The Cost of Unknown Diameter in Dynamic Networks

Haifeng Yu Yuda Zhao Irvan Jahja

National University of Singapore
Republic of Singapore

National University of Singapore

Example: Leader Election in Static Newtorks

- Example: Elect a leader from N nodes
  - Arbitrary static topology
  - Nodes have unique ids of $\Theta(\log N)$ bits, no failures, randomization allowed
  - Synchronous, each round a node either sends or receives
  - A sending node can send a message with $\Theta(\log N)$ bits, which will be received by all its receiving neighbors
- Smaller D results in smaller time complexity:
  - If D known, can elect a leader in $O(D \log N)$ rounds
  - If D unknown, can first estimate D in $O(D \log N)$ rounds

Background: Dynamic Networks

- Growing interesting in dynamic networks

[Censor-Hillel et al. PODC'11, Cornejo et al. PODC'12, Dutta et al. SPDA'13, Ghaffari et al. PODC'13, Haeupler et al. PODC'11, Haeupler et al. DISC'12, Kuhn et al. PODC'10, Kuhn et al. STOC'10, and etc.]

  - The nodes remain the same in all rounds
  - In each round, an adversary picks an arbitrary connected (undirected) graph as the topology for that round
  - (Dynamic) Diameter D defined as the number of rounds needed for a message to reach all node, when it is flooded from the worst-case node

Background: Dynamic Networks with Unknown Diameters

- Example: Elect a leader from N nodes, which form a ~~static network~~ dynamic network
  - If D is known, can still do $O(D \log N)$ rounds
  - If D is unknown, then ???
  - No existing efficient way to estimate D either…
- Unfortunately, D is often unknown for dynamic networks
  - Most protocols in the literature pessimistically assume $D = N$

Central Question #1

- What is the cost of unknown diameter in dynamic networks?
  - Namely, if a problem's time complexity is α when D is known, and β when D is unknown, what is the gap between α and β?

Our Central Novel Result #1

Cost of unknown diameter in dynamic networks can be exponential for many natural problems (we call them sensitive problems)

Leader-election

Consensus

Confirmed-flooding
A certain node V needs to propagate a token of O(log N) size to all nodes. V terminates once it has confirmed that all nodes have received the token.

Computing various globally-sensitive functions such as MAX and SUM

Our Central Novel Result #1

For all our sensitive problems

- If D known: $O(D\log N)$ rounds time complexity
- If D unknown: $\Omega(\sqrt[4]{N/\log N})$ rounds even when D turns out to be O(1).

First ever such lower bounds – no prior lower bounds.

- For small $D = O(\mathrm{polylog} N)$, this gap is exponential
- Fundamentally, the gap arises because the protocol is forced to worrying about the possibility of a large D…

Central Question #2

- Is there a way to avoid such a cost caused by the lack of the knowledge of the diameter in dynamic networks (e.g., by giving the nodes some other information)?

Our Central Novel Results #2

- For Leader-election and Consensus:
- The cost of unknown diameter can avoided if the protocol knows a good estimate N' of N
  - Let ε be the relative error in N'
  - $\varepsilon \leq \frac{1}{3} - c \Rightarrow$ $O(D\log N)$ rounds sufficient regardless of whether D is known – by our novel upper bound protocol
  - $\varepsilon \geq \frac{1}{3}$ $\Rightarrow$ $O(D\log N)$ rounds if D is known, and $\Omega(\sqrt[4]{N/\log N})$ rounds if D is not known

Roadmap

- Background √
- Summary of our novel results √
- Proof for our $\Omega(\sqrt[4]{N/\log N})$ lower bound in dynamic networks with unknown diameters
- Proof for our $O(D\log N)$ upper bound when a good estimate of N is given
- Conclusion

Overview of Our Lower Bound Proof

- Based on reduction from the DisjCP two-party communication complexity problem

Alice with input $\boldsymbol{X}$ ⟷ Bob with input $\boldsymbol{Y}$

Goal: Compute DisjCP($\boldsymbol{X}$, $\boldsymbol{Y}$) while minimizing communication

- $\Omega(n/q^2) - O(\log n) \leq$ DisjCP
  - [Chen et al. JACM 2014]
- DisjCP $\leq$ Confirmed-flooding
  - Our focus in the next
- Lower bounds on other problems
  - Similar but more complex, see paper for details…

