J. Parallel Distrib. Comput. 74 (2014) 2016–2026

Contents lists available at ScienceDirect

J. Parallel Distrib. Comput.

journal homepage: www.elsevier.com/locate/jpdc

Causality, influence, and computation in possibly disconnected synchronous dynamic networks☆,☆☆

Othon Michail [a,*], Ioannis Chatzigiannakis [a], Paul G. Spirakis [a,b]

[a] *Computer Technology Institute & Press "Diophantus" (CTI), Patras, Greece*

[b] *Department of Computer Science, University of Liverpool, UK*

HIGHLIGHTS

- We study computation in possibly disconnected dynamic distributed systems.
- We replace continuous connectivity by minimal temporal connectivity conditions.
- We propose metrics capturing the speed of information spreading in dynamic networks.
- We give efficient protocols for the counting and token dissemination problems.

ARTICLE INFO

Article history:
Received 1 March 2013
Received in revised form
17 July 2013
Accepted 29 July 2013
Available online 7 August 2013

Keywords:
Dynamic graph
Mobile computing
Worst-case dynamicity
Adversarial schedule
Temporal connectivity
Termination
Counting
Information dissemination
Optimal protocol

ABSTRACT

In this work, we study the propagation of influence and computation in dynamic distributed computing systems that are *possibly disconnected* at every instant. We focus on a *synchronous message-passing* communication model with *broadcast* and bidirectional links. Our network dynamicity assumption is a *worst-case dynamicity* controlled by an adversary scheduler, which has received much attention recently. We replace the usual (in worst-case dynamic networks) assumption that the network is connected at every instant by minimal *temporal connectivity* conditions. Our conditions only require that *another causal influence occurs within every time window of some given length.* Based on this basic idea, we define several novel metrics for capturing the speed of information spreading in a dynamic network. We present several results that correlate these metrics. Moreover, we investigate *termination criteria* in networks in which an upper bound on any of these metrics is known. We exploit our termination criteria to provide efficient (and optimal in some cases) protocols that solve the fundamental *counting* and *all-to-all token dissemination* (or *gossip*) problems.

1. Introduction

Distributed computing systems are becoming increasingly dynamic. The static and relatively stable models of computation can no longer represent the plethora of recently established and rapidly emerging information and communication technologies. In recent years, we have seen a tremendous increase in the number of new mobile computing devices. Most of these devices are equipped with some sort of communication, sensing, and mobility capabilities. Even the Internet has become mobile. The design is now focused on complex collections of heterogeneous devices that should be robust, adaptive, and self-organizing, possibly moving around and serving requests that vary with time. Delay-tolerant networks are highly dynamic, infrastructureless networks whose essential characteristic is a possible absence of end-to-end communication routes at any instant. Mobility may be *active*, when the devices control and plan their mobility pattern (e.g. mobile robots), or *passive*, in opportunistic-mobility networks, where mobility stems from the mobility of the carriers of the devices (e.g. humans carrying cell phones) or a combination of both (e.g. the devices have partial control over the mobility pattern, like for example when GPS devices provide route instructions to their carriers). Thus, it can vary from being completely predictable to being completely

☆ Supported in part by the project "Foundations of Dynamic Distributed Computing Systems" (FOCUS) which is implemented under the "ARISTEIA" Action of the Operational Programme "Education and Lifelong Learning" and is co-funded by the European Union (European Social Fund) and Greek National Resources..

☆☆ A preliminary version of the results in this paper has appeared in [24].

* Correspondence to: Computer Technology Institute & Press "Diophantus" (CTI), N. Kazantzaki Str., Patras University Campus, Rio, P.O. Box 1382, 26504, Greece.

E-mail addresses: michailo@cti.gr, michailo@ceid.upatras.gr (O. Michail), ichatz@cti.gr (I. Chatzigiannakis), spirakis@cti.gr (P.G. Spirakis).

0743-7315/$ – see front matter © 2013 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.jpdc.2013.07.007

unpredictable. Gossip-based communication mechanisms, e-mail exchanges, peer-to-peer networks, and many other contemporary communication networks all assume or induce some sort of highly dynamic communication network.

The formal study of dynamic communication networks is hardly a new area of research. There is a huge amount of work in distributed computing that deals with causes of dynamicity such as failures and changes in the topology that are rather slow and usually eventually stabilize (like, for example, in self-stabilizing systems [11]). However, the low rate of topological changes that is usually assumed there is unsuitable for reasoning about truly dynamic networks. Even graph-theoretic techniques need to be revisited: the suitable graph model is now that of a *dynamic graph* (also known as a *temporal graph* or *time-varying graph*) (see e.g. [25,15,8]), in which each edge has an associated set of time labels indicating availability times. Even fundamental properties of classical graphs do not easily carry over to their temporal counterparts. For example, Kempe, Kleinberg, and Kumar [15] found that there is no analog of Menger's theorem (see e.g. [7] for a definition) for arbitrary temporal networks with one label on every edge, which additionally renders the computation of the number of node-disjoint s–t paths **NP**-complete. Very recently, the authors of [25] achieved a reformulation of Menger's theorem which is valid for all temporal networks, and additionally they introduced several interesting cost-minimization parameters for optimal temporal network design and gave some first results on them. Even the standard network diameter metric is no longer suitable, and it has to be replaced by a dynamic/temporal version. In a dynamic star graph in which all leaf nodes but one go to the center one after the other in a modular way, any message from the node that last enters the center to the node that never enters the center needs $n-1$ steps to be delivered, where n is the size (number of nodes) of the network; that is, the *dynamic diameter* is $n-1$ while, one the other hand, the classical diameter is just 2 [3] (see also [18]).

