Deterministic subgraph detection in broadcast CONGEST

Janne H. Korhonen · Aalto University
Joel Rybicki · University of Helsinki

Aalto University
UNIVERSITY OF HELSINKI

1.

Introduction

Introduction:

CONGEST model

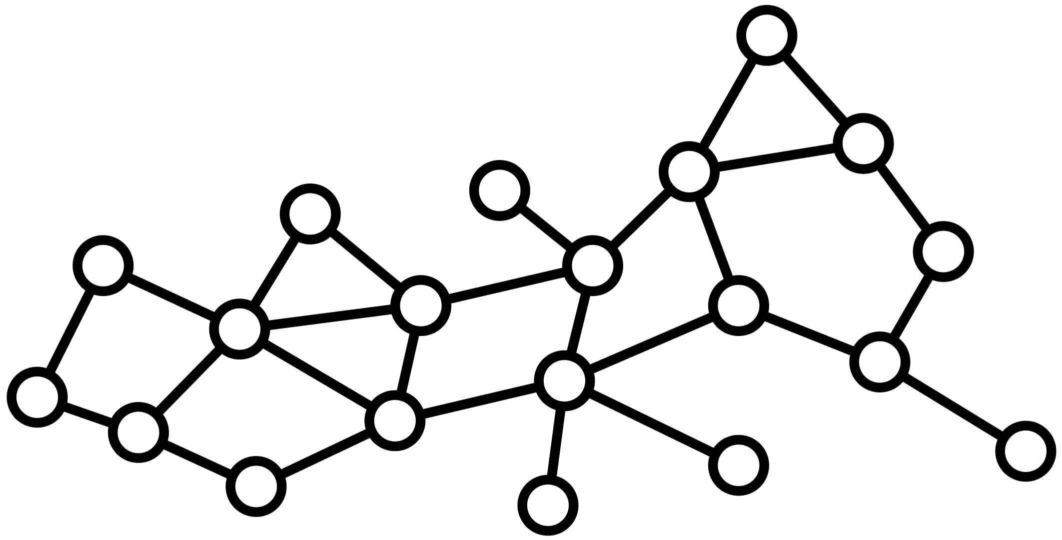

- **CONGEST model**
  - ***n*** nodes, connected by communication links
  - unique identifiers, synchronous communication
  - unlimited local computation
  - message size **$O(\log n)$** bits/round
  - time measure: ***number of rounds***

Introduction:

CONGEST model

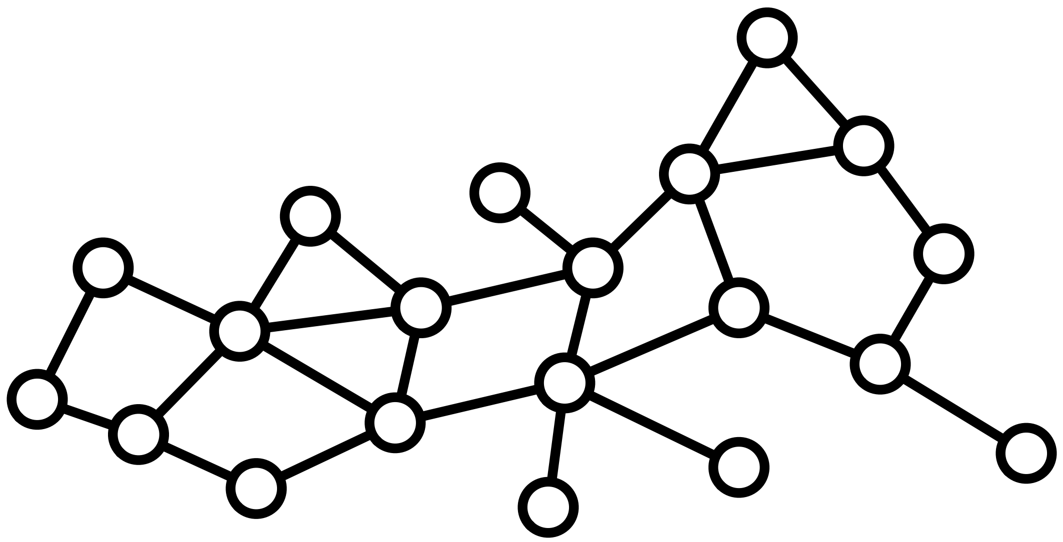

- **CONGEST model**
  - *n* nodes, connected by communication links
  - unique identifiers, synchronous communication
  - unlimited local computation
  - message size $O(\log n)$ bits/round
  - time measure: ***number of rounds***

- **Upper bounds:** broadcast CONGEST
- **Lower bounds:** unicast CONGEST

Introduction:

Subgraph detection

- **H-subgraph detection problem**
  - given a fixed *pattern graph* ***H*** on ***k*** nodes
  - does the network ***G*** contain ***H*** as a subgraph?

- *triangle detection, cycle detection, clique detection, …*

H

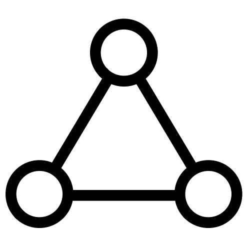

G

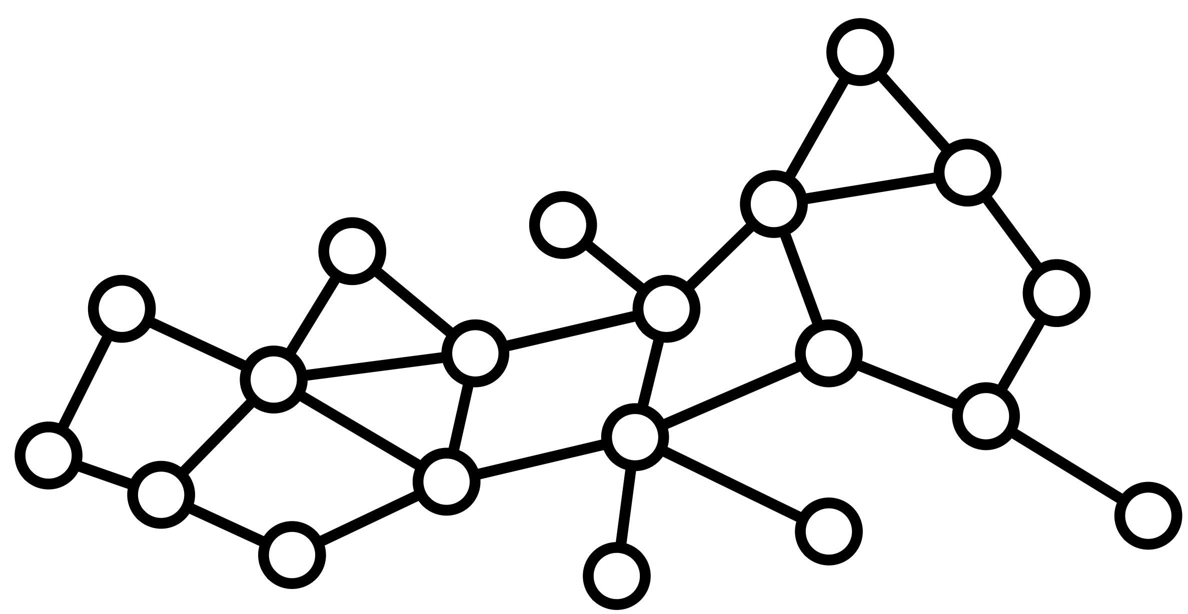

Introduction:

Subgraph detection

- **Detection:**
  - if node belongs to a copy of ***H***, output *one* copy of ***H***

- **Listing/enumeration:**
  - *all* copies of ***H*** are a part of *some* node's output

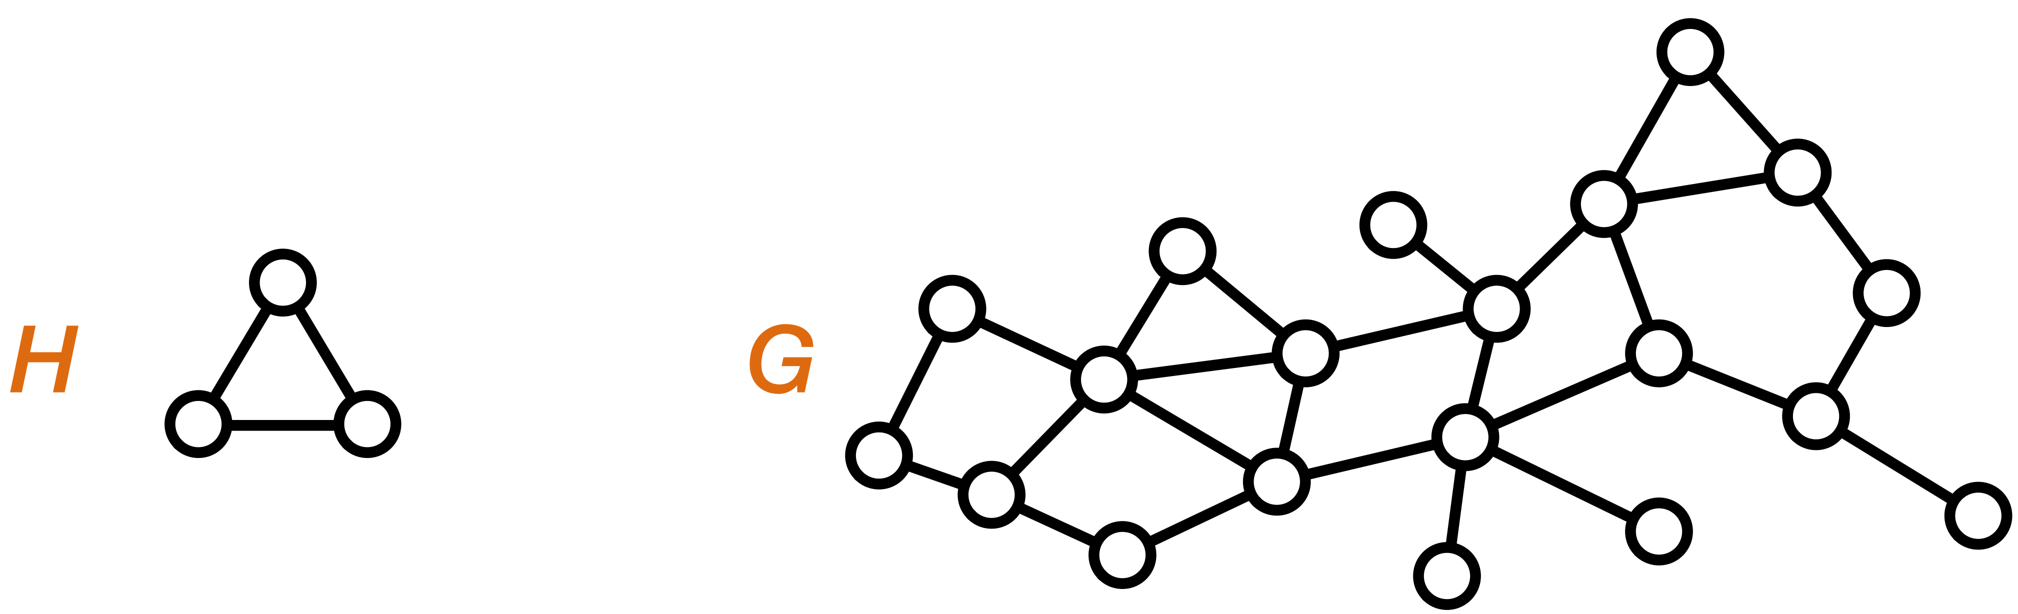

Introduction:

Subgraph detection

- ***H* has constant size *k***
  - In LOCAL: ***O*(1)** for any ***H*** trivially
  - In CONGEST: trivial upper bound $O(n^2)$

