On Constant-Factor Approximable Valued Constraint Satisfaction Problems

Andrei Krokhin
Durham University, UK

Joint work with
Víctor Dalmau (University Pompeu Fabra, Barcelona)
Rajsekar Manokaran (IIT Madras / KTH Stockholm)

Constraint Satisfaction Problems (CSPs)

- CSP(Γ): given $R_1(\mathbf{x}_1), \ldots, R_q(\mathbf{x}_q)$ over V, all $R_i \in \Gamma$, is there $\varphi : V \to A$ satisfying all constraints?
  - Example: CSP($\{\neq_2\}$) is 2-Colourability
- Max CSP(Γ): maximise $\sum_{i=1}^{q} w_i \cdot R_i(\mathbf{x}_i)$
  - Example: Max CSP($\{\neq_2\}$) is Max Cut
- Min CSP(Γ): minimise $\sum_{i=1}^{q} w_i \cdot (1 - R_i(\mathbf{x}_i))$
  - Example: Min CSP($\{\neq_2\}$) is MinUnCut
- complexity classification for finding optimal solutions for Min CSP is known [Thapper, Živný'12]
- In this talk: finding approximately optimal solutions

(Min/Max) CSP Instance Example

$V = \{x, y, z\}$, $A = \{a, b\}$, $\mathcal{C} = \{x \neq y, y \neq z, x \neq z\}$.

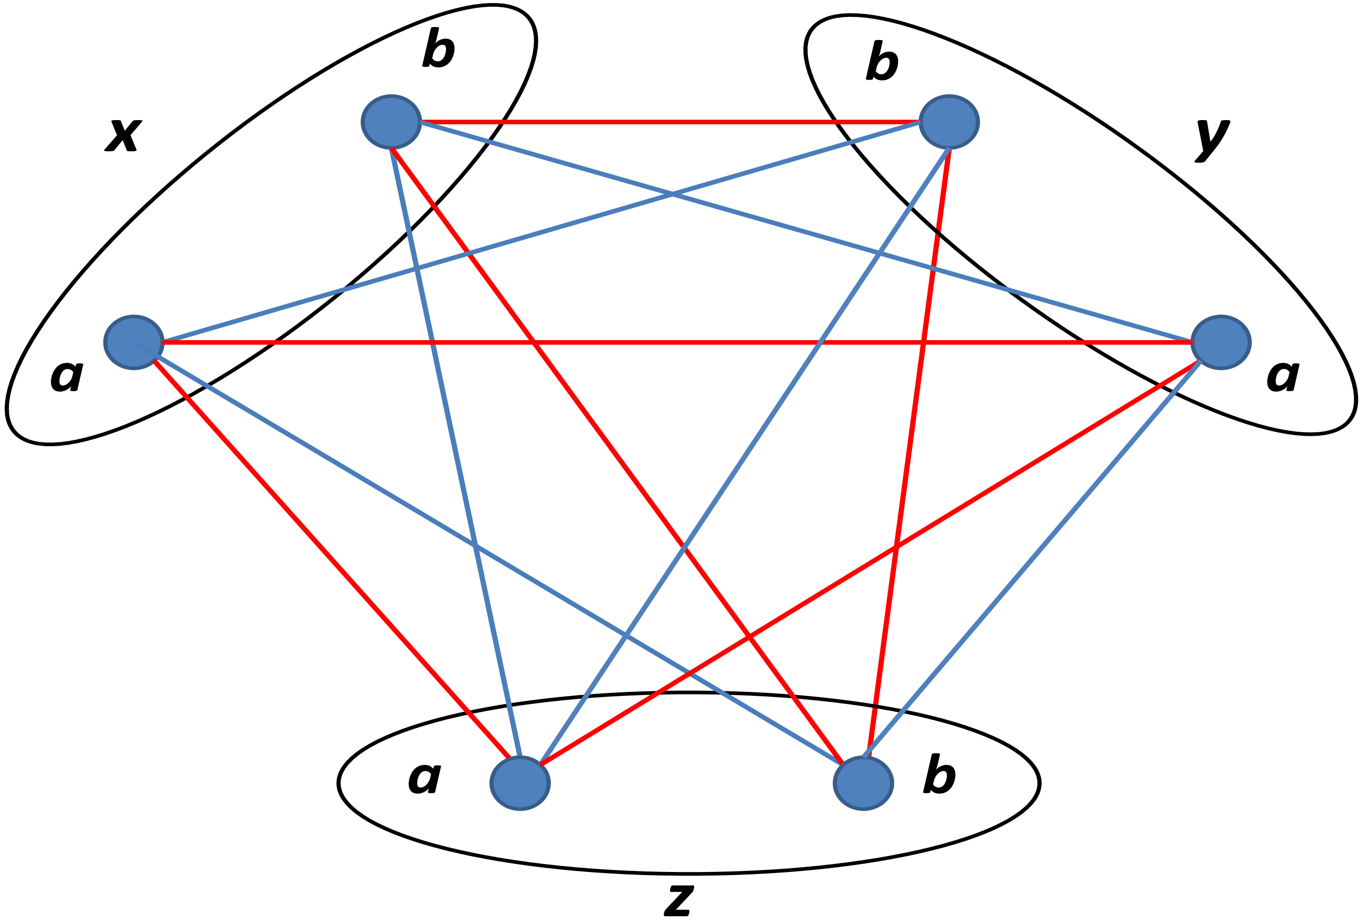

Min/Max CSP Solution Example

$V = \{x, y, z\}$, $A = \{a, b\}$, $\mathcal{C} = \{x \neq y, y \neq z, x \neq z\}$.

