

Availability and Locality Measurements of P2P File Systems

Jacky Chu Kevin Labonte
Brian Neil Levine

Department of Computer Science
University of Massachusetts – Amherst



Supported in part by National Science Foundation awards ANI-033055 and EIA-0080199, and in part by Sprint Advanced Technology Labs.

Motivation

- Not much knowledge about existing P2P network dynamics.
 - Characterization requires statistical models and measurements.
 - Need to model network characteristics along with properties of the network content.
- Results predict caching would be beneficial
 - Significant transfer locality.
 - Short availability to serve other peers.
- We expect results useful for queuing models.

Previous Work

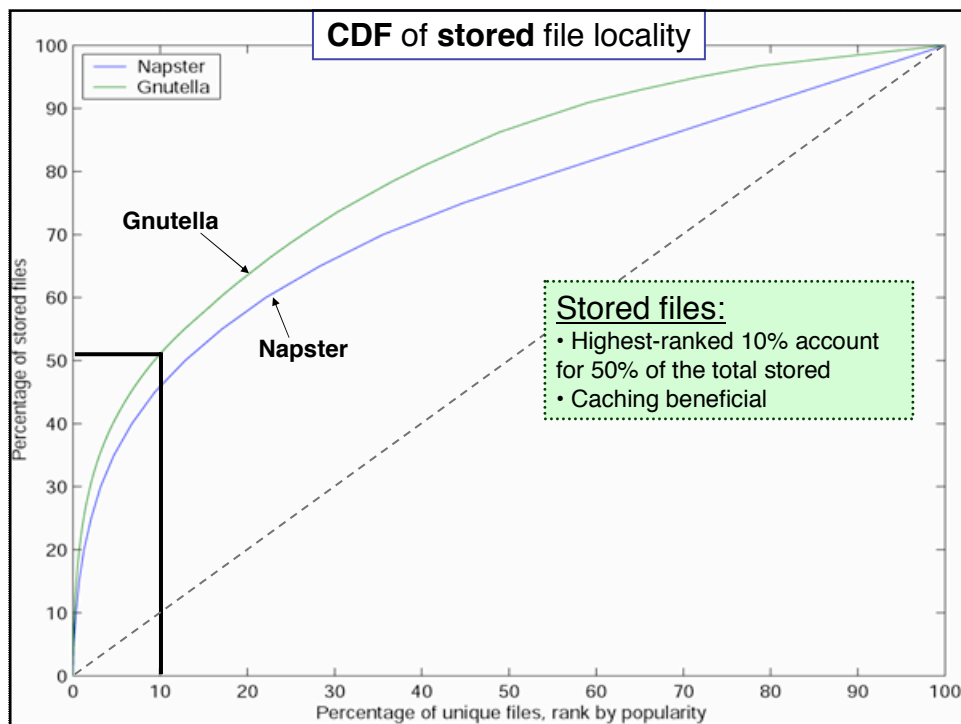
- *Ripeanu et al.* – Gnutella topology does not match underlying network topology.
- *Markatos* – A simple, query caching scheme can reduce query traffic by a factor of two.
- *Saroiu et al.* – Gnutella bandwidth, latency, and node availability over a 60-hour period.
- *Adar and Huberman* – A free-rider study, using Gnutella's QueryHit messages to infer peer downloads.

