J. Parallel Distrib. Comput. 99 (2017) 28–50

Contents lists available at ScienceDirect

J. Parallel Distrib. Comput.

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

Highly intensive data dissemination in complex networks

Gabriele D'Angelo *, Stefano Ferretti

Department of Computer Science and Engineering, University of Bologna, Italy

HIGHLIGHTS

- We study the performance of dissemination protocols on complex networks.
- We analyze dissemination environments that are highly intensive.
- We propose both an analytic and a simulation-based analysis.
- We use a scalable parallel and distributed simulation approach.

ARTICLE INFO

Article history:
Received 9 December 2015
Received in revised form
28 April 2016
Accepted 28 August 2016
Available online 4 September 2016

Keywords:
Data dissemination
Simulation
Complex networks
Performance evaluation

ABSTRACT

This paper presents a study on data dissemination in unstructured Peer-to-Peer (P2P) network overlays. The absence of a structure in unstructured overlays eases the network management, at the cost of non-optimal mechanisms to spread messages in the network. Thus, dissemination schemes must be employed that allow covering a large portion of the network with a high probability (e.g. gossip based approaches). We identify principal metrics, provide a theoretical model and perform the assessment evaluation using a high performance simulator that is based on a parallel and distributed architecture. A main point of this study is that our simulation model considers implementation technical details, such as the use of caching and Time To Live (TTL) in message dissemination, that are usually neglected in simulations, due to the additional overhead they cause. Outcomes confirm that these technical details have an important influence on the performance of dissemination schemes and that the studied schemes are quite effective to spread information in P2P overlay networks, whatever their topology. Moreover, the practical usage of such dissemination mechanisms requires a fine tuning of many parameters, the choice between different network topologies and the assessment of behaviors such as free riding. All this can be done only using efficient simulation tools to support both the network design phase and, in some cases, at runtime.

1. Introduction

Unstructured Peer-to-Peer (P2P) systems have been recognized as a good practice to build effective distributed applications. This is particularly evident when peers composing the network are dynamic, with frequent arrivals and departures. In fact, in this case, the use of agile attachment strategies to create an overlay network (i.e. the network composed of links representing an interaction/connection between nodes), plus the use of a simple dissemination protocol to let nodes interact, offer an easy way to manage interaction substrate, on top of which it is possible executing distributed applications. Nodes create links based on an attachment process that does not depend on the "type" of involved nodes. Thus, for instance, if we are dealing with a content management system, links are not created based on the contents owned by peers; rather, links are established based on other criteria, e.g. arbitrarily.

As concerns the spread of information, an interesting solution is based on gossip. This epidemic dissemination strategy uses randomized communication that distributes contents without a specific, content based, routing scheme. Gossip has been recognized as a robust and scalable communication paradigm to be employed in large-scale distributed environments [11,37,16]. In fact, although it has communication costs usually higher than other, optimized solutions, e.g. tree-based protocols, a gossip-based dissemination scheme is intrinsically fault tolerant.

There is a vast literature on gossip. Related studies are mainly theoretical, since their aim is to prove that large-scale networks can reliably and effectively employ these strategies to disseminate information [22,34,2,21,28]. For instance, in [2] it is shown

* Correspondence to: Department of Computer Science and Engineering, University of Bologna, Mura Anteo Zamboni 7, I-40127, Bologna, Italy.

E-mail addresses: g.dangelo@unibo.it (G. D'Angelo), s.ferretti@unibo.it (S. Ferretti).

http://dx.doi.org/10.1016/j.jpdc.2016.08.004
0743-7315/

Acronyms

LUNES	Large Unstructured NEtwork Simulator
FP	Fixed Probability (dissemination protocol)
PB	Probabilistic Broadcast (dissemination protocol)
DDF1	Degree Dependent Function 1 (dissemination protocol)
DDF2	Degree Dependent Function 2 (dissemination protocol)

that, depending on the system model, certain gossip-based protocols can achieve a message complexity around $O(n \log^3 n)$, or even $O(n)$, with high probability. It is worth mentioning that in the proposed models the behavior of nodes is usually simplified, and several practical issues are not considered, that instead should be took into consideration when building a distributed system. Some other works exploit simulation to evaluate epidemic strategies [22,33,37,16]. Also in these works, nodes have a very simple behavior. The rationale behind this choice is twofold, usually. Firstly, often a simple behavior of nodes/agents allows verifying quite easily if an interesting emergent behavior occurs at the whole system/network level. Secondly, these simplifications allow having a lightweight simulation model that enables the simulator to scale up to large networks. However, also in this case, while a general main result is obtained, there is a lack of technical details, which are instead important during the real deployment of these strategies.

