

Peer to peer: beyond file sharing

Anne-Marie Kermarrec

ASAP, INRIA



A case for decentralized systems

- Economical reasons
- Performance
- Some applications are intrinsically distributed
- Enhanced availability
- Resource sharing (data, storage, bandwidth) and aggregation
- Flexibility (load balancing)
- Incremental growth

Growing need of working collaboratively, sharing and aggregating distributed (geographically) distributed.



The new deal in distributed computing

Distributed systems are evolving

- **Scale shift**
- **Dynamics**

Traditional algorithms are no longer efficient : scalability

Peer to peer communication paradigm fills this gap

- Fully decentralized
- Self-organizing/enhanced availability
- Symmetric peers/load balancing
- Local knowledge of the system/global convergence




What makes P2P interesting?

- End-nodes are promoted to active components!
- Nodes **participate**, **interact**, **contribute** to the services they use: nodes share benefits AND duties
- Harness huge pools of resources available at the edge of the Internet
- Irregularities and unpredictability considered as the norm


Peer-to-Peer Systems





The core: overlay networks



Peer to peer overlay networks



Unstructured networks



Fully structured networks

- Provide various functionalities/performance: search, dissemination, etc
- Common characteristics
 - Self-organizing
 - Local knowledge
 - Resource aggregation
- Resulting properties
 - Scalability
 - Resilience to churn

Impact of the structure on search

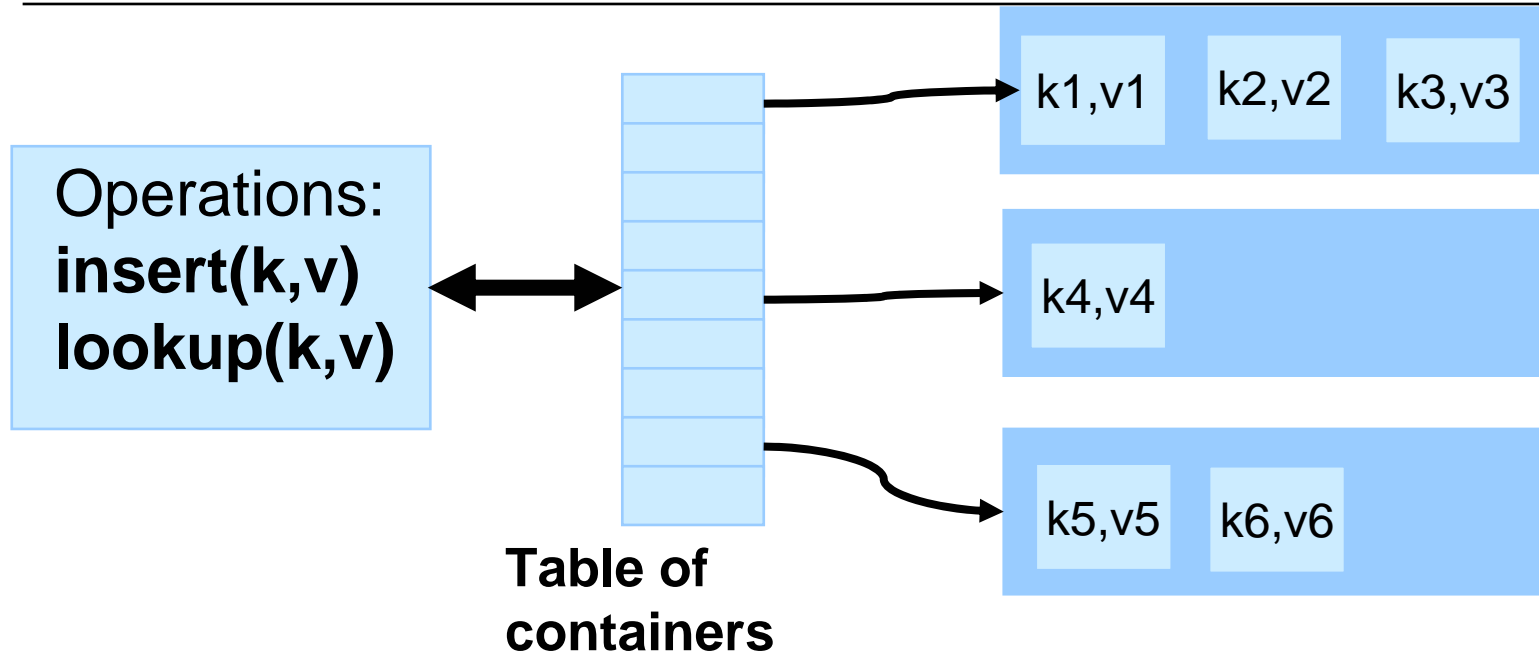
- Several ways of organizing a P2P overlay network
 - Search techniques: flooding versus routing
 - Expressiveness
 - Completeness
- Structured P2P overlay: DHT functionality
 - Support for exact search
- Unstructured gossip-based P2P overlays
 - Support for keyword-based search or range queries
- Weakly structured gossip-based overlays
 - Improve search efficiency upon fully unstructured overlays

