Complexity of Non-Uniform CSP

Andrei A. Bulatov
Simon Fraser University

NSAC, 2017

Constraint Satisfaction Problem

Definition: **CSP($\mathcal{A}$)**

Instance: $(V;\mathcal{A};\mathcal{C})$ where

- V is a finite set of variables
- $\mathcal{A}$ is a finite set of similar finite algebras
- $\mathcal{C}$ is a set of constraints $\{R_1(s_1),\dots,R_q(s_q)\}$ where each R_i is a subalgebra of a direct product of algebras from $\mathcal{A}$

Question: whether there is $h\colon V \to \bigcup \mathcal{A}$ such that, for any i, $R_i(h(s_i))$ is true

Constraint Satisfaction Problem

Definition:

Instance: $(V;\mathcal{A};\mathcal{C})$ where

- V is a finite set of variables
- $\mathcal{A}$ is a set of finite domains
- $\mathcal{C}$ is a set of constraints $\{R_1(s_1),\ldots,R_q(s_q)\}$ where each

  R_i is a relation over a Cartesian product of sets from $\mathcal{A}$

Question: whether there is $h\colon V \to \bigcup\mathcal{A}$ such that, for any i,

$R_i(h(s_i))$ is true

CSP and Friends

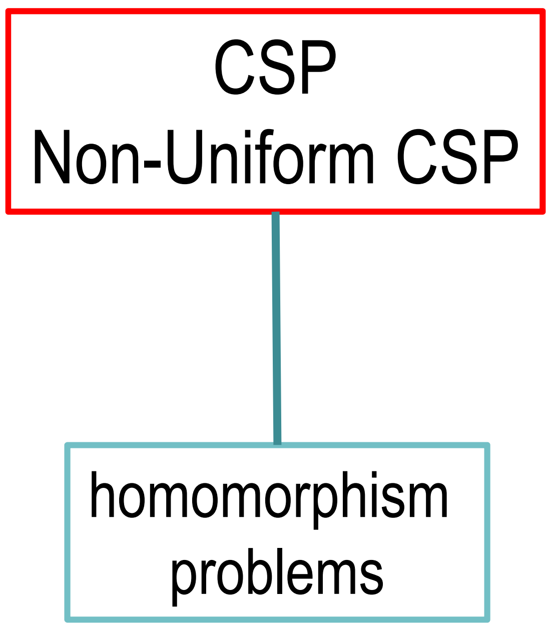

Homomorphism Problems

Homomorphism Problem:

Given relational structures G and H of the same type, decide, whether or not $G \to H$

Equivalent to CSP:

- G: elements are variables, tuples are constraint scopes
- H: elements are elements, relations are (constraint) relations

H-Coloring: (H is a fixed structure)

Given G, decide whether $G \to H$

Example: Graph Homomorphism, H-Coloring

k-Coloring:

Instance: A graph G.

Objective: Is there a k-coloring of G?

Is there a homomorphism from G to K_k?

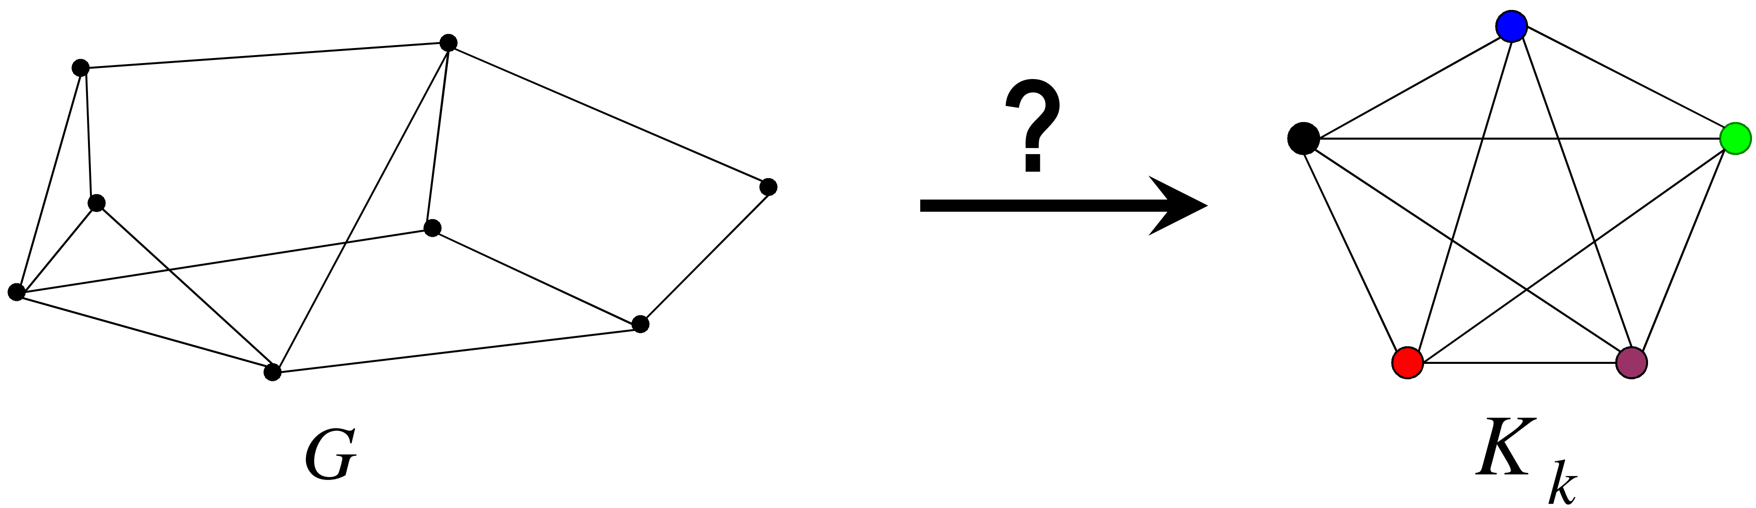

Homomorphism Problems II

Instead of fixing H, restrict possible G
Example: Problems on planar graphs

Vardi:
- Query complexity: fix H
- Data complexity: restrict G

CSP and Friends

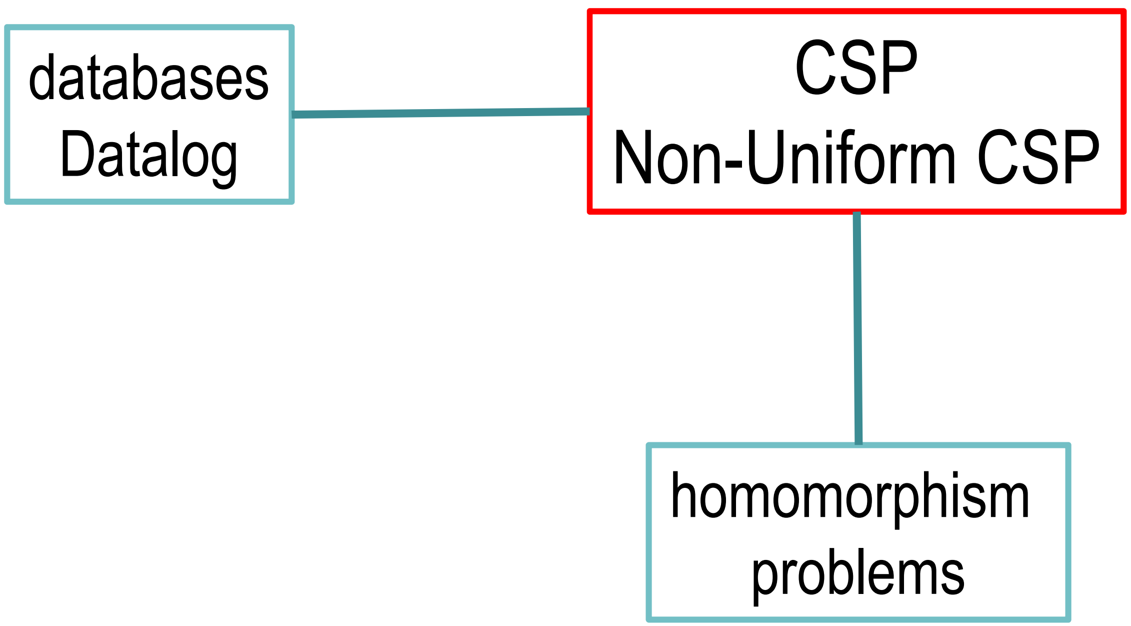

Databases

(Relational) Database:
A bunch of relations

Query:
A logic formula Φ. Enumerate all models of Φ in the database

Conjunctive query:

$$R_1(x, y) \wedge R_2(z, x, x) \wedge \ldots$$

Conjunctive queries = (enumeration) CSP

Databases: Query Containment and Equivalence

Conjunctive query is a homomorphism problem

$\Phi \rightarrow$ B

How about C.Q. Φ_1, Φ_2?