DisjCP ≤ Confirmed-flooding

- DisjCP$_{n,q}$
  - Alice's input $\boldsymbol{X}$ = 02021, Bob's input $\boldsymbol{Y}$ = 11032
  - $\boldsymbol{X}$ and $\boldsymbol{Y}$ must satisfy the cycle promise
  - DisjCP($\boldsymbol{X}$, $\boldsymbol{Y}$) = 0 if exists i where $\boldsymbol{X_i}=\boldsymbol{Y_i}=0$
  - DisjCP($\boldsymbol{X}$, $\boldsymbol{Y}$) = 1 otherwise
- Confirmed-flooding
  - A certain node V needs to propagate a token of O(log N) size to all nodes. V terminates once it has confirmed that all nodes have received the token
- Communication complexity vs time complexity

DisjCP ≤ Confirmed-flooding

- Let $\mathcal{P}$ be any black-box protocol for solving Confirmed-flooding on dynamic networks with unknown diameter

- We construct $\mathcal{Q}$ to solve DisjCP$_{n,q}$ $(\boldsymbol{X},\boldsymbol{Y})$.
  - In $\mathcal{Q}$, Alice and Bob together simulate the execution of $\mathcal{P}$ over a certain dynamic network $\mathcal{G}$
  - $\mathcal{G}$ is a function of $(\boldsymbol{X},\boldsymbol{Y})$
  - Alice and Bob does not know $\mathcal{G}$, but they will simulate the execution of $\mathcal{P}$ over $\mathcal{G}$ – this is the key challenge – but ignore for now, will discuss later..

The network $\mathcal{G}$ when DisjCP($\boldsymbol{X}$,$\boldsymbol{Y}$) = 1

FIRST ROUND

Left part: $\Theta(nq)$ nodes, $O(1)$ diameter

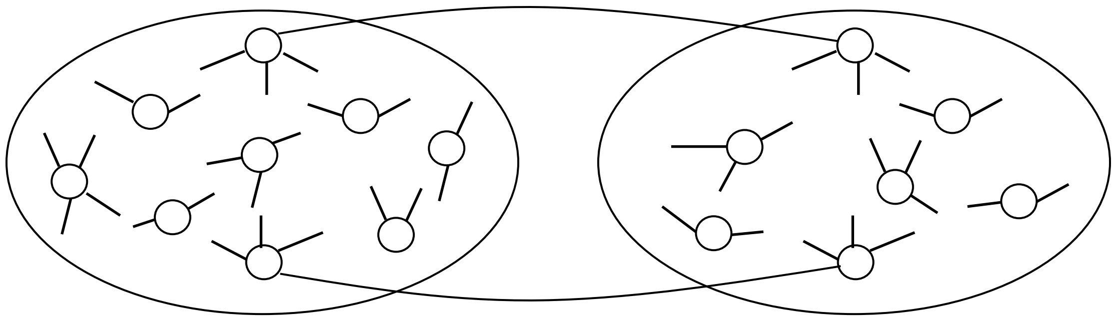

Right part: $\Theta(nq)$ nodes, $O(1)$ diameter

LATER ROUNDS

Each round: Always has $O(1)$ (static) diameter

Across rounds: Always has $O(1)$ dynamic diameter

Assume that $\mathcal{P}$ terminates with s rounds on this $\mathcal{G}$.

