The complexity of quantified constraints: collapsibility, switchability and the algebraic formulation

CATARINA CARVALHO, Department of Physics, Astronomy and Mathematics, University of Hertfordshire, UK
FLORENT MADELAINE, LACL, Université Paris-Est Créteil, France
BARNABY MARTIN, Department of Computer Science, Durham University, UK
DMITRIY ZHUK, Department of Mechanics and Mathematics, Lomonosov Moscow State University, Russia

Let $\mathbb{A}$ be an idempotent algebra on a finite domain. By mediating between results of Chen [1] and Zhuk [2], we argue that if $\mathbb{A}$ satisfies the polynomially generated powers property (PGP) and $\mathcal{B}$ is a constraint language invariant under $\mathbb{A}$ (that is, in $\mathrm{Inv}(\mathbb{A})$), then $\mathrm{QCSP}(\mathcal{B})$ is in NP. In doing this we study the special forms of PGP, switchability and collapsibility, in detail, both algebraically and logically, addressing various questions such as decidability on the way.

We then prove a complexity-theoretic converse in the case of infinite constraint languages encoded in propositional logic, that if $\mathrm{Inv}(\mathbb{A})$ satisfies the exponentially generated powers property (EGP), then $\mathrm{QCSP}(\mathrm{Inv}(\mathbb{A}))$ is co-NP-hard. Since Zhuk proved that only PGP and EGP are possible, we derive a full dichotomy for the QCSP, justifying what we term the Revised Chen Conjecture. This result becomes more significant now the original Chen Conjecture (see [3]) is known to be false [4].

Switchability was introduced by Chen in [1] as a generalisation of the already-known collapsibility [5]. There, an algebra $\mathbb{A} := (\{0, 1, 2\}; r)$ was given that is switchable and not collapsible. We prove that, for all finite subsets Δ of $\mathrm{Inv}(\mathbb{A})$, $\mathrm{Pol}(\Delta)$ is collapsible. The significance of this is that, for QCSP on finite structures, it is still possible all QCSP tractability (in NP) explained by switchability is already explained by collapsibility. At least, no counterexample is known to this.

CCS Concepts: • **Theory of computation** → *Design and analysis of algorithms*; *Logic*; *Computational complexity and cryptography*.

Additional Key Words and Phrases: quantified constraints, constraint satisfaction, logic, universal algebra, computational complexity

ACM Reference Format:
Catarina Carvalho, Florent Madelaine, Barnaby Martin, and Dmitriy Zhuk. 2018. The complexity of quantified constraints: collapsibility, switchability and the algebraic formulation. *ACM Trans. Comput. Logic* 37, 4, Article 111 (August 2018), 25 pages. https://doi.org/10.1145/1122445.1122456

Authors' addresses: Catarina Carvalho, c.carvalho2@herts.ac.uk, Department of Physics, Astronomy and Mathematics, University of Hertfordshire, Hatfield, UK; Florent Madelaine, florent.madelaine@gmail.com, LACL, Université Paris-Est Créteil, Créteil, France; Barnaby Martin, barnabymartin@gmail.com, Department of Computer Science, Durham University, Durham, UK; Dmitriy Zhuk, zhuk.dmitriy@gmail.com, Department of Mechanics and Mathematics, Lomonosov Moscow State University, Moscow, Russia.

1529-3785/2018/8-ART111 $15.00
https://doi.org/10.1145/1122445.1122456

1 INTRODUCTION

A large body of work exists from the past twenty years on applications of universal algebra to the computational complexity of *constraint satisfaction problems* (CSPs) – see for example the surveys [6–8] – and a number of celebrated results have been obtained through this approach. One considers the problem CSP($\mathcal{B}$) in which it is asked whether an input sentence φ holds on $\mathcal{B}$, a *constraint language* (equivalently, relational structure), where φ is *primitive positive*, that is using only $\exists$, $\land$ and $=$. The CSP is one of a wide class of model-checking problems obtained from restrictions of first-order logic. For almost all of these classes, we can give a complexity classification [9]. Chief among these celebrated results are the proofs of the Feder-Vardi "Dichotomy" Conjecture for CSPs [10–12]. The only outstanding class (other than its natural dual) is *quantified CSPs* (QCSPs) for positive Horn sentences – where $\forall$ is also present – which is used in Artificial Intelligence to model non-monotone reasoning or uncertainty [13].

It is well-known in folklore that the complexity classification for QCSPs embeds the classification for CSPs: if $\mathcal{B}+1$ is $\mathcal{B}$ with the addition of a new isolated element not appearing in any relations, then CSP($\mathcal{B}$) and QCSP($\mathcal{B}+1$) are polynomially equivalent. Thus the classification for QCSPs may be considered a project at least as hard as that for CSPs.

The algebraic approach to (Q)CSPs comes from a certain interplay between operations and relations. We say that a k-ary operation f *preserves* an m-ary relation R, whenever $(x_1^1, \dots, x_1^m), \dots, (x_k^1, \dots, x_k^m)$ in R, then also $(f(x_1^1, \dots, x_k^1), \dots, f(x_1^m, \dots, x_k^m))$ in R. The relation R is called *an invariant* of f, and the operation f is called *a polymorphism* of R. An operation f is a polymorphism of $\mathcal{B}$ if it preserves every relation from $\mathcal{B}$. Likewise, a relation R is an invariant of an algebra $\mathbb{A}$ if it is preserved by every operation of $\mathbb{A}$. We can also think of an invariant R as a subalgebra of a direct power of $\mathbb{A}$. We denote the set of polymorphisms of $\mathcal{B}$ by Pol($\mathcal{B}$) and the set of invariants of $\mathbb{A}$ as Inv($\mathbb{A}$).

For a finite-domain algebra $\mathbb{A}$ we associate a function $f_{\mathbb{A}} : \mathbb{N} \to \mathbb{N}$, giving the cardinality of the minimal generating sets of the sequence $\mathbb{A}, \mathbb{A}^2, \mathbb{A}^3, \dots$ as $f_{\mathbb{A}}(1), f_{\mathbb{A}}(2), f_{\mathbb{A}}(3), \dots$, respectively. A subset Λ of A^m is a generating set for $\mathbb{A}^m$ exactly if, for every $(a_1, \dots, a_m) \in A^m$, there exists a k-ary term operation f of $\mathbb{A}$ and $(b_1^1, \dots, b_m^1), \dots, (b_1^k, \dots, b_m^k) \in \Lambda$ so that $f(b_1^1, \dots, b_1^k) = a_1, \dots, f(b_m^1, \dots, b_m^k) = a_m$. We may say $\mathbb{A}$ has the g-GP if $f_{\mathbb{A}}(m) \leq g(m)$ for all m. The question then arises as to the growth rate of $f_{\mathbb{A}}$ and specifically regarding the behaviours constant, logarithmic, linear, polynomial and exponential. Wiegold proved in [14] that if $\mathbb{A}$ is a finite semigroup then $f_{\mathbb{A}}$ is either linear or exponential, with the former prevailing precisely when $\mathbb{A}$ is a monoid. This dichotomy classification may be seen as a gap theorem because no growth rates intermediate between linear and exponential may occur. We say $\mathbb{A}$ enjoys the *polynomially generated powers* property (PGP) if there exists a polynomial p so that $f_{\mathbb{A}} = O(p)$ and the *exponentially generated powers* property (EGP) if there exists a constant $b > 1$ so that $f_{\mathbb{A}} = \Omega(g)$ where $g(i) = b^i$.

The following is the merger of Conjectures 6 and 7 in [3] which we call the *Chen Conjecture.*

Conjecture 1 (Chen Conjecture). *Let $\mathcal{B}$ be a finite relational structure expanded with constants naming all the elements. If Pol($\mathcal{B}$) has PGP, then QCSP($\mathcal{B}$) is in NP; otherwise QCSP($\mathcal{B}$) is Pspace-complete.*

Conjecture 6 in [3] gives the NP membership and Conjecture 7 in [3] gives the Pspace-completeness. The first contribution of this paper is to prove that the NP membership of Conjecture 6 is indeed true. We do this by proving equivalent two notions of switchability that allows to combine known results from [1] and [2]. On the way we develop the notions of non-degenerate and projective adversaries that enable us to prove our result as well as particular observations on the existing notions of switchability and collapsibility. Let us recall that the Chen Conjecture is now known to be false [4].

The second contribution of this paper is Theorem 2 below, but note that we permit infinite signatures (languages) although our domains remain finite. This will involve deciding how to encode relations of $\text{Inv}(\mathbb{A})$ and will be discussed in detail later.

Theorem 2 (Revised Chen Conjecture). *Let $\mathbb{A}$ be an idempotent algebra on a finite domain A. If $\mathbb{A}$ satisfies PGP, then QCSP($\text{Inv}(\mathbb{A})$) is in NP. Otherwise, QCSP($\text{Inv}(\mathbb{A})$) is co-NP-hard.*

Note that, with infinite languages, the NP-membership for Theorem 2 requires a little extra work.
he third contribution of this paper, concerns another variant we dub the Alternative Chen Conjecture which was not posed by Chen himself but is nonetheless natural.

Conjecture 3 (Alternative Chen Conjecture). *Let $\mathbb{A}$ be an idempotent algebra on a finite domain A. If $\mathbb{A}$ satisfies PGP, then for every finite subset $\Delta \subset \text{Inv}(\mathbb{A})$, QCSP($\Delta$) is in NP. Otherwise, there exists a finite subset $\Delta \subset \text{Inv}(\mathbb{A})$ so that QCSP(Δ) is co-NP-hard.*

In Proposition 42 we present an example that refutes the second part of the Alternative Chen Conjecture.

In proving Theorem 2 we are saying that the complexity of QCSPs, with all constants included, is classified modulo the complexity of (infinite signature) CSPs, a subject to which we will return later. The following is a corollary to Theorem 2.

Corollary 4. *Let $\mathbb{A}$ be an idempotent algebra on a finite domain A. Either QCSP($\text{Inv}(\mathbb{A})$) is co-NP-hard or QCSP($\text{Inv}(\mathbb{A})$) has the same complexity as CSP($\text{Inv}(\mathbb{A})$).*

In this manner, our result follows in the footsteps of the similar result for the Valued CSP, which has also had its complexity classified modulo the CSP, as culminated in the paper [15].

In Chen's [1], a new link between algebra and QCSP was discovered. Chen's previous work in QCSP tractability largely involved the special notion of *collapsibility* [5], but in [1] this was extended to a *computationally effective* version of the PGP. For a finite-domain, idempotent algebra $\mathbb{A}$, call *simple k-collapsibility*[1] that special form of the PGP in which the generating set for $\mathbb{A}^m$ is constituted of all tuples $(x_1, \ldots, x_m)$ in which at least $m - k$ of these elements are equal. *Simple k-switchability* will be another special form of the PGP in which the generating set for $\mathbb{A}^m$ is constituted of all tuples $(x_1, \ldots, x_m)$ in which there exist $a_i < \ldots < a_{k'}$, for $k' \leq k$, so that

$$(x_1, \ldots, x_m) = (x_1, \ldots, x_{a_1}, x_{a_1+1}, \ldots, x_{a_2}, x_{a_2+1}, \ldots, \ldots, \ldots, x_{a_{k'}}, x_{a_{k'+1}}, \ldots, x_m),$$

where $x_1 = \ldots = x_{a_1-1}$, $x_{a_1} = \ldots = x_{a_2-1}$, $\ldots$, $x_{a_{k'}} = \ldots = x_m$. Thus, $a_1, a_2, \ldots, a_{k'}$ are the indices where the tuple switches value. We say that $\mathbb{A}$ is simply collapsible (switchable) if there exists k such that it is simply k-collapsible (k-switchable). We note that Zhuk uses this form of simple switchability, in [2], where he proves that the only kind of PGP for finite-domain algebras is simple switchability.

Our first contribution shows k-collapsibility, whose definition is deferred until adversaries are introduced in Section 2, and simple k-collapsibility, coincide. The same applies to k-switchability and simple k-switchability, and we will dwell on these distinctions no longer. For any finite algebra, k-collapsibility implies k-switchability, and for any 2-element algebra, k-switchability implies k-collapsibility (this latter fact is only known a posteriori).

Switchability was introduced by Chen in [1] as a generalisation of the already-known collapsibility [5] when he discovered a 4-ary operation r on the three-element domain so that $(\{0, 1, 2\}; r)$ has the PGP (switchability) but is not collapsible. Thus it seemed that collapsibility was not enough to

[1]We want to use a name different from "collapsibility" alone in order to differentiate this from Chen's original definition. In [16] we used capitalisation, with a leading capital letter for Chen's original version and all small letters for what we here designate simple.

explain membership of QCSP in NP. What we prove as our fourth contribution is that $\mathrm{Inv}(\{0,1,2\};r)$ is not finitely related, and what is more, every finite subset Δ of $\mathrm{Inv}(\{0,1,2\};r)$ is such that $\mathrm{Pol}(\Delta)$ is collapsible. Note that the parameter k of collapsibility is unbounded over these increasing finite subsets while the parameter of switchability clearly remains bounded.

1.1 Infinite languages

Our use of infinite languages (i.e. infinite signatures, since we work on a finite domain) is a controversial part of our discourse and merits special discussion. We wish to argue that a necessary corollary of the algebraic approach to (Q)CSP is a reconciliation with infinite languages. The traditional approach to consider arbitrary finite subsets of $\mathrm{Inv}(\mathbb{A})$ is unsatisfactory in the sense that choosing this way to escape the – naturally infinite – set $\mathrm{Inv}(\mathbb{A})$ is as arbitrary a choice as the choice of encoding required for infinite languages. However, the difficulty in that choice is of course the reason why this route is often eschewed. The first possibility that comes to mind for encoding a relation in $\mathrm{Inv}(\mathbb{A})$ is probably to list its tuples, while the second is likely to be to describe the relation in some kind of "simple" logic. Both these possibilities are discussed in [17], for the Boolean domain, where the "simple" logic is the propositional calculus. For larger domains, this would be equivalent to quantifier-free propositions over equality with constants. Both Conjunctive Normal Form (CNF) and Disjunctive Normal Form (DNF) representations are considered in [17] and a similar discussion in [18] exposes the advantages of the DNF encoding. The point here is that testing non-emptiness of a relation encoded in CNF may already be NP-hard, while for DNF this will be tractable. Since DNF has some benign properties, we might consider it a "nice, simple" logic while for "simple" logic we encompass all quantifier-free sentences, that include DNF and CNF as special cases. The reason we describe this as "simple" logic is to compare against something stronger, say all first-order sentences over equality with constants. Here recognising non-emptiness becomes Pspace-hard and since QCSPs already sit in Pspace, this complexity is unreasonable.

