Counting with Bounded Treewidth

Yitong Yin
Nanjing University

Joint work with **Chihao Zhang**
(SJTU ⟶ ?)

Spin Systems

graph $G=(V,E)$

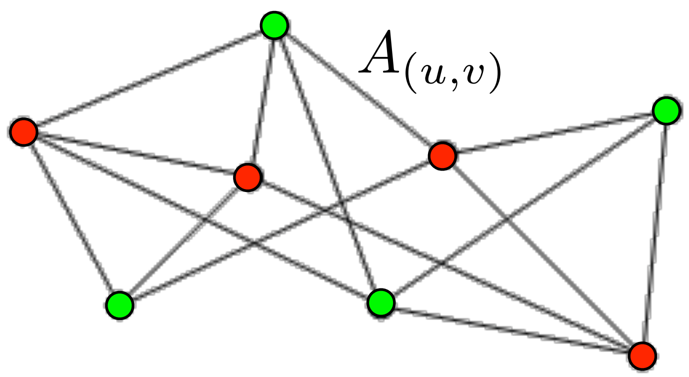

vertices are variables (domain $[q]$)

edge are constraints:

$\forall e=(u,v) \in E$

$$A_e : [q] \times [q] \to \mathbb{C}$$

configuration $\sigma \in [q]^V$

partition function:

$$Z = \sum_{\sigma \in [q]^V} \prod_{e=(u,v) \in E} A_e(\sigma_u, \sigma_v)$$

Holant Problem

graph $G=(V,E)$

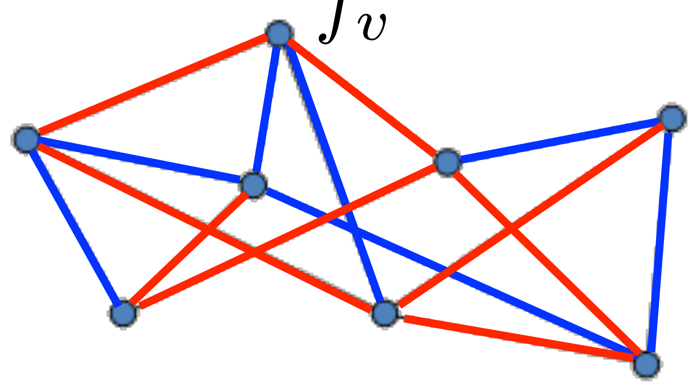

edges are variables (domain $[q]$)

vertices are constraints (*signatures*)

$$\forall v \in V, \quad f_v : [q]^{\deg(v)} \to \mathbb{C}$$

configuration $\sigma \in [q]^E$

$$\text{holant} = \sum_{\sigma \in [q]^E} \prod_{v \in V} f_v\left(\sigma\,|_{E(v)}\right)$$

$E(v) = (e_1, \ldots, e_d)$
incident edges of v

q^m configurations where $m=|E|$

graph $G=(V,E)$

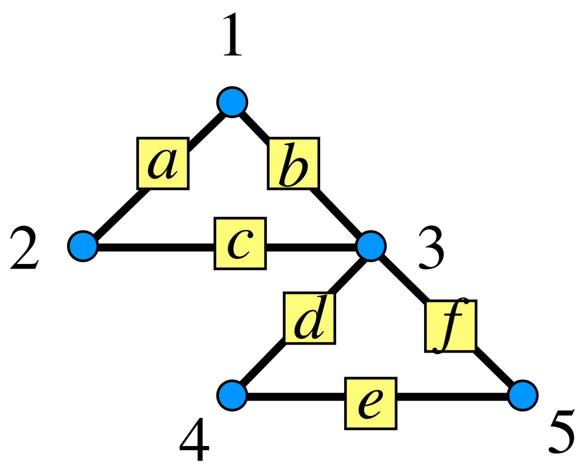

incidence graph $B(V,E,F)$

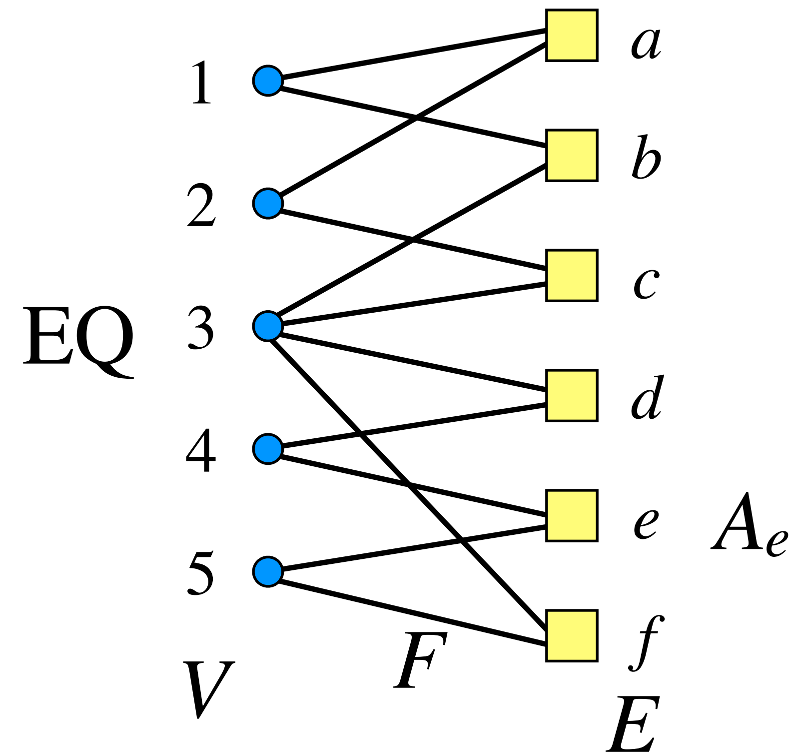

$\forall \sigma \in [q]^V$

$$w(\sigma) = \prod_{e=(u,v)\in E} A_e(\sigma_u, \sigma_v)$$

$$Z = \sum_{\sigma\in[q]^V} w(\sigma)$$

where

$$\mathrm{EQ}(x_1, \ldots, x_d) = \begin{cases} 1 & x_1 = \cdots = x_d \\ 0 & \text{otherwise} \end{cases}$$

$$\text{holant} = \sum_{\sigma\in[q]^F} \prod_{v\in V\cup E} f_v\left(\sigma\mid_{F(v)}\right)$$

graph $G=(V,E)$

Boolean symmetric $f : \{0,1\}^d \to \mathbb{C}$ ($q=2$)

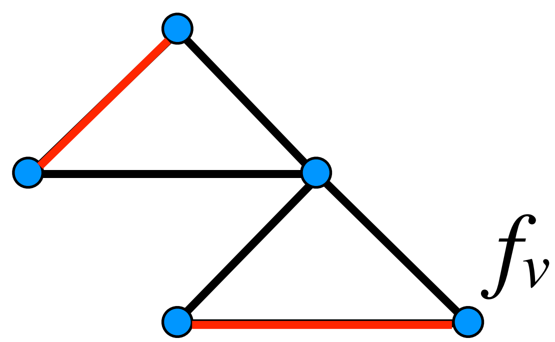

$f = [f_0, f_1, \ldots, f_d]$

where $f_i = f(x)$ for $\sum_j x_j = i$

counting matchings:

at every vertex $v \in V$,

$f_v = [1, 1, 0, 0 \ldots, 0]$

the at-most-one function

$$\text{holant} = \sum_{\sigma \in \{0,1\}^E} \prod_{v \in V} [1, 1, 0, \ldots, 0] \left(\sigma \mid_{E(v)}\right)$$

graph G=(V,E)

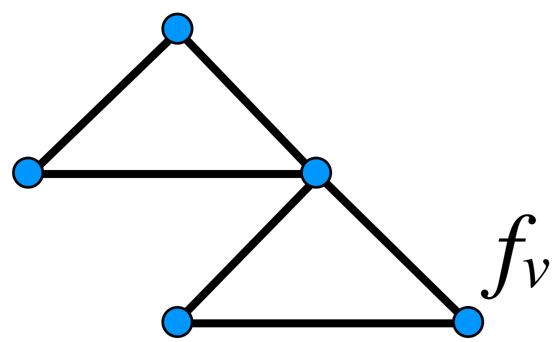

Boolean symmetric $f : \{0,1\}^d \to \mathbb{C}$ (q=2)

$f = [f_0, f_1, \ldots, f_d]$

where $f_i = f(x)$ for $\sum_j x_j = i$

counting *perfect* matchings:

at every vertex $v \in V$,

$f_v = [0, 1, 0, 0 \ldots, 0]$

the exact-one function

$$\text{holant} = \sum_{\sigma \in \{0,1\}^E} \prod_{v \in V} [0, 1, 0, 0, \ldots, 0] \left(\sigma \mid_{E(v)}\right)$$

graph $G=(V,E)$ Boolean symmetric $f:\{0,1\}^d \to \mathbb{C}$ ($q=2$)