We say Φ_1 is contained in Φ_2 $(\Phi_1 \leq \Phi_2)$ if every answer to Φ_1 is an answer to Φ_2

Queries Φ_1, Φ_2 are equivalent if $\Phi_1 \leq \Phi_2$ and $\Phi_2 \leq \Phi_1$

Chandra-Merlin:

$\Phi_1 \leq \Phi_2$ iff $\Phi_1 \rightarrow \Phi_2$

Datalog

Datalog is `logic language' simulating the `least fixed point' operator

P(x,y) :- E(x,y)
P(x,y) :- P(x,z), E(z,t), E(t,y)
R(x) :- P(x,x)

Datalog gives CSPs solvable by local propagation algorithms

Barto-Kozik: For non-uniform CSPs being solvable by Datalog is equivalent to a nice algebraic condition

CSP and Friends

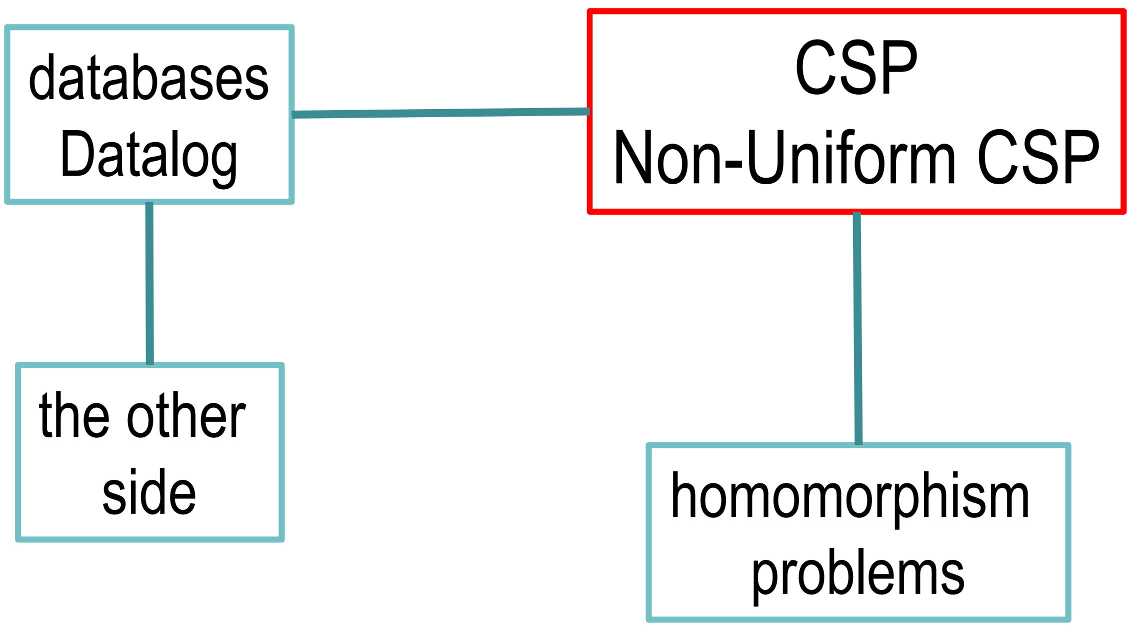

The Other Side

Let $\mathcal{G}$ be a class of structures

CSP($\mathcal{G}$,*):

Given $G \in \mathcal{G}$ and any H, decide whether $G \to H$

Grohe: For a class $\mathcal{G}$ of structures of bounded arity
CSP($\mathcal{G}$,*) is poly time iff the cores of structures from $\mathcal{G}$ have bounded treewidth (mod some complexity assumptions)

This condition can also be expressed through some logic games, homomorphism duality, etc.

The Other Side II

Marx: For a class *G* of structures **CSP(*G*,*)**

- is poly time if G has bounded fractional hypertree width
- is `fixed parameter tractable' if G has bounded submodular width
- `very hard' otherwise (mod some complexity assumptions)

CSP and Friends

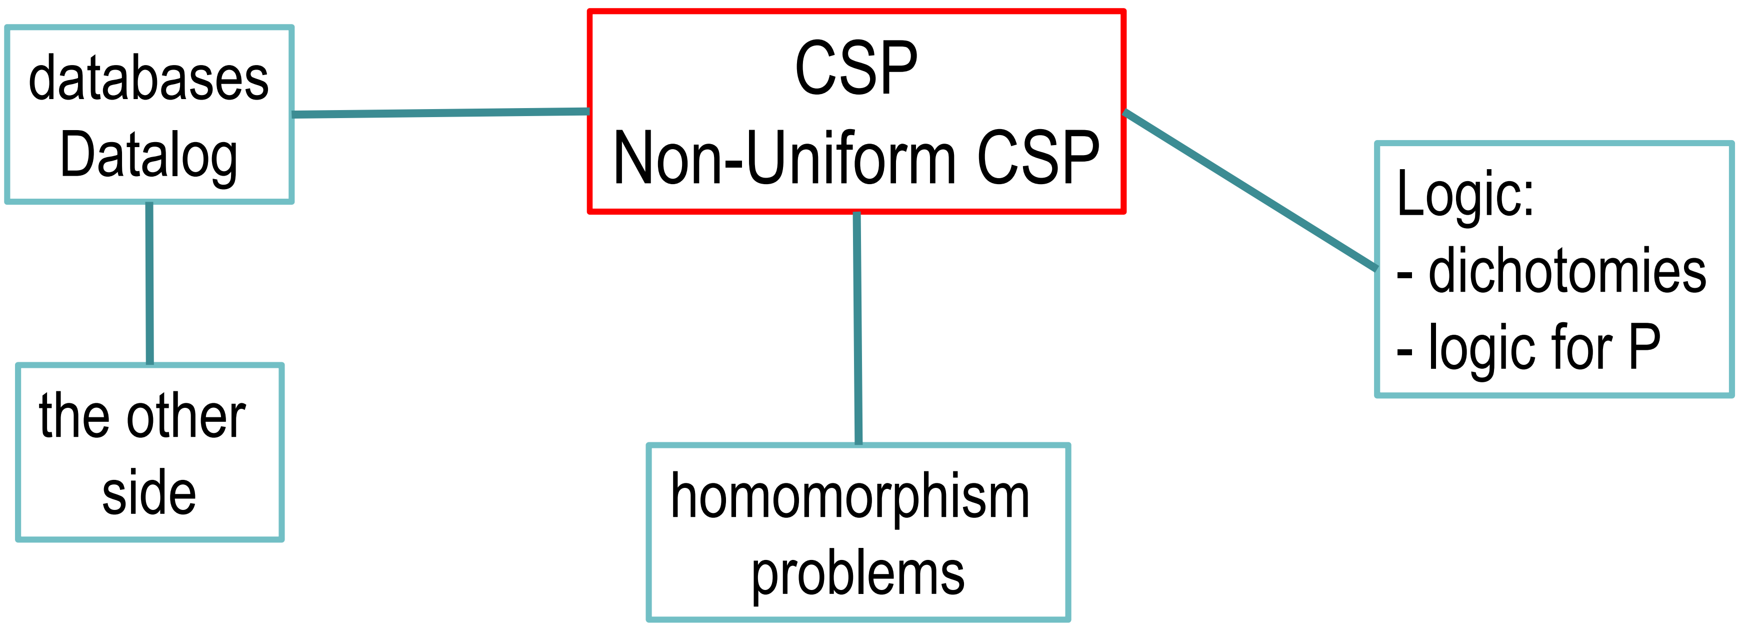

CSP vs. NP

Fagin: NP is the class of problems expressible in the existential second order logic (ESO)

If $P \neq NP$ there are infinitely many intermediate complexity classes (no dichotomy)

How much do we need to restrict NP to have a dichotomy?

