The Constraint Satisfaction Dichotomy Theorem for Beginners

Tutorial – Part 3

Ross Willard

University of Waterloo

BLAST 2019
CU Boulder, May 24, 2019

Recap:

Fix $\mathbb{A}$ (a finite Taylor algebra)

A <u>CSP instance compatible with $\mathbb{A}$</u> consists of

- a family $(\mathbb{A}_{x_i})$ of subalgebras of $\mathbb{A}$ (indexed by variables $x_i \in X$), and
- a set $\{C_t : 1 \leq t \leq m\}$ of "constraints," each of the form $R_t(x_{i_1}, \ldots, x_{i_k})$ where

$$R_t \leq_{sd} \mathbb{A}_{x_{i_1}} \times \cdots \times \mathbb{A}_{x_{i_k}}.$$

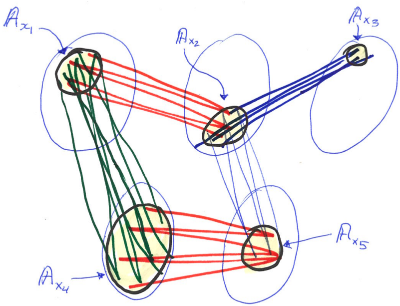

A <u>solution</u> to the instance is a function $\sigma : X \to A$ satisfying $\sigma(x_i) \in A_{x_i}$ for all $x_i \in X$ and $(\sigma(x_{i_1}), \ldots, \sigma(x_{i_k})) \in R_t$ for all $R_t(x_{i_1}, \ldots, x_{i_k})$.

Linear equations over $\mathbb{Z}_p$ (for various p) can be encoded in individual constraints in the right circumstances.

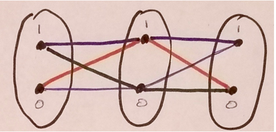

$\mathbb{Z}_2$

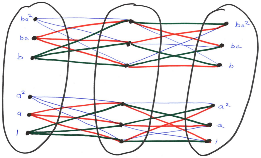

$\mathbb{S}_3$

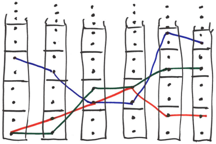

$\mathbb{SL}(2,5)$

Potatoes need not all be the same ...

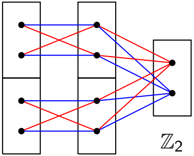

$\mathbb{Z}_4$ $\mathbb{Z}_4$ $\mathbb{Z}_2$

nor subdirectly irreducible.

Con $\mathbb{A}_x$ =

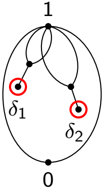

Weds. TCT Theorem (fork-free version, improved)

Suppose $\mathbb{A}_1, \ldots, \mathbb{A}_n$ are finite algebras in an (~~idempotent~~) Taylor variety with $n \geq 3$. Assume $R \leq_{sd} \mathbb{A}_1 \times \cdots \times \mathbb{A}_n$ and also

- R is critical (indecomposable and meet-irreducible), and
- R is fork-free,

and let R^* be its unique upper cover. Then

1. Each $\mathbb{A}_i$ is SI (subdirectly irreducible) with abelian monolith μ_i.
2. R^* is the $\mu_1 \times \cdots \times \mu_n$-closure of R.
3. $\exists$ prime p such that each μ_i-class ($\forall i$) can naturally be identified with a vector space over $\mathbb{Z}_p$.
4. With respect to these identifications, the restriction of R to any strand in R^* is defined by linear equations over $\mathbb{Z}_p$.
5. If $(0 : \mu_i) = 1$ for some (all) i, then there exists a simple affine algebra $\mathbb{M}$ such that (unimportant).

Contrary to Miklos's advice, we allow potatoes to be not SI.

$\mathrm{Con}\,\mathbb{A}_x =$ 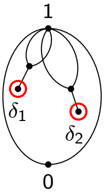

So we need a relativized version of the previous theorem.

Weds. TCT Theorem (relativized, with improvements)

Suppose $\mathbb{A}_1, \ldots, \mathbb{A}_n$ are finite algebras in a (not necess. idempotent) Taylor variety. Assume $R \leq \mathbb{A}_1 \times \cdots \times \mathbb{A}_n$ is subdirect, critical, and rectangular.[1]

- I.e., $\exists\, \delta_i^R \in \operatorname{Con} \mathbb{A}_i$ such that R is the pullback of some fork-free $\overline{R} \leq \mathbb{A}_1/\delta_1^R \times \cdots \times \mathbb{A}_n/\delta_n^R$.

