Dependence Logic, Spring 2011
Supplementary notes

Juha Kontinen

May 4, 2011

1 Supplementary material to the lecture of 31.3

We discuss the proof of the following equivalence from the lectures.

Proposition 1.1. $\exists x_n \exists x_m \phi \equiv^* \exists x_m \exists x_n \phi$.

Proof. Let $\mathcal{M}$ be a structure and X a team. We will show that if $(\exists x_n \exists x_m \phi, X, 1) \in \mathcal{T}_{\mathcal{M}}$, then $(\exists x_m \exists x_n \phi, X, 1) \in \mathcal{T}_{\mathcal{M}}$. The converse implication follows by symmetry. Furthermore, the equivalence

$$(\exists x_n \exists x_m \phi, X, 0) \in \mathcal{T}_{\mathcal{M}} \Leftrightarrow (\exists x_m \exists x_n \phi, X, 0) \in \mathcal{T}_{\mathcal{M}},$$

follows from the fact that

$$(\forall x_n \forall x_m \neg\phi, X, 1) \in \mathcal{T}_{\mathcal{M}} \Leftrightarrow (\forall x_m \forall x_n \neg\phi, X, 1) \in \mathcal{T}_{\mathcal{M}},$$

which has been shown in the lectures.

Let us then assume that $(\exists x_n \exists x_m \phi, X, 1) \in \mathcal{T}_{\mathcal{M}}$. This implies that $(\phi, X(F/x_n)(G/x_m), 1) \in \mathcal{T}_{\mathcal{M}}$ for some $F\colon X \to M$, and $G\colon X(F/x_n) \to M$. In order to show that $(\exists x_m \exists x_n \phi, X, 1) \in \mathcal{T}_{\mathcal{M}}$, it suffices to define analogues F^* and G^* of F and G, respectively, such that

$$X(G^*/x_m)(F^*/x_n) \subseteq X(F/x_n)(G/x_m), \tag{1}$$

since then, by the Closure Test, we get

$$(\phi, X(G^*/x_m)(F^*/x_n), 1) \in \mathcal{T}_{\mathcal{M}},$$

and $(\exists x_m \exists x_n \phi, X, 1) \in \mathcal{T}_{\mathcal{M}}$.

We will next define the functions F^* and G^*. Define $G^*\colon X \to M$ as follows: for $s \in X$,

$$G^*(s) = G(s(F(s)/x_n)). \tag{2}$$

It is easy to see that G^* is well defined. Let us then define $F^*\colon X(G^*/x_m) \to M$. Now, we are faced with the possibility that, for $s \in X(G^*/x_m)$, there might be several $s' \in X$ for which

$$s'(G^*(s')/x_m) = s. \tag{3}$$

The idea is now to choose for each $s \in X(G^*/x_m)$ one such s' and to define F^* using the function F. In other words, we define $F^*\colon X(G^*/x_m) \to M$ as follows: for $s \in X(G^*/x_m)$, we pick $s' \in X$ such that (3) holds and define

$$F^*(s) = F(s'). \tag{4}$$

Again, it is obvious that F^* is well-defined.

We still need to show that (1) holds. Let $t \in X(G^*/x_m)(F^*/x_n)$, $s = t \upharpoonright \mathrm{dom}(X) \cup \{x_m\}$, and $s' \in X$ the assignment such that $F^*(s) = F(s')$ (see (4)). Note that, because of this, $t(x_n) = F^*(s) = F(s')$. On the other hand, by (3) it holds that $G^*(s') = s(x_m)$ and thus by (2) we get that $t(x_m) = s(x_m) = G^*(s') = G(s'(F(s')/x_n))$. This shows that $t = s'(F(s')/x_n)(G(r)/x_m)$, where $r = s'(F(s')/x_n)$, and $t \in X(F/x_n)(G/x_m)$. □

The following example illustrates the fact that the inclusion in (1) can be strict.

