

Perspective on Overlay Networks

Panel: Challenges of Computing
on a Massive Scale

Ben Y. Zhao
FuDiCo 2002

Why Overlays

- ✦ Changes to standards is hard
- ✦ Possible benefits of overlays
 - Extensible communication models
(more than best-effort point to point)
 - More easily deployed services / infrastructures
 - More adaptable, fault-tolerant ...

Massive Scale

✦ It's coming...

- Driven by global scale resource sharing apps
- File-sharing, data dissemination, shared computation
- Viral deployment (Kazaa / alt.net)
- More thin clients going online: cellphones, PDAs, fridges, toasters

June 4, 2002

ravenben@eecs.berkeley.edu

Implications of Scale

✦ Pushing bounds of infrastructure scalability

- Logarithmic bounds are not enough!
- Worse case: massive congestion, network partitions, data / service unavailability

✦ Scrutinize infrastructure scalability

- Re-examine system level primitives
 - Control traffic, fault-detection, replication overhead
- Dampen impact of operations on infrastructure
 - Localize bandwidth / storage requirements, faults

June 4, 2002

ravenben@eecs.berkeley.edu

Highly Dynamic

■ It's here...

- Small half-life for P2P apps (leeching)
- Connectivity loss from congestion, flash crowds, misconfiguration, hardware faults, lower MTBF
 - BGP convergence 3-15 mins
- Mobile clients in semi-connected mode
 - E.g. chatting / file-sharing on your IPAQ / Palm

June 4, 2002

ravenben@eecs.berkeley.edu

Implications of Change

- Have to reexamine stability assumptions
 - Overlay maintenance under instability
 - Without care, worst case: sub-optimal paths, network partitions, broken invariants → loss of availability
- Proactive overlay construction and maintenance
 - Integrate optimal routes from square 1
 - Estimate global knowledge w/ distributed algorithms at node insertion
 - Run-time heuristics are too late
 - Constant fault-detection and self-repair
 - Use of redundancy to mask self-repair

June 4, 2002

ravenben@eecs.berkeley.edu

Extending Scope of Overlays

- Broader utility than file storage and retrieval
 - General routing layer (efficiency, fault-tolerance)
 - Detour, RAMP, RON, Tapestry
 - Flexible data dissemination layer
 - overlay multicast, anycast, mobility
 - Collaborative applications
 - Shared annotations, mail filtering
 - Data transcoding
 - Shared workspaces, games
 - Distributed computation
- Barter incentives
 - (Mojos, ads, access to resources)

June 4, 2002

ravenben@eecs.berkeley.edu

Introspective Tuning for Reality

- Exploit knowledge of the environment
- Network latency discontinuity
 - Order of magnitude latency jump from LAN to WAN
 - Generic solution:
 - Brocade: minimize wide-area hops via 2nd overlay
- Non-uniformity in node resources
 - PDA < laptops < desktops < servers
 - Segregate supernodes, infrastructure routers, pure clients, storage devices

June 4, 2002

ravenben@eecs.berkeley.edu

Many Security Challenges

■ Open P2P = Infrastructure for attacks?

- Misreporting data → unfair resource allocation
- DDoS attacks on infrastructure
- Sybil attack → enabling scale for attacker
 - Detecting maliciousness made harder

■ Scalable defense?

- Cryptographic methods for authentication, privacy, etc.
- Byzantine agreement
- Redundancy (ID aliasing, multiple namespaces)
 - Misuse leads to resource based DDoS

■ Access control?

June 4, 2002

ravenben@eecs.berkeley.edu