

CS-438

Decentralized Systems  
Engineering

Fall 2022

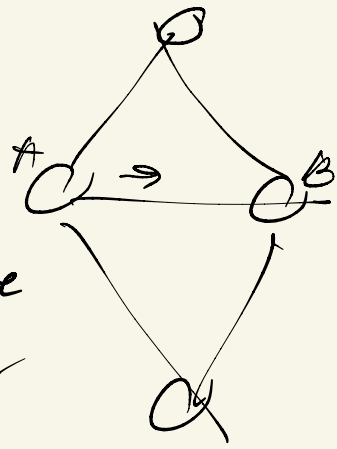
Week 4

# Gossip efficiency

- I HAVE / SEND ME (UseNet/MNTP)

- message content sent only 1x per node

- still requires interaction, sending message - IDs redundantly




- Can we minimize redundant bandwidth usage without interaction?

- Algorithms: Rumor-mongering  
(Randomized) Anti-entropy

# Rumor - mongering $n$ nodes

- if maximum degree of any node =  $d$
- naive gossip:  $nd$  messages  $M$
- goal: reduce to  $O(n)$



- rumor mongering:

1. pick random neighbor, send  $M$   
neighbor replies 1 bit: new?  
if  $M$  was new to neighbor: repeat  
else: flip a coin: repeat w/ prob 50%.  
else stop

- efficient at quickly delivering  $M$  to most nodes

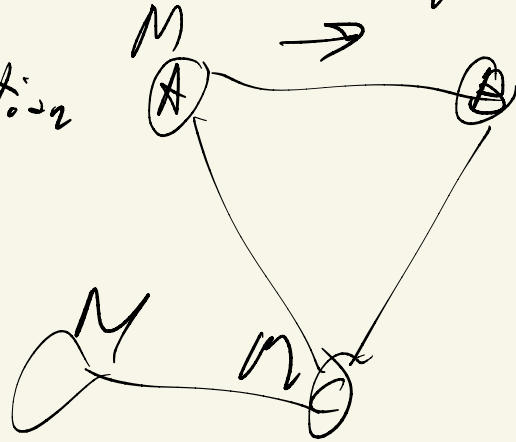
Ant: - entropy

- periodically (when timer fires):  $O(n)$

1. pick a neighbor at random

2. send "any thing new?" exchange news  
reduce entropy (diff. in message sets)

- ensures complete dissemination



# Search algorithms (distributed)

---

- unstructured - instantly adaptive to churn
- structured - can be much more efficient but brittle, require maintenance

# Unstructured search

- Napster, Gnutella


- flexible searches:  
(X and Y) or (Z)

- standard basic algorithm:

Flooding

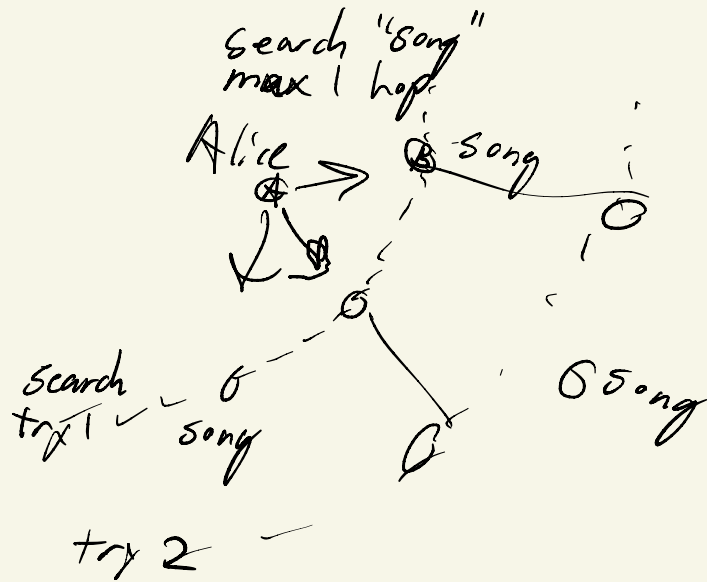
issues:

- unpredictable delays
- no all-to-all connectivity for replies
- efficiency: all nodes see,  
process all searches



# Optimizations

- Expanding-ring search uses limited flooding
- pragmatically - works
- increases latency
- not asymptotically better
- worse  $O(n \log n)$



# Bubble Storm

"mostly-unstructured"

- general search terms

- asymptotically more efficient:  $O(\sqrt{n})$

- leverages birthday paradox - "random meet-in-the-middle"