1-Structured P2P networks

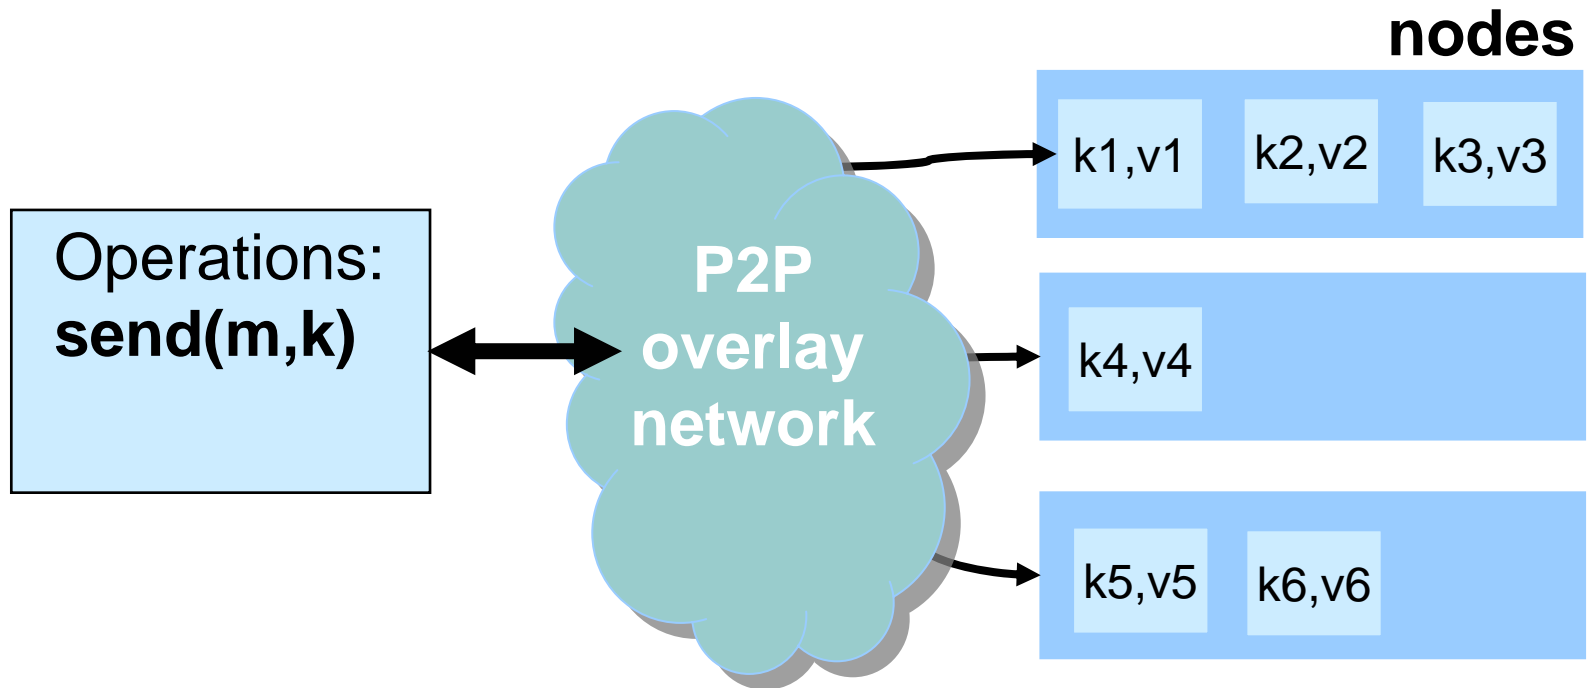
- Basic functionality: distributed hash table
- Applications
 - Content-delivery networks
 - Storage systems, Caching
 - Naming services
 - Multicast

