P6P: A Peer-to-Peer Approach to Internet Infrastructure

Lidong Zhou, Robbert van Renesse

Presented by: Dimitris Psaltoulis

P6P Overview

- P6P is a new network infrastructure
- Distinguishes end sites from the core Internet
- Decoupling of addresses as identifiers from addresses as locators for routing

Outline

- P6P Architecture
- P6P Routing
- Evaluation
- Comments
- Discussion

P6P Architecture (1)

- P6P breaks Internet in two parts
- End sites – P6P sites
- Core Internet

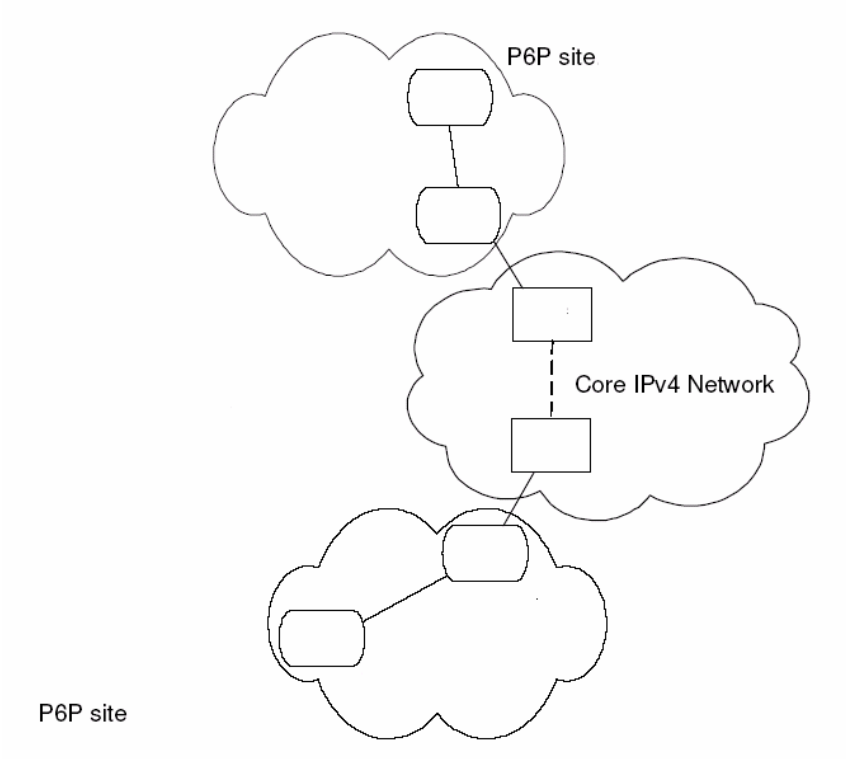

P6P Architecture (2)

- Each end site is assigned an IPv6 address called P6P address
- Each site has a unique site identifier
  - Common in the site
  - Location-independent
  - 48bit prefix of P6P address

