Lower bounds
JCRAA 2017

Aurélie Lagoutte

According to *Linear-in-Δ Lower bounds in the LOCAL model* [Göös, Hirvonen, Suomela]
and *No sublog-time approx scheme for Bipartite Vertex Cover* [Göös, Suomela]

LIMOS, Univ. Clermont Auvergne

October 3, 2017
Grenoble

Locality of a distributed algorithm

In LOCAL model with run-time t:

- Graph $G = (V, E)$ a network
- Nodes synchronously send messages to their neighbours
- In t rounds: the output at node v depends only on the t-neighborhood of v
  /!\Edge e incident to v: at distance 1
- Could be deterministic, randomized....

Different models of networks

- ID model: each node has a unique identifier from $\mathbb{N}$.
- OI model (*Order-invariance*): total order $<$ on the nodes
- PO model (*Port numbering and Orientation*): each node has a numbering of its incident edges + edges are oriented
- EC model (*Edge-Colouring*): proper edge-colouring

PO and EC models are anonymous

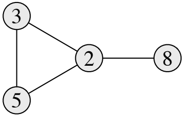

ID

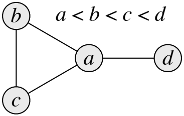

OI

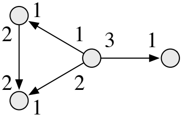

PO

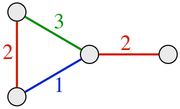

EC

PO can be seen as "special" edge-colouring:

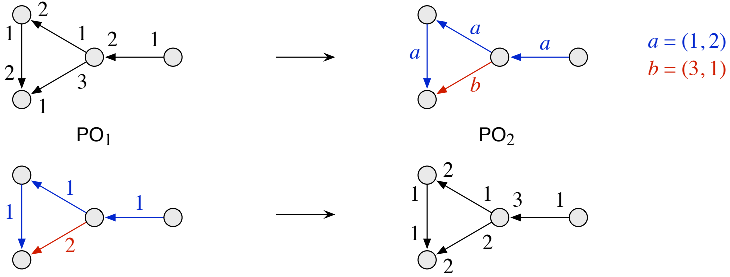

Lifts

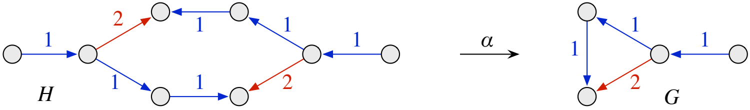

H is a lift of G via covering map α: α preserves degree and edge-colours

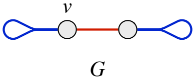

$d(v) = 1 + 1 = 2$

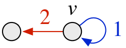

$d^+(v) = 1 + 1$ $\qquad$ $d^-(v) = 1$

In an anonymous model: an algorithm cannot distinguish between a graph and its lifts.

Two special lifts

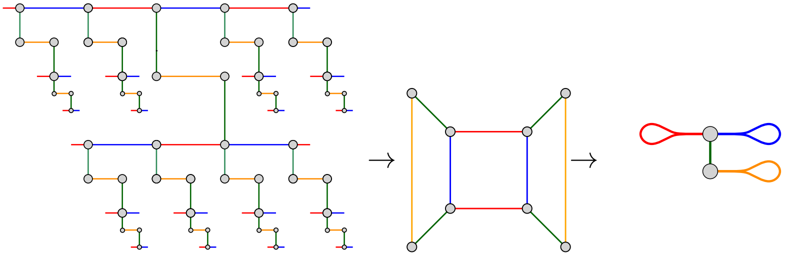

Universal cover U_G G Factor graph F_G

Maximal vs maximum Fractional Matching

Fractional matching: weight function $w : E(G) \to [0,1]$ such that:

$$\forall v \in V(G) \qquad \sum_{\substack{e \in E: \\ v \text{ incid. to } e}} w(e) \leq 1$$

A vertex v is *saturated* if its inequality is tight (=1).

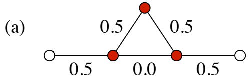

Maximum:
total weight=2

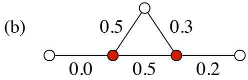

Maximal:
every edge has ≥ 1
endpoint saturated

Maximal Frac. Matching and Loopiness

Definition: loopiness

G is *k-loopy* if every node has $\geq k$ loops.
loopy=1-loopy

Maximal Frac. Matching and Loopiness

Definition: loopiness

G is *k-loopy* if every node has $\geq k$ loops.
loopy=1-loopy

Observation:
If a node v has a loop: it must be saturated by any *EC*-algo.