Then:

1. All the δ_i^R are meet-irreducible (say with covers μ_i^R).
2. If $n \geq 3$ then
   - Each μ_i^R is abelian modulo δ_i^R.
   - $\exists$ prime p such that the restriction of R to any strand (product of μ_i^R classes in R^*) encodes (modulo the δ_i^R's) linear equation(s) over $\mathbb{Z}_p$.
   - If $(\delta_i^R : \mu_i^R) = 1$ for some (all) i, then (unimportant).
3. If $n = 2$ then R is the graph of an isomorphism modulo the δ_i^R's.

[1] Zhuk's term. Kearnes and Szendrei use "the $(1,n-1)$ parallelogram property."

Similarity

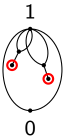

On Wednesday I alluded to a notion of similarity.

Definition

Suppose $\mathbb{A}_1, \mathbb{A}_2$ are finite SI algebras with abelian monoliths in a Taylor (not necessarily idempotent) variety.

We say that $\mathbb{A}_1$ is similar to $\mathbb{A}_2$ if $\exists$ finite $\mathbb{A}_3 \in \mathrm{HSP}(\mathbb{A}_1, \mathbb{A}_2)$ and $R \leq \mathbb{A}_1 \times \mathbb{A}_2 \times \mathbb{A}_3$ such that R is subdirect, critical, and fork-free.

Intuition: $\mathbb{A}_1$ and $\mathbb{A}_2$ are similar if their monolith classes can jointly participate in linear equations.

Examples:

- $\mathbb{Z}_4$ and $\mathbb{Z}_2$ are similar.
- $\mathbb{S}_3$ and $\mathbb{Z}_3$ are not similar.

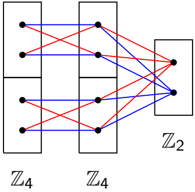

Summary:

A constraint $R(x_1, \ldots, x_n)$ (compatible with Taylor $\mathbb{A}$) encodes linear equations on strands when R is subdirect, critical and rectangular.

In this case R associates to potato $\mathbb{A}_{x_i}$ a meet-irreducible congruence $\delta^R_{x_i}$.

Two such constraints $R(\mathbf{x})$ and $S(\mathbf{y})$ can have their linear equations shared in a common system $\iff$ their potatoes modulo their δ's are similar SIs.

Plan

Goal: to explain the statement of the following technical theorem of Zhuk.

Theorem

Suppose Θ is a CSP instance compatible with the Taylor algebra $\mathbb{A}$.

If Θ morally should have solutions but doesn't, then this is sorta explained by annihilator=1 linear equations.

Notation, projections, relaxations

In the context of CSP instances:

- All tuples $\mathbf{x}$ of variables are assumed to have pairwise distinct entries.
- If $\mathbf{x} = (x_1, \ldots, x_n)$ and $\mathbf{y} = (y_1, \ldots, y_m)$ with $|\mathbf{y}| = m$, then $\mathbf{y} \subseteq \mathbf{x}$ means $\{y_1, \ldots, y_m\} \subseteq \{x_1, \ldots, x_n\}$. (Thus $|\mathbf{y}| \leq |\mathbf{x}|$.)
- If R is an n-ary relation, $|\mathbf{x}| = n$, and $\mathbf{y} \subseteq \mathbf{x}$, then

$$\begin{aligned} \mathrm{proj}^{\mathbf{x}}_{\mathbf{y}}(R) &= \{(a_{y_1}, \ldots, a_{y_m}) : (a_{x_1}, \ldots, a_{x_n}) \in R\} \\ \mathrm{proj}_{\mathbf{y}}(R(\mathbf{x})) &= S(\mathbf{y}) \quad \text{where } S = \mathrm{proj}^{\mathbf{y}}_{\mathbf{x}}(R). \end{aligned}$$

- If $R(\mathbf{x})$ and $S(\mathbf{y})$ are constraints, then we write

$$R(\mathbf{x}) \models S(\mathbf{y})$$

and say $S(\mathbf{y})$ is a relaxation of $R(\mathbf{x})$ if $\mathbf{y} \subseteq \mathbf{x}$ and $S \supseteq \mathrm{proj}^{\mathbf{x}}_{\mathbf{y}}(R)$.

Closure and subinstance

Suppose Θ is a CSP instance compatible with $\mathbb{A}$.

Definition.
The closure of Θ, written $\overline{\Theta}$, is the CSP instance with the same variable set and same potatoes as Θ, and whose constraints are all relaxations of constraints in Θ which are compatible with $\mathbb{A}$.

(Note that Θ and $\overline{\Theta}$ have the same solutions.)

