P-DONAS: A Prototype for a P2P-based Domain Name System in Access Networks

Peter Danielis, Jan Skodzik, Carsten Lorenz, Dirk Timmermann
University of Rostock
Institute of Applied Microelectronics and Computer Engineering
18051 Rostock, Germany
Tel./Fax: +49 (381) 498-7272 / -1187251
Email: peter.danielis@uni-rostock.de
Web: http://www.imd.uni-rostock.de/networking

Thomas Bahls, Daniel Duchow
Nokia Siemens Networks GmbH & Co. KG
Broadband Access Division
17489 Greifswald, Germany
Tel./Fax: +49 (3834) 555-642 / -602
Email: {thomas.bahls;daniel.duchow}@nsn.com

***Abstract*—The P2P-based Domain Name System (P-DONAS) is presented. It is implemented on a Xilinx Virtex-4 board and is highly configurable regarding the functional spectrum and many parameters concerning individual functionalities. A GUI has been developed for testing and demonstrating P-DONAS.**

I. Introduction

The Domain Name System (DNS), which represents one of the Internet core services, has not been designed for the dimension and complexity of today's Internet. Thus, scalability is the major issue [1]. Nowadays, Internet Service Providers (ISPs) need to provide and operate regional DNS server farms for resilience and load sharing. Addressing scalability properly means investment in extra equipment, i.e., additional DNS server farms, efforts to operate and manage them, and energy to have server farms run 24/7.

To facilitate cost reductions for ISPs, P-DONAS—a Peer-to-Peer (P2P)-based Domain Name System—has been developed. Access nodes of an ISP's access network are organized into a distributed hash table-based P2P network. These access nodes have a certain available storage and computing capacity, which may be used at no extra costs, assuming some idle time or storage capacity being left in an average access node. Each access node acts as traditional DNS server, solely stores a piece of DNS data, and shares it with all other access nodes. DNS requests issued to an access node are resolved via P2P lookups while maintaining full compatibility with traditional DNS. By deploying P-DONAS as both complement and replacement for traditional DNS, additional DNS server farms can be saved. P-DONAS is implemented as hardware/software prototype on a Xilinx Virtex-4 ML405 development board. Thereby, the development board emulates an access node with P-DONAS functionality. To meet different demands, P-DONAS is highly configurable at compilation time, e.g., regarding the actual functional spectrum and many parameters concerning the individual functionalities. Furthermore, a GUI has been developed to test and demonstrate the P-DONAS prototype.

Section II presents the P-DONAS prototype. In Section III, the demonstration and test tool is introduced. Section IV briefly sketches a demonstration scenario before the paper concludes in Section V.

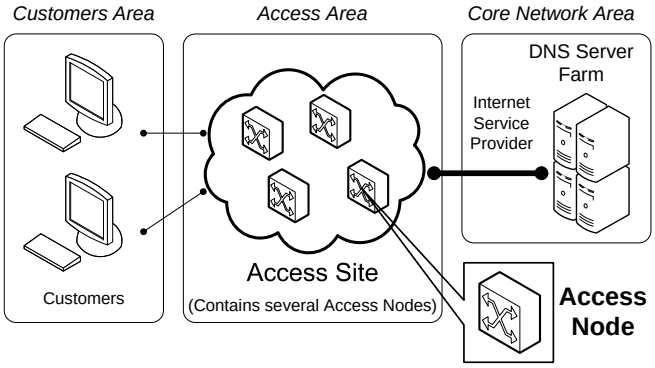

Fig. 1. Access site containing several access nodes connected to a DNS server farm.

II. P-DONAS – The P2P-based Domain Name System

DNS server farms are located at some of the ISPs' access sites in the access area (see Figure 1). Access sites comprise several access nodes.