Example 1.2. *Let $\mathcal{M}$ be a model with $M = \{0,1\}$. Consider the following team X of M with domain $\{x_1, x_2\}$:*

	x_1	x_2
s_1	0	1
s_2	0	0

(5)

Define $F\colon X \to M$ so that $s_1 \mapsto 1$ and $s_2 \mapsto 0$. Define then $G\colon X(F/x_1) \to M$ so that $G(s) = 0$ for all s. It is easy to verify that $X(F/x_1)(G/x_2)$ is the team:

	x_1	x_2
s_3	1	0
s_4	0	0

(6)

Let us now try to define F^* and G^* which would allow us to swap to order of supplementation. We want to define $G^*\colon X \to M$ that agrees with G. Note *that since $s(x_2) = 0$ for all $s \in X(F/x_1)(G/x_2)$, we have to define G^* so that $G^*(s) = 0$ for all $s \in X$. Then $X(G^*/x_2)$ is the team*

	x_1	x_2
s_4	0	0

(7)

and now, for F^* we have to possibilities; we can map $s_4 \in X(G^*/x_2)$ either to 0 or 1 and in both cases

$$X(G^*/x_2)(F^*/x_1) \subsetneq X(F/x_1)(G/x_2).$$

2 Supplementary material to the lecture of 21.4

The main result (Theorem 6.15 in [Vää07]) of this lecture is the following:

Theorem 2.1. *For every sentence $\phi \in \Sigma_1^1[L]$ there is a sentence $\phi^* \in \mathcal{D}[L]$ such that, for all structures $\mathcal{M}$:*

$$\mathcal{M} \models \phi \Leftrightarrow \mathcal{M} \models \phi^*.$$

Proof. We discuss the last part of the proof and, for the first parts, refer to the course textbook [Vää07]. First of all, we may assume that ϕ is in Skolem normal form:

$$\phi := \exists f_1 \ldots \exists f_n \forall x_1 \ldots \forall x_m \psi, \tag{8}$$

where ψ is quantifier-free. Furthermore, we may assume that for every function symbol $g_r \in \{f_1, \ldots, f_n\} \cup L$ appearing in ψ there is a unique tuple $x_1^r, \ldots, x_{k_r}^r$ of distinct variables such that all occurrences of g_r in ψ are of the form $g(x_1^r, \ldots, x_{k_r}^r)$ (see pages 96-97 in [Vää07] on how to achieve this).

We are now ready to translate the sentence (8) into dependence logic:

$$\phi^* := \forall x_1 \ldots \forall x_m \exists x_{m+1} \ldots \exists x_{m+n} (\bigwedge_{1 \le j \le n} {=}(x_1^j, \ldots, x_{k_j}^j, x_{m+j}) \wedge \psi'), \tag{9}$$

where ψ' is obtained from ψ by replacing all occurrences of the term $f_j(x_1^j, \ldots, x_{k_j}^j)$ by the variable x_{m+j}. We will next show that the sentences ϕ and ϕ^* are indeed logically equivalent. Note first that it suffices to show the following:

- $\star$ For all $\mathcal{M}$ and $F_i \colon M^{k_i} \to M$, for $1 \le i \le n$:

$$(\mathcal{M}, F_1, \ldots, F_n) \models \forall x_1 \ldots \forall x_m \psi \Leftrightarrow \mathcal{M} \models_{X(G_1/x_{m+1}) \cdots (G_n/x_{m+n})} \psi',$$

where $X = \{\emptyset\}(M/x_1)\cdots(M/x_m)$ (essentially the set of all m-tuples of M) and the supplement function

$$G_i\colon X(G_1/x_{m+1})\cdots(G_{i-1}/x_{m+i-1}) \to M$$

is defined using F_i as follows:

$$G_i(s) = F(s(x_1^i), \ldots, s(x_{k_i}^i)). \tag{10}$$

It is important to note that there is a one-to-one correspondence between the functions F_i and the supplement functions G_i that satisfy the dependence formula

$$=(x_1^i, \ldots, x_{k_i}^i, x_{m+i}).$$

In particular, G_i was defined in (10) in such a way that this dependence atom will be satisfied.

