# Distributed Computation and Reconfiguration in Actively Dynamic Networks

#### Othon Michail George Skretas Paul G. Spirakis

Department of Computer Science, University of Liverpool, UK Computer Engineering and Informatics Department, University of Patras, Greece

#### ACM Symposium on Principles of Distributed Computing 3-6 August 2020



- Networks whose structure may change
- Usually represented by a graph
- Edges and/or nodes come and go
- Dynamics:
  - Active or Passive
  - Centralized or Distributed



- Networks whose structure may change
- Usually represented by a graph
- Edges and/or nodes come and go

Dynamics:

- Active or Passive
- Centralized or Distributed



- Networks whose structure may change
- Usually represented by a graph
- Edges and/or nodes come and go

Dynamics:

- Active or Passive
- Centralized or Distributed



- Networks whose structure may change
- Usually represented by a graph
- Edges and/or nodes come and go ynamics:
- Active or Passive
- Centralized or Distributed





- Networks whose structure may change
- Usually represented by a graph
- Edges and/or nodes come and go

Dynamics:

- Active or Passive
- Centralized or Distributed





- Networks whose structure may change
- Usually represented by a graph
- Edges and/or nodes come and go

Dynamics:

- Active or Passive
- Centralized or Distributed



- Networks whose structure may change
- Usually represented by a graph
- Edges and/or nodes come and go
- Dynamics:
  - Active or Passive
  - Centralized or Distributed





- Networks whose structure may change
- Usually represented by a graph
- Edges and/or nodes come and go Dynamics:
  - Active or Passive
  - Centralized or Distributed





- Networks whose structure may change
- Usually represented by a graph
- Edges and/or nodes come and go Dynamics:
  - Active or Passive
  - Centralized or Distributed





- Networks whose structure may change
- Usually represented by a graph
- Edges and/or nodes come and go Dynamics:
  - Active or Passive
  - Centralized or Distributed





- No central coordinator
- General Goal:
  - Transform the initial network  $G_s$  into a target network  $G_f$  from a family of target networks
  - Exploit some good properties of  $G_f$  in order to more efficiently solve a distributed task

#### • No central coordinator

- Transform the initial network  $G_s$  into a target network  $G_f$  from a family of target networks
- Exploit some good properties of  $G_f$  in order to more efficiently solve a distributed task

#### • No central coordinator

- Transform the initial network  $G_s$  into a target network  $G_f$  from a family of target networks
- Exploit some good properties of  $G_f$  in order to more efficiently solve a distributed task

#### • No central coordinator

- Transform the initial network  $G_s$  into a target network  $G_f$  from a family of target networks
- Exploit some good properties of  $G_f$  in order to more efficiently solve a distributed task

#### • No central coordinator

- Transform the initial network  $G_s$  into a target network  $G_f$  from a family of target networks
- Exploit some good properties of  $G_f$  in order to more efficiently solve a distributed task

• No central coordinator

- Transform the initial network  $G_s$  into a target network  $G_f$  from a family of target networks
- Exploit some good properties of  $G_f$  in order to more efficiently solve a distributed task

- Activate connections to new neighbors or deactivate connections
- No central coordinator
- General Goal:
  - Transform the initial network  $G_s$  into a target network  $G_f$  from a family of target networks
  - Exploit some good properties of G<sub>f</sub> in order to more efficiently solve a distributed task

- Activate connections to new neighbors or deactivate connections
- No central coordinator
- General Goal:
  - Transform the initial network  $G_s$  into a target network  $G_f$  from a family of target networks
  - Exploit some good properties of G<sub>f</sub> in order to more efficiently solve a distributed task



# • Temporal Graphs

- Initiated by Berman [Be09] and Kempe et al. [KKK00]
- Distributed Computation in Passively Dynamic Networks
  - Kuhn et al. [KLO10]
  - Population Protocols [AADFP06]
- Construction of Overlay Networks
  - Most similar to our model
  - Introduced by Stoica et al. [SMK01] and [AS07]
- Programmable Matter
  - Geometry plays a big role