The network $\mathcal{G}$ when DisjCP($\boldsymbol{X},\boldsymbol{Y}$) = 0

FIRST ROUND

Exist $\Omega(q)$ special red nodes

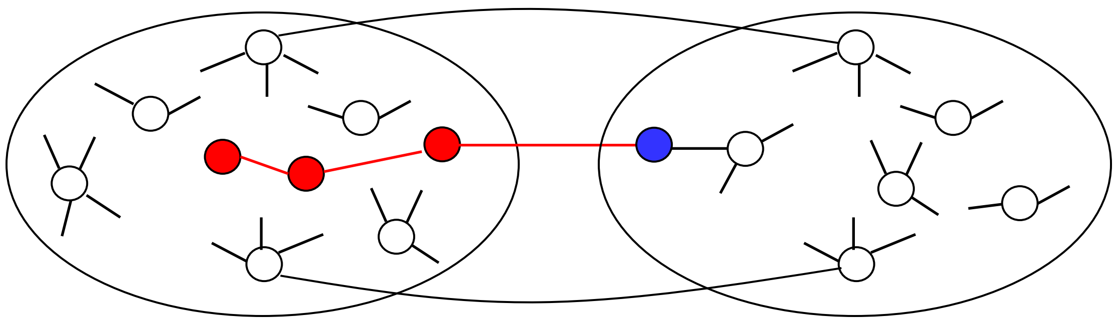

Exists some special blue node

LATER ROUNDS

Each round: Always $\Omega(q)$ (static) diameter

Across rounds: Always $\Omega(q)$ dynamic diameter

Furthermore: Needs $\Omega(q)$ rounds for the blue node to causally affect all other nodes in the right part -- necessary to enable Alice and Bob to later simulate

$\mathcal{P}$ needs $\Omega(q)$ rounds to terminate on this $\mathcal{G}$.