Valiant, Cai: for counting problems

Marx: combinatorial conditions

Feder/Vardi: MMSNP

Feder/Vardi, Kun:

MMSNP is poly time equivalent to CSP

Logic for P

No Fagin’s theorem for P

FO is very weak

LFP(FO) (think Datalog)

Gurevich: expresses all in P provided structures are ordered

otherwise does not work for linear algebra

LFP(FO)+counting quantifiers

Still does not express matrix rank

LFP(FO)+counting+rank operator

CSP and Friends

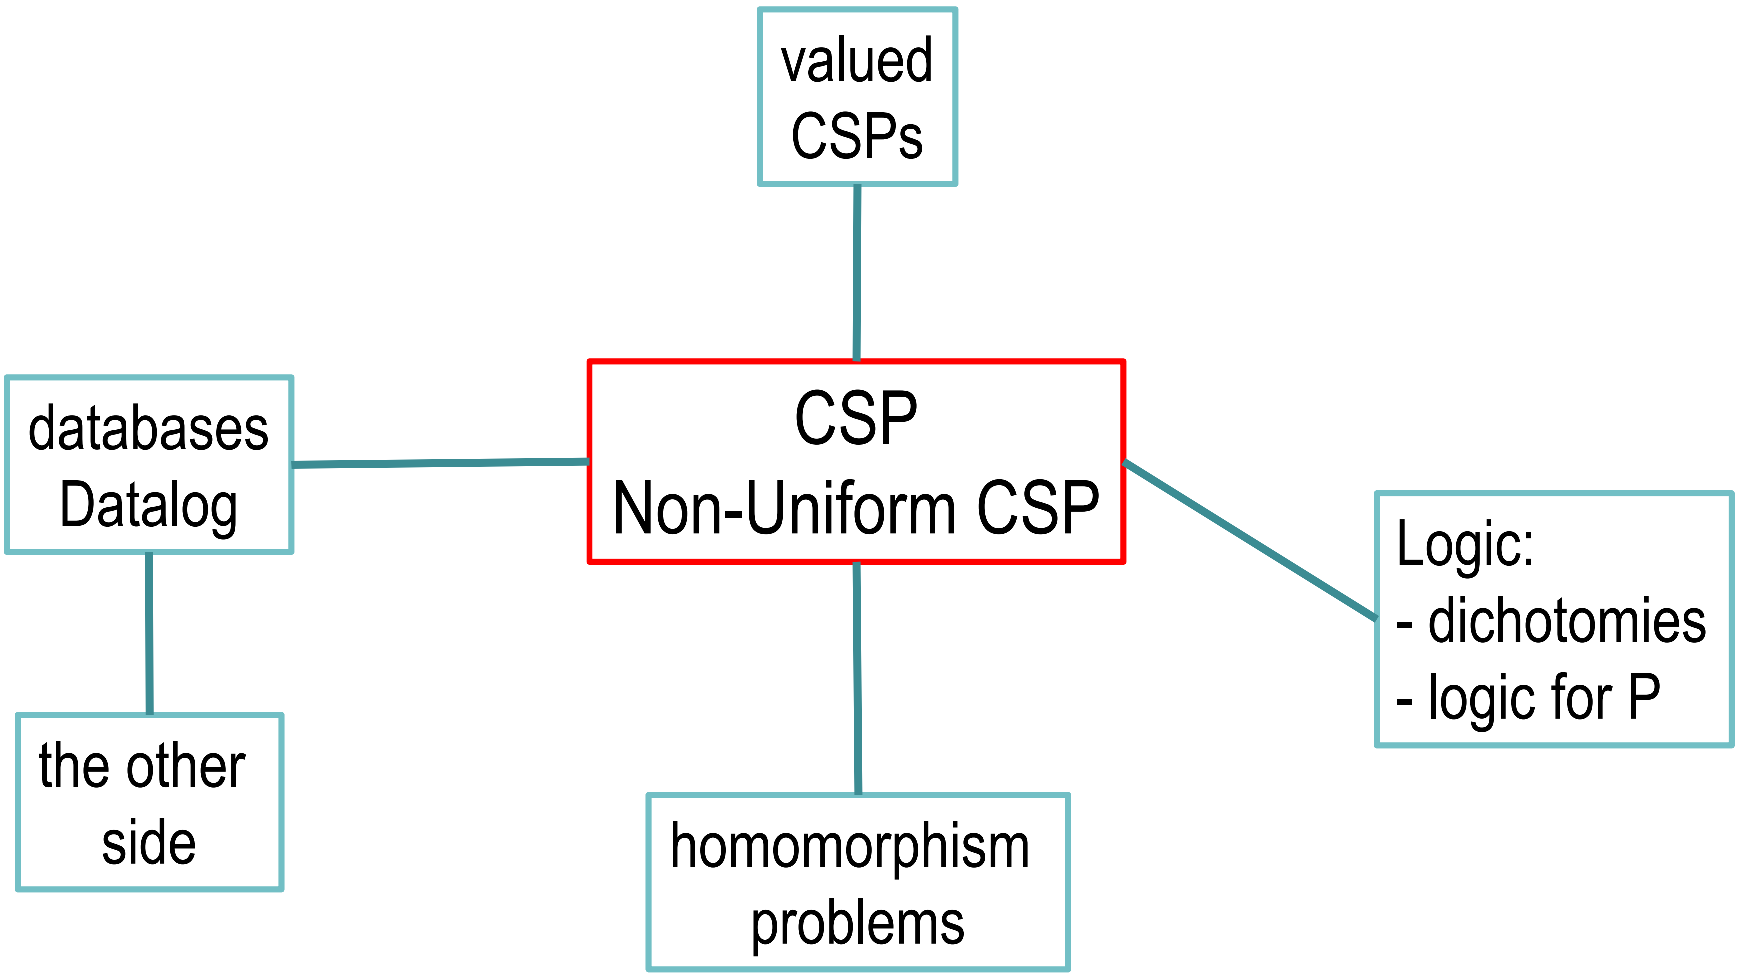

Valued CSPs

MaxCSP/MinCSP:

Given a CSP instance, satisfy as many constraints as possible / unsatisfy as few as possible

Valued CSPs:

Same as MinCSP, except every tuple in a constraint has a (numerical) value, and we need to minimize the total value of such tuples produced by an assignment

Valued CSP: Complexity

Zivny/Thapper:

Without crisp constraints, the only poly time algorithm is linear programming

Kolmogorov/Krokhin/Rolinek:

With crisp constraints, LP+whatever algorithm for CSP is the best that can be done

