Dynamic Networks Foundations

Paul G. Spirakis

joint work with
Othon Michail
Ioannis Chatzigiannakis

Computer Technology Institute & Press "Diophantus" (CTI)
Dept. of Computer Engineering & Informatics (CEID), Univ. of Patras

7th Athens Colloquium on Algorithms and Complexity (ACAC)
August 27-28, 2012
UoA, Athens, Greece

INTRODUCTION

Motivation

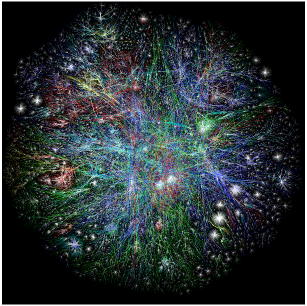

- Many natural and artificial networks are inherently dynamic
- e.g. animal, social, transportation, mobile (e.g. robot, sensor), wireless networks ...
- Mobile Internet
  - Almost 50% UK internet users are going online via mobile phone data connections, according to the UK Office for National Statistics.
- Figures: Internet and Global Airline Networks

Existing Work

- Population Protocols: strong global fairness, limited memory, stabilization

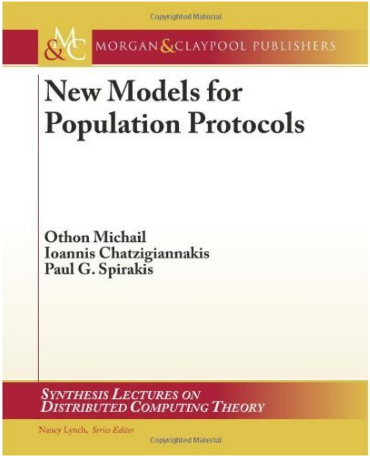

Existing Work

- Fault tolerance: dynamicity eventually ceases, restricted number, frequency, duration of network changes
- Dynamic Overlay Networks: nodes choose an overlay communication network themselves (e.g. by gossip), churn is typical
- Geometric and Random Mobility: unit disk graph model, birth-death edge processes, topology changes according to particular probability distributions

The Worst-case Perspective

- The network may change arbitrarily from time to time
- Adversary scheduler: controls the topology
- Nodes do not control the topology
- Necessary restriction: information must eventually spread
- Advantage: results hold for all possible dynamicity/mobility patterns and not just specific cases or distributions

THE MODEL

Dynamic Graph Model

- a.k.a. temporal or time-varying
- Each edge has a set of time-labels indicating availability times

Definition (Dynamic Graph)

A dynamic graph G is a pair (V, E), where V is a set of n nodes and $E : \mathbb{N}_{\geq 1} \rightarrow \mathcal{P}(\{\{u, v\} : u, v \in V\})$ is a function mapping a round number r to a set $E(r)$ of bidirectional links.

- An interesting result:
  - There is no analogue of Menger's theorem for arbitrary temporal networks [KKK00]

Menger's Theorem Violated

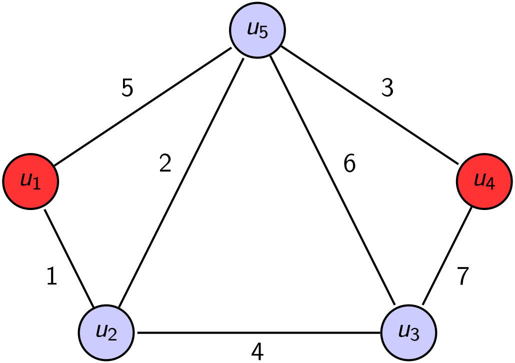

- Menger's Theorem: The max number of node-disjoint s-t paths is equal to the min number of nodes needed to separate s from t
- We now care for temporal paths (strictly increasing time-labels, a.k.a. journeys)
- There are no 2 disjoint temporal paths from u_1 to u_4 but
- After deleting any node (other than u_1 or u_4) there still remains a temporal u_1-u_4 path

Further Modeling Assumptions

- Set V of n nodes/processors
- Unlimited local storage
- Usually unique ids of size $O(\log n)$ bits
- Synchronous message passing
  - Discrete steps/rounds
  - Global clock available to the nodes
  - Communication via sending/receiving messages
- 2 types of message transmission
  1. Broadcast
  2. One-to-each

A Round

1. The adversary chooses the edges for the round
   - It can see the internal states of the nodes at the beginning of the round
2. At the same time and independently of the adversary's choice of edges each node generates its message(s) for the current round
   - No info about the internal state of neighbors when generating messages
   - Deterministic algorithms generate messages based solely on the internal state: the adversary can infer the messages
3. Messages are delivered to the sender's neighbors, as chosen by the adversary
4. The next round begins

TOOLS & METRICS

Causal Influence [Lam78]

- Crucial notion
- (u, r): the state of node u at time/round r
- $(u, r) \rightarrow (v, r+1)$ iff $u = v$ or $\{u, v\} \in E(r+1)$
- Causal order $\rightsquigarrow \subseteq (V \times \mathbb{N}_{\geq 0})^2$: the reflexive and transitive closure of $\rightarrow$
- $(u, r) \rightsquigarrow (v, r')$: node u's state in round r influences node v's state in round r'
  - u "influences" v through a chain of messages originating at u and ending at v (possibly going through other nodes in between)

2 Useful Sets

1. Past set of a time-node (u, t') from time t
   - $\mathrm{past}_{(u,t')}(t) := \{v \in V : (v, t) \rightsquigarrow (u, t')\}$
   - set of nodes whose t-state has causally influenced the t'-state of u
2. Future set of a time-node (u, t) at time t'
   - $\mathrm{future}_{(u,t)}(t') := \{v \in V : (u, t) \rightsquigarrow (v, t')\}$
   - set of nodes whose t'-state has been causally influenced by the t-state of u

Dynamic Diameter

- Classical diameter is not suitable for dynamic networks
  - A star graph can be made to have dynamic diameter $n-1$ while its diameter is just 2 [AKL08]

Dynamic diameter

- Upper bound on flooding time (time required for each node to causally influence every other node)
- Minimum $D \in \mathbb{N}$ s.t. for all times $t \geq 0$ and all $u, v \in V$ it holds that $(u, t) \rightsquigarrow (v, t + D)$
- Small dynamic diameter allows for fast dissemination of information
- Nodes do not know the dynamic diameter
- We shall allow minimal knowledge based on which nodes may infer upper bounds on the dynamic diameter