Let us first assume that $(\star)$ holds and show the equivalence of ϕ and ϕ^*. Let $\mathcal{M}$ be a structure and assume that $\mathcal{M} \models \phi$. Then there are functions $F_1, \ldots, F_n$ such that

$$(\mathcal{M}, F_1, \ldots, F_n) \models \forall x_1 \ldots \forall x_m \psi.$$

By $(\star)$, we get that

$$\mathcal{M} \models_{X(G_1/x_{m+1})\cdots(G_n/x_{m+n})} \psi',$$

where G_i is defined as in (10). By definition, the functions G_i satisfy the required dependencies, and hence we get that

$$\mathcal{M} \models_{X(G_1/x_{m+1})\cdots(G_n/x_{m+n})} \Big(\bigwedge_{1\le j\le n} =(x_1^j, \ldots, x_{k_j}^j, x_{m+j}) \wedge \psi'\Big).$$

By the semantics of the existential quantifier of dependence logic, we get that

$$\mathcal{M} \models_X \exists x_{m+1} \ldots \exists x_{m+n}\Big(\bigwedge_{1\le j\le n} =(x_1^j, \ldots, x_{k_j}^j, x_{m+j}) \wedge \psi'\Big),$$

and by the semantics of the universal quantifier:

$$\mathcal{M} \models \forall x_1 \ldots \forall x_m \exists x_{m+1} \ldots \exists x_{m+n}\Big(\bigwedge_{1\le j\le n} =(x_1^j, \ldots, x_{k_j}^j, x_{m+j}) \wedge \psi'\Big).$$

The converse implication is analogous. In other words, assuming $\mathcal{M} \models \phi^*$ we can show $\mathcal{M} \models \phi$ by reversing the steps above.

Let us then show that $(\star)$ holds. We will first show that for all $s \in X(G_1/x_{m+1})\cdots(G_n/x_{m+n})$

$$(\mathcal{M}, F_1, \ldots, F_n) \models_{s'} \psi \Leftrightarrow \mathcal{M} \models_s \psi', \tag{11}$$

where $s' = s \restriction \{x_1, \ldots, x_m\}$. We can prove this using induction on ψ for all quantifier-free formulas. The induction goes through since the interpretation $x_{m+i}\langle s\rangle$ of the variable x_{m+i} agrees with the interpretation of the term $f_i(x_1^i, \ldots, x_{k_i}^i)\langle s'\rangle$: $x_{m+i}\langle s\rangle = s(x_{m+i}) = G_i(s \restriction \{x_1, \ldots, x_{m+i-1}\}) = F_i(s(x_1^i), \ldots, s(x_{k_i}^i)) = F_i(s'(x_1^i), \ldots, s'(x_{k_i}^i)) = f_i(x_1^i, \ldots, x_{k_i}^i)\langle s'\rangle$.

Let us now prove $(\star)$. Assume $(\mathcal{M}, F_1, \ldots, F_n) \models \forall x_1 \ldots \forall x_m \psi$. Then for all $s' \colon \{x_1, \ldots, x_m\} \to M$ it holds that $(\mathcal{M}, F_1, \ldots, F_n) \models_{s'} \psi$. By (11), we get that for all $s \in X(G_1/x_{m+1})\cdots(G_n/x_{m+n})$: $\mathcal{M} \models_s \psi'$ and, since ψ' is a first-order formula of dependence logic

$$\mathcal{M} \models_{X(G_1/x_{m+1})\cdots(G_n/x_{m+n})} \psi'. \tag{12}$$

Let us then assume that (12) holds. By downward closure and the fact that ψ' is first-order, we get that for all $s \in X(G_1/x_{m+1})\cdots(G_n/x_{m+n})$: $\mathcal{M} \models_s \psi'$. Now we use again (11) to get that for all $s' \colon \{x_1, \ldots, x_m\} \to M$: $(\mathcal{M}, F_1, \ldots, F_n) \models_{s'} \psi$, and finally that $(\mathcal{M}, F_1, \ldots, F_n) \models \forall x_1 \ldots \forall x_m \psi$. This completes the proof of $(\star)$ as hence the claim of the theorem. □