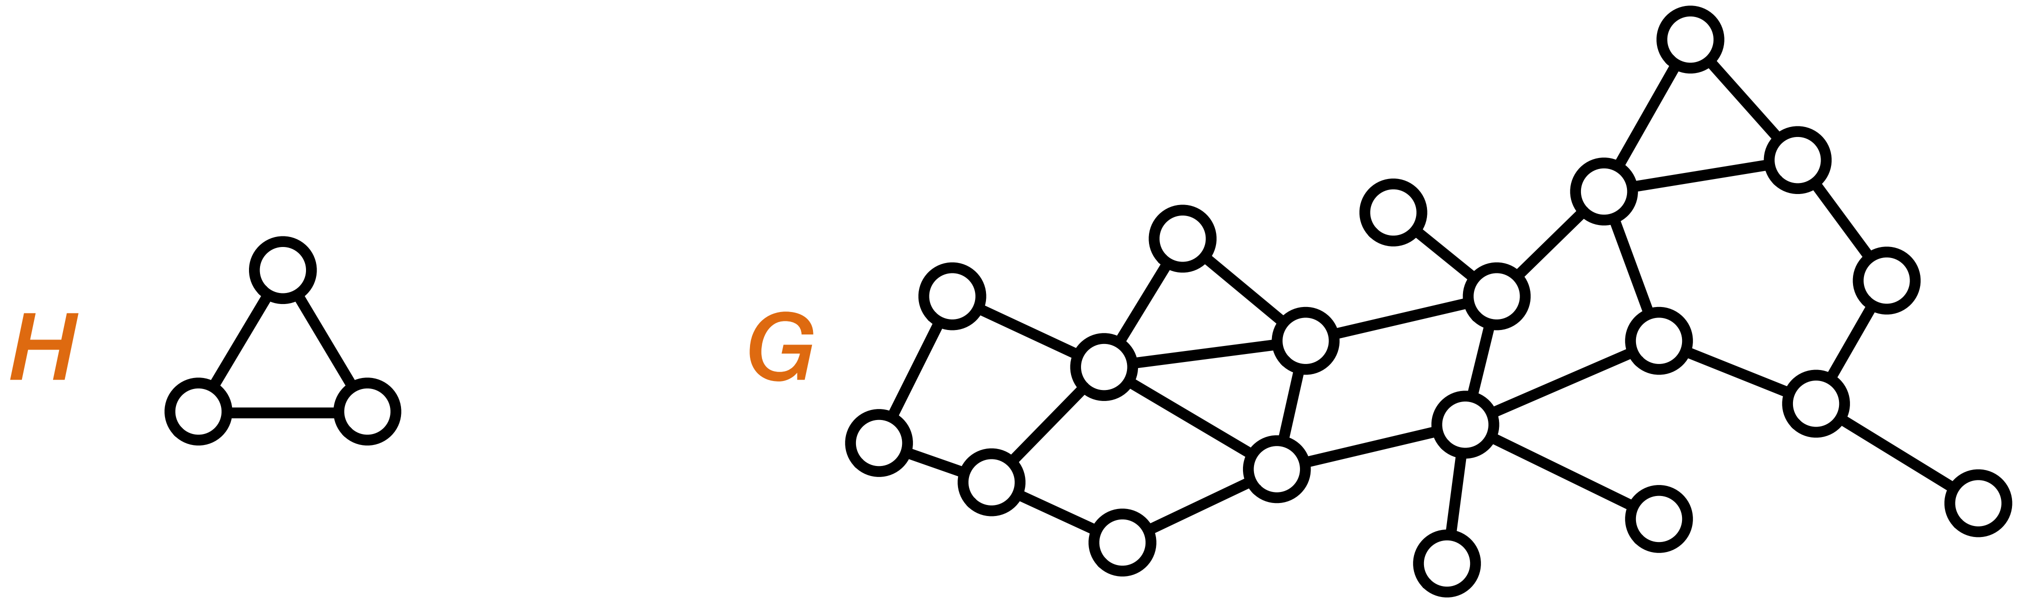

Introduction:

Prior work

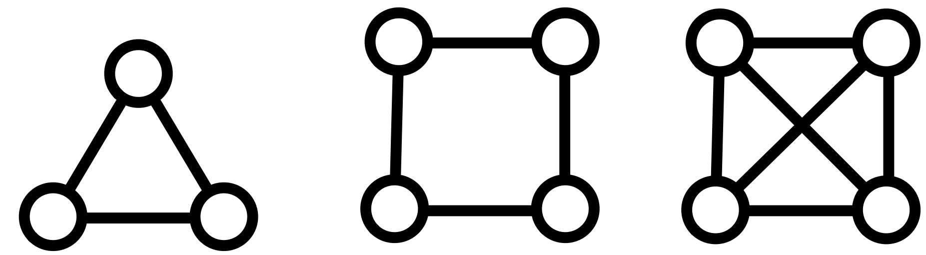

- **Upper bounds**
  - triangle finding in $\tilde{O}(n^{2/3})$ rounds [Izumi & Le Gall, PODC 2017]
  - triangle enumeration in $\tilde{O}(n^{3/4})$ rounds [Izumi & Le Gall, PODC 2017]
  - 4-cycle finding in $O(n^{1/2})$ rounds [Drucker, Kuhn, Ostmann, PODC 2014]
  - clique enumeration in $O(n)$ rounds (trivial)

- **Lower bounds**
  - k-cycles (k even) $\tilde{\Omega}(n^{2/k})$ rounds [Drucker, Kuhn, Ostmann, PODC 2014]
  - k-cycles (k odd, $k \geq 5$) $\tilde{\Omega}(n)$ rounds [Drucker, Kuhn, Ostmann, PODC 2014]
  - triangle enumeration $\tilde{\Omega}(n^{1/3})$ rounds [Izumi & Le Gall, PODC 2017]

Introduction:

Prior work, DISC 2017

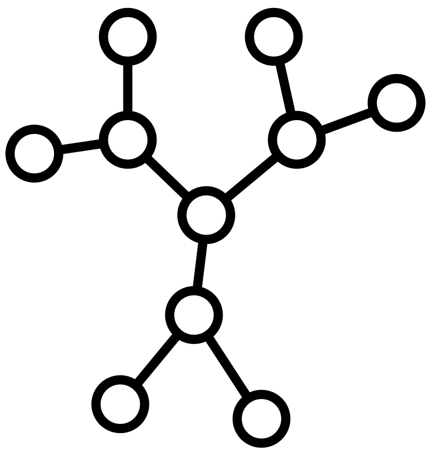

- **Guy Even, Reut Levi, and Moti Medina.**
  Faster and simpler distributed algorithms for testing and correcting graph properties in the CONGEST-model, 2017. arXiv:1705.04898 [cs.DC].
- **Orr Fischer, Tzlil Gonen, and Rotem Oshman.**
  Distributed property testing for subgraph-freeness revisited, 2017. arXiv:1705.04033 [cs.DS].
- **Pierre Fraigniaud, Pedro Montealegre, Dennis Olivetti, Ivan Rapaport, and Ioan Todinca.**
  Distributed subgraph detection, 2017. arXiv:1706.03996 [cs.DC].

Appearing together as *Three notes on distributed property testing*, DISC 2017.

- tree detection in $O(1)$ rounds

2.

Our Results:

Overview

Results 1:

Finding Trees and Cycles

- **Upper bounds**
  - k-trees in $O(1)$ rounds*
  - k-cycles in $O(n)$ rounds
  - k-pseudotrees (tree + 1 edge) in $O(n)$ rounds
- **Lower bounds**
  - k-cycles (k even) require $\Omega(n^{1/2}/\log n)$ rounds

Results 1:

Finding Trees and Cycles

- **Upper bounds**
  - k-trees in $O(k2^k)$ rounds*
  - k-cycles in $O(k2^k n)$ rounds
  - k-pseudotrees (tree + 1 edge) in $O(k2^k n)$ rounds
- **Lower bounds**
  - k-cycles (k even) require $\Omega(n^{1/2}/\log n)$ rounds

Results 1:

Finding Trees and Cycles

- **Some tight results…**
  - trees in $O(1)$ rounds
  - odd cycles are $\tilde{\Theta}(n)$
- **…and some not tight**
  - *gap* for even cycles between $O(n)$ and $\tilde{\Omega}(n^{1/2})$

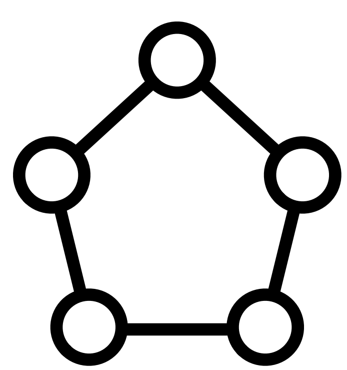

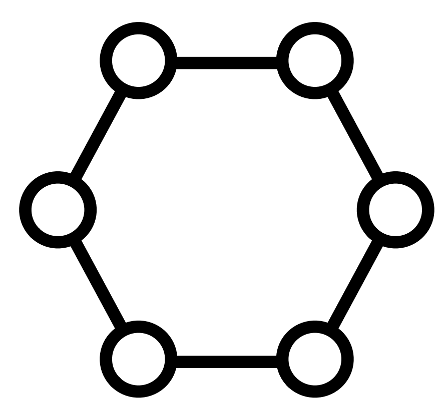

Results 2:

Enumeration in sparse graphs

- **does it help if the input graph *G* is sparse?**
- **notion of sparseness: bounded degeneracy**
  - input graph *G* with degeneracy *d*
  - degeneracy ≈ arboricity

Results 2:

Enumeration in sparse graphs

- **Upper bounds**
  - k-cliques and 4-cycles in $O(d + \log n)$ rounds
  - 5-cycles in $O(d^2 + \log n)$ rounds

- **Lower bounds**
  - finding 4-cycles and 5-cycles requires $\tilde{\Omega}(d)$ rounds
  - bounded degeneracy does not help with 6-cycles
    - need $\tilde{\Omega}(n^{1/2})$ rounds on graphs with degeneracy 2

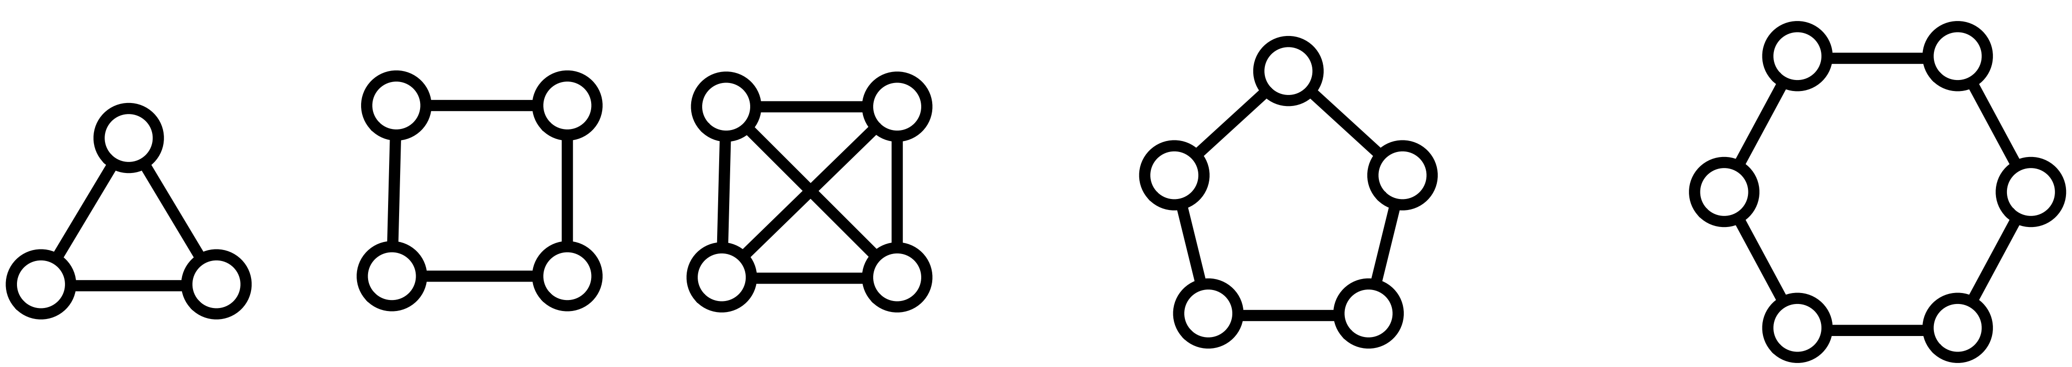

3.

Our Results:

Finding Trees and Cycles

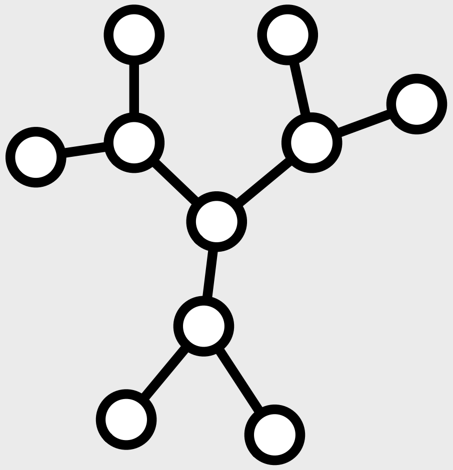

$O(1)$

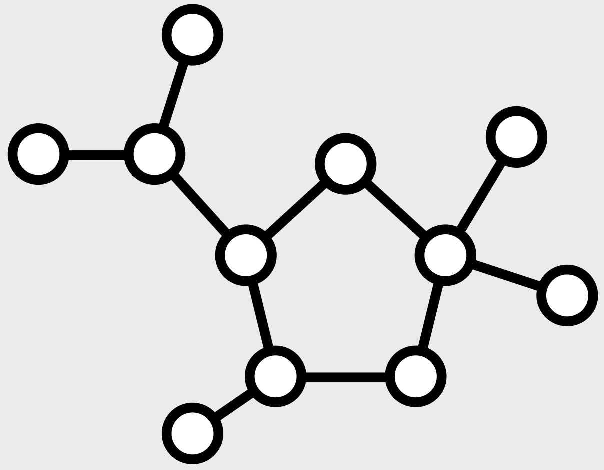

$O(n)$

Technical tool:

Representative families

- **Well-known algorithmic technique**
  - used in centralised *fixed-parameter* algorithms for subgraph detection
  - running times of type **$2^{O(k)}$ poly(n)**
  - compare with other FPT techniques: *colour-coding*, *polynomial sieving*,…

- **Pierre Fraigniaud, Pedro Montealegre, Dennis Olivetti, Ivan Rapaport, and Ioan Todinca.**
  Distributed subgraph detection, 2017. arXiv:1706.03996 [cs.DC].