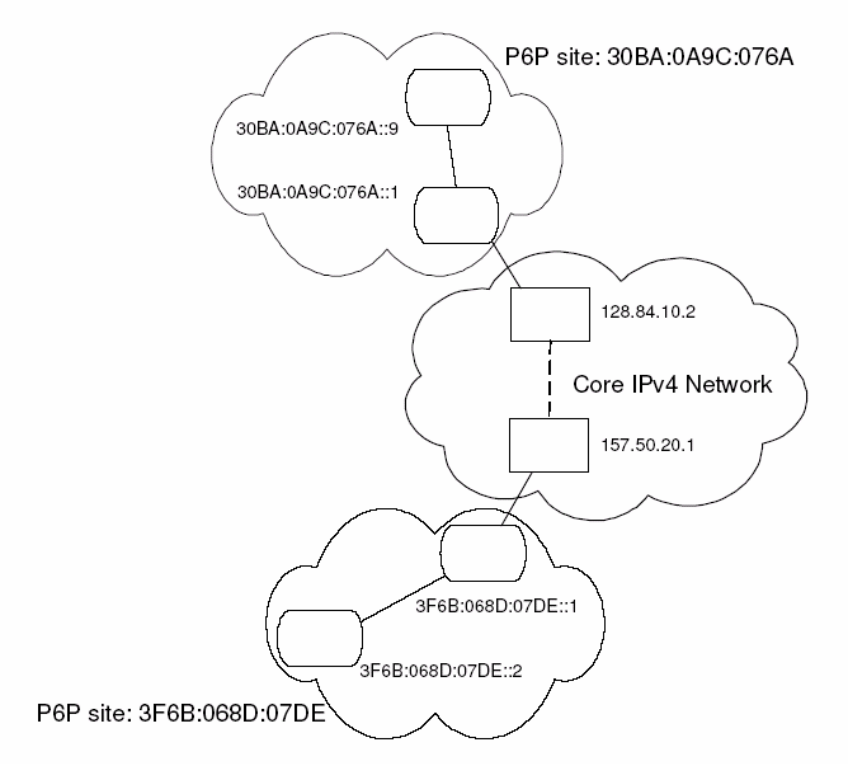

P6P Architecture (3)

- Routers in P6P are consisted logically of two types of compontents
  - Internal Gateway (IG)
  - External Gateway (EG)

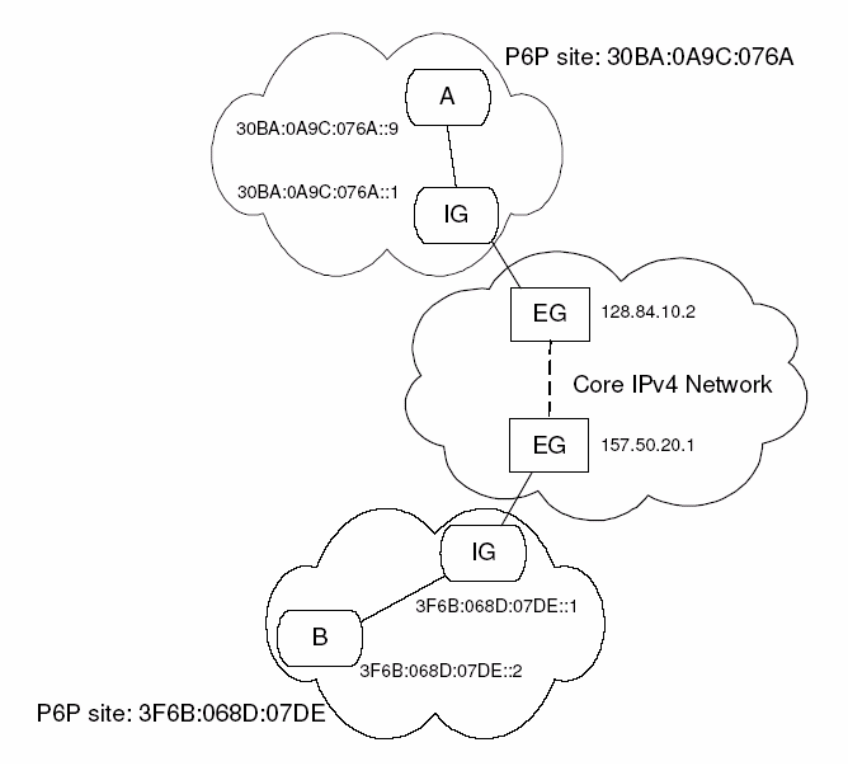

P6P Architecture (4)

