First-Order Logic vs. Fixed-Point Logic in Finite Set Theory*

Albert Atserias
Computer Science Department
University of California
Santa Cruz, CA 95064
atserias@cse.ucsc.edu

Phokion G. Kolaitis
Computer Science Department
University of California
Santa Cruz, CA 95064
kolaitis@cse.ucsc.edu

Abstract

The ordered conjecture *states that least fixed-point logic LFP is strictly more expressive than first-order logic FO on every infinite class of ordered finite structures. It has been established that either way of settling this conjecture would resolve open problems in complexity theory. In fact, this holds true even for the particular instance of the ordered conjecture on the class of BIT-structures, that is, ordered finite structures with a built-in BIT predicate. Using a well known isomorphism from the natural numbers to the hereditarily finite sets that maps BIT to the membership relation between sets, the ordered conjecture on BIT-structures can be translated to the problem of comparing the expressive power of FO and LFP in the context of finite set theory. The advantage of this approach is that we can use set-theoretic concepts and methods to identify certain fragments of LFP for which the restriction of the ordered conjecture is already hard to settle, as well as other restricted fragments of LFP that actually collapse to FO. These results advance the state of knowledge about the ordered conjecture on BIT-structures and contribute to the delineation of the boundary where this conjecture becomes hard to settle.*

1. Introduction and summary of results

The main goal of descriptive complexity theory is to investigate the connections between computational complexity and logic on classes of finite structures. As regards first-order logic, it is well known that every first-order definable query is computable in LOGSPACE. Moreover, research in descriptive complexity theory has established that essentially all major computational complexity classes can be characterized in terms of definability in natural extensions of first-order logic on classes of finite structures. In particular, Immerman [Imm86] and Vardi [Var82] showed that PTIME $=$ LFP on every class $\mathcal{C}$ of ordered finite structures, that is to say, if $\mathcal{C}$ is a class of ordered finite structures, then the class of *polynomial-time* computable queries on $\mathcal{C}$ coincides with the class of queries definable in *least fixed-point logic* on $\mathcal{C}$. Least fixed-point logic LFP is the extension of first-order logic FO obtained by augmenting the syntax and semantics with a least fixed-point operator for positive first-order formulas. As a general rule, least fixed-point logic is strictly more expressive than first-order logic. In particular, this holds true on the class $\mathcal{O}$ of all ordered finite structures over a fixed vocabulary, as well as on the class $\mathcal{F}$ of all (unordered) finite structures over a fixed vocabulary. There are, however, classes of unordered finite structures on which LFP collapses to FO, and both are properly contained in PTIME. McColm [McC90] was the first to focus attention on this phenomenon and to formulate a certain conjecture concerning necessary and sufficient conditions for the collapse of LFP to FO on an arbitrary class of finite structures. Although in its full generality McColm's conjecture was refuted by Gurevich, Immerman and Shelah [GIS94], it sparked a sequence of related investigations in finite model theory [KV92, DH95, KV96, DLW96]. Moreover, the following special case of McColm's conjecture still remains open:

Conjecture 1 *If $\mathcal{C}$ is an infinite class of ordered finite structures, then first-order logic* FO *is properly contained in least fixed-point logic* LFP *on $\mathcal{C}$.*

This conjecture, which is often called the *Ordered Conjecture,* was made by Kolaitis and Vardi [KV92]. In view of the aforementioned result of Immerman [Imm86] and Vardi [Var82], the ordered conjecture is equivalent to the assertion that FO $\neq$ PTIME on every infinite class of ordered finite structures. Thus, it enunciates an inherent limitation in the expressive power of first-order logic by stating that first-order logic can *never* capture polynomial time. All empirical evidence gathered to date supports it. At the same time, researchers have established that either way of settling the ordered conjecture will have sig-

*Research partially supported by NSF Grant CCR-9610257.

nificant complexity-theoretic consequences. Specifically, Dawar and Hella [DH95] showed that if the ordered conjecture is false, then $\mathrm{PTIME} \neq \mathrm{PSPACE}$. Furthermore, Dawar, Lindell and Weinstein [DLW96] pointed out that if the ordered conjecture holds, then $\mathrm{LINH} \neq \mathrm{ETIME}$, where LINH (the Linear-Time Hierarchy) is the class of languages computable by alternating Turing machines in linear time using a constant number of alternations, and ETIME is the class of languages computable by deterministic Turing machines in $2^{\mathrm{O}(n)}$ time. The separation of LINH from ETIME can be viewed as the linear version of the separation between the Polynomial-Time Hierarchy PH and the full Exponential Time EXPTIME; although it is widely believed that both these separations hold, neither has been established thus far.

Intuitively, the difficulty in proving the ordered conjecture arises from the potential presence of powerful built-in arithmetic predicates that may significantly enhance the expressive power of first-order logic on classes of ordered finite structures. One especially powerful such predicate is the binary relation BIT on the natural numbers $\omega = \{0, 1, 2, \ldots\}$, where $\mathrm{BIT}(k, m)$ holds if and only if the k-th bit of the binary expansion of m is equal to 1. The expressive power of first-order logic on finite structures with BIT as a built-in predicate has been investigated in depth by Immerman [Imm89] and by Barrington, Immerman and Straubing [BIS90]. In many respects, the presence of BIT embodies the difficulties encountered with the ordered conjecture; as a matter of fact, a particular instance of the ordered conjecture on finite structures with BIT turns out to be literally equivalent to open problems in complexity theory. More precisely, let $\mathcal{B} = \{\mathbf{B}_n : n \geq 1\}$ be the class of all ordered finite BIT-structures $\mathbf{B}_n = (\{0, 1, \ldots, n-1\}, \leq, \mathrm{BIT}_n)$, where $\leq$ is the standard linear order and $\mathrm{BIT}_n = \mathrm{BIT} \cap \{0, 1, \ldots, n-1\}^2$.

Question 1 *Is* $\mathrm{FO} \neq \mathrm{LFP}$ *on* $\mathcal{B}$*? In other words, does the ordered conjecture hold on* $\mathcal{B}$*?*

Gurevich, Immerman and Shelah [GIS94] raised this question and stated that it is a "fascinating question in complexity theory and logic related to uniformity of circuits and logical descriptions." Indeed, it can be shown that $\mathrm{FO} \neq \mathrm{LFP}$ on $\mathcal{B}$ if and only if log-time uniform AC^0 is different than polynomial-time uniform AC^0 (see [BIS90, Lin92] for the definitions of these circuit-complexity classes). Moreover, it can also be shown that $\mathrm{FO} \neq \mathrm{LFP}$ on $\mathcal{B}$ if and only if $\mathrm{LINH} \neq \mathrm{ETIME}$.

Our goal in this paper is to advance the state of knowledge about the ordered conjecture on $\mathcal{B}$ by seeking to delineate the "boundary" where this conjecture becomes hard to settle. To this effect, we study certain fragments of least fixed-point logic LFP on $\mathcal{B}$ and investigate the restriction of the ordered conjecture on these fragments. We first identify a natural proper fragment of LFP for which the ordered conjecture cannot be settled without resolving open problems in complexity at the same time. We then establish that the ordered conjecture actually *fails* when further restricted to certain fairly expressive fragments of LFP on $\mathcal{B}$, which means that these fragments collapse to first-order logic on $\mathcal{B}$. To isolate these fragments of LFP, it is more convenient to conduct our investigation in the context of finite set theory and to use set-theoretic concepts and methods. The starting point is a recent paper by Dawar, Doets, Lindell and Weinstein [DDLW98], where it was shown that the standard linear order is first-order definable using BIT. More precisely, let $\mathcal{BIT} = \{\mathbf{BIT}_n : n \geq 1\}$ be the class of (unordered) finite BIT-structures, where $\mathbf{BIT}_n = (\{0, 1, \ldots, n-1\}, \mathrm{BIT}_n)$. Dawar et al. [DDLW98] showed that there is a first-order formula over the vocabulary $\{\mathrm{BIT}\}$ that defines the standard linear order on the class $\mathcal{BIT}$. This rather surprising result was established by exploiting the existence of a well-known isomorphism (see [Bar75]) between (ω, BIT) and $(V_\omega, \in)$, where $V_\omega = \bigcup_{n \in \omega} V_n$ and the V_n's are the *finite ranks*, that is to say, finite initial segments of the universe of sets obtained by iterating the power-set operation: $V_0 = \emptyset, V_{n+1} = \mathrm{P}(V_n)$. The isomorphism between (ω, BIT) and $(V_\omega, \in)$ is the function $e : \omega \to V_\omega$ defined by the recursion:

$$e(0) = \emptyset, e(m) = \{e(k) : \mathrm{BIT}(k, m)\}$$

This isomorphism enables us to translate questions about the expressive power of logics on the class $\mathcal{BIT}$ to equivalent questions on its image class $\mathcal{BFR} = \{e(\mathbf{BIT}_n) : n \geq 1\}$, where $e(\mathbf{BIT}_n) = (\{e(0), e(1), \ldots, e(n-1)\}, \in)$. It is easy to see that the containments $\mathcal{FR} \subset \mathcal{BFR} \subset \mathcal{NFR}$ hold, where $\mathcal{FR} = \{(V_n, \in) : n \geq 1\}$ is the class of all finite ranks, and $\mathcal{NFR}$ is the class of all *near finite ranks*, that is, structures of the form $(M, \in)$ such that $V_n \subseteq M \subseteq V_{n+1}$ for some natural number n. Dawar et al. [DDLW98] showed that there is a first-order definable linear order on $\mathcal{NFR}$ such that its pre-image under the isomorphism e coincides with the standard linear order on ω. It follows that there is a first-order formula that defines the standard linear order on $\mathcal{BIT}$. Consequently, the ordered conjecture on $\mathcal{B}$ reduces to the question of whether $\mathrm{FO} \neq \mathrm{LFP}$ on $\mathcal{BIT}$, which, in turn, is equivalent to the question: is $\mathrm{FO} \neq \mathrm{LFP}$ on $\mathcal{BFR}$?