O. Michail, G. Skretas, P. G. Spirakis Distributed Computation and Reconfiguration in Actively Dynamic Networks 5 / 26

Our Model

- Initial connected network G<sub>s</sub>
- Static set V of n nodes
- Dynamic set E(i) of m active edges
- Standard synchronous message passing model
- Each node has a unique ID
- Nodes can only compare unique IDs
- Node *u* can activate an edge with node *v* if they have a common neighbor
- One edge between two nodes





- O. Michail, G. Skretas, P. G. Spirakis Distributed Computation and Reconfiguration in Actively Dynamic Networks 5 / 26

- Initial connected network G<sub>s</sub>
- Static set V of n nodes
- Dynamic set E(i) of m active edges
- Standard synchronous message passing model
- Each node has a unique ID
- Nodes can only compare unique IDs
- Node *u* can activate an edge with node *v* if they have a common neighbor
- One edge between two nodes





UNIVERSITY OF LIVERPOOL

- Initial connected network G<sub>s</sub>
- Static set V of n nodes
- Dynamic set E(i) of m active edges
- Standard synchronous message passing model
- Each node has a unique ID
- Nodes can only compare unique IDs
- Node *u* can activate an edge with node *v* if they have a common neighbor
- One edge between two nodes



UNIVERSITY OF LIVERPOOL

- Initial connected network G<sub>s</sub>
- Static set V of n nodes
- Dynamic set E(i) of m active edges
- Standard synchronous message passing model
- Each node has a unique ID
- Nodes can only compare unique IDs
- Node *u* can activate an edge with node *v* if they have a common neighbor
- One edge between two nodes





- Initial connected network G<sub>s</sub>
- Static set V of n nodes
- Dynamic set E(i) of m active edges
- Standard synchronous message passing model
- Each node has a unique ID
- Nodes can only compare unique IDs
- Node *u* can activate an edge with node *v* if they have a common neighbor
- One edge between two nodes



O. Michail, G. Skretas, P. G. Spirakis Distributed Computation and Reconfiguration in Actively Dynamic Networks 5 / 26

- Initial connected network G<sub>s</sub>
- Static set V of n nodes
- Dynamic set *E*(*i*) of *m* active edges
- Standard synchronous message passing model
- Each node has a unique ID
- Nodes can only compare unique IDs
- Node *u* can activate an edge with node *v* if they have a common neighbor
- One edge between two nodes





- Dynamic set E(i) of *m* active edges
  - Standard synchronous message passing model
  - Each node has a unique ID

• Initial connected network  $G_{s}$ 

• Static set V of n nodes

- Nodes can only compare unique IDs
- Node *u* can activate an edge with node *v* if they have a common neighbor
- One edge between two nodes



O. Michail, G. Skretas, P. G. Spirakis Distributed Computation and Reconfiguration in Actively Dynamic Networks 5 / 26

- Initial connected network G<sub>s</sub>
- Static set V of n nodes
- Dynamic set *E*(*i*) of *m* active edges
- Standard synchronous message passing model
- Each node has a unique ID
- Nodes can only compare unique IDs
- Node *u* can activate an edge with node *v* if they have a common neighbor
- One edge between two nodes





Our Model

- Initial connected network  $G_s$
- Static set V of n nodes
- Dynamic set E(i) of m active edges
- Standard synchronous message passing model
- Each node has a unique ID
- Nodes can only compare unique IDs
- Node *u* can activate an edge with node *v* if they have a common neighbor
- One edge between two nodes





- Initial connected network G<sub>s</sub>
- Static set V of n nodes
- Dynamic set E(i) of m active edges
- Standard synchronous message passing model
- Each node has a unique ID
- Nodes can only compare unique IDs
- Node *u* can activate an edge with node *v* if they have a common neighbor
- One edge between two nodes

