Replication Strategies in Unstructured Peer-to-Peer Networks

Edith Cohen and Scott Shenker

Presented By:

David Deschenes

April 10, 2003

Problem

How can file replication be used to improve the performance of search in unstructured peer-to-peer networks?

Unstructured P2P Networks

» Hosts form an overlay network
  - Each host maintains connections to “neighbor” hosts
  - Topology is unrelated to the location of data

» Hosts employ blind searching
  - Usually through limited broadcast
  - Query forwarding unrelated to query content
  - On average only as effective as random search

Peer-to-Peer Network Model

- Consists of n nodes, each with capacity □, and m files
- Total capacity $R = np$
- Let r_i denote the number of copies of the i^{th} file
- Let $p_i \equiv r_i / R$ be the fraction of total capacity allocated to the i^{th} file
- An allocation can be represented by the vector $\boldsymbol{p} = (p_1, p_2, p_3, \ldots, p_m)$
- A query rate distribution can be represented by the vector $\boldsymbol{q} = (q_1, q_2, q_3, \ldots, q_m)$
- An allocation strategy is a mapping from $\boldsymbol{p}$ to $\boldsymbol{q}$

Allocation Strategy Evaluation

- Measure the expected search size (ESS)
  - Number of nodes visited before the file is found
- Measure the maximum search size (L)
  - Number of nodes visited before the search is aborted

Uniform Allocation

- All files are replicated identically
- Minimizes L, therefore minimizing system resources spent on insoluable queries
- ESS = *m* / □

Proportional Allocation

- Improves queries for popular files at the expense of queries for rare files
- Presumably would improve overall performance
- Minimizes the maximum utilization rate, or the average rate of queries answered by a file copy
- ESS = $m / \square$

Square-Root Allocation

- For any two files the ratio of allocations is the square-root of the ratio of their corresponding query rates
- Minimizes ESS
- Approaches Uniform Allocation with respect to the minimization of L

Replication Approaches

- Copies are dynamically created and deleted
- After a successful search, the requesting host creates some number of copies, C, at randomly selected hosts
- $[C_i]$ is the average value of C for item i
- The approaches converge on a state where $[C_i] / [C_j]$ remains fixed over time