Locally consistent constraint satisfaction problems

Zdeněk Dvořák (KAM and ITI, Prague)
Daniel Král' (KAM and ITI, Prague and TU Berlin)
Ondřej Pangrác (KAM and ITI, Prague)

Constraint satisfaction problems

- constraint satisfaction problem (CSP)
  universe U
  constraint types $C_1 \subseteq U^{k_1}$, $C_2 \subseteq U^{k_2}$, $\ldots$
- instance
  set X of variables
  constraints: k_i-tuples of variables for the i-th type of constraint
- goal
  assignment $\varphi : X \to U$ such that each constraint is satisfied
- CSP includes satisfiability, graph 3-coloring, etc.
  in particular, the problem is NP-complete in general
- graph 3-coloring formulated as CSP:
  $U = \{$● ● ●$\}$, $X = V(G)$ and constraints correspond to edges

The problem and our results

- in practice, it is easy to identify small sets of inconsistent constraints
  3-coloring — check all subgraphs on at most 6 vertices
- CSP instance is k-consistent if any k constraints can be satisfied
- What can be said about k-consistent instances?

  bad news: definitely, not all constrained can be satisfied

  good news: sometimes a substantial fraction of them can be satisfied

- we focus on the case when the constraints are Boolean predicates
  design an assymptotically optimal approximation algorithm
  analuze assymptotical behaviour as k tends to infinity
  analyze completely predicates of small arity

Notation

- Π is a (finite) set of Boolean predicates
  $\Pi_{\mathrm{SAT}} = \{(x_1), (x_1 \vee x_2), (x_1 \vee x_2 \vee x_3), \ldots\}$
  $\Pi_{2-\mathrm{SAT}} = \{(x_1), (x_1 \vee x_2)\}$
- Σ is a finite set of predicates from Π
  $\Sigma = \{(x), (\neg x \vee y), (\neg y \vee z), (\neg z)\}$ for $\Pi = \Pi_{2-\mathrm{SAT}}$
- $\rho(\Sigma)$ denotes the fraction of predicates which can be satisfied
  $\rho(\Sigma) = 3/4$
- $\rho_k(\Pi) = \inf \rho(\Sigma)$ where Σ is a k-consistent system of predicates from Π
  $\rho_1(\Pi_{\mathrm{SAT}}) = \rho_1(\Pi_{2-\mathrm{SAT}}) = 1/2$
- $\rho_\infty(\Pi) = \lim_{k \to \infty} \rho_k(\Pi)$
- weighted variants $\rho_k^w(\Pi)$ and $\rho_\infty^w(\Pi)$

SAT problems

- $\rho_2^w(\Pi_{\mathrm{SAT}}) = \rho_2^w(\Pi_{2-\mathrm{SAT}}) = \frac{\sqrt{5}-1}{2} \approx 0.6180$ [Lieberherr, Specker 1981]
- $\rho_3^w(\Pi_{\mathrm{SAT}}) = \rho_3^w(\Pi_{2-\mathrm{SAT}}) = 2/3$ [Lieberherr, Specker 1982]
  a simple proof found by Yannakakis in 1992
- $\rho_4^w(\Pi_{\mathrm{SAT}}) \approx 0.6992$ but $\rho_4^w(\Pi_{2-\mathrm{SAT}}) > 0.6992$ [K. 2003]
- $\rho_\infty^w(\Pi_{\mathrm{SAT}}) \leq 3/4$ [Huang, Lieberherr 1985]
- $\rho_\infty^w(\Pi_{\mathrm{SAT}}) = \rho_\infty^w(\Pi_{2-\mathrm{SAT}}) = 3/4$ [Trevisan 1997]
- Trevisan also addressed the case when Π_k is a set of all k-ary Boolean predicates and showed: $\rho_\infty^w(\Pi_k) = \rho_\infty(\Pi_k) = 2^{1-k}$

