ARTICLE IN PRESS
Future Generation Computer Systems ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Future Generation Computer Systems

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

Improving the gossiping effectiveness with distributed strategic learning (Invited paper)

Christian Esposito [a,*], Aniello Castiglione [a], Francesco Palmieri [a], Massimo Ficco [b]

[a] Dipartimento di Informatica, Università degli Studi di Salerno, Via Giovanni Paolo II 132, I-84084 Fisciano (SA), Italy
[b] Dipartimento di Ingegneria Industriale e dell'Informazione, Seconda Università di Napoli, via Roma 29, I-81031 Aversa (CE), Italy

ARTICLE INFO

Article history:
Received 1 May 2016
Received in revised form
29 October 2016
Accepted 8 November 2016
Available online xxxx

Keywords:
Event-based communications
Gossiping
Reliable multicasting
Strategic learning
Game theory
Reinforcement learning

ABSTRACT

Gossiping is a widely known and successful approach to reliable communications, tolerating packet losses and link crashes. It has been extensively used in several middleware kinds, such as event notification services and application domains, like infrastructures for air traffic management, power grid control, health information exchange, just to cite some of them. Despite achieving a high loss-tolerance and scalability degrees, gossiping is affected by degraded performances and heavy traffic loads on the network. For this reason, it may be not optimal in applications where reliability must be provided jointly with timeliness and/or in congestion-prone networks. The crucial aspect for improving a gossiping scheme is deciding which nodes should receive a gossiping message, and our driving idea is to adopt a distributed strategic learning logic to determine such nodes in an efficient manner. This is able to resolve gossiping's weakness points and to achieve better performance and reduced traffic loads.

This paper describes how to introduced strategic learning in a gossip scheme so as to determine the best set of nodes that can be used to send gossip messages and to optimize their utility. Such a solution has been experimentally assessed through a set of simulations demonstrating the effectiveness of the proposal.

1. Introduction

Event notification is a communication pattern that can be found in several different application domains, such as infrastructures for air traffic management, power grid control, health information exchange, just to cite some of them. This is due to the recent increasing demand for an effective communication paradigm that can be used to interconnect several heterogeneous systems by achieving flexibility and scalability. In addition to their large scale, such systems perform critical operations, whose failure may lead to an enormous number of losses in terms of human lives and money. Therefore, the communication solutions realizing event notification assume great importance due to their intrinsic decoupling and efficiency guarantees that promote scalability. However, such solutions also require to provide a high reliability degree, by guaranteeing notification delivery despite failures within the network, and the current available options are lacking the required support of reliable notification solutions, which still remains an open issue. The current literature on reliable event notification encompasses several different techniques to augment the reliability degree offered by an event-based communication platform in order to tolerate the possible faults affecting them (interested readers can refer to [1] for more details on these techniques and a suitable fault model for event notification), among which gossiping is one of the most considered ones.

Gossiping [2] is a popular data distribution strategy for scalable and reliable multicasting over large-scale networks, tolerating the occurrence of network failures leading to severe message losses. Moreover, it represents a powerful solution to deliver a consistent view of certain pieces of data among asynchronous distributed processes despite of possible failures (which is theoretically unreachable due to the FLP impossibility proof [3]) by providing the so called eventual consistency [4]. In fact, the work in [5] demonstrated that gossiping is able to provide eventual consistency with high probability in large-scale and weakly connected systems. Thanks to these points of strength, gossiping has been proposed as a key building block for realizing several innovative systems, such as consolidating VMs in

* Corresponding author.
E-mail addresses: esposito@unisa.it (C. Esposito), castiglione@ieee.org (A. Castiglione), f.palmieri@unisa.it (F. Palmieri), massimo.ficco@unina2.it (M. Ficco).

http://dx.doi.org/10.1016/j.future.2016.11.006
0167-739X/

Please cite this article in press as: C. Esposito, et al., Improving the gossiping effectiveness with distributed strategic learning (Invited paper), Future Generation Computer Systems (2016), http://dx.doi.org/10.1016/j.future.2016.11.006

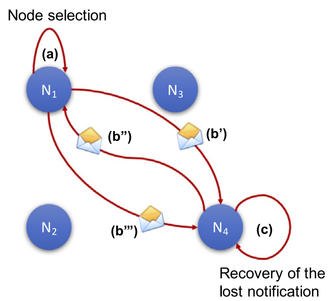

Fig. 1. Scheme of a gossiping approach to loss recovery.

large cloud data centers [6], Delay tolerant networks (DTNs) [7], distributed intrusion detection [8], or peer-to-peer live media streaming [9], just to cite some of them. Generally speaking, as depicted in Fig. 1, a gossiping scheme consists in nodes selecting with which other nodes to communicate (step a in the figure) and exchanging messages to detect possible losses (interactions b) and recover from them (step c). Such exchanged messages can contain some of the recently received notifications, or a list of these notifications. In the first case, called push mode, at the reception of a gossiped notification (represented by the interaction b', which is the only one to take place), the receiving node can decide to discard it if such a notification has been already received in the past, or to pass it to the consuming application, in order to recover a lost notification. In the second case, called pull mode, the list of recently received notifications is sent (interaction b') and it can be used to detect some losses and to trigger a retransmission (interaction b'' represents the request for a retransmission, which is performed in the interaction b'''), by recovering the lost notifications. So, more schematically, gossiping consists of two key phases: the node selection and the exchange of gossiping messages based on the adopted mode. Despite its advantages, gossiping presents a serious weakness, which we have studied and documented in a previous paper [10]: the first phase of node selection is generally performed according to a random nature, and this causes high number of unneeded messages being exchanged among the nodes.

