Peer-to-Peer Networking: the Operator's Perspective (*or rather an academic exercise*)

Antonio Liotta

University of Essex
Dept. Electronic Systems Engineering
aliotta@essex.ac.uk

Outline

- P2P overview
- What is P2P particularly good at?
- Wearing the shoes of the operator: myths, threats and opportunities
- Recommendations to the operator – what we should be looking at?

Peer-to-Peer in a nutshell

Client-Server computing

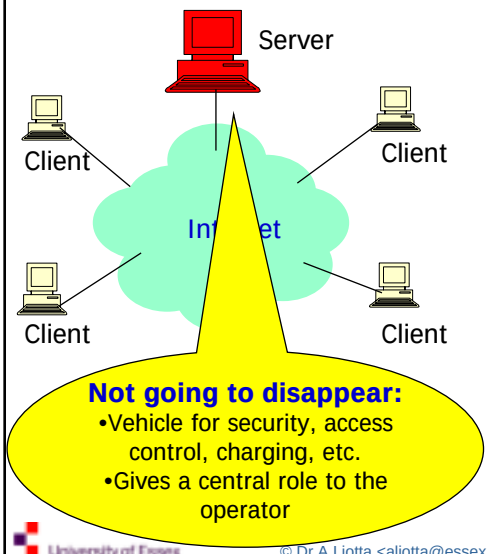

Client-server

- Server is data source
- Multiple clients per server
- Well-known, widely used (http, DNS, ftp, web services, etc)

Limitations

- Fault tolerance (single point of failure)
- Scalability hard to achieve
- Requires central administration
- Does not make the most of "Client" resources

The three flavours of P2P

- P2P communication
- P2P networks
- P2P computing / applications

P2P communication

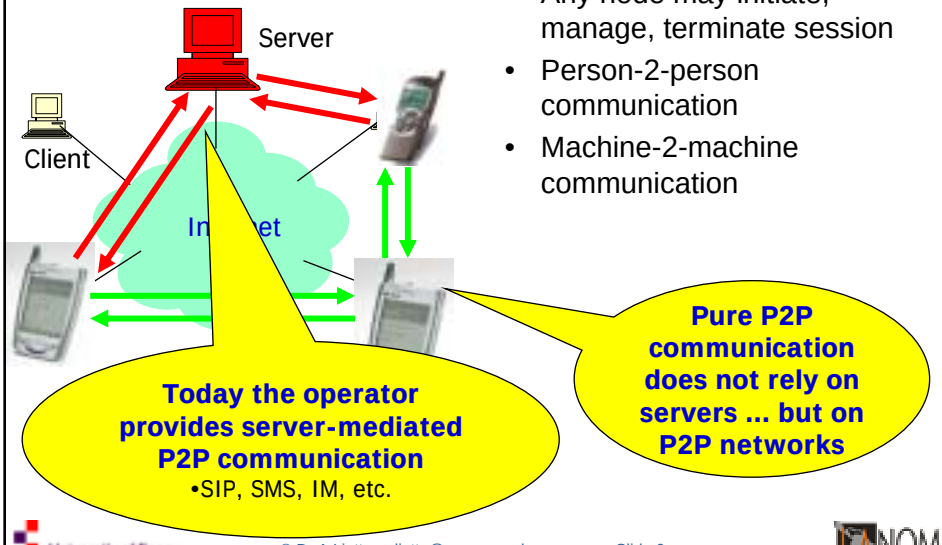

P2P communication

- Any node may initiate, manage, terminate session
- Person-2-person communication
- Machine-2-machine communication

P2P networks

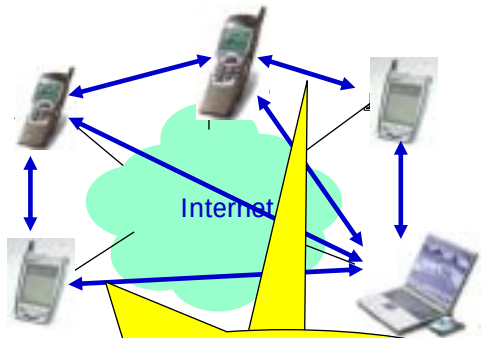

In Essence P2P consists of:
- Structured aggregation of resources
- Efficient discovery of resources
- Direct interaction among nodes (non-server mediated)

P2P networks

- Nodes are both clients and servers
- Nodes form a content/service network (application-level overlays)
- P2P networks are dynamic
- Autonomic self-managed (no centralised authority in theory!)