Our Model





- Leader Election: Elect a unique leader
- Token Dissemination: Each node has a unique piece of information. Every piece of information has to be disseminated to every node
- Depth-d Tree: Transform the initial network into a tree of depth d

#### • Leader Election: Elect a unique leader

- Token Dissemination: Each node has a unique piece of information. Every piece of information has to be disseminated to every node
- Depth-d Tree: Transform the initial network into a tree of depth d



- Leader Election: Elect a unique leader
- Token Dissemination: Each node has a <u>unique</u> piece of information. Every piece of information has to be disseminated to every node
- Depth-d Tree: Transform the initial network into a tree of depth d



- Leader Election: Elect a unique leader
- Token Dissemination: Each node has a <u>unique</u> piece of information. Every piece of information has to be disseminated to every node
- Depth-d Tree: Transform the initial network into a tree of depth d

#### UNIVERSITY OF LIVERPOOL

#### Very simple strategy:

- Activate an edge with every distance 2 neighbor
- Spanning clique, log *n* rounds
- Eliminate extra edges



# Very simple strategy:

- Activate an edge with every distance 2 neighbor
- Spanning clique, log *n* rounds
- Eliminate extra edges




- Activate an edge with every distance 2 neighbor
- Spanning clique, log *n* rounds
- Eliminate extra edges





- Activate an edge with every distance 2 neighbor
- Spanning clique, log *n* rounds
- Eliminate extra edges





- Activate an edge with every distance 2 neighbor
- Spanning clique, log *n* rounds
- Eliminate extra edges





- Activate an edge with every distance 2 neighbor
- Spanning clique, log *n* rounds
- Eliminate extra edges





- Activate an edge with every distance 2 neighbor
- Spanning clique, log *n* rounds
- Eliminate extra edges





- Activate an edge with every distance 2 neighbor
- Spanning clique, log *n* rounds
- Eliminate extra edges
- BUT





- Activate an edge with every distance 2 neighbor
- Spanning clique, log *n* rounds
- Eliminate extra edges

### BUT

• Activating and maintaining a connection does not come for free



- Total Edge Activations: The number of edges activated by the algorithm
- Maximum Activated Edges: The maximum number of activated edges by the algorithm per round
- Maximum Activated Degree: The maximum degree of the network based only on the activated edges by the algorithm



- Total Edge Activations: The number of edges activated by the algorithm
- Maximum Activated Edges: The maximum number of activated edges by the algorithm per round
- Maximum Activated Degree: The maximum degree of the network based only on the activated edges by the algorithm



- Total Edge Activations: The number of edges activated by the algorithm
- Maximum Activated Edges: The maximum number of activated edges by the algorithm per round
- Maximum Activated Degree: The maximum degree of the network based only on the activated edges by the algorithm



- Total Edge Activations: The number of edges activated by the algorithm
- Maximum Activated Edges: The maximum number of activated edges by the algorithm per round
- Maximum Activated Degree: The maximum degree of the network based only on the activated edges by the algorithm

Algorithm	Time	Total Edge Activations	Maximum Activated Edges	Maximum Activated Degree
GraphToStar	$O(\log n)$	$O(n \log n)$	<i>O</i> ( <i>n</i> )	<i>O</i> ( <i>n</i> )
GraphToWreath	$O(\log^2 n)$	$O(n \log^2 n)$	O(n)	O(1)
GraphToThinWreath	$O(\log^2 n / \log \log n)$	$O(n \log^2 n)$	<i>O</i> ( <i>n</i> )	$O(\log^2 n)$

### Table: Algorithms

Bounds	Time	Total Edge Activations	Maximum Activated Edges	Maximum Activated Degree
Centralized Lower	$\Omega(\log n)$	$\Omega(n)$	$\Omega(n/\log n)$	
Centralized Upper	$O(\log n)$	$\Theta(n)$		
Distributed Lower	$O(\log n)$	$\Omega(n \log n)$		

