A BGP-Based Mechanism for Lowest-Cost Routing

Joan Feigenbaum, Christos Papadimitriou, Rahul Sami, Scott Shenker

Presented by: Tony Z.C Huang

Theoretical Motivation

- Internet is comprised of separate administrative domains known as AS. Interdomain Routing occurs among different ASes.
- Each AS has its own economic incentive for routing behaviors.
- ASes do not behave simply as compliant or Byzantine Adversaries, they behave strategically to maximizes their returns and minimizes their costs.
  - They can do so by lying about the actual information they hold.
- Using game theory to design a mechanism that ensures every router behaves in certain way that optimizes overall network performance.

Mechanism Design

- n **agents**, $i \in \{1, \ldots n\}$;
- Each agent has some private information, t^i, called **type.**
- **Output specification**: mapping each type vector
  $t = (t^1, \ldots t^n)$, to a set of allowed **output**s.
  intuitively, output specification tells what a good planner would do if it has access to all private informations.
- **Valuation Function**: assigning, for a single agent i, a score for an output given the agent's private information, i.e, $v^i(t^i, o)$.
- **Mechanism**: Define,
  1) for each agent, a set of **strategies** A^i;
  2) each agent plays a strategy $a^i \in A^i$;
  3) for an input vector: $(a^1, \ldots, a^n)$, the mechanism outputs two things
     1. a single output, $o = o(a^1, \ldots a^n)$
     2. **payment** vector, $p = (p^1, \ldots p^n)$, where p^i is a payment function $= p^i(a^1, \ldots a^n)$;

Mechanism Design (Continues)

