CSPS AND THE AXIOM OF CHOICE

RILEY THORNTON

Some of this paper is based on joint work carried out with CMU undergraduates–Anthony Li, Isin Shah, Matthew Snodgrass, and Rui Zhou –as part of a SEMS summer research project.

1. Background

For a finite relational structure $\mathcal{D}$, the associated constraint satisfaction problem, CSP($\mathcal{D}$), is the problem of testing if there is a homomorphism from a given structure, $\mathcal{X}$, to $\mathcal{D}$. Specific CSPs include graph n coloring (with $\mathcal{D} = K_n$), $kSAT$ (with $\mathcal{D}$ the structure on $\{0, 1\}$ including all k-ary relations), and systems of linear equations over $\mathbb{F}_{p^n}$ (with $\mathcal{D}$ the structure on $\mathbb{F}_{p^n}$ with all affine relations).

In analogy with the computational problem CSP($\mathcal{D}$), Katay–Toth–Vidnyanszky [3] introduced a weak compactness theorem associated to $\mathcal{D}$:

Definition 1.1. For a finite relational structure $\mathcal{D}$, $\mathcal{K}(\mathcal{D})$ is the following statement: For any relational structure $\mathcal{X}$ in the same signature as $\mathcal{D}$, if every finite substructure of $\mathcal{X}$ has a homomorphism to $\mathcal{D}$, then $\mathcal{X}$ has a homomorphism to $\mathcal{D}$.

Finite structures come with a purely combinatorial notion of reduction called a pp-definition. If $\mathcal{D}$ pp-defines $\mathcal{E}$, then CSP($\mathcal{E}$) polytime reduces to CSP($\mathcal{D}$), and $\mathcal{K}(\mathcal{D})$ implies $\mathcal{K}(\mathcal{E})$ over ZF.

Definition 1.2. A structure $\mathcal{E}$ **pp-defines** a relation R if there are formulas α_i that assert either an equality or a relation from $\mathcal{E}$ so that

$$R(\vec{x}) \ :\Leftrightarrow \ (\exists \vec{z}) \ \bigwedge_{i=1}^{n} \alpha_i(\vec{x}, \vec{z})$$

We say a structure $\mathcal{D}$ pp-defines a structure $\mathcal{E}$ on the same domain if $\mathcal{D}$ pp-defines every relation in $\mathcal{E}$. In this case, we write $\mathcal{E} \leq_{pp} \mathcal{D}$.

There are various ways to relax pp-definition and extend it structures on different domains. For sake of time and simplicity, we will omit the full story. The relation $\leq_{pp}$ is dual to an algebraic relation.

Definition 1.3. A polymorphism of $\mathcal{D}$ is a homomorphism from $\mathcal{D}^n$ to $\mathcal{D}$. We write $Pol(\mathcal{D})$ for the algebra of polymorphisms of $\mathcal{D}$ (where we view each polymorphism as an operation).

Theorem 1.4. *For finite relational structures $\mathcal{D}$ and $\mathcal{E}$ on the same domain, $\mathcal{E} \leq_{pp} \mathcal{D}$ iff $Pol(\mathcal{D}) \subseteq Pol(\mathcal{E})$.*

If $\mathcal{E} \leq_{pp} \mathcal{D}$, then CSP($\mathcal{E}$) polynomial time reduces to CSP($\mathcal{D}$). So, the theorem above implies there must be some algebraic description of the NP complete CSPs. Bulatov and Zhuk gave such a characterization (resolving a conjecture of Bulatov–Jeavons–Krohkin and Feder–Vardi).

Theorem 1.5 ([1][7]; $P \neq NP$)**.** *For a finite relational structure $\mathcal{D}$,* CSP($\mathcal{D}$) *is in P iff $\mathcal{D}$ has a polymorphism f satisfying*

$$(\forall a, e, r) \;\; f(r, a, r, e) = f(a, r, e, a).$$

And, CSP($\mathcal{D}$) *is NP-complete otherwise.*

Analogously, pp-interpretations give reductions between compactness principles.

Proposition 1.6. *If $\mathcal{E} \leq_{pp} \mathcal{D}$, then*

$$ZF \vdash \mathcal{K}(\mathcal{D}) \Rightarrow \mathcal{K}(\mathcal{E}).$$

Katay–Toth–Vidnyanszky gave a characterization of the compactness principals of maximal strength.

