Peer-to-Peer Networking

Acknowledgements: Most materials presented in the slides are based on the tutorial slides made by Dr. Raouf Boutaba, Dr. Keith W. Ross, and Dr. Dan Rubenstein.

<u>Client/Server Computing Model</u>



Client: Process wishing to access data, use resources or perform operations on a different computer
 Server: Process managing data and all other shared resources amongst servers and clients, allows clients access to resource and performs computation
 Interaction: invocation / result message pairs

The Peer-to-Peer Model

Applications based on peer processes

ONot Client-Server

Oprocesses that have largely identical functionality



<u>The Peer-to-Peer Mania</u>

Started in the middle of 2000

O When the Internet fallen into predictable patterns

O Computer field shocks:

Napster, SETI@home, Freenet, Gnutella, Jabber, ...

many early P2P projects have an overtly political mission

Emergence of sporadically connected Internet nodes (laptops, handhelds, cell phones, appliances, ...)

These developments return content, choice, and control to ordinary users. Tiny endpoints on the Internet, sometimes without even knowing each other, exchange information and form communities.

The Peer-to-Peer Mania (cont.)

• A new energy erupting in the computing field

- Yet, P2P is the oldest architecture in the world of communications
 - Telephones are peer-to-peer
 - Usenet implementation of UUCP
 - Routing in the Internet
 - Internet endpoints have historically been peers

P2P technologies return the Internet to its original vision, in which everyone creates as well as consumes

<u>Outline</u>

O General overview

- Definition
- Overlay networks
- 🔿 Goals
- O Classification of P2P systems

OP2P Search activities

- Applications
- O Search
- Security
- O Resource management

OCase Studies

- File-sharing systems
- O Distributed Hash Tables (DHT)
- O Trust and reputation management

<u>Outline</u>

OGeneral overview

- O Definition
- Overlay networks
- O Goals
- O Classification of P2P systems

OP2P research activities

OCase Studies

<u>Definitions</u>

- Everything except the client/server model
- Network of nodes with equivalent capabilities/responsibili ties (symmetrical)
- Nodes are both Servers and clients called "Servents"
- Direct exchange of information between hosts at the edge of the Internet



Definitions (cont.)

• A transient network that allows a group of computer users to connect with each other and collaborate by sharing resources (CPU, storage, content).

The connected peers construct a virtual overlay network on top of the underlying network infrastructure

- Examples of overlays:
 - BGP routers and their peering relationships
 - Content distribution networks (CDNs)
 - Application-level multicast
 - → And P2P apps !



OGeneral overview

- O Definition
- Overlay networks
- O Goals
- O Classification of P2P systems

OP2P research activities

OCase Studies

Overlay Networks

- An overlay network is a set of logical connections between end hosts
- Overlay networks can be unstructured or structured
- Proximity not necessarily taken into account
- Overlay maintenance is an issue



Overlay edge

Overlays: All in the application layer

ODesign flexibility

 Topology
 Protocol
 Messaging over TCP, UDP, ICMP

Underlying physical net is transparent to developer



<u>Outline</u>

OGeneral overview

- O Definition
- Overlay networks
- 🔿 Goals
- O Classification of P2P systems

OP2P research activities

OCase Studies

OFuture research directions

Goals

Cost reduction through cost sharing

- Client/Server: Server bears most of the cost
- → P2P: Cost spread over all the peers (+Napster, ++Gnutella,...)
- aggregation of otherwise unused resources (e.g., seti@home)

Improved scalability/reliability

resource discovery and search algorithms (eg. Chord, CAN, ...)

Interoperability

→ for the aggregation of diverse resources (storage, CPU, ...)

O Increased autonomy

independence from servers, hence providers (e.g., A way around censorship, licensing restrictions, etc.)

Goals (cont.)

