Overcast: Application-Level, Reliable Multicast in an Overlay Network

John Jannotti, David K. Gifford, Kirk L. Johnson,

M. Frans Kaashoek, and James W. O'Toole, Jr.

Cisco Systems

Overcast solves three problems

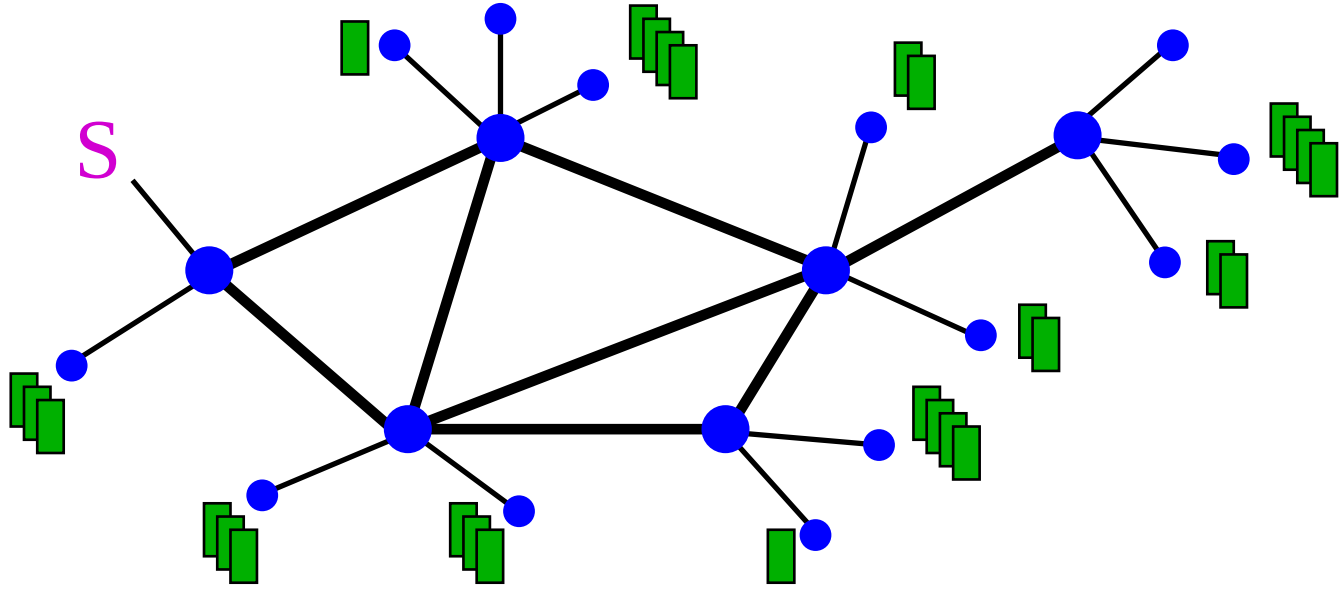

- Web servers must send data independently to all clients.
- High-bandwidth content requires users to wait.
- Live content must be carefully bandwidth controlled.

IP Multicast - Live content

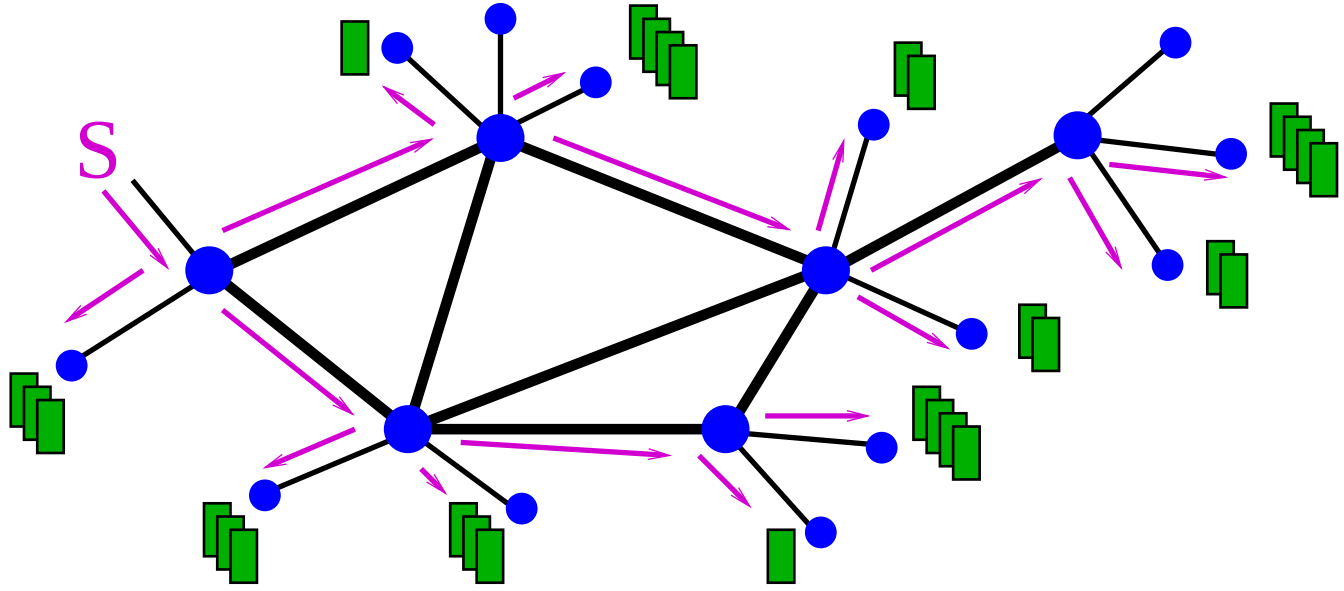

- Provides unreliable, live delivery only
- Requires router, server, and client modifications.
- Many difficult details - security, naming, billing.