The above set-theoretic framework makes it possible to isolate and study variants of the ordered conjecture for fragments of LFP that are obtained by applying the least fixed-point operator to collections of set-theoretic formulas with special syntactic properties. In this paper, we focus on the collection of Δ_0 formulas; these are the first-order formulas over a vocabulary containing $\{\in\}$ such that every occurrence of a quantifier is of the form $(\exists y \in z)$ or $(\forall y \in z)$. The collection of Δ_0 formulas has played a fundamental role in the development of set theory for two reasons: Δ_0 formulas possess desirable structural properties, known as absoluteness properties, and they can define many frequently encountered set-theoretic predicates. As Barwise [Bar75, page 10] puts it, "The importance of Δ_0 formulas rests in the metamathematical fact that any predicate defined by a Δ_0 formula is absolute, and the empirical fact that many predicates occurring in nature can be defined by Δ_0 formulas." Both these facts will be of use to us in the sequel. In particular, the proofs of most of our results will rely heavily on the preservation of Δ_0 formulas under end extensions.

Let $\mathrm{LFP}(\Delta_0)$ be the fragment of least fixed-point LFP that consists of the least fixed-points $\mathrm{LFP}_{\overline{x},R}\varphi(\overline{x},R)$ of all Δ_0 formulas $\varphi(\overline{x},R)$ that are positive in the relation symbol R. Consider now the following variant of the ordered conjecture: does $\mathrm{LFP}(\Delta_0)$ collapse to FO on $\mathcal{BFR}$? Clearly, a negative answer to this question will imply that the ordered conjecture holds on $\mathcal{B}$ and, thus, it is at least as hard to establish as the ordered conjecture on $\mathcal{B}$ itself. On the other hand, one may speculate with some reason that the answer to this question is positive, since Δ_0 formulas constitute a rather small (and well-behaved) fragment of first-order logic. Our first main result establishes that if $\mathrm{LFP}(\Delta_0) \subseteq \mathrm{FO}$ on $\mathcal{BFR}$, then $\mathrm{PTIME} \subseteq \mathrm{LINH}$, which, in turn, implies that $\mathrm{PTIME} \neq \mathrm{PSPACE}$. Thus, the collapse of $\mathrm{LFP}(\Delta_0)$ to FO on $\mathcal{BFR}$ can neither be affirmed nor refuted without resolving long-standing open problems in complexity theory at the same time.

The above result motivates the further study of the ordered conjecture for fragments of $\mathrm{LFP}(\Delta_0)$ on $\mathcal{BFR}$. To this effect, we isolate a fairly expressive collection of Δ_0 formulas and establish that the corresponding fragment of $\mathrm{LFP}(\Delta_0)$ collapses indeed to FO on the class $\mathcal{NFR}$ of near finite ranks and, consequently, on the class $\mathcal{BFR}$ as well. This fragment is inspired from the work of Dawar et al. [DDLW98], who showed that a linear order can be defined in first-order logic on $\mathcal{NFR}$. Their proof proceeds by first defining a linear order on $\mathcal{NFR}$ in a natural way as the least fixed-point $\mathrm{LFP}_{x,y,S}\psi(x,y,S)$ of a certain positive Δ_0 formula $\psi(x,y,S)$, and then showing that $\mathrm{LFP}_{x,y,S}\psi(x,y,S)$ can actually be expressed in first-order logic on $\mathcal{NFR}$. An inspection of that particular Δ_0 formula $\psi(x,y,S)$ reveals that it has the following special syntactic property: every occurrence of the binary relation symbol S involves bound variables only. We turn this property into a concept and define the class of *restricted* Δ_0 as the collection of Δ_0 formulas $\varphi(x_1,\ldots,x_k,R)$ such that every occurrence of the relation symbol R involves bound variables only. We then show that if $\varphi(x_1,\ldots,x_k,R)$ is an arbitrary restricted Δ_0 formula that is positive in R, then its least fixed point $\mathrm{LFP}_{x_1,\ldots,x_k,R}\varphi(x_1,\ldots,x_k,R)$ is first-order definable on $\mathcal{NFR}$. This generalizes the results in [DDLW98] and also provides a tool for easily showing that several other basic queries, such as *(finite) ordinal addition*, are first-order definable on $\mathcal{NFR}$.

After this, we consider fragments of $\mathrm{LFP}(\Delta_0)$ obtained by restricting the number of free variables in the Δ_0 formulas under consideration. We observe that if R is a unary relation symbol and $\varphi(x,R)$ is a Δ_0 formula that is positive in R, then the least fixed-point of $\varphi(x,R)$ coincides with the least fixed point of some restricted Δ_0 formula $\varphi^*(x,R)$. It follows that unary $\mathrm{LFP}(\Delta_0)$ collapses to FO on the class $\mathcal{NFR}$ of near finite ranks and, hence, on the class $\mathcal{BFR}$ as well. Clearly, this raises the question whether any similar collapses can be proved for $\mathrm{LFP}(\Delta_0)$ formulas of higher arities, while keeping in mind that we cannot hope to show that $\mathrm{LFP}(\Delta_0)$ formulas of arbitrary arities collapse to FO on $\mathcal{BFR}$ without simultaneously showing that $\mathrm{PTIME} \neq \mathrm{PSPACE}$. Our main result along these lines is that binary $\mathrm{LFP}(\Delta_0)$ collapses to FO on $\mathcal{BFR}$.

Finally, we derive tight polylogarithmic bounds for the growth of the *closure functions* of arbitrary positive Δ_0 formulas on $\mathcal{FR}$, where the closure function of a positive formula gives the number of iterations needed to reach the least fixed-point of the formula. As a corollary, we show that every $\mathrm{LFP}(\Delta_0)$-definable query on $\mathcal{FR}$ is a member of the complexity class NC of queries computable in polylogarithmic time using a polynomial number of processors. This result appears to be of independent interest and suggests the pursuit of descriptive complexity in the context of finite set theory.

2. Preliminaries

Let σ be a relational vocabulary, R a k-ary relation symbol not in σ, and $\varphi(x_1,\ldots,x_k,R)$ a first-order formula over the vocabulary $\sigma \cup \{R\}$. For every σ-structure $\mathbf{M}$ with universe M and every k-ary relation A on M, we write $\varphi^{\mathbf{M}}(A)$ for the k-ary relation on M defined by φ and A, that is,

$$\varphi^{\mathbf{M}}(A) = \{(a_1,\ldots,a_k) \in M^k : \mathbf{M} \models \varphi[a_1,\ldots,a_k,A]\}.$$

The relation A is a *pre fixed-point* of φ on $\mathbf{M}$ if $\varphi^{\mathbf{M}}(A) \subseteq A$; in a dual manner, A is a *post fixed-point* of φ on $\mathbf{M}$ if $A \subseteq \varphi^{\mathbf{M}}(A)$. Finally, A is a *fixed-point* of φ on $\mathbf{M}$ if $A = \varphi^{\mathbf{M}}(A)$. It is well known that if $\varphi(x_1,\ldots,x_k,R)$ is positive in R (which means that every occurrence of R in φ is within an even number of negation symbols), then $\varphi(x_1,\ldots,x_k,R)$ has both a *least fixed-point* $I_\varphi(\mathbf{M})$ and a *greatest fixed-point* $J_\varphi(\mathbf{M})$ on $\mathbf{M}$. As a matter of fact, the least fixed-point $I_\varphi(\mathbf{M})$ is the intersection of all pre fixed-points, whereas the greatest fixed-point $J_\varphi(\mathbf{M})$ is the union of all post fixed-points. Moreover, they can be computed via transfinite iterations, as follows. For every ordinal $\gamma \geq 0$, let $I^\gamma_\varphi(\mathbf{M}) = \varphi^{\mathbf{M}}(\bigcup_{\delta<\gamma} I^\delta_\varphi(\mathbf{M}))$ and

