Dichotomy Theorems for Counting Constraint Satisfaction Problems

Xi Chen
Columbia University

Constraint Satisfaction Problem

- Let $D = \{1, \ldots, d\}$ be a domain.
- A language is a finite set of relations $\Gamma = \{\Theta_1, \ldots, \Theta_h\}$.
  An instance of CSP(Γ) consists of a set of variables $x_1, \ldots, x_n$ and a set of constraints from Γ. It defines an n-ary relation $R \subseteq D^n$, where $(x_1, \ldots, x_n) \in R$ if all constraints are satisfied.

$$R \subseteq D^4 : \Theta_1(x_1, x_3, x_2) \wedge \Theta_2(x_4, x_3) \wedge \Theta_2(x_2, x_3)$$

- Decide if R is empty or not.

Examples

- d-COLORING: $D = \{1, \ldots, d\}$ and $\Gamma = \{\Theta\}$, where

$$\Theta = \big\{(i,j) : i,j \in D \text{ and } i \neq j\big\}$$

- INDEPENDENT SET: $D = \{1, 2\}$ and $\Gamma = \{\Theta\}$, where

$$\Theta = \big\{(1,1),(1,2),(2,1)\big\}$$

- 2-SAT: $D = \{0, 1\}$ and

$$\Gamma = \big\{x_1 \vee x_2,\ \overline{x_1} \vee x_2,\ x_1 \vee \overline{x_2},\ \overline{x_1} \vee \overline{x_2}\big\}$$

- 3-SAT ...

One of the most important classes of problems in TCS:

- Decision: whether a solution exists?
  The CSP dichotomy conjecture of Feder and Vardi is open.
- Optimization: satisfy as many constraints as possible, and more generally, the valued constraint satisfaction problem to find an assignment to maximize the total weight.
- Counting: This talk.

Counting Constraint Satisfaction Problem

- Let $D = \{1, \ldots, d\}$ be a domain.
- A language is a finite set of relations $\Gamma = \{\Phi_1, \ldots, \Phi_h\}$. An instance of $\#\text{CSP}(\Gamma)$ consists of variables $x_1, \ldots, x_n$ and a set of constraints from Γ. It defines an n-ary relation $R \subseteq D^n$, where $(x_1, \ldots, x_n) \in R$ if all constraints are satisfied.
- Compute $|R|$.

Examples

- $\#\,d$-COLORING: $D = \{1, \ldots, d\}$ and $\Gamma = \{\Theta\}$, where

$$\Theta = \big\{(i,j) : i,j \in D \text{ and } i \neq j\big\}.$$

- $\#$ INDEPENDENT SET: $D = \{1, 2\}$ and $\Gamma = \{\Theta\}$, where

$$\Theta = \big\{(1,1), (1,2), (2,1)\big\}.$$

- $\#$ 2-SAT: $D = \{0, 1\}$ and

$$\Gamma = \Big\{x_1 \vee x_2,\ \overline{x_1} \vee x_2,\ x_1 \vee \overline{x_2},\ \overline{x_1} \vee \overline{x_2}\Big\}$$

- $\#$ 3-SAT …

Weighted #CSP

- A weighted language $\mathcal{L} = \{g_1, \ldots, g_h\}$ with $g_i : D^{r_i} \to \mathbb{C}$.
- An instance of #CSP($\mathcal{L}$) consists of variables $x_1, \ldots, x_n$ over D and a set of functions from $\mathcal{L}$. It defines an n-ary function F: for any assignment $\mathbf{x} = (x_1, \ldots, x_n) \in D^n$, $F(\mathbf{x})$ is the product of the constraint function evaluations. E.g.,

$$F(x_1, x_2, x_3, x_4) = g_1(x_1, x_3, x_2) \cdot g_2(x_2, x_4) \cdot g_2(x_3, x_2)$$

- Compute $\sum_{\mathbf{x} \in D^n} F(\mathbf{x})$.

