2024.1.8 SODA 2024 @ Alexandria, Virginia, U.S.

Gap Amplification for Reconfiguration Problems

Naoto Ohsaka

(CygerAgent, Inc.)

Intro of reconfiguration

Imagine **connecting** a pair of feasible solutions (of NP problem)
under a particular adjacency relation

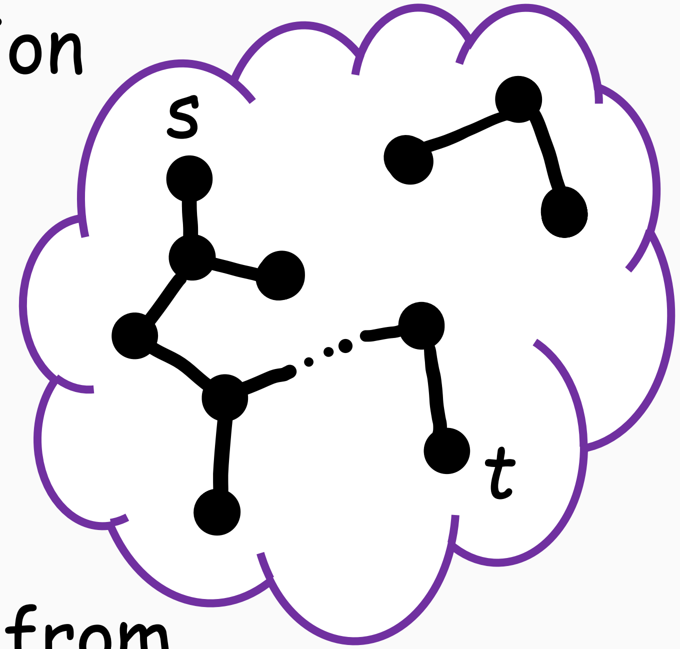

Q. Is a pair of solutions reachable to each other?
Q. If so, what is the shortest transformation?
Q. If not, how can the feasibility be relaxed?

Many reconfiguration problems have been derived from

Satisfiability, Coloring, Vertex Cover, Clique, Dominating Set, Feedback Vertex Set, Steiner Tree, Matching, Spanning Tree, Shortest Path, Set Cover, Subset Sum, …

See [Ito-Demaine-Harvey-Papadimitriou-Sideri-Uehara-Uno. Theor. Comput. Sci. 2011] [Nishimura. Algorithms 2018] [van den Heuvel. Surv. Comb. 2013] [Hoang. https://reconf.wikidot.com/]

Example

3-SAT Reconfiguration

[Gopalan-Kolaitis-Maneva-Papadimitriou. SIAM J. Comput. 2009]

- **Input:** 3-CNF formula φ & satisfying σ_s, σ_t
- **Output:** $\boldsymbol{\sigma} = \langle \sigma^{(0)}=\sigma_s, \ldots, \sigma^{(\ell)}=\sigma_t \rangle$ (reconf. sequence) s.t.

  $\sigma^{(i)}$ satisfies φ (feasibility)

  $\mathrm{Ham}(\sigma^{(i-1)}, \sigma^{(i)}) = 1$ (adjacency on hypercube)

YES case

$\varphi = (\bar{x} \vee \bar{y} \vee z) \wedge (\bar{x} \vee y \vee \bar{z}) \wedge (x \vee \bar{y} \vee \bar{z})$

$\sigma_s = (1,0,0)$

$\sigma_t = (0,1,0)$

⚠ Length of $\boldsymbol{\sigma}$ can be $2^{\Omega(\text{input size})}$

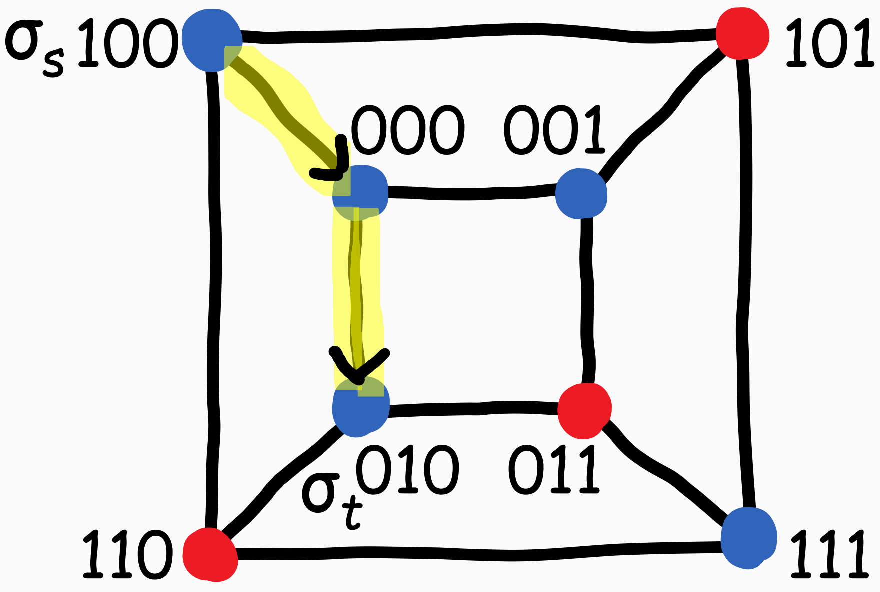

Example

3-SAT Reconfiguration

[Gopalan-Kolaitis-Maneva-Papadimitriou. SIAM J. Comput. 2009]

- **Input**: 3-CNF formula φ & satisfying σ_s, σ_t
- **Output**: $\boldsymbol{\sigma} = \langle \sigma^{(0)}=\sigma_s, \ldots, \sigma^{(\ell)}=\sigma_t \rangle$ (reconf. sequence) s.t.
  $\sigma^{(i)}$ satisfies φ (feasibility)
  $\mathrm{Ham}(\sigma^{(i-1)}, \sigma^{(i)}) = 1$ (adjacency on hypercube)

NO case

$\varphi = (\bar{x} \vee \bar{y} \vee z) \wedge (\bar{x} \vee y \vee \bar{z}) \wedge (x \vee \bar{y} \vee \bar{z})$