Lower bound

Theorem

There is no LOCAL algorithm that finds a maximal fractional matching in $o(\Delta)$ rounds (where Δ is the max. degree).

Lower bound

Theorem

There is no LOCAL algorithm that finds a maximal fractional matching in $o(\Delta)$ rounds (where Δ is the max. degree).

1. The maximal FM problem cannot be solved in time $o(\Delta)$ on loopy EC-graphs of maximum degree Δ.
2. $ID \rightsquigarrow OI \rightsquigarrow PO \rightsquigarrow EC$

Lower bound in EC

Theorem: step 1

The maximal FM problem cannot be solved in time $o(\Delta)$ on loopy EC-graphs of maximum degree Δ.

Proof:
Let $\mathcal{A}$ be an EC-algo computing maximal FM.
For $i = 0, \ldots, \Delta - 2$, exhibit a pair (G_i, H_i) such that:

- $\exists\ g_i \in V(G_i)$ and $h_i \in V(H_i)$ having the same i-neighbourhood...
- ... but g_i and h_i must have an incident loop e on which $\mathcal{A}$ disagrees
- G_i, H_i are trees apart from loops, and $\Delta - 1 - i$ loopy

Base case

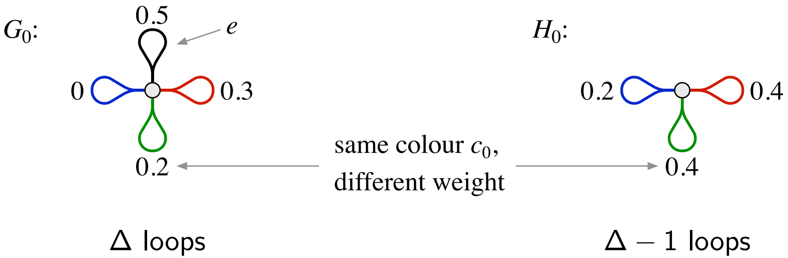

- ⇒ G_0, H_0 are trees apart from loops, and $\Delta - 1$ loopy
- ⇒ g_0 and h_0 have the same 0-neighborhood
- ⇒ g_0 and h_0 have an incident loop e on which $\mathcal{A}$ disagrees

Inductive case

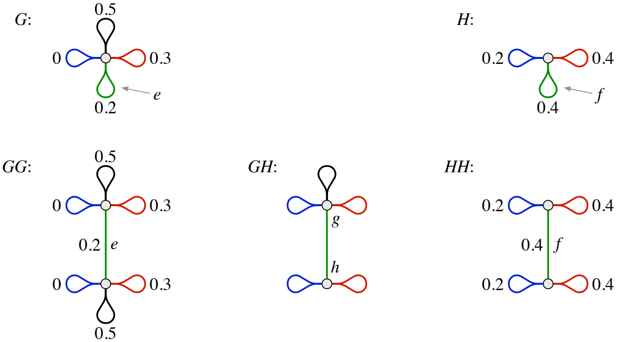

- Unfold $G \to GG$, unfold $H \to HH$
- Unfold & mix $G, H \to GH$
- Choose to keep GH and one of GG, HH

$\Rightarrow G_{i+1}, H_{i+1}$ are trees apart from loops, and $\Delta - 1 - (i+1)$ loopy

Inductive case

Find g_{i+1} and h_{i+1}

- with same $i+1$ neighborhood
- but incident loop e on which $\mathcal{A}$ disagrees

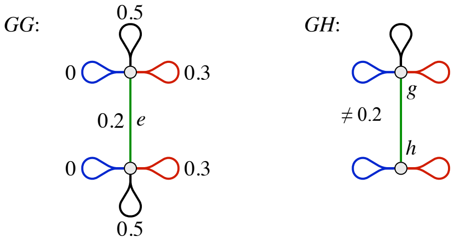

Inductive case

Find g_{i+1} and h_{i+1}

- with same $i+1$ neighborhood
- but incident loop e on which $\mathcal{A}$ disagrees

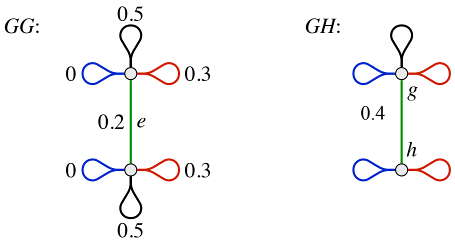

Inductive case

Find g_{i+1} and h_{i+1}

- with same $i+1$ neighborhood
- but incident loop e on which $\mathcal{A}$ disagrees

