Brief Announcement: Highly Dynamic and Fully Distributed Data Structures

John Augustine

augustine@cse.iitm.ac.in

Antonio Cruciani

antonio.cruciani@aalto.fi

A!

Aalto University
School of Science

Iqra Altaf Gillani

iqraaltaf@nitsri.ac.in

Peer to Peer Networks

Prevailing Definition: A network of peers, ideally fully decentralized

Key Challenges:

- Highly dynamic
- High **churn**

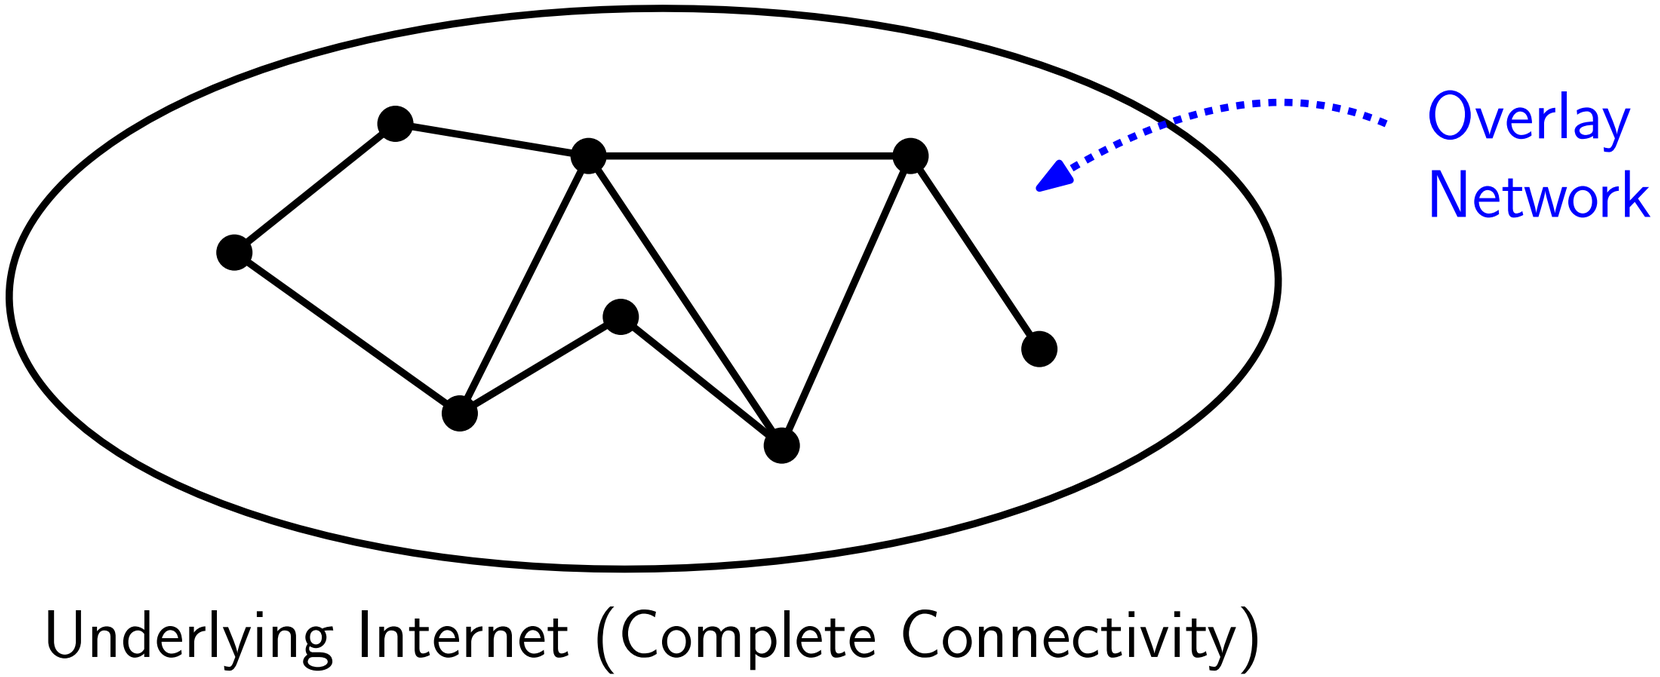

Model: Dynamic Network with Churn (DNC)

Synchronous: All nodes follows the same clock. In each round $r = 1, 2, 3, \ldots$

- each node sends/receives polylog(n) messages per round
- the message size is polylog(n)
- nodes perform local computations

Adversarial Dynamism:

An **oblivious** adversary (knows the algorithm but not the coin toss outcomes) designs the churn

$$\mathcal{G} = \left(G^0, G^1, \ldots, G^r, \ldots\right)$$

Our Problem

Goal: Given a churn rate of up to $\mathcal{O}(n/\log n)$ per round, maintain a distributed *approximate* data structure *efficiently*