$J^\gamma_\varphi(\mathbf{M}) = \varphi^{\mathbf{M}}(\bigcap_{\delta<\gamma} J^\delta_\varphi(\mathbf{M}))$. Then $I_\varphi(\mathbf{M}) = \bigcup_\gamma I_\varphi(\mathbf{M})$ and $J_\varphi(\mathbf{M}) = \bigcap_\gamma J_\varphi(\mathbf{M})$. Furthermore, there is an ordinal ξ such that $I^\xi_\varphi(\mathbf{M}) = \bigcup_{\delta<\xi} I^\delta_\varphi(\mathbf{M})$ and hence $I_\varphi(\mathbf{M}) = \bigcup_{\delta<\xi} I^\delta_\varphi(\mathbf{M})$. The least such ordinal is called the *closure ordinal* of φ on $\mathbf{M}$ and is denoted by $\mathrm{cl}_\varphi(\mathbf{M})$. Note that if $\mathbf{M}$ is a finite structure, then $\mathrm{cl}_\varphi(\mathbf{M})$ is a positive integer less than or equal than $|M|^k$.

Least fixed-point logic LFP is the extension of first-order logic FO obtained by augmenting the syntax with a new formula $\mathrm{LFP}_{x_1,\ldots,x_k,R}\varphi(x_1,\ldots,x_k,R)$, for every positive in R first-order formula $\varphi(x_1,\ldots,x_k,R)$; naturally, on every structure $\mathbf{M}$ this new formula is interpreted by the least fixed-point $I_\varphi(\mathbf{M})$. It is well known that this syntax gives rise to a robust collection of queries on finite structures. In particular, LFP-definable queries on finite structures are closed under iterated and nested applications of the least fixed-point operator for positive formulas (see [Imm86, GS86]).

A Δ_0 formula over the vocabulary $\sigma \cup \{\in\}$ is a first-order formula such that all occurrences of quantifiers are of the form $(\forall x \in y)$ and $(\exists x \in y)$. We let $\mathrm{LFP}(\Delta_0)$ denote the fragment of LFP that consists of the least fixed-points of positive Δ_0 formulas; this means that every $\mathrm{LFP}(\Delta_0)$ formula is of the form $\mathrm{LFP}_{x_1,\ldots,x_k,R}\varphi(x_1,\ldots,x_k,R)$, where $\varphi(x_1,\ldots,x_k,R)$ is a Δ_0 formula that is positive in the k-ary relation symbol R and has $x_1,\ldots,x_k$ as free variables. Note that every $\mathrm{LFP}(\Delta_0)$-formula involves a single application of the least fixed-point operator, that is, it contains no nested or iterated least fixed-points. Furthermore, no additional first-order and second-order parameters are allowed in $\mathrm{LFP}(\Delta_0)$-formulas.

In a series of papers, Sazonov has studied definability in a variant of $\mathrm{LFP}(\Delta_0)$ on the infinite structure $(V_\omega, \in)$ of all hereditarily finite sets (see [Saz97] for a survey). Here, we study uniform definability in $\mathrm{LFP}(\Delta_0)$ on the collection $\mathcal{BFR}$ of all finite structures that are images of the BIT-structures $\mathbf{BIT}_n = (\{0,1,\ldots,n-1\}, \mathrm{BIT}_n)$ under the isomorphism $e : \omega \to V_\omega$. In particular, we investigate the question: does $\mathrm{LFP}(\Delta_0)$ collapse to FO on $\mathcal{BFR}$? In the process of this investigation, we also study uniform definability in $\mathrm{LFP}(\Delta_0)$ on the classes $\mathcal{NFR}$ of near finite ranks and $\mathcal{FR}$ of finite ranks that envelop $\mathcal{BFR}$ from above and from below.

3. Complexity-theoretic aspects of $\mathrm{LFP}(\Delta_0)$

As explained in the introduction, separating FO from LFP on $\mathcal{BFR}$ is literally equivalent to separating LINH from ETIME, an open problem in complexity theory. In this section, we show that even the collapse of $\mathrm{LFP}(\Delta_0)$ to FO on $\mathcal{BFR}$ would yield important results in complexity theory. Hence, it is difficult to either refute or confirm that $\mathrm{LFP}(\Delta_0)$ collapses to FO on $\mathcal{BFR}$.

Theorem 1 *If* $\mathrm{LFP}(\Delta_0) \subseteq \mathrm{FO}$ *on* $\mathcal{BFR}$, *then* $\mathrm{PTIME} \subseteq \mathrm{LINH}$, *which in turn implies that* $\mathrm{PTIME} \neq \mathrm{PSPACE}$.

Proof (sketch): Let L be a language in PTIME over the alphabet $\{0,1\}$. We will show that the language $L' = \{1^{2^n}\#w\#\tilde{w} : w \in L,\ |w| = n\}$ over $\{0,1,\#\}$ is in $\mathrm{ATIME}(O(\log n), O(1))$; then a de-padding argument puts L in $\mathrm{ATIME}(O(n), O(1)) = \mathrm{LINH}$. Here, $\tilde{w}$ stands for the dual word of w obtained by interchanging 0's and 1's. By Immerman [Imm86] and Vardi [Var82], there is a sentence φ of LFP over the vocabulary $\{P, \leq, \mathrm{BIT}, 0, \max\}$, where P is a unary relation symbol, that defines L on the class of ordered binary words with BIT. We can assume that $\varphi = (\mathrm{LFP}_{\overline{x},R}\psi(\overline{x}, R))(\overline{0})$ with ψ first-order by the normal-form theorem for LFP. Moreover, $\leq$ does not occur in ψ, since by [DDLW98] it is first-order definable from BIT. Similarly, 0 and $\max$ need not occur in ψ, because they are first-order definable as well. We turn $\psi(x_1,\ldots,x_k,R)$ into a Δ_0 formula $\psi'(x_1,\ldots,x_k,y_1,y_2,S)$, where S is a $(k+2)$-ary relation variable. Replace each occurrence $\mathrm{BIT}(z_i,z_j)$ in ψ by $z_i \in z_j$, each positive occurrence $P(z_i)$ in ψ by $z_i \in y_1$, each negative occurrence $\neg P(z_i)$ in ψ by $z_i \in y_2$, and each occurrence $R(z_{i_1},\ldots,z_{i_k})$ by $S(z_{i_1},\ldots,z_{i_k},y_1,y_2)$. Finally, replace subformulas of the form $(\exists z)\theta$ by $(\exists z \in y_1)\theta \vee (\exists z \in y_2)\theta$, and subformulas of the form $(\forall z)\theta$ by $(\forall z \in y_1)\theta \wedge (\forall z \in y_2)\theta$. Using the isomorphism $e : (\omega, \mathrm{BIT}) \cong (V_\omega, \in)$, it is not difficult to prove by induction on the construction of ψ that for every $w \in \{0,1\}^n$ we have that $w \models \mathrm{LFP}\psi(\overline{x}, R)[\overline{0}]$ if and only if $e(\mathbf{BIT}_{2^n}) \models \mathrm{LFP}\psi'(\overline{x}, y_1, y_2, S)[\overline{\emptyset}, e(\mathrm{b}(w)), e(\mathrm{b}(\tilde{w}))]$, where $\mathrm{b}(w)$ stands for the natural number represented in binary by w. Since ψ' is a Δ_0 formula, the hypothesis of the theorem implies that the least fixed-point of ψ' is definable by a first-order formula on $\mathcal{BFR}$. It follows that the query $Q(\mathbf{BIT}_{2^n}) = \{(\mathrm{b}(w), \mathrm{b}(\tilde{w})) : w \in A,\ |w| = n\}$ is first-order definable on $\mathcal{BIT}$. A result in [BIS90] implies then that a suitable encoding of Q is computable in $\mathrm{ATIME}(O(\log n), O(1))$. It turns out that $L' = \{1^{2^n}\#w\#\tilde{w} : w \in L,\ |w| = n\}$ can serve as this encoding. This completes the sketch of the proof that $\mathrm{PTIME} \subseteq \mathrm{LINH}$. Since $\mathrm{LINH} \subseteq \mathrm{DSPACE}(n^2)$, the space-hierarchy theorem implies that $\mathrm{PTIME} \neq \mathrm{PSPACE}$. $\square$

It should be noted that from a result of Dawar and Hella [DH95] it follows that if $\mathrm{LFP} \subseteq \mathrm{FO}$ on $\mathcal{BFR}$, then $\mathrm{PTIME} \neq \mathrm{PSPACE}$. The preceding Theorem 1 shows that the separation of PTIME from PSPACE can be derived from the weaker hypothesis that $\mathrm{LFP}(\Delta_0) \subseteq \mathrm{FO}$ on $\mathcal{BFR}$.

