On Distributed Lower Bound Techniques and a Gap in the Distributed Complexity Landscape

joint work with: Alkida Balliu, Yi-Jun Chang, Jan Grebík, Ole Gabsdil, Tim Göttlicher, Christoph Grunau, Dennis Olivetti, Vasek Rozhoň, Jukka Suomela, Zoltan Vidnyánszky

Funded by the European Union

European Research Council
Established by the European Commission

This work is supported by ERC grant OLA-TOPSENS (grant agreement number 101162747) under the Horizon Europe funding programme.

LOCAL model

On Distributed Lower Bound Techniques and a Gap in the Distributed Complexity Landscape

joint work with: Alkida Balliu, Yi-Jun Chang, Jan Grebík, Ole Gabsdil, Tim Göttlicher, Christoph Grunau, Dennis Olivetti, Vasek Rozhoň, Jukka Suomela, Zoltan Vidnyánszky

Funded by the European Union

European Research Council
Established by the European Commission

This work is supported by ERC grant OLA-TOPSENS (grant agreement number 101162747) under the Horizon Europe funding programme.

Locally Checkable Problems

<u>Locally Checkable:</u>

Output correctness is defined via <u>local</u> (= $O(1)$-hop) constraints.

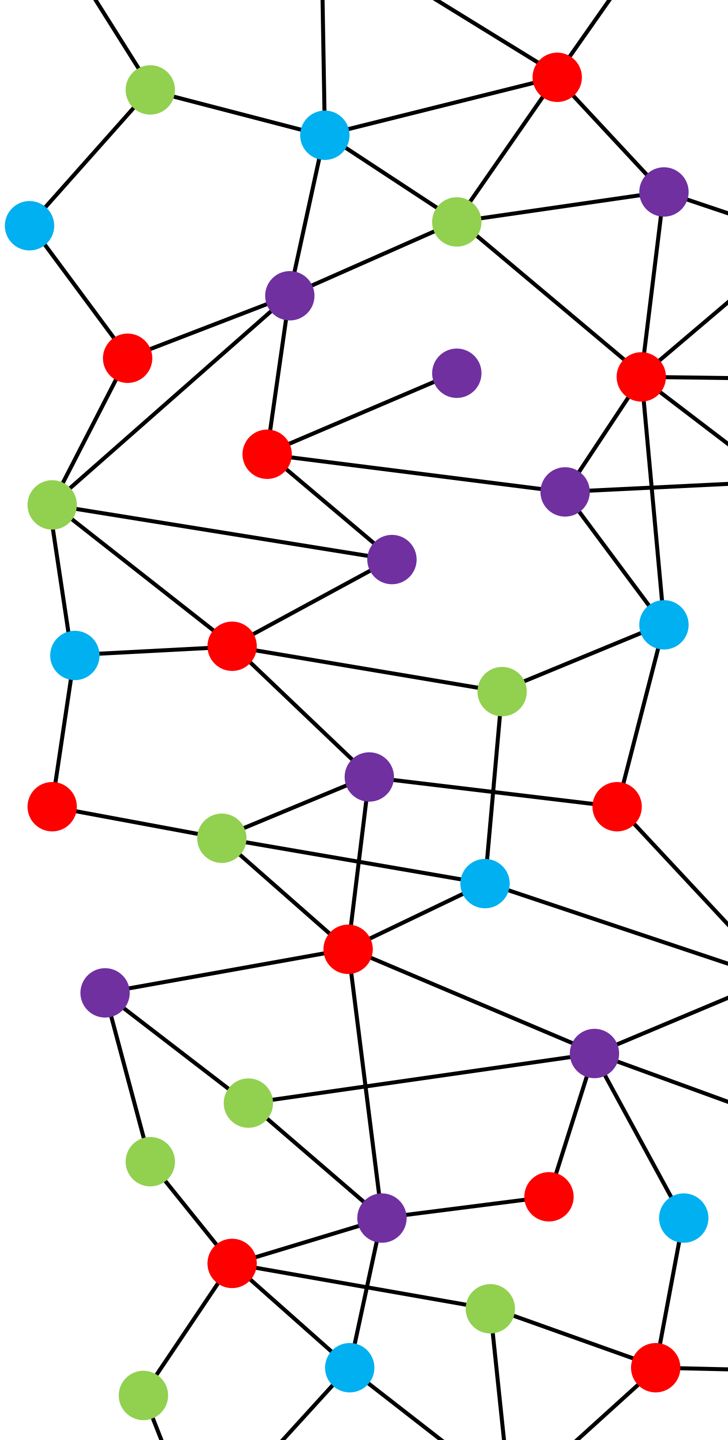

Locally Checkable Problems

<u>Locally Checkable:</u>

Output correctness is defined via <u>local</u> (= $O(1)$-hop) constraints.

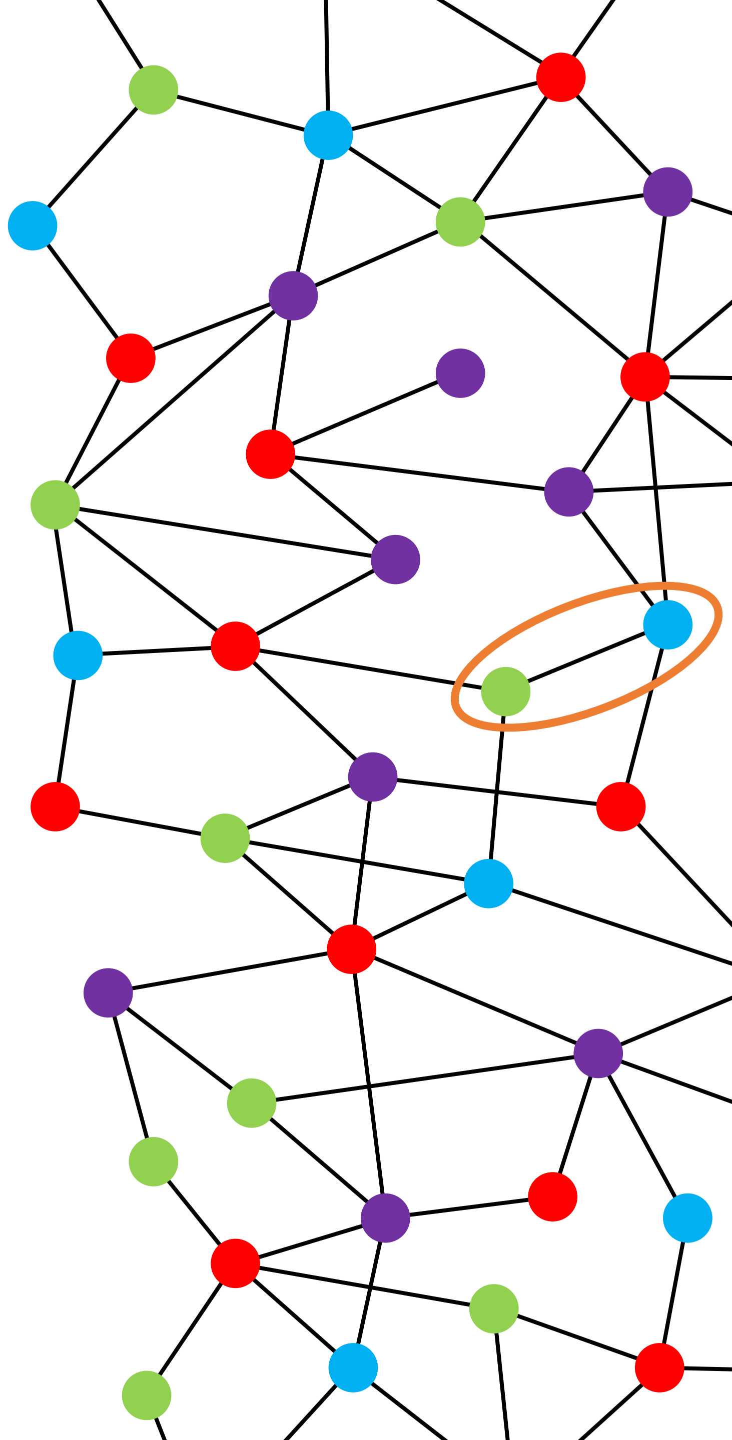

Locally Checkable Problems

<u>Locally Checkable:</u>

Output correctness is defined via <u>local</u> (= $O(1)$-hop) constraints.

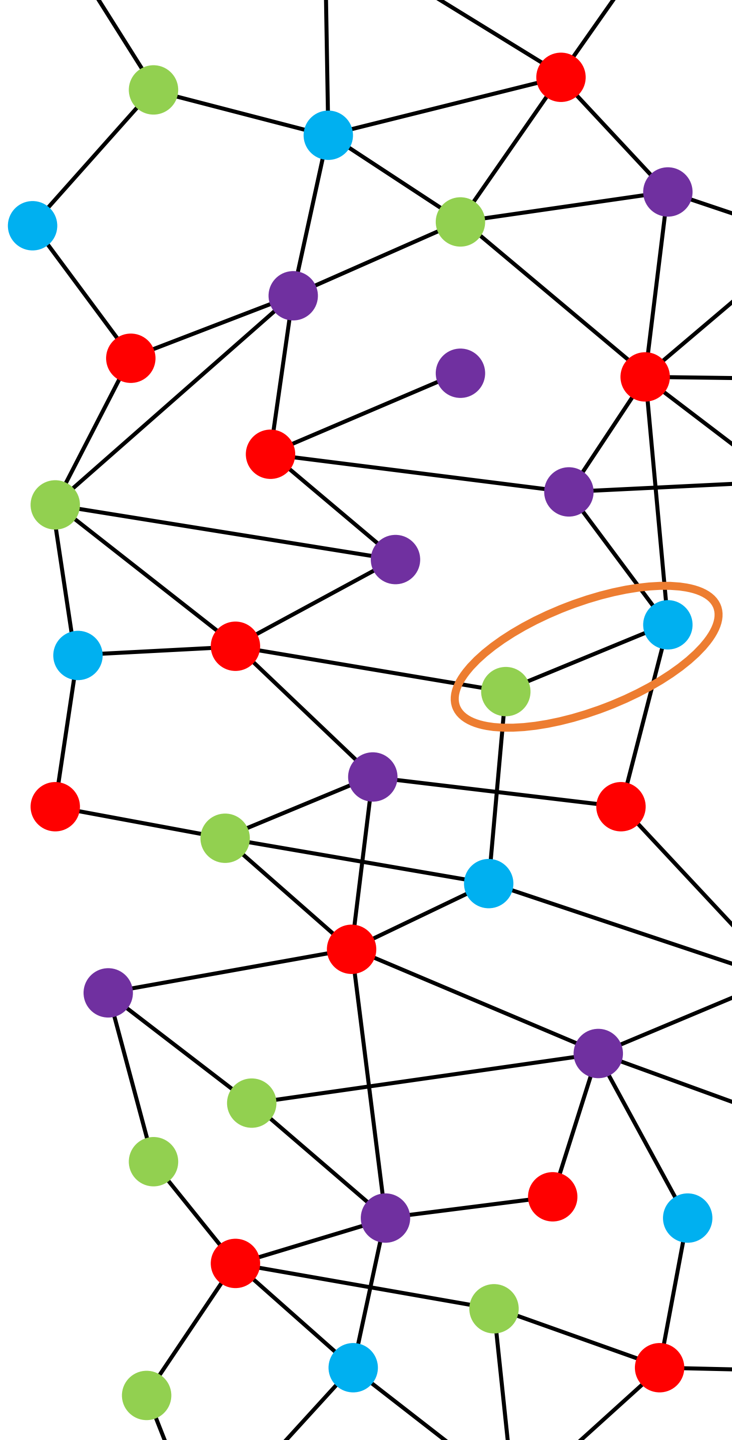

<u>Throughout the talk:</u>

- ❖ constant-degree trees
- ❖ deterministic algorithms
- ❖ constant label set

A Dichotomy

Any locally checkable problem

either can be solved in $O(\log^* n)$ rounds

or requires $\Omega(\log n)$ rounds.

[Chang, Pettie, FOCS'17]

A Dichotomy

Any locally checkable problem

either can be solved in $O(\log^* n)$ rounds

or requires $\Omega(\log n)$ rounds.

[Chang, Pettie, FOCS'17]

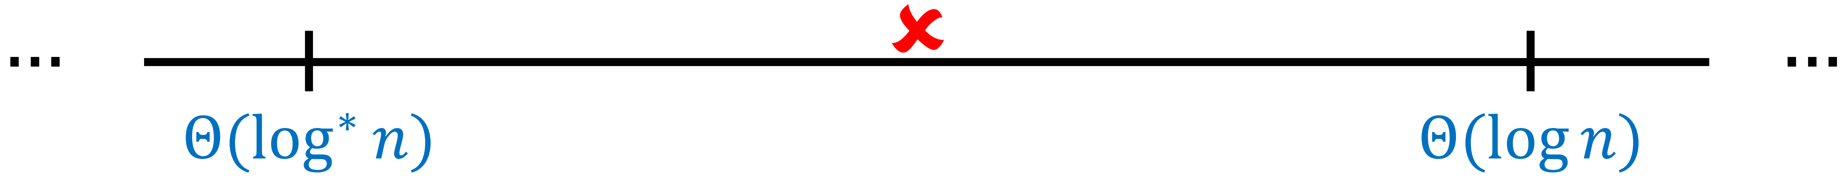

A Dichotomy

Any locally checkable problem

either can be solved in $O(\log^* n)$ rounds

or requires $\Omega(\log n)$ rounds.

[Chang, Pettie, FOCS'17]

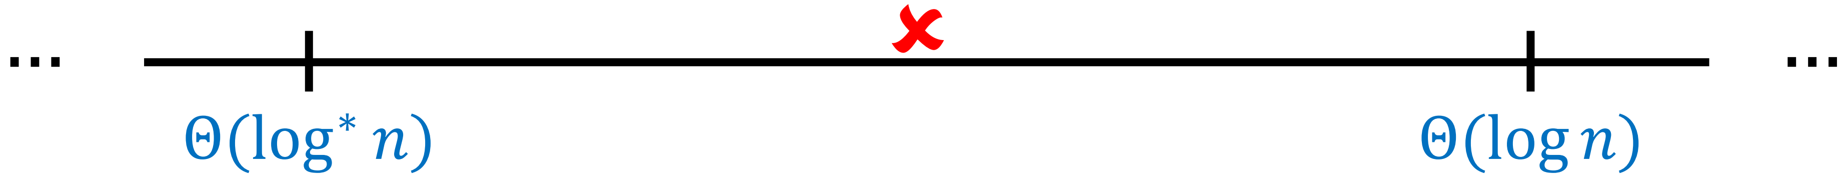

Is it decidable on which side of the gap a problem lies?

A Dichotomy

Any locally checkable problem

either can be solved in $O(\log^* n)$ rounds

or requires $\Omega(\log n)$ rounds.

[Chang, Pettie, FOCS'17]

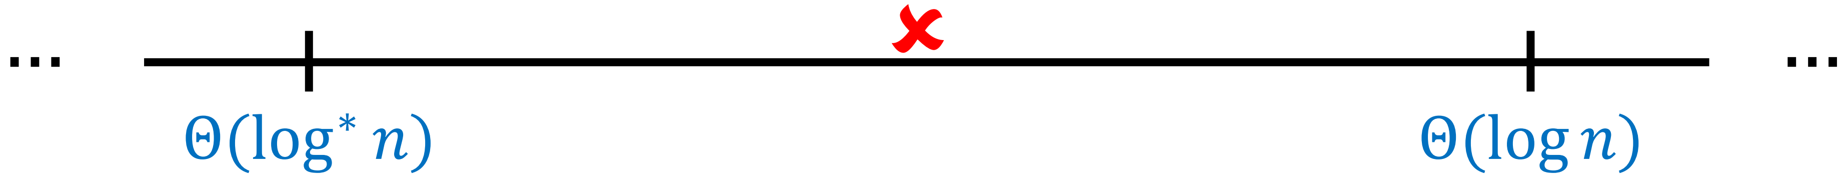

semi-decidable

Is it decidable on which side of the gap a problem lies?

A Dichotomy

Any locally checkable problem

either can be solved in $O(\log^* n)$ rounds

or requires $\Omega(\log n)$ rounds.

[Chang, Pettie, FOCS'17]

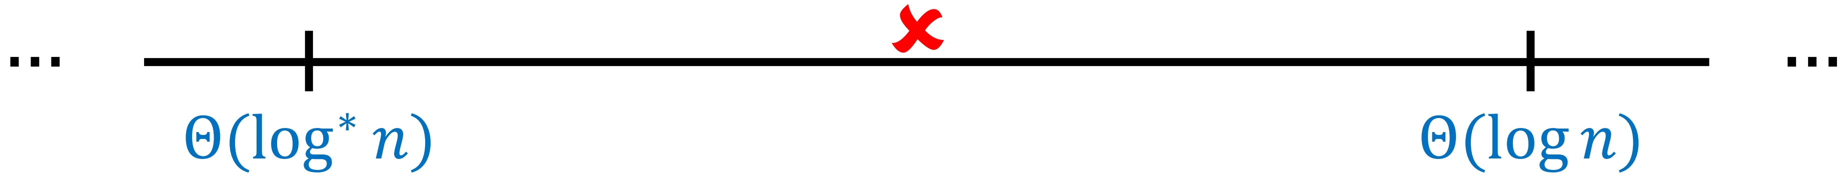

semi-decidable

Is it decidable on which side of the gap a problem lies?

?

Round Elimination

Round Elimination

For any problem Π,
we can automatically find a problem $\mathcal{R}(\Pi)$
that is exactly 1 round easier.

[Brandt, PODC'19]

Round Elimination

$$\Pi \xrightarrow{\mathcal{R}} \Pi_1 \xrightarrow{\mathcal{R}} \Pi_2 \xrightarrow{\mathcal{R}} \dots$$

$$T \qquad T_1 = T - 1 \qquad T_2 = T - 2$$

For any problem Π,
we can automatically find a problem $\mathcal{R}(\Pi)$
that is exactly 1 round easier.

Round Elimination

$$\Pi \xrightarrow{\mathcal{R}} \Pi_1 \xrightarrow{\mathcal{R}} \Pi_2 \xrightarrow{\mathcal{R}} \dots \xrightarrow{\mathcal{R}} \Pi_k \xrightarrow{\mathcal{R}} \dots$$

$$T \qquad T_1 = T - 1 \qquad T_2 = T - 2$$

For any problem Π,
we can automatically find a problem $\mathcal{R}(\Pi)$
that is exactly 1 round easier.

Round Elimination

If all problems Π_i are nontrivial,
then Π has complexity $\Omega(\log n)$.

$$\Pi \xrightarrow{\mathcal{R}} \Pi_1 \xrightarrow{\mathcal{R}} \Pi_2 \xrightarrow{\mathcal{R}} \dots \xrightarrow{\mathcal{R}} \Pi_k \xrightarrow{\mathcal{R}} \dots$$

For any problem Π,
we can automatically find a problem $\mathcal{R}(\Pi)$
that is exactly 1 round easier.

Round Elimination

If all problems Π_i are nontrivial, then Π has complexity $\Omega(\log n)$.

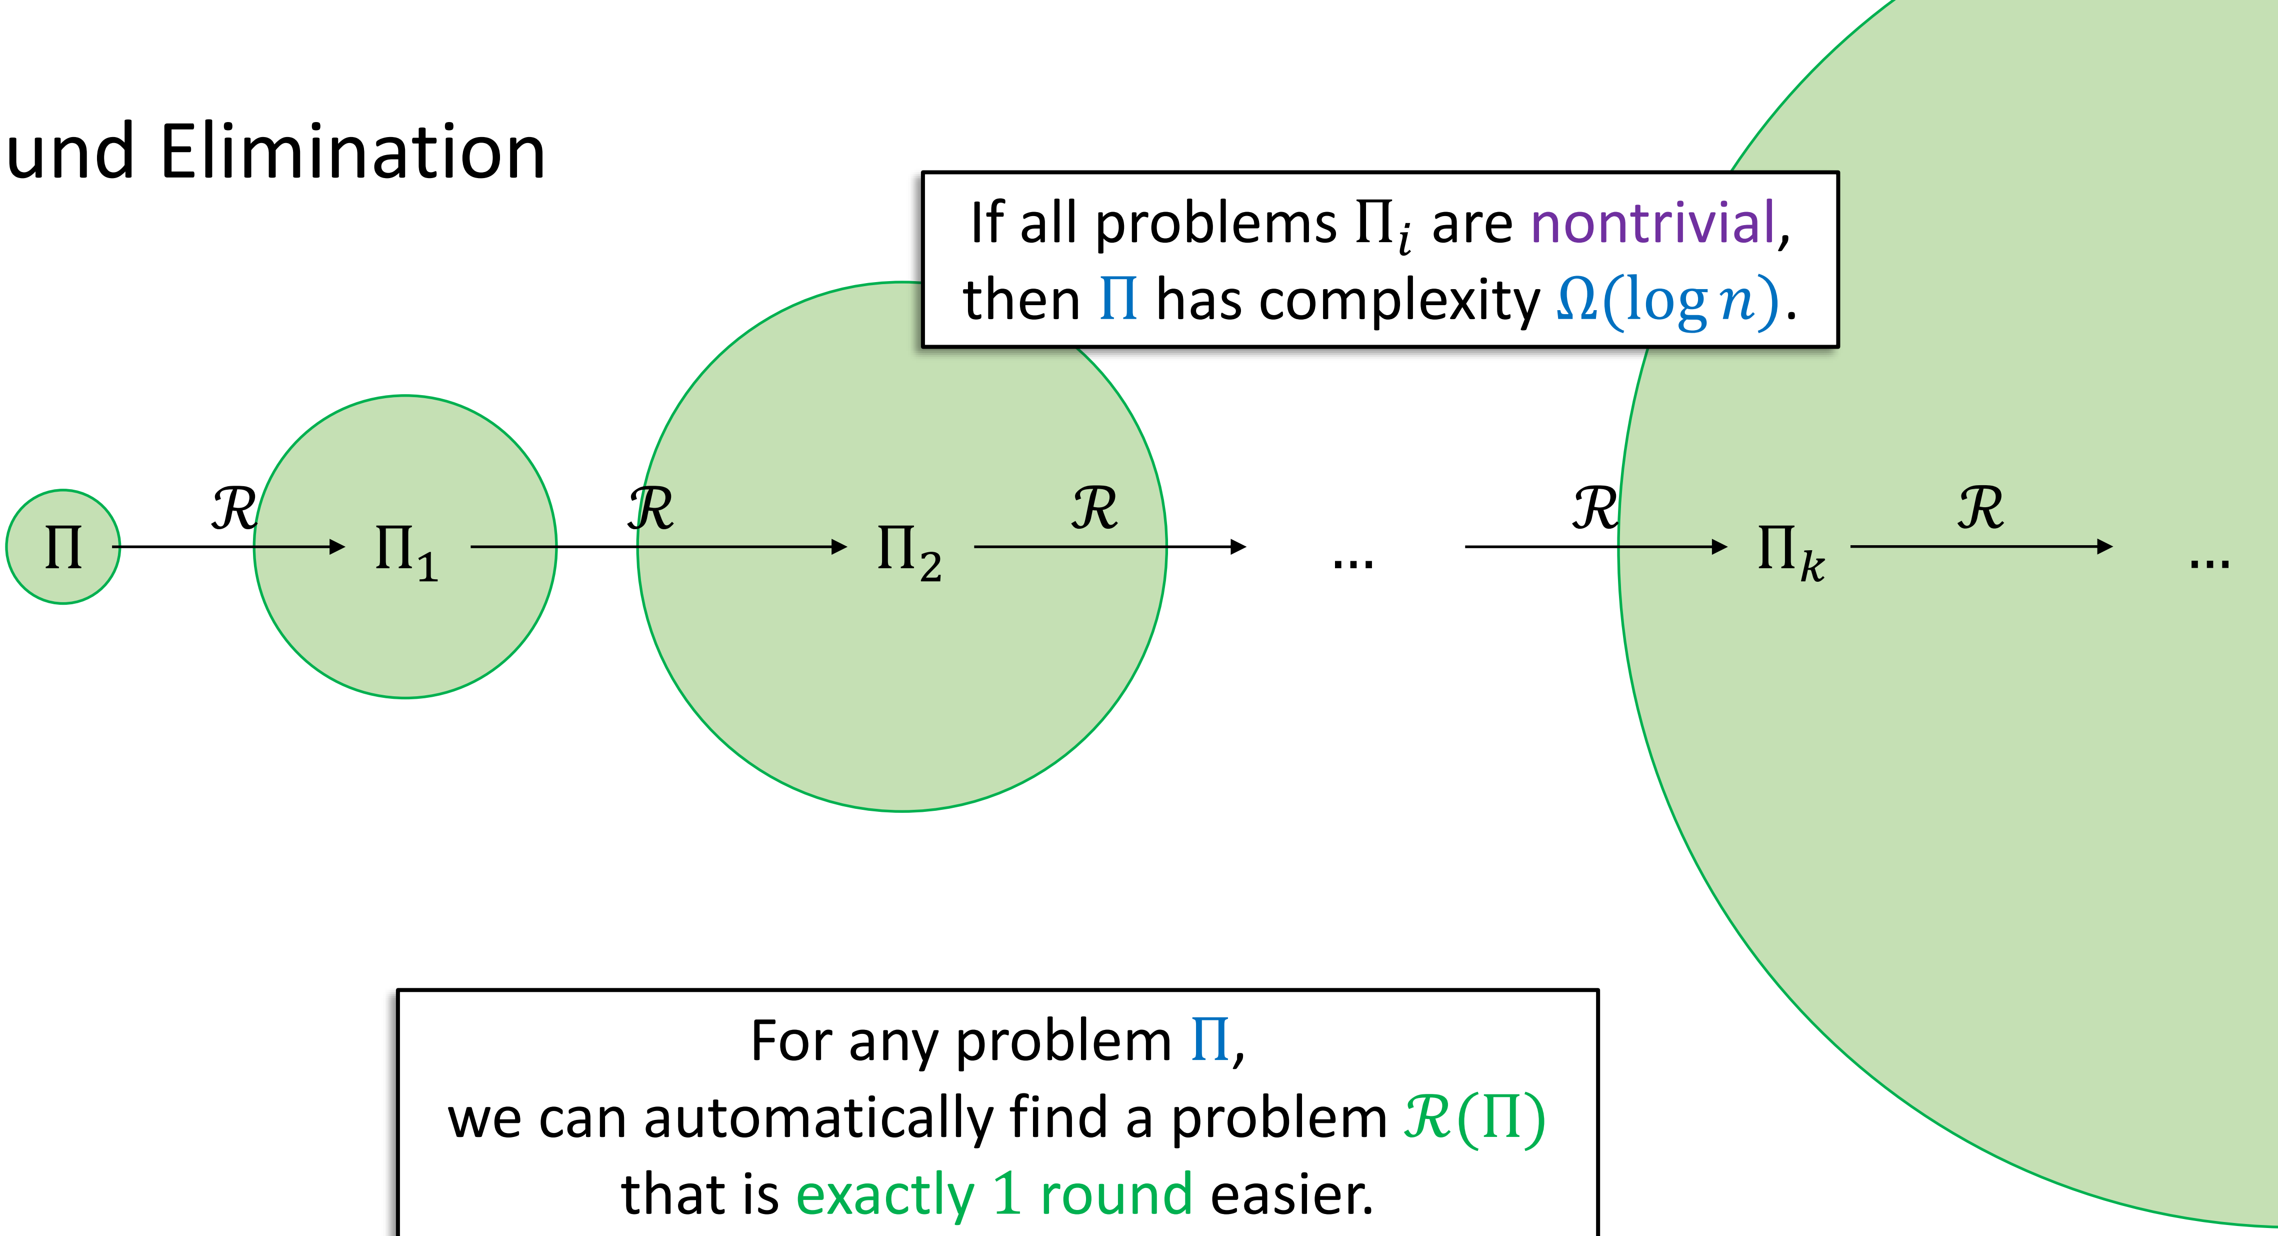

For any problem Π,
we can automatically find a problem $\mathcal{R}(\Pi)$
that is exactly 1 round easier.