Data Structure: Skip List

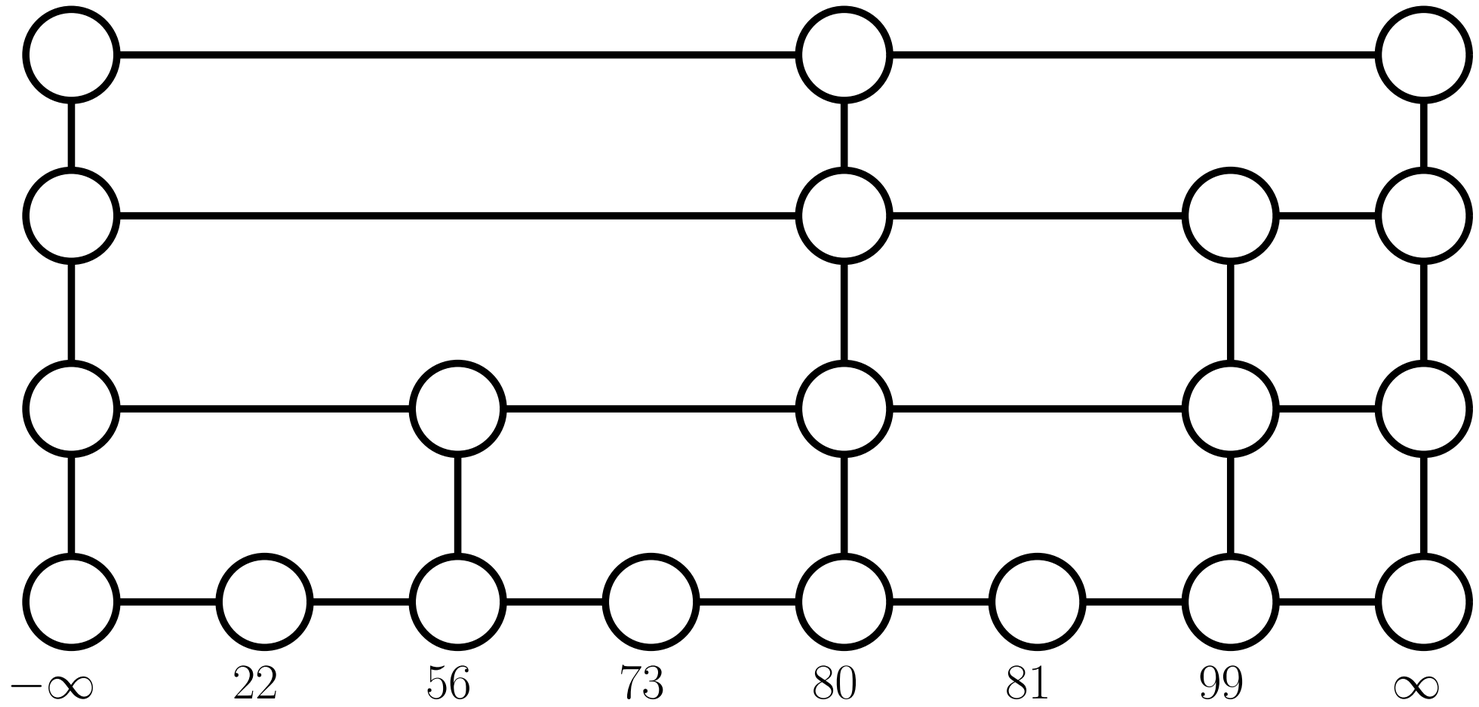

Dynamic Resource Competitiveness

$$\underbrace{\#CreatedEdges + \#Messages}_{Workload} \leq \beta \cdot Churn(t_{start} - \alpha, t_{end})$$

$\alpha = \mathcal{O}(\log n)$

$\beta =$ polylog(n)