Caching & Replication is insufficient

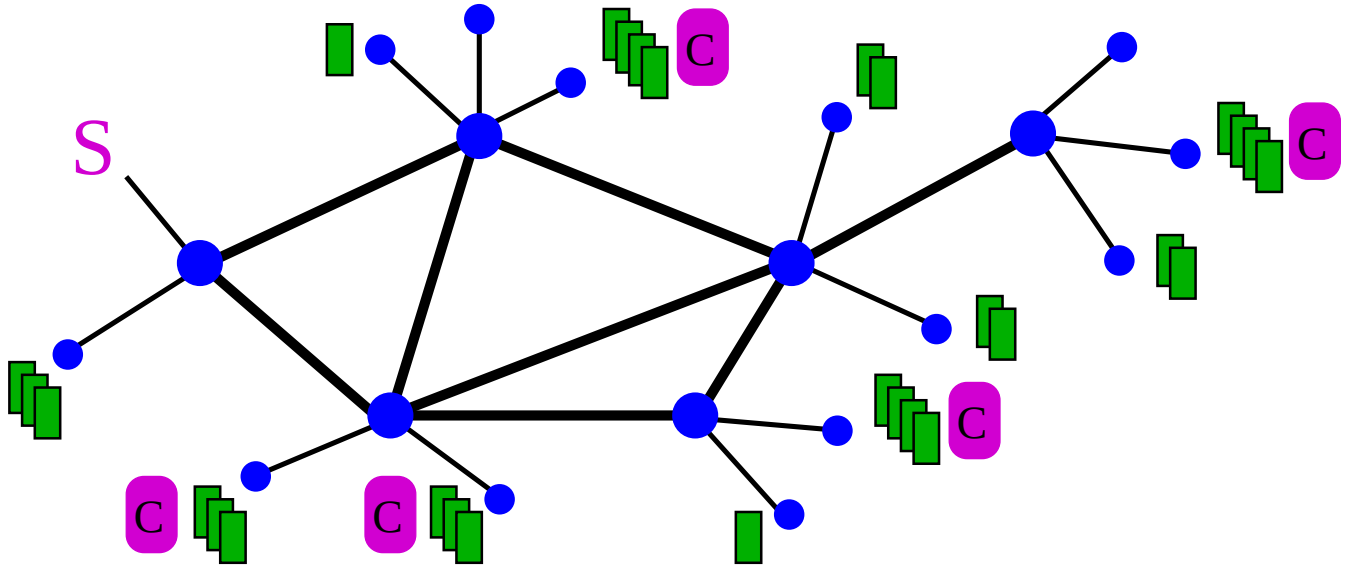

- Caching - Providers lose control. Cold caches discourage use.
- Replication - Scaling and/or maintenance difficulties.
- No live content.

Overcast unifies today's solutions

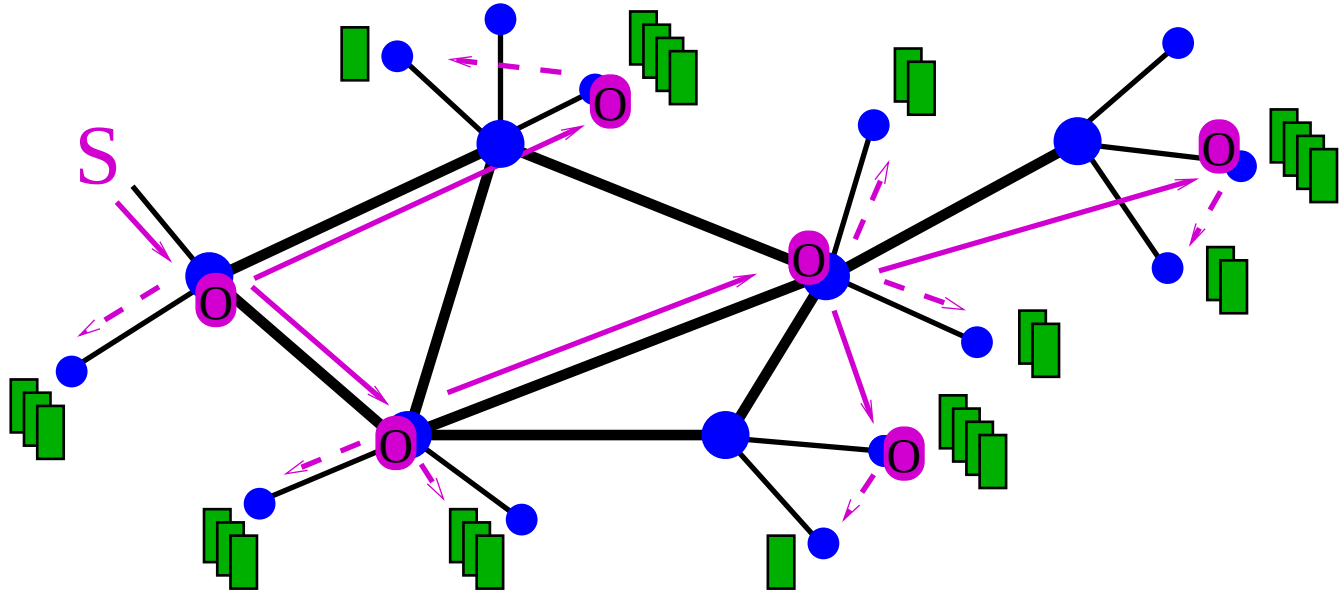

Overcast is application-level multicast in an overlay network.

- Overcast nodes are sprinkled around the Internet
- Content is multicast through a self-organized topology.
- Content is stored for optional delayed playback.
- HTTP clients may join Overcast groups.

Benefits and Challenges

Overlay networks are. . .

- Incrementally Deployable
- Application Specific

but there are challenges. . .

- Information Loss
- Inefficiency

Organizing an overlay network

Constructing an efficient topology is difficult because...

