


The Case For Cooperative Networking

Venkat Padmanabhan, Microsoft Research

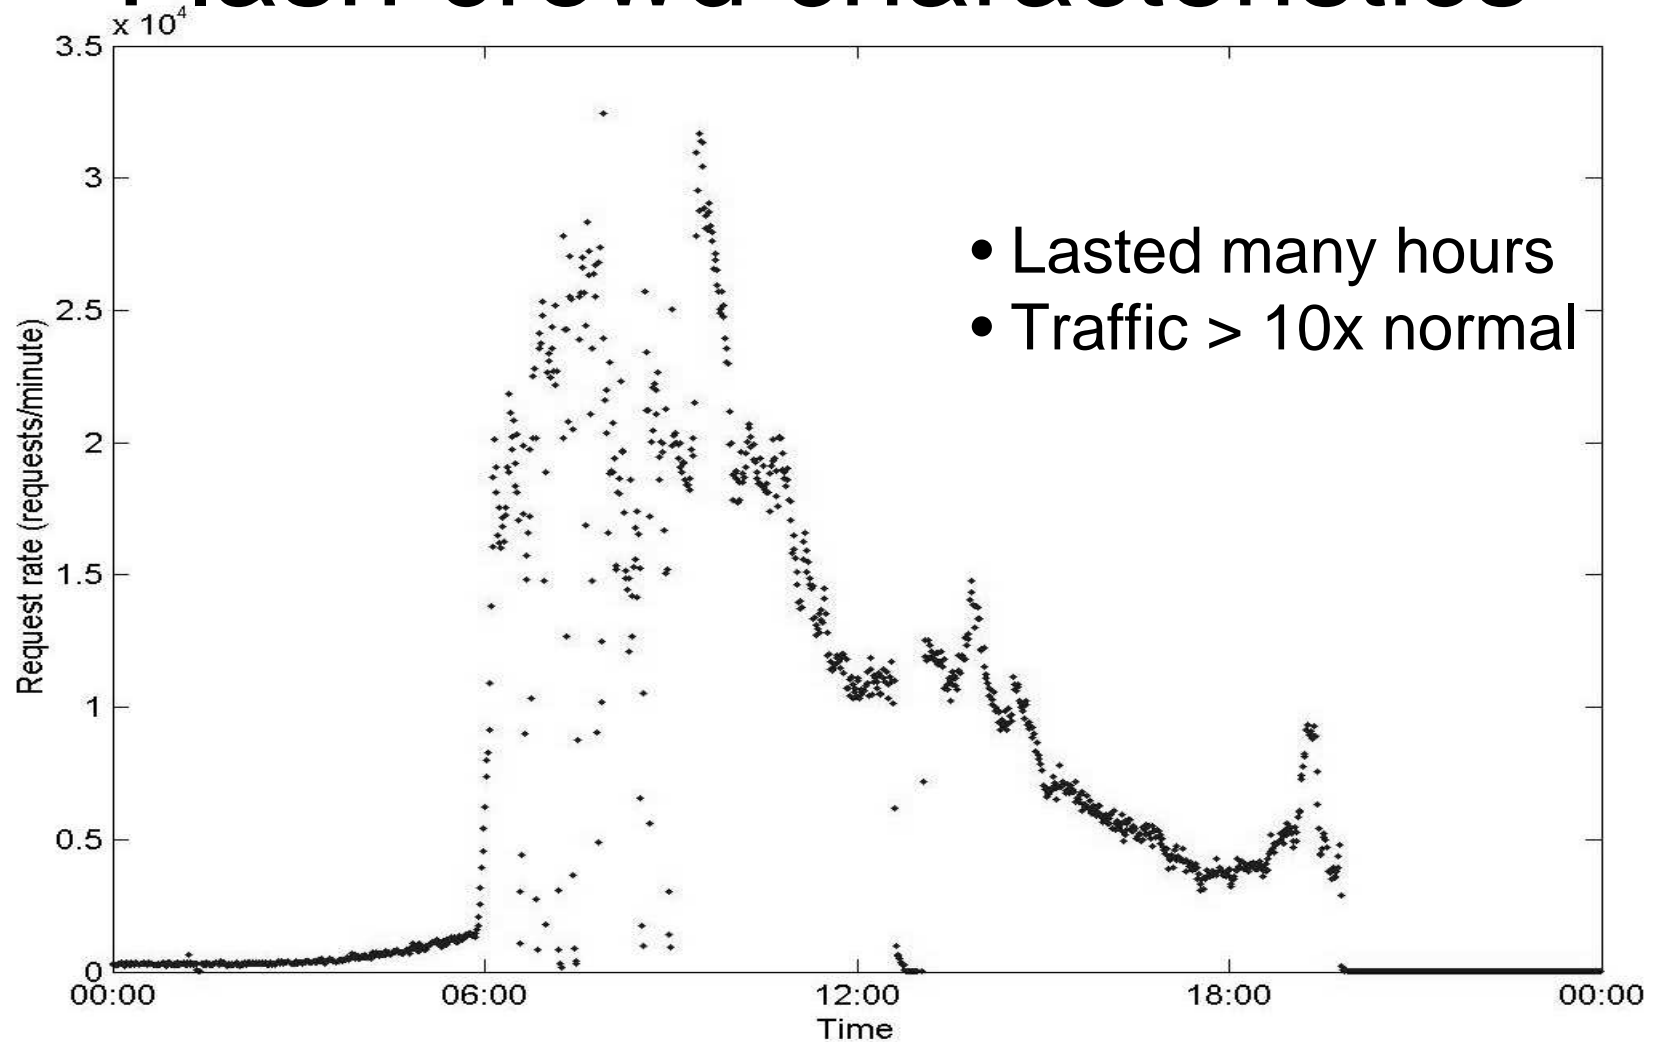
Kay Sripanidkulchai, CMU

Cooperative Networking



- Address problem case of client-server communication
- Complement traditional client-server communication
- Minimal assumption about peer participation
- Study in the context of the flash crowd problem


Flash crowd characteristics




Where is the bottleneck?

- Disk?
 - No, usually requests are for popular content
 - 90% of requests are for 100 files
- CPU?
 - Perhaps for dynamic content
- Bandwidth?
 - Yes, most likely close to the server
 - 65% of servers have bottleneck bandwidths of less than 1.5 Mbps (Stefan Saroiu, U.W.)