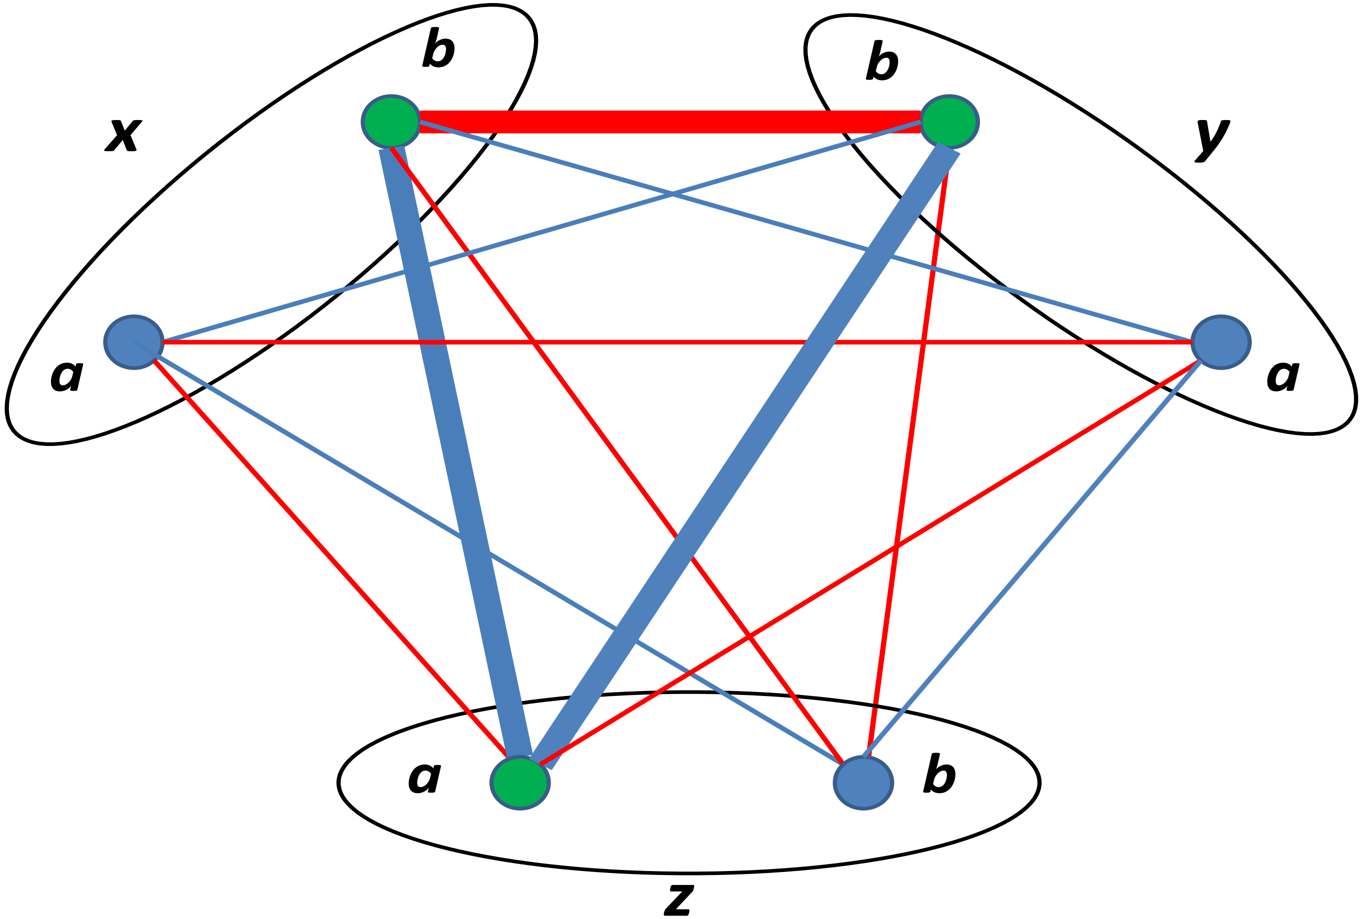

Approximation algorithms for Max CSP(Γ)

Definition 1 *Call* ALG *a c-approximation algorithm for* Max CSP(Γ) *if it runs in poly-time in* $|I|$ *and for each* I*, it finds a solution with value* ALGVal(I) *such that*

$$\mathrm{OPT}(I) \leq c(|I|) \cdot \mathrm{ALGVal}(I).$$

Fact 1 *Each* Max CSP(Γ) *belongs to* **APX***, i.e. has a c-approximation algorithm with constant c.*

- The algorithm assigns values uniformly at random.
- Can be derandomized by a standard method.
- Much research into locating optimal c.

Approximation Algorithms for Min CSP(Γ)

Definition 2 *Call* ALG *a c-approximation algorithm for* VCSP(Γ) *if it runs in poly-time in* $|I|$, *and for each* I, *it finds a solution with value* ALGVal(I) *such that*

$$\mathrm{ALGVal}(I) \leq c(|I|) \cdot \mathrm{OPT}(I).$$

Fact 2 *c-approx algo for* Min CSP(Γ) $\Rightarrow$ CSP(Γ) $\in$ **P**.

Problem 1 *Which problems* Min CSP(Γ) *belong to complexity class* **APX***?*

- Long-standing open problem: is MinUnCut there?
- Currently best answer: no, unless the UGC fails.

Some Known Results

k-HORN clauses: (x), $(\overline{x}_1 \vee \ldots \vee \overline{x}_{\leq k})$, $(x_1 \vee \overline{x}_2 \vee \ldots \vee \overline{x}_{\leq k})$.

k-IHBS clauses: (x), $(\overline{x}_1 \vee \ldots \overline{x}_{\leq k})$, $(x_1 \vee \overline{x}_2)$.

- Min CSP(k – IHBS) is in **APX** [Khanna et al'01]
- Min CSP(3 – HORN) is **NP**-hard to constant-factor approximate [Guruswami, Lee'14]
- MinUnCut has $O(\sqrt{\log n})$-approximation algorithm [Agarwal et al'06]
- MinUnCut is not in **APX** unless the UGC fails [Khot et al'07]
- Detailed classification for $A = \{0, 1\}$ [Khanna et al'01]

Algebra Works

MIN CSP(Γ) in **APX** - studied in (Dalmau, AK'13) as "CSP(Γ) that are robustly tractable with linear loss"

- One class of problems MIN CSP(Γ) in **APX** is found.
- Standard algebraic machinery works when $\Gamma \supseteq \{=\}$.
  - polymorphisms, algebras, idempotence, varieties
- Which algebraic properties lead to **APX**?

Fractional Solution Example

$V = \{x, y, z\}$, $D = \{a, b\}$, $\mathcal{C} = \{x \neq y, y \neq z, x \neq z\}$.