4. Restricted $\mathrm{LFP}(\Delta_0)$ collapses to FO

As mentioned earlier, Dawar et al. [DDLW98] showed that there is a first-order formula over the vocabulary $\{\in\}$

that defines a linear order on the class $\mathcal{NFR}$ of near finite ranks, that is, structures of the form $(M, \in)$ such that $V_n \subseteq M \subseteq V_{n+1}$. For this, they considered the following Δ_0 formula $\psi(x, y, S)$

$$(\exists y' \in y)(\forall x' \in x)(y' \notin x \wedge (x' \notin y \rightarrow S(x', y')))$$

and showed that its least fixed-point is definable by a first-order formula on $\mathcal{NFR}$. Observe that the occurrence of the relation symbol S in ψ involves only the bound variables x' and y'. We now abstract from this observation and introduce the following concept.

Definition 1 A Δ_0 formula $\varphi(x_1, \ldots, x_k, R)$ is *restricted* if every occurrence of the relation symbol R involves only bound variables of φ.

The main result of this section is that the least fixed-point of every positive restricted Δ_0 formula is first-order definable on $\mathcal{NFR}$ and, hence, on $\mathcal{BFR}$ as well. In fact, we show that it is definable by a first-order formula of low syntactic complexity. The class of Σ formulas is the smallest collection of formulas containing the Δ_0 formulas and closed under conjunction, disjunction, bounded quantifications $(\exists y \in x)$, $(\forall y \in x)$, and existential quantification $\exists y$. The collection of Π formulas is defined dually by allowing closure under universal quantification $\forall y$. We say that a query Q on a class $\mathcal{C}$ is Δ*-definable* if it is definable on $\mathcal{C}$ by a Σ formula and by a Π formula.

Theorem 2 *Let $\varphi(x_1, \ldots, x_k, R)$ be a restricted Δ_0 formula that is positive in the k-ary relation symbol R. The least fixed-point of φ is first-order definable on $\mathcal{NFR}$. In fact, it is Δ-definable on $\mathcal{NFR}$.*

The proof of the above theorem is inspired by the argument of Dawar et [DDLW98] showing that the least fixed-point of the Δ_0 formula $\psi(x, y, S)$ is first-order definable on $\mathcal{NFR}$. We need, however, to establish certain absoluteness properties of arbitrary Δ_0 formulas, as well as certain structural properties of arbitrary restricted Δ_0 formulas that will be used heavily in the sequel. We begin with a basic definition from set theory (see [Bar75, page 34]).

Let $\mathbf{M}$ and $\mathbf{N}$ be two structures over the vocabulary $\sigma \cup \{\in\}$. We say that $\mathbf{N}$ is an *end extension* of $\mathbf{M}$, and write $\mathbf{M} \subseteq_{\text{end}} \mathbf{N}$, if $\mathbf{M}$ is a substructure of $\mathbf{N}$ and for every $a \in M$ it is the case that $\{b \in N : b \in^{\mathbf{N}} a\} = \{b \in M : b \in^{\mathbf{M}} a\}$.

Lemma 1 (Absoluteness of Δ_0 formulas [Bar75, page 35]) *If $\varphi(x_1, \ldots, x_k)$ is a Δ_0 formula and $\mathbf{M} \subseteq_{\text{end}} \mathbf{N}$, then for every $(a_1, \ldots, a_k) \in M^k$ we have that $\mathbf{M} \models \varphi[a_1, \ldots, a_k]$ if and only if $\mathbf{N} \models \varphi[a_1, \ldots, a_k]$.*

If $\varphi(x_1, \ldots, x_k, R)$ is a Δ_0 formula, then the set of *R-free indices* of φ, denoted by $U(\varphi)$, is the set of indices of variables that are free in φ and appear in at least one occurrence of R in φ. For example, if $\varphi(x_1, x_2, x_3)$ is the formula $(\forall x_4 \in x_3) R(x_1, x_2, x_4)$, then $U(\varphi) = \{1, 2\}$. Note that a Δ_0 formula φ is restricted if and only if $U(\varphi) = \emptyset$. From now on, we identify structures $(M, \in)$ over $\in$ with their universe M. In the following two lemmas, we establish certain important properties of Δ_0 and of restricted Δ_0 formulas on near finite ranks.

Lemma 2 *Let $\varphi(x_1, \ldots, x_s, R)$ be a Δ_0 formula over $\{\in, R\}$, where R is a k-ary relation symbol, and let M be a near finite rank such that $V_n \subseteq M \subseteq V_{n+1}$. For every $m \leq n$, every relation $A \subseteq M^k$, and every tuple $\overline{a} = (a_1, \ldots, a_s) \in V_{m+1}^s \cap M^s$, we have that $M \models \varphi[\overline{a}, A]$ if and only if $M \models \varphi[\overline{a}, A \cap (V_m \cup \{a_i : i \in U(\varphi)\})^k]$. In particular, if φ is a restricted Δ_0 formula, then $M \models \varphi[\overline{a}, A]$ if and only if $M \models \varphi[\overline{a}, A \cap V_m^k]$.*

Proof: We proceed by induction on the construction of φ. The only non-trivial case is when φ is of the form $(\exists x_i \in x_j)\psi$. In this case, $M \models \varphi[\overline{a}, A]$ if and only if there is some $a \in M$ such that $a \in a_j$ and $M \models \psi[\overline{b}, A]$, where $\overline{b} = (a_1, \ldots, a_{i-1}, a, a_{i+1}, \ldots, a_s)$. Therefore, by induction hypothesis, $M \models \varphi[\overline{a}, A]$ if and only if there is some $a \in M$ such that $a \in a_j$ and $M \models \psi[\overline{b}, A \cap (V_m \cup \{b_l : l \in U(\psi)\})^k]$. Since $a_j \in V_{m+1}$, we have that $a_j \subseteq V_m$ and, hence, for every $a \in a_j$ it is the case that $V_m \cup \{b_l : l \in U(\psi)\} = V_m \cup \{b_l : l \in U(\psi),\ l \neq i\} = V_m \cup \{a_l : l \in U(\varphi)\}$. Consequently, $M \models \varphi[\overline{a}, A]$ if and only if there is some $a \in M$ such that $a \in a_j$ and $M \models \psi[\overline{b}, A \cap (V_m \cup \{a_i : i \in U(\varphi)\})^k]$, which means that $M \models \varphi[\overline{a}, A \cap (V_m \cup \{a_i : i \in U(\varphi)\})^k]$, as required. □

The next lemma yields an absoluteness property of Δ_0 inductions on near finite ranks and also reveals that every positive restricted Δ_0 formula has a unique fixed-point on near finite ranks.

Lemma 3 *Let $\varphi(x_1, \ldots, x_k, R)$ be a Δ_0 formula over $\{\in, R\}$ that is positive in the k-ary relation symbol R, and let M be a near finite rank such that $V_n \subseteq M \subseteq V_{n+1}$. For every $m \leq n$ and every $t \geq 0$, we have that $I^t_\varphi(M) \cap V^k_m = I^t_\varphi(V_m)$ and $J^t_\varphi(M) \cap V^k_m = J^t_\varphi(V_m)$. If, in addition, φ is a restricted Δ_0 formula, then $I_\varphi(M) = J_\varphi(M)$.*

Proof (sketch): The first statement is proved by induction on t using absoluteness. For the second statement, if φ is a positive restricted Δ_0 formula, then it can be shown that $J_\varphi(M) \subseteq I_\varphi(M)$ by induction on the maximum $\in$-rank of $a_1, \ldots, a_k$, where $(a_1, \ldots, a_k) \in J_\varphi(M)$, and using Lemma 2. □

We now have all the necessary tools to show that restricted $\text{LFP}(\Delta_0)$ collapses to FO on the class $\mathcal{NFR}$ of near finite ranks.

Proof of Theorem 2: The key idea of the proof is that there is a first-order formula $\theta(x_1, \dots, x_k)$ that *approximates* the greatest fixed-point $J_\varphi(M)$ on near finite ranks M; the formula θ is based on the characterization of $J_\varphi(M)$ as the union of all post fixed-points of φ. This approximation can be improved to yield an exact first-order definition of $J_\varphi(M)$ by first replacing every occurrence of R in φ by θ, and then iterating φ a constant number of times. The result will then follow from Lemma 3, which asserts that the greatest fixed-point of the restricted Δ_0 formula φ coincides with its least fixed-point. This type of argument was used by Dawar et al. [DDLW98] to show that the least fixed-point of the Δ_0 formula $\psi(x, y, S)$ in the beginning of Section 4 is first-order definable on $\mathcal{NFR}$. Here, we have to deploy the machinery of Lemmas 1, 2 and 3 to show that the argument can be extended and applied to every restricted Δ_0 formula.