The special case when $\mathcal{L}$ consists of a single symmetric binary function [Dyer and Greenhill 00], [Bulatov and Grohe 05], [Goldberg, Grohe, Jerrum and Thurley 09], [Cai, C and Lu 11].

Complexity Dichotomies Arise for #CSP

Theorem (Bulatov 08)

#CSP(Γ) either can be solved in P-time or is #P-complete.

Theorem (Dyer and Richerby 10)

An alternative proof; the tractability criterion is decidable in NP.

Further extended to nonnegative rational languages [Bulatov, Dyer, Goldberg, Jalsenius, Jerrum and Richerby 10], and nonnegative algebraic languages [Cai, C and Lu 11].

Theorem (Cai and C 12)

A dichotomy for #CSP($\mathcal{L}$) with complex weights.

Plan of the Talk

1. Dichotomy for Unweighted #CSP:
    - Tractability criterion: Strong balance
    - Mal'tsev polymorphisms and Witness functions
    - The main counting algorithm.
2. Dichotomy for Nonnegative and Complex #CSP

Definition

An $n \times m$ nonnegative matrix is rectangular if $A_{i,k}, A_{i,\ell}, A_{j,k} > 0$ imply $A_{j,\ell} > 0$ (or block-diagonal where every block is all positive).

$$\begin{pmatrix} 1 & 2 & 0 & 0 & 0 \\ 2 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 2 & 2 & 0 \\ 0 & 0 & 3 & 3 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Definition

A nonnegative matrix is block-rank-1 *if one can permute its rows and columns to make it block-diagonal and every block is rank* 1.

$$\begin{pmatrix} 1 & 2 & 0 & 0 & 0 \\ 2 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 2 & 2 & 0 \\ 0 & 0 & 3 & 3 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

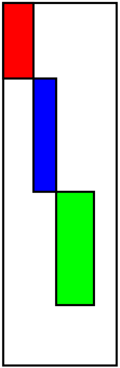

Strong Rectangularity and Strong Balance

- Let $\Gamma = \{\Theta_1, \ldots, \Theta_h\}$ be a language over D.
- Given an n-ary relation $R \subseteq D^n$ derived by an instance of $\#$CSP(Γ) and integers k, ℓ, r such that $k + \ell + r = n$, we are interested in the following $|D|^k \times |D|^\ell$ matrix $\mathbf{M}$:

$$M(\mathbf{u}, \mathbf{v}) = \Big|\big\{\mathbf{w} \in D^r : (\mathbf{u}, \mathbf{v}, \mathbf{w}) \in R\big\}\Big|,$$

with rows indexed by $\mathbf{u} \in D^k$, columns indexed by $\mathbf{v} \in D^\ell$.

Strong Rectangularity and Strong Balance

Definition (Dyer and Richerby 10)

Γ *is strongly rectangular if every such matrix* $\mathbf{M}$ *is rectangular;*
Γ *is strongly balanced if every such matrix* $\mathbf{M}$ *is block-rank-*1.

Strong balance implies strong rectangularity.
equivalent to congruence singularity [Bulatov 08].

Theorem (Bulatov 08)

If Γ is congruence singular, then $\#CSP(\Gamma)$ is solvable in P-time; otherwise $\#CSP(\Gamma)$ is #P-hard.

Theorem (Dyer and Richerby 10)

If Γ is strongly balanced, then $\#CSP(\Gamma)$ is solvable in P-time; otherwise $\#CSP(\Gamma)$ is #P-hard.

Proof of the Hardness Part:

Gadget construction: A reduction from EVAL($\mathbf{A}$) to $\#CSP(\Gamma)$ for a nonnegative $\mathbf{A}$ that violates the condition of [Bulatov-Grohe 05].