$\sigma_s = (1,0,0)$

$\sigma_t = (1,1,1)$

⚠ Length of $\boldsymbol{\sigma}$ can be $2^{\Omega(\text{input size})}$

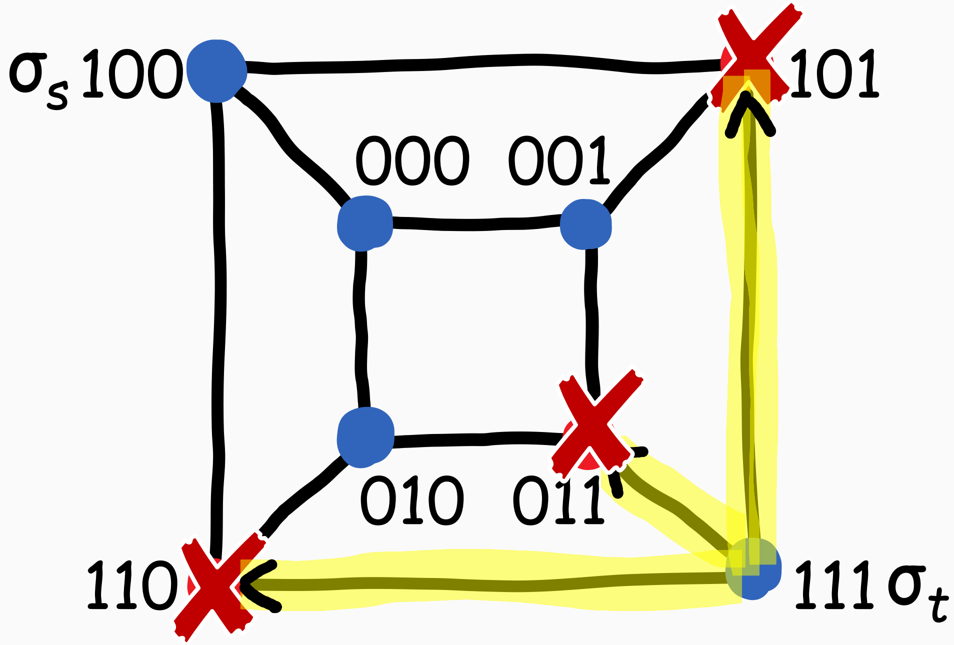

Optimization variants of reconfiguration problems

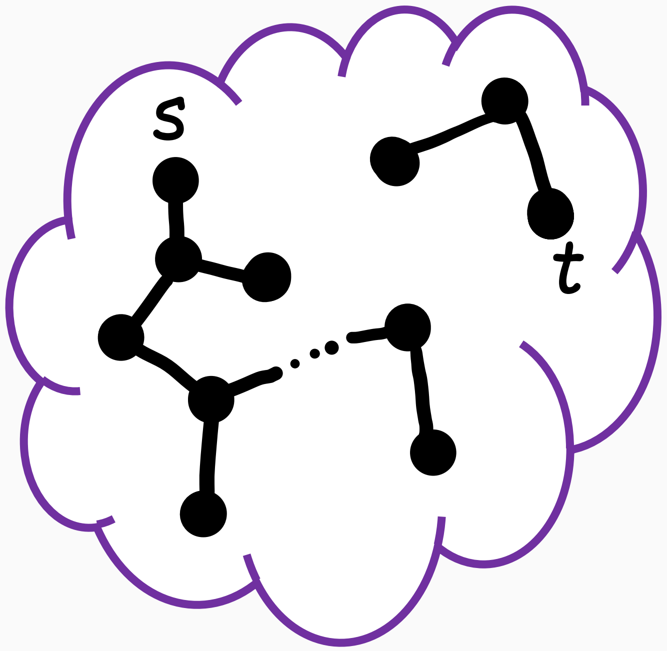

Even if…

- 😖NOT reconfigurable! and/or
- 😖many problems are **PSPACE**-complete!

Still want an "approximate" reconf. sequence

(e.g.) made up of almost-satisfying assignments

⬇

Relax feasibility to obtain approximate reconfigurability

e.g. Set Cover Reconf. [Ito-Demaine-Harvey-Papadimitriou-Sideri-Uehara-Uno. Theor. Comput. Sci. 2011]

Subset Sum Reconf. [Ito-Demaine. J. Comb. Optim. 2014]

Submodular Reconf. [O.-Matsuoka. WSDM 2022]

Example+

Maxmin 3-SAT Reconfiguration

[Ito-Demaine-Harvey-Papadimitriou-Sideri-Uehara-Uno. Theor. Comput. Sci. 2011]

- **Input:** 3-CNF formula φ & satisfying σ_s, σ_t
- **Output:** $\boldsymbol{\sigma} = \langle \sigma^{(0)}=\sigma_s, \dots, \sigma^{(\ell)}=\sigma_t \rangle$ (reconf. sequence) s.t.
  ~~$\sigma^{(i)}$ satisfies φ~~ (feasibility)
  $\mathrm{Ham}(\sigma^{(i-1)}, \sigma^{(i)}) = 1$ (adjacency on hypercube)
- **Goal:** $\max_{\boldsymbol{\sigma}} \mathrm{val}_{\varphi}(\boldsymbol{\sigma}) \stackrel{\text{def}}{=} \min_i$ (frac. of satisfied clauses by $\sigma^{(i)}$)

$\varphi = (\bar{x} \vee \bar{y} \vee z) \wedge (\bar{x} \vee y \vee \bar{z}) \wedge (x \vee \bar{y} \vee \bar{z})$

- $\sigma_s = (1,0,0)$
- $\sigma_t = (1,1,1)$

→ $\mathrm{val}_{\varphi}(\boldsymbol{\sigma}) = \min\{1, \frac{2}{3}, 1\} = \frac{2}{3}$

⚠ Length of $\boldsymbol{\sigma}$ can be $2^{\Omega(\text{input size})}$

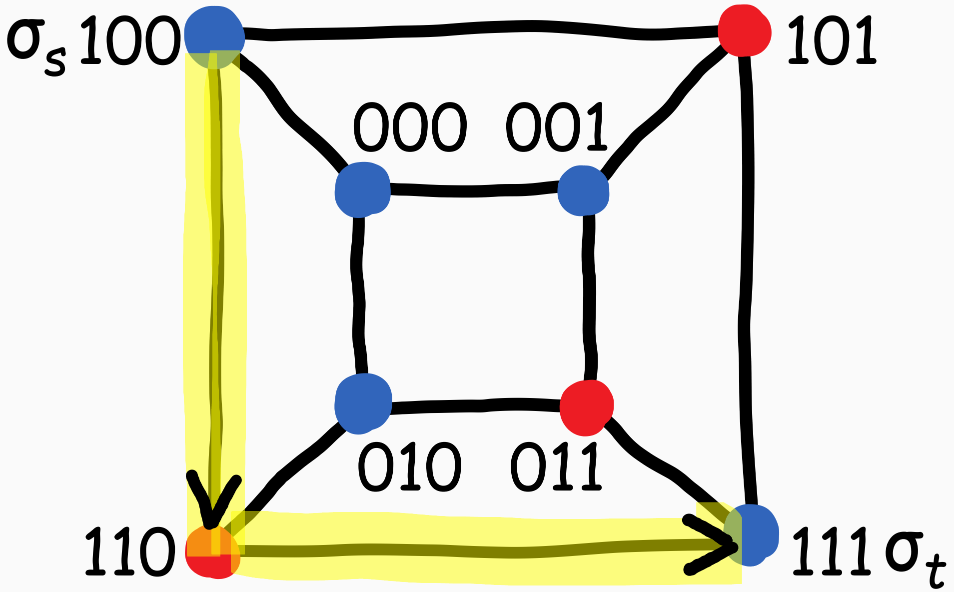

Known results on hardness of approximation

NP-hardness of approx. for Maxmin SAT & Clique Reconfiguration

[Ito-Demaine-Harvey-Papadimitriou-Sideri-Uehara-Uno. Theor. Comput. Sci. 2011]

- Not optimal ∵ SAT & Clique Reconf. are **PSPACE**-complete
- Rely on **NP**-hardness of approximating Max SAT & Max Clique

Significance of showing **PSPACE**-hardness

- no polynomial-time algorithm (**P** $\neq$ **PSPACE**)
- no polynomial-length sequence (**NP** $\neq$ **PSPACE**)

(probabilistically checkable proof)

Reconfiguration analogue of the PCP theorem

[Arora-Lund-Motwani-Sudan-Szegedy. J. ACM 1998] [Arora-Safra. J. ACM 1998]

Our working hypothesis [O. STACS 2023]

Reconfiguration Inapproximability Hypothesis (**RIH**)

Binary CSP G & satisfying ψ_s, ψ_t, **PSPACE**-hard to distinguish btw.

