Peer to Peer Systems and Blockchains
Academic Year 2019/2020
Mid Term
Distributed Hash Tables
Deadline 17-04-2020

The mid term consists in carrying out both the following two assignments.

1 DHT Attacks

- **Attacking the Chord DHT**

  Consider a Chord DHT in which an attacking node, Node A, wants to attack another node, Node B. Node A cannot control more than 50% of the nodes of the DHT. The strategy of Node A is to eclipse Node B from the rest of the network. Describe what Node A may do:

  - to eclipse all outgoing messages from Node B to other nodes, by intercepting these messages
  - to eclipse all ingoing messages to Node B.

- **Attacking the Kademlia DHT**

  Node B, the attacked node, decides that Chord is not a secure DHT so it switches a Kademlia network. Node A notices this and decides to try to eclipse B again on the Kademlia DHT. Again, Node A has the ability to create up to 50% nodes of the network. What is, in this case, the best strategy

  - to intercept Bob's outgoing messages to other nodes (eclipse outgoing links)
  - to prevent messages towards Bob from other nodes (eclipse ingoing links).

- why is the Eclipse attack is more easy in Chord with respect to Kademlia?

2 Analysis of Pastry

Pastry [1] is a Distributed Hash Table (DHT) defined by Rowstron (Microsoft Research) and Druschel (Rice University) in 2001. The Pastry DHT shares several characteristics with the Kademlia DHT. Its more salient feature is that it is highly scalable. Read carefully the paper [1] which introduces Pastry, then

- present a brief summary (at most one page) of what you consider the main characteristics of the Pastry DHT.

- Figure 1 presents a snapshot of a Pastry DHT. The dark nodes are currently belonging to the DHT, while the node X, with identifier 23300, is a new node just joining the overlay. The positions of the nodes in the Figure reflect the nodes' distance in the underlying networ: nodes which are close in the underlying network, i.e. their Round Trip Time is low, are close in the snapshot. Suppose that the size of both the leaf set, L, and of the neighbour set, M, is 2.
  - describe the join process carefully. It is required to show how the joining node X builds its routing tables, its Leaf set and its Neighborood set
  - after having built its routing table, node 23300 sends a message to node 03323. Describe the routing of the message, using the routing tables defined in the previous step.

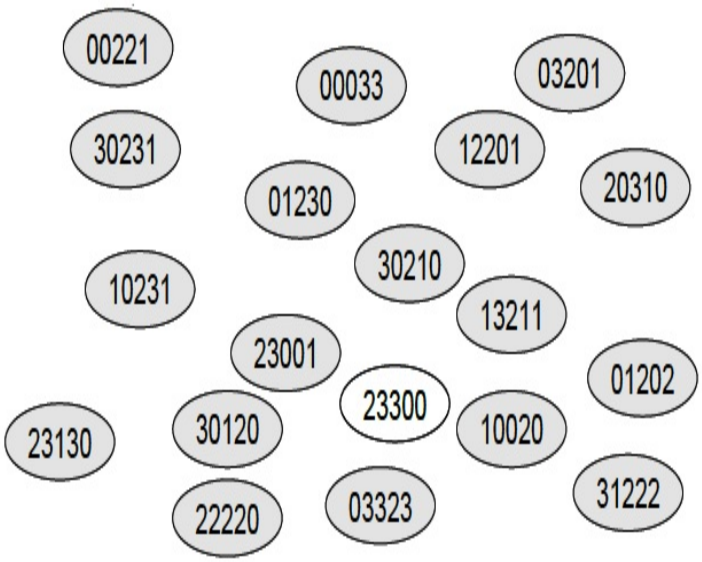

Figure 1: A Snapshot of a Pastry overlay

Justify in the final report each assumption you made, which is not explicitely stated in the text of this assignment. [1]

3 Requirements

The mid term must be done individually and the deadline is 17 April 2020. If the evaluation of both the mid and of the final term will be positive, the student will be relieved from the oral exam. The mid term is not mandatory, if it is not presented, the student will be required to pass the oral exam on the first part of the course and the student will be allowed to submit the final term. The student must submit through Moodle a report including the solutions of the two assignments. The rating of the assignment will

[1] Discussion on the Moodle forum of the course is greatly encouraged

be published on the Moodle as well and discussed during the question time (likely through Skype/Meet call).

References

[1] *Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems*, A. Rowstron, P. Druschel, IFIP/ACM International Conference on Distributed Systems Platforms and Open Distributed Processing, Springer 2001.