Going Viral: Flash Crowds in an Open CDN

IMC 2011 (Short Paper)

Patrick Wendell, U.C. Berkeley
Michael J. Freedman, Princeton University

What is a Flash Crowd?

- "Slashdot Effect", "Going Viral"

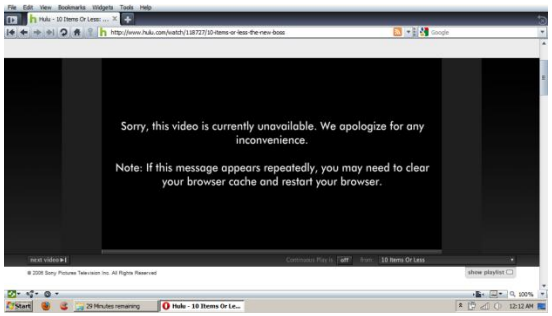

- Exponential surge in request rate (precisely defined in paper)

Key Questions

- What are primary drivers of flash crowds?
- How effective is cache cooperation during crowds against CDNs?
- How quickly do we need to provision resources to meet crowd traffic?

CoralCDN

- Network of ~300 distributed caching proxies

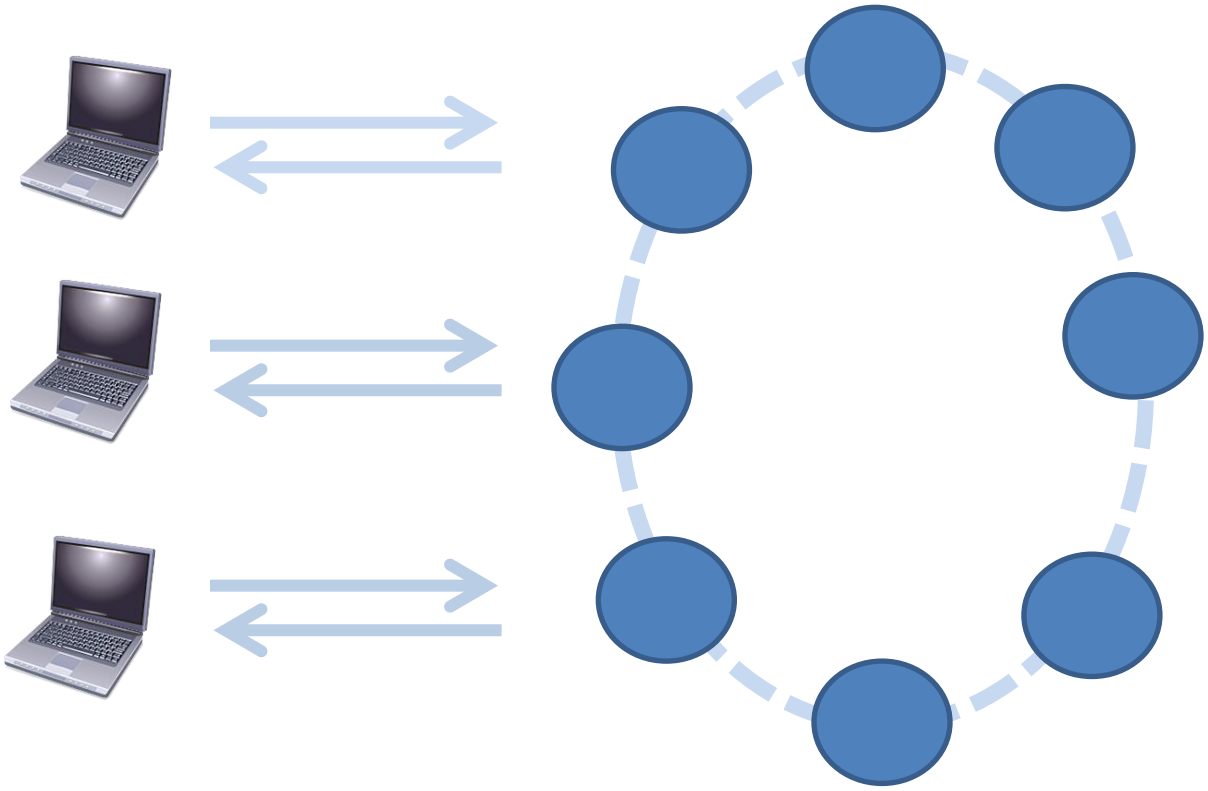

HTTP Clients

CoralCDN Proxies

Origin Server

CoralCDN

- Network of ~300 distributed caching proxies

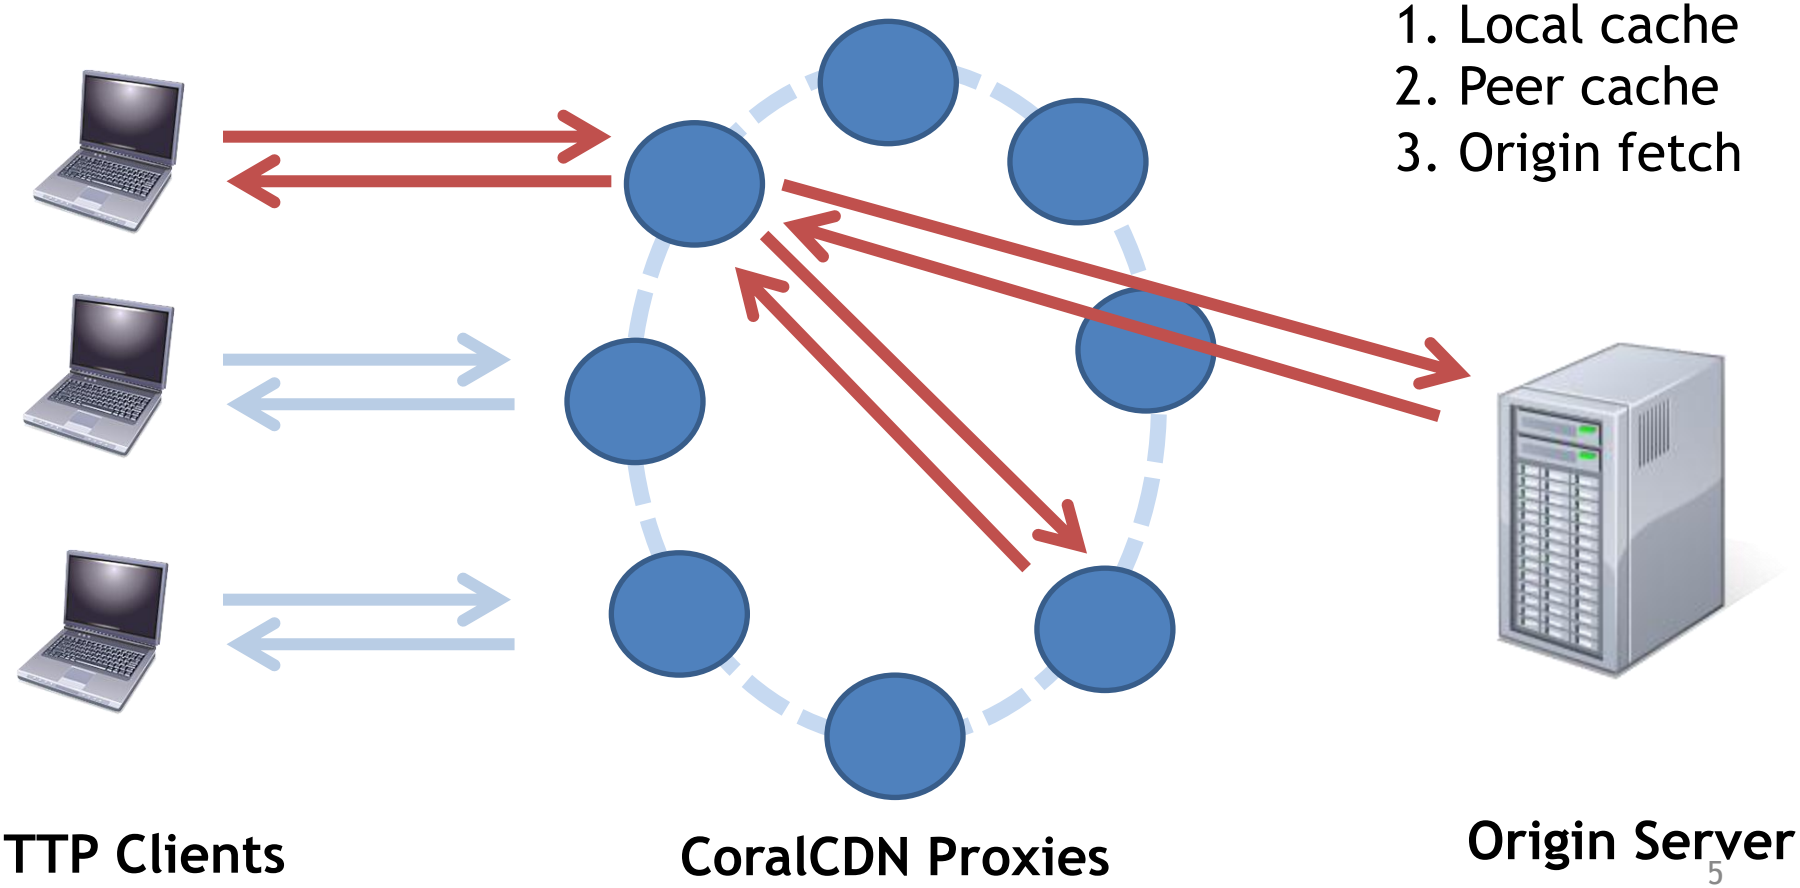