Polymorphisms

Definition

Let $\Theta \subseteq D^r$ *be an r-ary relation, and* $\psi : D^3 \to D$ *be a map. Then we say* ψ *is a* *polymorphism* *of* Θ *if* $\mathbf{u}, \mathbf{v}, \mathbf{w} \in \Theta$ *implies*

$$\Big(\psi(u_1, v_1, w_1), \psi(u_2, v_2, w_2), \ldots, \psi(u_r, v_r, w_r)\Big) \in \Theta.$$

$$\begin{array}{cccc} u_1 & u_2 & \ldots & u_r \\ v_1 & v_2 & \ldots & v_r \\ w_1 & w_2 & \ldots & w_r \\ \hline \psi(u_1, v_1, w_1) & \psi(u_2, v_2, w_2) & \ldots & \psi(u_r, v_r, w_r) \end{array}$$

Mal'tsev Polymorphisms

Definition

Map ψ is a Mal'tsev polymorphism of Θ if it also satisfies

$$\psi(a,b,b) = \psi(b,b,a) = a, \quad \text{for all } a,b \in D.$$

We say ψ is a Mal'tsev polymorphism of $\Gamma = \{\Theta_1, \dots, \Theta_h\}$ if ψ is a Mal'tsev polymorphism of every relation Θ_i.

Observation

If ψ is a Mal'tsev polymorphism of Γ, then it is also a Mal'tsev polymorphism of every relation R derived by a $\#CSP(\Gamma)$ instance.

$$R \subseteq D^4 : \Theta_1(x_1, x_2, x_3) \wedge \Theta_2(x_3, x_4) \wedge \Theta_2(x_4, x_2)$$

$$\begin{array}{cccc} u_1 & u_2 & u_3 & u_4 \\ v_1 & v_2 & v_3 & v_4 \\ w_1 & w_2 & w_3 & w_4 \\ \hline \psi(u_1, v_1, w_1) & \psi(u_2, v_2, w_2) & \psi(u_3, v_3, w_3) & \psi(u_4, v_4, w_4) \end{array}$$

Theorem

Γ is strongly rectangular iff it has a Mal'tsev polymorphism ψ.

Proof of the easier direction $\Leftarrow$.

Let R be a relation derived by a #CSP(Γ) instance. Let $\mathbf{u}, \mathbf{u}' \in D^k$ $\mathbf{v}, \mathbf{v}' \in D^\ell$. If the $(\mathbf{u}', \mathbf{v}), (\mathbf{u}, \mathbf{v}), (\mathbf{u}, \mathbf{v}')$ entries of $\mathbf{M}$ are positive:

$$\begin{array}{ccccccccccc} u_1' & \dots & u_k' & v_1 & \dots & v_\ell & w_1 & \dots & w_r & \in R \\ u_1 & \dots & u_k & v_1 & \dots & v_\ell & w_1' & \dots & w_r' & \in R \\ u_1 & \dots & u_k & v_1' & \dots & v_\ell' & w_1'' & \dots & w_r'' & \in R \\ \hline u_1' & \dots & u_k' & v_1' & \dots & v_\ell' & w_1^* & \dots & w_r^* & \in R \end{array}$$

This implies that the $(\mathbf{u}', \mathbf{v}')$ entry of $\mathbf{M}$ is also positive. □

Assume that ψ is given.

Let ψ be a Mal'tsev polymorphism of $R \subseteq D^n$. For $i \in [n]$:

1. $\mathrm{Pr}_i\, R$: the projection of R on the ith coordinate, i.e., $a \in \mathrm{Pr}_i\, R$ if there exists $\mathbf{u} \in R$ with $u_i = a$ (called a witness).
2. $\sim_i$ over $\mathrm{Pr}_i\, R$: $a \sim_i b$ if there exist $\mathbf{u} \in D^{i-1}$, $\mathbf{w}, \mathbf{w}' \in D^{n-i}$:

$$(\mathbf{u}, a, \mathbf{w}) \in R \quad \text{and} \quad (\mathbf{u}, b, \mathbf{w}') \in R.$$

Lemma

If R has a Mal'tsev polymorphism, $\sim_i$ is an equivalence relation.

Proof.