- (Completeness) $\exists \psi \; \mathrm{val}_G(\psi) = 1$ (some sequence violates no constraint)
- (Soundness) $\forall \psi \; \mathrm{val}_G(\psi) < 1\text{-}\varepsilon$ (any sequence violates >ε-frac. of constraints)

- ⟹ **PSPACE**-hard to approx. Maxmin Binary CSP Reconf.
- True if "**NP**-hard" is used [Ito et al. Theor. Comput. Sci. 2011]

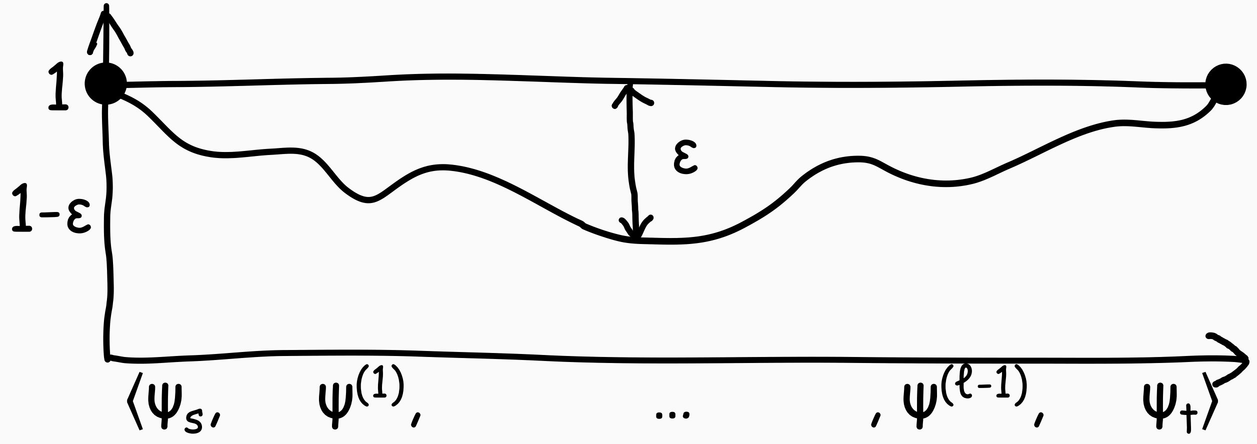

Under **RIH**, many problems are **PSPACE**-hard to approximate via gap-preserving reductions [O. STACS 2023]

Limitation of [O. STACS 2023]

😖 Inapprox. factors are not explicitly shown

Recall from [O. STACS 2023]

- **RIH** claims "∃ε Gap[1 vs. 1-ε] Binary CSP Reconf. is **PSPACE**-hard"
- Can reduce to Gap[1 vs. 1-δ] ∗∗ Reconf.

⚠ δ (as well as ε) can be arbitrarily small, because...

- δ depends on ε (e.g., $\delta = \varepsilon^2$)
- **RIH** doesn't specify any value of ε

➜ May not rule out 0.999⋯999-approx. algorithm

🎯 Wanna say Gap[1 vs. 0.999] ∗∗ Reconf. is **PSPACE**-hard only assuming **RIH**

🎯Our target: **Gap amplification**

- (Polynomial-time) reduction that makes a tiny gap into a larger gap

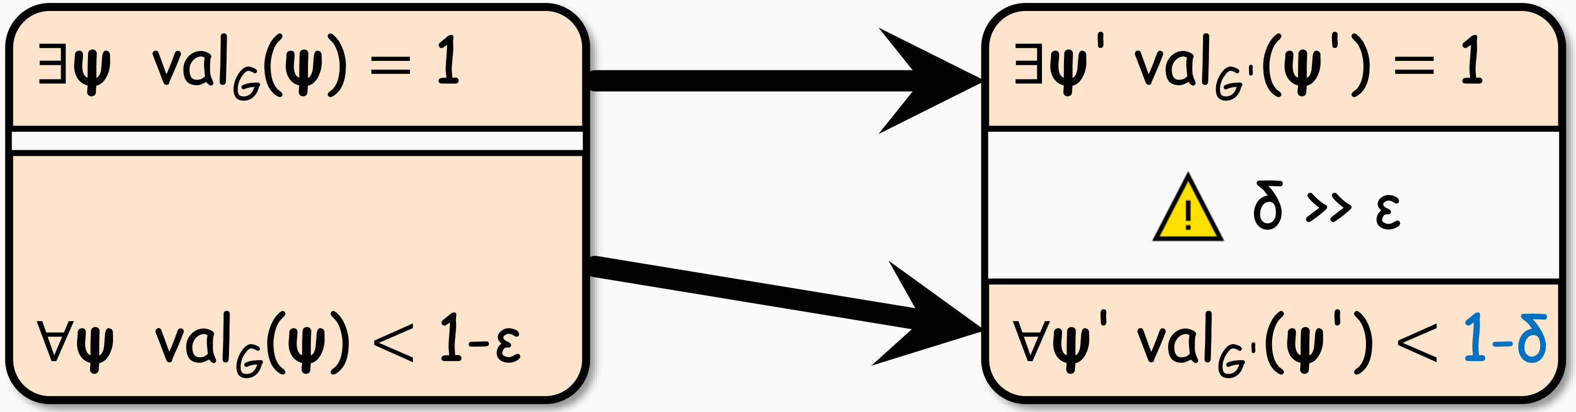

In **NP** world…

The parallel repetition theorem [Raz. SIAM J. Comput. 1998]

➜😋Gap[1 vs. 0.000⋯001] Binary CSP is **NP**-hard (i.e. gap ≈ 1)

In ***reconfiguration*** world…

😖Naïve parallel repetition fails to amplify gap ε of Gap[1 vs. 1-ε] Binary CSP Reconf. [O. arXiv 2023]

🎯Our target: **Gap amplification**

- (Polynomial-time) reduction that makes a tin… …p into a larger gap

$\exists\psi$ $val_G(\psi)$ … $val_{G'}(\psi')$

In NP w…

→ … (i.e. gap ≈ 1)

Can we derive explicit factors of
PSPACE-hardness of approx.
only assuming RIH?

In **reconfigur**… world…

😣Naïve parallel repetition fails to amplify gap ε of Gap[1 vs. 1-ε] Binary CSP Reconf. [O. arXiv 2023]

Our results

😋Can derive explicit inapproximability factors only assuming **RIH**!!

	Maxmin Binary CSP Reconfiguration	Minmax Set Cover Reconfiguration
PSPACE-hardness under **RIH**	**0.9942** (this paper)	**1.0029** (this paper)
NP-hardness rely on parallel repetition theorem [Raz. SIAM J. Comput. 1998]	>0.75 (this paper) 0.993 [Ito et al. Theor. Comput. Sci. 2011] [O. STACS 2023]	1.0029 (this paper)
approximability	≈0.25 [O. arXiv 2023]	2 [Ito et al. Theor. Comput. Sci. 2011]

Main result

Gap amplification for Binary CSP Reconf.

- We prove gap amplification à la Dinur [Dinur. J. ACM 2007]

(Informal) For any small const. $\varepsilon \in (0,1)$,

gap	alphabet size	degree	spectral expansion
1 vs. $1-\varepsilon$	W	d	λ
1 vs. $1-0.0058$	$W' = W^{d^{O(\varepsilon^{-1})}}$	$d' = \left(\frac{d}{\varepsilon}\right)^{O(\varepsilon^{-1})}$	$\lambda' = O\left(\frac{\lambda}{d}\right)d'$

- 😋 Can make λ'/d' arbitrarily small by decreasing λ/d
- 😖 Alphabet size W' gets gigantic depending on ε^{-1}

Related work

Probabilistically checkable debates — PCP-like charact. of **PSPACE**
[Condon-Feigenbaum-Lund-Shor. Chic. J. Theor. Comput. Sci.'95]

- $\Longrightarrow$ Quantified Boolean Formula is **PSPACE**-hard to approx.

Other optimization variants of reconfiguration (orthogonal to this study)

- Shortest sequence finding
  [Bonamy-Heinrich-Ito-Kobayashi-Mizuta-Mühlenthaler-Suzuki-Wasa. STACS 2020]
  [Ito-Kakimura-Kamiyama-Kobayashi-Okamoto. SIAM J. Discret. Math. 2022]
  [Kamiński-Medvedev-Milanič. Theor. Comput. Sci. 2011]
  [Miltzow-Narins-Okamoto-Rote-Thomas-Uno. ESA 2016]
- Incremental optimization
  [Blanché-Mizuta-Ouvrard-Suzuki. IWOCA 2020]
  [Ito-Mizuta-Nishimura-Suzuki. J. Comb. Optim. 2022]
  [Yanagisawa-Suzuki-Tamura-Zhou. COCOON 2021]

In the remainder of this talk...

