Computer Networks
ICS 651

- Distributed Hash Tables, DHTs
- bitcoin
- allnet

Distributed Hash Table Goals

- DHT
- P2P self-organizing storage network – and sometimes also a communications network
- access to one node of the network should give access to the entire content of the network
- data is identified by keys which are fixed-length bitstrings

DHT ID space and routing

- each node selects a random k-bit ID
- each node keeps a "routing table" with up to k entries
- the i-th entry in the routing table, if any, has a route to a node that has the same first i-1 bits of the ID, and a different i-th bit
- in this way, the first entry in the routing table is in the other half of the ID space
- the second entry in the routing table is in the same half of the ID space, but the other quarter, and so on
- to route a packet from node x to ID y, where the first i bits of x and y are the same, the packet is sent to the node (if any) in the ith position of the routing table -- this node will be closer to the destination
- if the network has N nodes, the routing table only needs log N entries (or as many entries as the key has bits)

DHT storage

- each node stores data for its own ID, and all IDs up to n-1 nodes later in ID space (maybe n=4)
- a request is routed until there is no route to a closer node. At this point, the requested data should be on the node
- when nodes leave the network (**churn**), the data should still be available redundantly in other nodes -- a data redistribution scheme then tries to have every item in at least n nodes

DHT performance

- at least one node must be reachable in order to connect with the network
- messages for new data must go through an average of log N hops, where each hop may cross the entire Internet
- when there is a lot of churn, routing tables must be recomputed frequently and data stored in the network must be redistributed frequently

Bitcoin

- distributed, reliable ("hard to falsify") time-stamping network
- each time-stamp record ("block") includes existing transactions, and a summary (hash from Merkle tree) of all previous transactions
- finding valid blocks require compute power,
  - first n bits of hash must all be zero
  - can only be found at random (we think!)
  - n evolves to keep the number of blocks generated at a near constant rate
- "block miners" get a reward for finding a block with a suitable hash
- each block keeps track of all transactions since the last block by including a hash of the most recent block
- P2P network distributes the blocks to all peers

Bitcoin transactions

- each transaction has a list of input transactions, adding up to an amount B
- each transaction has a list of output amounts and public keys, adding up to an amount B' <= B
- if B' < B, the difference B - B' belongs to the block miner as a transaction fee
- the input transactions must provide evidence of having a private key equivalent to the public key of the corresponding output transaction

Bitcoin block chain

- each bitcoin miner produces and hashes as many blocks as possible until it finds a hash with the first n bits zero
- the miner then broadcasts the block as fast as possible
- every other miner includes this block into a new block they try to generate
- two (or more) miners may find blocks A and B approximately simultaneously, and broadcast both throughout the network
- every miner has a choice of which block to include in the block they are working on
- the next successful block C will have picked one of the winners of the previous round, A or B
- almost all miners will now build their block on C, confirming either A or B, and one of them will produce a new block D
- eventually, only one of A or B is confirmed by the growing block chain

Other applications of blockchains

- distributed, reliable network for time-stamping transactions
- might be useful anywhere distributed recordkeeping is needed
  - e.g. real-estate transactions
- but proof-of-work is expensive!
  - proof-of-stake may be a less energy-intensive alternative
  - bankers may like the idea of a distributed network limited to only carefully-selected participants
- in any case, there must be incentive for the miners
  - e.g. in bitcoin, miners get newly-minted coins and transaction fees

AllNet

- See this 4-minute talk at the March 30th, 2016, Wetware Wednesday

  http://www2.hawaii.edu/~esb/2016spring.ics651/allnet-talk.pdf

- See also

  http://alnt.org/

AllNet status, 2018

- ad-hoc communication between Linux systems of ad-hoc 802.11 WiFi
- DHT used for communication over the Internet
- bandwidth-limited to an average 8KB/second
  - priorities let us forward own packets even when bandwidth is limited
- trace mechanism for debugging, similar to ping+traceroute

```
b1.01/16     0 hop
b3.00/16     1 hop
b4.00/16     2 hop
b5.00/16     3 hop      2.519748s rtt
```