Reconnecting the Internet with ariba: Self-Organizing Provisioning of End-to-End Connectivity in Heterogeneous Networks

Christian Hübsch, Christoph P. Mayer, Sebastian Mies,
Roland Bless, Oliver P. Waldhorst, and Martina Zitterbart
Institute of Telematics – Universität Karlsruhe (TH) – Germany
{huebsch,mayer,mies,bless,waldhorst,zit}@tm.uka.de

ABSTRACT

End-to-End connectivity in today's Internet can no longer be taken for granted. Middleboxes, mobility, and protocol heterogeneity complicate application development and often result in application-specific solutions. In our demo we present *ariba*: an overlay-based approach to handle such network challenges and to provide consistent homogeneous network primitives in order to ease application and service development.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Network communications*

General Terms

Design, Experimentation

Keywords

Overlays, Peer-to-Peer Networking, Heterogeneity

1. INTRODUCTION

The evolution of the Internet complicates establishing end-to-end connections between arbitrary devices through restricted connectivity by middleboxes (NATs, firewalls, etc.), or growing heterogeneity induced by new protocols like IPv6. Furthermore, established connections are hard to maintain in face of increasingly present device mobility. Several solutions are available to re-enable end-to-end connection establishment (e. g., NAT traversal mechanisms, MobileIP, application-specific overlays), but imply several disadvantages: First, they are not self-contained and therefore require additional dedicated components like home agents, rendezvous or relay servers. Second, they are not self-organizing and require manual configuration. Third, they do not consider protocol and network heterogeneity (for example IPv4, IPv6, or future protocols) but assume direct end-to-end connectivity.

In recent years, overlay-based applications have become popular since they can add functionality missing in the Internet without infrastructure support. Such overlays require establishing end-to-end connections between all participating nodes. Current applications either provide custom solutions to establish end-to-end connections by handling middleboxes, mobility, or heterogeneity, while others cannot be deployed in such networks at all.

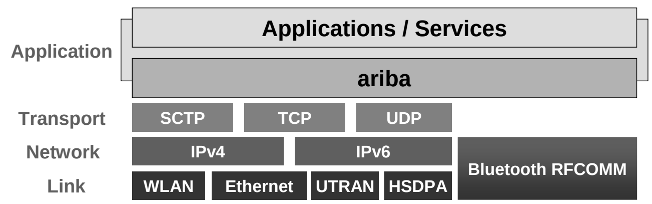

Figure 1: Integration of *ariba* into the network stack

To ease creation of overlay-based applications we present *ariba*: a generic solution to provide consistent per-application end-to-end connectivity employing identifier-based addressing [1]. *ariba* is a part of the *Spontaneous Virtual Networks* (SpoVNet) architecture [4], enabling spontaneous and flexible creation of overlay-based applications and services on top of heterogeneous networks. It is fully self-organizing and self-configuring, end-system-based, and does not require infrastructure support. On the one hand, *ariba* provides a homogeneous, mobility-invariant network substrate to application developers [2] (cf. Figure 1). On the other hand, self-configuration relieves end-users from error-prone manual configuration. Furthermore, *ariba* is adaptive, i. e., it can handle changing network settings. *ariba* itself uses an overlay to provide basic connectivity over heterogeneous networks. Applications can easily build new (transport) overlays on top of *ariba* without dealing with network challenges like NAT, mobility, or heterogeneity.

Our demonstration illustrates *ariba*'s main feature: application-specific provision of end-to-end connectivity in face of heterogeneous networks with mobile devices, easing both development and deployment of overlay-based applications. The demonstrator shows that *ariba* (re-)establishes application-specific end-to-end connectivity despite middleboxes, dynamic network changes, node mobility, and protocol heterogeneity.

2. OBJECTIVE OF THE DEMO

We used *ariba* to implement a group chat application based on an application layer multicast (ALM) service. Both the application and the ALM service were easily implemented without having to care about challenges like node mobility or network heterogeneity.

In the demo we consider an exemplary scenario (as shown in Figure 2) that consists of two LANs, one running IPv4 and one running IPv6, respectively. Furthermore, one IPv4

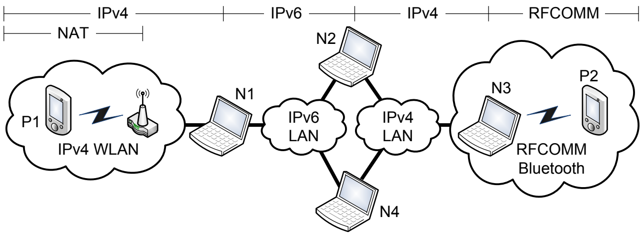

Figure 2: Example scenario consisting of IPv4 and IPv6 LANs, private WLAN, and a Bluetooth point-to-point serial connection.

