots, a_k \in \mathcal{U}^k$ guess sufficient information regarding the interpretation of terms $t(\bar{x})$ and atoms $A(\bar{x})$ occurring in φ. (Notice that the truth value of, e.g., $P(1,2)$ under $\mathcal{A}$ need not be guessed, if P occurs in φ only in atoms $P(x,x)$, say.)

(2) Verify that $\mathcal{A} \models \varphi$.

Finite and infinite models

Theorem 1.2

Let $\mathcal{C}$ be any class of FO sentences. If $\mathcal{C}$ enjoys the finite model property, then we can decide satisfiability for all sentences in $\mathcal{C}$.

The proof is based on Lemma 1.1.

But why don't we need an upper bound on the size m of smallest models of a given sentence φ to invoke the Lemma?

Enumerating upper bounds $m = 1, 2, 3, \ldots$ only yields only a semi-decision procedure!

What is the missing piece?

$\rightsquigarrow$ We have refutationally complete calculi for FOL, e.g. superposition. That is, we have a semi-decision procedure for *un*satisfiability, which complements the above procedure.

Proving the finite model property for BS

Bernays–Schönfinkel fragment:

all $\exists^*\forall^*$ prenex FO sentences without non-constant functions and without $\approx$.

How can we show the finite model property for BS?

Let φ be a BS sentence and let ψ result from φ by exhaustive Skolemization. Then, φ and ψ are satisfiable over the same domains. By Herbrand's Theorem, ψ is satisfiable if and only if there is a Herbrand model for ψ. As the Herbrand domain for ψ, i.e. the domain of all terms built from ψ's signature, is finite, we are done.

Proving the finite model property for BS

An alternative proof requires the notion of *substructures*:

Definition (Substructure)

Let $\Sigma = (\Pi, \Omega)$ be an FO signature and let $\mathcal{A}, \mathcal{B}$ be Σ-structures. We call $\mathcal{B}$ a *substructure* of $\mathcal{A}$ if

(a) $\mathcal{U}^{\mathcal{B}} \subseteq \mathcal{U}^{\mathcal{A}}$,

(b) for every $P \in \Pi$ of arity m we have $P^{\mathcal{B}} = P^{\mathcal{A}} \cap (\mathcal{U}^{\mathcal{B}})^m$,

(c) for every $f \in \Pi$ of arity m and all $a_1, \ldots, a_m \in \mathcal{U}^{\mathcal{B}}$ we have $f^{\mathcal{B}}(a_1, \ldots, a_m) = f^{\mathcal{A}}(a_1, \ldots, a_m)$.

Lemma 1.3 (Substructure Lemma, Lemma III.5.7 in [EFT94])

Let φ be any prenex FO sentence without existential quantifiers. If $\mathcal{A}$ is a model of φ, then every substructure $\mathcal{B}$ of $\mathcal{A}$ is also a model of φ.

Substructure example

$\mathcal{U}^{\mathcal{A}}$:

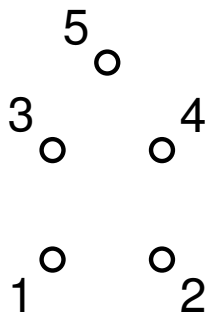

$E^{\mathcal{A}}$:

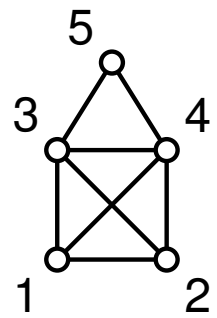

symmetric

$s^{\mathcal{A}}$:

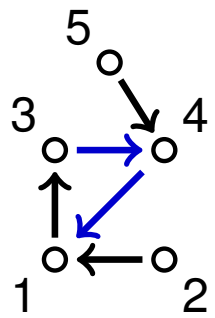

Changed for nicer sub-structures!

We observe $\mathcal{A} \models \forall x \exists y.\, E(x,y)$ $\quad \mathcal{A} \models \forall xy.\, E(x,y) \rightarrow E(y,x)$

$\mathcal{A} \models \exists z \forall x.\, s(x) \not\approx z$ $\quad \mathcal{A} \models \forall x.\, s(x) \not\approx x \wedge s(s(x)) \not\approx x$

Substructure example

$\mathcal{U}^{\mathcal{A}}$:

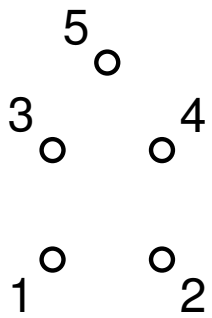

$E^{\mathcal{A}}$:

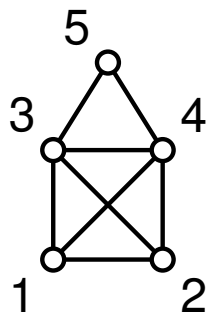

symmetric

$s^{\mathcal{A}}$:

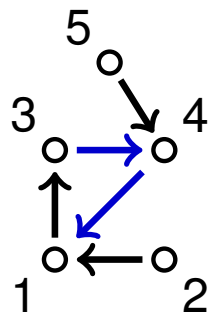

Changed for nicer sub-structures!

We observe $\mathcal{B} \models \forall x \exists y.\, E(x,y)$ $\quad \mathcal{B} \models \forall xy.\, E(x,y) \rightarrow E(y,x)$

$\mathcal{B} \not\models \exists z \forall x.\, s(x) \not\approx z$ $\quad \mathcal{B} \models \forall x.\, s(x) \not\approx x \wedge s(s(x)) \not\approx x$

$\mathcal{U}^{\mathcal{B}}$:

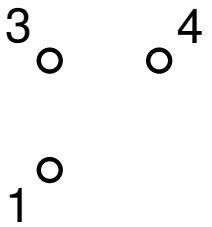

$E^{\mathcal{B}}$:

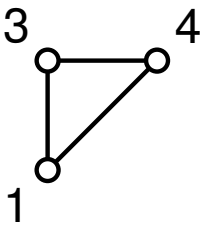

symmetric

$s^{\mathcal{B}}$:

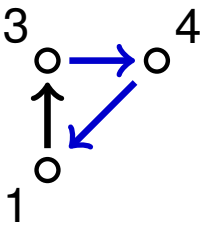

Substructure example

$\mathcal{U}^{\mathcal{A}}$:

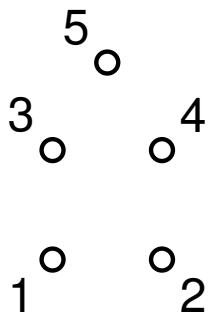

$E^{\mathcal{A}}$:

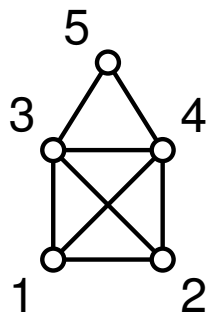

symmetric

$s^{\mathcal{A}}$:

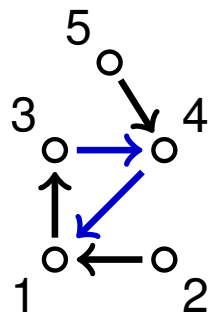

Changed for nicer sub-structures!

We observe $\mathcal{B} \models \forall x \exists y.\, E(x,y)$ $\quad \mathcal{B} \models \forall xy.\, E(x,y) \rightarrow E(y,x)$

$\mathcal{B} \not\models \exists z \forall x.\, s(x) \not\approx z$ $\quad \mathcal{B} \models \forall x.\, s(x) \not\approx x \wedge s(s(x)) \not\approx x$

Coinci-dence!

$\mathcal{U}^{\mathcal{B}}$:

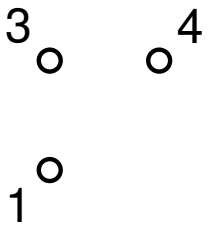

$E^{\mathcal{B}}$:

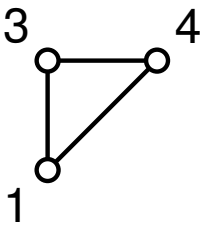

symmetric

$s^{\mathcal{B}}$:

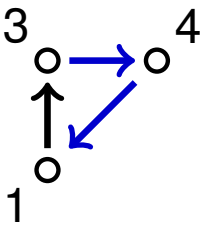

Proving the finite model property for BS (cont'd)

Bernays–Schönfinkel fragment:

$\exists^*\forall^*$ prenex sentences w/o non-constant functions and w/o $\approx$

Finite model property via Substructure Lemma:

Let φ be a BS sentence and let ψ result from φ by exhaustive Skolemization. Suppose ψ has a model $\mathcal{A}$ (over ψ's signature), possibly with infinite domain. Let $c_1, \ldots, c_k$ be the constants occurring in ψ. Consider the following structure $\mathcal{B}$ with

$\mathcal{U}^{\mathcal{B}} := \{c_1^{\mathcal{A}}, \ldots, c_k^{\mathcal{A}}\}$,
$P^{\mathcal{B}} := P^{\mathcal{A}} \cap (\mathcal{U}^{\mathcal{B}})^m$ for every m-ary predicate in ψ,
$c^{\mathcal{B}} := c^{\mathcal{A}}$ for every constant in ψ.

As $\mathcal{B}$ is a substructure of $\mathcal{A}$, the Substructure Lemma entails $\mathcal{B} \models \psi$, which entails $\mathcal{B} \models \varphi$.

Proving the finite model property for BS

Notice that the above proof also works in the presence of equality.

In fact, the Substructure Lemma entails a stronger result:

Lemma 1.4

Let $\varphi := \exists \bar{v} \forall \bar{x}.\, \psi$ be any BS sentence with quantifier-free ψ and k constant symbols. Suppose, there is a model $\mathcal{A} \models \varphi$. For any integer ℓ with $1 \leq k + |\bar{v}| \leq \ell \leq |\mathcal{U}^{\mathcal{A}}|$ there is a model $\mathcal{B}$ of φ with $|\mathcal{U}^{\mathcal{B}}| = \ell$. If $|\mathcal{U}^{\mathcal{A}}|$ is infinite, ℓ is not bounded from above.

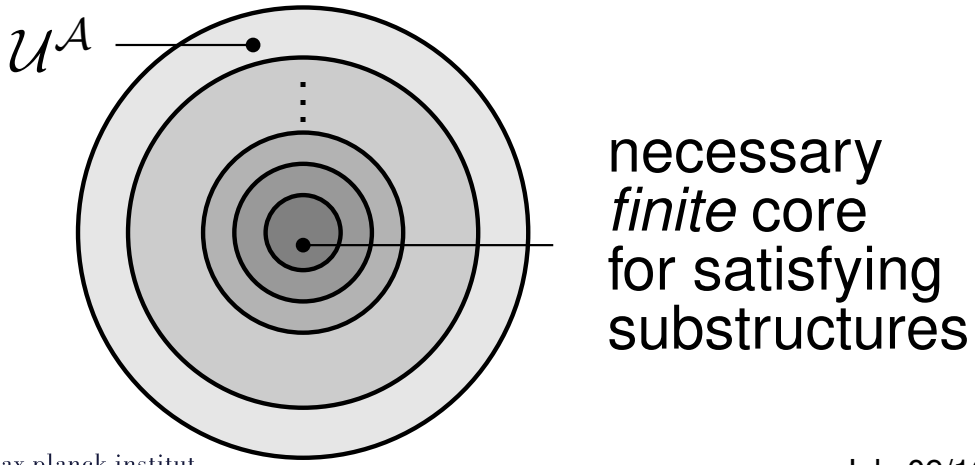

Proving the finite model property for BS

Theorem

The satisfiability problem for BS sentences is complete for NExpTime (nondet. exponential time).

Membership in NExpTime follows from Lemmas 1.1 and 1.4. NExpTime-hardness was shown by Lewis [Lew80], see also Theorem 6.2.21 in [BGG97].

Domain constraints in BS

BS sentences can impose lower bounds on the size of models:

$$\exists v_1, \ldots, v_k.\ \bigwedge_i \Big(P_i(v_i)\ \wedge\ \bigwedge_{j \neq i} \neg P_j(v_i) \wedge \neg P_i(v_j) \Big)$$

BS sentences cannot impose upper bounds! In fact, no satisfiable FOL sentence without equality can (see next slide).

