



# Distributed Systems Part II

## Exercise Sheet 9

### 1 DHT

As we have discussed in the lecture, most distributed hash tables (DHTs) are built on the same idea: a binary search tree. Each leaf of the tree is represented by a peer, nodes (including the root) do not really exist. A peer knows only a small subset of all the other existing peers. But if a peer knows the address of another peer it can always contact the other peer. The network is unreliable, messages can be lost, altered or arrive out of order.

Introduce new messages, protocols, restrictions or other ideas to protect a DHT from various Byzantine attacks. If you need to make assumptions, write them down. Each answer should be about 5-10 sentences.

- a) *Wrong lookup*: A search for a key roughly requires  $O(\log n)$  steps. In each step one new peer is included in the search. One way to search is to send a message through the DHT. The message contains the searched key and the address of the peer that started the search. The message is forwarded from one peer to the next such that it always gets a bit nearer to its (yet unknown) destination. Once the destination is reached, the final peer answers.  
A Byzantine peer sends the message either in the wrong direction, or to a non-existing peer.
- b) *Incorrect routing updates*: Each peer maintains a routing table containing the addresses of about  $\log n$  other peers. The peers send update messages to each other in order to keep their routing tables up to date.  
A Byzantine peer sends false updates, e.g. it tries to place dead links in a routing table.
- c) *Partitioning*: If a peer wants to join a DHT it has to make contact with a peer that is already part of the DHT. The new peer can then ask the old peer about other nodes of the DHT and insert itself at an appropriate place.  
A Byzantine peer builds up his own private DHT by sending wrong messages to joining peers. It gives joining peers only the addresses of peers in his own isolated net.

### 2 P2P File Sharing

- a) Tit-for-tat (T4T) denotes the strategy in a repeated two-player game to reply to the opponent's action with the same action. E.g., if the opponent previously was cooperative, the agent is cooperative. If not, the agent is not. Explain the problems that arise when we want to apply the T4T strategy to a P2P file sharing system.
- b) Successful file sharing protocols such as BitTorrent exploit the fact that peers downloading the same file (or collection of files) can be organized so that they can trade different parts of the same file among each other. The set of peers that are in the process of downloading

the same file  $f$  is called the *swarm* of  $f$ . If a peer wants to download a file  $f$  it enters the swarm of  $f$ , gets some trading capital, and exchanges parts of  $f$  with the other peers in the swarm.

In the following, let  $T_{ij}$  denote the number of file blocks transferred from peer  $i$  to peer  $j$  in the given P2P file sharing system. Furthermore, you may assume for simplicity that the file blocks are at any time distributed well among the peers. In particular, assume that every peer who is contacted by a freeloader always has a file block in which the freeloader is interested. Moreover, assume that the freeloader can trade any previously gathered file block  $b$  for a new one with a peer  $i$  unless it has received  $b$  from  $i$  earlier.

- i. The *pairwise  $\Delta$ -strict T4T* strategy is the strategy of a peer  $i$  in a file sharing system to only upload a file block to a peer  $j$  if  $T_{ij} - T_{ji} < \Delta$ . Consider a swarm of a file with  $m$  blocks where all peers trade according to a pairwise  $\Delta$ -strict T4T strategy. How large does a swarm need to be at least such that a freeloader can download the entire file for free?
- ii. Assume that a trusted authority keeps a global balance of each swarm participant's contribution, i.e., it announces any peer  $i$ 's contribution  $\beta_i = \sum_j (T_{ij} - T_{ji})$  upon request. The *global  $\Delta$ -strict T4T* strategy is the strategy of a peer  $i$  in a file sharing system to only upload a file block to a peer  $j$  if  $\beta_j < \Delta$ . If all peers employ a global  $\Delta$ -strict T4T strategy, how large can the file size  $m$  be at maximum, respectively how large must the swarm size  $n$  be at minimum, such that a freeloader can get the entire file for free?
- iii. Explain why a trusted authority is hard to implement in a real P2P system, and give an alternative to keeping a global balance in P2P file-sharing systems!