For the QCSP over infinite languages $\mathrm{Inv}(\mathbb{A})$, Chen and Mayr [19] have declared for our first, tuple-listing, encoding. In this paper we will choose the "simple" logic encoding, occasionally giving more refined results for its "nice, simple" restriction to DNF. Our choice of the "simple" logic encoding over the tuple-listing encoding will ultimately be justified by the (Revised) Chen Conjecture holding for "simple" logic yet failing for tuple-listings. Since the original Chen Conjecture is known now to be false [4], our result becomes more remarkable. However, there are some surprising consequences. It follows from [4] that there exists a finite and 3-element $\mathcal{B}$ with constants, so that $\mathrm{QCSP}(\mathrm{Inv}(\mathrm{Pol}(\mathcal{B})))$, under our encoding, and $\mathrm{QCSP}(\mathcal{B})$ have different complexities: the former being co-NP-hard while the latter is in P.

The Feder-Vardi Conjecture for CSPs is known to hold for infinite languages [20] but the proofs are based on the tuple-listing encoding. We cannot say whether the polynomial cases are preserved under the DNF encoding.

Let us consider examples of our encodings. For the domain $\{1,2,3\}$, we may give a binary relation either by the tuples $\{(1,2),(2,1),(2,3),(3,2),(1,3),(3,1),(1,1)\}$ or by the "simple" logic formula $(x \neq y \vee x = 1)$. For the domain $\{0,1\}$, we may give the ternary (not-all-equal) relation by the tuples $\{(1,0,0),(0,1,0),(0,0,1),(1,1,0),(1,0,1),(1,1,0)\}$ or by the "simple" logic formula $(x \neq y \vee y \neq z)$. In both of these examples, the simple formula is also in DNF.

Nota Bene. The results of this paper apply for the "simple" logic encoding as well as the "nice, simple" encoding in DNF except where specifically stated otherwise. These exceptions are Proposition 40 and Corollary 41 (which uses the "nice, simple" DNF) and Proposition 43 (which uses the tuple-listing encoding).

1.2 Related work

This is the journal version of [21] and [16]. The majority of the proofs were omitted from these conference papers but the section numbers are preserved in the arxiv versions. However, several parts of those papers have become superseded or otherwise outdated. This applies to Sections 3 and 5 of [21], leaving Section 4 appearing in its entirety (as Section 2 in this paper). From [16] we give Section 3 in its entirety but only the most interesting part of Section 4. Section 5 is omitted.

On the other hand, the canonical example of projective and non-degenerate adversaries is now known to be switchability [2]. This has raised the importance of Section 4 of [21] as the bridge between two forms of switchability and a necessary part of proving that PGP yields a QCSP in NP.

1.3 Some comment on notation

We use calligraphic notation $\mathcal{A}$ for constraint languages over domain A. Constraint languages can be seen as a set of relations over the same domain or as first-order relational structures and we rather conflate the two (already in the abstract). Sets such as $\text{Inv}(\mathbb{A})$ can be seen as infinite constraint languages and we might talk of (finite) subsets of this as a constraint language or a (finite-signature) reduct. We similarly conflate algebras with sets of operations on the same domain.

Algebras are indicated in blackboard notation $\mathbb{A}$. We may drop brackets around singleton sets. For example, if f is an operation, then we may write $\text{Inv}(f)$ as a shorthand for $\text{Inv}(\{f\})$. All domains in this paper are finite. We write pH to indicate positive Horn.

2 THE PGP: COLLAPSIBILITY AND BEYOND

Throughout this section, we will be concerned with a constraint language $\mathcal{A}$ that may or may not have some constants naming the elements. We will be specific when we require constants naming elements. In Chen's [1, 5], the assumption of constants naming elements is often implicit, e.g. through idempotency, but several of his theorems apply in the general case, and are reproduced here in generality.

Later in this section we will use Fraktur notation (e.g. $\mathfrak{A}$) for constraint languages embellished with additional constants (different from any basic constants just naming elements) that we ultimately use to denote universal variables.

2.1 Games, adversaries and reactive composition

For a primitive positive sentence φ we associate the structure $\mathcal{D}_\varphi$ whose elements a_v are variables v of φ and whose relational tuples are the atoms of φ. That is, an atom $R(v_1, \ldots, v_k)$ in φ becomes a tuple $(a_{v_1}, ..., a_{v_k}) \in R$ in $\mathcal{D}_\varphi$. We then say that $\mathcal{D}_\varphi$ is the *canonical database* of φ, and φ is the *canonical query* of $\mathcal{D}_\varphi$. We recall some terminology due to Chen [1, 5], for his natural adaptation of the model checking game to the context of pH-sentences. We shall not need to explicitly play these games but only to handle strategies for the existential player. An *adversary* $\mathscr{B}$ of length $m \geq 1$ is an m-ary relation over A. When $\mathscr{B}$ is precisely the set $B_1 \times B_2 \times \ldots \times B_m$ for some non-empty subsets $B_1, B_2, \ldots, B_m$ of A, we speak of a *rectangular adversary*. Let φ have universal variables $x_1, \ldots, x_m$ and quantifier-free part ψ. We write $\mathcal{A} \models \varphi_{\restriction \mathscr{B}}$ and say that *the existential player has a winning strategy in the* $(\mathcal{A}, \varphi)$*-game against adversary* $\mathscr{B}$ iff there exists a set of Skolem functions $\{\sigma_x : `\exists x' \in \varphi\}$ such that for any assignment π of the universally quantified variables of φ to A, where $\big(\pi(x_1), \ldots, \pi(x_m)\big) \in \mathscr{B}$, the map h_π is a homomorphism from $\mathcal{D}_\psi$ (the canonical database) to $\mathcal{A}$, where

$$h_\pi(x) := \begin{cases} \pi(x) & \text{if } x \text{ is a universal variable; and,} \\ \sigma_x(\pi|_{Y_x}) & \text{otherwise.} \end{cases}$$

(Here, Y_x denotes the set of universal variables preceding x and $\pi|_{Y_x}$ the restriction of π to Y_x.) Clearly, $\mathcal{A} \models \varphi$ iff the existential player has a winning strategy in the $(\mathcal{A}, \varphi)$-game against the so-called *full (rectangular) adversary* $A \times A \times \ldots \times A$ (which we will denote hereafter by A^m). We say that an adversary $\mathscr{B}$ of length m *dominates* an adversary $\mathscr{B}'$ of length m when $\mathscr{B}' \subseteq \mathscr{B}$. Note that $\mathscr{B}' \subseteq \mathscr{B}$ and $\mathcal{A} \models \varphi_{\restriction \mathscr{B}}$ implies $\mathcal{A} \models \varphi_{\restriction \mathscr{B}'}$. We will also consider sets of adversaries of the same length, denoted by uppercase Greek letters as in Ω_m; and, sequences thereof, which we denote with bold uppercase Greek letters as in $\mathbf{\Omega} = (\Omega_m)_{m \in \mathbb{N}}$. We will write $\mathcal{A} \models \varphi_{\restriction \Omega_m}$ to denote that $\mathcal{A} \models \varphi_{\restriction \mathscr{B}}$ holds for every adversary $\mathscr{B}$ in Ω_m. We call *width* of Ω_m and write $\text{width}(\Omega_m)$ for $\sum_{\mathscr{B} \in \Omega_m} |\mathscr{B}|$. We say that $\mathbf{\Omega}$ is *polynomially bounded* if there exists a polynomial $p(m)$ such that for every $m \geq 1$, $\text{width}(\Omega_m) \leq p(m)$. We say that $\mathbf{\Omega}$ is *effective* if there exists a polynomial $p'(m)$ and an algorithm that outputs Ω_m for every m in total time $p'(\text{width}(\Omega_m))$.

Let f be a k-ary operation on A and $\mathscr{A}, \mathscr{B}_1, \ldots, \mathscr{B}_k$ be adversaries of length m. We say that $\mathscr{A}$ is *reactively composable* from the adversaries $\mathscr{B}_1, \ldots, \mathscr{B}_k$ via f, and we write $\mathscr{A} \trianglelefteq f(\mathscr{B}_1, \ldots, \mathscr{B}_k)$ iff there exist partial functions $g_i^j : A^i \to A$ for every i in $[m]$ and every j in $[k]$ such that, for every tuple $(a_1, \ldots, a_m)$ in adversary $\mathscr{A}$ the following holds.

- for every j in $[k]$, the values $g_1^j(a_1), g_2^j(a_1, a_2), \ldots, g_m^j(a_1, a_2, \ldots, a_m)$ are defined and the tuple $\left(g_1^j(a_1), g_2^j(a_1, a_2), \ldots, g_m^j(a_1, a_2, \ldots, a_m)\right)$ is in adversary $\mathscr{B}_j$; and,
- for every i in $[m]$, $a_i = f\left(g_i^1(a_1, a_2, \ldots, a_i), g_i^2(a_1, a_2, \ldots, a_i), \ldots, g_i^k(a_1, a_2, \ldots, a_i)\right)$.

We write $\mathscr{A} \trianglelefteq \{\mathscr{B}_1, \ldots, \mathscr{B}_k\}$ if there exists a k-ary operation f such that $\mathscr{A} \trianglelefteq f(\mathscr{B}_1, \ldots, \mathscr{B}_k)$

Remark 5. *We will never show reactive composition by exhibiting a function f and partial functions g_j^i that depend on all their arguments. We will always be able to exhibit partial functions that depend only on their last argument.*

Reactive composition allows to interpolate complete Skolem functions from partial ones.

Theorem 6 ([1, Theorem 7.6]). *Let φ be a pH-sentence with m universal variables. Let $\mathscr{A}$ be an adversary and Ω_m a set of adversaries, both of length m.*

If $\mathcal{A} \models \varphi_{\restriction \Omega_m}$ and $\mathscr{A} \trianglelefteq \Omega_m$ then $\mathcal{A} \models \varphi_{\restriction \mathscr{A}}$.

Proof. We sketch the proof for the sake of completeness. Let $\Omega_m := \{\mathscr{B}_1, \ldots, \mathscr{B}_k\}$ and f and g_j^i be as in the definition of reactive composition and witnessing that $\mathscr{A} \trianglelefteq f(\mathscr{B}_1, \ldots, \mathscr{B}_k)$. Assume also that $\mathcal{A} \models \varphi_{\restriction \Omega_m}$. Given any sequence of play of the universal player according to the adversary $\mathscr{A}$, that is v_1 is played as $a_1 \in A_1$, v_2 is played as $a_2 \in A_2$, etc., we "go backwards through f" via the maps g_j^i to pinpoint *incrementally* for each $j \in [k]$ a sequence of play $v_1 = g_j^1(a_1)$, $v_2 = g_j^2(a_1, a_2)$ etc., thus yielding eventually a tuple that belongs to adversary $\mathscr{B}_j$. After each block of universal variables, we lookup the winning strategy for the existential player against each adversary $\mathscr{B}_j$ and "going forward through f", that is applying f to the choice of values for an existential variable against each adversary, we obtain a consistent choice for this variable against adversary $\mathscr{A}$ (this is because f is a polymorphism and the quantifier-free part of the sentence φ is conjunctive positive). Going back and forth we obtain eventually an assignment to the existential variables that is consistent with the universal variables being played as $a_1, a_2, \ldots, a_m$. $\square$

As a concrete example of an interesting sequence of adversaries, consider the adversaries for the notion of *p-collapsibility*. Let $p \geq 0$ be some fixed integer. For x in A, let $\Upsilon_{m,p,x}$ be the set of all rectangular adversaries of length m with p coordinates that are the set A and all the other that are the fixed singleton $\{x\}$. For $B \subseteq A$, let $\Upsilon_{m,p,B}$ be the union of $\Upsilon_{m,p,x}$ for all x in B. Let $\mathbf{Y}_{p,B}$ be the sequence of adversaries $\left(\Upsilon_{m,p,B}\right)_{m \in \mathbb{N}}$. Chen's original definition [5] for a structure $\mathcal{A}$ to be *p-collapsible from source B* was that for every m and for all pH-sentence φ with m universal variables, $\mathcal{A} \models \varphi_{\restriction \Upsilon_{m,p,B}}$ implies $\mathcal{A} \models \varphi$.

Let us consider now the adversaries for the notion of *p-switchability*. Let $p \geq 0$ be some fixed integer. Let $\Xi_{m,p}$ be the set of all tuples $(x_1, \dots, x_m)$ in which there exists $a_i < \dots < a_{k'}$, for $k' \leq p$, so that

$$(x_1, \dots, x_m) = (x_1, \dots, x_{a_1}, x_{a_1+1}, \dots, x_{a_2}, x_{a_2+1}, \dots, \dots, x_{a_{k'}}, x_{a_{k'+1}}, \dots, x_m),$$

where $x_1 = \dots = x_{a_1-1}, x_{a_1} = \dots = x_{a_2-1}, \dots, x_{a_{k'}} = \dots = x_{a_m}$. Let Ξ_p be the sequence of adversaries $\left(\Xi_{m,p}\right)_{m \in \mathbb{N}}$. Chen originally defined [1] a constraint language $\mathcal{A}$ to be *p-switchable* iff for every m and for all pH-sentences φ with m universal variables, $\mathcal{A} \models \varphi_{\restriction \Xi_{m,p}}$ implies $\mathcal{A} \models \varphi$. We will contrast the different definitions once again in the key forthcoming theorem "In Abstracto" (Theorem 19), where we will finally prove them equivalent.