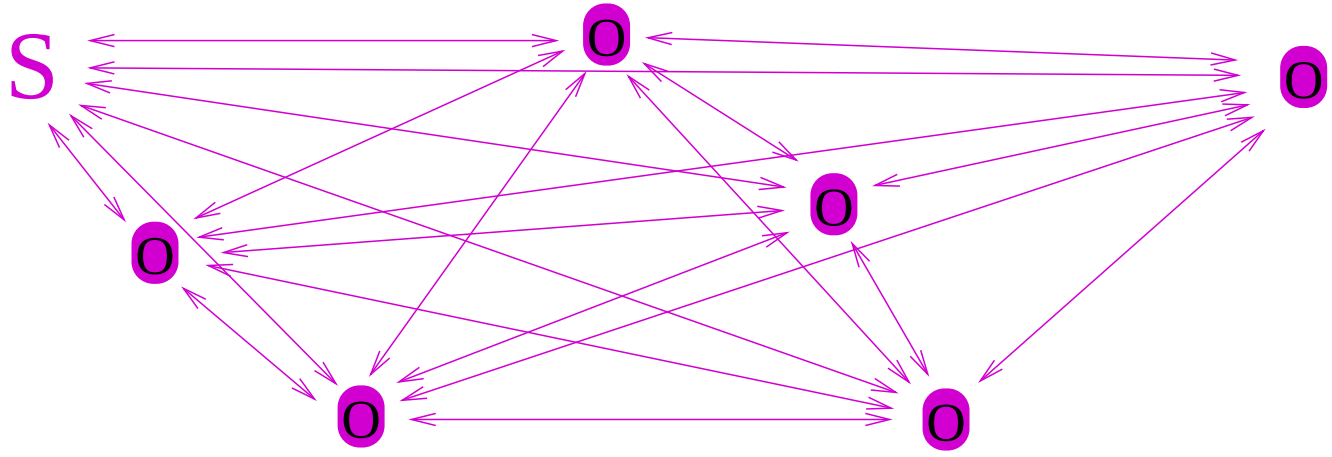

- the underlying topology is unknown.
- it must be built incrementally.
- “independent” links may not be.
- the property to be optimized is not obvious.

Overcast cares about bandwidth

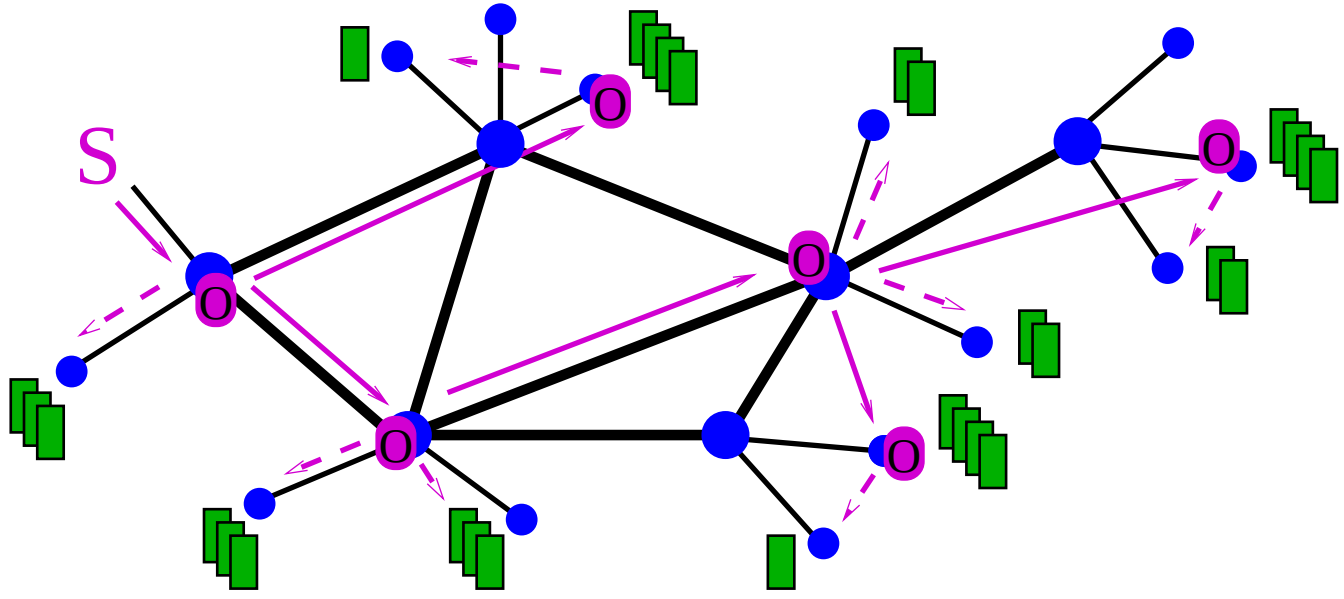

- Latency “doesn’t matter” for our applications.
- A tree rooted at the source is simple and effective.

Building an efficient distribution tree

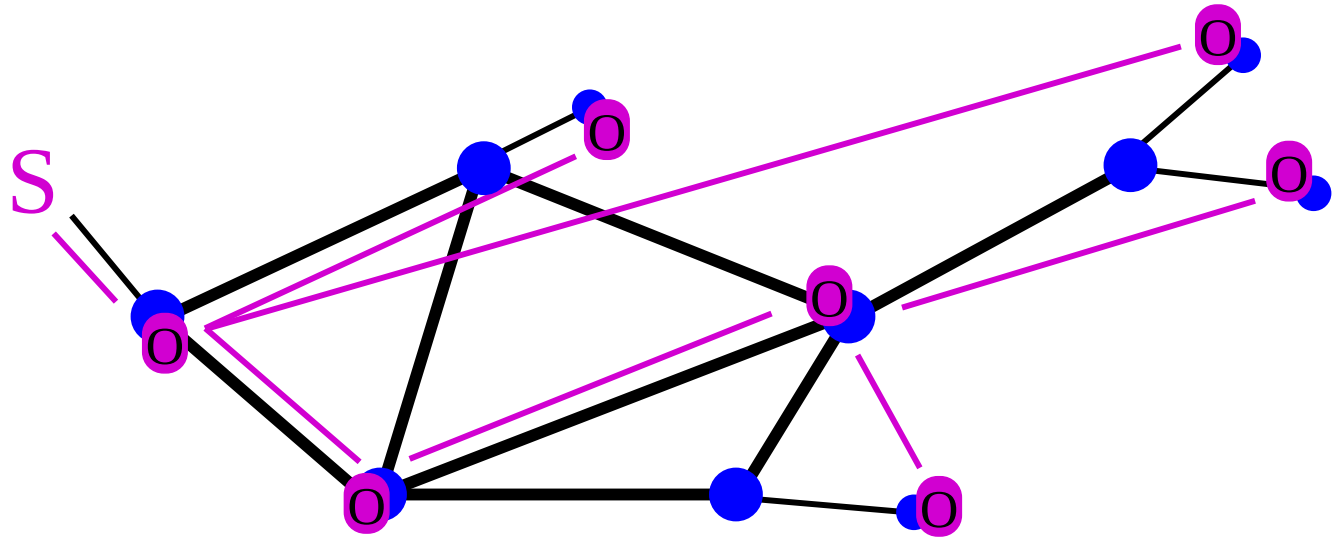

- New nodes attach at the root.

