The complexity of quantified constraints using the algebraic formulation

Catarina Carvalho[1], Barnaby Martin[2], and Dmitriy Zhuk[3]

1 **School of Physics, Astronomy and Mathematics, University of Hertfordshire, U.K.**
2 **School of Engineering and Computing Sciences, Durham University, U.K.**
3 **Moscow State University, 119899 Moscow, Russia.**

Abstract

Let $\mathbb{A}$ be an idempotent algebra on a finite domain. We combine results of Chen [7], Zhuk [20] and Carvalho et al. [5] to argue that if $\mathbb{A}$ satisfies the polynomially generated powers property (PGP), then QCSP(Inv($\mathbb{A}$)) is in NP. We then use the result of Zhuk to prove a converse, that if Inv($\mathbb{A}$) satisfies the exponentially generated powers property (EGP), then QCSP(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 the moral correctness of what we term the Chen Conjecture (see [8]).

We examine in closer detail the situation for domains of size three. Over any finite domain, the only type of PGP that can occur is switchability. Switchability was introduced by Chen in [7] as a generalisation of the already-known Collapsibility [6]. For three-element domain algebras $\mathbb{A}$ that are Switchable, we prove that for every finite subset Δ of Inv($\mathbb{A}$), Pol(Δ) is Collapsible. The significance of this is that, for QCSP on finite structures (over three-element domain), all QCSP tractability explained by Switchability is already explained by Collapsibility.

Finally, we present a three-element domain complexity classification vignette, using known as well as derived results.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases Quantified Constraints, Computational Complexity, Universal Algebra, Constraint Satisfaction

Digital Object Identifier 10.4230/LIPIcs.MFCS.2017.27

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) and a number of celebrated results have been obtained through this method. One considers the problem CSP($\mathcal{B}$) in which it is asked whether an input sentence φ holds on $\mathcal{B}$, where φ is *primitive positive*, that is using only $\exists$, $\wedge$ and $=$. The CSP is one of a wide class of model-checking problems obtained from restrictions of first-order logic. For almost every one of these classes, we can give a complexity classification [14]: the two outstanding classes are CSPs and its popular extension *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 [11].

The outstanding conjecture in the area is that all finite-domain CSPs are either in P or are NP-complete, something surprising given these CSPs appear to form a large microcosm of NP, and NP itself is unlikely to have this dichotomy property. This Feder-Vardi conjecture [12], given more concretely in the algebraic language in [4], remains unsettled, but is now

42nd International Symposium on Mathematical Foundations of Computer Science (MFCS 2017).
Editors: Kim G. Larsen, Hans L. Bodlaender, and Jean-Francois Raskin; Article No. 27; pp. 27:1–27:13
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

known for large classes of structures. It is well-known 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 following is the merger of Conjectures 6 and 7 in [8] which we call the *Chen Conjecture*.

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

In [8], Conjecture 6 gives the NP membership and Conjecture 7 the Pspace-completeness. We now know from [20] and [5] that the NP membership of Conjecture 6 is indeed true. The most interesting result of this paper is Theorem 1 below, but note that we permit infinite signatures (languages) although our domains remain finite. This aspect of our work will be discussed in detail later.

▶ **Theorem 1** (Revised Chen Conjecture). *Let $\mathbb{A}$ be an idempotent algebra on a finite domain A. If $\mathbb{A}$ satisfies PGP, then QCSP*(Inv($\mathbb{A}$)) *is in NP. Otherwise, QCSP*(Inv($\mathbb{A}$)) *is co-NP-hard.*

Zhuk has previously proved [20] that only the cases PGP and EGP may occur, even in the non-idempotent case. With infinite languages, the NP-membership for Theorem 1 is no longer immediate from [5], but requires a little extra work. We are also able to refute the following form.

▶ Conjecture 2 (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 $\Delta \subseteq \text{Inv}(\mathbb{A})$, QCSP($\Delta$) is in NP. Otherwise, there exists a finite subset $\Delta \subseteq \text{Inv}(\mathbb{A})$ so that QCSP(Δ) is co-NP-hard.

In proving Theorem 1 we are saying that the complexity of QCSPs, with all constants included, is classified modulo the complexity of CSPs.