- Anonymity/privacy
 - Difficult to ensure with a central server
 - Required by users who do not want a server/provider to know their involvement in the system
 - ➔ Freenetis a prime example
- Oynamism and Ad hoc communications
 - Resources (e.g., compute nodes) enter and leave the system continuously
 - Mechanisms are required to avoid polling (e.g., "buddy lists" in Instant messaging)
 - P2P systems typically do not rely on an established infrastructure
 - they build their own, e.g. logical overlay in CAN

<u>Outline</u>

OGeneral overview

- O Definition
- Overlay networks
- O Goals
- O Classification of P2P systems

OP2P research activities

OCase Studies

P2P Classification

- Degree of P2P decentralization
 - O Hybrid decentralized P2P
 - O Purely decentralized P2P
 - O Partially centralized P2P
- Degree of P2P structure
 - O Structured P2P
 - O Loosely structured P2P
 - O Unstructured P2P

Hybrid decentralized P2P

- Central server facilitates the interaction b/w peers.
- Central server performs the lookups and identifies the nodes of the network.
- 🔿 example: Napster
- (-) Single point of failure, scalability?, ...



Purely decentralized P2P

- O network nodes perform the same tasks (Servents)
- O no central coordination activity
- 🔿 examples: original Gnutella, Freenet
- O (-) data consistency?, Manageability?, Security?, Comm. overhead



Partially centralized P2P

Some of the nodes assume a more important role

O Supernodes act as local central indexes

🔿 examples: Kazaa, recent Gnutella



Unstructured P2P

• data is distributed randomly over the peers and broadcasting mechanisms are used for searching.

- O placement of data is unrelated to the overlay topology.
- 🔿 examples: Napster, Gnutella, KaZaa



Structured P2P

- Network topology is tightly controlled and files are placed at precisely specified locations.
- Provide a mapping between the file identifier and location
- Examples: Chord, CAN, PAST, Tapestry, Pastry, etc.



Loosely Structured P2P

- O Between structured and unstructured
- File locations are affected by routing hints, but they are not completely specified.



P2P classification summary

	Unstructured Networks	Loosely Structured Networks	Structured Networks
Hybrid Decentralized	Napster		
Pure Decentralized	Gnutella	Freenet	Chord, CAN, Tapestry
Partially Centralized	KaZaa, new- Gnutella		



OGeneral overview

P2P research activities

- O Applications
- O Search
- O Security
- Resource Management

OCase Studies

OFuture research directions



OGeneral overview

P2P research activities

- O Applications
- O Search
- O Security
- O Resource Management

OCase Studies

P2P Applications

- **O**File Sharing
- OCommunication
- Collaboration
- Computation
- ODatabases



P2P File Sharing

- File exchange: Killer application!
- O(+) Potentially unlimited file exchange areas
- O(+) High available safe storage: duplication and redundancy
- (+) Anonymity : preserve anonymity of authors and publishers
- O(+) Manageability
- O(-) Network bandwidth consumption
- O(-) Security
- O(-) Search capabilities

P2P File Sharing (cont.)

Examples of P2P file sharing applications:
ONapster

disruptive; proof of concept

OGnutella

→open source

O KaZaA

at some point, more KaZaAtraffic than Web traffic!
OeDonkey

→ popular in Europe

➔ Some implementations use Kademlia(DHT) for search

OBitTorrent:

≥53% of all P2P traffic in June 2004 was BitTorrenttraffic
Oand many others...

How do you explain this success?

P2P Communication

OInstant Messaging (IM)

- User A runs IM client on her PC
- Intermittently connects to Internet; gets new IP address for each connection
- Registers herself with "system"
- Learns from "system"that user B in her "buddy list"is active
- User A initiates direct TCP connection with User B: P2P
- User A and User B chat.
- Can also be voice, video and text.
- OAudio-Video Conferencing

Example: Voice-over-IP (Skype)

P2P Collaboration

• Application-level user collaboration

- **O**Shared Applications
 - Shared file editing (eg. Distributed Powerpoint)
 - Example: Groove
- Online games
 - Multi-players, distributed
 - Example: Descent (<u>www.planetdescent.com</u>)
- OTechnical challenges
 - Locating of peers
 - Fault tolerance
 - Real-time constraints

P2P Computation

O Achieves processing scalability by aggregating the resources of large number of individual PC.