Figure 3: Real demonstration setup

WLAN access point attached to notebook N1 and a Bluetooth device connected to notebook N3 are deployed. The WLAN uses NAT to multiplex the single IP address of the access point to multiple wireless devices. Furthermore, we employ native serial RFCOMM for communication between N3 and P2 (note that P2 is not using IP). Notebooks N2 and N4 are dual-stacked and connected to both the IPv4 and IPv6 LAN.

All end-systems in this scenario run the group chat application that requires end-to-end connectivity. Two nodes are *directly connected*, if they can communicate through a common subset of protocols and bidirectional packet flow is not inhibited by middleboxes. In the exemplary scenario, shown in Figure 2, nodes N1 and N4 are directly connected, whereas N1 and N3 are not. Using conventional approaches to achieve end-to-end connectivity between all nodes requires usually lots of additional mechanisms, as detailed in Section 1. Furthermore, if the network setting is changed, manual re-configuration is necessary to re-establish connectivity. During this time-consuming re-configuration process—which is usually error-prone and complex—end-to-end connectivity is unavailable.

ariba addresses these problems by using a generic self-organizing approach: It uses an overlay with an *identifier*-based addressing scheme to overcome network heterogeneity in addressing: Nodes using the same application are connected by a logical overlay structure that allows forwarding packets using node identifiers (e. g., using a Chord key-based routing protocol). Furthermore, it does not rely on homogeneous addressing or protocols in the underlay, in fact, *ariba* exploits different protocols to construct an overlay path upon heterogeneous networks where each overlay hop can run different transport- and network-layer protocols. Experiments have shown that this can be done efficiently and with low overhead even if packets must be passed up to the application layer for protocol conversion. Additionally, *ariba* considers that network settings are dynamic and may change over time: For example, notebook N4 may get disconnected from the IPv6 LAN, so N1 has to use N2 as relay to reach N4. In this case *ariba* adapts and re-establishes connectivity automatically.

Initially, viable relay paths are discovered during the join procedure. Relay paths may nevertheless fail if the network setting is changed. In this case the node can re-establish relay paths either by invoking the stabilization procedures of the overlay or by partially repeating the join phase. In the demonstrator, relay paths may not be optimally chosen, but concepts for their optimization exist [3].

The basic demo settings shown in Figure 2 can be modified by connecting and removing relaying nodes, as well as connecting nodes to different networks interactively. *ariba* will automatically sustain connectivity between nodes and handle splitting and merging of overlay partitions. To visualize internal protocol functionality the application additionally shows its local view of the network: relay paths traversing the node, logical overlay neighbors, and the ALM distribution tree.

In future work we will address the following open issues: First, advanced handling of overlay-partitioning and optimization of relay-paths is under current research. Second, we currently examine transparent support for legacy applications that we see as an important step for deployment of *ariba*. Third, porting of *ariba* to further platforms—like OpenWRT, and possibly infrastructure routers—is future work to allow for in-network support of overlay-based services.

Acknowledgements

This work was partially funded as part of the *Spontaneous Virtual Networks (SpoVNet)* project by the Landesstiftung Baden-Württemberg within the BW-FIT program and as part of the Young Investigator Group *Controlling Heterogeneous and Dynamic Mobile Grid and Peer-to-Peer Systems (CoMoGriP)* by the *Concept for the Future* of Karlsruhe Institute of Technology *(KIT)* within the framework of the German Excellence Initiative.

3. REFERENCES

[1] Ariba. http://www.ariba-underlay.org.

[2] R. Bless, C. Hübsch, S. Mies, and O. Waldhorst. The Underlay Abstraction in the Spontaneous Virtual Networks (SpoVNet) Architecture. In *Proc. 4th EuroNGI Conf. on Next Generation Internet Networks (NGI 2008)*, Apr. 2008. CD-ROM.

[3] S. Mies, O. Waldhorst, and H. Wippel. Towards End-to-End Connectivity for Overlays across Heterogeneous Networks. In *Proc. Int. Workshop on the Network of the Future (Future-Net 2009), co-located with IEEE ICC 2009*, Dresden, Germany, June 2009.

[4] O. Waldhorst, C. Blankenhorn, D. Haage, R. Holz, G. Koch, B. Koldehofe, F. Lampi, C. Mayer, and S. Mies. Spontaneous Virtual Networks: On the Road towards the Internet's Next Generation. *it — Information Technology Special Issue on Next Generation Internet*, 50(6):367–375, Dec. 2008. http://www.spovnet.de.