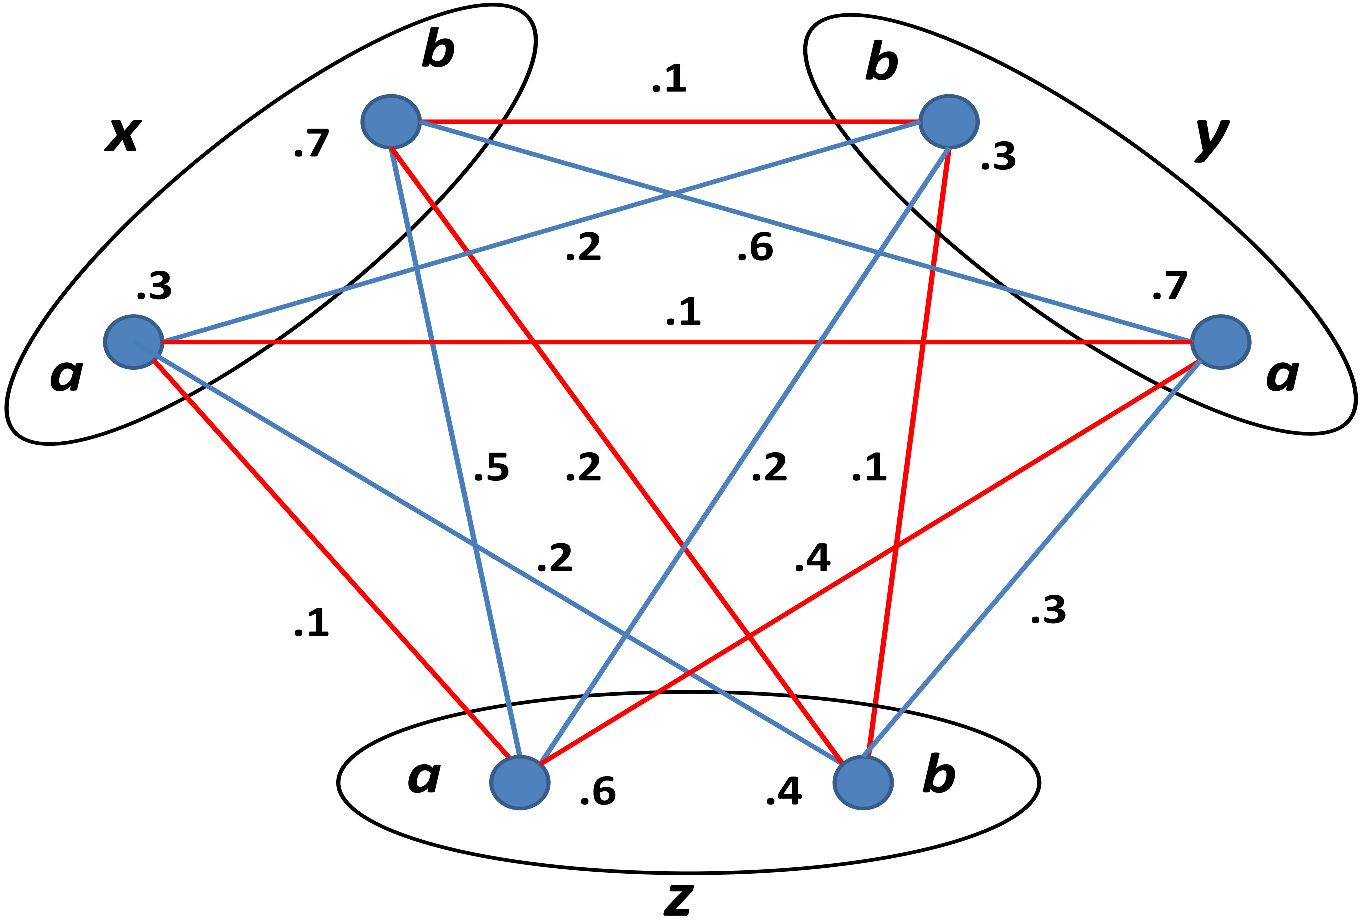

Consistent Marginals Example

$V = \{x, y, z\}$, $D = \{a, b\}$, $\mathcal{C} = \{x \neq y, y \neq z, x \neq z\}$.