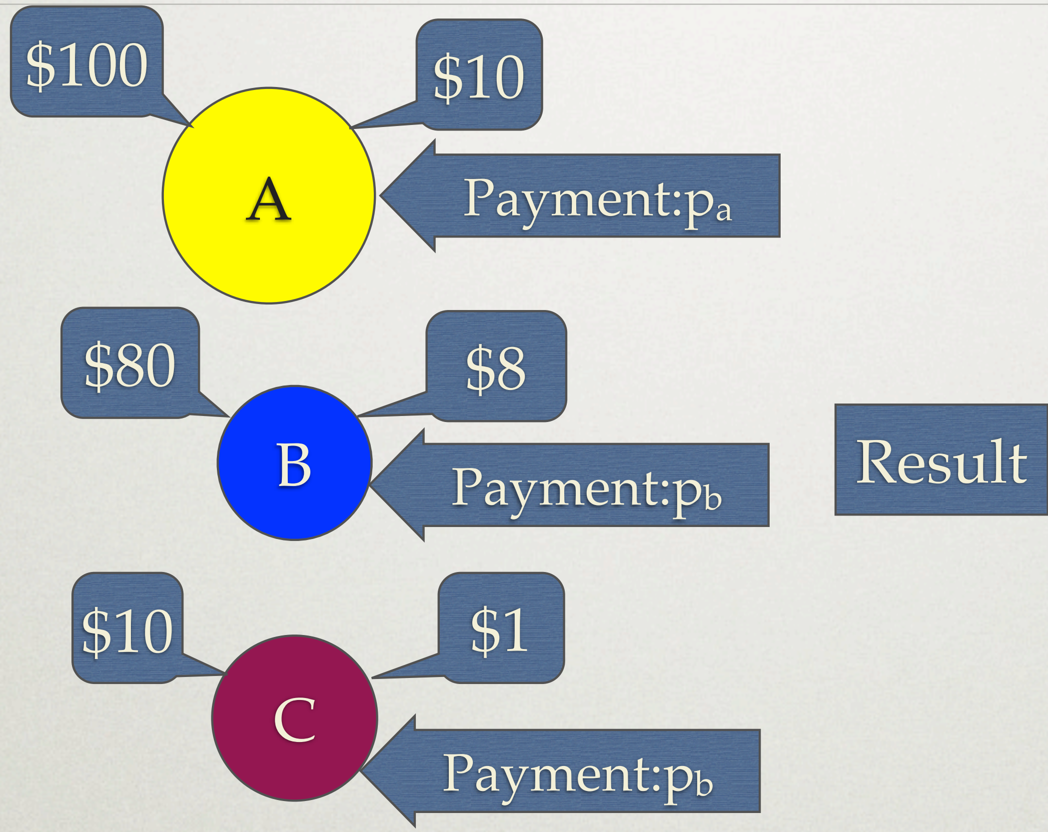

Output Specification: [100,80,10] --> [8,10]

Mechanism Design (Continues)

- Agent's **Utility**: $\tau_k = vi(ti, o) + pi$;
  - Represents the overall rewards to the agent.
  - Agent K's only goal is to maximize this utility.
- **Strategyproof Mechanism**: Mechanism where types are part of the strategy space Ai, each agent maximizes his utility by giving his type ti as input.
  $v^i(t^i, o(a^{-i}, t^i)) + p^i(a^{-i}, t^i) \geq v^i(t^i, o(a^{-i}, a^i)) + p^i(a^{-i}, a^i))$
  where, a^{-i} means $(a^1, \ldots a^{i-1}, a^{i+1}, \ldots a^n)$
- The mechanism wants each agent to report his private type truthfully, and pay agents in order to provide incentives for them to do so.

VCG Family Mechanism

- Consider an example: Vickrey's Second Price Auction
- Single Item is for sale on ebay among n buyers, with each buyer having a scalar value w_i that he is willing to pay (i.e, it's private information). Each people is bidding independently.
- Assume the price of final sale to be p.
- Utility in this case is
  $\tau_i = w_i - p$ (if i wins the item)
  $= 0$ (someone else wins)
- A natural social choice is to give the item to the player who values it the most.
- **Second Price Auction**: Player with highest prices wins, but pays 2nd-highest price.

VCG Family Mechanism (continues)

- Example:
  Player A bids \$10; (cost of -10)
  Player B bids \$8;   (cost of -8)
  Player C bids \$6;   (cost of -6)
- Under VCG: Highest wins, pays second highest: \$8;
  Under non-VCG: (Pay your bid), A bids \$10, pays \$10;
  (No payment), A bids \$10, pays none;
- Under non-VCG, agents have incentive to lie about their bids in order to gain more utility. For example, A would be better of stating its bid as \$9 for "Pay your bid", or ∞ under no payment.
- Under VCG, A can never be better off if it lies about its bids.
- The actual payment would be:
  payment = [declare "cost" (u_k)] + [cost without A (h_k)] - [cost with A (V)]
  = -10 + -8 - (-10) = -8.
  A pays -8, with utility of (10-8) = 2.

Problem Formulation

- The nodes are the strategic agents;
- Each As incurs a per-packet cost for carrying traffic which is the private information for each agent.
- The network is Biconnected, so there is no monopoly.
- Goal is to maximize the network efficiency by routing packet alongs the LCP.
- Takes in n AS numbers and constructs LCPs for all source-destination pairs.
- Compute the routes and payments with a distributed protocol based on BGP.
- Admitted Disadvantages:
  - Per-packet costs are not the best model.
  - LCP routing is not necessary the preferable routing policy.

Problem Formulation (continues)

- A network has a set of nodes N, n = |N|.
- A set L of bidirectional links between nodes in N.
- The network, called AS Graph, is biconnected, i.e, it's connected and nonseparable. Failure of a link would not separates the network.
- For any two nodes i,k ∊ N, T_{ij} is the traffic intensity (i.e, # of packets) being send from i to j.
- Node k incurs a transit costs c_k for each transit packet it carries, which is also the private information for this node.
- C_k assumes to be independent of which neighbor k received the packet from and which neighbor the packet is send to.
- Each node k is given a payment p_k to compensate it for carrying transit traffic.
- Payment can depends on the traffic and network topology, but the only assumption is that nodes receive no payment if carrying no traffic.