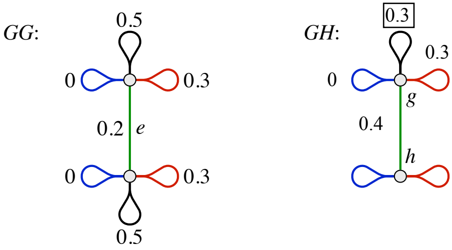

Inductive case

Find g_{i+1} and h_{i+1}

- with same $i+1$ neighborhood
- but incident loop e on which $\mathcal{A}$ disagrees

Push from g_i towards copy of G following edge disagreement.

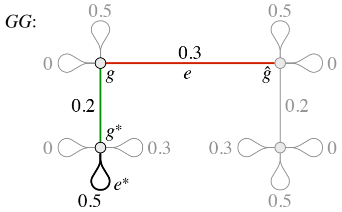

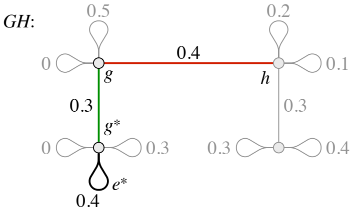

Reductions

$$ID \rightsquigarrow OI \rightsquigarrow PO \rightsquigarrow EC$$

Theorem: step 2

If $\exists\ \mathcal{A}_{ID}$ running in t-time for max FM on any graph of max. degree $\Delta \Rightarrow \exists\ \mathcal{A}_{EC}$ running in t-time computing max. FM on any loopy EC-graph of max. degree $\Delta/2$.

Reductions

$$ID \rightsquigarrow OI \rightsquigarrow PO \rightsquigarrow EC$$

Theorem: step 2

If $\exists\ \mathcal{A}_{ID}$ running in t-time for max FM on any graph of max. degree $\Delta \Rightarrow \exists\ \mathcal{A}_{EC}$ running in t-time computing max. FM on any loopy EC-graph of max. degree $\Delta/2$.

$$PO \rightsquigarrow EC$$

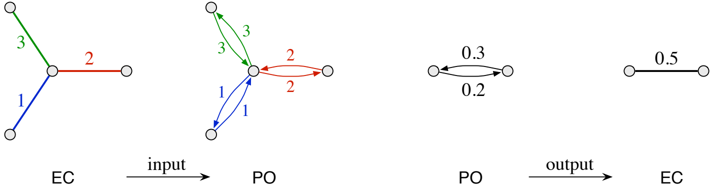

Let T be the infinite $2d$-regular d-edge-coloured PO-tree where d denote the maximum number of edge colours appearing in the input PO-graphs that have max. degree Δ.

Lemma

There is a total order $\prec$ on $V(T)$ such that all the ordered neighborhoods $(T, \prec, v)$ for $v \in V(T)$ are pairwise isomorphic (up to any radius).