The Data

- Complete CoralCDN trace over 4 years
- 33 Billion HTTP requests
- Per-request logging
  - <Time, URL, client IP, proxy IP, content cached?, ...>

Finding Crowds

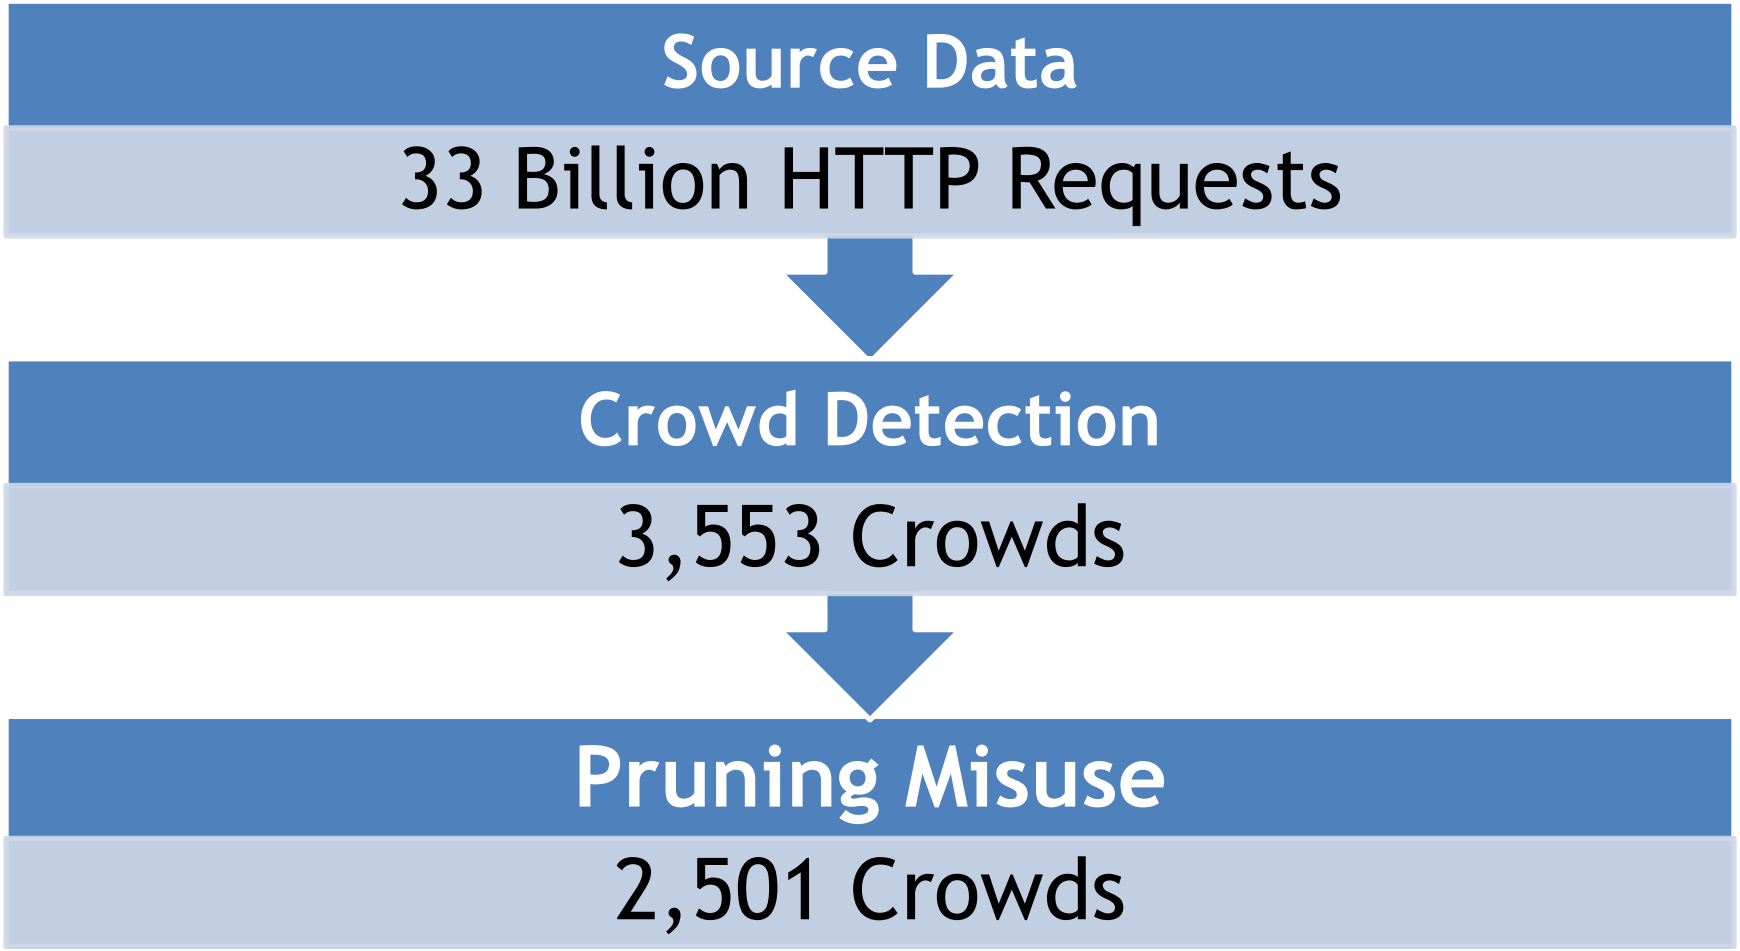

Crowd Sources

Common Referrers

Referrer	# Crowds
digg.com	123
reddit.com	20
stumbleupon.com	15
google.com	11
facebook.com	10
dugmirror.com	8
duggback.com	4
twitter.com	4

Common Referrers

Referrer	# Crowds
digg.com	123
reddit.com	20
stumbleupon.com	15
google.com	11
facebook.com	10
dugmirror.com	8
duggback.com	4
twitter.com	4

Common Referrers

Referrer	# Crowds
digg.com	123
reddit.com	20
stumbleupon.com	15
google.com	11
facebook.com	**10**
dugmirror.com	8
duggback.com	4
twitter.com	**4**