2. Related work

Distributed systems with worst-case dynamicity were first studied in [26]. Their outstanding novelty was to assume a communication network that may change arbitrarily from time to time subject to the condition that each instance of the network is connected. They studied asynchronous communication and considered nodes that can detect local neighborhood changes; these changes cannot happen faster than it takes for a message to transmit. They studied *flooding* (in which one node wants to disseminate one piece of information to all nodes) and *routing* (in which the information need only reach a particular destination node t) in this setting. They described a uniform protocol for flooding that terminates in $O(Tn^2)$ rounds using $O(\log n)$ bit storage and message overhead, where T is the maximum time it takes to transmit a message. They conjectured that, without identifiers (IDs), flooding is impossible to solve within the above resources. Finally, a uniform routing algorithm was provided that delivers to the destination in $O(Tn)$ rounds using $O(\log n)$ bit storage and message overhead.

Computation under worst-case dynamicity was further and extensively studied in a series of works by Kuhn et al. in the synchronous case. In [16], the network was assumed to be *T-interval connected*, meaning that any time window of length T has a static connected spanning subgraph (persisting throughout the window). Among others, *counting* (in which nodes must determine the size of the network) and *all-to-all token dissemination* (in which n different pieces of information, called tokens, are handed out to the n nodes of the network, each node being assigned one token, and all nodes must collect all n tokens) were solved in $O(n^2/T)$ rounds using $O(\log n)$ bits per message, almost-linear-time randomized approximate counting was established for $T=1$, and two lower bounds on token dissemination were given. Several variants of *coordinated consensus* in 1-interval connected networks were studied in [17]. Two interesting findings were that, in the absence of a good initial upper bound on n, eventual consensus is as hard as computing deterministic functions of the input, and that *simultaneous consensus* can never be achieved in less than $n-1$ rounds in any execution. [13] is a recent work that presents information-spreading algorithms in worst-case dynamic networks based on *network coding*. An *open* setting (modeled as high churn) in which nodes constantly join and leave has very recently been considered in [4]. For an excellent introduction to distributed computation under worst-case dynamicity, see [18]. Two very thorough surveys on dynamic networks are [27,8].

Another notable model for dynamic distributed computing systems is the *population protocol* model [1]. In that model, the computational agents are passively mobile and interact in ordered pairs, and the connectivity assumption is a *strong global fairness condition* according to which all events that may always occur, occur infinitely often. These assumptions give rise to some sort of structureless interacting automata model. The usually assumed *anonymity* and *uniformity* (i.e. n is not known) of protocols only allow for commutative computations that eventually stabilize to a desired configuration. Most computability issues in this area have now been established. Constant-state nodes on a complete interaction network (and several variations) compute the *semilinear predicates* [2]. Semilinearity persists up to $o(\log\log n)$ local space but not more than this [10]. If constant-state nodes can additionally leave and update fixed-length pairwise marks, then the computational power dramatically increases to the commutative subclass of $\mathbf{NSPACE}(n^2)$ [21]. For a very recent introductory text see [22].

3. Contribution

In this work, we study worst-case dynamic networks that are *free of any connectivity assumption about their instances*. Our dynamic network model is formally defined in Section 4.1. We only impose some *temporal connectivity* conditions on the adversary guaranteeing that *another causal influence occurs within every time window of some given length*, meaning that, in that time, another node first hears of the state that some node u had at some time t (see Section 4.3 for a formal definition of *causal influence*). Note that our temporal connectivity conditions are minimal assumptions that allow for bounded end-to-end communication in any dynamic network including those that have disconnected instances. Based on this basic idea, we define several novel generic metrics for capturing the speed of information spreading in a dynamic network. In particular, we define the *outgoing influence time* (oit) as the maximal time until the state of a node *influences* the state of another node, the *incoming influence time* (iit) as the maximal time until the state of a node *is influenced by* the state of another node, and the *connectivity time* (ct) as the maximal time until the two parts of any cut of the network become connected. These metrics are defined in Section 5, where also several results that correlate these metrics to themselves and to standard metrics, like for example the dynamic diameter, are presented.

In Section 5.1, we present a simple but very fundamental dynamic graph based on alternating matchings that has oit 1 (equal to that of instantaneous connectivity networks) but at the same time is *disconnected in every instance*. In Section 6, we exhibit another dynamic graph additionally guaranteeing that edges take maximal time to reappear. That graph is based on a geometric edge-coloring method due to Soifer for coloring a complete graph of even order n with $n-1$ colors [28]. Similar results have appeared before, but to the best of our knowledge only in probabilistic settings [9,6].