Sketch of gap amplification

1. Preprocessing step

- Degree reduction [O. STACS 2023]
- Expanderization (skipped)

2. Powering step

- Simple appl. of [Dinur. J. ACM 2007] [Radhakrishnan. ICALP 2006] to Binary CSP Reconf. looses perfect completeness
- TRICK: **Alphabet squaring** [O. STACS 2023] & **modified verifier**

Recap: Max Binary CSP
Dinur's gap amplification [Dinur. J. ACM 2007]

- **Input:** Binary CSP $G = (V,E,\Sigma,\Pi=(\pi_e)_{e\in E})$
- **Output:** $\psi: V\to\Sigma$

  ψ satisfies (v,w) if $(\psi(v), \psi(w)) \in \pi_{(v,w)}$
- **Goal:** $\max_\psi \mathrm{val}_G(\psi) \stackrel{\text{def}}{=}$ (frac. of edges satisfied by ψ)

Example

- 3-Coloring: $\Sigma = \{R,G,B\}$, $\pi_e = \{(R,G), (G,R), (G,B), (B,G), (B,R), (R,B)\}$
- 2-SAT: $\Sigma = \{0,1\}$, π_C = {asgmt. satisfying 2-literal clause C}

(Completeness) $\exists\psi\ \mathrm{val}_G(\psi) = 1 \implies \exists\psi'\ \mathrm{val}_{G'}(\psi') = 1$

(Soundness) $\forall\psi\ \mathrm{val}_G(\psi) < 1\text{-}\varepsilon \implies \forall\psi'\ \mathrm{val}_{G'}(\psi') < 1-\Omega(T\cdot\varepsilon)$

const. parameter

Recap: Dinur's gap amplification [Dinur. J. ACM 2007]

Powering step

Say 3-Coloring $\Sigma=\{R,G,B\}$

Original $G = (V,E,\Sigma,\Pi=(\pi_e)_{e\in E})$ ➜ New $G' = (V,E',\Sigma',\Pi')$

Asgmt. $\psi: V\to\Sigma$ ➜ Asgmt. $\psi': V\to\Sigma^V$

for simplicity

⚠ G must be EXPANDER

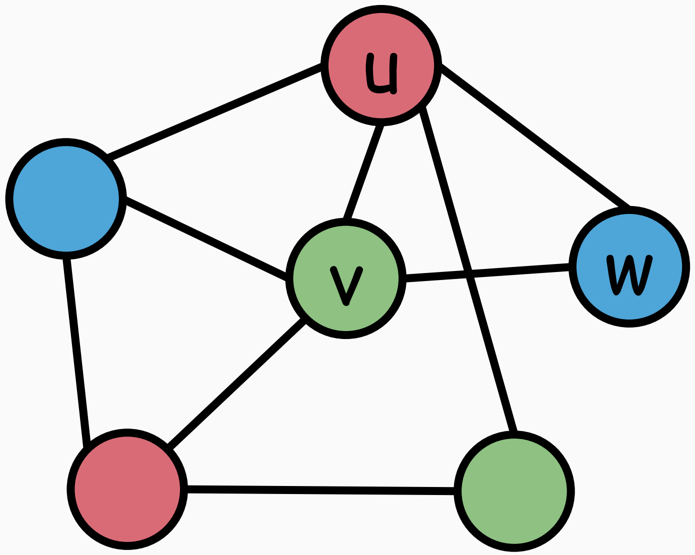

- $\psi'(x)[v] \stackrel{\text{def}}{=}$ "opinion" of $\psi'(x)$ about the value of v
- edge of G' = a length-T random walk over G

const. parameter

Recap: Dinur's gap amplification [Dinur. J. ACM 2007]

Powering step

Say 3-Coloring $\Sigma=\{R,G,B\}$

Original $G = (V,E,\Sigma,\Pi=(\pi_e)_{e\in E})$ ➜ New $G' = (V,E',\Sigma',\Pi')$

Asgmt. $\psi: V\to\Sigma$ ➜ Asgmt. $\psi': V\to\Sigma^V$

for simplicity

⚠ G must be EXPANDER

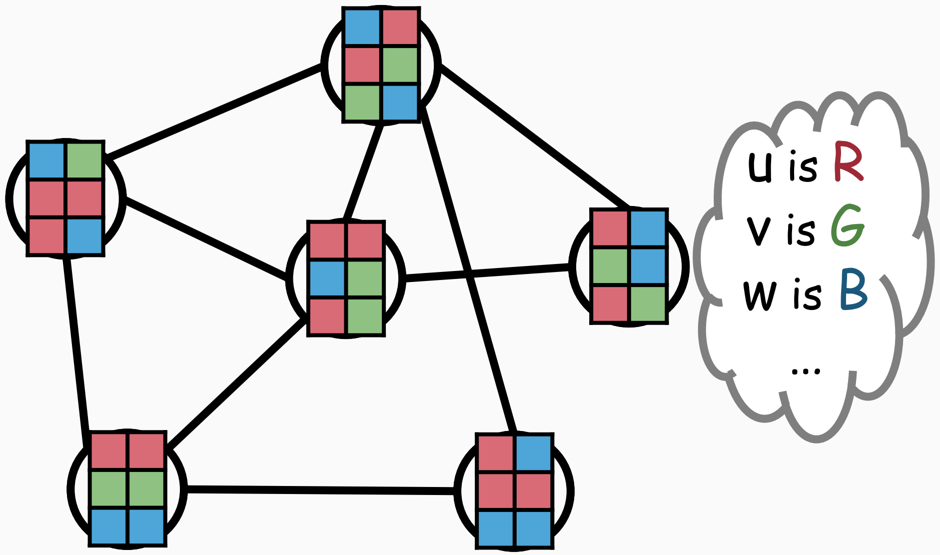

- $\psi'(x)[v] \overset{\text{def}}{=}$ "opinion" of $\psi'(x)$ about the value of v
- edge of G' = a length-T random walk over G

const. parameter

Recap: Dinur's gap amplification [Dinur. J. ACM 2007]

Powering step

Say 3-Coloring $\Sigma=\{R,G,B\}$

Original $G = (V,E,\Sigma,\Pi=(\pi_e)_{e\in E})$ ➜ New $G' = (V,E',\Sigma',\Pi')$

Asgmt. $\psi: V\to\Sigma$ ➜ Asgmt. $\psi': V\to\Sigma^V$

for simplicity

⚠ G must be EXPANDER

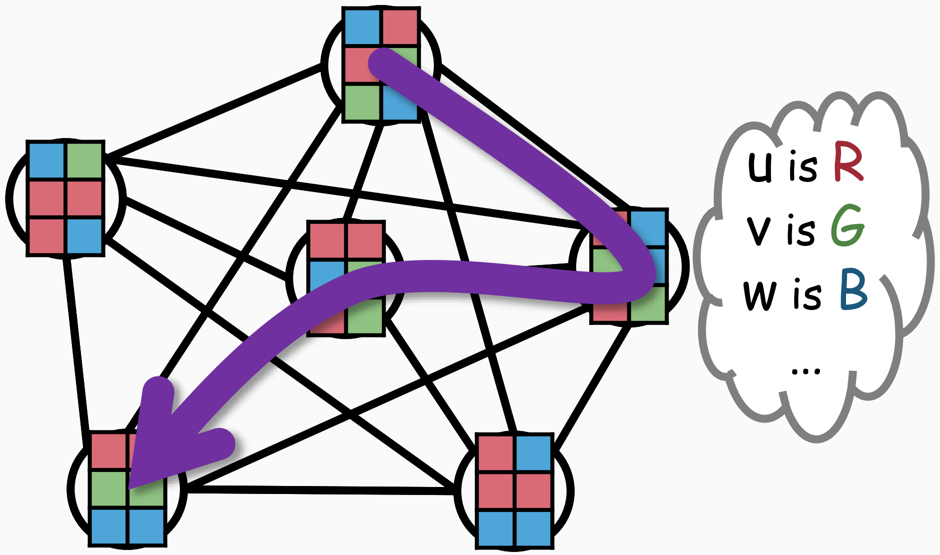

- $\psi'(x)[v] \overset{\text{def}}{=}$ "opinion" of $\psi'(x)$ about the value of v
- edge of G' = a length-T random walk over G

const. parameter

Recap: Dinur's gap amplification [Dinur. J. ACM 2007]

