Three Challenges in Distributed Optimization

Keren Censor-Hillel
Technion

Workshop on Local Algorithms WOLA 2019

This project has received funding from the European Union's Horizon 2020
Research and Innovation Programme under grant agreement no. 755839

Optimization problems

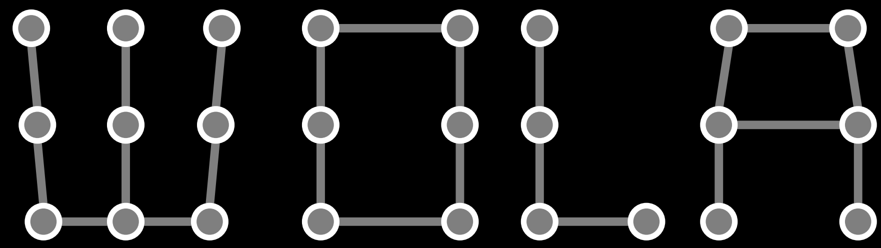

Optimization problems

Minimum vertex cover

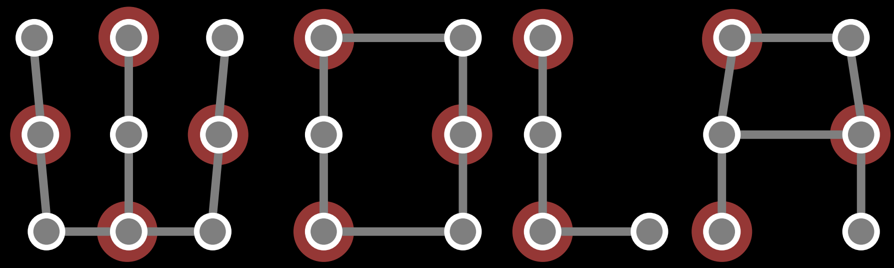

Optimization problems

Minimum vertex cover

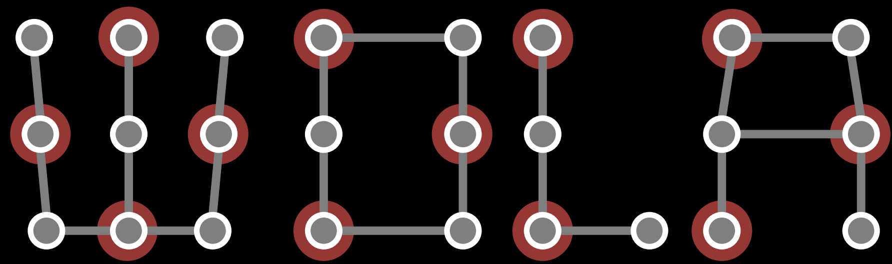

NP-hard

(2-ε)-approximation is UG-hard

(some algs. with a smaller-than-2 apx)

Distributed optimization

Minimum vertex cover

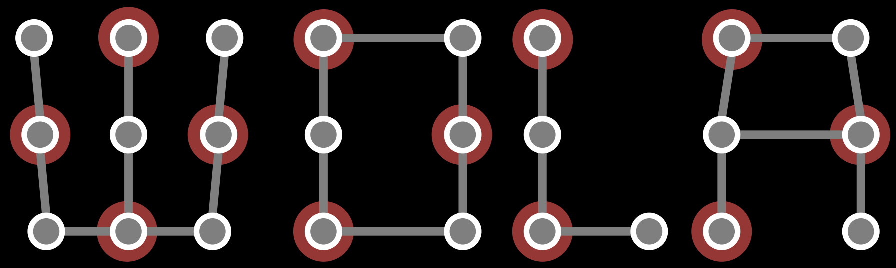

Complexity ?

Distributed graph algorithms

#Nodes = n
#Bandwidth = B
(typically O(logn))

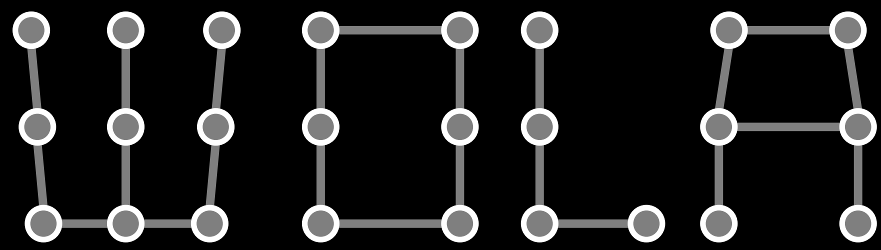

Knowledge of neighbors

CONGEST

#Rounds = ?

Distributed graph algorithms

#Nodes = n
#Bandwidth = B
~~(typically O(logn))~~

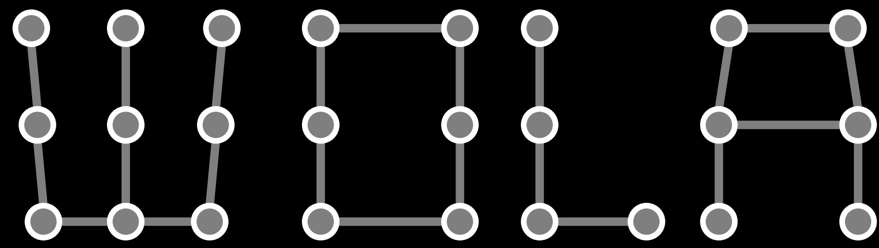

Knowledge of neighbors

Models:
CONGEST
LOCAL

#Rounds = ?

Distributed graph algorithms

#Nodes = n
#Bandwidth = B
~~(typically O(logn))~~

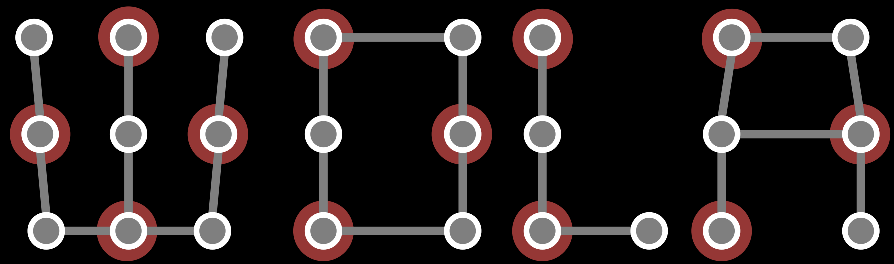

Knowledge of neighbors

#Rounds = ?

Models:

CONGEST
LOCAL

Distributed graph algorithms

#Nodes = n
#Bandwidth = B
~~(typically O(logn))~~

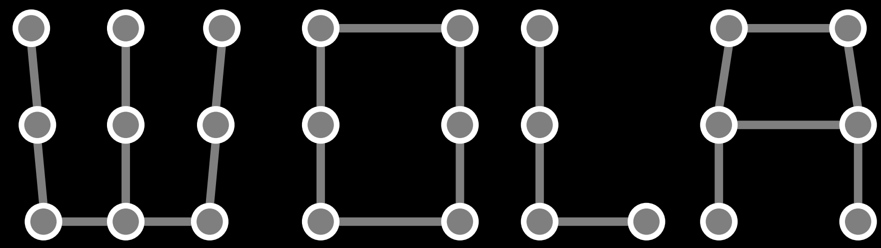

Knowledge of neighbors

Models:

CONGEST	$O(m) = O(n^2)$
LOCAL	$O(D)$

#Rounds = ?

Challenge 1:

Challenge 1: Distances

Distances

Exact minimum vertex cover: $\Omega(D)$ rounds

Distances

Exact minimum vertex cover: $\Omega(D)$ rounds

Approximations:

LOCAL: $(1+\epsilon)$-approximation in $O(poly\,(\log n\,/\epsilon))$ rounds

[Ghaffari, Kuhn, Maus '17]

Distances

Exact minimum vertex cover: $\Omega(D)$ rounds

Approximations:

LOCAL: $(1+\epsilon)$-approximation in $O(poly\,(\log n\,/\epsilon))$ rounds
[Ghaffari, Kuhn, Maus '17]

$\Omega(\log \Delta / \log\log \Delta)$, $\Omega(\sqrt{(\log n\,/\log\log n)})$ rounds
[Kuhn, Moscibroda, Wattenhofer '04]

$\Omega(1/\epsilon)$ [Ben-Basat, Kawarabayashi, Schwartzman '18]

Optimal solutions for pieces of the graph

LOCAL: $(1+\epsilon)$-approximation in $O(poly\ (\log n\ /\epsilon))$ rounds
[Ghaffari, Kuhn, Maus '17]

Come to Yannic's talk
tomorrow!

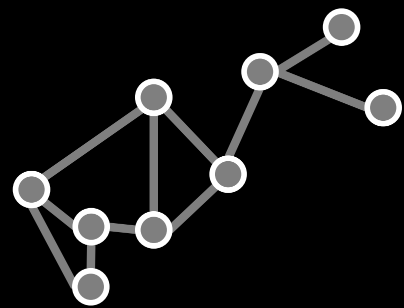

Optimal solutions for pieces of the graph

LOCAL: $(1+\epsilon)$-approximation in $O(poly\,(\log n\,/\epsilon))$ rounds
[Ghaffari, Kuhn, Maus '17]

Come to Yannic's talk
tomorrow!

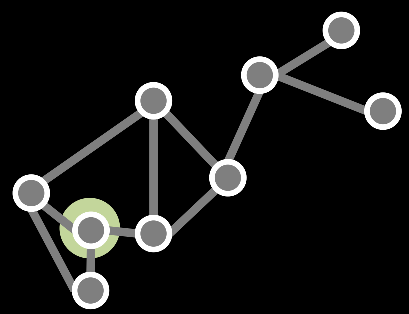

Optimal solutions for pieces of the graph

LOCAL: $(1+\epsilon)$-approximation in $O(poly\,(\log n\,/\epsilon))$ rounds
[Ghaffari, Kuhn, Maus '17]

Come to Yannic's talk tomorrow!

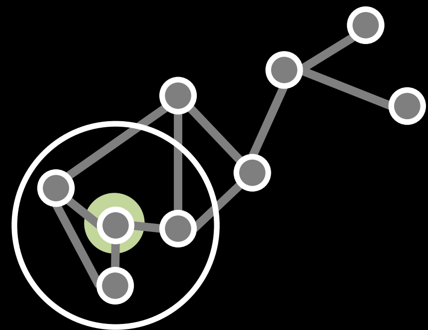

Optimal solutions for pieces of the graph

LOCAL: $(1+\epsilon)$-approximation in $O(poly(\log n/\epsilon))$ rounds
[Ghaffari, Kuhn, Maus '17]

Come to Yannic's talk
tomorrow!

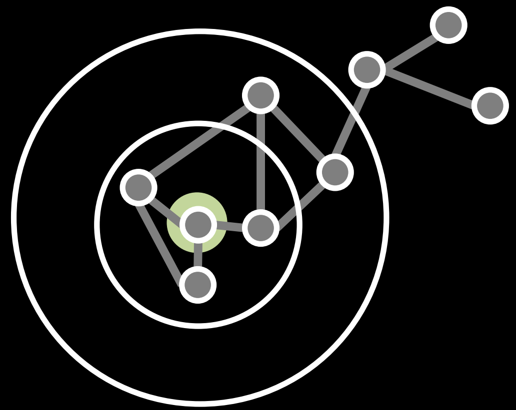

Optimal solutions for pieces of the graph

LOCAL: $(1+\epsilon)$-approximation in $O(poly\,(\log n\,/\epsilon))$ rounds
[Ghaffari, Kuhn, Maus '17]

Come to Yannic's talk
tomorrow!

