Computer Networks 106 (2016) 109–121

Contents lists available at ScienceDirect

Computer Networks

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

Towards limited scale-free topology with dynamic peer participation

Xiaoyan Lu[a], Eyuphan Bulut[b,*], Boleslaw Szymanski[a,c]

[a] Department of Computer Science, Rensselaer Polytechnic Institute, Troy, NY, 12180, United States
[b] Department of Computer Science, Virginia Commonwealth University, Richmond, VA, 23284, United States
[c] Wroclaw University of Technology, 50-370 Wroclaw, Poland

ARTICLE INFO

Article history:
Received 5 April 2016
Revised 17 June 2016
Accepted 17 June 2016
Available online 20 June 2016

Keywords:
Scale-free networks
Peer-to-Peer networks
Scalable
Overlay networks

ABSTRACT

Growth models have been proposed for constructing the scale-free overlay topology to improve the performance of unstructured peer-to-peer (P2P) networks. However, previous growth models are able to maintain the limited scale-free topology when nodes only join but do not leave the network; the case of nodes leaving the network while preserving a precise scaling parameter is not included in the solution. Thus, the full dynamic of node participation, inherent in P2P networks, is not considered in these models. In order to handle both nodes joining and leaving the network, we propose a robust growth model E-SRA, which is capable of producing the perfect limited scale-free overlay topology with user-defined scaling parameter and hard cut-off. Scalability of our approach is ensured since no global information is required to add or remove a node. E-SRA is also tolerant to individual node failure caused by errors or attacks. Simulations have shown that E-SRA outperforms other growth models by producing topologies with high adherence to the desired scale-free property. Search algorithms, including flooding and normalized flooding, achieve higher efficiency over the topologies produced by E-SRA.

1. Introduction

In addition to the specific search strategies and resource allocation methods, the overlay topologies (i.e. logical connectivity graph) have significant impact on the performance of unstructured peer-to-peer (P2P) networks. It has been shown that the scale-free topology with the power-law distribution of node degrees is by definition well suited for P2P networks [1]. This is because such topology has a logarithmically scaled diameter [5] (ranging from $O(\ln N)$ to $O(\ln\ln N)$), is highly tolerant to random failures [13] and network congestion [15], and it also is highly synchronizable [14].

In the original Barabasi-Albert (BA) model [12], when a new node joins the network, its likelihood to connect to a node increases with this node's degree. This behavior, also known as "Preferential Attachment", generates a network with scale-free features. The BA model, however, does not consider some important features of P2P applications, including the hard cut-off on degree and dynamic peer participation [11]. The hard cut-off, which restricts the feasible topologies to limited scale-free networks, is required because users in P2P networks usually are not willing to serve as hubs (i.e. nodes with high degree) because of the high bandwidth required to serve the ensuing traffic. Moreover, users usually keep joining and leaving the network periodically, which affects the degree distribution in an unpredictable manner. Efforts have been made to build the limited scale-free overlay topology for P2P networks. Previous studies [6–8] have proposed cost-efficient growth models to construct the limited scale-free topologies with hard cut-offs with nodes constantly joining the network. However, the dynamic of nodes' removal [11] is another inherent property of P2P networks and it is not taken into account in these growth models. In previous studies [9,10], although nodes are allowed to leave, a precise scaling parameter can not be produced as specified.

Scale-free topologies are relatively robust in face of the removal of randomly chosen nodes but are very vulnerable to the removal of hubs [17,18]. In P2P networks, a selfish hub may quit and rejoin the network to avoid high communication costs, potentially distorting the degree distribution of the topologies. Admittedly, the degree distribution of a scale-free topology is not significantly influenced by infrequent node removal. However, when node removals are significant fraction of node additions over some period, the accumulated effects will eventually destroy the scale-free topology. In order to preserve the scale-free topology, a mechanism is needed to keep the power-law degree distribution regardless of the frequency of node removals.

One challenge to preserve a scale-free topology is to avoid using global knowledge because the communication cost required to obtain it grows linearly with the size of network. Yet, degree

* Corresponding author.
E-mail addresses: lux5@rpi.edu (X. Lu), ebulut@vcu.edu (E. Bulut), szymab@rpi.edu (B. Szymanski).

http://dx.doi.org/10.1016/j.comnet.2016.06.019
1389-1286/

distribution is a global property. In addition, when a node is removed, each of its neighbors has one connection terminated. Thus, scale-free feature must be controlled with a cost-efficient approach which does not require global information. On the other hand, it is important to preserve an exact scaling parameter for a power-law scale-free topology because the best performance of particular P2P application is achieved when a specific scaling parameter is chosen. But previous decentralized protocols can only produce overlays with an inexact and constant scaling parameter, depending on the size of the network.

