Peer-to-Peer Systems
Exercise 3
Published: 20.11.2014
Presentation date: 27.11.2014

HEINRICH HEINE
UNIVERSITÄT DÜSSELDORF

Peer-to-Peer Systems – Exercise
Winter Term 2014/2015

General Remarks

Welcome to the exercise for the lecture Peer-to-Peer Systems.

Please follow the general remarks regarding the organization of the exercise.

- The lecture's website is to be found here: http://tsn.hhu.de/teaching/lectures/2014ws/p2p.html
- For further inquiries, please contact the lecturer under the following email address: graffi@cs.uni-duesseldorf.de

Problem 3.1 - FreePastry

In the following we will have a closer look on one of the most commonly used DHTs in research called FreePastry. In order to answer the following questions, please read the short paper about FreePastry which can be obtained at http://www.freepastry.org/PAST/overview.pdf.

a) Routing table

How does the routing table of a FreePastry node look like? How large is the routing table of a FreePastry node?

b) Routing

How is routing done in FreePastry? Please describe briefly the routing scheme.

c) ID ranges

Have a look at figure 1. For which ID range is the node 10233102 in FreePastry responsible? In the case that two nodes have the same distance to an ID, the node with the smaller ID shall be responsible.

NodeId 10233102

Leaf set

	SMALLER	LARGER	
10233033	10233021	10233120	10233122
10233001	10233000	10233230	10233232

Routing table

-0-2212102	1	-2-2301203	-3-1203203
0	1-1-301233	1-2-230203	1-3-021022
10-0-31203	10-1-32102	2	10-3-23302
102-0-0230	102-1-1302	102-2-2302	3
1023-0-322	1023-1-000	1023-2-121	3
10233-0-01	1	10233-2-32	
0		102331-2-0	
		2	

Neighborhood set

13021022	10200230	11301233	31301233
02212102	22301203	31203203	33213321

Abbildung 1: Example of a routing table for node 10233102.

d) FreePastry and Chord

Now compare FreePastry with Chord. What are the main similarities and differences between these two DHTs?

e) Robustness

Which of the both DHTs (Chord or Freepastry) do you consider to be more robust against churn? Please explain your answer.

Problem 3.2 - Kademlia

Kademlia is one of the most used p2p overlays, let us have a look at it.

a) XOR metric

Prove that the XOR metric of Kademlia fulfills the triangle inequality, i.e. that following holds for all IDs x, y, z:

$$d(x, y) + d(y, z) \geq d(x, z).$$

b) Example Kademlia network

Consider a Kademlia network with IDs ranging from 0 to $2^6 - 1$, i.e. being 6 bits long. Following peers participate in the network:

101100 110110 010101 000001 001110 011001

A) Determine the shortest unique prefixes of the IDs. I.e. note the unique initial bit combinations of each node ID.

B) Draw the ID tree from the view of peer 010101, use the prefixes only and mark the other subtrees seen from the peer.

C) Assume that the network is further populated, and peer 010101 knows all other peers. In which k-buckets are the nodes listed previously?

c) Populating the buckets

Assume a Kademlia network with ID size of 8 bits. The bucket size is $k = 4$. The k-buckets of the peer with ID 11001010 are as follows:
k-Bucket 7: 01001111, 00110011, 01010101, 00000010
k-Bucket 6: 10110011, 10111000, 10001000
k-Bucket 5: 11101010, 11101110, 11100011, 11110000
k-Bucket 4: 11010011, 11010110
k-Bucket 3: 11000111
k-Bucket 2:
k-Bucket 1:
k-Bucket 0:

A) Messages from following nodes arrive in this given order: 01101001, 10111000, 11110001, 10101010, 11100011, 11111111 How do the buckets, the orderings in the buckets and the waiting lists change?

B) Now the node detects that peer 11101110 cannot be reached anymore, what is the reaction?

C) Which addresses would the peer reply to a lookup looking for ID 11010010?