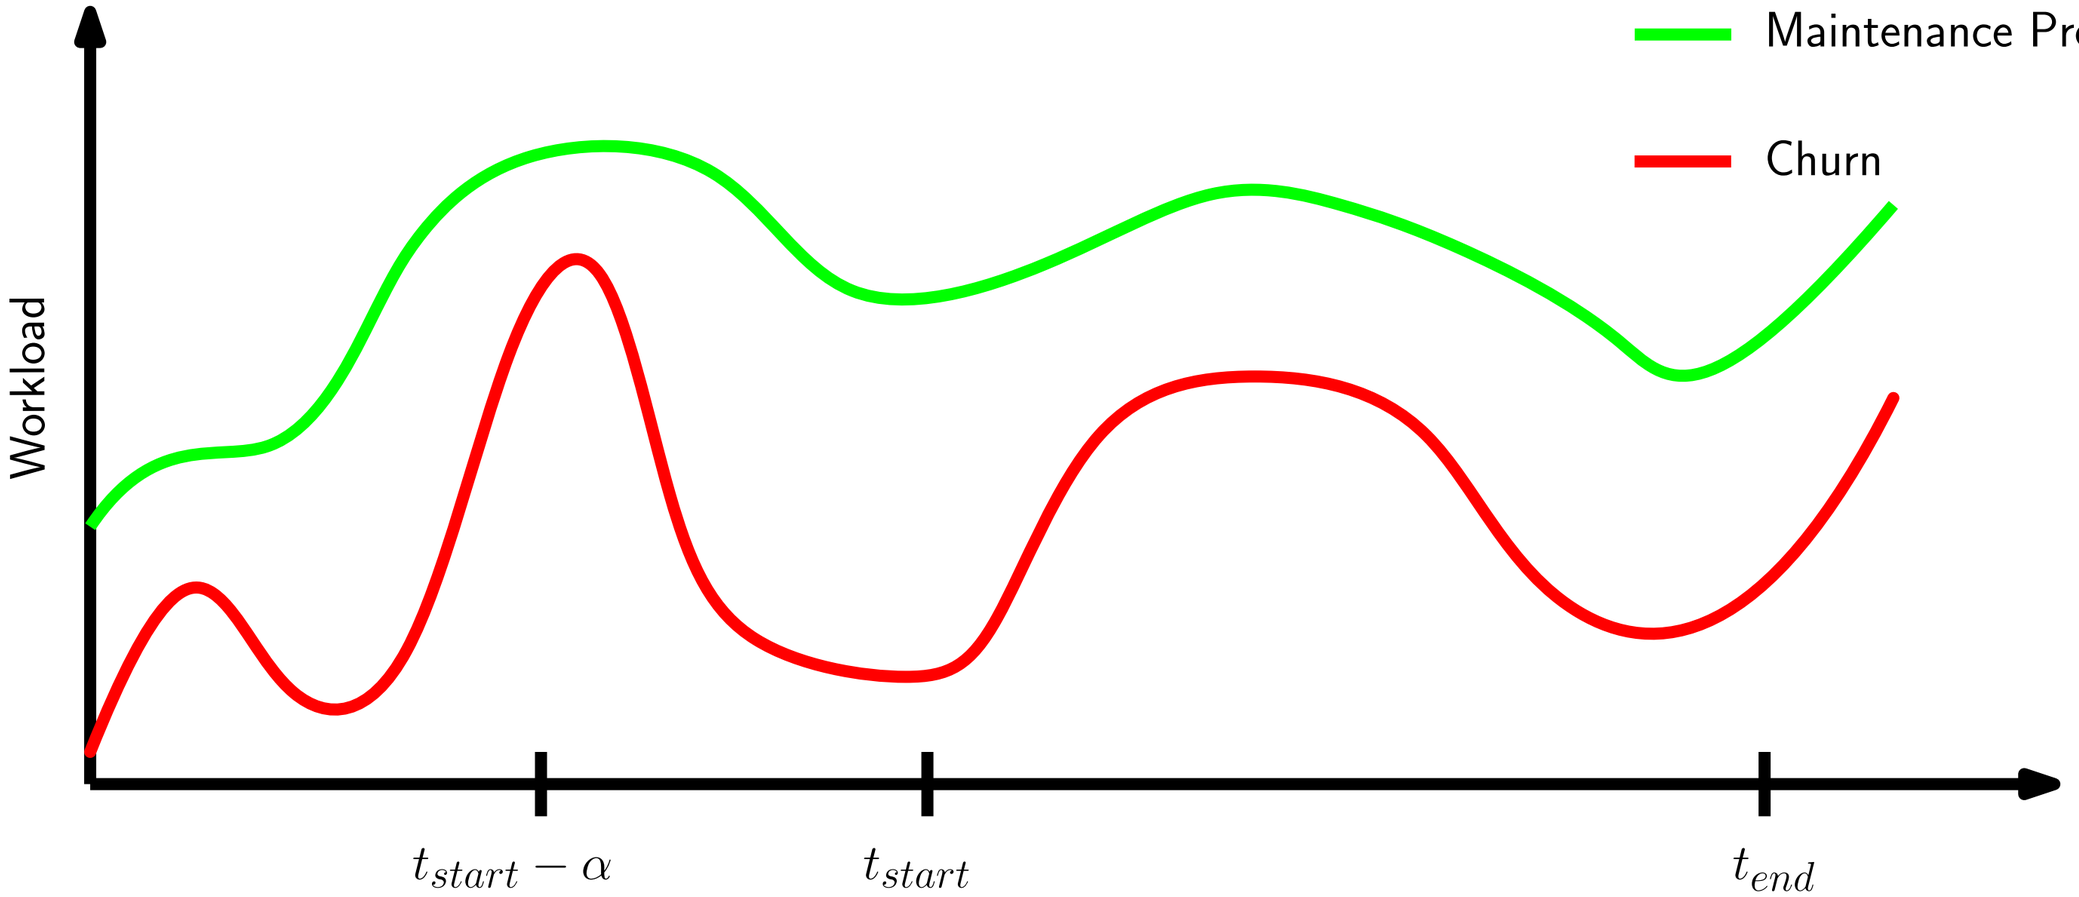

Overview of the maintenance phase

$\mathcal{O}(\log n)$ rounds Bootstrap Phase

Algorithm initialization

Adversary wakes up

Churn rate of up to $\mathcal{O}(n/\log n)$ per round

Maintenance Phase

We need to cope with the churn

The Overlay Network

Commitee

Clique of $\Theta(\log n)$ random nodes

Nodes *periodically* change their committee (randomly)

Cannot be destroyed whp

SPARTAN : Overlay Network ([Augustine et al. J.P.D.C. 2021]

