MODELING PEER-TO-PEER NETWORK TOPOLOGIES THROUGH "SMALL-WORLD" MODELS AND POWER LAWS

Mihajlo Jovanović ECECS Department, University of Cincinnati Cincinnati, OH 45221

I INTRODUCTION

The recent emergence of novel network applications such as Gnutella, Freenet, and Napster has reincarnated the familiar peer-to-peer (P2P) architecture model of the original Internet in new and innovative ways in an effort to facilitate worldwide sharing of information [1]. As a result, there is an everincreasing need for distributed protocols that would allow peer-to-peer applications to scale to a large community of users. The main difficulty in designing such protocols is that currently, very little is known about the nature of networks on which they would be operating. Current protocols were designed without any knowledge about the underlying network topology. The end result is that even simple protocols, as in the case of Gnutella, result in complex interactions that can adversely affect system performance.

To study these interactions, we first need an accurate model of the network topology. We model the topology of P2P networks by undirected graph G, where nodes represent hosts and edges represent connections between those hosts. In this paper we point out several important structural characteristics of the P2P network topology graph, namely the "smallworld" properties and several power law distributions of various graph metrics. We report measurements through experimental studies of a large P2P network application known as Gnutella. To obtain the Gnutella topology data, a network crawler that allows topology discovery to be performed in parallel was developed. Upon analysing the obtained topology data, we discovered it exhibits strong "small-world" properties. More specifically, the properties of small diameter and clustering were observed. In addition, we report evidence of four different power laws previously observed in other technological networks, such as the Internet and the WWW. Recent research results show that the "small-world" and power-law properties of the underlying network topology can significantly impact the performance of algorithms such as those for routing and searching [2, 3]. Therefore the existence of these properties in P2P networks presents an important issue to consider when designing new, more-scalable application-level protocols.

II "SMALL-WORLD" PROPERTIES

The term "small-world" originated with a famous social experiment conducted by Stanley Milgram in the late 1960s. Watts and Strogatz in [4] showed that other networks, such as those occurring in nature and technology, also exhibit "small-world" behavior. Throughout this work we use the term "small-world" loosely to mean that the network possesses both small diameter and is also highly clustered.

Upon analysing the Gnutella network topology data obtained by our crawler, we discovered both the small diameter and the clustering properties characteristic of "small-world" networks. To test for clustering, we generalised the definition for the clustering coefficient proposed in [4] to allow us to measure clustering at different levels of depth:

Definition 1. Characteristic coefficient $C(l)_v$ of vertex v is calculated by dividing the number of cross edges in a BFS-tree of depth l and rooted at v, by the maximum possible

number of cross edges given by $\binom{k}{2} - (k-1)$, where k is the

number of vertices in the BFS-tree. Clustering coefficient C(l) of a graph is defined as the average of $C(l)_v$ over all vertices v.

For the characteristic path length we used the definition used by Watts and Strogatz as follows:

Definition 2. Characteristic path length L of a graph G is defined as the number of edges in the shortest path between two vertices u and v, averaged over all pairs of vertices.

	Gnutella	BA	WS	G(n,p)	2D torus
11/13	0.0224	0.015	0.0373	0.004	0.0606
11/16	0.0089	0.0096	0.0372	0.0025	0.0606
12/20	0.0301	0.0179	0.0537	0.0062	0.0606
12/27	0.0206	0.0185	0.0539	0.0062	0.0606
12/28	0.0207	0.0174	0.0536	0.0056	0.0606

Table 1. Clustering coefficients for the Gnutella,Barabási-Albert, Watts-Strogatz, random graph, and the2D torus topologies (*l=2*)

	Gnutella	BA	WS	G(n,p)	2D mesh
11/13	3.7230	3.4749	4.5971	4.4873	20.6667
11/16	4.4259	4.0754	4.6116	5.5372	21.3333
12/20	3.3065	3.1902	4.2249	3.6649	22.0000
12/27	3.3036	3.1805	4.1917	3.7100	21.3333
12/28	3.3282	3.2075	4.2520	3.7688	22.6667