For BS with equality, consider the following examples:

$$\forall xy.\ x \approx y$$

$$\exists v_1, \ldots, v_k \forall x.\ \bigvee_i x \approx v_i$$

$$\forall xy.\ \Big(\bigwedge_{1 \leq i \leq k} \big(P_i(x) \leftrightarrow P_i(y) \big) \Big) \rightarrow x \approx y$$

What are the imposed size bounds?

Domain constraints in FOL

Theorem (Upward Löwenheim-Skolem Thm. for FOL w/o $\approx$)

Let φ be any satisfiable FO sentence without equality and let $\mathcal{U}$ be any set. Then, there is a model $\mathcal{B} \models \varphi$ whose domain $\mathcal{U}^{\mathcal{B}}$ is a superset of $\mathcal{U}$.

Proof: Let $\mathcal{A}$ be a model of φ. Fix some element $a_0 \in \mathcal{U}^{\mathcal{A}}$. We define $\mathcal{B}$ such that $\mathcal{U}^{\mathcal{B}}$ is the disjoint union of $\mathcal{U}^{\mathcal{A}}$ and $\mathcal{U}$. Let τ be the mapping $\mathcal{U}^{\mathcal{B}} \to \mathcal{U}^{\mathcal{A}}$ with $\tau(a) = a$ for every $a \in \mathcal{U}^{\mathcal{A}}$ and $\tau(a) = a_0$ for every $a \in \mathcal{U}$. For every m-ary predicate P we set

$$P^{\mathcal{B}} := \{(a_1, \ldots, a_m) \in \mathcal{U}^{\mathcal{B}} \mid (\tau(a_1), \ldots, \tau(a_m)) \in P^{\mathcal{A}}\}.$$

For every m-ary function f and all $a_1, \ldots, a_m \in \mathcal{U}^{\mathcal{B}}$ we set

$$f^{\mathcal{B}}(a_1, \ldots, a_m) := f^{\mathcal{A}}(\tau(a_1), \ldots, \tau(a_m)).$$

It is not hard to show that $\mathcal{B} \models \varphi$ follows from $\mathcal{A} \models \varphi$. (Exercise!)

Domain constraints in FOL

Theorem (Löwenheim-Skolem Thm., from finite to infinite)

Let Φ be a set of FO sentences (with $\approx$) such that for every integer n there exists a finite model $\mathcal{A}_n \models \Phi$ with at least n domain elements. Then, Φ has an infinite model.

Proof: For every positive n let ψ_n be a satisfiable FO sentence whose models all have size $\geq n$ (see previous slides for an example). Consider the formula sets $\Phi_m := \Phi \cup \{\psi_n \mid 2 \leq n \leq m\}$ and $\Phi' := \bigcup_{m \geq 2} \Phi_m$. Since Φ is satisfiable over arbitrarily large finite structures, each set Φ_m is satisfiable, too. Since each finite subset of Φ' is contained in some Φ_m, *compactness of FOL* entails that Φ' is satisfiable as well. For any model $\mathcal{B} \models \Phi'$ we get $\mathcal{B} \models \Phi$ and $\mathcal{B} \models \{\psi_n \mid n \geq 2\}$. Due to the latter, $\mathcal{B}$'s domain $\mathcal{U}^{\mathcal{B}}$ must be infinite.

Domain constraints in FOL

The above theorem has interesting consequences. For instance, it entails some limitations regarding the expressiveness of FOL.

Proposition (FOL cannot express finiteness)

There is no signature Σ and Σ-sentence φ such that for all Σ-structures $\mathcal{A}$ we have $\mathcal{A} \models \varphi$ if and only if $\mathcal{U}^{\mathcal{A}}$ is finite.

Proof: Exercise!

FOL cannot control infinite domains

Theorem (Löwenheim-Skolem Thm., from infinite to larger)

Let φ be any FO sentence (with $\approx$) that is satisfied by some structure with an infinite domain. Let $\mathcal{U}$ be any set. Then, there is some model $\mathcal{A} \models \varphi$ whose domain is a superset of $\mathcal{U}$.

Theorem (Downward Löwenheim-Skolem Theorem)

Consider any signature $\Sigma = (\Pi, \Omega)$ with countable Π and Ω.
(i) Every satisfiable set of Σ-sentences without equality has an infinite countable model.
(ii) Every satisfiable set of Σ-sentences with equality has a (finite or infinite) countable model.

For proofs, see [EFT94], Chapter VI, or [End01], Section 2.6, or [Hod97], Corollaries 3.1.4 and 5.1.4.

FOL cannot control infinite domains

Due to the Löwenheim–Skolem Theorems, it becomes clear that first-order logic is not expressive enough to characterize the natural numbers, the integers, the rationals, or the reals.

Proposition

There is no countable first-order signature Σ and no set Φ of Σ-sentences such that all models of Φ are isomorphic to $\mathbb{N}$. The same holds for $\mathbb{Z}, \mathbb{Q}, \mathbb{R}$.

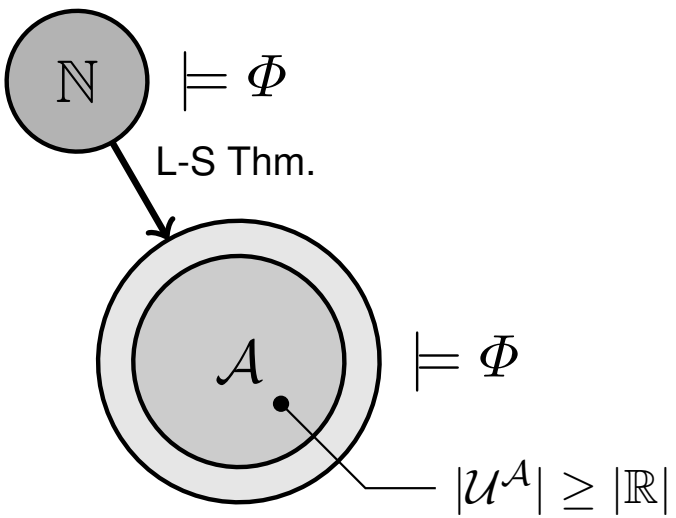

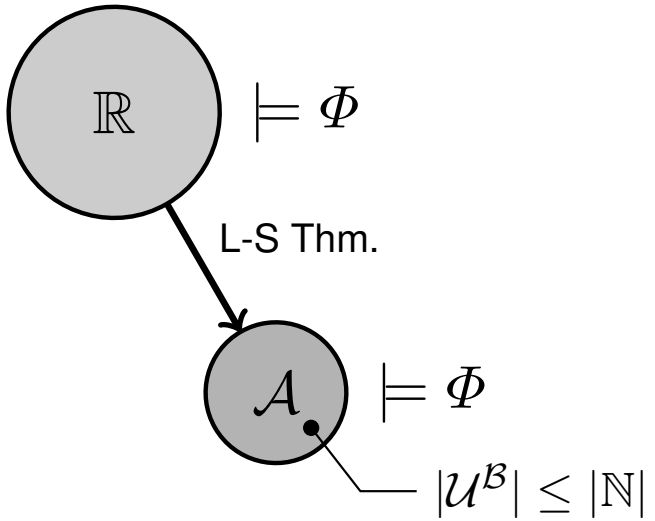

Proving the finite model property for MFO

Monadic first-order fragment (MFO):
only unary predicates, no $\approx$, no non-constant functions

Consider any satisfiable MFO sentence φ with k distinct unary predicates $P_1, \ldots, P_k$. Let $\mathcal{A} \models \varphi$.

How can φ distinguish two domain elements $a, b \in \mathcal{U}^{\mathcal{A}}$?
Only by some P_i such that $\mathcal{A} \models P_i(a)$ and $\mathcal{A} \not\models P_i(b)$ or vice versa.

Set $a \sim b$ iff $\mathcal{A} \models P_i(a) \leftrightarrow P_i(b)$ for all i. Define $\mathcal{U}^{\mathcal{B}} := \mathcal{U}^{\mathcal{A}}/_{\sim}$.

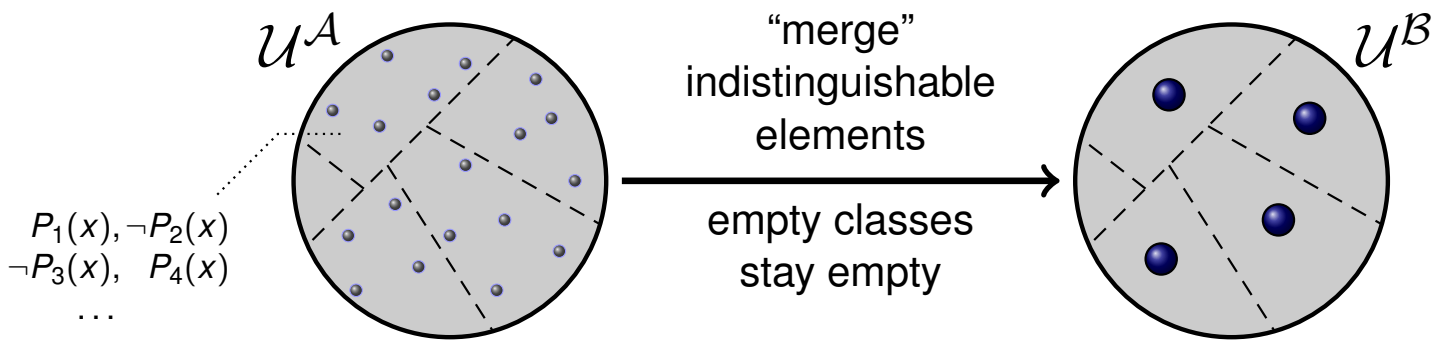

Prove $\mathcal{B} \models \varphi$. (Exercise!) $\mathcal{U}^{\mathcal{B}}$ contains at most 2^k elements.

Proving the finite model property for MFO

Lemma 1.5 (Finite models for MFO with $\approx$ and constants)

Let φ be a satisfiable FO sentence with k predicates, all unary, m constants, and ℓ quantifiers. There is a model $\mathcal{A} \models \varphi$ with at most $(m + \ell) \cdot 2^k$ domain elements.

Proof: Since we allow equality, it is in general not sufficient to keep only one representative for every equivalence class in $\mathcal{U}^{\mathcal{A}}/_{\sim}$. For each such class we pick $(m + \ell)$ distinct elements (if available in $\mathcal{A}$; otherwise we select all the available ones) and put them into $\mathcal{U}^{\mathcal{B}}$. Their membership w.r.t. $P_i^{\mathcal{B}}$ is defined like in $\mathcal{A}$. After defining the constants $c^{\mathcal{B}}$ appropriately, $\mathcal{B} \models \varphi$ follows. (Exercise!)

Theorem

Satisfiability for MFO sentences is NExpTime-complete.

Membership: L 1.1 and 1.5. Hardness: see Thm 6.2.13 in [BGG97].

References

[BGG97] Börger, Grädel, Gurevich. *The Classical Decision Problem.* Springer, 1997

[EFT94] Ebbinghaus, Flum, Thomas. *Mathematical Logic.* Second edition. Springer, 1994

[End01] Enderton. *A Mathematical Introduction to Logic.* Second edition. Academic Press, 2001

[Hod97] Hodges. *A Shorter Model Theory.* Cambridge University Press, 1997

[Lew80] Lewis. *Complexity Results for Classes of Quantificational Formulas.* J. of Computer and System Sciences 21(3), 1980

Acknowledgment:

Many of the shown proofs follow the outline of similar proofs from the course *Advanced Logics*, held by Christel Baier at Dresden University in summer 2011. For any typos and errors in the present slides solely the author of the slides is to be blamed.

Why all this??

Next week, we shall consider decidable BS(LRA) fragments extending BS with Simple Bounds.

In order to show decidability, we will re-use some of the methods we have seen today.