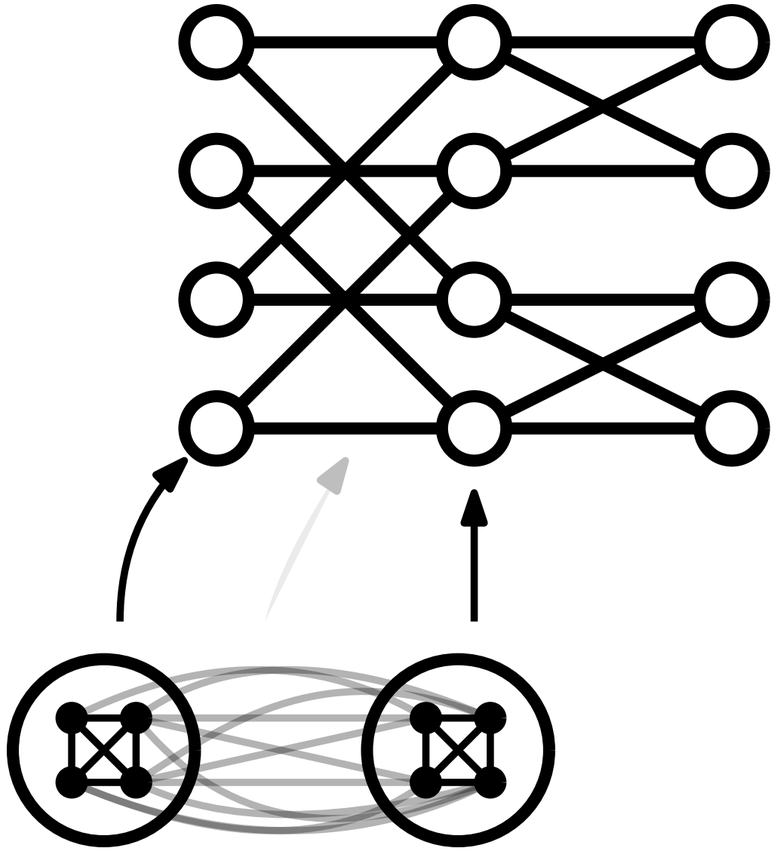

The Maintenance Cycle

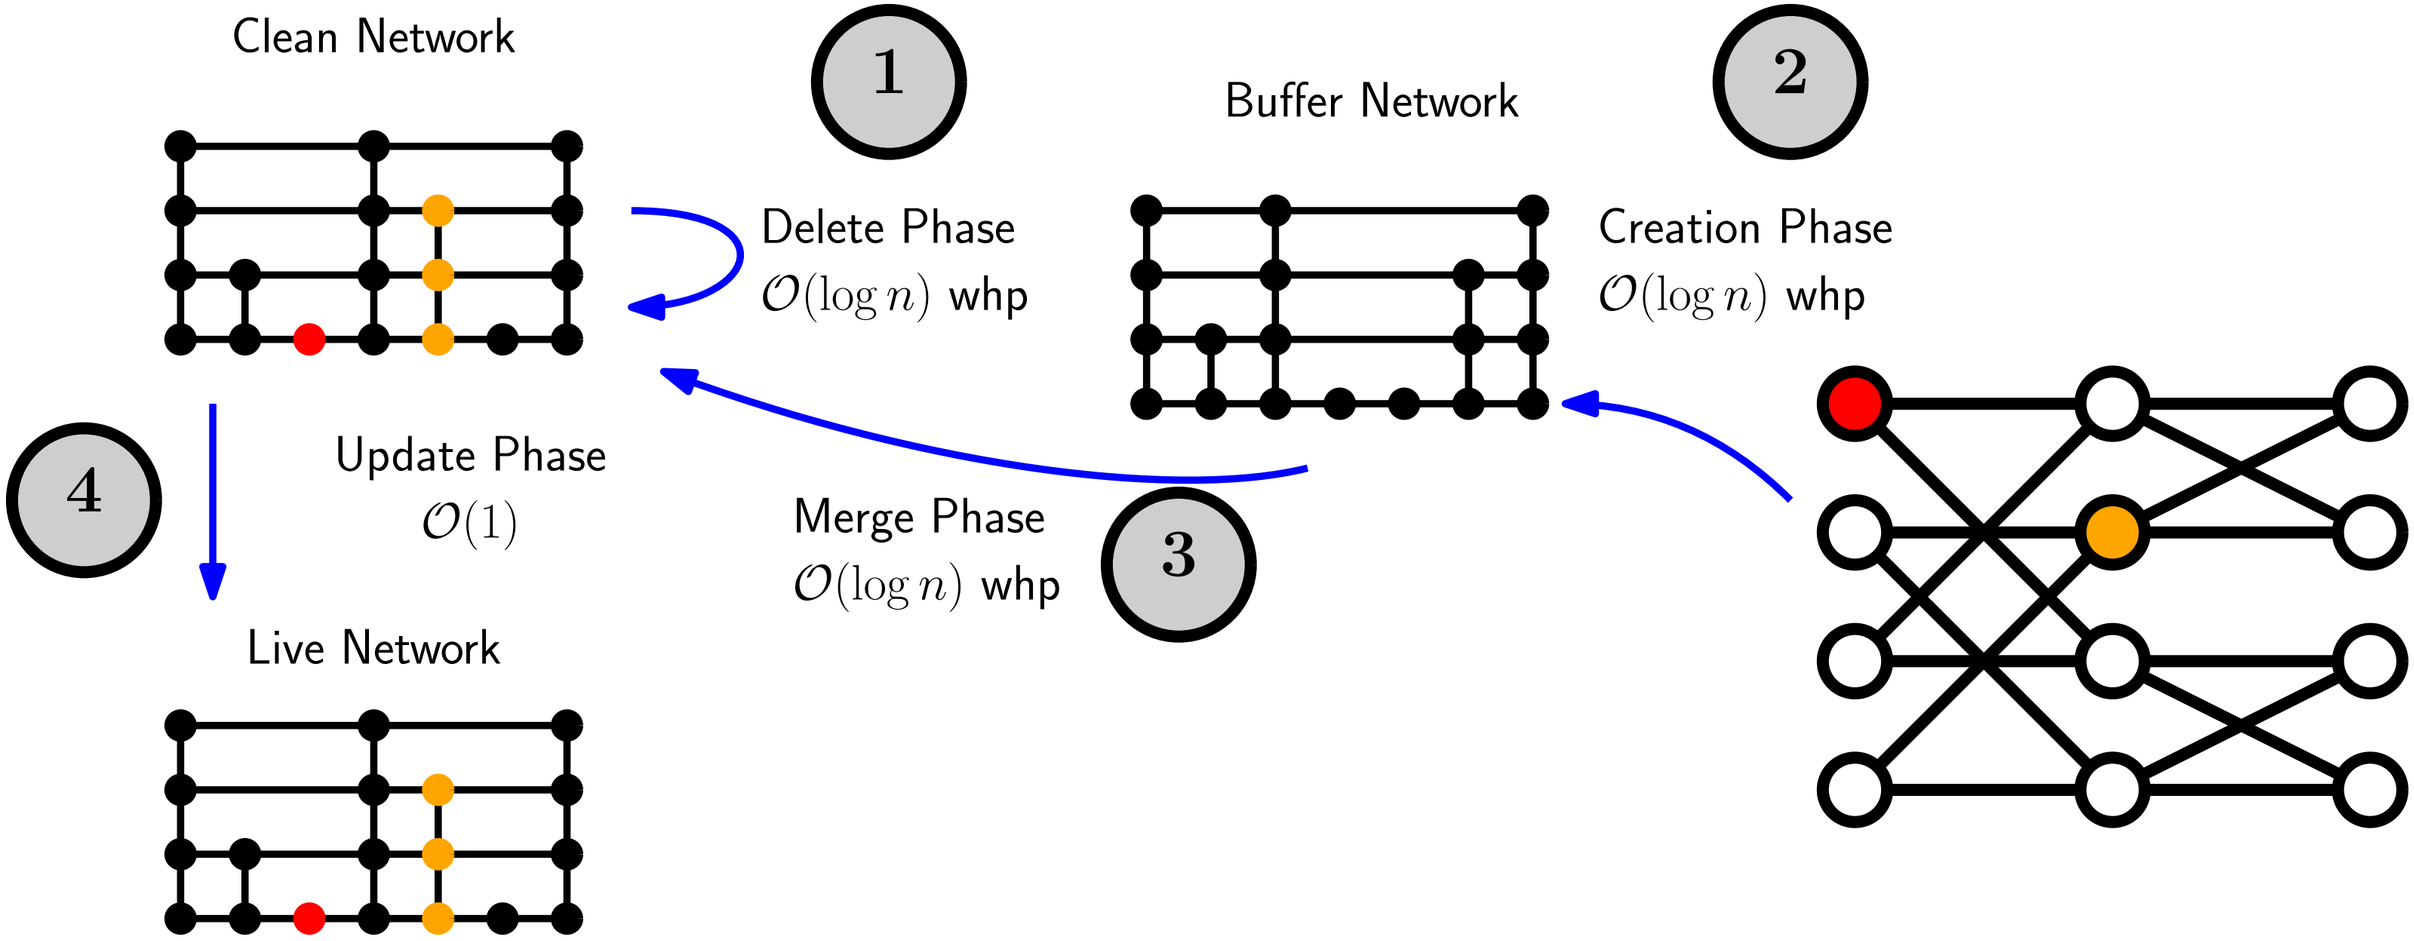

Overview of the results

Theorem 1 *A distributed skip-list (or skip-graph) can be maintained against a churn rate of* $\mathcal{O}(n/\log n)$.

Corollary 1 *Each node can hold* $k =$*polylog*(n) *elements of the data structure.*

Theorem 2 *Any data structure* $\mathcal{D}$ *that admits a* T*-rounds maintenance cycle can be maintained using our framework against a churn rate of* $\mathcal{O}(n/T)$

