


Meta Service Discovery


Presented at IEEE 3rd Intl Workshop on Mobile P2P Computing (MP2P 06)

John Buford
Panasonic Digital Networking Lab

Alan Brown, Mario Kolberg
Dept of Computer Science
Stirling University


March 17, 2006





Service Discovery

- A Service Discovery Mechanism (SDM) is a way for networked devices to locate services offered by other nodes
- Many SDMs exist, including
 - IETF SLP, UPnP SSDP, Salutation, Bluetooth SDP, Web services UDDI
- To support wide-area service discovery, several approaches are being proposed
 - Multi-cast connected server hierarchies (SSDS and a proposed extension to Web Services)
 - Peer-to-peer service overlays such as INS/Twine



Meta Service Discovery

- Concept
 - In the future, many SDMs likely to co-exist in same environment
 - Service overlays and their associated SDMs will be tailored to different application and semantic requirements
 - Different network administrative domains want to control their own SDMs
- What problem is solved
 - How can a node determine what SDMs are available in given network domain
 - How can a node determine what SDMs are available in a given geographic area?
 - Both of these problems are important for roaming devices
- Relationship to existing SDMs
 - Existing SDMs are not open with respect to advertising to and allowing discovery of other SDMs
 - Existing SDMs can be used with meta SD



SDM Description

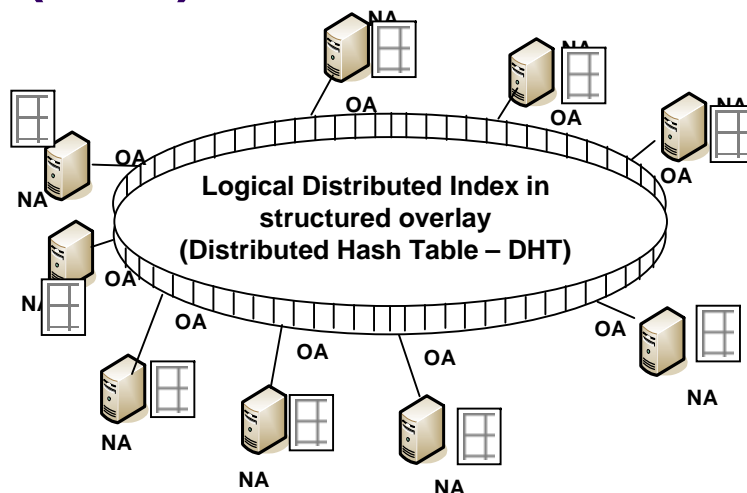
- Need an SDM description for advertisement and discovery
- Example

```
<SDM>
  <Domain>example.com</Domain>
  <Protocol>SLP</Protocol>
  <Location>New York</Location>
  <Operator>T. Smith</Operator>
  <Address>slp://services.example.com</Address>
  <Stub>http://services.example.com/slp.jar</Stub>
</SDM>
```

- Similar to, but simpler than, Service Descriptions used in many SDMs
- Other information may be in SDM, but may not need to be indexed
 - Protocol version, dependencies on other protocols
 - Security mechanisms/requirements
 - Other meta data characterizing scope of services
 - E.g., employee services, travel services, information services, ...



P2P Distributed Index Using Distributed Hash Table (DHT) in Structured Overlay



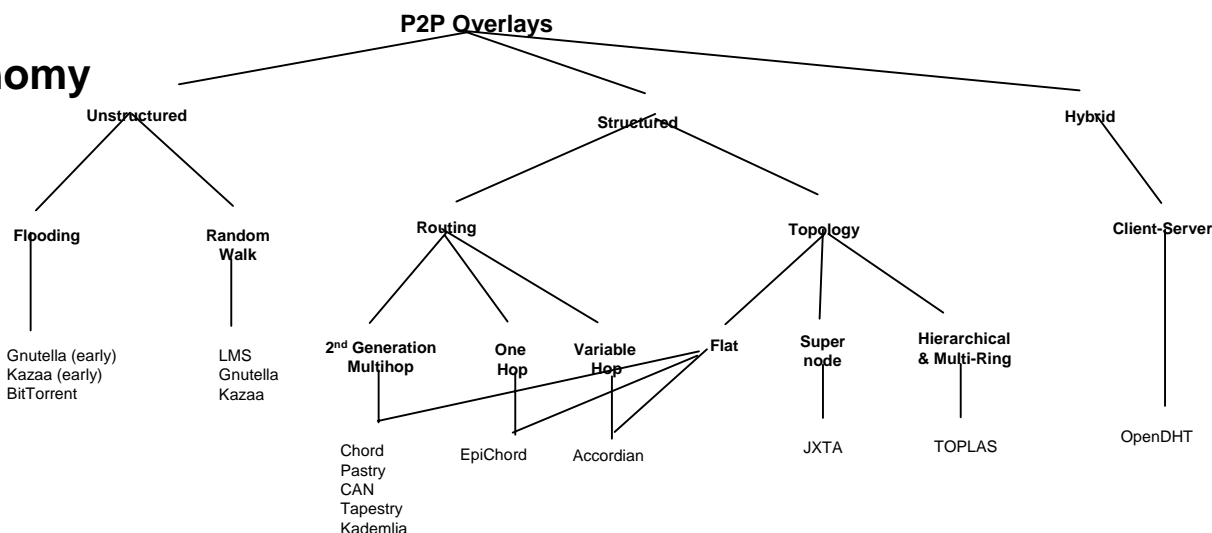
OA = Overlay Address
NA = network address