Dynamic Diameter

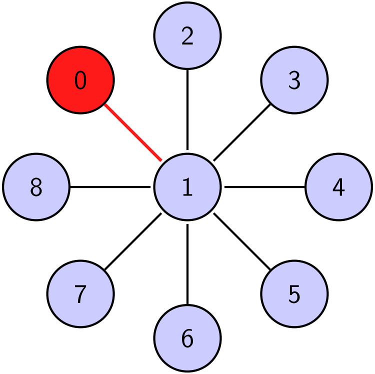

Dynamic Diameter

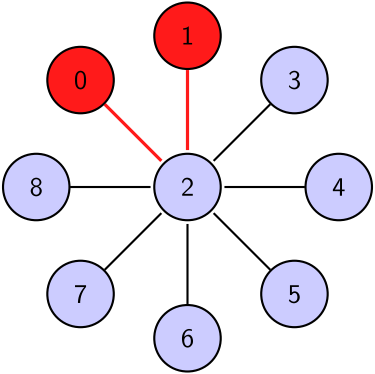

Dynamic Diameter

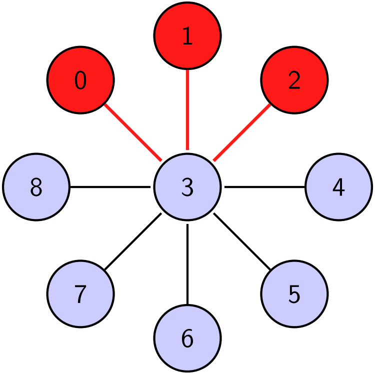

Dynamic Diameter

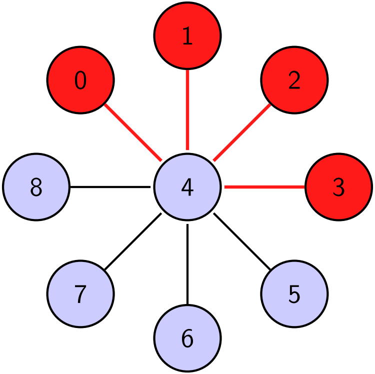

Dynamic Diameter

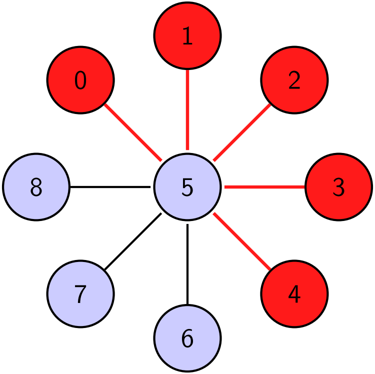

Dynamic Diameter

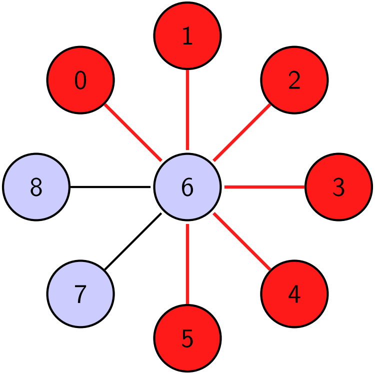

Dynamic Diameter

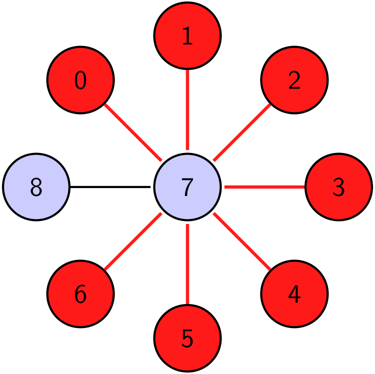

Dynamic Diameter

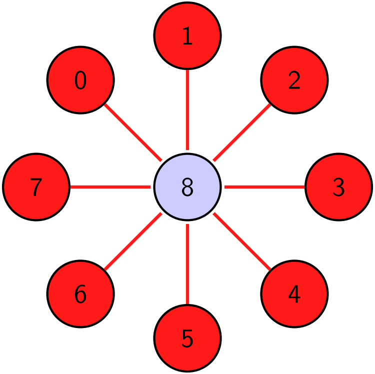

Dynamic Diameter

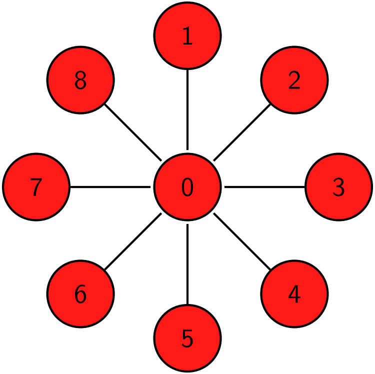

Dynamic Diameter

- Diameter = 2

while on the other hand . . .

- Dynamic diameter = 8 (i.e. $n - 1$)
  - Node 8 first "hears of" node 0 in round 8

T-interval Connected Dynamic Graphs [KLO10]

- Represent dynamic networks that are connected at every instant
- T represents the rate of connectivity changes

Definition

A dynamic graph $G = (V, E)$ is said to be T-interval connected, for $T \geq 1$, if, for all $r \in \mathbb{N}$, the static graph $G_{r,T} := (V, \bigcap_{i=r}^{r+T-1} E(r))$ is connected.

For example

- In 1-interval connected the underlying connected spanning subgraph may change arbitrarily from round to round
- In ∞-interval connected a connected spanning subgraph is preserved forever

T-interval Connected Dynamic Graphs

- Allow for constant propagation of information
- There is always an edge in every cut

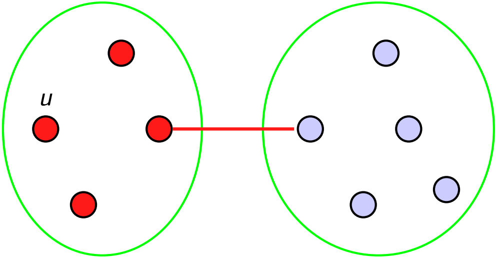

Nodes that have heard of u

Nodes that have not heard of u

T-interval Connected Dynamic Graphs

Lemma ([KLO10])

For any node $u \in V$ and time $r \geq 0$ we have

1. $|\{v \in V : (u, 0) \rightsquigarrow (v, r)\}| \geq \min\{r + 1, n\}$,
2. $|\{v \in V : (v, 0) \rightsquigarrow (u, r)\}| \geq \min\{r + 1, n\}$.

- The dynamic diameter is small
  - At most linear in n
- For all times $t \geq 0$ the t-state of any node influences the $(t + n - 1)$-state of every other node

Instantaneous Connectivity: State-of-the-art

- The idea first appears in [OW05] in an asynchronous setting
- They studied flooding and routing
- Flooding was solved in $O(Tn^2)$ rounds using $O(\log n)$ bit storage and message overhead
  - T is the maximum time it takes to transmit a message
- Routing was solved in $O(Tn)$ rounds using $O(\log n)$ bit storage and message overhead

Instantaneous Connectivity: State-of-the-art

- T-interval connectivity was proposed in [KLO10]
- The synchronous case was studied for the first time
- Counting and all-to-all token dissemination were solved in
  1. $O(n)$ rounds using $O(n \log n)$ bits per message
  2. $O(n^2/T)$ rounds using $O(\log n)$ bits per message
- They also gave the following lower bound:
  - Any deterministic centralized algorithm for k-token dissemination in 1-interval connected graphs requires at least $\Omega(n \log k)$ rounds to complete in the worst case

Possibly Disconnected Dynamic Networks

Not all dynamic networks have connected instances

- Most natural dynamic networks are only temporally connected
- There are dynamic networks with always disconnected instances in which information spreads as fast as in those with always connected instances
- No results were known for this type of worst-case dynamic networks
- Note: All results that follow are from [MCS12a] and [MCS12b]

Metrics for Disconnectivity

1. Outgoing Influence Time (oit)
2. Incoming Influence Time (iit)
3. Connectivity Time (ct)

oit

- Maximal time until the state of a node influences the state of another node
- Minimum $k \in \mathbb{N}$ s.t. for all $u \in V$ and all times $t, t' \geq 0$ s.t. $t' \geq t$ it holds that

$$|\mathrm{future}_{(u,t)}(t'+k)| \geq \min\{|\mathrm{future}_{(u,t)}(t')| + 1, n\}$$

- Example: the oit of a T-interval connected graph is 1
- If a dynamic graph $G = (V, E)$ has oit 1 then every instance has at least $\lceil n/2 \rceil$ edges.