#### Table: Bounds for Depth-d tree

### • Nodes are partitioned into committees

- Committees organized into gadget networks
- Each node forms its own committee
- Committees compete and merge
- Final network has a single committee
- Logarithmic running time
- Time: phases\*gadgetdiameter

LIVERPOO

- Nodes are partitioned into committees
- Committees organized into gadget networks
- Each node forms its own committee
- Committees compete and merge
- Final network has a single committee
- Logarithmic running time
- Time: phases\*gadgetdiameter

UVERPOO

UNIVERSITY OF LIVERPOOL

- Nodes are partitioned into committees
- Committees organized into gadget networks
- Each node forms its own committee
- Committees compete and merge
- Final network has a single committee
- Logarithmic running time
- Time: phases\*gadgetdiameter





- Nodes are partitioned into committees
- Committees organized into gadget networks
- Each node forms its own committee
- Committees compete and merge
- Final network has a single committee
- Logarithmic running time
- Time: phases\*gadgetdiameter





- Nodes are partitioned into committees
- Committees organized into gadget networks
- Each node forms its own committee
- Committees compete and merge
- Final network has a single committee
- Logarithmic running time
- Time: phases\*gadgetdiameter



10 / 26







- Nodes are partitioned into committees
- Committees organized into gadget networks
- Each node forms its own committee
- Committees compete and merge
- Final network has a single committee
- Logarithmic running time
- Time: phases\*gadgetdiameter







- Nodes are partitioned into committees
- Committees organized into gadget networks
- Each node forms its own committee
- Committees compete and merge
- Final network has a single committee
- Logarithmic running time
- Time: phases\*gadgetdiameter



LIVERPOO



- Nodes are partitioned into committees
- Committees organized into gadget networks
- Each node forms its own committee
- Committees compete and merge
- Final network has a single committee
- Logarithmic running time
- Time: phases\*gadgetdiameter



LIVERPOO





## • Gadget network: Star

• TreeToStar subroutine

TreeToStar

- Transforms any initial oriented Tree graph into a spanning Star graph in log *n* time
- Activates an edge with grandparent
- Deactivate an edge with parent

UNIVERSITY OF LIVERPOOL

- Gadget network: Star
- TreeToStar subroutine
- TreeToStar
  - Transforms any initial oriented Tree graph into a spanning Star graph in log *n* time
  - Activates an edge with grandparent
  - Deactivate an edge with parent



- Gadget network: Star
- TreeToStar subroutine

### TreeToStar

- Transforms any initial oriented Tree graph into a spanning Star graph in log *n* time
- Activates an edge with grandparent
- Deactivate an edge with parent





- Gadget network: Star
- TreeToStar subroutine
- TreeToStar
  - Transforms any initial oriented Tree graph into a spanning Star graph in log *n* time
  - Activates an edge with grandparent
  - Deactivate an edge with parent





- Gadget network: Star
- TreeToStar subroutine
- TreeToStar
  - Transforms any initial oriented Tree graph into a spanning Star graph in log *n* time
  - Activates an edge with grandparent
  - Deactivate an edge with parent





- Gadget network: Star
- TreeToStar subroutine
- TreeToStar
  - Transforms any initial oriented Tree graph into a spanning Star graph in log *n* time
  - Activates an edge with grandparent
  - Deactivate an edge with parent





- Gadget network: Star
- TreeToStar subroutine
- TreeToStar
  - Transforms any initial oriented Tree graph into a spanning Star graph in log *n* time
  - Activates an edge with grandparent
  - Deactivate an edge with parent





- Gadget network: Star
- TreeToStar subroutine
- TreeToStar
  - Transforms any initial oriented Tree graph into a spanning Star graph in log *n* time
  - Activates an edge with grandparent
  - Deactivate an edge with parent