Round Elimination

If all problems Π_i are nontrivial,
then Π has complexity $\Omega(\log n)$.

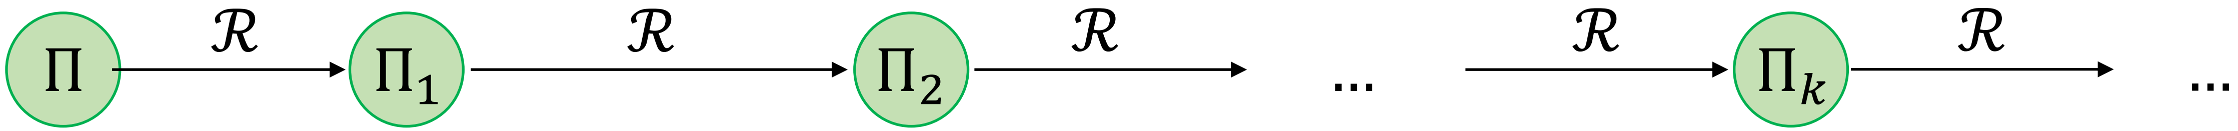

For any problem Π,
we can automatically find a problem $\mathcal{R}(\Pi)$
that is exactly 1 round easier.

Fixed Points

If all problems Π_i are nontrivial,
then Π has complexity $\Omega(\log n)$.

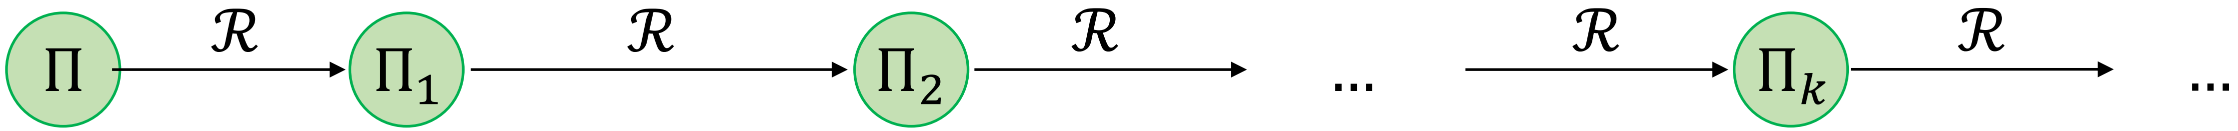

Fixed point Π:
$\mathcal{R}(\Pi) = \Pi$.

Fixed Points

If all problems Π_i are nontrivial,
then Π has complexity $\Omega(\log n)$.

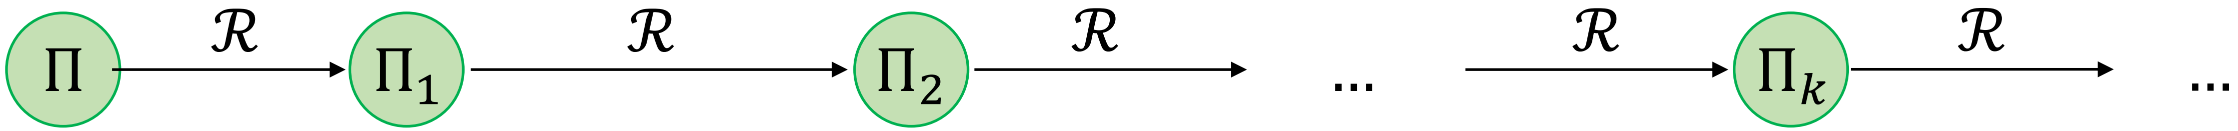

Fixed point Π:
$\mathcal{R}(\Pi) = \Pi$.

If Π is a nontrivial fixed point,
then Π has complexity $\Omega(\log n)$.

Fixed Points

If all problems Π_i are nontrivial,
then Π has complexity $\Omega(\log n)$.

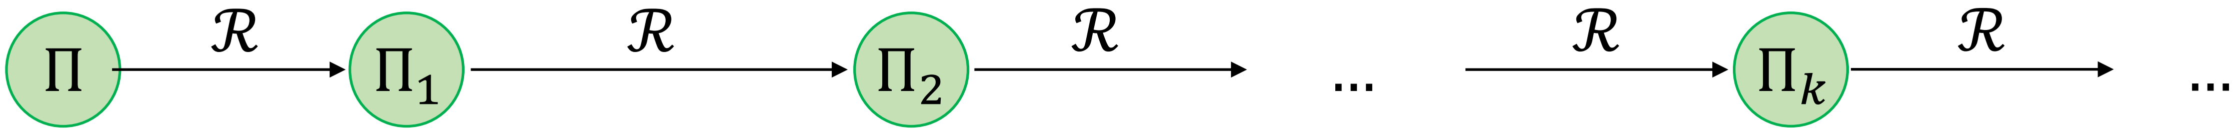

Fixed point Π:
$\mathcal{R}(\Pi) = \Pi$.

Sinkless Orientation:

Orient the edges such that
no node of degree ≥ 3 is a sink.

If Π is a nontrivial fixed point,
then Π has complexity $\Omega(\log n)$.

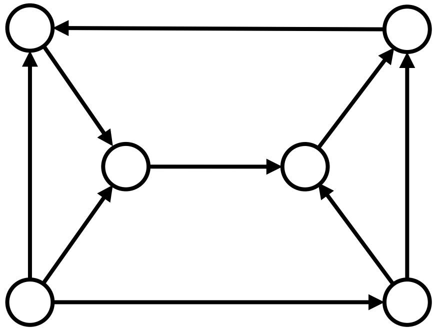

Fixed Points

If all problems Π_i are nontrivial,
then Π has complexity $\Omega(\log n)$.

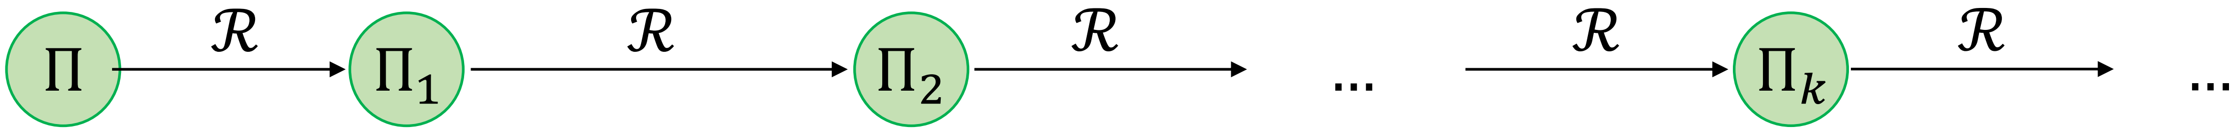

Rare!

Fixed point Π:
$\mathcal{R}(\Pi) = \Pi$.

Sinkless Orientation:

Orient the edges such that
no node of degree ≥ 3 is a sink.

If Π is a nontrivial fixed point,
then Π has complexity $\Omega(\log n)$.

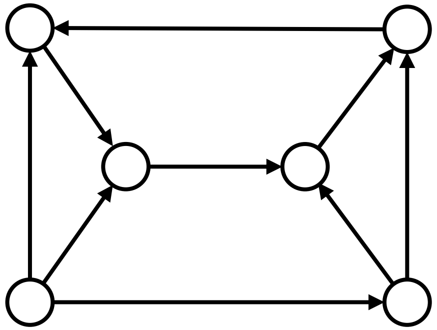

Fixed Points

If all problems Π_i are nontrivial,
then Π has complexity $\Omega(\log n)$.

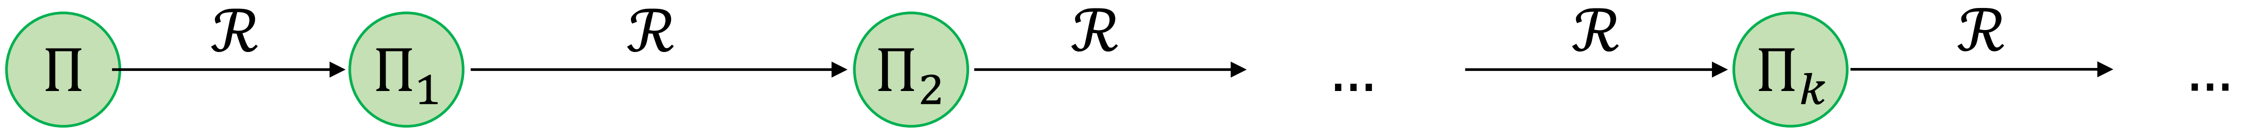

Rare! Fixed point Π:
$\mathcal{R}(\Pi) = \Pi$.

Sinkless Orientation:

Orient the edges such that
no node of degree ≥ 3 is a sink.

If Π can be relaxed to a nontrivial fixed point,
then Π has complexity $\Omega(\log n)$.

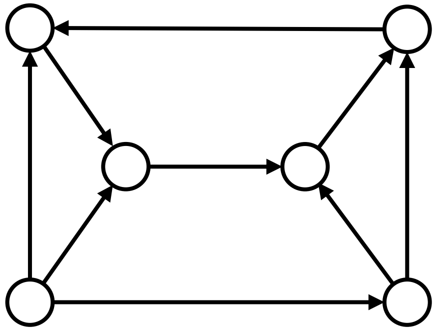

Fixed Points

If all problems Π_i are nontrivial,
then Π has complexity $\Omega(\log n)$.

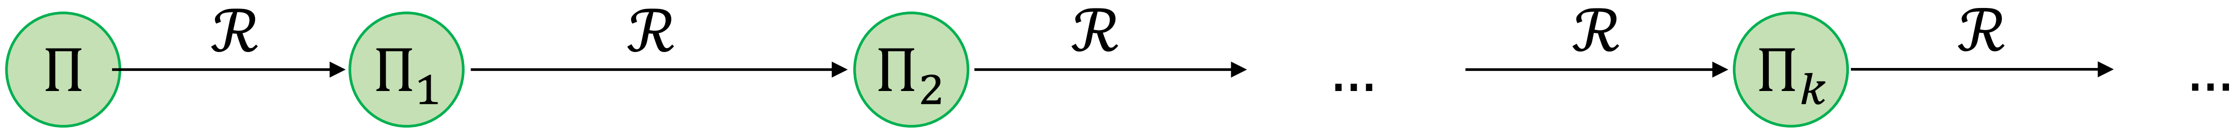

Rare!

Fixed point Π:
$\mathcal{R}(\Pi) = \Pi$.

Sinkless Orientation:

Orient the edges such that
no node of degree ≥ 3 is a sink.

0-round
reduction

If Π can be relaxed to a nontrivial fixed point,
then Π has complexity $\Omega(\log n)$.

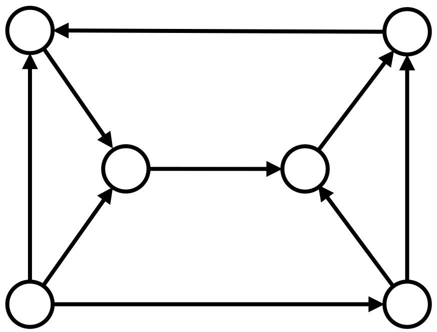

Fixed Points

If all problems Π_i are nontrivial, then Π has complexity $\Omega(\log n)$.

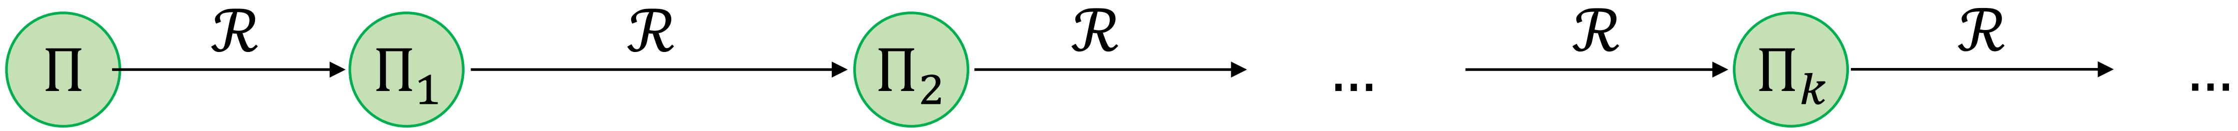

Rare!

Fixed point Π: $\mathcal{R}(\Pi) = \Pi$.

Sinkless Orientation:

Orient the edges such that no node of degree ≥ 3 is a sink.

0-round reduction

If Π can be relaxed to a nontrivial fixed point, then Π has complexity $\Omega(\log n)$.

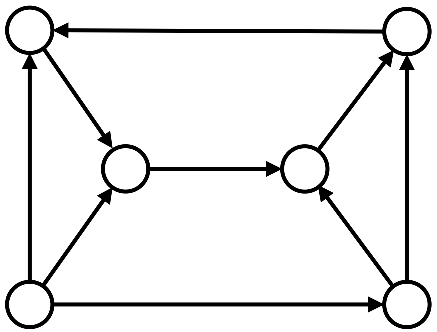

Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?

Fixed Points

If all problems Π_i are nontrivial, then Π has complexity $\Omega(\log n)$.

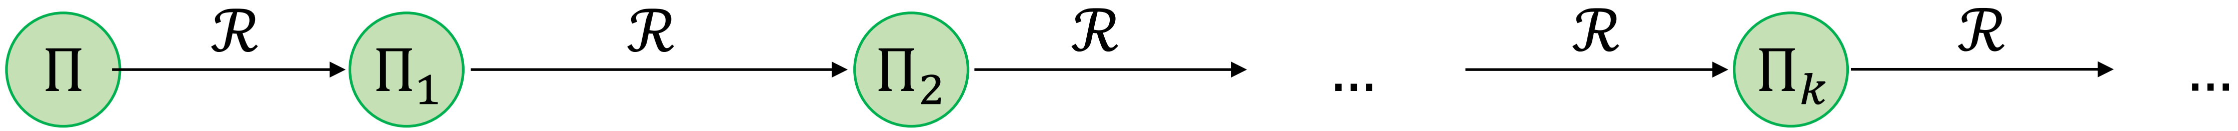

Rare!

Fixed point Π: $\mathcal{R}(\Pi) = \Pi$.

0-round reduction

If Π can be relaxed to a nontrivial fixed point, then Π has complexity $\Omega(\log n)$.

Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?

if yes ...

decidability between $O(\log^* n)$ / $\Omega(\log n)$!

Fundamental Questions

- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?

Fundamental Questions

- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?

- How can we find nontrivial fixed point relaxations?

The Setting with Input

<u>Assumption:</u> an (adversarial) solution to some fixed locally checkable problem $\mathcal{I}$ is given as input

The Setting with Input

Assumption: an (adversarial) solution to some fixed locally checkable problem $\mathcal{I}$ is given as input

"setting with input $\mathcal{I}$"

The Setting with Input

<u>Assumption:</u> an (adversarial) solution to some fixed locally checkable problem $\mathcal{I}$ is given as input

"setting with input $\mathcal{I}$"

Does Δ-coloring have an $\Omega(\log n)$-round lower bound in 4-regular graphs with indegree = outdegree = 2?

The Setting with Input

<u>Assumption:</u> an (adversarial) solution to some fixed locally checkable problem $\mathcal{I}$ is given as input

"setting with input $\mathcal{I}$"

Reductions
(great for UBs)

Does Δ-coloring have an $\Omega(\log n)$-round lower bound in 4-regular graphs with indegree = outdegree = 2?

The Setting with Input

<u>Assumption:</u> an (adversarial) solution to some fixed locally checkable problem $\mathcal{J}$ is given as input

"setting with input $\mathcal{J}$"

Reductions
(great for UBs)

"Round elimination still works"
(great for LBs)

Does Δ-coloring have an $\Omega(\log n)$-round lower bound in 4-regular graphs with indegree = outdegree = 2?

The Setting with Input

<u>Assumption:</u> an (adversarial) solution to some fixed locally checkable problem $\mathcal{J}$ is given as input

"setting with input $\mathcal{J}$"

Reductions
(great for UBs)

"Round elimination still works"
(great for LBs)

If Π can be relaxed to a nontrivial fixed point, then Π has complexity $\Omega(\log n)$.

Does Δ-coloring have an $\Omega(\log n)$-round lower bound in 4-regular graphs with indegree = outdegree = 2?

The Setting with Input

Assumption: an (adversarial) solution to some fixed locally checkable problem $\mathcal{J}$ is given as input

"setting with input $\mathcal{J}$"

Reductions
(great for UBs)

"Round elimination still works"
(great for LBs)

If Π can be relaxed to a nontrivial fixed point, then Π has complexity $\Omega(\log n)$.

LLL has complexity $\Omega(\log n)$

[Brandt, Fischer, Hirvonen, Keller, Lempiäinen, Rybicki, Suomela, Uitto, STOC'16]

Fundamental Questions

- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?

- How can we find nontrivial fixed point relaxations?

Fundamental Questions

- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?

- How can we <u>find</u> nontrivial fixed point relaxations?
  - ... together with a suitable input $\mathcal{I}$?

Fundamental Questions

- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?
  - ... in the setting with input?

- How can we <u>find</u> nontrivial fixed point relaxations?
  - ... together with a suitable input $\mathcal{I}$?

Some Answers

- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?
  - ... in the setting with input?

We remove all known obstacles.

[Balliu, Brandt, Gabsdil, Olivetti, Suomela, SODA'26]

- How can we <u>find</u> nontrivial fixed point relaxations?
  - ... together with a suitable input $\mathcal{J}$?

Some Answers

Marks' Technique

- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?
  - ... in the setting with input?

We remove all known obstacles.

- How can we find nontrivial fixed point relaxations?
  - ... together with a suitable input $\mathcal{I}$?

Marks' Technique

Marks' Technique

assume that there is a
$o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player
games based on $\mathcal{A}$

↓

show: each possible distribution
of winning strategies implies
that $\mathcal{A}$ is incorrect

Marks' Technique

assume that there is a
$o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player
games based on $\mathcal{A}$

↓

show: each possible distribution
of winning strategies implies
that $\mathcal{A}$ is incorrect

Marks' Technique

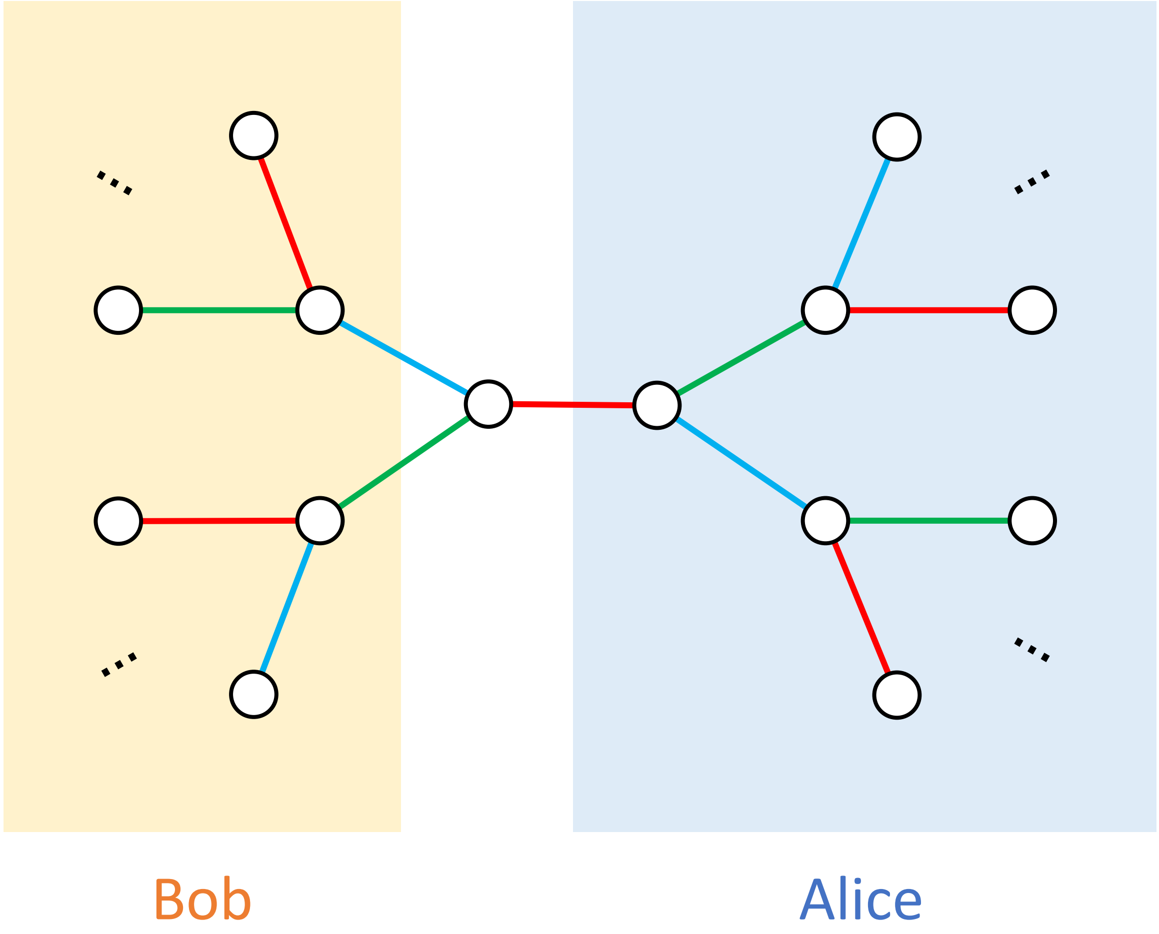

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

Marks' Technique

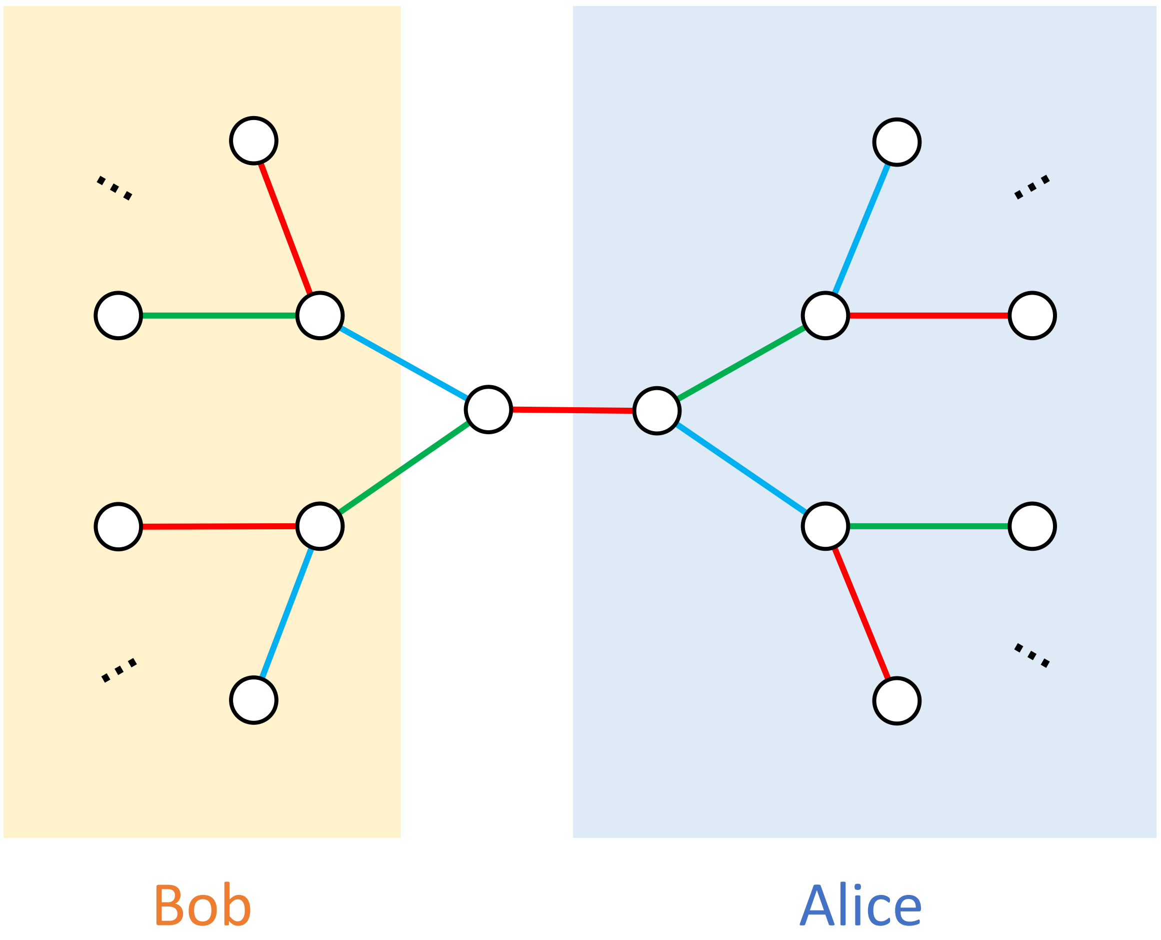

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

One game for each pair (ID, color).