Alternating Matchings

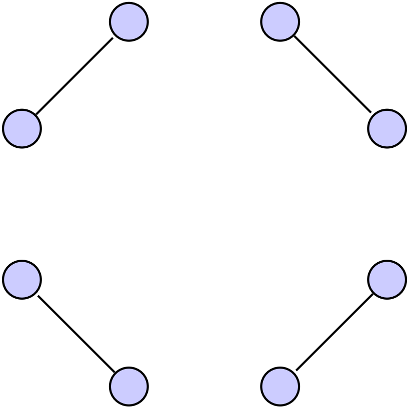

Alternating Matchings

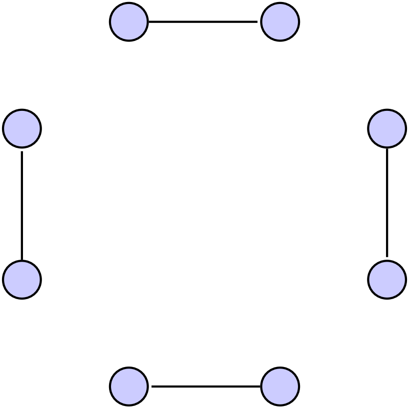

Alternating Matchings

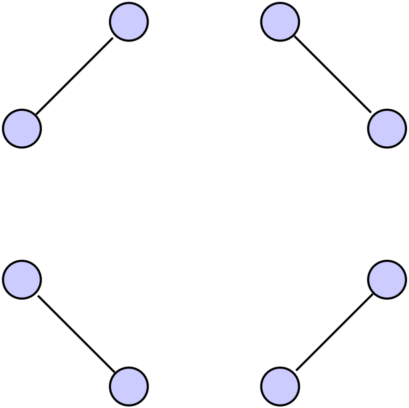

Alternating Matchings

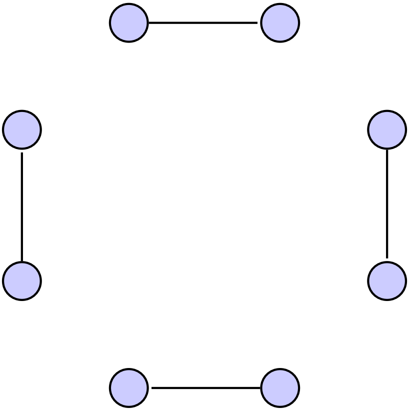

Soifer's Dynamic Graph

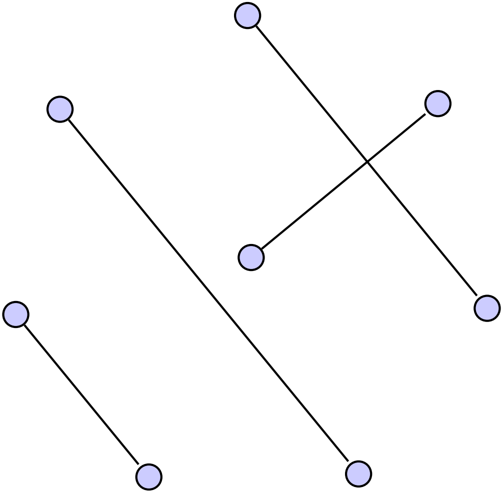

Soifer's Dynamic Graph

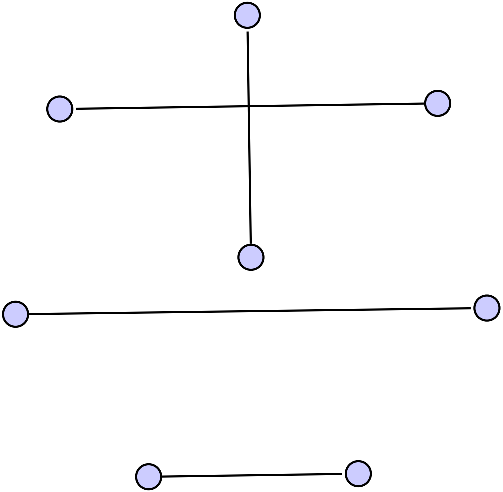

Soifer's Dynamic Graph

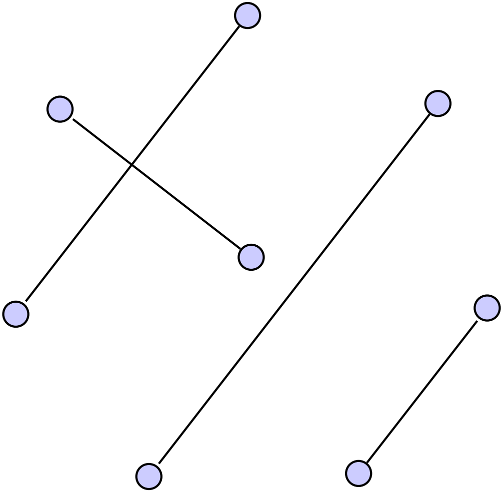

Soifer's Dynamic Graph

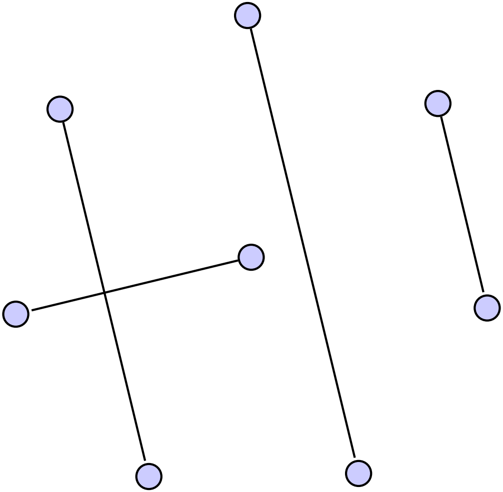

Soifer's Dynamic Graph

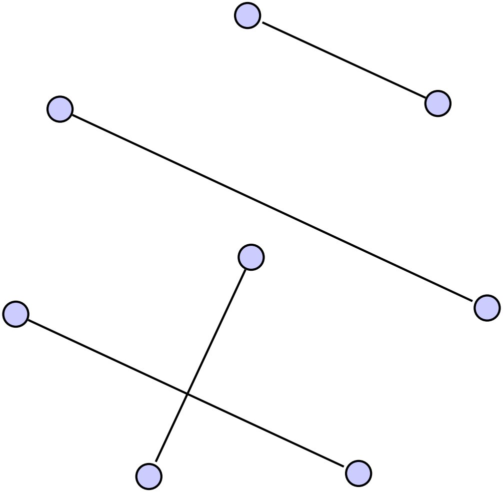

Soifer's Dynamic Graph

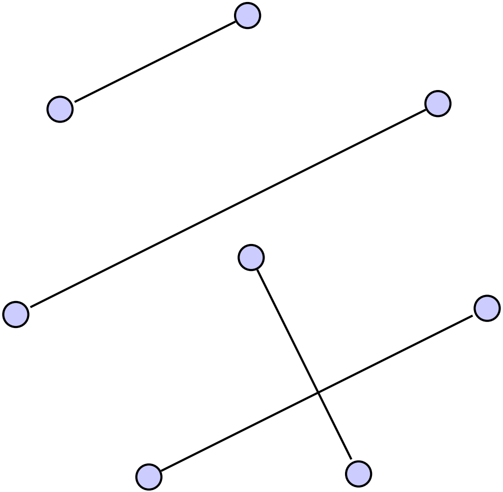

Soifer's Dynamic Graph

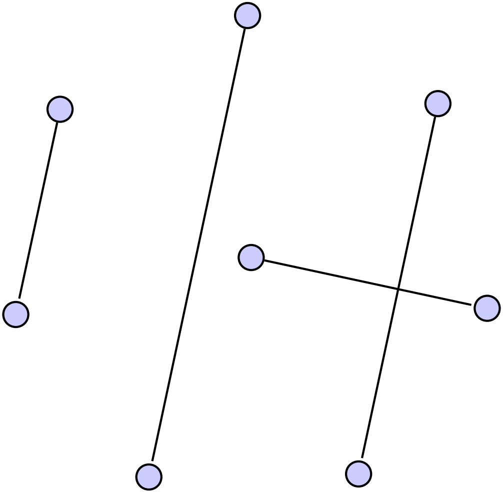

Soifer's Dynamic Graph

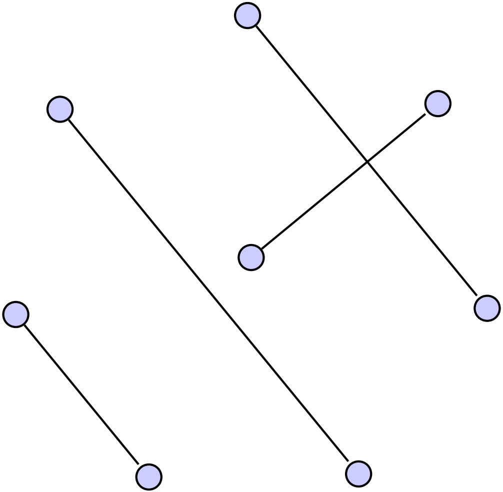

Possibly Disconnected Dynamic Networks

- Both examples
  - have disconnected instances
  - but have unit oit (thus, dynamic diameter linear in n)