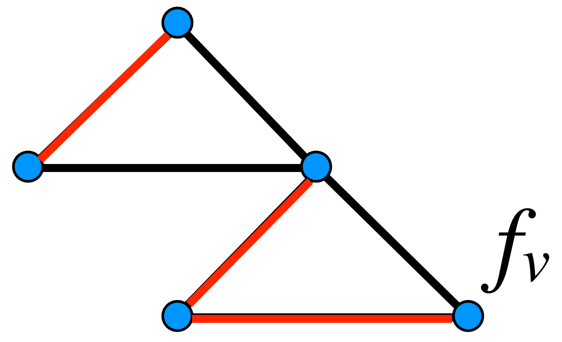

$f = [f_0, f_1, \ldots, f_d]$

where $f_i = f(x)$ for $\Sigma_j x_j = i$

counting edge covers:

at every vertex $v \in V$,

$f_v = [0, 1, 1, 1\ldots, 1]$

the at-least-one function

$$\text{holant} = \sum_{\sigma \in \{0,1\}^E} \prod_{v \in V} [0,1,1,1,\ldots,1]\left(\sigma\mid_{E(v)}\right)$$

graph $G=(V,E)$ Boolean symmetric $f:\{0,1\}^d \to \mathbb{C}$ ($q=2$)

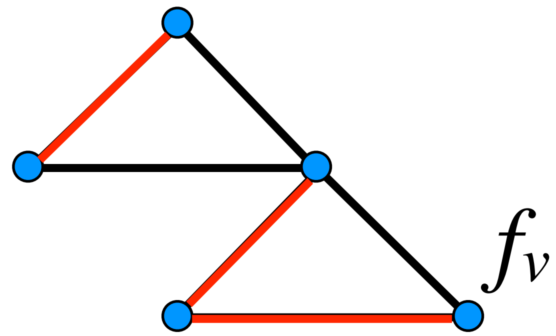

$f = [f_0, f_1, \ldots, f_d]$

where $f_i = f(x)$ for $\sum_j x_j = i$

$f_v = [1, \mu, 1, \mu, 1, \mu, \ldots]$

subgraphs world (Jerrum-Sinclair'90):

$$Z = \sum_{X \subseteq E} \mu^{\mathrm{odd}(X)}$$

odd(X) = #odd-degree vertices in X

$$\text{holant} = \sum_{\sigma \in \{0,1\}^E} \prod_{v \in V} [1, \mu, 1, \mu, \ldots] \left(\sigma|_{E(v)}\right)$$

graph $G=(V,E)$

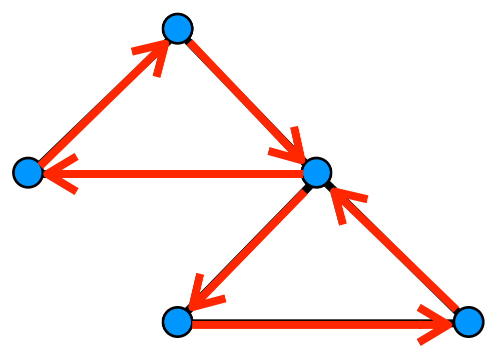

incidence graph $B(V,E,F)$

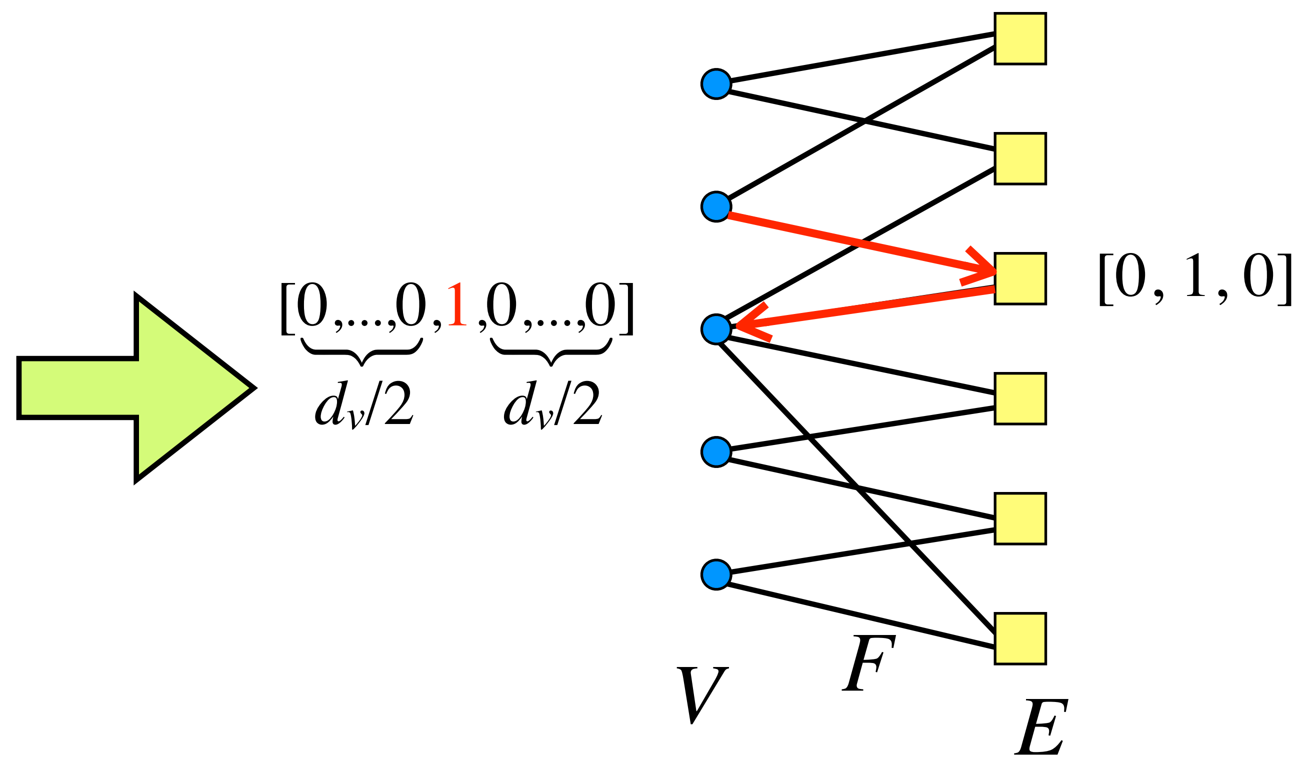

counting Eulerian orientations:

$$\text{holant} = \sum_{\sigma\in\{0,1\}^F}\prod_{v\in V}[\underbrace{0,\ldots,0}_{\deg(v)/2},1,\underbrace{0,\ldots,0}_{\deg(v)/2}]\left(\sigma\mid_{F(v)}\right)\prod_{e=(e_1,e_2)\in E}[\sigma(e_1)\neq\sigma(e_2)]$$

Treewidth

- tw(G): *treewidth* of graph G
- measures how much a graph is like a tree:
  - tw(tree) = 1
  - tw(k×k grid) = k
  - tw(K_n) = n-1

- q-state spin system on G with n vertices:
  - **Courcelle's Theorem:** $f(q, \mathrm{tw})\mathrm{poly}(n)$ time [Courcelle'90]
  - Junction-tree belief propagation: $q^{O(\mathrm{tw})}\mathrm{poly}(n)$ time
- Holant (*tensor networks*) on G with O(1) max-degree: $q^{O(\mathrm{tw})}\mathrm{poly}(n)$ [Markov, Shi'08] [Arad, Landau'10]

Treewidth

graph $G=(V,E)$

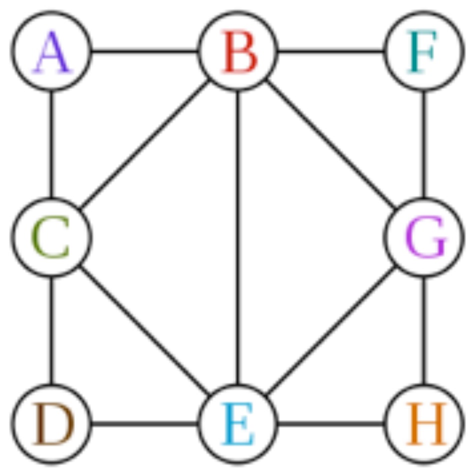

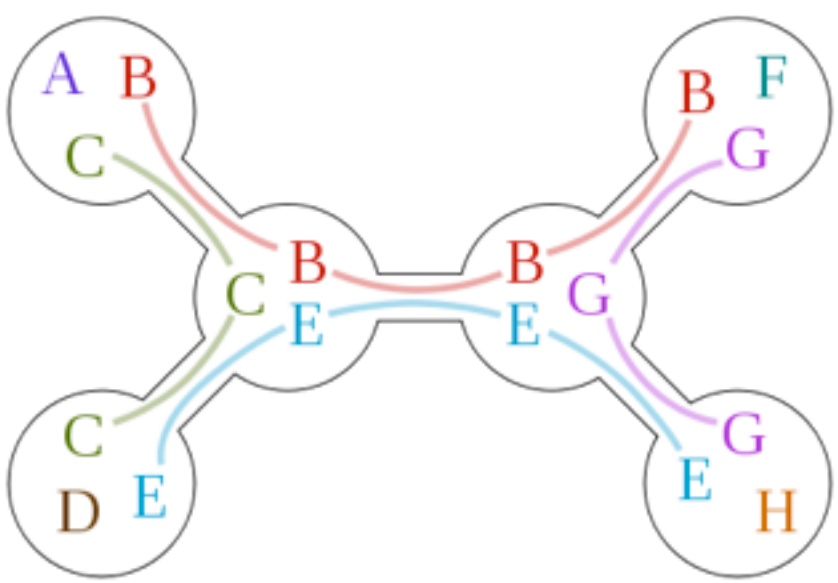

tree-decomposition

a tree of “bags” of vertices:

1. Every vertex is in some bag.
2. Every edge is in some bag.
3. If two bags have a same vertex, then all bags in the path between them have that vertex.

width: max bag size-1

treewidth: width of optimal tree decomposition

Separator Decomposition

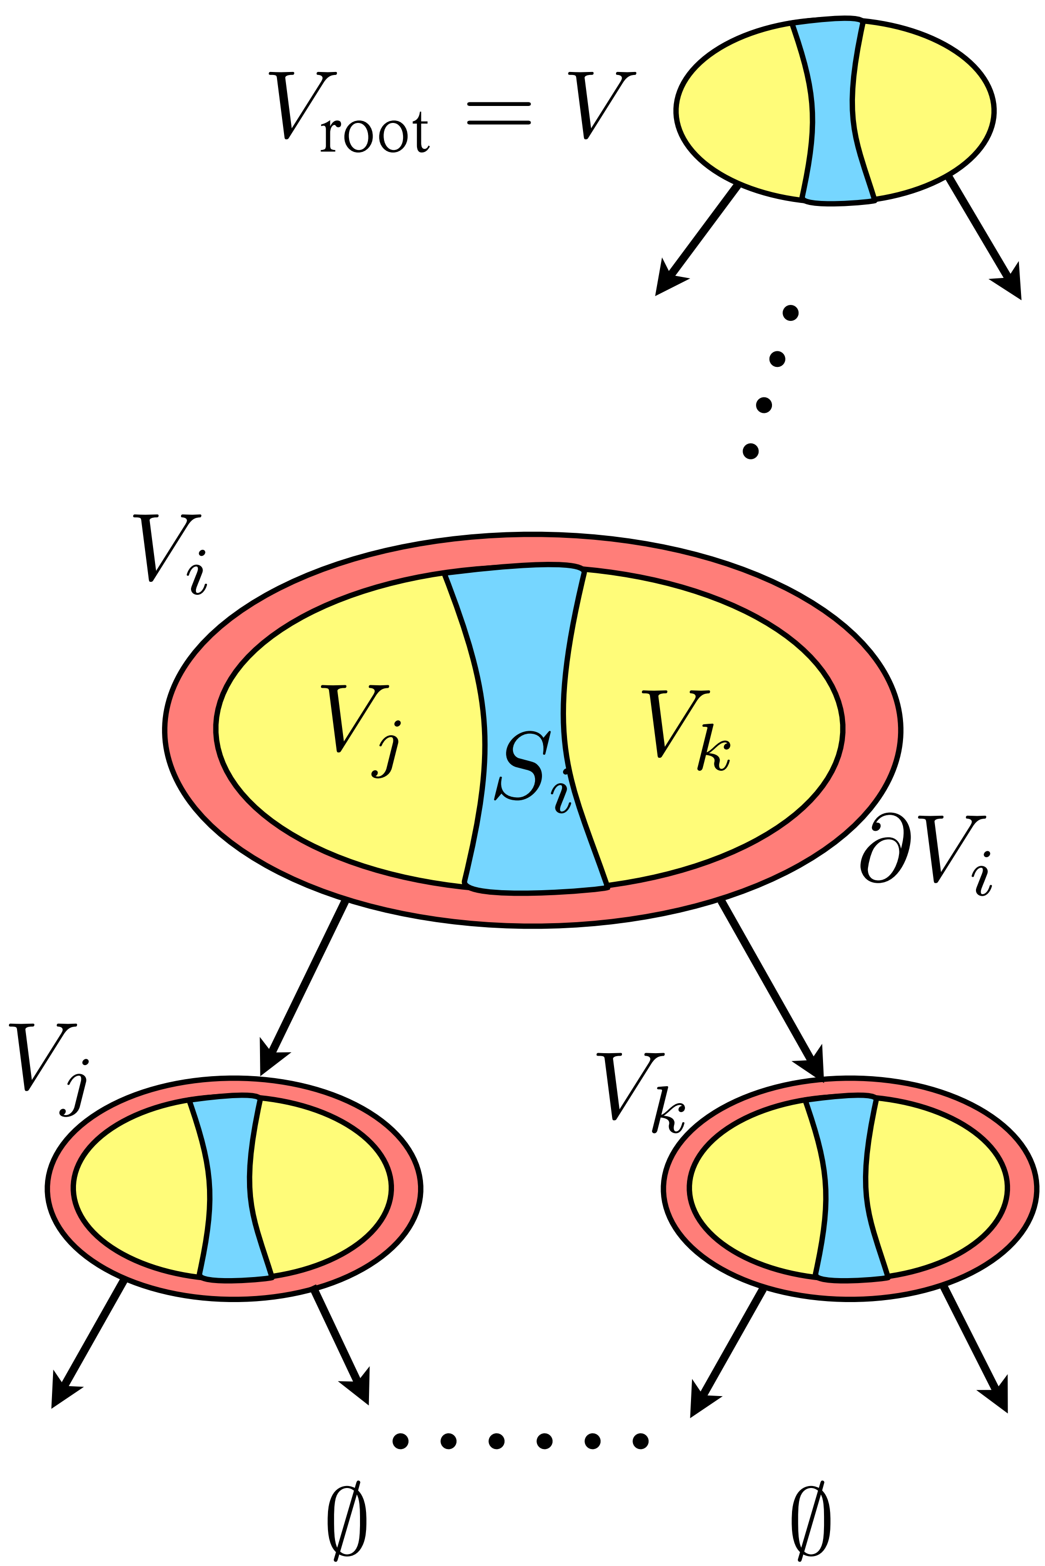

given a graph G=(V,E)

a binary tree T_G of ≤ n nodes:

1. Every node is a vertex set $V_i \subseteq V$. The root is V. Every leaf is $\emptyset$.
2. Every node V_i has a separator $S_i \neq \emptyset$ in $G[V_i]$ separating V_i into V_j and V_k.
3. V_j and V_k are children of V_i.

width of T_G = $\max_i \{|S_i|, |\partial V_i|\}$

∂V_i : vertex boundary of V_i in G

sw(G): width of optimal T_G

1. sw(G) = θ(tw(G))
2. T_G can be constructed in $2^{O(\text{tw})}$poly(n) time

FPT Algorithm for Spin System

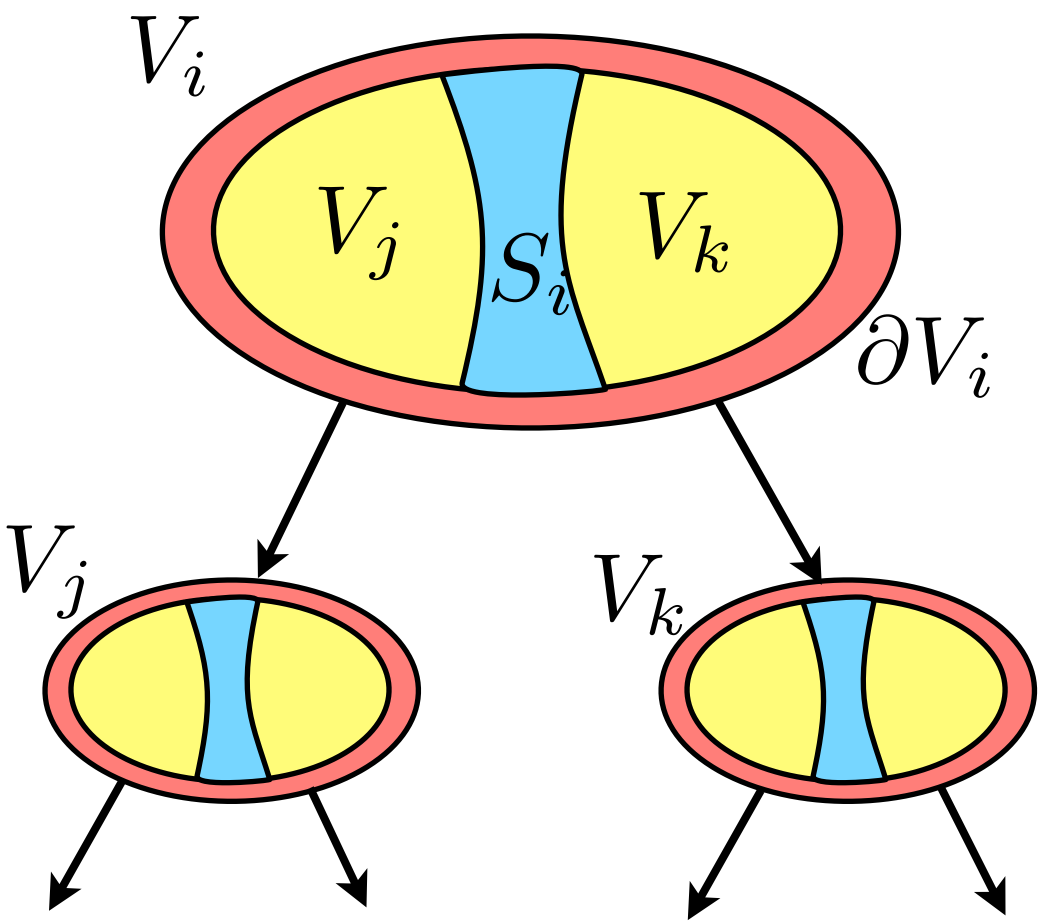

max $\{|S_i|, |\partial V_i|\} = O(\text{tw}(G))$

for q-state spin systems

$$Z = \sum_{\sigma \in [q]^V} \prod_{e=(u,v)\in E} A_e(\sigma_u, \sigma_v)$$

dynamic programming:

- table size: $q^{O(\text{tw}(G))}n$
- running time: $q^{O(\text{tw}(G))}n$

for $\tau \in [q]^{\partial V_i}$:

$$Z(V_i, \tau) = \sum_{\substack{\text{feasible} \\ \sigma \in [q]^{S_i}}} Z(V_j, \tau \cup \sigma \mid_{\partial V_j}) Z(V_k, \tau \cup \sigma \mid_{\partial V_k}) Z(S_i \cup \partial V_i, \sigma \cup \tau)$$

also works for bounded-degree Holant

“Fine-grained” Classification

classify the computational complexity of

$$\mathrm{Holant}(\mathcal{G}, \mathcal{F})$$

in terms of graph family $\mathcal{G}$ and function family $\mathcal{F}$

Classify $\mathrm{Holant}(\mathcal{G}, \mathcal{F})$:

- $\mathcal{G}$ = all graphs: $\mathbf{2^{O(\mathit{n})}}$ time
- $\mathcal{G}$ = graphs with treewidth k: $\mathbf{2^{O(\mathit{k})}poly(\mathit{n})}$ time
- $\mathcal{G}$ = planar graphs:
  - PTAS for *log-Holant* (log-partition function)
  - FPTAS for Holant assuming *strong spatial mixing*

Pinning & Peering

symmetric $f : [q]^d \to \mathbb{C}$ fix any $\tau \in [q]^{d\text{-}k}$

$$\mathrm{Pin}_\tau(f) : [q]^k \to \mathbb{C}$$

Pinning: $\forall \sigma \in [q]^k$

$$\mathrm{Pin}_\tau(f)(\sigma) = f(\tau_1, \tau_2, \ldots, \tau_{d\text{-}k}, \underbrace{\sigma_1, \sigma_2, \ldots, \sigma_k}_{k})$$

Peering: equivalent relation between $\tau, \tau' \in [q]^{d\text{-}k}$

$$\tau \sim \tau' \text{ if } \mathrm{Pin}_\tau(f) = \mathrm{Pin}_{\tau'}(f)$$

f is *r*-regular if $r = \max_{0 \le k \le d} \left| \left\{ \mathrm{Pin}_\tau(f) \mid \tau \in [q]^{d-k} \right\} \right|$
not violating f

Boolean symmetric $f : \{0,1\}^d \to \mathbb{C}$

$\tau \in \{0,1\}^{i+j}$ τ has i 1's and j 0's

$$f = [~~f_0, f_1, \ldots, f_{i-1}~~, [f_i, \ldots, f_{d-j-1}], ~~f_{d-j}, \ldots, f_d~~]$$

$[f_i, \ldots, f_{d-j-1}] = \mathrm{Pin}_\tau(f)$ where $f_i = f(x)$ for $\Sigma_j x_j = i$

> f is r-regular if $r = \max_{0 \le k \le d} \left| \{ \mathrm{Pin}_\tau(f) \mid \tau \in [q]^{d-k} \} \right|$
>
> not violating f

- 2-state spin systems: EQ=[1,0,...,0,1] is 2-regular, A_e is 3-regular
- matchings & PMs: [1,1,0,...,0] and [0,1,0,...,0] are 2-regular
- edge covers: [0,1,1,...,1] is 2-regular
- subgraphs world: $[1, \mu, 1, \mu, \ldots]$ is 2-regular
- Eulerian orientations: [0, ... ,0, 1, 0, ... , 0] is $(d/2+1)$-regular

symmetric $f : \{0,1,2\}^d \to \mathbb{C}$

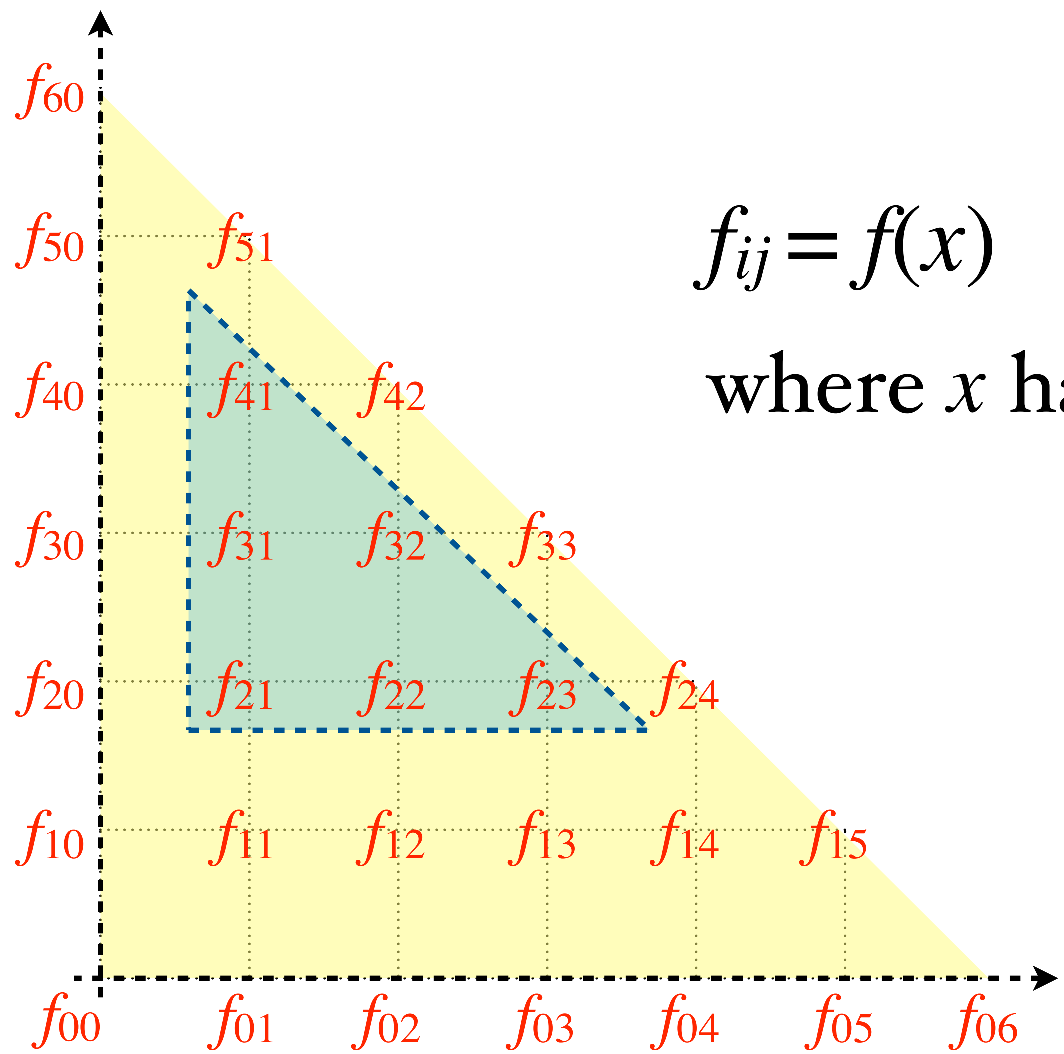

$\mathrm{Pin}_{\tau}(f)$

τ =0112

$f_{ij} = f(x)$

where x has i 0's and j 1's

- all d-ary symmetric function is at most $\binom{d+q-1}{q-1}$-regular
- bounded-degree Holant is O(1)-regular

Theorem

For any constant domain size $q \geq 2$, Holant of any graph $G=(V, E)$ with r-regular symmetric signatures can be computed in time:

- $r^{O(n)}$ where $n=|V|$
- $r^{O(\mathrm{tw}(G))}\, n + 2^{O(\mathrm{tw}(G))}\,\mathrm{Poly}(n)$

- extendable to asymmetric signatures (under proper assumption for evaluating asymmetric functions with unbounded arity);
- implications in approximate counting on planar graphs:
  - PTAS for log-holant
  - FPTAS for holant assuming strong spatial mixing