Problem Formulation (continues)

- Let $I_k(c;i,j)$ be the indication function for LCP from i to j;
  i.e $I_k(c;i,j)$ = 1 (k is a node on LCP from i to j)
  0 (otherwise)
  Note: $I_i(c;i,j) = I_j(c;i,j) = 0$.
- Total cost uk incurred by transit node k:
  $u_k(c) = c_k \sum_{i,j \in N} T_{ij} I_k(c;i,j)$
- **Total Cost** to the society:
  $V(c) = \sum_k u_k(c)$.
- We want to minimize the total cost to the society, and it's equivalent to minimizing, for every $i,j \in N$, the cost of the path between i and j.

Problem Formulation (continues)

- The only way we can assured of minimizing V(•) is for agents to inputs their true costs. And we must rely on pricing scheme to incentivize agents to do so.
- The mechanism must be strategyproof, i.e for all x
  $\tau_k(c) > \tau_k(c \mid {}^k x)$,
  where $c \mid {}^k x = c_i \; (i \neq k)$
  $x \; (i = k)$

Pricing Mechanism

Theorem 1 *When routing picks lowest-cost paths, and the network is biconnected, there is a unique strategyproof pricing mechanism that gives no payment to nodes that carry no transit traffic. The payments to transit nodes are of the form* $p^k = \sum_{i,j \in N} T_{ij} p^k_{ij}$, *where*

$$p^k_{ij} = c_k I_k(c;i,j) + \left[\sum_{r \in N} I_r(c|^k\infty;i,j)c_r - \sum_{r \in N} I_r(c;i,j)c_r \right].$$

What pricing Scheme Really says...

- Payment to agent k for carrying traffic from i to j =
  (1) 0 if k does not carry any traffic, otherwise
  (2) payment = k's declared cost + k's helpfulness
  = k's declared cost
  + (cost to the network without k)
  - (cost to the network with k)
- This is provably the only payment scheme that guarantees:
  (1) optimal social welfare.
  (2) truthful report of private information from agents.

Implementation

- By using the current BGP protocol to distribute information about the cost function in the network.
- To know the payment, we need to know the
  (1) private information;
  (2) cost to the network when routing through me;
  (3) cost to the network when routing through an alternative paths;
- (1) is obvious.
- (2) is computed by summing the information carried in the BGP header along its AS-path, and is stored in the BGP routing table.
- (3) can not be calculated immediately, but it can be estimated by keeping a upper bound and update this upper bound by the information carried in BGP headers, and have it converges to a stable correct value.
  - Big Assumption for this result: the network topology is static.

Implementation (continues)

- Details of how this algorithm is derived is omitted here
- The key result is that the estimation is tight, and it converges to the correct LCP and prices after a certain number of rounds.

- If a is i's parent in $T(j)$, then i scans the incoming array and updates its own values if necessary:

$$p_{ij}^{v_r} = \min(p_{ij}^{v_r}, p_{aj}^{v_r}) \quad \forall r \leq s-1$$

- If a is a child of i in $T(j)$, i updates its payment values using

$$p_{ij}^{v_r} = \min(p_{ij}^{v_r}, p_{aj}^{v_r} + c_i + c_a) \quad \forall r \leq s$$

- If a is neither a parent nor a child, i first scans a's updated path to find the nearest common ancestor v_t. Then i performs the following updates:

$$\forall r \leq t \;\; p_{ij}^{v_r} = \min(p_{ij}^{v_r}, p_{aj}^{v_r} + c_a + c(a,j) - c(i,j))$$
$$\forall r > t \;\; p_{ij}^{v_r} = \min(p_{ij}^{v_r}, c_k + c_a + c(a,j) - c(i,j))$$

Discussion?

- Known issues to VCG mechanism.
- Over payment?
- Unrealistic assumptions made by author about the network.
- Does re-feedback fall into this theoretical category?