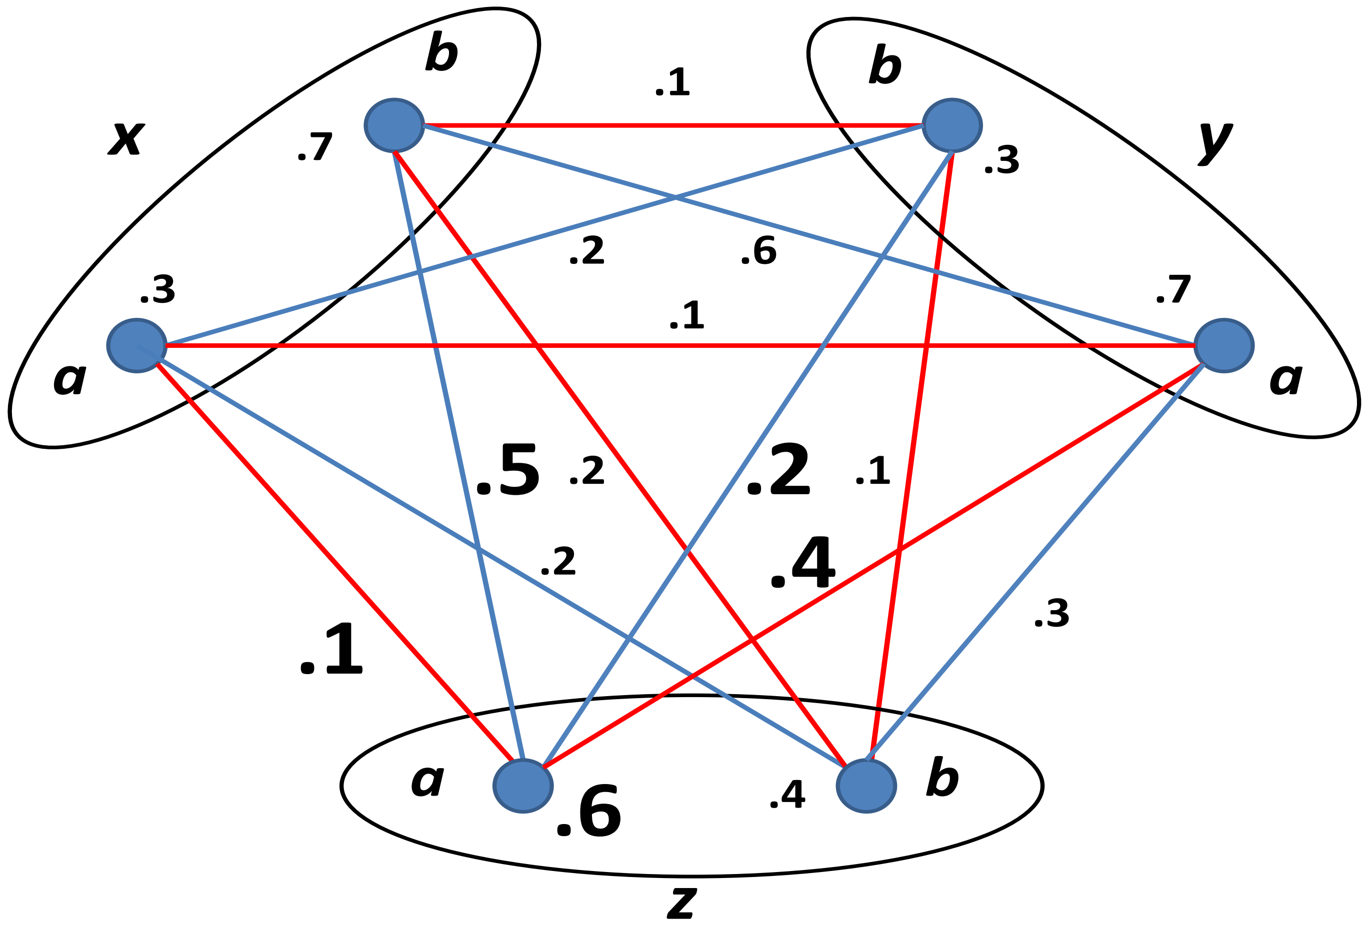

Marginal Distributions Example

$V = \{x, y, z\}$, $D = \{a, b\}$, $\mathcal{C} = \{x \neq y, y \neq z, x \neq z\}$.

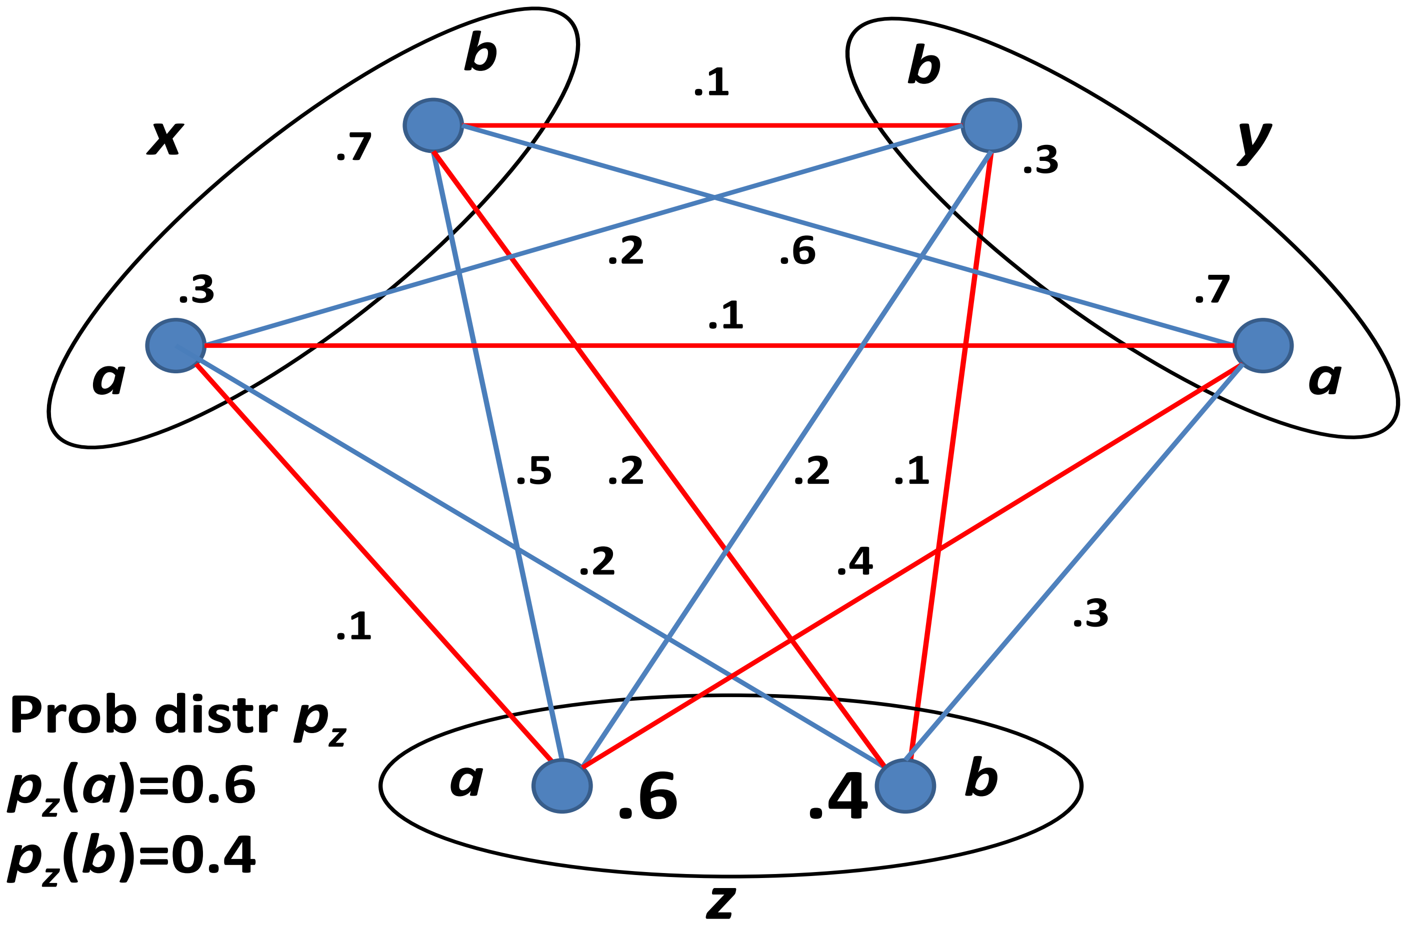

Basic LP Relaxation for Min CSP(Γ)