Marks' Technique

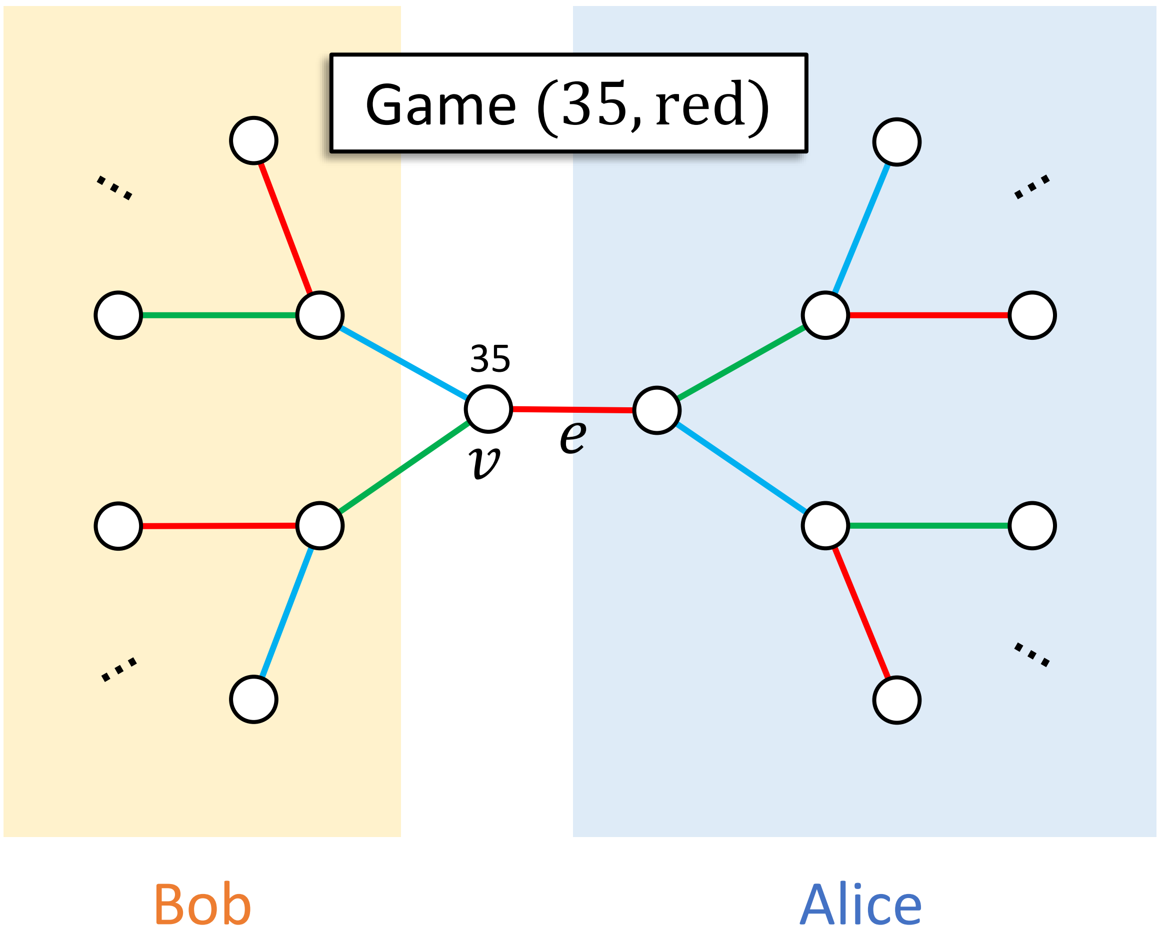

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

One game for each pair (ID, color).

Marks' Technique

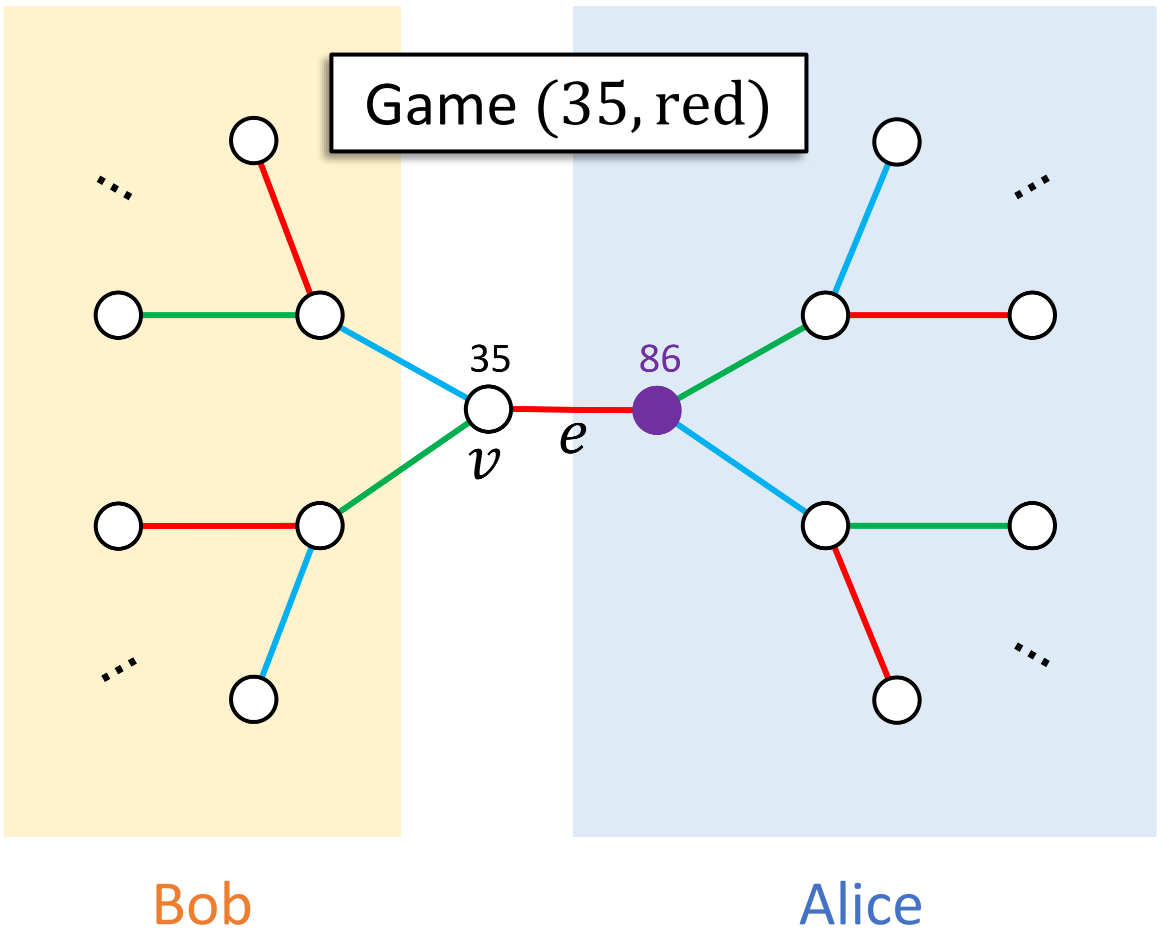

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

One game for each pair (ID, color).

Marks' Technique

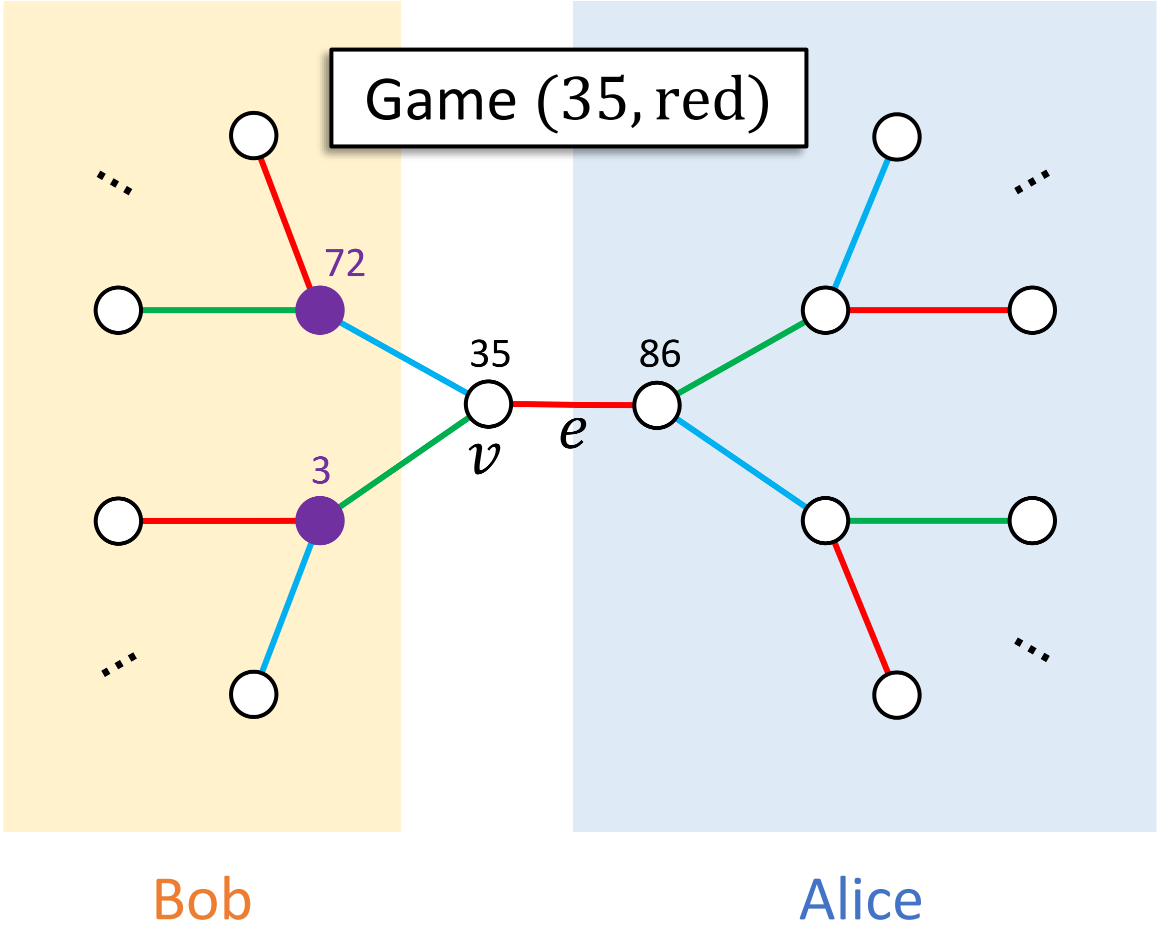

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

One game for each pair (ID, color).

Marks' Technique

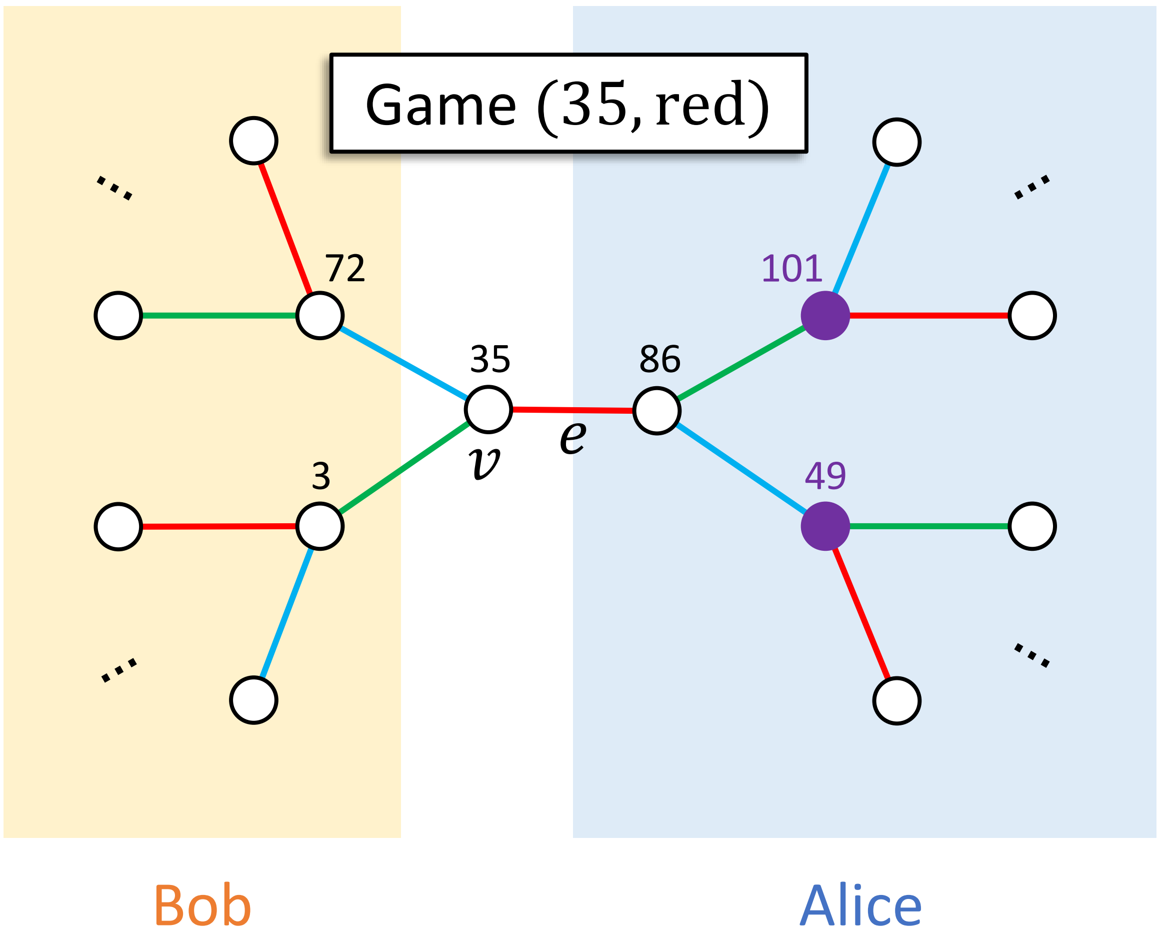

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

One game for each pair (ID, color).

Marks' Technique

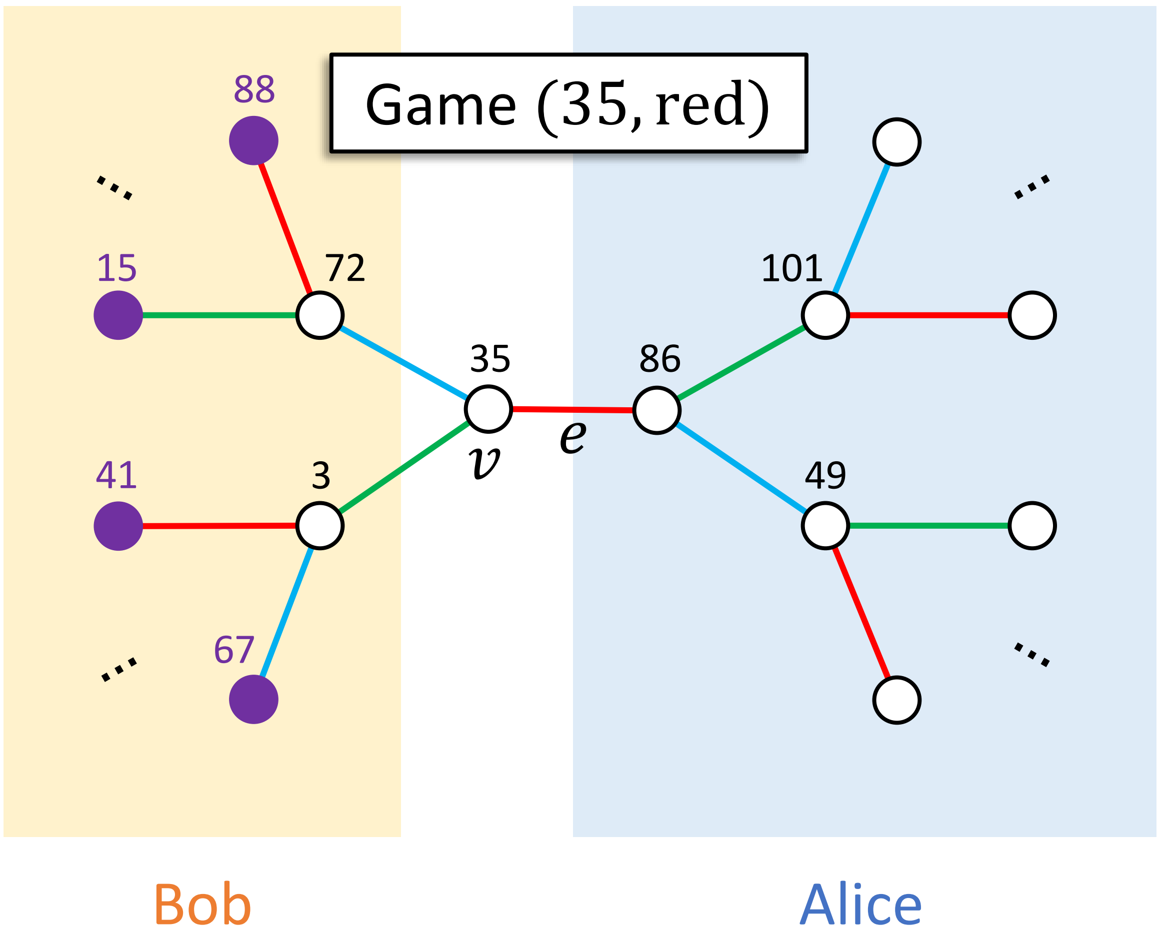

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

One game for each pair (ID, color).

Marks' Technique

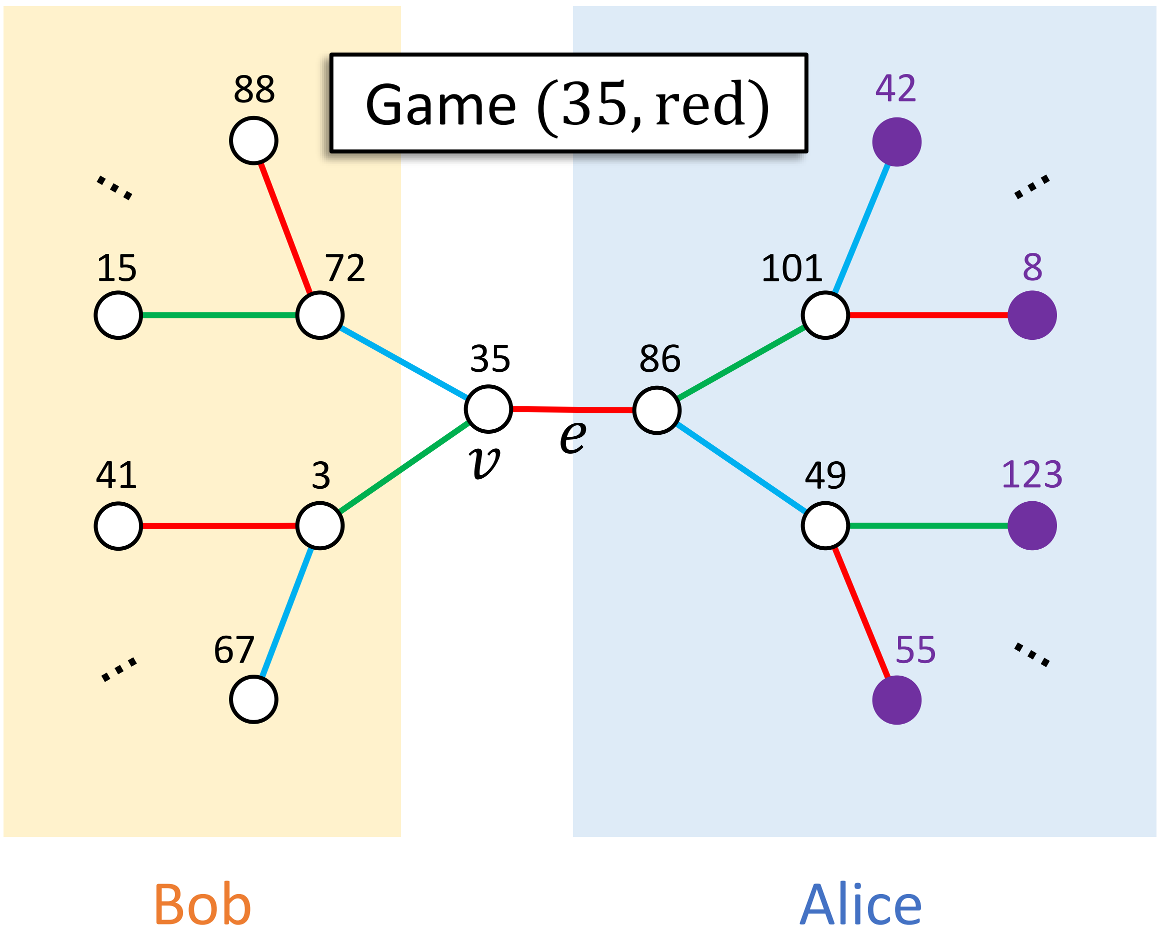

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

One game for each pair (ID, color).

Marks' Technique

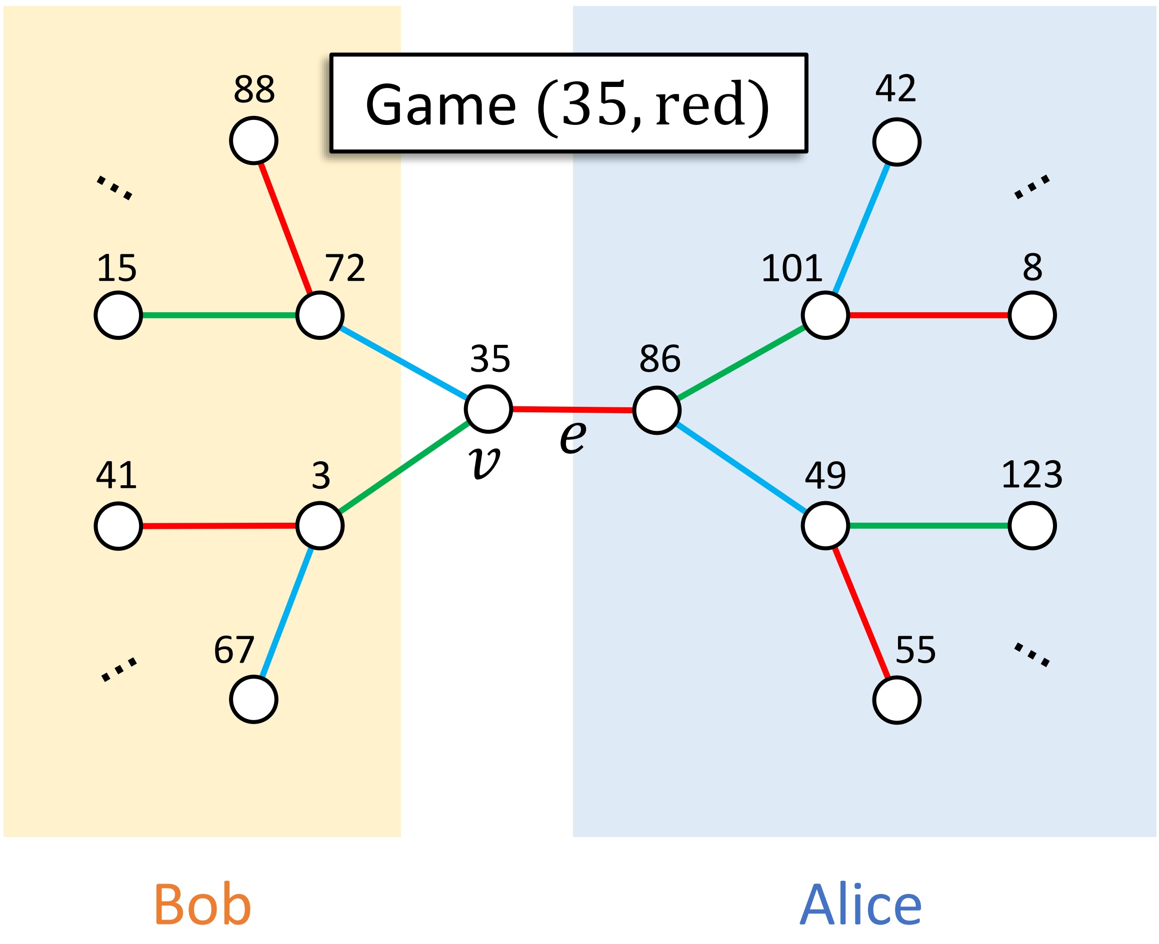

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

One game for each pair (ID, color).

Marks' Technique

Sinkless Orientation

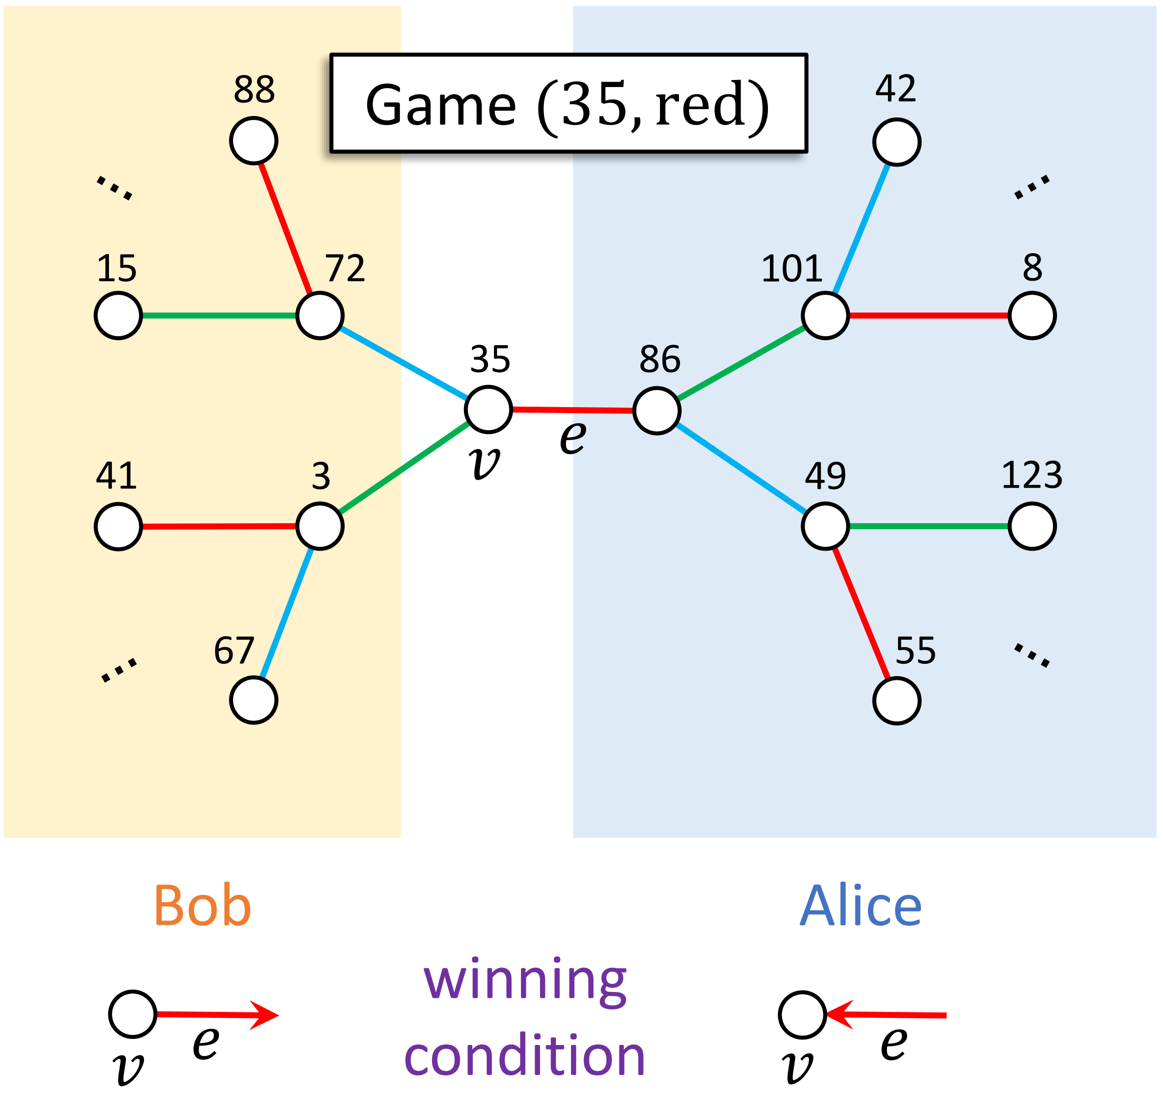

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

One game for each pair (ID, color).