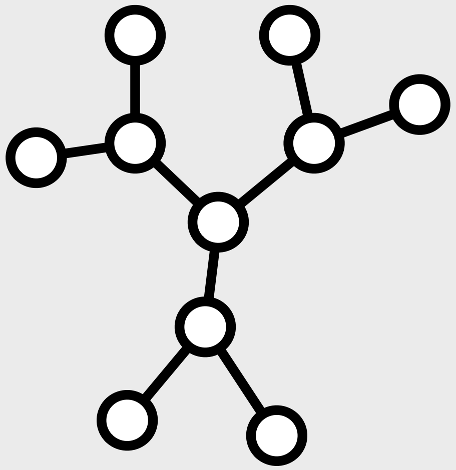

$O(k2^k)$

explicit construction of all partial subtrees

+

“filtering” with representative families

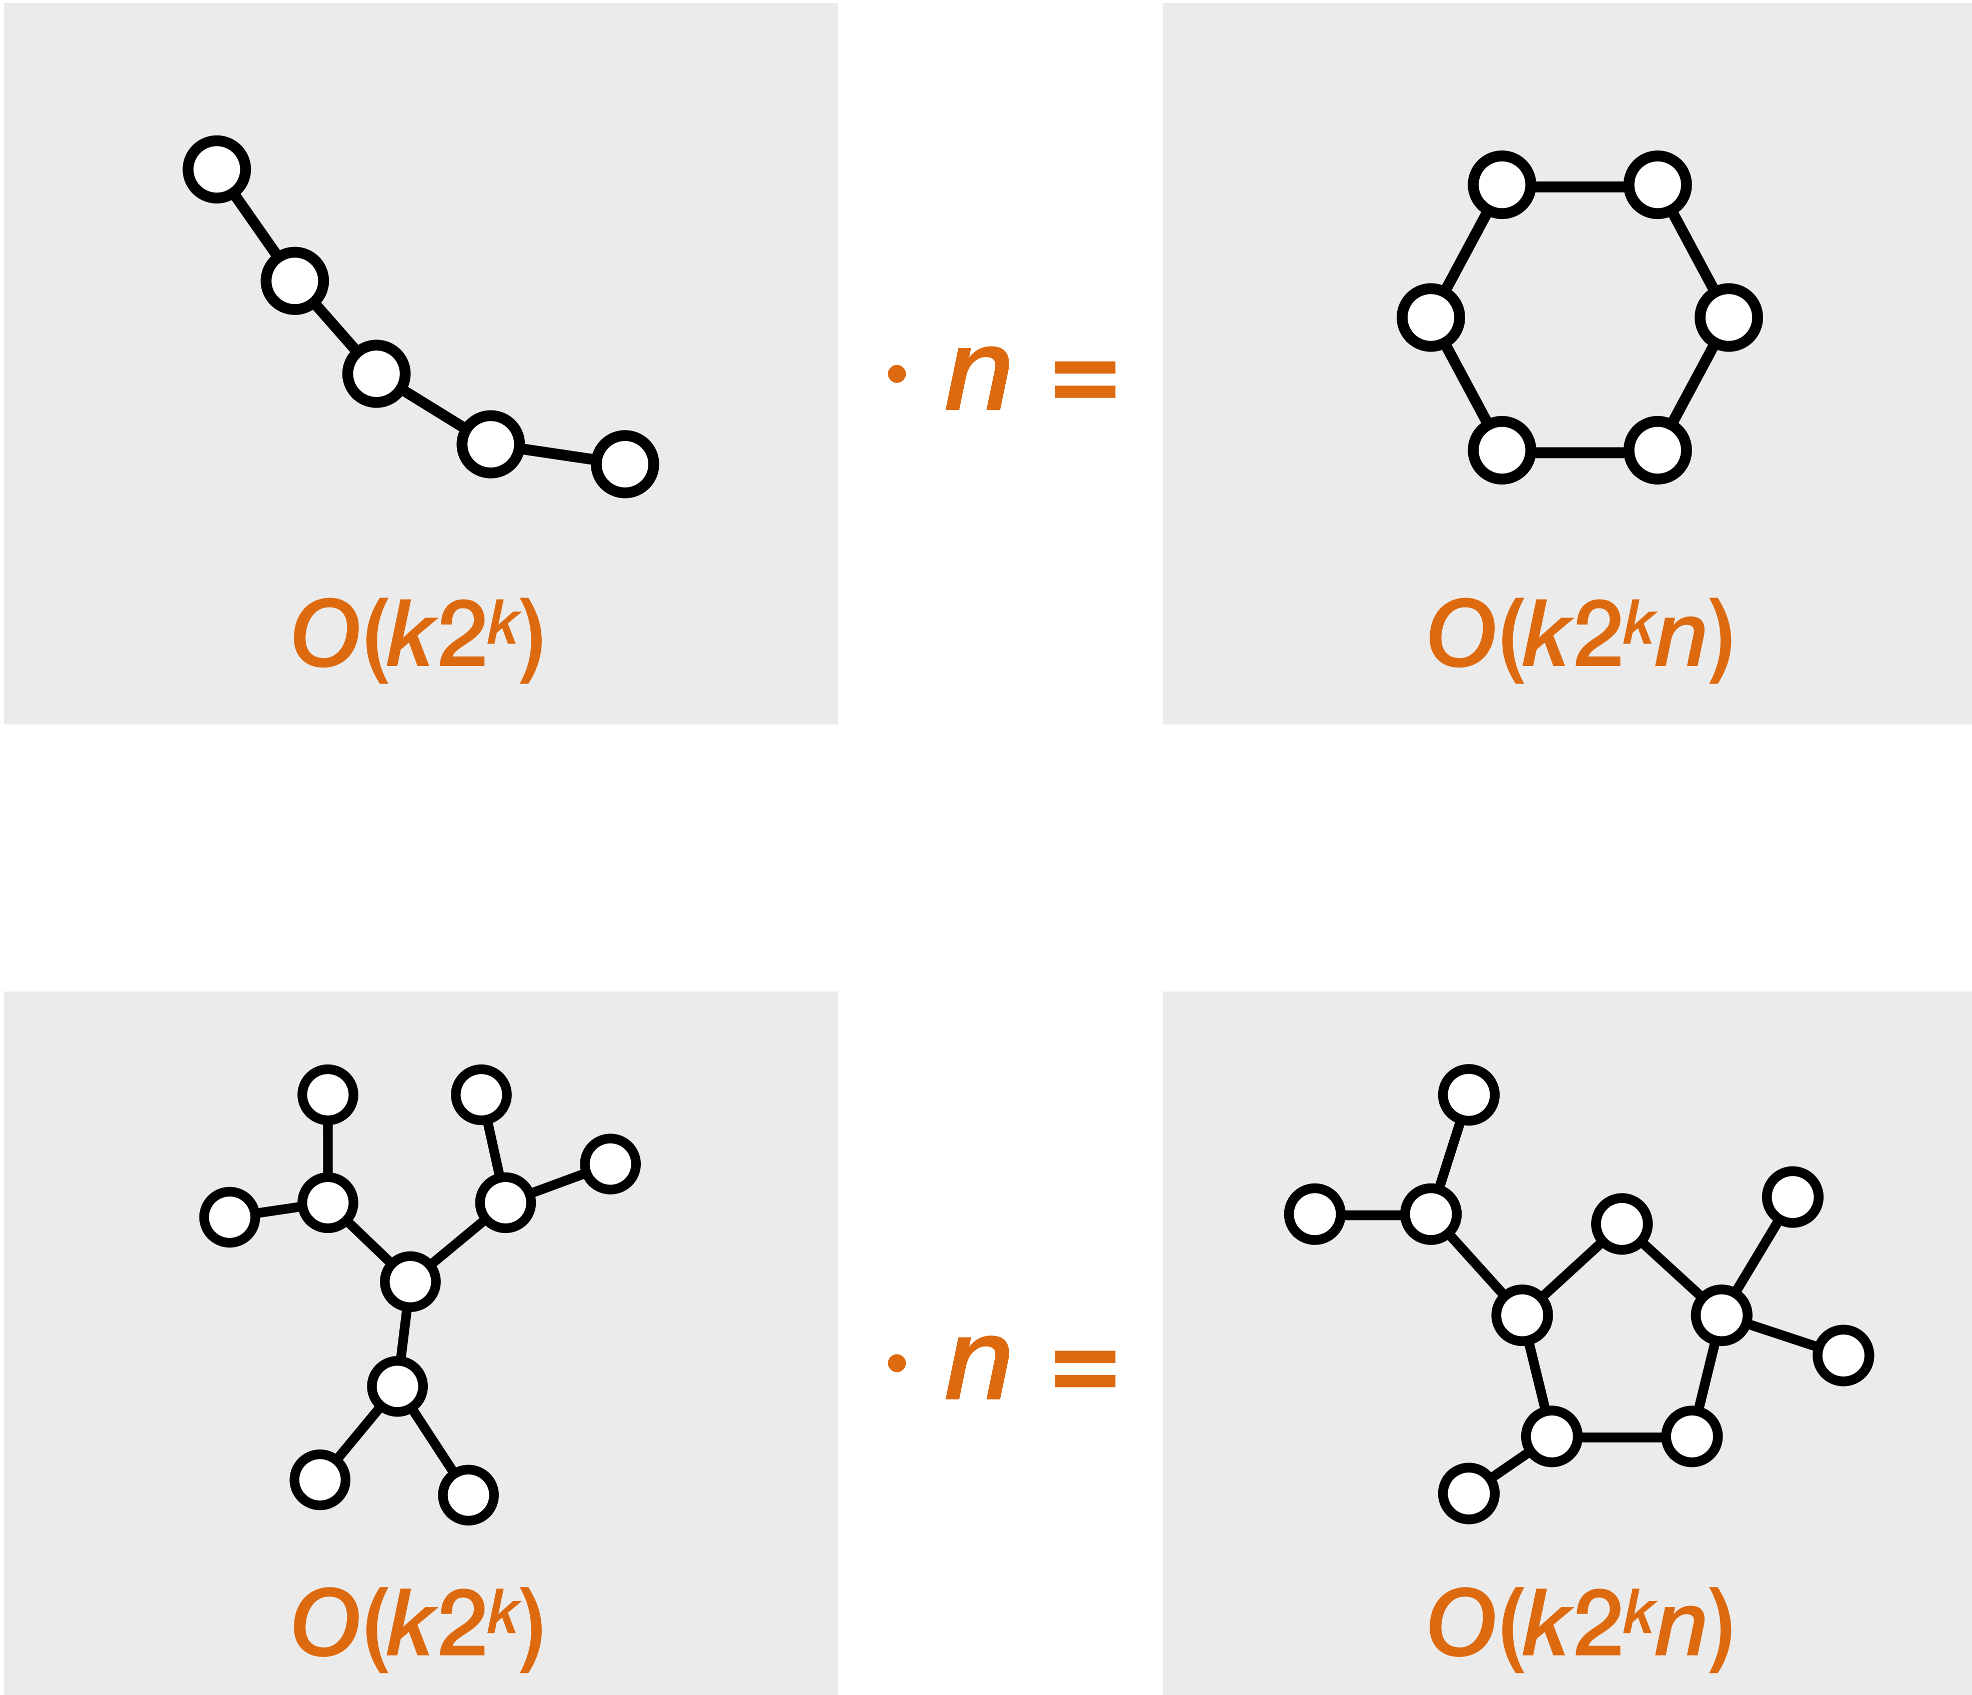
$O(k2^k)$
$\cdot\, n =$
$O(k2^k n)$
$O(k2^k)$
$\cdot\, n =$
$O(k2^k n)$

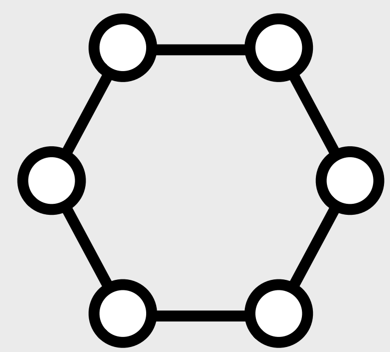

$\Omega(n^{1/2}/\log n)$

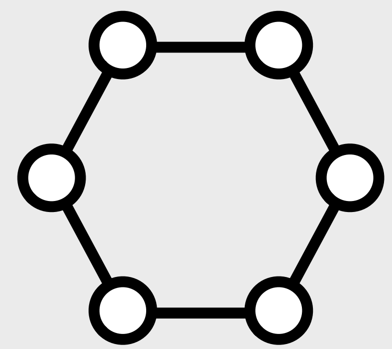

$\Omega(n^{1/2}/\log n)$

very standard communication complexity reduction

4.