We will construct a Σ formula $\psi(x_1, \dots, x_k)$ that defines the greatest fixed-point $J_\varphi(M)$ on every $M \in \mathcal{NFR}$ such that $V_n \subseteq M \subseteq V_{n+1}$ for some $n > 2k - 1$ (for near finite ranks M below V_{2k} we can obtain a Δ_0 definition of $J_\varphi(M)$ by iterating φ a sufficient number of times). The construction of ψ is carried out in three steps. For the first step, define a Δ_0 formula $\mathrm{post}(y)$ expressing that y is the encoding of a post fixed-point of φ on M. We use the standard encoding of a pair $\langle x, y\rangle$ by the set $\{\{x\}, \{x, y\}\}$, and of a tuple $\langle x_1, \dots, x_k\rangle$ by $\langle\langle x_1, \dots, x_{k-1}\rangle, x_k\rangle$. Then $\mathrm{post}(y)$ is the formula $\mathrm{rel}_k(y) \wedge (\forall \langle z_1, \dots, z_k\rangle \in y)(\widehat{\varphi}(z_1, \dots, z_k, R/y))$, where $\mathrm{rel}_k(y)$ is a Δ_0 formula expressing that y is a set of encodings of k-tuples, and $\widehat{\varphi}$ is obtained from φ by replacing atomic formulas $R(z_{i_1}, \dots, z_{i_k})$ by $\langle z_{i_1}, \dots, z_{i_k}\rangle \in y$. Let $A \subseteq M^k$ be such that the set $\langle A\rangle = \{\langle a_1, \dots, a_k\rangle : (a_1, \dots, a_k) \in A\}$ is in M; it is not hard to show that $M \models \mathrm{post}[\langle A\rangle]$ if and only if A is a post fixed-point of $\varphi^M(A)$ on M. For the second step, let $\theta(x_1, \dots, x_k)$ be the Σ formula $(\exists y)(\mathrm{post}(y) \wedge \langle x_1, \dots, x_k\rangle \in y)$ expressing that $(x_1, \dots, x_k)$ belongs to some post fixed-point of φ.

Claim 1: $J_\varphi(M) \cap V^k_{n-2k+1} \subseteq \theta^M \subseteq J_\varphi(M)$.

Proof of Claim 1: Since φ is Δ_0, Lemma 3 implies that $J_\varphi(M) \cap V^k_{n-2k+1} = J_\varphi(V_{n-2k+1})$. Thus, for the first inclusion it suffices to show that if $\overline{a} = (a_1, \dots, a_k) \in J_\varphi(V_{n-2k+1})$, then $M \models \theta[\overline{a}]$. We need a witness r for the quantifier $(\exists y)$ in θ. Put $A = J_\varphi(V_{n-2k+1})$ and $r = \langle A\rangle$. Since $A \subseteq V^k_{n-2k+1}$, we have that $r \in V_n \subseteq M$. Moreover, since $\overline{a} \in A$, we have that $\langle\overline{a}\rangle \in r$; hence, $M \models (\langle\overline{x}\rangle \in y)[\overline{a}, r]$. We can now show that $M \models \mathrm{post}(y)[r]$ using Lemma 1 and the fact that $V_{n-2k+1} \subseteq_{\mathrm{end}} M$. For the second inclusion, we use the fact that $J_\varphi(M)$ is the union of all post fixed-points of φ on M.

Thus, θ yields an approximation of the greatest fixed-point $J_\varphi(M)$ of φ. For the third step, we iterate φ a number of times (in fact, $2k$ times) to obtain progressively better approximations. Define $\theta_0(x_1, \dots, x_k) := \theta(x_1, \dots, x_k)$ and $\theta_i(x_1, \dots, x_k) := \varphi(x_1, \dots, x_k, R/\theta_{i-1})$, for $i \in \{1, \dots, 2k\}$. Note that each θ_i is a Σ formula.

Claim 2: $J_\varphi(M) \cap V^k_{n-2k+1+i} \subseteq \theta^M_i \subseteq J_\varphi(M)$, for every $i = 0, \dots, 2k$.

Proof of Claim 2: Since φ is Δ_0, Lemma 3 implies that $J_\varphi(M) \cap V^k_{n-2k+1+i} = J_\varphi(V_{n-2k+1+i})$. We proceed by induction on i. Claim 1 takes care of the case $i = 0$. Assume that Claim 2 holds for $i - 1$. For the first inclusion, if $\overline{a} = (a_1, \dots, a_k) \in J_\varphi(M) \cap V^k_{n-2k+1+i}$, then $M \models \varphi[\overline{a}, J_\varphi(M)]$. Since $\overline{a} \in V^k_{n-2k+1+i} \cap M^k$ and φ is a restricted Δ_0 formula, Lemma 2 implies that $M \models \varphi[\overline{a}, J_\varphi(M) \cap V^k_{n-2k+i}]$. Note that this is the first place in this proof where we use the assumption that the Δ_0 formula φ is restricted. By induction hypothesis, $J_\varphi(M) \cap V^k_{n-2k+i} \subseteq \theta^M_{i-1}$; therefore, $M \models \varphi[\overline{a}, \theta^M_{i-1}]$, since φ is monotone. It follows that $M \models \varphi(\overline{x}, R/\theta_{i-1})[\overline{a}]$ and $M \models \theta_i[\overline{a}]$. The second inclusion can be proved using the induction hypothesis and the monotonicity of φ; the details are left to the reader.

Finally, let ψ be the Σ formula $\theta_{2k}(x_1, \dots, x_k)$. Claim 2 implies that ψ defines the greatest fixed-point $J_\varphi(M)$ of φ on every near finite rank M such that $V_n \subseteq M \subseteq V_{n+1}$ for some $n > 2k - 1$. Let $\tilde{\varphi}(\overline{x}, R)$ be the dual formula $\neg\varphi(\overline{x}, \neg R)$ of φ. Note that $\tilde{\varphi}$ is also a restricted Δ_0 formula. Moreover, $J_{\tilde{\varphi}}(M) = M - I_\varphi(M)$ (see [Mos74]). Therefore, $I_\varphi(M)$ is Π definable by taking the negation of the Σ formula that defines $J_{\tilde{\varphi}}(M)$. The Δ definability of $I_\varphi(M)$ follows immediately, since $I_\varphi(M) = J_\varphi(M)$ by Lemma 3. This completes the proof of Theorem 2. □

Example 1: In the full paper we show that, in spite of the stringent syntactic requirements imposed on restricted Δ_0 formulas, several natural queries can be expressed using restricted $\mathrm{LFP}(\Delta_0)$ formulas. Thus, Theorem 2 provides a versatile tool for showing that such queries are first-order definable on $\mathcal{NFR}$. Here, we illustrate this technique by considering the query *(finite) ordinal addition* Q_{add}, where if M is a near finite rank, then $Q_{\mathrm{add}}(M) = \{(\alpha, \beta, \gamma) \in M^3 : \alpha, \beta, \gamma$ are ordinals such that $\gamma = \alpha + \beta\}$. The standard recursive specification of ordinal addition does not lead to a restricted Δ_0 formula. Consider, however, the following alternate recursive specification: if α, β, γ are ordinals, then $\gamma = \alpha + \beta$ if and only if

$$\begin{aligned}&(\alpha = 0 \wedge \gamma = \beta) \vee (\beta = 0 \wedge \gamma = \alpha) \vee\\&(\exists\alpha' \in \alpha)(\exists\beta' \in \beta)(\exists\gamma' \in \gamma)(\exists\delta' \in \gamma')\\&(\alpha = \alpha' + 1 \wedge \beta = \beta' + 1 \wedge \gamma = \gamma' + 1 \wedge\\&\gamma' = \delta' + 1 \wedge \delta' = \alpha' + \beta').\end{aligned}$$

This recursive specification can easily be transformed into the least fixed-point of a positive restricted Δ_0 formula that defines Q_{add}; to this end, we use the fact that being an ordinal is Δ_0 definable, and that $\alpha + 1 = \alpha \cup \{\alpha\}$.

5. Unary $\mathrm{LFP}(\Delta_0)$ and binary $\mathrm{LFP}(\Delta_0)$

For every positive integer k, let *k-ary* $\mathrm{LFP}(\Delta_0)$ denote the fragment of $\mathrm{LFP}(\Delta_0)$ that allows the formation of least fixed-points of positive Δ_0 formulas $\varphi(x_1, \ldots, x_k, R)$ such that R is a relation symbol of arity k. The following simple observation reveals that the smallest of these fragments collapses to FO on $\mathcal{NFR}$.