- Application areas
 Financial applications
 Biotechnology
 Astronomy,...
- Related projects
 - ⇒seti@home
 - Avaki
 - Entropia
 - →Gridella

P2P Databases

- Fragments large database over physically distributed nodes
- Overcomes limitations of distributed DBMS
 - Static topology
 - Heavy administration work
- O Dissemination of data sources over the Internet
 - Each peer is a node with a database
 - Set of peers changes often (site availability, usage patterns)
- Examples:
 - AmbientDB (http://homepages.cwi.nl/~boncz/ambientdb.html)
 - > XPeer: self-organizing XML DB
- 🔿 No global schema



OGeneral overview

P2P research activities

- O Applications
- O Search
- O Security
- O Resource Management

OCase Studies

Search

- Challenges:
 - O Efficiency
 - Measure: resource (mostly bandwidth) consumed
 - O Autonomy:.
 - Level of control at individual node on data/index placement, connectivity & message routing
 - O Robustness
 - Stability in the presence of transient population of peers
 - O Scalability
 - Accommodate millions of nodes without degrading efficiency
 - O Expressiveness
 - of query expression: e.g. key lookup, keyword lookup, range queries etc.
 - O Completeness:
 - Guarantee on successful search and retrieval of all matches





- Blind search
- Informed search
 - OLocal Index
 - ➢Proactive
 - ➤Reactive
 - ○Distributed Index>DHT-based systems
Search taxonomy

Various Blind Search Methods:

OBFS/Flooding (Gnutella)

- forward incoming query to every neighbor.
- Flooding is restricted by TTL (time-to-live)

O Modified BFS

- Flood only a portion of the neighbors,
- Reduces search traffic at the cost of lower hit rate.

ORandom walk

- forward query message to randomly chosen neighbors
- Uses a number of walkers in parallel

OIterative Deepening

- uses consecutive BFS at increasing depth.
- Suitable when number of hits required is pre-specified and the possibility of matches at neighbors within a small radius.

Search Taxonomy

Local index: example proactive methods

Neighborhood Signature Search

- Each node maintains a set of bloom-filters (hash summary) of the keywords stored at peers within a small radius (r)
- Query is forwarded to a peer out of neighborhood radius in case no match is found in neighborhood



Search Taxonomy

Local index: example proactive methods OAssociative Search

- Nodes are arranged into interest groups.
- Query is flooded in one or more target interest group(s)



Associative Search

Search Taxonomy

Local Index: example reactive methods

Objective Distributed Resource Location Protocol (DRLP)

- Uses random walk when no information is available
- In case of a hit, the document location is stored at each node along the reverse path to the requester

OIntelligent BFS:

- Informed version of modified-BFS
- Query routing is guided by past routing decision
- Nodes store query-neighborID tuples
- In case of a hit, routing information at all nodes on the reverse path is updated.



Hash Table

Odata structure that maps "keys" to "values"

Interface
 put(key, value)
 get(key)

Distributed Hash Table (DHT)
 similar, but spread across the Internet
 challenge: locate content

What is a DHT? (cont.)

Single-node hash table:

Key = hash (data) put(key, value) get(key)->value

Distributed Hash Table (DHT):

Key = hash (data) Lookup (key) -> node-IP@ Route (node-IP@, PUT, key, value) Route (node-IP@, GET, key) -> value

🖵 Idea:

- Assign particular nodes to hold particular content (or reference to content)
- Every node supports a routing function (given a key, route messages to node holding key)

What is a DHT? (cont.)