Limitations of pure P2P networks

- Hard to secure, charge etc
- Best-effort service
- Not mobile-friendly
- Not interoperable

P2P computing – enabler of innovative applications

- Distributed file sharing (where it all started!)
- New forms of content distribution and delivery
- Instant messaging including pervasive devices communication
- Collaborative work/development
- Online communities & gaming / thematic web networks
- Distributed search and indexing – ‘deep’ searches of internet content
- Sharing CPU & storage to better utilise capital investment

The impact of P2P: technical viewpoint

- No centralised repository means:
  - Robustness: no single point of failure
  - Scalability: bypassing processing and network bottlenecks
- Structured aggregation/organisation of resources/data/services means:
  - Dynamic distribution (publish)
  - Dynamic distributed search (discovery)
- Data replication means:
  - Resilience to volatility of resources (intermittent presence, fault-tolerance, etc.)
  - Reduced (distributed) search/retrieval times
  - Autonomic disaster recovery
- Distributed resource sharing/pooling means:
  - Ubiquitous access to virtually unlimited CPU cycles and storage from any terminal (including lightweight terminals)

The impact of P2P: some facts

- Currently 60-80% of Internet traffic is P2P
- BitTorrent traffic accounts for over 1/3 of IP traffic
- Estimated 40 millions BitTorrent users by 2006
- Strong trend towards distributed, multiparty communications (skype, VoIP, etc)
- Over 50M skype registered users; 3M concurrent users; 10Billion minutes served in 1 year
- New P2P applications represent a widely acknowledge opportunity to add value to network services

Primary P2P research question

How can we efficiently and accurately publish/discover resources and services in a P2P network?

Solution 1: Introduce some centralisation

Solution 2: Introduce some structure

"centralisation" and "structure" define two dimensions for classifying P2P networks

Examples of P2P with centralisation

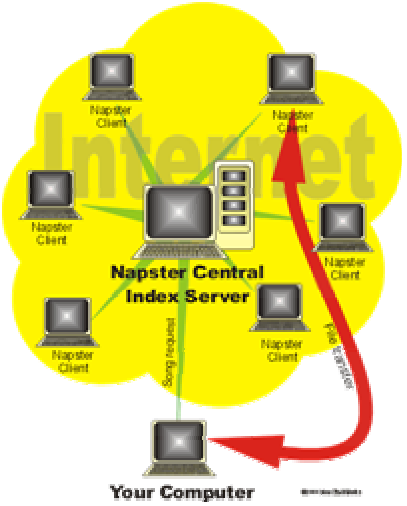

Figure 1. Typical request sequence. The request moves through the network from node to node, backing out of a dead-end (step 3) and a loop (step 7) before locating the desired file.

Examples of structured P2P: Chord

- Emphasis on efficient location / retrieval of data
- Each node stores information about only a small number of other nodes
- Each node knows more about nearby nodes than far-away nodes
- Example: a network with 10,000 peers has a 6-hop diameter

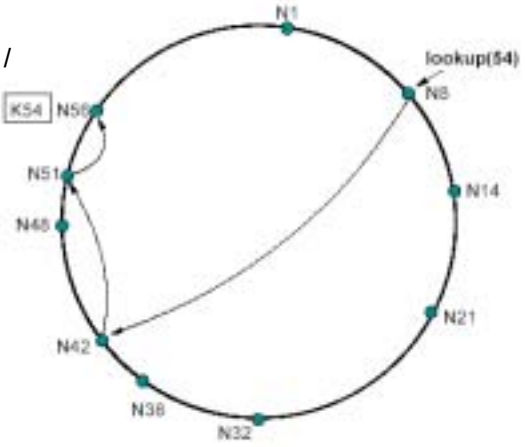

What is P2P particularly good at?