A Simple $r^{O(n)}$-time Algorithm

given a Holant instance: $G=(V,E)$ $V = \{v_1, v_2, \ldots, v_n\}$

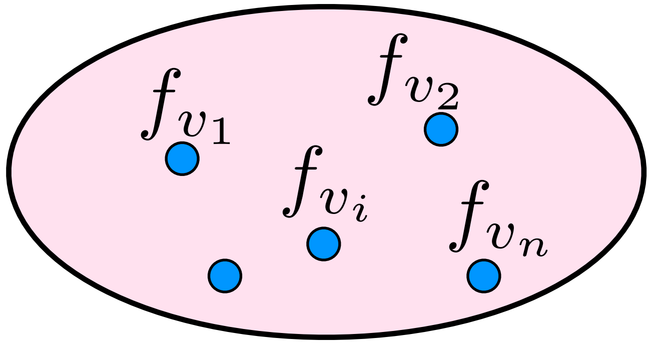

all $f_{v_i} : [q]^{\deg(v_i)} \to \mathbb{C}$

are $\leq r$-regular

f is r-regular if $r = \max_{0\leq k\leq d} \left|\{\mathrm{Pin}_\tau(f) \mid \tau \in [q]^{d-k}\}\right|$

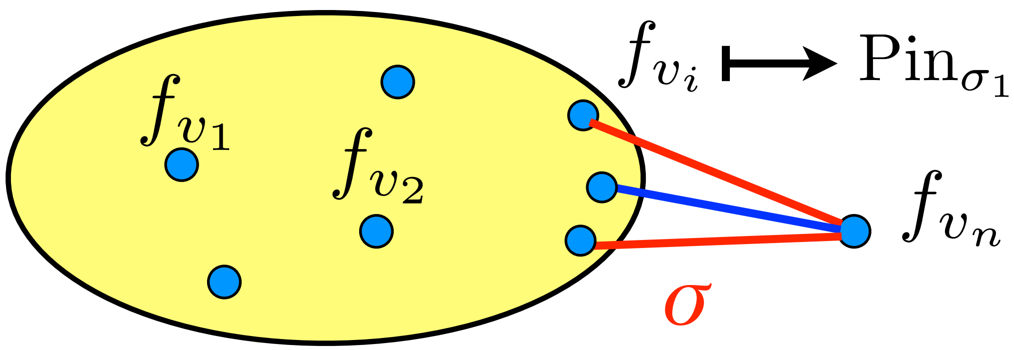

dynamic programming:

- table size: $r^n n$
- running time: $r^n q^n = r^{O(n)}$

$$\mathrm{holant}(G, \{f_{v_1}, \ldots, f_{v_n}\}) = \sum_{\substack{\text{feasible} \\ \sigma\in[q]^{\deg(v_n)}}} f_{v_n}(\sigma) \cdot \mathrm{holant}\left(G \setminus \{v_n\}, \{\text{new } f^\sigma_{v_1}, \ldots, f^\sigma_{v_i}\}\right)$$

where $f^\sigma_{v_i} = \begin{cases} f_{v_i} & v_i \not\sim v_n \\ \mathrm{Pin}_{\sigma_j}(f_{v_i}) & v_i \text{ is } j\text{'s nbr. of } v_n \end{cases}$

Oracles

given a Holant instance: $G=(V,E)$ $V = \{v_1, v_2, \ldots, v_n\}$

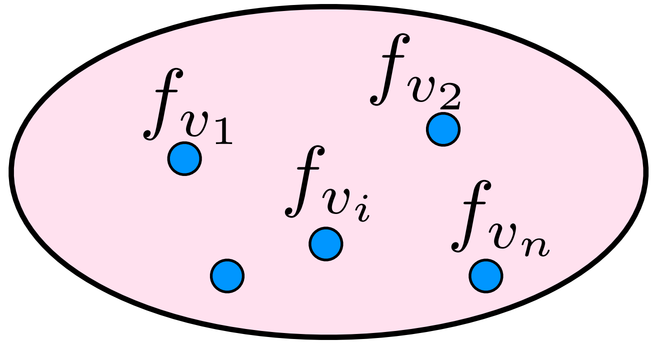

all $f_{v_i} : [q]^{\deg(v_i)} \to \mathbb{C}$

are $\leq r$-regular

fix f_{v_i} and arity k: there are $\leq r$ distinct $\mathrm{Pin}_\tau(f_{v_i}) : [q]^k \to \mathbb{C}$

denoted as $f^k_{i,1}, f^k_{i,2}, \ldots, f^k_{i,r}$

Evaluation oracle: Given (i, j, k) and $\sigma \in [q]^k$, returns $f^k_{i,j}(\sigma)$

Pinning oracle: Given (i, j, k) and $\tau \in [q]^{k-l}$, returns j' that $f^l_{i,j'} = \mathrm{Pin}_\tau(f^k_{i,j})$

For symmetric function $f : [q]^d \to \mathbb{C}$ with constant q,
these oracles can be implemented using:

- Poly(d) preprocessing time
- Poly(d) space
- O(1) query time

for Holant with finitely many signatures:
this can be "hard-wired" into the algorithm

An FPT Algorithm

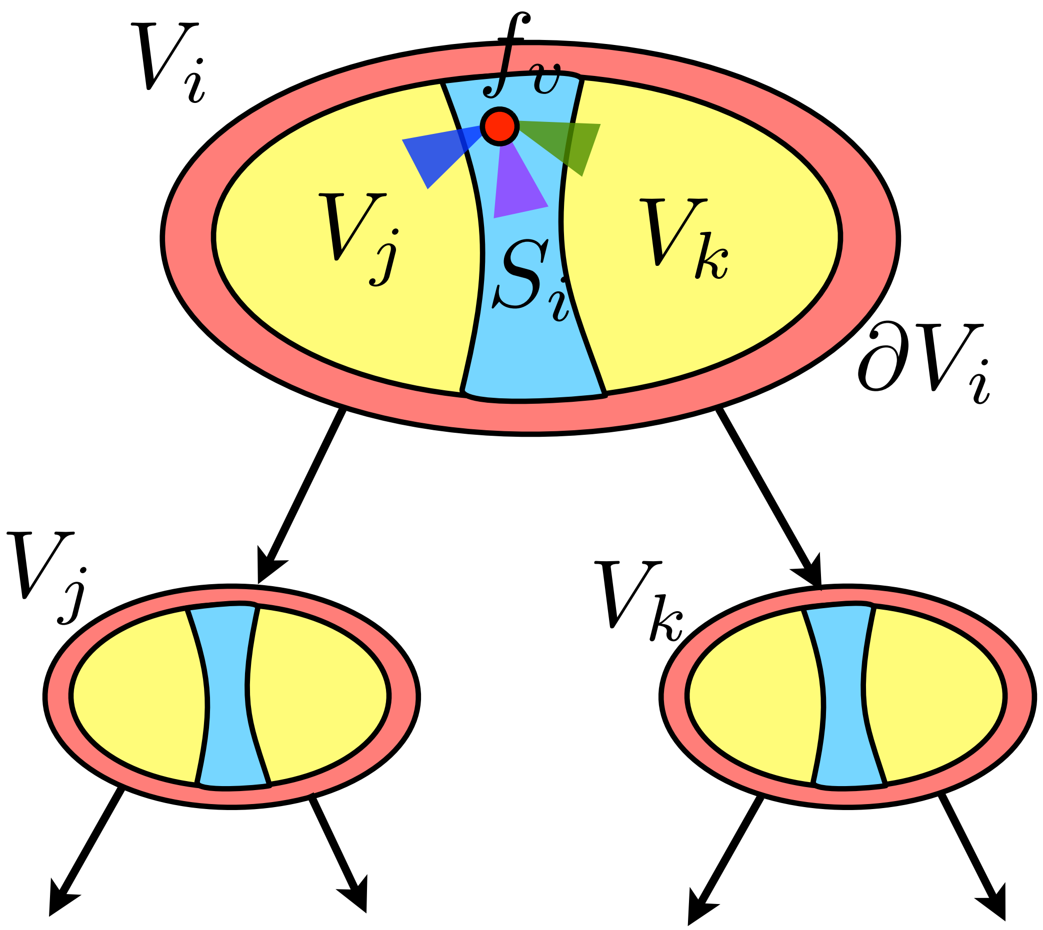

max $\{|S_i|, |\partial V_i|\}$ = O(tw(G))

for Holant problems:

$$\text{holant} = \sum_{\sigma \in [q]^E} \prod_{v \in V} f_v\left(\sigma\mid_{E(v)}\right)$$

each $v \in S_i$ has unbounded degree
but only has 3 classes of edges

Peering: $\tau \sim \tau'$ if $\text{Pin}_{\tau}(f) = \text{Pin}_{\tau'}(f)$

Observation:

$f(\varrho\sigma\tau)$ is determined by the equivalent classes for ϱ, σ, τ

enumerate equivalent classes of local configurations!