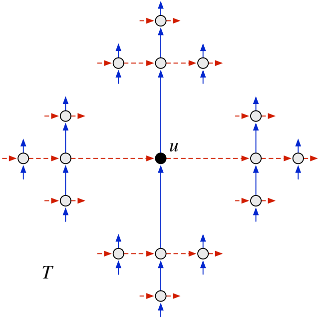

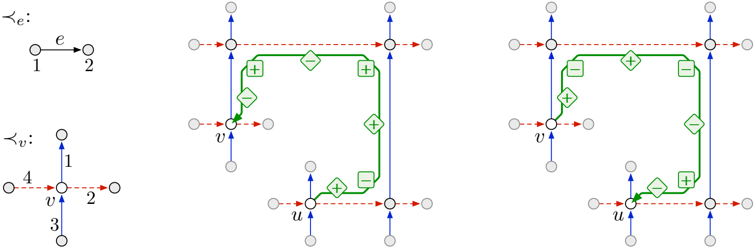

$$u \prec v \Leftrightarrow [\![u \rightsquigarrow v]\!] > 0$$

Clearly antisymmetric and total
Also transitive

$OI \rightsquigarrow PO$

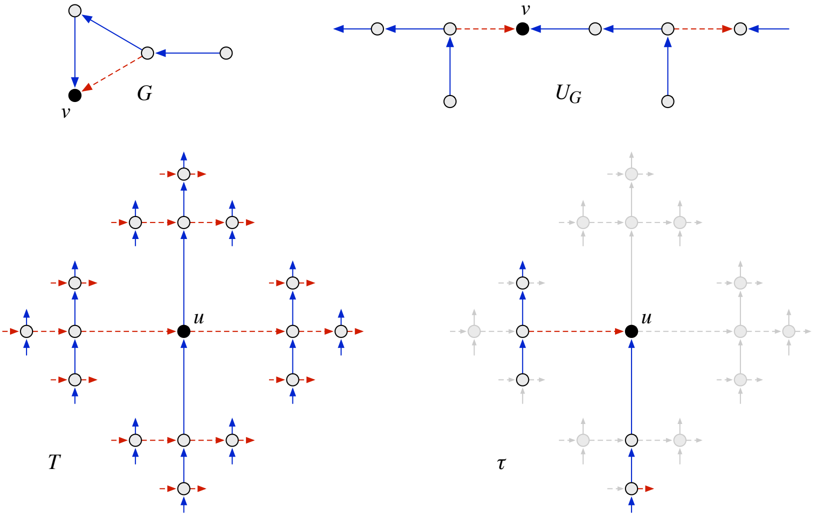

$$ID \rightsquigarrow OI$$

Lemma [Naor and Stockmeyer]

If $\mathcal{A}$ is an ID-algo that outputs finitely many values, there is an infinite set $I \subseteq \mathbb{N}$ such that $\mathcal{A}$ is an OI-algo when restricted to graphs whose identifiers are in I.

Problem: $\mathcal{A}$ any ID-algo for maximal FM $\rightarrow$ infinitely many possible outputs.
Trick: let $\mathcal{A}^*$ simulate $\mathcal{A}$ and outputs 1 if v is saturated, 0 otherwise.

Conclusion

Theorem: step 1

The maximal FM problem cannot be solved in time $o(\Delta)$ on loopy EC-graphs of maximum degree Δ.

Theorem: step 2

If $\exists\ \mathcal{A}_{ID}$ running in t-time for max FM on any graph of max. degree $\Delta \Rightarrow \exists\ \mathcal{A}_{EC}$ running in t-time computing max. FM on any loopy EC-graph of max. degree $\Delta/2$.

Theorem

There is no LOCAL algorithm that finds a maximal fractional matching in $o(\Delta)$ rounds (where Δ is the max. degree).

Another problem, another lower bound

2-Vertex-Cover
Input: a 2-coloured graph G
Output: a minimum-size vertex cover

Model: LOCAL randomised

Theorem: Inapproximability of 2-VC

There exists a $\delta > 0$ such that no randomised distributed algorithm with run-time $o(\log n)$ can find an expected $(1+\delta)$-approx of 2-Vertex-Cover on graphs of maximum degree $\Delta = 3$.

Context

Max matching

$$\max \sum_{e \in E} x_e$$

subj. to: $\sum_{e: v \in e} x_e \leq 1, \forall v$
$x \geq 0$

Vertex cover

$$\min \sum_{v \in V} y_v$$

subj. to: $\sum_{v: v \in e} y_v \geq 1, \forall e$
$y \geq 0$

- Primal and Dual LP admit local $(1+\varepsilon)$-approximation in time $\mathcal{O}_\varepsilon(1)$
- The integral primal problem admits a local constant-time approximation scheme when Δ is bounded.
- Here: $\Omega(\log n)$ lower bound for dual integral problem even with $\Delta = 3$

Toy model & Auxiliary problem

- Deterministic LOCAL
- but anonymous.
- Edges are oriented (not for communication)

RECUT problem
Input: Each node has a color $\ell(v) \in \{\text{red, blue}\}$
Output: An output labelling ℓ_{out} such that

$$\partial \ell_{out} = \frac{\text{nb of edges crossing red-blue}}{|E|}$$

is minimized subject to:
an all-red (resp. all-blue) ℓ must lead to all-red (resp. all blue) ℓ_{out}