- Each EG maintains a routing table

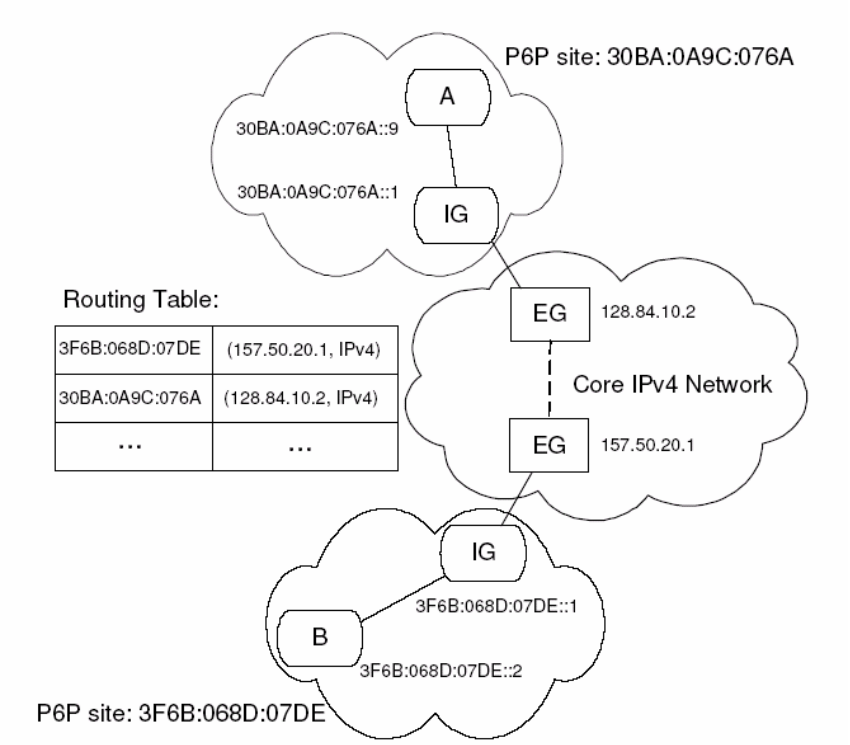

P6P Routing (1)

- P6P site identifiers are not location-dependent
- Non-hierarchical mapping is necessary
- The protocol used in P6P is based on a DHT

P6P Routing (2)

- Each EG runs a DHT agent
  - This solution is too expensive if done for every packet
- Each EG maintains a routing table that caches a subset of the mapping

P6P Routing (3)

- DHT(id) is the EG that the site identifier *id* maps to in the DHT
- An INSTALL message contains a routing entry (a mapping)
- A LOOKUP request contains an *id* and returns
  - Nothing if *id* does not exist in the EG's routing table
  - An INSTALL message with the corresponding entry if *id* exists in EG's routing table

P6P Routing (4)

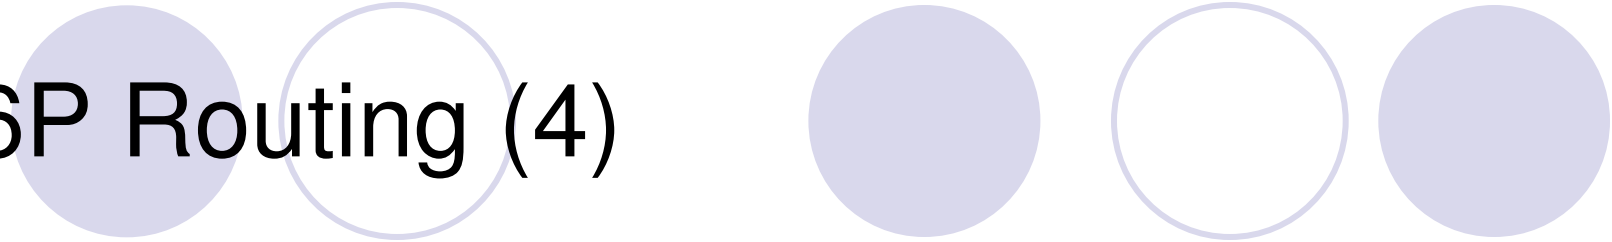

IG forwards a packet with destination *id* to EG

if (mapping for *id* exists in routing table)

    EG routes the packet

else

    EG sends LOOKUP request to DHT(*id*)

if (id exists in DHT(id))

    returns the mapping in an INSTALL message

P6P Routing (5)

- Optimization: Each EG piggybacks the K most recently looked up entries in its routing table on each routing message (LOOKUP or INSTALL) it sends.
- Address records are updated every T_{expire} seconds
- P6P uses public key cryptography to protect P6P sites from being hijacked
  - Routing table entries consist of a pair of X.509 certificates
    - Owner certificate
    - Map certificate

Evaluation

- Simulation runs for 1000 rounds and every round has *N* microrounds, where *N* is the number of sites
- In each microround:
  - A random source site is chosen
  - A random destination site is chosen
  - A packet is sent from a randomly chosen address within the source site to a randomly chosen address within the destination site

Evaluation – Average miss ratio (1)