This implies both a considerable traffic load on the network, which may eventually cause congestion phenomena, and a non-optimal recovery time of lost data by wasting gossip messages sent toward nodes that do not require them. In a previous publication [10], we have proposed a solution to alleviate this problem by limiting the typical random behavior of gossiping when selecting nodes, and biasing the node selection by means of a proper heuristics, *i.e.*, by assigning a weight to each possible gossip candidate and selecting only those ones with the highest weights. In this manner, we forced the algorithm to prefer gossiping only with the nodes that are expected to require a recovery action in order to tolerate the occurrence of a message loss. The above solution, however, did not demonstrated to cause a considerable improvement in the quality of the gossip algorithm. It implied a reduction of the protocol overhead, *i.e.*, the mean number of messages that a node received during a gossip round. However, if we measure the utility of the gossiping messages as the percentage of those received messages that have been used to recover a lost notification, we can observe that this utility is not close to the optimal value. In addition, we have noticed that our weighting approach causes some issues when recovering notification for the nodes with low weights, and this implies a lower reliability capability of the gossiping. In this paper, we aim at further improving our previous approach by applying game theory with a proper strategic selection of the nodes which a given node has to gossip with. This diverges from other approaches that have tried to improve gossiping by using alternative means, such as the ones based on location information [11] or structured approaches [12]. The novel contribution of this work can be summarized as follows:

- Formalize by means of an optimization problem the explosion of unnecessary messages, which affect the overhead and quality of a gossip scheme;
- Propose a solution of such a problem by using non-cooperative game theory;
- Characterize the pay-off in the proposed game by using distributed strategic learning;

In the following, we review the literature on gossiping and provide some preliminary concepts in Section 2. Section 3 describes the key characteristics of gossiping and its drawbacks in a quantitative manner by means of results we have obtained from simulations. Section 4 presents some related works on the topic and how our work is compared to them. Our approach applying distributed strategic learning to optimize and improve the efficiency and effectiveness of gossiping is presented in Section 5. In Section 6, we experimentally evaluate the achieved reliability, performance and overhead of our approach. We conclude with some final remarks and indicate a possible plan for future works in Section 7.

2. Background

2.1. Gossiping

Gossiping [2] is one of the most known and studied reactive method for reliable communications in large-scale systems, and it has been extensively used to build reliable multicasting. Basically, it implements a distributed loss tolerance scheme by means of retransmissions, where lost messages are recovered thanks to temporal redundancy. The classic retransmission-based method for tolerating message losses is the Automatic Request reply (ARQ) scheme [13], where a data source queues the sent messages and makes a retransmission in case of a negative acknowledge of a missing message (due to a loss detection made by a destination), or a missing positive acknowledgment for a certain message (determined as a missing delivery of the positive acknowledgment before the expiration of a proper timeout). On the contrary to ARQ, gossiping does not centralize the recovery duties in a single node of a given communication infrastructure, and this greatly helps in having a scalable approach to loss tolerance. Fig. 1 schematically describes the behavior of a generic gossiping scheme: a node stores a sent/received new message in a buffer with a size b, and starts a so-called gossip round, which consists in exchanging certain information about these messages to a randomly-selected set of nodes of size f_{out} (called fanout). Thanks to these message exchanges, retransmission takes place and any detected message loss is recovered. The messages are not stored in a node queue indefinitely, but only for a limited number of rounds f_{cap} (called capacity), after which they are dequeued and deleted.

Many variants of the gossiping algorithms exist within the current literature [1], which can be categorized as depicted in Fig. 2. The first classification is based on the communication protocol adopted during gossiping, *i.e.*, when starting a gossip round. In the *Push Approaches* represented in the Fig. 2, gossip messages are forwarded to the other nodes as soon as a new message is received (*i.e.*, capacity is constant and implicitly equal to 1). On the other hand, in *Pull Approaches* represented in the Fig. 2, after the expiration of a proper timeout t, nodes periodically send a given gossip message to the other selected nodes. Another classification criterion is based on the content of the exchanged gossip messages. The first one is that gossip messages hold the content of the exchanged application messages, so as to properly implement a timely retransmission scheme at the expenses of heavy network traffic. The second one is that a gossip message contains a digest, which can be positive, *i.e.*, the list of the recently received notifications, or negative, *i.e.*, the list of missed notifications. If a

Please cite this article in press as: C. Esposito, et al., Improving the gossiping effectiveness with distributed strategic learning (Invited paper), Future Generation Computer Systems (2016), http://dx.doi.org/10.1016/j.future.2016.11.006