Distributed Hash Table (DHT)

containers




Distributed Hash Table




Mapping

■ key ● nodeld



Identifier Space


Sending messages to keys



Pastry: Routing

Exemple

#(Irisatech-10/10/08)
= d46a1c



Properties

- $\log_{16} N$ hops
- Size of the state maintained (routing table): $O(\log M)$

Pastry: Routing table(#65a1fcx)

Line 0

<i>0</i>	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	<i>5</i>		<i>7</i>	<i>8</i>	<i>9</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>
<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>		<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>

Line 1

<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>		<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>
<i>0</i>	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>		<i>6</i>	<i>7</i>	<i>8</i>	<i>9</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>
<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>		<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>

Line 2


<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>		<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>
<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>		<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>
<i>0</i>	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	<i>5</i>	<i>6</i>	<i>7</i>	<i>8</i>	<i>9</i>		<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>
<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>		<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>

Line 3

<i>6</i>		<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>	<i>6</i>
<i>5</i>		<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>	<i>5</i>
<i>a</i>		<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>	<i>a</i>
<i>0</i>		<i>2</i>	<i>3</i>	<i>4</i>	<i>5</i>	<i>6</i>	<i>7</i>	<i>8</i>	<i>9</i>	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>	<i>e</i>	<i>f</i>
<i>x</i>		<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>	<i>x</i>

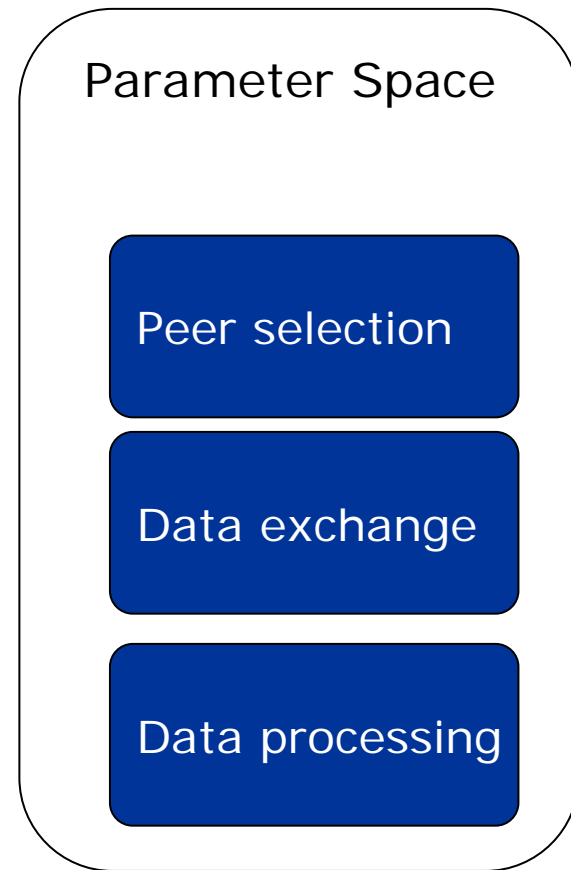

1-Unstructured P2P networks

- Flexible infrastructure
- Applications
 - Video streaming
 - Content-based search
 - Multicast