- Gadget network: Star
- TreeToStar subroutine

TreeToStar

- Transforms any initial oriented Tree graph into a spanning Star graph in log *n* time
- Activates an edge with grandparent
- Deactivate an edge with parent





- Gadget network: Star
- TreeToStar subroutine
- TreeToStar
  - Transforms any initial oriented Tree graph into a spanning Star graph in log *n* time
  - Activates an edge with grandparent
  - Deactivate an edge with parent






























Theorem: For any initial connected graph  $G_s$ , the GraphToStar algorithm solves the Depth-1 Tree problem in  $O(\log n)$  time with at most  $O(n \log n)$  total edge activations and O(n) active edges per round

### Correctness

- Committees keep merging
- Can't get isolated indefinitely
- Time Complexity
  - Committees "double" in size
  - Isolated
- Edge Complexity
  - Omitted



LIVERPOO

### • Gadget network: Wreath

• LineToCompleteBinaryTree subroutine

#### LineToCompleteBinaryTree

- Transforms any initial oriented line into a spanning Complete Binary Tree in log *n* time
- Activates an edge with grandparent
- Deactivate an edge with parent
- Stop when grandparent has two children



- Gadget network: Wreath
- LineToCompleteBinaryTree subroutine
- LineToCompleteBinaryTree
  - Transforms any initial oriented line into a spanning Complete Binary Tree in log *n* time
  - Activates an edge with grandparent
  - Deactivate an edge with parent
  - Stop when grandparent has two children





15 / 26

- Gadget network: Wreath
- LineToCompleteBinaryTree subroutine
- LineToCompleteBinaryTree
  - Transforms any initial oriented line into a spanning Complete Binary Tree in log *n* time
  - Activates an edge with grandparent
  - Deactivate an edge with parent
  - Stop when grandparent has two children





- Gadget network: Wreath
- LineToCompleteBinaryTree subroutine
- LineToCompleteBinaryTree
  - Transforms any initial oriented line into a spanning Complete Binary Tree in log *n* time
  - Activates an edge with grandparent
  - Deactivate an edge with parent
  - Stop when grandparent has two children





- Gadget network: Wreath
- LineToCompleteBinaryTree subroutine
- LineToCompleteBinaryTree
  - Transforms any initial oriented line into a spanning Complete Binary Tree in log *n* time
  - Activates an edge with grandparent
  - Deactivate an edge with parent
  - Stop when grandparent has two children





15 / 26

- Gadget network: Wreath
- LineToCompleteBinaryTree subroutine
- LineToCompleteBinaryTree
  - Transforms any initial oriented line into a spanning Complete Binary Tree in log *n* time
  - Activates an edge with grandparent
  - Deactivate an edge with parent
  - Stop when grandparent has two children





15 / 26

- Gadget network: Wreath
- LineToCompleteBinaryTree subroutine
- \_ineToCompleteBinaryTree
  - Transforms any initial oriented line into a spanning Complete Binary Tree in log *n* time
  - Activates an edge with grandparent
  - Deactivate an edge with parent
  - Stop when grandparent has two children





- Gadget network: Wreath
- LineToCompleteBinaryTree subroutine

### LineToCompleteBinaryTree

- Transforms any initial oriented line into a spanning Complete Binary Tree in log *n* time
- Activates an edge with grandparent
- Deactivate an edge with parent
- Stop when grandparent has two children





15 / 26

- Gadget network: Wreath
- LineToCompleteBinaryTree subroutine
- LineToCompleteBinaryTree
  - Transforms any initial oriented line into a spanning Complete Binary Tree in log *n* time
  - Activates an edge with grandparent
  - Deactivate an edge with parent
  - Stop when grandparent has two children





15 / 26

- Gadget network: Wreath
- LineToCompleteBinaryTree subroutine

### LineToCompleteBinaryTree

- Transforms any initial oriented line into a spanning Complete Binary Tree in log *n* time
- Activates an edge with grandparent
- Deactivate an edge with parent
- Stop when grandparent has two children