- In Soifer's dynamic graph edges take maximal time to reappear

iit

- Maximal time until the state of a node is influenced by the state of another node
- Minimum $k \in \mathbb{N}$ s.t. for all $u \in V$ and all times $t, t' \geq 0$ s.t. $t' \geq t$ it holds that
$$|\mathrm{past}_{(u,t'+k)}(t)| \geq \min\{|\mathrm{past}_{(u,t')}(t)| + 1, n\}$$
- Example: the iit of a T-interval connected graph can be up to $n - 2$

ct

- Maximal time until the two parts of any cut of the network become connected
- Minimum $k \in \mathbb{N}$ s.t. for all times $t \in \mathbb{N}$ the static graph $(V, \bigcup_{i=t}^{t+k-1} E(i))$ is connected
- If the ct is 1 then we obtain a 1-interval connected graph
- Greater ct allows for disconnected instances

Alternating Matchings (ct=2)

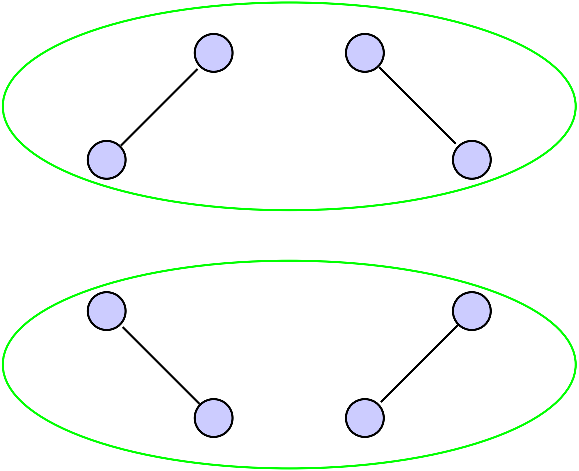

Alternating Matchings (ct=2)

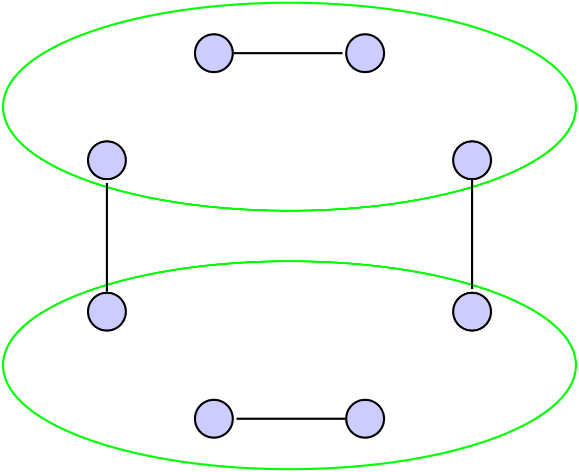

Alternating Matchings (ct=2)

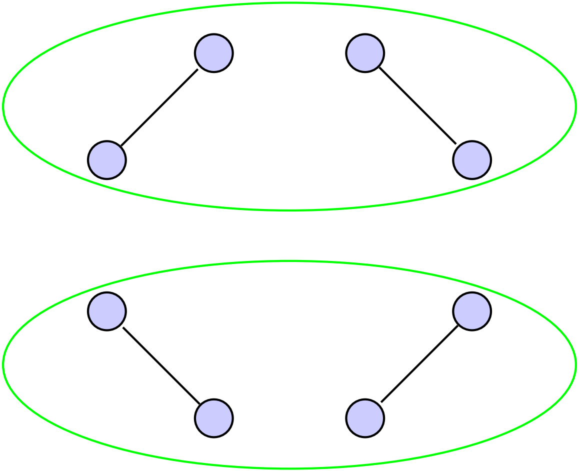

Alternating Matchings (ct=2)