For instance, we will identify a finite set of equivalence classes of tuples of reals that are indistinguishable by the available arithmetic atoms.

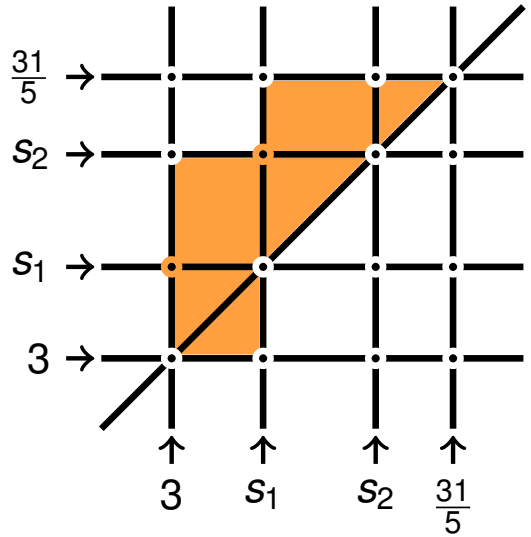

2. Decidable Fragments of Bernays–Schönfinkel Modulo Linear Rational Arithmetic

We have already seen *BS clauses with simple bounds*:

$$\forall xy.\ \underbrace{0 \leq x \wedge x < 5 \wedge y \neq 3}_{\text{LRA atoms } x \,\#\, k \text{ with } k \in \mathbb{Z} \text{ and } \# \in \{<, \leq, =, \neq, \geq, >\}} \ \|\ \underbrace{P(x,y) \wedge Q(x) \rightarrow P(y,x)}_{\textit{free} \text{ atoms in clauses (here written as implications)}}$$

"$\|$" can be read as $\wedge$, i.e. $(\Lambda \wedge \Gamma) \rightarrow \Delta$

Now: *BS with simple linear rational constraints BS(SLR)*:

$$\exists cd\, \forall xy.\ c \neq d \wedge x > c + 2d - 3 \wedge x < y \ \|\ Q(x,y) \rightarrow T(x) \vee Q(y,x)$$

We allow LRA atoms $s \ \# \ t$ where

- $\# \in \{<, \leq, =, \neq, \geq, >\}$,
- LRA terms with operators $+$, $-$, etc. don't contain univ. variables,
- univ. variables may be compared to univ. variables, i.e. $x \ \# \ y$.

BS(SLR) normal form

Definition (BS(SLR) normal form)

A BS(SLR) clause $\Lambda \| \Gamma \to \Delta$ is in *normal form* if every univ. variable in Λ also occurs in Γ or in Δ.

A BS(SLR) clause set N is in *normal form* if

(a) All clauses in N are in normal form and pairwise variable disjoint.

(b) N can be divided into two parts $N_{\mathbb{Q}}$ and N_{BS} such that

(b1) every clause in $N_{\mathbb{Q}}$ is a unit clause containing exactly one positive LRA literal, and

(b2) for every clause $\Lambda \| \Gamma \to \Delta$ in N_{BS} either Γ or Δ is nonempty and any LRA atom $s \# t$ in Λ is such that s and t are either background-sort variables or Skolem constants, respectively.

We assume that N_{BS} contains at least one free-sort constant.

BS(SLR) normal form

Lemma (BS(SLR) normal form)

For every BS(SLR) clause set N there is an equisatisfiable BS(SLR) clause set in BS(SLR) normal form.

Example:

$$N = \{\ c \neq d \ \wedge\ x > c + 2d - 3 \ \wedge\ x < y \parallel\ Q(x, y) \rightarrow\ T(x) \ \vee\ Q(y, x)\ \}$$

Equisatisfiable BS(SLR) normal form:

$N' = N'_{\mathbb{Q}} \cup N'_{\mathrm{BS}}$ with

$N'_{\mathbb{Q}} = \{\ e = c + 2d - 3\ \}$ for a fresh Skolem constant e and

$N'_{\mathrm{BS}} = \{\ c \neq d \ \wedge\ x > e \ \wedge\ x < y \parallel Q(x, y) \rightarrow\ T(x) \ \vee\ Q(y, x)\ \}$

$\rightsquigarrow$ Similar to abstraction in SMT and Nelson–Oppen context.

BS(SLR) normal form

Lemma

For every finite BS(SLR) clause set there is an equisatisfiable finite BS(SLR) clause set that is in BS(SLR) normal form.

The normal form can be established using a quantifier-elimination procedure for LRA (e.g. Fourier–Motzkin or Loos–Weispfenning) and methods similar to abstraction techniques used in the context of combinations of theories.

For the next couple of slides, we shall concentrate on the N_{BS}-parts of BS(SLR) clause sets. That is, complex LRA terms are ignored.

Expressiveness of SLR constraints

Which pairs can be distinguished by the following constraints?

- $x \# 3$
- $y \# \frac{31}{5}$
- $x \# y$
- $x \# c$
- $y \# d$

$\rightsquigarrow$ Depends on how c, d are interpreted!

Under $\mathcal{A}$ with $3 < c^{\mathcal{A}} < d^{\mathcal{A}} < \frac{31}{5}$

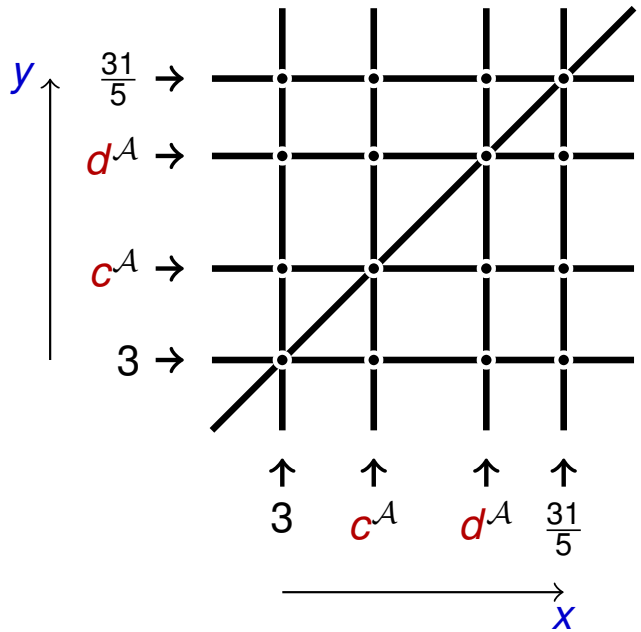

$\rightsquigarrow$ Two points on the same line segment / within the same white area cannot be distinguished.

Expressiveness of SLR constraints

From now on all structures interpret LRA symbols and terms in the standard LRA semantics.

Definition ($\mathcal{A}$-induced partition of $\mathbb{Q}$)

Consider any finite BS(SLR) clause set N and any structure $\mathcal{A}$. Let $r_1 < \ldots < r_k$ be the values of all distinct rationals assigned to LRA constants from N under $\mathcal{A}$. By $\mathcal{J}_\mathcal{A}$ we denote the following partition of $\mathbb{Q}$ into $2k+1$ intervals:

$$\{(-\infty, r_1), \{r_1\}, (r_1, r_2), \{r_2\}, \ldots, \{r_k\}, (r_k, +\infty)\}.$$

Illustration for $k = 5$ rational values:

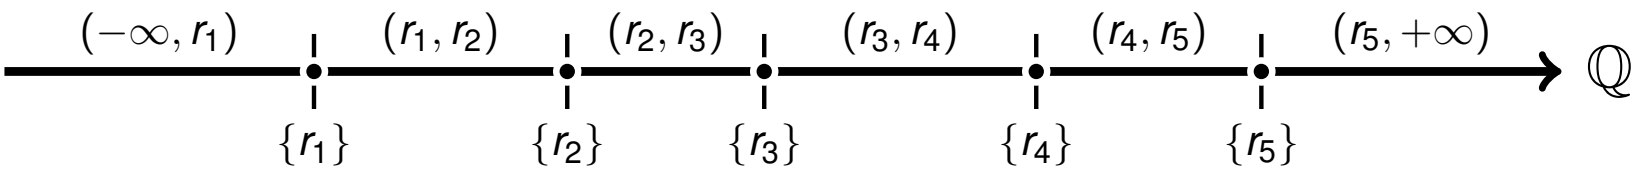

Expressiveness of SLR constraints

Definition ($\mathcal{J}_\mathcal{A}$-equivalence, $\sim_{\mathcal{J}_\mathcal{A}}$)

Let $\mathcal{A}$ be any structure and let m be any positive integer. Two tuples $\bar{r}, \bar{q} \in \mathbb{Q}^m$ are called *$\mathcal{J}_\mathcal{A}$-equivalent*, denoted $\bar{r} \sim_{\mathcal{J}_\mathcal{A}} \bar{q}$, if

(i) for every index i and every interval $J \in \mathcal{J}_\mathcal{A}$ we have $r_i \in J$ iff $q_i \in J$, and

(ii) for all i, j we have $r_i < r_j$ iff $q_i < q_j$.

The induced equivalence relation over $\mathbb{Q}$-tuples is $\sim_{\mathcal{J}_\mathcal{A}}$.

For $m = 2$ and $r_1 < \ldots < r_4$ we get the following equiv. classes of $\mathbb{Q}^2$:

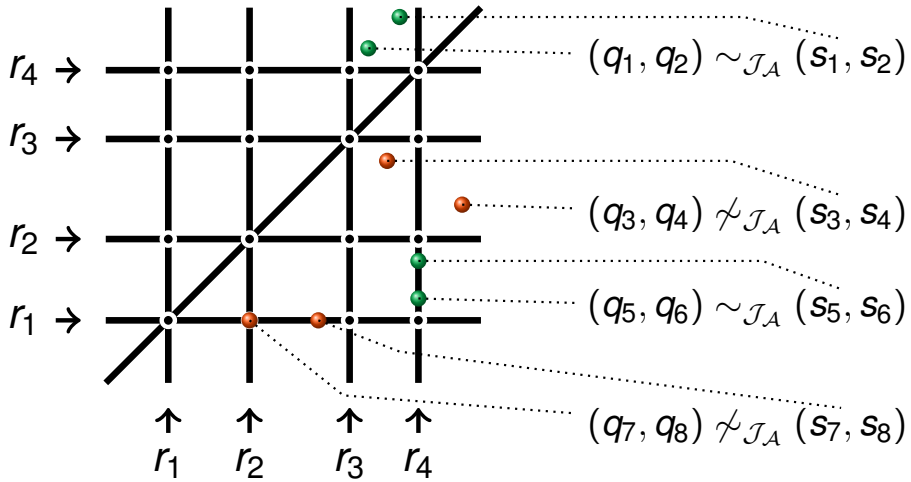

Expressiveness of SLR constraints

Proposition

Let $\Lambda(\bar{x})$ be any conjunction of SLR atoms with variables from $\bar{x}$. Let $\mathcal{A}$ be any structure. For any two $\mathbb{Q}$-tuples $\bar{r}, \bar{q}$ we observe that $\bar{r} \sim_{\mathcal{J}_\mathcal{A}} \bar{q}$ entails

$$\mathcal{A} \models \Lambda(\bar{r}) \quad \text{if and only if} \quad \mathcal{A} \models \Lambda(\bar{q}).$$

In other words: SLR constraints cannot distinguish $\mathcal{J}_\mathcal{A}$-equivalent $\mathbb{Q}$-tuples.

$\mathcal{J}_{\mathcal{A}}$-uniformity

For what follows we fix a finite BS(SLR) clause set N and two nonnegative integers ℓ, m and assume that all predicates P in N have the sort $P : \mathcal{S}^{\ell} \times \mathbb{Q}^{m}$, where $\mathcal{S}$ denotes the (single) free sort occurring in N.