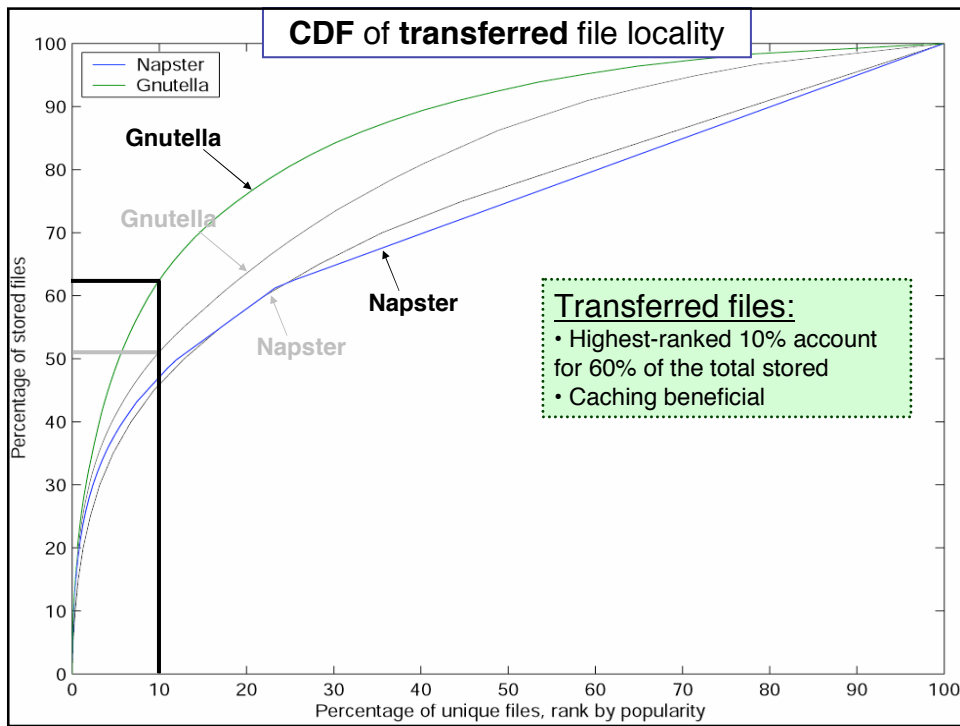
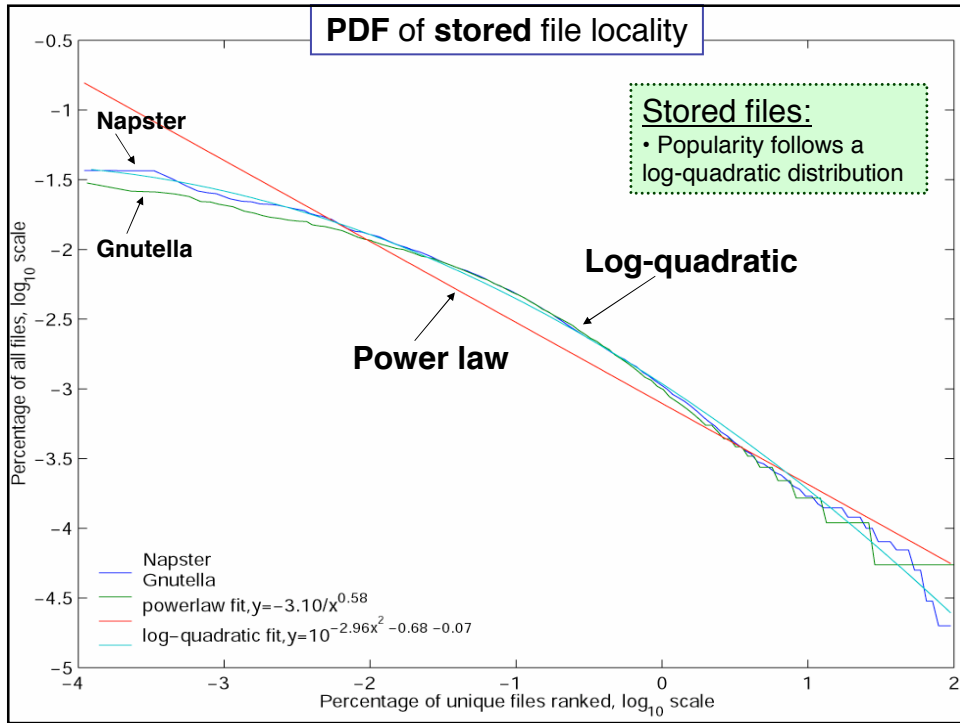
Our Study

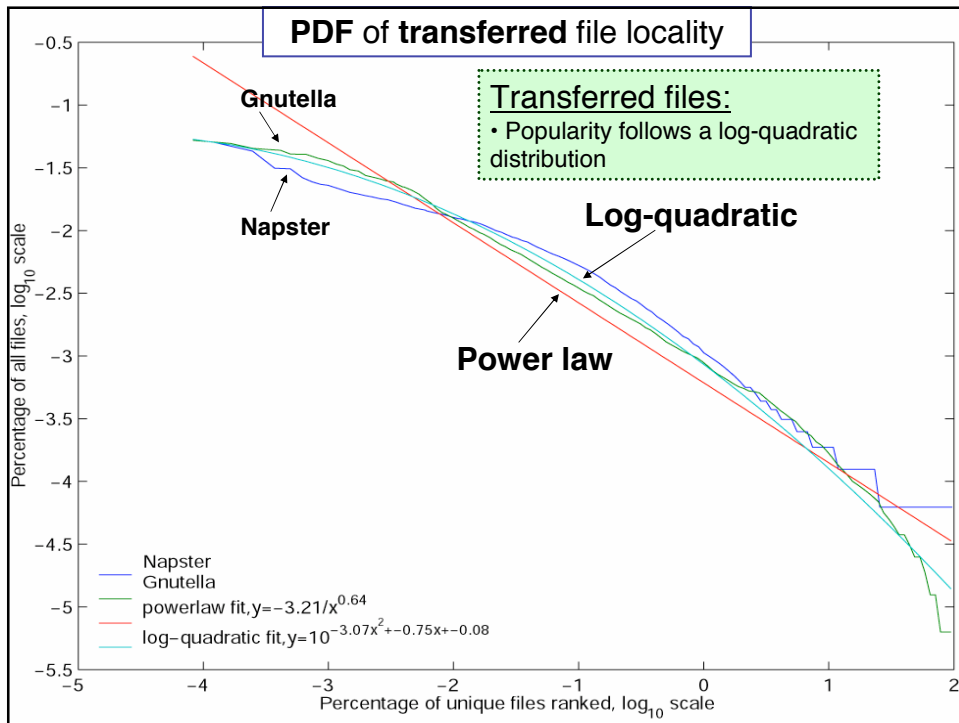
- Part I: File Locality
 - A more accurate estimation of file downloads between peers than previous work.
- Part II: Node Uptime
 - A longer-duration study on node availability, extending past results.