Marks' Technique

Sinkless Orientation

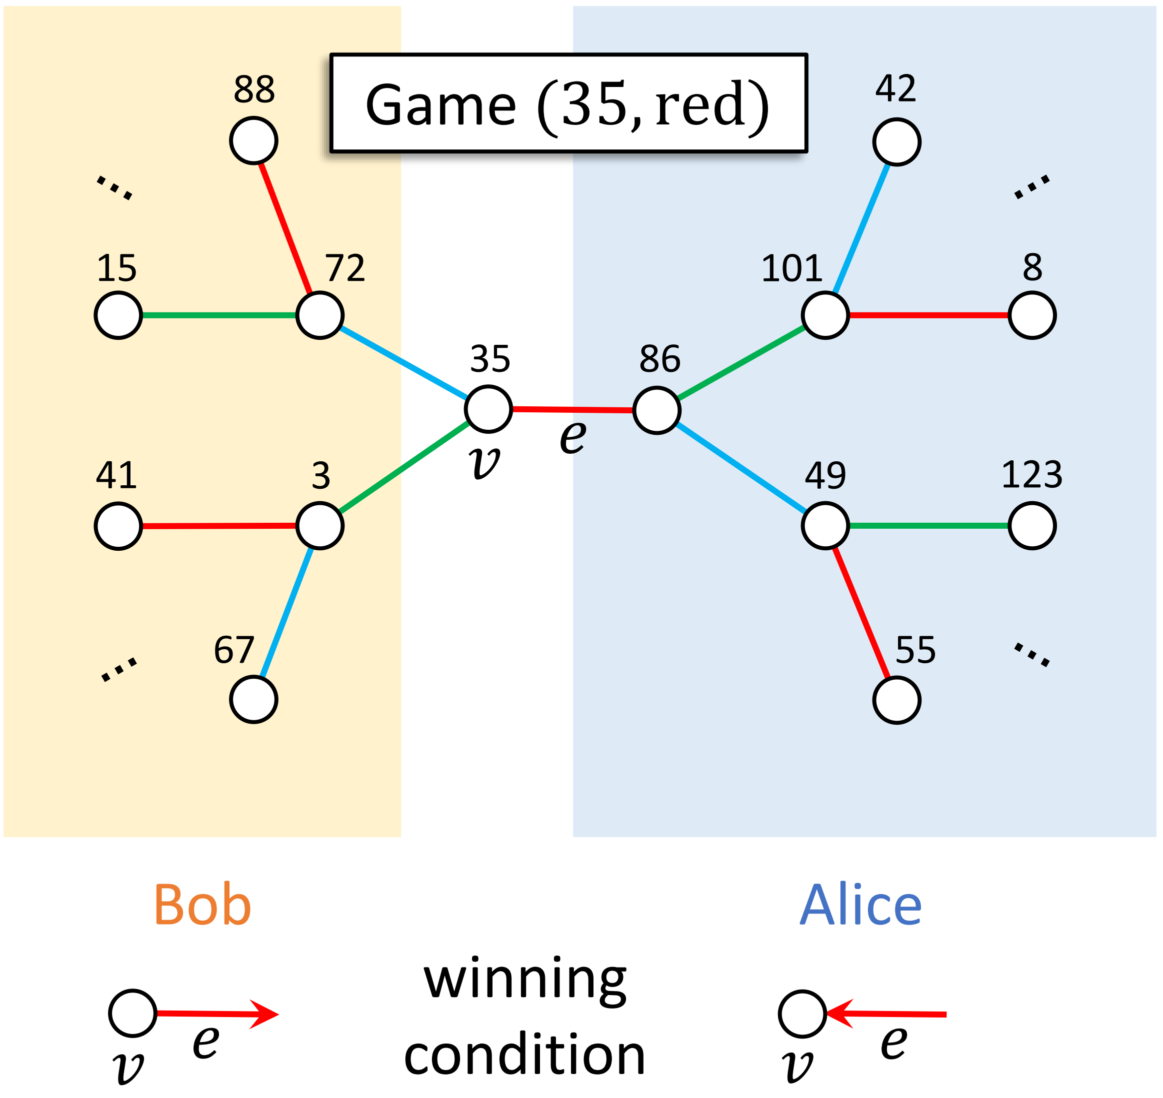

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

One game for each pair (ID, color).

Marks' Technique

Sinkless Orientation

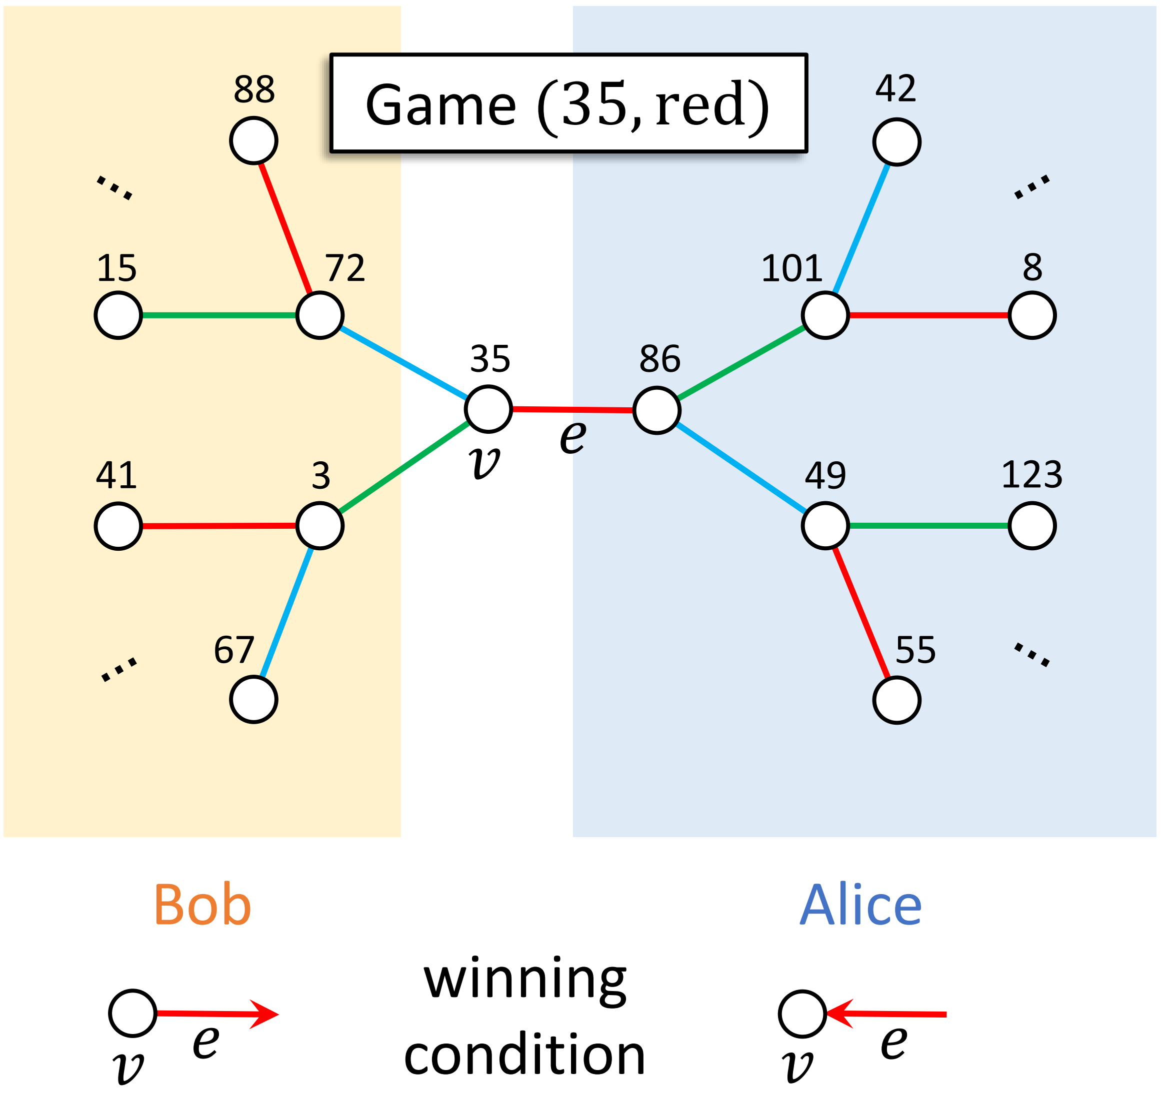

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

Can Alice win all of $(35, \text{red})$, $(35, \text{green})$, and $(35, \text{blue})$?

Marks' Technique

Sinkless Orientation

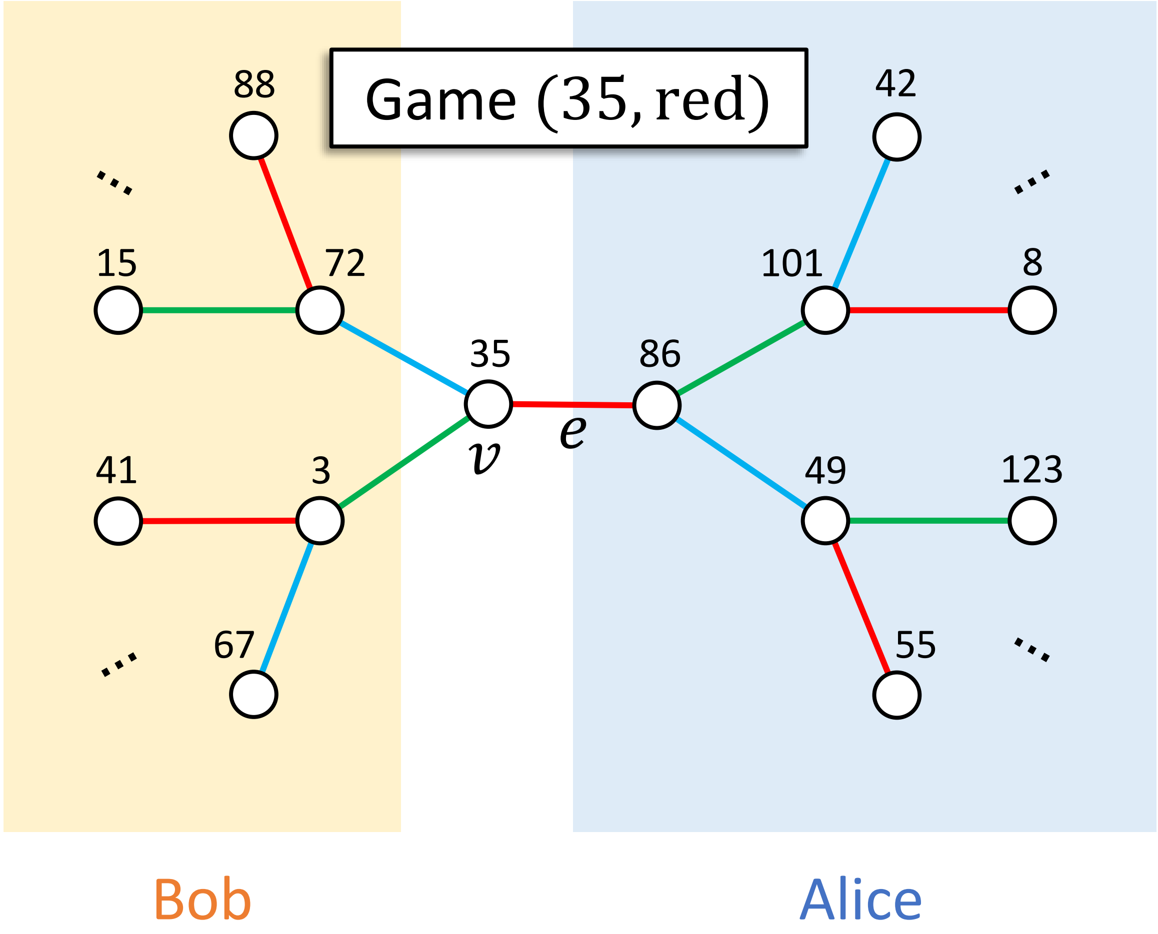

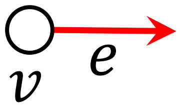

winning condition

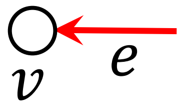

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

Can Alice win all of $(35, \text{red})$, $(35, \text{green})$, and $(35, \text{blue})$?

Let's play Alice's strategies against each other!

Marks' Technique

Sinkless Orientation

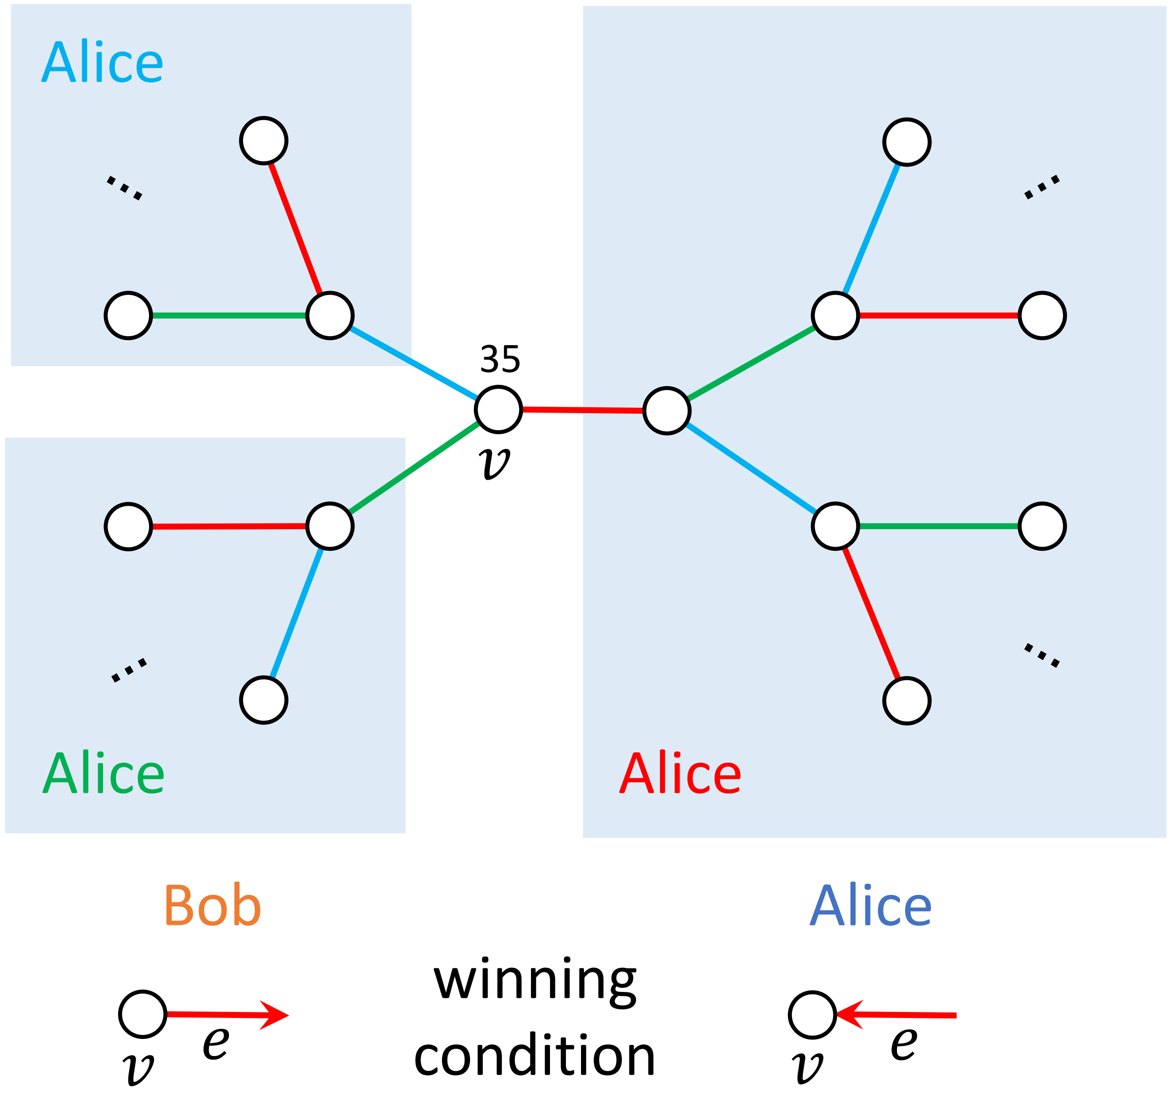

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

Can Alice win all of $(35, \text{red})$, $(35, \text{green})$, and $(35, \text{blue})$?

Let's play Alice's strategies against each other!

Marks' Technique

Sinkless Orientation

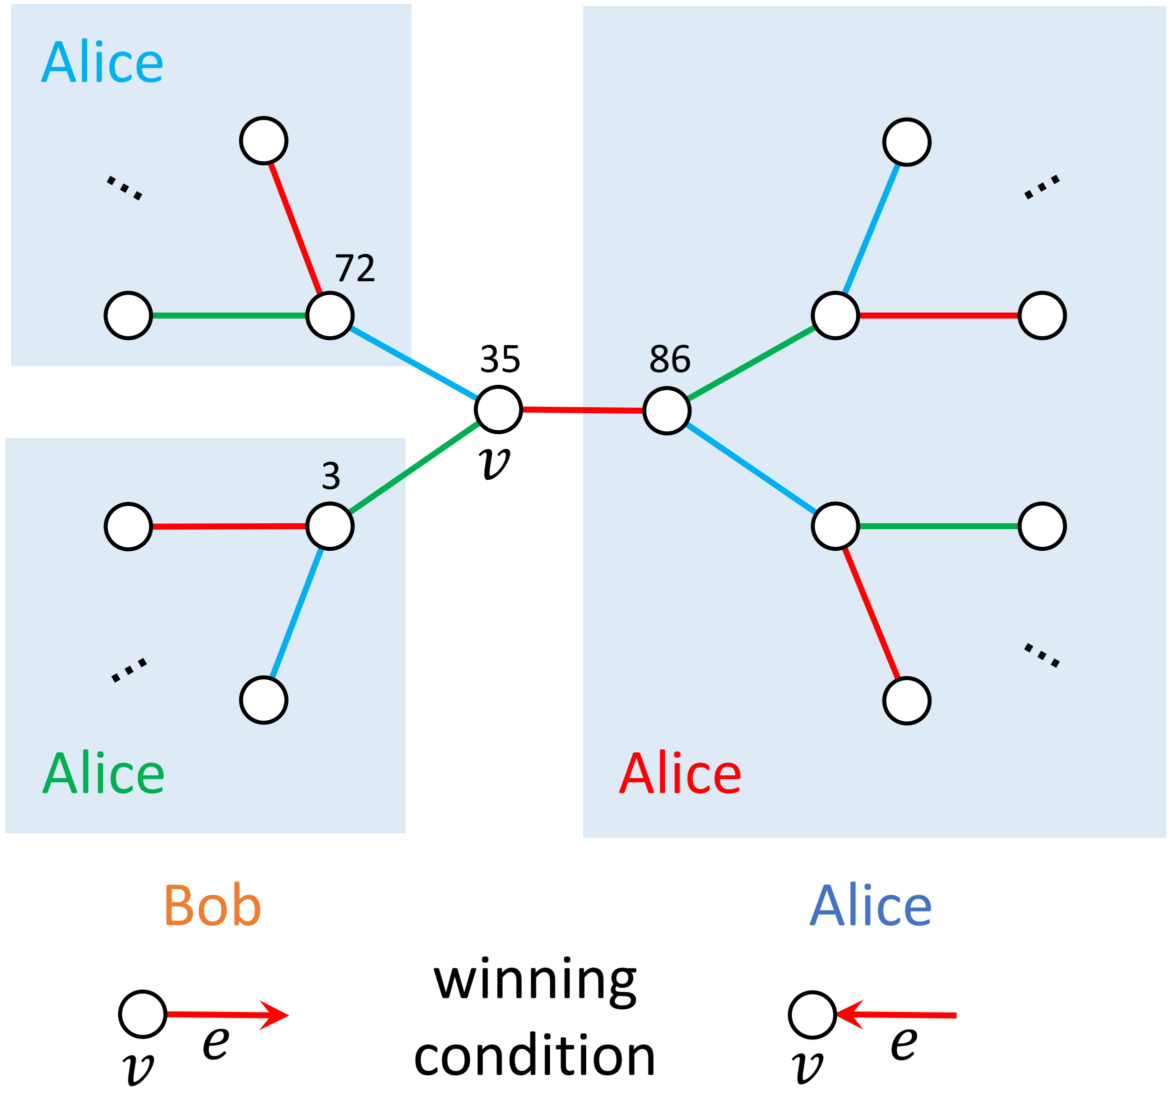

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

Can Alice win all of $(35, \text{red})$, $(35, \text{green})$, and $(35, \text{blue})$?

Let's play Alice's strategies against each other!

Marks' Technique

Sinkless Orientation

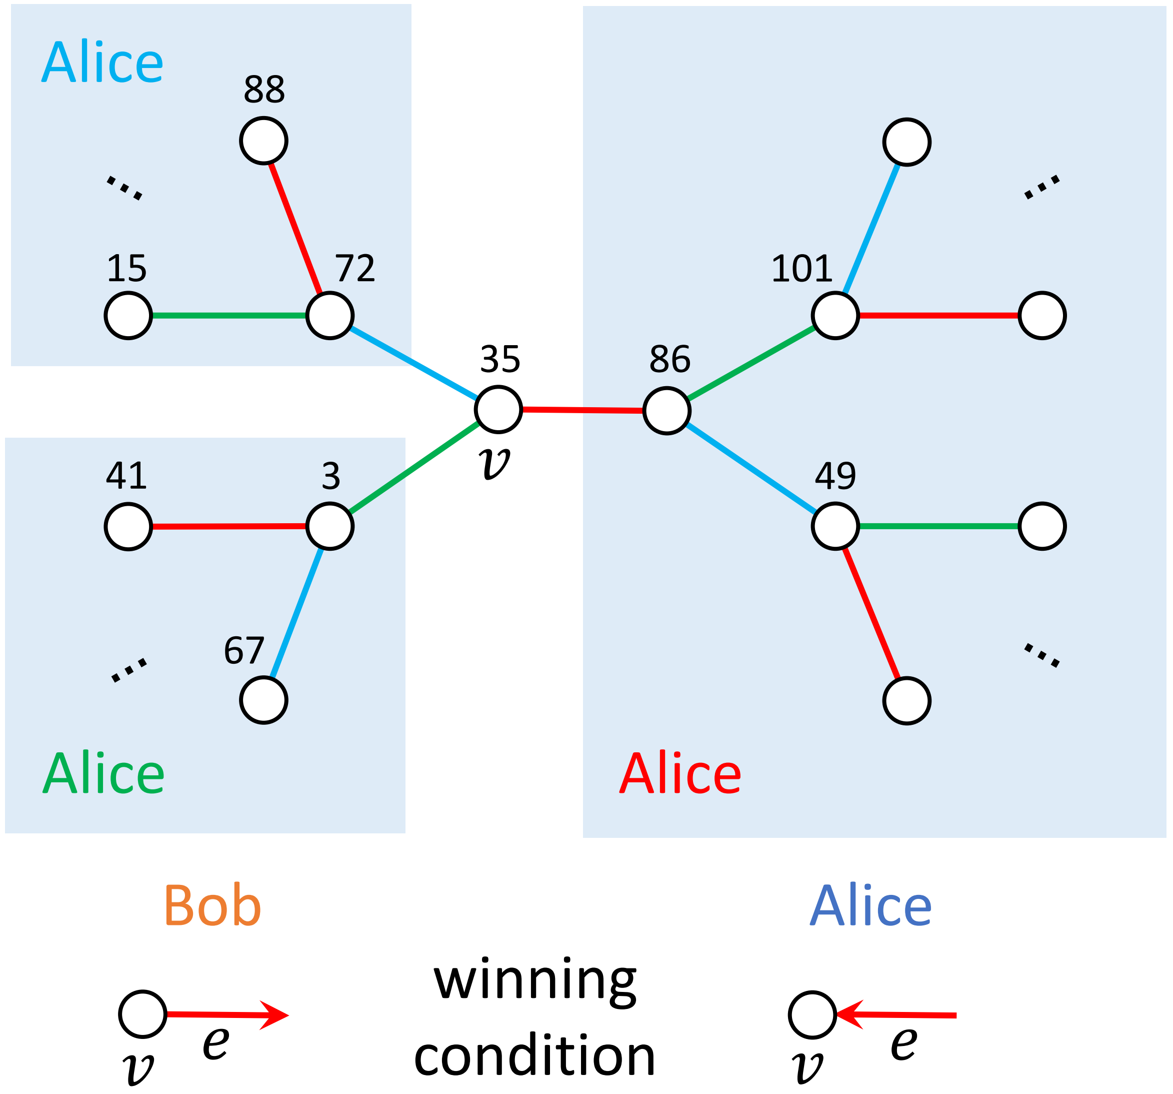

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

Can Alice win all of $(35, \text{red})$, $(35, \text{green})$, and $(35, \text{blue})$?

Let's play Alice's strategies against each other!