Theorem 1.7 ([3]; $Con(ZF)$)**.** $\big(ZF \vdash \mathcal{K}(\mathcal{D}) \Leftrightarrow BPI\big)$ *if and only if $\mathcal{D}$ has no polymorphism f satisfying*

$$(\forall a, e, r) \;\; f(r, a, r, e) = f(a, r, e, a).$$

Corollary 1.8 ($P \neq NP, Con(ZF)$)**.** CSP($\mathcal{D}$) *is NP-complete if and only if $ZF \vdash \mathcal{K}(\mathcal{D}) \Leftrightarrow BPI$.*

We build on this work in two ways. First, we characterize the compactness principles at the opposite end of the spectrum, namely the provable $\mathcal{K}(\mathcal{D})$. Second, we make progress towards classifying $\mathcal{K}(\mathcal{D})$ for $\mathcal{D}$ a structure on 2 elements. In particular, we answer a question of KTV and show that $\mathcal{K}(K_2)$ is strictly weaker than $\mathcal{K}(2SAT)$ or $\mathcal{K}(\mathbb{F}_2)$.

2. Provable compactness principles

The structures whose associated compactness principle is provable from ZF turn out to coincide with a well-studied class from computer science.

Definition 2.1. A structure is **width 1** if there is a homomorphism $f : P(\mathcal{D}) \to \mathcal{D}$ where $P(D)$ is the structure with domain $\{A \subseteq \mathcal{D} : A \neq \emptyset\}$ a relations defined by

$$R^{\mathcal{P}(\mathcal{D})}(\vec{A}) \;:\Leftrightarrow\; (\forall a \in A_i)(\exists \vec{e} \in R^{\mathcal{D}}) \; e_i = a \text{ and } \vec{e} \in \prod_i A_i.$$

For example K_2 is not width 1. An example of a structure with width 1 is $HornSAT$, defined as follows: The domain of $HornSAT$ is $\{0, 1\}$ and it has all singleton unary relations as well as all relations of the form

$$(x_1 \wedge x_2 \wedge ... \wedge x_n) \to y.$$

There are many characterizations of width-1 CSPs. We briefly sketch one characterization due to the corresponding author.

Proposition 2.2. [6] *For $\mathcal{D}$ not width 1, there are instances of $\mathcal{Y}$ and $\mathcal{X}$ of $\mathcal{D}$ and $x_0 \in \mathcal{X}$ so that the following hold: (1) any acyclic structure with a homomorphism to $\mathcal{X}$ has a homomorphism to $\mathcal{D}$, (2) $\mathcal{Y}$ is acyclic and has a homomorphism f to $\mathcal{X}$, and (3) there is no homomorphism from $\mathcal{Y}$ to $\mathcal{D}$ that is constant on $f^{-1}[x_0]$.*

Now, if $\mathbb{Y}_0$ is the profinite limit of a generic sequence of acylic structures that map to $\mathcal{Y}$, any Baire measurable map from $\mathbb{X}_0$ to $\mathcal{D}$ will have to be constant on a copy of $f^{-1}[x_0]$, so will not be a homomorphism. Get the following:

Theorem 2.3 ([6, Theorem 5.7]; DC)**.** *If $\mathcal{D}$ is not width 1, then $\mathcal{D}$ has a Borel instance $\mathcal{X}$ so that every finite substructure of $\mathcal{X}$ has a solution, but $\mathcal{X}$ has no Baire measurable solution.*

The converse is true as well, though we don't need it here. As a corollary

Theorem 2.4 ($Con(ZF)$)**.** *If $\mathcal{D}$ is not width 1, ZF does not prove $\mathcal{K}(\mathcal{D})$*

Proof. Shelah has shown that $ZF + DC +$"all sets are Baire measurable" is consistent relative to ZF [5]. □

For the converse of the corollary,

Proposition 2.5. *If $\mathcal{D}$ is width 1, then ZF proves $\mathcal{K}(\mathcal{D})$.*

proof sketch. It suffices to show (in ZF) that if every substructure of $\mathcal{X}$ has a homomorphism to $\mathcal{D}$, then $\mathcal{X}$ has a homomorphism to $P(\mathcal{D})$. One can be constructed by induction as follows: Let $f_0(x) = \mathcal{D}$. Call an element a of $f_i(x)$ good for x if, for any constraint $\vec{e} \in R^{\mathcal{X}}$ with $e_i = x$ there is $\vec{a} \in R^{\mathcal{D}} \cap \prod_j f(e_j)$ with $a_i = a$. Set $f_{i+1}(x) = \{a \in f_i(x) : a \text{ is good for } x\}$.