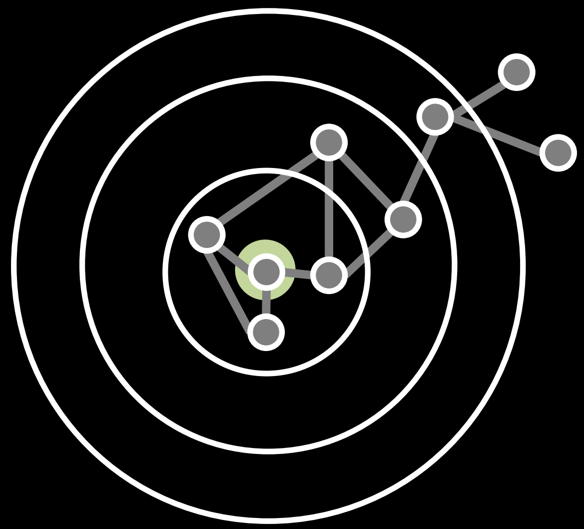

Challenge 2:

Challenge 2: Congestion

Congestion

LOCAL: $(1+\epsilon)$-approximation in $O(poly\ (\log n\ /\epsilon))$ rounds
[Ghaffari, Kuhn, Maus '17]

Cannot collect dense neighborhoods quickly

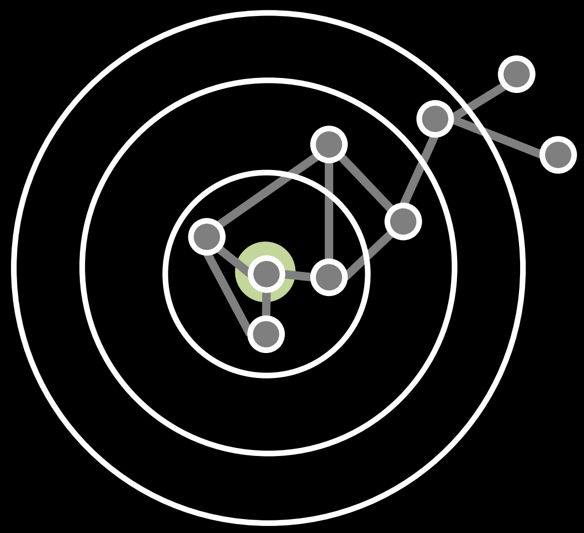

CONGEST

(2+ε)-approximation in $O(\log\Delta/\log\log\Delta)$ rounds
[Bar-Yehuda, C., Schwartzman '16]

$\Omega(\log\Delta/\log\log\Delta), \Omega(\sqrt{(\log n/\log\log n)})$ rounds for approximation
[Kuhn, Moscibroda, Wattenhofer '04]

CONGEST

(2+ε)-approximation in $O(\log \Delta / \log\log \Delta)$ rounds
[Bar-Yehuda, C., Schwartzman '16]

2-approximation [Ben-Basat , Even, Kawarabayashi, Schwartzman '18]

$o(n^2)$, (2-ε)-approximation [Ben-Basat , Kawarabayashi, Schwartzman '18]

$\Omega(\log \Delta / \log\log \Delta)$, $\Omega(\sqrt{(\log n / \log\log n)})$ rounds for approximation
[Kuhn, Moscibroda, Wattenhofer '04]

CONGEST

(2+ε)-approximation in $O(\log\Delta/\log\log\Delta)$ rounds
[Bar-Yehuda, C., Schwartzman '16]

2-approximation [Ben-Basat , Even, Kawarabayashi, Schwartzman '18]

o(n^2), (2-ε)-approximation [Ben-Basat , Kawarabayashi, Schwartzman '18]

$\Omega(\log\Delta/\log\log\Delta)$, $\Omega(\sqrt{(\log n / \log\log n)})$ rounds for approximation
[Kuhn, Moscibroda, Wattenhofer '04]

$\Omega(n^2/poly\log n)$ rounds for exact [C., Khoury, Paz '17]

Minimum vertex cover in CONGEST

#rounds

$\widetilde{\Omega}(n^2)$
[C., Khoury, Paz '17]

[Ben-Basat, Even, Kawarabayashi, Schwartzman '18]

[Kuhn, Moscibroda, Wattenhofer '04]

[Bar-Yehuda, C., Schwartzman '16]

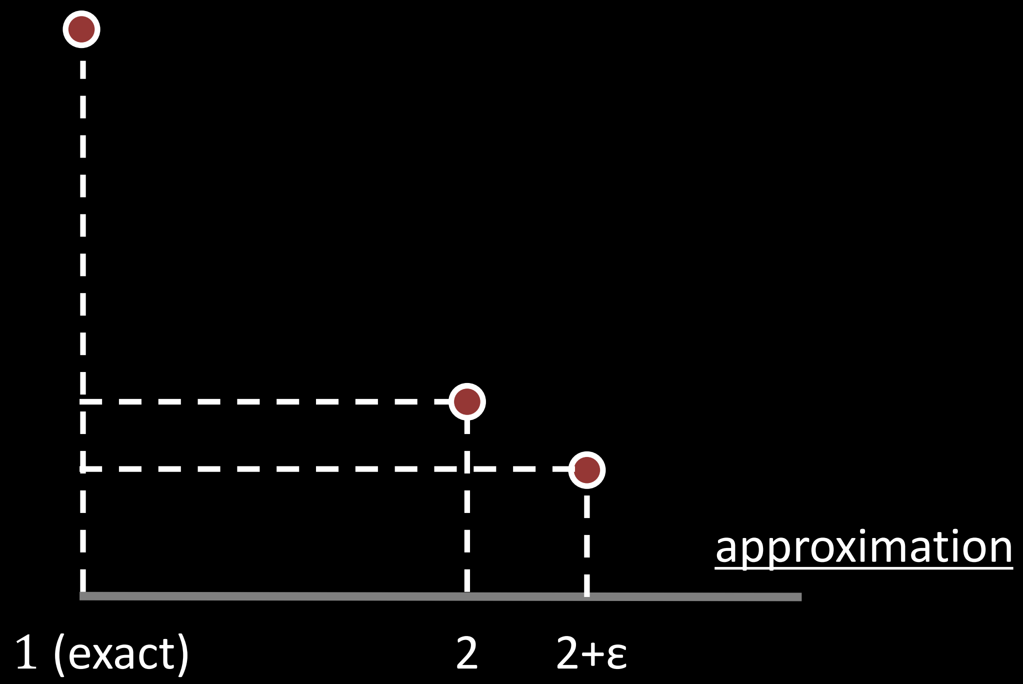

Minimum vertex cover in CONGEST

#rounds

$\widetilde{\Omega}(n^2)$
[C., Khoury, Paz '17]

[Ben-Basat, Even, Kawarabayashi, Schwartzman '18]

[Kuhn, Moscibroda, Wattenhofer '04]

[Bar-Yehuda, C., Schwartzman '16]

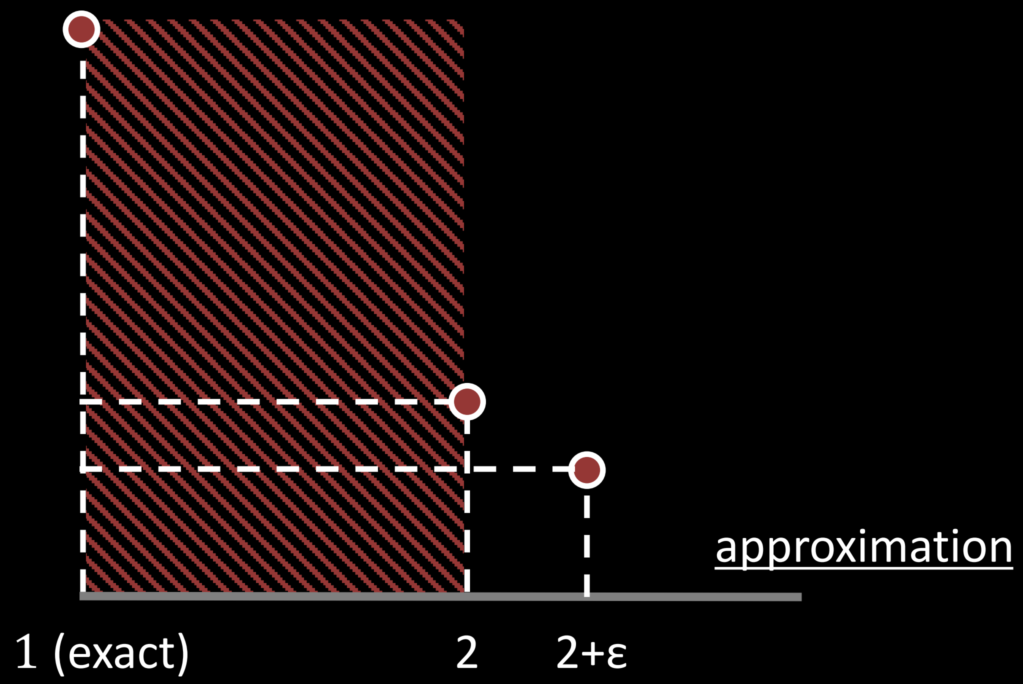

Minimum vertex cover in CONGEST

#rounds

$\tilde{\Omega}(n^2)$
[C., Khoury, Paz '17]

[Bachrach, C., Dory, Efron, Leitersdorf, Paz '19]

[Ben-Basat, Even, Kawarabayashi, Schwartzman '18]

[Kuhn, Moscibroda, Wattenhofer '04]

[Bar-Yehuda, C., Schwartzman '16]

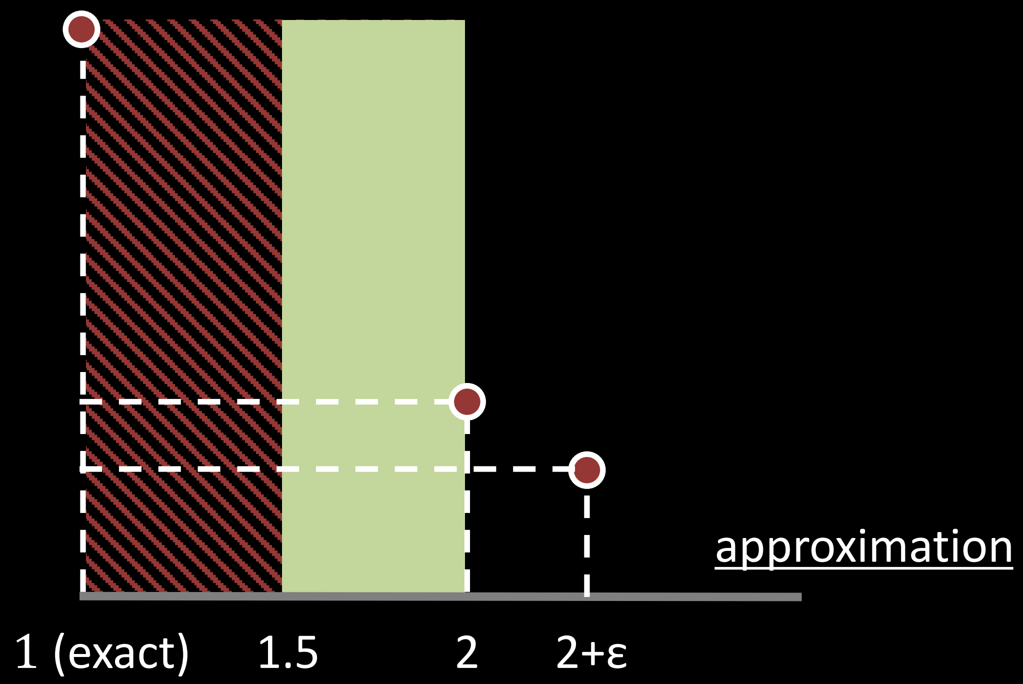

Minimum vertex cover in CONGEST

#rounds

[Bachrach, C., Dory, Efron, Leitersdorf, Paz '19]