2.2 The Π_2-case

For a Π_2-pH sentence, i.e. with prefix $\forall^* \exists^*$, the existential player knows the values of all universal variables beforehand, and it suffices for her to have a winning strategy for each instantiation (and perhaps no way to reconcile them as should be the case for an arbitrary sentence). This also means that considering a set of adversaries of the same length is not really relevant in this Π_2-case as we may as well consider the union of these adversaries or the set of all their tuples .

Lemma 7 (**principle of union**). *Let Ω_m be a set of adversaries of length m and φ a Π_2-sentence with m universal variables. Let $\mathscr{O}_{\cup \Omega_m} := \bigcup_{\mathscr{O} \in \Omega_m} \mathscr{O}$ and $\Omega_{tuples} := \{\{t\} | t \in \mathscr{O}_{\cup \Omega_m}\}$. We have the following equivalence.*

$$\mathcal{A} \models \varphi_{\restriction \Omega_m} \iff \mathcal{A} \models \varphi_{\restriction \mathscr{O}_{\cup \Omega_m}} \iff \mathcal{A} \models \varphi_{\restriction \Omega_{tuples}}$$

The forward implications

$$\mathcal{A} \models \varphi_{\restriction \Omega_m} \implies \mathcal{A} \models \varphi_{\restriction \mathscr{O}_{\cup \Omega}} \implies \mathcal{A} \models \varphi_{\restriction \Omega_{\text{tuples}}}$$

of Lemma 7 hold clearly for arbitrary pH-sentences. The proof is trivial and is a direct consequence of the following obvious fact.

Fact 8. *Let Ω_m be a set of adversaries of length m and φ a Π_2-sentence with m universal variables.*

$$\mathcal{A} \models \varphi_{\restriction \Omega_m}$$
$$\Updownarrow$$
$$\forall \mathscr{O} \in \Omega_m \forall t = (a_1, \dots, a_m) \in \mathscr{O}\ \mathcal{A} \models \varphi_{\restriction \{t\}}$$

Remark 9 (following Lemma 7). *For a sentence that is not Π_2, this does not necessarily hold. For example, consider $\forall x \forall y \exists z \forall w\ E(x,z) \wedge E(y,z) \wedge E(w,z)$ on the irreflexive 4-clique $\mathcal{K}_4$. The sentence is not true, but for all individual tuples (x_0, y_0, w_0), we have $\exists z\ E(x_0,z) \wedge E(y_0,z) \wedge E(w_0,z)$.*

Let $\mathscr{A}$ be an adversary and Ω_m a set of adversaries, both of length m. We say that Ω_m *generates* $\mathscr{A}$ iff for any tuple t in $\mathscr{A}$, there exists a k-ary polymorphism f_t of $\mathcal{A}$ and tuples $t_1, \dots, t_k$ in $\mathscr{O}_{\cup \Omega_m}$ such that $f_t(t_1, \dots, t_k) = t$. We have the following analogue of Theorem 6.

Proposition 10. *Let φ be a Π_2-pH-sentence with m universal variables. Let $\mathscr{A}$ be an adversary and Ω_m a set of adversaries, both of length m.*

If $\mathcal{A} \models \varphi_{\restriction \Omega_m}$ and Ω_m generates $\mathscr{A}$ then $\mathcal{A} \models \varphi_{\restriction \mathscr{A}}$.

Proof. The hypothesis that Ω_m generates $\mathscr{A}$ can be rephrased as follows : for each tuple t in $\mathscr{A}$, $\{t\} \trianglelefteq f_t(t_1, t_2, \dots, t_k)$, where $t_1, t_2, \dots, t_k$ belong to $\mathscr{O}_{\cup \Omega_m}$. To see this, it remains to note that the suitable g_i^j's from the definition of composition are induced trivially as there is no choice: for every j in $[k]$ and every i in $[m]$ pick $g_i^j(a_1, a_2, \dots, a_i) = t_{i,j}$ where $t_{i,j}$ is the ith element of t_j. So by Theorem 6, if $\mathcal{A} \models \varphi_{\upharpoonright \Omega_{\text{tuples}}}$ then $\mathcal{A} \models \varphi_{\upharpoonright \{t\}}$. As this holds for any tuple t in $\mathscr{A}$, via the principle of union, it follows that $\mathcal{A} \models \varphi_{\upharpoonright \mathscr{A}}$. □

We will construct a *canonical* Π_2*-sentence to assert that an adversary is generating*. Let $\mathscr{O}$ be some adversary of length m. Let $\sigma^{(m)}$ be the signature σ expanded with a sequence of m constants. For a map μ from $[m]$ to A, we write $\mu \in \mathscr{O}$ as shorthand for $(\mu(1), \mu(2), \ldots, \mu(m)) \in \mathscr{O}$. For some set Ω_m of adversaries of length m, we consider the following $\sigma^{(m)}$-structure:

$$\bigotimes_{\mathscr{O} \in \Omega_m} \bigotimes_{\mu \in \mathscr{O}} \mathfrak{A}_\mu$$

where the $\sigma^{(m)}$-structure $\mathfrak{A}_\mu$ denotes the expansion of $\mathcal{A}$ by m constants as given by the map μ, and $\otimes$ denotes the direct product. Let $\varphi_{\Omega_m, \mathcal{A}}$ be the Π_2-pH-sentence[2] created from the canonical query of the σ-reduct of this $\sigma^{(m)}$-structure with the m constants c_j becoming variables w_j, universally quantified outermost, when all *constants are pairwise distinct*. Otherwise, we will say that Ω_m is *degenerate*, and not define the canonical sentence. An example of this construction is furnished in Example 11. Ω_m is degenerate precisely if there exist $i, j \in [m]$ so that, for all μ in $\mathscr{O}_{\cup \Omega_m}$, $\mu(i) = \mu(j)$.

Example 11. $\varphi := \forall w_1, w_2, w_3 \, \exists y_1^1, y_1^2, y_1^3, y_1^4 \, E(y_1^1, w_1) \wedge E(w_1, y_1^2) \wedge E(w_1, y_1^3) \wedge E(y_1^3, y_1^2) \wedge E(y_1^4, w_2) \wedge E(w_3, y_1^4)$.

The sentence φ, depicted on the left, comes from the $\sigma^{(3)}$-structure depicted on the right.

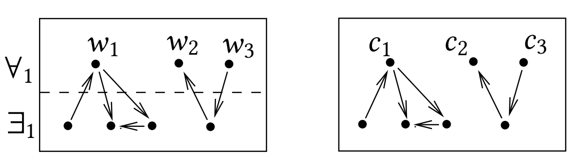

Note that adversaries such as $\Upsilon_{m,p,B}$ corresponding to p-collapsibility are not degenerate for $p > 0$, and degenerate for $p = 0$.

Proposition 12. *Let Ω_m be a set of adversaries of length m that is not degenerate. The following are equivalent.*

(i) *for any Π_2-pH sentence ψ, $\mathcal{A} \models \psi_{\upharpoonright \Omega_m}$ implies $\mathcal{A} \models \psi$.*
(ii) *for any Π_2-pH sentence ψ, $\mathcal{A} \models \psi_{\upharpoonright \mathscr{O}_{\cup \Omega_m}}$ implies $\mathcal{A} \models \psi$.*
(iii) *for any Π_2-pH sentence ψ, $\mathcal{A} \models \psi_{\upharpoonright \Omega_{tuples}}$ implies $\mathcal{A} \models \psi$.*
(iv) $\mathcal{A} \models \varphi_{\mathscr{O}_{\cup \Omega_m}, \mathcal{A}}$
(v) $\mathcal{A} \models \varphi_{\Omega_{tuples}, \mathcal{A}}$
(vi) *Ω_m generates A^m.*

Proof. The first three items are equivalent by Lemma 7 (these implications have the same conclusion and equivalent premises). The fourth and fifth items are trivially equivalent since $\varphi_{\mathscr{O}_{\cup \Omega_m}, \mathcal{A}}$ and $\varphi_{\Omega_{\text{tuples}}, \mathcal{A}}$ are the same sentence.

We show the implication from the third item to the fifth. By construction, $\varphi_{\Omega_{\text{tuples}}, \mathcal{A}}$ is Π_2 and it suffices to show that there exists a winning strategy for $\exists$ against any adversary $\{t\}$ in Ω_{tuples}. This is true by construction. Indeed, note that there exists a winning strategy for $\exists$ in the $(\mathcal{A}, \varphi_{\Omega_{\text{tuples}}, \mathcal{A}})$-game against adversary $\{t\}$ iff there is a homomorphism from the $\sigma^{(m)}$-structure $\bigotimes_{t' \in \Omega_{\text{tuples}}} \mathfrak{A}_{\mu_{t'}}$

[2]For two constraint languages $\mathcal{A}$ and $\mathcal{B}$, when Ω_m is A^m and m is $|A|^B$, $\mathcal{B}$ models this canonical sentence iff QCSP$(\mathcal{A}) \subseteq$ QCSP$(\mathcal{B})$ [22]

to the $\sigma^{(m)}$-structure $\mathfrak{A}_{\mu_t}$, where $\mu_t : [m] \longrightarrow A$ is the map induced naturally by t. The projection is such a homomorphism.

The penultimate item implies the last one: instantiate the universal variables of $\varphi_{\Omega_{\text{tuples}},\mathcal{A}}$ as given by the m-tuple t and pick for f_t the homomorphism from the product structure witnessing that $\exists$ has a winning strategy.

Finally, the last item implies the first one by Proposition 10. □

2.3 The unbounded case

Let n denote the number of elements of the structure $\mathcal{A}$. Let $\mathscr{B}$ be an adversary from $\Omega_{n\cdot m}$. We will denote by Proj$\mathscr{B}$ the set of adversaries of length m induced by projecting over some arbitrary choice of m coordinates, one in each block of size n; that is $1 \leq i_1 \leq n, n+1 \leq i_2 \leq 2\cdot n, \dots, n\cdot(m-1)+1 \leq i_m \leq n\cdot m$. Of special concern to us are *projective sequences of adversaries* Ω satisfying the following for every $m \geq 1$,

$$\forall \mathscr{B} \in \Omega_{n\cdot m}\ \exists \mathscr{A} \in \Omega_m \bigwedge_{\widetilde{\mathscr{B}}\in \mathrm{Proj}\mathscr{B}} \widetilde{\mathscr{B}} \subseteq \mathscr{A} \quad (m\textbf{-projectivity})$$

As an example, consider the adversaries for collapsibility.

FACT 13. *Let $B \subseteq A$ and $p \geq 0$. The sequence of adversaries $\Upsilon_{p,B}$ are projective.*

Example 14. *For a concrete illustration consider $A = \{0, 1, 2\}$ (thus $n = 3$). We illustrate the fact that $\Upsilon_{p=2,B=\{0\}}$ is projective for $m = 4$ and some adversary $\mathscr{B} \in \Omega_{n\cdot m} = \Upsilon_{p=2,B=\{0\},3\cdot4=12}$. Adversaries are depicted vertically with horizontal lines separating the blocks.*

$\mathscr{B} \in \Omega_{n\cdot m}$	Proj$\mathscr{B}$				$\mathscr{A} \in \Omega_m$
A	A	A		$\cancel{A}$	
0	$\cancel{0}$	$\cancel{0}$	…	$\cancel{0}$	A
0	$\cancel{0}$	$\cancel{0}$		0	
0	0	0		$\cancel{0}$	
0	$\cancel{0}$	$\cancel{0}$	…	$\cancel{0}$	0
0	$\cancel{0}$	$\cancel{0}$		0	
0	0	0		$\cancel{0}$	
0	$\cancel{0}$	$\cancel{0}$	…	$\cancel{0}$	0
0	$\cancel{0}$	$\cancel{0}$		0	
0	0	$\cancel{0}$		$\cancel{0}$	
A	$\cancel{A}$	A	…	$\cancel{A}$	A
0	$\cancel{0}$	$\cancel{0}$		0	

The adversary $\mathscr{A}$ dominates any adversary obtained by projecting the original larger adversary $\mathscr{B}$ by keeping a single position per block.

We could actually consider w.l.o.g. sequences of *singleton* adversaries.

FACT 15. *If Ω is projective then so is the sequence $\left(\bigcup_{\mathscr{O}\in\Omega_m} \mathscr{O}\right)_{m\in\mathbb{N}}$.*

A *canonical sentence for composability for arbitrary pH-sentences* with m universal variables may be constructed similarly to the canonical sentence for the Π_2 case, except that it will have $m\cdot n$ universal variables, which we view as m blocks of n variables, where n is the number of elements of the structure $\mathcal{A}$. Let $\mathscr{O}$ be some adversary of length m. Let $\sigma^{(n\cdot m)}$ be the signature σ expanded with a sequence of $n\cdot m$ constants $c_{1,1}, \dots, c_{n,1}, c_{1,2} \dots, c_{n,2}, \dots c_{1,m} \dots, c_{n,m}$. We say that a map μ from $[n] \times [m]$ to A is *consistent* with $\mathscr{O}$ iff for every $(i_1, i_2, \dots, i_m)$ in $[n]^m$, the tuple $(\mu(i_1, 1), \mu(i_2, 2), \dots, \mu(i_m, m))$ belongs to the adversary $\mathscr{O}$. We write $A^{[n\cdot m]}_{\restriction\mathscr{O}}$ for the set of such consistent maps. For some set Ω_m of adversaries of length m, we consider the following

$\sigma^{(n\cdot m)}$-structure:

$$\bigotimes_{\mathscr{O}\in\Omega_m}\bigotimes_{\mu\in A^{[n\cdot m]}_{\restriction\mathscr{O}}}\mathfrak{A}_{\mathscr{O},\mu}$$

where the $\sigma^{(n\cdot m)}$-structure $\mathfrak{A}_{\mathscr{O},\mu}$ denotes the expansion of $\mathcal{A}$ by $n\cdot m$ constants as given by the map μ. Let $\varphi_{n,\Omega_m,\mathcal{A}}$ be the Π_2-pH-sentence created from the canonical query of the σ-reduct of this $\sigma^{(n\cdot m)}$ product structure with the $n\cdot m$ constants c_{ij} becoming variables w_{ij}, universally quantified outermost. As for the canonical sentence of the Π_2-case, this sentence is not well defined if constants are not pairwise distinct, which occurs precisely for degenerate adversaries.