File Locality: Experiment Details

- **Monitoring peers:**
 - Napster: ~40 days of data, starting December 2000.
 - Gnutella: ~30 days of data, starting February 2002.
- **Method:**
 - Periodically monitored files shared by 20,000 peers.
 - Inferred file transfers by differences in file lists between subsequent checks.
- **Results:**
 - Locality of stored and transferred files closely follows a log-quadratic distribution.



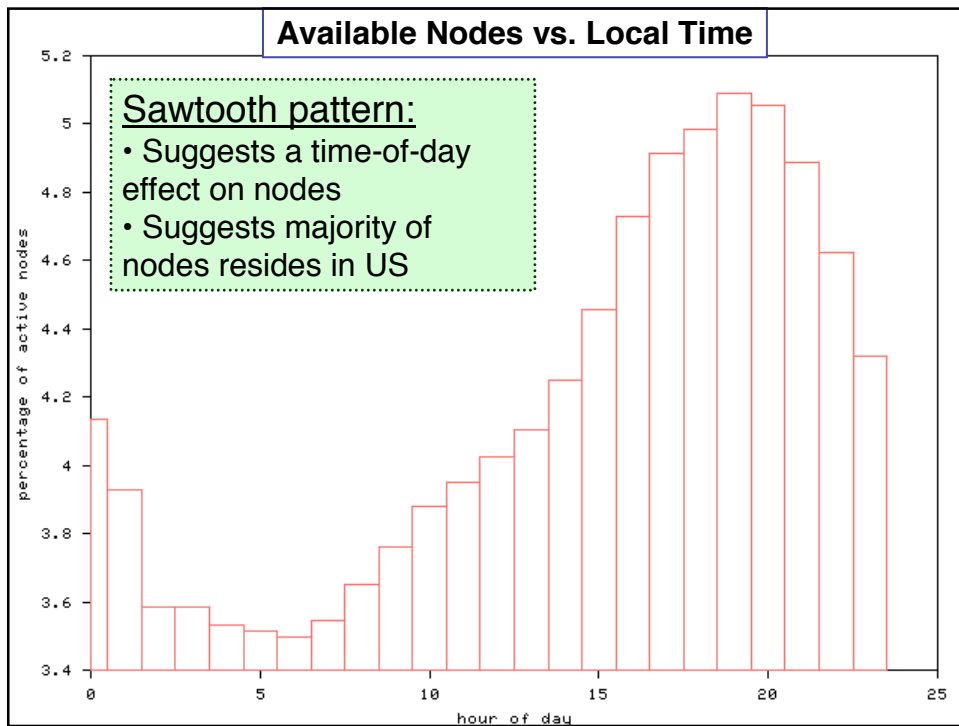
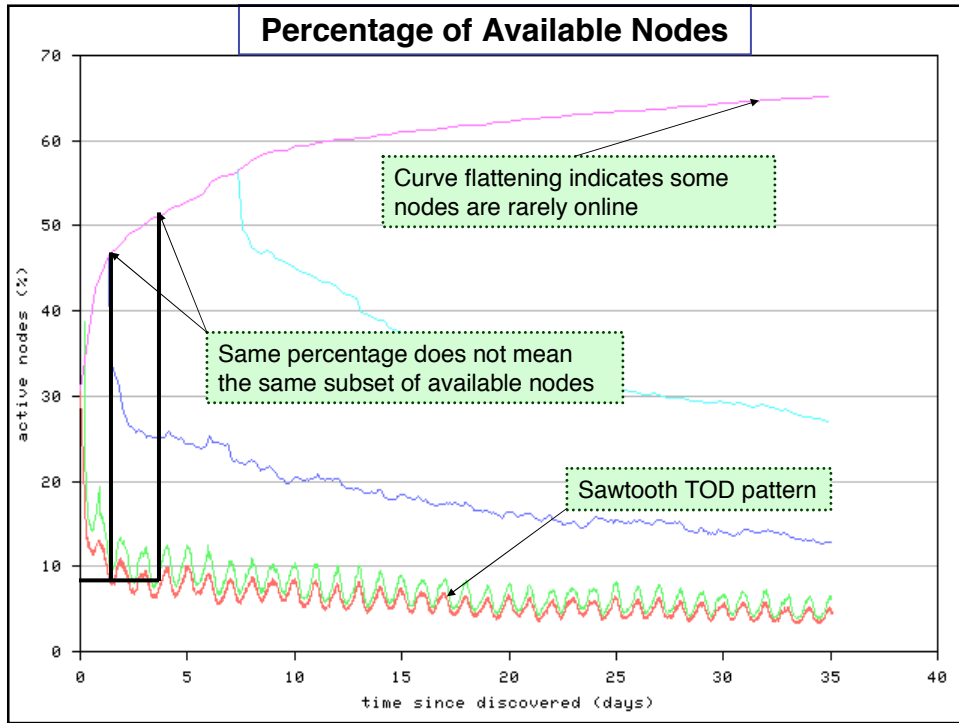