$\widetilde{\Omega}(n^2)$
[C., Khoury, Paz '17]

[Ben-Basat, Even, Kawarabayashi, Schwartzman '18]

[Kuhn, Moscibroda, Wattenhofer '04]

[Bar-Yehuda, C., Schwartzman '16]

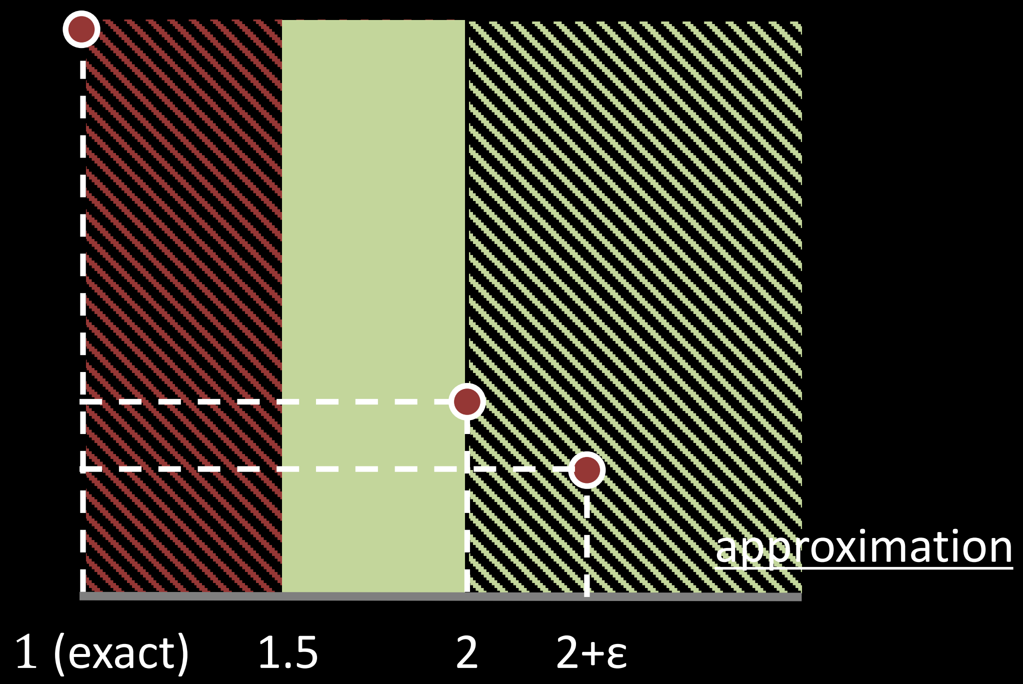

Minimum vertex cover in CONGEST

#rounds

$\widetilde{\Omega}(n^2)$

[C., Khoury, Paz ‘17]

[Bachrach, C., Dory, Efron, Leitersdorf, Paz ‘19]

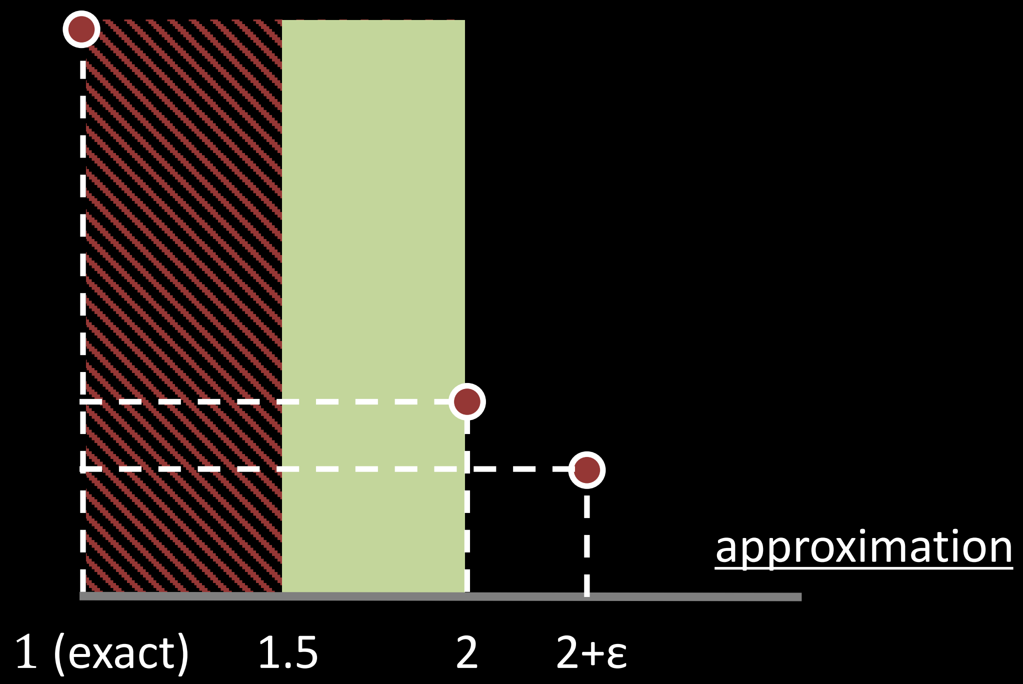

[Ben-Basat, Even, Kawarabayashi, Schwartzman ‘18]

[Kuhn, Moscibroda, Wattenhofer ‘04]

[Bar-Yehuda, C., Schwartzman ‘16]

2-Party communication

[Yao '79]

History

History

[Peleg, Rubinovich ‘97] [Lotker, Patt-Shamir, Peleg ’01] [Elkin ‘04] [Das-Sarma, Holzer, Kor, Korman, Nanongkai, Pandurangan, Peleg, Wattenhofer ‘11] [Frischknecht, Holzer, Wattenhofer ‘12] [Ghaffari, Kuhn ’13] [Drucker, Kuhn, Oshman ‘14] [Nanongkai, Das-Sarma, Pandurangan ‘14] [Das-Sarma, Molla, Pandurangan ‘15] [Holzer, Pinsker, ‘15] [Pandurangan, Peleg, Scquizzato ‘16] [Pandurangan, Robinson, Scquizzato ‘16] [C., Kavitha, Paz, Yehudayoff ’16] [Fischer, Gonen, Kuhn, Oshman ‘18] [C., Dory ‘17] [C., Dory ’18]

2-Party communication

Alice: $x=x_1,\ldots,x_k$ Bob: $y=y_1,\ldots,y_k$

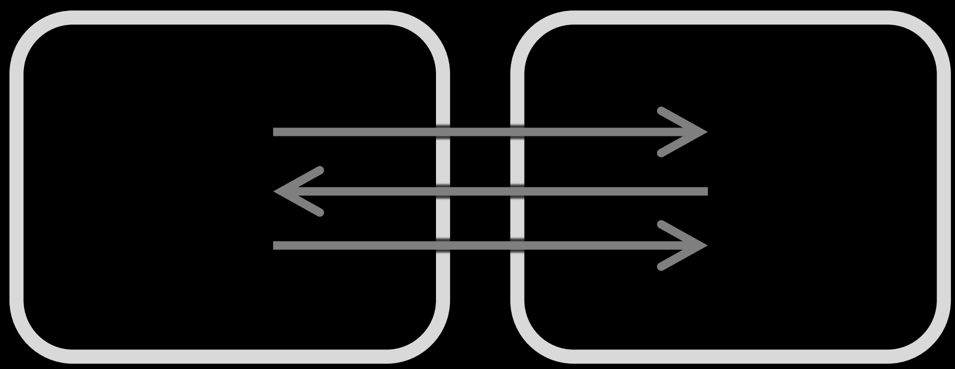

Goal: $f(x,y)$

2-Party communication

Alice: $x=x_1,\ldots,x_k$ Bob: $y=y_1,\ldots,y_k$

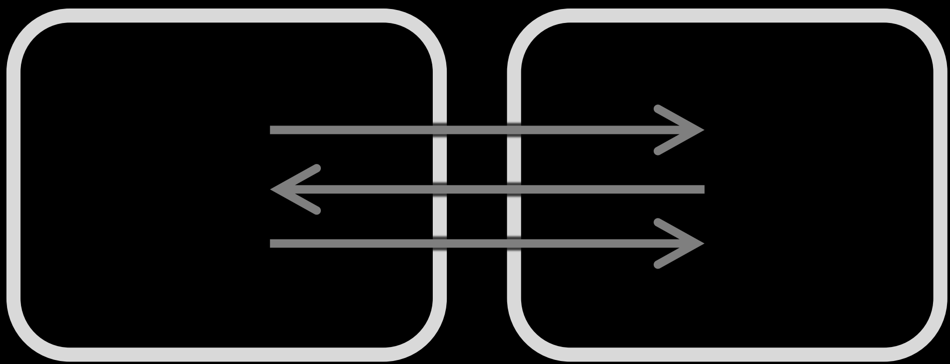

Goal: $f(x,y)$

Set-Disjointess:

$\exists\, i: x_i=y_i=1$?

cost=$\Omega(k)$

[Kalyanasundaram and Schnitger '87]
[Razborov '90]
[Bar-Yossef et al. '04]

2-Party communication → CONGEST

Alice: $x=x_1,\ldots,x_k$ Bob: $y=y_1,\ldots,y_k$

Graph property P of $G_{x,y}$ determines $f(x,y)$

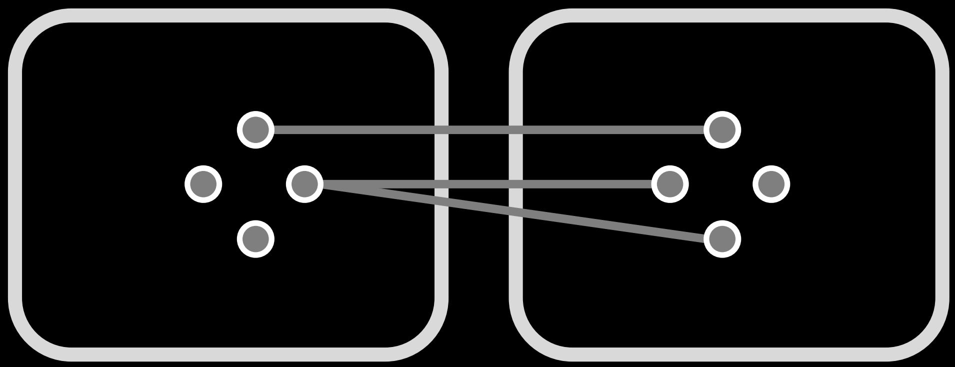

2-Party communication → CONGEST

Alice: $x=x_1,\ldots,x_k$ Bob: $y=y_1,\ldots,y_k$

Graph property P of $G_{x,y}$ determines $f(x,y)$

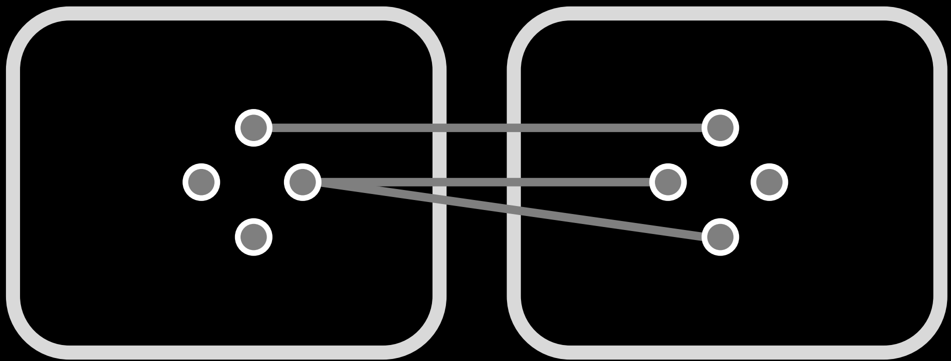

Rounds•cut•B = $\Omega(\text{cost}(f(k)))$

2-Party communication ➔ CONGEST

Alice: $x=x_1,\ldots,x_k$ Bob: $y=y_1,\ldots,y_k$

Graph property P of $G_{x,y}$ determines $f(x,y)$

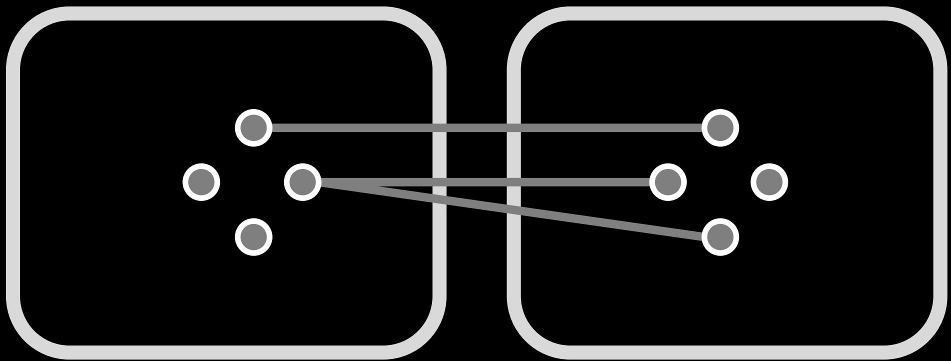

Rounds•**cut**•B = Ω(cost(f(**k**)))