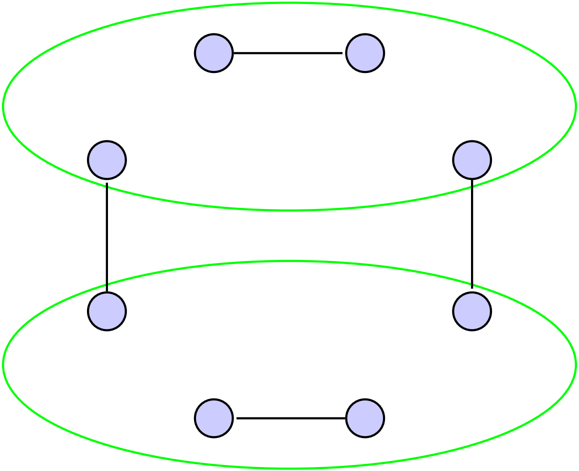

oit vs ct

Proposition

1. *oit* $\leq$ *ct but*
2. *there is a dynamic graph with oit 1 and ct* $= \Omega(n)$.

oit $= 1$ and ct $= \Omega(n)$

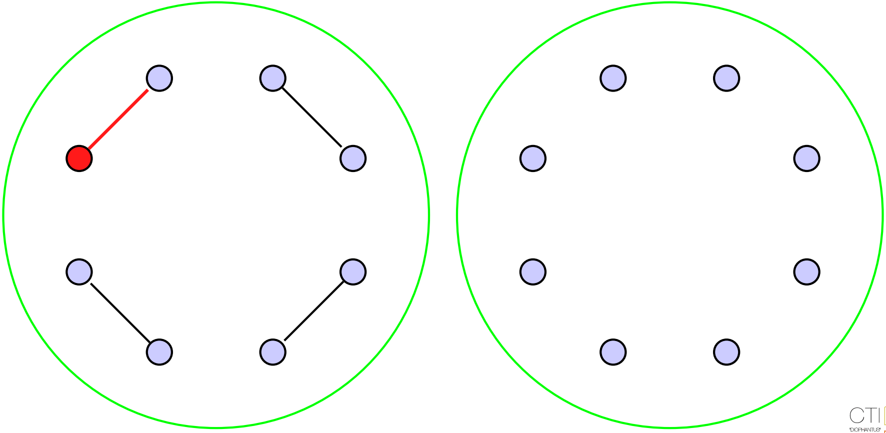

oit = 1 and ct = $\Omega(n)$

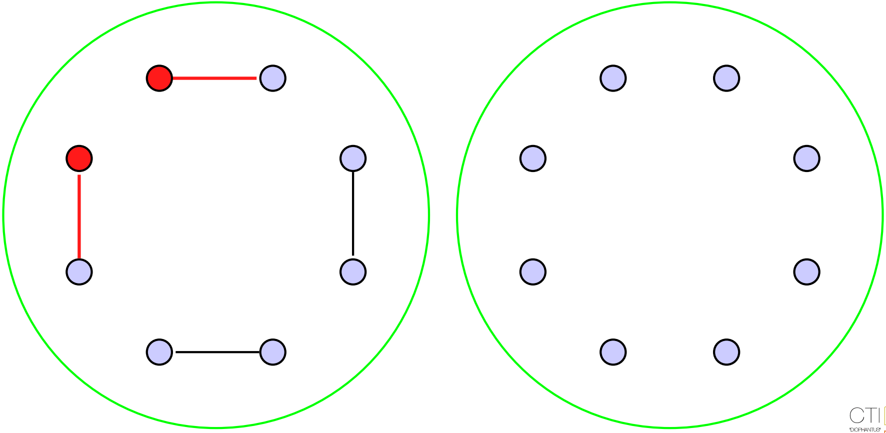

oit $= 1$ and ct $= \Omega(n)$

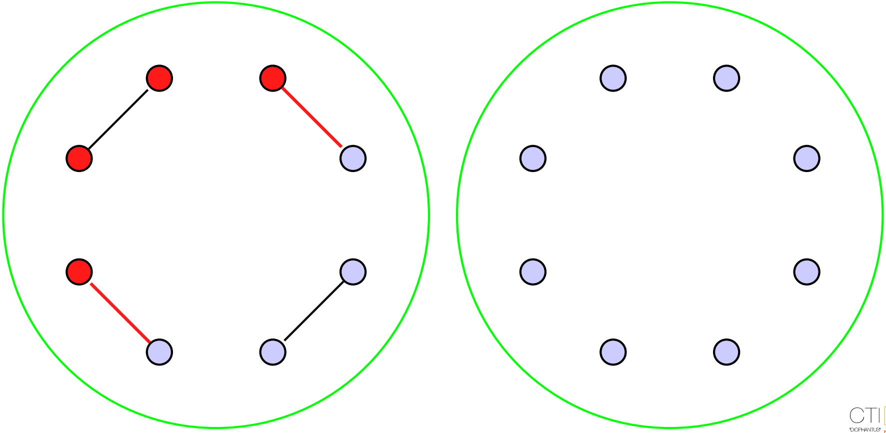

oit = 1 and ct = $\Omega(n)$

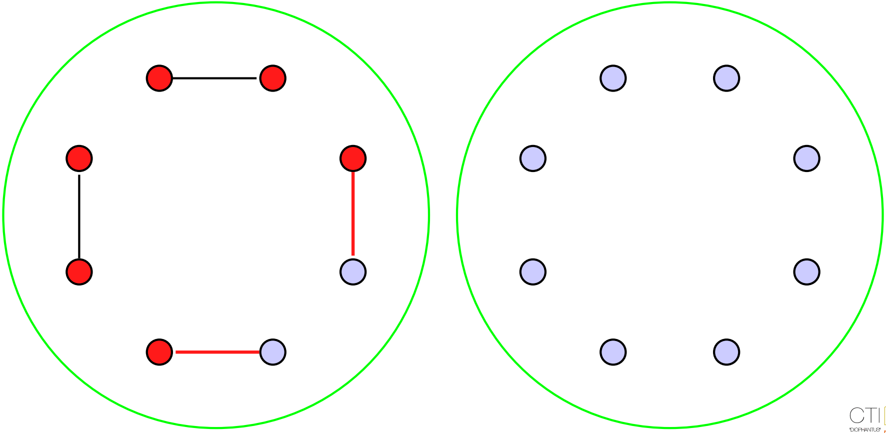

oit = 1 and ct = $\Omega(n)$

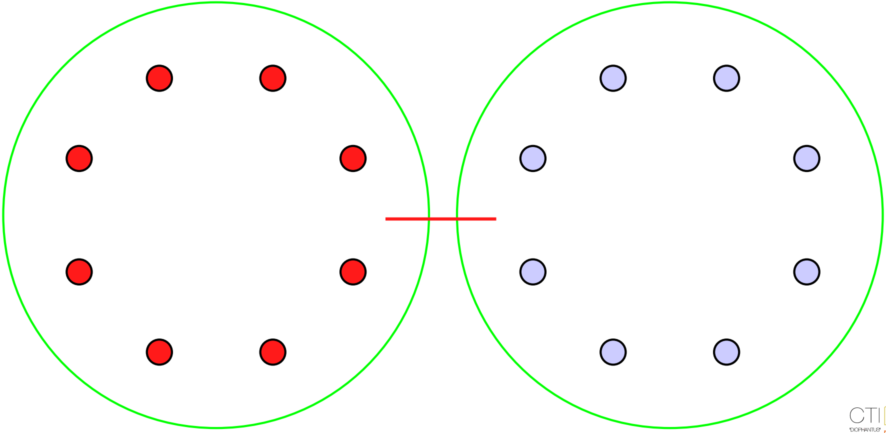

TERMINATION AND COMPUTATION

Termination Criteria

- To perform global (terminating) computation, nodes must be able determine for all times $0 \leq t \leq t'$ whether $\text{past}_{(u,t')}(t) = V$
- If nodes know n, then a node can determine at time t' whether $\text{past}_{(u,t')}(t) = V$ by counting all different t-states that it has heard of so far
- If n is not known: the subject of our work
- Termination criterion: any locally verifiable property that can be used to determine whether $\text{past}_{(u,t')}(t) = V$

Problems/Tasks

- Counting: Nodes must determine the network size n
- All-to-all Token Dissemination (or Gossip): each node is provided with a unique token, and all nodes must collect all n tokens
- Functions on Inputs: each nodes gets an input symbol from some set X and the goal is to have all nodes compute some function f on the distributed input (e.g. min,max,avg)

Problems/Tasks

Termination criteria can be used to solve these problems

- Nodes constantly broadcast all initial states that they have heard of so far
- If a node knows at round r that it has been causally influenced by the initial states of all other nodes, then to solve
- Counting: output $|\mathrm{past}_{(u,r)}(0)|$
- All-to-all Dissemination: output $\mathrm{past}_{(u,r)}(0)$
- Functions: locally compute f on the input symbols of all $u \in \mathrm{past}_{(u,r)}(0)$

Known Upper Bound on the ct

- Nodes know some upper bound T on the ct
- We give an optimal termination criterion
- This gives a protocol for counting, all-to-all token dissemination, and functions on inputs which is optimal, requiring $O(D + T)$ rounds in any dynamic network with dynamic diameter D

CTI "DIOPHANTUS"

Optimal Protocol

Theorem (Repeated Past)

Node u knows at time t that $\mathrm{past}_{(u,t)}(0) = V$ *iff* $\mathrm{past}_{(u,t)}(0) = \mathrm{past}_{(u,t)}(T)$.

Proof

- *"If":* $|\mathrm{past}_{(u,t)}(0)| \geq \min\{|\mathrm{past}_{(u,t)}(T)| + 1, n\}$.

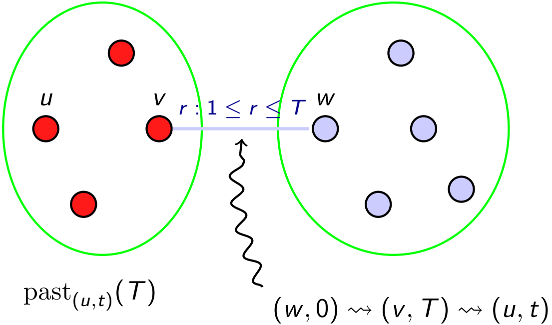

Optimal Protocol

- *"Only if"*:
  - $v \in \mathrm{past}_{(u,t)}(0) \backslash \mathrm{past}_{(u,t)}(T)$
  - *u has not heard from v since time $r < T$*