▶ **Corollary 2.** *Let $\mathbb{A}$ be an idempotent algebra on a finite domain A. Either QCSP*(Inv($\mathbb{A}$)) *is co-NP-hard or QCSP*(Inv($\mathbb{A}$)) *has the same complexity as CSP*(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 [13].

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 $\mathbb{A}, \mathbb{A}^2, \mathbb{A}^3, \ldots$ as $f_{\mathbb{A}}(1), f_{\mathbb{A}}(2), f_{\mathbb{A}}(3), \ldots$, respectively. A subset Λ of A^m is a generating set for $\mathbb{A}^m$ exactly if, for every $(a_1, \ldots, a_m) \in A^m$, there exists a k-ary term operation f of $\mathbb{A}$ and $(b^1_1, \ldots, b^1_m), \ldots, (b^k_1, \ldots, b^k_m) \in \Lambda$ so that $f(b^1_1, \ldots, b^k_1) = a_1, \ldots, f(b^1_m, \ldots, b^k_m) = a_m$. 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 [19] 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 so that $f_{\mathbb{A}} = \Omega(g)$ where $g(i) = b^i$.

In Hubie Chen's [7], a new link between algebra and QCSP was discovered. Chen's previous work in QCSP tractability largely involved the special notion of *Collapsibility* [6], but in [7] this was extended to a *computationally effective* version of the PGP. For a finite-domain, idempotent algebra $\mathbb{A}$, *k-collapsibility* may be seen as 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. *k-switchability* may be seen as 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 exists $a_i < \ldots < a_{k'}$, for $k' \leq k$, so that

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

where $x_1 = \ldots = x_{a_1-1}$, $x_{a_1} = \ldots = x_{a_2-1}$, $\ldots$, $x_{a_{k'}} = \ldots = x_{a_m}$. Thus, $a_1, a_2, \ldots, a_{k'}$ are the indices where the tuple switches value. Note that these are not the original definitions, which we will see shortly, but they are proved equivalent to the original definitions (at least for finite signatures) in [5]. Moreover, these are the definitions that we will use. We say that $\mathbb{A}$ is collapsible (switchable) if there exists k such that it is k-collapsible (k-switchable). We note that Zhuk uses this definition of switchability in [20] in which he proved that the only kind of PGP for finite-domain algebras is switchability.

Let us capitalise Collapsibility and Switchability to indicate Chen's original definitions from [7] are used, following an example for *arithmetic* versus *Arithmetic* by Raymond Smullyan in [18]. There is the potential for confusion at the start of the sentence but, as was the case with Smullyan, the two will transpire to be interchangeable throughout our discourse. It is straightforward to see that k-Switchability implies k-switchability and k-Collapsibility implies k-collapsibility. The converses, for finite signatures, also hold, but this requires rather more work [5]. For any finite algebra, k-Collapsibility implies k-Switchability, and for any 2-element algebra, k-Switchability implies k-Collapsibility. Chen originally introduced Switchability because he found a 3-element algebra that enjoyed the PGP but was not Collapsible [7]. He went on to prove that Switchability of $\mathbb{A}$ implies that the corresponding QCSP is in P, what one might informally state as QCSP(Inv($\mathbb{A}$)) in P, where Inv($\mathbb{A}$) can be seen as the structure over the same domain as $\mathbb{A}$ whose relations are precisely those that are preserved by (invariant under) all the operations of $\mathbb{A}$. However, the QCSP was traditionally defined only on finite sets of relations (else the question arises as to encoding), thus a more formal definition might be that, for any finite subset Δ of Inv($\mathbb{A}$), QCSP(Δ) is in P. What we prove in this paper is that, as far as the QCSP is concerned, Switchability on a three-element algebra $\mathbb{A}$ is something of a mirage. What we mean by this is that when $\mathbb{A}$ is Switchable, for all finite subsets Δ of Inv($\mathbb{A}$), already Pol(Δ) is Collapsible. Thus, for QCSP complexity for three-element structures, we do not need the additional notion of Switchability to explain tractability, as Collapsibility will already suffice. Since these notions were originally introduced in connection with the QCSP this is particularly surprising. Note that the parameter k of Collapsibility is unbounded over these increasing finite subsets Δ while the parameter of Switchability clearly remains bounded. In some way we are suggesting that Switchability itself might be seen as a limit phenomenon of Collapsibility.