- Information
- Bandwidth
- Computing resources

P2P for ‘overlay network’ management

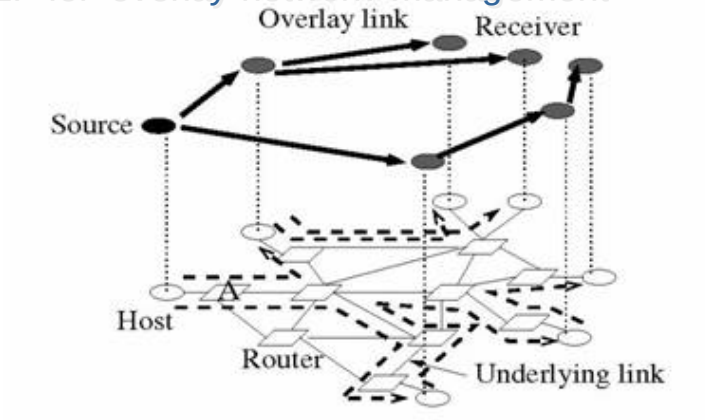

- P2P networks are very dynamic with very-high peer turnover rate (in systems such as Gnutella half of the peers are replaced within an hour)
- Maintaining connectivity is non-trivial
- P2P can manage large-scale dynamics thanks to distributed algorithms

P2P and the ‘small-worlds’ theory

- The ‘small-worlds’ theory says that a hidden pattern is the key to how networks interact an exchange information
- “In any well-designed network there are only 6 degrees of separations between any two nodes”
- Can P2P overlays bring the Internet down to a 6-hop network?
  - Lots of work on how to build low-diameter P2P networks
  - Recent findings on P2P networks with constant node degree and logarithmic diameter

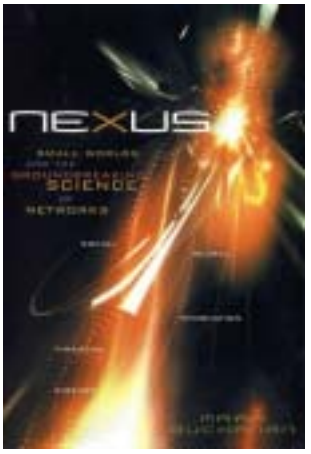

P2P for presence management

- PRESENCE is a dynamic profile of the user, visible to others and used to control services
  - Personal status or phone profile (available, busy, in meeting, on holiday)
  - Terminal status (out of coverage, engaged on voice call, switched off)
  - Location (roaming abroad, in this locality)
  - Terminal capabilities (supports SMS, MMS, chat, instant messaging, video)
  - Terminal selection ('route incoming voice calls now to this terminal')
  - Other information (personal logo, mood, etc)
- Presence in IMS is handled via a Presence Server → not scalable
- In P2P presence information is published/discovered using distributed mechanisms

P2P for group management

- Group: a collection of peers sharing a common interest (a peer may join several peer groups)
- Groups allow to create:
  - A secure environment (with security policies)
  - A scoping environment (local domain of specialisation – e.g. for document or CPU sharing)
  - A monitoring environment (for peers to monitor other peers for any special purpose – e.g. accounting)

- P2P systems provide the natural means for supporting the 'Group' concept
- In JXTA groups provide a set of core services:
  - Discovery service
  - Membership service
  - Access service
  - Pipe service
  - Resolver service
  - Monitoring service

P2P: extra reliability via data redundancy

- P2P provides means for transparently handling data replication
- A peer in 'client' mode acts also a 'server', publishing information as soon as it becomes available
- The more a file is hit the more replicas will be created

- Each node knows IP addresses of next *r* nodes
- Each key is replicated at next *r* nodes

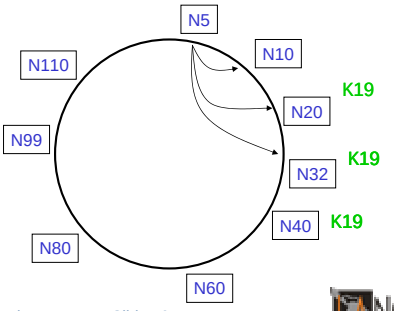

Wearing the shoes of the operator: myths, threats and opportunities

Displacing myths

"All P2P services can be done via CS"