  - *Arbitrarily many nodes connected to no node until time $r-1$ and only to v thereafter*
  - *Thus, arbitrarily many nodes may be concealed from u*
  - *Implies that even if $\mathrm{past}_{(u,t)}(0) = V$ node u cannot know it*

□

Optimal Protocol:

- Nodes constantly forward all 0-states and T-states that they have heard of
- They halt as soon as $\mathrm{past}_{(u,t)}(0) = \mathrm{past}_{(u,t)}(T)$ and give the desired output

Known Upper Bound on the oit

- Nodes know some upper bound K on the oit
- We give a termination criterion which, though being far from the dynamic diameter, is optimal if a node terminates based on its past set
- We then develop a novel technique that gives an optimal termination criterion based on the future set of a node

Inefficiency of Hearing the Past

Theorem

In any given dynamic graph with oit upper bounded by K, take a node u and a time t and denote $|\mathrm{past}_{(u,t)}(0)|$ by l. It holds that $|\{v : (v, 0) \rightsquigarrow (u, t + Kl(l + 1)/2)\}| \geq \min\{l + 1, n\}$.

- That is, if a node u has at some point heard of l nodes, then u hears of another node in $O(Kl^2)$ rounds (if an unknown one exists)
- The bound is locally computable
  - Both the upper bound on the oit (K) and the number of existing incoming influences (l) are known
- Straightforward translation to protocols for our problems
- Poor time complexity: $O(Kn^2)$
- However, some sense of optimality: a node cannot obtain a better upper bound based solely on K and l

Inefficiency of Hearing the Past

- Even the "Repeated Past" criterion, that is optimal in the ct case, does not work in the oit case
- Essentially, for any t', while u has not been yet causally influenced by all initial states its past set from time 0 may become equal to its past set from time t'

Theorem

For any time t' (which can only depend on the upper bound K on the oit) there is a dynamic graph with oit upper bounded by K, a node u, and a time $t \in \mathbb{N}$ s.t. $\mathrm{past}_{(u,t)}(0) = \mathrm{past}_{(u,t)}(t')$ while $\mathrm{past}_{(u,t)}(0) \neq V$.

Hearing the Future

- Termination criterion:
  - If $\mathrm{future}_{(u,0)}(t) = \mathrm{future}_{(u,0)}(t+K)$ then $\mathrm{future}_{(u,0)}(t) = V$
- Fundamental goal: Allow a node know its future set
- Novelty: instead of hearing the past, a node now directly keeps track of its future set and is informed by other nodes of its progress

Hearing the Future

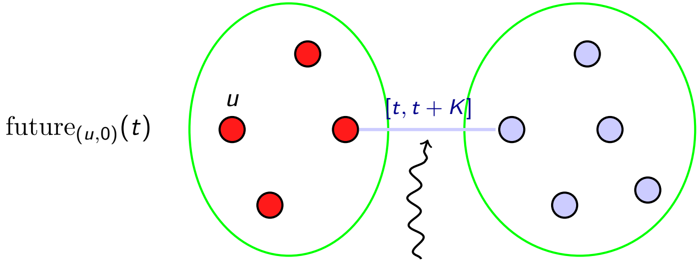

- An outgoing influence must occur in at most K rounds
- u keeps track of $\mathrm{future}_{(u,0)}(t)$
- checks whether it has increased by time $t + K$
- If not, no further nodes can exist