DHT in action



DHT in action: put()



DHT in action: put()



DHT in action: put()



DHT in action: get()



Iterative vs. Recursive Routing



<u>Peers vs Infrastructure</u>

Peer Based DHT:

Application users provide nodes for DHT
 Examples: file sharing, etc

Infrastructure Based DHT
 Set of managed nodes provide DHT service
 Perhaps serve many applications

DHT Design Goals

An "overlay" network with:

- Decentralization and self-organization, i.e. no central authority, local routing decisions
 Flexibility in mapping keys to physical nodes and routing
 Robustness to joining/leaving
 Scalability, i.e. low communication overhead
 Efficiency, i.e. low latency
- A consistent "storage" mechanism with
 - No guarantees on persistence
 - Maintenance via soft state

Comparison to Other Facilities

Facility	Abstraction	Easy Use/Prg	Scalability	Load-Balance
DHT	high	high	high	yes
Centralized Lookup	medium	medium	low	no
P2P flooding queries	medium	high	low	no
Distributed FS	low	medium	medium	no

Facility	Fault-Tolerance	Self-Org	Admin
DHT	high	yes	low
Centralized Lookup	low	no	medium
P2P flooding queries	depends	yes	low
Distributed FS	medium	no	high



OGeneral overview

P2P research activities

- O Applications
- O Search
- O Security
- O Resource Management
- OCase Studies

<u>Security</u>

- Protection against:
 - Routing attacks
 - Redirect queries in wrong direction or to non-existing nodes
 - Misleading updates
 - Partition
 - O Storage/Retrieval attacks
 - Node responsible for holding data item does not store or deliver it as required
 - Inconsistent behavior
 - Node sometimes behaves and sometimes does not
 - Identity spoofing
 - > Node claims to have an identity that belongs to another node
 - Node delivers bogus content
 - DoS attacks
 - → Let legitimate (authenticated) users through
 - Rapid joins/leaves

In P2P systems, security issues are raised by the presence of malicious peers

O Trust assessment and management



Need for providing levels of:

- O Data availability
- Privacy
- Confidentiality
- Integrity
- Authenticity

Challenges:

- Secure storage
- Secure routing
- O Access control, authentication and identity management
- Anonymity
- Reputation management



Secure storage: cryptographic algorithms and protocols for securing the published and stored content

Self-certifying data (integrity)

OInformation Dispersal (availability)

OSHA (confidentiality and availability)



Secure routing: relies on 3 primitives

OSecure assignment of IDs to nodes

OSecure maintenance of routing tables

OSecure forwarding of messages



Access control, authentication and identity management: often ignored but essential to counter attacks like Sybil attacks, denial of service, etc.

OCentral authentication/identification authority

• Access keys and author-based access control list

Intellectual property and digital rights management

<u>Anonymity</u>

O The degree to which a system allows for anonymous transactions

O not want someone to identify author, publisher, reader, server, or document on the systems

Project	Types and techniques of Anonymity					
	Publisher	Reader	Server	Document		
Gnutella	Multicasting, covert paths	N/A	N/A	N/A		
Freenet	covert paths, identity spoofing	covert paths	Non-voluntary placement	encryption		
APFS	covert paths	covert paths	N/A	N/A		
FreeHaven	Covert paths (remailer)	covert paths	broadcast	Encryprion/split files into shares		
Publius	Covert paths (remailer)	N/A	Non-voluntary placement	Encryption/split key		
PAST	N/A	N/A	Non-voluntary placement	encryption		



OGeneral overview

P2P research activities

- O Applications
- O Search
- O Security
- O Resource Management

OCase Studies

<u>Resource Management</u>

Focus here is on p2p content distribution systems

Main resources to be managed:

Ocontent

OStorage capacity

OBandwidth

Resource Management (cont.)