1.1 Infinite languages

Our use of infinite languages (i.e. signatures, since we work on a finite domain) is the only 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 Inv($\mathbb{A}$) is unsatisfactory in the sense that choosing this way to escape the – naturally infinite – set Inv($\mathbb{A}$) is as arbitrary 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 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 [10], 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 [10] and a similar discussion in [2] 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 [9] 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. Note that our demonstration of the (Revised) Chen Conjecture for infinite languages with the "simple" logic encoding does not resolve the original Chen Conjecture for finite languages $\mathcal{B}$ with constants because $\mathrm{QCSP}(\mathrm{Inv}(\mathrm{Pol}(\mathcal{B})))$ could conceivably have higher complexity than $\mathrm{QCSP}(\mathcal{B})$ due to a succinct representation of relations in $\mathrm{Inv}(\mathrm{Pol}(\mathcal{B}))$. Indeed, this belies one justification for the preferential study of finite subsets of $\mathrm{Inv}(\mathrm{Pol}(\mathcal{B}))$, since for finite signature $\mathcal{B}$ we can then say $\mathrm{QCSP}(\mathcal{B})$ and $\mathrm{QCSP}(\mathrm{Inv}(\mathrm{Pol}\mathcal{B}))$ must have the same complexity. Note that for finite relational bases $\mathcal{B}', \mathcal{B}''$ of $\mathrm{Inv}(\mathrm{Pol}(\mathcal{B}))$, $\mathrm{QCSP}(\mathcal{B}')$ and $\mathrm{QCSP}(\mathcal{B}'')$ must have the same complexity. Further, we do not know of any concrete finite $\mathcal{B}$ with constants, so that $\mathrm{QCSP}(\mathrm{Inv}(\mathrm{Pol}(\mathcal{B})))$ and $\mathrm{QCSP}(\mathcal{B})$ have different complexity.

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 2 and Corollary 10 (which uses the "nice, simple" DNF) and Proposition 4 (which uses the tuple encoding).

Related work This paper is the merger of [16, 15], neither of which was submitted for publication, considerably extended.

2 Preliminaries

Let $[k] := \{1, \dots, k\}$. A *k-ary polymorphism* of a relational structure $\mathcal{B}$ is a homomorphism f from $\mathcal{B}^k$ to $\mathcal{B}$. Let $\mathrm{Pol}(\mathcal{B})$ be the set of polymorphisms of $\mathcal{B}$ and let $\mathrm{Inv}(\mathbb{A})$ be the set of relations on A which are invariant under (each of) the operations of some finite algebra $\mathbb{A}$. $\mathrm{Pol}(\mathcal{B})$ is an object known in Universal Algebra as a *clone*, which is a set of operations containing all projections and closed under composition (superposition). A *term operation* of an algebra $\mathbb{A}$ is an operation which is a member of the clone generated by $\mathbb{A}$.

We will conflate sets of operations over the same domain and algebras just as we do sets of relations over the same domain and constraint languages (relational structures). Indeed, the only technical difference between such objects is the movement away from an ordered signature, which is not something we will ever need. A *reduct* of a relational structure $\mathcal{B}$ is a relational structure $\mathcal{B}'$ over the same domain obtained by forgetting some of the relations. If Δ is some finite subset of $\text{Inv}(\mathbb{A})$, then we may view Δ a being a finite reduct of the structure (associated with) $\text{Inv}(\mathbb{A})$.

A k-ary operation f over A is a *projection* if $f(x_1, \dots, x_k) = x_i$, for some $i \in [k]$. When α, β are strict subsets of A so that $\alpha \cup \beta = A$, then a k-ary operation f on A is said to be *$\alpha\beta$-projective* if there exists $i \in [k]$ so that if $x_i \in \alpha$ (respectively, $x_i \in \beta$), then $f(x_1, \dots, x_k) \in \alpha$ (respectively, $f(x_1, \dots, x_k) \in \beta$).