Protocol *Hear_from_known*

- A unique leader l (can be dropped)
- r denotes the current round
- Each node u keeps
  - $Infl_u$: keeps track of all nodes that first heard of $(l, 0)$ by u
  - A_u: keeps track of the $Infl_v$ sets that u is aware of, initially set to $(u, Infl_u, 1)$
  - *timestamp*: initially set to 1

Protocol *Hear_from_known*

- u broadcasts in every round (u, A_u) and if it has heard of $(l, 0)$ also broadcasts $(l, 0)$
- If $r > \max_{(v \neq u, r') \in Infl_u}\{r'\} + K$ then u adds (u, r) in $Infl_u$
    - As r is the maximum known time until which u has performed no further propagations of $(l, 0)$
- If u modifies $Infl_u$, it also sets $timestamp \leftarrow r$
- u updates A_u by storing in it the most recent $(v, Infl_v, timestamp)$ triple of each node v that it has heard of
- u clears multiple (w, r) records from the $Infl_v$ lists of A_u

Protocol *Hear_from_known*

- *tmax*: the maximum timestamp appearing in A_l
  - the maximum time for which the leader knows that some node was influenced by $(l, 0)$ at that time
- *R*: the set of nodes that the leader knows to have been influenced by $(l, 0)$
- If at some round r it holds at the leader that for all $u \in R$ there is a $(u, Infl_u, timestamp) \in A_l$ s.t.
  1. $timestamp \geq tmax + K$ and
  2. $\max_{(w \neq u, r') \in Infl_u} \{r'\} \leq tmax$

  then the leader halts
- Intuitively, all influenced nodes did not perform further influences for K rounds, thus no other nodes exist

Protocol *Hear_from_known*

Theorem

Protocol Hear_from_known solves counting and all-to-all dissemination in $O(D + K)$ rounds by using messages of size $O(n \log Kn)$, in any dynamic network with dynamic diameter D, and with oit upper bounded by some K known to the nodes.

- This is optimal w.r.t. time
- To drop the leader assumption:
  - all nodes begin as leaders
  - nodes prefer the leader with the smallest id that they have heard of so far
  - keep an $Infl_{(u,v)}$ only for the smallest v that they have heard of so far

Improving Message Size

- The leader initiates individual conversations with the nodes that it already knows to have been influenced by its initial state
- Sends an invitation to a particular node which is forwarded by all nodes
- A node that receives an invitation replies with the necessary data
  - this message is now preferred and forwarded by all nodes until it gets to the leader
- To make nodes prefer a particular message
  - we accompany messages with timestamps of creation-time and
  - have all nodes prefer the data with the most recent timestamps
- Terminates in $O(Dn^2 + K)$ rounds by using messages of size $O(\log D + \log n)$

Some Nice Reductions

Theorem

Assume that the oit or the iit of a dynamic graph, $G = (V, E)$, is upper bounded by K. Then for all times $t \in \mathbb{N}$ the graph $(V, \bigcup_{i=t}^{t+K\lfloor n/2 \rfloor -1} E(i))$ is connected.

Corollary

Any $f(n)$-time protocol that is correct on 1-interval-connected graphs has an equivalent $K\lfloor n/2 \rfloor f(n)$-time protocol for graphs with either oit or iit upper bounded by K and known to the nodes.

Proof.

The dynamic graph $G' = (V, E')$, where $E'(t) = \bigcup_{i=(t-1)K\lfloor n/2 \rfloor +1}^{tK\lfloor n/2 \rfloor} E(i)$, $t \geq 1$, is 1-interval connected. □

These results also carry over to the ct case as,

- if the ct of a dynamic graph is upper bounded by T, then, for all times $t \in \mathbb{N}$, the dynamic graph $(V, \bigcup_{i=t}^{t+T-1} E(i))$ is connected

ANONYMOUS DYNAMIC NETWORKS

Counting and Naming in Anonymous Dynamic Networks

- Counting: Compute n
- Naming: End up with unique identities
- Due to our reduction it suffices to focus on 1-interval connected networks
- Anonymity: Nodes do not initially have any ids,
- Unknown network: Nodes do not know the topology or the size of the network
  - We allow minimal knowledge when necessary

Static Networks with Broadcast

- Counting: Impossible to solve without a leader
- Naming: Impossible to solve even with a leader and even if nodes know n
- These impossibilities carry over to dynamic networks as well
- With a leader we solve counting in linear time using $O(\log n)$ bits per message
  - Nodes labeled with their distance from the leader
  - Each node u knows the number of upper level neighbors $up(u)$
  - Each lowest-level node u sends to the upper level $1/up(u)$
  - Intermediate nodes v sum up the values received from the lower level and send the result devided by $up(v)$ (which will be only processed by the upper level)
  - The count arrives in small parts to the leader, that computes it by summing up
  - A preprocessing step computes the eccentricity of the leader that is necessary for termination

Dynamic Networks with Broadcast

- Conjecture: Nontrivial computation is impossible
  - it is impossible to compute (even with a leader) the predicate "exists an a in the input".
- Thus, assume a unique leader that knows an upper bound
  1. d on maximum degree ever to appear in the dynamic network or
  2. e on the maximum expansion (maximum number of concurrent new influences ever occuring)
- Naming is still impossible
- We have devised protocols that obtain $O(d^n)$ and $O(n \cdot e)$ upper bounds on the count

Dynamic Networks with One-to-Each

- One-to-each message transmission
  - In every round r, each node u generates a different message $m_{u,v}(r)$ to be delivered to each current neighbor v
  - We relax broadcast in order to avoid the previous impossibilities
- Unique leader
  - Without it impossibility of naming persists even under one-to-each

Protocol *Dynamic_Naming*

- Already named nodes assign unique ids and acknowledge their id to the leader
  - Initially only the leader
- All nodes constantly forward all assigned ids that they have heard of so that they eventually reach the leader
- At some round r, the leader knows a set of assigned ids $K(r)$
- The termination criterion
  - If $|K(r)| \neq |V|$: in at most $|K(r)|$ additional rounds the leader must hear from a node outside $K(r)$
  - If $|K(r)| = |V|$: no new info will reach the leader in the future and the leader may terminate after the $|K(r)|$-round waiting period ellapses

Protocol *Dynamic_Naming*

Theorem

Dynamic_Naming solves the naming problem in anonymous unknown dynamic networks under the assumptions of one-to-each message transmission and of a unique leader. All nodes terminate in $O(n)$ rounds and use messages of size $\Theta(n^2)$.