An FPT Algorithm

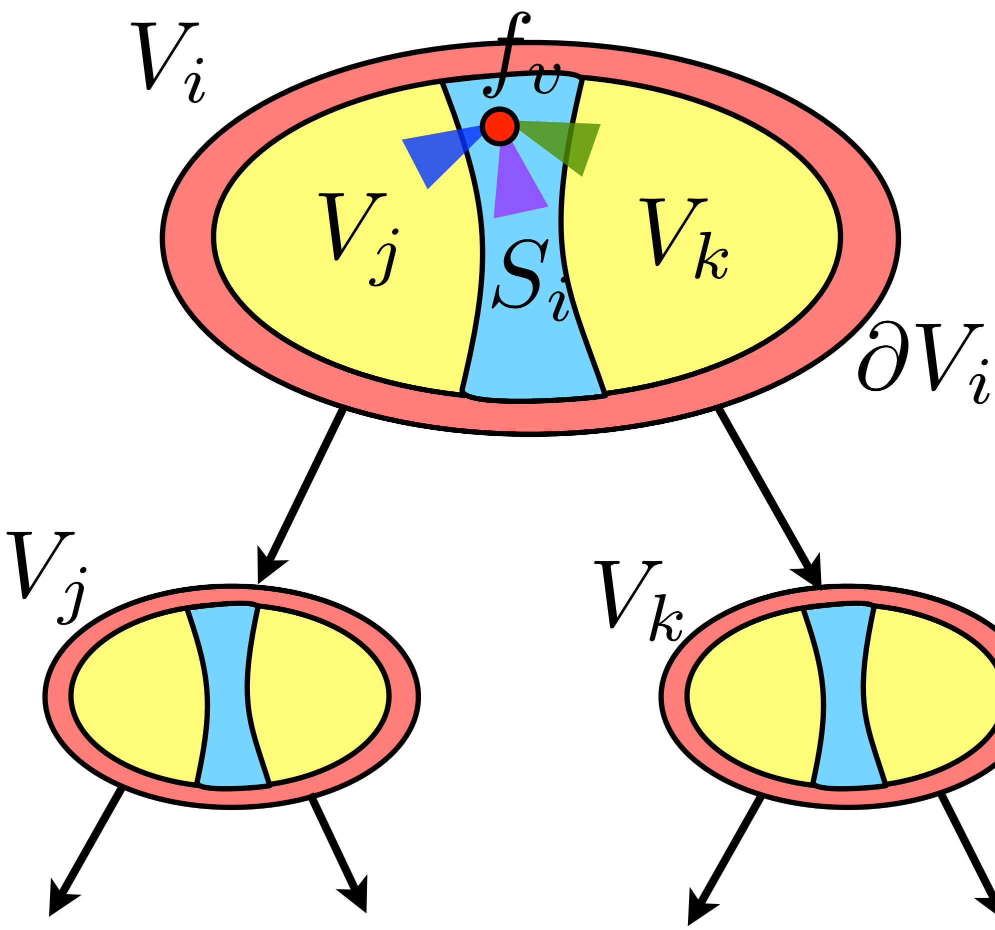

each $v \in S_i$ has unbounded degree
but only has 3 classes of edges

$\tau \sim \tau'$ if $\mathrm{Pin}_\tau(f) = \mathrm{Pin}_{\tau'}(f)$

$f(\varrho\sigma\tau)$ is determined by the equivalent classes for ϱ, σ, τ

$$\mathrm{Peer}_\tau(f)(\sigma) = \begin{cases} 1 & \mathrm{Pin}_\sigma(f) = \mathrm{Pin}_\tau(f) \\ 0 & \text{o.w.} \end{cases}$$

f is r-regular ⇒ $\mathrm{Peer}_\tau(f)$ is $\leq r$-regular

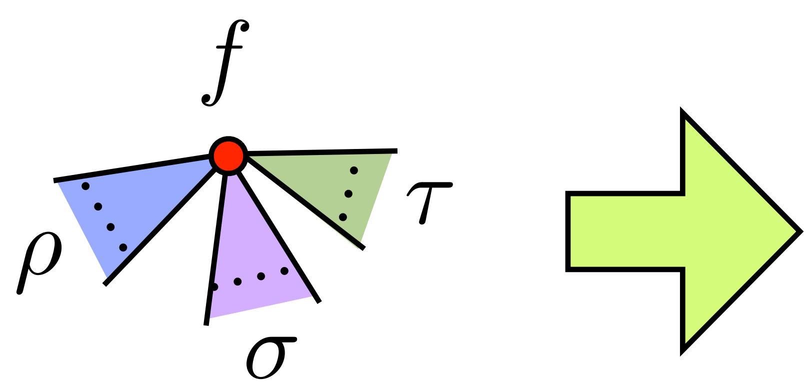

locally enumerate
equivalent classes ϱ, σ, τ

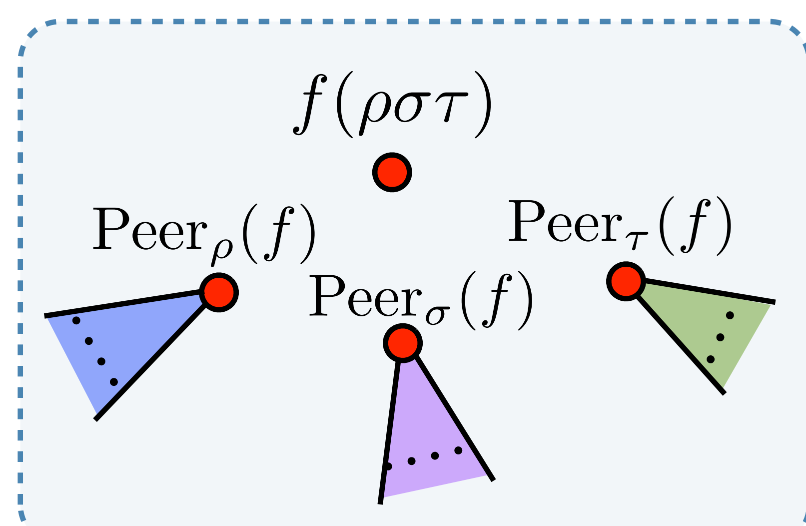

An FPT Algorithm

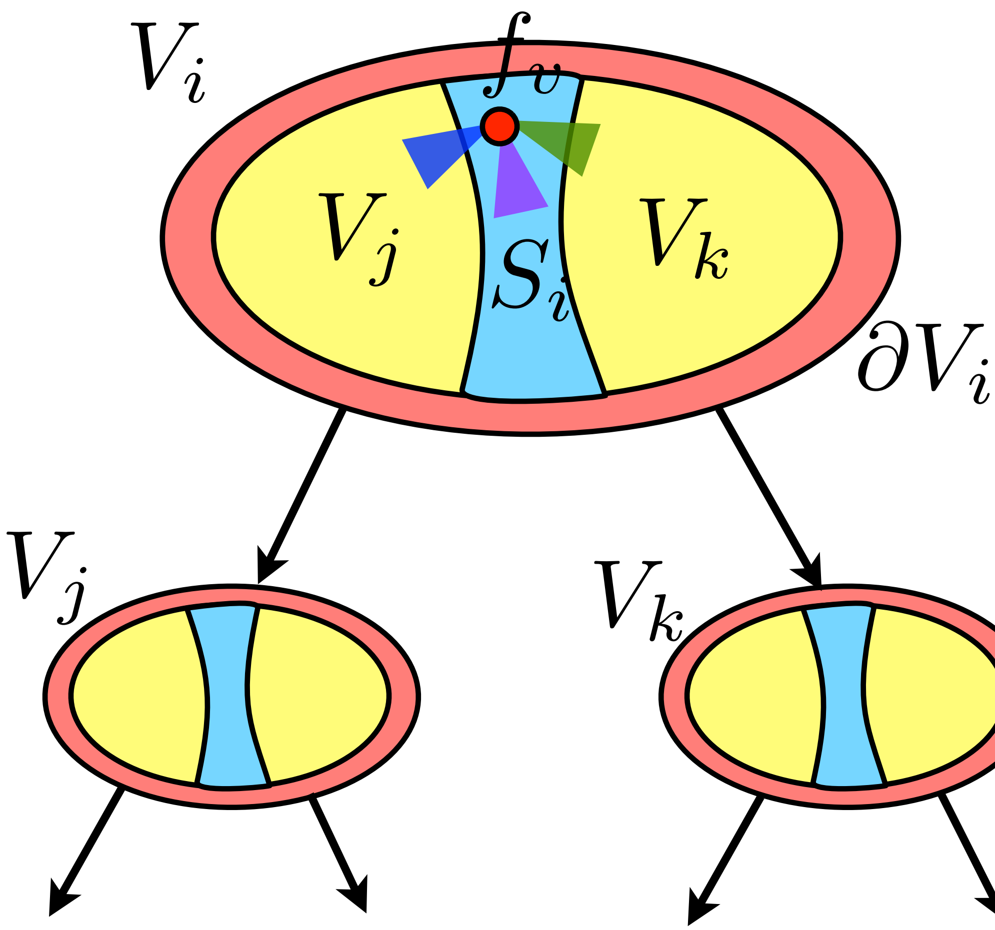

each $v \in S_i$ has unbounded degree but only has 3 classes of edges

$\tau \sim \tau'$ if $\mathrm{Pin}_\tau(f) = \mathrm{Pin}_{\tau'}(f)$

$f(\varrho\sigma\tau)$ is determined by the equivalent classes for ϱ, σ, τ

$$\mathrm{Peer}_\tau(f)(\sigma) = \begin{cases} 1 & \mathrm{Pin}_\sigma(f) = \mathrm{Pin}_\tau(f) \\ 0 & \text{o.w.} \end{cases}$$

$$\text{Holant}\,(V_i, \text{boundary signatures } \mathrm{Peer}_\pi(f_v) \text{ at } v \in \partial V_i)$$

$$= \sum_{\substack{\text{equiv. classes } \rho_v, \sigma_v, \tau_v \\ \text{at every } v \in S_i}} \text{Holant}\left(V_j, \begin{array}{c}\text{old boundary signatures of } V_j \\ \text{and } \mathrm{Peer}_{\rho_v}(f_v) \text{ at } v \in S_i\end{array}\right) \cdot \text{Holant}\left(V_k, \begin{array}{c}\text{old boundary signatures of } V_k \\ \text{and } \mathrm{Peer}_{\tau_v}(f_v) \text{ at } v \in S_i\end{array}\right) \cdot \text{Holant}\left(S_i \cup \partial V_i, \begin{array}{c}\text{old boundary signatures of } V_i \\ \text{and } \mathrm{Peer}_{\sigma_v}(f_v) \text{ at } v \in S_i\end{array}\right) \cdot \prod_{v \in S_i} f_v(\rho_v, \sigma_v, \tau_v)$$