To supersede additional ISP provided DNS servers farms, P-DONAS has been developed as both complement and replacement for traditional DNS. P-DONAS is realized by organizing access nodes of an ISP's access sites into a Kad-based DHT. Kad is the distributed hash table-based P2P network, which has been chosen for the implementation of P-DONAS [2]. The access nodes save URL-IP address pairs in a structured way (see Figure 2) based on hash values. Thereby, each access node acts as traditional DNS server towards customers but solely stores a piece of all DNS data.

A. System Architecture

The system architecture of a P-DONAS node is depicted in Figure 3. The main components are the cache memory (1), a memory with DNS entries (2), a routing table (3), the Kad block (4), and a block processing functionality for DNS packets (5). The cache serves as a temporary memory for URL-IP address pairs. In the cache, popular URL-IP address pairs are saved, which the P-DONAS node is not necessarily responsible for based on its hash value. The memory with

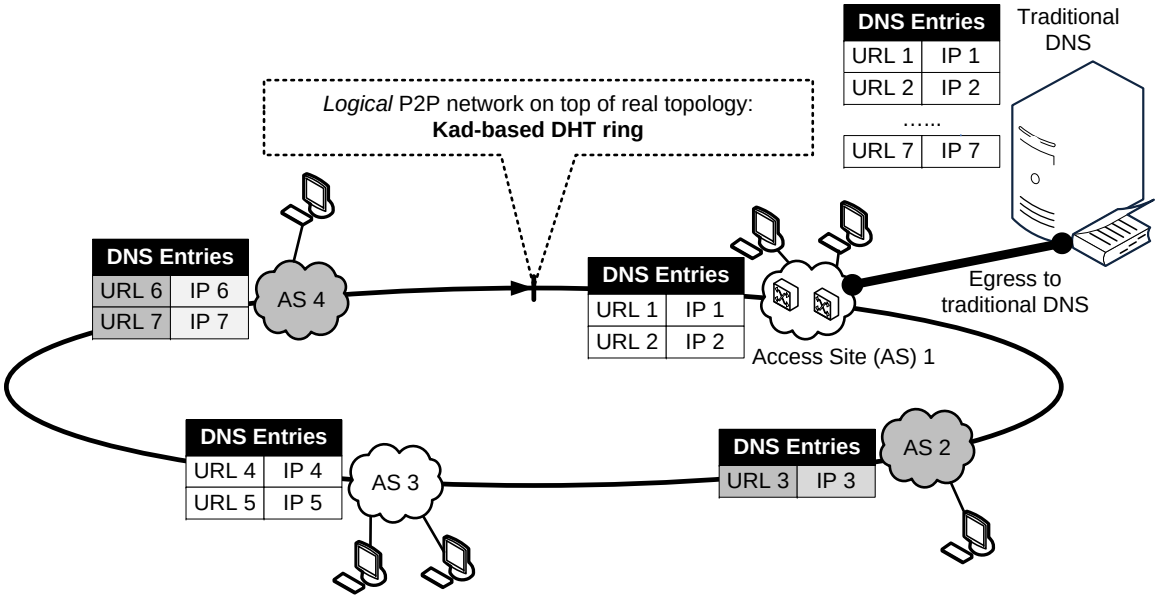

Fig. 2. Kad-based distributed hash table (DHT) ring for structured storage of URL-IP address pairs. The distributed hash table ring has connection to traditional DNS to maintain compatibility.

DNS entries stores URL-IP address pairs the P-DONAS node is responsible for based on its hash value. The routing table contains contact information about other P-DONAS nodes to be able to communicate with them. The Kad block realizes the functionality of the Kad protocol. If a P-DONAS node shall execute an operation on the distributed hash table ring, the Kad block is advised to contact P-DONAS nodes from its routing table.

III. The Demonstration and Test Tool

To demonstrate and test P-DONAS' functionality, a graphical user interface (GUI) has been developed. The screenshot in Figure 4 shows the GUI of the tool's latest version. It is developed in C++ using the open source version of Trolltech's Qt [3]. The first tab "Configuration" of the GUI allows for the configuration of a DNS packet to be sent such as DNS requests. The second tab "Sent Data" is dedicated to show sent DNS packets. The third tab "Received Data" depicts received DNS packets.