Marks' Technique

Sinkless Orientation

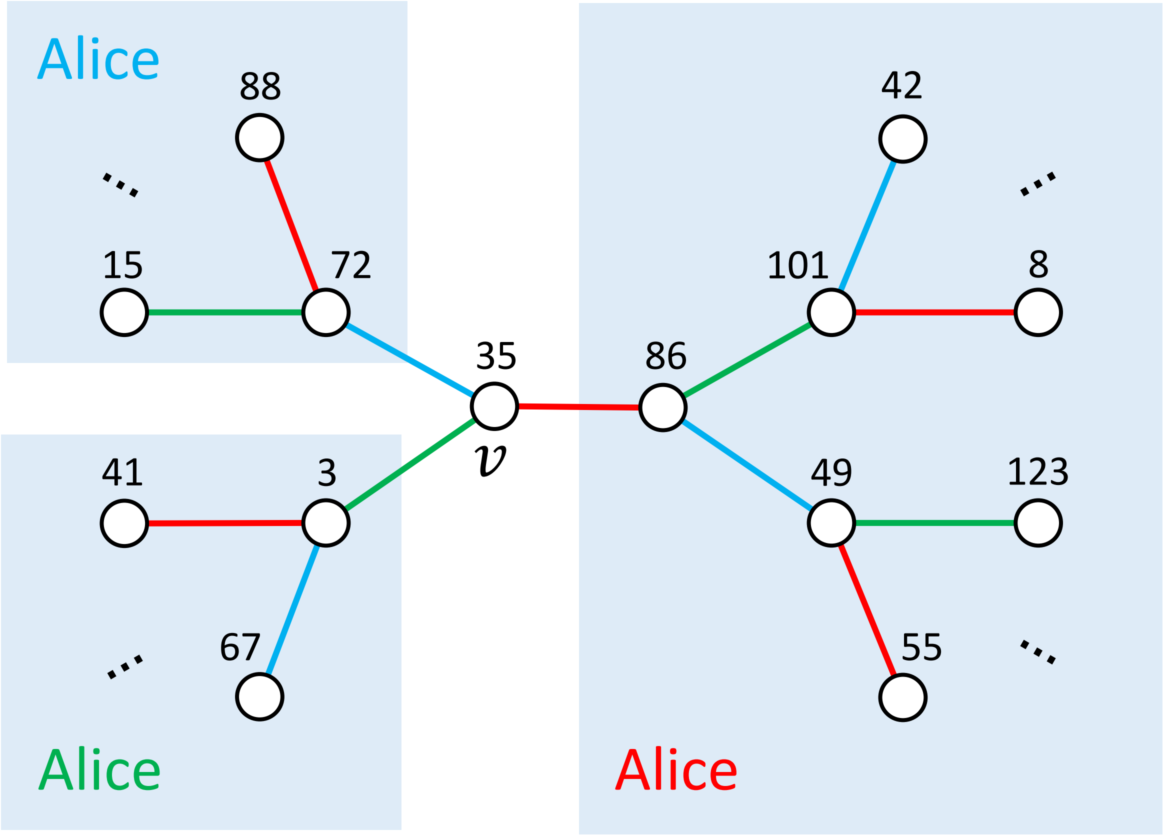

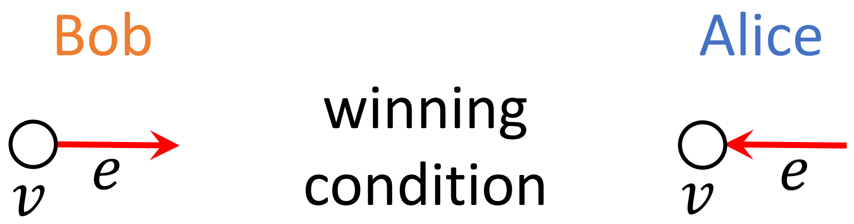

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

Can Alice win all of $(35, \mathrm{red})$, $(35, \mathrm{green})$, and $(35, \mathrm{blue})$?

Let's play Alice's strategies against each other!

Marks' Technique

Sinkless Orientation

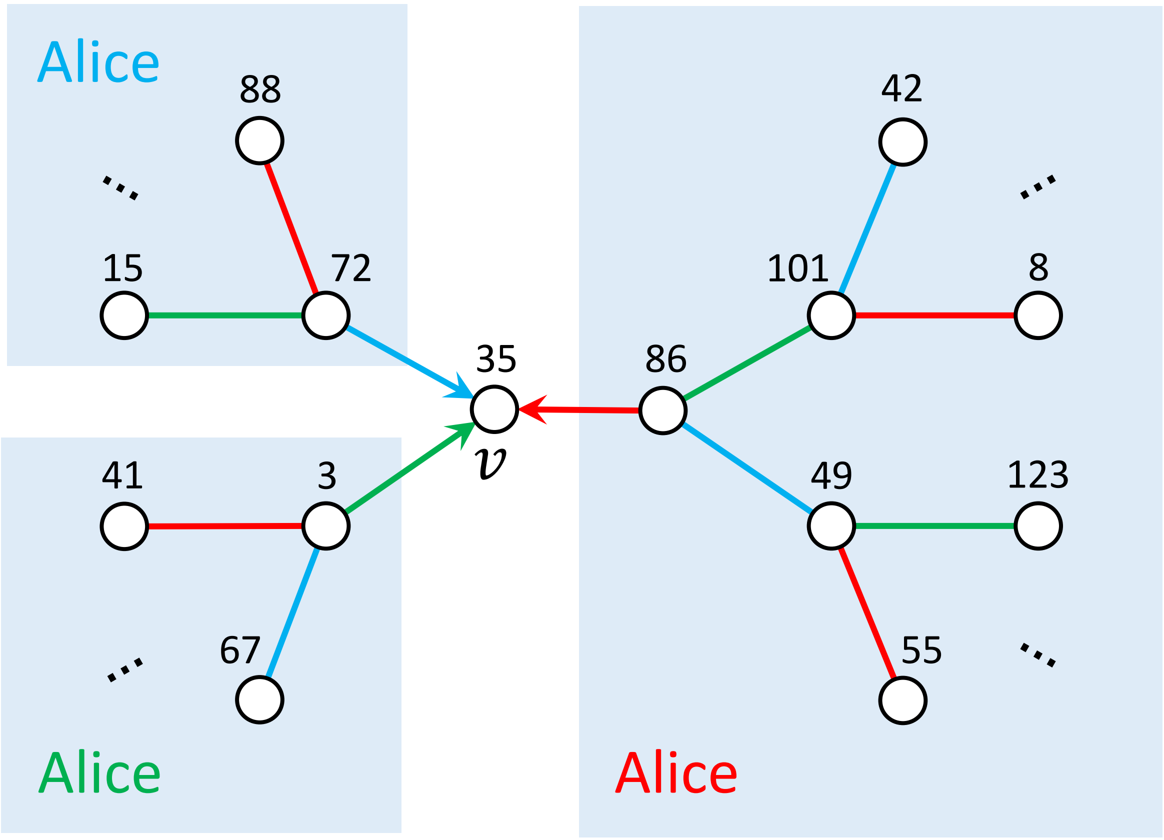

Bob

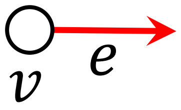

winning condition

Alice

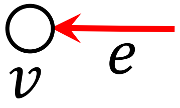

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

Can Alice win all of $(35, \text{red})$, $(35, \text{green})$, and $(35, \text{blue})$?

Let's play Alice's strategies against each other!

Marks' Technique

Sinkless Orientation

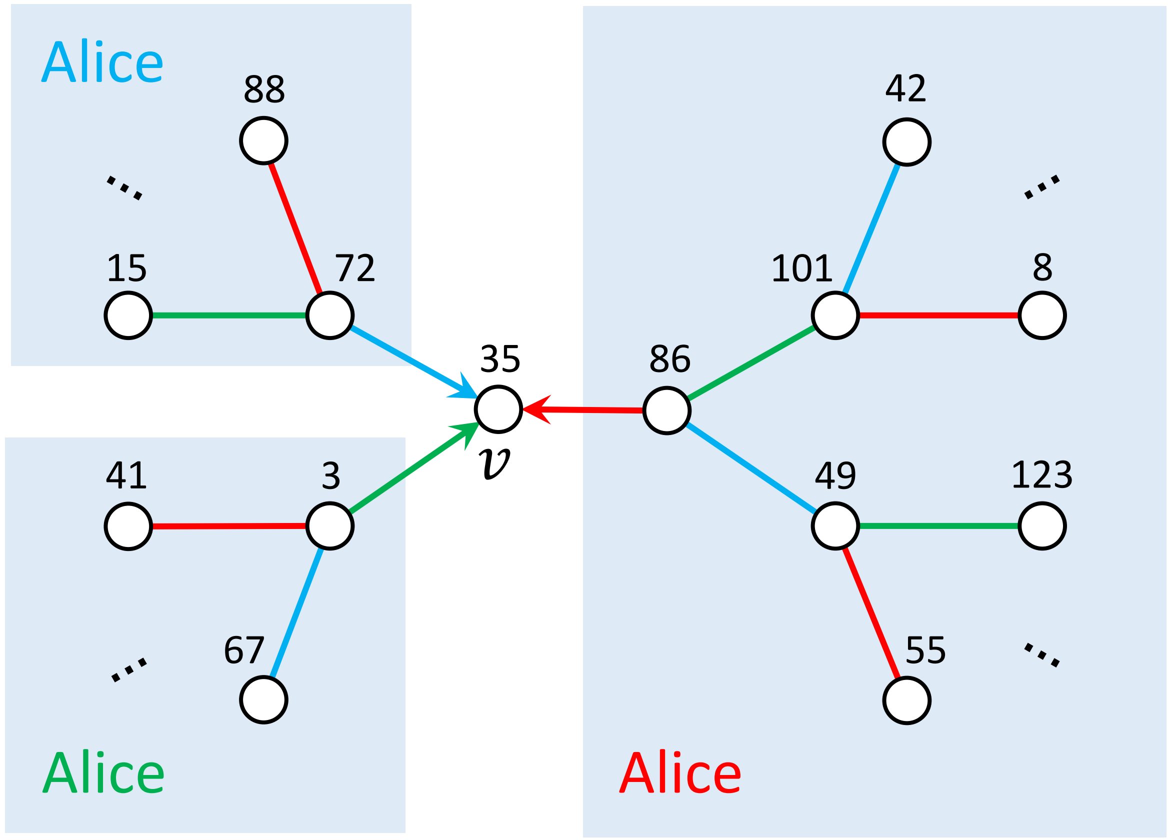

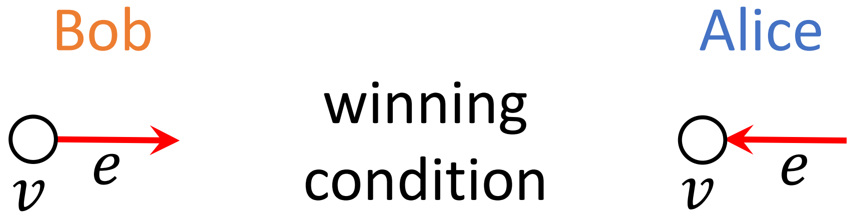

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

Can Alice win all of $(35, \mathrm{red})$, $(35, \mathrm{green})$, and $(35, \mathrm{blue})$?

No!

Let's play Alice's strategies against each other!

Marks' Technique

Sinkless Orientation

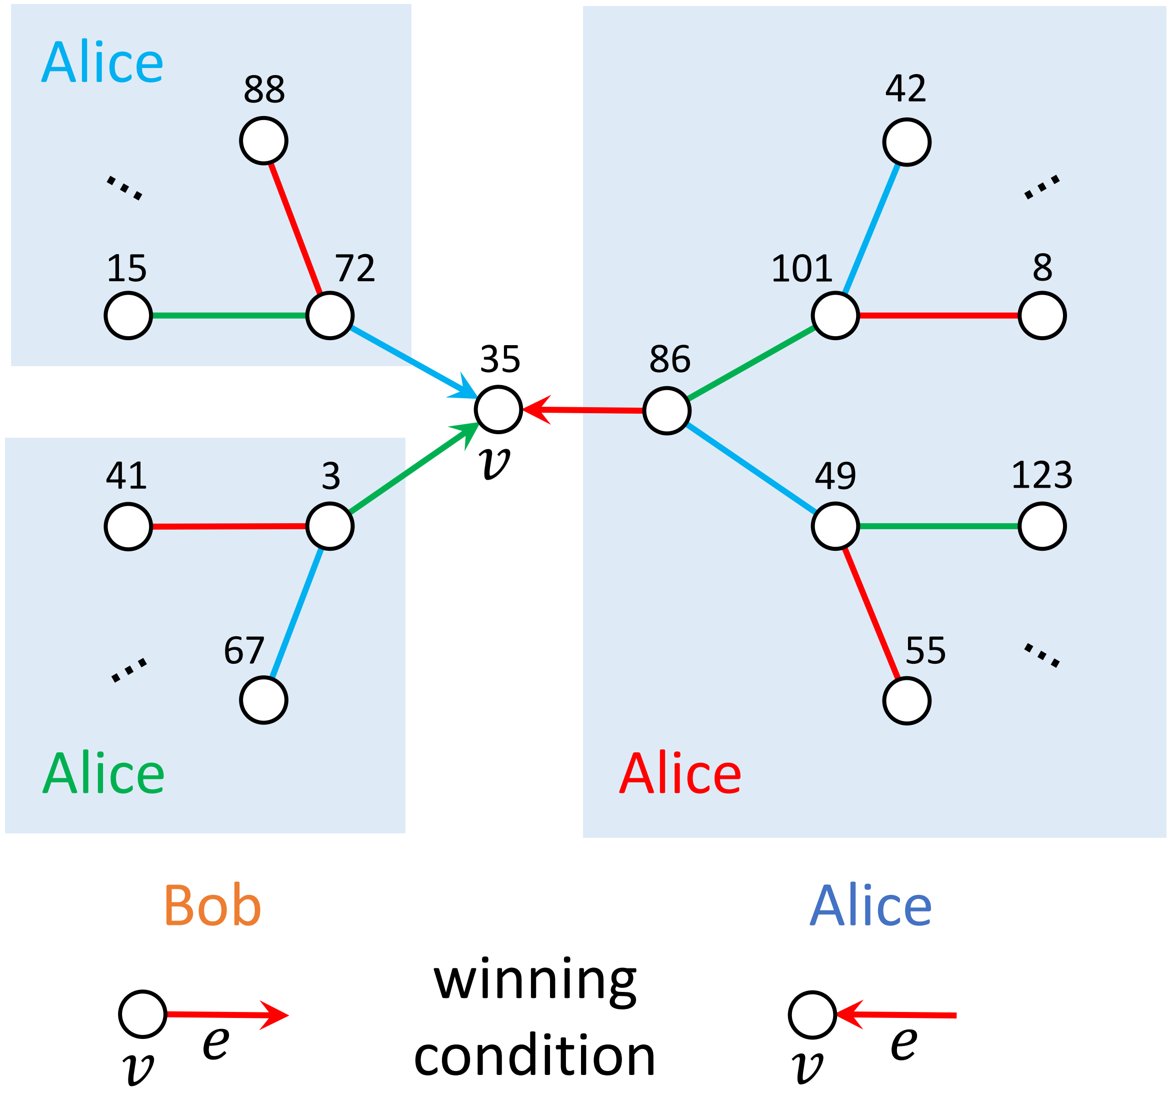

	red	green	blue
⋮			
35			
36			
37			
38			
⋮			

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

Can Alice win all of $(35, \text{red})$, $(35, \text{green})$, and $(35, \text{blue})$?

No!

Let's play Alice's strategies against each other!

Marks' Technique

Sinkless Orientation

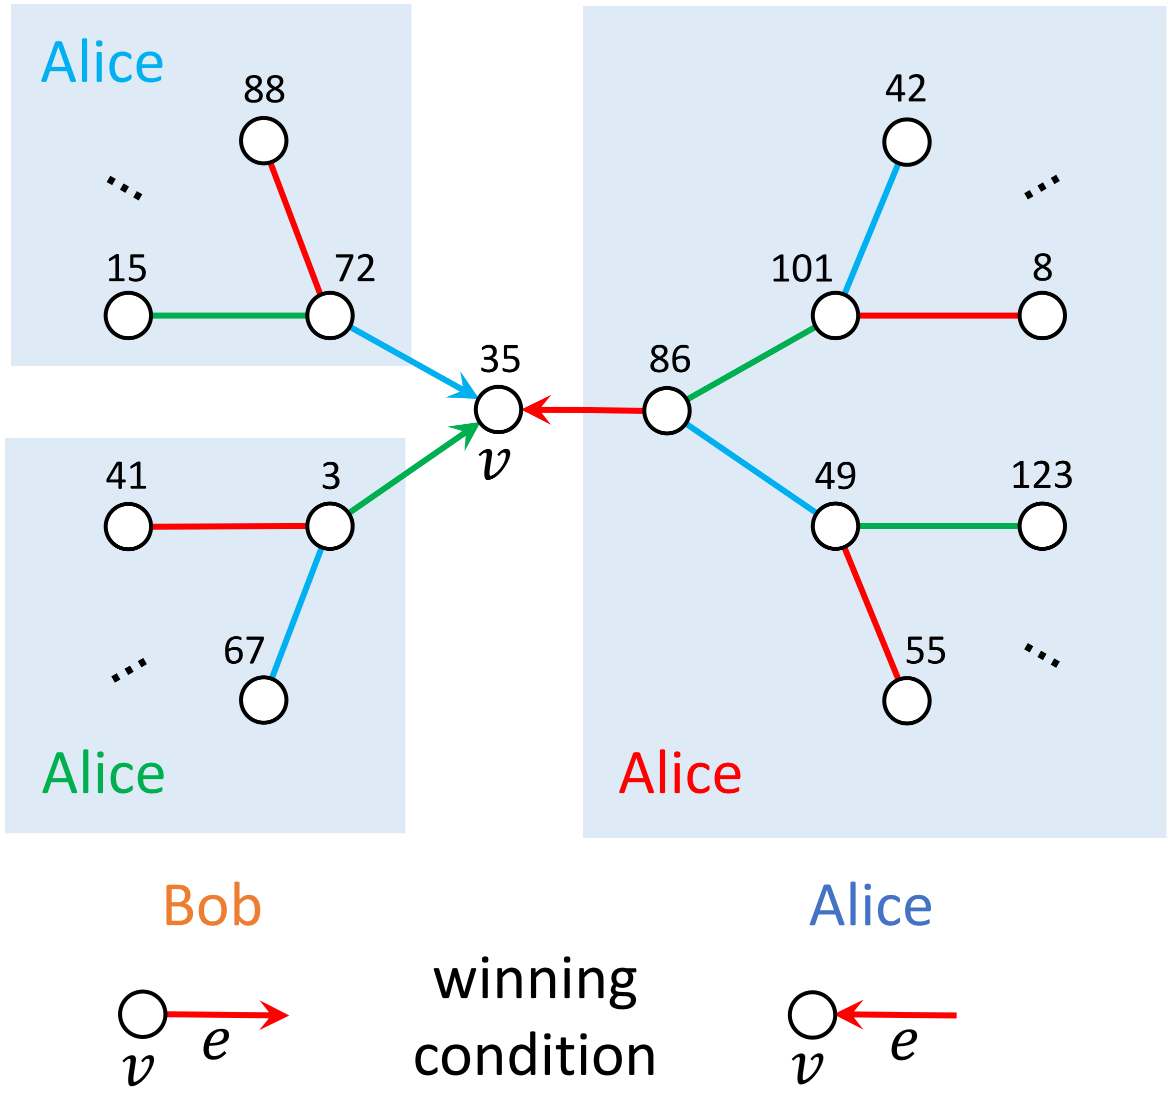

	red	green	blue
⋮			
35	B	A	A
36	A	A	B
37	A	B	A
38	B	A	A
⋮			

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

Can Alice win all of $(35, \text{red})$, $(35, \text{green})$, and $(35, \text{blue})$?

No!

Let's play Alice's strategies against each other!

Marks' Technique

Sinkless Orientation

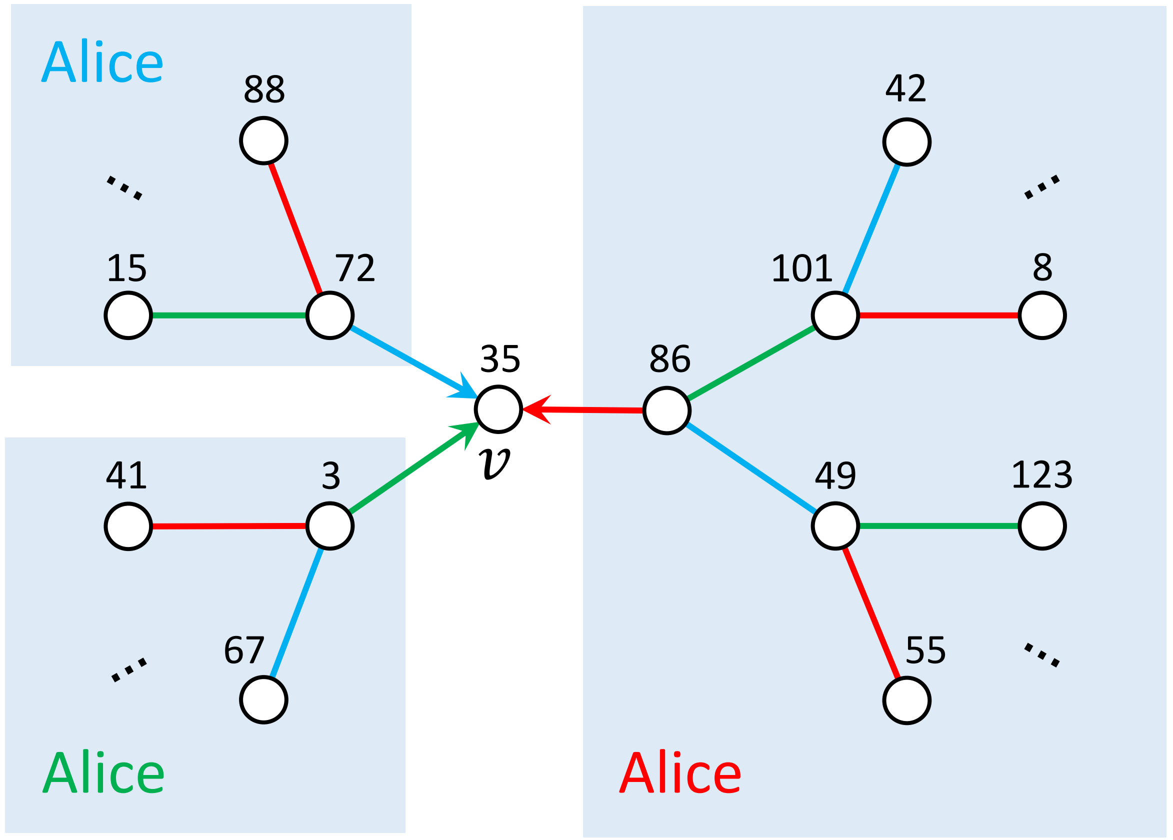

Bob

Alice

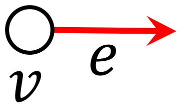

winning condition

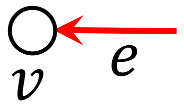

	red	green	blue
⋮			
35	B	A	A
36	A	A	B
37	A	B	A
38	B	A	A
⋮			

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

Can Alice win all of $(35, \text{red})$, $(35, \text{green})$, and $(35, \text{blue})$?

No!

Let's play Alice's strategies against each other!

Marks' Technique

Sinkless Orientation

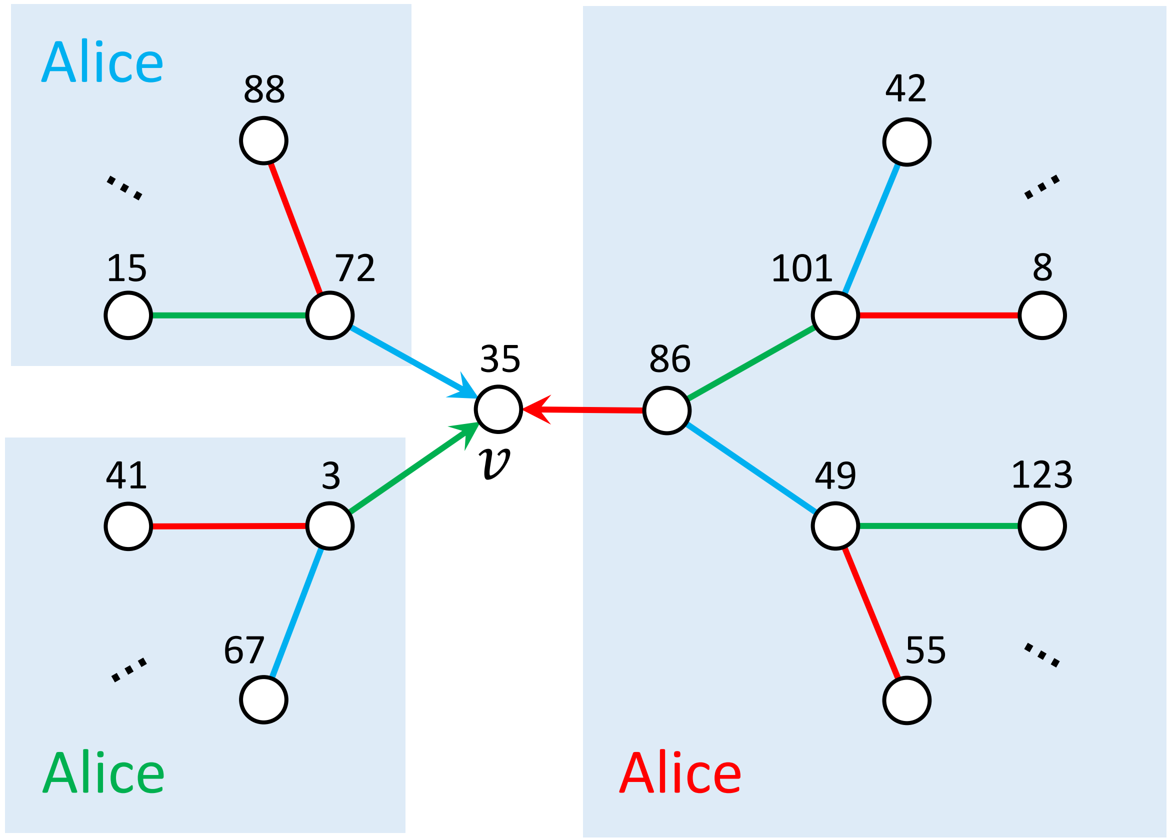

Bob

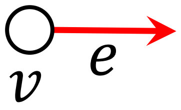

winning condition

Alice

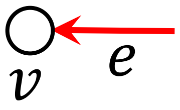

	red	green	blue
⋮			
35	B	A	A
36	A	A	B
37	A	B	A
38	B	A	A
⋮			

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

Can Bob win both of $(35, \mathrm{red})$ and $(38, \mathrm{red})$?