Proposition 1 *If $\varphi(x, R)$ is a unary Δ_0 formula that is positive in R, then there is a unary restricted Δ_0 formula $\varphi^*(x, R)$ that is positive in R and such that $I_\varphi(\mathbf{M}) = I_{\varphi^*}(\mathbf{M})$ on every structure $\mathbf{M}$. Consequently, unary $\mathrm{LFP}(\Delta_0)$ collapses to FO on $\mathcal{NFR}$.*

Proof: Let $\varphi^*(x, R)$ be the restricted Δ_0 formula obtained from $\varphi(x, R)$ by replacing every occurrence of $R(x)$ by $x \neq x$, while preserving all occurrences $R(z)$ with z a bound variable. An easy induction shows that for every structure $\mathbf{M}$ and every ordinal γ we have that $I^\gamma_\varphi(\mathbf{M}) = I^\gamma_{\varphi^*}(\mathbf{M})$; thus $I_\varphi(\mathbf{M}) = I_{\varphi^*}(\mathbf{M})$. Theorem 2 implies then that unary $\mathrm{LFP}(\Delta_0)$ collapses to FO on $\mathcal{NFR}$. □

According to Theorem 1, if $\mathrm{LFP}(\Delta_0)$ collapses to FO on $\mathcal{BFR}$, then $\mathrm{PTIME} \neq \mathrm{PSPACE}$. The proof of this theorem makes a crucial use of the hypothesis that k-ary $\mathrm{LFP}(\Delta_0)$ collapses to FO for *every* $k \geq 1$. In view of Proposition 1, one may investigate whether k-ary $\mathrm{LFP}(\Delta_0)$ collapses to FO on $\mathcal{BFR}$ for *particular* values of k bigger than 1. The main result of this section is that binary $\mathrm{LFP}(\Delta_0)$ also collapses to FO on $\mathcal{BFR}$.

Theorem 3 *Let $\varphi(x_1, x_2, R)$ be a Δ_0 formula that is positive in the binary relation symbol R. The least fixed-point of φ and the greatest fixed-point of φ are first-order definable on $\mathcal{BFR}$. In fact, the least fixed-point is Π-definable, whereas the greatest fixed-point is Σ-definable on $\mathcal{BFR}$.*

Proof (sketch): For simplicity, we only sketch the proof for the collapse of binary $\mathrm{LFP}(\Delta_0)$ to FO on the smaller class $\mathcal{FR}$ of finite ranks. After presenting the sketch, we outline how the proof can be modified to establish the collapse on the class $\mathcal{BFR}$.

As in the proof of Theorem 2, the first key idea is that the greatest fixed-point $J_\varphi(V_n)$ of φ on V_n can be approximated by a first-order formula. In fact, we can start with the Σ formula $\theta(x_1, x_2)$ featured in that proof, because Claim 1 uses only the hypothesis that φ is a Δ_0 formula (and not the additional hypothesis in Theorem 2 that φ is restricted). Thus $J_\varphi(V_n) \cap V_{n-3} \subseteq \theta^{V_n} \subseteq J_\varphi(V_n)$. Moreover, it is not hard to verify that $\theta^{V_n} = J_\varphi(V_{n-3})$, because V_n is an actual finite rank (instead of a near finite rank). Our goal is to improve on this approximation of $J(V_n)$ by defining Σ formulas θ_i such that $J_\varphi(V_{n-3+i}) = \theta_i^{V_n}$, for $i = 1, 2, 3$. Since $\varphi(x_1, x_2, R)$ may not be a restricted formula, a difficulty arises from the potential presence in φ of subformulas of the form $R(x_i, x_j)$, $R(x_i, z)$, and $R(z, x_i)$, where $i, j \in \{1, 2\}$ and z is a bound variable of φ. Actually, in building the formulas θ_i, the most serious difficulty is caused by the subformulas $R(x_i, z)$ and $R(z, x_i)$. Note that, since φ is Δ_0 and z is a bound variable of φ, every element of V_n witnessing z must be in V_{n-1}. Therefore, for every choice of x_i, the set of elements of V_n witnessing $R(x_i, z)$ (or $R(z, x_i)$) is a subset of V_{n-1} and, hence, a member of V_n. In turn, this makes it possible to use first-order existential quantifiers over V_n to quantify the set of all such elements witnessing z. Note that this would not be possible, if we had to work with an arbitrary near finite rank M, as the above set may not be in M.

We will build the desired Σ formula θ_i from θ_{i-1}, for $i = 1, 2, 3$, where we take θ_0 to be θ. Let S be the set $\{11, 12, 21, 22\}$. For each $T \subseteq S$ containing 12, we will define a formula $\delta_T(x_1, x_2)$. Then each $\theta_i(x_1, x_2)$ will be the formula

$$\beta_{3-i}(x_1) \wedge \beta_{3-i}(x_2) \wedge \Big(\bigvee_{T \in \mathcal{P}} \delta_T(x_1, x_2)\Big),$$

where each $\beta_k(x_i)$ is the formula $(\exists z_1) \ldots (\exists z_k)(x_i \in z_1 \wedge \bigwedge_{l=1}^{k-1} z_l \in z_{l+1})$ and $\mathcal{P}$ is the set of subsets of S containing 12. The first two subformulas of $\theta_i(x_1, x_2)$ are introduced to ensure that x_1 and x_2 belong to V_{n-3+i}. Let O be the set $\{10, 01, 20, 02\}$ and let $y_{j_1 j_2}$ be a new variable for each $j_1 j_2 \in O$. From now on, we use the abbreviations $\overline{x} = (x_1, x_2)$, $\overline{y} = (y_{10}, y_{01}, y_{20}, y_{02})$, and $\overline{z} = (z_1, z_2)$. Let $\delta_T(\overline{x})$ be the formula $(\exists \overline{y})(\delta'_T(\overline{x}, \overline{y}) \wedge \delta''_T(\overline{x}, \overline{y}))$ where

$$\begin{aligned} \delta'_T &\equiv \textstyle\bigwedge_{j_1 j_2 \in O} \rho_{j_1 j_2}(\overline{x}, \overline{y}), \\ \delta''_T &\equiv \textstyle\bigwedge_{i_1 i_2 \in T} \varphi(x_{i_1}, x_{i_2}, R(\overline{z})/\sigma_T(\overline{z}, \overline{x}, \overline{y})), \end{aligned}$$

and $\rho_{j_1 j_2}(\overline{x}, \overline{y})$ is the formula

$$(\forall x_0 \in y_{j_1 j_2})\varphi(x_{j_1}, x_{j_2}, R(\overline{z})/\sigma_T(\overline{z}, \overline{x}, \overline{y})),$$

while $\sigma_T(\overline{z}, \overline{x}, \overline{y})$ is the formula $\theta_{i-1}(\overline{z}) \vee \sigma'_T(\overline{z}, \overline{x}, \overline{y}) \vee \sigma''_T(\overline{z}, \overline{x}, \overline{y})$ where

$$\begin{aligned} \sigma'_T &\equiv \textstyle\bigvee_{j_1 j_2 \in O}(\exists x_0 \in y_{j_1 j_2})(z_1 = x_{j_1} \wedge z_2 = x_{j_2}), \\ \sigma''_T &\equiv \textstyle\bigvee_{i_1 i_2 \in T}(z_1 = x_{i_1} \wedge z_2 = x_{i_2}), \end{aligned}$$

with x_0 a fresh variable. Note that $\rho_{j_1 j_2}$ says that $y_{j_1 j_2}$ is the set of witnesses for bound variables that relate to x_1 or to x_2. Since θ_0 is a Σ formula, it is easy to see that each θ_i is also a Σ formula.

Claim 3: $J_\varphi(V_{n-3+i}) = \theta_i^{V_n}$ for $i = 0, \ldots, 3$.