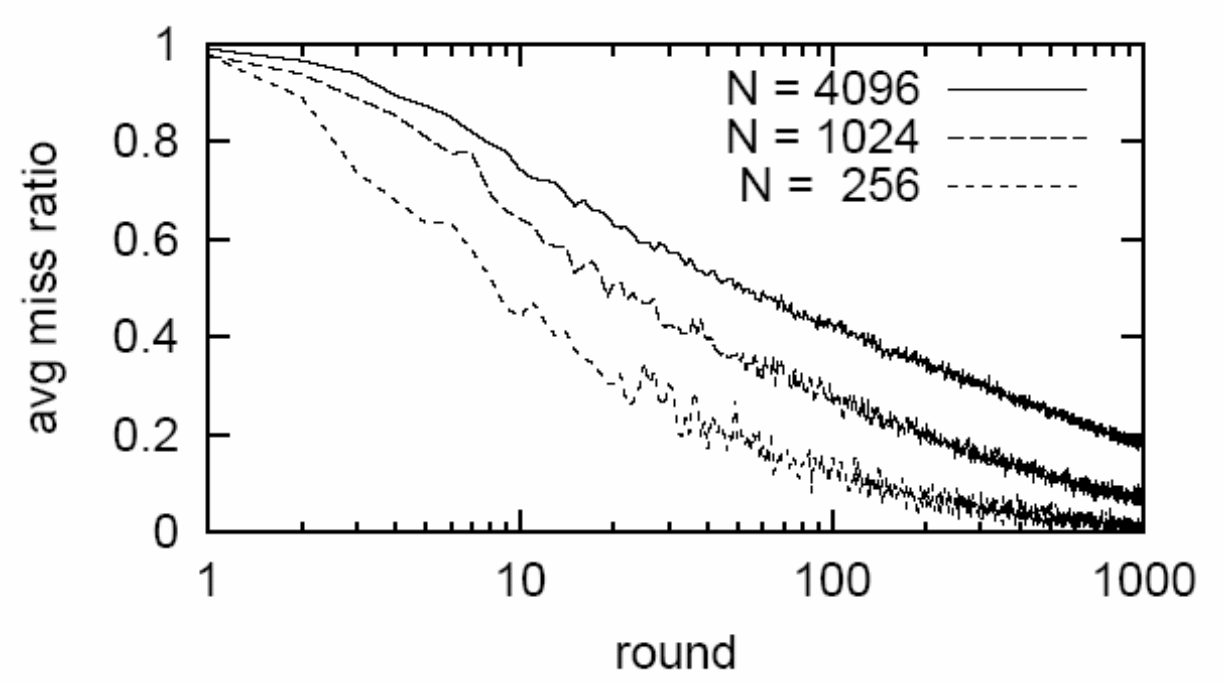

- K=8

Evaluation – Average miss ratio (2)