If some $f_i(x)$ is empty, then this is witnessed by a finite substructure of X. But this contradicts our assumption about finite substructures of x. Thus, $f(x) = \bigcap_i f_i(x)$ defines a homomorphism to $\mathcal{D}$. □

3. Boolean structures

We call a structure on $\{0, 1\}$ Boolean. In 1941, Post [4] classified the possible polymorphism algebras of Boolean structures. See Figure 1.

Let maj be the majority operator ($\text{maj}(x, y, z)$ is the repeated value among x, y, z), and let $\oplus(x, y, z) = x + y + z \mod 2$. The only algebras pictured here that will be relevant for us are

- $\top = \bigcup_i 2^{2^i} = Pol(\emptyset)$
- $D = \langle \text{maj}, \neg \rangle = Pol(K_2)$, $DP = \langle \text{maj}, \oplus \rangle = Pol(K_2 \cup \{\{0\}\})$
- $DM = \langle \text{maj} \rangle = Pol(2SAT)$
- $AP = \langle \oplus \rangle = Pol(\mathbb{F}_2)$, $AD = \langle \oplus, \neg \rangle$
- $\bot = \langle \emptyset \rangle = pol(3SAT)$, $UD = \langle \neg \rangle$.

Every other algebra (highlighted in blue in Figure 1) contains either $\vee, \wedge$, or a constant function, and so is width 1. The picture simplifies to Figure 2.

Note that $\langle \text{maj}, \neg \rangle = \langle \text{maj}, \oplus, \neg \rangle$ corresponds to K_2, and $\langle \text{maj}, \oplus \rangle$ corresponds to K_2 with a singleton unary predicate. One can check that for the structures we're considering, adding a unary predicate doesn't change the strength of $\mathcal{K}(\mathcal{D})$. Algebraically, this means unary automorphisms $Pol(\mathcal{D})$ don't change complexity (this is generally the case for so-called cores). So, we are left with 5 compactness principles pictured at the right in Figure 2 (note that the image has flipped since algebra containment is dual to reductions).

By Theorem 1.7, $BPI \equiv \mathcal{K}(3SAT)$ is strictly stronger than the compactness principles for $HornSAT$, $2SAT$, $\mathbb{F}_2$, and K_2. It is well-known (see [2]) that $\mathcal{K}(K_2)$ is equivalent to the axiom of choice for pairs and is not provable from ZF. KTV asked whether $\mathcal{K}(K_2)$ was equivalent to either of the two remaining compactness principles. We show that this is not the case.

Theorem 3.1 (Con(ZF))**.** *$\mathcal{K}(K_2)$ does not prove $\mathcal{K}(2SAT)$ or $\mathcal{K}(\mathbb{F}_2)$.*

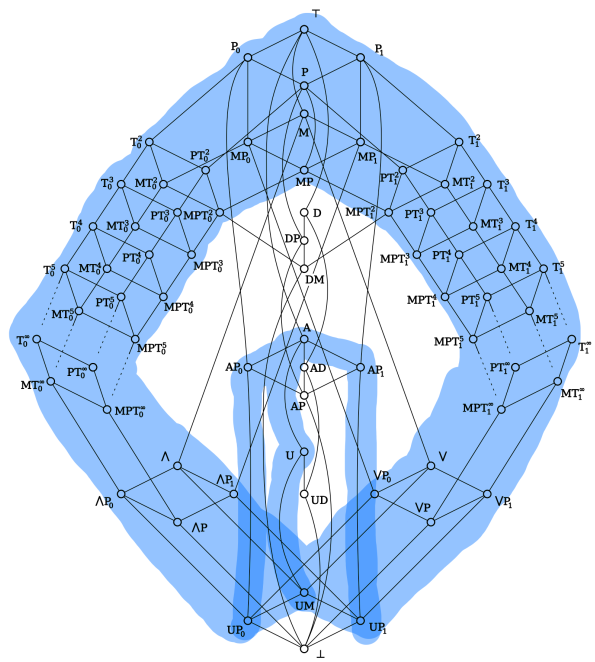

FIGURE 1. Post's lattice of algebras ordered by $\subseteq$ with the polymorphism algebras of width 1 structures highlighted.