Lemma 16. *Let Ω_m be a set of adversaries of length m that is not degenerate. Let $\mathcal{A}$ be a structure of size n. If $\mathcal{A}$ models $\varphi_{n,\Omega_m,\mathcal{A}}$ then the full adversary A^m is reactively composable from Ω_m. That is, $\mathcal{A}\models\varphi_{n,\Omega_m,\mathcal{A}} \implies A^m\trianglelefteq\Omega_m$*

Proof. We let each block of n universal variables of the canonical sentence $\varphi_{n,\Omega_m,\mathcal{A}}$ enumerate the elements of A. That is, given an enumeration $a_1,a_2,\ldots,a_n$ of A, we set $w_{i,j}=a_i$ for every j in $[m]$ and every i in $[n]$.

The assignment to the existential variables provides us with a k-ary polymorphism (the sentence being built as the conjunctive query of a product of k copies of $\mathcal{A}$) together with the desired partial maps. A coordinate r in $[k]$ corresponds to a choice of some adversary $\mathscr{O}$ of Ω_m and some map μ_r from $[n]\times[m]$ to A, consistent with this adversary. The partial map $g^r_\ell : A^\ell\to A$ with ℓ in $[m]$ (and r in $[k]$) is given by μ_r as follows: $g^r_\ell(a_{i_1},\ldots,a_{i_\ell})$ depends only on the last coordinate a_{i_ℓ} and takes value $\mu_r(i,\ell)$ if $a_{i_\ell}=a_i$. By construction of the sentence and the property of consistency of such μ_r with the adversary $\mathscr{O}$, these partial functions satisfy the properties as given in the definition of reactive composition. □

Lemma 17. *Let Ω be a sequence of sets of adversaries that has the m-projectivity property for some $m\geq 1$ such that $\Omega_{n\cdot m}$ is not degenerate. The following holds.*

(i) *$\mathcal{A}\models\psi_{\restriction\Omega_{\mathrm{n\cdot m}}}$, where $\psi=\varphi_{n,\Omega_{\mathrm{m}},\mathcal{A}}$*
(ii) *If for every Π_2-sentence ψ with $m\cdot n$ universal variables, it holds that $\mathcal{A}\models\psi_{\restriction\Omega_{\mathrm{m\cdot n}}}$ implies $\mathcal{A}\models\psi$, then $\mathcal{A}\models\varphi_{n,\Omega_{\mathrm{m}},\mathcal{A}}$.*

Proof. The second statement is a direct consequence of the first one. The proof of the first statement generalises an argument used in the proof of Proposition 12. Consider any adversary $\mathscr{O}$ in $\Omega_{n\cdot m}$. For convenience, we name the positions of this adversary in a similar fashion to the universal variables of the sentence, namely by a pair (i,j) in $[n]\times[m]$. By projectivity, there exists an adversary $\mathscr{O}'$ in Ω_m which dominates any adversary $\tilde{\mathscr{O}}$ in Proj$\mathscr{O}$ (obtained by projecting over an arbitrary choice of one position in each of the m blocks of size n). In the product structure underlying the formula $\varphi_{n,\Omega_m,\mathcal{A}}$, we consider the following structure:

$$\bigotimes_{\mu\in A^{[n\cdot m]}_{\restriction\mathscr{O}'}}\mathfrak{A}_{\mathscr{O}',\mu}$$

An instantiation of the universal variables of $\varphi_{n,\Omega_m,\mathcal{A}}$ according to some tuple t from the adversary $\mathscr{O}$ corresponds naturally to a map μ_t from $[n]\times[m]$ to A. Observe that our choice of $\mathscr{O}'$ ensures that this map μ_t is consistent with $\mathscr{O}'$. An instantiation of the universal variables by μ_t induces a $\sigma^{(n\cdot m)}$-structure $\mathfrak{A}_{\mu_t}$ and a winning strategy for $\exists$ amounts to a homomorphism from the product $\sigma^{(n\cdot m)}$-structure underlying the sentence to this $\mathfrak{A}_{\mu_t}$. Since the component $\mathfrak{A}_{\mathscr{O}',\mu_t}$ of this product structure is isomorphic to $\mathfrak{A}_{\mu_t}$, we may take for a homomorphism the corresponding projection. This shows that $\mathcal{A}\models\psi_{\restriction\Omega_{\mathrm{n\cdot m}}}$ where $\psi=\varphi_{n,\Omega_{\mathrm{m}},\mathcal{A}}$. □

Theorem 18. *Let Ω be a sequence of sets of adversaries that has the m-projectivity property for some $m \geq 1$ such that $\Omega_{n \cdot m}$ is not degenerate. The following chain of implications holds*

$$\text{(i)} \implies \text{(ii)} \implies \text{(iii)} \implies \text{(iv)}$$

where,

(i) *For every Π_2-pH-sentence ψ with $m \cdot n$ universal variables, $\mathcal{A} \models \psi_{\restriction \Omega_{m \cdot n}}$ implies $\mathcal{A} \models \psi$.*
(ii) $\mathcal{A} \models \varphi_{n,\Omega_m,\mathcal{A}}$.
(iii) $A^m \trianglelefteq \Omega_m$.
(iv) *For every pH-sentence ψ with m universal variables, $\mathcal{A} \models \psi_{\restriction \Omega_m}$ implies $\mathcal{A} \models \psi$.*

Proof. The first implication holds by the previous lemma (second item of Lemma 17, this is the step where we use projectivity). The second implication is Lemma 16. The last implication is Theorem 6. □

Thus, in the projective case, when an adversary is good enough in the Π_2-case, it is good enough in general. This can be characterised logically via canonical sentences or "algebraically" in terms of reactive composition or the weaker and more usual composition property (see (vi) below).

Theorem 19 (In abstracto). *Let $\Omega = \left(\Omega_m\right)_{m \in \mathbb{N}}$ be a projective sequence of adversaries, none of which are degenerate. The following are equivalent.*

(i) *For every $m \geq 1$, for every pH-sentence ψ with m universal variables, $\mathcal{A} \models \psi_{\restriction \Omega_m}$ implies $\mathcal{A} \models \psi$.*
(ii) *For every $m \geq 1$, for every Π_2-pH-sentence ψ with m universal variables, $\mathcal{A} \models \psi_{\restriction \Omega_m}$ implies $\mathcal{A} \models \psi$.*
(iii) *For every $m \geq 1$, $\mathcal{A} \models \varphi_{n,\Omega_m,\mathcal{A}}$.*
(iv) *For every $m \geq 1$, $\mathcal{A} \models \varphi_{\mathcal{O}_{\cup \Omega_m},\mathcal{A}}$.*
(v) *For every $m \geq 1$, $A^m \trianglelefteq \Omega_m$.*
(vi) *For every $m \geq 1$, Ω_m generates A^m.*

Proof. Propositions 12 establishes the equivalence between (ii), (iv) and (vi) for fixed values of m (numbered there as (i), (iv) and (vi), respectively).

To lift these relatively trivial equivalences to the general case, i.e. from Π_2 to unbounded, the method of our current proof no longer preserves the parameter m. The chain of implications of Theorem 18 translates here, once the parameter is universally quantified, to the chain of implications

$$\text{(ii)} \implies \text{(iii)} \implies \text{(v)} \implies \text{(i)}$$

The fact that (i) implies (ii) is trivial[3], which concludes the proof. □

Remark 20. *The above equivalences can be read along two dimensions:*

	general	Π_2
logical interpolation	(i)	(ii)
canonical sentences	(iii)	(iv)
algebraic interpolation	(v)	(vi)

Chen's original definitions of collapsibility and switchability correspond with item (i), while the definitions given in the introduction correspond with item (vi). For example, it is the formulation (i) that provides Chen's original proof that switchability yields a QCSP in NP (Theorem 7.11 in [1]).

[3]We note in passing and for purely pedagogical reason that the implication (v) to (vi) is also trivial, while the natural implication (iii) to (iv) will appear as an evidence to the reader once the definition of the canonical sentences is digested.

In that same paper, the property of switchability as defined in the introduction is only shown to yield that the m-alternation-QCSP (allow only inputs in Π_m prenex form, where m is fixed) is in NP (Proposition 3.3 in [1]). Let $\text{CSP}_c(\mathcal{A})$ and $\text{QCSP}_c(\mathcal{A})$ be the versions of CSP($\mathcal{A}$) and QCSP($\mathcal{A}$), respectively, in which constants naming the elements of $\mathcal{A}$ may appear in instances. The following ostensibly generalises Theorem 7.11 [1] to effective and "projective" PGP, though we now know from [2] via Theorem 19 that switchability explains all finite-domain algebra PGP.

Corollary 21. *Let $\mathcal{A}$ be a constraint language. Let Ω be a sequence of non degenerate adversaries that is effective, projective and polynomially bounded such that Ω_m generates A^m for every $m \geq 1$.*

Let $\mathcal{A}'$ be the constraint language $\mathcal{A}$, possibly expanded with constants naming elements, at least one for each element that occurs in Ω. The problem QCSP($\mathcal{A}$) reduces in polynomial time to CSP($\mathcal{A}'$). In particular, if $\mathcal{A}$ has all constants, the problem QCSP($\mathcal{A}$) reduces in polynomial time to CSP($\mathcal{A}$).

Proof. To check whether a pH-sentence φ with m universal variables holds in $\mathcal{A}$, by Theorem 19, we only need to check that $\mathcal{A} \models \varphi_{\restriction \mathcal{B}}$ for every $\mathcal{B}$ in Ω_m. The reduction proceeds as in the proof of [1, Lemma 7.12], which we outline here for completeness.

Pretend first that we reduce $\mathcal{A} \models \varphi_{\restriction \mathcal{B}}$ to a collection of CSP instances, one for each tuple t of $\mathcal{B}$, obtained by instantiation of the universal variables with the corresponding constants. If x is an existential variable in φ, let x_t be the corresponding variable in the CSP instance corresponding to t. We will in fact enforce equality constraints via renaming of variables to ensure that we are constructing Skolem functions. For any two tuples t and t' in $\mathcal{B}$ that agree on their first ℓ coordinates, let Y_ℓ be the corresponding universal variables of φ. For every existential variable x such that Y_x (the universally quantified variables of φ preceding x) is contained in Y_ℓ, we identify x_t with $x_{t'}$. $\square$

Since Zhuk has proved that all cases of PGP in finite algebras come from switchability, the most important cases of In Abstracto (Theorem 19) and Corollary 21 involve the already introduced adversaries $\Xi_{m,p}$ (for some p) substituted for the placeholder Ω_m. Note that when $p > 0$, $\Xi_{m,p}$ is non-degenerate, and the sequence $\left(\Xi_{m,p}\right)_{m \in \mathbb{N}}$ is readily seen to be projective. We have resisted giving In Abstracto (Theorem 19) only for switchability in order to emphasise that the proof comes alone from non-degenerate and projective. However, let us state its consequence nonetheless.

Corollary 22. *Let $\mathcal{A}$ be a finite constraint language, with constants naming all of its elements, so that Pol($\mathcal{A}$) is switchable. Then QCSP($\mathcal{A}$) reduces to a polynomial number of instances of CSP($\mathcal{A}$) and is in NP.*

2.4 Studies of collapsibility

Let $\mathcal{A}$ be a constraint language, $B \subseteq A$ and $p \geq 0$. Recall the structure $\mathcal{A}$ is *p-collapsible with source B* when for all $m \geq 1$, for all pH-sentences φ with m universal quantifiers, $\mathcal{A} \models \varphi$ iff $\mathcal{A} \models \varphi_{\restriction \Upsilon_{m,p,B}}$. Collapsible structures are very important: to the best of our knowledge, they are in fact the only examples of structures that enjoy a form of polynomial QCSP to CSP reduction. This is different if one considers structures with infinitely many relations where the more general notion of *switchability* crops up [1]. Our abstract results of the previous section apply to both switchability and collapsibility but we concentrate here on the latter. This result applies since the underlying sequence of adversaries are projective (see Fact 13), as long as $p > 0$ (non degenerate case).

Corollary 23 (**In concreto**). *Let $\mathcal{A}$ be a structure, $\emptyset \subsetneq B \subseteq A$ and $p > 0$. The following are equivalent.*

(i) *$\mathcal{A}$ is p-collapsible from source B.*
(ii) *$\mathcal{A}$ is Π_2-p-collapsible from source B.*

(iii) *For every m, the structure $\mathcal{A}$ satisfies the canonical Π_2-sentence with $m \cdot |A|$ universal variables $\varphi_{n,\Upsilon_{m,p,B},\mathcal{A}}$.*
(iv) *For every m, the structure $\mathcal{A}$ satisfies the canonical Π_2-sentence with m universal variables $\varphi_{\mathcal{U},\mathcal{A}}$, where $\mathcal{U} = \bigcup_{O \in \Upsilon_{m,p,B}} O$.*
(v) *For every m, there exists a polymorphism f of $\mathcal{A}$ witnessing that $A^m \trianglelefteq \Upsilon_{m,p,B}$.*
(vi) *For every m, for every tuple t in A^m, there is a polymorphism f_t of $\mathcal{A}$ of arity k at most $\binom{m}{p} \cdot |B|$ and tuples $t_1, t_2, \ldots, t_k$ in $\Upsilon_{m,p,B}$ such that $f_t(t_1, t_2, \ldots, t_k) = t$.*

Remark 24. *When $p = 0$, we obtain degenerate adversaries and this is due to the fact that if a QCSP is permitted equalities, then 0-collapsibility can never manifest (think of $\forall x, y\ x = y$).*

Suppose $\mathcal{A}$ is expanded with constants naming all the elements. Then in [5], Case (v) of Corollary 23 is equivalent to Pol($\mathcal{A}$) being p-collapsible (in the algebraic sense: Definition 3.11 in [5]). It is proved in [5] that if Pol($\mathcal{A}$), is k-collapsible (in the algebraic sense), then $\mathcal{A}$ is k-collapsible (in the relational sense: Definition 5.1 in [5]). We note that Corollary 23 proves the converse, finally tying together the two forms of collapsibility (algebraic and relational) that appear in [5] .