We recall QCSP($\mathcal{B}$), where $\mathcal{B}$ is some structure on a finite-domain, is a decision problem with input ϕ, a pH-sentence (i.e. using just $\forall$, $\exists$, $\wedge$ and $=$) involving (a finite set of) relations of $\mathcal{B}$, encoded in propositional logic with equality and constants. The yes-instances are those ϕ for which $\mathcal{B} \models \phi$. If the input sentence is restricted to have alternation Π_k then the corresponding problem is designated Π_k-CSP($\mathcal{B}$).

2.1 Games, adversaries and reactive composition

We now recall some terminology due to Chen [6, 7], 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. This will enable us to give the original definitions for Collapsibility and Switchability. An *adversary* $\mathcal{B}$ of length $m \geq 1$ is an m-ary relation over A. When $\mathcal{B}$ is precisely the set $B_1 \times B_2 \times \ldots \times B_m$ for some non-empty subsets $B_1, B_2, \dots, B_m$ of A, we speak of a *rectangular adversary* (we will sometimes specify this as a tuple rather than a product). Let ϕ be a pH-sentence with universal variables $x_1, \dots, x_m$ and quantifier-free part ψ. We write $\mathcal{A} \models \phi_{\restriction \mathcal{B}}$ and say that *the existential player has a winning strategy in the* $(\mathcal{A}, \phi)$*-game against adversary* $\mathcal{B}$ iff there exists a set of Skolem functions $\{\sigma_x : \text{`}\exists x\text{'} \in \phi\}$ such that for any assignment π of the universally quantified variables of ϕ to A, where $\big(\pi(x_1), \dots, \pi(x_m)\big) \in \mathcal{B}$, 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, $\mathcal{A} \models \phi$ iff the existential player has a winning strategy in the $(\mathcal{A}, \phi)$-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 $\mathcal{B}$ of length m *dominates* an adversary $\mathcal{B}'$ of length m when $\mathcal{B}' \subseteq \mathcal{B}$. Note that $\mathcal{B}' \subseteq \mathcal{B}$ and $\mathcal{A} \models \phi_{\restriction \mathcal{B}}$ implies $\mathcal{A} \models \phi_{\restriction \mathcal{B}'}$. We will also consider sets of adversaries of the same length, denoted by uppercase Greek letters as in Ω_m (here the length is m); and, sequences thereof, which we denote with bold uppercase Greek letters as in $\mathbf{\Omega} = \big(\Omega_m\big)_{m \in \mathbb{N}}$. We will write $\mathcal{A} \models \phi_{\restriction \Omega_m}$ to denote that $\mathcal{A} \models \phi_{\restriction \mathcal{B}}$ holds for every adversary $\mathcal{B}$ in Ω_m.

We now introduce reactive composition as a means to obtain larger adversaries from a number of smaller adversaries. Let f be a k-ary operation of $\mathcal{A}$ and $\mathcal{A}, \mathcal{B}_1, \dots, \mathcal{B}_k$ be adversaries of length m. We say that $\mathcal{A}$ is *reactively composable* from the adversaries $\mathcal{B}_1, \dots, \mathcal{B}_k$ via f, and we write $\mathcal{A} \trianglelefteq f(\mathcal{B}_1, \dots, \mathcal{B}_k)$ 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 $(a_1, \dots, a_m)$ in adversary $\mathcal{A}$ the following holds.

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

We write $\mathcal{A} \trianglelefteq \{\mathcal{B}_1,\dots,\mathcal{B}_k\}$ if there exists a k-ary operation f such that $\mathcal{A} \trianglelefteq f(\mathcal{B}_1,\dots,\mathcal{B}_k)$ Reactive composition allows to interpolate complete Skolem functions from partial ones.

▶ **Theorem 3** ([7, Theorem 7.6]). *Let ϕ be a pH-sentence with m universal variables. Let $\mathcal{A}$ be an adversary and Ω_m a set of adversaries, both of length m.*

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

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 co-ordinates that are the set A and all the others 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 $\boldsymbol{\Upsilon}_{p,B}$ be the sequence of adversaries $(\Upsilon_{m,p,B})_{m\in\mathbb{N}}$. We will define a structure $\mathcal{A}$ to be *p-Collapsible from source B* iff for every m and for all pH-sentence ϕ with m universal variables, $\mathcal{A} \models \phi_{\restriction \Upsilon_{m,p,B}}$ implies $\mathcal{A} \models \phi$.