RECUT problem

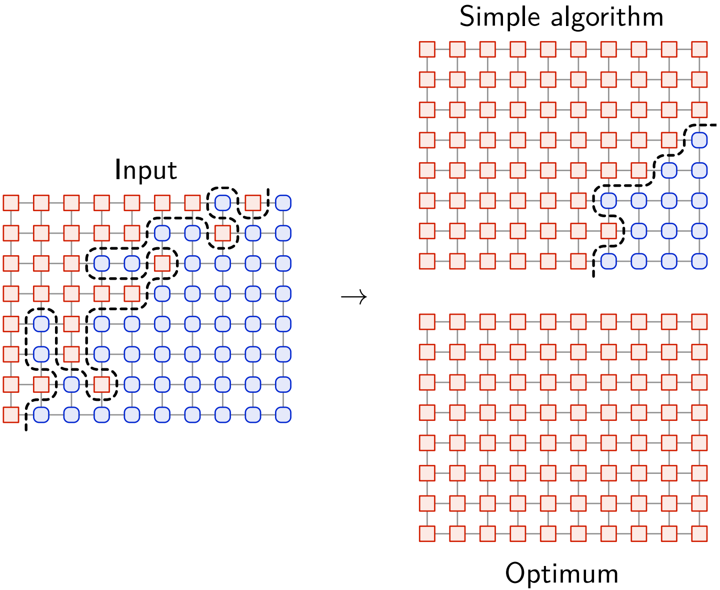

Reduction

$\exists\ \mathcal{A}$ computing a $(1+\varepsilon)$-approx of 2-VC in t rounds when $\Delta = 3$

$\Rightarrow$ $\exists\ \mathcal{A}'$ computing a recut $\ell_{\mathcal{A}}$ of size $\partial \ell_{\mathcal{A}} = O(\varepsilon)$
on balanced 4-regular tree-like digraphs in t rounds

$\Delta = 4$

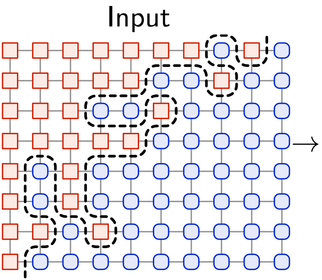

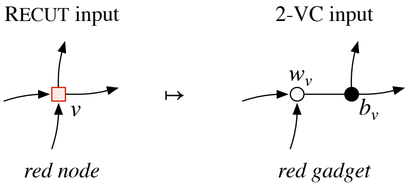

Fig. 2 Red gadget for $\Delta = 4$.

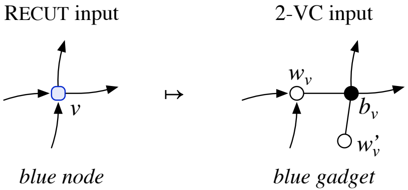

Fig. 3 Blue gadget for $\Delta = 4$ (assuming an all-red input produces an all-white output).

Output

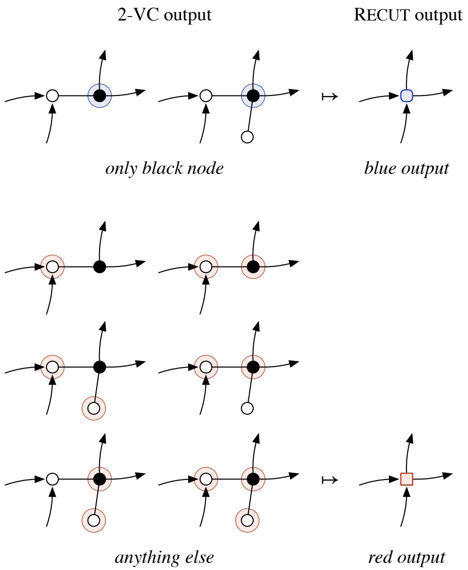

Fig. 4 Mapping the output of $\mathcal{A}$ to a solution of the RECUT problem.

Analysis: bad nodes

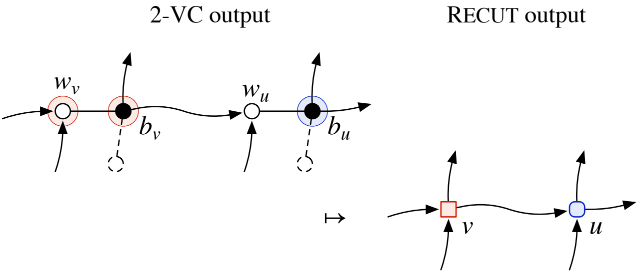

Fig. 5 A bad node: v is red and its out-neighbour u is blue.

Analysis: bad nodes

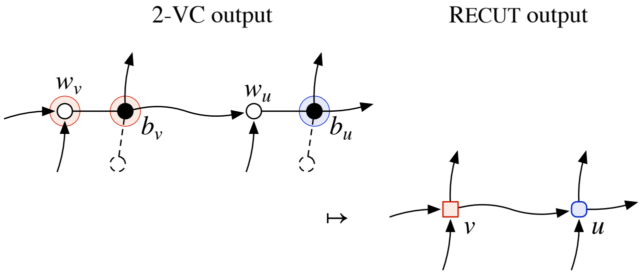

Fig. 5 A bad node: v is red and its out-neighbour u is blue.