The basic LP relaxation for instance I with constraints $\mathcal{C}$.
The variables are

- $p_v(a) \in [0,1]$ for each $v \in V, a \in A$;
- $p_C(\mathbf{t}) \in [0,1]$ for each constraint C in I and $\mathbf{t} \in A^{ar(C)}$.

$$\text{minimize} \quad \sum_{C=(\mathbf{x},R)\in\mathcal{C}} w_C \cdot \sum_{R(\mathbf{t})=0} p_C(\mathbf{t}) \quad \text{subject to:}$$

- p_v, p_C - probability distributions for all $v \in V, C \in \mathcal{C}$
- consistent marginals

Since Γ is fixed, this relaxation has polynomial size (in I).

Optimality of BLP

Rounding: converting fractional solution to proper solution

The integrality gap of BLP for Min CSP(Γ) is

$$\alpha = \sup_{\text{instance } I} \frac{\mathrm{OPT}(I)}{\mathrm{BLPVal}(I)}$$

Meaning: α is best poss approx factor from rounding BLP.

Theorem 1 (Ene, Vondrak, Wu'13)
For any $\Gamma \supseteq \{=\}$, *if* Min CSP(Γ) *has a c-factor approx algorithm with* $c < \alpha$ *then the UGC fails. In particular, if* $\alpha = \infty$ *then* Min CSP(Γ) $\notin$ **APX** *(unless the UGC fails).*

Meaning: enough to consider BLP-based approx algorithms

The Standard Simplex

Let $\Delta(X)$ = {probability distributions on a set X}.

The standard $(k-1)$-dimensional simplex where $k = |X|$

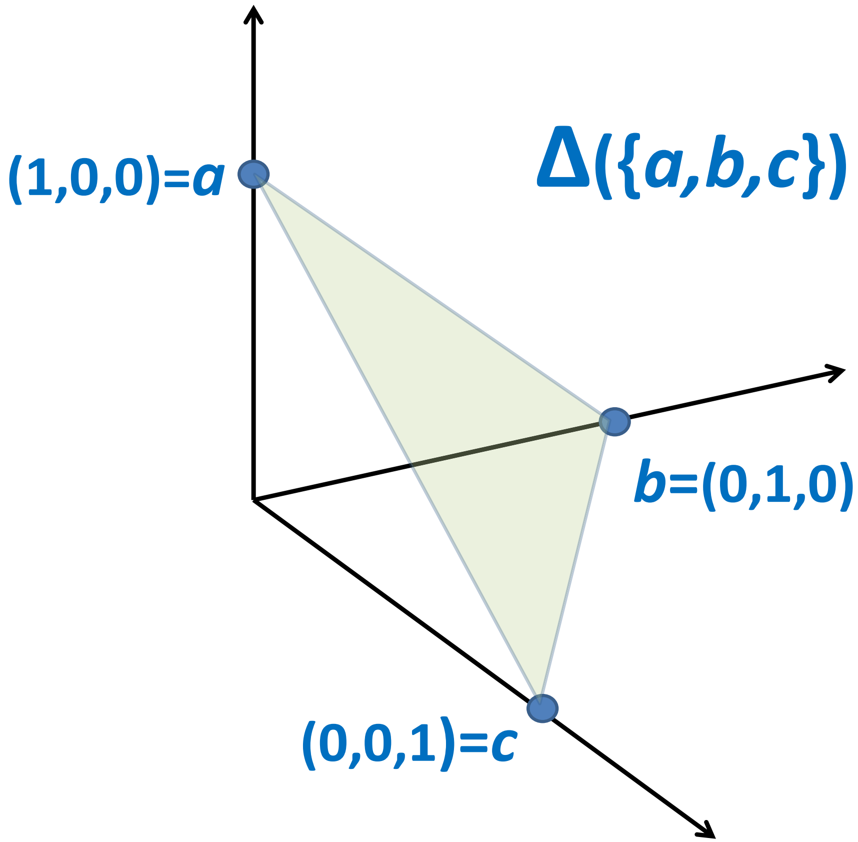

Simplex Discretized

Let $\Delta_n(X) = \{p \in \Delta(X) \mid \forall x \in X \;\; p(x) \in n^{-1}\mathbb{Z}\}$.