Theorem 2.1 gives us a normal form for sentences of dependence logic:

Corollary 2.2. *Every sentence $\phi \in \mathcal{D}$ is logically equivalent to a sentence $\phi^* \in \mathcal{D}$ of the form*

$$\forall x_1 \ldots \forall x_m \exists x_{m+1} \ldots \exists x_{m+n} \psi,$$

where ψ is quantifier-free.

Proof. Let $\phi \in \mathcal{D}$ be a sentence. Then there is a sentence $\tau_{1,\phi}$ of Σ_1^1 that is logically equivalent to ϕ. By Theorem 2.1, there is a sentence $\phi^* \in \mathcal{D}$ of the required form that is logically equivalent to $\tau_{1,\phi}$, and thus to ϕ. □

It is important to note that Theorem 2.1 applies only to sentences of dependence logic and it does not characterize formulas of dependence logic with free variables. We have seen that for every formula $\phi(x_1, \ldots, x_k) \in \mathcal{D}[L]$ there is a sentence $\tau_{1,\phi}(S) \in \Sigma_1^1[L \cup \{S\}]$ such that for all $\mathcal{M}$ and teams X with domain $\{x_1, \ldots, x_k\}$

$$\mathcal{M} \models_X \phi \Leftrightarrow (\mathcal{M}, rel(X)) \models \tau_{1,\phi}(S).$$

However, this result only gives us an upperbound for the complexity of open formulas of dependence logic. We also know that all formulas of dependence logic satisfy downward closure and hence not all sentences of $\Sigma^1_1[L \cup \{S\}]$ correspond to formula of dependence logic. On the syntactic level, downward closure is reflected by the fact that the relation symbol S appears in $\tau_{1,\phi}(S)$ only negatively, that is, all subformulas of $\tau_{1,\phi}(S)$ involving S are of the form $\neg S(t_1, ..., t_k)$.We will next show that S appearing only negatively is the syntactic counterpart of downwards monotonicity (in other words, downwards closure).

Definition 2.3. *Let $\phi \in \Sigma^1_1[L \cup \{R\}]$, where ϕ is in negation normal form and R is a k-ary relation symbol.*

- *We say that ϕ is downwards monotone with respect to R, if for all $L \cup \{R\}$ structures $(\mathcal{M}, A)$ (where $A \subseteq M^k$ interprets R) and $B \subseteq A$:*

$$(\mathcal{M}, A) \models \phi \Rightarrow (\mathcal{M}, B) \models \phi.$$

- *We say that R appears in ϕ only negatively, if all subformulas of ϕ involving R are of the form $\neg R(t_1, ..., t_k)$.*

The following now holds:

Proposition 2.4. *Let $\phi \in \Sigma^1_1[L \cup \{R\}]$, where ϕ is in negation normal form. Then*

- *if R appears only negatively in ϕ, then ϕ is downwards monotone with respect to R.*
- *if ϕ is downwards monotone with respect to R, then there is a Σ^1_1-formula ϕ^* logically equivalent to ϕ in which R appears only negatively.*

The next theorem shows that formulas of dependence logic correspond exactly to the negative (downwards monotone) fragment of Σ^1_1:

Theorem 2.5. *[KV09] For every sentence $\psi \in \Sigma^1_1[L \cup \{R\}]$, where R k-ary and in which R appears only negatively, there is $\phi(x_1, \ldots, x_k) \in \mathcal{D}[L]$ such that, for all $\mathcal{M}$ and teams X with domain $\{x_1, \ldots, x_k\}$:*

$$\mathcal{M} \models_X \phi \iff (\mathcal{M}, rel(X)) \models \psi \vee \forall \overline{y} \neg R(\overline{y}).$$

The disjunct $\forall \overline{y} \neg R(\overline{y})$ is needed since $X = \emptyset$ satisfies all formulas of $\mathcal{D}$ but ψ need not always be true if $rel(X) = \emptyset$.

