Peer-to-peer Networking Research

Bobby Bhattacharjee
Joseph JaJa
Michael A. Marsh

P2P Research Focus at UMIACS

- unstructured, decentralized distributed systems
- distributed digital archiving
- emergent behaviors

Unstructured Decentralized Systems

S. Bhattacharjee, M.A. Marsh
with Jonathan Katz, Aravind Srinivasan

decentralized distributed systems useful
- pub/sub, PKI, privacy-presrv. distribution

important components: naming, lookup
- efficient structured lookup exists
- we introduce efficient unstructured lookup

Decentralized PKI

centralized PKI

- does not scale
- requires trusted infrastructure

decentralized PKI

- (potentially) scales well
- trust needed along query/response path

trust relations not likely to adhere to particular structure $\Rightarrow$ efficient unstructured lookup

Trust Networks

S. Bhattacharjee, M.A. Marsh
with J. Katz, A. Srinivasan

trusted peers
- more likely to store data
- retrieved data more likely correct

“small-world” phenomenon in social networks
- high-trust path likely for any two individuals

good behavior rewarded, bad behavior punished
- game theory

Trust Inference

peers that know each other have direct trust values

most peers will not have direct trust values
- have to infer indirect trust values

is the peer reliable?
which peer is more reliable?
- multiple public keys presented for one name