Here, we introduce the **E**nhanced **S**emi-**R**andomized Growth **A**lgorithm (E-SRA), which preserves the power-law degree distribution in an overlay topology with nodes dynamically joining and leaving the network. It allows arbitrary nodes to be removed from a network and does not require collecting global information about the topology. Our approach assumes only local information is available, i.e. degree of the neighbors of the removed node. One advantage of our approach is that it provides partial tolerance to failing nodes. If a single node[1] fails due to attacks or errors, the neighbors are able to detect the failure and work to preserve the power-law distribution. In terms of the message complexity, one broadcast is needed to remove one node. If nodes randomly leave the network, the average number of the point-to-point messages sent by a single node to preserve the power-law degree distribution is linear to the maximum (also called hard cut-off) degree of the nodes. Combined with the growth model proposed in previous work [8], scale-free overlay networks can be maintained efficiently with nodes freely joining and leaving. As discussed in [8], it is important to adhere to exact scaling parameters because only then full advantage of properties such as slowly growing diameter hold. Moreover, capability to control post-construction parameters gives the flexibility to have a desired γ that will give the best efficiency in the search algorithm in use.

Simulations have shown that our approach can generate the perfect limited scale-free topologies with different patterns of adding/removing nodes while other models presented in the literature [6–8] have failed to handle node removal, especially when nodes with high degrees leave the network frequently. Moreover, the overlay topologies constructed by E-SRA provide better search performance in various settings; here, we considered search algorithms including flooding, where messages are forwarded to all neighbors, and normalized flooding, where messages are forwarded to k (i.e. the minimum degree) randomly selected neighbors.

The major properties of E-SRA include the following:

- Tolerance to dynamic peer participation: Nodes with any degree are free to leave and join the network while the power-law distribution is preserved.
- Partial tolerance to failures: the topology can recover from a single node failure as long as the neighbors of the failing node are alive at the moment of the failure.
- Scalability: no global information is required to add or remove a node.
- Flexible parameter settings: a topology with user-defined parameters (i.e. the minimum/maximum degree, scaling parameter γ) is created to ensure its optimal performance in applications such as search algorithms.

Section 2 introduces the previous work on growth models for limited scale-free topologies. Section 3 presents the detailed algorithms and its analysis. Section 4 describes simulation results. The discussion is presented in Section 5 and the conclusions are included in Section 6.

[1] We assume the nodes crash incrementally. When a group of connected nodes crash at the same time, some global restoration algorithm would be more appropriate.

2. Related work

The scale-free property is shown to exist in many natural or artificial systems, such as protein-protein interaction networks [19], the Internet [2], the World Wide Web [3], and scientific collaboration networks [4]. The degree distribution in these networks follows the power-law: $P(i) \sim i^{-\gamma}$, where $P(i)$ is the fraction of nodes with degree i and γ is the scaling parameter which varies between different types of networks ($2 \leq \gamma \leq 3$ in most cases). In a limited scale-free topology, only nodes with degrees smaller than the hard cut-off (i.e. the maximum) degree have degree distribution that follows the power-law.

The scale-free topology has some good properties, including high tolerance to random attacks [13], high synchronizability [14] and resistance to congestion [15]. For this reason, several growth models are proposed to construct the scale-free overlay topology. The BA model [12] manages to explain the evolution of scale-free topologies by a core principle named “Preferential Attachment”. But it is not practical in real P2P applications because the global information is required to maintain it. To address this issue, HAPA [6], Gaian [7], subPA [9] and SRA [8] algorithms were introduced to construct the scale-free overlay topology with partial or no global information.

“Preferential Attachment” [12] means a new node is more likely to connect to heavily linked nodes when it joins the network. The BA model has some disadvantages as the growth model for the overlay topology of P2P networks. Firstly, it does not provide hard cut-offs. Since a heavily linked node uses a lot of bandwidth in P2P networks, nodes usually are not willing to maintain high degrees. For this reason, a user-defined hard cut-off (i.e. the maximum) degree is imposed lower than the natural cut-off arising in the BA model. The imposed hard cut-offs restrict the feasible overlays to the limited scale-free topologies, which are more practical. Moreover, the BA model also requires the global information about the topology to add connections when a new node joins. In real-world applications, however, the communication cost of obtaining the global information is prohibitive. Therefore, a distributed approach that constructs the topology without global information is desired.

In [6], authors study the construction of limited scale-free overlay topologies for unstructured peer-to-peer networks. In the Hop-and-Attempt Preferential Attachment (HAPA) algorithm [6], a new node joining the network connects to k (i.e. the minimum degree) nodes in a random route starting from a randomly chosen node. This scheme works because high degree nodes are more likely to occur in a random route than nodes with low degree. The hard cut-off is used to avoid “superhubs”, which are nodes with degrees linear to the network size. HAPA algorithm produces the topologies with degree distribution approximately following the power-law with scaling parameter $\gamma = 3$.

Gaian [7] algorithm is proposed for distributed database systems whose efficiency and reliability depends on the overlay topology. In this algorithm, a new node broadcasts a message when it joins the network using computing with time principle [20]. Each receiver computes the maximum time of delay, t_v, which is proportional to the inverse of its degree, and chooses the time of delay t_d in the interval $[0, t_v]$ uniformly randomly. Instead of replying to the sender instantly, the receiver waits the time of delay t_d and then replies. In this way, nodes with higher degrees are likely to wait shorter period. The new node connects to the first k responders. It gives a better chance to the new node to connect to nodes with high degrees. This mechanism, which reduces the communication overhead by allowing nodes to self-select themselves according to