Theorem 2.5 characterizes the expressive power of formulas of dependence logic with free variables. It is proved using the same idea as Theorem 2.1 and it implies, in particular, that the normal form of Corollary 2.2 holds for all formulas of dependence logic.

3 Supplementary material to the lecture of 28.4

In this lecture we review some basics of computational complexity theory.

In computational complexity theory algorithmic problems are encoded as languages over finite alphabets Σ.

Definition 3.1. *An alphabet Σ is a finite set of symbols. The set of all finite Σ-strings with positive length is denoted by Σ^+. A language L is a subset of Σ^+.*

The complexity of an algorithmic problem (encoded as a language) can be measured by the resources a Turing machine needs to decide the problem.

Definition 3.2. *Let Σ be a finite alphabet and $L \subseteq \Sigma^+$.*

- *A deterministic Turing machine N* decides *L, if for all $w \in \Sigma^+$: if $w \in L$ then N accepts w and if $w \notin L$ then N rejects w.*
- *A nondeterministic Turing machine N* decides *L, if for all $w \in \Sigma^+$: if $w \in L$ then at least one of N's computations with input w accepts and if $w \notin L$ then all of N's computations with input w reject.*
- *For $f : \mathbb{N} \to \mathbb{N}$, A TM N decides L in time f, if N decides L and for each $w \in \Sigma^+$ all computations of N with input w stop after at most $f(|w|)$ many computation steps.*
- *For $f : \mathbb{N} \to \mathbb{N}$, A TM N decides L in space f, if N decides L and for each $w \in \Sigma^+$ all computations of N use at most $f(|w|)$ cells on the worktape(s) of N.*

The class P (NP) contains those languages L, for which there is a deterministic (nondeterministic) TM N and a polynomial $p \in \mathbb{N}[x]$ s.t. N decides L in time p. Analogously, the class L (NL) contains those languages L, for which there is a deterministic (nondeterministic) TM N and a constant c such that N decides L in space $c \log(|w|)$. It is known that L $\subseteq$ NL $\subseteq$ P $\subseteq$ NP.

In order to compare the complexity of different problems, we need the notion of reduction.

Definition 3.3. *Let $L_1 \subseteq \Sigma^+$ and $L_2 \subseteq T^+$.*

1. *We say that L_1 is log-space reducible to L_2 if there is a function $f : \Sigma^+ \to T^+$, computable by a log-space bounded deterministic TM, such that for all $w \in \Sigma^+$:*
$$w \in L_1 \Leftrightarrow f(w) \in L_2.$$

2. *A language L is complete for a complexity class $\mathcal{C}$ if $L \in \mathcal{C}$ and all languages $L' \in \mathcal{C}$ can be reduced to L.*

Example 3.4. *We consider propositional formulas F in conjunctive normal form, i.e., F is of the form*

$$\bigwedge_i \alpha_i,$$

where $\alpha_i := \bigvee_j l_j$ and l_j is of the form p_k or $\neg p_k$ for some propositional symbol p_k. We consider the following problems:

$$\textit{2-SAT} := \{F \mid F \textit{ satisfiable and } |\alpha_i| \leq 2\}$$

$$\textit{3-SAT} := \{F \mid F \textit{ satisfiable and } |\alpha_i| \leq 3\}$$

It is easy to check that, e.g., $(p_1 \vee p_2) \wedge (p_3 \vee \neg p_1)$ is an instance of 2-SAT. It is known that 2-SAT is complete for NL and 3-SAT for NP.

3.1 Descriptive complexity theory

Computational complexity was originally defined by measuring the time and space used in a computation. In 1974, Ronald Fagin [Fag74] showed that the complexity class NP, the problems computable in non-deterministic polynomial time, is exactly the set of problems describable in existential second-order logic. This result showed that the complexity of a problem can be understood as the richness of a formal language needed to specify the problem. The research program which has followed Fagin's seminal result is referred to as Descriptive Complexity Theory.