Routing Table

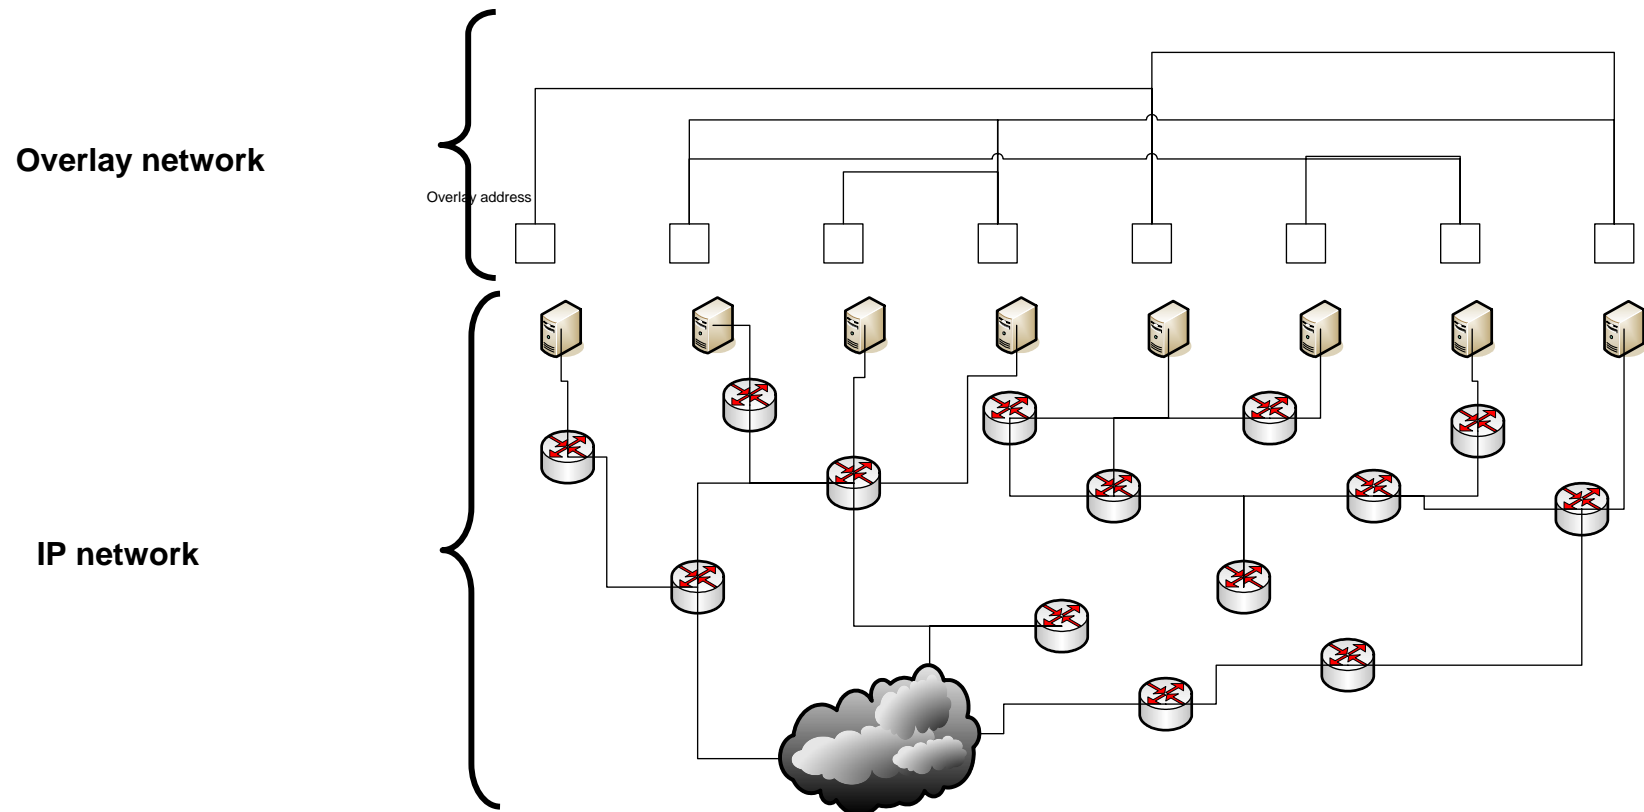
Peer to Peer Overlay Taxonomy

For an overview of P2P Overlays, see
J. Buford and K. Ross.
P2P Overlay Design
Overview. P2P SIP ad
Hoc Meeting. IETF 64.
Nov 2005