Definition ($\mathcal{A}$-colors, $\mathcal{A}$-colorings)

Let $e_1, \ldots, e_n$ be all free-sort constants occurring in N. Let $\mathcal{A}$ be a structure and set $\widehat{\mathcal{S}} := \left\{ \mathsf{a} \in \mathcal{S}^{\mathcal{A}} \;\middle|\; \mathsf{a} = e_i^{\mathcal{A}} \text{ for some } e_i \right\}$.
An *$\mathcal{A}$-color* is any set of expressions $P\bar{\mathsf{a}}$ where P is some predicate from N and $\bar{\mathsf{a}} \in \widehat{\mathcal{S}}^{\ell}$.
An *$\mathcal{A}$-coloring* of $\mathbb{Q}^m$ is a total mapping $\chi_{\mathcal{A}}$ assigning $\mathcal{A}$-colors to $\mathbb{Q}$-tuples such that for each $\bar{r} \in \mathbb{Q}^m$ we have

$$P\bar{\mathsf{a}} \in \chi_{\mathcal{A}} \quad \text{if and only if} \quad \mathcal{A} \models P(\bar{\mathsf{a}}, \bar{r}), \quad \text{i.e. } (\bar{\mathsf{a}}, \bar{r}) \in P^{\mathcal{A}}.$$

$\rightsquigarrow$ Since N and $\widehat{\mathcal{S}}$ are finite, there are only finitely many $\mathcal{A}$-colors.

$\mathcal{J}_\mathcal{A}$-uniformity

Definition ($\mathcal{J}_\mathcal{A}$-uniformity)

A structure $\mathcal{A}$ is *$\mathcal{J}_\mathcal{A}$-uniform* if $\chi_\mathcal{A}$ colors each and every $\sim_{\mathcal{J}_\mathcal{A}}$-equivalence class uniformly, i.e. for all pairs $\bar{r} \sim_{\mathcal{J}_\mathcal{A}} \bar{q}$ we have $\chi_\mathcal{A}(\bar{r}) = \chi_\mathcal{A}(\bar{q})$.
More precisely, for all pairs $\bar{r} \sim_{\mathcal{J}_\mathcal{A}} \bar{q}$, all predicates P in N and all free-sort tuples $\bar{a} \in (S^\mathcal{A})^\ell$ we have

$$\mathcal{A} \models P(\bar{a}, \bar{r}) \quad \text{if and only if} \quad \mathcal{A} \models P(\bar{a}, \bar{q}).$$

For all P, all $\bar{a}$ and all $\mathcal{J}_\mathcal{A}$-equivalence classes C, we either have

$$\mathcal{A} \models P(\bar{a}, \bar{r})$$

for all $\bar{r} \in C$ or for none.

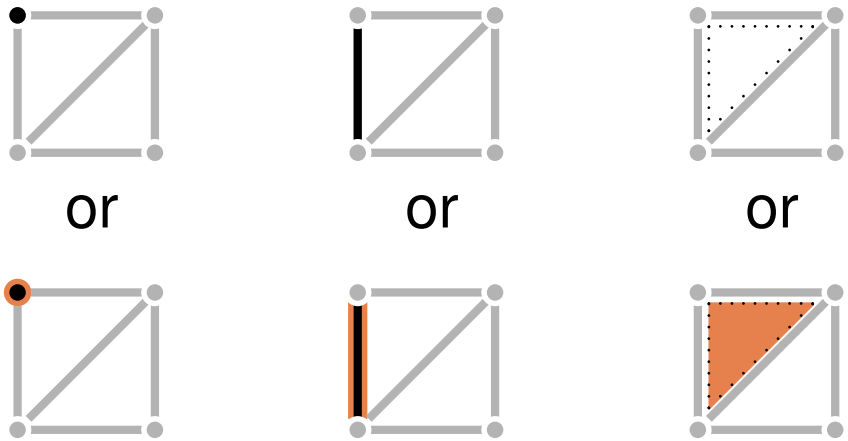

$\mathcal{J}_\mathcal{A}$-uniformity

Theorem 2.1 (Existence of uniform models)

Any satisfiable finite BS(SLR) clause set has a model $\mathcal{A}$ that is $\mathcal{J}_\mathcal{A}$-uniform and that interprets the free sort $\mathcal{S}$ with a finite domain.

Proof: Central proof ideas will be outlined in what follows. □

Since $\sim_{\mathcal{J}_\mathcal{A}}$ induces only finitely many equivalence classes, each describable by finite means, any $\mathcal{J}_\mathcal{A}$-uniform model $\mathcal{A}$ can be described by finite means, too.

This property is akin to the finite model property (cf. last lecture) and leads to a decidable satisfiability problem!

Constructing $\mathcal{J}_{\mathcal{A}}$-Uniform Structures

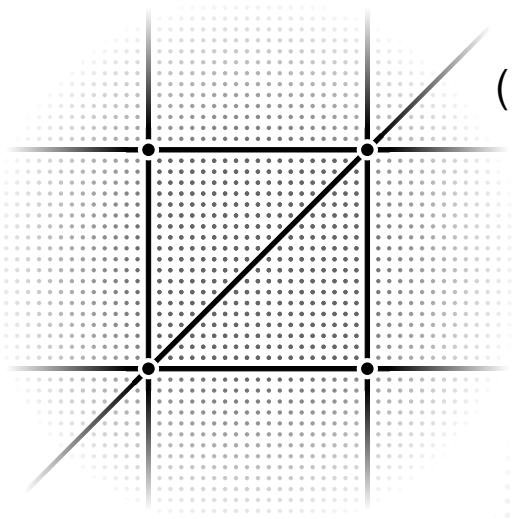

(1) Divide $\mathbb{Q}^m$ into finitely many equiv. classes

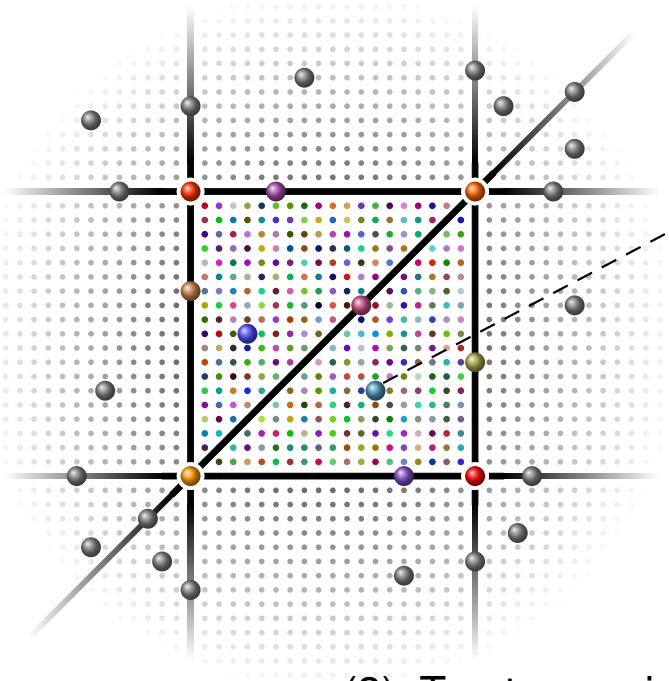

$$P_1(\bar{a}, x, y), \quad \neg P_2(\bar{a}, x, y),$$
$$\dots$$
$$\neg P_1(\bar{b}, x, y), \quad P_2(\bar{b}, x, y),$$
$$\dots$$

(2) Pick a ***suitable*** representative from every equiv. class

(3) Treat equivalence classes uniformly under $\mathcal{A}$, based on representatives

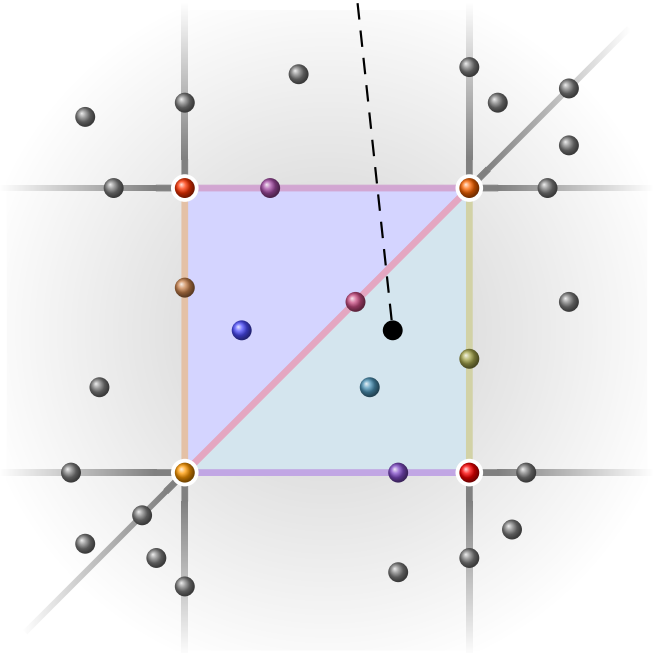

Constructing $\mathcal{J}_\mathcal{A}$-Uniform Structures

What are **suitable** representatives?

Recall that we have fixed a clause set N whose predicates all have the sort $\mathcal{S}^\ell \times \mathbb{Q}^m$. Suppose N is satisfiable.
Consider any model $\mathcal{A} \models N$ and let $J_0, \{r_1\}, J_1, \ldots, \{r_k\}, J_k$ be all intervals in $\mathcal{J}_\mathcal{A}$ in ascending order.

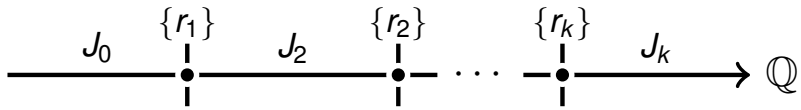

Let λ be the max. number of distinct univ. quantified variables in any clause in N. (In case of $\lambda < m$, set $\lambda = m$).

We need a collection of finite subsets $Q_i \subseteq J_i$ with $|Q_i| \geq \lambda$ such that for all $\mathcal{J}_\mathcal{A}$-equivalent $\bar{r}, \bar{q} \in Q^m$, where $Q := \bigcup_i Q_i \cup \bigcup_j \{r_j\}$, we have $\chi_\mathcal{A}(\bar{r}) = \chi_\mathcal{A}(\bar{q})$.

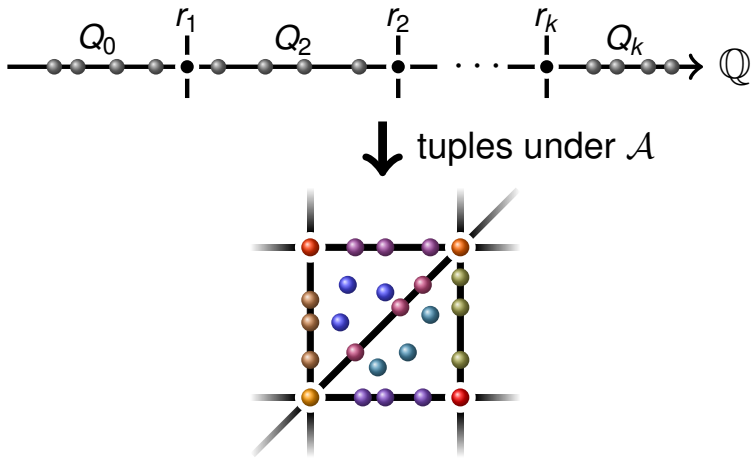

Constructing $\mathcal{J}_{\mathcal{A}}$-Uniform Structures

Lemma

Such finite sets $Q_i \subseteq J_i$ always exist for any model $\mathcal{A} \models N$.

The proof of this result is based on methods from Ramsey theory [GRS90]. Details are available in [Voi17].

Lemma 2.2

Let N, $\mathcal{A}$, λ, $r_1, \ldots, r_k$, and $Q_1, \ldots, Q_k$ be defined as above. We can construct a model $\mathcal{B} \models N$ that is $\mathcal{J}_{\mathcal{B}}$-uniform and that interprets $\mathcal{S}$ with a finite set.