Operation of CoopNet




Operation of CoopNet




Operation of CoopNet




Operation of CoopNet




Operation of CoopNet



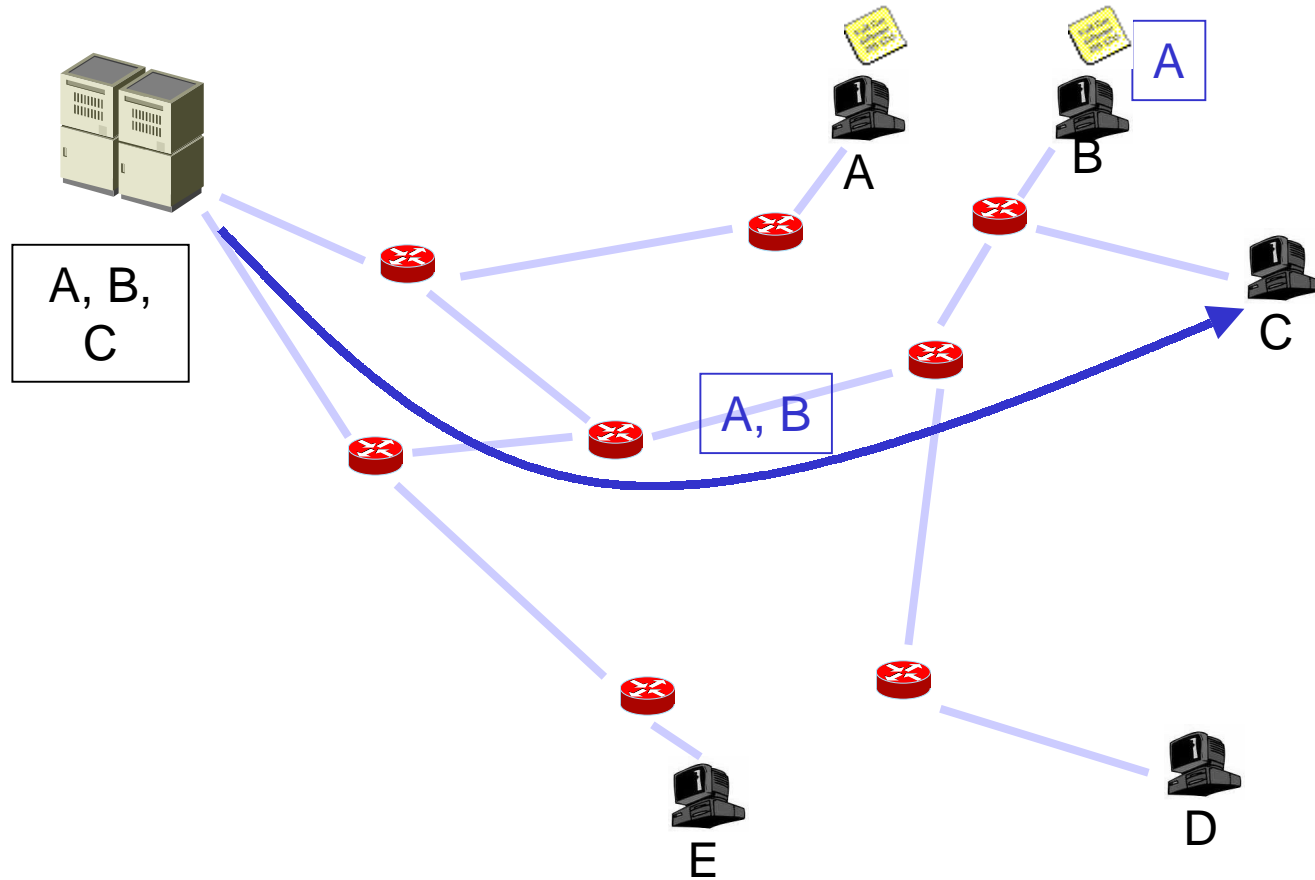
Operation of CoopNet




Operation of CoopNet



Operation of CoopNet




Operation of CoopNet



Operation of CoopNet

- Availability of centralized server simplifies lookups
 - Similar to Napster
- Server bandwidth savings is $\sim 100x$
 - 200 B redirection vs. 20 kB webpage
- Effective for substantial flash crowds

Implementation



- HTTP Pragma to signal between client and server
- Server caches IP addresses of CoopNet clients that have recently accessed each file
- Server includes small peer list in redirection response
- Implementation details
 - Server-side: ISAPI filter and extension for Microsoft IIS Server
 - Client-side: local proxy

Peer selection

- BGP prefix clusters
 - Light-weight, but crude
- Delay coordinates
- Match bottleneck bandwidth
 - Peers report bottleneck bandwidth in request
 - Server returns peers with matching bandwidth
 - No incentive for under-reporting
- Concurrent downloads

