

A Backward-Compatible Protocol for Inter-routing over Heterogeneous Overlay Networks

Giang Ngo Hoang, Luigi Liquori, Vincenzo Ciancaglini, Petar Maksimovic,
Hung Nguyen Chan

► **To cite this version:**

Giang Ngo Hoang, Luigi Liquori, Vincenzo Ciancaglini, Petar Maksimovic, Hung Nguyen Chan. A Backward-Compatible Protocol for Inter-routing over Heterogeneous Overlay Networks. 28th Annual ACM Symposium on Applied Computing, SAC 13, Coimbra, Portugal, ACM press, 2013., Mar 2013, coimbra, Portugal. ACM, pp.649-651, 2013, <10.1145/2480362.2480485>. <hal-00906835>

HAL Id: hal-00906835

<https://hal.inria.fr/hal-00906835>

Submitted on 13 May 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Backward-Compatible Protocol for Inter-routing over Heterogeneous Overlay Networks

Giang Ngo Hoang¹²
giang.ngo_hoang@inria.fr

Luigi Liquori³
Luigi.Liquori@inria.fr

Vincenzo Ciancaglini¹³
Vincenzo.Ciancaglini@inria.fr

Petar Maksimovic^{14*}
Petar.Maksimovic@inria.fr

Hung Nguyen Chan⁵
hungnc@vielina.com

¹University of Nice Sophia-Antipolis, France

²Hanoi University of Science and Technology, Vietnam

³The National Institute for Research in Computer Science and Control, France

⁴Mathematical Institute of the Serbian Academy of Sciences and Arts

⁵Vietnam Research Institute of Electronics, Informatics and Automation, Vietnam

ABSTRACT

Overlay networks are logical networks running on the highest level of the OSI stack: they are applicative networks used by millions of users everyday. In many scenarios, it would be desirable for peers belonging to overlays running different protocols to communicate with each other and exchange certain information. However, due to differences in their respective protocols, this communication is often difficult or even impossible to be achieved efficiently, even if the overlays are sharing common objectives and functionalities. In this paper, we address this problem by presenting a new overlay protocol, called OGP (*Overlay Gateway Protocol*), allowing different existing networks to route messages between each other in a backward-compatible fashion, by making use of specialized peers joined together into a super-overlay. Experimental results on a large scale Grid5000 infrastructure show that having only a small number of nodes running the OGP protocol is sufficient for achieving efficient routing between heterogeneous overlay networks.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: [Network communications, Network topology]; C.2.2 [Network Protocols]: [Routing protocols]; C.2.6 [Inter-networking]: [Routers, Standards].

General Terms

Algorithms, Performance, Experimentation.

Keywords

Inter-overlay routing, Backward compatibility, Overlay protocols, Network co-operation, Broadcast, Multicast.

1. INTRODUCTION

Nowadays, many distributed applications, such as those involving cloud computing, peer-to-peer networks, and Voice over IP (VoIP) applications are built on top of various overlay networks, which differ from each other in many aspects. For one, the topology of overlays such as Chord [1] and Kademlia [2] is structured, while that of Gnutella or Freenet [3] is unstructured. Overlay networks can also use different types of message routing, querying - simple, wildcard, or range querying - and different algorithms for

*This work has been partially supported by the Ministry of Science, Technology and Development, Republic of Serbia, projects ON174026 and III044006.

encoding messages, as well. These variations, as a rule, result in the overall incompatibility of overlay networks. Conversely, there exist clear advantages of having different overlay networks co-operate, such as increased search space, simple localization of participating overlays, and easily achievable content redundancy.

Several solutions for co-operation of overlays have been proposed. Some of them merge same-type overlays into a larger one, while others introduce dedicated gateways as bridges across overlays, or exploit co-located nodes to forward messages between overlays. However, these solutions are either costly in terms of power processing or generated traffic, or suffer from the single-point-of-failure problem. Additionally, most of them require some modification of already existing peers, which is usually near to impossible.

Our main contribution is the Overlay Gateway Protocol (OGP) - a new protocol for efficient inter-overlay co-operation. Our approach is inspired by the Border Gateway Protocol (BGP) which allows routing among Autonomous System (AS) in the Internet. OGP enables routing among independent overlay networks, and is run by a small number of peers from each of the standard overlay networks, in addition to their native protocols. These peers form a super-overlay we call the *OGP overlay*, and is equipped with efficient algorithms for unicast, broadcast, and multicast of messages from one standard overlay to the rest. Peers in the OGP overlay can reach across standard overlays of which they are not members, and act as gateways for other peers, such as peers created for taking advantage of OGP, as well as native peers unaware of OGP. Our approach ensures backward-compatibility, meaning that native peers can continue to operate normally, without any modifications, in the presence of OGP.

2. OGP PROTOCOL OVERVIEW

2.1 Classification of peers

In OGP, the peers are classified into the following categories:

Blind peers. Blind peers are peers belonging to only one standard overlay, are not aware of the existence of the OGP protocol, and use only the protocol native to their overlay.