- True in principle but ...
- P2P is the natural paradigm for distribution, sharing, collaboration
- P2P scales better (it makes the most of edge resources)
- P2P catalyses new value-added services

"P2P replaces CS"

- Not true!
- Major P2P issues are tackled via various flavours of centralisation or hierarchy
- P2P is complementary to CS

"P2P is a disruptive technology (enables operator bypass)"

- Partially true but ...
- Same threat posed by any IP-based service (C-S, mobile code, Web services, etc.)
- Not a good reason for ignoring the threat
- The operator is in a unique position to turn this threat in opportunity (value-added P2P services)
- The operator is the natural provider of P2P QoS, control, security, interoperation

Operator Bypass

- Conventional telco networks rely on:
  - Well-known revenue sharing models
  - Interconnection principles
  - Charging & billing mechanism
- Operator bypass (or de-intermediation) refers to the threats posed to the operator's revenue streams by the evolution of terminals and services towards IP
- Fundamental threat:
  - services obtained by 3rd-parties bypassing the operator
  - The operator is reduced to a bit pipe provider
  - Off-net services not controllable by the operator

Technical threats of operator bypass

- Standardisation
  - IMS + regulatory constraints = facilitate operator bypass
- Network
  - The user can create different bearers for different services
- Terminal
  - 3rd-party clients create a direct relationship between user and service provider
  - Midlets allow applications to run transparently from operator (all they need is http/sip)
- Security
  - Faulty or malicious clients may lead to fraud, security breaches, service disruption, etc
  - The user may misperceive operator's responsibility

Way around operator bypass

- Playing with tariffs (disadvantaging off-net traffic)
- Fiddling with network QoS (impairing off-net traffic)
- Tampering with terminal (avoid open architecture)
- Striking deals with other operators and 3rd-party providers

- Competing: exploiting their advantageous position, the operator may offer better, more timely value-added services

Turning ‘threats’ into ‘opportunities’

- The operator has unique advantages
  - Relevant know-how
  - Standardised service-centric infrastructure (OSA, IMS) speeds up deployment of new services and reduces management costs
  - Ability to expose/control advanced capabilities (billing, QoS, etc)
  - Large customer base (trusted relationship, security, etc.)
  - Partnerships with 3rd-party service providers, terminal manufacturers, and operators’ federation
  - Marketing in place
- Hence the operator can offer better and more timely P2P services than anybody else
- Operator-mediated P2P offers:
  - a single solution to managed P2P services
  - the additional benefit of security and QoS
- P2P facilitates resource sharing and collaboration (new value-added services)

P2P Opportunities: added value

- Dependability
  - Scalability
  - Fault-tolerance
  - Security
  - Anonymity
- QoS
- Free riding (incentive schemes)
- Interoperability
- ‘Deep’ search
- Mobility
- ...

Why P2P should go ‘mobile’

- An integrated MP2P core increases scalability of:
  - P2P communication
  - P2P services (distributed publish/search, size of peer groups)
- MP2P connects the mobile user to other P2P communities
- MP2P will spark unprecedented traffic volumes

“MP2P is a unique selling point for a mobile operator”