Verifier's test on G' [Radhakrishnan. ICALP 2006]

Pick a random walk **W** $= \langle e_1, \ldots, e_T \rangle$ from x to y

$\psi'(x)$ & $\psi'(y)$ pass the test at $e_i = (v,w)$ if

x & y agree on color of (v,w)
opinions about (v,w) satisfy $\pi_{(v,w)}$

ψ' satisfies π'_W $\overset{\text{def}}{\iff}$ $\psi'(x)$ & $\psi'(y)$ pass test at every edge in **W**

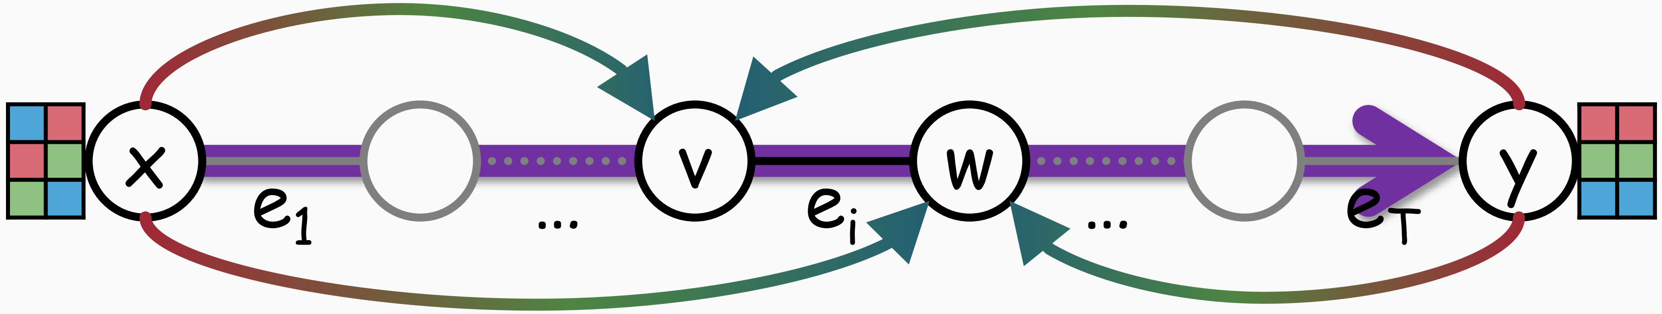

Recap: Dinur's gap amplification [Dinur. J. ACM 2007]

Verifier's test on G' [Radhakrishnan. ICALP 2006]

Pick a random walk **W** $= \langle e_1, \ldots, e_T \rangle$ from x to y

ψ'(x) & ψ'(y) pass the test at $e_i = (v,w)$ if

- ψ'(x)[v] = ψ'(y)[v]
- ψ'(x)[w] = ψ'(y)[w]
- (ψ'(x)[v], ψ'(x)[w]) satisfies e_i

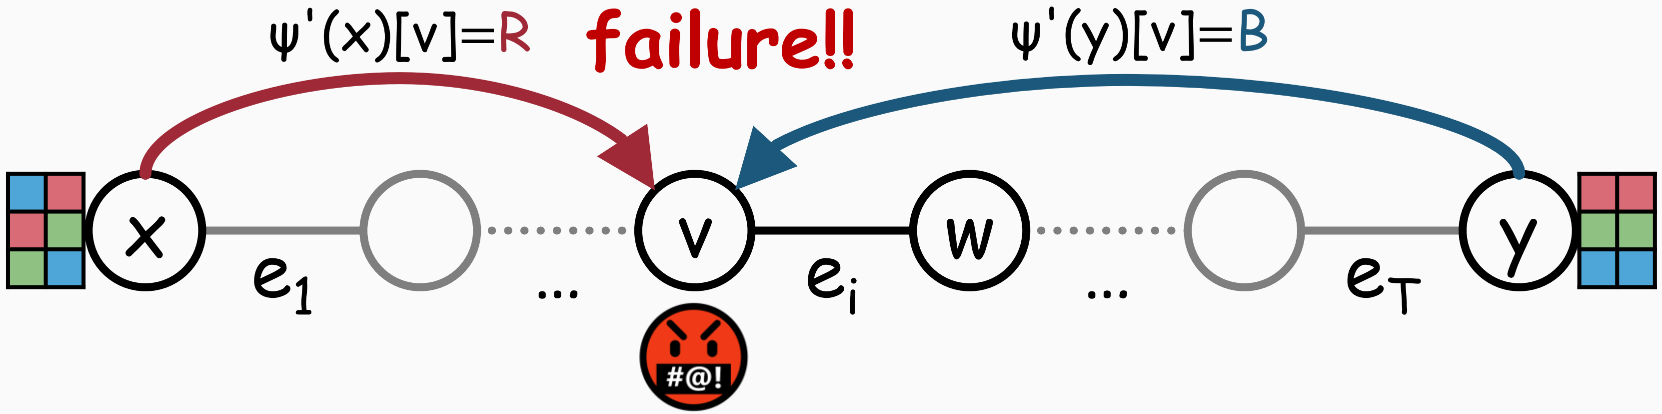

Recap: Dinur's gap amplification [Dinur. J. ACM 2007]

Verifier's test on G' [Radhakrishnan. ICALP 2006]