In order to state the result of Fagin, we first need to fix an ecoding of finite structures into strigs.

Let $L = \{R_1, \ldots, R_m\}$ be a vocabulary. Fix a L-structure $\mathcal{M}$. We assume that $\mathrm{Dom}(\mathcal{M}) = \{0, \ldots, n-1\}$ for some n (other possibility is to let the domain be any finite set and assume that the structure is equipped with an ordering). Now each relation $R_i^{\mathcal{M}}$ can be encoded by a binary string $bin(R_i^{\mathcal{M}})$ of length n^{r_i}, where r_i is the arity of R_i, such that "1" in a given position indicates that the corresponding tuple in the lexicographic ordering of $\mathrm{Dom}(\mathcal{M})^{r_i}$ is in $R_i^{\mathcal{M}}$. The binary encoding $bin(\mathcal{M})$ of $\mathcal{M}$ is defined as the concatenation of the bit strings coding its relations:

$$bin(\mathcal{M}) = bin(R_1^{\mathcal{M}}) \cdots bin(R_m^{\mathcal{M}}).$$

Given a class K of L-structures, we write

$$L_K = \{bin(\mathcal{M}) \mid \mathcal{M} \in K\}$$

for the language corresponding to K. Now that we have encoded classes of structures to languages over alphabet $\{0,1\}$, we can formulate the result ($\Sigma_1^1 \equiv \mathrm{NP}$) of Fagin more precisely. Fagin showed that for all L and all classes K of L-structures (closed under isomorphisms) the following are equivalent:

1. K is the class of models of some L-sentence $\phi \in \Sigma_1^1$,

2. $L_K \in \mathrm{NP}$.

Since on the level of sentences the logics Σ_1^1 and $\mathcal{D}$ are equivalent, we get that

$$\mathcal{D} \equiv \mathrm{NP}.$$

4 Supplementary material to the lecture of 2.5

In this lecture we discuss the complexity of quantifier-free formulas of dependence logic. The material discussed here is contained in [Kon10]. We first introduce a notion called coherence that is useful in classifying quantifier-free formulas according to their complexity. The quantifier-free formulas considered below may freely use the connectives $\wedge$ and $\vee$ but negation is only allowed in front of atomic formulas.

Definition 4.1. *Let $\phi(x_1, \ldots, x_n) \in \mathcal{D}$ be quantifier-free. We say that ϕ is k-coherent ($k \in \mathbb{N}$) if for all structures $\mathcal{M}$ and teams X with domain $\{x_1, \ldots, x_n\}$:*

$$\mathcal{M} \models_X \phi \Leftrightarrow \textit{for all } Y \subseteq X (\textit{if } |Y| = k, \textit{ then } \mathcal{M} \models_Y \phi).$$

Note that the direction " $\Rightarrow$ " above holds for any formula ϕ by downwards closure. Note further that if $|X| < k$, then a k-coherent formula is trivially satisfied.

The following proposition lists some observations about coherence.

Proposition 4.2. *Let ϕ and ψ be quantifier-free formulas of $\mathcal{D}$.*

1. *If ϕ is a first-order formula, then ϕ is 1-coherent,*

2. *If ϕ is of the form $=\!(t_1, \ldots, t_n)$, then ϕ is 2-coherent,*

3. *If ϕ is k-coherent and ψ is l-coherent, for some $l \leq k$, then $\phi \wedge \psi$ is k-coherent,*

4. *If ϕ is 1-coherent and ψ is k-coherent, then $\phi \vee \psi$ is k-coherent,*

5. *If ϕ is of the form $\vee_{1\le j\le k}\psi_j$, where ψ_j is the formula $=(t_1,\ldots,t_n)$, then ϕ is $k+1$-coherent.*

Next we show that if ϕ is k-coherent for some $k\in\mathbb{N}$, then the Σ^1_1-sentence $\tau_{1,\phi}$ corresponding to ϕ can be replaced by a FO-sentence.