IV. The Demonstration Scenario

A P-DONAS node is implemented on a Xilinx Virtex4 FX20 development board (ML-405). The typical, complex environment of P-DONAS on several access nodes cannot be rebuilt. Instead, a scenario has been prepared, which presents the interaction between two P-DONAS nodes.

In the demonstration, a customer issues a DNS request to the first P-DONAS node A. Node A tries to find the requested data in its cache memory but fails as it has not stored any data yet. Via the Kad network, A asks the second node B for the requested data as B has been set to be responsible for this data. Node B fails in finding the requested data as it has not stored it yet either.

Automatically, the DNS request is forwarded to an external DNS name server by A and an answer is received. This answer is stored in A's cache memory and via a store operation in the Kad network, B receives and stores the data as well.

Further DNS requests for the same data can be answered by the first or the second node directly.

V. Conclusion

The working prototype of the P2P-based Domain Name System called P-DONAS has been presented. A graphical tool for the test and demonstration of the P-DONAS prototype on an FPGA development board has been proposed as well. Furthermore, the authors briefly introduced a demonstration scenario.

Acknowledgment

The authors would like to thank the Broadband Access Division of Nokia Siemens Networks GmbH & Co. KG in Greifswald, Germany for their inspiration and continued support in this project. This work is partly granted by Nokia Siemens Networks.

References

[1] V. Ramasubramanian and E. G. Sirer, "The Design and Implementation of a Next Generation Name Service for the Internet." ACM SIGCOMM, 2004, pp. 331–342.
[2] P. Maymounkov and D. Mazieres, "Kademlia: A peer-to-peer information system based on the xor metric." IPTPS, 2002.
[3] Qt by Trolltech. [Online]. Available: http://www.trolltech.com

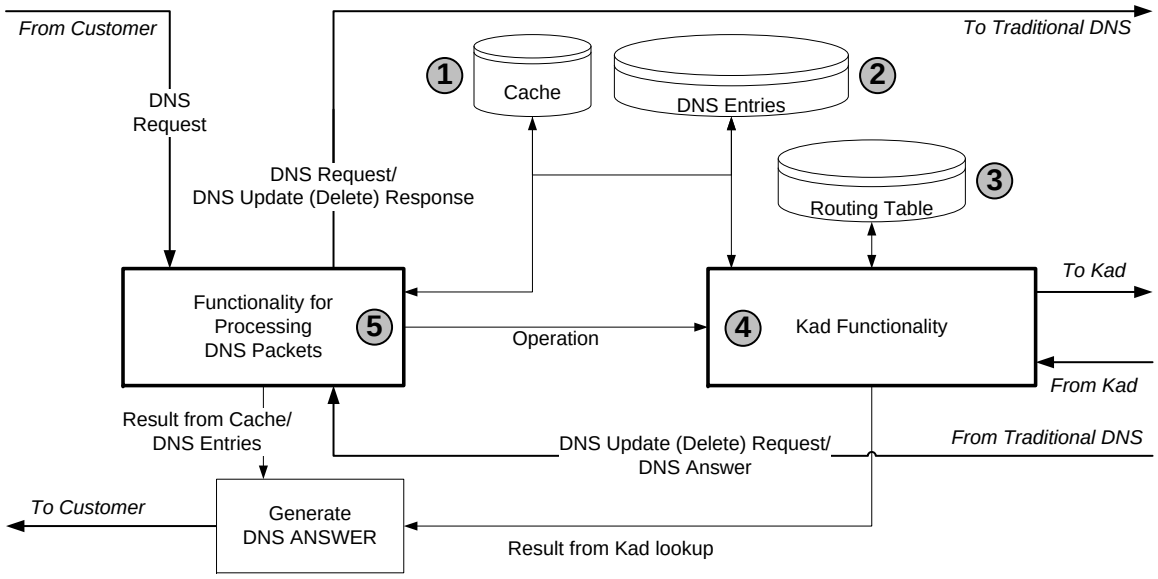

Fig. 3. System architecture of a P-DONAS node.