Full OGP peers. Full OGP peers simultaneously belong to one standard overlay and the OGP overlay. They run both the native protocol of the standard overlay they belong to, and the OGP protocol of the OGP overlay. They route messages from one standard overlay to another via the OGP overlay and can also serve as gateways for lightweight OGP peers and some kinds of blind peers.

Lightweight OGP peers. Lightweight OGP peers take advantage of the inter-overlay routing provided by OGP overlay. They belong only to one standard overlay and do not participate in the OGP overlay but also keep a list of full OGP peers. They run both the native protocol of the standard overlay they belong to, and the lightweight OGP protocol for communicating with full OGP peers.

2.2 Scenarios of usage

In Figure 1, three scenarios are shown to illustrate the routing of three lookup queries, in which full OGP peers, lightweight OGP peers and blind peers interact in order to reach across overlays using the OGP super-overlay. The three smaller ovals represent standard overlays, while the largest oval represents the OGP super-overlay, forwarding messages back and forth between standard overlays. The black squares B, C, G, N and P represent full OGP peers, the black circles A, D and F represent lightweight OGP peers, while the white circles E, H, and M represent blind peers. Solid lines represent requests, while dashed lines represent responses.

First scenario. In the first scenario, the full OGP peer P is looking for some information located on the blind peer M: it forwards, using the OGP protocol, the message to N which is a full OGP peer in Overlay 3. Upon receiving the message, N converts the message in accordance with the native protocol of Overlay 3, and forwards it to M. The return path takes us back through N and then to P, following the native protocol of Overlay 3 first, and then the OGP protocol.

Second scenario. In the second scenario, the lightweight OGP peer A, belonging to Overlay 1, is looking for some information which is located at the lightweight OGP peer D in Overlay 2: it looks up its list of full OGP peers and then forwards the message to B, using the OGP protocol. B then forwards the request asking for this information to the full OGP peer C, belonging to Overlay 2, via the OGP overlay. Upon receiving the message request, C reconstructs the request to be in accordance with the possibly different format defined by the native protocol of Overlay 2, and then forwards it to D. Then D sends the response back to C, which reconstructs the response to be in accordance with the format of the OGP protocol, and finally sends it back to A via B.




Figure 1: The OGP Topology.

Third scenario. In the third scenario, the blind peer E of Overlay 2 is looking for some information contained at the blind peer H of Overlay 3. E launches this request on Overlay 2 which then touches the lightweight OGP peer F by chance. F sends a request, using the OGP protocol, asking full OGP peer G to lookup the information, which, in turn, contacts the blind peer H. Between G and H, the routing process follows the native protocol of Overlay 3. Upon receiving and reconstructing the response, G forwards it to peer F, via the OGP overlay, and finally the routing proceeds to E, following the native protocol of Overlay 2. Here, the native protocols must be such that their request is sent clear, i.e the information within the request which is required for the request reconstruction happen in lightweight OGP peer is either available in clear, or can be recovered through a decryption operation of some sort.

The above scenarios provide an insight into how OGP peers, lightweight OGP peers and blind peers whose requests are made in clear, can exploit the routing infrastructure of the OGP super-overlay across heterogeneous overlays.

2.3 Inter-routing schemes

The OGP protocol provides three schemes of inter-routing between standard overlays: OGP unicast, OGP multicast and OGP broadcast. A full OGP peer can use any of these schemes.

OGP Unicast. In the OGP unicast scheme, the full OGP peers route requests into only one standard overlay different from the one the request originated from.

OGP Multicast. With the multicast strategy, we can selectively choose multiple destination overlays who will receive the request according to the needs of the distributed application. For instance, in the first scenario, by choosing Overlay 2 and Overlay 3 as destination, the routing could proceed, in addition to the unicast path, also via the full OGP peer C into Overlay 2, possibly with a different result. All of the results are returned to the original sender.

OGP Broadcast. With the broadcast strategy, all standard overlays are chosen as destination, and all of the obtained results are returned to the sender, just like with the multicast. For example, in the second scenario, the routing could proceed, in addition to the unicast path, also via the full OGP peer N into Overlay 3.

3. CONCLUSIONS AND FUTURE WORK

In this paper, we have introduced a new Kademlia-like protocol, named OGP, which provides new and efficient unicast, multicast and broadcast algorithms for inter-routing between heterogeneous overlays. It is used within the OGP super-overlay, which is designed to integrate seamlessly with existing overlays, while preserving backward-compatibility.

The experiments performed on the French Grid5000 (G5K) platform have shown that the OGP super overlay routes a message from one overlay to others with a high successful rate and low path length. The drawback of our model is the need to maintain the OGP overlay, in the sense that the traffic generated by a full OGP peer on it is much larger than the traffic generated by a standard peer.

Our further work on this topic includes a thorough investigation of the efficiency of the OGP protocol on much larger systems, and an adaptation of OGP for interconnection of streaming systems.

4. REFERENCES

- [1] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup service for internet applications," in *In proc. of SIGCOMM*, 2001.
- [2] P. Maymounkov and D. Mazières, "Kademlia: A Peer-to-peer Information System Based on the XOR Metric," in *Proc. of IPTPS*, 2002.
- [3] The freenet project website. [Accessed November 19, 2012]. [Online]. Available: <https://freenetproject.org/>.