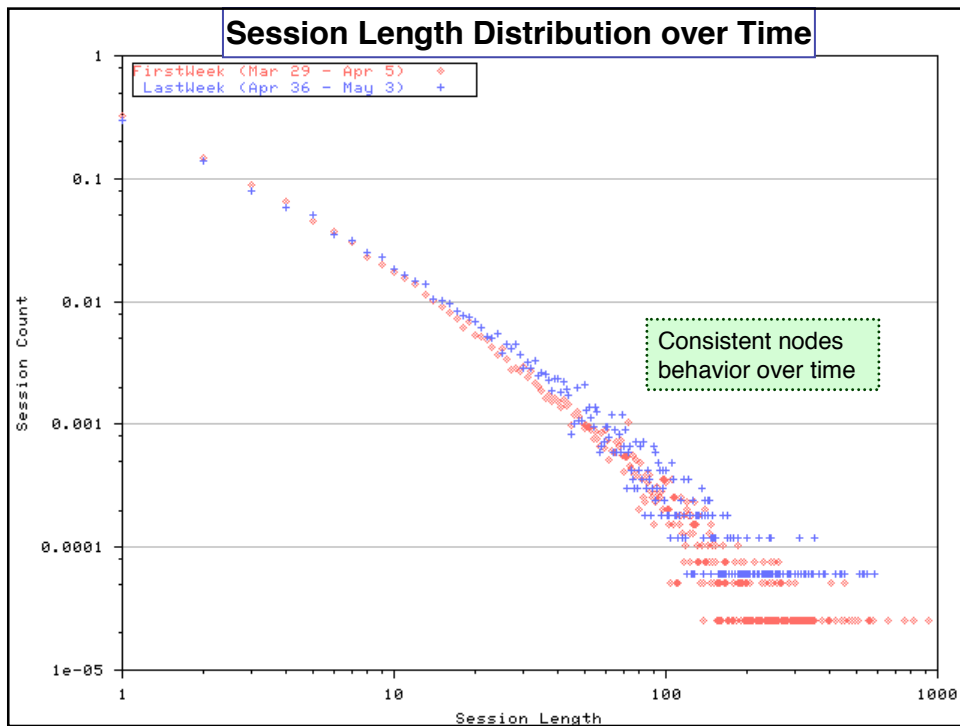
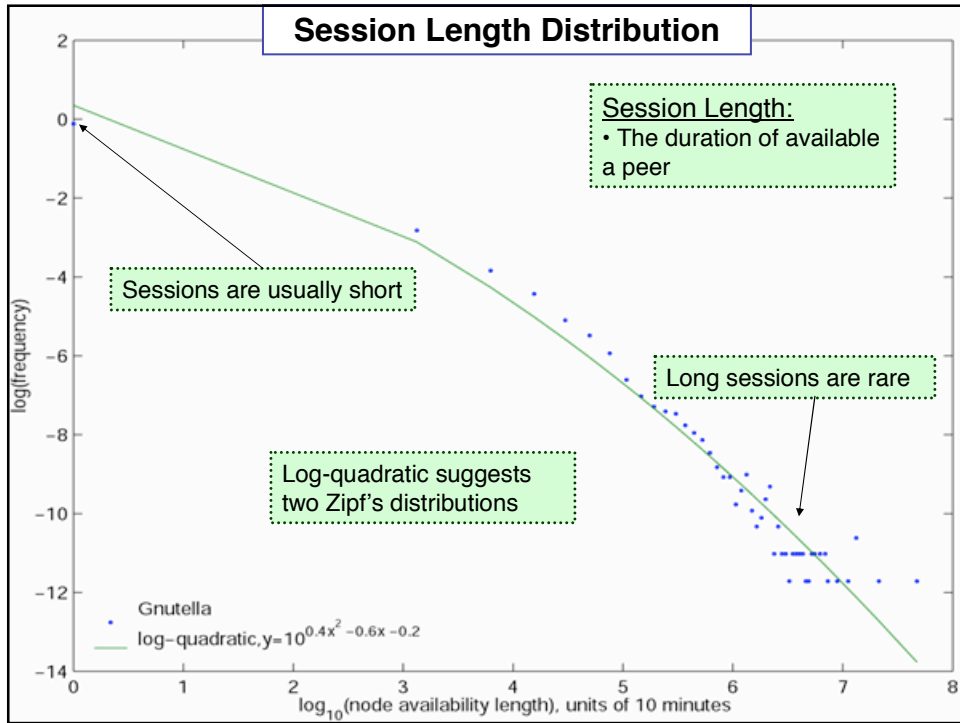