CSP and Friends

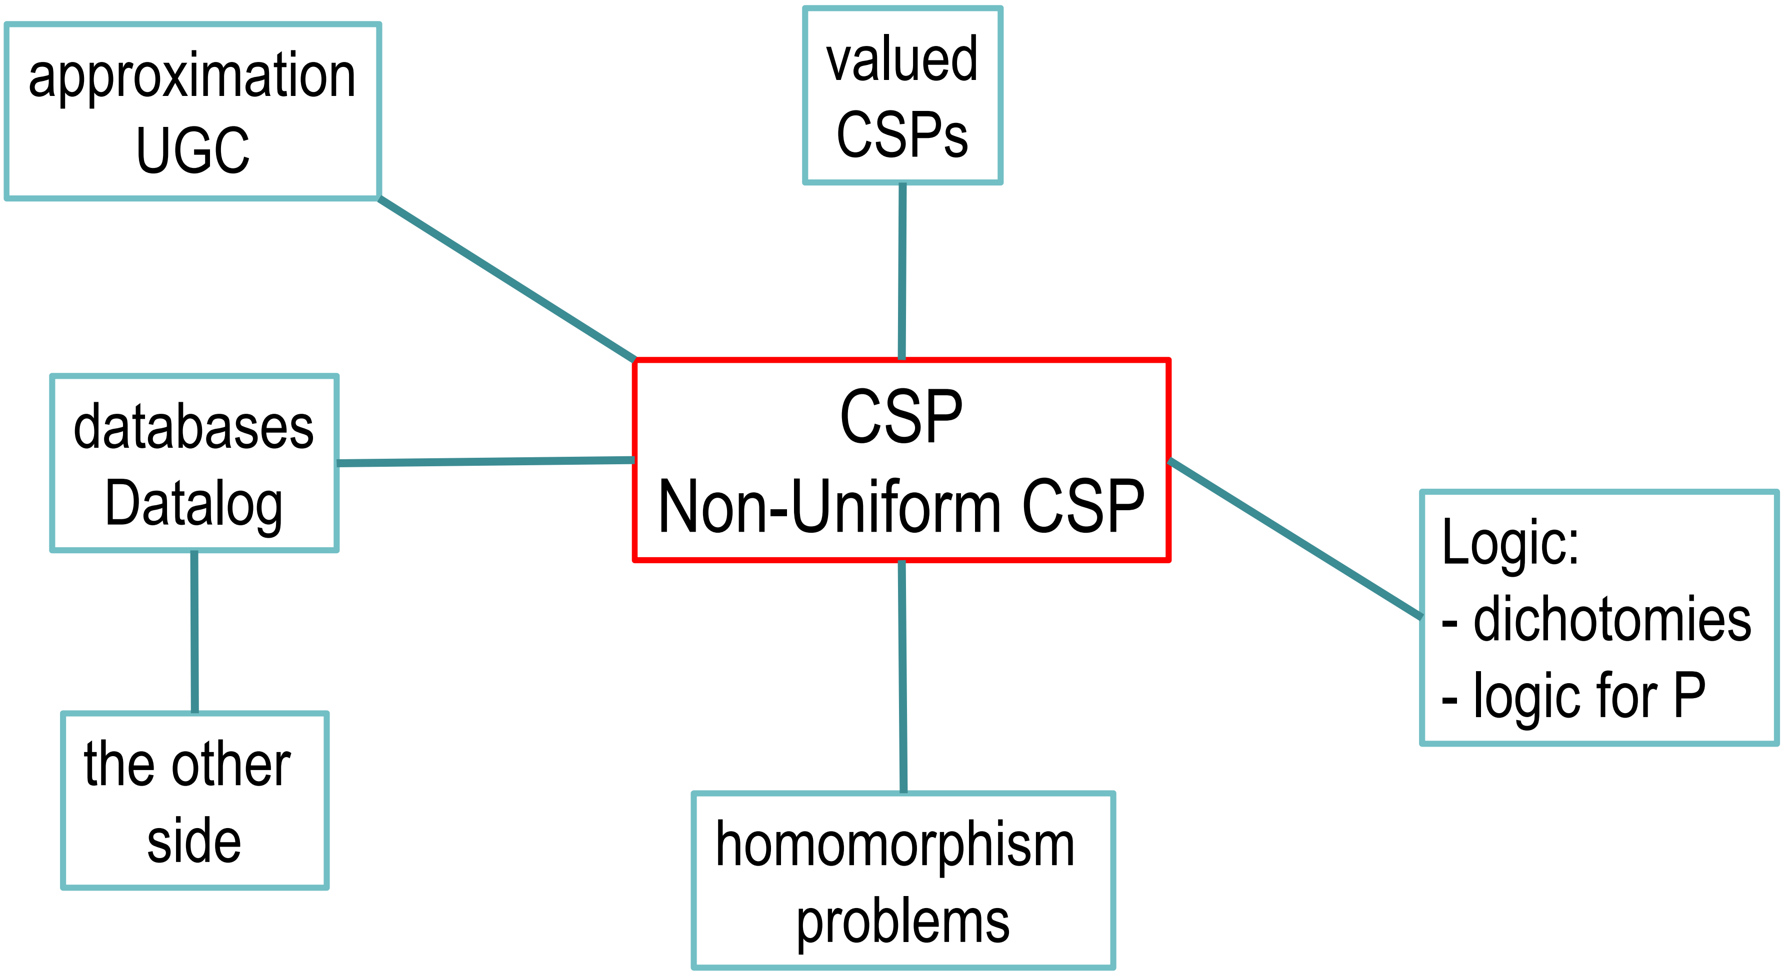

Approximation

Approximation algorithms and complexity is a big area

Often we are talking about approximating a MaxCSP or a Valued CSP

Approximation: Unique Games Conjecture

Consider a CSP with binary constraints

$$R_1(x, y) \wedge R_2(z, x) \wedge \ldots$$

where each relation is the graph of a permutation

Unique Games Conjecture (Khot):
Such a CSP is absolutely impossible to approximate

Raghavendra:
Assuming UGC, an optimal approximation algorithm for any CSP without crisp constraint