Table 2. Characteristic path length for the Gnutella,Barabási-Albert, Watts-Strogatz, random graph, and the2D mesh topologies

The statistics for the clustering coefficient and the characteristic path length for five different snapshots of the Gnutella topology obtained during the months of November and December of 2000 are presented in tables 1 and 2. The values for the Gnutella topology graphs are benchmarked



against two widely used "small-world" models, the Watts-Strogatz [4] and the Barabási-Albert model [5], the random graph and the 2-D torus. The parameters for these models were chosen in a way so that the number of nodes and average degree of the resulting graph is approximately equal to that of the original Gnutella topology. For random graphs, average values out of 100 trials are shown.

As can be seen, all of the Gnutella topology snapshots demonstrate the "small-world" phenomenon: characteristic path length is comparable to that of a random graph, while the clustering coefficient is an order of magnitude higher. We report consistent results for the clustering coefficient with l =3. These results clearly indicate strong "small-world" properties of the Gnutella network topology. We argue that this is an important issue to consider when modelling P2P networks such as Gnutella. More specifically, an accurate P2P model must inevitably generate topologies exhibiting the described "small-world" properties. Our discovery can also aid in analysing and predicting the performance of algorithms such as those for routing and searching operating on such topologies. For example, Gnutella's current broadcast routing strategy is clearly not likely to work well on a clustered topology of a "small-world" network, as it would generate large amounts of duplicate messages. This would result in poor utilisation of network bandwidth and hinder scaling - a phenomenon observed in practice during the summer of 2000.

III POWER-LAWS

The major limitation of the described "small-world" models is due to increasing evidence of various power-laws of the form $y = x^a$, governing distribution of various graph metrics for many large, self-organising networks [5, 6].

The authors in [6] discovered four of these power-laws characterising topology of the Internet at both inter-domain and router level. These power-laws are defined as follows:

Power-Law 1 (rank exponent \mathcal{R}): The outdegree d_v of a node v is proportional to the rank of the node r to the power of a constant \mathcal{R} : $d_v \propto r_v^{\mathcal{R}}$. The rank r of a node v is defined as its index in the order of decreasing outdegree.

Power-Law 2 (out-degree exponent *O*): The frequency f_d of the out-degree *d* is proportional to the out-degree to the power of a constant $O: f_d \propto d^O$.

Power-Law 3 (hop-plot exponent \mathcal{H}): The total number of pairs of nodes P(h) within h hops is proportional to the number of hops to the power of a constant \mathcal{H} : $P(h) \propto h^{\mathcal{H}}$, when $h \ll \delta$, the diameter. The number of pairs P(h) is the total number of pairs of nodes within less or equal to h hops, including self-pairs, and counting all other pairs twice.

Power-Law 4 (eigen exponent \mathcal{E}): The eigenvalues λ_i of a graph are proportional to the order *i* to the power of a constant \mathcal{E} : $\lambda_i \propto i^{\mathcal{E}}$.

Upon analysing the Gnutella topology data obtained using our network crawler, we discovered it obeys all four of these power-laws. We present the results for power-laws 2 and 3 in figures 1 and 2, respectively. The charts for the other two power-laws are omitted due to space constraints.

Power-laws relationships between variables are typically plotted on a logarithmic scale, since their plot should, by definition appear linear. Power-law exponents can then be viewed as the slope of this linear plot. Linear regression was used to fit a line in a set of two-dimensional points using the least-square error method. To quantify the quality of the approximation, with each figure we included the absolute value of the correlation coefficient r ranging between -1 and 1. An |r| value of 1 indicated perfect linear correlation.

Power-law 2 is the one that is the most frequently cited in the recent studies of large network topologies. Figure 1 shows this power-law holds for the Gnutella topology with out-degree exponent O = -1.4 and a high correlation coefficient of 0.94. For comparison, plots for both a snapshot of the Gnutella network topology and a simple connected random graph of the same size are presented. For the random topology, this property clearly doesn't hold.

