

196:001 Peer-to-peer Networks: assignment 3

Sample solution

Question 1. (20 points)

Answer:

BitTorrent works without a designated tracker by adding a DHT (Distributed Hash Table) layer on top of the BitTorrent network, which allows nodes to store and retrieve contact information for peers in a torrent. In effect, each peer becomes a tracker. The implementation detail is discussed in [1]. The following is an overview.

BitTorrent uses a "distributed sloppy hash table" (DHT) for storing peer contact information for "trackerless" torrents. The protocol is based on Kademila [2] and is implemented over UDP.

DHT and Routing Table

Each node has a globally unique identifier known as the "node ID." Node IDs are chosen at random from 160-bit. Nodes must maintain a routing table containing contact information of many short-distance neighbors and a handful of long distance neighbors. XOR is used as the distance metric and the result is interpreted as an unsigned integer. The routing table is subdivided into "buckets" that each covers a portion of the name space.

Query


When a node wants to find peers for a torrent, it searches its routing table and then contacts the neighbors with IDs closest to the infohash and asks for the contact information of peers currently downloading the torrent. If a contacted node knows about peers for the torrent, the peer's contact information is returned with the response. Otherwise, the contacted node must respond with the contact information of the nodes in its routing table that are closest to the infohash of the torrent. The original node iteratively queries nodes that are closer to the target infohash until it cannot find any closer nodes. After the search is exhausted, the node then inserts the peer contact information into its routing table.

The return value for a query for peers includes an opaque value known as the "token." For a node to announce that its controlling peer is downloading a torrent, it must present the token received from the same queried node in a recent query for peers. When a node attempts to "announce" a torrent, the queried node checks the token against the querying node's IP address. This is to prevent malicious hosts from signing up other hosts for torrents. Tokens must be accepted for a reasonable amount of time after they have been distributed.

Question 2. (30 points)

Answer:


1. The social optimum for a system of 6 nodes is a star.



The total cost is: $3 * 5 + 2 * 5 + 2 * (6 * (6 - 1) - 2 * 5) = 65$.

2. The sequence of possible moves is:

- node E adds an edge to node B;
- node C removed the edge to node D.



The total cost in this configuration is **75**. The Price of Anarchy for this system is $75/65=1.15$.

References:

[1] Andrew Loewenstern <drue at bittorrent.com>

http://http://www.bittorrent.org/beps/bep_0005.html

[2] Peter Maymounkov, David Mazieres, "Kademlia: A Peer-to-peer Information System Based on the XOR Metric", *IPTPS 2002*.

<http://www.cs.rice.edu/Conferences/IPTPS02/109.pdf>