CSP and Friends

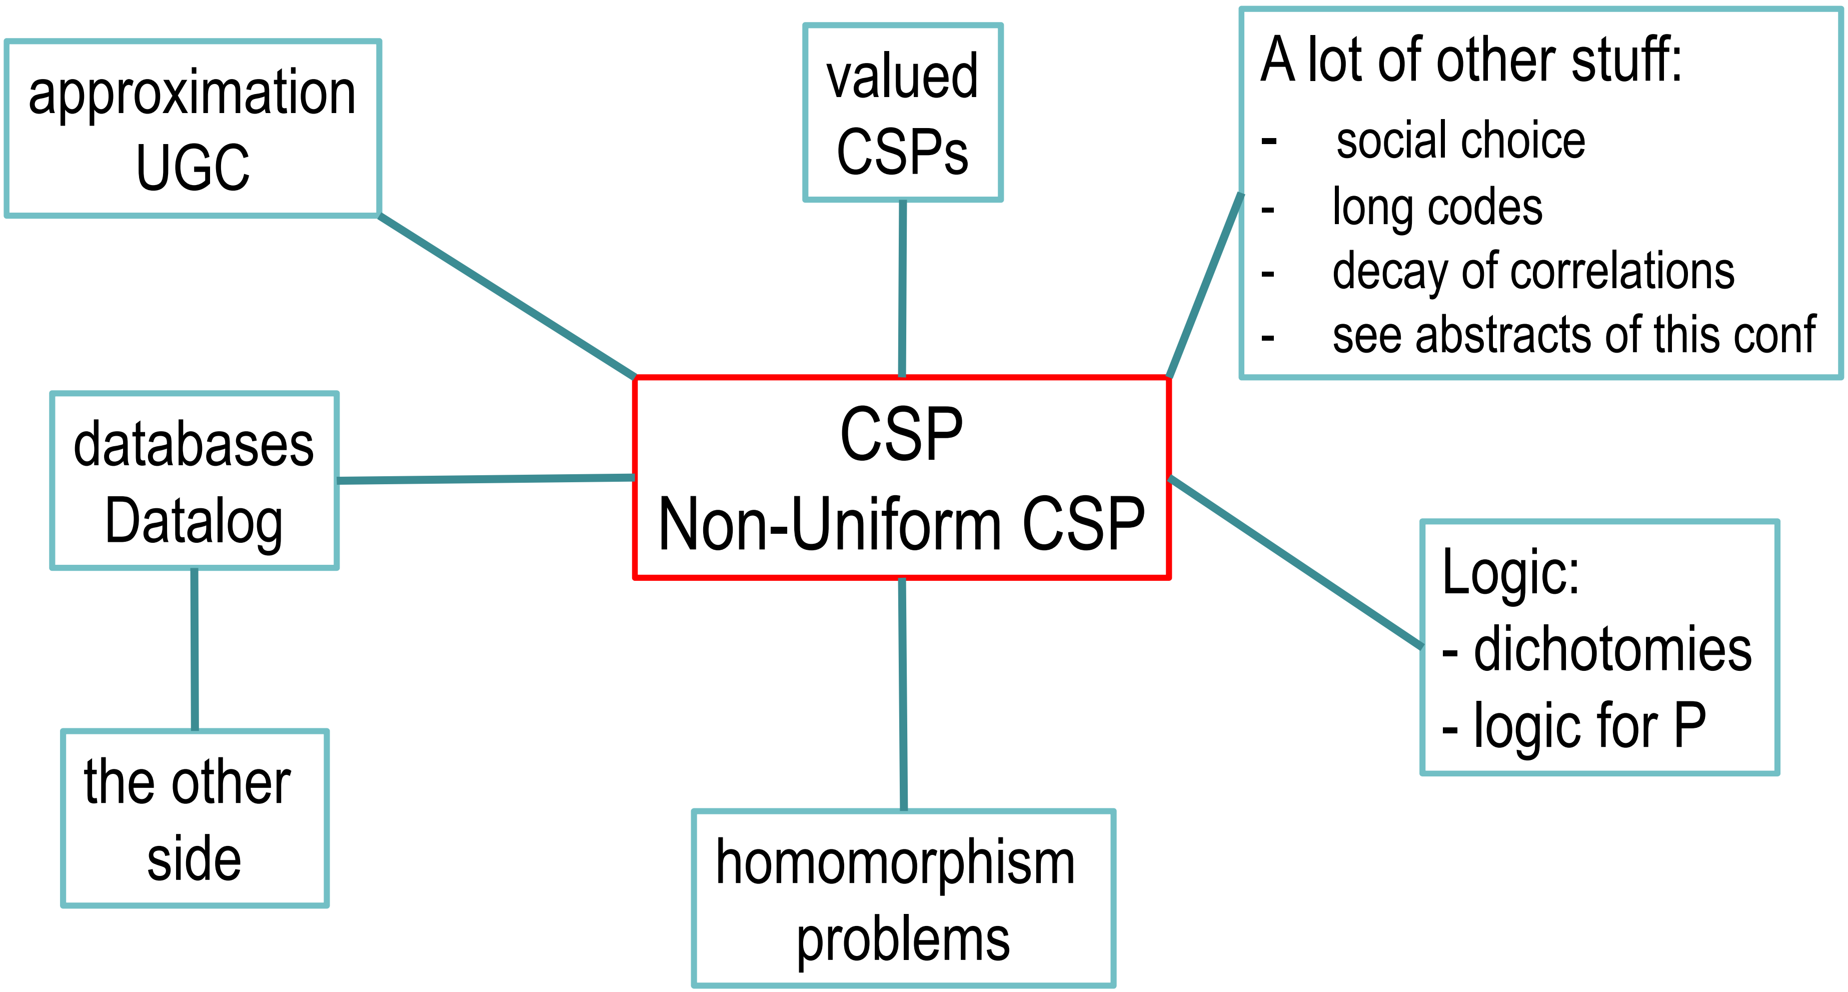

Now the talk begins

Dichotomy conjecture and theorem

Theorem
For any finite class $\mathcal{A}$ of finite similar algebras the problem **CSP($\mathcal{A}$)** is either solvable in polynomial time or NP-complete.

It suffices to prove the theorem for idempotent algebras

Theorem
For any finite class $\mathcal{A}$ of finite similar idempotent algebras the problem **CSP($\mathcal{A}$)** is solvable in polynomial time if $\mathcal{A}$ has a WNU. It is NP-complete otherwise.

Two Main Algorithms

- Local propagation algorithms: Datalog (Vardi, Kolaitis, Dalmau, Barto, Kozik, B., …)

- Few subalgebras: edge term, generating set for solutions (B., Dalmau, Berman, Idziak, Marković, McKenzie, Valeriote, Kearns, Szendrei)

Ingredients

- Separation of prime congruence intervals
- Semilattice edges
- Algorithm

Separation of prime congruence intervals

Let $\boldsymbol{R}$ be a subdirect product of $\boldsymbol{A}_1 \times \cdots \times \boldsymbol{A}_n$,
let $i, j \in \{1, \ldots, n\}$ and $\alpha \prec \beta$, $\gamma \prec \delta$ prime intervals in $Con(\boldsymbol{A}_i)$ and $Con(\boldsymbol{A}_j)$, respectively

We say that $\alpha \prec \beta$ *can be separated* from $\gamma \prec \delta$, if there is a polynomial f of $\boldsymbol{R}$ such that $f(\beta) \nsubseteq \alpha$ while $f(\delta) \subseteq \gamma$

Coherent Sets

Let P = (V,$\mathbf{A}$,$\mathcal{C}$) be an instance.
Let $v \in V$ and $\alpha \prec \beta$ a prime interval in $Con(\mathbf{A}_v)$

The set $W = W(v, \alpha, \beta)$ of all $w \in V$ such that $Con(\mathbf{A}_w)$ contains $\gamma \prec \delta$ such that $\alpha \prec \beta, \gamma \prec \delta$ cannot be separated is called a coherent set

Coherent Sets II

Let $P = (V, \mathcal{A}, \mathcal{C})$ be an instance.

$\boldsymbol{P}_W$ is a restricted problem $(W, \mathcal{A}, \mathcal{C}|_W)$:

$$R(s) \rightarrow \; pr_{s \cap W} R(s \cap W)$$

Condition (QC): some commutator-like condition of a prime interval in a congruence lattice

Theorem

If $\alpha \prec \beta$ does not satisfy Condition (QC) then $\boldsymbol{P}_W$ can be decomposed into a constant number of instances over smaller domains

Splitting Instances

Let $\boldsymbol{P}_W$ be as before and $\alpha_w \prec \beta_w$ prime interval in $Con(\boldsymbol{A}_w)$ such that $\alpha_w \prec \beta_w$ cannot be separated from $\alpha_u \prec \beta_u$ for any $u, w \in W$

There are $\theta_w \in Con(\boldsymbol{A}_w)$ such that $\boldsymbol{P}_W$ is *$\bar{\theta}$-linked*, that is, for any $u, w \in W$ and $\bar{a}, \bar{b} \in \boldsymbol{P}_{u,w}$ if $(a_u, b_u) \in \theta_u$ then $(a_w, b_w) \in \theta_w$

