


Peer-to-peer overview

Antony Rowstron
Microsoft Research, Cambridge, UK.

Peer-to-peer systems

- Nodes have symmetric roles
 - No centralised components
- Two basic types
 - Unstructured
 - E.g. Gnutella
 - Structured
 - CAN, Chord, Kademia, Pastry, Tapestry, SkipGraphs, SkipNets, Viceroy

Structured P2P overlay networks



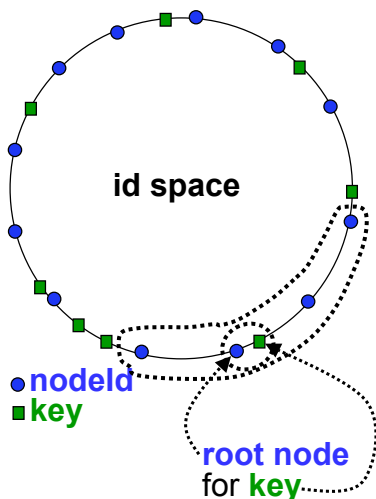
- structured **overlay network maps keys to nodes**
- **routes messages to keys**; can implement hash table

[CAN, Chord, Kademlia, Pastry, Skipnets, Tapestry, Viceroy]

NeXtworking'03 June 23-25, 2003, Chania, Crete, Greece
The First COST-IST(EU)-NSF(USA) Workshop on EXCHANGES & TRENDS IN NETWORKING

Rowstron 3

Mapping keys to nodes




- large **id space** (128-bit integers)
- **node ids** picked randomly from space
- **keys** picked randomly from space
- key is managed by its **root node**:
 - live node with id closest to the key
- key is replicated by its **replica roots**:
 - **r** nodes with ids closest to key

NeXtworking'03 June 23-25, 2003, Chania, Crete, Greece
The First COST-IST(EU)-NSF(USA) Workshop on EXCHANGES & TRENDS IN NETWORKING

Rowstron 4

Pastry: routing



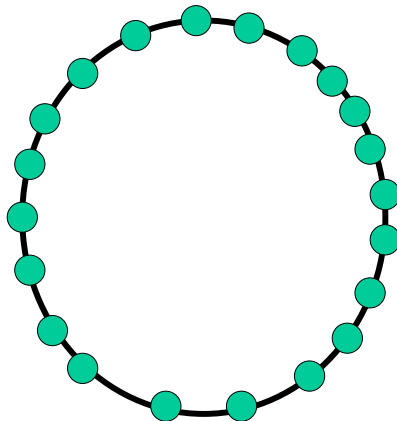
- prefix matching: each hop resolves an extra key digit
- neighbor set used to find root node in last hop
- properties: $\log_{16} N$ hops with low delay routes

NeXtworking'03 June 23-25,2003, Chania, Crete, Greece
The First COST-IST(EU)-NSF(USA) Workshop on EXCHANGES & TRENDS IN NETWORKING

Rowstron

5

Structured overlays (Pastry)-like




NeXtworking'03 June 23-25,2003, Chania, Crete, Greece
The First COST-IST(EU)-NSF(USA) Workshop on EXCHANGES & TRENDS IN NETWORKING


Rowstron

6


Unstructured overlays (Gnutella-like)



Unstructured overlays (Flooding)



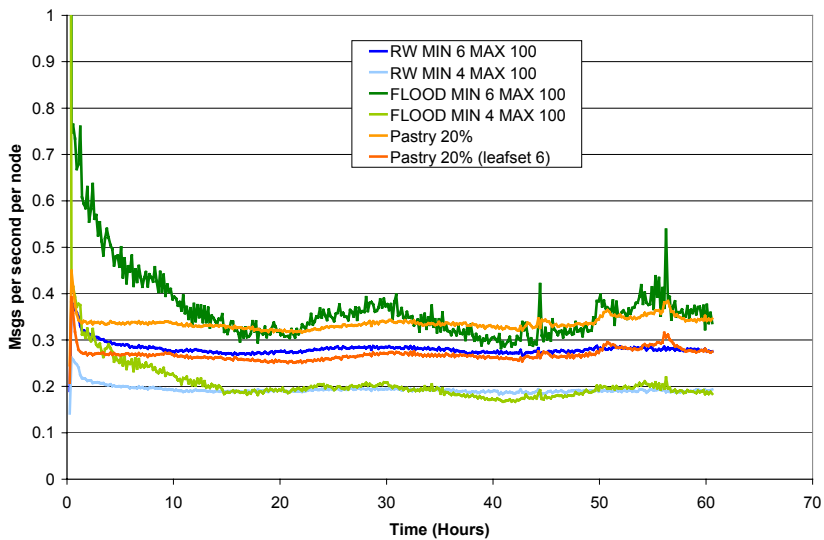
Unstructured overlays (Random Walk)