Pick a random walk **W** $= \langle e_1, \ldots, e_T \rangle$ from x to y

ψ'(x) & ψ'(y) pass the test at $e_i = (v,w)$ if

- ψ'(x)[v] = ψ'(y)[v]
- ψ'(x)[w] = ψ'(y)[w]
- (ψ'(x)[v], ψ'(x)[w]) satisfies e_i

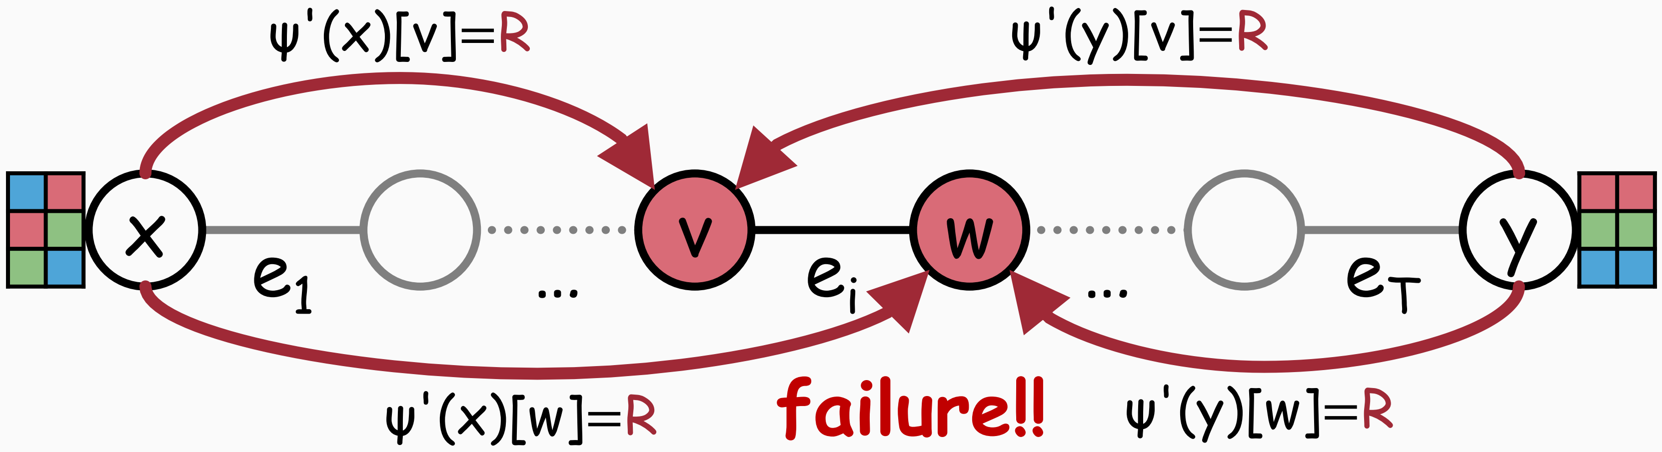

Recap: Dinur's gap amplification [Dinur. J. ACM 2007]

Verifier's test on G' [Radhakrishnan. ICALP 2006]

Pick a random walk **W** $= \langle e_1, \ldots, e_T \rangle$ from x to y

$\psi'(x)$ & $\psi'(y)$ pass the test at $e_i = (v,w)$ if

- $\psi'(x)[v] = \psi'(y)[v]$
- $\psi'(x)[w] = \psi'(y)[w]$
- $(\psi'(x)[v], \psi'(x)[w])$ satisfies e_i

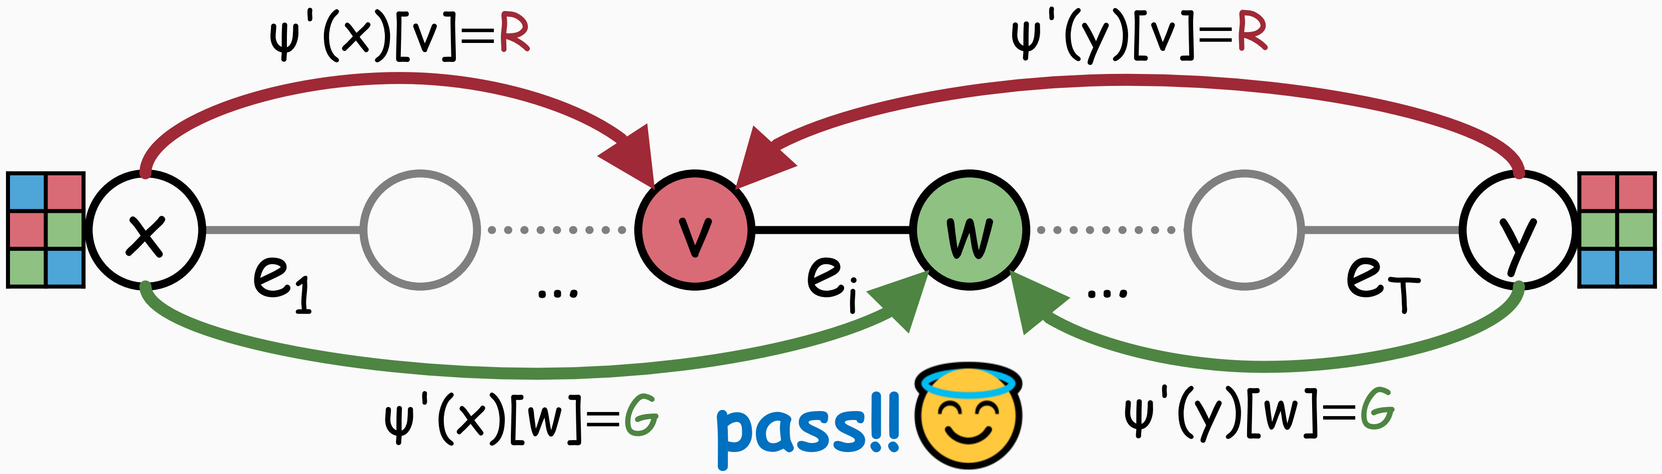

Recap: Maxmin Binary CSP Reconfiguration

[Ito et al. Theor. Comput. Sci. 2011] [O. STACS 2023]

- **Input**: Binary CSP $G = (V, E, \Sigma, \Pi=(\pi_e)_{e \in E})$ & satisfying $\psi_s, \psi_t \colon V \to \Sigma$
- **Output**: $\psi = \langle \psi^{(0)}=\psi_s, \ldots, \psi^{(\ell)}=\psi_t \rangle$ (reconf. sequence) s.t.

  ~~ψ satisfies all edges of G~~ (feasibility)

  $\mathrm{Ham}(\psi^{(i-1)}, \psi^{(i)}) = 1$ (adjacency on hypercube)
- **Goal**: $\max_\psi \mathrm{val}_G(\psi) \stackrel{\text{def}}{=} \min_i$ (frac. of edges satisfied by $\psi^{(i)}$)

  $\mathrm{OPT}_G(\psi_s \leftrightsquigarrow \psi_t) \stackrel{\text{def}}{=}$ max. value of ↑

Under **RIH**, $\exists \varepsilon$ Gap[1 vs. 1-ε] Binary CSP Reconf. is **PSPACE**-hard:

- $\mathrm{OPT}_G(\psi_s \leftrightsquigarrow \psi_t) = 1$ ($\exists \psi$ every $\psi^{(i)}$ satisfies all edges), or
- $\mathrm{OPT}_G(\psi_s \leftrightsquigarrow \psi_t) < 1-\varepsilon$ ($\forall \psi$ some $\psi^{(i)}$ violates ε-frac. of edges)

Barrier of gap amplification for Binary CSP Reconf.

😭Loosing perfect completeness

🎯**Goal:** $\mathrm{OPT}_G(\psi_s \leftrightsquigarrow \psi_t) = 1$ ❌⇒ $\mathrm{OPT}_{G'}(\psi'_s \leftrightsquigarrow \psi'_t) = 1$