The operator’s working agenda

Integration of P2P in IMS

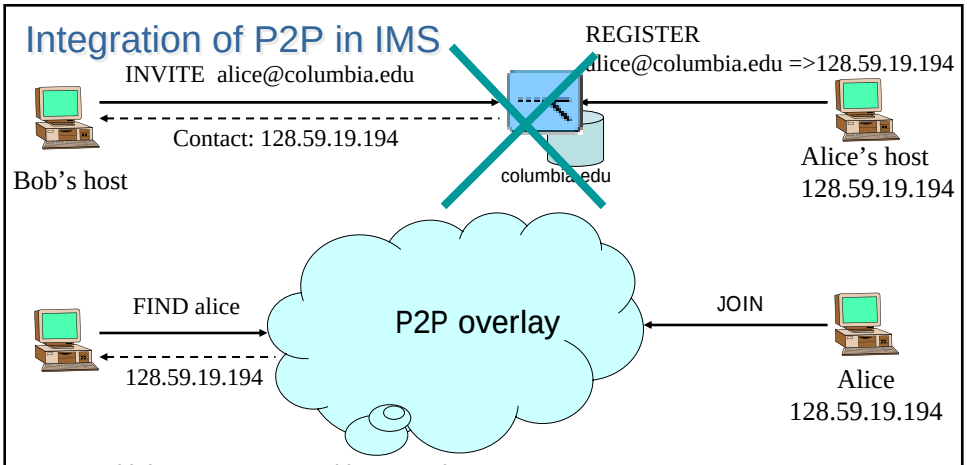

- Initial step represented by P2P sip
  - Reduce costly back-end telephony servers
  - No servers → no maintenance or configuration
  - Interoperability (unlike skype which is proprietary)
  - Handling heterogeneous and mobile nodes
  - Eliminating SIP server we increase scalability and robustness but also call set up latency (avg lookup 6 times higher than C-S SIP)

Integration of P2P in IMS (cont'd)

- Security, authentication, trust
  - Conventional P2P implements a no-trust relationship among nodes
  - User may impersonate other users
  - Not easy to integrate into P2P but essential and of great appeal
- QoS
  - P2P scales better in principle but is best-effort
  - Need to look at how QoS can be incorporated in P2P via OSA service-centric network

P2P incentive schemes / new business models

- P2P upsets the operator's economic models completely
  - The caller-pay model does not suite P2P
  - Who pays for traffic / computation when a peer is on 'server' mode?
  - Pay-as-you go does not suite P2P
  - Pay-flat-rate may saturate network / last mile bottleneck
- Need to look at suitable charging/billing models for operator-mediated P2P services
- It may be easier for the operator to offer/enforce incentive to P2P. E.g.:
  - You accumulate free SMS when your terminals acts as 'server'

MP2P is not straightforward

Existing P2P frameworks are not 'mobile-friendly'

- Too heavyweight to run on thin terminals
- Intolerant to intermittent connectivity (P2P pipes break)
- Too slow to sustain the dynamics of mobile networks
- Discovery time of mobile/intermittent resources is unpredictable (our study shows times ranging from minutes to days)
- Clash with mobile operator *modus operandi*

Existing P2P frameworks are not 'operator-friendly'

- Charging is difficult if not impossible
- Security, authentication and trust challenges
- P2P Anonymity is not always acceptable
- QoS is difficult if not impossible
- Lack of interoperability, standardisation

Some theoretical fundamental research issues

- The trade-off between routing table size and network diameter
- Effective autonomic overlay management
  - Keep it up-to-date
  - Keep it small
  - Keep it stable (node degree)
- Semantic ‘deep’ search
- Legal issues (copyright management): how do we integrate DRM into MP2P over IMS?

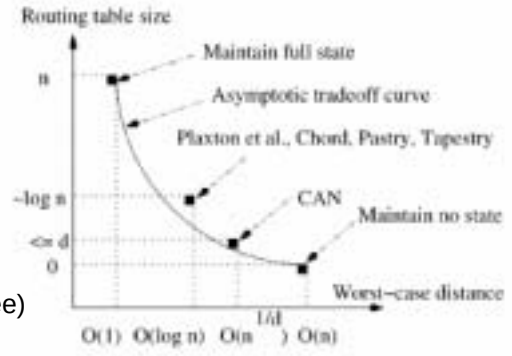

Source: J.Xu et al. IEEE JSAC Jan 2004.

Concluding remarks

- P2P: a communication means well beyond P2P file exchanging application
- Quality/scalability of P2P communication measurably better than CS-SIP (e.g. skype Vs. X-lite)
- Is P2P a threat to the operator?
  - No, provided that the operator stops its anti-disruptive attitude and starts integrating P2P
- Is the operator a threat to P2P? Maybe!
  - The operator can play dirty in the bit pipe
  - Current P2P economic models are rudimentary
  - Lack of security/trust hamper open P2P frameworks
- Turning threats into opportunities requires integration of P2P into IMS
- The management challenge: P2P overlay management
- The business challenge: incentive-based models