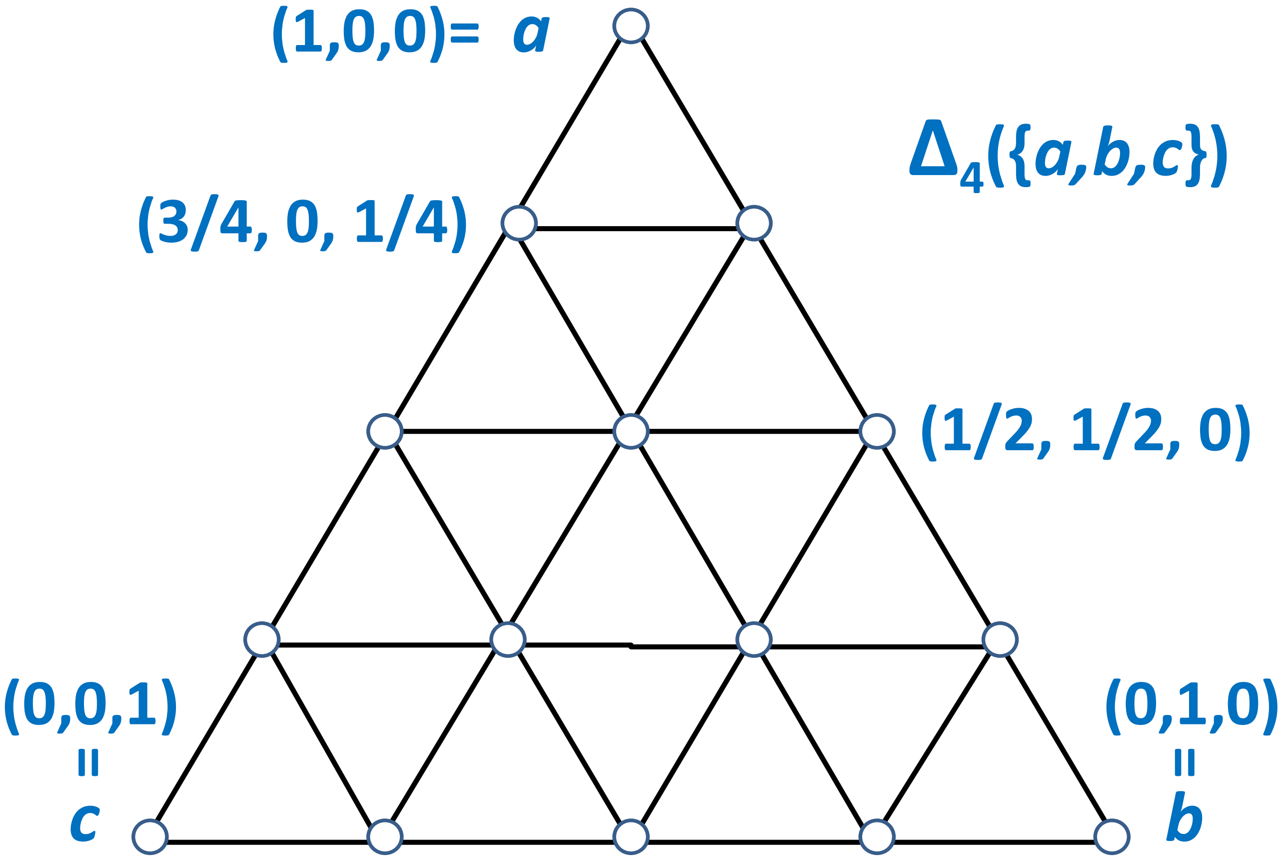

Rounding BLP Solution

- Let s be an optimal solution for BLP(I). Can assume there is n such that s gives $p_v \in \Delta_n(A)$ for each $v \in V$.
- Any map $g : \Delta_n(A) \to A$ can be used to round s for I; as follows: $v \mapsto g(p_v)$. Good $g \Rightarrow$ good approximation.
- $\Delta_n(A) \quad \leftrightarrow \quad$ multisets on A of size n
  - $p \in \Delta_n(A) \leftrightarrow$ [$a \in A$ appears $p(a) \cdot n$ times]
- An operation $f : A^n \to A$ is symmetric if, $\forall \pi \in S_n$,

$$f(x_1, \ldots, x_n) = f(x_{\pi(1)}, \ldots, x_{\pi(n)}).$$

- n-ary symmetric operations $\equiv$ mappings $\Delta_n(A) \to A$.

Symmetric Operation Example

This is a 4-ary (idempotent) symmetric operation f

For example, $f(a, c, a, a) = a$ and $f(b, b, a, a) = a$

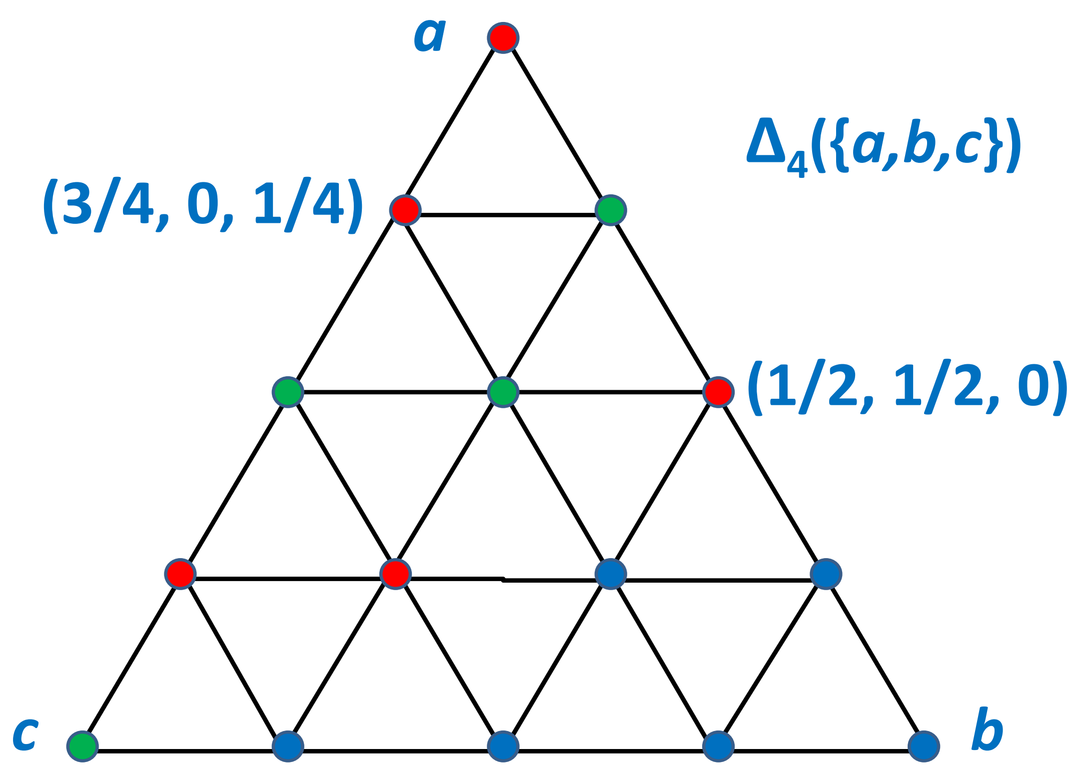

Deciding CSP(Γ) by BLP

Theorem 2 (Kun et al '11) *For any* Γ, *TFAE*

1. *BLP decides* $\mathrm{CSP}(\Gamma)$, *i.e.* $\mathrm{BLPVal}(I) = 0 \Rightarrow I$ *is sat.*
2. *For each* n, Γ *has an* n*-ary symmetric polymorphism.*

Let I be an instance of $\mathrm{CSP}(\Gamma)$ with $\mathrm{BLPVal}(I) = 0$ and let s be an optimal solution to $\mathrm{BLP}(I)$. Can assume $\exists n$

- s gives $p_v \in \Delta_n(A)$ for each $v \in V$.
- s gives $p_C \in \Delta_n(A^{ar(C)})$ for each $C \in \mathcal{C}$.