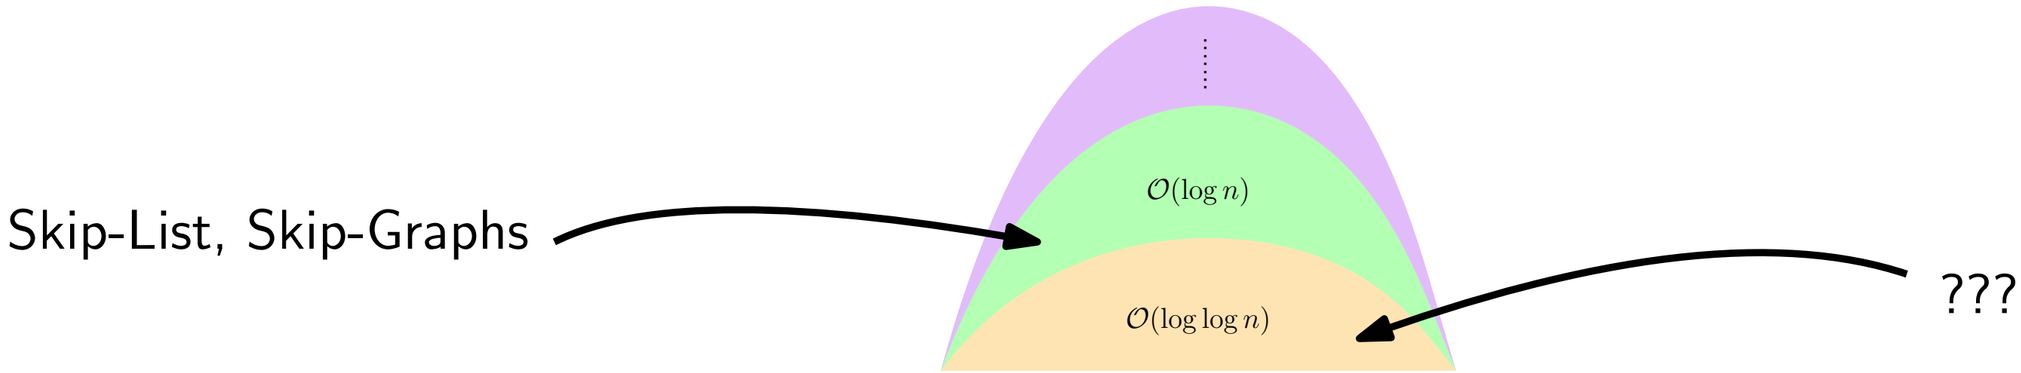

Thank You

The Live and Clean Networks

Clean Network

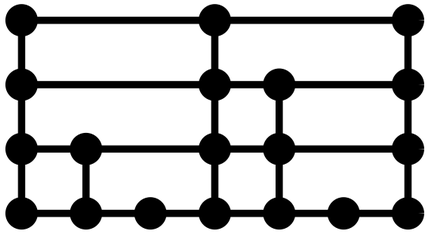

Live Network

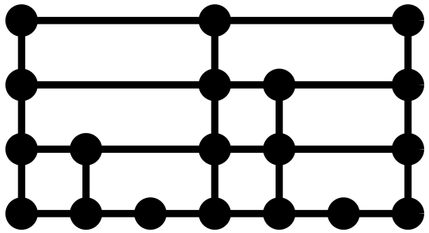

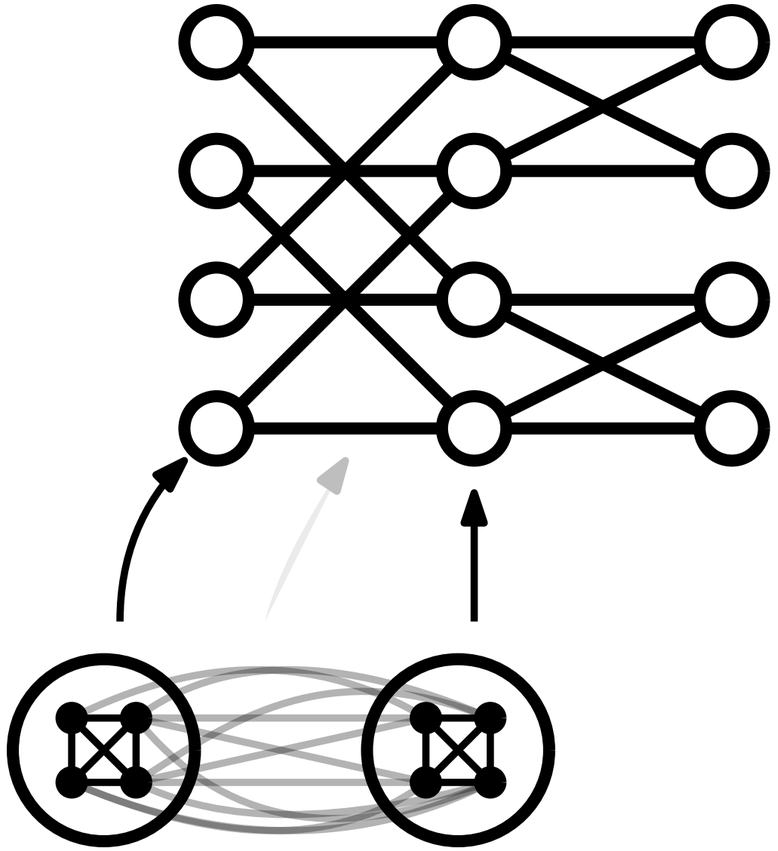

Covering for Churned Nodes

Clean Network

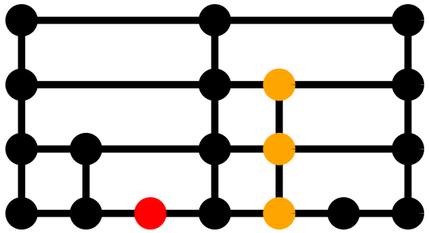

Live Network

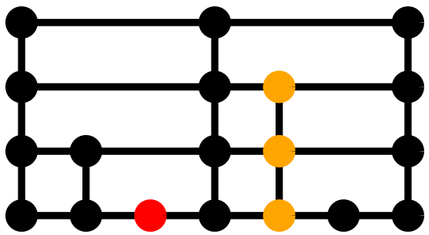

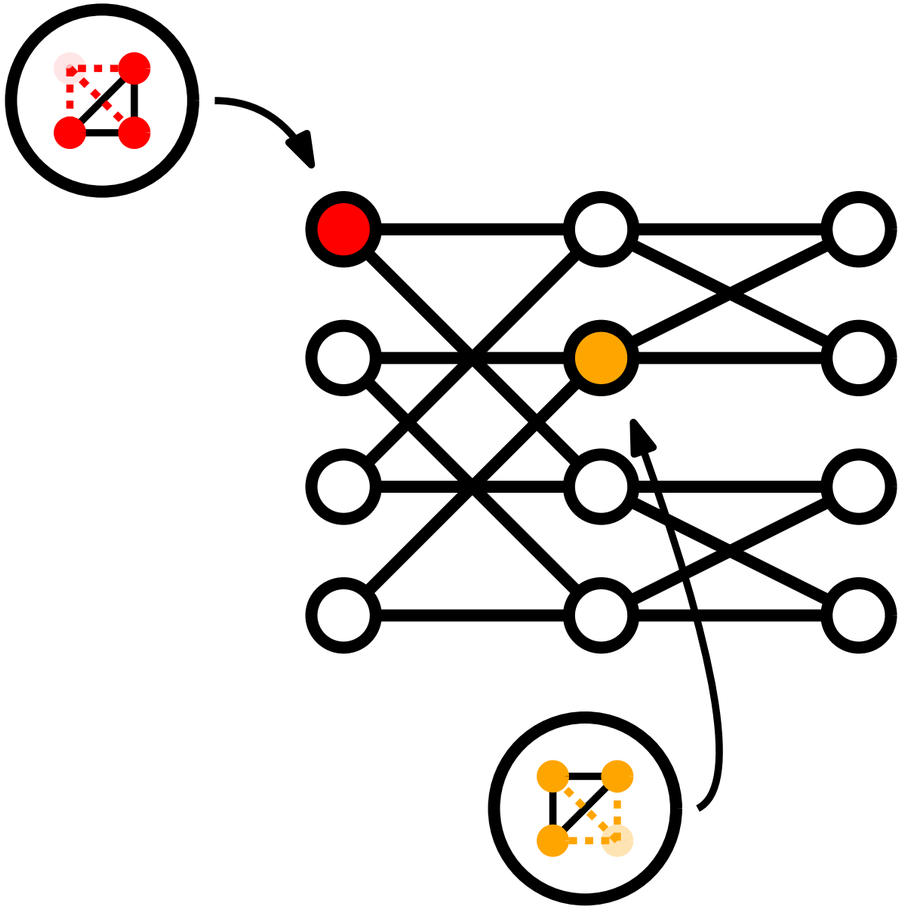

The Maintenance Cycle

Clean Network

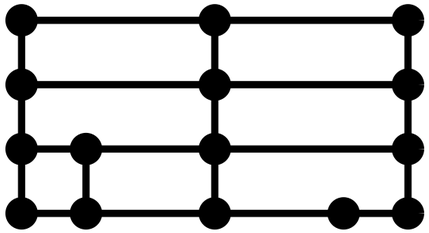