All vertices should have the SAME opinion about v's value

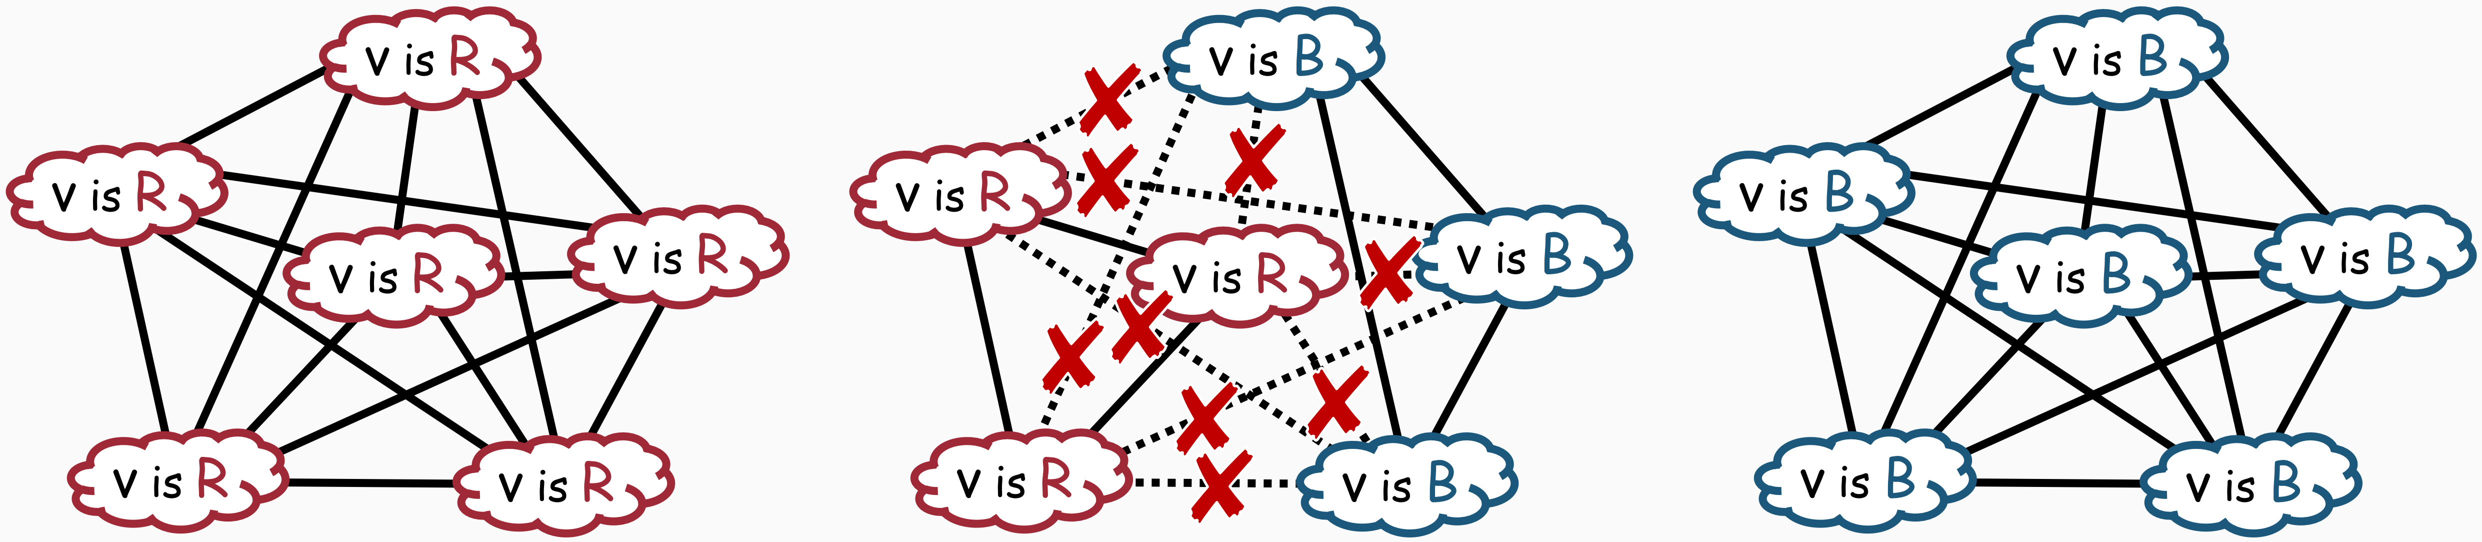

$\forall x \ \ \psi'_s(x)[v] \stackrel{\text{def}}{=} R$ $\exists x,y \ \ \psi'^{(i)}(x)[v] \neq \psi'^{(i)}(y)[v]$ $\forall x \ \ \psi'_t(x)[v] \stackrel{\text{def}}{=} B$

😖Verifier rejects⤴

Our solution

Alphabet squaring trick [O. STACS 2023]

🎯Think as if opinion could take a pair of values!

- Original $\Sigma = \{R, G, B\}$
- New $\Sigma_{sq} = \{R, G, B, \textbf{RG}, \textbf{GB}, \textbf{BR}\}$
- α & β are **consistent** ⇔ α⊆β or α⊇β

	R	RG	G	GB	B	BR
R	●	●				●
RG	●	●	●			
G		●	●	●		
GB			●	●	●	
B				●	●	●
BR	●				●	●

Our solution

Modifying verifier's test

🎯Think as if opinion could take a pair of values!

- Original $\Sigma = \{R, G, B\}$
- New $\Sigma_{sq} = \{R, G, B, \mathbf{RG}, \mathbf{GB}, \mathbf{BR}\}$
- α & β are **consistent** $\Leftrightarrow$ $\alpha \subseteq \beta$ or $\alpha \supseteq \beta$

	R	RG	G	GB	B	BR
R	●	●				●
RG	●	●	●			
G		●	●	●		
GB			●	●	●	
B				●	●	●
BR	●				●	●

Pick RW **W** $= \langle e_1, \ldots, e_T \rangle$ from x to y as before

$\psi'(x)$ & $\psi'(y)$ pass modified test at $e_i = (v,w)$ if

opinions of x & y are **consistent** at (v,w)

opinions about (v,w) satisfy $\pi_{(v,w)}$

Our solution

Modifying verifier's test

🎯Think as if opinion could take a pair of values!

- Original $\Sigma = \{R, G, B\}$
- New $\Sigma_{sq} = \{R, G, B, RG, GB, BR\}$
- α & β are **consistent** $\Leftrightarrow$ $\alpha \subseteq \beta$ or $\alpha \supseteq \beta$

	R	RG	G	GB	B	BR
R	●	●				●
RG	●	●	●			
G		●	●	●		
GB			●	●	●	
B				●	●	●
BR	●				●	●

Pick RW **W** $= \langle e_1, \ldots, e_T \rangle$ from x to y as before

$\psi'(x)$ & $\psi'(y)$ <u>pass modified test</u> at $e_i = (v,w)$ if

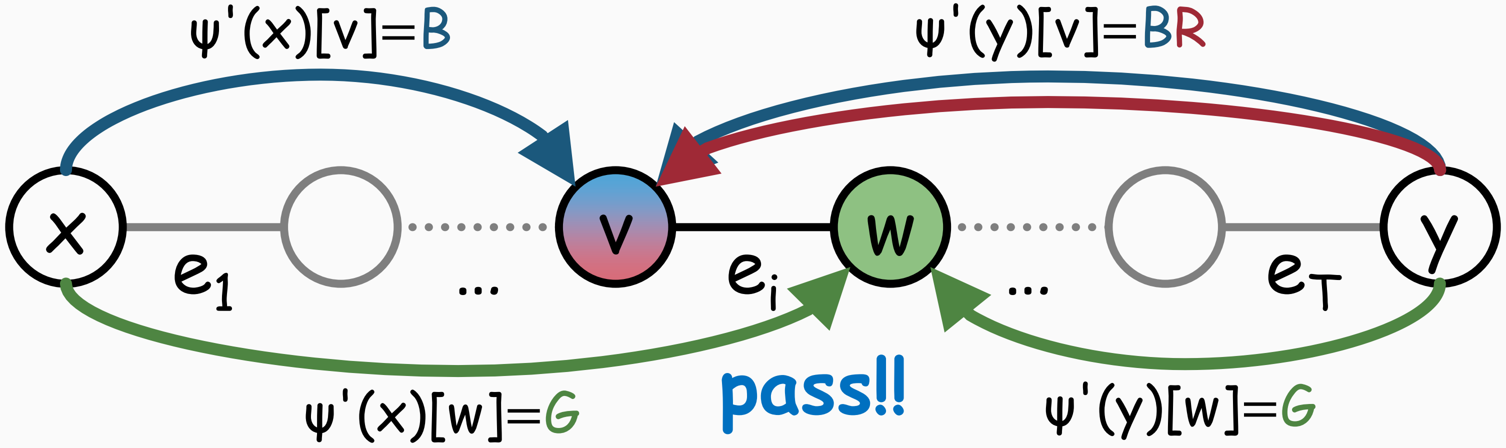

Our solution

Modifying verifier's test

🎯Think as if opinion could take a pair of values!