Gossip-based generic substrate

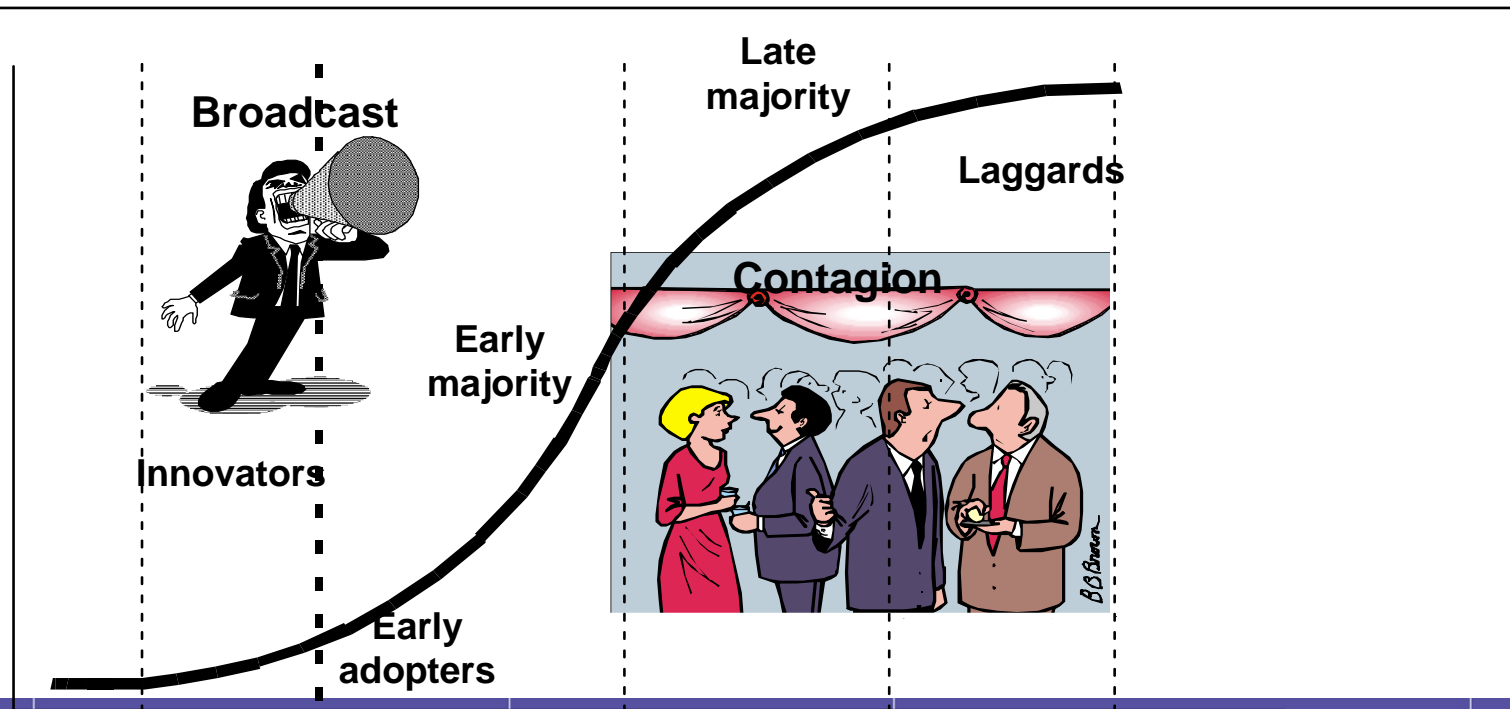
- Each node maintains a set of neighbours (c entries)
- Periodic peerwise exchange of information
- Each process runs an active and passive threads




Overlay maintenance

Peer selection	Random	Oldest	Closest
Data exchange	List of neighbours	½ List of neighbours	List of neighbours
Data processing	Random merging	Age-based merging	Proximity Based merging