Rounds = Ω(cost(f(**k**))/**cut**•B)

Warm-up: MVC, CONGEST

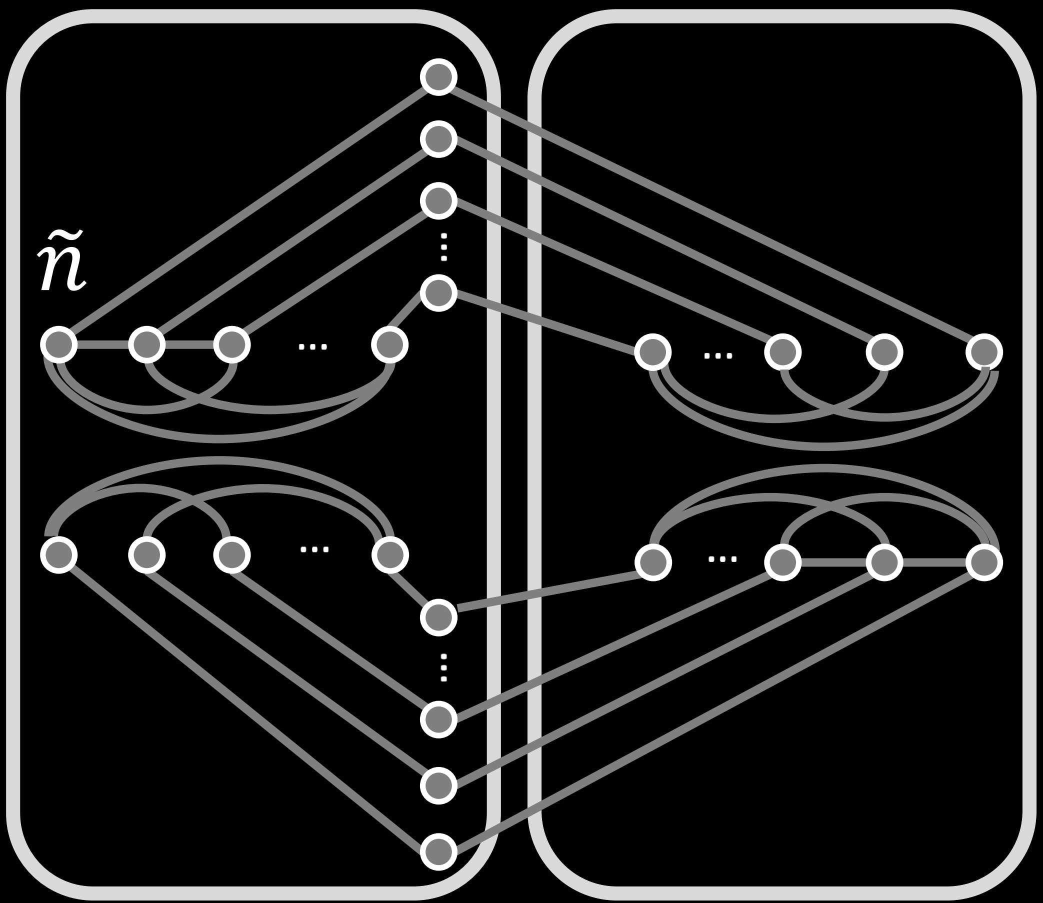

cut = $\tilde{n} = n/6$

Warm-up: MVC, CONGEST

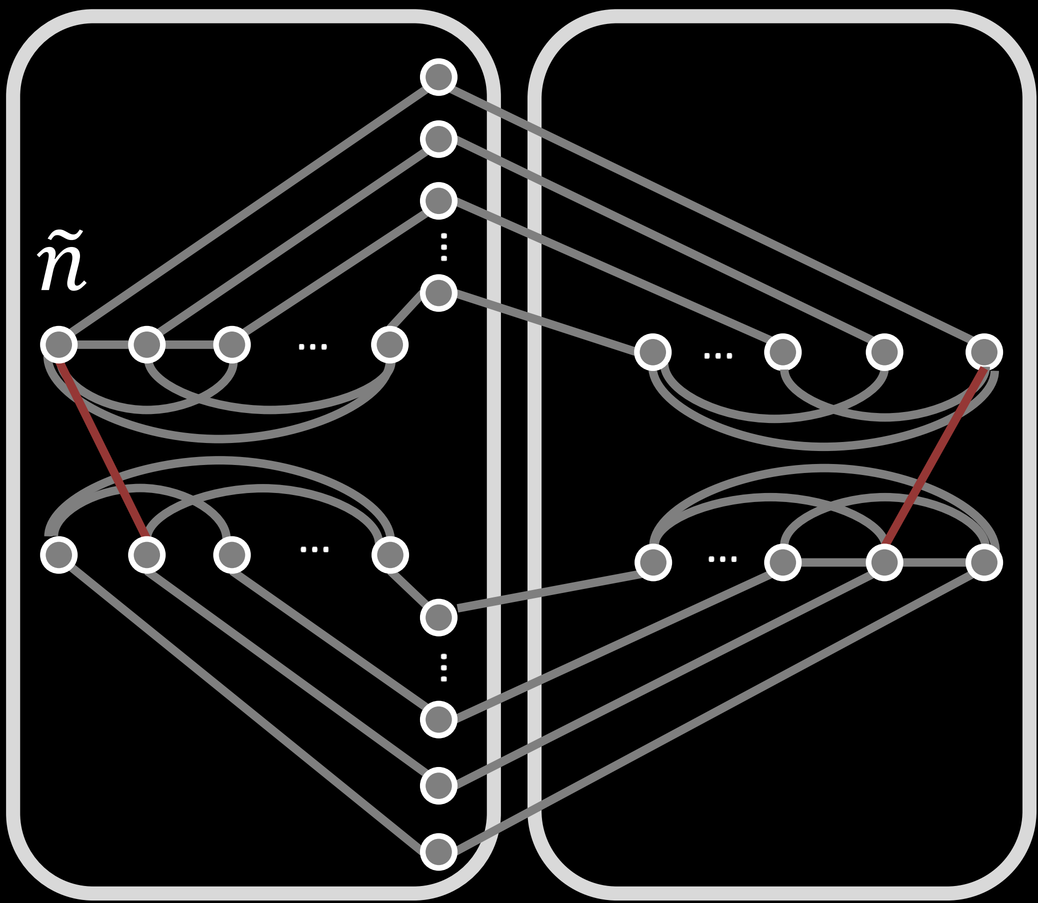

Add edge iff input is 0

cut = $\tilde{n} = n/6$

Warm-up: MVC, CONGEST

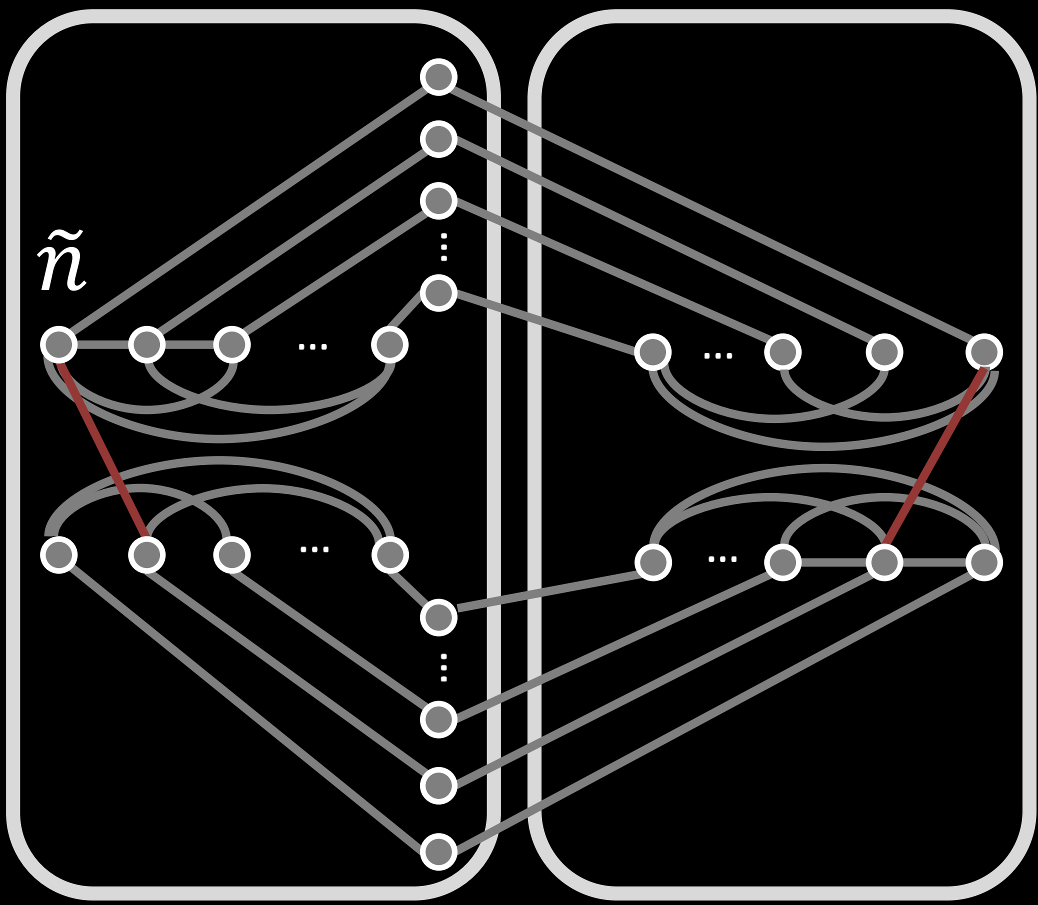

Add edge iff input is 0

cut = $\tilde{n} = n/6$

$k = \tilde{n}^2$

Warm-up: MVC, CONGEST