We must remark that a group called Clip2 independently discovered this particular power-law for the Gnutella network topology. However they reported the power-law exponent of-2.3, in disagreement with our result.



hops (power-law 3) for four snapshots of the Gnutella topology

We believe the reason for this discrepancy is due to the fact that our results are based on the network crawls performed during December of 2000, while the other result dates back to the summer of the same year. Since that time, the Gnutella network has undergone significant changes in terms of its structure and size. While the values of the node degree exponent for all of the Gnutella topology instances obtained during the month of December are consistently around -1.4, we have observed values of -1.6 for the data obtained in November. This may be taken as indication of a highly dynamic, evolving state of the Gnutella network.

It has been shown that power-laws 3 and 4 hold for almost all types of topologies, including random, regular, and hierarchical [7]. Power-law three by definition holds for regular topologies such as a ring topology and a 2-D mesh, with hop-plot exponents of 1 and 2, respectively, for $h \ll \delta$. It is therefore not surprising that we have also observed these power-laws in the Gnutella network topology. However a case has been made that, while the mere presence of these two power-laws is not a distinguishing property of a graph, the values of their exponents can be. To show this, we compare the hop-plots for four different Gnutella topology snapshots. For each one, we approximated only the first four

hops. Clearly, power-law 3 holds for all four snapshots with very high correlation coefficients of 0.99. More importantly, the hop-plot exponents seam to be clustered tightly around the value of 3.5. Notice that this value lies right between the exponent values reported for the inter-domain and router level topology instances of the Internet [6]. Like the authors in [6, 7], we must concede that the results for this particular power-law may be misleading given such small number of data points. This limitation is imposed by the fact that these graphs have a small diameter.

An example application of power-law 3 that seams particularly applicable to Gnutella was suggested in [4]. They introduced a concept of the effective diameter δ_{ef} , which is essentially the number of hops required to reach a "sufficiently large" portion of a network. In other words, any two nodes are within hops of each other with high probability. We present the definition below for convenience.

Definition 3. (effective diameter) Given a graph G with N nodes, E edges, and \mathcal{H} hop-plot exponent, the effective diameter, δ_{ef} , is defined as:

$$\delta_{ef} = \left(\frac{N^2}{N+2E}\right)^{1/p}$$

Substituting the values for the Gnutella topology snapshot from December 28, 2000, we get that, during that time, a more cost-effective value for the maximum TTL would have been 4 (instead of 7, which is the default specified by the Gnutella protocol).

IV CRAWLER ARCHITECTURE

Gnutella is a highly dynamic network in which topology is constantly changing as hosts join and leave the network, establish new connections, and close the existing ones. Therefore discovering topology of the Gnutella network really means taking a snapshot of the topology at a specific point in time. Clearly, this imposes an additional requirement for any topology discovery algorithm to be efficient, since the accuracy of a topology map is inversely proportional to the time used to obtain it. To address this requirement, we have designed our network crawler to operate in parallel under the distributed computing model. We felt that the task at hand would be perfectly suited for a distributed model, since it requires very little inter-processor communication. In fact, in our design, communication only occurs at the beginning of the process, to distribute input, and at the end, to gather the output at a central location. The crawler was implemented using Java remote method invocation (RMI) and deployed on a network of workstations (NOW). In our implementation, crawling a subset of the Gnutella network is provided as a service residing on various remote locations throughout our network. In other words, our crawling algorithm is implemented as a distributed object residing on remote machines. Our system includes an object serving as the "brain" of the entire computation. This central object is responsible for "bootstrapping" the entire topology discovery process by distributing the initial list of Gnutella hosts to other remote objects. Upon receiving the input, each remote object performs topology discovery of its portion of the network, and subsequently returns a graph object representing network topology to the central object. The central object is then responsible for merging all the output graphs into a single one representing topology of the entire Gnutella network. Given the fact that, at the time of our experiments, the size of the Gnutella network rarely exceeded two thousand hosts, in practice, we have been able to obtain the complete topology data in constant time using only a handful of processors. Visualisations of the Gnutella network topology using data obtained by our crawler are omitted due to space constraints.