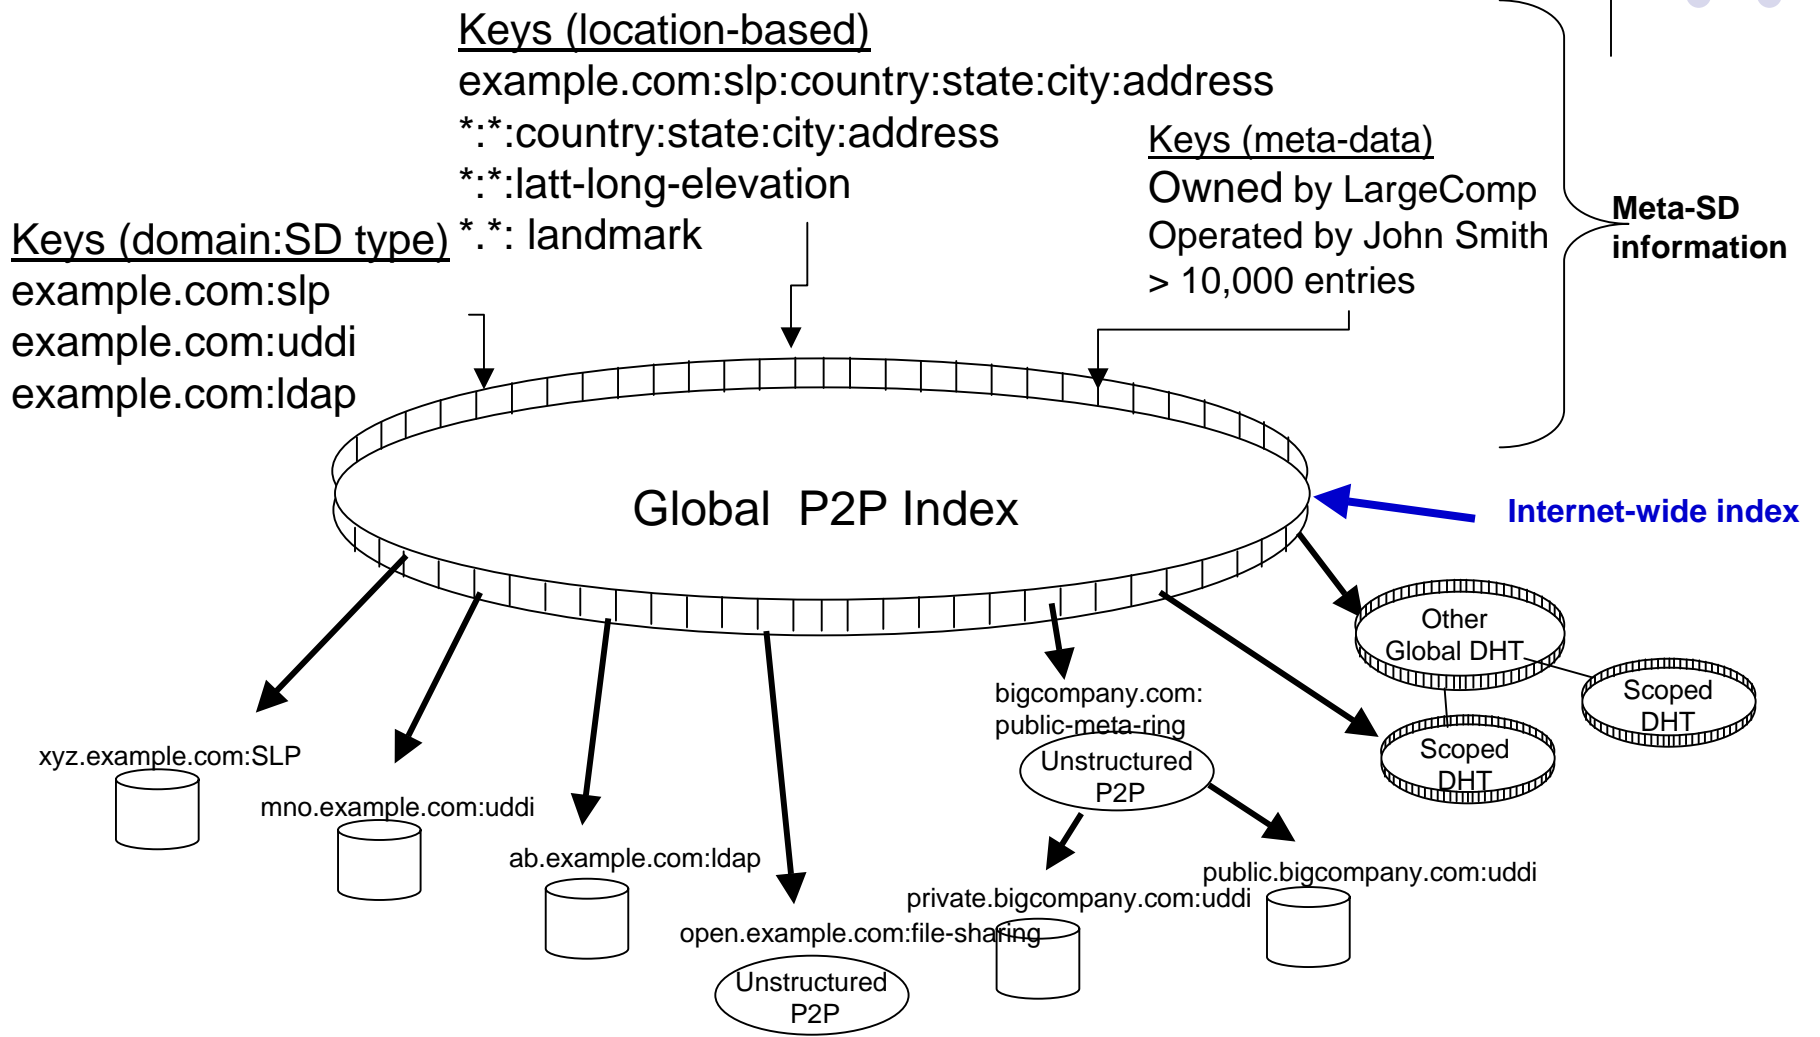


Overlay Network





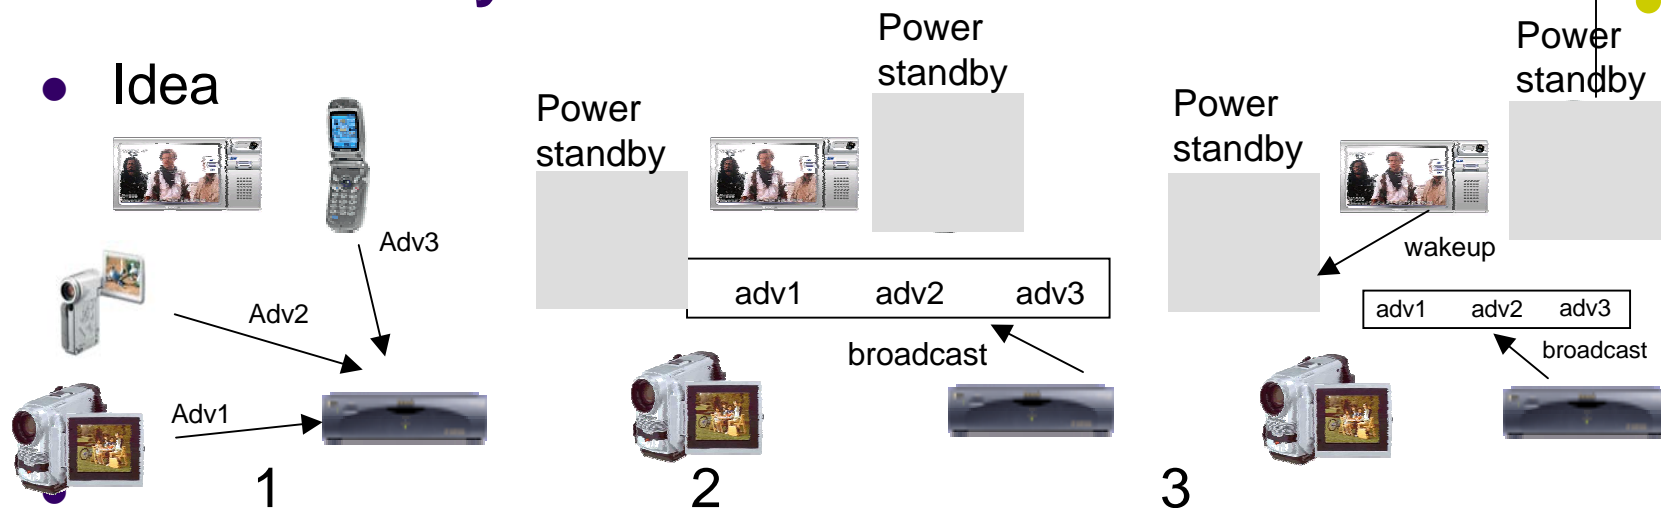
Meta Discovery Using P2P Index at Top Level





Meta Discovery in DDSD

- Idea



- Broadcast format examples (no groups)