1-EXTENDABLE PREDICATES

- a Boolean predicate P is 1-extendable if the following holds:
  if one of its arguments is fixed, the remaining ones can be choosen
  in such a way that the predicate is satisfied.
  example: $P_1(x, y, z) = (x \vee y \vee z)$ vs. $P_2(x, y, z) = x \wedge (y \Leftrightarrow z)$
- $\sigma(P)$ is the fraction of satisfiing assignments of P
  example: $\sigma(P_1) = 7/8$
- $\rho^w_\infty(\{P\}) = \rho_\infty(\{P\}) = \sigma(P)$ if the arity of P is at least two
- upper bound: take a random assignment
- lower bound:
  take a random hypergraph of high girth and associate predicates with its edges
  random $\Longrightarrow$ you cannot do better than in random assigment
  high girth $\Longrightarrow$ you do not have small sets of inconsistent predicates

Predicates of arity at most three

$\sigma(\{P\})$	P	$k=1$	$k=2$	$k=3$	$k=4$	$k=5$	$k\geq 6$	$k\to\infty$
1	$x \wedge y$	1/4	1	1	1	1	1	1
2	$x \Leftrightarrow y$	1/2	1/2	1/2	1/2	1/2	1/2	1/2
3	$x \vee y$	3/4	3/4	3/4	3/4	3/4	3/4	3/4
1	$x \wedge y \wedge z$	1/8	1	1	1	1	1	1
2	$x \Leftrightarrow y \Leftrightarrow z$	1/4	1/4	1/4	1/4	1/4	1/4	1/4
	$x \wedge (y \Leftrightarrow z)$	1/4	8/27	1/2	1/2	1/2	1/2	1/2
3	exactly one	3/8	3/8	3/8	3/8	3/8	3/8	3/8
	$x \wedge (y \vee z)$	3/8	$\frac{2\sqrt{3}}{9}$	1/2	$\frac{\sqrt{5}-1}{2}$	2/3	$\rho_{l-2}(\Pi_{2-\text{SAT}})$	3/4
	$(x \Leftrightarrow y) \wedge (x \Rightarrow z)$	3/8	3/8	3/8	3/8	3/8	3/8	3/8
4	$x \Rightarrow y \Rightarrow z$	1/2	1/2	1/2	1/2	1/2	1/2	1/2
	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
	$(x \wedge y) \Leftrightarrow z$	1/2	1/2	1/2	1/2	1/2	1/2	1/2
5	$x \vee (y \wedge z)$	5/8	5/8	5/8	5/8	5/8	5/8	5/8
	$\ldots$	$\ldots$	$\ldots$	$\ldots$	$\ldots$	$\ldots$	$\ldots$	$\ldots$
6	$\neg(x \Leftrightarrow y \Leftrightarrow z)$	3/4	3/4	3/4	3/4	3/4	3/4	3/4
	$x \vee (y \Leftrightarrow z)$	3/4	3/4	3/4	3/4	3/4	3/4	3/4
7	$x \vee y \vee z$	7/8	7/8	7/8	7/8	7/8	7/8	7/8

General case

- a finite set Π, want to determine $\rho^w_\infty(\Pi)$ and $\rho_\infty(\Pi)$
- a restriction P' of a predicate P is obtained by fixing some of the arguments
  restriction can be described by a vector $\tau \in \{0, 1, \star\}^r$
  $P'(x, y) = P(x, y, 0) = x \wedge y$ where $P(x, y, z) = x \wedge (y \vee z)$ and $\tau = \star\star 0$
- $\pi_{P,\tau}(p) : \langle 0, 1\rangle \to \langle 0, 1\rangle$ is the probability that P is satisfied if x_i is true with the probability $1 - p$, p and $1/2$, if τ_i is 0, 1 and $\star$, respectively.
  $\pi_{P,\tau}(p)$ is a polynomial in p
- $\pi(\Pi)$ is the set of all $\pi_{P,\tau}$ for $P \in \Pi$ where the restriction of P is 1-extendable
- $\overline{\pi(\Pi)}$ is the set convex hull of $\pi(\Pi)$

$$\rho^w_\infty(\Pi) = \min_{f \in \overline{\pi(\Pi)}} \max_{p \in [0,1]} f(p)$$

The same holds for $\rho_\infty(\Pi)$ if Π does not contain a predicate of arity one.