Alice and Bob Solve DisjCP

- DisjCP(X,Y) = 1

$\mathcal{P}$ terminates in s rounds

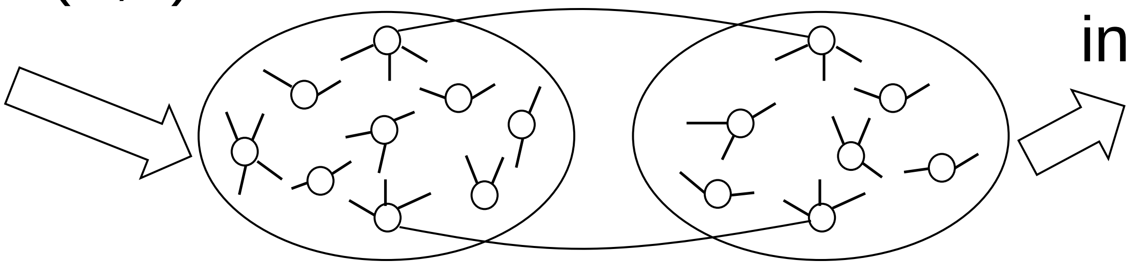

- DisjCP(X,Y) = 0

$\mathcal{P}$ terminates in $\Omega(q)$ rounds

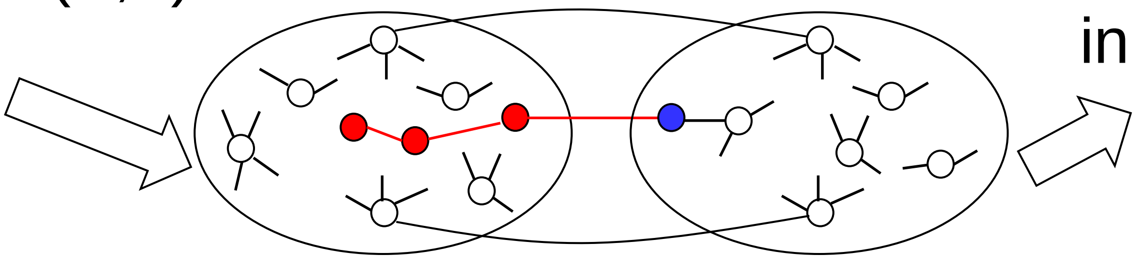

Choose q so that s is smaller than $\Omega(q)$

Alice and Bob Solve DisjCP

- DisjCP($\boldsymbol{X}$,$\boldsymbol{Y}$) = 1

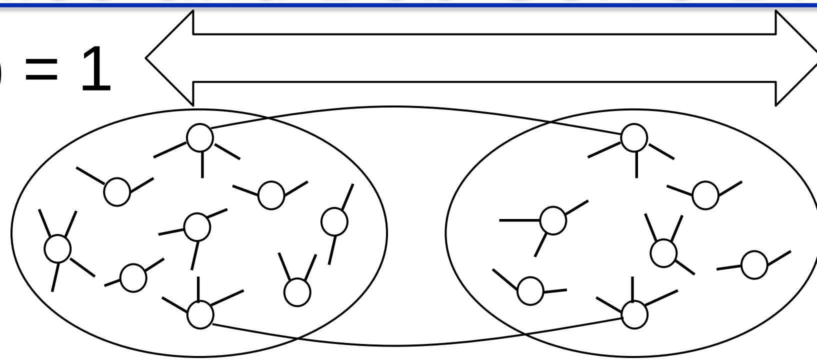

$\mathcal{P}$ terminates in s rounds

- DisjCP($\boldsymbol{X}$,$\boldsymbol{Y}$) = 0

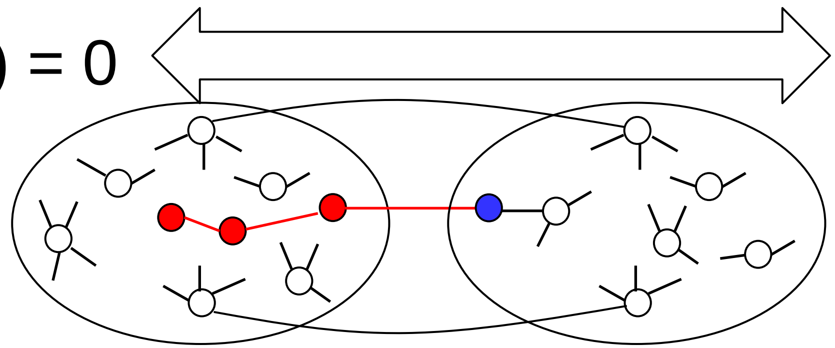

$\mathcal{P}$ terminates in $\Omega(q)$ rounds

- Alice and Bob simulate the black-box $\mathcal{P}$ for Confirmed-flooding:
  - $\mathcal{P}$ terminates in s rounds $\Rightarrow$ DisjCP($\boldsymbol{X}$,$\boldsymbol{Y}$) = 1
  - $\mathcal{P}$ does not terminate in s rounds $\Rightarrow$ DisjCP($\boldsymbol{X}$,$\boldsymbol{Y}$) = 0

From communication complexity to time complexity

- When simulating $\mathcal{P}$, Alice and Bob needs to communicate
  - To simulate for s rounds, needs $O(s\log N)$ bits of communication
- Existing lower bound on DisjCP:
  - Solving DisjCP needs $\Omega(n/q^2) - O(\log n)$ bits
- Solving $O(s\log N) = \Omega(n/q^2) - O(\log n)$ gives the lower bound on s
  - Recall that s is the number of rounds needed by $\mathcal{P}$ on networks with O(1) diameter

The devil is in the details.

- Key Challenge: Alice and Bob does not know G, but they will simulate the execution of P over G
  - Let U be the set of nodes simulated by Alice and V be the set of the nodes simulated by Bob

- We employs a variety of involved techniques:
  - Leverage the cycle promise
  - U and V may change over time
  - Union of U and V may not cover all nodes
  - U and V may intersect : Alice and Bob may disagree on their "view" of G

Roadmap

- Background √
- Summary of our novel results √
- Proof for our $\Omega(\sqrt[4]{N/\log N})$ lower bound in dynamic networks with unknown diameters √
- Proof for our $O(D\log N)$ upper bound when a good estimate of N is given
- Conclusion

Upper bound leveraging a good estimate of N

- Theorem: If an estimate N' for N is given to the protocol where the relative error is less than $\frac{1}{3} - c$, then we can solve Consensus and Leader-election within $O(D \log N)$ round, even if D is unknown

- High-level idea:
  - Keep guessing D
  - Try to lock a majority – roll back if unsuccessful
  - Use counting with one-sided error, together with N', to determine majority
  - See details in the paper…

Conclusions

- In dynamic networks, many common problems are sensitive to the knowledge of the diameter
  - Leader-election, Consensus, globally-sensitive functions, Confirmed-flooding
  - If D known: $O(D\log N)$ rounds time complexity
  - If D unknown: $\Omega(\sqrt[4]{N/\log N})$ rounds even when D turns out to be O(1)
- For some problems, the knowledge of D can be replaced by a good estimate of the system size
  - Leader-election, Consensus
  - $O(D\log N)$ rounds regardless of whether D is known