For p-Switchability, the set of adversaries will be of the form $\Xi_{m,p}$, where each adversary is built from the set of tuples that have some $k' < p$ switches at specific points $0 < a_1 < \dots < a_{k'} \leq m$.

For rectangular adversaries, such as $\Upsilon_{m,p,x}$, reactive composition is rather simpler than in the definition above, becoming just (ordinary) composition, as follows. $\mathcal{A}$ is *composable* from the adversaries $\mathcal{B}_1,\dots,\mathcal{B}_k$ via f if $f(B_1^i,\dots,B_i^k) \supseteq A^i$, where $\mathcal{A} = (A^1,\dots,A^m)$ and each $\mathcal{B}_j = (B_j^1,\dots,B_j^m)$. Reactive composition plays a key role in the proof of our main theorem but its use appears only in other papers that we will cite. Ordinary composition is the only type of reactive composition that will be used in this paper.

3 The Chen Conjecture

3.1 NP-membership

We need to revisit the main result of [5] to show that it holds not just for finite signatures but for infinite signatures also. In its original the following theorem discussed "projective sequences of adversaries, none of which are degenerate". This includes Switching adversaries and we give it in this latter form. We furthermore remove some parts of the theorem that are not currently relevant to us.

▶ **Theorem 4** (In abstracto [5]). *Let $\boldsymbol{\Omega} = \left(\Omega_m\right)_{m\in\mathbb{N}}$ be the sequence of the set of all (k-)Switching m-ary adversaries over the domain of $\mathcal{A}$, a finite structure. 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 $\mathrm{Pol}(\mathcal{A})^m$.

▶ **Corollary 5** (In abstracto levavi). *Let $\boldsymbol{\Omega} = \left(\Omega_m\right)_{m\in\mathbb{N}}$ 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 $\mathrm{Pol}(\mathcal{A})^m$.

Proof. We know from Theorem 4 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 essentially a corollary of the works of Chen and Zhuk [7, 20] via [5].

▶ **Theorem 6.** *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.*

Proof. We know from Theorem 7 in [20] that $\mathbb{A}$ is Switchable, whereupon we apply Corollary 5, $(vi) \Rightarrow (i)$. By considering instances whose universal variables involve only the polynomial number of tuples from the Switching Adversary, one can see that QCSP$(\mathrm{Inv}(\mathbb{A}))$ reduces to a polynomial number of instances of CSP$(\mathrm{Inv}(\mathbb{A}))$ and is therefore in NP. Further details of the NP algorithm are given in Corollary 38 of [5] but the argument here follows exactly Section 7 from [7], in which it was originally proved that Switchability yields the corresponding QCSP in NP. ◀

Note that Chen's original definition of Switchability, based on adversaries and reactive composability, plays a key role in the NP membership algorithm in Theorem 6. It is the result from [5] that is required to reconcile the two definitions of switchability as equivalent, and indeed Corollary 5 is needed in this process for infinite signatures. If we were to use just our definition of switchability then it is only possible to prove, à la Proposition 3.3 in [7], that the bounded alternation Π_n-CSP$(\mathrm{Inv}(\mathbb{A}))$ is in NP. Thus, using just the methods from [7] and [20], we cannot prove the Revised Chen Conjecture, but rather some bounded alternation (re)revision.

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 7.** *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 satisfiable 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 8.** *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 [20] 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 7, 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 [17] 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 $\neg\phi$ 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 $\neg\phi$ by transforming the quantifier-free part $x_1 = y_1 = z_1 \vee \ldots \vee x_k = y_k = z_k$ to $\tau_k = \rho'(x_1, y_1, z_1) \vee \ldots \vee \rho'(x_k, y_k, z_k)$.

($\neg\phi \in$ 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 $\neg\phi$ now gives the corresponding disjunct that will be satisfied in τ_k.