Our Results:

Enumeration in sparse graphs

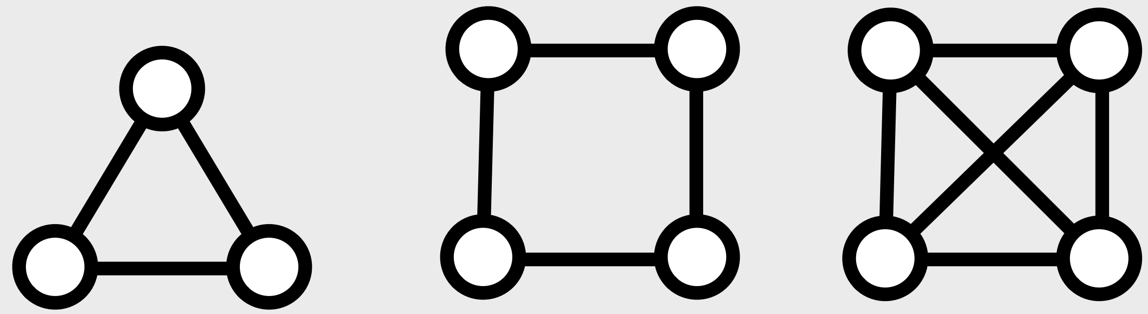

$O(d + \log n)$

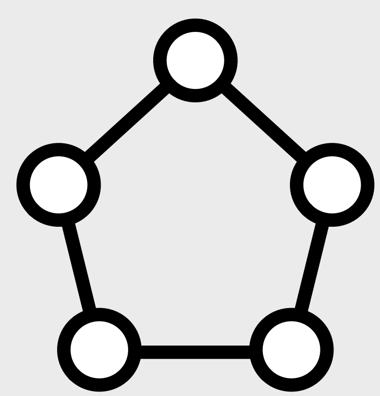

$O(d^2 + \log n)$

Preliminaries:

Degeneracy

- **The following are equivalent:**
  - graph ***G*** has degeneracy ***d***
  - graph ***G*** has acyclic orientation with out-degree ***d***

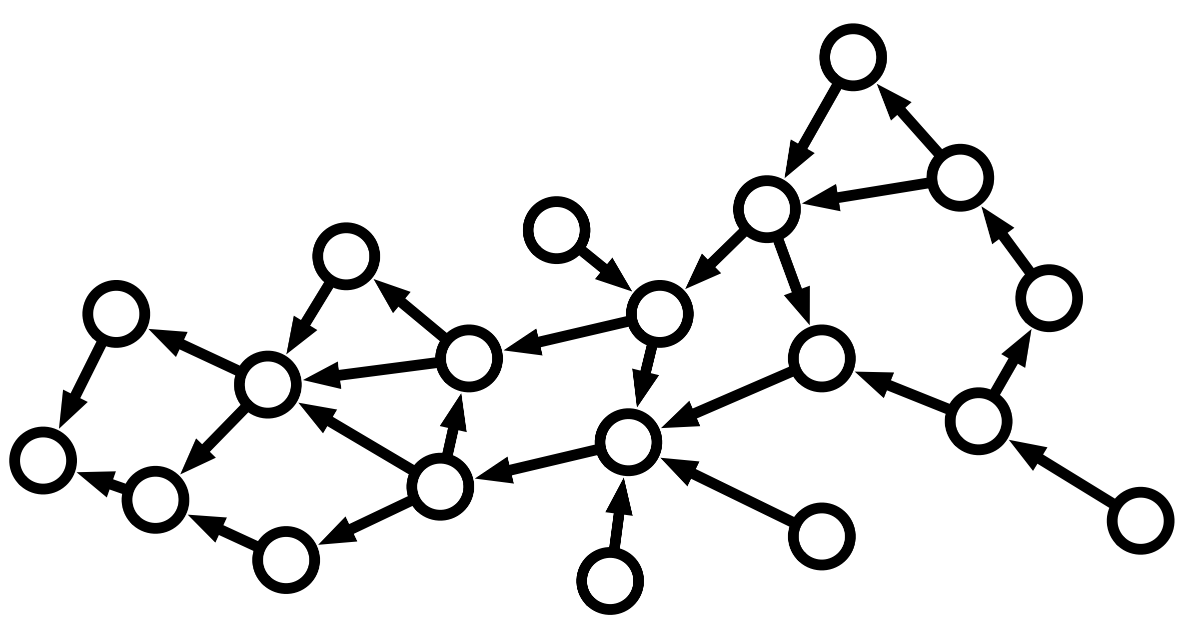

Preliminaries:

Degeneracy

- **The following are equivalent:**
  - graph ***G*** has degeneracy ***d***
  - graph ***G*** has acyclic orientation with out-degree ***d***

- acyclic orientation with out-degree **$O(d)$** can be found in **$O(\log n)$** rounds [Barenboim & Elkin 2010]

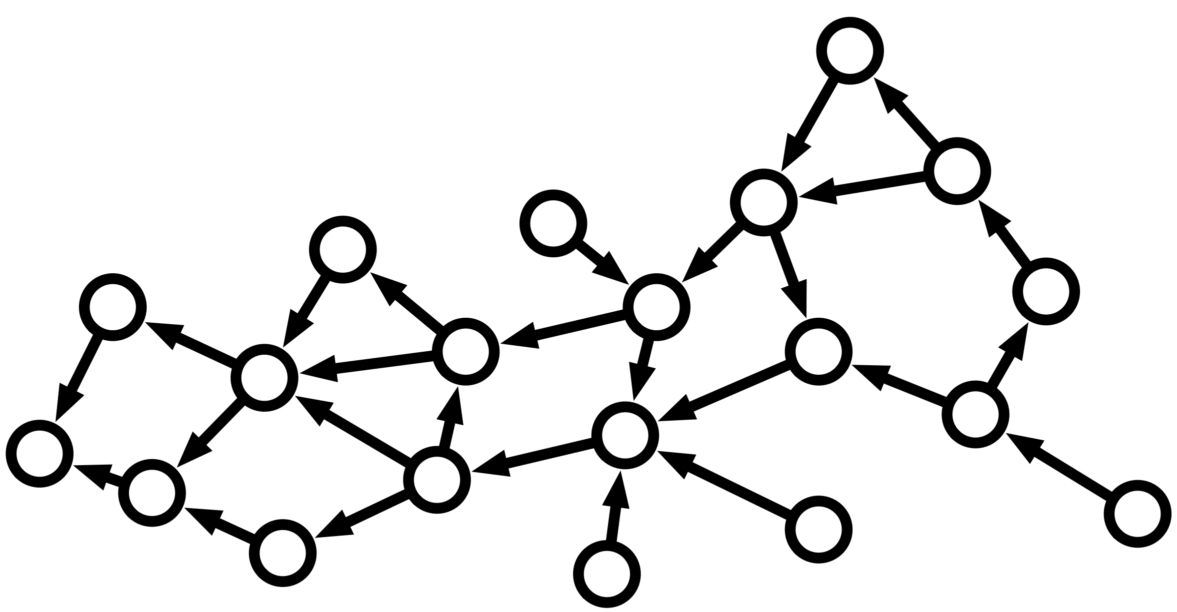

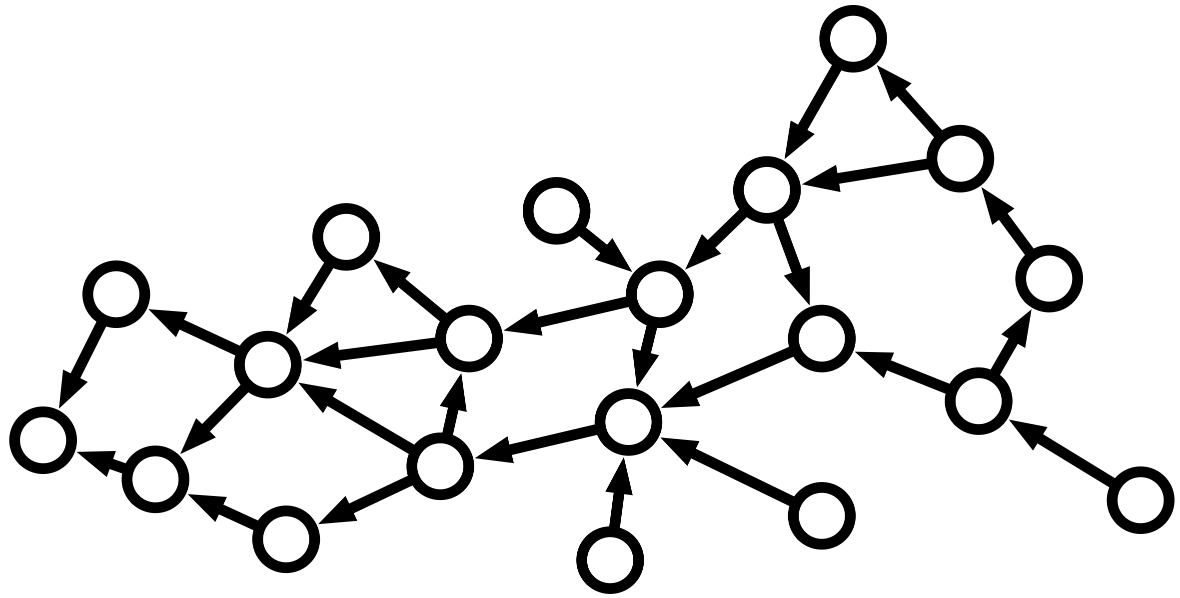

Basic idea: all nodes broadcast their outgoing edges ($O(d)$ rounds)

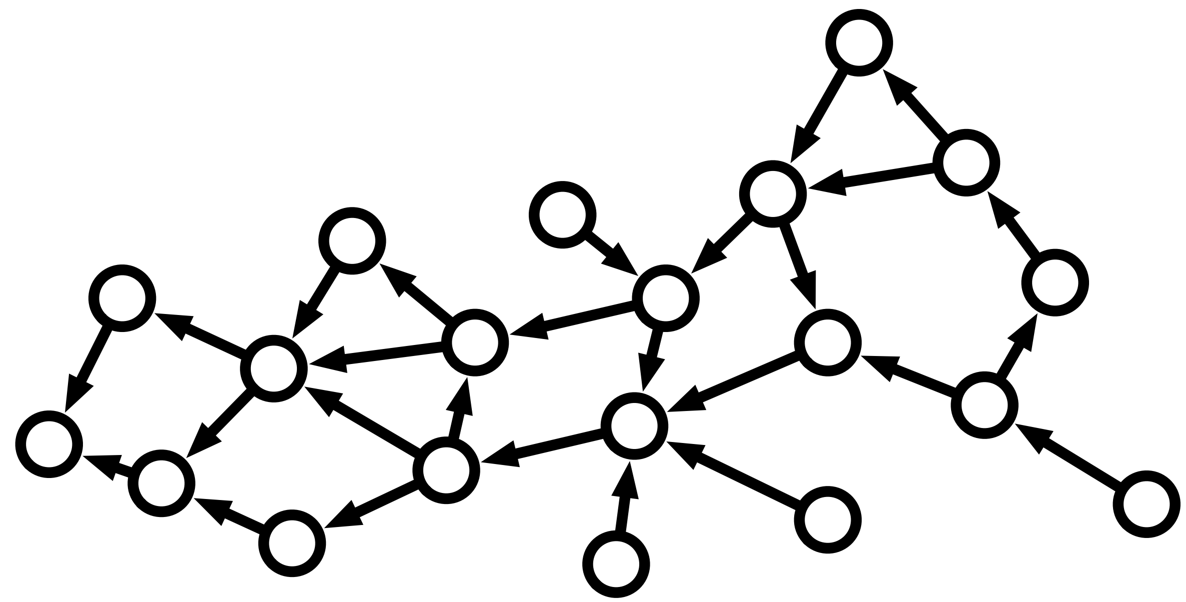

Basic idea: all nodes broadcast their outgoing edges
(***O*(*d*)** rounds)