Building an efficient distribution tree

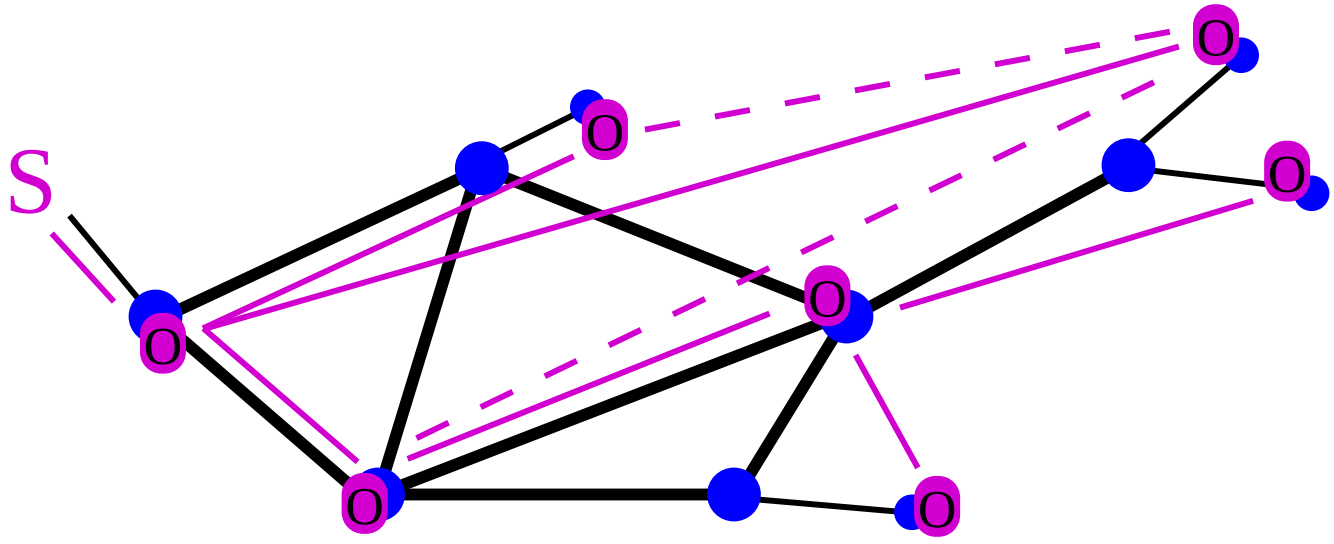

- New nodes attach at the root.
- Periodically consider new parents.

Building an efficient distribution tree

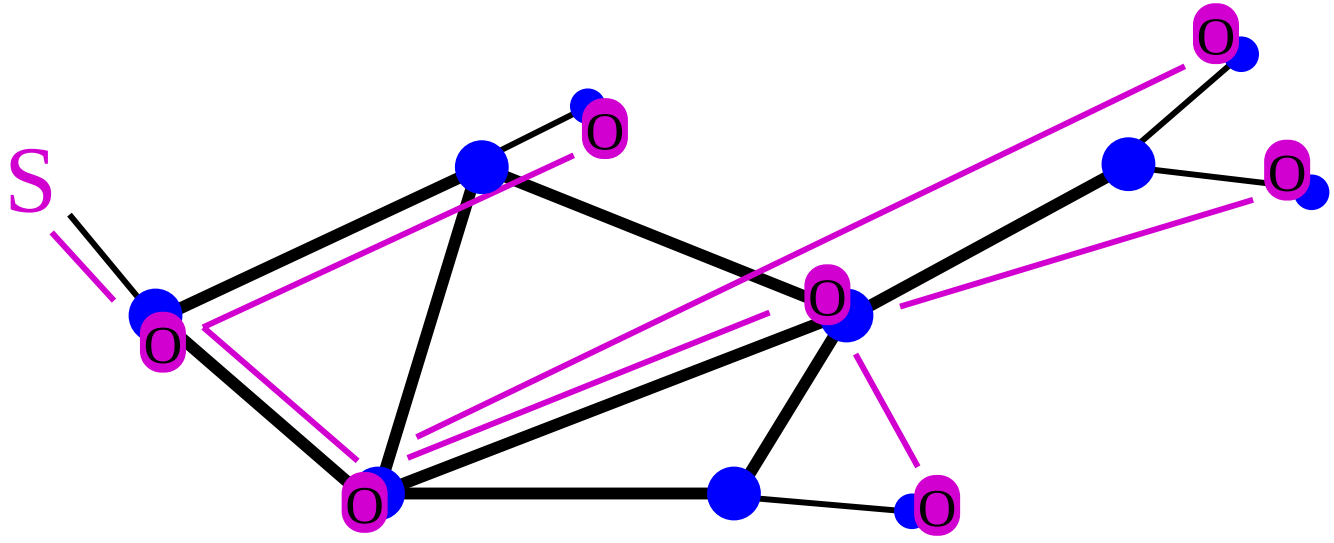

- New nodes attach at the root.
- Periodically consider new parents.
- Ties are broken by hop count, latency.