Proof of Claim 3: The claim holds for $i = 0$, since θ_0 is θ. Fix $i \in \{1, 2, 3\}$ and assume that the claim holds for $i - 1$; we show that it holds for i. We show first that $\theta_i^{V_n} \subseteq J_\varphi(V_{n-3+i})$. Let $\overline{a} = (a_1, a_2) \in V_n^2$ be such that

$V_n \models \theta_i[\overline{a}]$. From the definition of θ_i, it follows that $\overline{a} \in V^2_{n-3+i}$, and $V_n \models \delta_T[\overline{a}]$ for some $T \subseteq S$ with $12 \in T$. For every $j_1j_2 \in O$, let $b_{j_1j_2} \in V_n$ be a witness for the quantifier $\exists y_{j_1j_2}$ in δ_T, and let $\overline{b}$ be the corresponding witness for $\overline{y}$. Since $12 \in T$, we have that $V_n \models \varphi[a_1, a_2, \sigma_T^{V_n}(\overline{a}, \overline{b})]$. It suffices to show that $\sigma_T^{V_n}(\overline{a}, \overline{b})$ is a post fixed-point of φ on V_n; then, the monotonicity of φ and Lemma 3 will imply that $\overline{a}$ is in $J_\varphi(V_{n-3+i})$. We need to verify that $\sigma_T^{V_n}(\overline{a}, \overline{b}) \subseteq \varphi^{V_n}(\sigma_T^{V_n}(\overline{a}, \overline{b}))$. Let $\overline{c} \in V_n^2$ be such that $V_n \models \sigma_T[\overline{c}, \overline{a}, \overline{b}]$. Then $\overline{c}$ must satisfy one of the three disjuncts of σ_T. Assume first that $V_n \models \theta_{i-1}[\overline{c}]$. Using Lemmas 2 and 3 and the induction hypothesis that $\theta_{i-1}^{V_n} = J_\varphi(V_{n-3+i-1})$, it can be verified that $\theta_{i-1}^{V_n}$ is a post fixed-point of φ on V_n. Thus $V_n \models \varphi[\overline{c}, \theta_{i-1}^{V_n}]$. But $\theta_{i-1}^{V_n} \subseteq \sigma_T^{V_n}(\overline{a}, \overline{b})$ and hence, by monotonicity, $V_n \models \varphi[\overline{c}, \sigma_T^{V_n}(\overline{a}, \overline{b})]$, as required. Assume now that $\overline{c}$ satisfies the second disjunct. Then, there exist $j_1j_2 \in O$ and $a_0 \in b_{j_1j_2}$ such that $\overline{c} = (a_{j_1}, a_{j_2})$. From the choice of $b_{j_1j_2}$ and the definition of $\rho_{j_1j_2}$, we know that $V_n \models \varphi[a_{j_1}, a_{j_2}, \sigma_T^{V_n}(\overline{a}, \overline{b})]$; therefore, $V_n \models \varphi[\overline{c}, \sigma_T^{V_n}(\overline{a}, \overline{b})]$, as required again. The case of the third disjunct is handled in a similar manner.

Consider next the inclusion $J_\varphi(V_{n-3+i}) \subseteq \theta_i^{V_n}$. If $\overline{a} \in J_\varphi(V_{n-3+i})$, then $V_n \models \varphi[\overline{a}, J_\varphi(V_{n-3+i})]$. Let T be the biggest subset of S containing 12 and such that for every $i_1i_2 \in T$ it is the case that $V_n \models \varphi[a_{i_1}, a_{i_2}, J_\varphi(V_{n-3+i})]$. Let A be the set $J_\varphi(V_{n-3+i}) \cap (V_{n-3+i-1} \cup \{a_1, a_2\})^2$. By Lemma 2, for every $i_1i_2 \in T$ we have that $V_n \models \varphi[a_{i_1}, a_{i_2}, A]$. We construct witnesses $b_{j_1j_2}$ for $y_{j_1j_2}$ in δ_T such that $\sigma_T^{V_n}(\overline{a}, \overline{b}) = A$, which will prove the claim. Fix $j_1j_2 \in O$ and define $b_{j_1j_2}$ as follows: if $j_1 = 0$, then $b_{j_1j_2} = \{a \in V_{n-3+i-1} : (a, a_{j_2}) \in A\}$; if $j_2 = 0$, then $b_{j_1j_2} = \{a \in V_{n-3+i-1} : (a_{j_1}, a) \in A\}$. Using the induction hypothesis and the definition of A, it can be checked that the sets $b_{j_1j_2}$ have the aforementioned properties.

This concludes the proof that the Σ formula $\theta_3(x_1, x_2)$ defines the greatest fixed-point $J_\varphi(V_n)$ of $\varphi(x_1, x_2, R)$ on $\mathcal{FR}$. By applying the same argument to the dual formula $\tilde{\varphi}(x_1, x_2, R)$, we establish that the least fixed-point $I_\varphi(V_n)$ of $\varphi(x_1, x_2, R)$ is Π definable on $\mathcal{FR}$.

Finally, we comment on the modifications needed to extend the proof to the class $\mathcal{BFR}$. The first modification is that we have to add an extra iteration in the construction of the first-order formula that defines the greatest fixed-point of $\varphi(x_1, x_2, R)$ on $\mathcal{BFR}$. Specifically, instead of four formulas, we will need five formulas $\theta_0, \theta_1, \theta_2, \theta_3, \theta_4$; the greatest fixed-point will be defined by the last formula θ_4.

In proving that binary $\mathrm{LFP}(\Delta_0)$ collapses to FO on $\mathcal{FR}$, we used twice the assumption that we were working with structures of the form V_n for some $n \geq 1$, instead of structures $M \in \mathcal{BFR}$ with $V_n \subseteq M \subseteq V_{n+1}$ for some $n \geq 1$. The first time was to ensure that $\theta^M = J_\varphi(V_{n-3})$. This can be taken care of by considering the formula $\tilde{\theta}(x_1, x_2) \equiv \tilde{\beta}_3(x_1) \wedge \tilde{\beta}_3(x_2) \wedge \theta(x_1, x_2)$. Here, $\tilde{\beta}_i(x)$ is a Σ formula similar to the formula $\beta_i(x)$ defined in the proof, whose intended meaning is that $\tilde{\beta}_i^M = V_{n-i}$ for every $M \in \mathcal{BFR}$ such that $V_n \subseteq M \subset V_{n+1}$. For the rest of the proof, $\tilde{\beta}_i(x)$ should be used in place of $\beta_i(x)$.

The second time we used the assumption that $M = V_n$ for some $n \geq 1$ was to ensure that the existentially quantified variables $y_{10}, \ldots, y_{02}$ can be witnessed by elements of the universe of the structure M. As mentioned earlier, these witnesses may not exist within an arbitrary $M \in \mathcal{NFR}$, or even an arbitrary $M \in \mathcal{BFR}$. When showing that binary $\mathrm{LFP}(\Delta_0)$ collapses to FO on $\mathcal{BFR}$, this difficulty is only encountered in the last iteration, that is, in the correctness of the formula θ_4. This obstacle, however, can be overcome by using the fact that the witnesses can be restricted to be subsets of the *transitive closure* of the arguments x_1, x_2, where the transitive closure of a set a is defined inductively as follows: $\mathrm{TC}(a) = a \cup \bigcup\{\mathrm{TC}(b) : b \in a\}$. The reason is that, since φ is a Δ_0 formula, the interpretation of every bound variable of φ can be restricted to the transitive closure of the sets that interpret the free variables of φ. Consequently, we may split the witnesses b_{ij} defined in the proof into three sets $b_{ij}^1, b_{ij}^2, b_{ij}^3$ with the following interpretations: $b_{ij}^1 = b_{ij} \cap a_1$, $b_{ij}^2 = b_{ij} \cap a_2$, and $b_{ij}^3 = b_{ij} \cap ((\mathrm{TC}(a_1) \cup \mathrm{TC}(a_2)) - (a_1 \cup a_2))$. Observe that if $M \in \mathcal{BFR}$ and $a_1, a_2 \in M$, then $b_{ij}^1, b_{ij}^2 \in M$, because M is closed under taking subsets (if $a \subseteq b$, then $e^{-1}(a) \leq e^{-1}(b)$); observe also that $b_{ij}^3 \in M$, because its rank is less than the maximum of the ranks of a and b. The sets b_{ij}^1, b_{ij}^2, and b_{ij}^3 capture all relevant information about b_{ij}; therefore, only these elements need to be existentially quantified. Clearly, several changes in the formulas have to be made; the technical details will be included in the full version of the paper. □

The inquisitive reader may wonder whether the argument of Theorem 3 extends to higher arities. The answer is that it does *not*, for the reason that we cannot encode binary relations on V_{n-1} by elements of V_n, whereas it is possible to encode unary relations (sets) on V_{n-1} by elements of V_n. We note, however, that Theorem 3 extends to binary Δ_0 formulas over an expanded vocabulary that, in addition to $\in$, contains other relation symbols, as long as they are interpreted by Δ-definable queries on $\mathcal{BFR}$.

Example 2: The aforementioned extension of Theorem 3 can be used to show that the *Even Cardinality* query

$$Q_{\text{even}}(M) = \{a \in M : \text{the cardinality of } a \text{ is even}\}$$

is first-order definable on $\mathcal{BFR}$. For this, one can write a Δ_0 formula $\varphi(x, y, R)$ over the vocabulary $\{\in, <, R\}$, whose least fixed-point defines the binary query "y is an even element of x with respect to the linear order $<$" on

$\mathcal{BFR}$. Here, $<$ is interpreted by the Δ-definable linear order on $\mathcal{BFR}$ described in the beginning of Section 4.