Goal: $a \sim_i b$ and $b \sim_i c$ imply $a \sim_i c$.

$a \sim_i b \Rightarrow$ there exist $\mathbf{u}, \mathbf{v}, \mathbf{v}'$ such that $(\mathbf{u}, a, \mathbf{v}), (\mathbf{u}, b, \mathbf{v}') \in R$.

$b \sim_i c \Rightarrow$ there exist $\mathbf{u}', \mathbf{w}, \mathbf{w}'$ such that $(\mathbf{u}', b, \mathbf{w}), (\mathbf{u}', c, \mathbf{w}') \in R$.

$$\begin{array}{cccc} \mathbf{u} & a & \mathbf{v} & \in R \\ \mathbf{u} & b & \mathbf{v}' & \in R \\ \mathbf{u}' & b & \mathbf{w} & \in R \\ \hline \mathbf{u}' & a & \mathbf{w}^* & \in R \end{array}$$

Since $(\mathbf{u}', c, \mathbf{w}') \in R$, we have $a \sim_i c$. □

Witness Function: A Succinct Representation

Definition (Dyer and Richerby 10)

Suppose $R \subseteq D^n$ has a Mal'tsev polymorphism ψ. We say

$$\omega : [n] \times D \to D^n \cup \{\text{NIL}\}$$

is a witness function of R if for every $i \in [n]$:

1. $a \notin \Pr_i R \;\Rightarrow\; \omega(i,a) = \text{NIL}$;
2. $a \in \Pr_i R$, $\omega(i,a) \in R$ *and its* i*th entry is* a;
3. *If* $a \sim_i b$, $\omega(i,a)$ *and* $\omega(i,b)$ *share the same* $(i-1)$*-prefix.*

Similar to the compact representation of [Bulatov and Dalmau 06].
A witness function ω of $R \subseteq D^n$ is of polynomial length.

Witness Function: A Succinct Representation

Lemma

Suppose $R \subseteq D^n$ has a Mal'tsev polymorphism ψ. Given ω and a tuple $\mathbf{u} \in D^n$, one can decide if $\mathbf{u} \in D^n$ or not in P-time.

Lemma

Suppose $R \subseteq D^n$ has a Mal'tsev polymorphism ψ. Given ω and $\mathbf{u} \in D^t$ for some $t \leq n$, one can decide if $\mathbf{u} \in \mathrm{Pr}_{[t]}\, R$ in P-time.

Round 1: Check if $u_1 \in \text{Pr}_1 R$; if so find a witness.

- If $\omega(1, u_1) = \text{NIL}$, reject.
  Otherwise, let $\omega(1, u_1) = (u_1, v_2, \mathbf{w}) \in R$ (a witness).

Let $\omega(1, u_1) = (u_1, v_2, \mathbf{w}) \in R$.

Round 2: Check if $(u_1, u_2) \in \Pr_{[2]} R$; if so find a witness.

1. If $\omega(2, u_2) = \text{NIL}$, reject.
2. If $\omega(2, u_2)$ and $\omega(2, v_2)$ have different first entries, reject. As $(u_1, u_2) \in \Pr_{[2]} R$ would imply that $u_2 \sim_2 v_2$.
3. Otherwise, let $\omega(2, u_2) = (w_1, u_2, \mathbf{w}')$, $\omega(2, v_2) = (w_1, v_2, \mathbf{w}^*)$.

$$\begin{array}{cccc} u_1 & v_2 & \mathbf{w} & \in R \\ w_1 & v_2 & \mathbf{w}^* & \in R \\ w_1 & u_2 & \mathbf{w}' & \in R \\ \hline u_1 & u_2 & \mathbf{w}'' & \in R \end{array}$$

The result $(u_1, u_2, \mathbf{w}'')$ is a witness for $(u_1, u_2) \in \Pr_{[2]} R$.

Repeat for t rounds ...