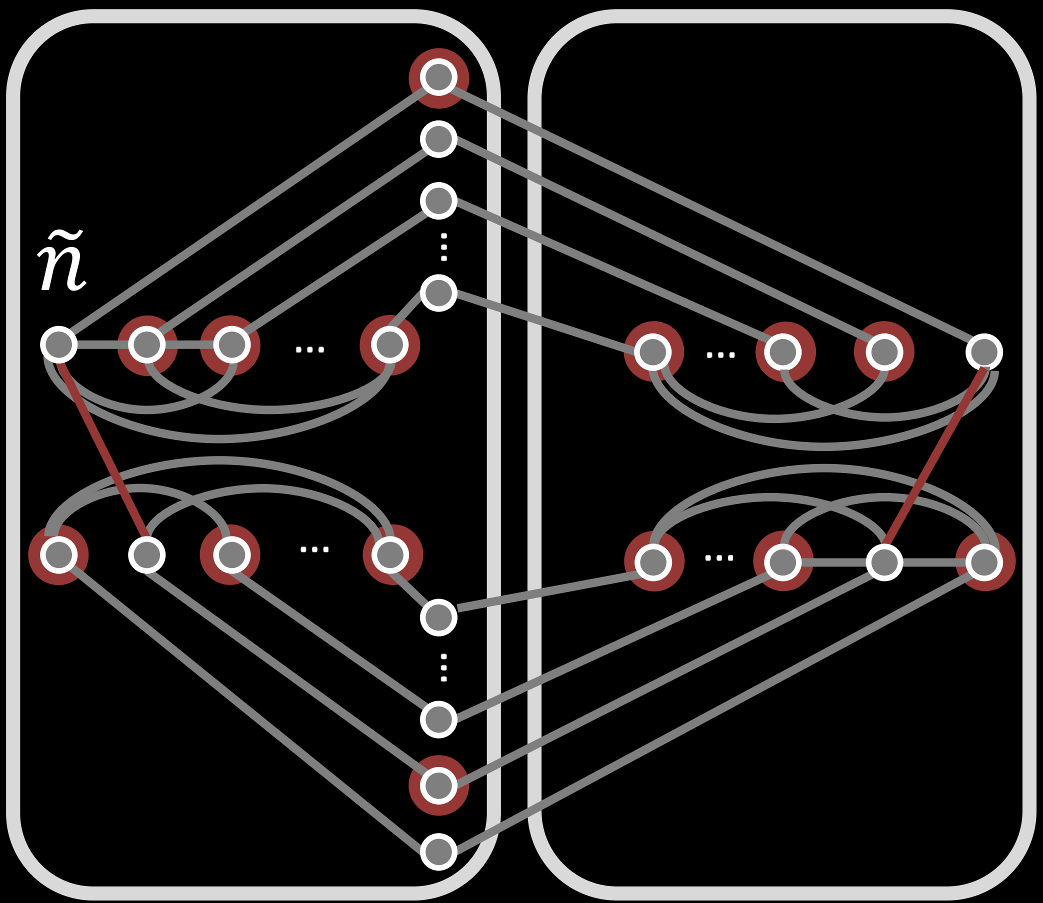

MVC = $4\tilde{n}$-2

inputs not disjoint

Warm-up: MVC, CONGEST

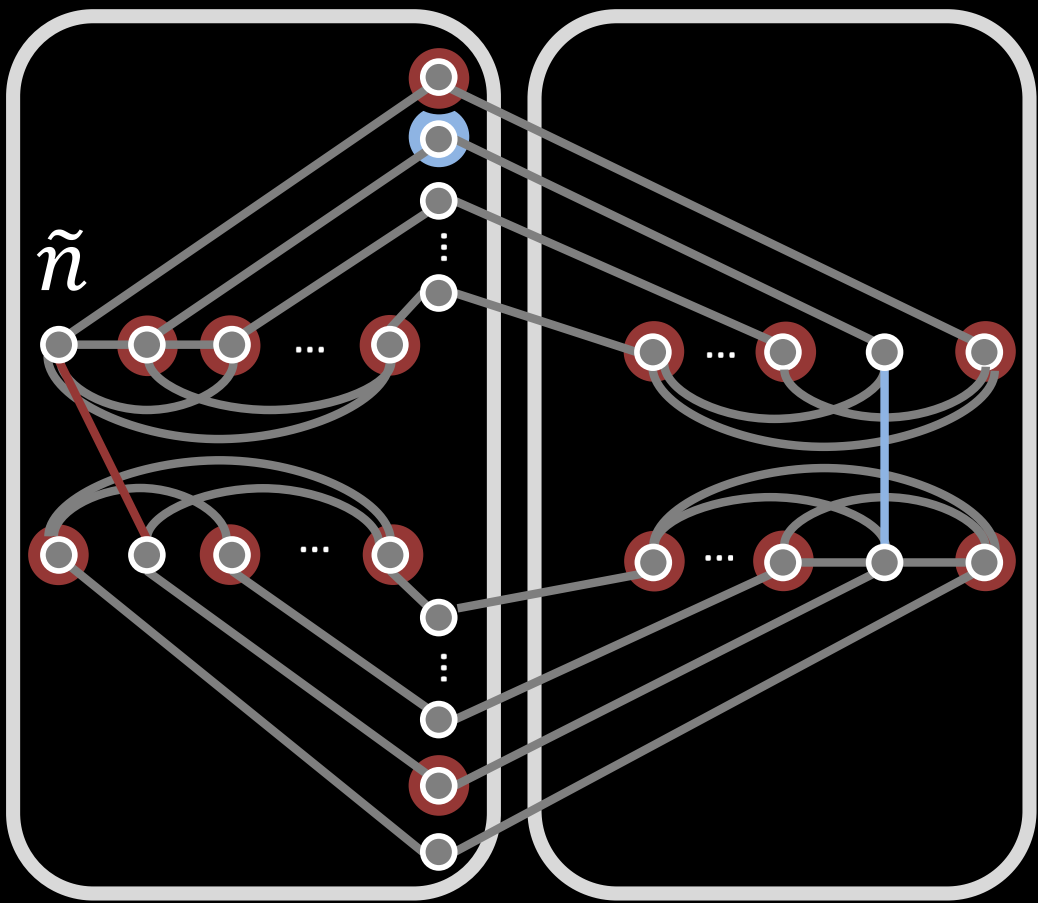

MVC = $4\tilde{n}$-2 ⟺

inputs not disjoint

MVC ≥ $4\tilde{n}$-1 ⟺

inputs disjoint

Warm-up: MVC, CONGEST

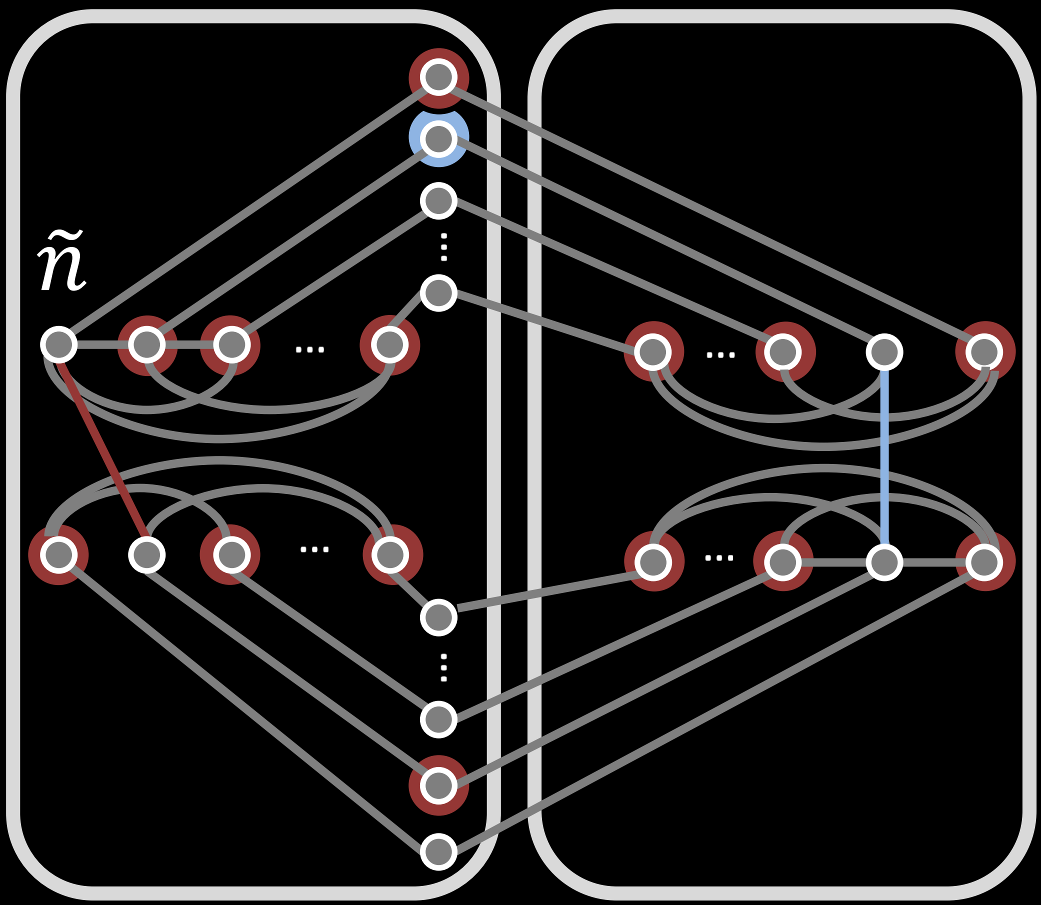

MVC = 4$\tilde{n}$-2 ⟺ inputs not disjoint

MVC ≥ 4$\tilde{n}$-1 ⟺ inputs disjoint

Lower bound:

$\Omega(k/n \log n) = \widetilde{\Omega}(n)$

rounds

Minimum vertex cover in CONGEST

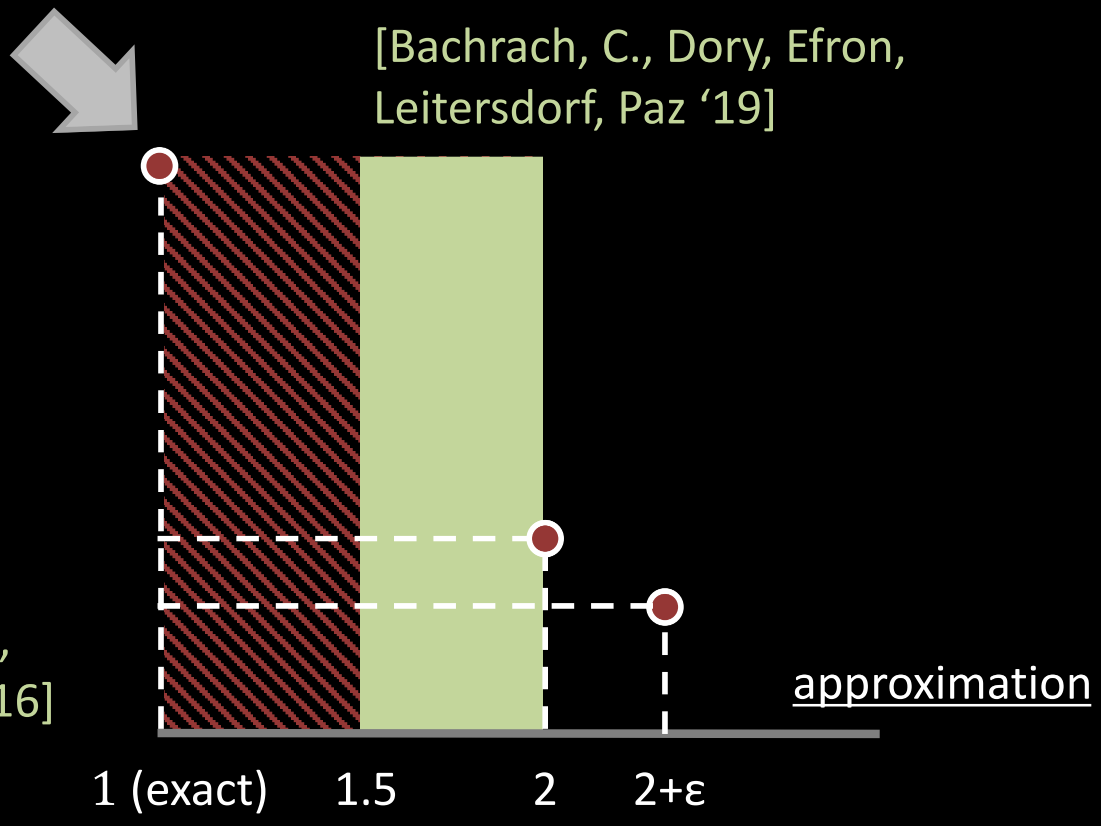

Small cuts for MVC

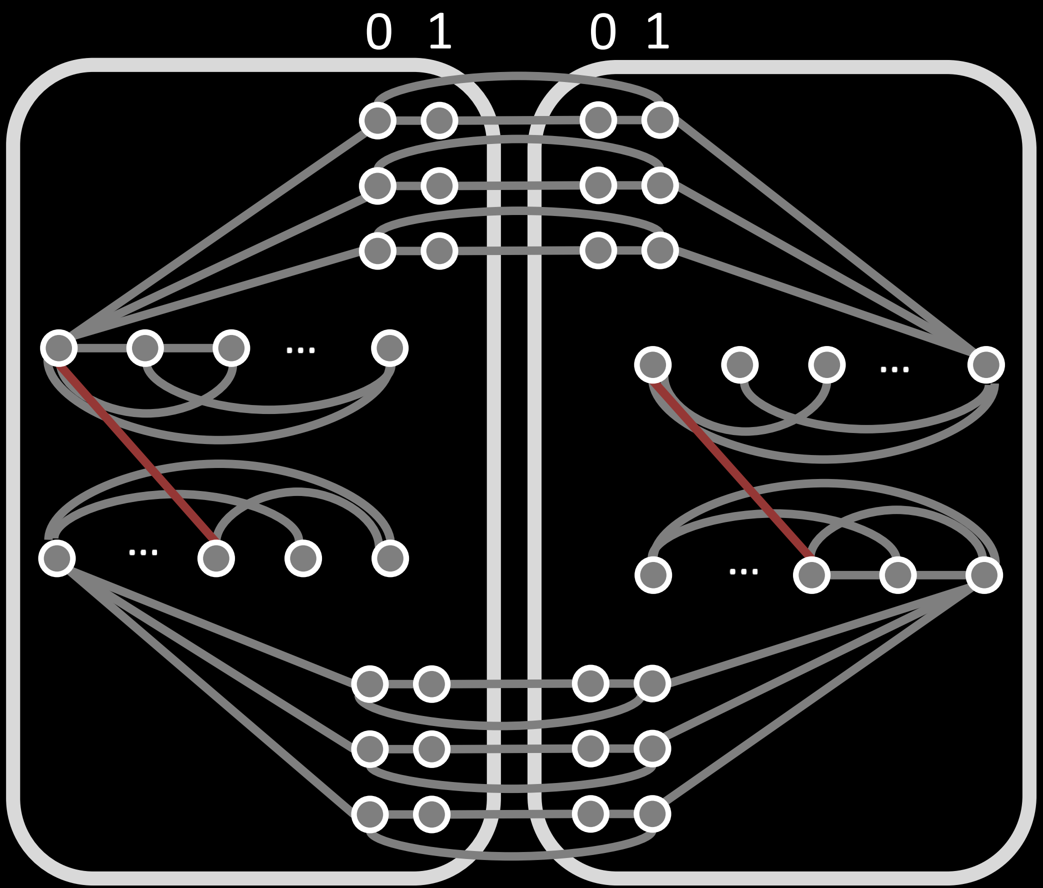

[C., Khoury, Paz '17]

Bit-gadget used for, e.g., diameter in [Abboud, C., Khoury '16]

Small cuts for MVC

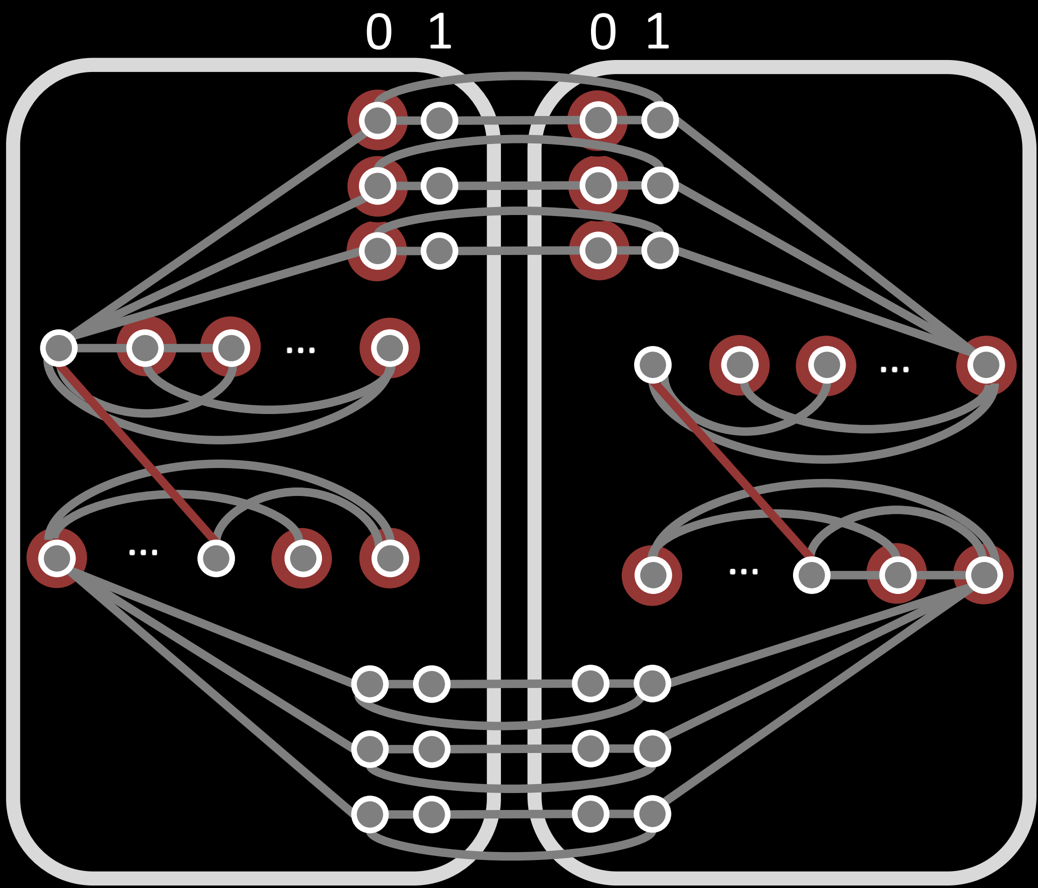

[C., Khoury, Paz '17]

MVC = 4($\tilde{n}$-1+log $\tilde{n}$)

⇕

inputs not disjoint

Small cuts for MVC

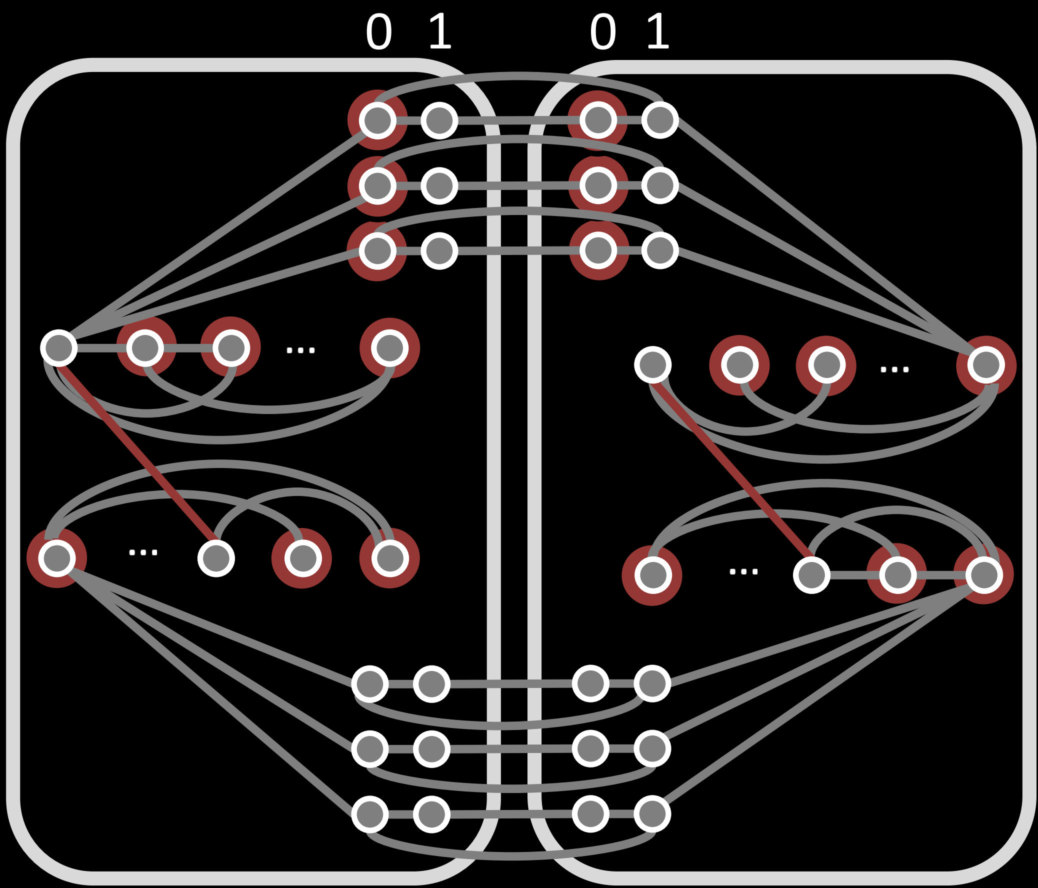

[C., Khoury, Paz '17]

k$=\Theta(n^2)$, **cut**$=\Theta(\log n)$

Rounds

$= \Omega(cost(f(\mathbf{k}))/\mathbf{cut} \cdot \log n)$

$= \Omega(n^2/\log^2 n)$

Minimum vertex cover in CONGEST

#rounds

$\widetilde{\Omega}(n^2)$
[C., Khoury, Paz '17]

[Bachrach, C., Dory, Efron, Leitersdorf, Paz '19]

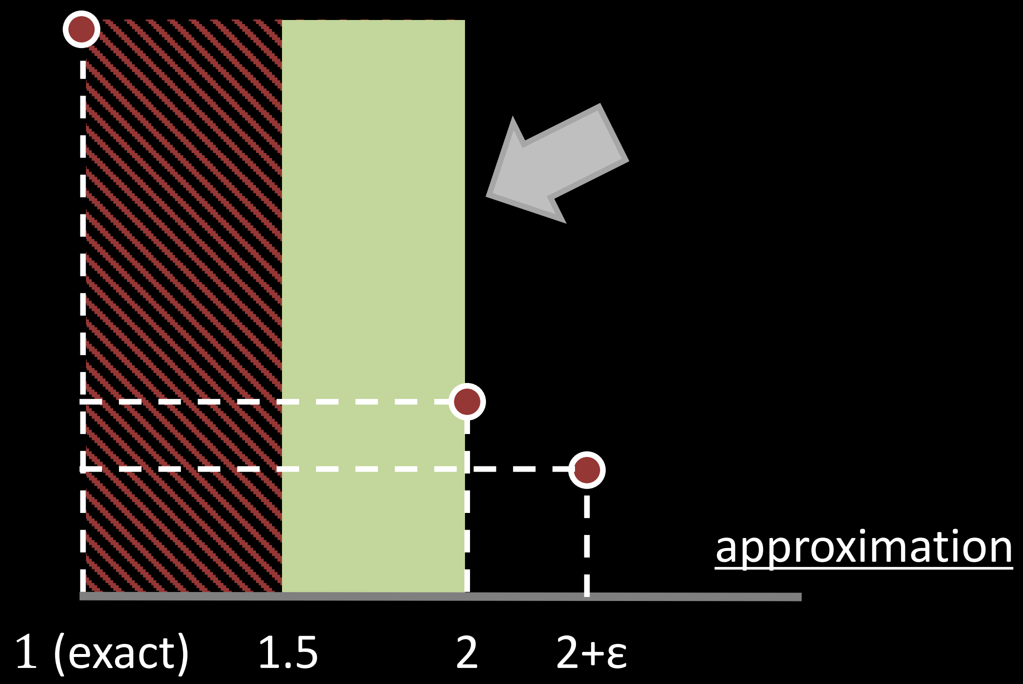

[Ben-Basat, Even, Kawarabayashi, Schwartzman '18]

[Kuhn, Moscibroda, Wattenhofer '04]

[Bar-Yehuda, C., Schwartzman '16]

Why not for 1.5-approximation?