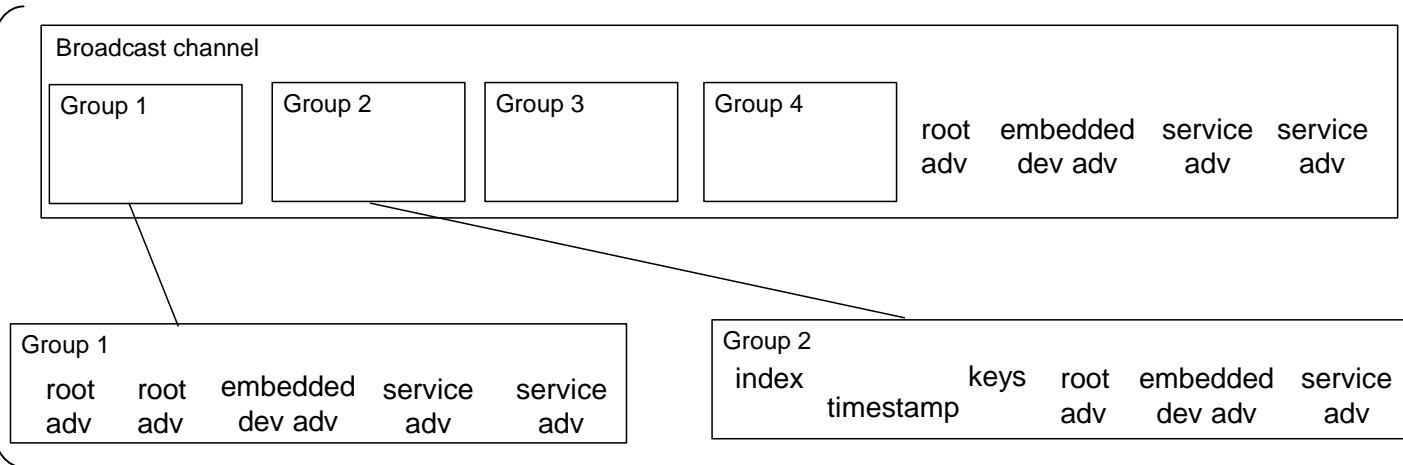
Broadcast channel		
adv	adv	adv

Broadcast channel								
index	timestamp of last change	root adv	root adv	embedded dev adv	embedded dev adv	service adv	service adv	