Lemma (Dyer and Richerby 10)

Suppose Γ *has a Mal'tsev polymorphism* ψ*. Given an #CSP instance, a witness function for its relation can be built in P-time.*

Tractability

Assume that $\Gamma = \{\Theta_1, \ldots, \Theta_h\}$ is strongly balanced and thus, has a Mal'tsev polymorphism ψ.

Given a $\#$CSP(Γ) instance with n variables that defines $R \subseteq D^n$:

- For each $t = 1, \ldots, n$, let

$$F^{[t]}(x_1, \ldots, x_t) = \left|\left\{\mathbf{w} \in D^{n-t} : (x_1, \ldots, x_t, \mathbf{w}) \in R\right\}\right|.$$

View each $F^{[t]}$, $t \geq 2$, as a $d^{t-1} \times d$ matrix.

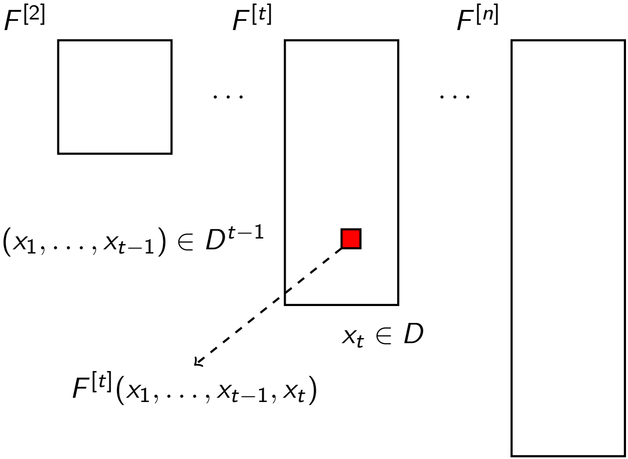
$F^{[2]}$
$F^{[t]}$
$F^{[n]}$
$\ldots$
$\ldots$
$(x_1, \ldots, x_{t-1}) \in D^{t-1}$
$x_t \in D$
$F^{[t]}(x_1, \ldots, x_{t-1}, x_t)$

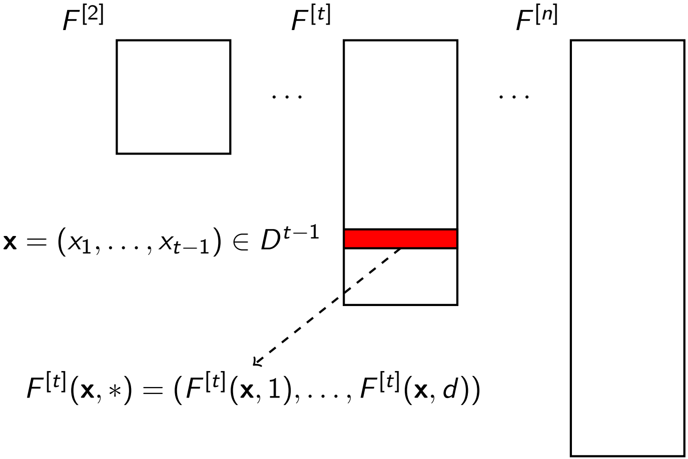
$F^{[2]}$
$F^{[t]}$
$F^{[n]}$
$\mathbf{x} = (x_1, \ldots, x_{t-1}) \in D^{t-1}$
$F^{[t]}(\mathbf{x}, *) = (F^{[t]}(\mathbf{x}, 1), \ldots, F^{[t]}(\mathbf{x}, d))$

For each $t \geq 2$, access to a data structure for $F^{[t]}$:

- One can send a $(t-1)$-tuple $\mathbf{x} \in D^{t-1}$ to the data structure. If $F^{[t]}(\mathbf{x}, *) = \mathbf{0}$, return $\mathbf{v} = \mathbf{0}$; otherwise, return a nonzero vector $\mathbf{v}$ linearly dependent with $F^{[t]}(\mathbf{x}, *)$, in P-time.