Proof:

Claim I: Let μ be any integer $1 \leq \mu \leq \lambda$. For each of the equiv. class in $\mathbb{Q}^{\mu}/_{\sim_{\mathcal{J}_{\mathcal{A}}}}$ we find one representative lying in Q^{μ}, where $Q := \bigcup_i Q_i \cup \bigcup_j \{r_j\}$. $\Diamond$

Constructing $\mathcal{J}_{\mathcal{A}}$-Uniform Structures

Proof of Lemma 2.2 (continued):

Let $\widehat{\mathcal{S}} := \left\{ a \in \mathcal{S}^{\mathcal{A}} \;\middle|\; a = e^{\mathcal{A}} \text{ for some free-sort constant } e \text{ in } N \right\}$.

We construct $\mathcal{B}$ as follows. $\mathcal{S}^{\mathcal{B}} := \widehat{\mathcal{S}}$ and $e^{\mathcal{B}} := e^{\mathcal{A}}$ for every free-sort constant. (Notice that this entails $\sim_{\mathcal{J}_{\mathcal{B}}} = \sim_{\mathcal{J}_{\mathcal{A}}}$.)
For all predicates $P : \mathcal{S}^{\ell} \times \mathbb{Q}^m$ in N and for all $\bar{a} \in (\mathcal{S}^{\mathcal{B}})^m$ and all $\bar{r} \in \mathbb{Q}^m$ we pick some $\bar{q} \in Q^m$ with $\bar{q} \sim_{\mathcal{J}_{\mathcal{A}}} \bar{r}$ (exists by Claim I !) and define $P^{\mathcal{B}}$ so that

$$(\bar{a}, \bar{r}) \in P^{\mathcal{B}} \quad \text{if and only if} \quad (\bar{a}, \bar{q}) \in P^{\mathcal{A}}.$$

Claim II: The structure $\mathcal{B}$ is $\mathcal{J}_{\mathcal{B}}$-uniform. (Holds by construction of $\mathcal{B}$ and the properties of Q). $\Diamond$

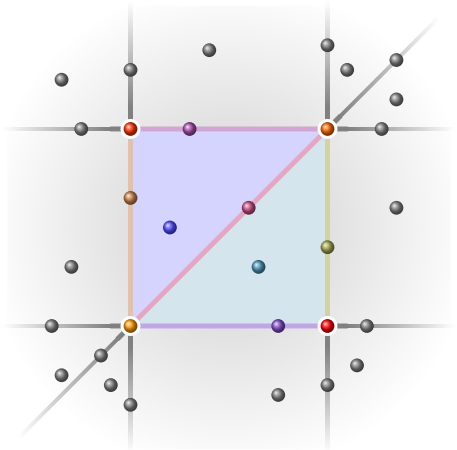

Constructing $\mathcal{J}_{\mathcal{A}}$-Uniform Structures

Proof of Lemma 2.2 (continued):

It remains to show $\mathcal{B} \models N$.

Consider any clause $C = \Lambda \| \Gamma \to \Delta$ from N with variables $x_1, \ldots, x_\mu$ of sort $\mathbb{Q}$. Let β be any (sort-respecting) variable assignment over $\mathcal{S}^{\mathcal{B}} \cup \mathbb{Q}$. By Claim I, there is some variable assignment γ that coincides with β on all free-sort variables and for which

$$\big(\gamma(x_1), \ldots, \gamma(x_\mu)\big) \sim_{\mathcal{J}_{\mathcal{B}}} \big(\beta(x_1), \ldots, \beta(x_\mu)\big).$$

As $\mathcal{A} \models N$, we get $\mathcal{A}, \gamma \models C$. We can show $\mathcal{B}, \beta \models C$ by case distinction on why $\mathcal{A}, \gamma \models C$ holds. (Exercise!)

This finally entails $\mathcal{B} \models N$. □

Existence of $\mathcal{J}_{\mathcal{A}}$-Uniform Structures

Theorem 2.1 is a corollary of Lemma 2.2.

Theorem 2.1 (Existence of uniform models – Restated)

Any satisfiable finite BS(SLR) clause set has a model $\mathcal{A}$ that is $\mathcal{J}_{\mathcal{A}}$-uniform and that interprets the free sort $\mathcal{S}$ with a finite domain.

Intuitively, the theorem holds due to the following observation:

> Uninterpreted predicates in BS(SLR) clause sets need not distinguish what arithmetic constraints cannot distinguish.

Recall that $\sim_{\mathcal{J}_{\mathcal{A}}}$ induces only finitely many equivalence classes over $\mathbb{Q}^m$. Therefore, Theorem 2.1 constitutes a property similar to the finite model property.

$\rightsquigarrow$ We shall exploit this for deciding satisfiability.

Reduction from BS(SLR) to two-sorted BS

Consider a finite clause set $N_{\mathbb{Q}} \cup N_{\mathrm{BS}}$ in BS(SLR) normal form.

Recall that

(a) $N_{\mathbb{Q}}$ contains only pure LRA clauses with additional Skolem constants of sort $\mathbb{Q}$ and

(b) N_{BS} contains BS(SLR) clauses but only LRA atoms $s \# t$ where s, t are either variables or Skolem constants of sort $\mathbb{Q}$.

Define the clause set $N_{<,\leq}$ stipulating

(1) irreflexivity, transitivity, and totality of $<$, and

(2) $\forall xy.\, x \leq y \;\leftrightarrow\; x \approx y \vee x < y$.

Reduction from BS(SLR) to two-sorted BS

Let $c_1, \ldots, c_k$ be all constants of sort $\mathbb{Q}$ in N_{BS} and let λ be the max. number of variables in any clause in N_{BS}.

Given any total preorder $\preceq$ (a reflexive and transitive relation) on the $c_1, \ldots, c_k$, we define the following clause set $N_{\preceq}$.
Suppose that $c_{j_1} \prec \ldots \prec c_{j_{k'}}$ is a max. $\prec$-chain, where $a \prec b$ means $a \preceq b$ and $a \not\approx b$.
We define $N_{\preceq}$ so that it stipulates

$$d_{0,1} < \ldots < d_{0,\lambda} < c_{j_1} < d_{1,1} < \ldots < d_{1,\lambda} < c_{j_2} < \ldots$$
$$< c_{k'-1} < d_{k'-1,1} < \ldots < d_{k'-1,\lambda} < c_{j_{k'}} < d_{k',1} < \ldots < d_{k',\lambda}$$

for freshly added constants $d_{i,j}$ and

$$c_j \approx c_{j'}$$

whenever $c_j \preceq c_{j'}$ and $c_{j'} \preceq c_j$.

Decision Procedure for finite BS(SLR) clause sets

Recap: We have $N = N_{\mathbb{Q}} \cup N_{\mathrm{BS}}$ in BS(SLR) normal form,

$N_{<,\leq}$ defining the semantics of $<, \leq$, and

$N_{\preceq}$ aligning the c_i and additional constants $d_{i,j}$ in accordance with $\preceq$.

Decision procedure for N:

(1) Nondeterministically fix a preorder $\preceq$ on $\mathbb{Q}$-sort Skolem constants $c_1, \ldots, c_k$ in N_{BS}.

(2) Check whether there is an assignment $\gamma : \{c_1, \ldots, c_k\} \to \mathbb{Q}$ s.t.
$$\mathbb{Q}, \gamma \models N_{\mathbb{Q}} \cup \{c_i \leq c_j \mid c_i \preceq c_j\}.$$

(3) Check whether the two-sorted BS clause set $N_{\mathrm{BS}} \cup N_{<,\leq} \cup N_{\preceq}$ is satisfied by some "$\mathcal{J}_{\mathcal{A}}$-uniform" structure $\mathcal{A}$, pretending $<, \leq$ are free predicates.

(4) If both checks succeed, then the BS(SLR) clause set N satisfiable.

Decision Procedure for finite BS(SLR) clause sets

Theorem 2.3

Satisfiability problem for finite BS(SLR) clause sets is NExpTime-complete.

NExpTime-hardness follows from the sat. problem for BS being NExpTime-hard. Proving membership in NExpTime requires appropriate upper bounds regarding

(a) the blowup for the BS(SLR)-normal-form transformation,
(b) the dec. procedure for LRA,
(c) the length of the clause sets $N_{<,\leq}$ and $N_{\preceq}$.

Having all these, we employ Lemma 1.1 from last lecture.

Decision Procedure for finite BS(SLR) clause sets

Reconsider Step (3) of the decision procedure:

> *Check whether $N_{\mathsf{BS}} \cup N_{<,\leq} \cup N_{\preceq}$ is satisfied by some "$\mathcal{J}_{\mathcal{A}}$-uniform" structure $\mathcal{A}$.*

We need to require $\mathcal{J}_{\mathcal{A}}$-uniformity of $\mathcal{A}$ due to upper bounds on the model size that are potentially imposed by N_{BS}, if equality is allowed (cf. last lecture).

Exercise: Construct a finite BS clause set N using FO equality $\approx$ and $<$ (treated as a free predicate of sort $< : \mathcal{S} \times \mathcal{S}$) such that $N \cup N_{<,\leq}$ is satisfied by a finite model but not when the domain $\mathbb{Q}$ is used and $<$ is interpreted as usual over $\mathbb{Q}$.

Restricting SLR Constrains Further

Exercise:

What would change if we disallow SLR constraints of the form $x \neq y$ and $x < y$ for universally quantified variables x, y?

Things get simpler and less complex. We need simpler kinds of representatives / fewer of them. Rather simple instantiation methods suffice for decision procedures, e.g. based on superposition modulo LRA.

Any ideas about the details?

BS(SLR) as a combination of theories

Once again, consider $N_{\mathbb{Q}} \cup N_{\mathrm{BS}}$ in BS(SLR) normal form:

(a) $N_{\mathbb{Q}}$ contains only pure LRA clauses with additional Skolem constants of sort $\mathbb{Q}$ and

(b) N_{BS} contains BS(SLR) clauses but only LRA atoms $s \# t$ where s, t are either variables or Skolem constants of sort $\mathbb{Q}$.

Alternative view of BS(SLR):

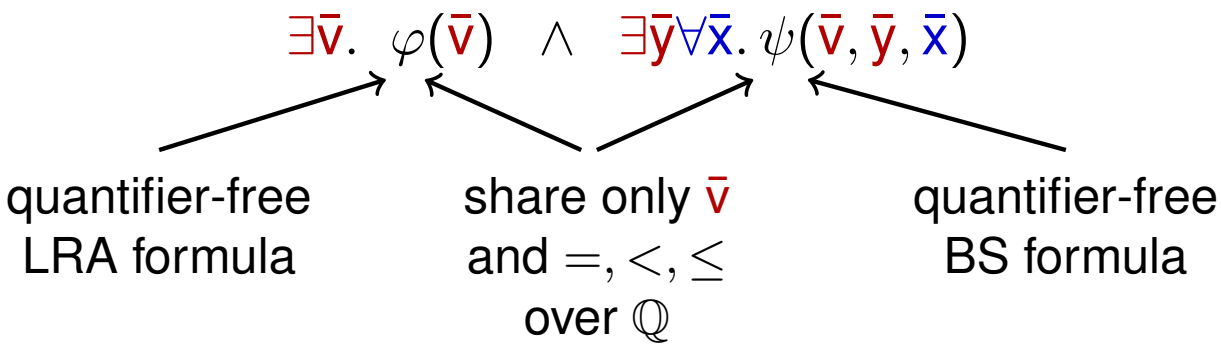

$\rightsquigarrow$ A solver for $\varphi(\bar{v})$ plus a dec. procedure for $\exists \bar{y} \forall \bar{x}.\ \psi(\bar{r}, \bar{y}, \bar{x})$ can be combined into a dec. procedure for BS(SLR).