Building an efficient distribution tree

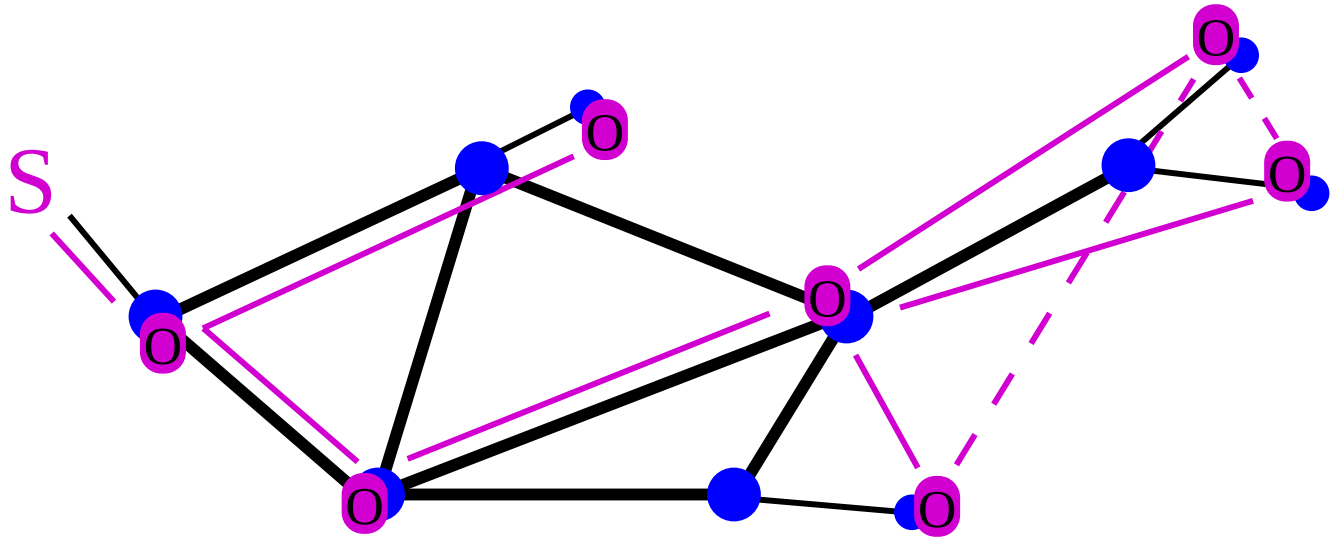

- New nodes attach at the root.
- Periodically consider new parents.
- Ties are broken by hop count, latency.
- Repeat, repeat, repeat. . . repeat

HTTP clients can join Overcast groups

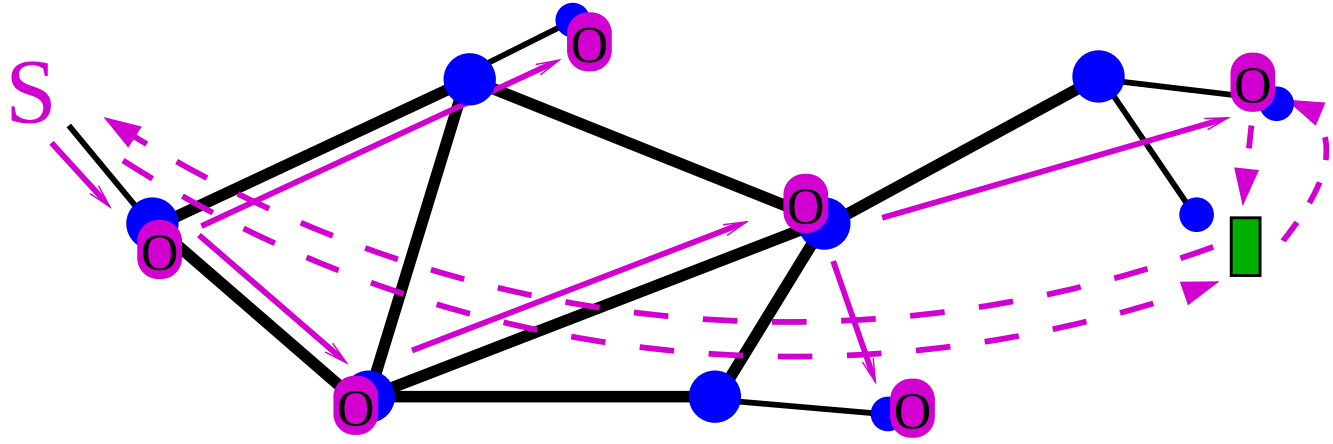

- URLs name groups
- http://{overcast root}//{name/of/group}?{playback details}
- http://overcast.weather.com/radar/usa/ny?10:00am
- “Redirection” to a nearby Overcast node may occur in HTTP or DNS