Theorem 4.3. *Let $L=\{R_1,\ldots,R_m\}$ and let $\phi(x_1,\ldots,x_n)\in\mathcal{D}[L]$ be quantifier-free. Assume that ϕ is k-coherent for some $k\in\mathbb{N}$. Then there is a sentence $\phi^*\in\mathrm{FO}[L\cup\{S\}]$ such that for all structures $\mathcal{M}$ and teams X with domain $\{x_1,\ldots,x_n\}$:*

$$\mathcal{M}\models_X\phi\Leftrightarrow(\mathcal{M},rel(X))\models\phi^*.$$

Before we prove Theorem 4.3, we discuss the following lemma used in the proof. For the lemma, we need the notion of a substructure:

Definition 4.4. *Let $L=\{R_1,\ldots,R_m\}$ and let $\mathcal{M}$ be a L-structure. For any set $B\subseteq M$ we denote by $\mathcal{M}|B$ the substructure of $\mathcal{M}$ generated by B, that is, the domain of $\mathcal{M}|B$ is B and $R_i^{\mathcal{M}|B}=R_i^{\mathcal{M}}\cap B^{l_i}$, where l_i is the arity of R_i.*

Let $\mathcal{M}$ be a L-structure and X a team of M with domain $\{x_1,\ldots,x_n\}$. Let $A_X\subseteq M$ be defined as follows

$$A_X=\{a\in M\mid s(x_i)=a,\text{ for some } s\in X\text{ and }1\le i\le n\}.$$

Lemma 4.5. *Let ϕ be as in Theorem 4.3. Then for all structures $\mathcal{M}$ and teams X with domain $\{x_1,\ldots,x_n\}$:*

$$\mathcal{M}\models_X\phi\Leftrightarrow\mathcal{M}|A_X\models_X\phi.$$

Proof. We prove using induction on ϕ that for all $\mathcal{M}$, X and sets $A_X\subseteq B\subseteq M$:

$$\mathcal{M}\models_X\phi\Leftrightarrow\mathcal{M}|B\models_X\phi.$$

The claim is trivial for atomic and negated atomic formulas. Also the case of conjunction is straightforward. We consider the case ϕ is of the form $\psi_1\vee\psi_2$. Assume that $\mathcal{M}\models_X\phi$ and let B such that $A_X\subseteq B\subseteq M$. Then there is a partition $X=Y\cup Z$ such that

$$\mathcal{M}\models_Y\psi_1\text{ and }\mathcal{M}\models_Z\psi_2.$$

By the induction hypothesis, we get that $\mathcal{M}|B\models_Y\psi_1$ and $\mathcal{M}|B\models_Z\psi_2$ since obviously $B\supseteq A_X$ contains both of the sets A_Y and A_Z. Therefore, it holds that $\mathcal{M}|B \models_X \phi$. Let us then assume that $\mathcal{M}|B \models_X \phi$. Then there is a partition $X = Y \cup Z$ such that

$$\mathcal{M}|B \models_Y \psi_1 \text{ and } \mathcal{M}|B \models_Z \psi_2.$$

By the induction hypothesis

$$\mathcal{M} \models_Y \psi_1 \text{ and } \mathcal{M} \models_Z \psi_2,$$

and hence also $\mathcal{M} \models_X \phi$. □

We are now ready for the proof of Theorem 4.3:

Proof of Theorem 4.3. Since ϕ is k-coherent, for all structures $\mathcal{M}$ and teams X with domain $\{x_1, \ldots, x_n\}$:

$$\mathcal{M} \models_X \phi \Leftrightarrow \text{for all } Y \subseteq X (\text{if } |Y| = k, \text{ then } \mathcal{M} \models_Y \phi).$$