Delete Phase
$\mathcal{O}(\log n)$ whp

Live Network

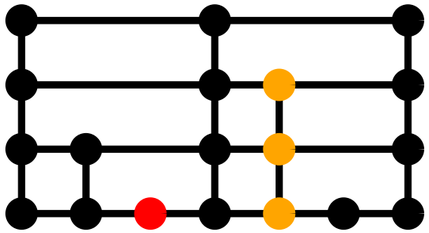

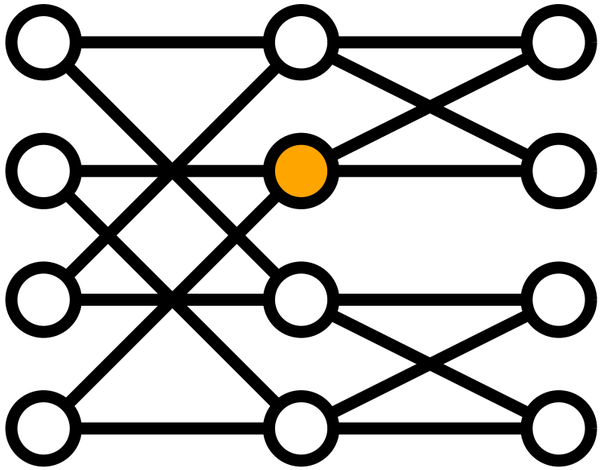

The Maintenance Cycle

Clean Network

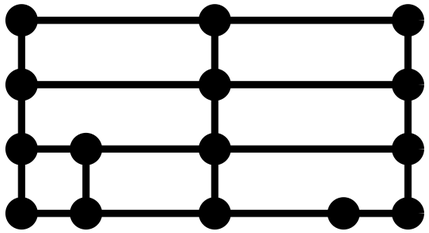

Buffer Network

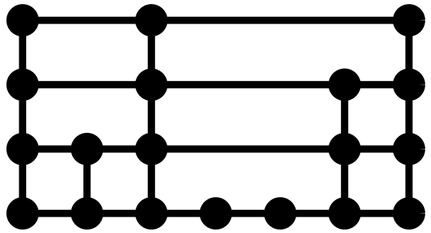

Creation Phase
$\mathcal{O}(\log n)$ whp

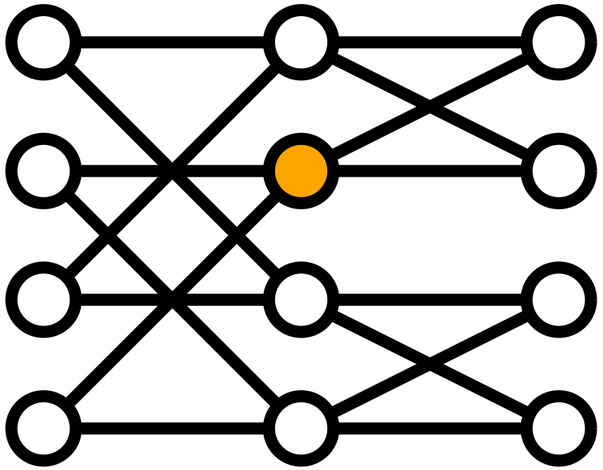

Live Network

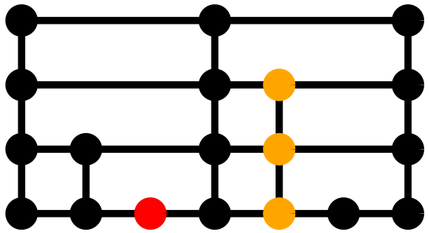

The Maintenance Cycle

Clean Network

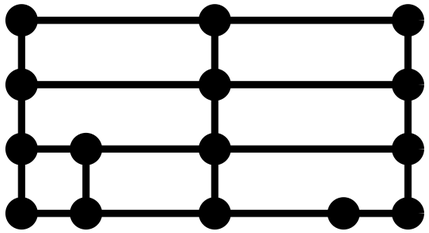