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

The demonstration of co-NP-hardness in the previous theorem was inspired by a similar proof in [1]. Note that an alternative proof that τ_k is in Inv($\mathbb{A}$) is furnished by the observation that it is preserved by all $\alpha\beta$-projections (see [20]). We note surprisingly that co-NP-hardness in Theorem 8 is optimal, in the sense that some (but not all!) of the cases just proved co-NP-hard are also in co-NP.

▶ Proposition 1. Let α, β strict subsets of $A := \{a_1, \dots, a_n\}$ so that $\alpha \cup \beta = A$ and $\alpha \cap \beta \neq \emptyset$. Then QCSP($A; \{\tau_k : k \in \mathbb{N}\}, a_1, \dots, a_n$) 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($A; \{\tau_k : k \in \mathbb{N}\}, a_1, \dots, a_n$). We will now seek to eliminate atoms $v = a$ ($a \in \{a_1, \dots, a_n\}$) 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 $\neg\phi'$ 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. Call $\mathcal{B}$ *existentially trivial* if 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 $x_1, \dots, x_k \in B$, whenever $(x_1, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_k) \in R^{\mathcal{B}}$ then also $(x_1, \dots, x_{i-1}, c, x_{i+1}, \dots, x_k) \in R^{\mathcal{B}}$. 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 9.** *Let α, β be strict subsets of $A := \{a_1, \dots, a_n\}$ so that $\alpha \cup \beta = A$ and $\alpha \cap \beta \neq \emptyset$. The set of relations pp-definable in $(A; \{\tau_k : k \in \mathbb{N}\}, a_1, \dots, a_n)$ is almost existentially trivial.*

Proof. Consider a formula with a pp-definition in $(A; \{\tau_k : k \in \mathbb{N}\}, a_1, \dots, a_n)$. 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. Indeed, their atoms τ_n may now be removed since they will always be satisfied. Thus we are left with a conjunction of equalities and atoms τ_n, and the result follows. ◀

▶ Proposition 2. 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 1 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 10.** *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 1 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.

▶ Proposition 3. Let α, β 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 7 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 $(3m+1)$-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 $(3m+1)$ 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 $(3m+1)$-tuples that have $3m$ instances of the same element. By the pigeonhole principle, the tuple $\bar{t}$ must appear as one of the $(3m+1)$ m-tuples in σ_i and this is clearly a contradiction.

It follows from [6] that QCSP($\mathcal{B}$) reduces to a polynomially bounded ensemble of $\binom{n}{3m} \cdot n \cdot n^{3m}$ instances CSP($\mathcal{B}$), and the result follows. ◀

3.3 The question of the tuple encoding

▶ Proposition 4. 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 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 1 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 with is $O(n^4)$. ◀

The proof of Proposition 4 suggests an alternative proof of Proposition 3, but placing the corresponding QCSP in P instead of NL. Proposition 4 shows that Chen's Conjecture fails for the tuple encoding in the sense that it provides a language $\mathcal{B}$, expanded with constants, so that Pol($\mathcal{B}$) has EGP, yet QCSP($\mathcal{B}$) is in P under the tuple 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 encoding.

4 Switchability, Collapsability and the three-element case

An algebra $\mathbb{A}$ is a *G-set* if its domain is not one-element and every of its operations f is of the form $f(x_1, \ldots, x_k) = \pi(x_i)$ where $i \in [k]$ and π is a permutation on A. An algebra $\mathbb{A}$ contains a G-set as a *factor* if some homomorphic image of a subalgebra of $\mathbb{A}$ is a G-set. A *Gap Algebra* [6] is a three-element idempotent algebra that omits a G-set as a factor and is not Collapsible.

Our first task is the deduction of the following theorem, whose lengthy proof is omitted. For each of the following two theorems, α and β are chosen such that α, β are strict subsets of $\{0,1,2\}$, $\alpha \cup \beta = \{0,1,2\}$ and $\alpha \cap \beta \neq \emptyset$.

▶ **Theorem 11.** *Suppose $\mathbb{A}$ is a Gap Algebra that is not $\alpha\beta$-projective. Then, for every finite subset of Δ of Inv($\mathbb{A}$), Pol(Δ) is Collapsible.*

Our second task is the deduction of the following theorem, whose lengthy proof is omitted.