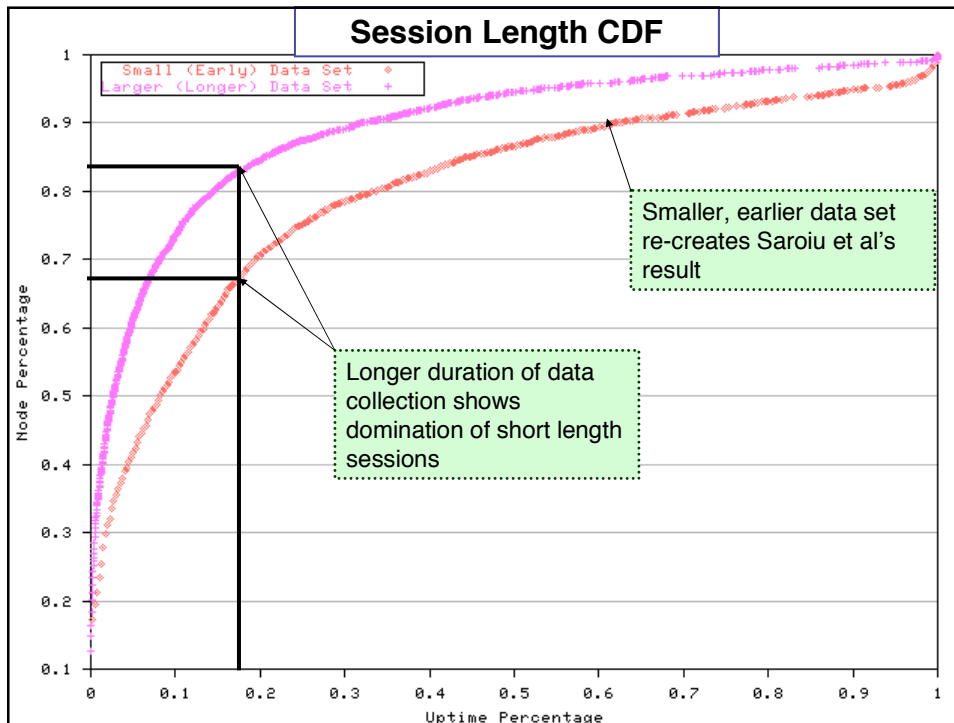
Node Availability – Experiment Setup

Defining Monitoring Set

- A custom Gnutella client extracted (IP,port) pairs from QueryHit messages routed through our node.
- Monitored 5,000 nodes.
- After each node is collected, it is scanned immediately with *nmap*.
- Every 10 minutes, all 5,000 nodes are scanned; 5-week period total.







Conclusions

- Locality of total stored file and also transferred file favors caching.
- Difficult to rely on a small set of “always” online normal user nodes to maintain system availability. Supernodes can be introduced to be the backbone of system. (e.g. Limewire’s router)
- Time-of-day effects suggest a geographical pattern of users.
- We are continuing to record traffic.

<http://signl.cs.umass.edu>