- Original $\Sigma = \{R, G, B\}$
- New $\Sigma_{sq} = \{R, G, B, RG, GB, BR\}$
- α & β are **consistent** $\Leftrightarrow$ $\alpha \subseteq \beta$ or $\alpha \supseteq \beta$

	R	RG	G	GB	B	BR
R	●	●				●
RG	●	●	●			
G		●	●	●		
GB			●	●	●	
B				●	●	●
BR	●				●	●

Pick RW **W** $= \langle e_1, \ldots, e_T \rangle$ from x to y as before

$\psi'(x)$ & $\psi'(y)$ pass modified test at $e_i = (v,w)$ if

(C1) $\psi'(x)[v]$ & $\psi'(y)[v]$ are **consistent**

(C2) $\psi'(x)[w]$ & $\psi'(y)[w]$ are **consistent**

(C3) $(\psi'(x)[v] \cup \psi'(y)[v]) \times (\psi'(x)[w] \cup \psi'(y)[w]) \subseteq \pi_{(v,w)}$

⚠️This verifier is "much weaker" than before

Our solution

😋 Σ_{sq} preserves perfect completeness

🎯 **Goal:** $\mathrm{OPT}_G(\psi_s \leftrightsquigarrow \psi_t) = 1 \implies \mathrm{OPT}_{G'}(\psi'_s \leftrightsquigarrow \psi'_t) = 1$

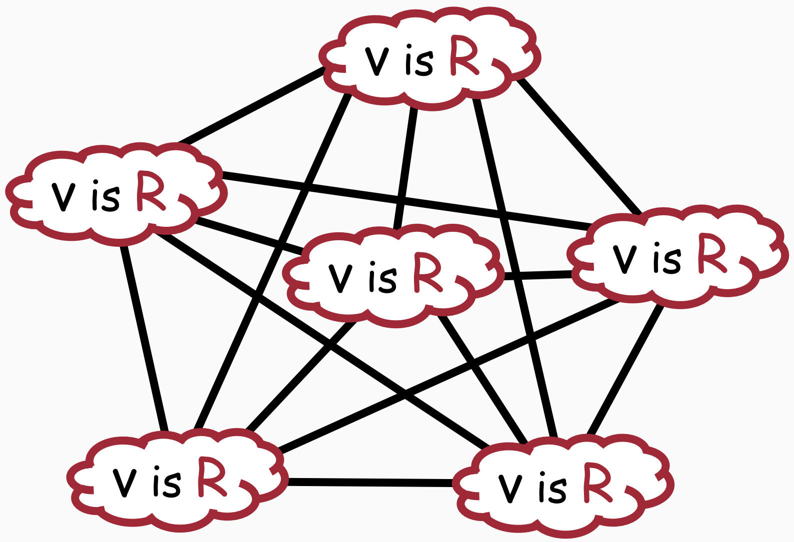

Can transform all R opinions into all B opinions via BR's

Our solution

😋 Σ_{sq} preserves perfect completeness

🎯 **Goal:** $\mathrm{OPT}_G(\psi_s \leftrightsquigarrow \psi_t) = 1 \quad \Longrightarrow \quad \mathrm{OPT}_{G'}(\psi'_s \leftrightsquigarrow \psi'_t) = 1$

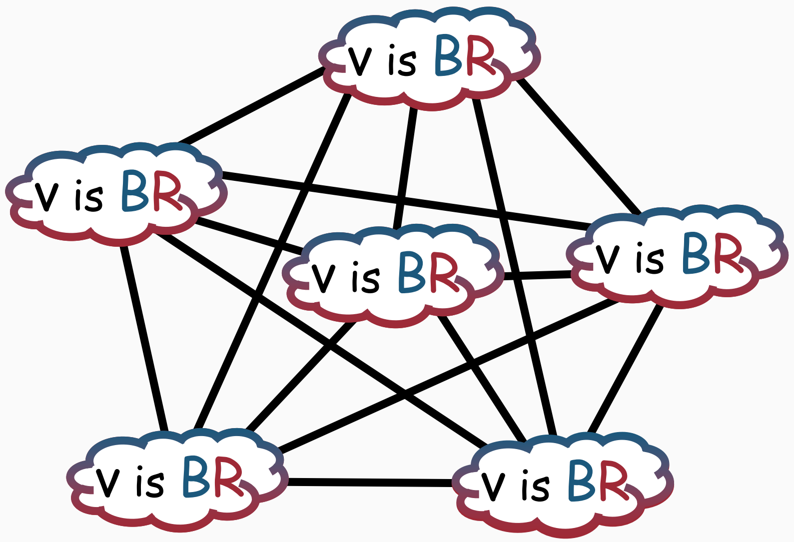

Can transform all R opinions into all B opinions via BR's

Our solution

😋 Σ_{sq} preserves perfect completeness

🎯 **Goal:** $\mathrm{OPT}_G(\psi_s \leftrightsquigarrow \psi_t) = 1 \implies \mathrm{OPT}_{G'}(\psi'_s \leftrightsquigarrow \psi'_t) = 1$

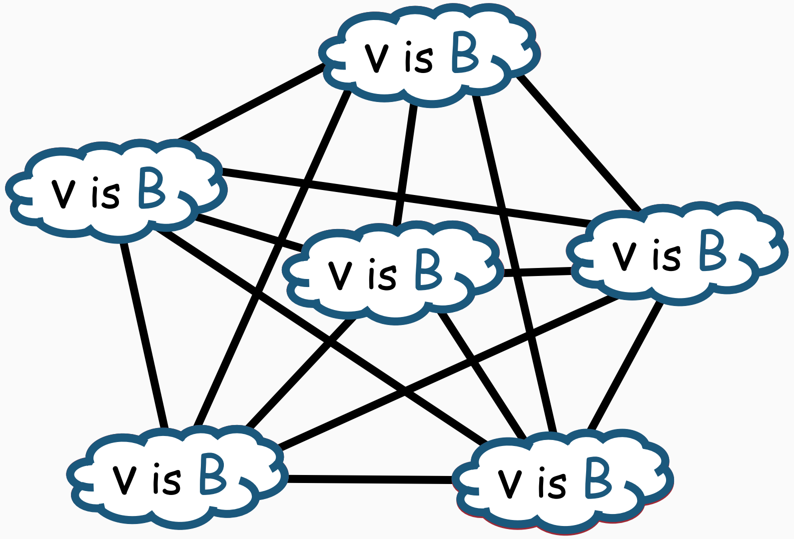

Can transform all R opinions into all B opinions via BR's

Our solution

Soundness STILL works

🎯 **Goal:** $OPT_G(\psi_s \leftrightsquigarrow \psi_t) < 1\text{-}\varepsilon \Longrightarrow OPT_{G'}(\psi'_s \leftrightsquigarrow \psi'_t) < 1 - \Omega(T\cdot\varepsilon)$

$\psi = \langle \psi^{(0)}, \ldots, \psi^{(\ell)} \rangle$ ⇠ Optimal $\psi' = \langle \psi'^{(0)}, \ldots, \psi'^{(\ell)} \rangle$

plurality vote

- We KNOW "$\exists i \; val_G(\psi^{(i)}) < 1\text{-}\varepsilon\text{+}o(1)$"
- Suppose $\psi^{(i)}$ violates (v,w) of G

Pr[$\psi'^{(i)}$ fails modified test at (v,w) | **W** touches (v,w)] = $\Omega(1)$

⚠ DIFFERENT from [Radhakrishnan. ICALP 2006]

$\because \psi'^{(i)}: V \to (\Sigma_{sq})^V$ but $\psi^{(i)}: V \to \Sigma$

{R, G, B, RG, GB, BR} {R, G, B}

Conclusions & open problems

- Gap amplification for Binary CSP Reconf. à la Dinur

- Optimal inapproximability? Maybe $\frac{1}{2}$
- Other gap amplification techniques?
- Alphabet reduction?