6. Closure functions of Δ_0 formulas

Let $\varphi(x_1, \ldots, x_k, R)$ be an arbitrary positive first-order formula. From the preliminaries, recall that if $\mathbf{M}$ is a finite structure, then $\mathrm{cl}_\varphi(\mathbf{M})$ is the smallest integer m such that $I^m_\varphi(\mathbf{M}) = \bigcup_{m'<m} I^{m'}_\varphi(\mathbf{M})$. In general, the rate of growth of $\mathrm{cl}_\varphi(\mathbf{M})$ can be as high as a polynomial in the cardinality of the universe of $\mathbf{M}$. Here, we analyze the rate of growth of the closure function of positive Δ_0 formulas on finite ranks $(V_n, \in)$, $n \geq 1$. We establish that if φ is a positive Δ_0 formula, then $\mathrm{cl}_\varphi(V_n)$ is bounded by a polylogarithm of the cardinality $|V_n|$ of V_n. Moreover, we show that if φ is a restricted Δ_0 formula, then $\mathrm{cl}_\varphi(V_n)$ is bounded by the iterated logarithm of $|V_n|$.

Theorem 4 *Let $\varphi(x_1, \ldots, x_k, R)$ be a Δ_0 formula that is positive in the k-ary relation symbol R. Then, for all sufficiently large $n > 0$, we have*

$$\mathrm{cl}_\varphi(V_n) \leq nk^{2k+1}|V_{n-1}|^{k-1} \leq \log^k(|V_n|).$$

Moreover, if φ is a restricted Δ_0 formula, then $\mathrm{cl}_\varphi(V_n) \leq n \leq 1 + \log^*(|V_n|)$.

Proof (Sketch): We outline the proof of the first statement only; the second has an easier proof. We show that $\mathrm{cl}_\varphi(V_n) \leq \mathrm{cl}_\varphi(V_{n-1}) + k^{2k+1}|V_{n-1}|^{k-1}$ holds for all $n \geq 1$. The result will follow by expanding this recurrence and using the fact that $nk^{2k+1}|V_{n-1}|^{k-1} \leq \log^k(|V_n|)$ for all sufficiently large n. Put $t = k^{2k+1}|V_{n-1}|^{k-1}$, and let m be the closure ordinal of φ on V_{n-1}. It is enough to show that $I^{m+t+1}_\varphi(V_n) \subseteq I^{m+t}_\varphi(V_n)$. So, let us assume that $\overline{a} \in I^{m+t+1}_\varphi(V_n)$, which means that $V_n \models \varphi[\overline{a}, I^{m+t}_\varphi(V_n)]$. Lemma 2 implies that $V_n \models \varphi[\overline{a}, I^{m+t}_\varphi(V_n) \cap (V_{n-1} \cup \{a_1, \ldots, a_k\})^k]$. We claim that $I^{m+t}_\varphi(V_n) \cap (V_{n-1} \cup \{a_1, \ldots, a_k\})^k \subseteq I^{m+t-1}_\varphi(V_n)$. This will prove our goal, because the monotonicity of φ implies that $V_n \models \varphi[\overline{a}, I^{m+t-1}_\varphi(V_n)]$ and, therefore, $\overline{a} \in I^{m+t}_\varphi(V_n)$.

Using Lemma 3 and the choice of m, it is easy to show that $I^{m+t}_\varphi(V_n) \cap V^k_{n-1} \subseteq I^{m+t-1}_\varphi(V_n)$. So, to prove the claim, it remains to consider those tuples in $I^{m+t}_\varphi(V_n) \cap (V_{n-1} \cup \{a_1, \ldots, a_k\})^k$ that are not in $I^{m+t}_\varphi(V_n) \cap V^k_{n-1}$. For each $s \geq 0$, let A^s be the set $I^{m+s}_\varphi(V_n) \cap (V_{n-1} \cup \{a_1, \ldots, a_k\})^k - I^{m+s}_\varphi(V_n) \cap V^k_{n-1}$. First note that $|A^t| \leq |(V_{n-1} \cup \{a_1, \ldots, a_k\})^k - V^k_{n-1}|$. A simple counting argument shows that the cardinality of the set $(V_{n-1} \cup \{a_1, \ldots, a_k\})^k - V^k_{n-1}$ is

$$\sum_{i=1}^{k} \binom{k}{i} k^i |V_{n-1}|^{k-i} \leq \sum_{i=1}^{k} k^k k^k |V_{n-1}|^{k-1} = t.$$

Therefore $|A^t| \leq t$. Now let $r \geq 0$ be the smallest integer such that $A^r = A^{r+1}$. Such an r exists, because $I^{m+s}_\varphi(V_n)$ eventually reaches the fixed-point $I_\varphi(V_n)$. If $r \geq t$, then the sequence of proper inclusions $\emptyset \subset A^0 \subset A^1 \subset \cdots \subset A^t$ holds. Hence $|A^t| > t$, which contradicts the fact that $|A^t| \leq t$. Thus, $r < t$ must hold, from which it follows that $A^t = A^{r+1} = A^r = A^{t-1} \subseteq I^{m+t-1}_\varphi(V_n)$. $\Box$

In the full paper, we provide examples showing that the above bounds are tight. As regards restricted Δ_0 formulas, the formula $\psi(x, y, S)$ for the linear order is such an example.

We conclude by pointing out that the above Theorem 4 implies that every $\mathrm{LFP}(\Delta_0)$-definable query on $\mathcal{FR}$ is in NC, the parallel complexity class consisting of all queries computable in polylogarithmic time using a polynomial number of processors (see [Pap94, GHR95, BDG90] for a thorough coverage of NC). This suggests the systematic pursuit of descriptive complexity in the context of finite set theory.

Acknowledgment: We are grateful to Moshe Y. Vardi for his constructive criticism and comments on an earlier draft of this paper.

References

[Bar75] J. Barwise. *Admissible Sets and Structures*. Springer-Verlag, 1975.

[BDG90] J. L. Balcázar, J. Díaz, and J. Gabarró. *Structural Complexity II*. Springer-Verlag, 1990.

[BIS90] D. A. M. Barrington, N. Immerman, and H. Straubing. On uniformity within NC^1. *Journal of Computer and System Sciences*, 41:274–306, 1990.

[DDLW98] A. Dawar, K. Doets, S. Lindell, and S. Weinstein. Elementary properties of finite ranks. *Mathematical Logic Quarterly*, 44:349–353, 1998.

[DH95] A. Dawar and L. Hella. The expressive power of finitely many generalized quantifiers. *Information and Control*, 123:172–184, 1995.

[DLW96] A. Dawar, S. Lindell, and S. Weinstein. First order logic, fixed point logic and linear order. In *Computer Science Logic '95*, volume 1092 of *Lecture Notes in Computer Science*, pages 161–177. Springer-Verlag, 1996.

[GHR95] R. Greenlaw, H. J. Hoover, and W. L. Ruzzo. *Limits to parallel computation : P-completeness theory*. Oxford University Press, 1995.

[GIS94] Y. Gurevich, N. Immerman, and S. Shelah. McColm's conjecture. In *Proc. 7th IEEE Symp. on Logic in Comp. Sci.*, pages 10–19, 1994.

[GS86] Y. Gurevich and S. Shelah. Fixed-point extensions of first-order logic. *Annals of Pure and Applied Logic*, 32(3):265–280, 1986.

[Imm86] N. Immerman. Relational queries computable in polynomial time. *Information and Control*, 68:86–104, 1986.

[Imm89] N. Immerman. Expressibility and parallel complexity. *SIAM Journal on Computing*, 18:625–638, 1989.

[KV92] Ph. G. Kolaitis and M. Y. Vardi. Fixpoint logic vs. infinitary logic in finite-model theory. In *Proc. 6th IEEE Symp. on Logic in Comp. Sci.*, pages 46–57, 1992.

[KV96] Ph. G. Kolaitis and M. Y. Vardi. On the expressive power of variable-confined logics. In *Proc. 9th IEEE Symp. on Logic in Comp. Sci.*, pages 348–359, 1996.

[Lin92] S. Lindell. A purely logical characterization of circuit uniformity. In *Proc. 7th Annual IEEE Structure in Complexity Theory Conference*, pages 185–192, 1992.

[McC90] G. L. McColm. When is arithmetic possible? *Annals of Pure and Applied Logic*, 50:29–51, 1990.

[Mos74] Y. N. Moschovakis. *Elementary Induction on Abstract Structures*. North Holland, 1974.

[Pap94] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley Publishing Company, 1994.

[Saz97] V. Yu. Sazonov. On bounded set theory. In M.L. Dalla Chiara et al., editor, *Logic and Scientific Methods, Volume One of the Tenth ICLMPS*, pages 85–103. Kluwer, 1997.

[Var82] M. Y. Vardi. The complexity of relational query languages. In *Proc. 14th ACM Symp. on Theory of Computing*, pages 137–146, 1982.