Extending BS(SLR)

SF(polyR): $\exists \bar{v}.\ \varphi(\bar{v}) \wedge \psi(\bar{v})$

- $\varphi(\bar{v})$: arithmetic formula with quantifiers, with polynomials over $\mathbb{R}$, e.g. $x^3y + 3x^2y^2$
- $\wedge$: share only $\bar{v}$ and $=, <, \leq$ over $\mathbb{R}$
- $\psi(\bar{v})$: relational formula, equivalent to $\exists \bar{y}' \forall \bar{x}'.\ \psi'(\bar{v}, \bar{y}', \bar{x}')$

$\rightsquigarrow$ We only need two additional components:

(1) solver for $\varphi(\bar{v})$ proposing an *arrangement* $\bigwedge_{i,j} [\neg] v_i < v_j \wedge [\neg] v_i = v_j$

(2) procedure transforming $\psi(\bar{v})$ into $\exists \bar{y}' \forall \bar{x}'.\ \psi'(\bar{v}, \bar{y}', \bar{x}')$

$\rightsquigarrow$ (1) is available for arithmetic with polynomials over $\mathbb{R}$ (virt. substs.),

$\rightsquigarrow$ (2) is available for the *separated fragment (SF)* [SVW16]:

$\exists \bar{z}\, \forall \bar{x}_1 \exists \bar{y}_1 \ldots \forall \bar{x}_n \exists \bar{y}_n.\ \psi$, no atom contains blue *and* red variables

$\rightsquigarrow$ SF subsumes BS and MFO

We have seen:

BS clauses with simple bounds:

$$\forall xy.\, 0 \leq x \wedge x < 5 \wedge y \neq 3 \parallel P(x,y) \wedge Q(x) \rightarrow P(y,x)$$

BS with simple linear rational constraints BS(SLR):

$$\exists cd\, \forall xy.\ c \neq d \wedge x > c + 2d - 3 \wedge x < y \parallel Q(x,y) \rightarrow T(x) \vee Q(y,x)$$

Now: *BS with bounded difference constraints BS(BD)*:

$$\forall xyz.\ x - y < 3 \wedge -2 \leq x, y \leq 2 \wedge x < z \wedge Q(x,y) \rightarrow T(x) \vee Q(y,x)$$

We allow LRA atoms $x - y \,\#\, r$ and $x \,\#\, q$ where

- $r, q \in \mathbb{Q}$ are numerical values, $\# \in \{<, \leq, =, \neq, \geq, >\}$,
- atoms $x - y \,\#\, r$ need to be conjoined with bounds $\ell \leq x \leq u$ and $\ell' \leq y \leq u'$ with $\ell, u, \ell', u' \in \mathbb{Q}$,
- other rational variables, e.g. z above, need not be bounded.

BS(BD): Distinguishable Regions of $\mathbb{Q}^2$

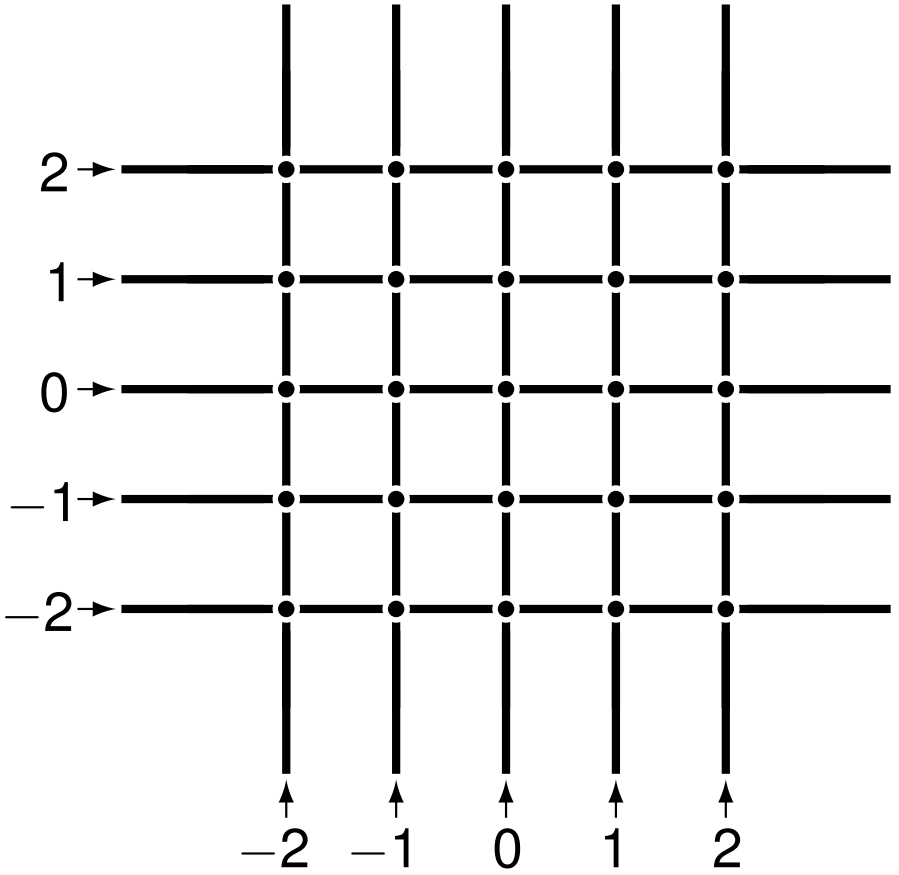

Constraints $x \# r$
with $r \in \{-2, -1, 0, 1, 2\}$
can distinguish
the grid regions.

BS(BD): Distinguishable Regions of $\mathbb{Q}^2$

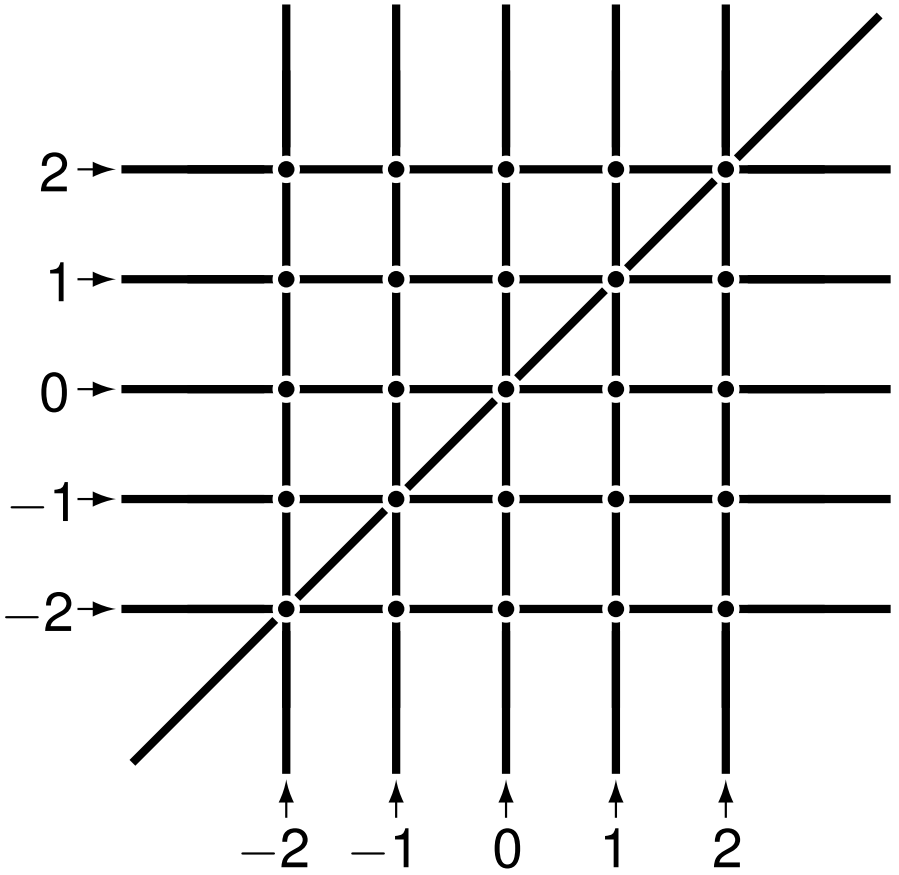

Constraints $x \# r$ with $r \in \{-2, -1, 0, 1, 2\}$ can distinguish the grid regions.

Constraints $x \# y$ can distinguish the simple diagonal.

BS(BD): Distinguishable Regions of $\mathbb{Q}^2$

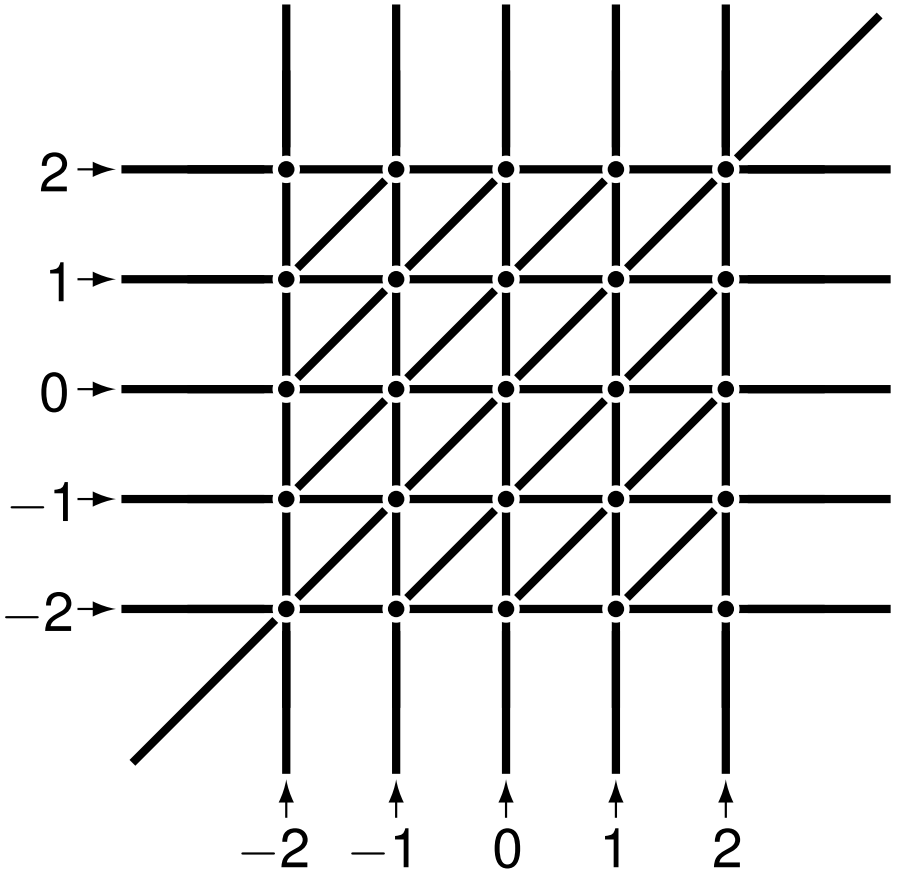

Constraints $x \# r$
with $r \in \{-2, -1, 0, 1, 2\}$
can distinguish
the grid regions.

Constraints $x \# y$
can distinguish
the simple diagonal.

Difference constraints
$x - y \# r$
with bounds
$-2 \leq x, y \leq 2$
can distinguish
more triangles.