Given access to such data structures, compute

$$|R| = \sum_{a_1 \in D} F^{[1]}(a_1).$$

The Main Counting Algorithm

To compute $F^{[1]}(a_1)$ for some $a_1 \in D$:

1. send (a_1) to the data structure for $F^{[2]}$
2. receive $\mathbf{v}$ that is linearly dependent with $F^{[2]}(a_1, *)$
3. if $\mathbf{v} = \mathbf{0}$, $F^{[2]}(a_1, *) = \mathbf{0} \;\Rightarrow\; F^{[1]}(a_1) = 0$
4. otherwise, let v_{a_2} be a nonzero entry of $\mathbf{v}$, $a_2 \in D$:

$$F^{[1]}(a_1) = \sum_{b \in D} F^{[2]}(a_1, b) = F^{[2]}(a_1, a_2) \left(\frac{1}{v_{a_2}} \sum_{b \in D} v_b \right)$$

To compute $F^{[2]}(a_1, a_2)$:

1. send (a_1, a_2) to the data structure for $F^{[3]}$
2. receive $\mathbf{w}$ that is linearly dependent with $F^{[3]}((a_1, a_2), *)$
3. if $\mathbf{w} = \mathbf{0}$, then $F^{[2]}(a_1, a_2) = 0$
4. so $\mathbf{w} \neq \mathbf{0}$; let w_{a_3} be a nonzero entry of $\mathbf{w}$, $a_3 \in D$:

$$F^{[2]}(a_1, a_2) = F^{[3]}(a_1, a_2, a_3) \left(\frac{1}{w_{a_3}} \sum_{b \in D} w_b \right)$$

The Main Counting Algorithm

After $n - 1$ steps, the algorithm reduces $F^{[1]}(a_1)$ to

$$F^{[n]}(a_1, a_2, \ldots, a_n)$$

for some appropriate $a_2, \ldots, a_n \in D$. $F = F^{[n]}$ is easy to evaluate.

Rest of the tractability proof: *How to build the data structures?*

Strong balance $\Rightarrow F^{[t]}$ is a block-rank-1 matrix.

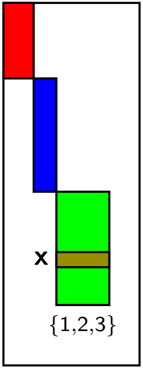

Each equivalence class $\mathcal{E}_i$ of $\sim_t$ corresponds to a block:

$$a \sim_t b \;\Rightarrow\; F^{[t]}(\mathbf{x}, a) > 0 \text{ and } F^{[t]}(\mathbf{x}, b) > 0 \text{ for some } \mathbf{x}$$

To build the data structure for $F^{[t]}$, it suffices to compute a representative vector $\mathbf{v}_j$ for each block (equivalently, each equivalence class $\mathcal{E}_j$ of $\sim_t$.

- For a query $\mathbf{x} \in D^{t-1}$, $\mathbf{x}$ is in the block of $\mathcal{E}_j$
$$\Updownarrow$$
$$F^{[t]}(\mathbf{x}, a) > 0 \text{ for some } a \in \mathcal{E}_j$$
$$\Updownarrow$$
$$(\mathbf{x}, a) \in \Pr_{[t]} R \text{ for some } a \in \mathcal{E}_j$$
- Return $\mathbf{v}_j$ if $\mathbf{x}$ is in the block of $\mathcal{E}_j$; return $\mathbf{0}$ if it does not belong to any block.

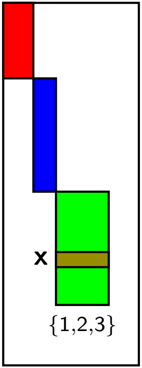

Initialization: The data structure for $F^{[n]}$.

- For each equivalence class $\mathcal{E}_j$ of $\sim_n$, pick an $a \in \mathcal{E}_j$. Let $\omega(n, a) = (\mathbf{x}, a) \in R$. Set the representative vector $\mathbf{v}_j$ of $\mathcal{E}_j$ to be $F^{[n]}(\mathbf{x}, *)$.

