Theoretical Computer Science 584 (2015) 80–90

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

A characterization of oblivious message adversaries for which Consensus is solvable

Étienne Coulouma[a,1], Emmanuel Godard[a,*,2], Joseph Peters[b,3]

[a] *LIF, Université Aix-Marseille and CNRS, France*
[b] *School of Computing Science, Simon Fraser University, Canada*

ARTICLE INFO

Article history:
Received 18 November 2013
Received in revised form 27 November 2014
Accepted 10 January 2015
Available online 20 January 2015

Keywords:
Consensus
Synchronous
Message passing
Omission faults
Dynamic networks
Message adversary
Oblivious faults

ABSTRACT

We consider the Consensus problem on arbitrary oblivious message adversaries. A message adversary models a communication network whose topology evolves from round to round. We make no assumptions about the possible topologies. A message adversary is oblivious if the set of possible topologies is the same at every round.

We give the first complete necessary and sufficient condition for message adversaries that admits a Consensus algorithm. For the necessary part we present a specialized bivalency proof. For the sufficiency part, we present a new algorithm that is based upon reconstructing a partial, but significant, view of the actual communications that occurred during the evolution of the network. This reconstruction might be of independent interest.

1. Introduction

1.1. Dynamic networks and message adversaries

Designing algorithms for static networks is an area that has been studied with numerous approaches (distributed/centralized, online/offline, …). This is one of the main themes of distributed computing. Designing algorithms for dynamic networks, where the network structure can be modified during the computation is less understood. Numerous research projects have studied systems where the origin of the dynamicity is from faults (resulting in deletions and additions of nodes and edges to the network). Indeed, fault-tolerance is probably one of the main endeavors in distributed computing. However, faults are in general of limited scope, of limited number and, above all, are considered to be anomalies with respect to the normal and correct behavior of the system. So, here we consider systems that are never stable, where the number of changes is not bounded and changes are continuously occurring, and where these changes are not considered anomalies but are an integral part of the system at hand. Such highly dynamic systems do exist, they are actually quite common, and they are becoming pervasive.

We consider communication networks in which the topology can evolve from round to round. A specific link can dynamically disappear and then appear again after an unpredictable number of rounds, and it can continue to alternate between

---

[*] Corresponding author.
E-mail addresses: emmanuel.godard@lif.univ-mrs.fr (E. Godard), peters@cs.sfu.ca (J. Peters).

[1] Research supported by ANR PANDA.

[2] Research supported by PIMS.

[3] Research supported by NSERC of Canada (31-611005).

http://dx.doi.org/10.1016/j.tcs.2015.01.024
0304-3975/

being present and absent in an unpredictable way. This variability can also be interpreted as *intermittent omission faults*. This synchronous presentation of the communications has also been shown to be good for *layered analysis* [11]. This model is more general than other models, such as *component failure* models, whose evolutions, once they appear somewhere, are located there permanently. Interestingly, we will show, using the notion of communication events, that this model is also closely related to a failure model that was introduced in [16], and was called the *mobile faults* or *dynamic faults* model. A similar notion was presented in [1], where the terminology "message adversary" was introduced.

An important property of the systems we are studying here is that the set of possible simultaneous communications is the same for each round. In some sense, the system has no "memory" of its previous evolution. The adversary, which changes the topology of communications, is oblivious. Real systems often exhibit such memory-less, or oblivious, behavior. Moreover, we do not restrict ourselves to networks with complete connectivity as is usually done. In this paper we consider the *most general case* of such systems, i.e., systems in which the set of possible simultaneous communications is arbitrary. This allows the modeling of any system in which communications can happen intermittently, in any arbitrary pattern, including systems in which the communications are not symmetric.

We investigate the Consensus problem in these networks. While it has long been known that solvability of the Broadcast problem implies solvability of the Consensus problem, we show here the precise relationship between those two problems. In [7], an impossibility proof for Consensus was presented in the context of omission faults. In [5], we showed that the necessary condition for solvability of Consensus in [7] is actually sufficient by presenting a new Consensus algorithm in the context of dynamic networks. In this paper, we give a common framework for these two results in the message adversary model, and we obtain Theorem 4.10 which is the first complete characterization of solvability of Consensus on arbitrary oblivious message adversaries.

1.1.1. The Consensus problem

The Consensus problem is a very well studied problem in the area of Distributed Algorithms. It is defined as follows. Each node of the network starts with an initial value, and all nodes of the network have to agree on a common value, which is one of the initial values. Many versions of the problem concern the design of algorithms for systems that are unreliable.

1.1.2. The Broadcast problem

Two of the most widely studied patterns of information propagation in communication networks are *broadcasting* and *gossiping*. A *broadcast* is the distribution of an initial value from one node of a network to every other node of the network. A *gossip* is a simultaneous broadcast from every node of the network. The Broadcast problem that we study in this paper is to find a node from which a broadcast can be successfully completed.

There are close relationships between broadcasting and gossiping, and the Consensus problem. Indeed, the Consensus problem can be solved by first gossiping and then applying a deterministic function at each node to the set of initial values. But a gossip is not actually necessary. If there exists a distinguished node v_0 in the network, then a Consensus algorithm can be easily derived from an algorithm that broadcasts from v_0. However the Broadcast problem and the Consensus problem are not equivalent, as will be made clear in Section 3.2.

1.2. Arbitrary patterns of dynamicity

In this paper, we study arbitrary message adversaries, that is *arbitrary patterns* of intermittent communications. Two main patterns of dynamicity have been studied. First, there is the case when there is a possibly infinite number of missing links and this number is only bounded by some integer f in a round. This is the model of [17] (where it originated from faults). In this setting, it is possible that the Consensus problem does not admit a solution. So the question is whether there is a Consensus Algorithm for this f, and if there exist solutions, what is the minimum number of rounds needed.

It is also possible that the number of missing communications is actually finite, and in this setting Consensus is obviously solvable (for example using a flooding algorithm) and the question is what are the best time and space complexities for a solution.

1.2.1. Our contributions

In this paper, we investigate oblivious message adversaries, that is, communication networks where the topology evolves arbitrarily from round to round and nodes do not know their neighbors at a given instant [1]. The topology of the network at a given instant is called a *communication event*. We give a necessary and sufficient condition for oblivious message adversaries for which Consensus is solvable.

Having the full characterization is important from a theoretical point of view, but this is also very interesting as both sides illustrate the problem at hand. Furthermore, we have here a constructive presentation of our tools, especially about the relation β (to be formally defined later) with Proposition 6.2.

A node from which it is possible to broadcast if the system is restricted to a given communication event is called a *source* for the communication event. We define an equivalence relation β on the communication events that is based on the collective local observations of the events by the sources. The characterization is as follows: *it is possible to solve Consensus if and only if for all β-classes $\mathcal{C}$, there is a node that can broadcast when only events from $\mathcal{C}$ can occur.* It is very simple to characterize Broadcastability (see Theorem 4.6), so we get very simple and efficient conditions about solving Consensus.