Merge Phase

$\mathcal{O}(\log n)$ whp

Buffer Network

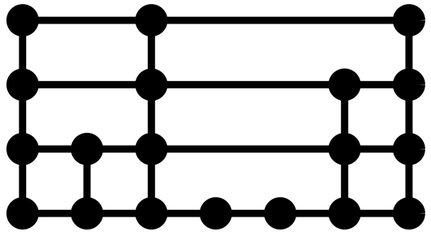

Live Network

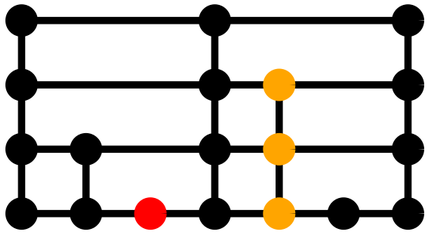

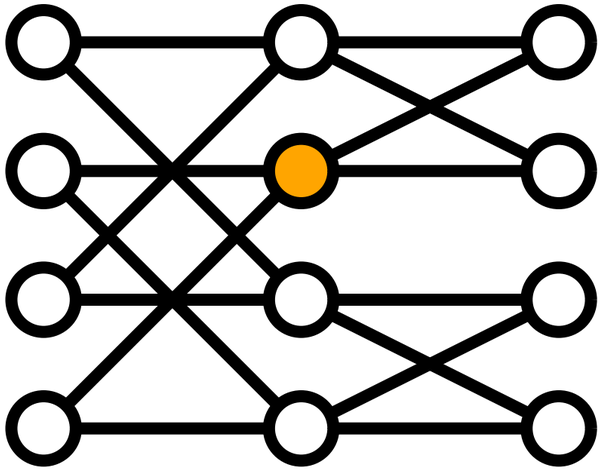

The Maintenance Cycle

Clean Network

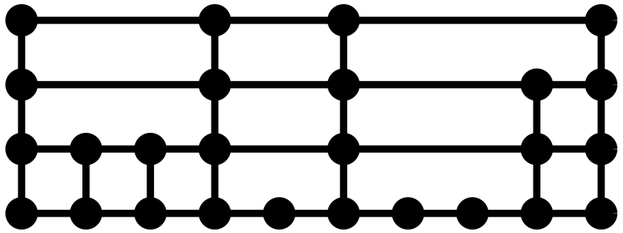

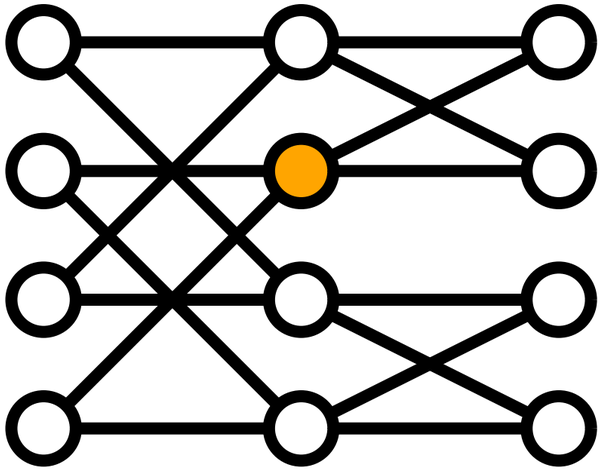

Live Network

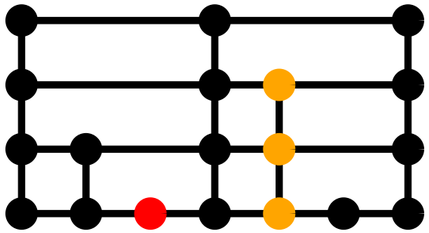

The Maintenance Cycle

Clean Network

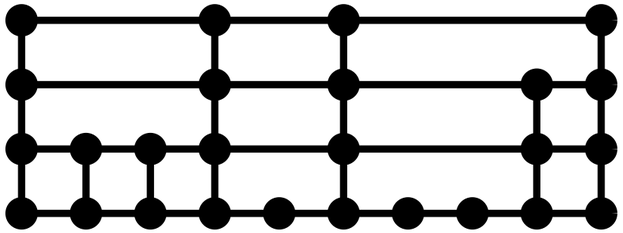

Update Phase

$\mathcal{O}(1)$

Live Network