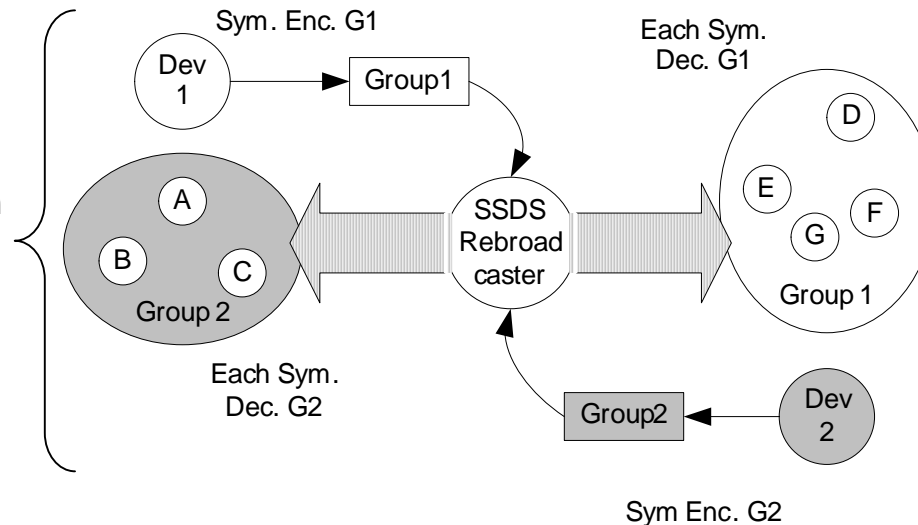


Meta Discovery in DDSD

Example channel organization for grouping advertisements




Example distribution of advertisements and invocation keys by group membership

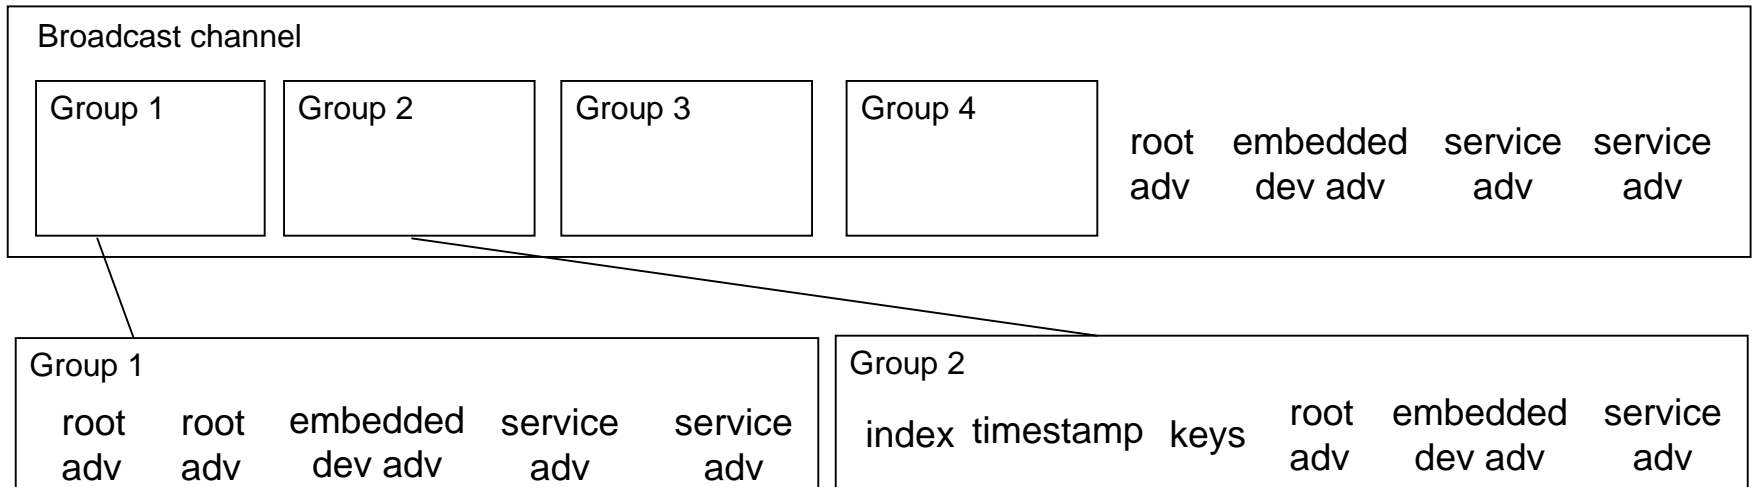





(a)