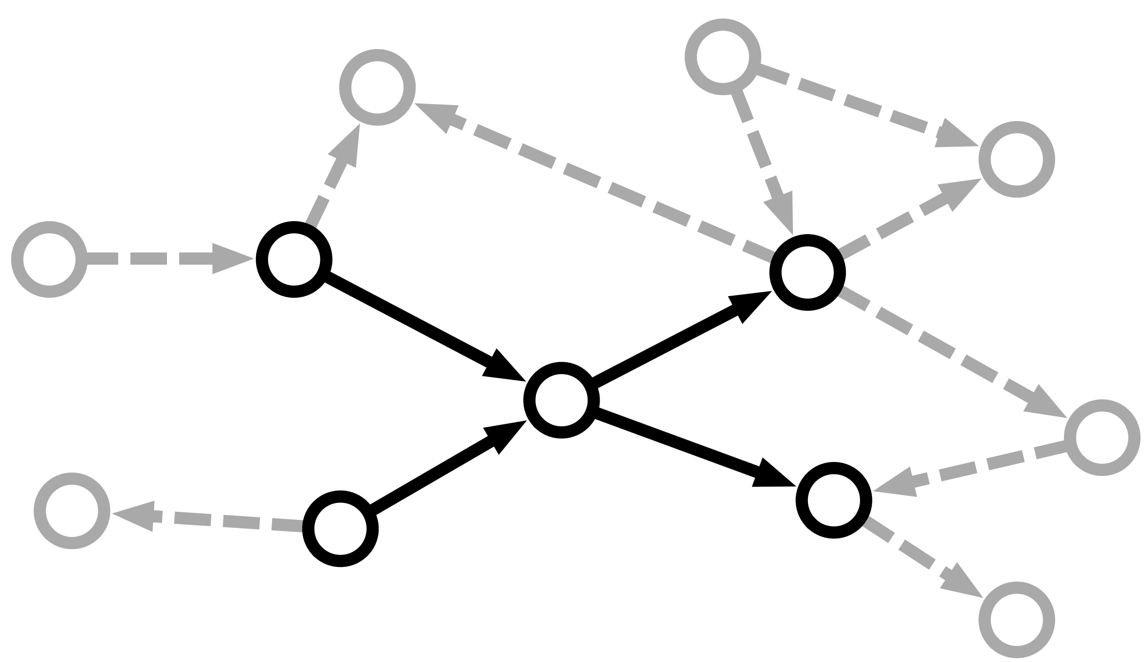

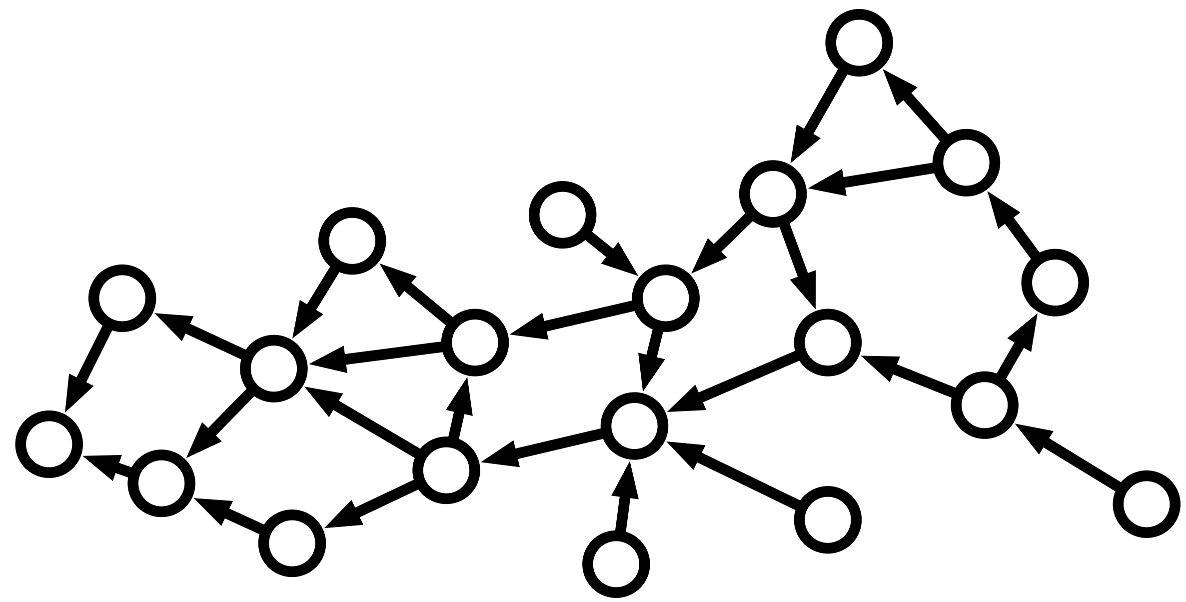

Basic idea: all nodes broadcast their outgoing edges ($O(d)$ rounds)

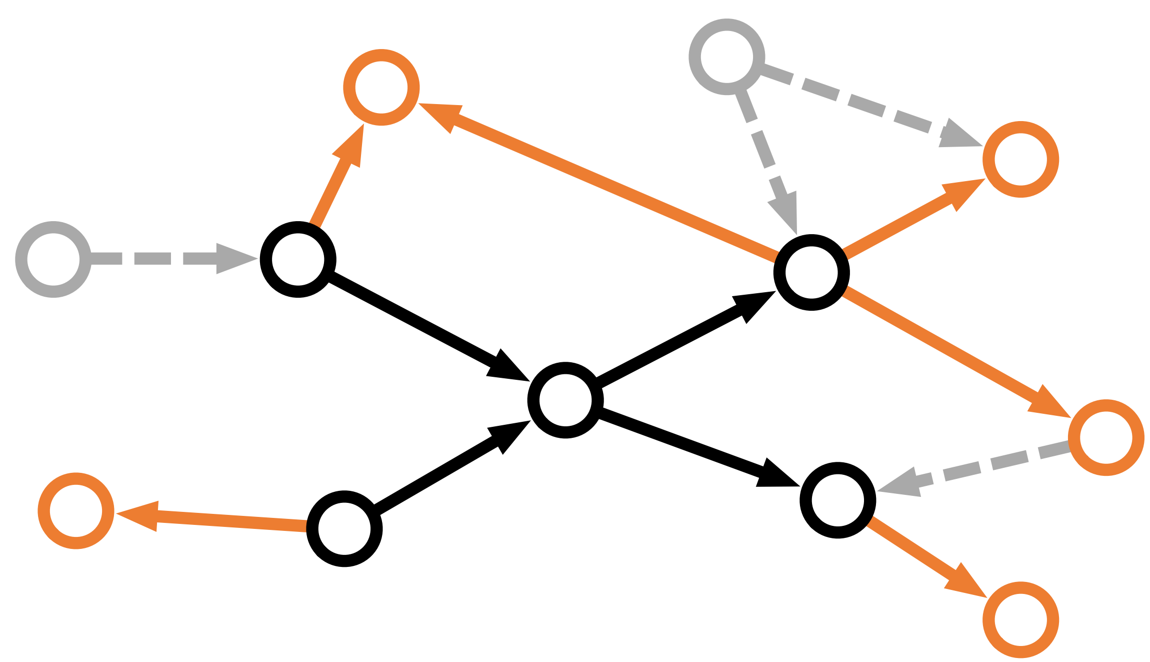

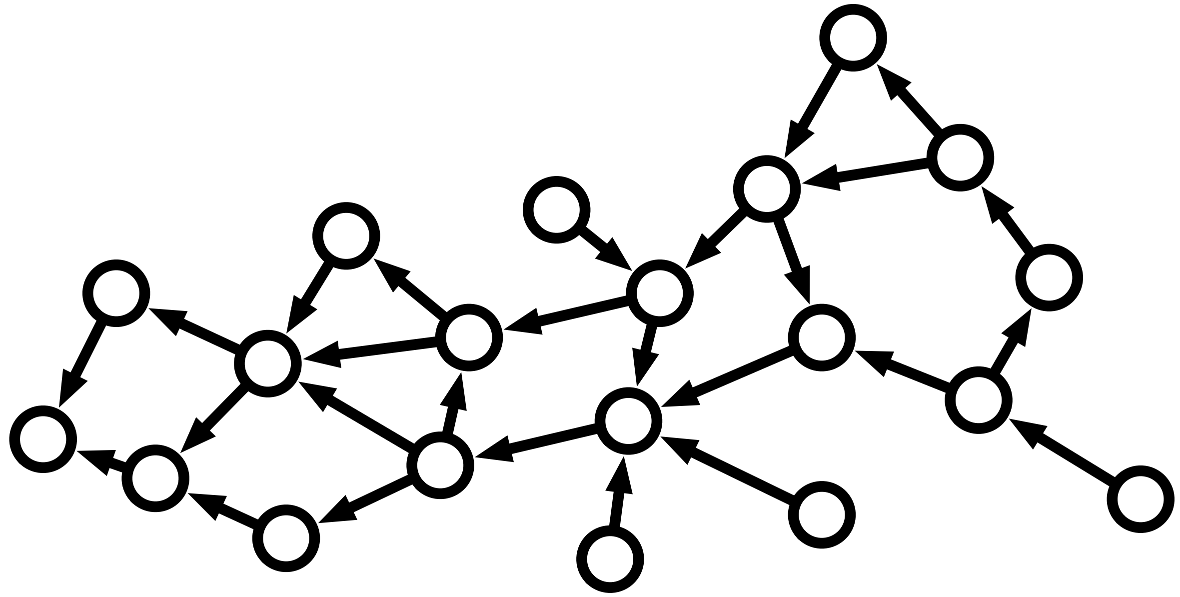

Basic idea: all nodes broadcast their outgoing edges
($O(d)$ rounds)

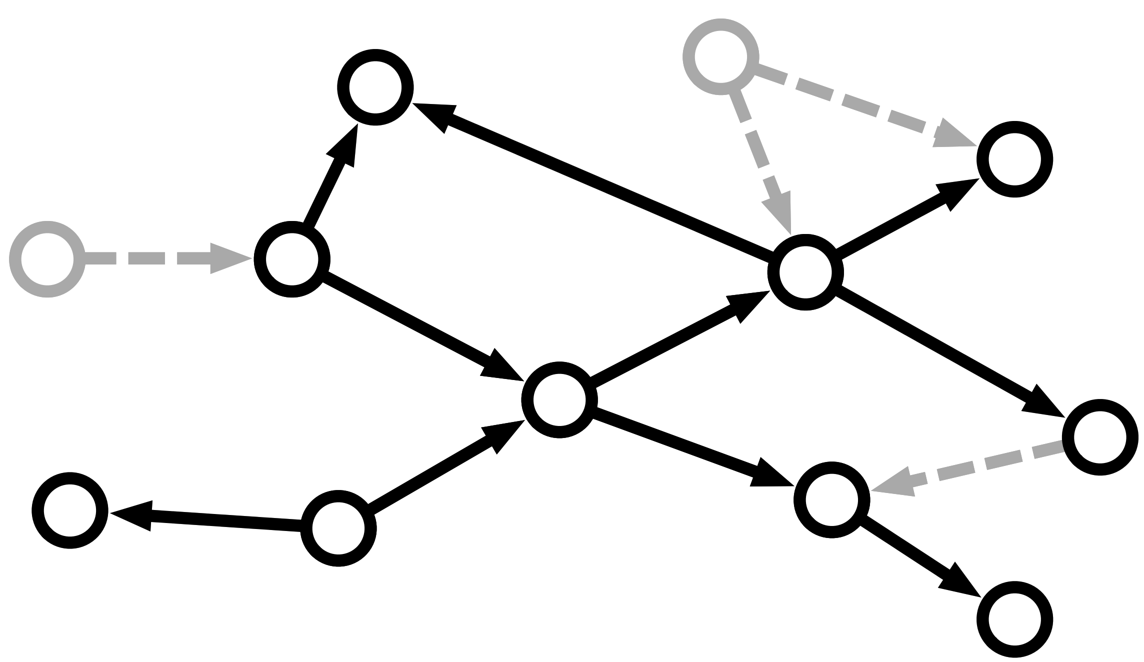

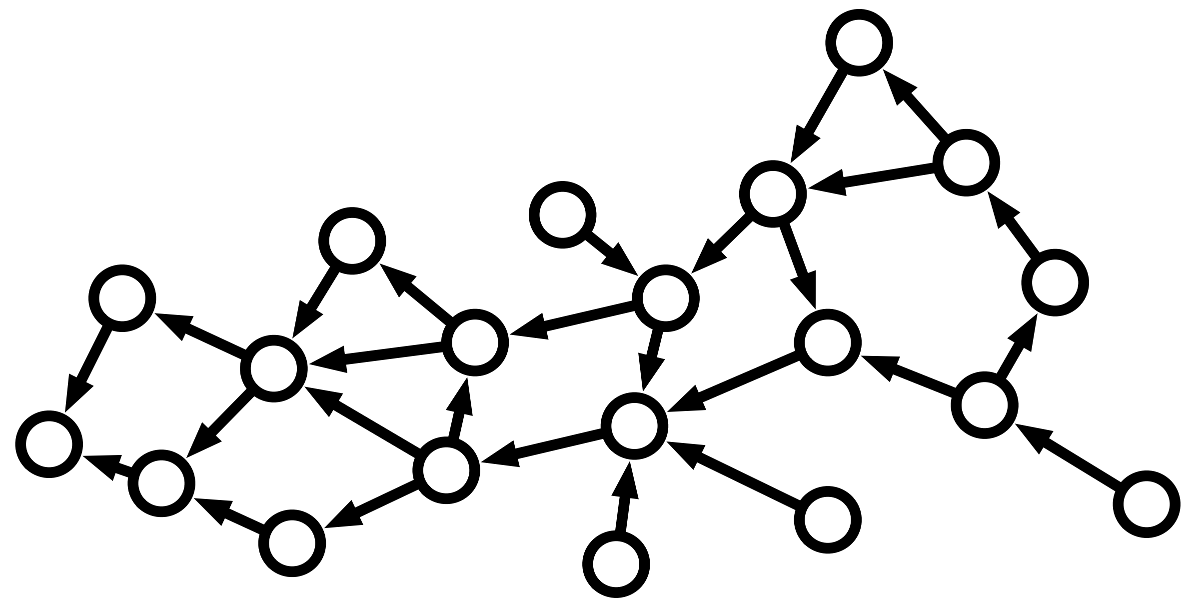

Basic idea: all nodes broadcast their outgoing edges ($O(d)$ rounds)