NeXtworking'03 June 23-25,2003, Chania, Crete, Greece
The First COST-IST(EU)-NSF(USA) Workshop on EXCHANGES & TRENDS IN NETWORKING

Rowstron 9

Maintenance costs



NeXtworking'03 June 23-25,2003, Chania, Crete, Greece
The First COST-IST(EU)-NSF(USA) Workshop on EXCHANGES & TRENDS IN NETWORKING

Rowstron 10

Security

- Gnutella and robustness
 - Information propagation
- Structured overlays
 - Initial work – adding constraints
 - Potentially expensive

Underlay meets overlay

- Should overlays exploit network locality?
 - What heuristics can be used?
- Exploiting network services in overlays

Controlling and managing resource utilisation

- Incentives and reputation systems
 - Some difficult problems – in particular auditing resource usage/contribution by nodes

Using p2p infrastructure

- Richer searching in DHT's
 - Peer-to-Peer Information Retrieval Using Self-Organizing Semantic Overlay Networks. *Chunqiang Tang (U Rochester), Zhichen Xu (HP Labs), Sandhya Dwarkadas (U Rochester)* Peer-to-Peer Information Retrieval Using Self-Organizing Semantic Overlay Networks. *Chunqiang Tang (U Rochester), Zhichen Xu (HP Labs), Sandhya Dwarkadas (U Rochester)*
- Interesting different applications
 - Content streaming
 - SplitStream

Content Streaming: SplitStream

Problems with a single multicast tree in a p2p environment:


- Interior nodes must invest upstream bandwidth many times the downstream bandwidth (unfair)
- Peers acting as interior nodes may fail (poor reliability)

SplitStream: basic approach


Stripe content across multiple *independent* multicast trees

- Fairness: interior node in one tree is leaf in other trees
- Reliability: failure of one tree causes drop in quality (at worst)
 - Erasure coding or multiple description video codecs
- Flexibility: can accommodate low bandwidth nodes

SplitStream: Multiple multicast trees



SplitStream: Multiple multicast trees



Other issues

- Simulators
 - Network topologies
 - Workloads
 - Scalability
- Testbeds etc.
- Complexity

Further information

- R. Mahajan, M. Castro and A. Rowstron, "Controlling the Cost of Reliability in Peer-to-peer Overlays", IPTPS'03, Berkeley, CA, February 2003.
- M. Castro, P. Druschel, A. Ganesh, A. Rowstron, and D. S. Wallach, "Security for structured peer-to-peer overlay networks". In Proceedings of the Fifth Symposium on Operating Systems Design and Implementation (OSDI'02), Boston, MA, December 2002.
- M. Castro, P. Druschel, Y. C. Hu and A. Rowstron, "Exploiting network proximity in peer-to-peer overlay networks", Technical report MSR-TR-2002-82, 2002.
- S. Iyer, A. Rowstron and P. Druschel, "SQUIRREL: A decentralized, peer-to-peer web cache", 12th ACM Symposium on Principles of Distributed Computing (PODC 2002), Monterey, California, USA, July 2002.
- "The Impact of DHT Routing Geometry on Resilience and Proximity". Krishna P. Gummadi (U Washington), Ramakrishna Gummadi (USC), Steven D. Gribble (U Washington), Sylvia Ratnasamy (Intel), Scott Shenker (ICSI), Ion Stoica (UC Berkeley), SIGCOMM 2003
- Internet Indirection Infrastructure. Ion Stoica, Daniel Adkins, Shelley Zhuang, Scott Shenker, Sonesh Surana (University of California, Berkeley), SIGCOMM 2003

- Workshop on Economics of Peer-to-Peer Systems
<http://www.sims.berkeley.edu/research/conferences/p2pecon/program.html>
- IPTPS'03 <http://iptps03.cs.berkeley.edu/>