We will now give an application of Corollary 23. We will work over partially reflexive paths which are paths in which some vertices are self-loops and others are loop-free. For a sequence $\beta \in \{0,1\}^*$, of length $|\beta|$, let $\mathcal{P}_\beta$ be the undirected path on $|\beta|$ vertices such that the i^{th} vertex has a loop iff the i^{th} entry of β is 1 (we may say that the path $\mathcal{P}$ is *of the form* β). A path $\mathcal{H}$ is *quasi-loop-connected* if it is of either of the forms

(i) $0^a 1^b \alpha$, for $b > 0$ and some α with $|\alpha| = a$, or
(ii) $0^a \alpha$, for some α with $|\alpha| \in \{a, a-1\}$.

A path whose self-loops induce a connected subgraph is further said to be *loop-connected.*

Application 25. *Let $\mathcal{A}$ be a partially reflexive path (no constants are present) that is quasi-loop connected. Then Pol($\mathcal{A}$) has the PGP.*

Proof. Indeed, a partially reflexive path $\mathcal{A}$ that is quasi-loop connected has the same QCSP as a partially reflexive path that is loop-connected $\mathcal{B}$ [23] since for some $r_a > 0$ there is a surjective homomorphism g from $\mathcal{A}^{r_a}$ to $\mathcal{B}$ and for some $r_b > 0$ there is a surjective homomorphism h from $\mathcal{B}^{r_b}$ to $\mathcal{A}$ (see main result of [22]). Indeed, this motivated the name quasi-loop connected itself. We also know that $\mathcal{B}$ admits a majority polymorphism m [24] and is therefore 2-collapsible from any singleton source [5] and that Theorem 23 holds for $\mathcal{B}$. Pick some arbitrary element a in $\mathcal{A}$ such that there is some b in $\mathcal{B}$ satisfying $g(a, a, \ldots, a) = b$. Use b as a source for $\mathcal{B}$.

We proceed to lift (vi) of Corollary 23 from structure $\mathcal{B}$ to $\mathcal{A}$, which we recall here for $\mathcal{B}$: for every m, for every tuple t in B^m, there is a polymorphism f_t of $\mathcal{B}$ of arity k and tuples $t_1, t_2, \ldots, t_k$ in $\Upsilon_{m,2,b}$ such that $f_t(t_1, t_2, \ldots, t_k) = t$.

Let g^k denote the surjective homomorphism from $(\mathcal{A}^{r_a})^k$ to $\mathcal{B}^k$ that applies g blockwise. Going back from t_i through g, we can find r_a tuples $t_{i,1}, t_{i,2}, \ldots, t_{i,r_a}$ all in $\Upsilon_{m,2,a}$ (adversaries based on the domain of $\mathcal{A}$) such that $g(t_{i,1}, t_{i,2}, \ldots, t_{i,r_a}) = t_i$. Thus, we can generate any $\widetilde{t}$ in $\mathcal{B}$ via $f_{\widetilde{t}} \circ (g^k)$ from tuples of $\Upsilon_{m,2,a}$.

Let $\hat{t}$ be now some tuple of $\mathcal{A}$. By surjectivity of h, let $\widetilde{t_1}, \widetilde{t_2}, \ldots, \widetilde{t_{r_b}}$ be tuples of $\mathcal{B}$ such that $h(\widetilde{t_1}, \widetilde{t_2}, \ldots, \widetilde{t_{r_b}}) = \hat{t}$. The polymorphism of $\mathcal{A}$ $(f_{\widetilde{t_1}} \circ (g^k), f_{\widetilde{t_2}} \circ (g^k), \ldots, f_{\widetilde{t_{r_b}}} \circ (g^k))$ shows that $\Upsilon_{m,2,a}$ generates $\hat{t}$. This shows that $\mathcal{A}$ is also 2-collapsible from a singleton source. □

The last two conditions of Corollary 23 provide us with a semi-decidability result: for each m, we may look for a particular polymorphism (v) or several polymorphisms (vi). Instead of a sequence of polymorphisms, we now strive for a better algebraic characterisation. We will only be able to do so for the special case of a singleton source, but this is the only case hitherto found in nature.

Let $\mathcal{A}$ be a structure with a constant x naming some element. Call a k-ary polymorphism of $\mathcal{A}$ such that f is surjective when restricted at any position to $\{x\}$ a *Hubie-pol in* $\{x\}$. Chen uses the following lemma to show 4-collapsibility of bipartite graphs and disconnected graphs [3, Examples 1 and 2]. Though, we know via a direct argument [25] that these examples are in fact 1-collapsible from a singleton source.

Lemma 26 (Chen's lemma [5, Lemma 5.13]). *Let $\mathcal{A}$ be a structure with a constant x naming some element, so that $\mathcal{A}$ has a k-ary Hubie-pol in $\{x\}$. Then $\mathcal{A}$ is $(k-1)$-collapsible from source $\{x\}$.*

Proof. We sketch the proof for pedagogical reasons. Via Corollary 23, it suffices to show that for any m, A^m is generated by $\Upsilon_{m,k-1,x}$ (instead of the notion of reactive composition).

Consider adversaries of length $m = k$ for now, that is from $\Upsilon_{k,k-1,x}$. If we apply the Hubie-pol f to these k adversaries, we generate the full adversary A^k. With a picture (adversaries are drawn as columns):

$$f\begin{pmatrix} \{x\} & A & A & \dots & A \\ A & \{x\} & A & \dots & A \\ \vdots & & \ddots & & \vdots \\ A & \dots & A & \{x\} & A \\ A & \dots & A & A & \{x\} \end{pmatrix} = \begin{pmatrix} A \\ A \\ \vdots \\ A \\ A \end{pmatrix} = A^k$$

Expanding these adversaries uniformly with singletons $\{x\}$ to the full length m, we may produce an adversary from $\Upsilon_{m,k,x}$. With a picture for *e.g.* trailing singletons:

$$f\begin{pmatrix} \{x\} & A & A & \dots & A \\ A & \{x\} & A & \dots & A \\ \vdots & & \ddots & & \vdots \\ A & \dots & A & \{x\} & A \\ A & \dots & A & A & \{x\} \\ \{x\} & \{x\} & \{x\} & \dots & \{x\} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \{x\} & \{x\} & \{x\} & \dots & \{x\} \end{pmatrix} = \begin{pmatrix} A \\ A \\ \vdots \\ A \\ A \\ \{x\} \\ \vdots \\ \{x\} \end{pmatrix}$$

Shifting the first additional row of singletons in the top block, we will obtain the family of adversaries from $\Upsilon_{m,k,x}$ with a single singleton in the first $k+1$ positions. It should be now clear that we may iterate this process to derive A^m eventually via some term f' which is a superposition of f and projections and is therefore also a polymorphism of $\mathcal{A}$. ◻

Remark 27. *An extended analysis of our proof should convince the careful reader that we may in the same fashion prove reactive composition (the polymorphism's action is determined for a row independently of the others). Thus, appealing to the previous section is not essential, though it does allow for a simpler argument.*

An interesting consequence of last section's formal work is a form of converse of Chen's Lemma, which allows us to give an algebraic characterisation of collapsibility from a singleton source.

Proposition 28. *Let x be a constant in $\mathcal{A}$. The following are equivalent:*

(i) *$\mathcal{A}$ is collapsible from $\{x\}$.*
(ii) *$\mathcal{A}$ has a Hubie-pol in $\{x\}$.*

Proof. Lemma 26 shows that (ii) implies collapsibility. We prove the converse.

Assume p-collapsibility. By Fact 13, we may apply Theorem 19. For $m = p + 1$, item (v) of this theorem states that there is a polymorphism f witnessing that $A^{p+1} \trianglelefteq \Upsilon_{p+1,p,x}$ (diagrammatically, we may draw a similar picture to the one we drew at the beginning of the previous proof). Clearly, f satisfies (ii). □

In the proof of the above, for $(i) \Rightarrow (ii) \Rightarrow (i)$, we no longer control the collapsibility parameter as the arity of our polymorphism is larger than the parameter we start with. By inspecting more carefully the properties of the polymorphism f we get as a witness that $\mathcal{A}$ models a canonical sentence, we may derive in fact p-collapsibility by an argument akin to the one used above in the proof of Chen's Lemma. We obtain this way a nice concrete result to counterbalance the abstract Theorem 19.

Theorem 29 (p-collapsibility from a singleton source). *Let x be a constant in $\mathcal{A}$ and $p > 0$. The following are equivalent.*

(i) *$\mathcal{A}$ is p-collapsible from $\{x\}$.*
(ii) *For every $m \geq 1$, the full adversary A^m is reactively composable from $\Upsilon_{m,p,x}$.*
(iii) *$\mathcal{A}$ is Π_2-p-collapsible from $\{x\}$.*
(iv) *For every $m \geq 1$, $\Upsilon_{m,p,x}$ generates A^m.*
(v) *$\mathcal{A}$ models $\varphi_{|A|,\Upsilon_{p+1,p,x},\mathcal{A}}$ (which implies that $\mathcal{A}$ admits a particularly well behaved Hubie-pol in $\{x\}$ of arity $(p+1)|A|^p$).*

Proof. Equivalence of the first four points appears in Corollary 23, as does the equivalence with the statement : For every $m \geq 1$, $\mathcal{A}$ models $\varphi_{n,\Upsilon_{m,p,x},\mathcal{A}}$. So they imply trivially the last point by selecting $m = p + 1$.

We show that the last point implies the penultimate one. The proof principle is similar to that of Chen's Lemma. As we have argued similarly before, the last point implies the existence of a polymorphism f. This polymorphism enjoys the following property (each column represents in fact n^p coordinates of A):

$$f\left(\begin{array}{c|c|c|c|c} \{x\} & A & A & \dots & A \\ A & \{x\} & A & \dots & A \\ \vdots & & \ddots & & \vdots \\ A & \dots & A & \{x\} & A \\ A & \dots & A & A & \{x\} \end{array}\right) = \begin{pmatrix} A \\ A \\ \vdots \\ A \\ A \end{pmatrix} = A^{p+1}$$

So arguing as in the proof of Chen's Lemma, we may conclude similarly that for all m, the full adversary A^m is composable from $\Upsilon_{m,p,x}$. □

Remark 30. *We say that a structure $\mathcal{A}$ is B-*conservative *where B is a subset of its domain iff for any polymorphism f of $\mathcal{A}$ and any $C \subseteq B$, we have $f(C, C, \dots, C) \subseteq C$. Provided that the structure is conservative on the source set B, we may prove a similar result for p-collapsibility from a conservative source.*

2.4.1 The curious case of 0*-collapsibility.* Expanding on Remark 24, we note that if we forbid equalities in the input to a QCSP, then we can observe the natural case of 0-collapsibility, to which now we turn. This is not a significant restriction in a context of complexity, since in all but trivial cases of a one element domain, one can propagate equality out through renaming of variables.

We investigated a similar notion in the context of positive equality-free first-order logic, the syntactic restriction of first-order logic that consists of sentences using only $\exists, \forall, \wedge$ and $\vee$. For this logic, relativisation of quantifiers fully explains the complexity classification of the model checking problem (a tetrachotomy between Pspace-complete, NP-complete, co-NP-complete and

Logspace) [26]. In particular, a complexity in NP is characterised algebraically by the preservation of the structure by a *simple A-shop* (to be defined shortly), which is equivalent to a strong form of 0-collapsibility since it applies not only to pH-sentences but also to sentences of positive equality free first-order logic. We will show that this notion corresponds in fact to 0-collapsibility from a singleton source. Let us recall first some definitions.

A *shop* on a set B, short for surjective hyper-operation, is a function f from B to its powerset such that $f(x) \neq \emptyset$ for any x in B and for every y in B, there exists x in B such that $f(x) \ni y$. An *A-shop*[4] satisfies further that there is some x such that $f(x) = B$. A *simple A-shop* satisfies further that $|f(x')| = 1$ for every $x' \neq x$. We say that a shop f is a *she* of the structure $\mathcal{B}$, short for *surjective hyper-endomorphism*, iff for any relational symbol R in σ of arity r, for any elements $a_1, a_2 \dots, a_r$ in B, if $R(a_1, \dots, a_r)$ holds in $\mathcal{B}$ then $R(b_1, \dots, b_r)$ holds in $\mathcal{B}$ for any $b_1 \in f(a_1), \dots, b_r \in f(a_r)$. We say that $\mathcal{B}$ *admits a (simple) A-she* if there is a (simple) A-shop f that is a she of $\mathcal{B}$.

Theorem 31. *Let $\mathcal{B}$ be a finite structure. The following are equivalent.*

(i) *$\mathcal{B}$ is 0-collapsible from source $\{x\}$ for some x in B for equality-free pH-sentences.*
(ii) *$\mathcal{B}$ admits a simple A-she.*
(iii) *$\mathcal{B}$ is 0-collapsible from source $\{x\}$ for some x in B for sentences of positive equality free first-order logic.*

Proof. The last two points are equivalent [27, Theorem 8] (this result is stated with A-she rather than simple A-she but clearly, $\mathcal{A}$ has an A-she iff it has a simple A-she). The implication (ii) to (i) follows trivially.

We prove the implication (i) to (ii) by contraposition. Assume that $A = [n] = \{1, \dots, n\}$ and suppose that $\mathcal{A}$ has no simple A-she. We will prove that $\mathcal{A}$ does not admit universal relativisation to x for pH-sentences. We assume also w.l.o.g. that $x = 1$. Let Ξ be the set of simple A-shops ξ s.t. $\xi(1) = [n]$. Since each ξ is not a she of $\mathcal{A}$, we have a quantifier-free formula with $2n-1$ variables R_ξ that consists of a single positive atom (not all variables need appear explicitly in this atom) such that $\mathcal{A} \models R_\xi(1, \dots, 1, 2, \dots, n)$,[5] but $\mathcal{A} \not\models R_\xi(\xi^1, \dots, \xi^n, \xi(2), \dots, \xi(n))$ for some $\xi^1, \dots, \xi^n \in [n] = \xi(1)$.