▶ **Theorem 12.** *Suppose $\mathbb{A}$ is a 3-element idempotent algebra that is not $\alpha\beta$-projective, containing a 2-element G-set as a subalgebra. Then, $\mathbb{A}$ is Collapsible.*

▶ **Corollary 13.** *Suppose $\mathbb{A}$ is a 3-element idempotent algebra that is not EGP, i.e. is Switchable. Then, for every finite subset of Δ of Inv($\mathbb{A}$), Pol(Δ) is Collapsible.*

Proof. Recall Lemma 11 in [20] that $\mathbb{A}$ has EGP iff there exists α and β such that α, β are strict subsets of D, $\alpha \cup \beta = D$, and all operations of $\mathbb{A}$ are $\alpha\beta$-projective.

If $\mathbb{A}$ does not contain a G-set as a factor, then $\mathbb{A}$ is a Gap Algebra and the result follows from Theorem 11. Otherwise, $\mathbb{A}$ contains a G-set as a factor. If $\mathbb{A}$ contains a G-set as a homomorphic image then $\mathbb{A}$ has EGP from [7]. Else, since $\mathbb{A}$ is 3-element, $\mathbb{A}$ contains a 2-element G-set as a subalgebra and we are in the situation of Theorem 12. ◀

5 A three-element vignette

We would love to be able to improve Theorem 1 to describe the boundary between those cases that are co-NP-complete and those that are Pspace-complete, if indeed such a result is true. However, even in the three-element case this appears challenging, but we are able to provide a variant vignette, whose proof is omitted.

▶ **Theorem 14.** *Let $\mathbb{A}$ be an idempotent algebra on a 3-element domain. Either*
- Π_k-*CSP*$(\mathrm{Inv}(\mathbb{A}))$ *is in NP, for all k; or*
- Π_k-*CSP*$(\mathrm{Inv}(\mathbb{A}))$ *is co-NP-complete, for all k; or*
- Π_k-*CSP*$(\mathrm{Inv}(\mathbb{A}))$ *is* Π_2^{P}*-hard, for some k.*

Note that the trichotomy of Theorem 14 does not hold for QCSP along the same boundary for, respectively, NP, co-NP-complete and Pspace-complete. For the semilattice-without-unit s it is known that Π_k-CSP(Inv(s)) is co-NP-complete, for all k, while QCSP(Inv(s)) is Pspace-complete [3].

6 Discussion

The major contribution of this paper is its discussion of the Chen Conjecture with two infinite-signature variants one of which is proved to hold (with encoding in "simple logic") and one of which fails (with the tuple listing).

In addition to this, the contribution is largely mathematical, examining the relationship between Switchability and Collapsibility in the three-element case. However, this mathematical study uncovers something of importance to the computer scientist who is not reconciled to infinite signatures! Since here it demonstrates that all three-element domain NP-memberships that may be shown by Switchability, may already be shown by Collapsibility.

The work associated with Theorem 11 is distinctly non-trivial and involves a new method, whereas the work associated with Theorem 12 uses known methods and involves mostly turning the handle with these. Similarly, the work involved with the three element vignette uses known methods on top of our earlier new results.

The Chen Conjecture in its original form remains open. As does the general question (for arbitrary finite domains) as to whether, if $\mathbb{A}$ is Switchable, all finite subsets $\mathcal{B}$ of Inv($\mathbb{A}$) are so that Pol($\mathcal{B}$) is Collapsible. However, to now prove the Chen Conjecture it is sufficient to prove, for any finite $\mathcal{B}$ expanded with all constants such that Pol($\mathcal{B}$) has EGP, that there exists polynomially (in i) computable pp-definitions (over $\mathcal{B}$) of the relations τ_i (where α and β are suitably chosen to witness EGP). A first step towards this is to establish whether there are even polynomially sized pp-definitions of these τ_i.

The appearance of a co-NP-complete QCSP is likely to be an anomaly of our introduction of infinite signatures. Such a QCSP is unlikely to exist with a finite signature (at least, nothing like this is hitherto known). Indeed, its presence might be used as an argument against the acceptance of infinite signatures, if it is interpreted as an aberration. For the reader in this mind, we ask to please review the earlier paean to infinite signatures.