Content management: deletion, update and versioning

- Often not supported for security, robustness to attacks, lack of synchronization between peers
- OUpdate and deletion provided to publishers (Publius)

Complex content history archival (OceanStore)

<u>Resource Management (cont.)</u>

Storage capacity management: content expiration, storage vs. contribution

Ocontracts with duration in FreeHaven; expired content is evicted

OUsed capacity storage in exchange for compensation (MojoNation)

OPeer contribution in exchange to the used disk space (PAST)

Resource Management (cont.)

Bandwidth management: routing performance, search facility

OStructured systems for an optimized network overhead

OIndirect files for an optimized network overhead in key-word based search

• Indirect files point to regular files



OGeneral overview

OP2P research activities

Case Studies

- File-sharing systems
- O Distributed Hash Tables (DHT)
- O Distributed Computing
- O Trust and reputation management



🔿 KazaA

- O Distributed Hash Tables (DHT)
- O Distributed Computing
- O Trust and reputation management



- Hybrid decentralized, unstructured
- Combination of client/server and P2P approaches
- A network of registered users running a client software, and a central directory server
- The server maintains 3 tables:
 - (File_Index, File_Metadata)
 (User_ID, User_Info)
 (User_ID, File_Index)



Download X.mp3

Napster (cont.)

History

- 5/99: Shawn Fanning (freshman, Northeasten U.) founds Napster Online music service
- 🔿 12/99: first lawsuit
- 3/00: 25% UWisc traffic Napster
- 2/01: US Circuit Court of Appeals: Napster knew users violating copyright laws
- 7/01: # simultaneous online users: Napster 160K, Gnutella: 40K, Morpheus (KaZaA): 300K





Napster (cont.)

judge orders Napster to pull plug in July '01

other file sharing apps take

napster

over!





Pure decentralized, unstructured

- Characteristic:
 - •Few nodes with high connectivity.
 - Most nodes with sparse connectivity.
- Goal: distributed and anonymous file sharing
- Each application instance (node) :
 - Ostores/serves files
 - Oroutes queries to its neighbors
 - Oresponds to request queries

<u>Gnutella (cont.)</u>

History:

- ○3/14/00: released by AOL, almost immediately withdrawn
- OBecame open source
- Many iterations to fix poor initial design (poor design turned many people off)

Issues

How much traffic does one query generate?
How many hosts can it support at once?
What is the latency associated with querying?
Is there a bottleneck?

<u>Gnutella (cont.)</u>


<u>Gnutella (cont.)</u>

Advantages:

- Robustness to random node failure
- Completeness (constrained by the TTL)

Disadvantages:

- OCommunication overhead
- ONetwork partition (controlled flooding)
- OSecurity

Free riding problem

Tradeoff:

- OLow TTL => low communication overhead
- OHigh TTL => high search horizon



Purely decentralized, loosely structured

- Goal: provide anonymous method for storing and retrieving data (files)
- **2** operations:
 - O Insertion
 - O Search
- Indexes (keys) are used for:
 - file clustering
 - Assisting in routing
 - Search optimization

Each node maintains local data-store (files with similar keys) and routing table

Freenet (cont.)

- Bootstrapping
 - No explicit solution
- File insertion
 - O User assigns a hash key to file
 - Sends an insert message to the user's own node
 - Node checks its data store for collision
 - Node looks up the closest key and forwards the message to another node
 - This is done recursively until TTL expires or collision detected.
 - If no collision, user sends data down the path established by the insert message.
 - Each node along the path stores it and creates a routing table entry



Search



Connection
 Data request
 Data reply
 Request failed
 Chain

Freenet (cont.)