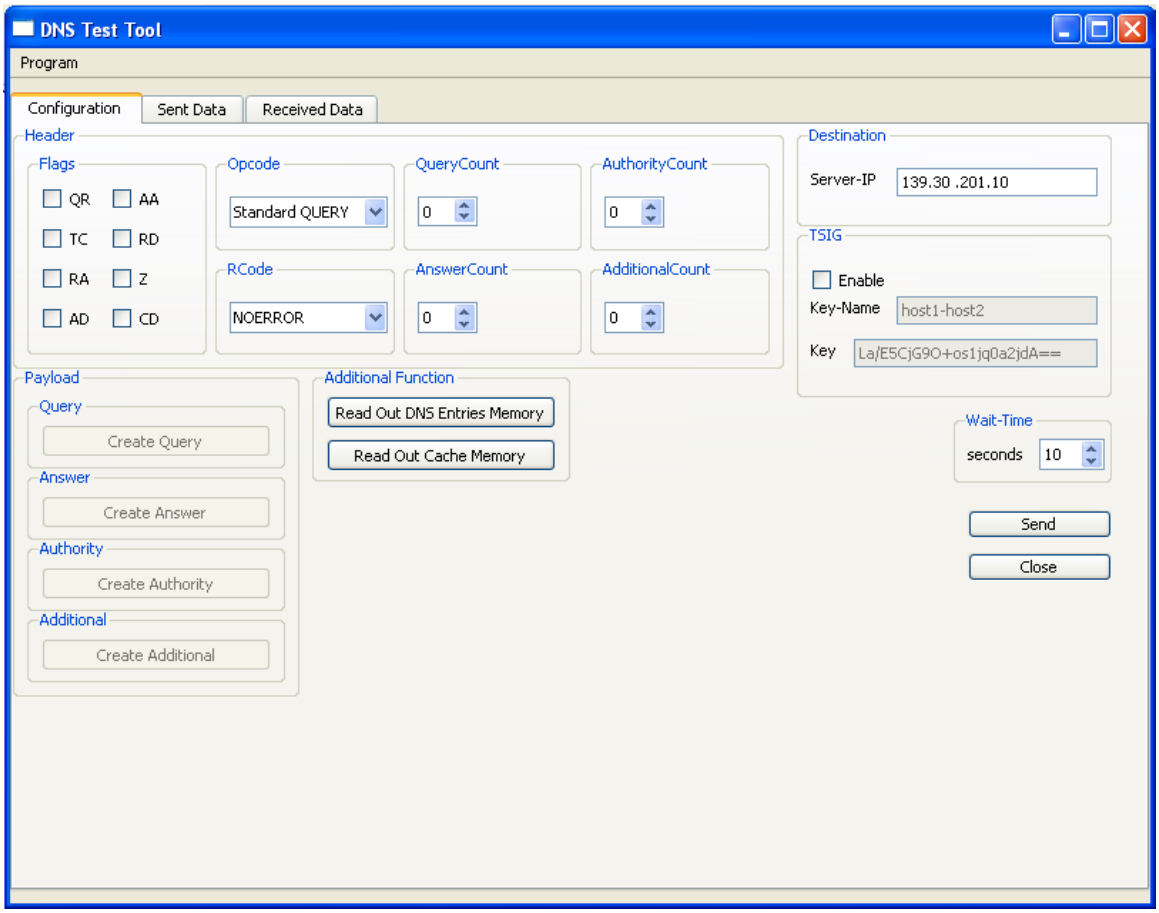

Fig. 4. P-DONAS demonstration and test tool.