V CONCLUSIONS

Modeling complex graph structures produced by modern P2P network applications is a difficult task. The main contribution of this work to the task at hand is in identifying two important ingredients of a uniform topology model for P2P networks, namely the "small-world" and the power-law properties. Recently we were able to utilise our topology model in order to study the effects of heterogeneous latencies on message reachability in distributed networks operating under flooding protocols. Specifically, we showed that heterogeneous latencies present in many large-scale P2P network applications, when combined with the standard protocol mechanisms of time-to-live (TTL) and unique message identification (UID) used to govern flooding message transmissions, can potentially have a devastating effect on the reachability of message broadcast. We call this combined effect "short-circuiting," Detailed measurements obtained through both network simulation studies and experimental studies can be found in [8]. Our results support the conclusion that, on average, the real effects of shortcircuiting are significant, but not devastating to the performance of an overall large-scale system. It is our hope that these results, as well as the ones presented in this paper, will prove useful in designing the next generation of application-level protocols for P2P network applications.

REFERENCES

- [1] The Gnutella Homepage, http://gnutella.wego.com.
- [2] Kleinberg, J., The small-world phenomenon: An algorithmic perspective, Technical Report 99-1776, Cornell University Computer Science Dept, Oct 1999.
- [3] Adamic, L. A., and Lukose, R. M., Puniyani, A. R., and Huberman, B. A., Search in power-law networks, http://www.parc.xerox.com/istl/groups/iea/papers/plsearc h, March, 2001.
- [4] Watts, D. J. and Strogatz, S. H., Collective dynamics of small-world networks, *Nature*, 393:440-442, June 1998.
- [5] Barabási, A. and Albert, R., Emergence of scaling in random networks, *Science*, 286:509-512, October 15, 1999.
- [6] Faloutsos, M., Faloutsos, P., and Faloutsos, C., On power-law relationships of the internet topology, In *SIGCOMM*, pages 251-262, 1999.
- [7] Medina, A., Matta, I., and Byers, J., On the origin of power laws in internet topologies, ACM Computer Communications Review, 30(2), April 2000.
- [8] Annexstein, F. S., Berman, K. A., and Jovanović, M., Latency effects on reachability in large-scale peer-to-peer networks, In *Thirteenth ACM Symposium on Parallel Algorithms and Architectures*, Crete Island, Greece, July 2001.

Acknowledgement: Research supported in part by NSF Grant No. CCR-9877139 and the University of Cincinnati Research Assistantship.

Sadržaj: U radu su predstavljene značajne strukturne karakteristike P2P mreža, tzv. "small-world" karakteristike i više stepenih raspodela grafičkih veličina, kroz istraživanje na primeru Gnutella mreže. Za sakupljanje topologičkih podataka mreže razvijen je kod koji omogucava paralelno odredjivanje topologije mreže. Nakon sto je topologija mreže odredjena, utvrdjeno je da ona poseduje, tzv. "small-world" osobine. Preciznije rečeno, utvrdjeno je da topologiju mreže karakteriše relativno mali prečnik i tendencija grupisanja. Pored toga, primećeno je prisustvo četiri različite stepene raspodele, koje su ranije uočene i u drugim tehnološkim mrežama, kao na primer na Internetu i na WWW. Autor smatra da su ova otkrića značajan korak ka uniformnom modelu P2P mrežnih aplikacija, koja će potencijalno doprineti analiziranju performansi postojećih algoritama, kao i dizajniranju novih, efikasnijih rešenja.

MODELOVANJE TOPOLOGIJE PEER-TO-PEER MREŽA KROZ "SMALL-WORLD" MODELE I ZAKONE STEPENE RASPODELE, Mihajlo Jovanović