BS(BD): Distinguishable Regions of $\mathbb{Q}^2$

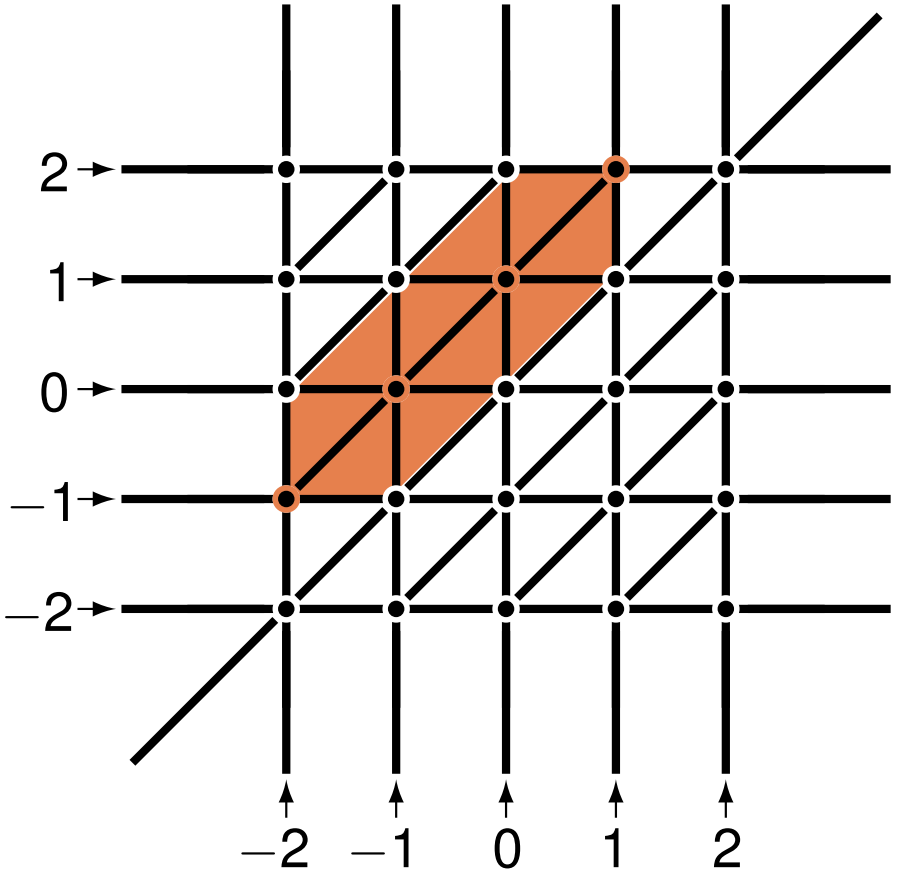

$$-2 < x - y < 0 \quad \wedge \quad -2 \leq x, y \leq 2 \quad \wedge \quad x \leq 1 \quad \wedge \quad -1 \leq y$$

Constraints $x \# r$ with $r \in \{-2, -1, 0, 1, 2\}$ can distinguish the grid regions.

Constraints $x \# y$ can distinguish the simple diagonal.

Difference constraints $x - y \# r$ with bounds $-2 \leq x, y \leq 2$ can distinguish more triangles.

Expressiveness of BD constraints

Definition ($\simeq_\kappa$-equivalence)

Let $\mathcal{A}$ be any structure and let m, κ be any positive integers. For any rational value r we write $\lfloor r \rfloor$ to address its integer part and $\mathrm{fr}(r)$ to address $r - \lfloor r \rfloor$.

Two tuples $\bar{r}, \bar{q} \in \mathbb{Q}^m$ are called $\simeq_\kappa$-*equivalent*, denoted $\bar{r} \simeq_\kappa \bar{q}$, if

(i) for every i either $r_i > \kappa$ and $s_i > \kappa$, or $r_i < -\kappa$ and $s_i < -\kappa$, or the following conditions are met:
  (i.a) $\lfloor r_i \rfloor = \lfloor s_i \rfloor$ and
  (i.b) $\mathrm{fr}(r_i) = 0$ if and only if $\mathrm{fr}(s_i) = 0$, and

(ii) for all i, j
  (ii.a) if $r_i, r_j > \kappa$ or $r_i, r_j < -\kappa$, then $r_i \leq r_j$ iff $s_i \leq s_j$,
  (ii.b) if $-\kappa \leq r_i, r_j \leq \kappa$, then $\mathrm{fr}(r_i) \leq \mathrm{fr}(r_j)$ iff $\mathrm{fr}(s_i) \leq \mathrm{fr}(s_j)$.

$\simeq_\kappa$-uniformity

We define $\simeq_\kappa$-*uniformity* in analogy to $\mathcal{J}_\mathcal{A}$-uniformity.

Lemma

Any satisfiable finite BS(BD) clause set N has a model $\mathcal{A}$ that is $\simeq_\kappa$-uniform and that interprets the free sort $\mathcal{S}$ with a finite domain, where κ is the smallest positive integer that is larger than the absolute value of any rational number occurring in N.

Theorem 2.4

Satisfiability problem for finite BS(BD) clause sets is NExpTime-complete.

The general proof outline is similar to the BS(SLR) case.

3. An Application for BS(BD):

Formalizing Reachability for Timed Automata

Reminder: Timed Automata [AD94, HNSY94]

Finite state machines equipped with *real*-valued clocks x, y, ...

Constraints: $x \# d$, $x - y \# d$, $\# \in \{<, \leq, =, \geq, >\}$, $d \in \mathbb{N}$

Operations: $x \leftarrow 0$

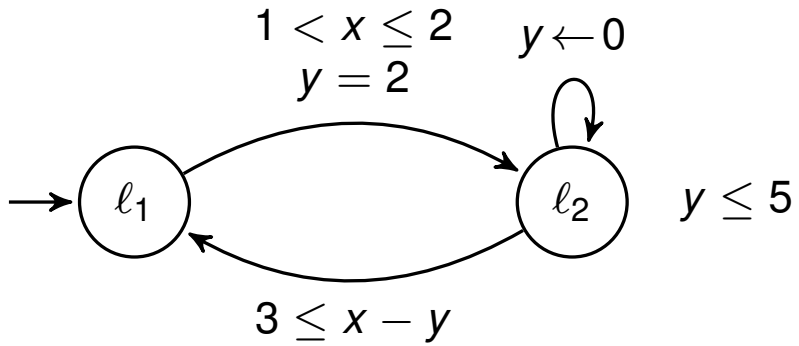

Semantics:
- states $\left\langle \ell, \begin{array}{l} x \mapsto r_1 \\ y \mapsto r_2 \end{array} \right\rangle$ with $r_1, r_2 \in \mathbb{R}$,
- transitions between locations (instantaneous), and
- progress of time (for all clocks simultaneously).

$\rightsquigarrow$ Reachability is PSPACE-complete.

TA Constraints: Distinguishable Regions of $\mathbb{Q}^2$

For two clocks x, y, TA constraints with constants $d \leq 2$ can distinguish regions as follows:

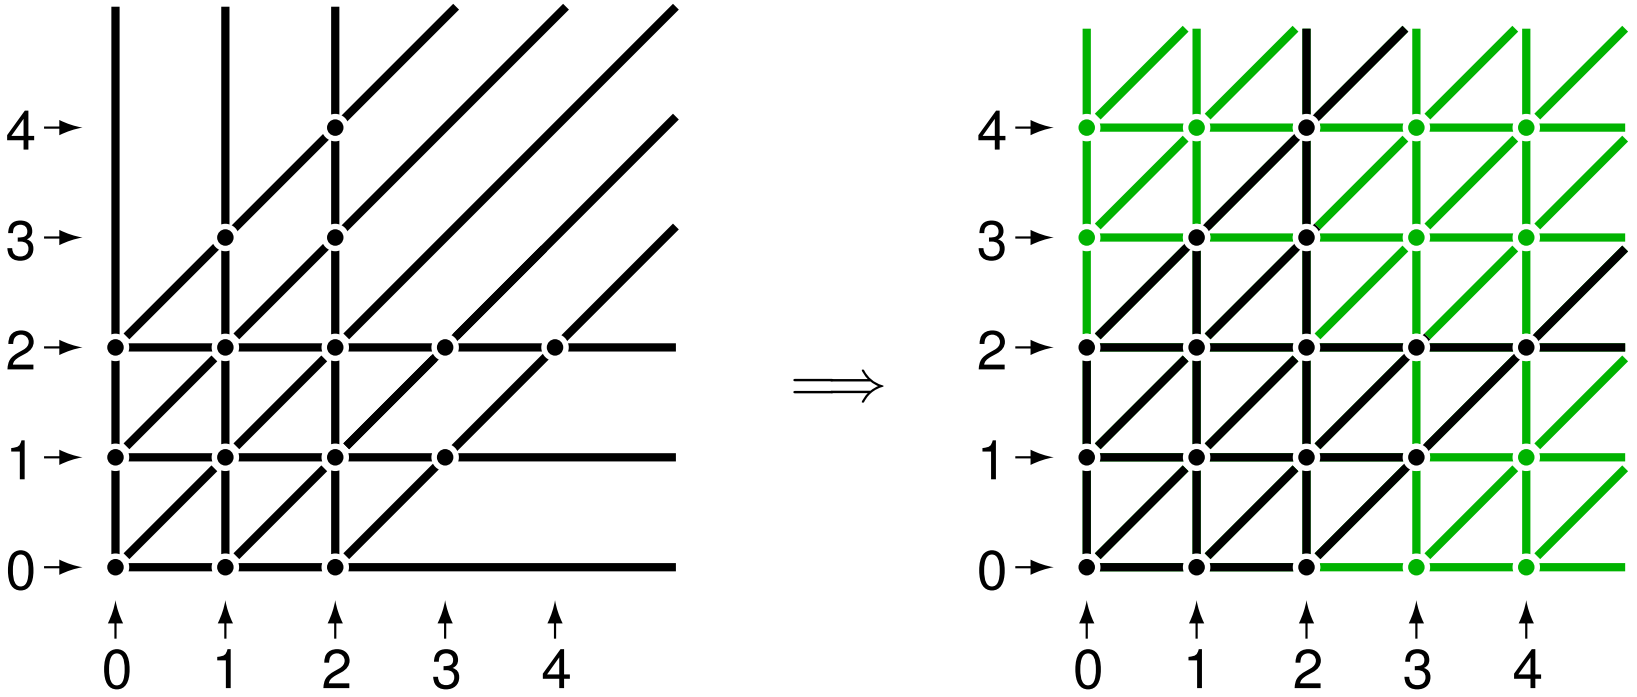

$\rightsquigarrow$ The BS(BD) regions for constraints with constants from $\{-4, \ldots, 0, \ldots, 4\}$ are a refinement of the TA regions.

Encoding Reachability for a Timed Automaton [FW12]

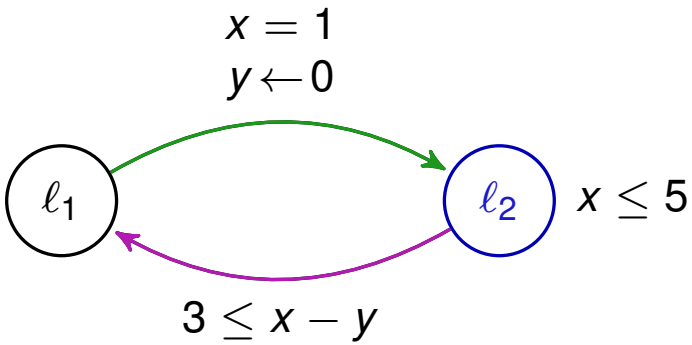

$$x = 1 \wedge y' = 0 \wedge x \leq 5 \wedge \mathrm{Reach}(\ell_1, x, y) \rightarrow \mathrm{Reach}(\ell_2, x, y')$$

$$(\exists t.\, t \geq 0 \wedge x' = x + t \wedge y' = y + t) \wedge x' \leq 5 \wedge \mathrm{Reach}(\ell_2, x, y) \rightarrow \mathrm{Reach}(\ell_2, x', y')$$

$$3 \leq x - y \wedge \mathrm{Reach}(\ell_2, x, y) \rightarrow \mathrm{Reach}(\ell_1, x, y)$$

Start clause: $x = 0 \wedge y = 0 \wedge \rightarrow \mathrm{Reach}(\ell_1, x, y)$

Query clause: $y = 4 \wedge \mathrm{Reach}(\ell_2, x, y) \rightarrow \square$

$\rightsquigarrow$ Saturation leads to $\square$ if the answer to the query is *YES*.