Today, the use of parallel and distributed simulation and the advent of multi/many-core processors make possible adding more details on the behavior of simulated entities. Adding such complexity in simulation corresponds to give more emphasis on the impact of some important algorithmic details and expedients that can affect the dissemination performance in a P2P overlay network. In this work, we assess the performance of different dissemination protocols on different P2P overlays, and study the impact of caching and of the Time To Live (TTL) to distribute messages. In our assessment, we employ parallel and distributed simulation. The metrics employed during the assessment are the coverage of the network and the delay for disseminating messages. Not only, a theoretical model is provided that, given a gossip protocol and the topology of the underlying network overlay, allows estimating the threshold values corresponding to a phase transition between the ability of a given gossip protocol to spread a message to a significant set of nodes in the net, and a local dissemination that reaches a limited neighborhood of a node, only.

The contributions of this work are the following.

- We study a *degree dependent* dissemination algorithm, that relays messages to nodes based on their degree. We employ different degree dependent functions in our simulations.
- We provide a theoretical model that based on the dissemination protocol and the degree distribution of nodes composing the network, is able to determine the threshold values for the parameters of the dissemination algorithm. Such a threshold identifies a phase transition: below the threshold a disseminated message reaches a small, local fraction of network nodes, while above the threshold the message reaches a giant component of the network, i.e. a set of nodes of the order of the network size.
- We study the impact of cache and TTL on data dissemination.
- We perform a parallel and distributed simulation of dissemination algorithms in large scale networks; different network topologies are employed (i.e. random graphs, scale-free networks, Watts–Strogatz small-world networks, k-regular). To the best of our knowledge, this is the first contribution that shows results of large scale simulations over different topologies, where the dissemination is so highly intensive and nodes behavior considers cache and TTL management. Moreover, contrary to other typical simulation studies, where a single node acts as the source generating messages, in our simulations all nodes generate messages, concurrently. From a simulation point of view, this a more complex problem that mimics network intensive networking applications, e.g. P2P online gaming and distributed virtual environments [13], P2P file sharing or wireless sensor networks.
- A preliminary evaluation of the impact of free riding [30] on data dissemination with different gossip protocols and network topologies is reported.
- To assess the performance of the considered dissemination protocols, we use a metric termed "overhead ratio", that measures the total number of delivered messages (for a given protocol) over the minimum number of messages needed to obtain a complete coverage (the lower bound), given the considered graph. The rationale behind this metric is to quantify the overhead, in terms of sent messages, for a communication protocol. It allows to compare the behavior of different dissemination strategies over different network topologies.
- Given a set of nodes in a P2P overlay, and the need to create a given overlay, an issue is how to set the network in order to guarantee certain communication properties. Our work permits to understand, during the design phase of a P2P overlay, how to set network parameters so as to obtain a certain overhead, that would guarantee a certain network coverage and delay. In most cases, this can be done only exploring the space of parameters of the available dissemination protocols. In practice, this requires the execution of a high number of simulations runs. This confirms, the need of scalable and efficient simulation tools.

The reminder of this paper is structured as follows. Section 2 presents some related works available in the literature. In Section 3, we discuss some background needed to understand the gossip algorithms and the performance assessment. Section 4 presents the considered dissemination algorithms. Section 5 presents the theoretical model and employs it to study the ability of the considered dissemination strategies to spread messages over an overlay network, given its degree distribution. In Section 6, the simulation testbed is described. Section 7 reports on simulation experiments we carried out. The main results from the performance evaluation are discussed more fully in Section 8. Finally, concluding remarks are reported in Section 9.

2. Related work

In this section, we review some works concerned with data dissemination in unstructured P2P networks. Since the considered schemes are based on gossip-style epidemic protocols, we focus on gossip-related approaches. Moreover, a main rationale for this choice is that gossiping poses some challenges when we try to guarantee a high (possibly full) network coverage and low delay, in scalable and highly intensive scenarios. Gossip is a simple, yet effective strategy to disseminate information. Its main feature is the use of randomization to propagate data. It has been proved that in certain contexts this provides better reliability and scalability than deterministic approaches [19].

Gossip-based communication can use either push, pull or push–pull schemes. According to a push based dissemination, it is the sender that decides which nodes will receive a message it is relaying through the network. Pull-based approaches let receivers trigger a communication with another node that will send some