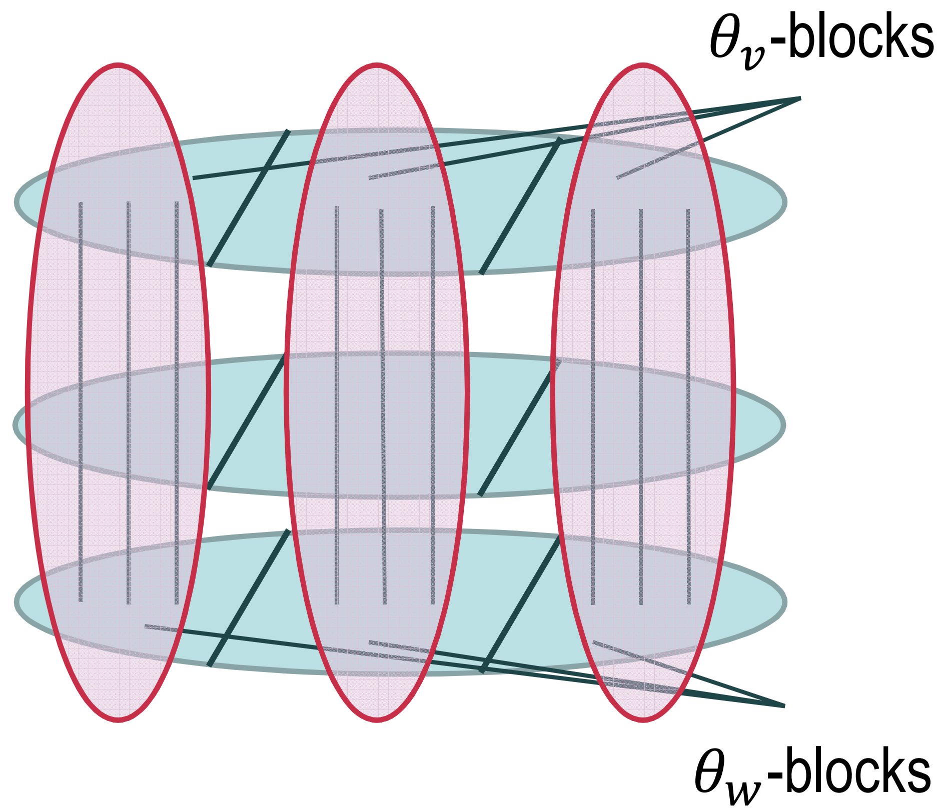

Ingredients

- Separation of prime congruence intervals
- Semilattice edges
- Algorithm

Semilattice Edges

Let $\boldsymbol{A}$ be an algebra.
A pair $a, b \in \boldsymbol{A}$ is said to be a semilattice edge if there is a term operation $\cdot$ of $\boldsymbol{A}$ which is semilattice on $\{a, b\}$, i.e.

- $a \cdot a = a$
- $a \cdot b = b \cdot a = b \cdot b = b$

Operation $\cdot$ can be chosen such that it is semilattice on all semilattice edges of all algebras from $\boldsymbol{A}$

Algebra $\boldsymbol{A}$ is semilattice free if it does not have a semilattice edge

Ingredients

- Separation of prime congruence intervals
- Semilattice edges
- Algorithm

Algorithm: Assumptions

Let $P = (V, \mathbf{A}, \mathcal{C})$ be an instance
We will assume:

- every non-semilattice free domain of P is subdirectly irreducible,
  let μ_v denote the monolith of $\boldsymbol{A}_v$

Algorithm: Max and Center

Let P = (V,**A**,C) be an instance

max(P) is the maximal size of domains of P with a semilattice edge

Max(P) $\subseteq V$ is the set of variables whose domains are not semilattice free and have size max(P)

Center(P) $\subseteq V$ is the set of variables $v \in V$ such that $0_v \prec \mu_v$ satisfies Condition (QC)

Algorithm: Cases

Let $P = (V,\mathcal{A},\mathcal{C})$ be an instance

Recursion on max(P)

We consider 3 cases

(A) All the domains in P are semilattice free

(B) $\text{Max}(P) \cap \text{Center}(P) = \varnothing$

(C) $\text{Max}(P) \cap \text{Center}(P) \neq \varnothing$

Algorithm: Case (A)

Theorem
Let A be a semilattice free algebra. Then A has few subpowers

Suppose all the domains in P are semilattice free
Then P can be solved by the few subpowers algorithm

Quotient Problem

Let $P = (V, \mathbf{A}, \mathbf{C})$ be an instance

$\boldsymbol{P}_W/\bar{\mu}$ is the problem $(V, \mathbf{A}/\bar{\mu}, \mathbf{C}/\bar{\mu})$, where

$$R(s) \rightarrow R/\bar{\mu}(s)$$

Algorithm: Block-Minimality

Let P = (V,$\boldsymbol{A}$,C) be an instance
It is called block-minimal, if
for every $v \in V$ and every $\alpha \prec \beta \in Con(\boldsymbol{A}_v)$

- if $\alpha \prec \beta$ does not satisfy Condition (QC), P_W, $W = W(v, \alpha, \beta)$, is minimal

- if $\alpha \prec \beta$ satisfies Condition (QC), then $P_W/\bar{\mu}$ is minimal

Observation: Establishing block minimality is done by solving polynomially many smaller instances

Algorithm: Case (B) - Empty Center

Theorem

Let $P = (V, A, C)$ be a block-minimal instance.

If Max(P) $\cap$ Center(P) = $\emptyset$ then P has a solution.

Algorithm: Case (C) - Nonempty Center

Let α_v^* be μ_v if $v \in Max(P) \cap Center(P)$, and 0_v otherwise

Theorem

Let $P = (V, A, C)$ be a block-minimal instance.

(1) There is a solution φ of P' = $P/\overline{\alpha^*}$ such that for every $v \in V$ for which A_v is not semilattice free, there is a α_v^*-block B_v such that $B_v, \varphi(v)$ is a semilattice edge.

(2) Instance P'' = P$\cdot$ φ is equivalent to P and such that max(P'') < max(P)

Thank you!

Ingredients

- Separation of prime congruence intervals
- Quasi-Centralizers
- Semilattice edges
- Strategies

Separation of prime congruence intervals

Let $\boldsymbol{A}$ be an algebra and $\alpha \prec \beta,\ \gamma \prec \delta$ prime intervals in $Con(\boldsymbol{A})$

We say that $\alpha \prec \beta$ *can be separated* from $\gamma \prec \delta$, if there is a polynomial f of $\boldsymbol{A}$ such that $f(\beta) \nsubseteq \alpha$ while $f(\delta) \subseteq \gamma$

Let $\boldsymbol{R}$ be a subdirect product of $\boldsymbol{A}_1 \times \cdots \times \boldsymbol{A}_n$,
let $i, j \in \{1, \ldots, n\}$ and $\alpha \prec \beta,\ \gamma \prec \delta$ prime intervals in $Con(\boldsymbol{A}_i)$ and $Con(\boldsymbol{A}_j)$, respectively
We say that $\alpha \prec \beta$ *can be separated* from $\gamma \prec \delta$, if there is a polynomial f of $\boldsymbol{R}$ such that $f(\beta) \nsubseteq \alpha$ while $f(\delta) \subseteq \gamma$

Collapsing polynomials

Let $\boldsymbol{R}$ be a subdirect product of $\boldsymbol{A}_1 \times \cdots \times \boldsymbol{A}_n$,
let $\alpha \prec \beta$ be a prime interval in $Con(\boldsymbol{A}_1)$ be such that $\alpha \prec \beta$ can be separated from EVERY interval $\gamma \prec \delta$ from $Con(A_j)$ for EVERY $j \neq 1$

Then there is a polynomial f of R such that

- $f(\beta) \nsubseteq \alpha$
- $|f(A_j)| = 1$ for every $j \neq 1$

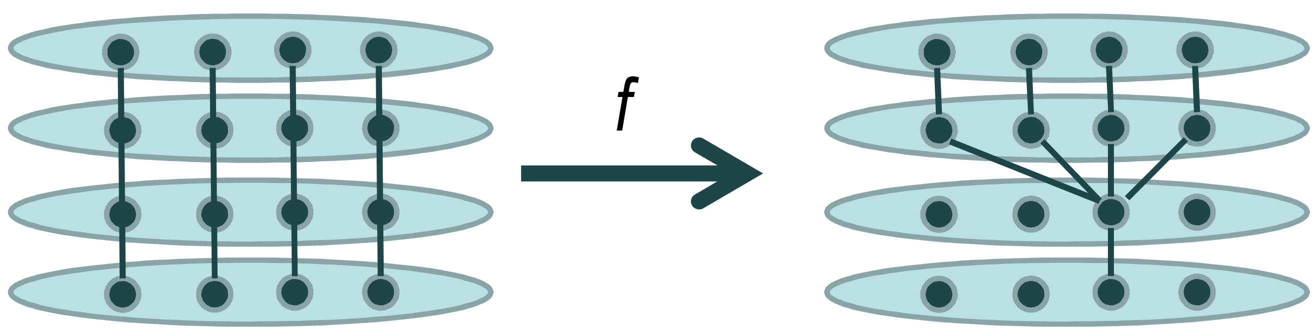

Quasi-Centralizers

Let $\boldsymbol{A}$ be an algebra and $\alpha \prec \beta$ prime intervals in $Con(\boldsymbol{A})$
$\chi(\alpha, \beta)$ denotes the binary relation on A given by:

$(a, b) \in \chi(\alpha, \beta)$ iff for any term $f(x, y, z_1, \dots, z_n)$ and any $c_1, \dots, c_n \in A$: $g(\beta) \subseteq \alpha \Leftrightarrow h(\beta) \subseteq \alpha$,
where $g(x) = f(x, a, c_1, \dots, c_n)$ and
$\mathrm{h}(x) = f(x, b, c_1, \dots, c_n)$

It is a congruence of $\boldsymbol{A}$

Splitting Relations

Let $\boldsymbol{R}$ be a subdirect product of $\boldsymbol{A}_1 \times \boldsymbol{A}_2$ and $\alpha \prec \beta, \gamma \prec \delta$ prime intervals in $Con(\boldsymbol{A}_1), Con(\boldsymbol{A}_2)$, respectively, such that they cannot be separated from each other.

Also, let $\theta_1 = \chi(\alpha, \beta), \theta_2 = \chi(\gamma, \delta)$

Then $\boldsymbol{R}$ is *$\bar{\theta}$-linked*, that is, for any $(a, b), (c, d) \in R$ if $(a, c) \in \theta_1$ then $(b, d) \in \theta_2$ and the other way round

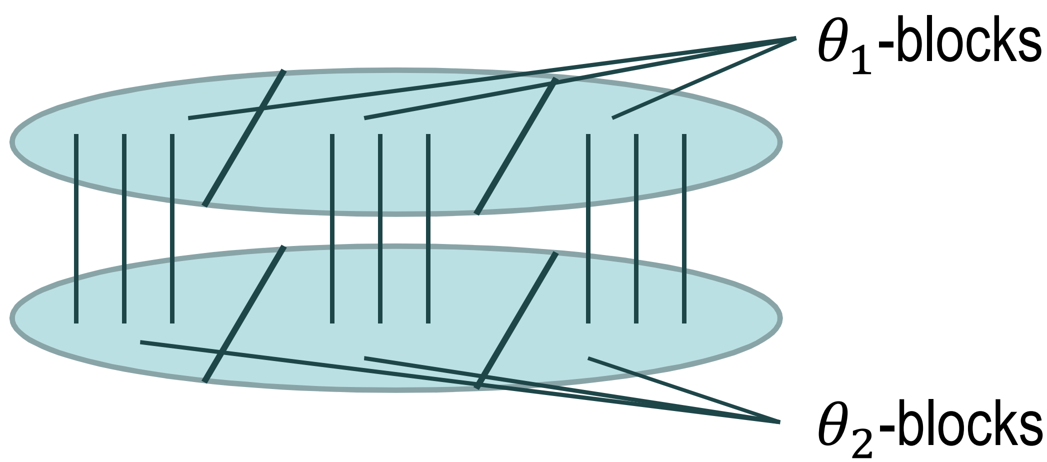

Splitting Relations II

Let ***R*** be a subdirect product of $\boldsymbol{A}_1 \times \cdots \times \boldsymbol{A}_n$ and $\alpha_i \prec \beta_i$ prime interval in $Con(\boldsymbol{A}_i)$ such that $\alpha_i \prec \beta_i$ cannot be separated from $\alpha_j \prec \beta_j$ for any i, j.

Also, let $\theta_i = \chi(\alpha_i, \beta_i)$,

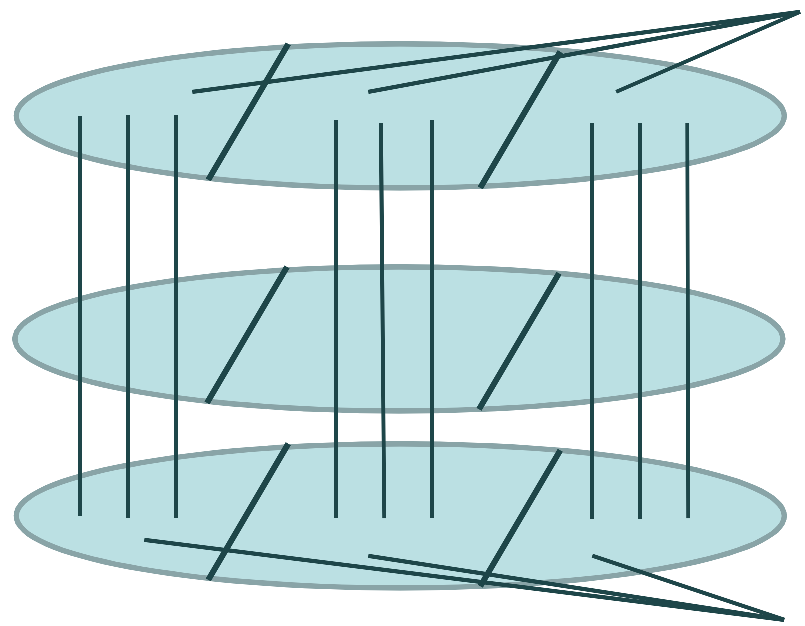

Then ***R*** is *$\bar{\theta}$-linked*, that is, for any $\bar{a}, \bar{b} \in R$ if $(a_i, b_i) \in \theta_i$ then $(a_j, b_j) \in \theta_j$ for any i, j

Coherent Sets

Let P = (V,***A***,***C***) be a (2,3)-minimal instance.
Let $v \in V$ and $\alpha \prec \beta$ a prime interval in $Con(\boldsymbol{A}_v)$
The set $W = W(v, \alpha, \beta)$ of all $w \in V$ such that $Con(\boldsymbol{A}_w)$ contains a prime interval $\gamma \prec \delta$ and $\alpha \prec \beta, \gamma \prec \delta$ cannot be separated from each other.

Theorem
If $\chi(\alpha, \beta)$ is not the full congruence, P_W can be decomposed into a constant number of instances over smaller domains

Semilattice Edges

Let $\boldsymbol{A}$ be an algebra.
A pair $a, b \in \boldsymbol{A}$ is said to be a semilattice edge if there is a term operation $\cdot$ of $\boldsymbol{A}$ which is semilattice on $\{a, b\}$, i.e.

- $a \cdot a = a$
- $a \cdot b = b \cdot a = b \cdot b = b$

Operation $\cdot$ can be chosen such that it is semilattice on all semilattice edges of all algebras from $\boldsymbol{A}$

For any $a, b \in \boldsymbol{A}$ either $a \cdot b = a$ or $a, a \cdot b$ is a semilattice pair

Semilattice Edges II

Theorem

Let $\boldsymbol{A}$ be an algebra and $\alpha \prec \beta \in Con(\boldsymbol{A})$ such that $\beta \leq \chi(\alpha,\beta)$. For any $a,b,c \in A$ such that $(b,c) \in \beta$ and $(a,b) \in \chi(\alpha,\beta)$, it holds $(a \cdot b, a \cdot c) \in \alpha$.