Tracking topology changes

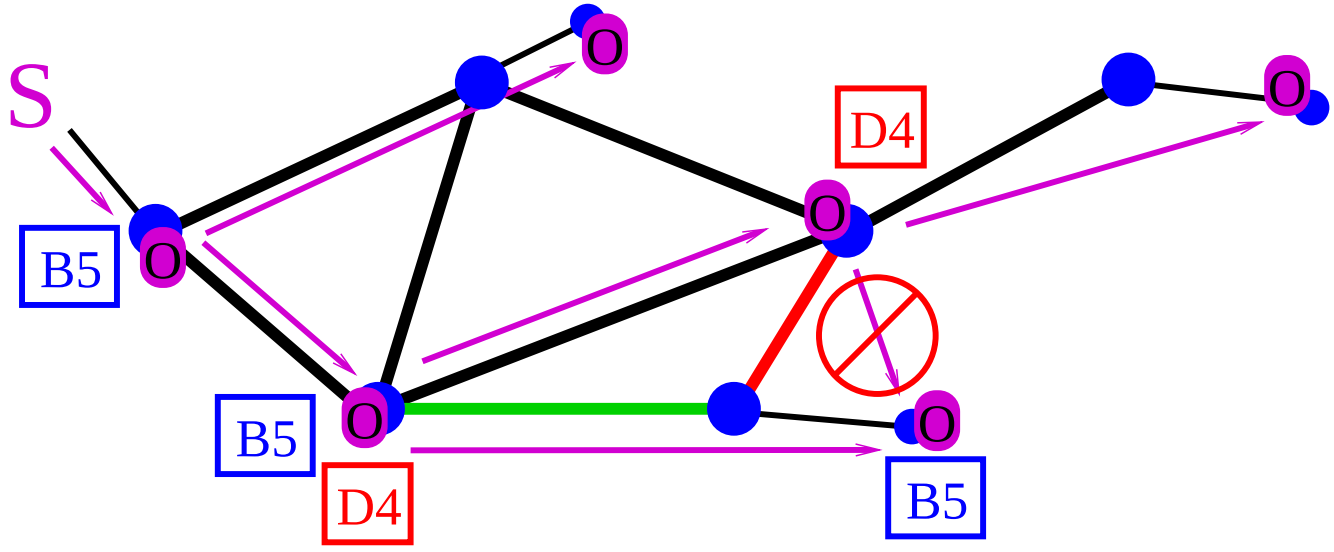

- **Hello** messages from each node to its parent.
- **Birth** and **Death** certificates propagate up the tree.
- Sequence numbers avoid race conditions and quash outdated certificates.

Overcast's "root" is highly available

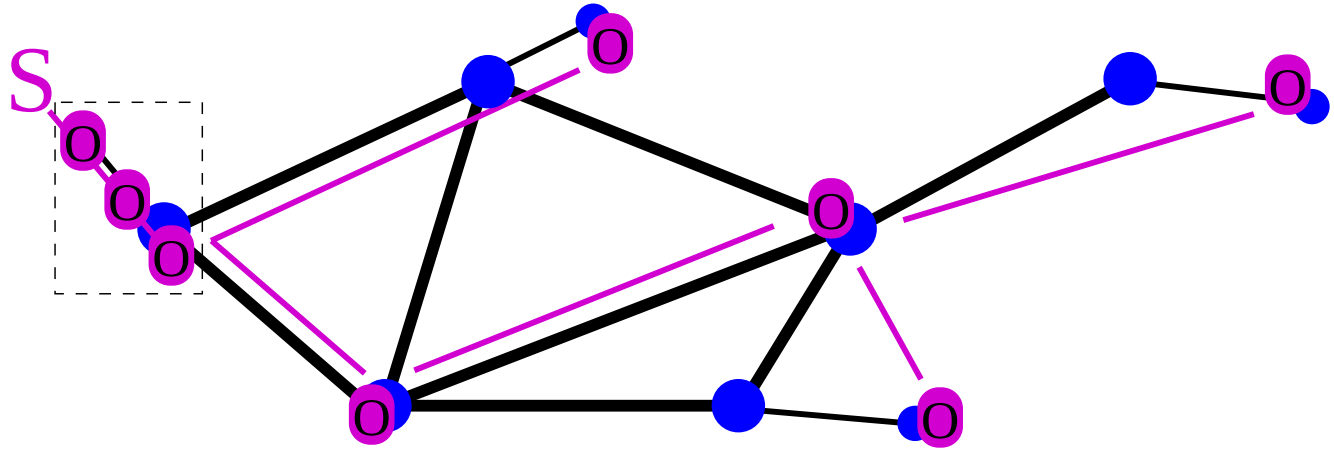

- The root denotes a *pool* of nodes.
- Any node from the pool may act as the root.
- Requests go to a "known live" root using DNS.

Our application of Overcast

We have built a product based on Overcast at Cisco.

- Deployed (up to 200 nodes, many behind NATs and firewalls)
- High quality video distribution on the web.
- Linux and Apache for base functionality.
- Builds distribution trees, predistributes content, redirects HTTP requests to local nodes, detects failures.
- Does not use all of the presented algorithms yet.

How well does Overcast work?

- How much bandwidth does Overcast supply?
- How resource intensive is Overcast?

We have simulated large Overcast networks to find out.

Simulation Scenario

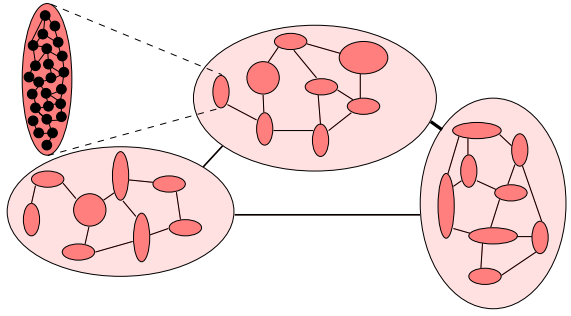

- 600 node “transit-stub” networks. Topologies from GT-ITM.
- Three transit networks, each with eight stub networks.
- We annotate these topologies with bandwidth information.
- Overcast nodes are placed randomly or “backbone-first”