An FPT Algorithm

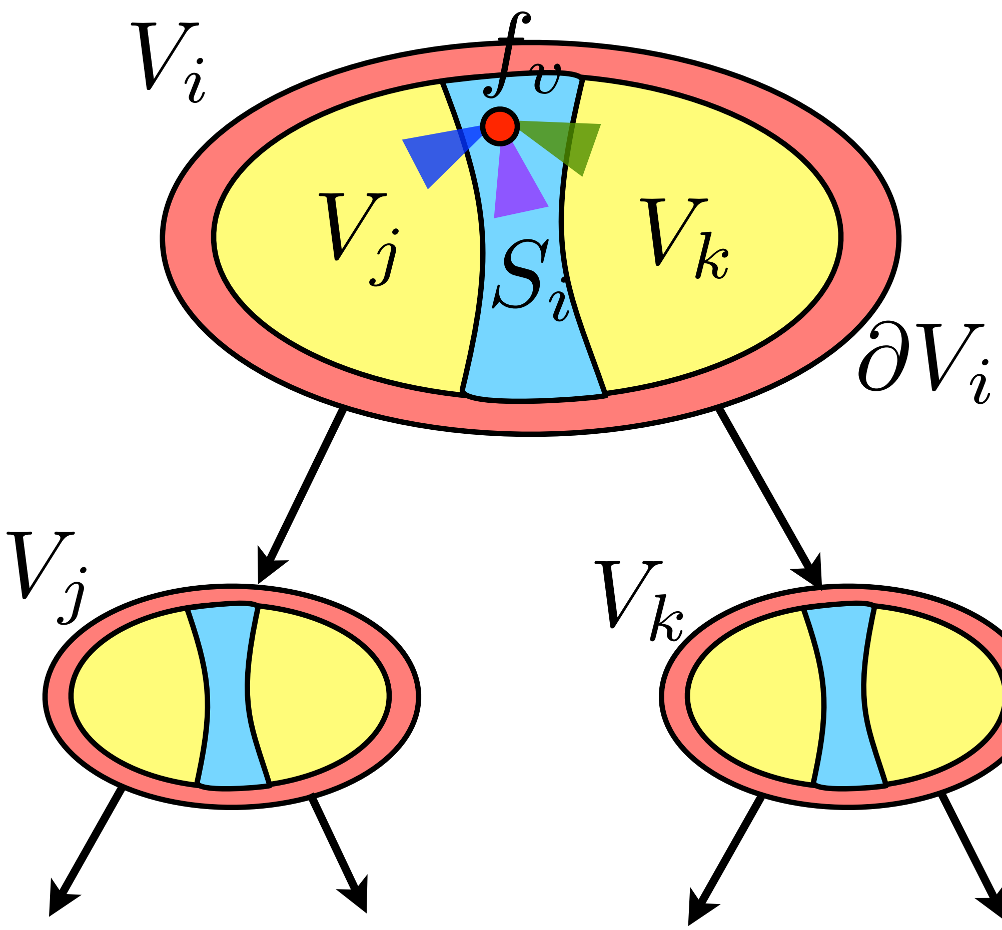

fix $f_{v_1}, f_{v_2}, \ldots, f_{v_n}$

each $v \in S_i$ has unbounded degree

but only has 3 classes of edges

$\tau \sim \tau'$ if $\text{Pin}_{\tau}(f) = \text{Pin}_{\tau'}(f)$

$f(\varrho\sigma\tau)$ is determined by the equivalent classes for ϱ, σ, τ

$$\text{Peer}_{\tau}(f)(\sigma) = \begin{cases} 1 & \text{Pin}_{\sigma}(f) = \text{Pin}_{\tau}(f) \\ 0 & \text{o.w.} \end{cases}$$

Peering oracle: Given i, (equivalent class) τ, returns $\text{Peer}_{\tau}(f_{v_i})$

Pinning oracle & Evaluation oracle: for $\text{Peer}_{\tau}(f_{v_i})$

Evaluation oracle: Given i, (equivalent classes) ϱ, σ, τ, returns $f_{v_i}(\rho\sigma\tau)$

For symmetric functions, can be implemented using Poly(n) preprocessing time, Poly(n) space, and O(1) query time

Theorem

For any constant domain size $q \geq 2$, Holant of any graph $G=(V, E)$ with r-regular signatures can be computed in time:

- $r^{O(n)}$ where $n=|V|$
- $r^{O(\mathrm{tw}(G))}|V| + 2^{O(\mathrm{tw}(G))}\mathrm{Poly}(|V|)$

assuming Peering, Pinning, & Evaluation oracles.

asymmetric $f : [q]^d \to \mathbb{C}$

fix any $\tau \in [q]^{d-k}$ and $S \in \binom{[d]}{d-k}$

$$\mathrm{Pin}_{S,\tau}(f) = f(\underbrace{\tau_1, \ldots, \tau_2, \ldots, \tau_3, \ldots, \tau_{d-k}}_{\text{at positions in } S}, \ldots)$$

f is r-regular if $r = \max_{0 \leq k \leq d} \max_{S \in \binom{[d]}{d-k}} \left|\left\{\mathrm{Pin}_{S,\tau(f)} \mid \tau \in [q]^{d-k}\right\}\right|$

Planar Decomposition

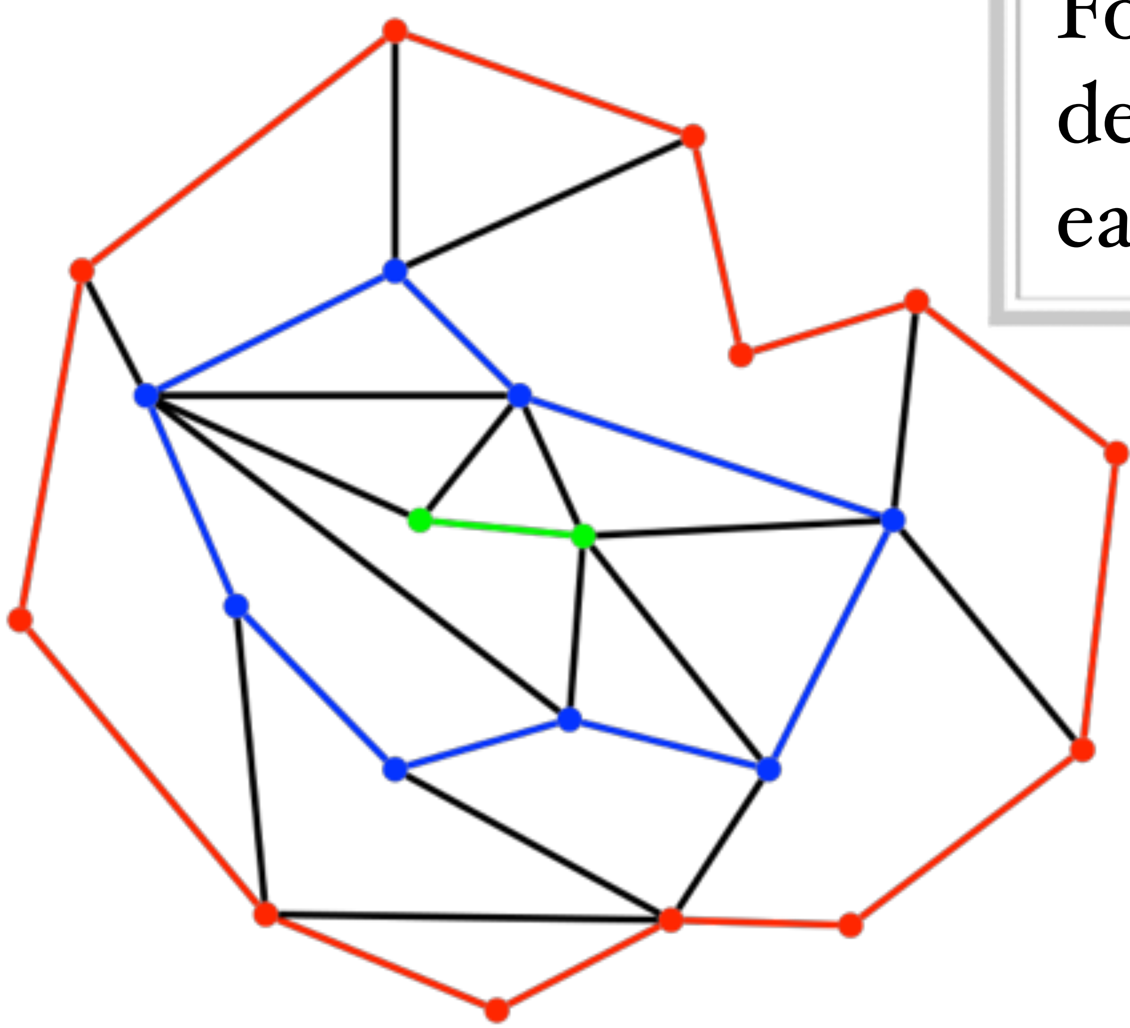

Baker's Decomposition (Baker'93):

For all k, a planar graph G can be decomposed into subgraphs G_1, ..., G_k each of treewidth $O(k)$.