Robust Efficient Lookup

S. Bhattacharjee, M.A. Marsh
with A. Srinivasan

identifier lookup

unstructured networks $\Rightarrow$ flooding $O(n)$

distributed hash tables $\Rightarrow$ $O(\log n)$

Local Minima Search on unstructured networks

better than flooding

almost as good as DHT (mostly) $\Rightarrow$ $O(n^{1/4})$

Finding Local Minima

virtualized namespace

peers “hashed” to uniformly-distributed IDs

objects hashed to same ID space

local minimum

peer with ID closest to object ID in its neighborhood

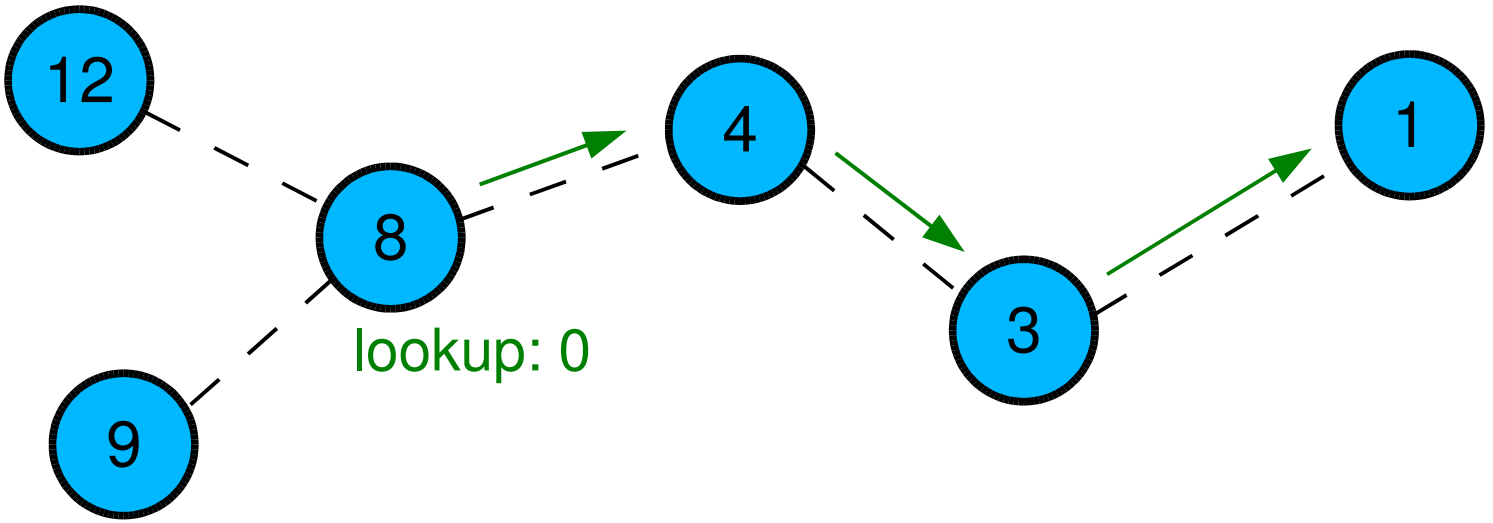

Random Local Minima

unique local min from any peer

basin of attraction for a local min

random walks ⇒ random basins

mult lookups ⇒ different local min

k distinct local minima

replicas at r

s search probes, Pr[fail] ≤ $e^{-rs/k}$

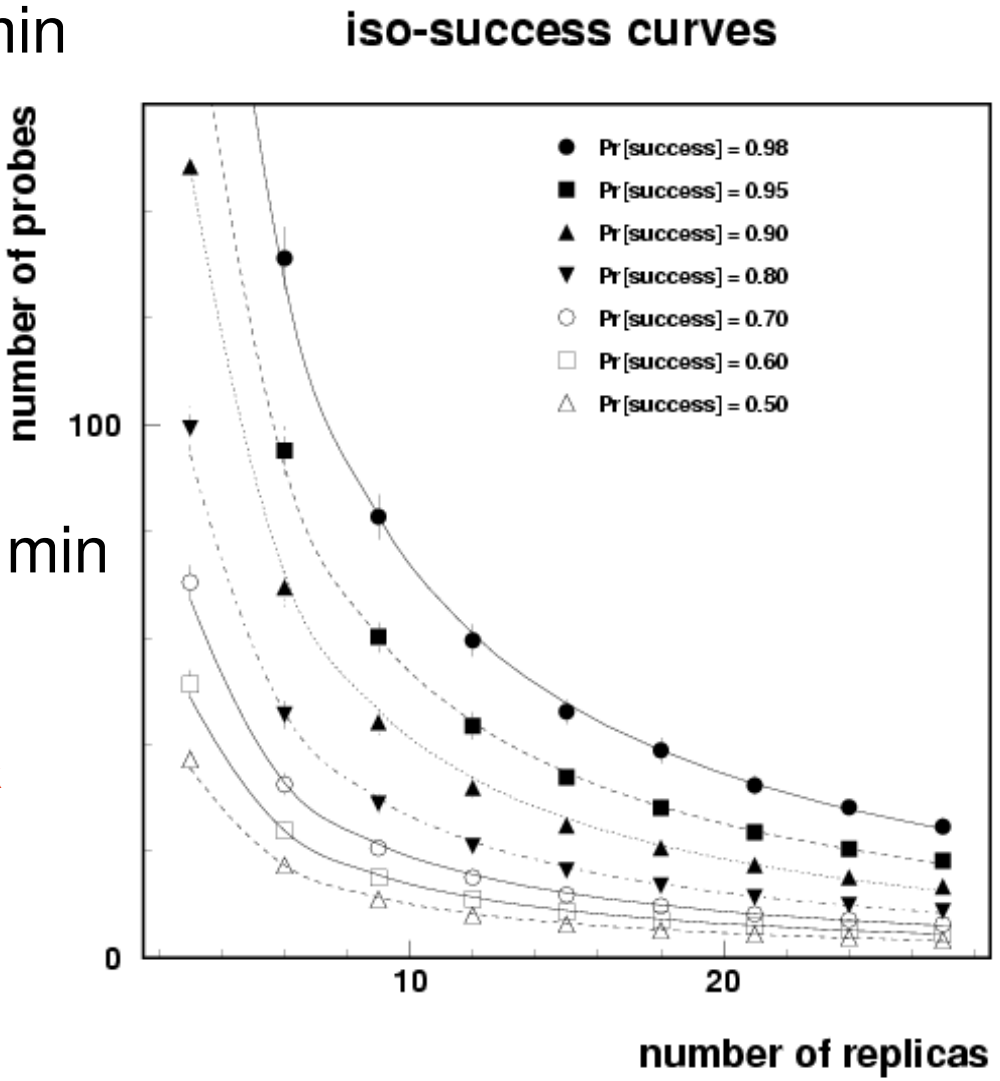

Adaptive Replication

peers – only local information

⇒ how many replicas to place?

use successful lookups

send s back to owner

compute average <s>

desired r = F(<s>)