Common Referrers

Referrer	# Crowds
digg.com	123
reddit.com	20
stumbleupon.com	15
google.com	11
facebook.com	10
dugmirror.com	8
duggback.com	4
twitter.com	4

CDN Caching Strategies

Cooperation in Caching

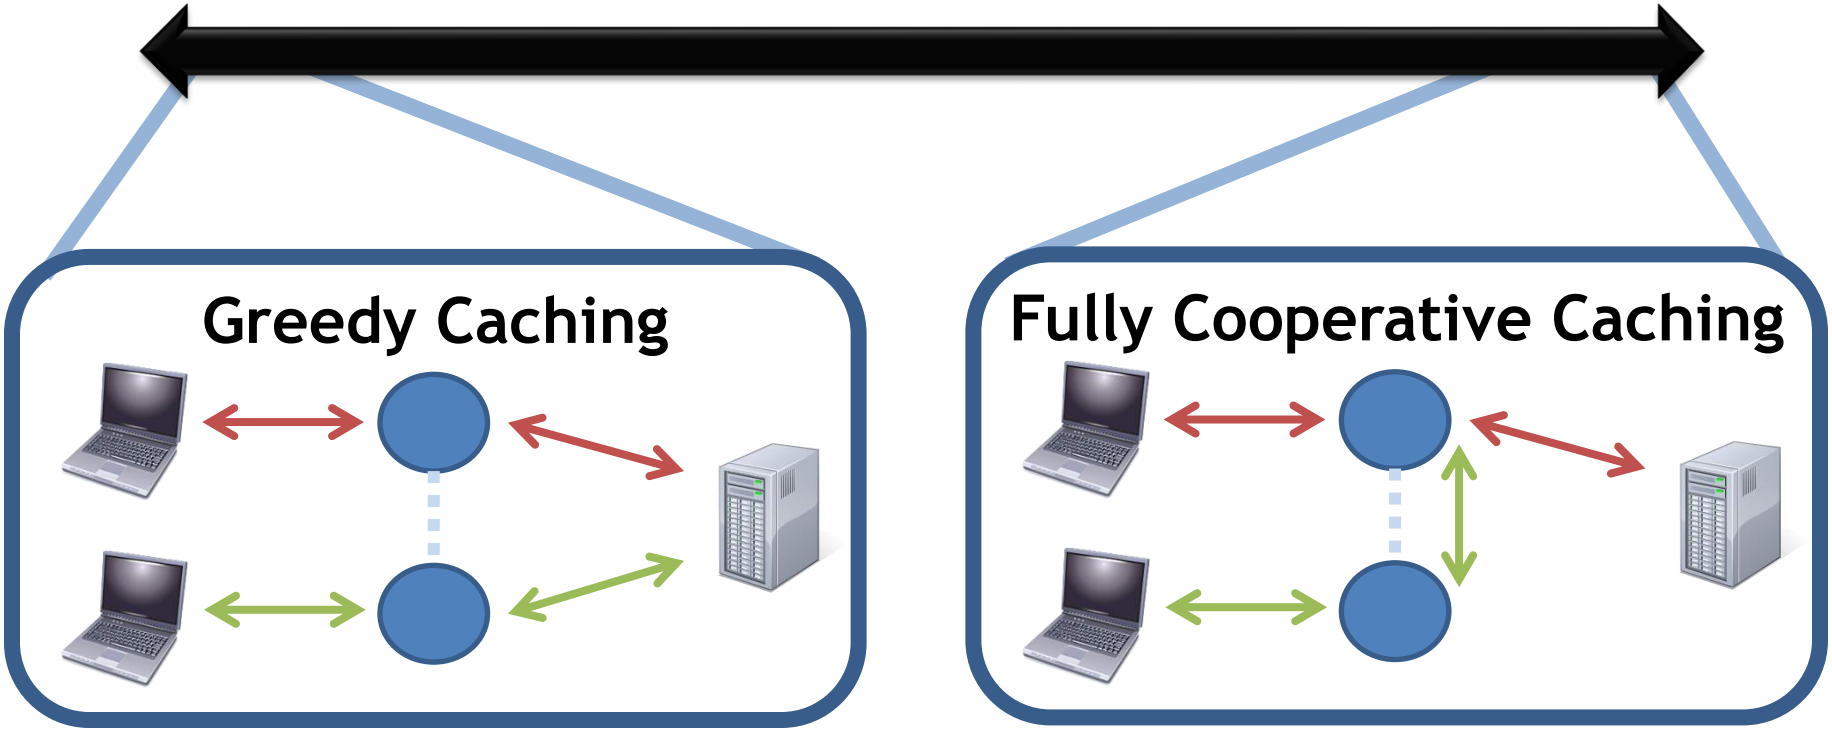

Benefits of Cooperation?

- Depends how clients distribute over proxies

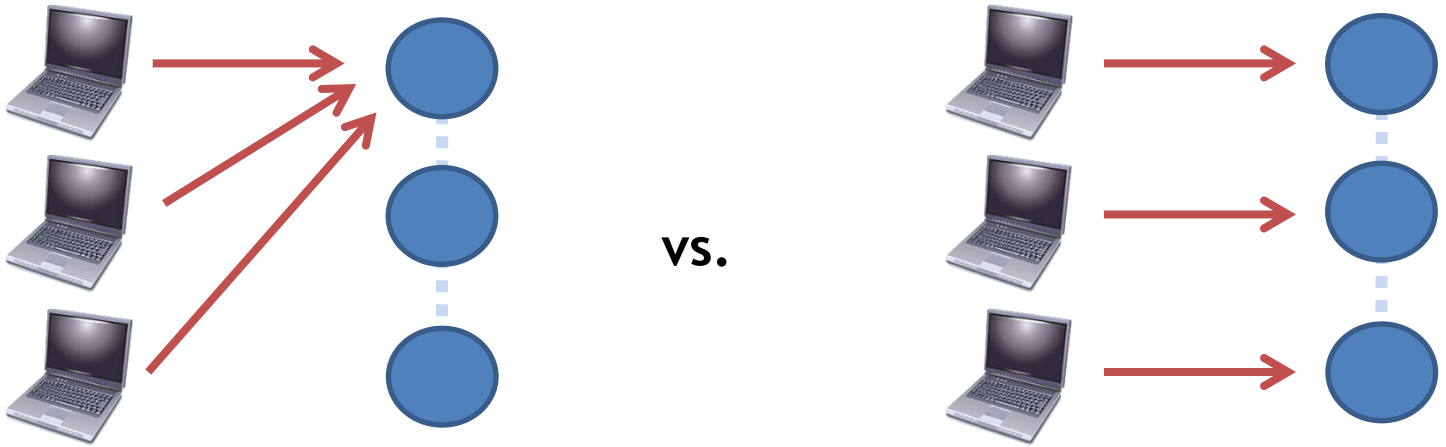

vs.