- Alice/Bob compute OPT_A and OPT_B for their parts
- The smaller is at most OPT/2

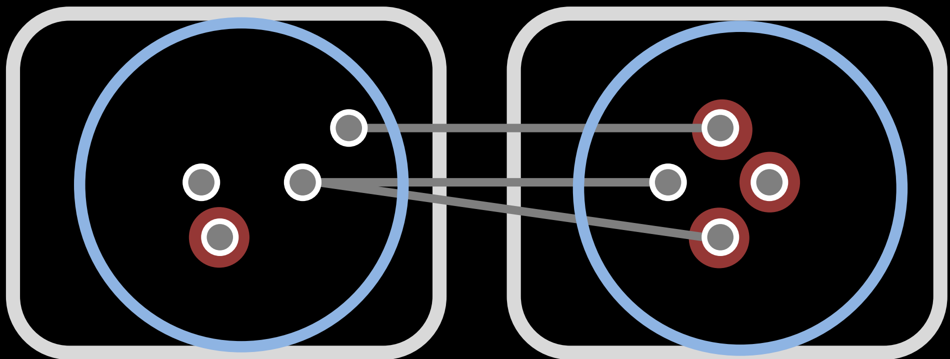

Why not for 1.5-approximation?

- Alice/Bob compute OPT_A and OPT_B for their parts
- The smaller is at most OPT/2

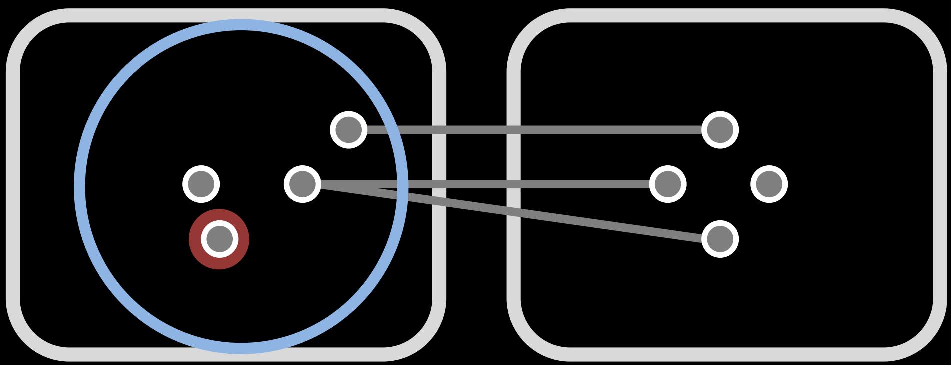

Why not for 1.5-approximation?

- Alice/Bob compute OPT_A and OPT_B for their parts
- The smaller is at most OPT/2
- The other side fills in the rest, at most OPT

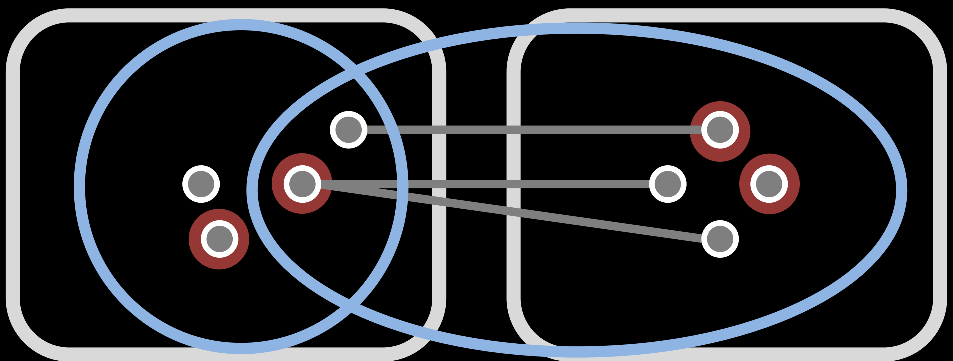

Optimization problems

Maximum independent set in CONGEST

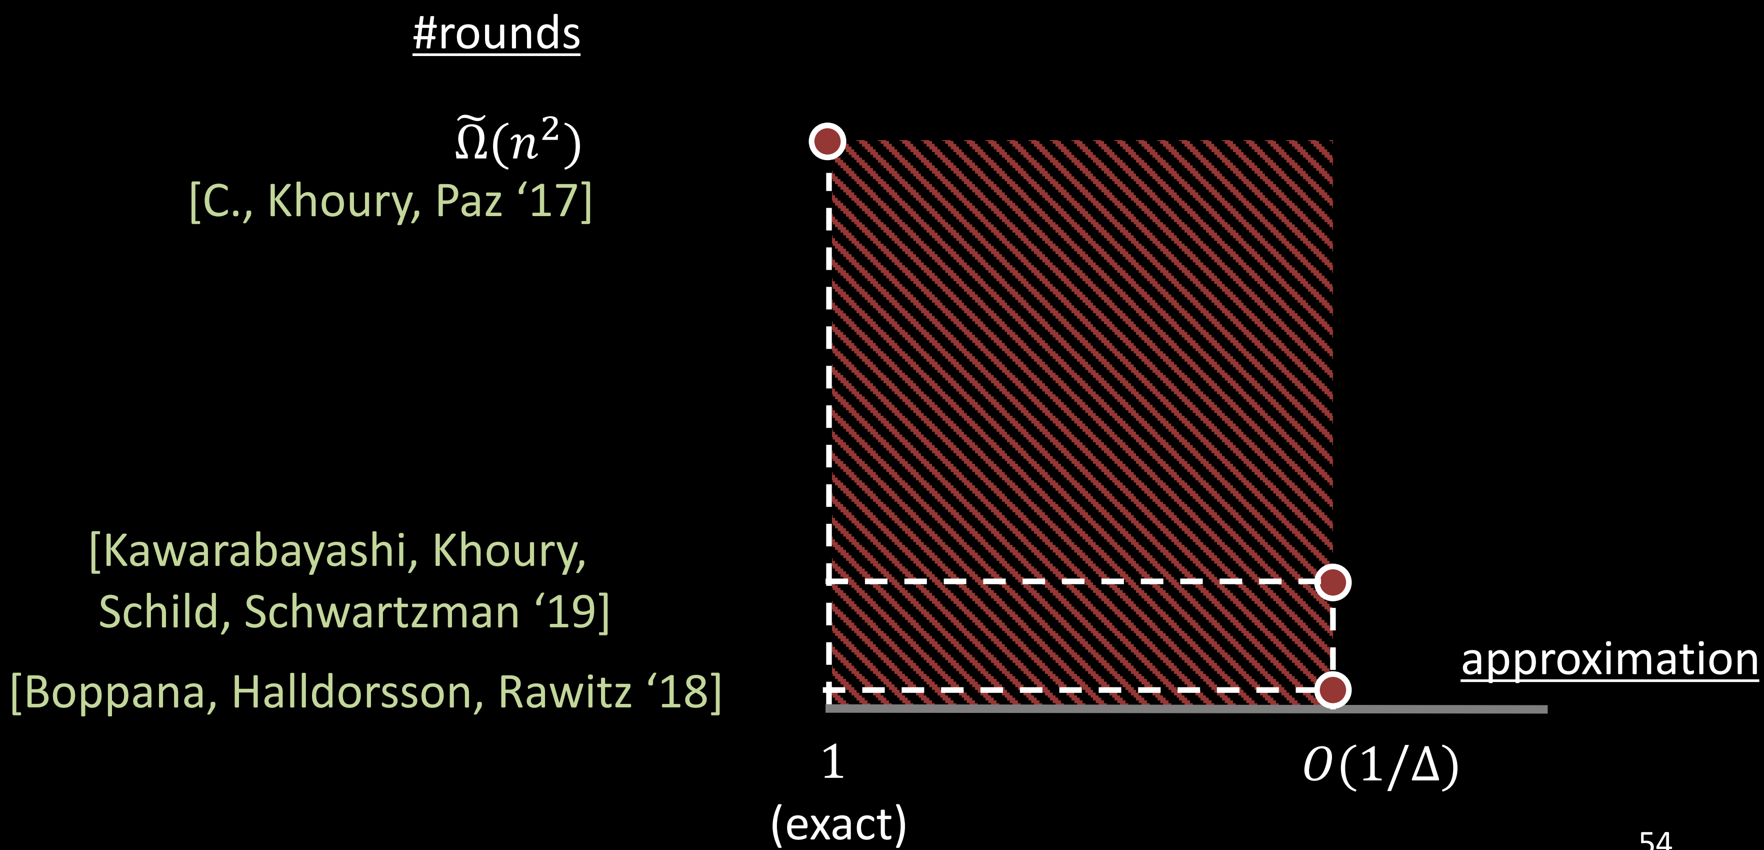

Maximum independent set in CONGEST

#rounds

[Bachrach, C., Dory, Efron, Leitersdorf, Paz '19]

$\widetilde{\Omega}(n^2)$

[C., Khoury, Paz '17]

[Kawarabayashi, Khoury, Schild, Schwartzman '19]

[Boppana, Halldorsson, Rawitz '18]