Acknowledgements

We thank Hubie Chen and Michał Wrona for many useful discussions, as well as several anonymous referees for their advising on two previous drafts.

References

1 Manuel Bodirsky and Hubie Chen. Quantified equality constraints. *SIAM J. Comput.*, 39(8):3682–3699, 2010.

2 Manuel Bodirsky and Jan Kára. The complexity of equality constraint languages. *Theory of Computing Systems*, 3(2):136–158, 2008. A conference version appeared in the proceedings of CSR'06.

3 Ferdinand Börner, Andrei A. Bulatov, Hubie Chen, Peter Jeavons, and Andrei A. Krokhin. The complexity of constraint satisfaction games and qcsp. *Inf. Comput.*, 207(9):923–944, 2009.

4 A. Bulatov, A. Krokhin, and P. G. Jeavons. Classifying the complexity of constraints using finite algebras. *SIAM Journal on Computing*, 34:720–742, 2005.

5 Catarina Carvalho, Florent R. Madelaine, and Barnaby 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.

6 Hubie Chen. The complexity of quantified constraint satisfaction: Collapsibility, sink algebras, and the three-element case. *SIAM J. Comput.*, 37(5):1674–1701, 2008. `doi:http://dx.doi.org/10.1137/060668572`.

7 Hubie Chen. Quantified constraint satisfaction and the polynomially generated powers property. *Algebra universalis*, 65(3):213–241, 2011. An extended abstract appeared in ICALP B 2008. URL: `http://dx.doi.org/10.1007/s00012-011-0125-4`, `doi:10.1007/s00012-011-0125-4`.

8 Hubie Chen. Meditations on quantified constraint satisfaction. In *Logic and Program Semantics - Essays Dedicated to Dexter Kozen on the Occasion of His 60th Birthday*, pages 35–49, 2012.

9 Hubie Chen and Peter Mayr. Quantified constraint satisfaction on monoids, 2016.

10 Nadia Creignou, Sanjeev Khanna, and Madhu Sudan. *Complexity Classifications of Boolean Constraint Satisfaction Problems*. SIAM Monographs on Discrete Mathematics and Applications 7, 2001.

11 Uwe Egly, Thomas Eiter, Hans Tompits, and Stefan 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*, pages 417–422. AAAI Press/ The MIT Press, 2000.

12 T. Feder and M. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory. *SIAM Journal on Computing*, 28:57–104, 1999.

13 Vladimir Kolmogorov, Andrei A. Krokhin, and Michal 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*, pages 1246–1258, 2015. URL: `http://dx.doi.org/10.1109/FOCS.2015.80`, `doi:10.1109/FOCS.2015.80`.

14 Florent R. Madelaine and Barnaby Martin. On the complexity of the model checking problem. *CoRR*, abs/1210.6893, 2012. Extended abstract appeared at LICS 2011 under the name "A Tetrachotomy for Positive First-Order Logic without Equality". URL: `http://arxiv.org/abs/1210.6893`.

15 Barnaby Martin. On the chen conjecture regarding the complexity of qcsps. *CoRR*, abs/1607.03819, 2016. URL: `http://arxiv.org/abs/1607.03819`.

16 Barnaby Martin and Dmitriy Zhuk. Switchability and collapsibility of gap algebras. *CoRR*, abs/1510.06298, 2015. URL: `http://arxiv.org/abs/1510.06298`.

17 Christos H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

18 R.M. Smullyan. *Godel's Incompleteness Theorems*. Oxford Logic Guides. Oxford University Press, 1992. URL: `https://books.google.co.uk/books?id=04zalcCdKZsC`.

19 James Wiegold. Growth sequences of finite semigroups. *Journal of the Australian Mathematical Society (Series A)*, 43:16–20, 8 1987. Communicated by H. Lausch. URL: `http://journals.cambridge.org/article_S1446788700028925`, `doi:10.1017/S1446788700028925`.

20 D. Zhuk. The Size of Generating Sets of Powers. *ArXiv e-prints*, April 2015. `arXiv:1504.02121`.