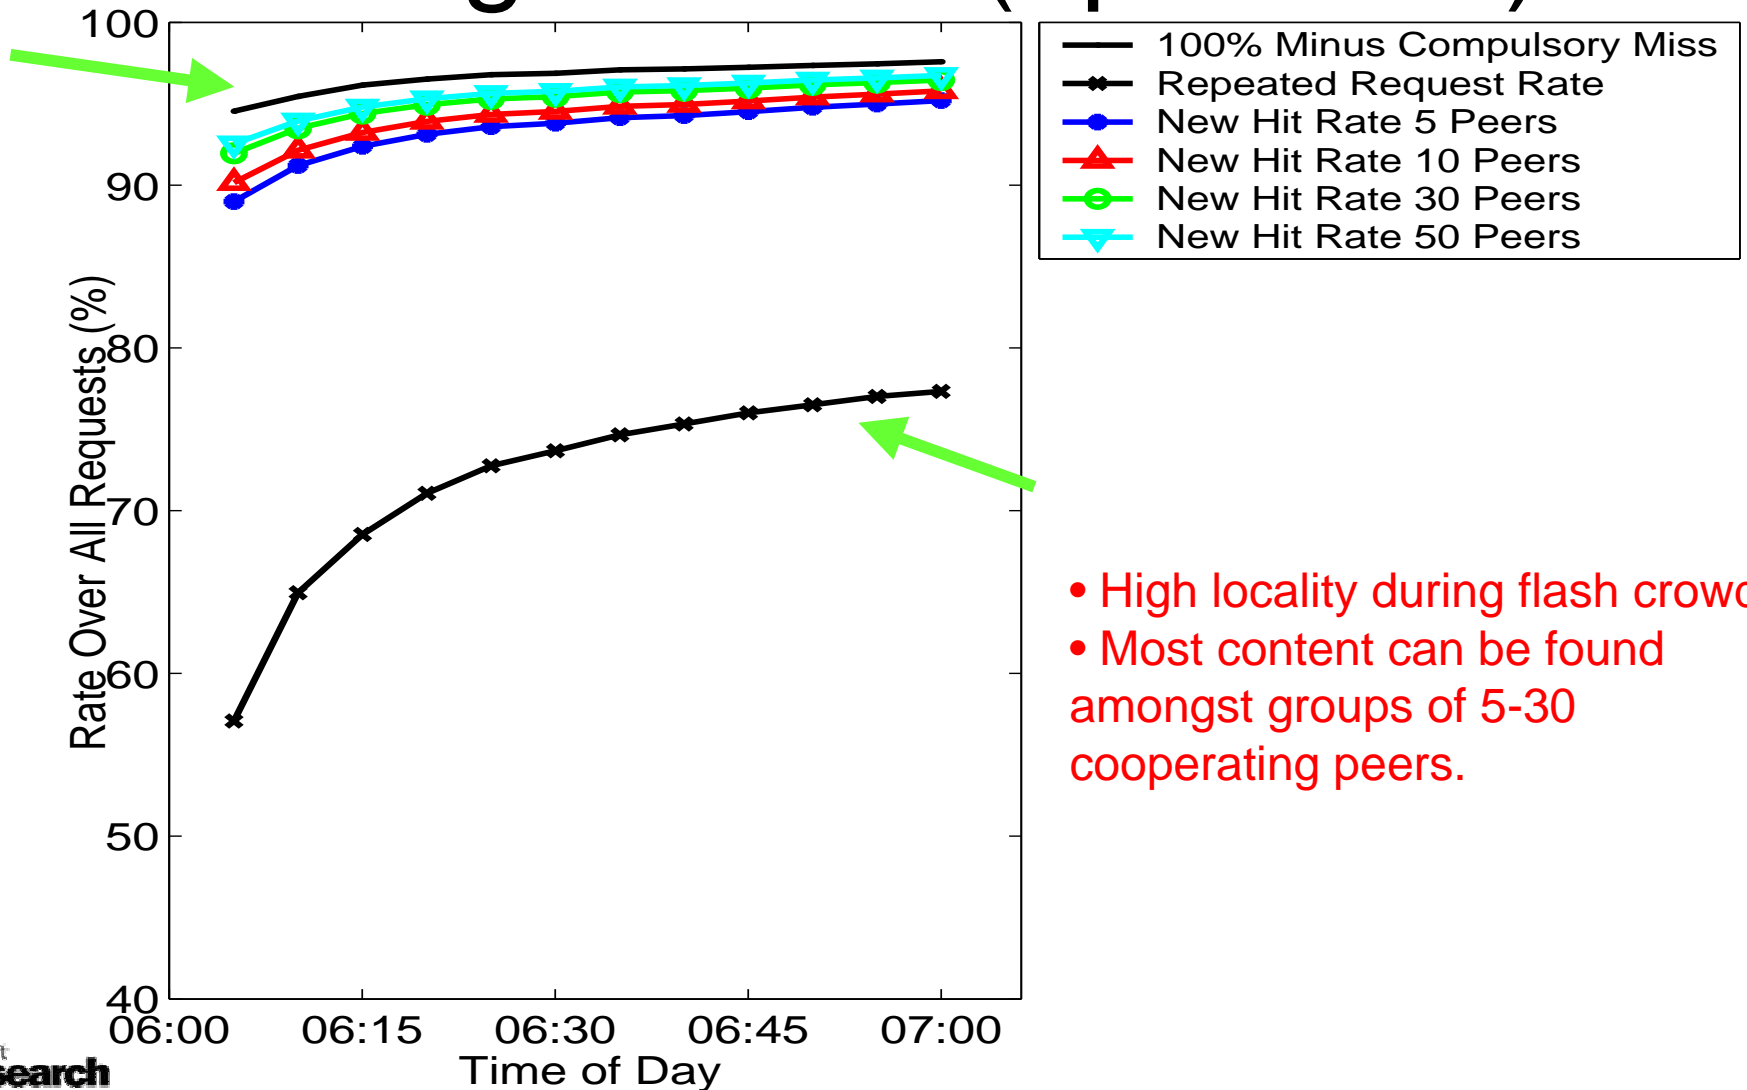
Initial evaluation

- Goal
 - To study the potential benefits and overhead of CoopNet
- Evaluation strategy
 - Trace-based simulations using MSNBC Sept 11, 2001 traces
 - Only 200 clients cooperate
 - Server returns 5-50 IP addresses of recent CoopNet clients