- Depends how many objects a crowd contains

Clients Use Many Proxies

- Clients globally distributed, even during crowds
- Most caches participate in most crowds

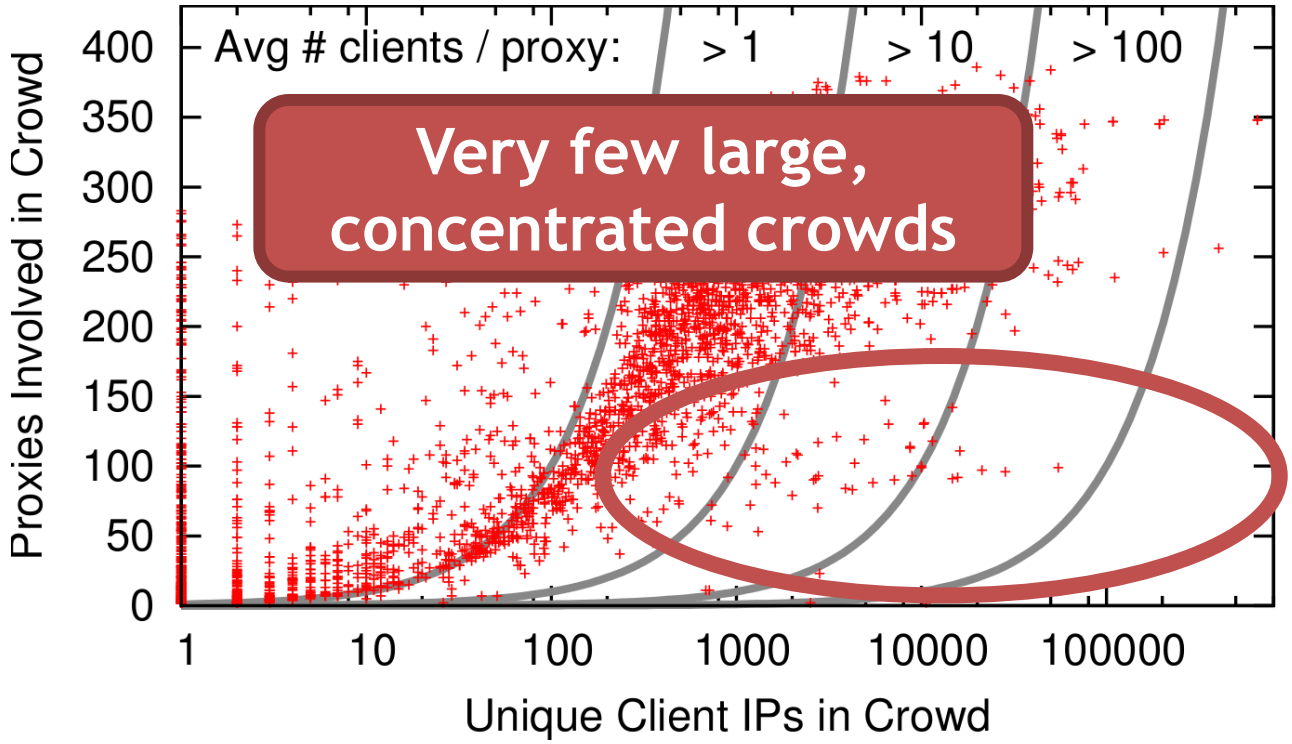

Crowds Contain Many Objects

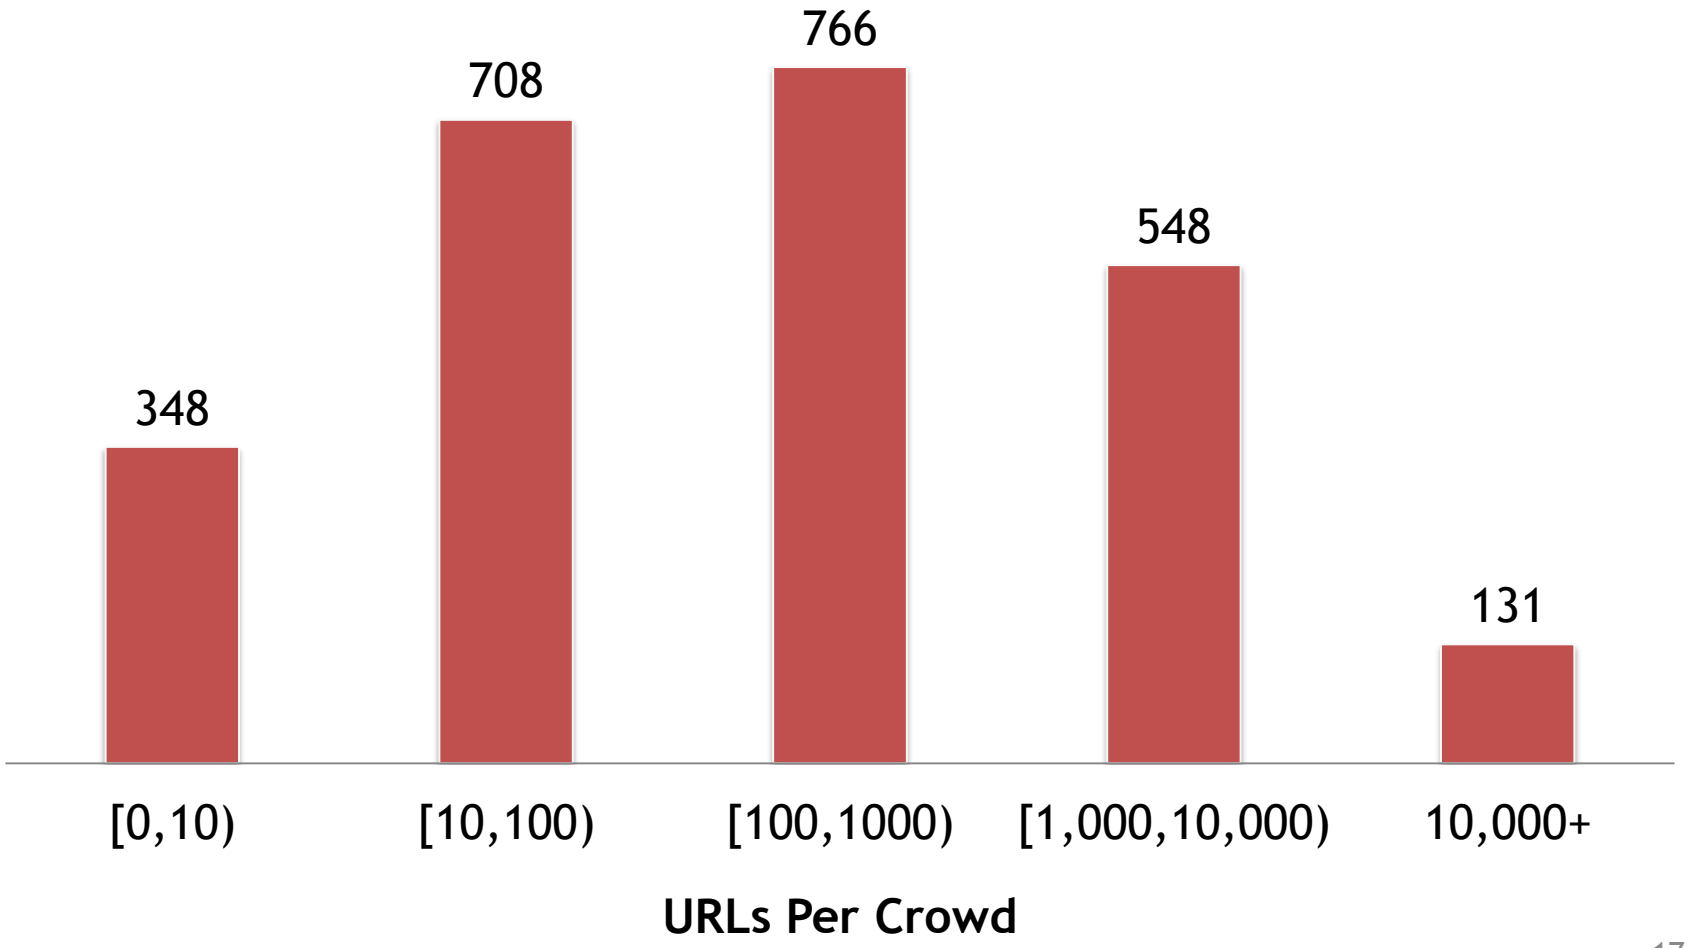

Benefits from Cooperation

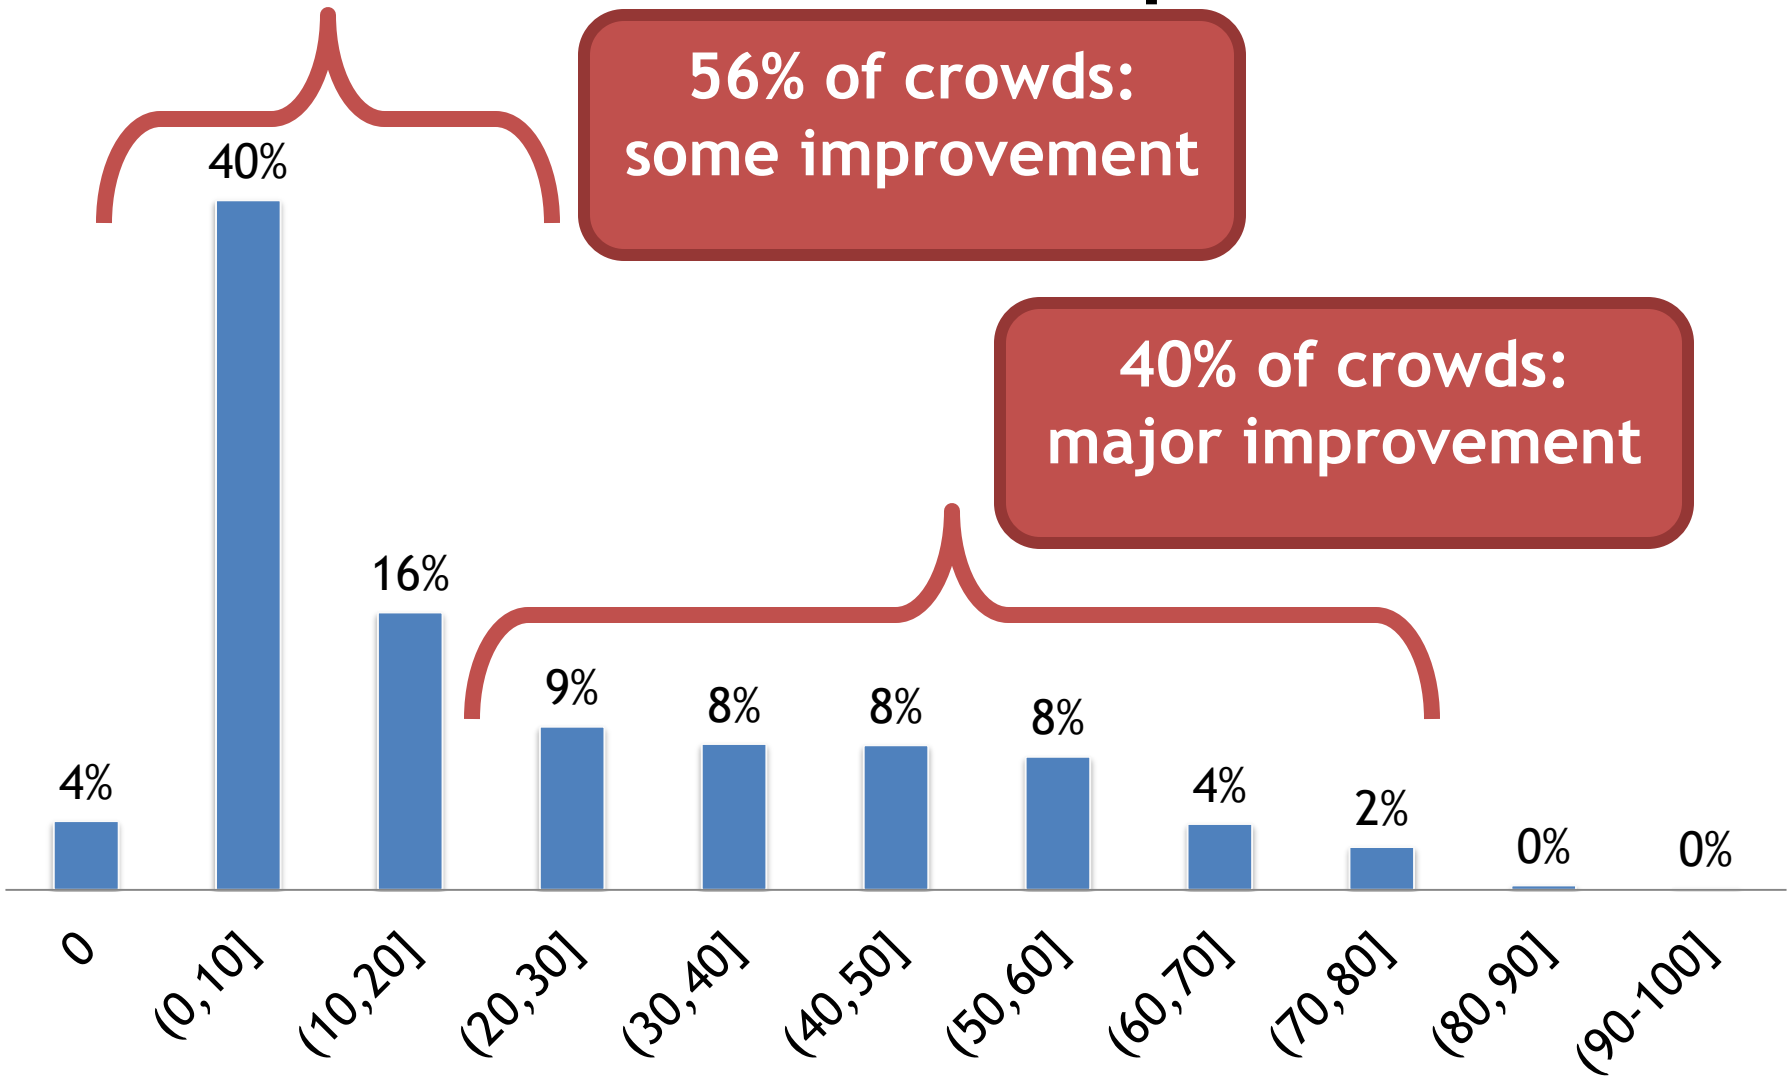

Provisioning Resources For Crowds

Examples of Resource Provisioning

- CDN: static content
  - Expand cache set for particular domain
  - Ω(Seconds)
- Cloud Computing Platform: dynamic service