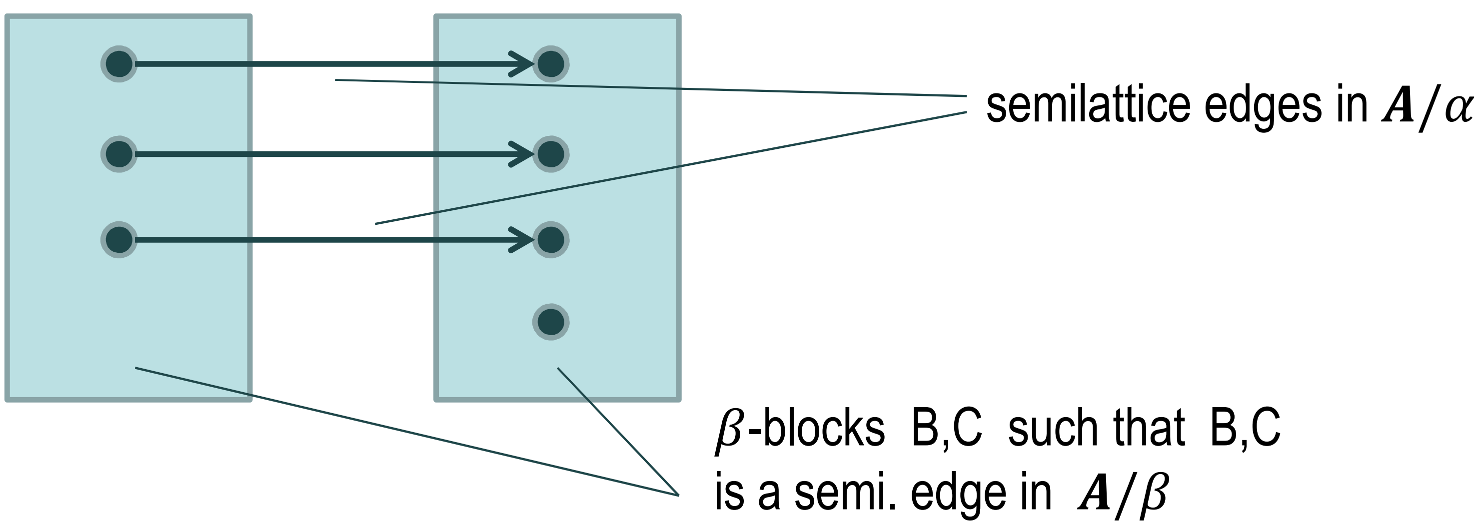

Algorithm: Standard Reductions

Let $P = (V,\mathbf{A},C)$ be an instance
We will assume:

- P is (2,3)-minimal
- every its domain is subdirectly irreducible
  let μ_v denote the monolith of $\boldsymbol{A}_v$

Algorithm: Max and Center

Let $P = (V, \mathbf{A}, C)$ be an instance

max(P) is the maximal size of semilattice free domains of P

Max(P)$\subseteq V$ is the set of variables whose domains are semilattice free and have size max(P)

Center(P) $\subseteq V$ is the set of variables $v \in V$ such that $\chi(0_v, \mu_v)$ is the full congruence

Algorithm: Cases

Let $P = (V, \mathcal{A}, \mathcal{C})$ be an instance

We consider 3 cases

(A) All the domains in P are semilattice free

(B) Max(P) $\cap$ Center(P) = $\emptyset$

(C) Max(P) $\cap$ Center(P) $\neq \emptyset$

Algorithm: Case (A)

Theorem
Let A be a semilattice free algebra. Then A has few subpowers

Suppose all the domains in P are semilattice free
Then P can be solved by the few subpowers algorithm

Algorithm: Block-Minimality

Let $P = (V, \boldsymbol{A}, \boldsymbol{C})$ be an instance
It is called block-minimal, if
for every $v \in V$ and every $\alpha \prec \beta \in Con(\boldsymbol{A}_v)$

- if $\chi(\alpha, \beta)$ is not the full congruence, P_W,
$W = W(v, \alpha, \beta)$, is minimal

- if $\chi(\alpha, \beta)$ is the full congruence, then $P_W / \bar{\mu}$ is minimal

Observation: Establishing block minimality is done by solving polynomially many smaller instances

Algorithm: Case (B) - Empty Center

Theorem

Let $P = (V, A, C)$ be a block-minimal instance.

If $\text{Max}(P) \cap \text{Center}(P) = \varnothing$ then P has a solution.

Algorithm: Case (C) - Nonempty Center

Let α_v^* be μ_v if $v \in Max(P) \cap Center(P)$, and 0_v otherwise

Theorem
Let $P = (V, \mathcal{A}, \mathcal{C})$ be a block-minimal instance.
(1) If P = P$/\overline{\alpha^*}$ is 1-minimal then there is a solution φ of P' such that for every $v \in V$ such that A_v is not semilattice free there is a α_v^*-block B_v such that $B_v, \varphi(v)$ is a semilattice edge.
(2) Instance P'' = P$\cdot\varphi$ is equivalent to P and such that max(P'') < max(P)

Strategies I

Theorem
Let $P = (V,\mathbf{A},\mathcal{C})$ be a block-minimal instance.
If Max(P) $\cap$ Center(P) = $\varnothing$ then P has a solution.

We show that for any $\beta_v \in Con(\mathbf{A}_v)$ there is a solution of $P/\bar{\beta}$.

If β_v is the full congruence, such a solution exists
If $\beta_v = 0_v$ then we have a solution of P

Strategies II

Let $\beta_v \in Con(\mathbf{A}_v)$ and B_v a β_v-block

$W(\bar{\beta})$ is the set of triples (v, α, β), where $v \in V,\ \alpha \prec \beta \leq \beta_v \in Con(\mathbf{A}_v)$

Let **R** be a collection of relations $R_{C,v,\alpha\beta}$ for each constraint $C = \langle s, R\rangle \in \boldsymbol{C}$ and $(v, \alpha, \beta) \in W(\bar{\beta})$

Let $S(C, v, \alpha\beta) = s \cap W(v, \alpha, \beta)$ be the set of its coordinate positions

A tuple $\boldsymbol{a} \in \prod_{x \in X} \boldsymbol{A}_x$ for $X \subseteq V$ is said to be **R**-compatible if for any $C = \langle s, R\rangle \in \boldsymbol{C}$ and $(v, \alpha, \beta) \in W(\bar{\beta})$ $pr_T a \in pr_T R_{C,v,\alpha\beta}$, where $T = X \cap S(C, v, \alpha\beta)$

Strategies III

R is said to be a $\bar{\beta}$-strategy with respect to $\bar{B}$ if for every $C = \langle s, R \rangle \in$ C and $(v, \alpha, \beta) \in W(\bar{\beta})$ the following conditions hold ($W = W(v, \alpha\beta)$):

(S1) the relations $R^{X,\mathcal{R}}$, $X \subseteq V$, $|X| \leq 2$, consisting of R-compatible tuples from R^X, form a nonempty (2,3)-strategy for P

(S2) for every $(w, \gamma, \delta) \in W(\bar{\beta})$ (let $U = W(v, \alpha, \beta)$) and every $\boldsymbol{a} \in pr_{s \cap W \cap U} R_{C,v,\alpha\beta}$ tuple $\boldsymbol{a}$ extends to an R-compatible solution of $\boldsymbol{P}_U$

Strategies IV

(S3) $R \cap \prod_w B_w \neq \emptyset$ and for any $I \subseteq s$ any R-compatible tuple $\boldsymbol{a} \in pr_I R$ extends to an R-compatible tuple from R

Tightening Strategies

Theorem

Let R be a $\bar{\beta}$-strategy with respect to $\bar{B}$.
Let $(v, \alpha, \beta) \in W(\bar{\beta})$ be such that $\alpha|_{B_v} \neq \beta|_{B_v}$ and $\beta = \beta_v$. Set $\beta'_v = \alpha$ and $\beta'_w = \beta_w$
Let $B'_v \subseteq B_v$ be an α-block.
Then there is a $\bar{\beta}'$-strategy with respect to $\bar{B}'$