Performance metrics

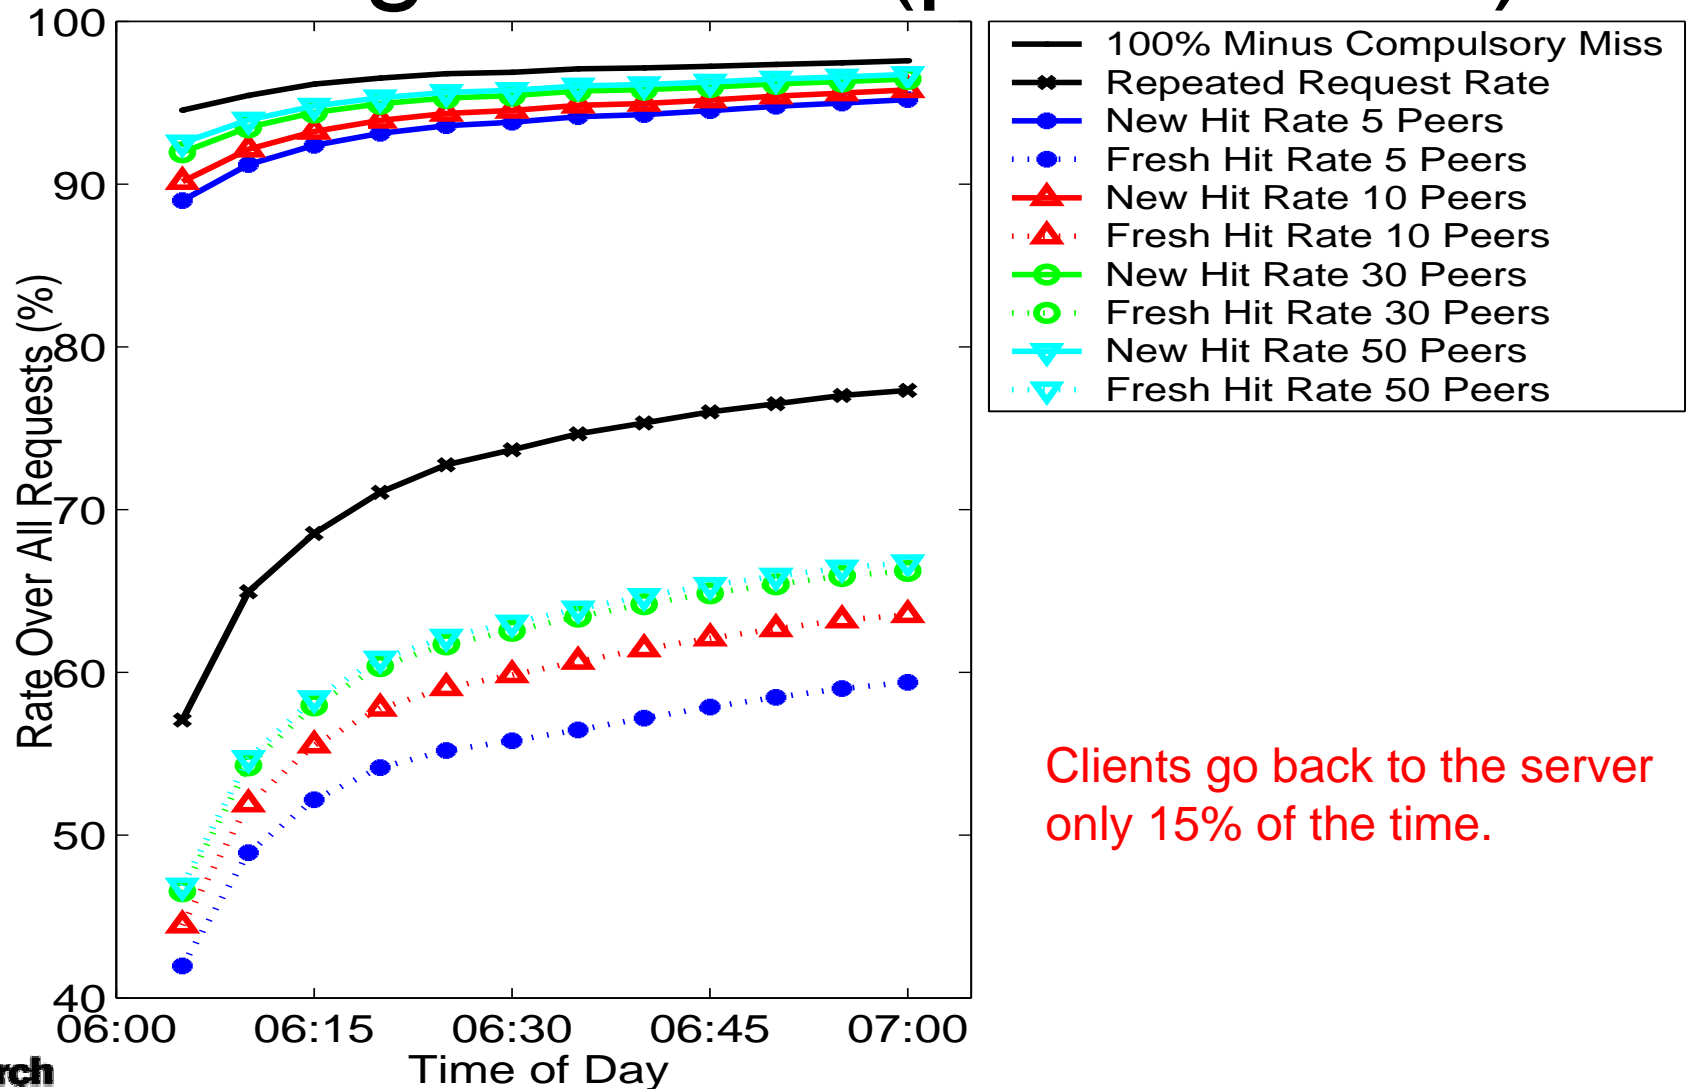
- How often content is found amongst peers
 - Optimistic: Files are not modified during the simulation
 - Pessimistic: Files are modified at every access
- Additional load incurred by participating
- How often cooperating peers are nearby