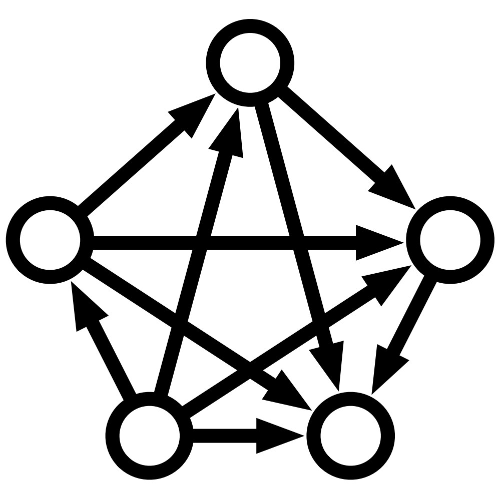

cliques: the sink will see all edges

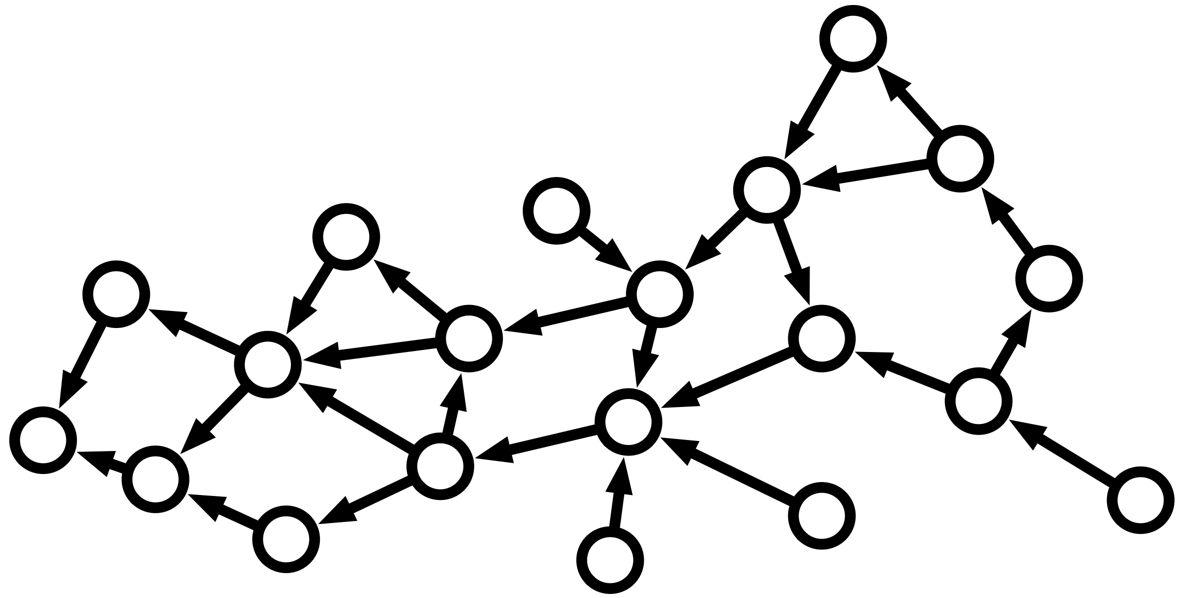

Basic idea: all nodes broadcast their outgoing edges ($O(d)$ rounds)

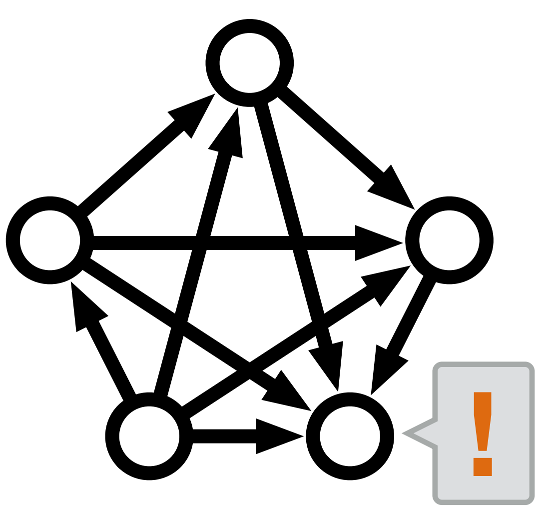

cliques: the sink will see all edges

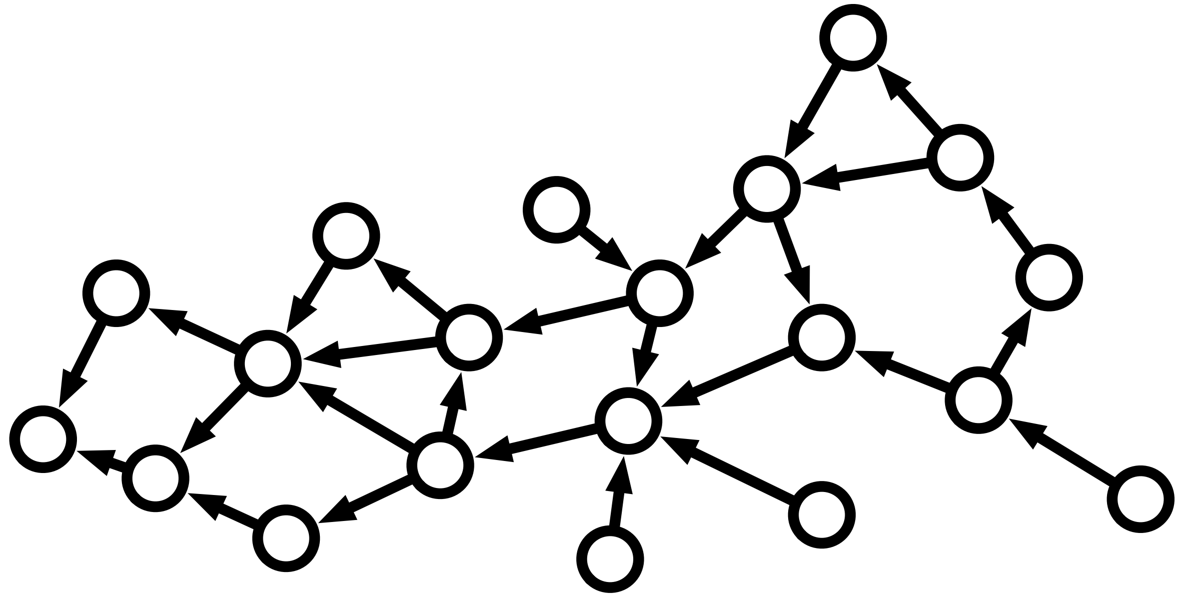

Basic idea: all nodes broadcast their outgoing edges ($O(d)$ rounds)

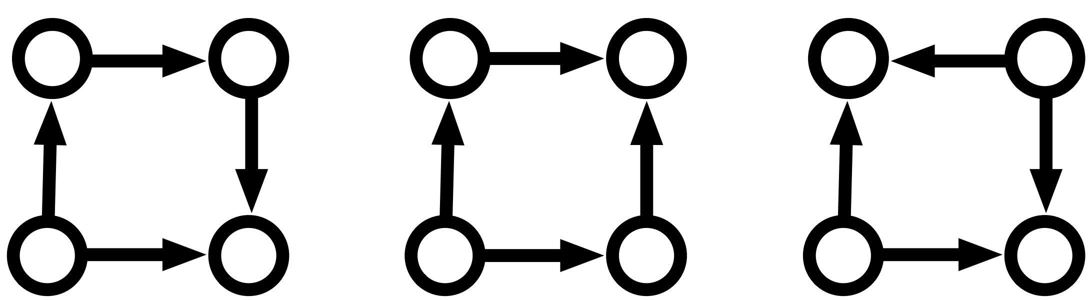

4-cycles: *some* node will see all edges (3 cases to consider)

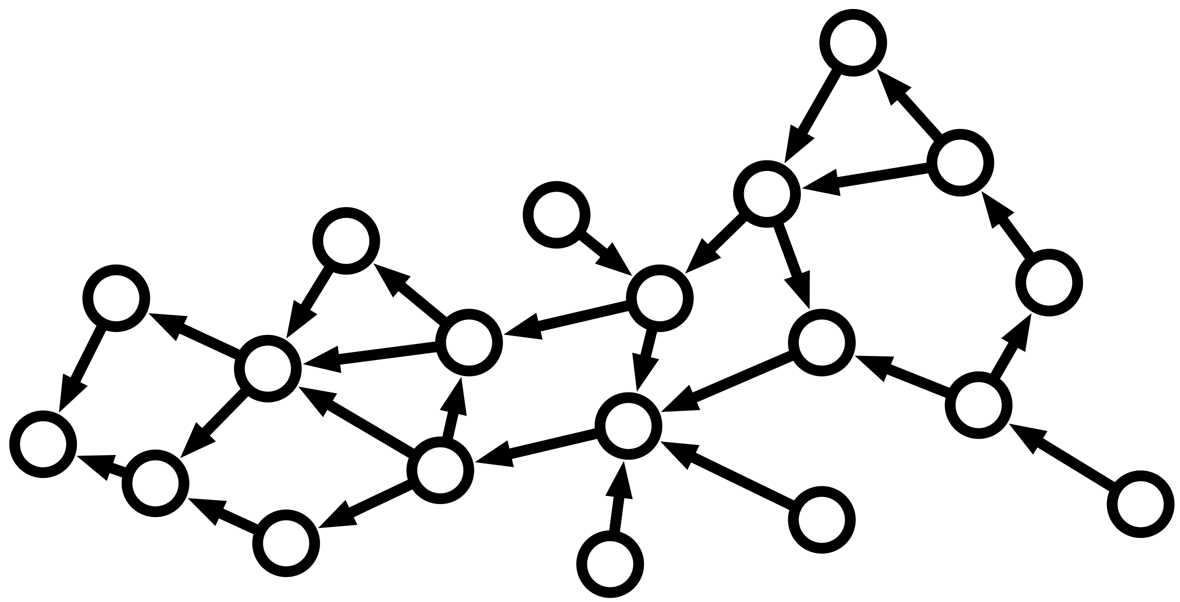

Basic idea: all nodes broadcast their outgoing edges ($O(d)$ rounds)

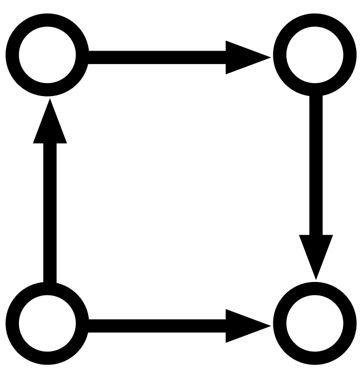

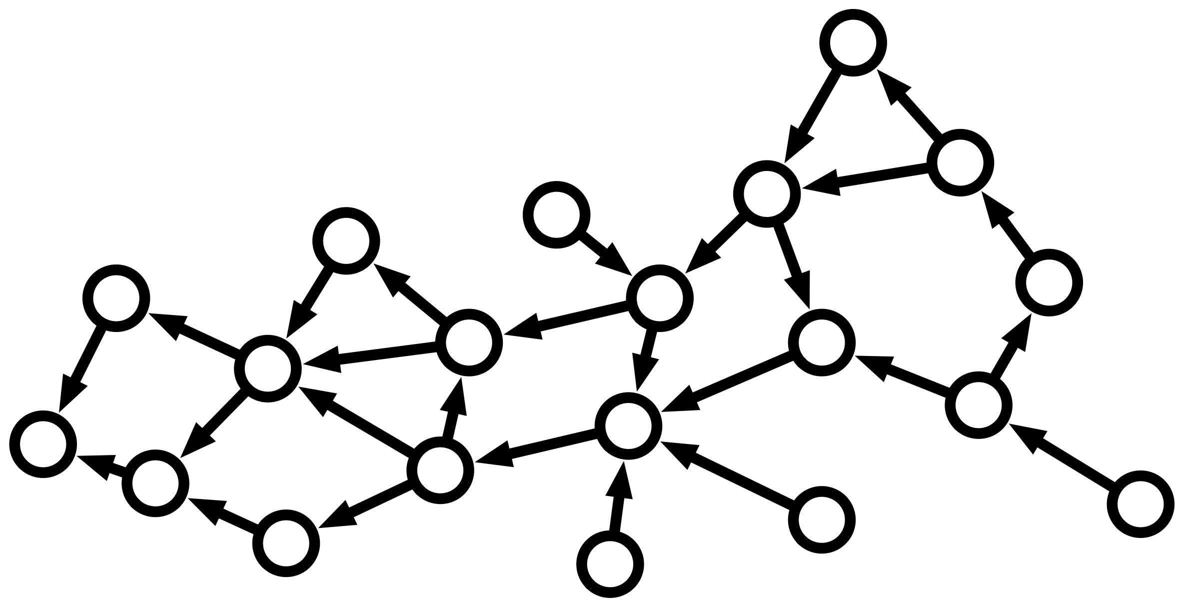

Basic idea: all nodes broadcast their outgoing edges ($O(d)$ rounds)