Advantages:

- Anonymity
- ORobustness to random node failure
- Low communication overhead
- OMore scalable
- Self-organized

Disadvantages:
 Poor routing decision
 Spam
 Security



- Partially centralized, unstructured
- Every peer is either a supernode(SN) or an ordinary node (ON) assigned to a supernode
- Each supernode knows where the other supernodes are (Mesh structure)





- Bootstrap:
 - Connection to a known SN
 - O Upload METADATA for shared files
 - →file name
 - → file size
 - file descriptor
 - → ContentHash
- Search:
 - ON -> SN (-> SN)*
 - Keyword-based
 - ContentHash in result

Download:

- O HTTP
- ContentHash-based
- Failure -> automatic search on ContentHash



Advantages:
 Scalability
 Efficiency
 Exploits heterogeneity of peer
 Fault-tolerance

Disadvantages
 Pollution
 DoS attacks on SN



BitTorrent (cont.)

File is broken into pieces

○ Typically piece is 256 KBytes

O Upload pieces while downloading pieces

Piece selection

Select rarest piece

• Except at beginning, select random pieces

Tit-for-tat

- O Bit-torrent uploads to at most four peers
- Among the uploaders, upload to the four that are downloading to you at the highest rates
- A little randomness too, for probing



• General overview

P2P research activities

- Case Studies
 - File-sharing systems
 - O Distributed Hash Tables (DHT)
 - Consistent hashing
 - O Kademlia
 - O Analysis of DHT systems
 - O Distributed Computing
 - Trust and reputation management



Consistent Hashing

OCAN (Content Addressable Network)

○Kademlia

Consistent Hashing

- Consistent hashing [Karger97]
- Overlay network is a circle
- Each node has randomly chosen id
 - Example: hash on the IP address
 - Keys in same id space
- Node's successor in circle is node with next largest id
 Each node knows IP address of its successor
- Key is stored in closest successor



Consistent Hashing (cont.)

Principles:

Node departures	Node joins			
Each node must track s ≥ 2 successors	You are a new node, id k			
	Ask any node n to find the node n ' successor for id k			
If your successor leaves, take next one	Get successor list from n'			
Ask your new successor for	Tell your predecessors to update their successor lists			
list of its successors; update your <mark>s</mark> successors	Thus, each node must track its predecessor			

neighbors = s+1: O(1)

O Routing table: (neighbor_id, neighbor_ip@)

Average # of messages to find key:

Can we do better?



Routing geometry: Hypercube

- A hash value = a point in a D-dimensional Cartesian space
- Each node responsible of a D-dimensional cube

Neighbors are nodes that "touch" in more than a point

- Example: D=2
 - 2,3,4,5 are 1's neighbors
 - 6 is a neighbor of 2

○ # neighbors: 0(D)

				2		
			З	1		5
				4		
					6	7

CAN (cont.)

- Routing:
 - Recursively, from (n1, ..., nD) to (m1, ..., mD), choose the closest neighbor to (m1, ..., mD)
 - expected # overlay hops: O(DN^{1/D})/4

🗖 Node Join:

- find some node in the CAN (via bootstrap process) (9)
- choose a point in the space uniformly at random (X)
- using CAN, inform the node that currently covers the space (8)
- O that node (8)splits its space in half
 - 1st split along 1st dimension, if last split along dimension i< D, next split along i+1st dimension
- keeps half the space and gives other half to joining node



CAN (cont.)

- Node Leave:
 - Ieaf (3) removed
 - Find a leaf node that is either
 - a sibling
 - descendant of a sibling where its sibling is also a leaf node (5)
 - \bigcirc (5) takes over (3)'s region (moves to (3)'s position on the tree)
 - (5)'s sibling (2)takes over (5)'s previous region
 - O The cube structure remains intact









Node / Peer

- Nodes are treated as leaves in a binary tree
- Node's position in the tree is determined by the shortest unique prefix of its ID
- A Node is responsible for all "closest" IDs, i.e. IDs having same prefix as itself
- Distance between ID x and y is measured as $d(x,y) = x \oplus y$
 - e.g. d(010101_b, 110001_b) = 100100_b⇔d(21₁₀, 49₁₀) = 36₁₀
 - Nodes/IDs tree (i.e. with longest common prefix) are closer

Kademlia: Index distribution



- For any node (say the red node with prefix 0011) the binary tree is divided into a series of maximal subtrees that do not contain the node.
- A node must know at least one node in each of these subtrees.