  - Spin up new VM instances
  - Ω(Minutes)
- If you squint, these are similar problems

Required Resource Spin-up Time

Spin-up	% Crowds Underprovisioned
10 Minutes	75%
1 Minute	50%
10 Seconds	

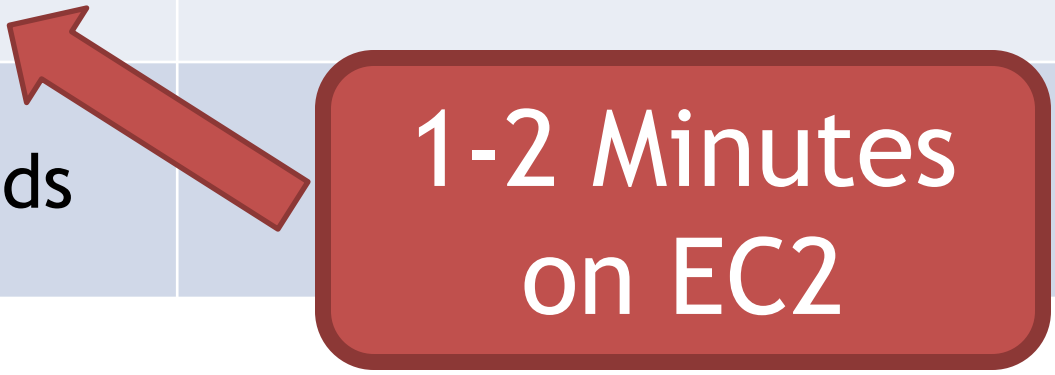

Conclusions

- What are primary drivers of flash crowds?
  - Aggregators and portals, but also social/search
- How effective is cache cooperation during crowds against CDNs?
  - Large benefit for 40% of crowds
- How fast do we need to provision resources during crowds?
  - Likely require sub-minute responsiveness

Questions?