Overcast provides bandwidth

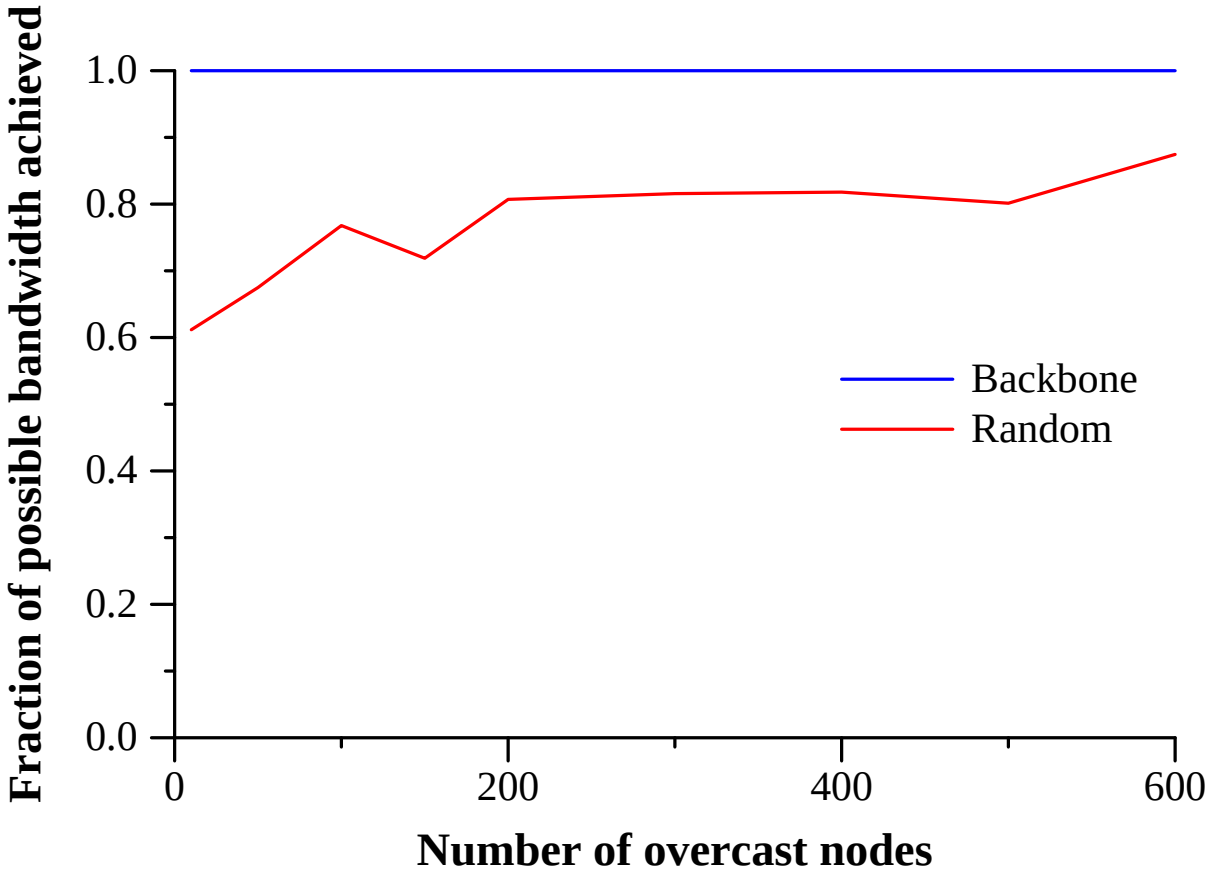

Overcast compared to a bandwidth optimal distribution tree.

Bandwidth isn't Everything

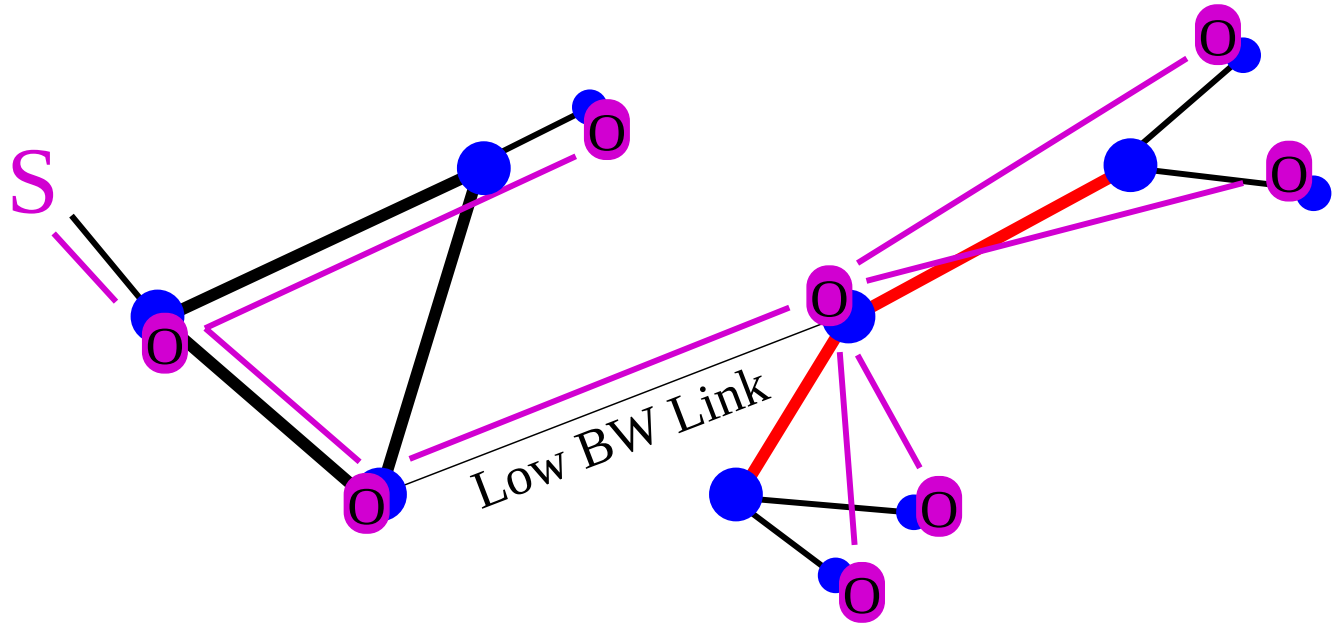

A "perfect" topology with respect to bandwidth may still waste resources.

Overcast reuses links “moderately”