Since $\mathcal{A}$ is a $(1+\varepsilon)$-approx: at most $\varepsilon|V|$ bad vertices.
So nb of edges from red to blue $\leq 2\varepsilon|V|$

$$\partial\ell_{\mathcal{A}} = \frac{\text{nb of edges red-blue}}{|E|} = \frac{2 \cdot \text{nb of edges from red to blue}}{|E|} \leq 2\varepsilon$$

$\Delta = 3$

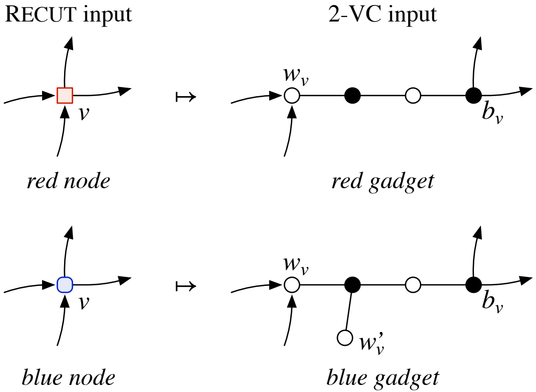

Fig. 6 Gadgets for $\Delta = 3$.

Nearly balanced recut

Lemma

If $\mathcal{A}$ runs in $o(\log n)$ rounds for RECUT then for each 4-regular graph $\exists$ input ℓ for which $\mathcal{A}$ computes a nearly balanced recut $|red| = n/2 \pm o(n)$.

Change of v_i's color seen only by $\leq 4^r + 1 = o(n)$ vertices

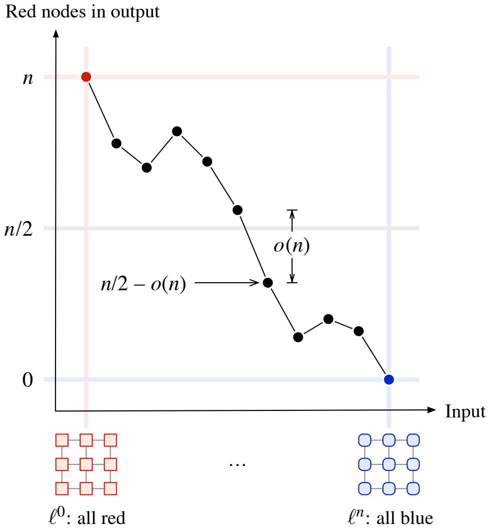

Finally

δ-Expander graph: edge expansion condition

$$e(S, V \setminus S) \geq \delta \cdot |S| \qquad \text{for all } S \subseteq V, |S| \leq n/2$$

- Take $\mathcal{F}$=family of Ramanujan graphs $\delta = (2 - \sqrt{3})$−expanders [Morgenstern 15] having girth $\theta(\log n) \Rightarrow$ tree-like for any $\mathcal{A}$ on $o(\log n)$ rounds.
- 4-regular $\Rightarrow$ Orient with an Euler tour to balanced 4-reg. digraphs
- Apply previous lemma: $\exists$ input ℓ for which $\mathcal{A}$ produces nearly balanced recut
- By expansion property, $\partial\ell_{\mathcal{A}} \geq \delta/4 - o(1)$
- Contrapositive of RECUT$\leq$ 2-VC reduction: no $(1 + \varepsilon)$-approx of 2-VC in $o(\log n)$ rounds

Back to randomised algorithm

$\exists\ \mathcal{A}$ randomised LOCAL algo for 2-VC in $o(\log n)$ rounds?

Back to randomised algorithm

$\exists\ \mathcal{A}$ randomised LOCAL algo for 2-VC in $o(\log n)$ rounds?

Without loss of generality:

- **Deterministic run-time:** Each node runs for at most $t = o(\log n)$ steps
- **Las Vegas algo. :** Never fails

Then with same simulation RECUT $\rightarrow$ 2-VC as before:

$$\exists \text{ input } \ell \text{ s.t. } \mathbb{E}(|\ell_{out}^{-1}(red)|) = n/2 - o(n)$$

Local Concentration Bound [Janson]:
with high proba this number is concentrated around its expectation.