cs.berkeley.edu/~pwendell

Extra Slides / Charts

Actual Spin-up Times on EC2

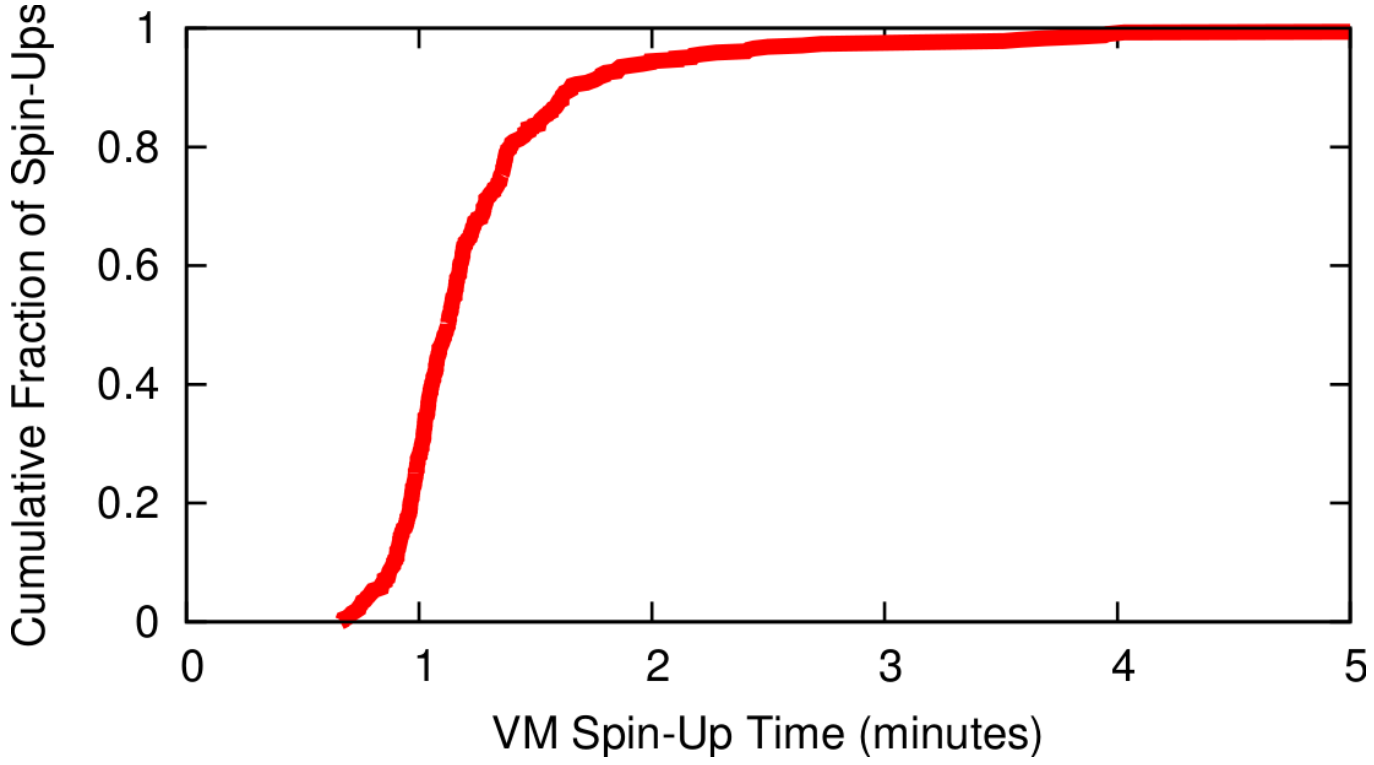

How Fast is Fast?