Finding content (optimistic)



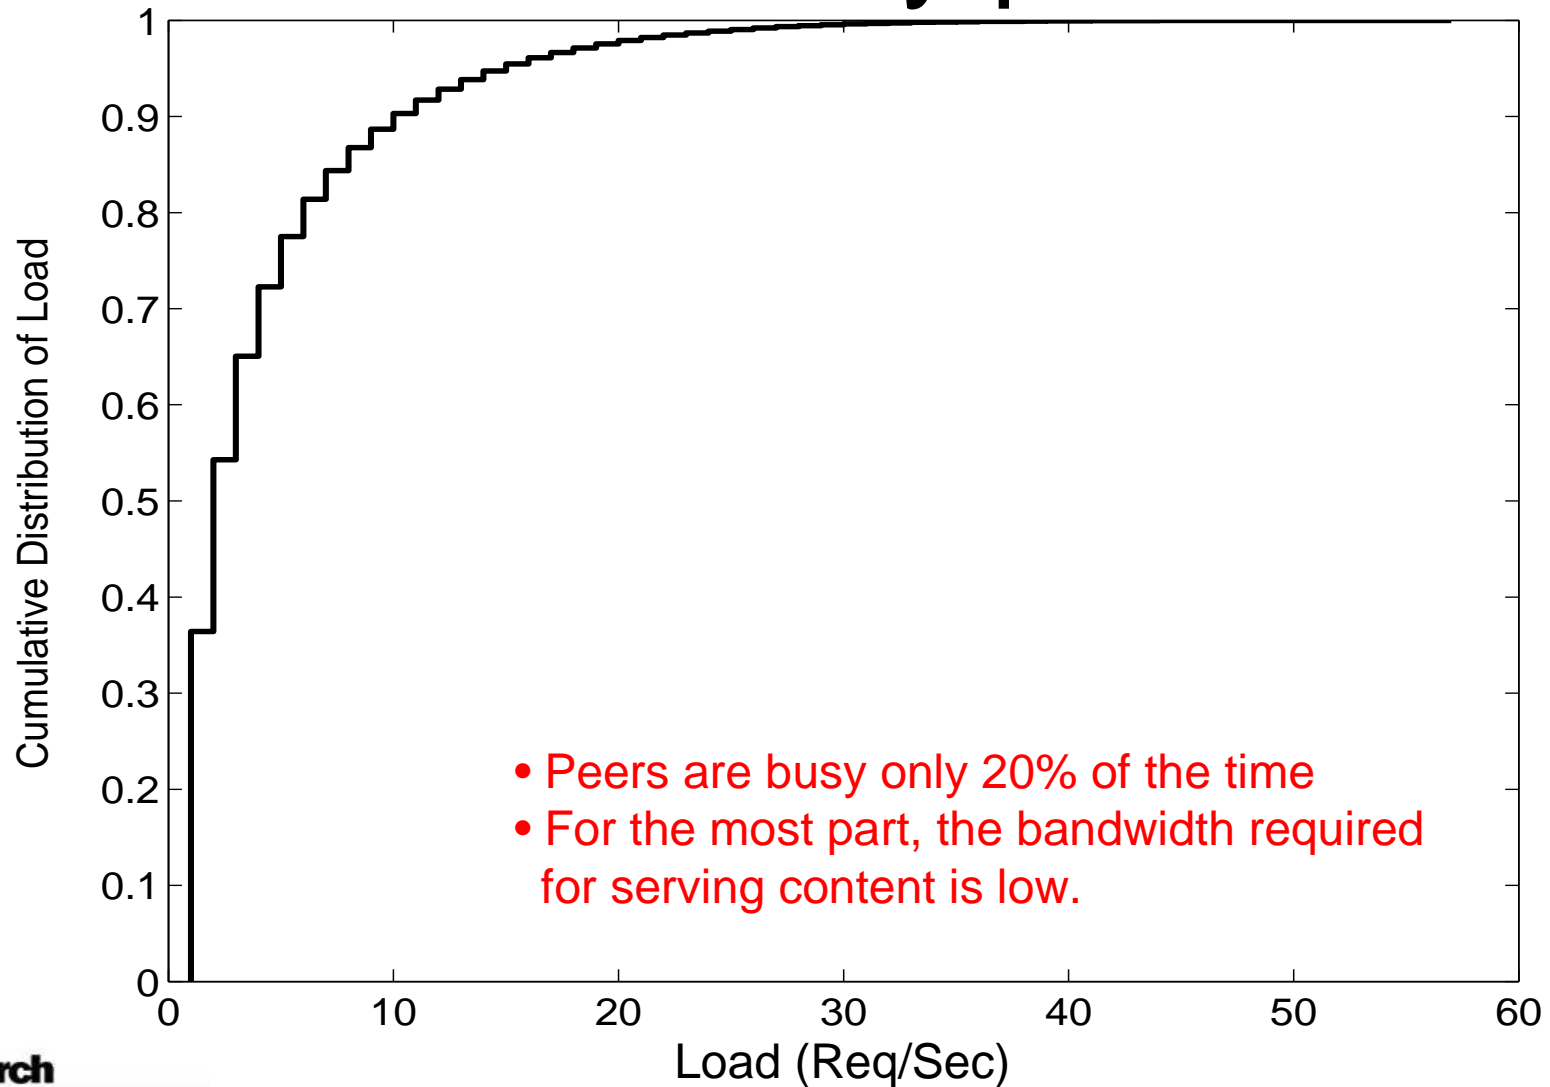
- High locality during flash crowd.
- Most content can be found amongst groups of 5-30 cooperating peers.