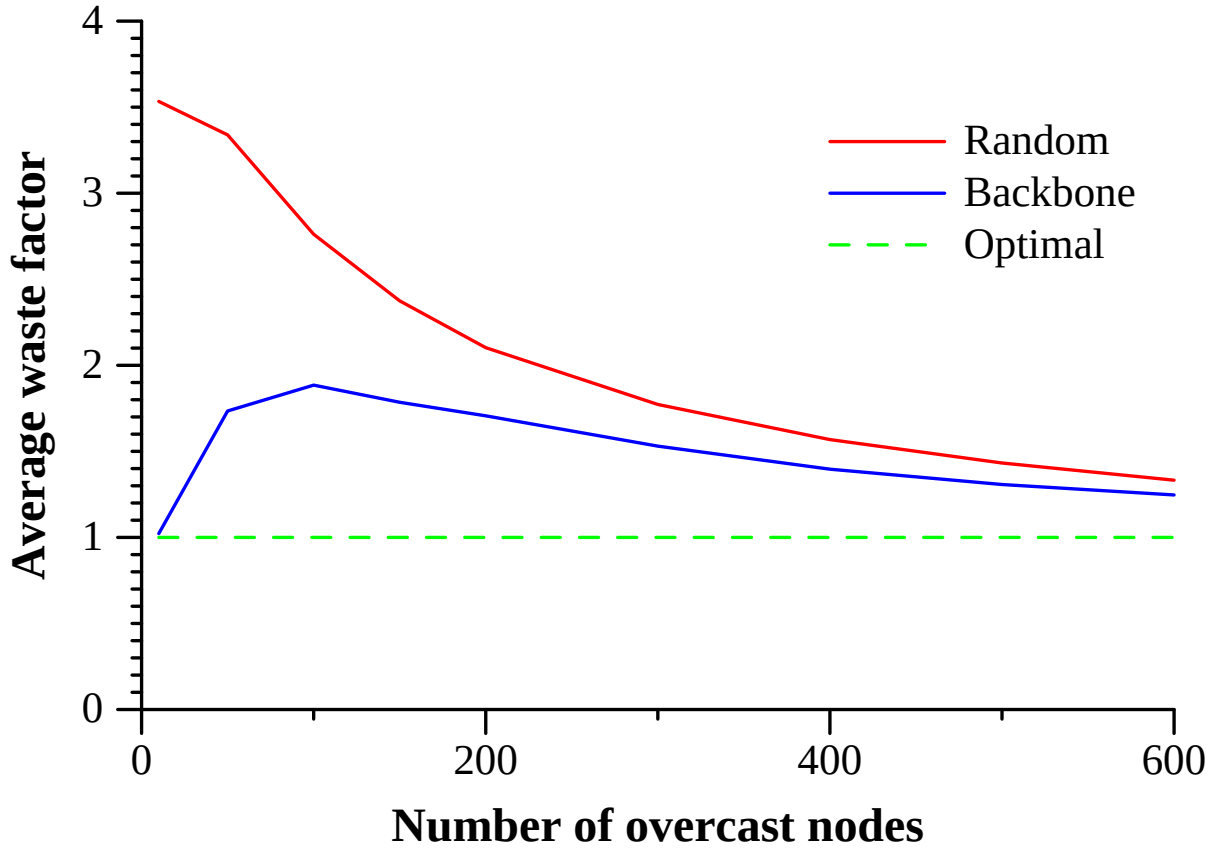

Waste may include links used more than once or “long” paths.

Adding Overcast Nodes

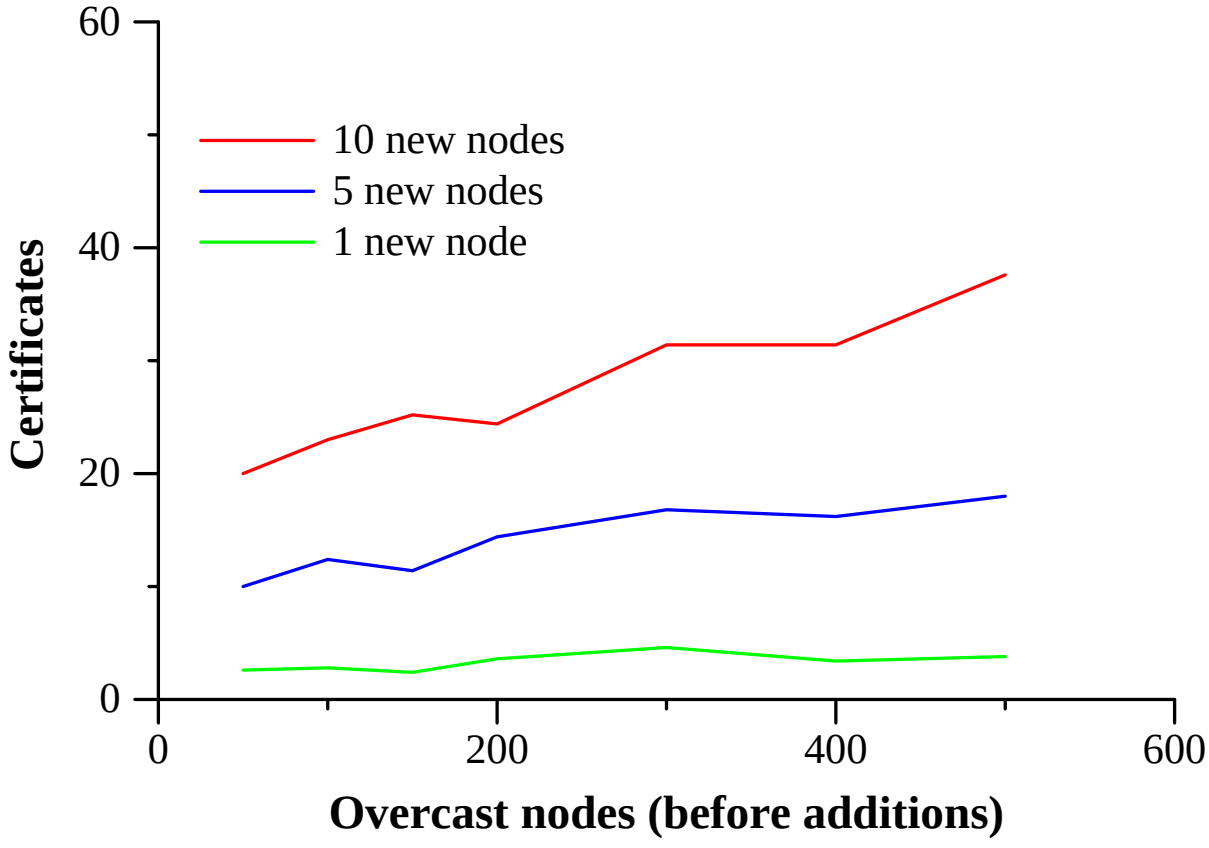

Certificates may be birth *or* death certificates due to reorganization.

Conclusions

- Overlay networks offer the opportunity to **deploy** interesting network services.
- Overcast builds efficient distribution trees without global knowledge.
- Overcast provides
  1. Live broadcasting
  2. “Time-shifting” or On Demand access
  3. Reliability