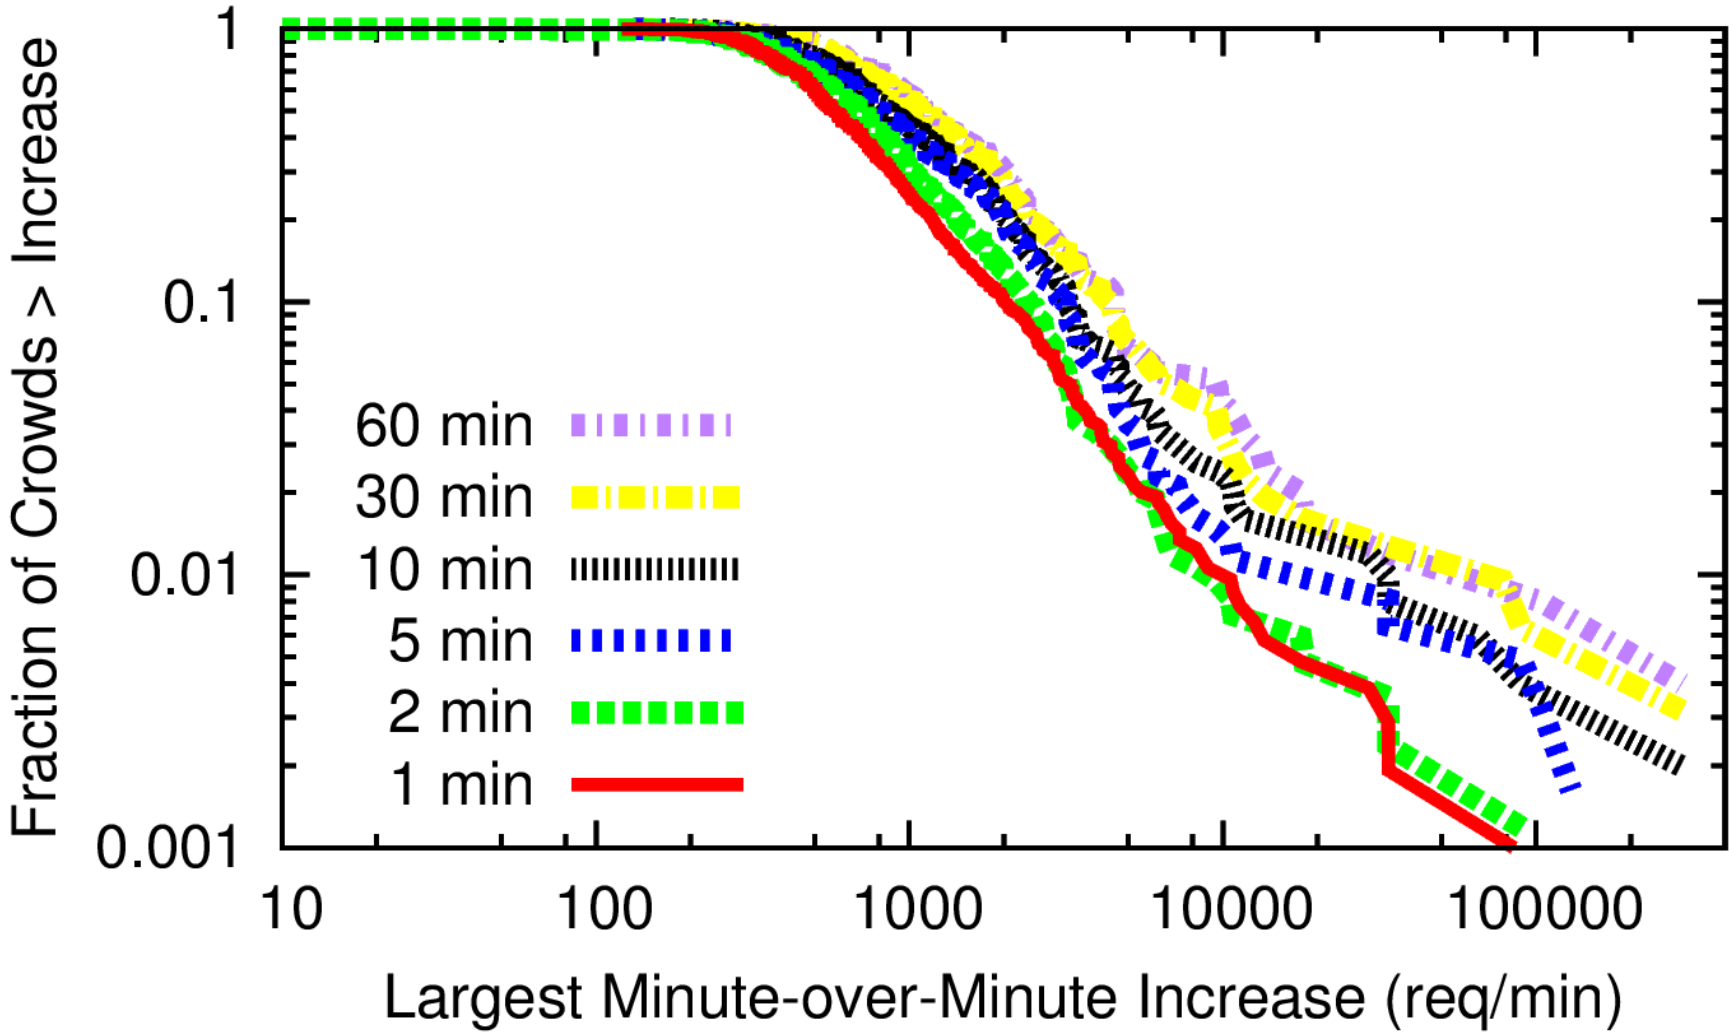

Origin Hits Saved by Cooperation

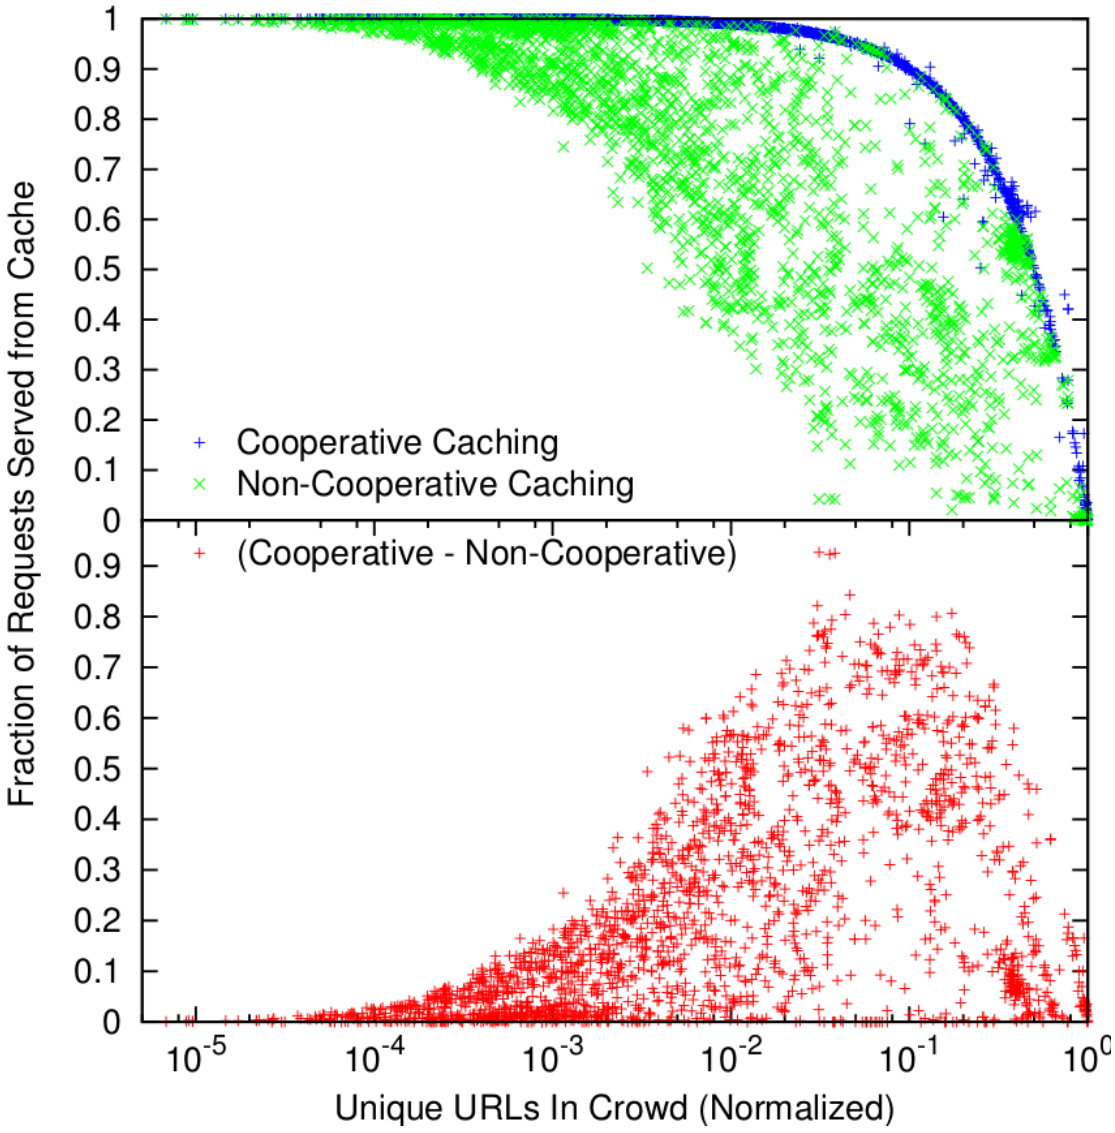

Bursty Redirection

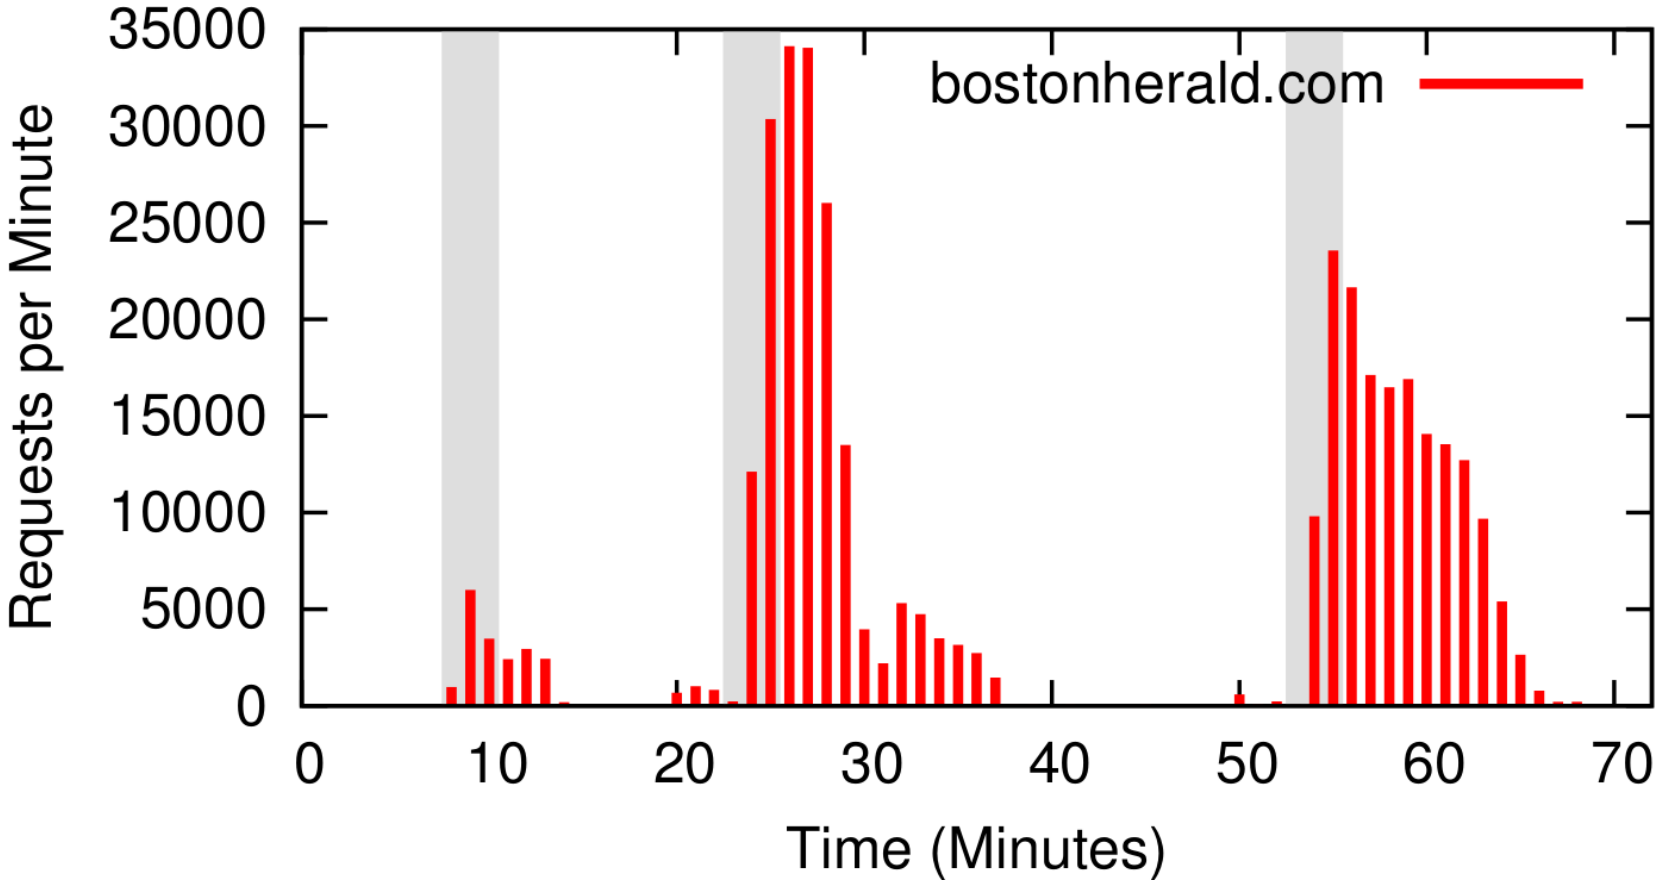

Clients Distributed Widely

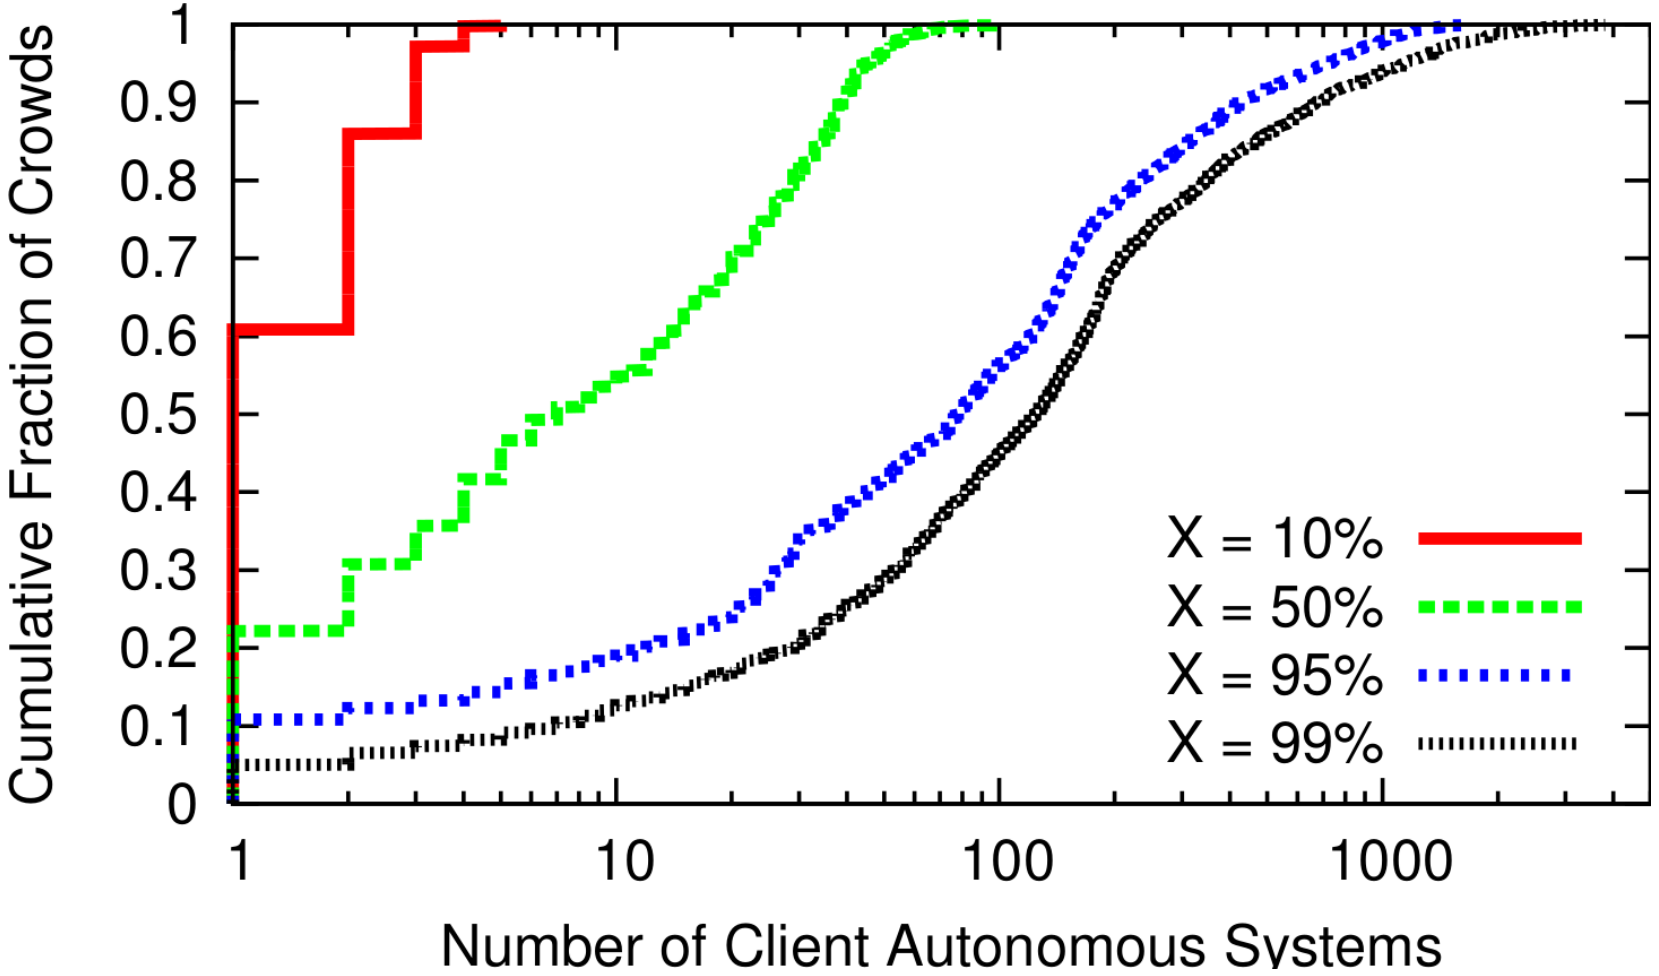

Detecting Crowds

1. Rapid surge in request rate

   $r_{i+1} > 2r_i$ for several i

2. High rate of traffic relative to inferred capacity

   $r_{max} > r_{avg} * 20$

Crowd Mitigation/Insurance

Content Mostly Static

Caching CDNs

Content Mostly Dynamic

Scalable Storage
and Computation