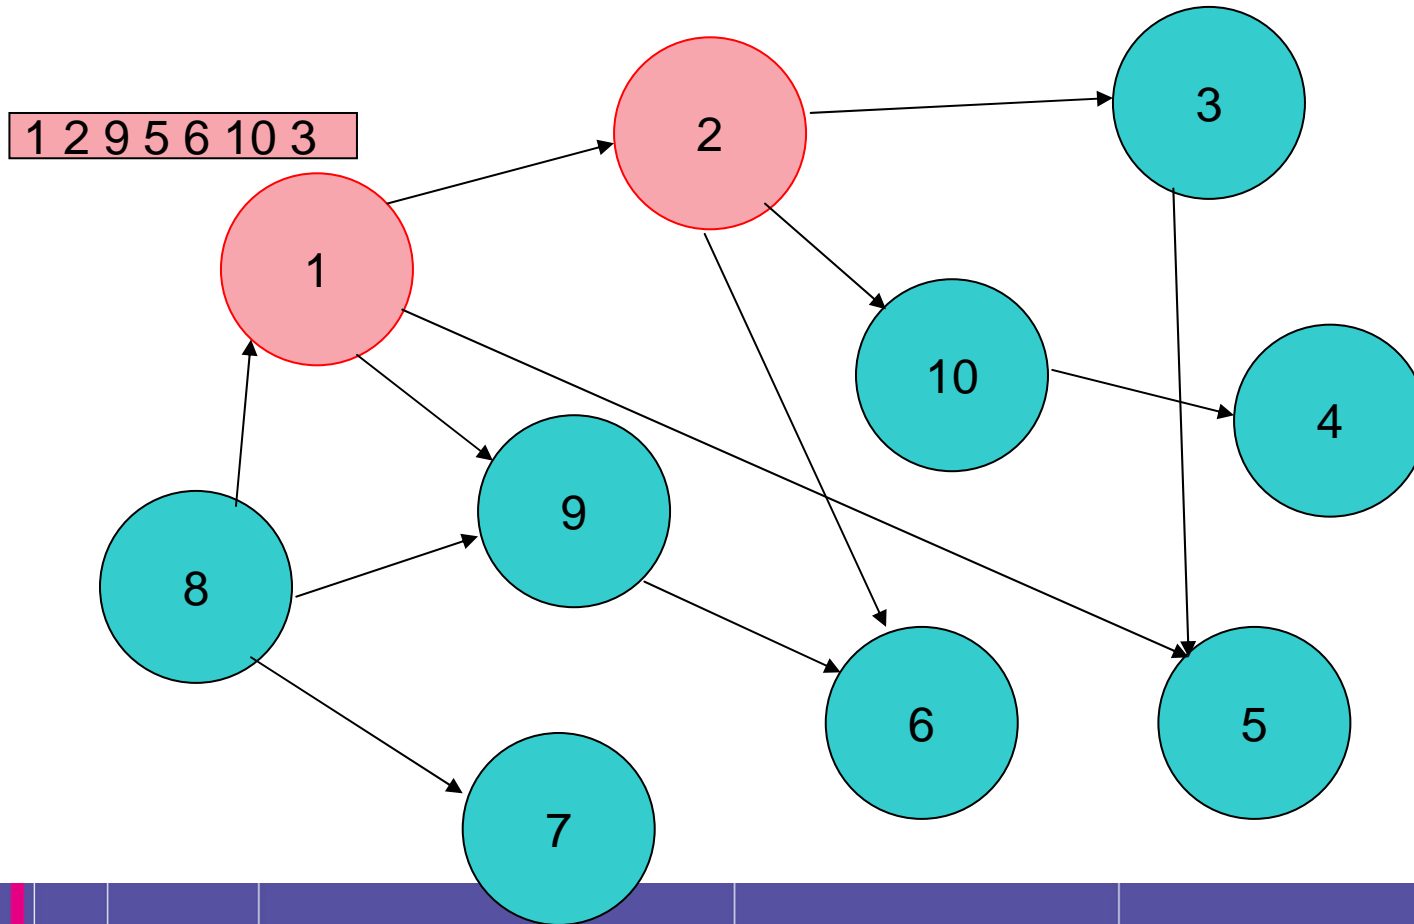
Why are we interested in building random graphs?




Example: Gossip-based generic protocol



Example: Gossip-based generic protocol



Example: Gossip-based generic protocol





Resulting graphs properties

Relationship « who knows who »

- Highly dynamic
- Capture quickly changes in the overlay networks

Flexible and powerful infrastructure

The take-away slide