By Lemma 4.5,

$$\mathcal{M} \models_Y \phi \Leftrightarrow \mathcal{M}|A_Y \models_Y \phi.$$

Note that since $|Y| = k$, the set A_Y has at most kn elements. Furthermore, since the vocabulary L is finite, there are only finitely many non-isomorphic $L \cup \{S\}$-structures of size at most kn. Let

$$\mathcal{A}_1, \ldots, \mathcal{A}_r,$$

list all isomorphism types of such structures and define I such that

$$I = \{1 \leq j \leq r \mid \mathcal{A}_j = (\mathcal{M}', rel(Y)) \text{ and } \mathcal{M}' \models_Y \phi\}.$$

We are now ready to define the sentence $\varphi^* \in \mathrm{FO}[L \cup \{S\}]$ (below we abbreviate $x_1^i \ldots x_n^i$ by $\overline{x}^i$):

$$\forall x_1^1 \ldots \forall x_n^1 \forall x_1^2 \ldots \forall x_n^2 \ldots \forall x_1^k \ldots \forall x_n^k ((\bigwedge_{1 \leq t \leq k} S(x_1^t, \ldots, x_n^t) \wedge \bigwedge_{i \neq j} \overline{x}^i \neq \overline{x}^j) \rightarrow \bigvee_{j \in I} \theta_j),$$

where θ_j expresses that the substructure generated by (the interpretations of) $\{x_1^1, \ldots, x_n^1, \ldots, x_1^k, \ldots, x_n^k\}$ with S interpreted as $\{\overline{x}^1, \ldots, \overline{x}^k\}$ is isomorphic to the structure $\mathcal{A}_j$. So the sentence ϕ^* expresses over $(\mathcal{M}, rel(X))$ that all the relevant small substructures $(\mathcal{M}', rel(Y))$ of this structure are such that $\mathcal{M}' \models_Y \phi$. By k-coherence of ϕ, ϕ and ϕ^* are logically equivalent. □

We end this section by noting that not all quantifier-free formulas are coherent [Kon10]:

Theorem 4.6. *The formula* $=(x_1, x_2) \vee =(x_3, x_4)$ *is not k-coherent for any* $k \in \mathbb{N}$.

4.1 The complexity of quantifier-free formulas of dependence logic

Definition 4.7. *Let* $L = \{R_1, \ldots, R_m\}$ *and let* $\phi(x_1, \ldots, x_n) \in \mathcal{D}[L]$ *be quantifier-free. Let*

$$K_\phi = \{(\mathcal{M}, rel(X)) \mid \mathcal{M} \models_X \phi\},$$

and let $L_\phi \subseteq \{0,1\}^+$ *be the language encoding the class* K_ϕ*.*

We can now meaningfully discuss the complexity of quantifier-free formulas ϕ of dependence logic by referring to the complexity of the language L_ϕ. Theorem 4.3, and the fact that FO $\subseteq$ L, immediately implies the following:

Corollary 4.8. *If* ϕ *is* k*-coherent for some* $k \in \mathbb{N}$*, then* $L_\phi \in \mathrm{L}$*, that is, it can de decided by a deterministic TM using only logarithmic space.*

Theorem 4.9. *Suppose that* ϕ *and* ψ *are* 2*-coherent formulas. Then* $L_{\phi \vee \psi} \in$ *NL.*

References

[Fag74] Ronald Fagin. Generalized first-order spectra and polynomial-time recognizable sets. In *Complexity of computation (Proc. SIAM-AMS Sympos. Appl. Math., New York, 1973)*, pages 43–73. SIAM–AMS Proc., Vol. VII. Amer. Math. Soc., Providence, R.I., 1974.

[Kon10] Jarmo Kontinen. *Coherence and Complexity in Fragments of Dependence Logic.* PhD thesis, University of Amsterdam, 2010.

[KV09] Juha Kontinen and Jouko Väänänen. On definability in dependence logic. *J. Log. Lang. Inf.*, 18(3):317–332, 2009.

[Vää07] Jouko Väänänen. *Dependence logic: A New Approach to Independence Friendly Logic*, volume 70 of *London Mathematical Society Student Texts.* Cambridge University Press, Cambridge, 2007.