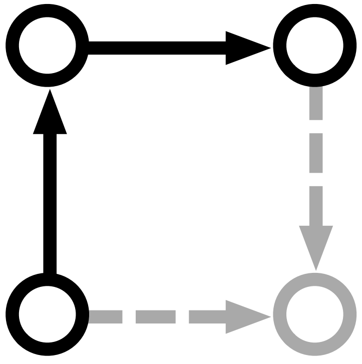

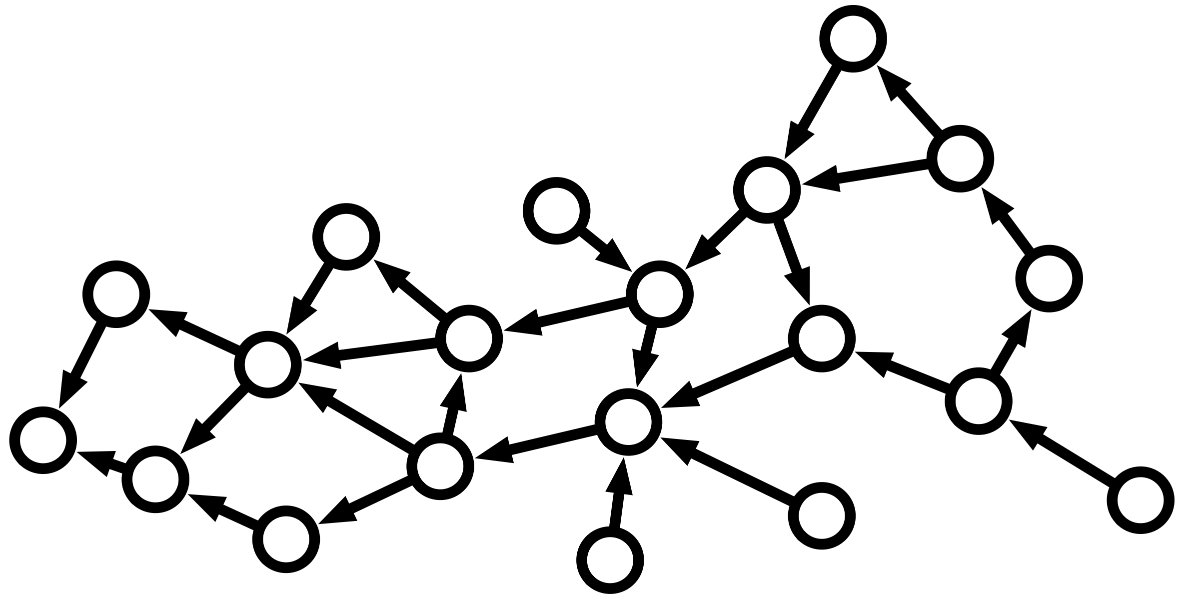

Basic idea: all nodes broadcast their outgoing edges ($O(d)$ rounds)

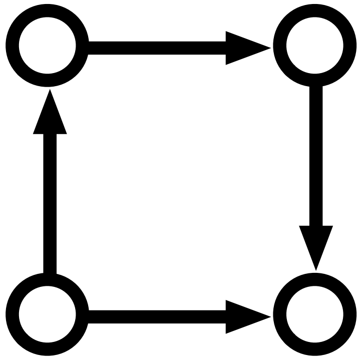

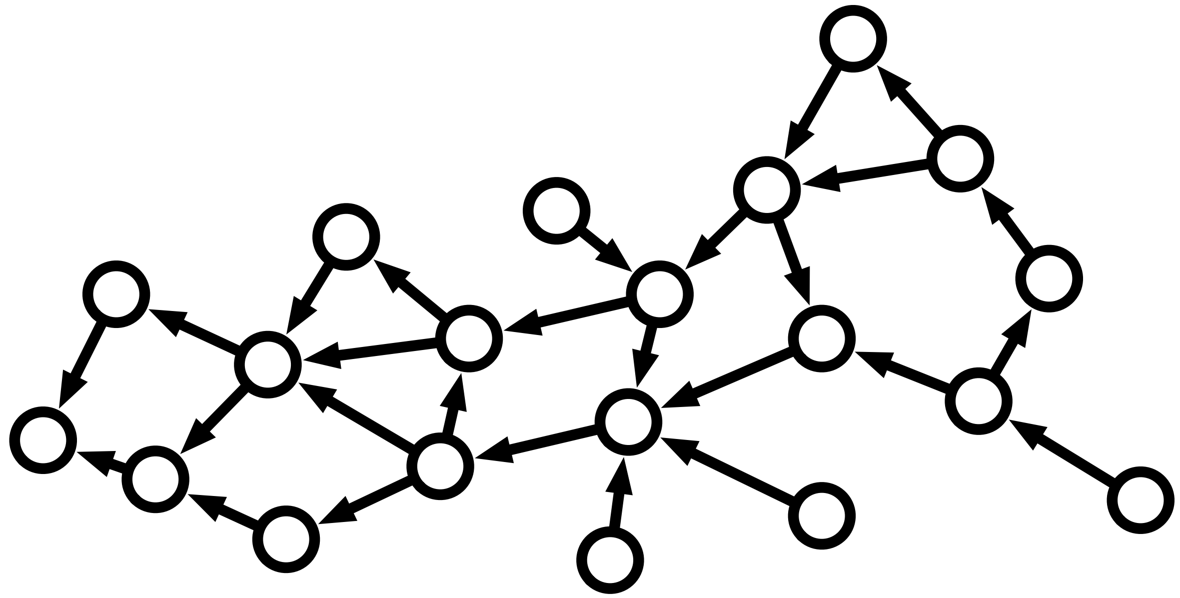

Basic idea: all nodes broadcast their outgoing edges ($O(d)$ rounds)

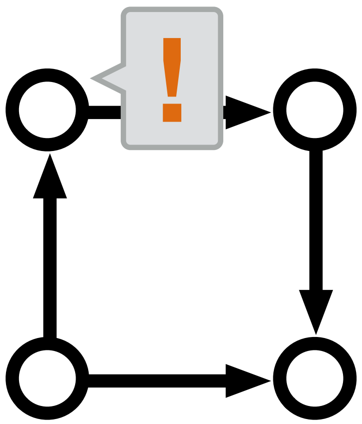

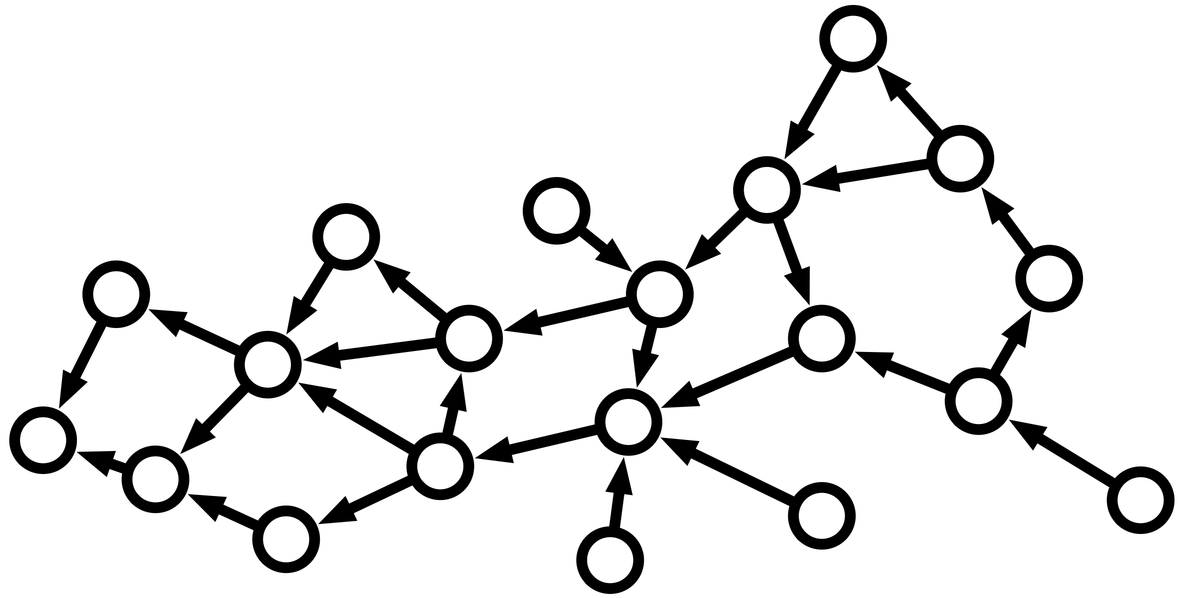

Basic idea: all nodes broadcast their outgoing edges ($O(d)$ rounds)

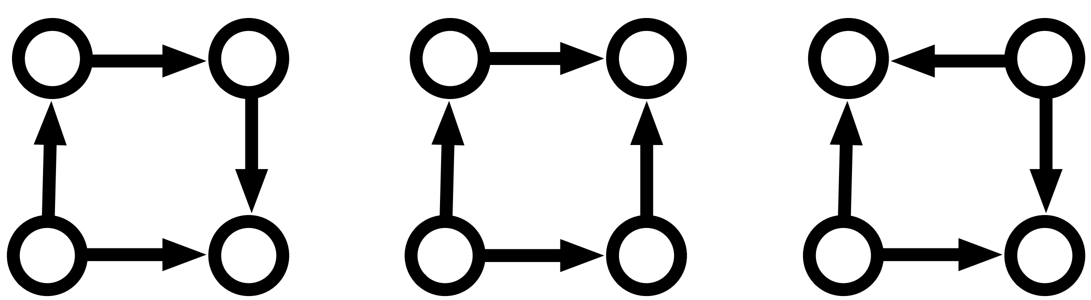

4-cycles: *some* node will see all edges (3 cases to consider)

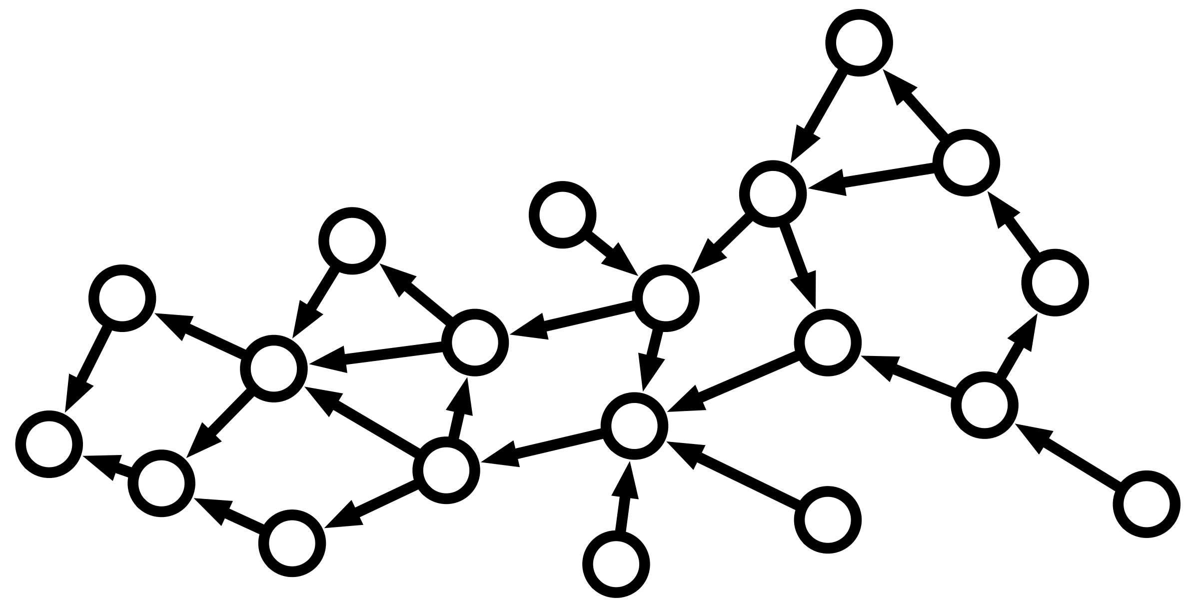

Basic idea: all nodes broadcast their outgoing edges ($O(d)$ rounds)

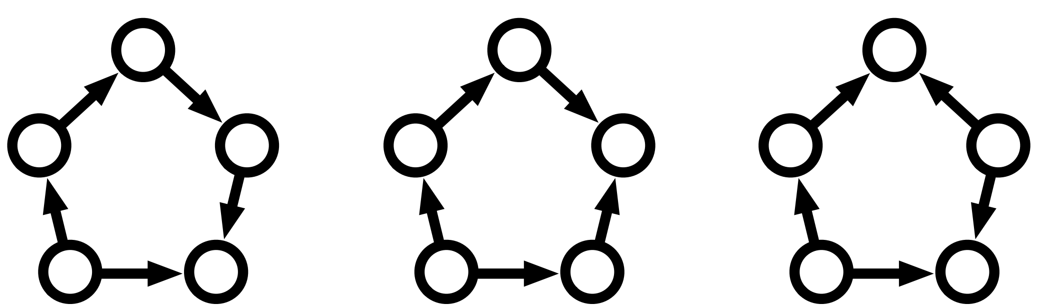

5-cycles: broadcast outgoing 2-paths ($O(d^2)$ rounds)