Encoding Reachability for a Timed Automaton [FW12]

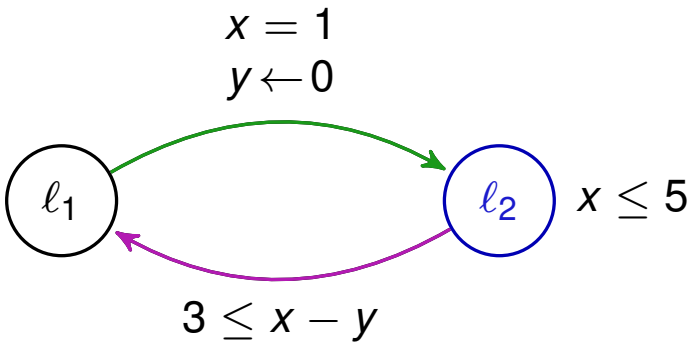

$$x = 1 \wedge y' = 0 \wedge x \leq 5 \wedge \mathrm{Reach}(\ell_1, x, y) \rightarrow \mathrm{Reach}(\ell_2, x, y')$$

$$\big(\exists t.\, t \geq 0 \wedge x' = x + t \wedge y' = y + t\big) \wedge x' \leq 5 \wedge \mathrm{Reach}(\ell_2, x, y) \rightarrow \mathrm{Reach}(\ell_2, x', y')$$

$$3 \leq x - y \wedge \mathrm{Reach}(\ell_2, x, y) \rightarrow \mathrm{Reach}(\ell_1, x, y)$$

Start clause: $x = 0 \wedge y = 0 \wedge \rightarrow \mathrm{Reach}(\ell_1, x, y)$

Query clause: $y = 4 \wedge \mathrm{Reach}(\ell_2, x, y) \rightarrow \square$

$\rightsquigarrow$ Saturation leads to $\square$ if the answer to the query is *YES*.

$\rightsquigarrow$ Syntax restrictions of BS(BD) are not met.

Encoding Reachability for a Timed Automaton [FW12]

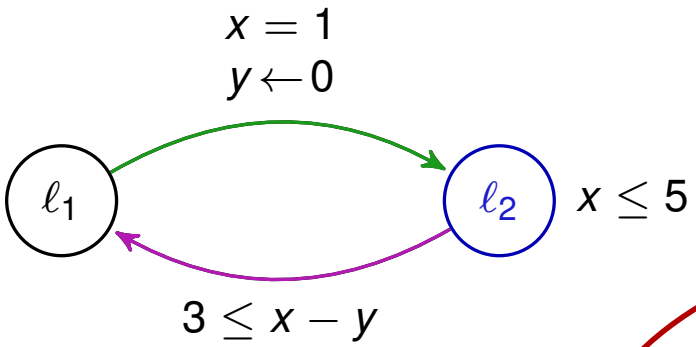

$$x = 1 \wedge y' = 0 \wedge x \leq 5 \wedge \mathrm{Reach}(\ell_1, x, y) \rightarrow \mathrm{Reach}(\ell_2, x, y')$$

$$(\exists t.\, t \geq 0 \wedge x' = x + t \wedge y' = y + t) \wedge x' \leq 5 \wedge \mathrm{Reach}(\ell_2, x, y) \rightarrow \mathrm{Reach}(\ell_2, x', y')$$

$$3 \leq x - y \wedge \mathrm{Reach}(\ell_2, x, y) \rightarrow \mathrm{Reach}(\ell_1, x, y)$$

$\exists t.\, t \geq 0 \wedge x' = x + t \wedge y' = y + t$

is equivalent to

$x' \geq x \wedge y' \geq y \wedge x' - x = y' - y$

$\rightsquigarrow$ Syntax restrictions of BS(BD) are not met.

De-Synchronizing Progression of Time

The constraint $x' \geq x \ \wedge \ y' \geq y \ \wedge \ x' - x = y' - y$

enforces *synchronous* progression of time.

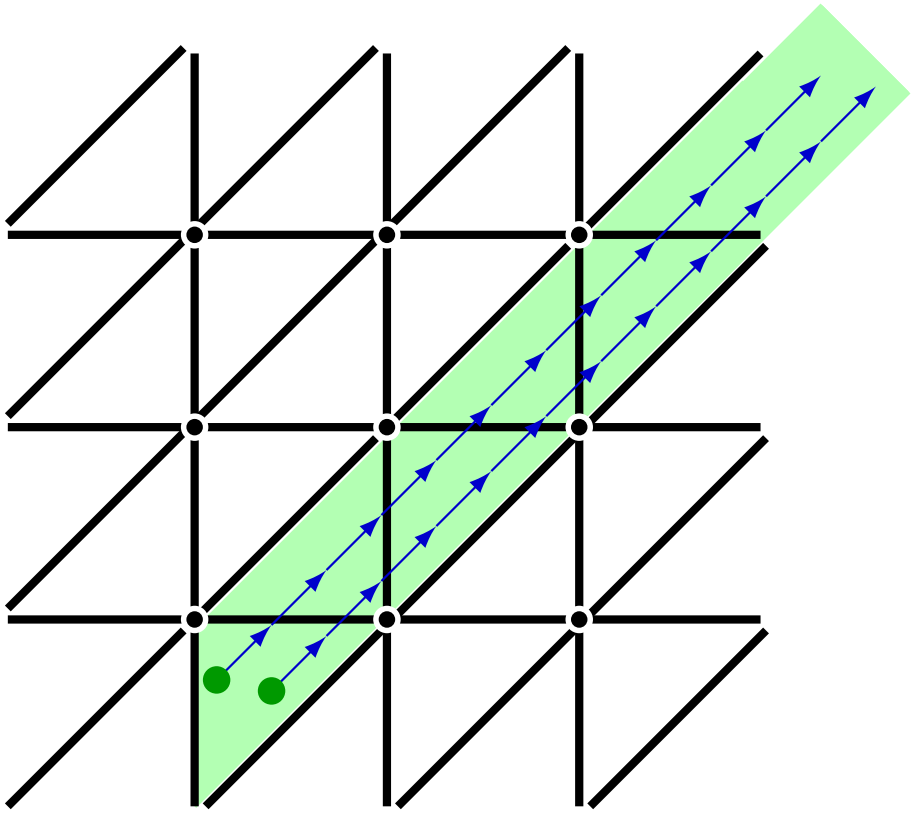

- $\rightsquigarrow$ Synchronous time progression from a reachable point yields a one-dimensional reachable area.
- $\rightsquigarrow$ Since other points in the same region are reachable, we obtain a whole reachable corridor.

De-Synchronizing Progression of Time

The constraint $x' \geq x \ \wedge \ y' \geq y \ \wedge \ x' - x = y' - y$

enforces *synchronous* progression of time.

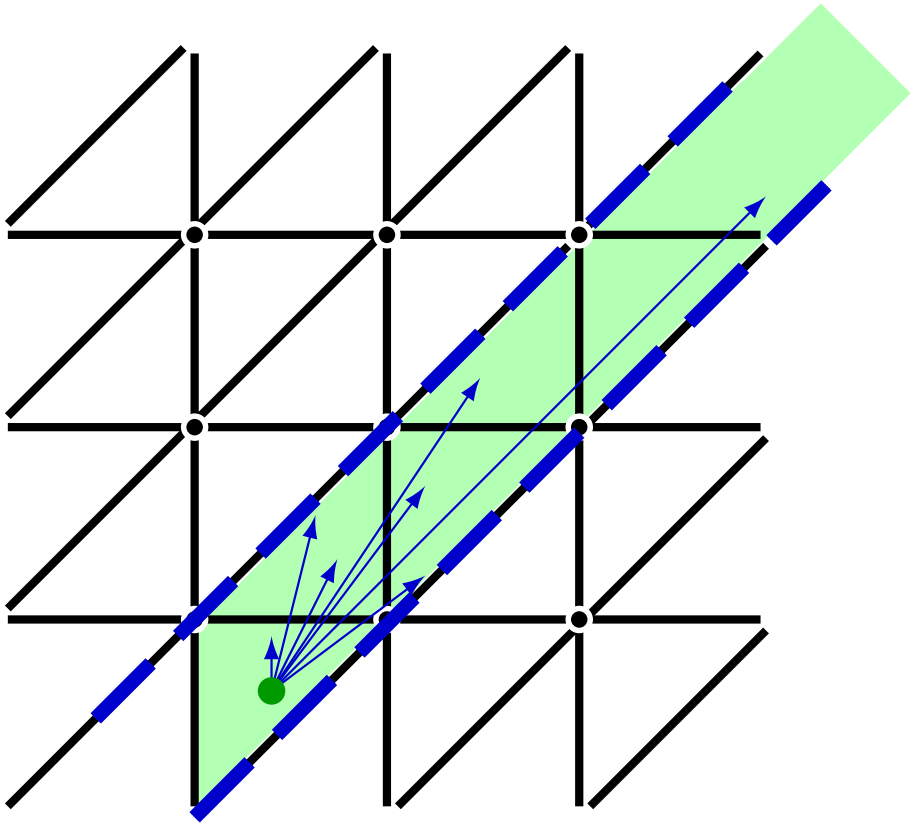

- $\rightsquigarrow$ Synchronous time progression from a reachable point yields a one-dimensional reachable area.
- $\rightsquigarrow$ Since other points in the same region are reachable, we obtain a whole reachable corridor.
- $\rightsquigarrow$ We can weaken the synchronicity requirement to

$$\bigwedge_{k \in \{-\kappa,\ldots,\kappa\}} \big(x - y \leq k \leftrightarrow x' - y' \leq k\big) \wedge \big(x - y \geq k \leftrightarrow x' - y' \geq k\big).$$

Encoding Reachability for a Timed Automaton

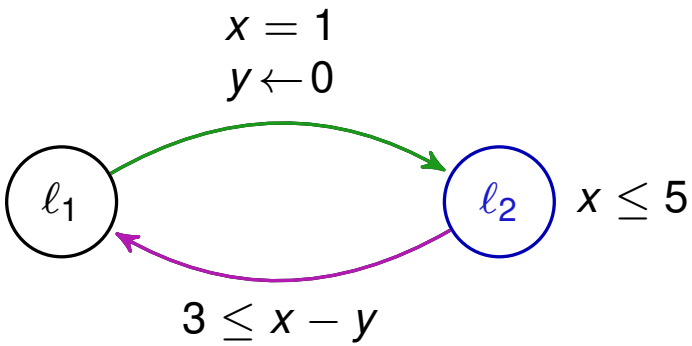

$$x = 1 \wedge y' = 0 \wedge x \leq 5 \wedge 0 \leq x, y, y' < 11 \wedge \text{Reach}(\ell_1, x, y) \rightarrow \text{Reach}(\ell_2, x, y')$$

$$\bigwedge_{k \in \{-10,\ldots,10\}} \Big((x - y \leq k \leftrightarrow x' - y' \leq k) \wedge (x - y \geq k \leftrightarrow x' - y' \geq k) \Big) \wedge 0 \leq x, y, x', y' \leq 11 \wedge x' \leq 5 \wedge \text{Reach}(\ell_2, x, y) \rightarrow \text{Reach}(\ell_2, x', y')$$

$$3 \leq x - y \wedge 0 \leq x, y, y' < 11 \wedge \text{Reach}(\ell_2, x, y) \rightarrow \text{Reach}(\ell_1, x, y)$$

Start clause: $x = 0 \wedge y = 0 \wedge \rightarrow \text{Reach}(\ell_1, x, y)$

Query clause: $y = 4 \wedge \text{Reach}(\ell_2, x, y) \rightarrow \square$

$\rightsquigarrow$ Syntax restrictions of BS(BD) are met.

$\rightsquigarrow$ Reachability for TA can be expressed with BS(BD).

References

[AD94] Alur, Dill. *A Theory of Timed Automata.* Theoretical Computer Science 126(2), 1994

[FW12] Fietzke, Weidenbach. *Superposition as a Decision Procedure for Timed Automata.* Mathematics in Computer Science 6(4), 2012

[GRS90] Graham, Rothschild, Spencer. *Ramsey Theory.* Second edition. Wiley, 1990

[HNSY94] Henzinger, Nicollin, Sifakis, Yovine. *Symbolic Model Checking for Real-Time Systems.* Information and Computation 111(2), 1994

[SVW16] Sturm, Voigt, Weidenbach. *Deciding First-Order Satisfiability when Universal and Existential Variables are Separated.* Logic in Computer Science (LICS'16). Extended version available under `https://arxiv.org/abs/1511.08999`

[Voi17] Voigt. *The Bernays–Schönfinkel–Ramsey Fragment with Bounded Difference Constraints over the Reals is Decidable.* Frontiers of Combining Systems (FroCoS'17). Extended version available under `https://arxiv.org/abs/1706.08504`