This means that for each $\eta : \{2, \dots, n\} \longrightarrow [n]$ there is some $2n-1$-ary "atom" R_η such that $\mathcal{A} \models R_\eta(1, \dots, 1, 1, 2, \dots, n)$,[6] but $\mathcal{A} \not\models R_\eta(\xi^1, \dots, \xi^n, \eta(2), \dots, \eta(n))$ for some $\xi^1, \dots, \xi^n \in [n]$. Let $\mathrm{E} = [n]^{[n-1]}$ denotes the set of ηs.

Suppose we had universal relativisation to 1. Then we know that

$$\mathcal{A} \models \bigwedge_{\eta \in \mathrm{E}} R_\eta(1, \dots, 1, 1, 2, \dots, n),$$

that is,

$$\mathcal{A} \models \exists y_1, \dots, y_n \bigwedge_{\eta \in \mathrm{E}} R_\eta(1, \dots, 1, y_1, y_2, \dots, y_n).$$

According to relativisation this means also that

$$\mathcal{A} \models \exists y_1, \dots, y_n \forall x_1, \dots, x_n \bigwedge_{\eta \in \mathrm{E}} R_\eta(x_1, \dots, x_n, y_1, y_2, \dots, y_n).$$

But we know

$$\mathcal{A} \models \forall y_1, \dots, y_n \exists x_1, \dots, x_n \bigvee_{\eta \in \mathrm{E}} \neg R_\eta(x_1, \dots, x_n, y_1, y_2, \dots, y_n),$$

since the ηs range over all maps $[n]$ to $[n]$. Contradiction. $\square$

[4]The A does not stand for the name of the set, it is short for *All*.
[5]There are n ones.
[6]There are n ones.

The above applies to singleton source only, but up to taking a power of a structure (which satisfies the same QCSP), we may always place ourselves in this singleton setting for 0-collapsibility.

THEOREM 32. *Let $\mathcal{B}$ be a structure. The following are equivalent.*

(i) *$\mathcal{B}$ is 0-collapsible from source C*
(ii) *$\mathcal{B}^{|C|}$ is 0-collapsible from some (any) singleton source x which is a (rainbow) $|C|$-tuple containing all elements of C.*

PROOF. Let $B = \{1, 2, \ldots, b\}$.

- (downwards). Let x be $|B|$-tuple containing all elements of B, w.l.o.g. $x = (1, 2, \ldots, b)$. Let φ be a pH sentence. Assume that $\mathcal{A}^{|B|} \models \varphi_{\restriction(x,x,\ldots,x)}$. Equivalently, for any i in B, $\mathcal{A} \models \varphi_{\restriction(i,i,\ldots,i)}$. Thus, 0-collapsibility from source B implies that $\mathcal{A} \models \varphi$. Since A and its power satisfy the same pH-sentences[22] we may conclude that $\mathcal{A}^{|B|} \models \varphi$.
- (upwards). Assume that for any i in B, $\mathcal{A} \models \varphi_{\restriction(i,i,\ldots,i)}$. Equivalently, $\mathcal{A}^{|B|} \models \varphi_{\restriction(x,x,\ldots,x)}$ where x is any $|B|$-tuple containing all elements of B. By assumption, $\mathcal{A}^{|B|} \models \varphi$ and we may conclude that $\mathcal{A} \models \varphi$.

□

2.5 Issues of decidability

The following is a corollary of Theorem 29.

COROLLARY 33. *Given $p \geq 1$, a structure $\mathcal{A}$ and x a constant in $\mathcal{A}$, we may decide whether $\mathcal{A}$ is p-collapsible from $\{x\}$.*

PROOF. We use Case (v) of Theorem 29. We construct $\varphi_{|A|,\Upsilon_{p+1,p,x},\mathcal{A}}$ explicitly then test if it is true on $\mathcal{A}$. □

We are not aware of a similar decidability result when the source is not a singleton. Neither are we aware of a decision procedure for collapsibility in general (when the p is not specified).

The case of switchability in general can be answered by [2]. Let α, β be strict subsets of A so that $\alpha \cup \beta = A$. A k-ary operation f on A is said to be *$\alpha\beta$-projective* if there exists $i \in [k]$ so that $f(x_1, \ldots, x_k) \in \alpha$, if $x_i \in \alpha$, and $f(x_1, \ldots, x_k) \in \beta$, if $x_i \in \beta$. A constraint language $\mathcal{A}$, expanded with constants naming all the elements, is switchable iff there exists some α and β, strict subsets of A, so that $\alpha \cup \beta = A$ and some polymorphism of $\mathcal{A}$ is not $\alpha\beta$-projective ([2]). If the maximal number of tuples in a relation of $\mathcal{A}$ is m then only polymorphisms of arity m need be considered.

3 THE CHEN CONJECTURE FOR INFINITE LANGUAGES

3.1 NP-membership

We need to revisit Theorem 19 in the case of infinite languages (signatures) and switchability. We omit parts of the theorem that are not relevant to us.

THEOREM 34 (IN ABSTRACTO LEVAVI). *Let $\Omega = \left(\Omega_m\right)_{m\in\mathbb{N}}$ be the sequence of the set of all (k-)switching m-ary adversaries over the domain of $\mathcal{A}$, a finite-domain structure with an infinite signature. The following are equivalent.*

(i) *For every $m \geq 1$, for every pH-sentence ψ with m universal variables, $\mathcal{A} \models \psi_{\restriction\Omega_m}$ implies $\mathcal{A} \models \psi$.*
(vi) *For every $m \geq 1$, Ω_m generates A^m.*

PROOF. We know from Theorem 19 that the following are equivalent.

(i') For every finite-signature reduct $\mathcal{A}'$ of $\mathcal{A}$ and $m \geq 1$, for every pH-sentence ψ with m universal variables, $\mathcal{A}' \models \psi_{\restriction \Omega_m}$ implies $\mathcal{A}' \models \psi$.

(vi') For every finite-signature reduct $\mathcal{A}'$ of $\mathcal{A}$ and every $m \geq 1$, Ω_m generates $\mathrm{Pol}(\mathcal{A}')^m$.

Since it is clear that both $(i) \Rightarrow (i')$ and $(vi) \Rightarrow (vi')$, it remains to argue that $(i') \Rightarrow (i)$ and $(vi') \Rightarrow (vi)$.

[$(i') \Rightarrow (i)$.] By contraposition, if (i) fails then it fails on some specific pH-sentence ψ which only mentions a finite number of relations of $\mathcal{A}'$. Thus (i') also fails on some finite reduct of $\mathcal{A}'$ mentioning these relations.

[$(vi') \Rightarrow (vi)$.] Let m be given. Consider some chain of finite reducts $\mathcal{A}_1, \ldots, \mathcal{A}_i, \ldots$ of $\mathcal{A}$ so that each $\mathcal{A}_i$ is a reduct of $\mathcal{A}_j$ for $i < j$ and every relation of $\mathcal{A}$ appears in some $\mathcal{A}_i$. We can assume from (vi') that Ω_m generates $\mathrm{Pol}(\mathcal{A}_i)^m$, for each i. However, since the number of tuples $(a_1, \ldots, a_m)$ and operations mapping Ω_m pointwise to $(a_1, \ldots, a_m)$, witnessing generation in $\mathrm{Pol}(\mathcal{A}')^m$, is finite, the sequence of operations $(f^i_1, \ldots, f^i_{|A|^m})$ (where f^i_j witnesses generation of the jth tuple in A^m) witnessing these must have an infinitely recurring element as i tends to infinity. One such recurring element we call $(f_1, \ldots, f_{|A|^m})$ and this witnesses generation in $\mathrm{Pol}(\mathcal{A})^m$. □

Note that in $(vi') \Rightarrow (vi)$ above we did not need to argue uniformly across the different $(a_1, \ldots, a_m)$ and it is enough to find an infinitely recurring operation for each of these individually.

The following result is the infinite language counterpoint to Corollary 22, that follows from Theorem 34 just as Corollary 22 followed from Theorem 19.

Theorem 35. *Let $\mathbb{A}$ be an idempotent algebra on a finite domain A. If $\mathbb{A}$ satisfies PGP, then QCSP*$(\mathrm{Inv}(\mathbb{A}))$ *reduces to a polynomial number of instances of CSP*$(\mathrm{Inv}(\mathbb{A}))$ *and is in NP.*

3.2 co-NP-hardness

Suppose there exist α, β strict subsets of A so that $\alpha \cup \beta = A$, define the relation $\tau_k(x_1, y_1, z_1 \ldots, x_k, y_k, z_k)$ by

$$\tau_k(x_1, y_1, z_1 \ldots, x_k, y_k, z_k) := \rho'(x_1, y_1, z_1) \vee \ldots \vee \rho'(x_k, y_k, z_k),$$

where $\rho'(x, y, z) = (\alpha \times \alpha \times \alpha) \cup (\beta \times \beta \times \beta)$. Strictly speaking, the α and β are parameters of τ_k but we dispense with adding them to the notation since they will be fixed at any point in which we invoke the τ_k. The purpose of the relations τ_k is to encode co-NP-hardness through the complement of the problem (monotone) 3-*not-all-equal-satisfiability* (3NAESAT). Let us introduce also the important relations $\sigma_k(x_1, y_1, \ldots, x_k, y_k)$ defined by

$$\sigma_k(x_1, y_1, \ldots, x_k, y_k) := \rho(x_1, y_1) \vee \ldots \vee \rho(x_k, y_k),$$

where $\rho(x, y) = (\alpha \times \alpha) \cup (\beta \times \beta)$.

Lemma 36. *The relation τ_k is pp-definable in σ_k.*

Proof. We will argue that τ_k is definable by the conjunction Φ of 3^k instances of σ_k that each consider the ways in which two variables may be chosen from each of the (x_i, y_i, z_i), i.e. $x_i \sim y_i$ or $y_i \sim z_i$ or $x_i \sim z_i$ (where $\sim$ is infix for ρ). We need to show that this conjunction Φ entails τ_k (the converse is trivial). We will assume for contradiction that Φ is satisfied but τ_k not. In the first instance of σ_k of Φ some atom must be true, and it will be of the form $x_i \sim y_i$ or $y_i \sim z_i$ or $x_i \sim z_i$. Once we have settled on one of these three, $p_i \sim q_i$, then we immediately satisfy 3^{k-1} of the conjunctions of Φ, leaving $2 \cdot 3^{k-1}$ unsatisfied. Now we can evaluate to true no more than one other among $\{x_i \sim y_i, y_i \sim z_i, x_i \sim z_i\} \setminus \{p_i \sim q_i\}$, without contradicting our assumptions. If we do evaluate this to true also, then we leave 3^{k-1} conjunctions unsatisfied. Thus we are now down to looking at variables with subscript other than i and in this fashion we have made the space one smaller, in total $k-1$. Now, we will need to evaluate in Φ some other atom of the form $x_j \sim y_j$ or $y_j \sim z_j$ or $x_j \sim z_j$, for $j \neq i$. Once we have settled on at most two of these three then we immediately satisfy 3^{k-2} of the conjunctions remaining of Φ, leaving 3^{k-2} still unsatisfied. Iterating this thinking, we arrive at a situation in which 1 clause is unsatisfied after we have gone through all k subscripts, which is a contradiction. □

Theorem 37. *Let $\mathbb{A}$ be an idempotent algebra on a finite domain A. If $\mathbb{A}$ satisfies EGP, then QCSP*$(\mathrm{Inv}(\mathbb{A}))$ *is co-NP-hard.*

Proof. We know from Lemma 11 in [2] that there exist α, β strict subsets of A so that $\alpha \cup \beta = A$ and the relation σ_k is in $\mathrm{Inv}(\mathbb{A})$, for each $k \in \mathbb{N}$. From Lemma 36, we know also that τ_k is in $\mathrm{Inv}(\mathbb{A})$, for each $k \in \mathbb{N}$.

We will next argue that τ_k enjoys a relatively small specification in DNF (at least, polynomial in k). We first give such a specification for $\rho'(x, y, z)$.

$$\rho'(x,y,z) := \bigvee_{a,a',a'' \in \alpha} x = a \wedge y = a' \wedge z = a'' \vee \bigvee_{b,b',b'' \in \beta} x = b \wedge y = b' \wedge z = b''$$

which is constant in size when A is fixed. Now it is clear from the definition that the size of τ_n is polynomial in n.

We will now give a very simple reduction from the complement of 3NAESAT to QCSP$(\mathrm{Inv}(\mathbb{A}))$. 3NAESAT is well-known to be NP-complete [28] and our result will follow.

Take an instance φ of 3NAESAT which is the existential quantification of a conjunction of k atoms $\mathrm{NAE}(x, y, z)$. Thus $\neg\varphi$ is the universal quantification of a disjunction of k atoms $x = y = z$. We build our instance ψ of QCSP$(\mathrm{Inv}(\mathbb{A}))$ from $\neg\varphi$ by transforming the quantifier-free part $x_1 = y_1 = z_1 \vee \ldots \vee x_k = y_k = z_k$ to $\tau_k = \rho'(x_1, y_1, z_1) \vee \ldots \vee \rho'(x_k, y_k, z_k)$.

($\neg\varphi \in$ co-3NAESAT implies $\psi \in$ QCSP$(\mathrm{Inv}(\mathbb{A}))$.) From an assignment to the universal variables $v_1, \ldots, v_m$ of ψ to elements $x_1, \ldots, x_m$ of A, consider elements $x'_1, \ldots, x'_m \in \{0, 1\}$ according to

- $x_i \in \alpha \setminus \beta$ implies $x'_i = 0$,
- $x_i \in \beta \setminus \alpha$ implies $x'_i = 1$, and
- $x_i \in \alpha \cap \beta$ implies we don't care, so w.l.o.g. say $x'_i = 0$.