• Consider a query for ID 1110... initiated by node 00111...

DHT Challenges

- Consistency
 - Soft-state publication
 - Tradeoff between consistency and communication overhead

Performance

- Load-balancing
- Query Response time (optimize # of hops)
- => Replication?

Reduce latency in routing

- Neighbors in a ring and far away in the internet
- => Location-awareness

Geometry

- Ring topology adds flexibility and helps in resilience
- Bootstrapping
 - Relies on known node

DHT Applications

OGlobal file systems [OceanStore, CFS, PAST, Pastiche, UsenetDHT]

- Onaming services [Chord-DNS, Twine, SFR]
- ODB query processing [PIER, Wisc]
- Internet-scale data structures [PHT, Cone, SkipGraphs]
- Communication services [i3, MCAN, Bayeux]
- File sharing [OverNet, eMule, eDonkey, etc.]
- OEvent notification [Scribe, Herald]

<u>Outline</u>

OGeneral overview OP2P research activities

OCase Studies

- O File-sharing systems
- O Distributed Hash Tables (DHT)
- O Distributed Computing
- O Trust and reputation management
- Measurements of existing P2P systems
- OFuture research directions



- Goal : to discover alien civilizations
- Analyzes the radio emission from the space 2 collected by the radio telescopes using processing power of millions of unused internet PCs



- Two major components : a database server and clients
- Supported platforms : Windows, Linux, Solaris, and HP-UX

<u>Outline</u>

OGeneral overview

OP2P research activities

Case Studies

O File-sharing systems

- O Distributed Hash Tables (DHT)
- O Distributed Computing
- O Trust and reputation management
- Measurements of existing P2P systems
- OFuture research directions

File Sharing in a P2P system



Need for a Reputation Management scheme

File Sharing in a Reputation-Based P2P system



Reputation-based Systems



<u>RM in partially-decentralized p2p</u> systems

KaZaA Reputation System

- Integrity Rating of Files :
 - Excellent, average, poor, delete
- Participation Level for rating peers:
 - (Uploads in MB/Downloads in MB)* 100
- O Drawbacks of KaZaA reputation system:
 - No distinction between good and malicious peers
 - Designed to reward who are practicing good behavior
 - Does not punish malicious peers

The Inauthentic Detector Algorithm (IDA)
The Malicious Detector Algorithm (MDA)

RM Design: Completely Decentralized



RM Design: Partially Decentralized

Proposed Approach: Use Supernodes for RM



RM Overhead



Effect of liar peers

No liar peers

- Good peers get high reputation values
- Malicious peers get low reputation values

With liar peers

- Good peers can have low reputation values
- Malicious peers can have high reputation values
- → Need for a mechanism to detect liar peers

IDA and MDA







Reputation systems have the following challenges:

- O Eliciting feedbacks:
 - Peers may not send feedbacks
 - It is difficult to ensure an honest feedback

O Distributing the feedback history is challenging:

- Peers may change their pseudonyms easily to erase prior feedback history.
- The lack of portability from system to system.
- The use of a more universal framework will be beneficial.

• Aggregating feedbacks is a difficult task



OGeneral overview

OP2P research activities

OCase Studies

• Future research directions
Future Research Directions

- P2P research is an exciting area with many open problems and opportunities, including the design of:
 - New distributed object placement and query routing
 - O New hash table data structures and algorithms
 - O Efficient security and privacy
 - Semantic grouping of information in P2P networks
 - Incentive mechanisms and reputation systems
 - Convergence of Grid and P2P systems
 - O Providing transactional and atomic guarantees on P2P

Future Research Directions

- P2P and Management:
 - Content management is key
 - O Autonomic behavior (e.g., self-organizing communities of interest)
 - Trust and reputation management is a challenge
 - Further exploiting P2P architectures for management apps