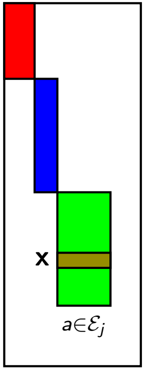

Induction: The data structure for $F^{[t]}$.

- For each equivalence class $\mathcal{E}_j$ of $\sim_i$, pick an $a \in \mathcal{E}_j$. Let $\omega(t, a) = (\mathbf{x}, a, \mathbf{v}) \in R$. Set the representative vector $\mathbf{v}_j$ to be $F^{[t]}(\mathbf{x}, *)$.
- Use data structures for $F^{[t+1]}, \ldots, F^{[n]}$ and the main counting algorithm to evaluate $F^{[t]}$.

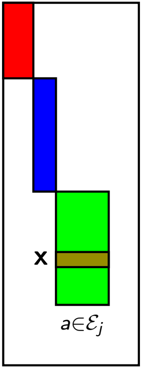

Strong Balance for Nonnegative Languages

- Let $\mathcal{L} = \{g_1, \ldots, g_h\}$ be a nonnegative language.
- Given an n-ary function F derived by an instance of $\#\text{CSP}(\mathcal{L})$ and integers k, ℓ, r such that $k + \ell + r = n$, we are interested in the following $|D|^k \times |D|^\ell$ matrix $\mathbf{M}$:

$$M(\mathbf{u}, \mathbf{v}) = \sum_{\mathbf{w} \in D^r} F(\mathbf{u}, \mathbf{v}, \mathbf{w}).$$

Definition

$\mathcal{L}$ is strongly balanced if every such matrix $\mathbf{M}$ is block-rank-1.

Theorem (Cai, C and Lu)

If a nonnegative language $\mathcal{L}$ is strongly balanced, then $\#CSP(\mathcal{L})$ is solvable in P-time; otherwise, it is $\#P$-hard.

Plan of the Talk

1. Dichotomy for Unweighted #CSP:
    - Tractability criterion: Strong balance
    - Mal'tsev polymorphisms and Witness functions
    - The counting algorithm
    - Generalization to Nonnegative #CSP
2. Dichotomy for #CSP with Complex Values

Cancellations ($\{\pm 1\}$ or even roots of unity) may sometimes lead to efficient algorithms and more tractable cases (e.g., Permanent vs Determinant and Holographic algorithms [Valiant 04]).

Let $\mathcal{L}$ be a complex-valued language, and $F : D^n \to D$ be an n-ary function derived by a #CSP($\mathcal{L}$) instance. Let $R \subseteq D^n$:

$$\mathbf{x} \in R \iff F(\mathbf{x}) \neq 0.$$

Even with a witness function ω of R, not clear how to use ω to decide efficiently if $F^{[t]}(x_1, \ldots, x_t) = 0$ or not, where

$$F^{[t]}(x_1, \ldots, x_t) = \sum_{\mathbf{w} \in D^{n-t}} F(x_1, \ldots, x_t, \mathbf{w}).$$

The Same Data Structures

For each $t \geq 2$, build a data structure for $F^{[t]}$:

- One can send a $(t-1)$-tuple $\mathbf{x} \in D^{t-1}$ to the data structure. If $F^{[t]}(\mathbf{x}, *) = \mathbf{0}$, return $\mathbf{v} = \mathbf{0}$; otherwise, return a nonzero vector $\mathbf{v}$ linearly dependent with $F^{[t]}(\mathbf{x}, *)$, in P-time.

Then a similar main counting algorithm can compute efficiently

$$\sum_{\mathbf{x} \in D^n} F(\mathbf{x}).$$

The First Difficulty

1. An $d^{t-1} \times d$ matrix may have d^{t-1} pairwise linearly independent rows. Cannot even afford to store this many representative vectors.