If $g \in \mathrm{SymPol}_n(\Gamma)$ then $v \mapsto g(p_v)$ satisfies all constraints.

Proof of Satisfaction

- Pick a constraint $C = R(\mathbf{x})$. Let $\mathbf{x} = (v_1, \ldots, v_m)$.
- Know $p_C(\mathbf{t}) > 0 \Rightarrow R(\mathbf{t}) = 1$. Recall: $p_C \in \Delta_n(A^m)$.
- Take $n \cdot p_C(\mathbf{t})$ copies of each tuple $\mathbf{t}$ with $R(\mathbf{t}) = 1$.
- Call them $\mathbf{a}_1 = (a_{11}, \ldots, a_{1m}), \ldots, \mathbf{a}_n = (a_{n1}, \ldots, a_{nm})$.

$$\begin{array}{cccccccc} & g & & g & & g & & \\ R(& a_{11} & , & \ldots & , & a_{1m} &) & = 1 \\ & \vdots & & \vdots & & \vdots & & \vdots \\ R(& a_{n1} & , & \ldots & , & a_{nm} &) & = 1 \\ \hline R(& g(p_{v_1}) & , & \ldots & , & g(p_{v_m}) &) & = 1 \end{array}$$

Stability and Integrality Gap

For $d_1, d_2 \in \Delta_n(A)$, let $\mathrm{dist}(d_1, d_2) = \max_{a \in A} |d_1(a) - d_2(a)|$

Let ϕ be a probability distribution on $\mathrm{SymPol}_n(\Gamma)$.

Say that ϕ is c-stable if, for all $d_1, d_2 \in \Delta_n(A)$,

$$\Pr_{g \sim \phi}\{g(d_1) \neq g(d_2)\} \leq c \cdot \mathrm{dist}(d_1, d_2).$$

Theorem 3 (Dalmau, AK, Manokaran)
For any $\Gamma \supseteq \{=\}$, TFAE

1. *BLP has finite integrality gap for* MIN CSP(Γ).
2. *There is $c \geq 1$ such that, for all n, Γ admits a c-stable probability distribution on* $\mathrm{SymPol}_n(\Gamma)$.