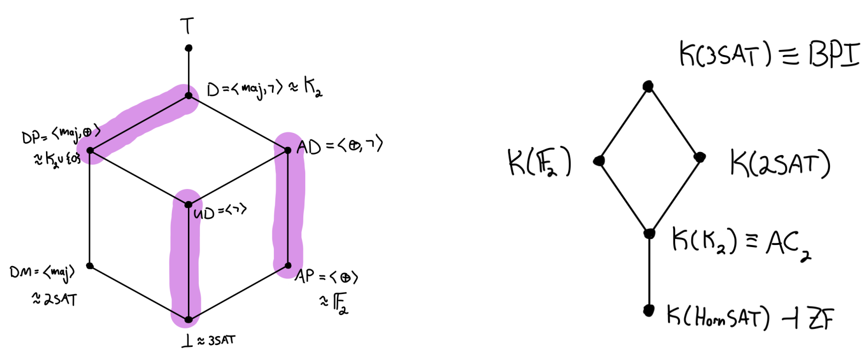

FIGURE 2. The non-width-1 algebras and the corresponding compactness principles

Proof sketch. The same model works to separate both $\mathcal{K}(\mathbb{F}_2)$ and $\mathcal{K}(2SAT)$ from $\mathcal{K}(K_2)$. Let $(\mathbb{B}, \leq)$ be the Fraïssé limit of the class of pairs $(\mathcal{B}, \prec)$ where $\mathcal{B}$ is a finite Boolean algebra and $\prec$ is a linear order with $0 \prec 1$. Let $\mathbb{M}$ be the corresponding symmetric model.

This structure admits weak elimination of imaginaries (this is not trivial). So, $\mathbb{M} \vDash (\forall A)(\exists \alpha \in On)\ |A| \leq [\mathbb{B}]^{<\omega} \times \alpha$. In particular, every set can be linearly ordered in $\mathbb{M}$. So, $\mathbb{M} \vDash \mathcal{K}(K_2)$.

Consider the instance of $2SAT$ in $\mathbb{M}$ defined by $\mathcal{X} = (\mathbb{B}, \to^{\mathcal{X}}, \neq^{\mathcal{X}})$ where $a \to^{\mathcal{X}} b$ if and only if $a \cup b = a$ and $a \neq^{\mathcal{X}} b$ if and only if $a = b^c$. One can check that this has no symmetric solution, though every finite subinstance a solution. Similarly, the instance of $\mathbb{F}_2$, $\mathcal{Y}$ with domain $\mathbb{B}$ given by

$$(x + y + z = 1)^{\mathcal{Y}} :\Leftrightarrow x \triangle y \triangle z = 1$$

shows that $\mathcal{K}(\mathbb{F}_2)$ is false in $\mathbb{M}$. □

All that remains for Boolean structures is to understand the relationship between $\mathcal{K}(\mathbb{F}_2)$ and $\mathcal{K}(2SAT)$.

References

[1] Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 319–330, 2017.

[2] Thomas J Jech. *The axiom of choice*. Courier Corporation, 2008.

[3] Tamás Katay, Laszlo Marton Toth, and Zoltan Vidnyanszky. The csp dichotomy, the axiom of choice, and cyclic polymorphisms, 2023.

[4] Emil Leon Post. *The Two-Valued Iterative Systems of Mathematical Logic*. London: Oxford University Press, 1941.

[5] Saharon Shelah. Can you take solovay's inaccessible away? *Israel Journal of mathematics*, 48:1–47, 1984.

[6] Riley Thornton. An algebraic approach to Borel CSPs. arxiv: 2203.16712, 2024.

[7] Dmitriy Zhuk. A proof of CSP dichotomy conjecture. *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 331–342, 2017.