

On the Structure of Unstructured Overlay Networks*

João Leitão
INESC-ID/IST
jleitao@gsd.inesc-id.pt

José Pereira
University of Minho
jop@di.uminho.pt

Luís Rodrigues
INESC-ID/IST
ler@ist.utl.pt

Abstract

Unstructured overlay networks are a key component of many peer-to-peer systems. These overlays exhibit a set of interesting properties that derive from their inherent randomness. In this fast abstract we briefly discuss the key aspects that need to be considered when attempting to bias the structure of unstructured networks, such that it becomes possible to improve the efficiency of applications and services executed at the top level, without impairing the correctness of the overlay.

1. Introduction

Unstructured, or random, overlay networks [2, 5] are a key component of many peer-to-peer systems. These logical overlays are built by establishing random neighboring relations across nodes.

Unstructured overlays networks are interesting as they present a low cost to build and maintain; moreover this cost is usually evenly shared among all nodes in the system. Also these overlays usually present a natural level of redundancy which makes them more resilient, both to node failure and also to message losses, as messages can also be redundantly transmitted among the several overlay links across nodes. In fact, it has been shown that it is possible to build random overlays that are able to remain connected even in face of node failures in the order of 80% [4].

On the other hand, pure random overlay networks do not allow to exploit the properties of the physical network (cartesian proximity), or of the content stored by the nodes (content proximity), among others. If these properties are taken into account, the performance of the applications that are run on top of these overlays may be substantially improved; in fact, the topic of improving topologies for gossip-based protocols was pointed as a relevant research area [1]. In this paper we discuss key aspects and guidelines for

adding some degree of structure to unstructured overlay networks.

2. Key Overlay Properties

We initiate our discussion by identifying some key properties of random overlays that must be preserved, even when techniques to bias its structure are applied.

Connectivity: the overlay is connected if there is a path that allows every node to reach every other node. This is the most important property as it ensures that nodes can rely on the overlay to communicate.

Uniform Degree Distribution: the degree of a node is the number of edges (or neighbors) of the node. This is a measure of both reachability and contribution of the node to maintain the overlay. To ensure efficiency and good fault tolerance properties, all nodes in the system should present a uniform degree.

Low Clustering Coefficient: the clustering coefficient of a node is the number of edges between that node's neighbors divided by the maximum possible number of edges across those neighbors. In order to avoid excessive redundant communication between nodes, but also to lower the diameter of the overlay and ensure high fault resilience, this value should be as low as possible.

For more details on these, and other related properties, the reader can refer to [3] or [4].

3. Introducing Structure

We now present a number of guidelines for building adaptive mechanisms that aim at introducing some degree of structure in random overlay networks, in order to ensure their scalability and to preserve the connectivity of the overlay.

– For scalability, each adaptation step should involve a *limited number of nodes*.

– Moreover, adaptation should *use limited information* which contributes both to lower the overhead of the protocol, reducing the communication required between peers, but also to ensure the scalability of the system.

*This work was partially supported by project "P-SON: Probabilistically Structured Overlay Networks" (POS_C/EIA/60941/2004).

– The adaptation mechanisms should not disrupt the overlay in face of *limited or incorrect information*. Notice that the scale of peer-to-peer systems can reach several thousands of nodes, meaning that some degree of inconsistency in the information gathered should be expected.

– The adaptation mechanisms should *strive to maintain the node degree* of all nodes involved in each topology adaptation, which ensures that the connectivity of each node in the overlay is constant.

– The adaptation mechanisms should *promote the overlay stability* which can be achieved by two distinct strategies: *i*) only bias nodes which are fully connected (*i.e.* that already have the target number of nodes) and, *ii*) reducing the number of topology adaptations to those that effectively contribute to improve the overlay.

– *Every node should maintain some unbiased neighbors*, allowing the overlay to improve itself while maintaining an acceptable low clustering coefficient. Note that low clustering coefficient is closely related with the randomness of unstructured overlays and can be lost if a node applies a bias to all its neighbors. This is expected, as usually nodes which become neighbors due to biasing of the topology, will probably select the same peers as neighbors, which increases the clustering of the overlay.

Finally, and because these overlays should be general-purpose (in the sense that several distinct applications and services can be executed on top of them), protocols which adapt their topology should also be flexible allowing, with minimal modifications, to apply a bias based on different criteria such as, link latency or content similarity.

4. Illustration




Figure 1. Clustering Coefficient

We now illustrate the use of the guidelines listed above to

develop adaptive mechanisms that strive to introduce some structure in an unstructured overlay network. In this example, the goal is to lower the overall cost of the overlay links (the cost of links is provided by a companion oracle).

In our prototype nodes maintain the list of neighbors ordered, biasing only their lowest cost neighbors (we always maintain the highest cost nodes). It operates by having periodically one node starting an *optimization cycle* in which 4 nodes try to switch 2 overlay links between them, ensuring that the overall cost of the links drops whenever a bias is effectively applied.

Figure 1 shows the evolution of clustering coefficient on the overlay, in a system with 10.000 nodes, where each node keeps at most 5 neighbors. Values are depicted for different numbers of unbiased neighbors from 0 (none) to 5 (all). By keeping a small number of unbiased neighbors (2 or 3) the cost of the overlay can be lowered to less than 75% of the cost of a pure random overlay, while maintaining a low clustering coefficient.

5. Conclusion

In this fast abstract we addressed the issue of introducing structure in unstructured overlay networks. We have enumerated key overlay properties that need to be preserved and listed a number of guidelines that should be followed when building adaptive mechanisms to introduce structure. By following these guidelines one can significantly improve the topology of these overlays without impairing their key properties. This ensures the efficiency and the correctness of applications and services that use the overlay.

References

- [1] K. Birman. The promise, and limitations, of gossip protocols. *SIGOPS Oper. Syst. Rev.*, 41(5):8–13, 2007.
- [2] A. J. Ganesh, A.-M. Kermarrec, and L. Massoulié. SCAMP: Peer-to-peer lightweight membership service for large-scale group communication. In *Networked Group Communication*, pages 44–55, 2001.
- [3] J. Leitão. Gossip-based broadcast protocols. Master’s thesis, University of Lisbon, 2007.
- [4] J. Leitão, J. Pereira, and L. Rodrigues. HyParView: A membership protocol for reliable gossip-based broadcast. In *DSN ’07: Proc. of the 37th Annual IEEE/IFIP Intl. Conf. on Dependable Systems and Networks*, pages 419–429, Edinburgh, UK, 2007. IEEE Computer Society.
- [5] S. Voulgaris, D. Gavidia, and M. Steen. Cyclon: Inexpensive membership management for unstructured p2p overlays. *Journal of Network and Systems Management*, 13(2):197–217, June 2005.