- The individual conversations technique can again reduce the message size to $\Theta(\log n)$ paying in $O(n^3)$ termination-time

Conclusions

- We studied for the first time worst-case dynamic networks that are free of any connectivity assumption about their instances
- To enable a quantitative study we proposed some novel generic metrics that capture the speed of information propagation in dynamic networks
- We proved that fast dissemination and computation are possible even under continuous disconnectivity
- We presented optimal termination conditions and protocols based on them for fundamental distributed computing problems

Open Problems

- Improve the $O(Dn^2 + K)$ upper bound on all-to-all dissemination (with messages of size $O(\log D + \log n)$ or just $O(\log n)$ if possible)
    - The square is due to the fact that a new influence may be performed at the same time by many nodes that are unaware of each other.
- Improve the $O(n^3)$ upper bound on naming (with logarithmic messages)
- Give lower bounds for these problems in possibly disconnected dynamic networks (e.g. for centralized algorithms)
    - The only known lower bounds for dynamic networks assume connected instances
- Asynchronous communication model for the disconnected case in which nodes can broadcast when there are new neighbors

Open Problems

- Information dissemination is only guaranteed under continuous broadcasting
  - How can the number of redundant transmissions be reduced in order to improve communication efficiency?
  - Is there a way to exploit visibility to this end?
  - Does predictability help (i.e. some knowledge of the future)?
- Use randomization to construct fast and symmetry-free protocols
- Dynamic networks in which nodes have some partial control over their mobility

Thank You!

References

D. Kempe, J. Kleinberg, and A. Kumar.
Connectivity and inference problems for temporal networks.
In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 504–513, 2000.

C. Avin, M. Koucký, and Z. Lotker.
How to explore a fast-changing world (cover time of a simple random walk on evolving graphs).
In *Proceedings of the 35th international colloquium on Automata, Languages and Programming, Part I*, pages 121–132, Berlin, Heidelberg, 2008. Springer-Verlag.

F. Kuhn, N. Lynch, and R. Oshman.
Distributed computation in dynamic networks.
In *Proceedings of the 42nd ACM symposium on Theory of computing*, STOC '10, pages 513–522, New York, NY, USA, 2010. ACM.

L. Lamport.
Time, clocks, and the ordering of events in a distributed system.
Commun. ACM, 21[7]:558–565, 1978.

References

O. Michail, I. Chatzigiannakis, P.G. Spirakis.
Causality, Influence, and Computation in Possibly Disconnected Synchronous Dynamic Networks.
Submitted Paper, 2012.

D. Angluin, J. Aspnes, Z. Diamadi, M. J. Fischer, and R. Peralta.
Computation in networks of passively mobile finite-state sensors.
Distributed Computing, pages 235–253, March 2006.

D. Angluin, J. Aspnes, D. Eisenstat, and E. Ruppert.
The computational power of population protocols.
Distributed Computing, 20[4]:279–304, November 2007.

J. Augustine, G. Pandurangan, P. Robinson, and E. Upfal.
Towards robust and efficient computation in dynamic peer-to-peer networks.
In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 551–569. SIAM, 2012.

References

H. Attiya and J. Welch.
Distributed computing: fundamentals, simulations, and advanced topics, volume 19.
Wiley-Interscience, 2004.

H. Baumann, P. Crescenzi, and P. Fraigniaud.
Parsimonious flooding in dynamic graphs.
In *Proceedings of the 28th ACM symposium on Principles of distributed computing*, PODC '09, pages 260–269. ACM, 2009.

B. Bollobás.
Modern Graph Theory.
Springer, corrected edition, July 1998.

A. Casteigts, P. Flocchini, W. Quattrociocchi, and N. Santoro.
Time-varying graphs and dynamic networks.
Ad-hoc, Mobile, and Wireless Networks, pages 346–359, 2011.

References

A. E. Clementi, C. Macci, A. Monti, F. Pasquale, and R. Silvestri.
Flooding time in edge-markovian dynamic graphs.
In *Proceedings of the twenty-seventh ACM symposium on Principles of distributed computing*, PODC '08, pages 213–222, New York, NY, USA, 2008. ACM.

I. Chatzigiannakis, O. Michail, S. Nikolaou, A. Pavlogiannis, and P. G. Spirakis.
Passively mobile communicating machines that use restricted space.
Theor. Comput. Sci., 412[46]:6469–6483, October 2011.

S. Dolev.
Self-stabilization.
MIT Press, Cambridge, MA, USA, 2000.

B. Haeupler.
Analyzing network coding gossip made easy.
In *Proceedings of the 43rd annual ACM symposium on Theory of computing*, pages 293–302. ACM, 2011.

References

F. Kuhn and R. Oshman.
Dynamic networks: models and algorithms.
SIGACT News, 42:82–96, March 2011.
Distributed Computing Column, Editor: Idit Keidar.

F. Kuhn, R. Oshman, and Y. Moses.
Coordinated consensus in dynamic networks.
In *Proceedings of the 30th annual ACM SIGACT-SIGOPS symposium on Principles of distributed computing*, pages 1–10, 2011.

N. A. Lynch.
Distributed Algorithms.
Morgan Kaufmann; 1st edition, 1996.

O. Michail, I. Chatzigiannakis, and P. G. Spirakis.
Mediated population protocols.
Theor. Comput. Sci., 412[22]:2434–2450, May 2011.

References

O. Michail, I. Chatzigiannakis, and P. G. Spirakis.
New Models for Population Protocols.
N. A. Lynch (Ed), Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool, 2011.

O. Michail, I. Chatzigiannakis, and P. G. Spirakis.
Naming and counting in anonymous unknown dynamic networks.
In *Proceedings of the 26th International Symposium on Distributed Computing*, DISC '12, 2012.

E. Mendelsohn and A. Rosa.
One-factorizations of the complete graph – a survey.
J. Graph Th., 9:43–65, 1985.

R. O'Dell and R. Wattenhofer.
Information dissemination in highly dynamic graphs.
In *Proceedings of the 2005 joint workshop on Foundations of mobile computing*, DIALM-POMC '05, pages 104–110, 2005.

References

C. Scheideler.
Models and techniques for communication in dynamic networks.
In *Proceedings of the 19th Annual Symposium on Theoretical Aspects of Computer Science*, STACS '02, pages 27–49, 2002.

A. Soifer.
The Mathematical Coloring Book: Mathematics of Coloring and the Colorful Life of its Creators.
Springer; 1st edition, 2009.

D. Angluin.
Local and global properties in networks of processors (extended abstract).
In *Proceedings of the twelfth annual ACM symposium on Theory of computing*, STOC '80, pages 82–93. ACM, 1980.

H. Attiya, M. Snir, and M. K. Warmuth.
Computing on an anonymous ring.
J. ACM, 35[4]:845–875, October 1988.

References

P. Fraigniaud, A. Pelc, D. Peleg, and S. Pérennes.
Assigning labels in unknown anonymous networks (extended abstract).
In *Proceedings of the nineteenth annual ACM symposium on Principles of distributed computing*, PODC '00, pages 101–111. ACM, 2000.

V. Kostakos.
Temporal graphs.
Physica A: Statistical Mechanics and its Applications, 388[6]:1007–1023, 2009.

N. Sakamoto.
Comparison of initial conditions for distributed algorithms on anonymous networks.
In *Proceedings of the eighteenth annual ACM symposium on Principles of distributed computing*, PODC '99, pages 173–179. ACM, 1999.

M. Yamashita and T. Kameda.
Computing on anonymous networks. i. characterizing the solvable cases.
Parallel and Distributed Systems, IEEE Transactions on, 7[1]:69–89, 1996.