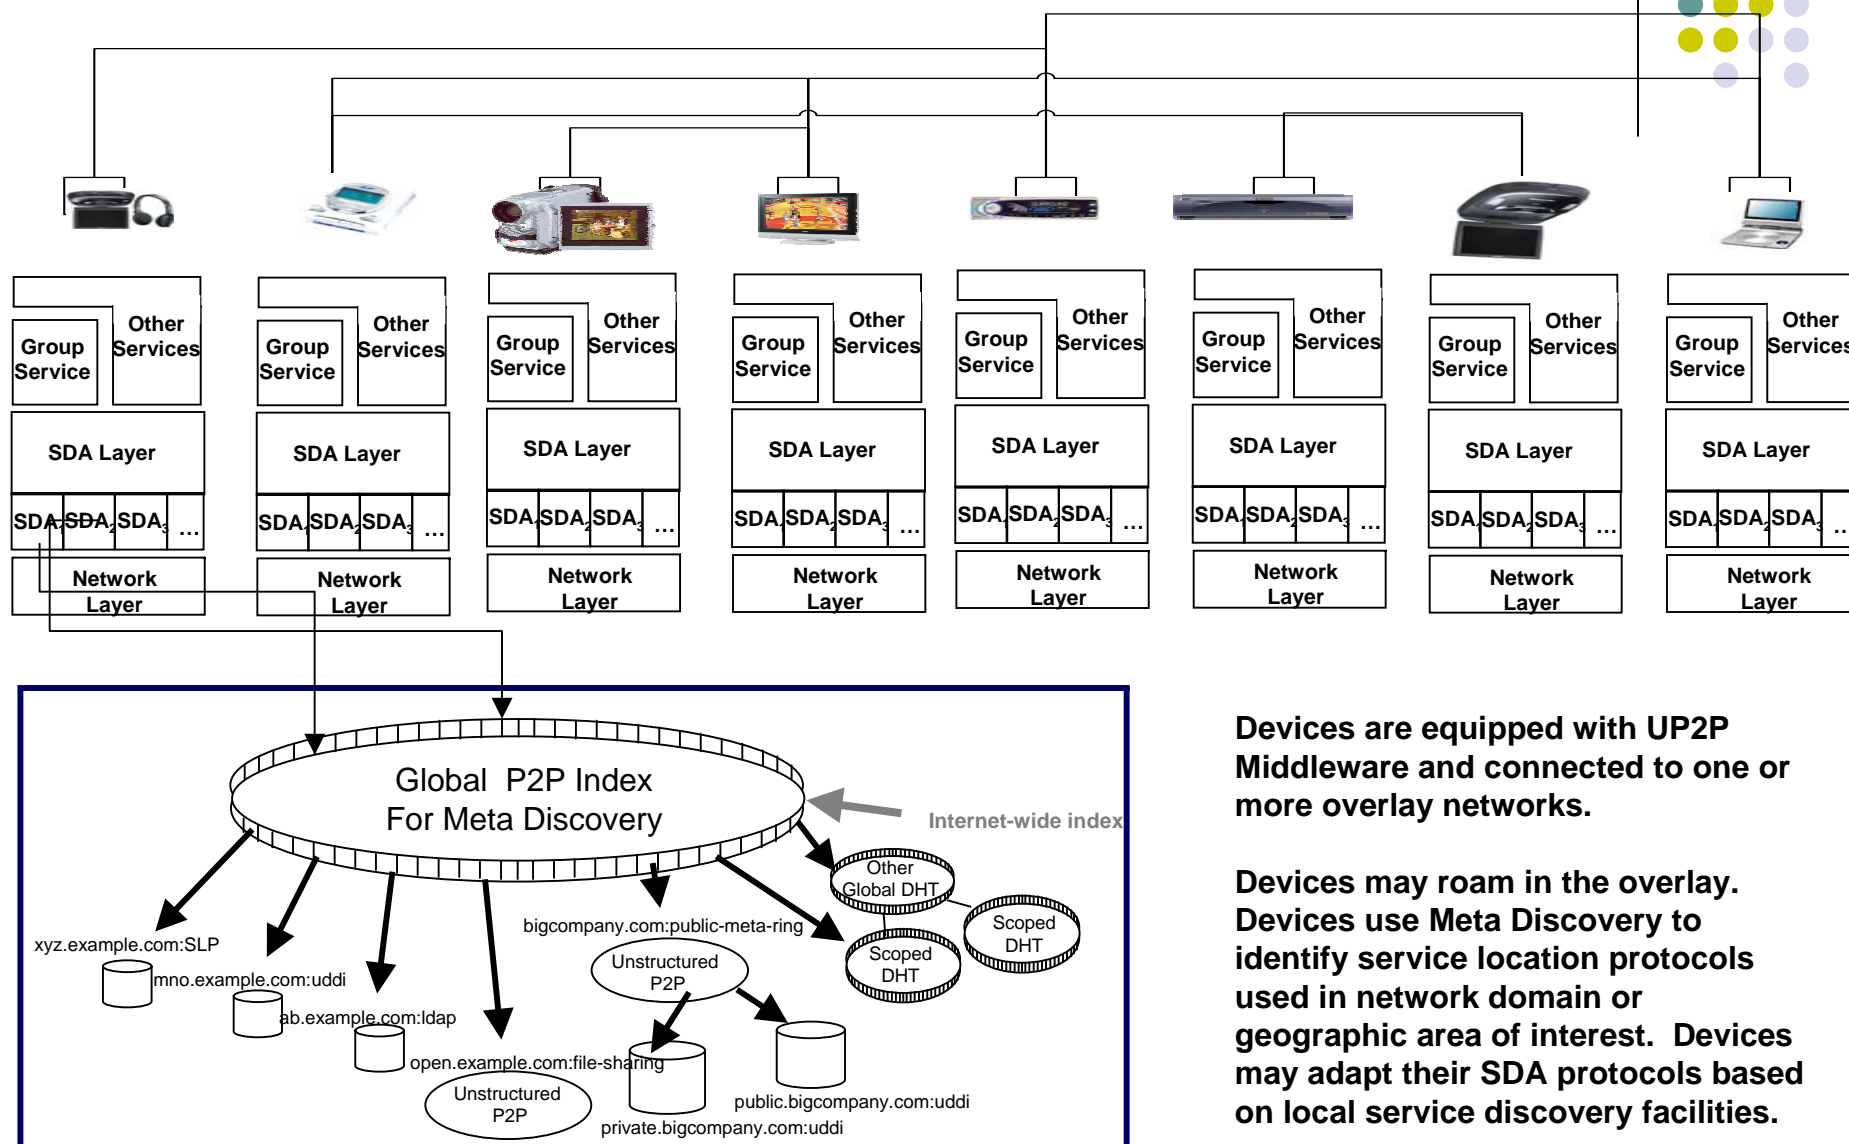



(b)




(c)





Devices are equipped with UP2P Middleware and connected to one or more overlay networks.

Devices may roam in the overlay. Devices use Meta Discovery to identify service location protocols used in network domain or geographic area of interest. Devices may adapt their SDA protocols based on local service discovery facilities.



Implementation Experiments

- Prototyped Meta-SD on 3 peer-to-peer overlays:
 - Pastry (Microsoft)
 - OpenDHT (Intel & Berkeley)
 - INS/Twine (MIT) (a service overlay built on Chord)
- Results
 - Storing an SDM description (or any structured document) in a DHT leads to multiple entries
 - About 15 for each SDM in our design
 - INS/Twine does this mapping automatically
 - OpenDHT had the most flexible insert mechanism of the 3 systems
- We were also interested in how many entries might be needed if scaled to Internet-wide, and how uniformly the keys would be distributed in the index