Fractional Symmetric Operation Example

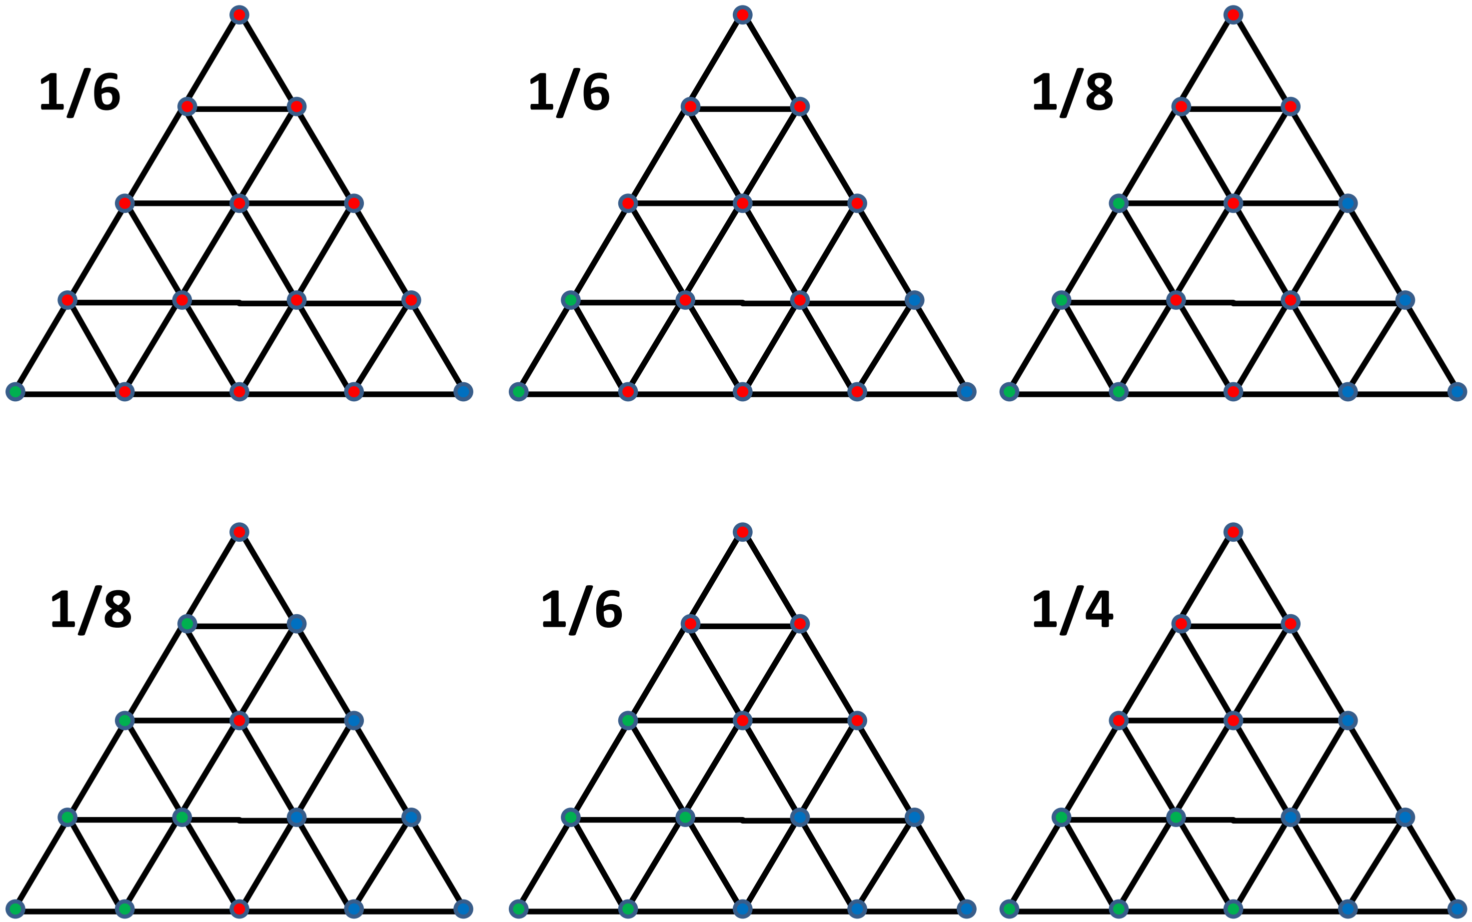

Examples

- Non-example: Take $3-\mathrm{HORN}$.
  - Only one n-ary symmetric polym $g(\mathbf{x}) = \bigwedge x_i$.
  - Take $d_1 = (1, 1, \ldots, 1)$ and $d_2 = (0, 1, \ldots, 1)$
  - Easy: $\mathrm{dist}(d_1, d_2) = 1/n$, but $\Pr[g(d_1) \neq g(d_2)] = 1$.
  - Hence infinite integrality gap and UG-hardness
- Example: Take $\Gamma = \{\leq, 0, 1\}$ on $A = \{0, 1\}$.
  - For $1 \leq j \leq n$, let $g_{n,j}(\mathbf{x}) = 1$ iff $|\{x_i : x_i = 1\}| \geq j$.
  - Each $g_{n,j}$ is monotone, so polymorphism
  - If $\mathrm{dist}(d_1, d_2) = r/n$ then $\Pr[g(d_1) \neq g(d_2)] \leq r/n$.
  - Hence 1-stability and finite integrality gap.

Rounding from Stable Distributions

- Let I be an instance of CSP(Γ), take an optimal solution to BLP(I), obtain $p_v \in \Delta_n(A)$ for each $v \in V$ and $p_C \in \Delta_n(A^{ar(C)})$ for each $C \in \mathcal{C}$.
- Draw g from the c-stable distribution ϕ_n; $v \mapsto g(p_v)$.
- This is a randomized $(2 \cdot maxar \cdot c)$-approx algorithm.
- Pick $C = R(\mathbf{x})$ and estimate $\Pr_{g\sim\phi_n}\{R(g(\mathbf{x})) = 0\}$.
- Modify p_C to q_C such that $q_C(\mathbf{t}) > 0 \Rightarrow R(\mathbf{t}) = 1$.
- For marginals q_i's of q_C, have $R(g(q_1), \ldots, g(q_m)) = 1$.
- Marginals of p_C and q_C are close, use c-stability of ϕ_n.
- Get $\Pr_{g\sim\phi_n}\{R(g(\mathbf{x})) = 0\} \leq 2 \cdot m \cdot c \cdot (1 - p_C(R))$.

A Positive Result

Theorem 4 (Dalmau,AK'13)
Assume that Γ *is hom-equivalent to a CL* Γ' *on some set* L *(of subsets) s.t.* Γ' *has polymorphism* $x \cap (y \cup z)$ *where* $(L, \cap, \cup)$ *is a distrib lattice. Then* Min CSP(Γ) $\in$ **APX**.

- There are other problems Min CSP(Γ) in **APX**.

NP-hardness Result

- Let Γ have c-stable distributions ϕ_n on its symmetric polymorphisms g. Wlog assume $\forall x\ g(x,\ldots,x)=x$.
- For n-tuples $d_1=(b,a,\ldots,a)$ and $d_2=(a,\ldots,a)$, have

$$\Pr_{g\sim\phi_n}\{g(d_1)\neq a\}=\Pr_{g\sim\phi_n}\{g(d_1)\neq g(d_2)\}\leq c{\cdot}\mathrm{dist}(d_1,d_2)=\frac{c}{n}.$$

- So, for $n>c\cdot|A|^2$, $supp(\phi_n)$ contains NU operations:

$$\forall x,y\ \forall i\ f(x,\ldots,x,\underset{i}{y},x,\ldots,x)=x.$$

Theorem 5 (Dalmau, AK, Manokaran)
If Γ has no NU polymorphism then it is **NP***-hard to constant-factor approximate* Min CSP(Γ).

From VCSP to Min CSP

- Valued constraint: $f(\mathbf{x})$ where $f : A^m \to [0, 1]$
- VCSP(Γ): minimise $\sum_{i=1}^{q} w_i \cdot f_i(\mathbf{x}_i)$ where all $f_i \in \Gamma$
- Min CSP is a special case of VCSP

Lemma 1 (Dalmau, AK, Manokaran)
For each valued CL Γ, there is a (non-valued) CL Γ' such that VCSP(Γ) *is in* **APX** *iff* Min CSP(Γ') *is in* **APX**.

Open Problems

- Use c-stability to get an efficient rounding algorithm.
- Improve more UG-hardness to NP-hardness.
- Get rid of the $\{=\} \subseteq \Gamma$ assumption (if possible).
- Study algebras with many symmetric operations.
- Decidability issues for symmetric polymorphisms.
- Link c-stability with Prague-like strategies.
- Extend results to non-constant c and/or SDP.

The Unique Games Conjecture (UGC)

For a permutation σ on A, let $\sigma^\circ = \{(x,y) \mid y = \sigma(x)\}$.
For $A = \{0,1,\ldots,k-1\}$, let $\Gamma_k = \{\sigma^\circ \mid \sigma \text{ a perm on } A\}$.

Conjecture 1 (Khot'02)
For each $\epsilon > 0$, there is $k = k(\epsilon)$ such that it is **NP***-hard to tell $(1-\epsilon)$-satisfiable from at most ϵ-satisfiable instances of* MAX CSP(Γ_k) *(aka* UNIQUE GAMES*).*

- One of the hottest conjectures in Theoretical CS
- If true, optimal approx algorithms for many classical problems, incl. all MAX CSP(Γ) [Raghavendra'08].
- If false, there is a new powerful approx technique