Trilogy
Re-Architecting the Internet

Routing tables as a *public good*
First steps/thoughts

Costas Kalogiros (AUEB)

February 25, 2009

Introduction

- Assume a *compact routing* scheme for identifiers (overlay routing)
  - Relies on an underlay routing scheme
  - In parallel to other overlay schemes

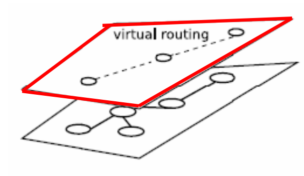

- Stretch vs. routing table size tradeoff
  - Users could influence system parameter k selection, but should be done according to their utility for fast routing

larger tables	smaller tables
faster routing	slower routing
lower k	higher k

- Purpose: Design a centralized mechanism that incentivizes users to declare their true utility and result in economic-aware routing schemes

The public good model

- Suppose system is composed of N users and let:
  - Q denote the total number of entries in *compact routing* tables
  - u(Q) be the maximum utility over all users
    - *Public good* notion: Q is non-rivalrous (it is not consumed when used by a user)
  - $\theta_i \varepsilon [0,1]$: user's i preference parameter (private information)
  - C(Q) is the total cost of supporting Q routing entries

θ_i u(Q) = user's i utility

Mechanism Design

- Mechanism designer announces rules and user i declares his θ_i
- Mechanism rules/outputs:
  - $Q(\theta)$ to be realized, where $\theta=(\theta_1, \theta_2, ..., \theta_n)$
    - affects compact routing parameter k
  - Access list of users $\pi(\theta) = (\pi_1(\theta), \ldots, \pi_n(\theta))$
    - User i is able to use system if $\pi_i(\theta)=1$
    - We assume that he cannot be serviced, but he is reachable through this routing scheme (that's why we don't assume this overlay routing scheme is the only one available)
  - Users' contribution fee $p(\theta)=(p_1(\theta), \ldots, p_n(\theta))$
    - can be number of routers

Mechanism Design

Aim is to maximize expected social welfare

$$\underset{\pi_1(\cdot),\ldots,\pi_n(\cdot),\,Q(\cdot)}{\text{maximize}} \; E\left[\sum_i \pi_i(\boldsymbol{\theta})\theta_i u(Q(\boldsymbol{\theta})) - c(Q(\boldsymbol{\theta}))\right]$$

subject to

budget balance

$$E\left[\sum_i \pi_i(\boldsymbol{\theta})p_i(\boldsymbol{\theta}) - c(Q(\boldsymbol{\theta}))\right] \geq 0$$

individual rationality

$$E\left[\pi_i(\theta_i,\boldsymbol{\theta}_{-i})\left\{\theta_i u(Q(\theta_i,\boldsymbol{\theta}_{-i}) - p_i(\theta_i,\boldsymbol{\theta}_{-i})\right\}\right] \geq 0$$

Incentive compatibility

$$E\left[\pi_i(\theta_i,\boldsymbol{\theta}_{-i})\left\{\theta_i u(Q(\theta_i,\boldsymbol{\theta}_{-i}) - p_i(\theta_i,\boldsymbol{\theta}_{-i})\right\}\right]$$
$$\geq E\left[\pi_i(\hat{\theta}_i,\boldsymbol{\theta}_{-i})\left\{\theta_i u(Q(\hat{\theta}_i,\boldsymbol{\theta}_{-i}) - p_i(\hat{\theta}_i,\boldsymbol{\theta}_{-i})\right\}\right]$$

Courcoubetis & Weber have shown that when N is large a simple **fixed fee** policy can achieve near-optimal Social Welfare

- Fee is independent of the peers' declarations

Future extension

- Supporting different levels of discovery for different users
  - This is possible in some compact routing schemes (like NICR). It would need an extra type of routing table for more direct routes.
  - Can this mechanism be applied to underlay routing as well?

source

destination (ID = 1482)

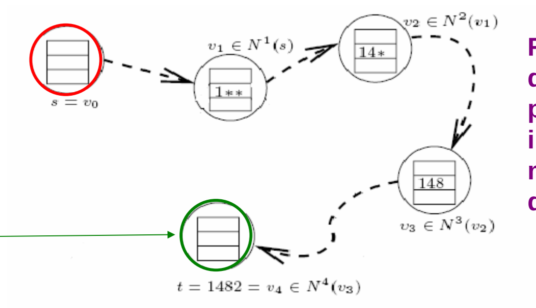

Routing to destination proceeds through increasing prefix matching of destination's ID