Example 1

- $\Pi = \{P_1, P_2\}$ where $P_1(x_1) = (x_1)$ and $P_2(x_1, x_2) = (x_1 \vee x_2)$
- a single restriction of P_1 is 1-extendable
  four restriction of P_2 are 1-extendable: 11, $1\star$, $\star 1$ and $\star\star$
- $\pi_{P_1,1}(p) = p$, $\pi_{P_2,11}(p) = 2p - p^2$, $\pi_{P_2,1\star}(p) = \pi_{P_2,\star 1}(p) = (p+1)/2$,
  $\pi_{P_2,\star\star}(p) = 3/4$.
- $\pi(\Pi) = \{p, 2p - p^2, (p+1)/2, 3/4\}$
- $\rho^w_\infty(\Pi) = \min_{f \in \overline{\pi(\Pi)}} \max_{p \in [0,1]} f(p)$
  each function in $\pi(\Pi)$ is at least $3/4$ for $p = 3/4$
  $\rho^w_\infty(\Pi) = 3/4$ [Trevisan 1997]

Example 2

- $\Pi = \{P\}$ where $P(x_1, x_2, x_3, x_4, x_5) = (x_1 \wedge (x_2 \vee x_3 \vee x_4 \vee x_5))$
- several 1-extendable restrictions, but each is symmetric to one of the following: $1\star\star\star\star$, $10\star\star\star$, $11\star\star\star$, $100\star\star$, $110\star\star$, $111\star\star$, $1100\star$, $1100\star$, $1110\star$, $1111\star$, 11000, 11100, 11110 and 11111.
- $\rho^w_\infty(\Pi) = \min_{f \in \overline{\pi(\Pi)}} \max_{p \in [0,1]} f(p)$
  $\pi_{P,100\star\star}(p) = p(1 - p^2/4)$
  $\max \pi_{P,100\star\star}(p) = 3/4$
  each function in $\pi(\Pi)$ is at least $3/4$ for $p = 3/4$
  $\rho^w_\infty(\Pi) = \rho_\infty(\Pi) = 3/4$

Sketch of the upper bound proof

- fix Π and $\varepsilon > 0$
- classify the variables according to the number of constraints forcing their values
  in the algorithm, one rather uses the depth of "forcing"
  $X_1, \ldots, X_{2R/\varepsilon}$ and Y
- this also partitions constraints to $2R/\varepsilon$ levels
- choose i such that the fraction of constraints in the i-th level is small
  it is at most $\varepsilon/2$
- fix variables from $X_1, \ldots, X_i$ to their forced values (0/1)
  keep remaining variables free ($\star$)
- this leads to a polynomial $f(p)$ from $\overline{\pi(\Pi)}$
- take a random assignment for p that maximizes $f(p)$

Algorithmic aspects

- depth of “forcing” can be easily computed in time $O(1/\varepsilon)$
  inconsitence found $\Longrightarrow$ the input is not sufficiently locally consistent
- the maximum of the polynomial $f(p)$ can be found with precision $\varepsilon/2$ in time $O(1/\varepsilon)$ (bounded derivatives)
- derandomization by the method of conditional expectations
- robust algorithm linear in the size of input and $1/\varepsilon$

Graph homomorphisms

- a single binary constraint $\Gamma \subseteq U^2$ for a large universe U
- the target oriented graph H
  $V(H) = U$ and $uv \in E(H)$ if $(u, v) \in \Gamma$
- instance can be viewed as graph G with $V(G) = X$ and
  arcs corresponding to input constraints
- k-consistent: every $G' \subseteq G$ of size at most k can be mapped to H ($G \xrightarrow{k} H$)
- find a largest subgraph $G_0 \subseteq G$ that can be mapped to H
- the corresponding ratio $\rho_k(H) = \min_{G \xrightarrow{k} H} \max_{G_0 \subseteq G, G_0 \to H} \frac{||G_0||}{||G||}$
- if H contains an oriented cycle, then $\rho_k(H)$ is the density of the densiest
  subgraph of H
  if H contains a loop, then $\rho_k(H) = 1$

Open problems

- determine $\rho_k(H)$ for acyclic oriented graph H
- determine $\rho_k(H)$ for oriented trees H
- determine $\rho_k(H)$ for consistently oriented paths H

General setting

- higher arity leads to hypergraph homomorphisms
- more constraints leads to edge-colored (hyper)graphs