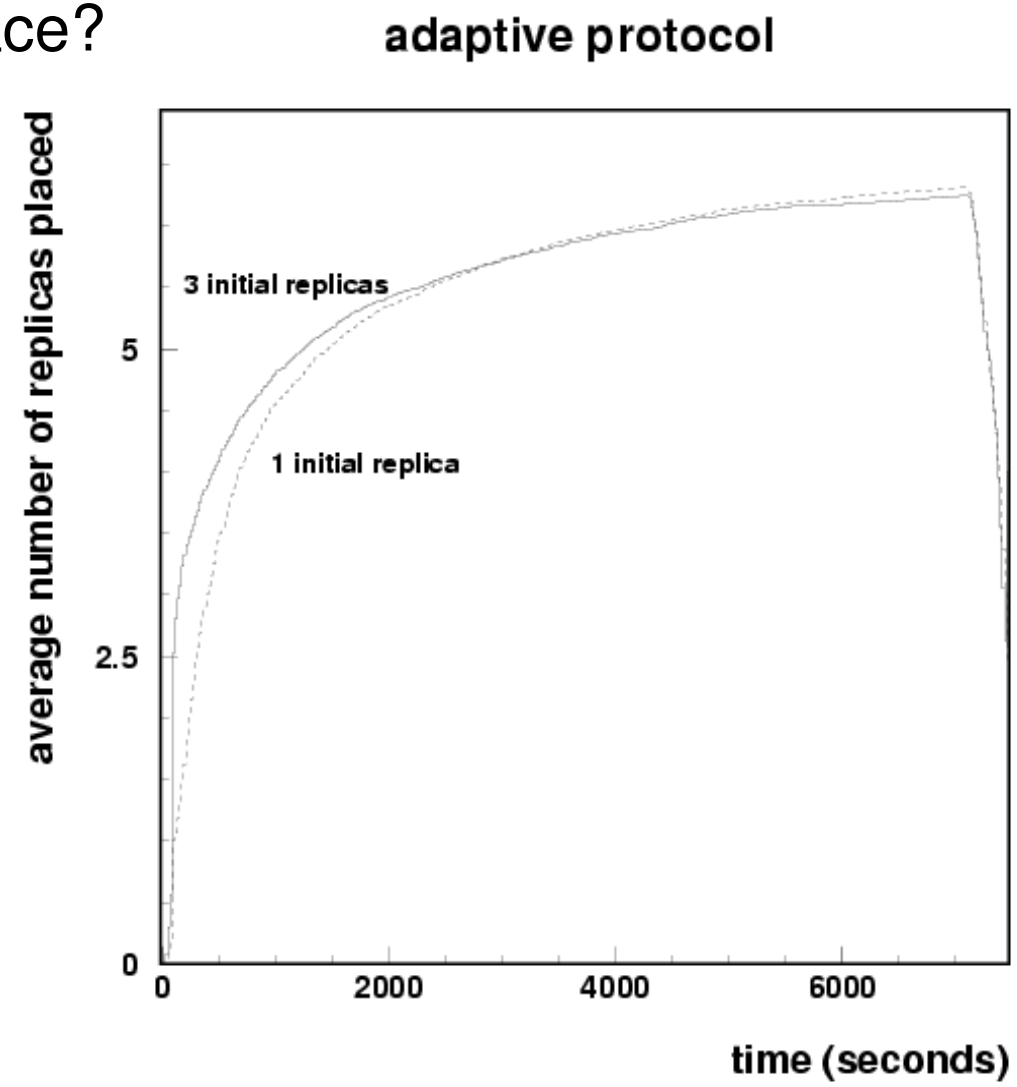

Distributed Digital Archiving

J. JaJa, Qingmin Shi

increasing digitization of documents

- good for short-term availability
- long-term risk losing cultural/historical items

leverage work on decentralized systems

- scale of problem

  $\Rightarrow$ decentralized, distributed solution

Distributed Digital Archiving

trust might not be significant issue

– even libraries do not always trust each other

equipment failures much more likely

⇒ replication critical

peers have limited resources

⇒ load balancing shifts data to less-used peers

Distributed Digital Archiving

O(log n) for all ops, including key range search

load balancing incurs amortized constant overhead

d-fold replication

- divide peers into d virtual peers
- each virtual peer has version index
- skip graph of virtual peers w/ particular version

pilot archive: UMD, SDSC, NARA

Decentralized Systems

digital archiving, unstructured networks work done independently

- seem to be converging
  - replication
  - efficient lookup
- combining efforts likely to benefit both

Emergent Network Behaviors

M.A. Marsh
with John Blake, Walt Kaechele, Dany Smith

large p2p networks common

peers run multiple protocols

TCP/IP, BitTorrent, Gnutella, . . .

$\Rightarrow$ complex interactions

self-similarity (eg, Internet traffic)

all point to complex systems

– often chaotic

– can lead to emergent behaviors

Emergent Network Behaviors

emergent behaviors not well studied

can be critical for security/stability of network

- spontaneous network failures
- induced long-range effects:
  - “ringing”
  - “seizures”

Emergent Network Behaviors

experiments on large-scale systems unfeasible

rely on simulations
- networks up to size of Internet
- timescales from μs to days

target for preliminary results ~6 months

simulation framework likely useful elsewhere