Finding content (pessimistic)



Clients go back to the server only 15% of the time.

Load at busy peers



Nearby peers

- Proximity criterion: matching BGP prefix
- First 30 minutes of the trace
 - 563,284 clients from 69,778 BGP prefix clusters
- High probability of finding another peer in the same prefix cluster
 - First 2 minutes: 80%
 - First 30 minutes: 90%

Open issues

- **Server may still be overwhelmed**
 - Large number of CoopNet peers may lead to large number of redirection messages
 - Small number of CoopNet peers does not reduce load at the server
- **Distributed content lookup**
 - Avoid the server to the extent possible, but can fall back to the server when necessary
 - Design algorithm for small groups of cooperating peers

Open issues

- What about distributed hash table schemes?
 - Necessary given small group sizes?
 - Performance under highly dynamic membership?
 - What is a key?

Alternative approaches

- Proxy caching
 - Deployment barriers
 - Not effective when clients are scattered across the Internet
- Commercial CDN's
 - Not cost-effective for small websites
- P2P system of servers
 - More feasible than client-client cooperation?

Summary

- Existing client-server applications can benefit from selective use of peer-to-peer communications
- Availability of a server simplifies system design
- Distributed lookups only needed for small groups
- CoopNet is a promising solution for flash crowds
- We are also studying CoopNet for streaming media content (NOSSDAV '02)
- Visit our webpage at www.research.microsoft.com/~padmanab/projects/CoopNet/