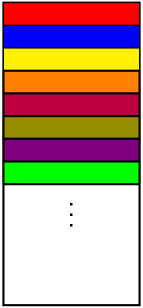

A Hint from Counting Graph Homomorphisms

Real matrices [Goldberg, Grohe, Jerrum and Thurley 09] and complex matrices [Cai, C and Lu 11]

$$\begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & 1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & \zeta & \zeta^2 & \zeta^3 & \zeta^4 \\ 1 & \zeta^2 & \zeta^4 & \zeta & \zeta^3 \\ 1 & \zeta^3 & \zeta & \zeta^4 & \zeta^2 \\ 1 & \zeta^4 & \zeta^3 & \zeta^2 & \zeta^1 \end{pmatrix}$$

Wishful thinking: What if any two rows of $F^{[t]}$ are either linearly dependent or orthogonal $\Rightarrow$ At most d representative vectors.

The Block Orthogonality condition

Let F be a function defined by a $\#\text{CSP}(\mathcal{L})$ instance. Then every two rows of $F^{[t]}$ are either linearly dependent or orthogonal.

Lemma

If $\mathcal{L}$ violates this condition, then $\#CSP(\mathcal{L})$ is $\#P$-hard.

The Second Difficulty

2. Let $\{\mathbf{v}_i\}$ be the representative vectors of $F^{[t]}$. Let S_j denote the set of rows $\mathbf{u} \in D^{t-1}$ that are linearly dependent with $\mathbf{v}_i$. Given a query $\mathbf{u} \in D^{t-1}$, how to decide if $\mathbf{u} \in S_i$ or not?

A witness function for R no longer helps!

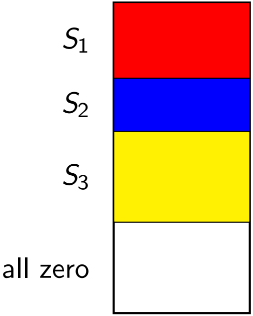

Wishful thinking: Every $S_j \subseteq D^{t-1}$ has a Mal'tsev polymorphism.
If so, one can hope to build a witness function for each S_j.

The Mal'tsev condition

Let F be a function defined by a $\#\mathrm{CSP}(\mathcal{L})$ instance. Then all such sets $S_j \subseteq D^{t-1}$ share a Mal'tsev polymorphism ψ.

Lemma

If $\mathcal{L}$ violates this condition, then $\#\mathrm{CSP}(\mathcal{L})$ is #P-hard.

The Last Difficulty

3. How to build the data structures for $F^{[t]}$ inductively? Need to compute the representative vectors (at most d many) and to compute a witness function ω_j for each $S_j \subseteq D^{t-1}$.

Type Partition condition

Manipulate relations that share a Mal'tsev polymorphism.

Lemma

If $\mathcal{L}$ violates the Type Partition condition, $\#CSP(\mathcal{L})$ is $\#P$-hard.

Putting the Pieces Together . . .

For $t = n, \ldots, 2$, build inductively a data structure for $F^{[t]}$:

1. Compute the representative vectors. Number of representative vectors is at most d: the Block Orthogonality condition.
2. Compute a witness function ω_j of $S_j \subseteq D^{t-1}$ for each representative vector $\mathbf{v}_j$. Existence: the Mal'tsev condition.

Computation of these objects uses the Type Partition condition.

Theorem

If $\mathcal{L}$ satisfies all these three conditions, $\#CSP(\mathcal{L})$ can be solved in polynomial time; otherwise, $\#CSP(\mathcal{L})$ is $\#P$-hard.

Open Problem

1. Determine the decidability of the tractability criterion:
   - Given a finite set of complex-valued functions $\mathcal{L}$, decide whether $\mathcal{L}$ satisfies the tractability criterion.
   - Counting graph homomorphisms: in P.
   - Dichotomy for nonnegative #CSP: in NP.
   - Dichotomy for complex #CSP: decidable?

   Jin-Yi's take: The land is logically conquered, but one does not really know what treasures lie within.

2. Possibility to apply the ideas elsewhere?

Thanks!