The disjunct that is satisfied in the quantifier-free part of $\neg\varphi$ now gives the corresponding disjunct that will be satisfied in τ_k.

($\psi \in$ QCSP$(\mathrm{Inv}(\mathbb{A}))$ implies $\neg\varphi \in$ co-3NAESAT.) From an assignment to the universal variables $v_1, \ldots, v_m$ of $\neg\varphi$ to elements $x_1, \ldots, x_m$ of $\{0, 1\}$, consider elements $x'_1, \ldots, x'_m \in A$ according to

- $x_i = 0$ implies x'_i is some arbitrarily chosen element in $\alpha \setminus \beta$, and
- $x_i = 1$ implies x'_i is some arbitrarily chosen element in $\beta \setminus \alpha$.

The disjunct that is satisfied in τ_k now gives the corresponding disjunct that will be satisfied in the quantifier-free part of $\neg\varphi$. □

The demonstration of co-NP-hardness in the previous theorem was inspired by a similar proof in [29]. Note that an alternative proof that τ_k is in $\mathrm{Inv}(\mathbb{A})$ is furnished by the observation that it is preserved by all $\alpha\beta$-projections (see [2]). We note surprisingly that co-NP-hardness in Theorem 37 is optimal, in the sense that some (but not all!) of the cases just proved co-NP-hard are also in co-NP.

Proposition 38. *Let α, β be strict subsets of $A := \{a_1, \ldots, a_n\}$ so that $\alpha \cup \beta = A$ and $\alpha \cap \beta \neq \emptyset$. Then QCSP$(A; \{\tau_k : k \in \mathbb{N}\}, a_1, \ldots, a_n)$ is in co-NP.*

Proof. Assume $|A| > 1$, i.e. $n > 1$ (note that the proof is trivial otherwise). Let φ be an input to QCSP$(A; \{\tau_k : k \in \mathbb{N}\}, a_1, \ldots, a_n)$. We will now seek to eliminate atoms $v = a$ ($a \in \{a_1, \ldots, a_n\}$)

from φ. Suppose φ has an atom $v = a$. If v is universally quantified, then φ is false (since $|A| > 1$). Otherwise, either the atom $v = a$ may be eliminated with the variable v since v does not appear in a non-equality relation; or φ is false because there is another atom $v = a'$ for $a \neq a'$; or $v = a$ may be removed by substitution of a into all non-equality instances of relations involving v. This preprocessing procedure is polynomial and we will assume w.l.o.g. that φ contains no atoms $v = a$. We now argue that φ is a yes-instance iff φ' is a yes-instance, where φ' is built from φ by instantiating all existentially quantified variables as any $a \in \alpha \cap \beta$. The universal φ' can be evaluated in co-NP (one may prefer to imagine the complement as an existential $\neg\varphi'$ to be evaluated in NP) and the result follows. □

In fact, this being an algebraic paper, we can even do better. Let $\mathcal{B}$ signify a set of relations on a finite domain but not necessarily itself finite. For convenience, we will assume the set of relations of $\mathcal{B}$ is closed under all co-ordinate projections and instantiations of constants at specified coordinates. Call $\mathcal{B}$ *existentially trivial* if (in addition to the closure property just described) there exists an element $c \in B$ (which we call a *canon*) such that for each k-ary relation R of $\mathcal{B}$ and each $i \in [k]$, and for every $x_1, \ldots, x_k \in B$, whenever $(x_1, \ldots, x_{i-1}, x_i, x_{i+1}, \ldots, x_k) \in R^{\mathcal{B}}$ then also $(x_1, \ldots, x_{i-1}, c, x_{i+1}, \ldots, x_k) \in R^{\mathcal{B}}$. We want to expand this class to *almost existentially trivial* by permitting conjunctions of the form $v = a_i$ or $v = v'$ with relations that are existentially trivial.

Lemma 39. *Let α, β be strict subsets of $A := \{a_1, \ldots, a_n\}$ so that $\alpha \cup \beta = A$ and $\alpha \cap \beta \neq \emptyset$. The set of relations pp-definable in $(A; \{\tau_k : k \in \mathbb{N}\}, a_1, \ldots, a_n)$ is almost existentially trivial.*

Proof. Let us first note that $(A; \{\tau_k : k \in \mathbb{N}\}, a_1, \ldots, a_n)$ is existentially trivial with canon any $c \in \alpha \cap \beta$. Consider a formula with a pp-definition in $(A; \{\tau_k : k \in \mathbb{N}\}, a_1, \ldots, a_n)$. We assume that only free variables appear in equalities since otherwise we can remove these equalities by substitution. Now existential quantifiers can be removed and their variables instantiated as the canon c. Thus we are left with a conjunction of equalities and atoms τ_n, and the result follows. □

Proposition 40. *If $\mathcal{B}$ is comprised exclusively of relations that are almost existentially trivial, then QCSP($\mathcal{B}$) is in co-NP under the* ***DNF encoding****.*

Proof. The argument here is quite similar to that of Proposition 38 except that there is some additional preprocessing to find out variables that are forced in some relation to being a single constant or pairs of variables within a relation that are forced to be equal. In the first instance that some variable is forced to be constant in a k-ary relation, we should replace with the $(k-1)$-ary relation with the requisite forcing. In the second instance that a pair of variables are forced equal then we replace again the k-ary relation with a $(k-1)$-ary relation as well as an equality. Note that projecting a relation to a single or two co-ordinates can be done in polynomial time because the relations are encoded in DNF. After following these rules to their conclusion one obtains a conjunction of equalities together with relations that are existentially trivial. Now is the time to propagate variables to remove equalities (or find that there is no solution). Finally, when only existentially trivial relations are left, all remaining existential variables may be evaluated to the canon c. □

Corollary 41. *Let α, β be strict subsets of $A := \{a_1, \ldots, a_n\}$ so that $\alpha \cup \beta = A$ and $\alpha \cap \beta \neq \emptyset$. Then QCSP*$(\mathrm{Inv}(\mathrm{Pol}(A; \{\tau_k : k \in \mathbb{N}\}, a_{,} \ldots, a_n)))$ *is in co-NP under the* ***DNF encoding****.*

This last result, together with its supporting proposition, is the only time we seem to require the "nice, simple" DNF encoding, rather than arbitrary propositional logic. We do not require DNF for Proposition 38 as we have just a single relation in the signature for each arity and this is easy to keep track of. We note that the set of relations $\{\tau_k : k \in \mathbb{N}\}$ is not maximal with the property that with the constants it forms a co-clone of existentially trivial relations. One may add, for example, $(\alpha \times \beta) \cup (\beta \times \alpha)$.

The following, together with our previous results, gives the refutation of the Alternative Chen Conjecture (Conjecture 3).

Proposition 42. *Let α, β be strict subsets of $A := \{a_1, \ldots, a_n\}$ so that $\alpha \cup \beta = A$ and $\alpha \cap \beta \neq \emptyset$. Then, for each finite signature reduct $\mathcal{B}$ of $(A; \{\tau_k : k \in \mathbb{N}\}, a_1, \ldots, a_n)$, QCSP($\mathcal{B}$) is in NL.*

Proof. We will assume $\mathcal{B}$ contains all constants (since we prove this case gives a QCSP in NL, it naturally follows that the same holds without constants). Take m so that, for each $\tau_i \in \mathcal{B}$, $i \leq m$. Recall from Lemma 36 that τ_i is pp-definable in σ_i. We will prove that the structure $\mathcal{B}'$ given by $(A; \{\sigma_k : k \leq m\}, a_1, \ldots, a_n)$ admits a $(2m+1)$-ary near-unanimity operation f as a polymorphism, whereupon it follows that $\mathcal{B}$ admits the same near-unanimity polymorphism. We choose f so that all tuples whose map is not automatically defined by the near-unanimity criterion map to some arbitrary $a \in \alpha \cap \beta$. To see this, imagine that this f were not a polymorphism. Then some $(2m+1)$ m-tuples in σ_i would be mapped to some tuple not in σ_i which must be a tuple $\bar{t}$ of elements from $(\alpha \setminus \beta) \cup (\beta \setminus \alpha)$. Note that column-wise this map may only come from $2m+1$-tuples that have $2m$ instances of the same element. By the pigeonhole principle, the tuple $\bar{t}$ must appear as one of the $(2m+1)$ m-tuples in σ_i and this is clearly a contradiction.

It follows from [5] that an instance of QCSP($\mathcal{B}$) with p universal variables reduces to a polynomially bounded ensemble of $\binom{p}{2m} \cdot n \cdot n^{2m}$ instances of CSP($\mathcal{B}$) (if $p < 2m$, this is just $n \cdot n^{2m}$), and the result follows. □

Let us note that it is now known there exists a finite Δ on 3 elements so that Pol(Δ) has EGP, yet QCSP(Δ) is in P [4].

3.3 The question of the tuple-listing encoding

Proposition 43. *Let $\alpha := \{0, 1\}$ and $\beta := \{0, 2\}$. Then, QCSP($\{0, 1, 2\}; \{\tau_k : k \in \mathbb{N}\}, 0, 1, 2$) is in P under the **tuple-listing encoding**.*

Proof. Consider an instance φ of this QCSP of size n involving relation τ_m but no relation τ_k for $k > m$. The number of tuples in τ_m is $> 3^m$. Following Proposition 38 together with its proof, we may assume that the instance is strictly universally quantified over a conjunction of atoms (involving also constants). Now, a universally quantified conjunction is true iff the conjunction of its universally quantified atoms is true. We can further say that there are at most n atoms each of which involves at most $3m$ variables. Therefore there is an exhaustive algorithm that takes at most $O(n \cdot 3^{3m})$ steps which is $O(n^4)$. □

The proof of Proposition 43 suggests an alternative proof of Proposition 42, but placing the corresponding QCSP in P instead of NL. Proposition 43 shows that Chen's Conjecture fails for the tuple encoding in the sense that it provides a language $\mathcal{B}$, expanded with constants naming all the elements, so that Pol($\mathcal{B}$) has EGP, yet QCSP($\mathcal{B}$) is in P under the tuple-listing encoding. However, it does not imply that the algebraic approach to QCSP violates Chen's Conjecture under the tuple encoding. This is because $(\{0, 1, 2\}; \{\tau_k : k \in \mathbb{N}\}, 0, 1, 2)$ is not of the form Inv($\mathbb{A}$) for some idempotent algebra $\mathbb{A}$. For this stronger result, we would need to prove QCSP(Inv(Pol($\{0, 1, 2\}; \{\tau_k : k \in \mathbb{N}\}, 0, 1, 2$))) is in P under the tuple-listing encoding. However, such a violation to Chen's Conjecture under the tuple-listing encoding is now known from [4].

4 FINITE SUBSETS Δ OF $\mathrm{Inv}(r)$ ARE SUCH THAT $\mathrm{Pol}(\Delta)$ IS COLLAPSIBLE

In this section, we assume that all relations are defined on the finite set $\{0, 1, 2\}$. We will consider the 4-ary idempotent operation r defined by Chen in [1].

$$\begin{array}{ccc} 0111 & & 1 \\ 1011 & r & 1 \\ 0001 & \mapsto & 0 \\ 0010 & & 0 \\ \text{else} & & 2. \end{array}$$

Chen proved that $(\{0, 1, 2\}; r)$ is 2-switchable but not k-collapsible, for any k [1]. We will prove that, for all finite subsets $\Delta \subset \mathrm{Inv(r)}$, $\mathrm{Pol}(\Delta)$ is collapsible.

Define $s(x, y) := r(x, x, y, y)$. Then s is a semilattice-without-unit operation and plays a pivotal role in the 3-element QCSP classification [4, 5, 30].

A relation ρ is called *essential* if it cannot be represented as a conjunction of relations with smaller arities. A tuple $(a_1, a_2, \ldots, a_n)$ is called *essential for a relation* ρ if $(a_1, a_2, \ldots, a_n) \notin \rho$ and for every $i \in \{1, 2, \ldots, n\}$ there exists $b \in A$ such that $(a_1, \ldots, a_{i-1}, b, a_{i+1}, \ldots, a_n) \in \rho$. Let us define a relation $\tilde{\rho}$ for every relation $\rho \subseteq D^n$. Put $\sigma_i(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n) := \exists y\ \rho(x_1, \ldots, x_i, y, x_{i+1}, \ldots, x_n)$ and let

$$\tilde{\rho}(x_1, \ldots, x_n) := \sigma_1(x_2, x_3, \ldots, x_n) \wedge \sigma_2(x_1, x_3, \ldots, x_n) \wedge \ldots \wedge \sigma_n(x_1, x_2, \ldots, x_{n-1}).$$

Lemma 44. *A relation ρ is essential iff there exists an essential tuple for ρ.*

Proof. (Forwards.) By contraposition, if ρ is not essential, then $\tilde{\rho}$ is equivalent to ρ, and there can not be an essential tuple.

(Backwards.) An essential tuple witnesses that a relation is essential via $\tilde{\rho}$. □

Note that ρ being essential is equivalent also to $\rho \neq \tilde{\rho}$.

Lemma 45. *Suppose $(2, 2, x_3, \ldots, x_n)$ is an essential tuple for ρ. Then ρ is not preserved by s.*

Proof. Since $(2, 2, x_3, \ldots, x_n)$ is an essential tuple, $(x_1, 2, x_3, \ldots, x_n)$ and $(2, x_2, x_3, \ldots, x_n)$ are in ρ for some x_1 and x_2. But applying s now gives the contradiction. □

For a tuple $\mathbf{y}$, we denote its ith co-ordinate by $\mathbf{y}(i)$. For $n \geq 3$, we define the arity $n+1$ idempotent operation f_n as follows

$$\begin{array}{c} f_n(0, 0 \ldots, 0, 0) = 0 \\ f_n(1, 1, \ldots, 1, 1) = 1 \\ f_n(1, 0, \ldots, 0, 0) = 0 \\ f_n(0, 1, \ldots, 0, 0) = 0 \\ \vdots \\ f_n(0, 0, \ldots, 1, 0) = 0 \\ f_n(0, 0, \ldots, 0, 1) = 0 \\ \text{else } 2 \end{array}$$

The functions f_n are very similar to partial near-unanimity functions.

