

CS-438

Decentralized Systems Engineering

Week 4

Search and Routing Protocols

- Unstructured: don't assume anything about network structure
- Structured: can build/maintain some kind of "map" information (BGP, intra-domain, ...)

Structured search: search for (node, user, file, volume, ...)

- Base case algorithm: Flooding search via broadcast/gossip/...

- 1. Announcement: name/ID of search target (e.g. "Bob", wildcard like "*eilish*", ...)
- 2. Answer: gossip back? Follow a reverse path built during announcement (source routing)

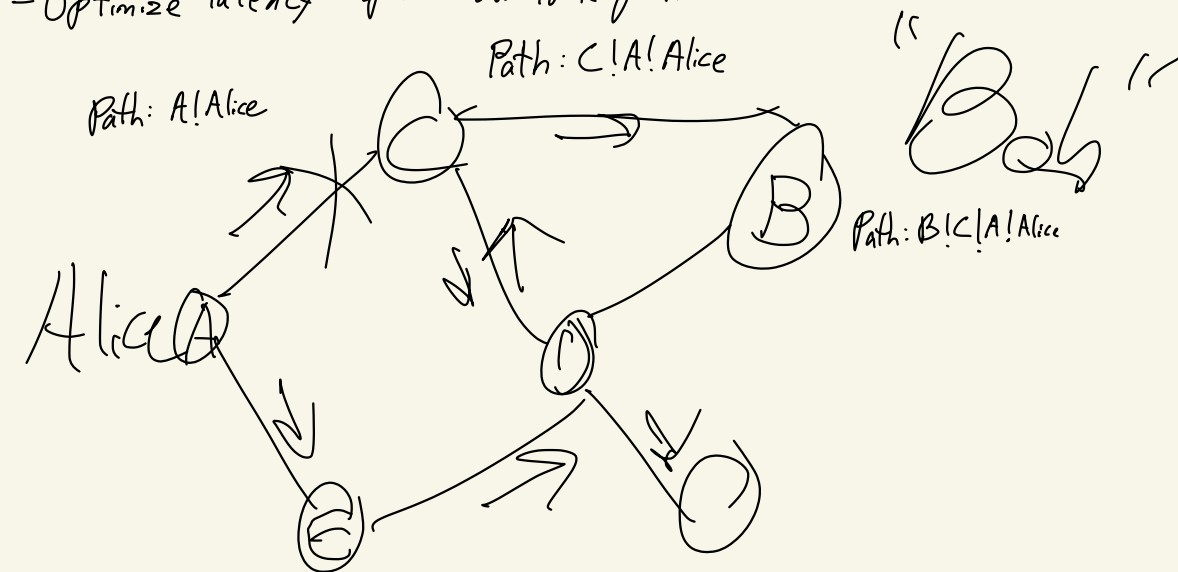
Upside: simple, robust if answers are fresh / Downside: BW cost of global searches $O(n^2)$

- Optimize via ^{space} locality: 1. set hop limit/TTL/message budget
2. increase limit (TTL+=1, msg budget *=2)
- Optimize via time locality: e.g., cache paths for some time
- Optimize latency: spend BW to keep fresh cache

DSDV
Destination
Sequenced
Distance
Vector

UseNet headers

Path:
Subject: ...
Date: ...
Message-ID:
From: alice@A.edu



Asymptotically better than $O(n^2)$ BW w/ unstructured search flexibility?

Yes: $\tilde{O}(\sqrt{n})$ per search

- BubbleStorm: 1. each node announces its metadata - send it to \sqrt{n} random nodes \leftarrow each takes $O(\sqrt{n})$ BW
2. search: query a random $c \cdot \sqrt{n}$ nodes

- How to sample $O(\sqrt{n})$ nodes at random efficiently

- Mark chains, random walks on expander, fast mixing properties

- Start at any point, take $O(\log n)$ steps "at random" - produce almost uniform random sample