Definition.
Suppose S is a nonempty subset of the set of constraints of Θ. The subinstance of Θ determined by S is the CSP instance Σ whose variables are the variables occurring in S, potatoes at those variables are the same as in Θ, and constraints are those in S.

Notation: $\Sigma \subseteq \Theta$.

Fragmented

Suppose Θ is a CSP instance compatible with $\mathbb{A}$.

Let Var(Θ) be the set of variables occurring in the constraints of Θ.

Definition.
Θ is fragmented if there exists a partition $\mathrm{Var}(\Theta) = X_1 \cup X_2$ such that no constraint mentions both a variable from X_1 and a variable from X_2.

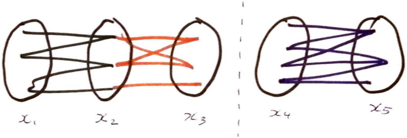

Linked relations

Assume Θ is not fragmented.

Definition.
For each $x \in \mathrm{Var}(\Theta)$, the linked relation at x is the equivalence relation $\equiv_x$ on A_x consisting of those pairs $(a, b) \in (A_x)^2$ which are in the same connected component of the potato diagram of Θ.

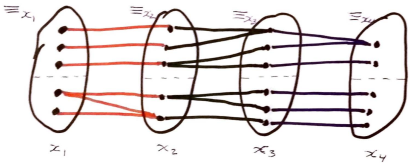

Note: if any one of the $\equiv_x$ equals $(A_x)^2$, then all are (by non-fragmented).

Definition. Θ is linked if every $\equiv_x$ is $(A_x)^2$.

Linked components

Assume Θ is compatible with $\mathbb{A}$ and is not fragmented.

Fact: if Θ is cycle-consistent, then:

1. Each $\equiv_x$ is a congruence of $\mathbb{A}_x$.
2. Hence each $\equiv_x$-block is a subuniverse of $\mathbb{A}_x$ (by idempotency).
3. Hence if no $\equiv_x$ is $(A_x)^2$, then Θ decomposes into at most $|A|$ disjoint CSP instances compatible with $\mathbb{A}$ (and with strictly smaller potatoes).

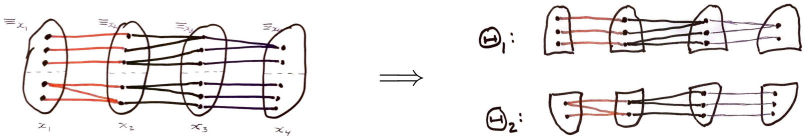

These smaller CSP instances are called the linked components of Θ.

Full subconsistency

Assume Θ is compatible with $\mathbb{A}$.

Definition.
Θ is fully consistent if for every $x \in \mathrm{Var}(\Theta)$ and every $a \in A_x$, Θ has a solution σ satisfying $\sigma(x) = a$.

Definition.
Θ is fully subconsistent if

1. Θ is cycle-consistent, and
2. For every subinstance $\Sigma \subseteq \overline{\Theta}$, if Σ is not fragmented and not linked, then each linked component of Σ is fully consistent.

(Zhuk says "irreducible.")

Claim: In solving CSP instances compatible with $\mathbb{A}$, we only need to consider fully subconsistent instances Θ.

"Proof." Suppose Θ fails (2). So for some linked component Λ of some such $\Sigma \subseteq \overline{\Theta}$ there exists an $x \in \text{Var}(\Lambda)$ and $a \in A_x$ such that

$$\boxed{\text{No solution of } \Lambda \text{ passes through } a \text{ at } x.}$$

1. Then the same is true of Θ.
2. We assume we have a CSP algorithm which can recursively solve CSP instances compatible with $\mathbb{A}$ having strictly smaller potatoes than Θ.
3. Applied to Λ, this algorithm can be used detect this defect at a, x.
4. Thus we can know to remove a from A_x in a preprocessing stage. □

Relaxed coverings

Fix a CSP instance Θ compatible with $\mathbb{A}$.

A <u>relaxed covering</u> of Θ is another CSP instance Ω compatible with $\mathbb{A}$, together with a map $\pi : \mathrm{Var}(\Omega) \to \mathrm{Var}(\Theta)$, satisfying:

1. $\forall y \in \mathrm{Var}(\Omega)$, $\mathbb{A}_y^{\Omega} = \mathbb{A}_{\pi(y)}^{\Theta}$.
2. $\forall$ constraint $S(\mathbf{y})$ in Ω, either
   - $\pi|_{\mathbf{y}}$ is injective and $S(\pi(\mathbf{y}))$ is a relaxation of a constraint in Θ, or
   - $\mathbf{y} = (y_1, y_2)$ and $\pi(y_1) = \pi(y_2) = x$, say, and S is a reflexive subuniverse of $\mathbb{A}_x \times \mathbb{A}_x$.