Lemma 46. *Consider $\mathbb{A} := (\{0, 1, 2\}; r)$. Then any relation $\rho \in \mathrm{Inv}(\mathbb{A})$ of arity $h < n+1$ is preserved by f_n.*

Proof. We prove this statement for a fixed n by induction on h. For $h = 1$ we just need to check that f_n preserves the unary relations $\{0, 2\}$ and $\{1, 2\}$, as these (and the full and empty relations) are the only unary relations that are in $\mathrm{Inv}(\mathbb{A})$.

Assume that ρ is not preserved by f_n. Then there exist tuples $\mathbf{y}_1, \ldots, \mathbf{y}_{n+1} \in \rho$ such that $f_n(\mathbf{y}_1, \ldots, \mathbf{y}_{n+1}) = \gamma \notin \rho$. We consider a matrix whose columns are $\mathbf{y}_1, \ldots, \mathbf{y}_{n+1}$. Let the rows of this matrix be $\mathbf{x}_1, \ldots, \mathbf{x}_h$.

By the inductive assumption every σ_i from the definition of $\widetilde{\rho}$ is preserved by f_n, which means that $\widetilde{\rho}$ is preserved by f_n, which means, since $\gamma \notin \rho$, that γ is an essential tuple for ρ.

We consider two cases. First, assume that γ doesn't contain 2. Then it follows from the definition that every $\mathbf{x}_i$ contains at most one element that differs from $\gamma(i)$. Since $n + 1 > h$, there exists $i \in \{1, 2, \ldots, n + 1\}$ such that $\mathbf{y}_i = \gamma$. This contradicts the fact that $\gamma \notin \rho$.

Second, assume that γ contains 2. Then by Lemma 45, γ contains exactly one 2. Without loss of generality, we assume that $\gamma(1) = 2$. It follows from the definition of f_n that $\mathbf{x}_i$ contains at most one element that differs from $\gamma(i)$ for every $i \in \{2, 3, \ldots, h\}$. Hence, since $n + 1 > h$, for some $k \in \{1, 2, \ldots, n + 1\}$ we have $\mathbf{y}_k(i) = \gamma(i)$ for every $i \in \{2, 3, \ldots, h\}$. Since $f_n(\mathbf{x}_1) = 2$, we have one of three subcases. First subcase, $\mathbf{x}_1(j) = 2$ for some j. We need one of the properties

	$\mathbf{y}_k$	$\mathbf{y}_j$	γ
$\mathbf{x}_1$	0	2	2
	0	0	0
	0	1	0
	1	1	1

	$\mathbf{y}_k$	$\mathbf{y}_j$	γ
$\mathbf{x}_1$	1	2	2
	0	0	0
	0	1	0
	1	1	1

depending on whether $\mathbf{y}_k(1)$ is 0 or 1 (it cannot be 2). We then obtain $\gamma = r(y_k, y_j, y_j, y_j) \in \rho$ – a contradiction.

Second subcase, $\mathbf{y}_k(1) = 1, \mathbf{y}_m(1) = 0$ for some $m \in \{1, 2, \ldots, n + 1\}$. We need the property

	$\mathbf{y}_k$	$\mathbf{y}_m$	γ
$\mathbf{x}_1$	1	0	2
	0	0	0
	0	1	0
	1	1	1

and we obtain $\gamma = r(y_k, y_k, y_k, y_m) \in \rho$ – a contradiction.

Third subcase, $\mathbf{y}_k(1) = 0, \mathbf{y}_m(1) = 1$ and $\mathbf{y}_l(1) = 1$ for $m, l \in \{1, 2, \ldots, n + 1\} \setminus \{k\}$, $m \neq l$. We need the property

	$\mathbf{y}_k$	$\mathbf{y}_m$	$\mathbf{y}_l$	γ
$\mathbf{x}_1$	0	1	1	2
	0	0	0	0
	0	0	1	0
	0	1	0	0
	1	1	1	1

and we can check that $\gamma = r(y_k, y_k, y_m, y_l) \in \rho$ – a contradiction. This completes the proof. □

Corollary 47. *Suppose* $\mathbb{A} := (\{0, 1, 2\}; r)$. *Then, for every finite subset* Δ *of Inv*$(\mathbb{A})$, *Pol*(Δ) *is collapsible.*

Proof. Let n be the maximal arity of the relations in Δ. Then f_n is a Hubie-pol in $\{1\}$, and the result follows from Lemma 26. □

5 CONCLUSION

One important application of our abstract investigation of PGP yields a nice characterisation in the concrete case of collapsibility, in particular in the case of a singleton source which we now know can be equated with preservation under a single polymorphism, namely a Hubie polymorphism. So far, this is the only known explanation for a complexity of a (finite signature) QCSP in NP which provokes the following two questions.

Question 48. *For a structure $\mathcal{A}$, is it the case that QCSP($\mathcal{A}$) is in NP iff $\mathcal{A}$ admits a Hubie polymorphism?*

Question 49. *For a structure $\mathcal{A}$, is it the case that QCSP($\mathcal{A}$) is in NP iff $\mathcal{A}$ is collapsible?*

Lurking between these questions is the question as to whether collapsibility is always existing from a singleton source (though a better parameter might be obtained from a larger source).

We can also phrase an important algebraic variant of these questions.

Question 50. *For an algebra $\mathbb{A}$ that is switchable, is it the case that for all finite subsets $\Delta \subset \mathrm{Inv}(\mathbb{A})$, $\mathrm{Pol}(\Delta)$ is collapsible?*

In the long version of this paper [31], we prove that this is true for all 3-element algebras. The possibility of an affirmative answer to this question justifies our continuing interest in collapsibility.

One can further wonder if the parameter p of collapsibility depends on the size of the structure $\mathcal{A}$. In particular, this would provide a positive answer to the following.

Question 51. *Given a structure $\mathcal{A}$, can we decide if it is p-collapsible for some p?*

Finally, since this paper was drafted, Zhuk has written a new, self-contained, short and elegant proof of Corollary 22, which can be found in [32].

ACKNOWLEDGMENT

The authors would like to thank several anonymous reviewers for their patience, stamina and very useful suggestions, which have been a great help to prepare the final version of this paper.

REFERENCES

[1] H. Chen, "Quantified constraint satisfaction and the polynomially generated powers property," *Algebra universalis*, vol. 65, no. 3, pp. 213–241, 2011, an extended abstract appeared in ICALP B 2008. [Online]. Available: http://dx.doi.org/10.1007/s00012-011-0125-4

[2] D. Zhuk, "The Size of Generating Sets of Powers," *Journal of Combinatorial Theory, Series A*, vol. 167, pp. 91–103, 2019.

[3] H. Chen, "Meditations on quantified constraint satisfaction," in *Logic and Program Semantics - Essays Dedicated to Dexter Kozen on the Occasion of His 60th Birthday*, 2012, pp. 35–49.

[4] D. Zhuk and B. Martin, "QCSP monsters and the demise of the Chen conjecture," in *Proccedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, K. Makarychev, Y. Makarychev, M. Tulsiani, G. Kamath, and J. Chuzhoy, Eds. ACM, 2020, pp. 91–104. [Online]. Available: https://doi.org/10.1145/3357713.3384232

[5] H. Chen, "The complexity of quantified constraint satisfaction: Collapsibility, sink algebras, and the three-element case," *SIAM J. Comput.*, vol. 37, no. 5, pp. 1674–1701, 2008.

[6] A. Bulatov, P. Jeavons, and A. Krokhin, "The complexity of constraint satisfaction: An algebraic approach (a survey paper)," *In: Structural Theory of Automata, Semigroups and Universal Algebra (Montreal, 2003), NATO Science Series II: Mathematics, Physics, Chemistry*, vol. 207, pp. 181–213, 2005.

[7] L. Barto, A. A. Krokhin, and R. Willard, "Polymorphisms, and how to use them," in *The Constraint Satisfaction Problem: Complexity and Approximability*, ser. Dagstuhl Follow-Ups, A. A. Krokhin and S. Zivný, Eds. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017, vol. 7, pp. 1–44. [Online]. Available: https://doi.org/10.4230/DFU.Vol7.15301.1

[8] L. Barto and M. Kozik, "Absorption in universal algebra and CSP," in *The Constraint Satisfaction Problem: Complexity and Approximability*, ser. Dagstuhl Follow-Ups, A. A. Krokhin and S. Zivný, Eds. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017, vol. 7, pp. 45–77. [Online]. Available: https://doi.org/10.4230/DFU.Vol7.15301.2

[9] F. R. Madelaine and B. Martin, "On the complexity of the model checking problem," *CoRR*, vol. abs/1210.6893, 2012, extended abstract appeared at LICS 2011 under the name "A Tetrachotomy for Positive First-Order Logic without Equality". [Online]. Available: http://arxiv.org/abs/1210.6893

[10] A. A. Bulatov, "A dichotomy theorem for nonuniform CSPs," in *Proceedings of FOCS'17*, 2017, arXiv:1703.03021.

[11] D. Zhuk, "A proof of CSP dichotomy conjecture," in *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, Oct 2017, pp. 331–342.

[12] D. Zhuk, "A proof of the CSP dichotomy conjecture," *J. ACM*, vol. 67, no. 5, pp. 30:1–30:78, 2020. [Online]. Available: https://doi.org/10.1145/3402029

[13] U. Egly, T. Eiter, H. Tompits, and S. Woltran, "Solving advanced reasoning tasks using quantified boolean formulas," in *Proc. 17th Nat. Conf. on Artificial Intelligence and 12th Conf. on Innovative Applications of Artificial Intelligence.* AAAI Press/ The MIT Press, 2000, pp. 417–422.

[14] J. Wiegold, "Growth sequences of finite semigroups," *Journal of the Australian Mathematical Society (Series A)*, vol. 43, pp. 16–20, 8 1987, communicated by H. Lausch. [Online]. Available: http://journals.cambridge.org/article_S1446788700028925

[15] V. Kolmogorov, A. A. Krokhin, and M. Rolinek, "The complexity of general-valued CSPs," in *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, 2015, pp. 1246–1258. [Online]. Available: http://dx.doi.org/10.1109/FOCS.2015.80

[16] C. Carvalho, B. Martin, and D. Zhuk, "The complexity of quantified constraints using the algebraic formulation," in *42nd International Symposium on Mathematical Foundations of Computer Science, MFCS 2017, August 21-25, 2017 - Aalborg, Denmark*, 2017, pp. 27:1–27:14. [Online]. Available: https://doi.org/10.4230/LIPIcs.MFCS.2017.27

[17] N. Creignou, S. Khanna, and M. Sudan, *Complexity Classifications of Boolean Constraint Satisfaction Problems.* SIAM Monographs on Discrete Mathematics and Applications 7, 2001.

[18] M. Bodirsky and J. Kára, "The complexity of equality constraint languages," *Theory of Computing Systems*, vol. 3, no. 2, pp. 136–158, 2008, a conference version appeared in the proceedings of CSR'06.

[19] H. Chen and P. Mayr, "Quantified constraint satisfaction on monoids," in *25th EACSL Annual Conference on Computer Science Logic, CSL 2016, August 29 - September 1, 2016, Marseille, France*, 2016, pp. 15:1–15:14. [Online]. Available: https://doi.org/10.4230/LIPIcs.CSL.2016.15

[20] D. Zhuk, "A modification of the CSP algorithm for infinite languages," *CoRR*, vol. abs/1803.07465, 2018. [Online]. Available: http://arxiv.org/abs/1803.07465

[21] C. Carvalho, F. R. Madelaine, and B. Martin, "From complexity to algebra and back: digraph classes, collapsibility and the PGP," in *30th Annual IEEE Symposium on Logic in Computer Science (LICS)*, 2015.

[22] H. Chen, F. Madelaine, and B. Martin, "Quantified constraints and containment problems," in *23rd Annual IEEE Symposium on Logic in Computer Science*, 2008, pp. 317–328.

[23] F. R. Madelaine and B. Martin, "Containment, equivalence and coreness from CSP to QCSP and beyond," in *Principles and Practice of Constraint Programming - 18th International Conference, CP 2012, Québec City, QC, Canada, October 8-12, 2012. Proceedings*, ser. Lecture Notes in Computer Science, M. Milano, Ed., vol. 7514. Springer, 2012, pp. 480–495. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-33558-7_36

[24] B. Martin, "QCSP on partially reflexive forests," in *Principles and Practice of Constraint Programming - 17th International Conference, CP 2011*, 2011.

[25] B. Martin and F. Madelaine, "Towards a trichotomy for quantified H-coloring," in *2nd Conf. on Computatibility in Europe, LNCS 3988*, 2006, pp. 342–352.

[26] F. R. Madelaine and B. Martin, "A tetrachotomy for positive first-order logic without equality," in *LICS*, 2011, pp. 311–320.

[27] ——, "The complexity of positive first-order logic without equality," in *LICS.* IEEE Computer Society, 2009, pp. 429–438.

[28] C. H. Papadimitriou, *Computational Complexity.* Addison-Wesley, 1994.

[29] M. Bodirsky and H. Chen, "Quantified equality constraints," *SIAM J. Comput.*, vol. 39, no. 8, pp. 3682–3699, 2010.

[30] F. Börner, A. A. Bulatov, H. Chen, P. Jeavons, and A. A. Krokhin, "The complexity of constraint satisfaction games and qcsp," *Inf. Comput.*, vol. 207, no. 9, pp. 923–944, 2009.

[31] C. Carvalho, B. Martin, and D. Zhuk, "The complexity of quantified constraints," *CoRR*, vol. abs/1701.04086, 2017. [Online]. Available: http://arxiv.org/abs/1701.04086

[32] D. Zhuk, "The complexity of the quantified csp having the polynomially generated powers property," *arXiv preprint arXiv:2110.09504*, 2021.