Let's play Bob's strategies against each other!

Marks' Technique

Sinkless Orientation

	red	green	blue
⋮			
35	B	A	A
36	A	A	B
37	A	B	A
38	B	A	A
⋮			

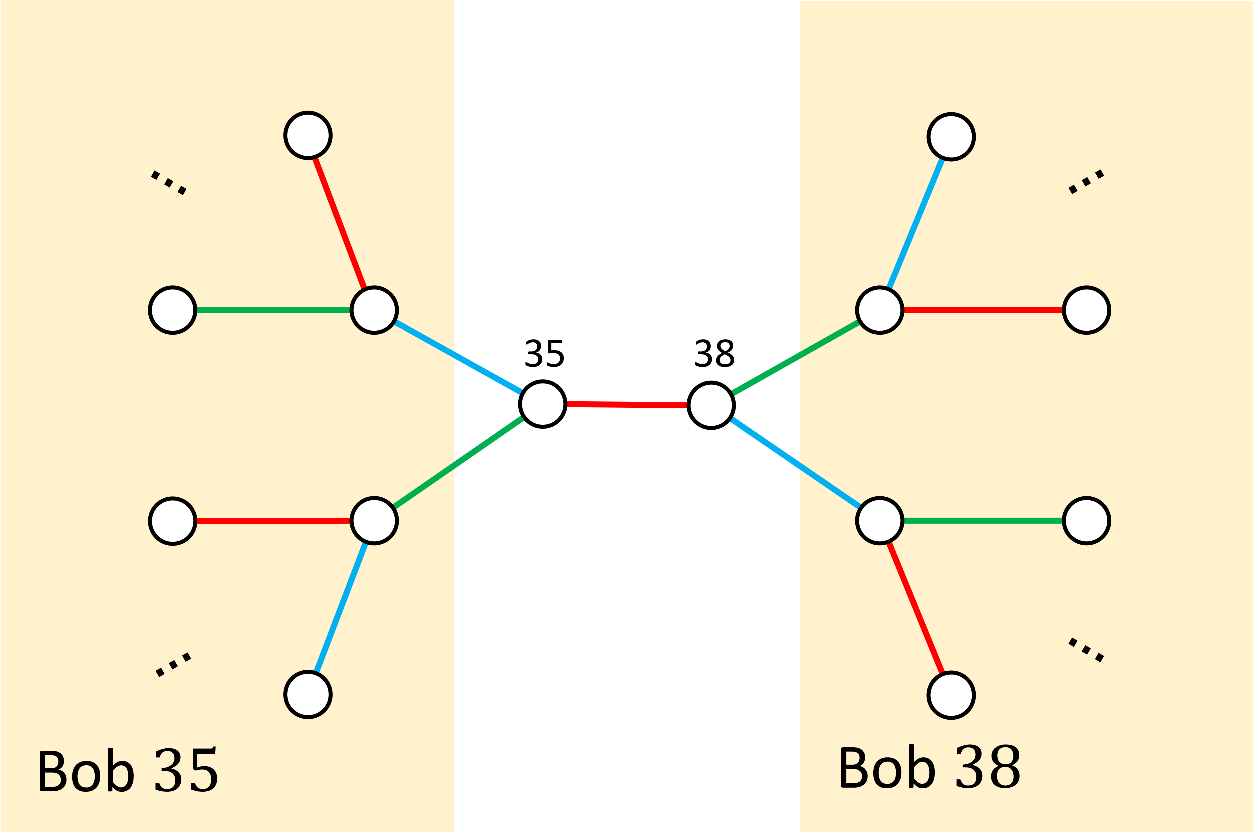

Bob

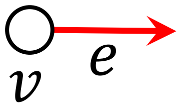

winning condition

Alice

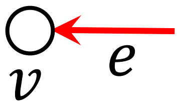

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

Can Bob win both of $(35, \mathrm{red})$ and $(38, \mathrm{red})$?

Let's play Bob's strategies against each other!

Marks' Technique

Sinkless Orientation

	red	green	blue
⋮			
35	B	A	A
36	A	A	B
37	A	B	A
38	B	A	A
⋮			

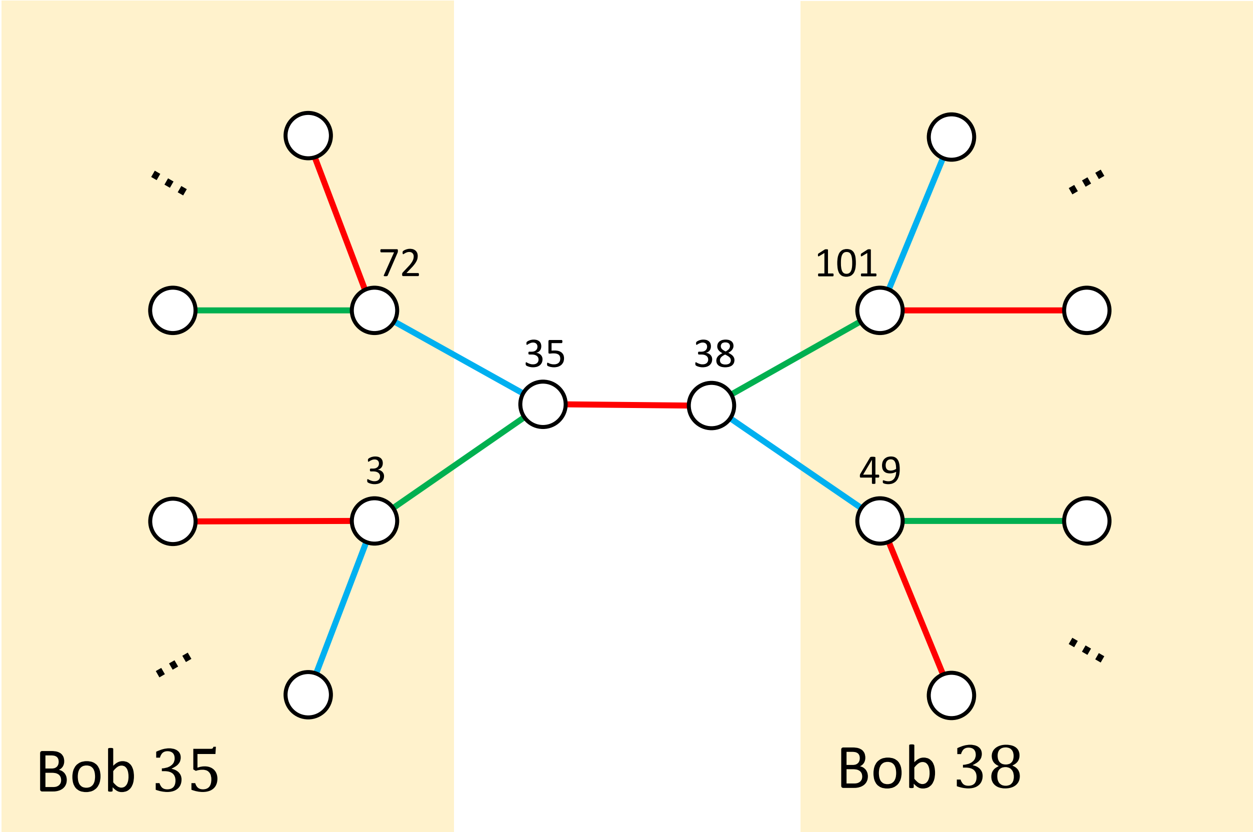

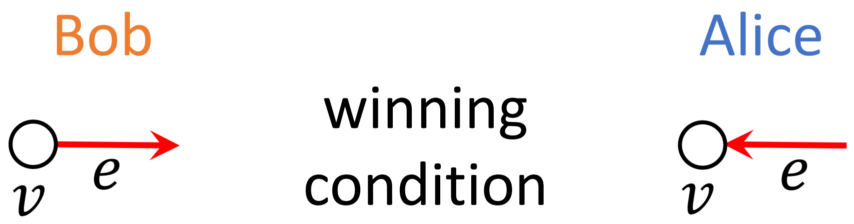

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

Can Bob win both of $(35, \mathrm{red})$ and $(38, \mathrm{red})$?

Let's play Bob's strategies against each other!

Marks' Technique

Sinkless Orientation

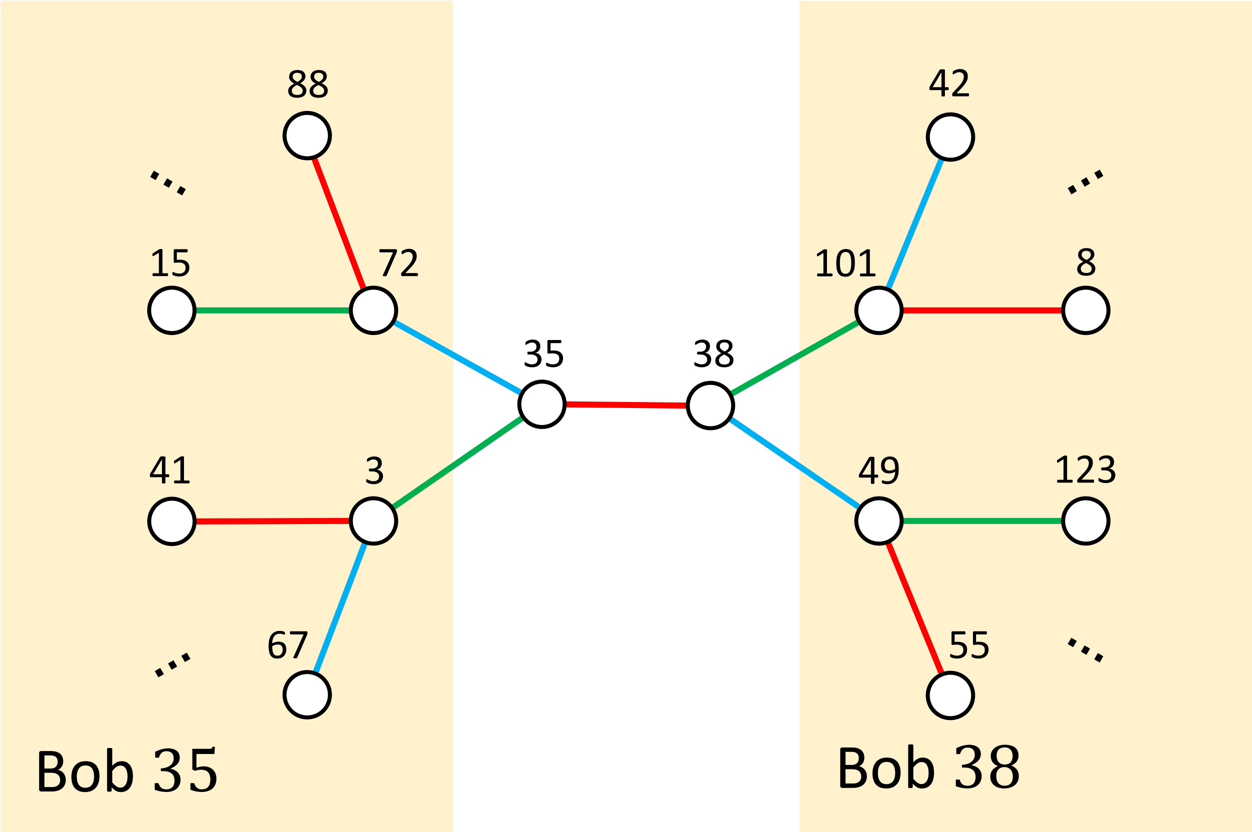

	red	green	blue
⋮			
35	B	A	A
36	A	A	B
37	A	B	A
38	B	A	A
⋮			

Bob

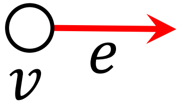

winning condition

Alice

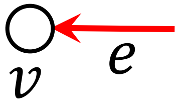

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

Can Bob win both of $(35, \text{red})$ and $(38, \text{red})$?

Let's play Bob's strategies against each other!

Marks' Technique

Sinkless Orientation

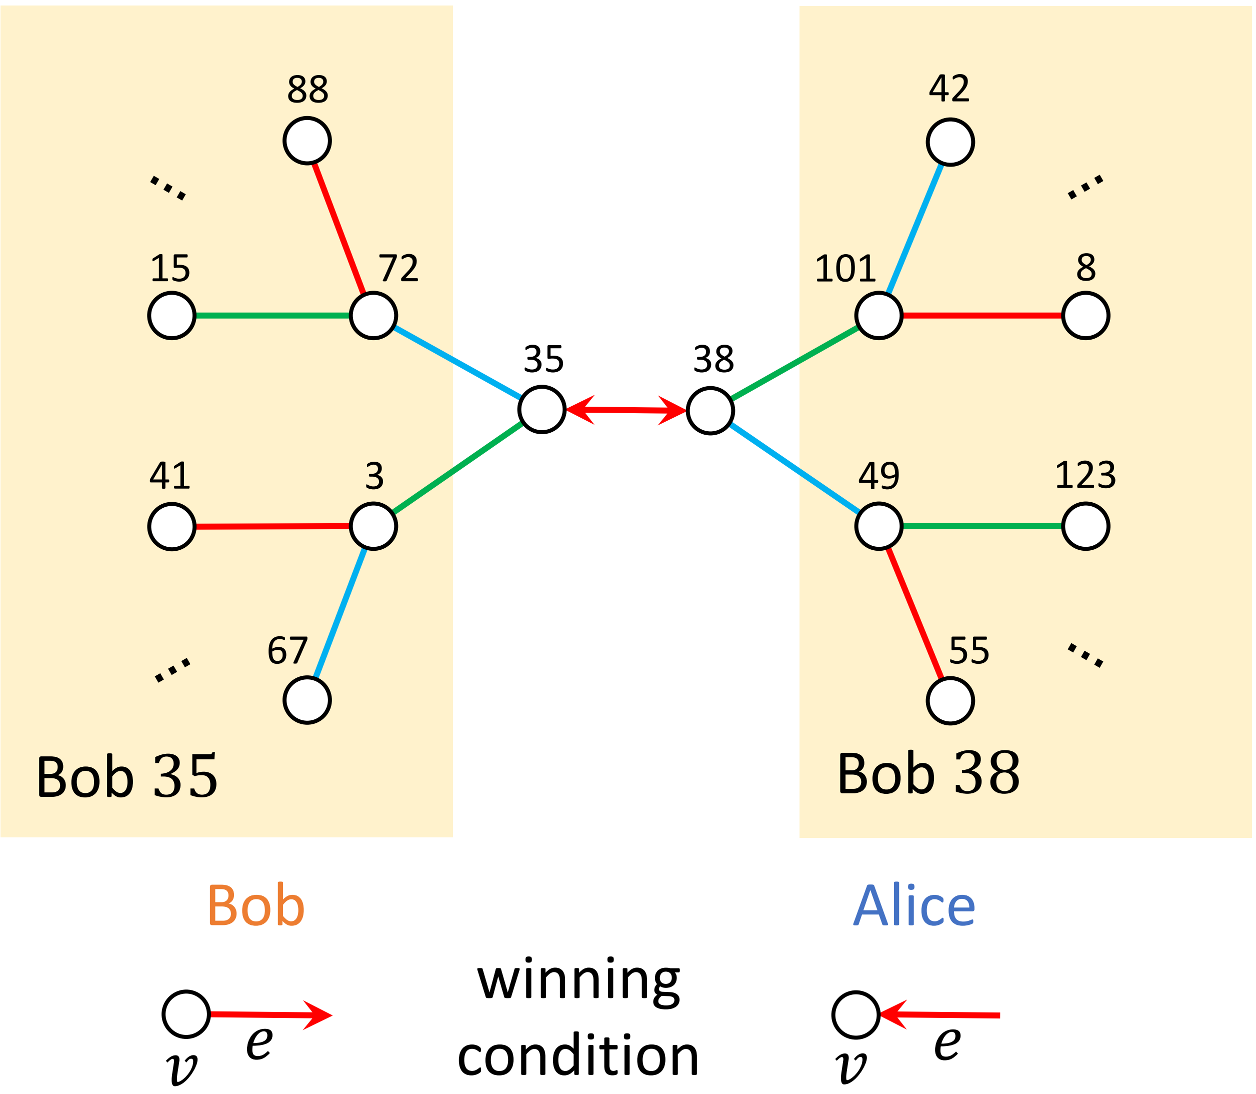

	red	green	blue
⋮			
35	B	A	A
36	A	A	B
37	A	B	A
38	B	A	A
⋮			

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

Can Bob win both of $(35, \mathrm{red})$ and $(38, \mathrm{red})$?

Let's play Bob's strategies against each other!

Marks' Technique

Sinkless Orientation

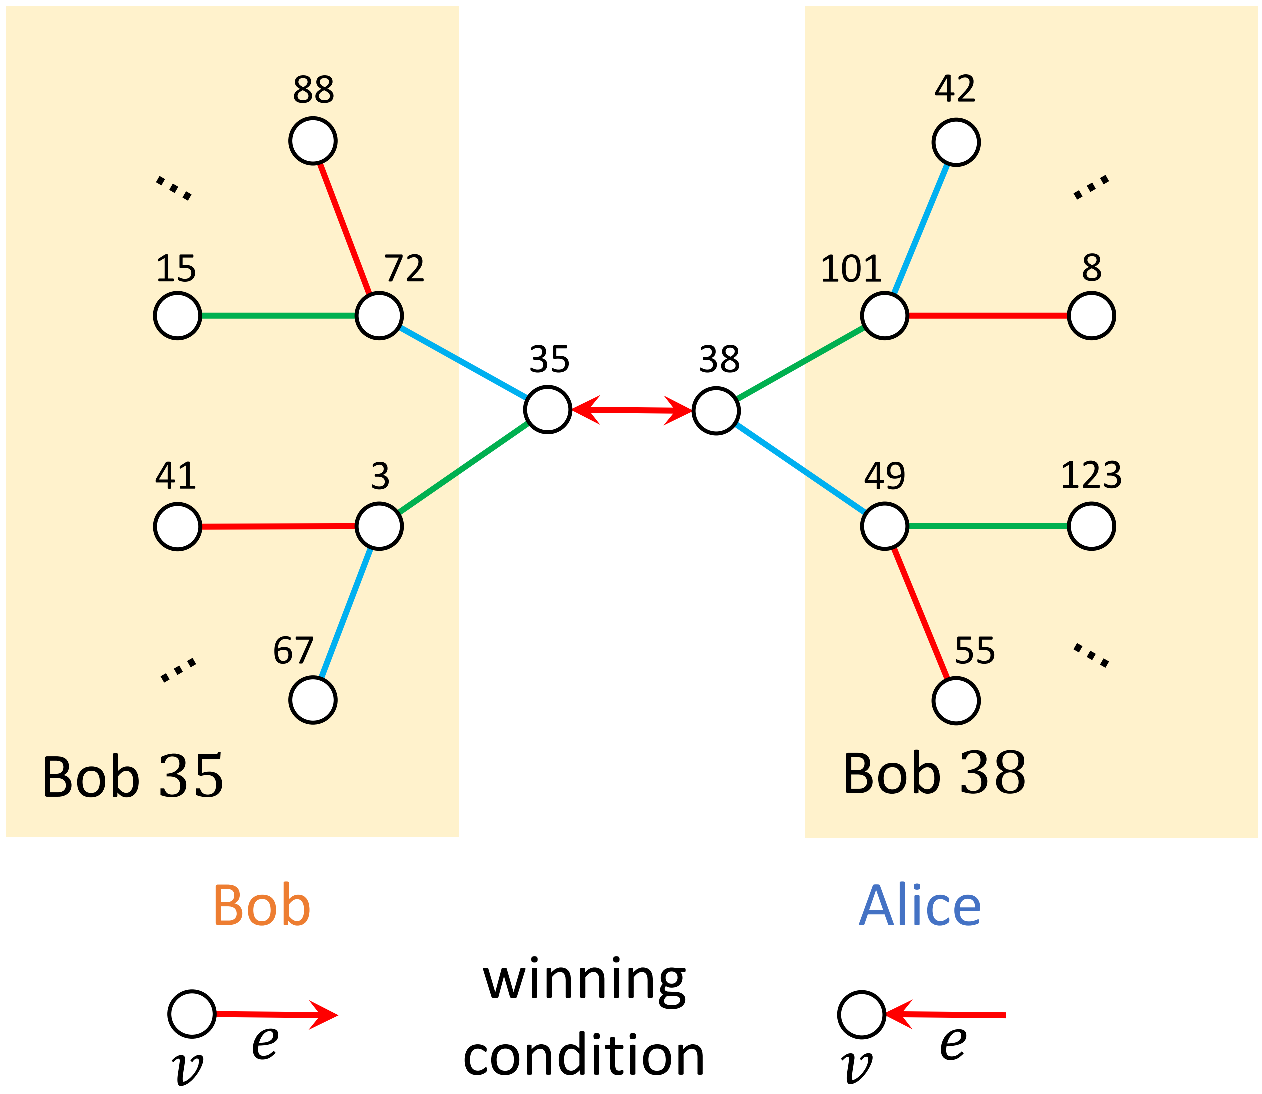

	red	green	blue
⋮			
35	B	A	A
36	A	A	B
37	A	B	A
38	B	A	A
⋮			

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect

Can Bob win both of $(35, \mathrm{red})$ and $(38, \mathrm{red})$?

No!

Let's play Bob's strategies against each other!

Marks' Technique

Sinkless Orientation

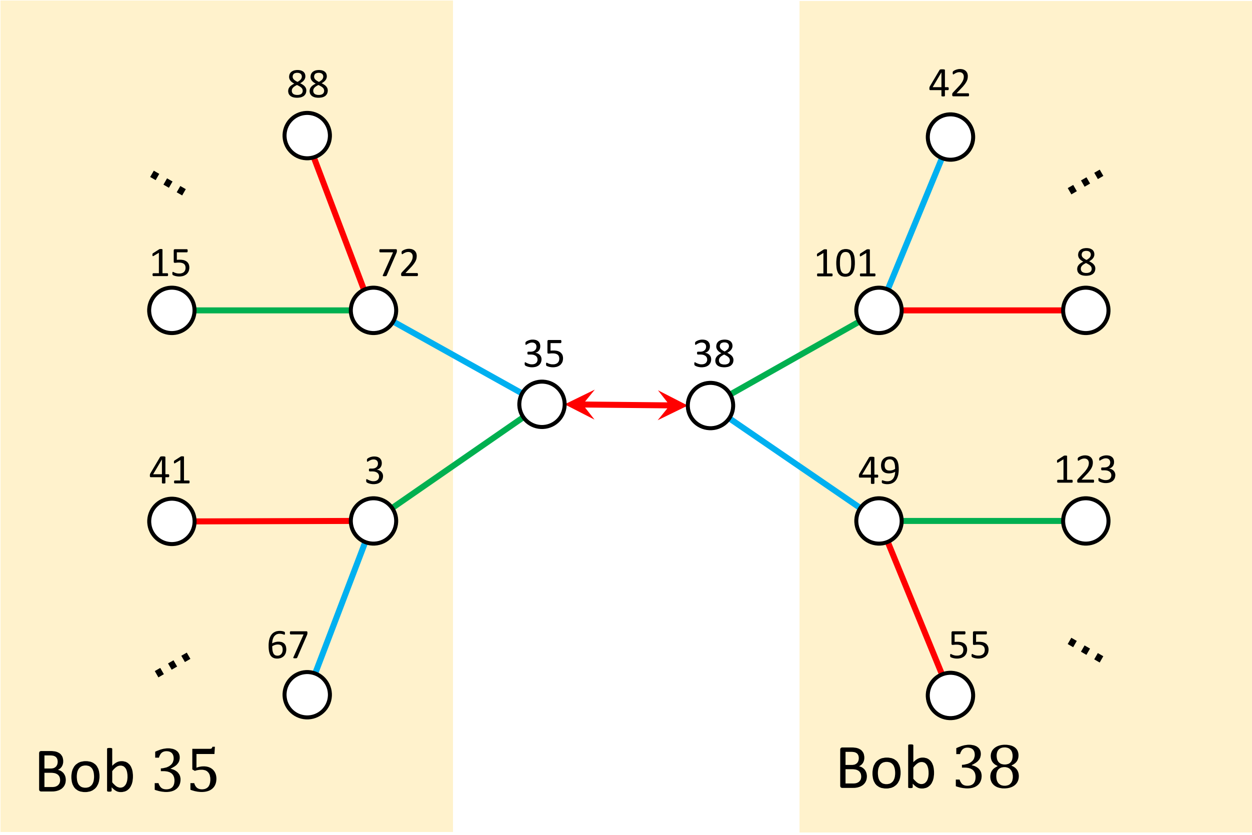

	red	green	blue
⋮			
35	B	A	A
36	A	A	B
37	A	B	A
38	B	A	A
⋮			

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect ✓

Can Bob win both of $(35, \text{red})$ and $(38, \text{red})$?

No!

Bob

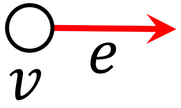

winning condition

Alice

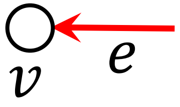

Let's play Bob's strategies against each other!