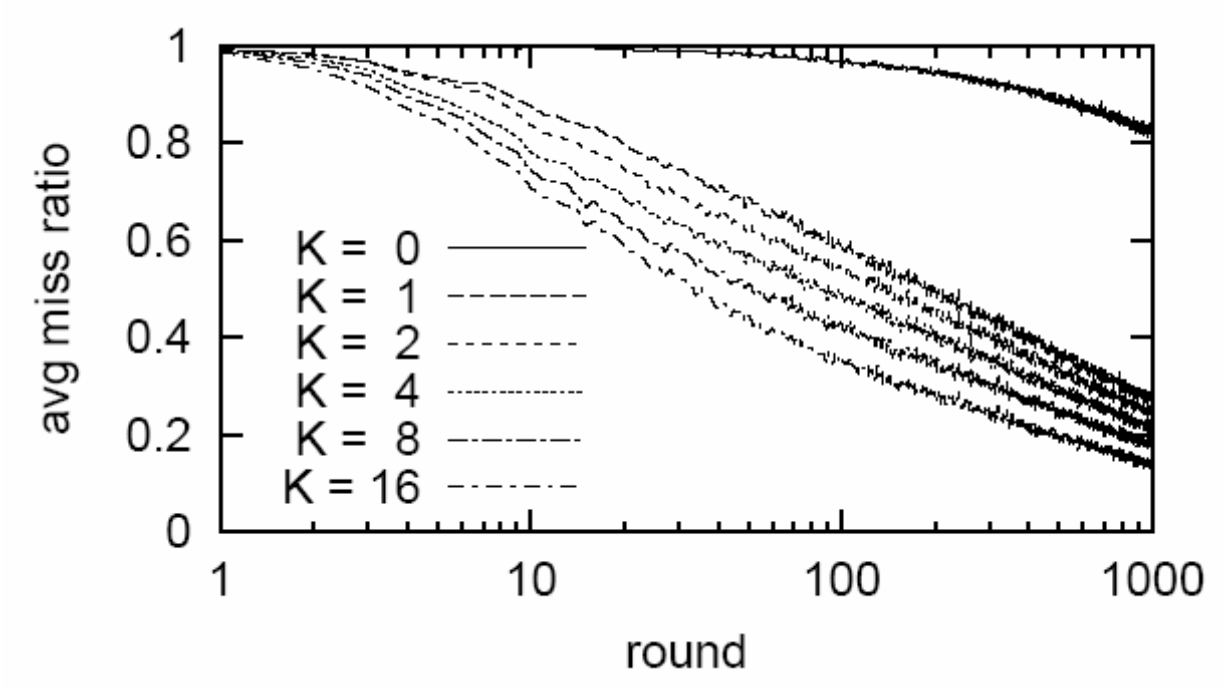

- N=4096

Evaluation – Average load on EGs

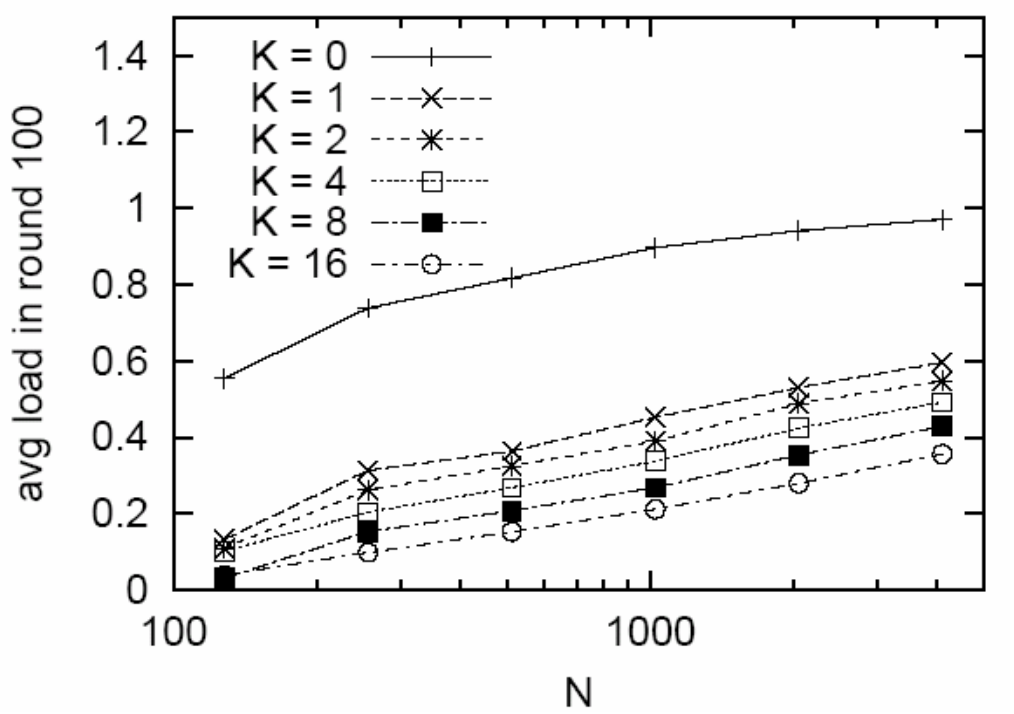

- Round 100

Comments

- Starts performing well after ~100*N microrounds. It would be interesting to see how long this is in a real life experiment
- No description of how the DHT agent works
- What happens if an EG fails?
- In order to introduce the use of IPv6 for end hosts, the authors address the need of associating each IPv6 address with an IPv4 address

Discussion

- Is DHT the right approach for mapping P6P site identifiers to core addresses? Other approaches (i.e. directory service)
- Is P6P scalable?
- Can it be deployed within a site?
- How can P6P support multihoming?