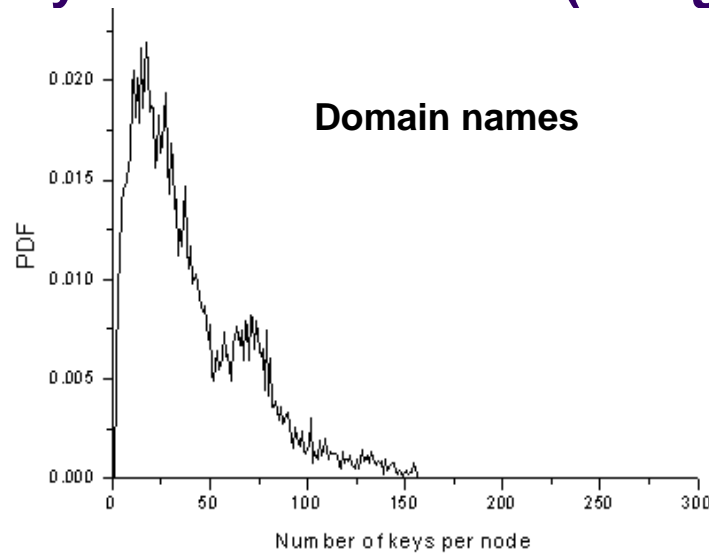


Internet-Scale Sizing

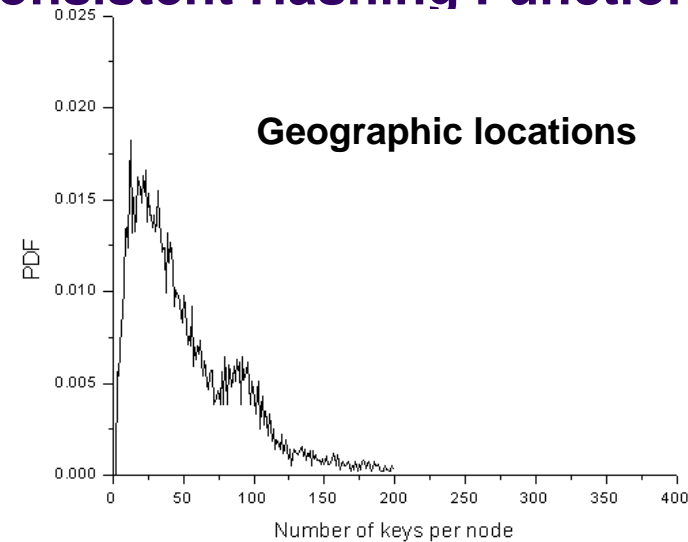
- Estimate the number of SDM entries needed if Meta Service Discovery were widely deployed.
 - assume that each domain that has type A DNS Resource Record advertises at least one SDM.
 - most recent measurement is **71,723,098** type A DNS entries.
- Each SDM might also be referenced by location
 - a location could be according to street address, landmark, or latitude-longitude (LL).
 - Street address and landmarks can be converted to a standard format and directly hashed, or they can be converted to a corresponding latitude-longitude.
 - LL can be normalized to decimal format and aligned to the nearest grid point.
 - The resulting grid point can be directly indexed, considerably simplifying lookup.
- Assume that SDM locations correspond to population densities
- Used two public data sets, one containing the LL position for the 2555 largest cities and another containing the square area for the 40 largest cities
 - Both are worldwide datasets.
 - For the 40 largest cities, assuming a grid spacing of 1 city block (about 200 meters), there are about **3.5M** grid points in the largest 40 cities
- **So for Internet-scale use, the index would have to store ~ 100M SDM descriptions, or about 1.5B entries**
 - **For an index served by 100K to 1M nodes, this would be up to 15K entries per node, a modest amount of data**



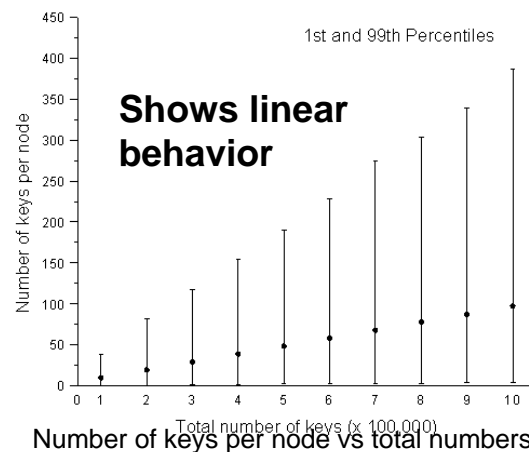
Key Distribution Tests (using Consistent Hashing Function)



PDF of key distribution of 405,097 real domain names in a 10,000 node network

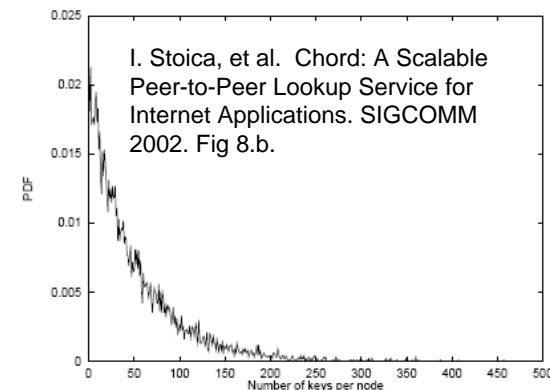



PDF for location based key distribution of 500,000 keys over 10,000 nodes



Number of keys per node vs total numbers of keys for 10,000 nodes

These key distributions are similar to random key insertions reported by Chord using the same hash function.





Conclusions

- Using the new Meta SD method, a node can determine what SDMs are available in given network domain and in a given geographic area
- An internet-scale Meta SD index is feasible using existing DHTs
 - Sizing and key distribution tests
- Has advantages compared to using an internet-scale service overlay
 - Indexing SDMs is more practical than indexing service descriptions