Marks' Technique

Sinkless Orientation

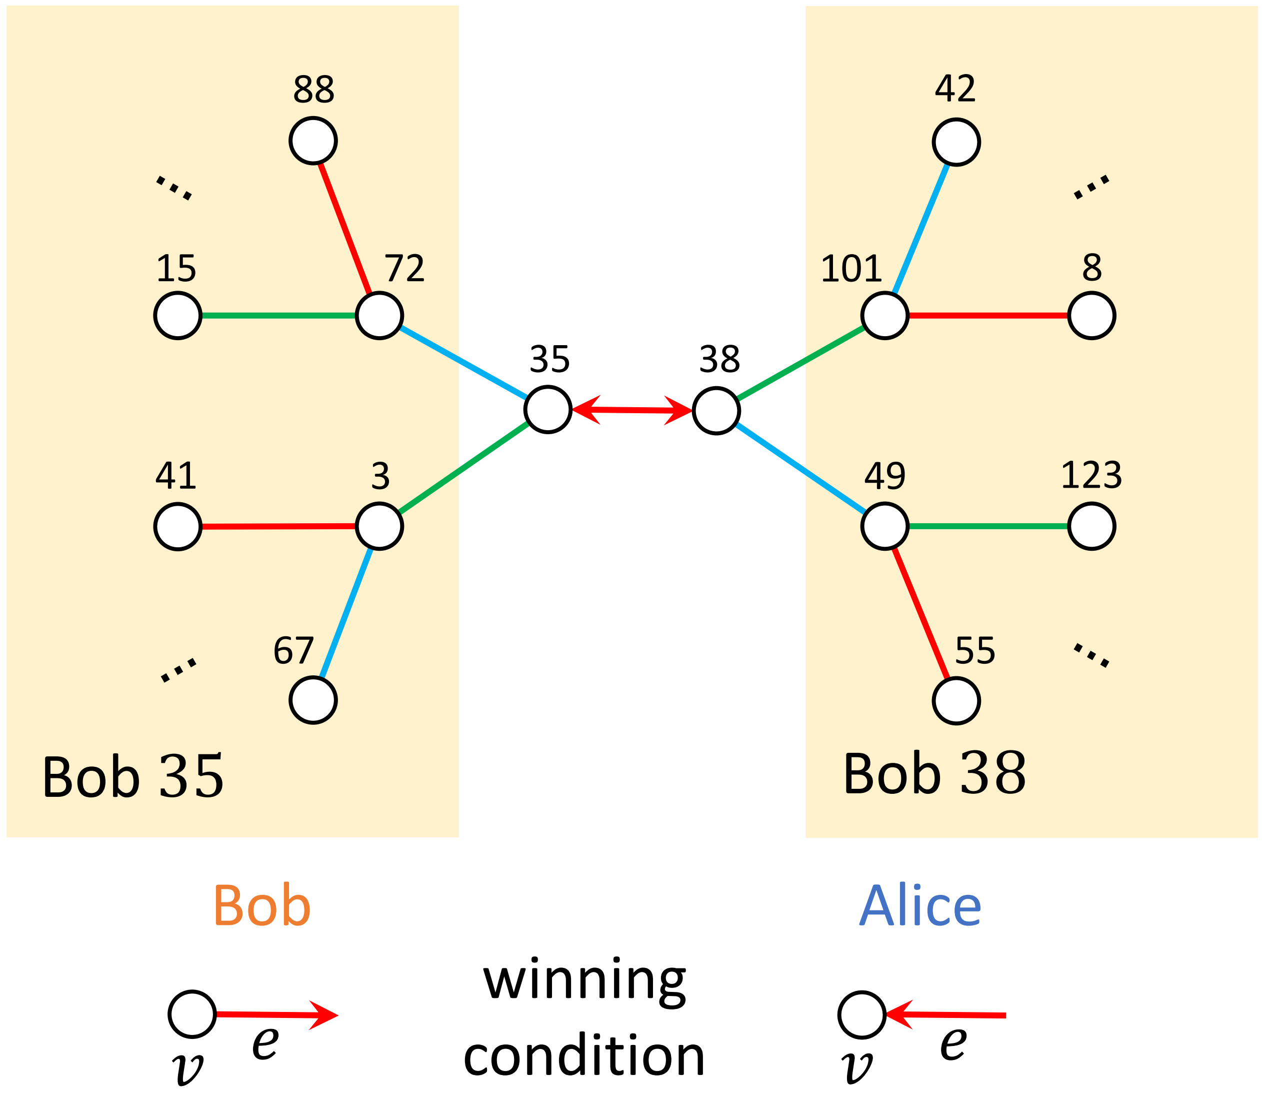

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect ✓

Can Bob win both of $(35, \text{red})$ and $(38, \text{red})$?

No!

Let's play Bob's strategies against each other!

Marks' Technique

Sinkless Orientation

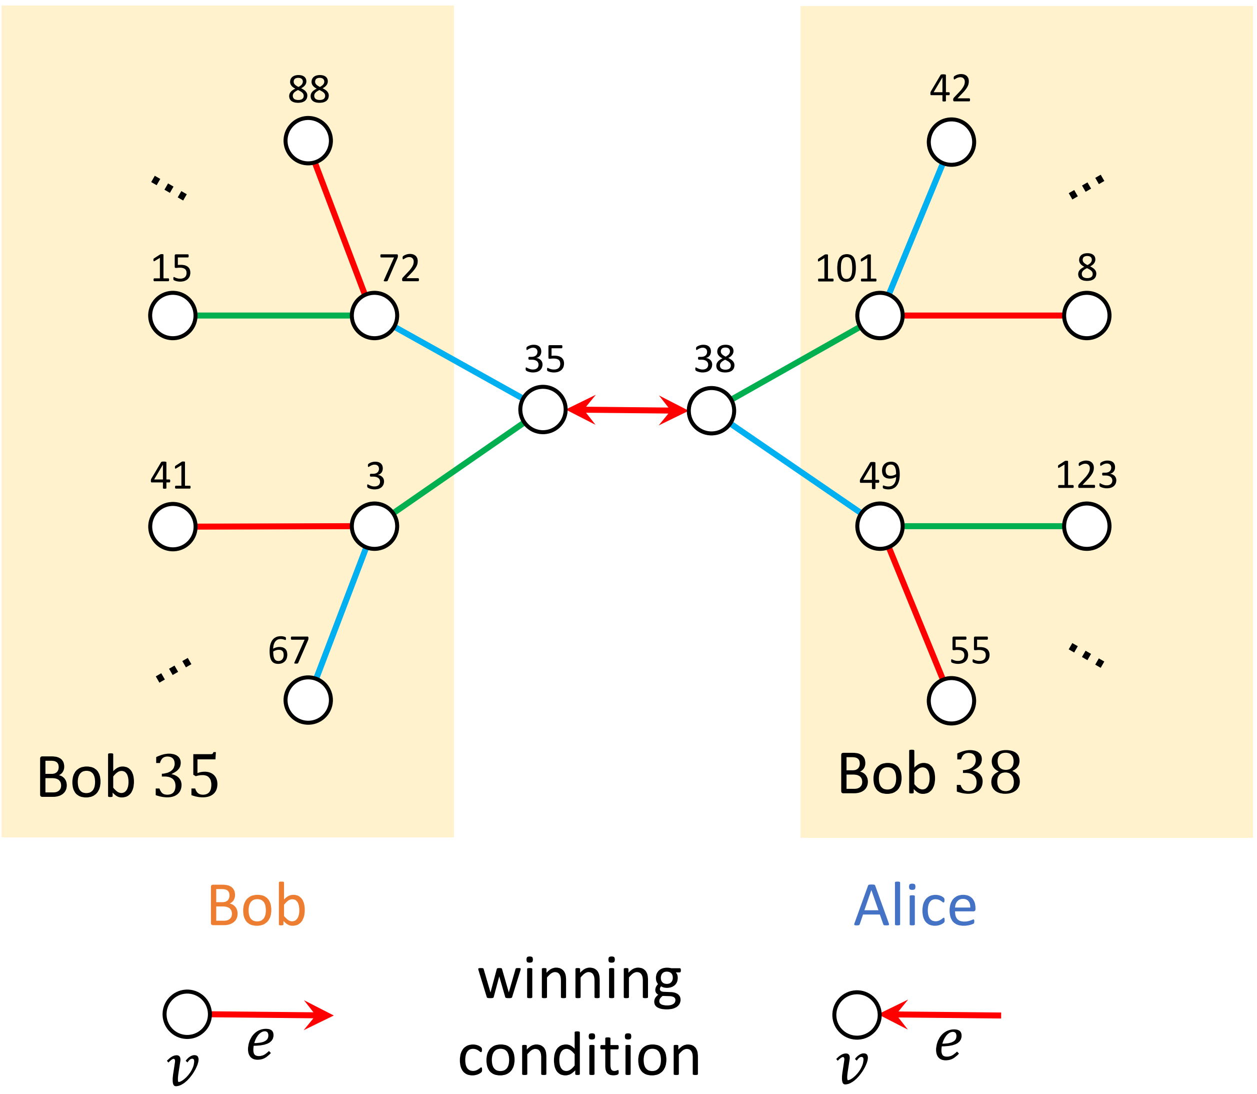

assume that there is a $o(\log n)$-round algorithm $\mathcal{A}$ false!

↓

define a set of two-player games based on $\mathcal{A}$

↓

show: each possible distribution of winning strategies implies that $\mathcal{A}$ is incorrect ✓

Can Bob win both of $(35, \text{red})$ and $(38, \text{red})$? No!

Let's play Bob's strategies against each other!

Lower Bounds

How can we prove complexity lower bounds in the LOCAL model?

Lower Bounds

How can we prove complexity lower bounds in the LOCAL model?

Round Elimination

Marks' Technique

Lower Bounds

How can we prove complexity lower bounds in the LOCAL model?

Round Elimination

Marks' Technique

[Brandt, Fischer, Hirvonen, Keller, Lempiäinen, Rybicki, Suomela, Uitto, STOC'16]

- Δ-coloring
- sinkless orientation

Lower Bounds

How can we prove complexity lower bounds in the LOCAL model?

Round Elimination

[Brandt, Fischer, Hirvonen, Keller, Lempiäinen, Rybicki, Suomela, Uitto, STOC'16]

- Δ-coloring
- sinkless orientation

Marks' Technique

[Marks, JAMS'16]
[Brandt, Chang, Grebík, Grunau, Rozhoň, Vidnyánszky, ITCS'22]

- Δ-coloring
- sinkless orientation

Lower Bounds

How can we prove complexity lower bounds in the LOCAL model?

Round Elimination ⟷ ? ⟷ Marks' Technique

Round Elimination

[Brandt, Fischer, Hirvonen, Keller, Lempiäinen, Rybicki, Suomela, Uitto, STOC'16]

- Δ-coloring
- sinkless orientation

Marks' Technique

[Marks, JAMS'16]
[Brandt, Chang, Grebík, Grunau, Rozhoň, Vidnyánszky, ITCS'22]

- Δ-coloring
- sinkless orientation

Lower Bounds

How can we prove complexity lower bounds in the LOCAL model?

?

Round Elimination

[Brandt, Fischer, Hirvonen, Keller, Lempiäinen, Rybicki, Suomela, Uitto, STOC'16]

- Δ-coloring
- sinkless orientation

? ⟷

works in the Borel context

Marks' Technique

[Marks, JAMS'16]
[Brandt, Chang, Grebík, Grunau, Rozhoň, Vidnyánszky, ITCS'22]

- Δ-coloring
- sinkless orientation

Some Answers

Marks' Technique

- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?
  - ... in the setting with input?

We remove all known obstacles.

- How can we find nontrivial fixed point relaxations?
  - ... together with a suitable input $\mathcal{I}$?

Some Answers

Marks' Technique

- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?
  - ... in the setting with input?

We remove all known obstacles.

- How can we find nontrivial fixed point relaxations?
  - ... together with a suitable input $\mathcal{I}$?

Marks' Technique $\cong$ Round Elimination fixed point relaxations with SO as input

[Balliu, Brandt, Gabsdil, Olivetti, 2026+]

Some Answers

Marks' Technique

- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?

  We remove all known obstacles.

  - ... in the setting with input?

    No!

- How can we find nontrivial fixed point relaxations?
  - ... together with a suitable input $\mathcal{I}$?

Some Answers

Marks' Technique

- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?

  We remove all known obstacles.

  - ... in the setting with input?

    No! SSO with input SO does not.

- How can we <u>find</u> nontrivial fixed point relaxations?
  - ... together with a suitable input $\mathcal{J}$?

Some Answers

- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?
  - ... in the setting with input?
- How can we <u>find</u> nontrivial fixed point relaxations?
  - ... together with a suitable input $\mathcal{I}$?

    (so far: guess some input $\mathcal{I}$ and somehow try to find a nontrivial fixed point relaxation w.r.t. $\mathcal{I}$)

Marks' Technique

We remove all known obstacles.

No! SSO with input SO does not.

Some Answers

- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?
  - ... in the setting with input?

We remove all known obstacles.

Marks' Technique

No! SSO with input SO does not.

- How can we find nontrivial fixed point relaxations?
  - ... together with a suitable input $\mathcal{I}$?

    (so far: guess some input $\mathcal{I}$ and somehow try to find a nontrivial fixed point relaxation w.r.t. $\mathcal{I}$)

For any Π and any fixed point $\mathcal{F}$, there is an easiest problem $\mathcal{I}$ such that $\mathcal{F}$ is a relaxation of Π in the setting with input $\mathcal{I}$.

Some Answers

Marks' Technique

- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?

  We remove all known obstacles.

  - ... in the setting with input?

    No! SSO with input SO does not.

- How can we <u>find</u> nontrivial fixed point relaxations?

  - ... together with a suitable input $\mathcal{J}$?

    (so far: guess some input $\mathcal{J}$ and somehow try to find a nontrivial fixed point relaxation w.r.t. $\mathcal{J}$)

For any Π and any fixed point $\mathcal{F}$, there is an easiest problem $\tau_{\mathcal{F}}(\Pi)$ such that $\mathcal{F}$ is a relaxation of Π in the setting with input $\tau_{\mathcal{F}}(\Pi)$.

Some Answers

- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?

  We remove all known obstacles.

  Marks' Technique

  - ... in the setting with input?

    No! SSO with input SO does not.

- How can we find nontrivial fixed point relaxations?

  - ... together with a suitable input $\mathcal{I}$?

    (so far: guess some input $\mathcal{I}$ and somehow try to find a nontrivial fixed point relaxation w.r.t. $\mathcal{I}$)

For any Π and any fixed point $\mathcal{F}$, there is an easiest problem $\tau_{\mathcal{F}}(\Pi)$ such that $\mathcal{F}$ is a relaxation of Π in the setting with input $\tau_{\mathcal{F}}(\Pi)$.

new: take a known fixed point $\mathcal{F}$ and check whether $\mathcal{F}$ is nontrivial w. input $\tau_{\mathcal{F}}(\Pi)$.

Some Answers

Marks' Technique

- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?

  We remove all known obstacles.

  - ... in the setting with input?

    No! SSO with input SO does not.

- How can we find nontrivial fixed point relaxations?

  - ... together with a suitable input $\mathcal{I}$?

    (so far: guess some input $\mathcal{I}$ and somehow try to find a nontrivial fixed point relaxation w.r.t. $\mathcal{I}$)

$\tau_{\mathcal{F}}(\tau_{\mathcal{F}}(\Pi))$ is a nontrivial fixed point relaxation of Π.

For any Π and any fixed point $\mathcal{F}$, there is an easiest problem $\tau_{\mathcal{F}}(\Pi)$ such that $\mathcal{F}$ is a relaxation of Π in the setting with input $\tau_{\mathcal{F}}(\Pi)$.

if yes ...

new: take a known fixed point $\mathcal{F}$ and check whether $\mathcal{F}$ is nontrivial w. input $\tau_{\mathcal{F}}(\Pi)$.

Open Questions and Research Directions

- Is it decidable on which side of the $O(\log^* n) - \Omega(\log n)$ gap a problem lies?

A Hierarchy of Models

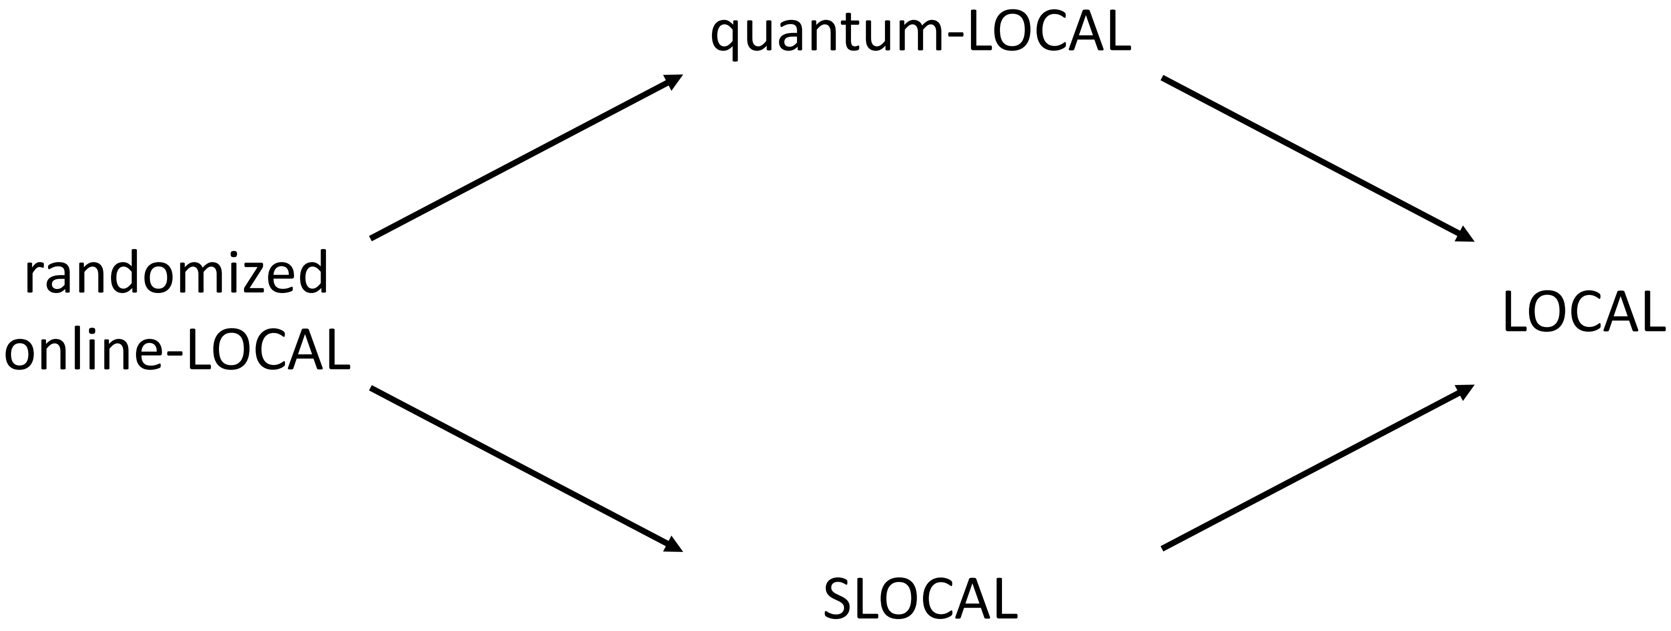

[Akbari, Coiteux-Roy, d'Amore, Le Gall, Lievonen, Melnyk, Modanese, Pai, Renou, Rozhoň, Suomela, STOC'25]