Jerrum-Goldberg-McQuillan'12:

For spin systems on planar graphs, if always holant ≥ $\exp(\Omega(n))$ then ∃ PTAS for log-holant.

Correlation Decay

strong spatial mixing (SSM): $\forall \sigma_B \in [q]^B$

$$\left|\Pr(\sigma(e) = c \mid \sigma_A) - \Pr(\sigma(e) = c \mid \sigma_A, \sigma_B)\right| \le \mathrm{poly}(|V|)\exp(-\Omega(t))$$

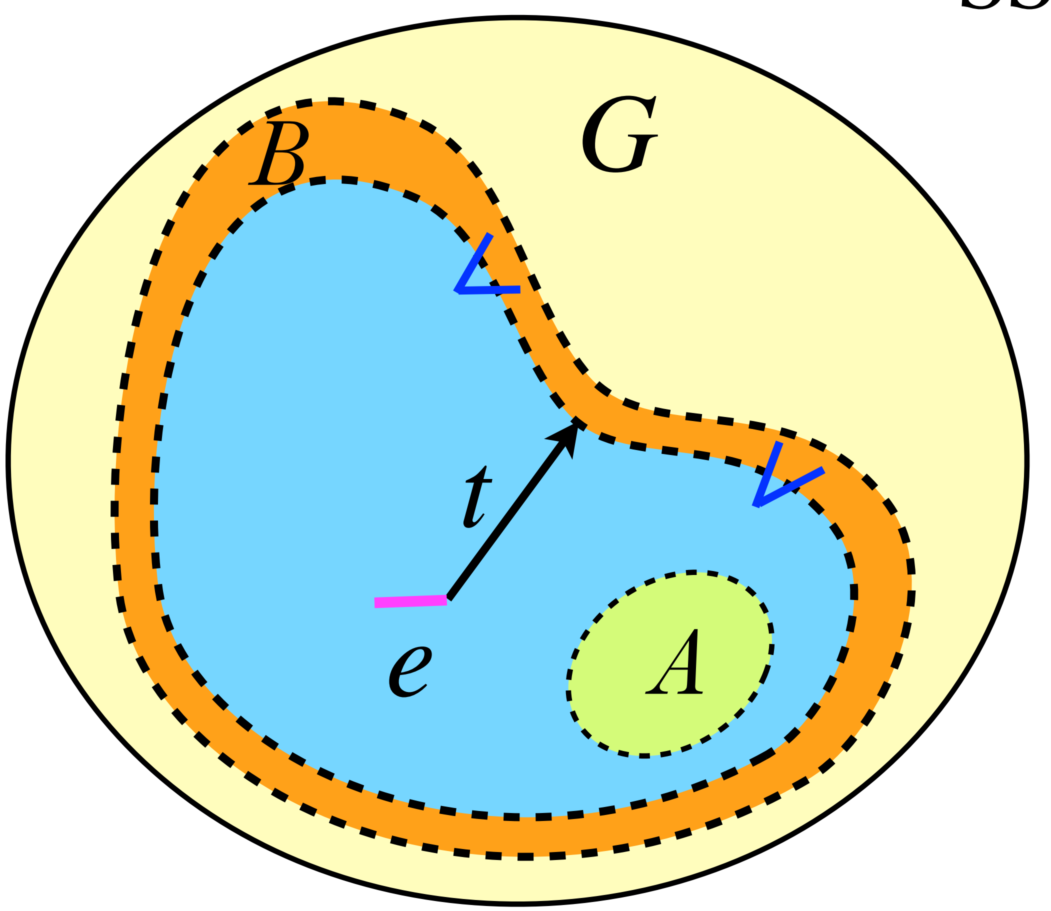

SSM: sufficiency of local information for approx. of $\Pr(\sigma(e) = c \mid \sigma_A)$

? ⇒ efficiency of local computation

Theorem (Demaine-Hajiaghayi'04)
For *apex-minor-free* graphs, treewidth of t-ball is O(t).

Correlation Decay

strong spatial mixing (SSM): $\forall \sigma_B \in [q]^B$

$$\left|\Pr(\sigma(e) = c \mid \sigma_A) - \Pr(\sigma(e) = c \mid \sigma_A, \sigma_B)\right| \leq \mathrm{poly}(|V|) \exp(-\Omega(t))$$

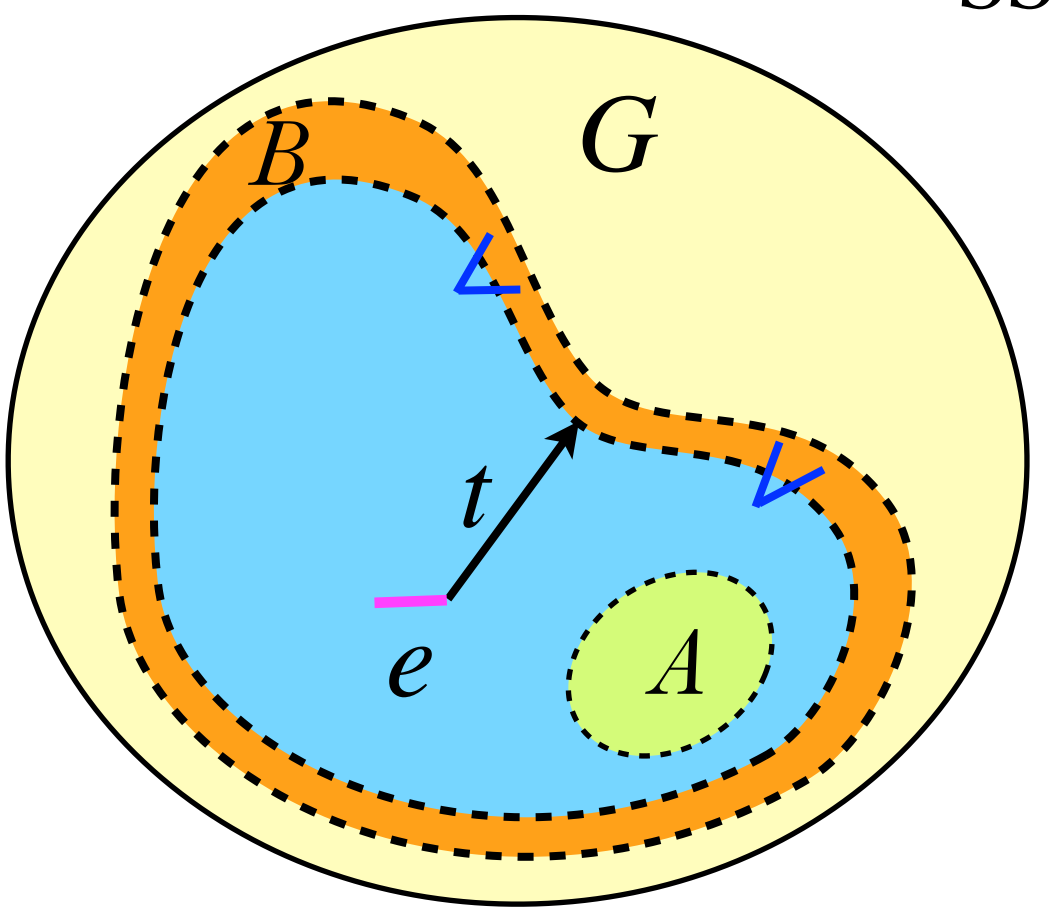

SSM: sufficiency of local information for approx. of $\Pr(\sigma(e) = c \mid \sigma_A)$

for planar graphs ⇒ efficiency of local computation

for self-reducible problems ⇒ FPTAS for Holant

Conclusion

- A class of Holant problems behave like bounded-degree #CSP:
  - $2^{O(n)}$- and $2^{O(\mathrm{tw})}\mathrm{Poly}(n)$-time algorithms.
- Regularity implies efficient DP algorithms, but:
  - is not robust to holographic transformation;
  - does not cover inversion-based algorithms.
- Open question: a Holant problem with Boolean-domain symmetric signatures (and is easy for decision) with $n^{\Omega(\mathrm{tw})}$ lower bound:
  - must have unbounded degree, must be irregular.

Thank you!