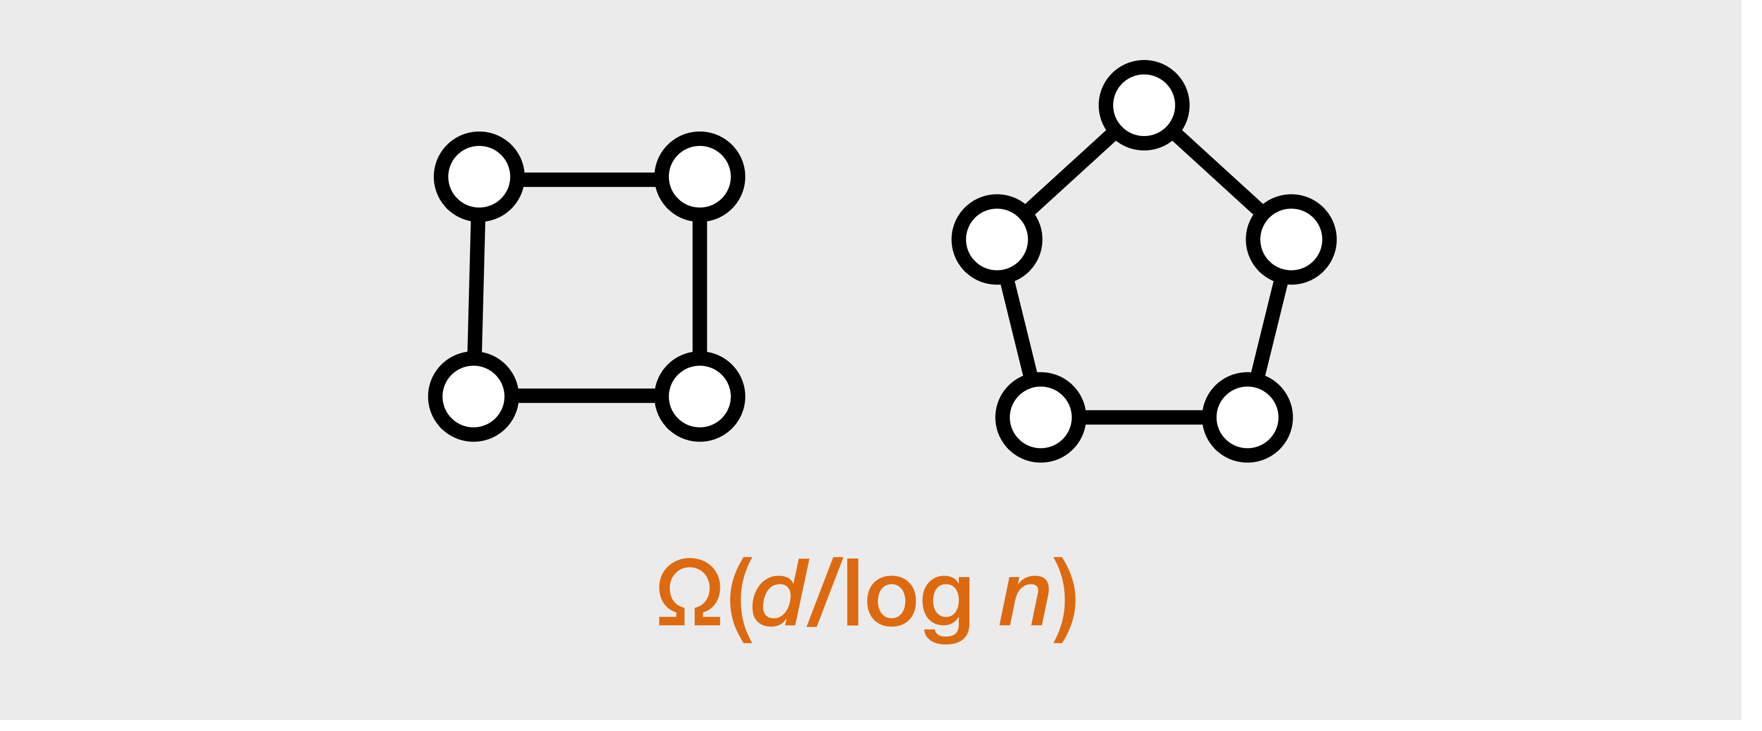
Ω(*d*/log *n*)

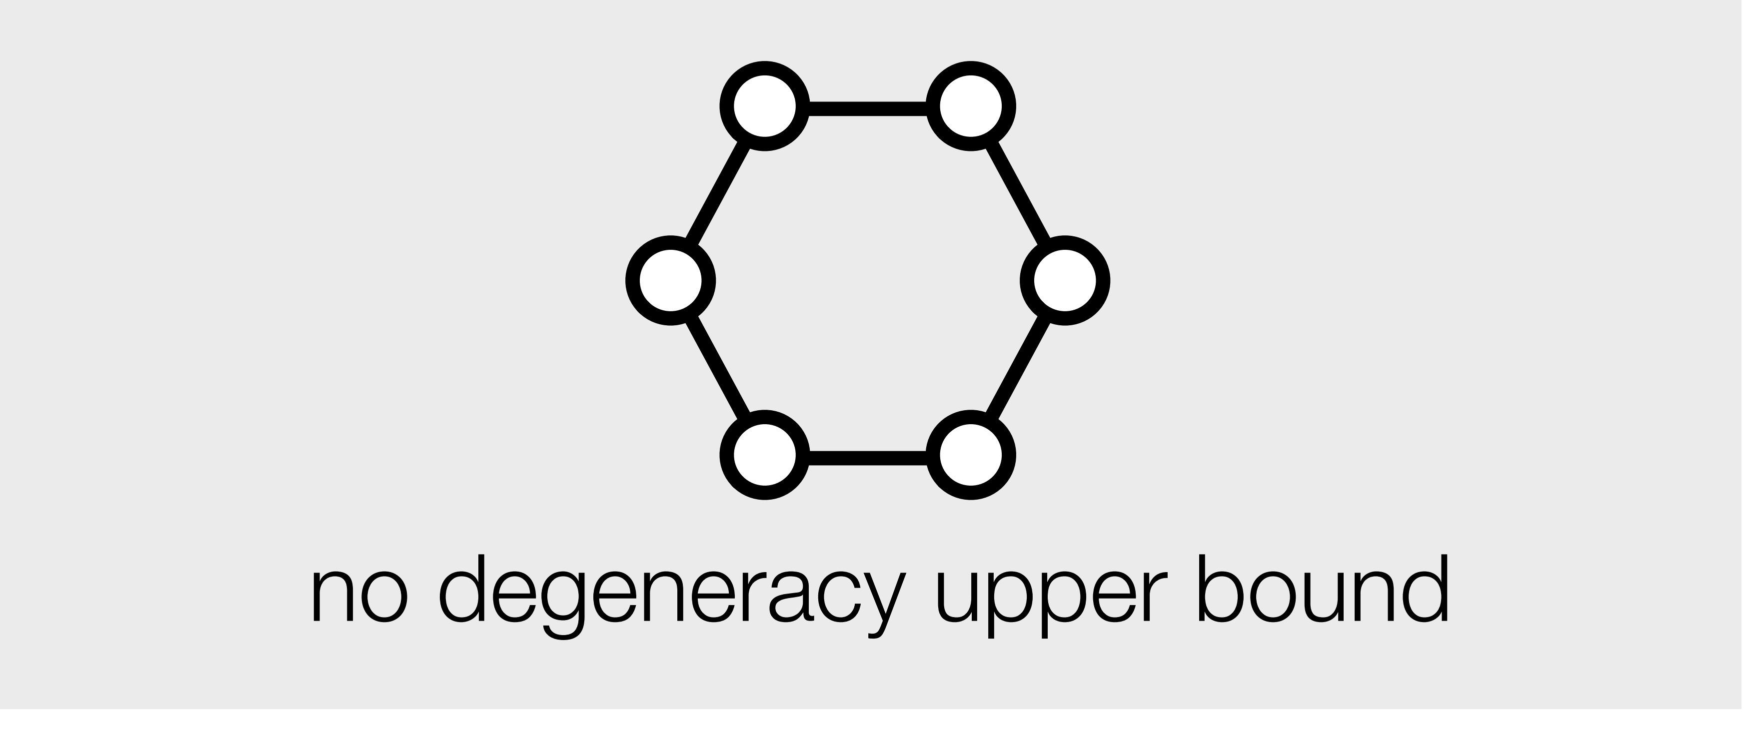
no degeneracy upper bound

5.

Conclusions

Conclusions:

General upper/lower bounds?

- **General question:** given arbitrary ***H***, what is the complexity of detecting ***H***?
  - general upper bound $O(n)$?
  - connection to **tree-width:** trees 1, cycles 2, …?

- **Special cases:**
  - **triangles:** ***???***
  - **even cycles:** gap between $O(n)$ and $\Omega(n^{1/2})$

Conclusions:

General upper/lower bounds?

- **Graphs requiring $\Omega(n^{2-\varepsilon})$ rounds for any $\varepsilon>0$**
  - diameter 3 [Fischer, Gonen & Oshman 2017]
  - tree-width 2 [our work]

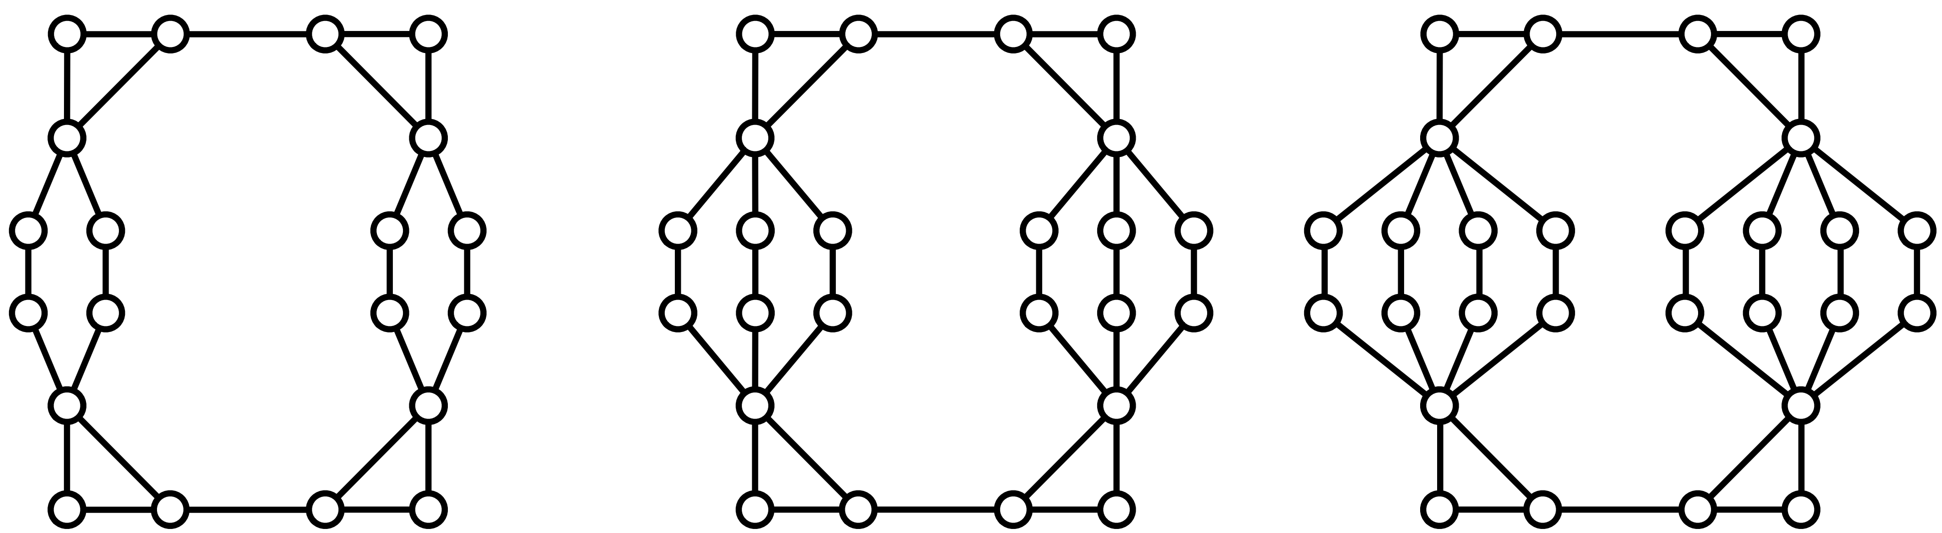

$\Omega(n^{2-1/2})$ $\Omega(n^{2-1/3})$ $\Omega(n^{2-1/4})$ ...

Conclusions:

General upper/lower bounds?

- **Graphs requiring $\Omega(n^{2-\varepsilon})$ rounds for any $\varepsilon>0$**
  - diameter 3 [Fischer, Gonen & Oshman 2017]
  - tree-width 2 [our work]

- **Corresponding upper bound?**
  - lower bound $\Omega(n^2/\text{polylog } n)$ does not seem possible with standard techniques
  - **conjecture:** for any H, some $O(n^{2-\varepsilon})$ upper bound

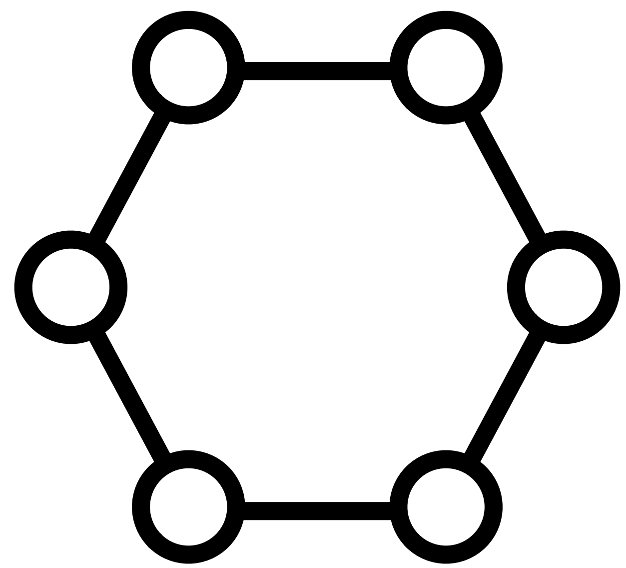

Thanks! Questions?