(See picture)

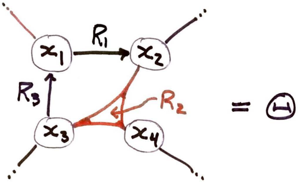

A relaxed covering of Θ

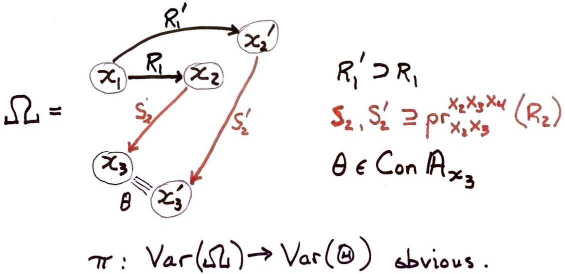

Point: every solution σ to Θ automatically gives the solution $\sigma \circ \pi$ to Ω. So if Ω has no solutions, neither does Θ.

Theorem 9.8 (Zhuk), weak version

Let $\mathbb{A}$ be a finite Taylor algebra, and Θ a CSP instance compatible with $\mathbb{A}$ which is fully subconsistent and yet is inconsistent (has no solutions).

Then $\exists$ a relaxed covering Ω of Θ and a subinstance $\Sigma \subseteq \Omega$ such that

1. Every constraint relation in Ω is critical and rectangular.
2. Neither Ω nor Σ is fragmented; both are linked.
3. Σ encodes annihilator=1 linear equations, meaning:
    1. $\forall\ S(\mathbf{y}) \in \Sigma$, $\forall\ y_i \in \mathbf{y}$, the δ-congruence $\delta^S_{y_i}$ and its unique upper cover $\mu^S_{y_i}$ satisfy $(\delta^S_{y_i} : \mu^S_{y_i}) = 1$. In particular, $\mathbb{A}_{y_i}/\delta^S_{y_i}$ has abelian monolith.
    2. All SI quotients $\mathbb{A}_{y_i}/\delta^S_{y_i}$ arising from the δ-congruences of constraints $S(\mathbf{y}) \in \Sigma$ are similar.
4. Ω is inconsistent. Σ is not fully consistent.
5. If any single constraint $S(\mathbf{y}) \in \Omega$ is replaced by $S^*(\mathbf{y})$ (where S^* is the unique upper cover of S), then the resulting Ω' is consistent.

Masterpiece

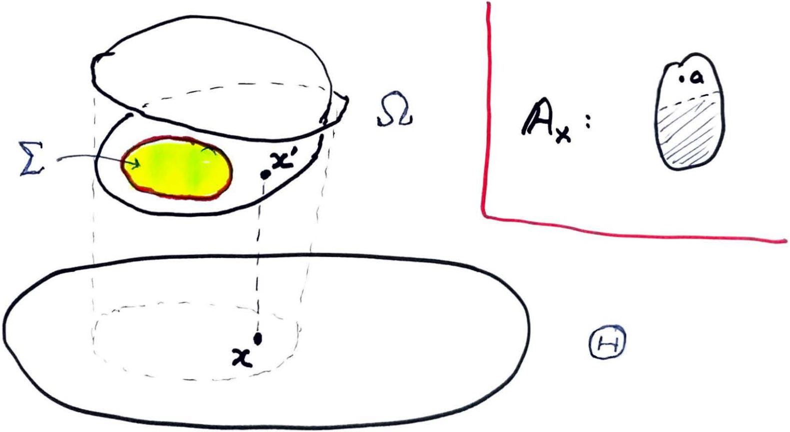

Concluding remarks

1. Zhuk actually proves a strengthened version in which, for particular families $(B_x : x \in \mathrm{Var}(\Theta))$ of subuniverses of the potatoes, the hypothesis that Θ is inconsistent is weakened to " Θ has no solution in the B_x's," and in the conclusion, the references to Ω or Ω' being (in)consistent are changed to their (not) having a solution in the B_y's (obtained from the B_x's via the covering map).
   - This strengthened version is THE key result in Zhuk's proof of the CSP Dichotomy Theorem.
2. Zhuk only proves this theorem for Taylor algebras $\mathbb{A}$ having a single basic operation which is a "special weak near-unanimity" operation.
   - I claim that the same proof (using TCT for the centralizer facts) works for any Taylor algebra.
3. This theorem (for arbitrary finite Taylor algebras) should be good for more than just the CSP Dichotomy Theorem!!!

Thank you!