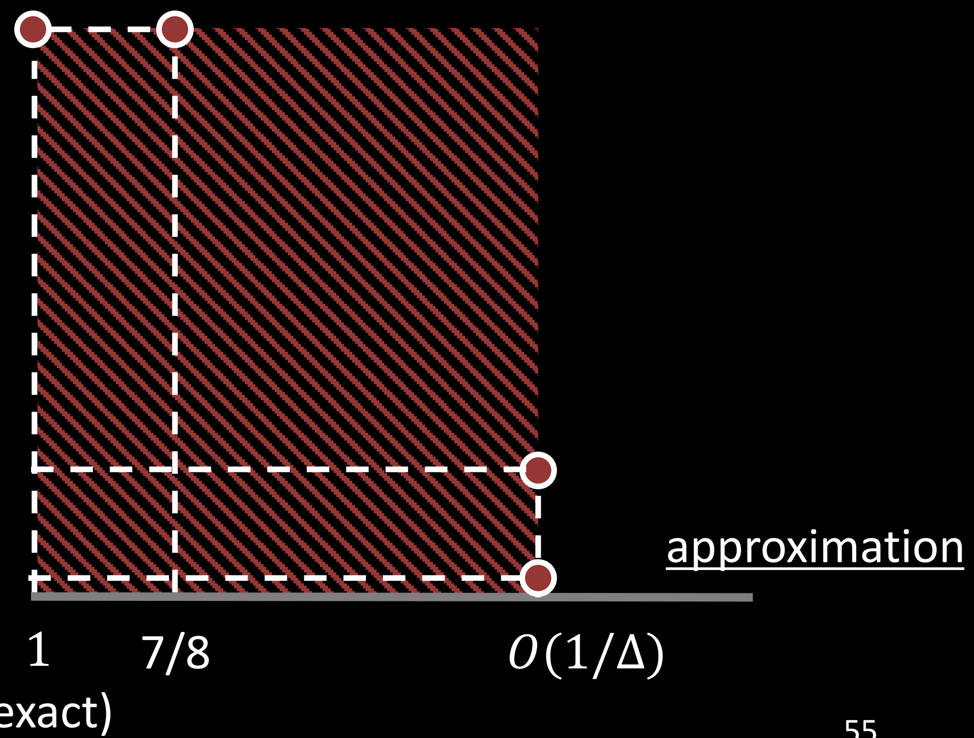

Maximum independent set in CONGEST

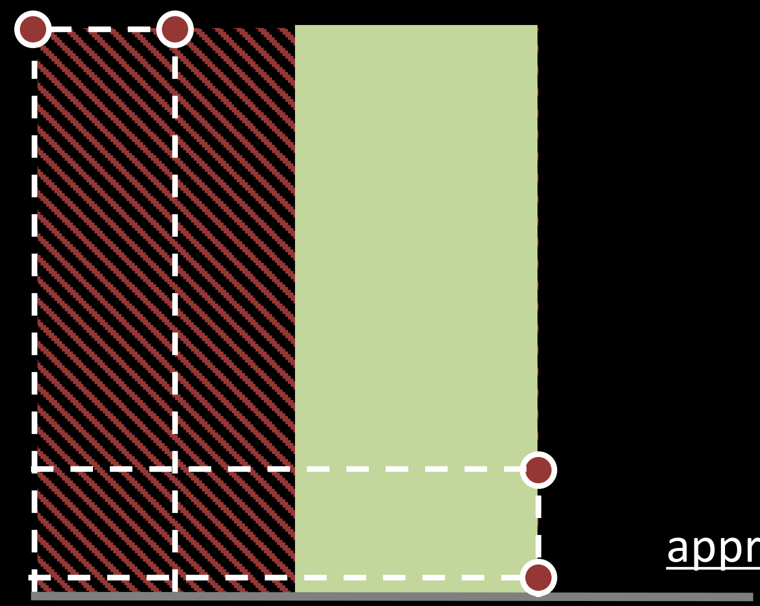

7/8-approximation for MaxIS

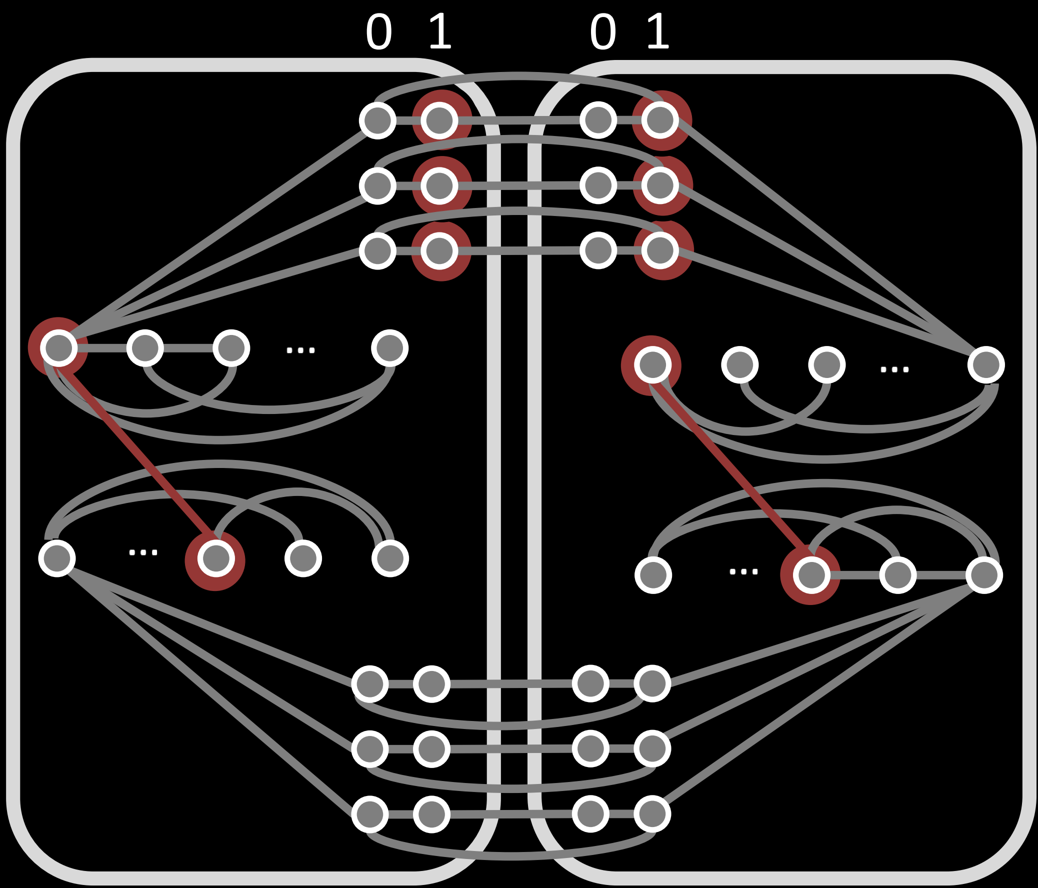

[Bachrach, C., Dory, Efron, Leitersdorf, Paz '19]

Bit-gadget

7/8-approximation for MaxIS

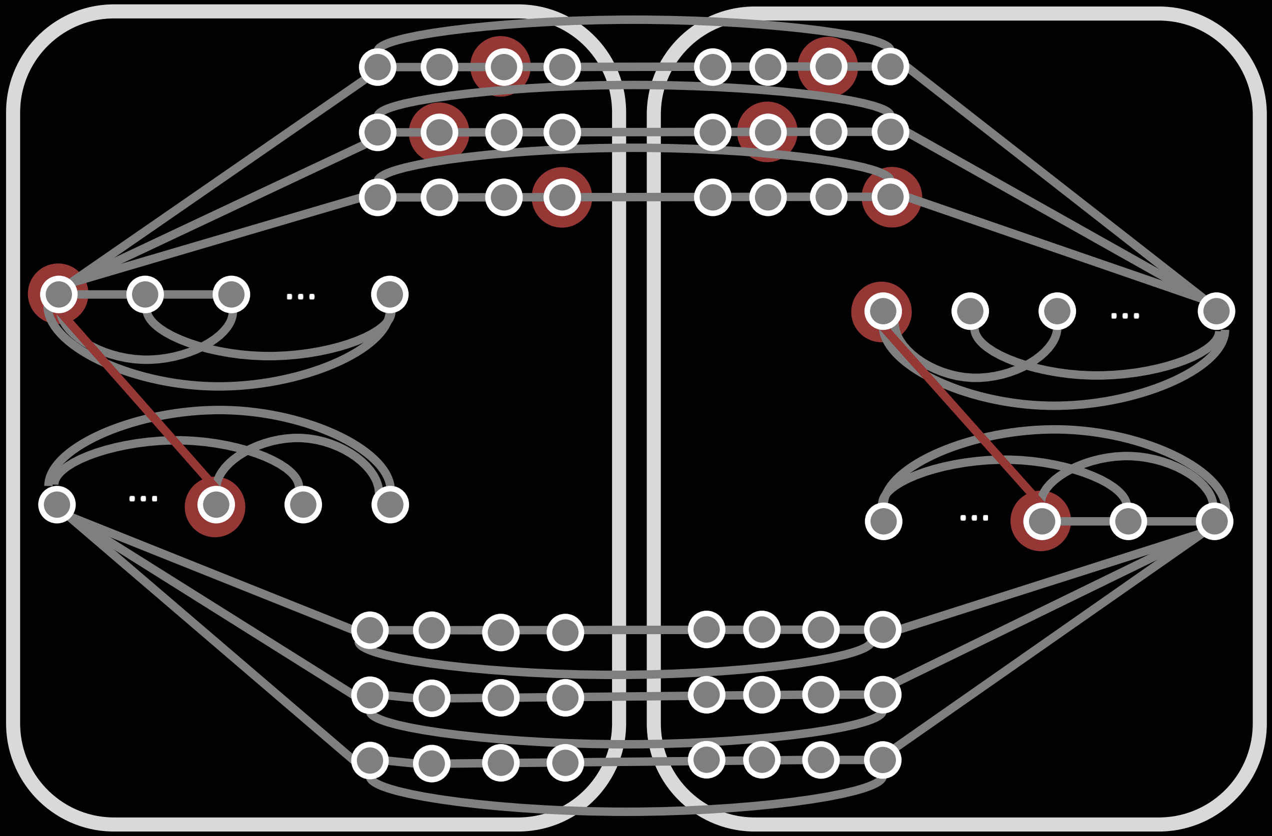

[Bachrach, C., Dory, Efron, Leitersdorf, Paz '19]

Code-gadget: non-binary codewords with a polylog distance

$\tilde{\Omega}(n^2)$ rounds

Optimization problems

- Minimum vertex cover
- Maximum independent set

Optimization problems

- Minimum vertex cover
- Maximum independent set
- Max cut

  [Bachrach, C., Dory, Efron, Leitersdorf, Paz '19]

  – Requires $\widetilde{\Omega}(n^2)$ rounds
  – $(1-\epsilon)$-approximation in $\tilde{O}(n)$ rounds [Zelke '09]

Optimization problems

- Minimum vertex cover
- Maximum independent set
- Max cut

  [Bachrach, C., Dory, Efron, Leitersdorf, Paz '19]

  – Requires $\widetilde{\Omega}(n^2)$ rounds
  – $(1-\epsilon)$-approximation in $\tilde{O}(n)$ rounds [Zelke '09]
- FT-BFS

  [Ghaffari, Parter '16]

  – 2-approximation in $\tilde{O}(D)$ rounds
  – sequential approximation implies set cover

Challenge 3:

Challenge 3: ???

Alice-Bob limitations

- For triangle detection [Drucker, Kuhn, Oshman ‘14]
  - [Izumi, Le Gall ‘17]: $\tilde{O}(n^{2/3})$
  - [Chang, Pettie, Zhang, ’18]: $\tilde{O}(n^{1/2})$
  - [Chang, Saranurak ‘19]: $\tilde{O}(n^{1/3})$
- For weighted APSP above $\Omega(n)$ [C., Khoury, Paz ’17]
  - [Elkin ’17]: $\tilde{O}(n^{5/3})$
  - [Huang, Nanongkai, Saranurak ’17]: $\tilde{O}(n^{5/4})$
  - [Agarwal, Ramachandran, King, Pontecorvi ’18]: $\tilde{O}(n^{3/2})$ det.
  - [Bernstein, Nanongkai ‘19]: $\tilde{O}(n)$
- For 4-clique detection above $\Omega(n^{1/2})$ [Czumaj, Konrad ‘18]
  - [Eden, Fiat, Fischer, Kuhn, Oshman ‘19]: first sublinear algorithm

Distributed optimization

1. Distances
2. Congestion
3. ???

Can we have better approximation factors in the CONGEST model?

Distributed optimization

1. Distances
2. Congestion
3. ???

Can we have better approximation factors in the CONGEST model?
CONGESTED CLIQUE? MPC? Streaming?

Distributed optimization

1. Distances
2. Congestion
3. ???

A challenge for WOLA 2020

Can we have better approximation factors in the CONGEST model?

CONGESTED CLIQUE? MPC? Streaming?

Distributed optimization

1. Distances
2. Congestion
3. ???

A challenge for WOLA 2020

Can we have better approximation factors in the CONGEST model?

CONGESTED CLIQUE? MPC? Streaming?

Thank you!