A Hierarchy of Models

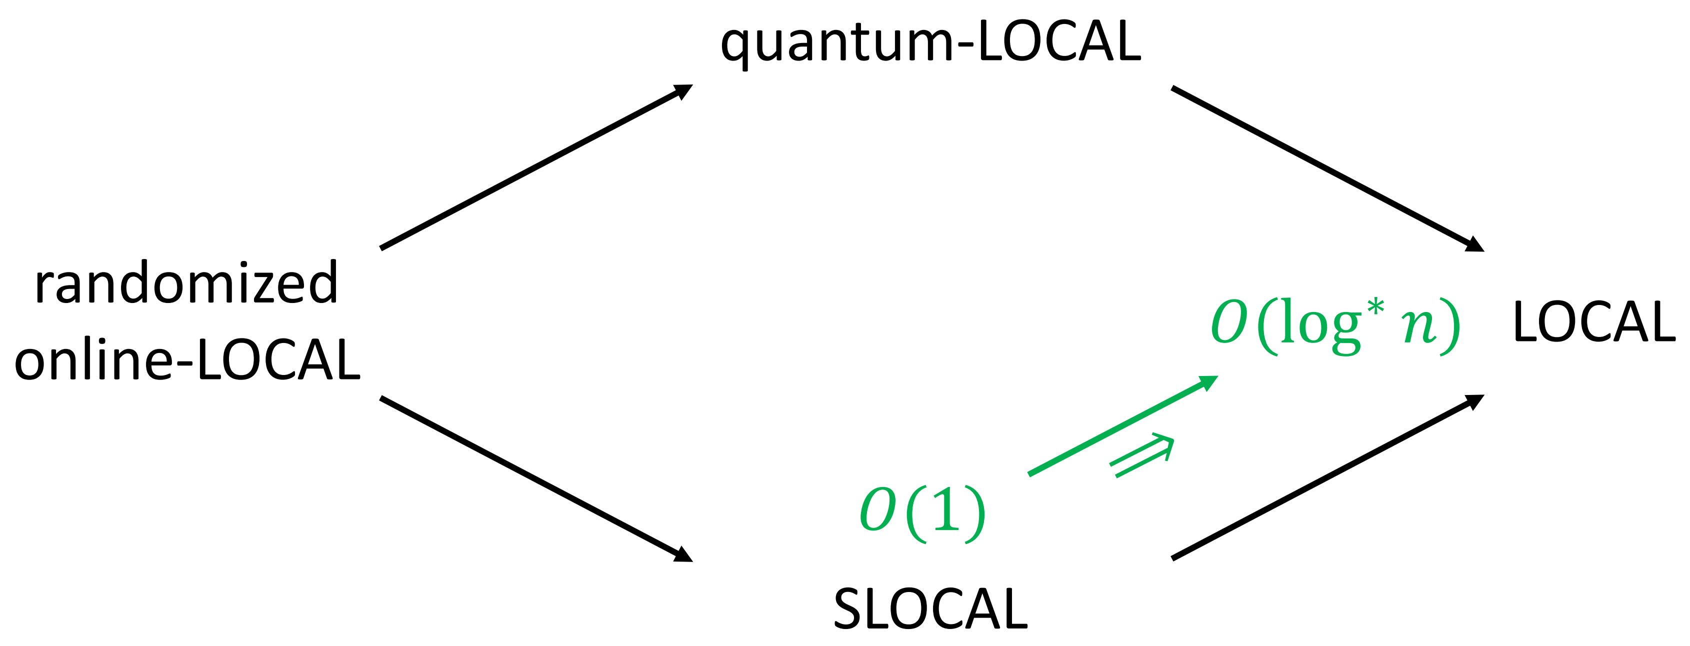

A Hierarchy of Models

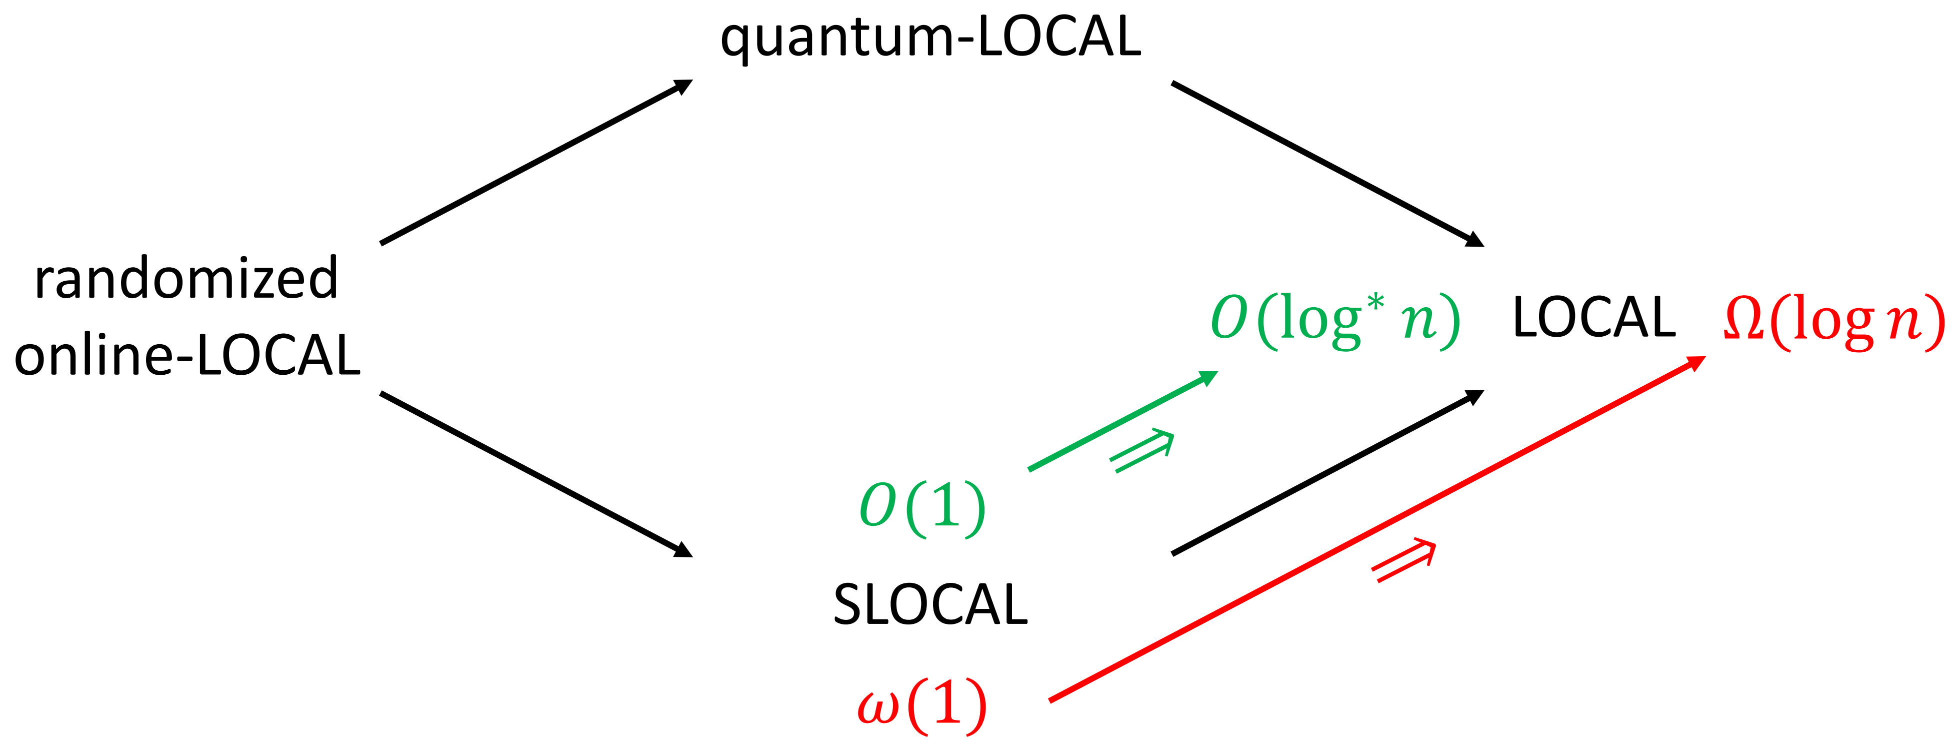

A Hierarchy of Models

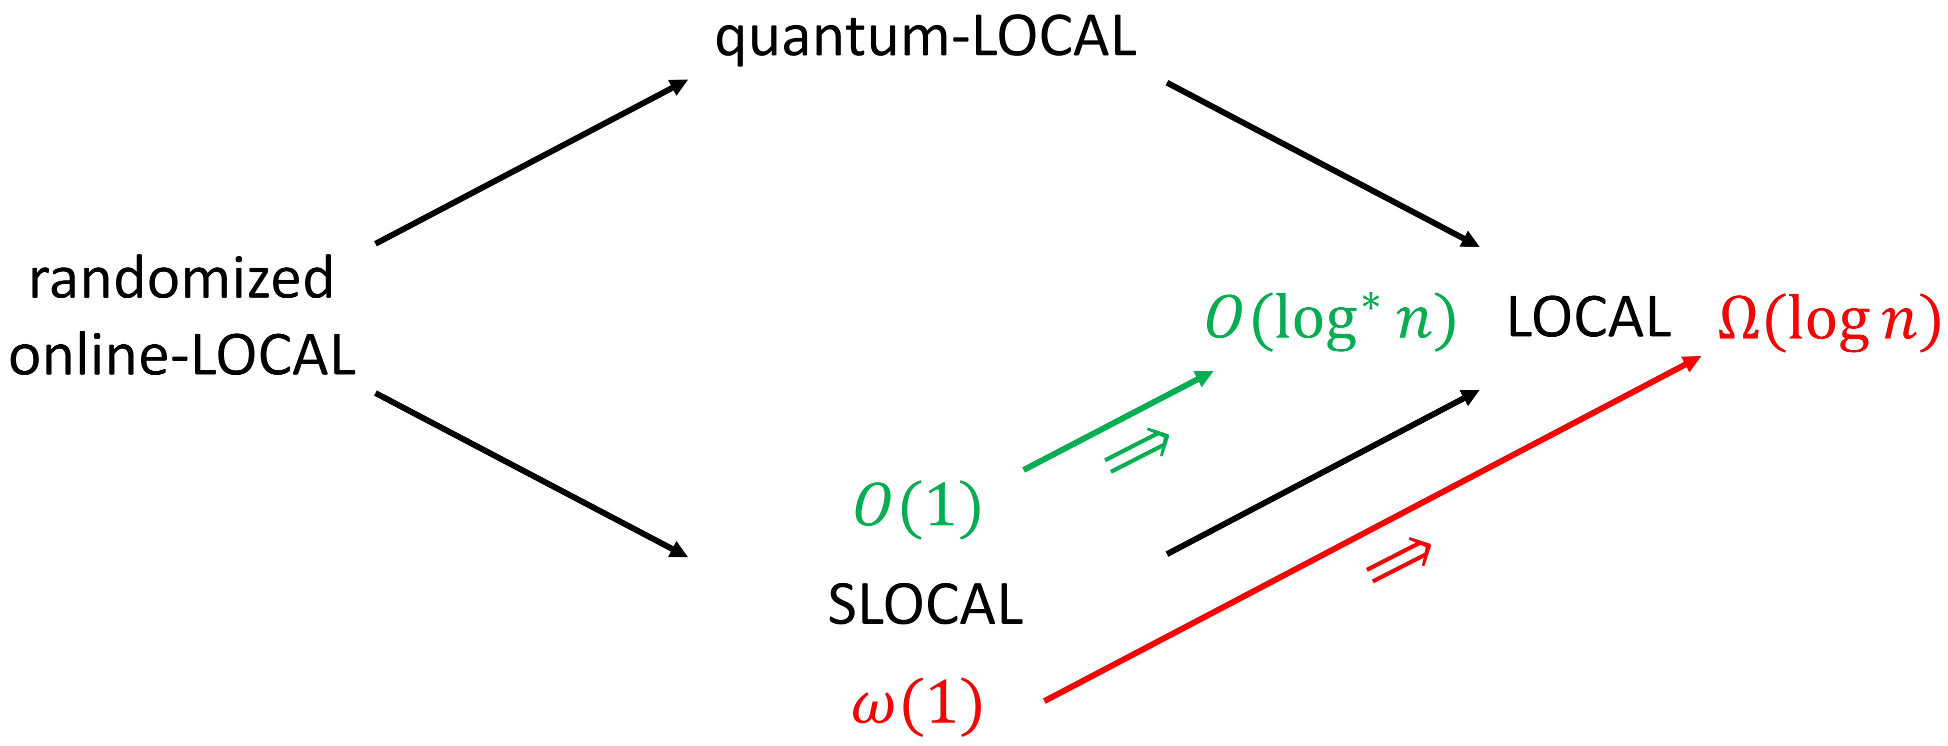

no model-specific lower bound technique known

A Hierarchy of Models

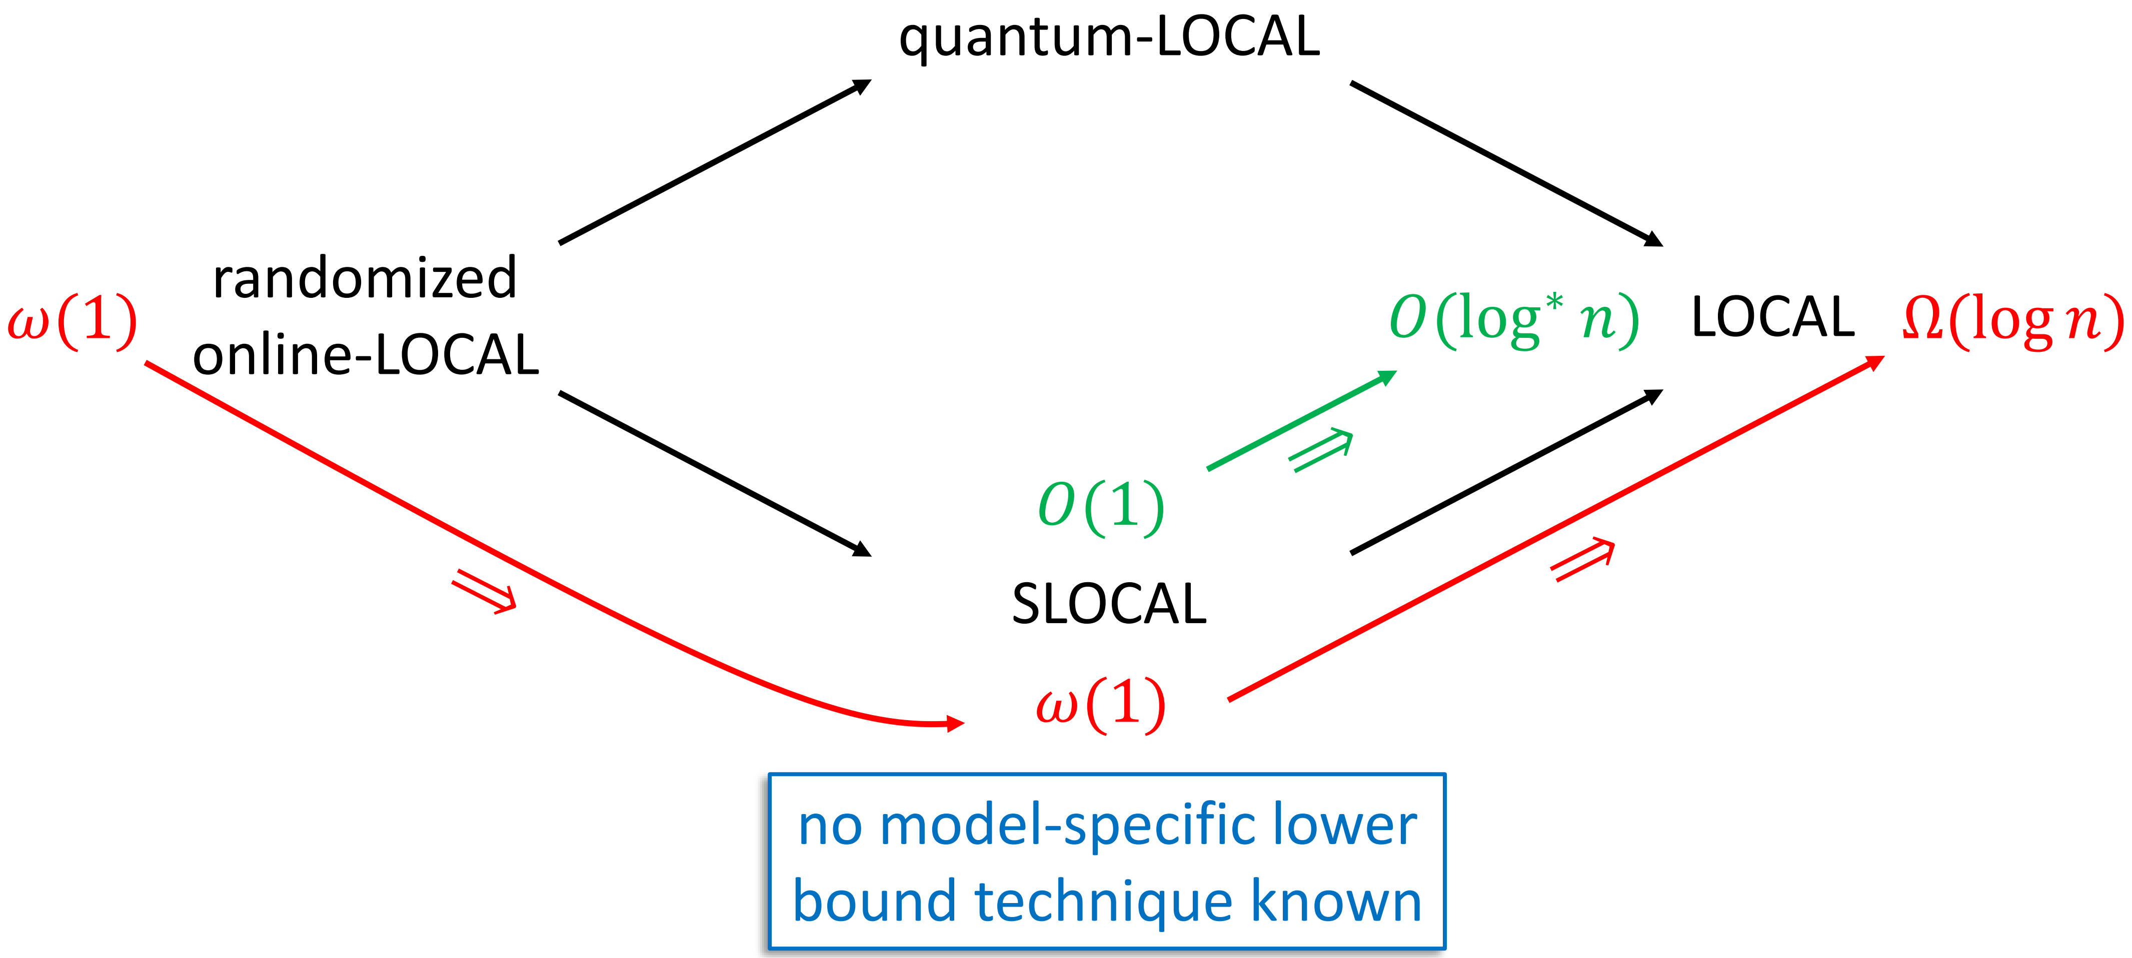

A Hierarchy of Models

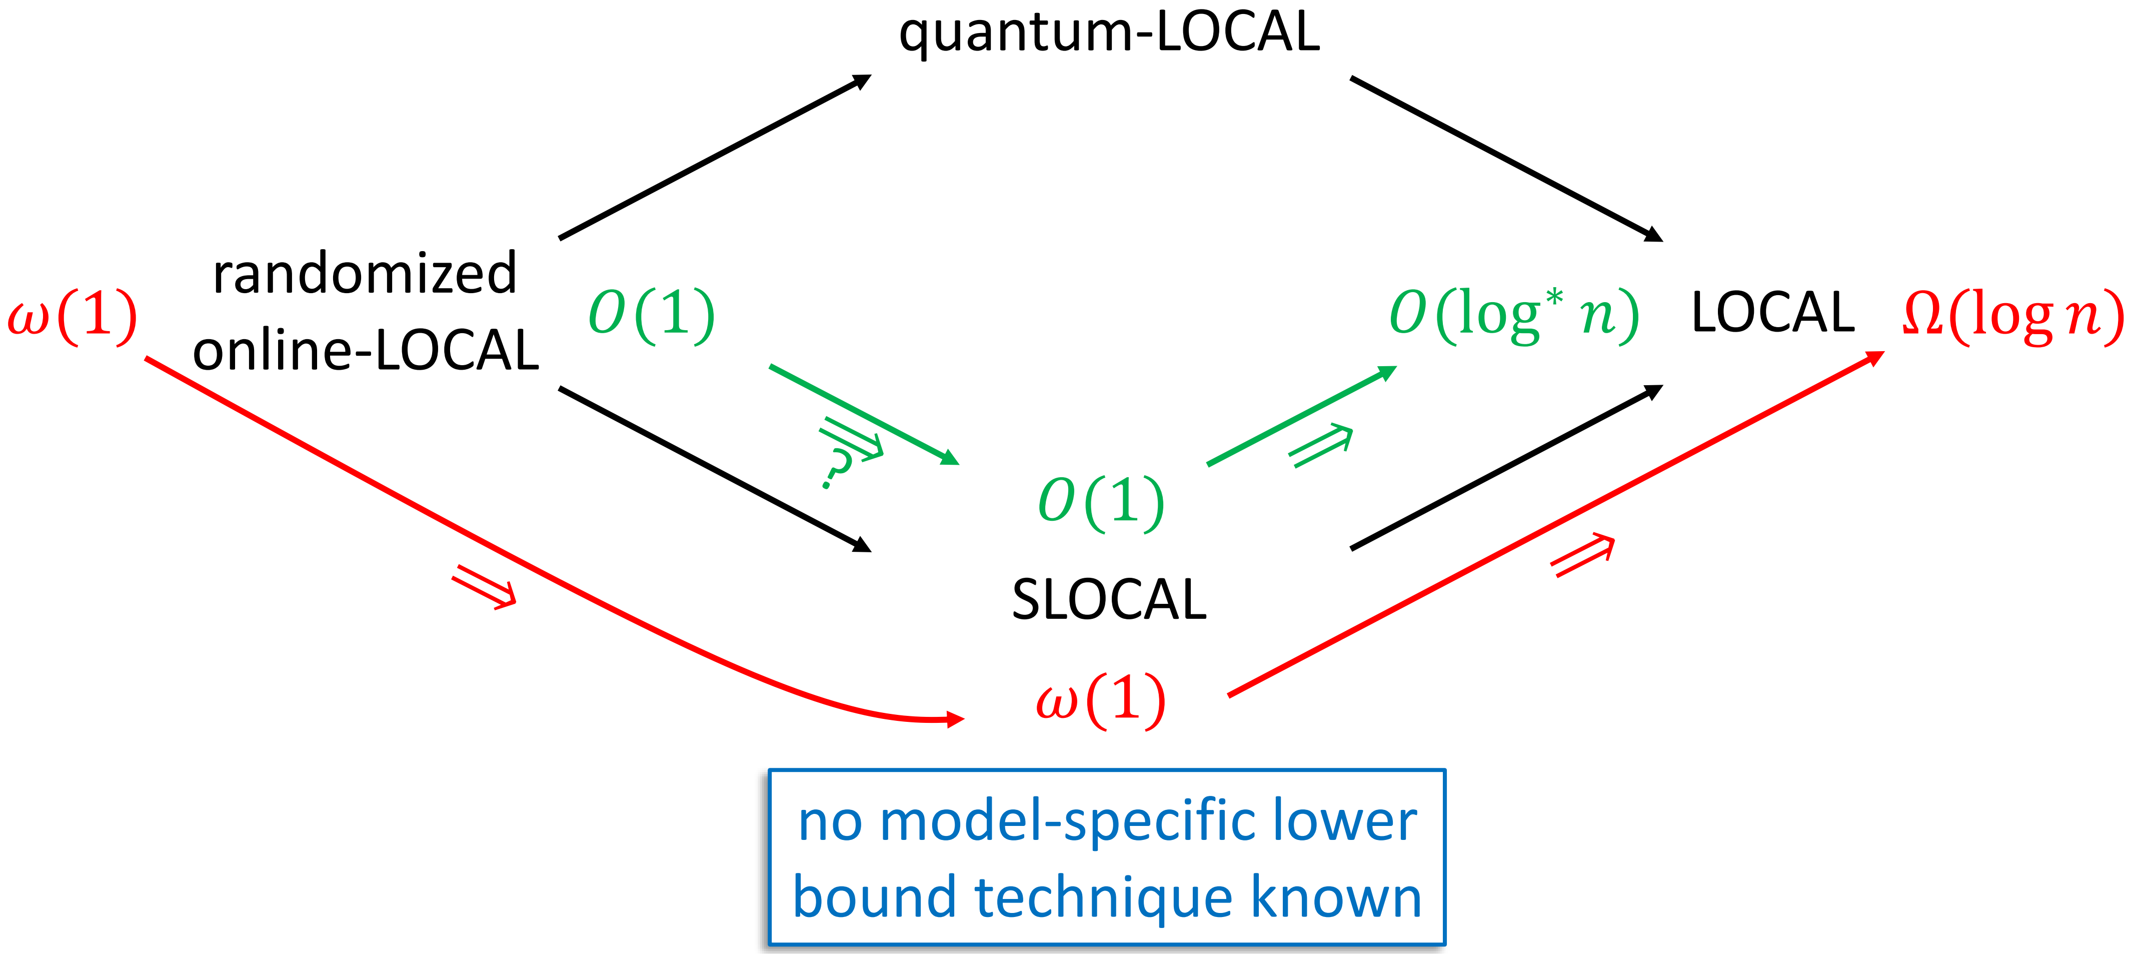

A Hierarchy of Models

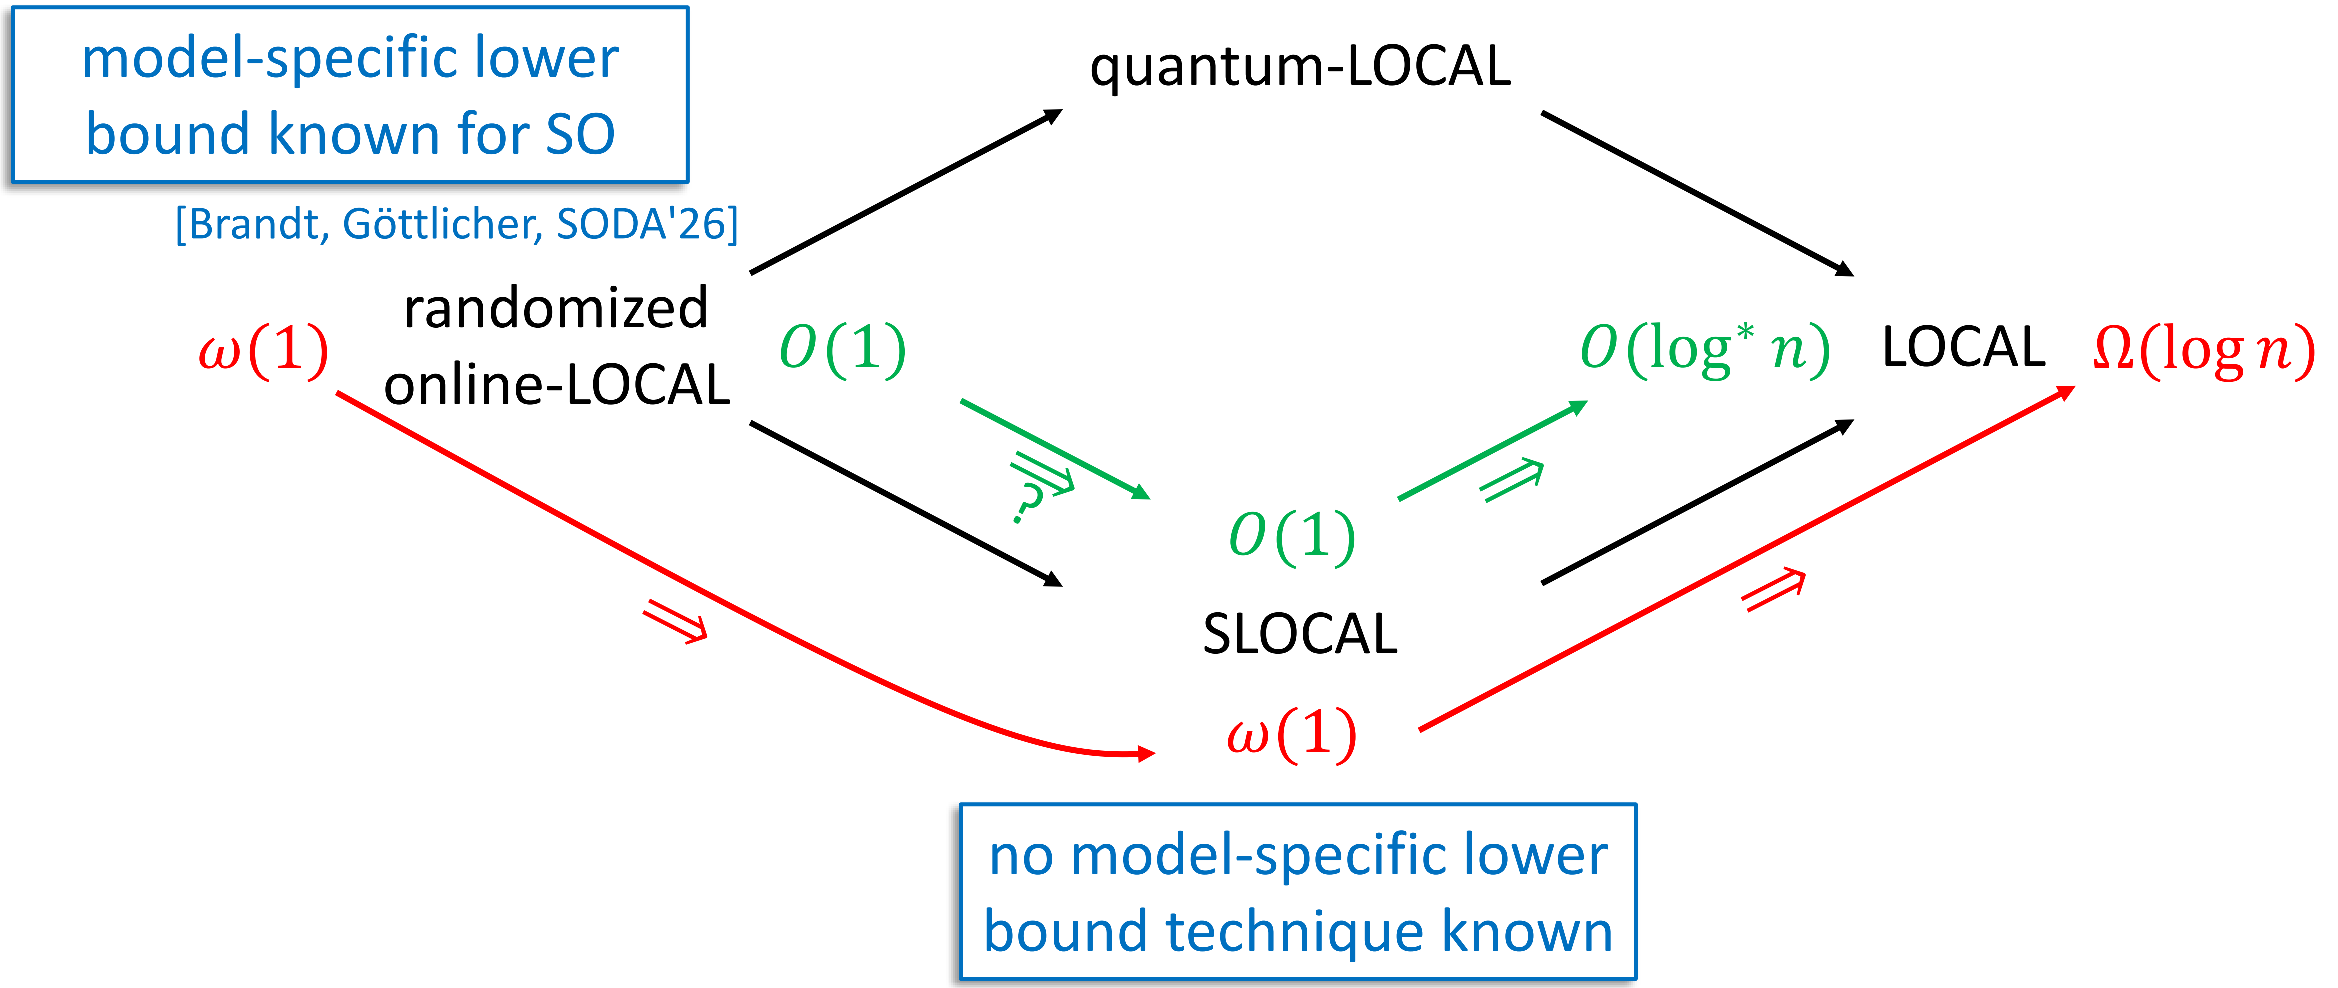

Open Questions and Research Directions

- Is it decidable on which side of the $O(\log^* n) - \Omega(\log n)$ gap a problem lies?

Open Questions and Research Directions

- Is it decidable on which side of the $O(\log^* n) - \Omega(\log n)$ gap a problem lies?

- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?

Open Questions and Research Directions

- Is it decidable on which side of the $O(\log^* n) - \Omega(\log n)$ gap a problem lies?

- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?

- Which relations exist between different (established) lower bound techniques?
  In which other models can (distributed) lower bound techniques be applied?

Open Questions and Research Directions

- Is it decidable on which side of the $O(\log^* n) - \Omega(\log n)$ gap a problem lies?
- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?
- Which relations exist between different (established) lower bound techniques?
  In which other models can (distributed) lower bound techniques be applied?
- Can we find a generalization/extension of Marks' technique corresponding to round elimination fixed points with some input that is different from SO?

Open Questions and Research Directions

- Is it decidable on which side of the $O(\log^* n) - \Omega(\log n)$ gap a problem lies?

- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?

- Which relations exist between different (established) lower bound techniques?
  In which other models can (distributed) lower bound techniques be applied?

- Can we find a generalization/extension of Marks' technique corresponding to round elimination fixed points with some input that is different from SO?

- Explore the setting with inputs (and reductions in general).

Open Questions and Research Directions

- Is it decidable on which side of the $O(\log^* n) - \Omega(\log n)$ gap a problem lies?
- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?
- Which relations exist between different (established) lower bound techniques?
  In which other models can (distributed) lower bound techniques be applied?
- Can we find a generalization/extension of Marks' technique corresponding to round elimination fixed points with some input that is different from SO?
- Explore the setting with inputs (and reductions in general).
- Is it decidable whether a problem can be solved in $O(1)$ rounds?

Open Questions and Research Directions

- Is it decidable on which side of the $O(\log^* n) - \Omega(\log n)$ gap a problem lies?
- Does every problem of complexity $\Omega(\log n)$ have a nontrivial fixed point relaxation?
- Which relations exist between different (established) lower bound techniques?
  In which other models can (distributed) lower bound techniques be applied?
- Can we find a generalization/extension of Marks' technique corresponding to round elimination fixed points with some input that is different from SO?
- Explore the setting with inputs (and reductions in general).
- Is it decidable whether a problem can be solved in $O(1)$ rounds?

Thanks!