- Gadget network: Wreath
- LineToCompleteBinaryTree subroutine

### LineToCompleteBinaryTree

- Transforms any initial oriented line into a spanning Complete Binary Tree in log *n* time
- Activates an edge with grandparent
- Deactivate an edge with parent
- Stop when grandparent has two children





- Gadget network: Wreath
- LineToCompleteBinaryTree subroutine
- LineToCompleteBinaryTree
  - Transforms any initial oriented line into a spanning Complete Binary Tree in log *n* time
  - Activates an edge with grandparent
  - Deactivate an edge with parent
  - Stop when grandparent has two children





- Gadget network: Wreath
- LineToCompleteBinaryTree subroutine
- LineToCompleteBinaryTree
  - Transforms any initial oriented line into a spanning Complete Binary Tree in log *n* time
  - Activates an edge with grandparent
  - Deactivate an edge with parent
  - Stop when grandparent has two children





- Gadget network: Wreath
- LineToCompleteBinaryTree subroutine

### LineToCompleteBinaryTree

- Transforms any initial oriented line into a spanning Complete Binary Tree in log *n* time
- Activates an edge with grandparent
- Deactivate an edge with parent
- Stop when grandparent has two children



































- Gadget network: ThinWreath
- LineToCompletePolylogarithmicTree subroutine
- Requires knowledge of the size of the network
- Changes in low and high level strategy



- Gadget network: ThinWreath
- LineToCompletePolylogarithmicTree subroutine
- Requires knowledge of the size of the network
- Changes in low and high level strategy



- Gadget network: ThinWreath
- LineToCompletePolylogarithmicTree subroutine
- Requires knowledge of the size of the network
- Changes in low and high level strategy



- Gadget network: ThinWreath
- LineToCompletePolylogarithmicTree subroutine
- Requires knowledge of the size of the network
- Changes in low and high level strategy



- Three Algorithms
- Trade off between time and edge complexity

- Improve time and degree together
- Change the property of the target network
- Lower Bounds



- Three Algorithms
- Trade off between time and edge complexity

- Improve time and degree together
- Change the property of the target network
- Lower Bounds



- Three Algorithms
- Trade off between time and edge complexity

- Improve time and degree together
- Change the property of the target network
- Lower Bounds



- Three Algorithms
- Trade off between time and edge complexity

- Improve time and degree together
- Change the property of the target network
- Lower Bounds



- Three Algorithms
- Trade off between time and edge complexity

- Improve time and degree together
- Change the property of the target network
- Lower Bounds



- Three Algorithms
- Trade off between time and edge complexity

- Improve time and degree together
- Change the property of the target network
- Lower Bounds



- Three Algorithms
- Trade off between time and edge complexity

- Improve time and degree together
- Change the property of the target network
- Lower Bounds

Algorithm	Time	Total Edge Activations	Maximum Activated Edges	Maximum Activated Degree
GraphToStar	$O(\log n)$	$O(n \log n)$	<i>O</i> ( <i>n</i> )	<i>O</i> ( <i>n</i> )
GraphToWreath	$O(\log^2 n)$	$O(n \log^2 n)$	<i>O</i> ( <i>n</i> )	O(1)
${\sf GraphToThinWreath}$	$O(\log^2 n / \log \log n)$	$O(n \log^2 n)$	<i>O</i> ( <i>n</i> )	$O(\log^2 n)$

#### Table: Algorithms for Depth-d Tree

Bounds	Time	Total Edge Activations	Maximum Activated Edges	Maximum Activated Degree
Centralized Lower	$\Omega(\log n)$	$\Omega(n)$	$\Omega(n/\log n)$	
Centralized Upper	$O(\log n)$	$\Theta(n)$		
Distributed Lower	$O(\log n)$	$\Omega(n \log n)$		

#### Table: Bounds for Depth-d Tree

# Thank you!!!