A Cost Model for Network Traffic (with an application to paid-peering)

Amogh Dhamdhere (CAIDA)
with
Murtaza Motiwala, Nick Feamster (Georgia Tech),
Anukool Lakhina (Guavus) [Cost Model]

Pierre Francois (UCL), Constantine Dovrolis (Georgia Tech)
[Paid-peering]

2nd Workshop on Internet Economics, UCSD, December 2011

Outline

- These are really two talks
- But they are related
- Part 1: Formalizing a cost model for network traffic

[CCR'12]

- Part 2: A value-based framework for peering agreements

[ITC'10, NANOG 49]

- The cost model can be useful for measuring peering "value"

Part 1: Formalizing a Cost Model for Network Traffic

Optimizing Network Costs

- Traffic-related costs contribute to the total cost of running a network
- Routing in recession: configuring routing in a network to minimize traffic-related costs
- Relatively easy: How do Individual elements contribute to costs? Harder: How much do individual ingress-egress flows cost?
- Need a holistic traffic cost model that can attribute costs to individual flows

Need for a Traffic Cost Model

Need for a Traffic Cost Model

Cost-based paths
selector

Need for a Traffic Cost Model

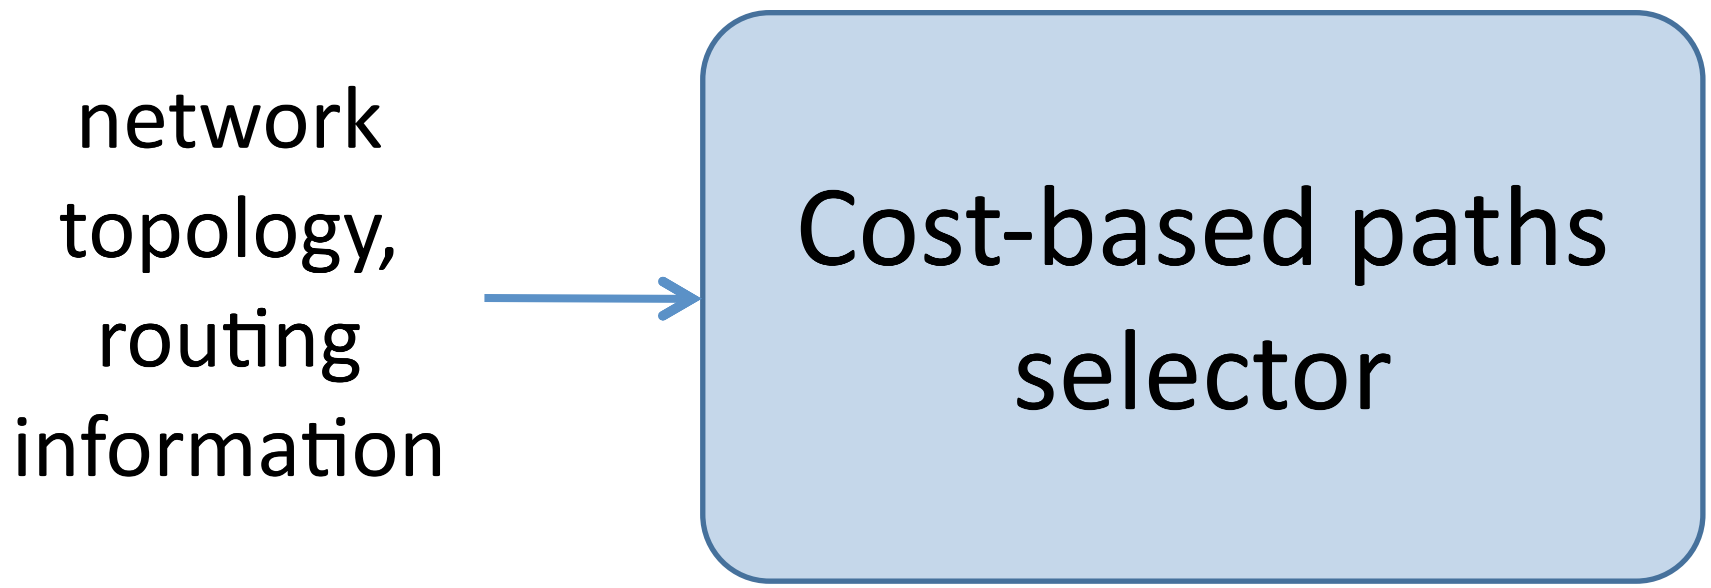

Need for a Traffic Cost Model

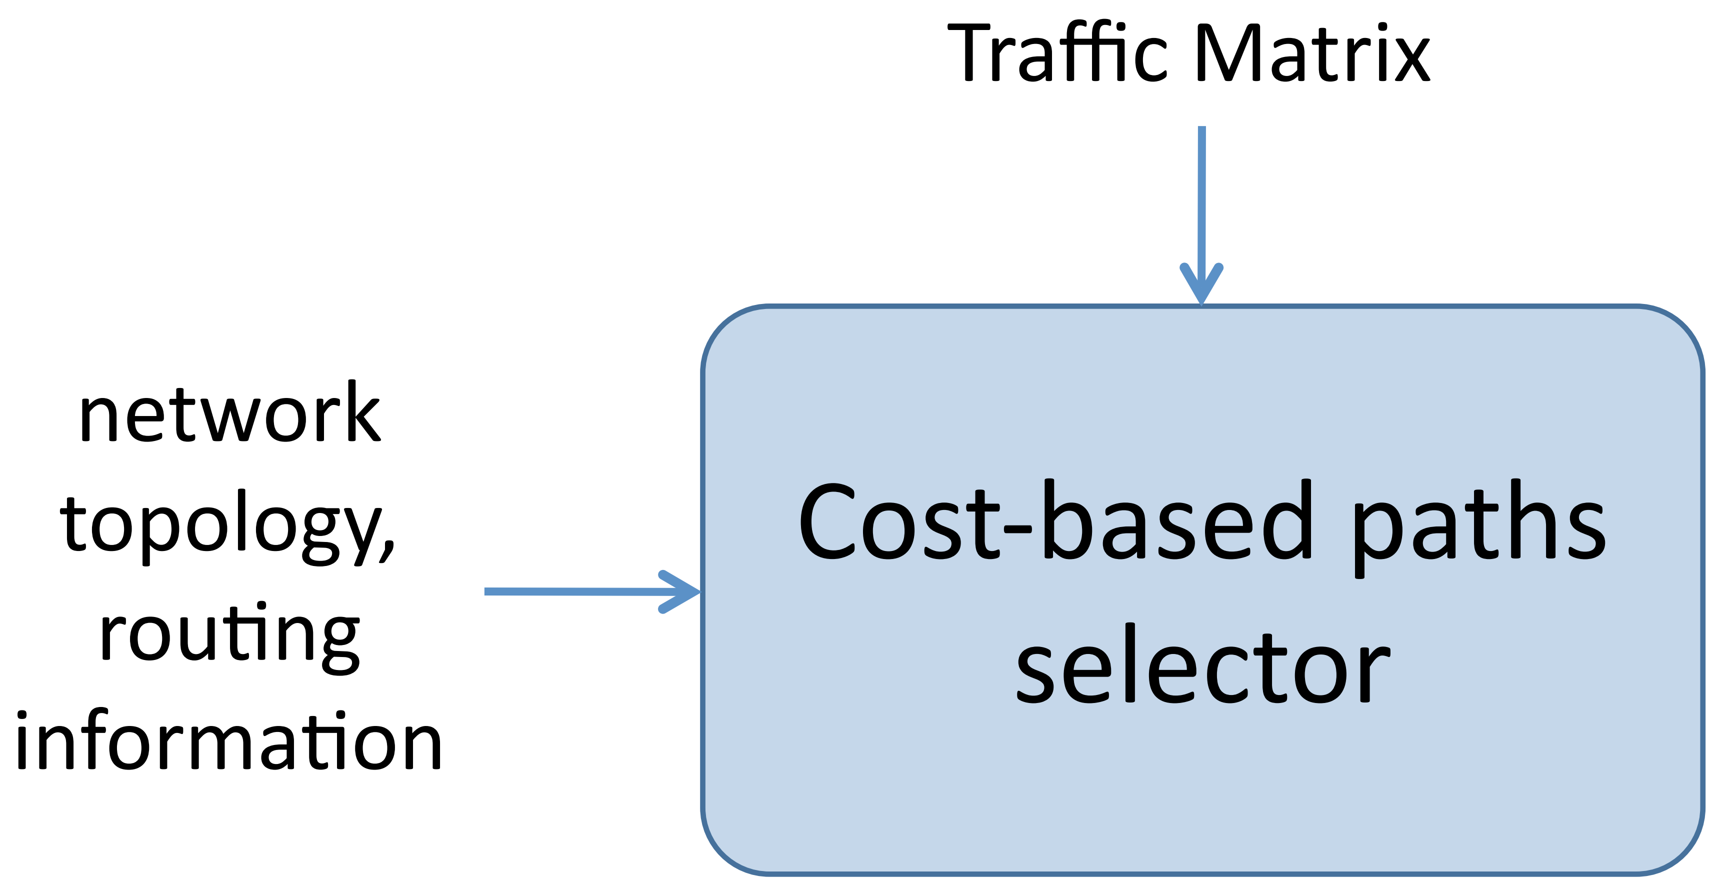

Need for a Traffic Cost Model

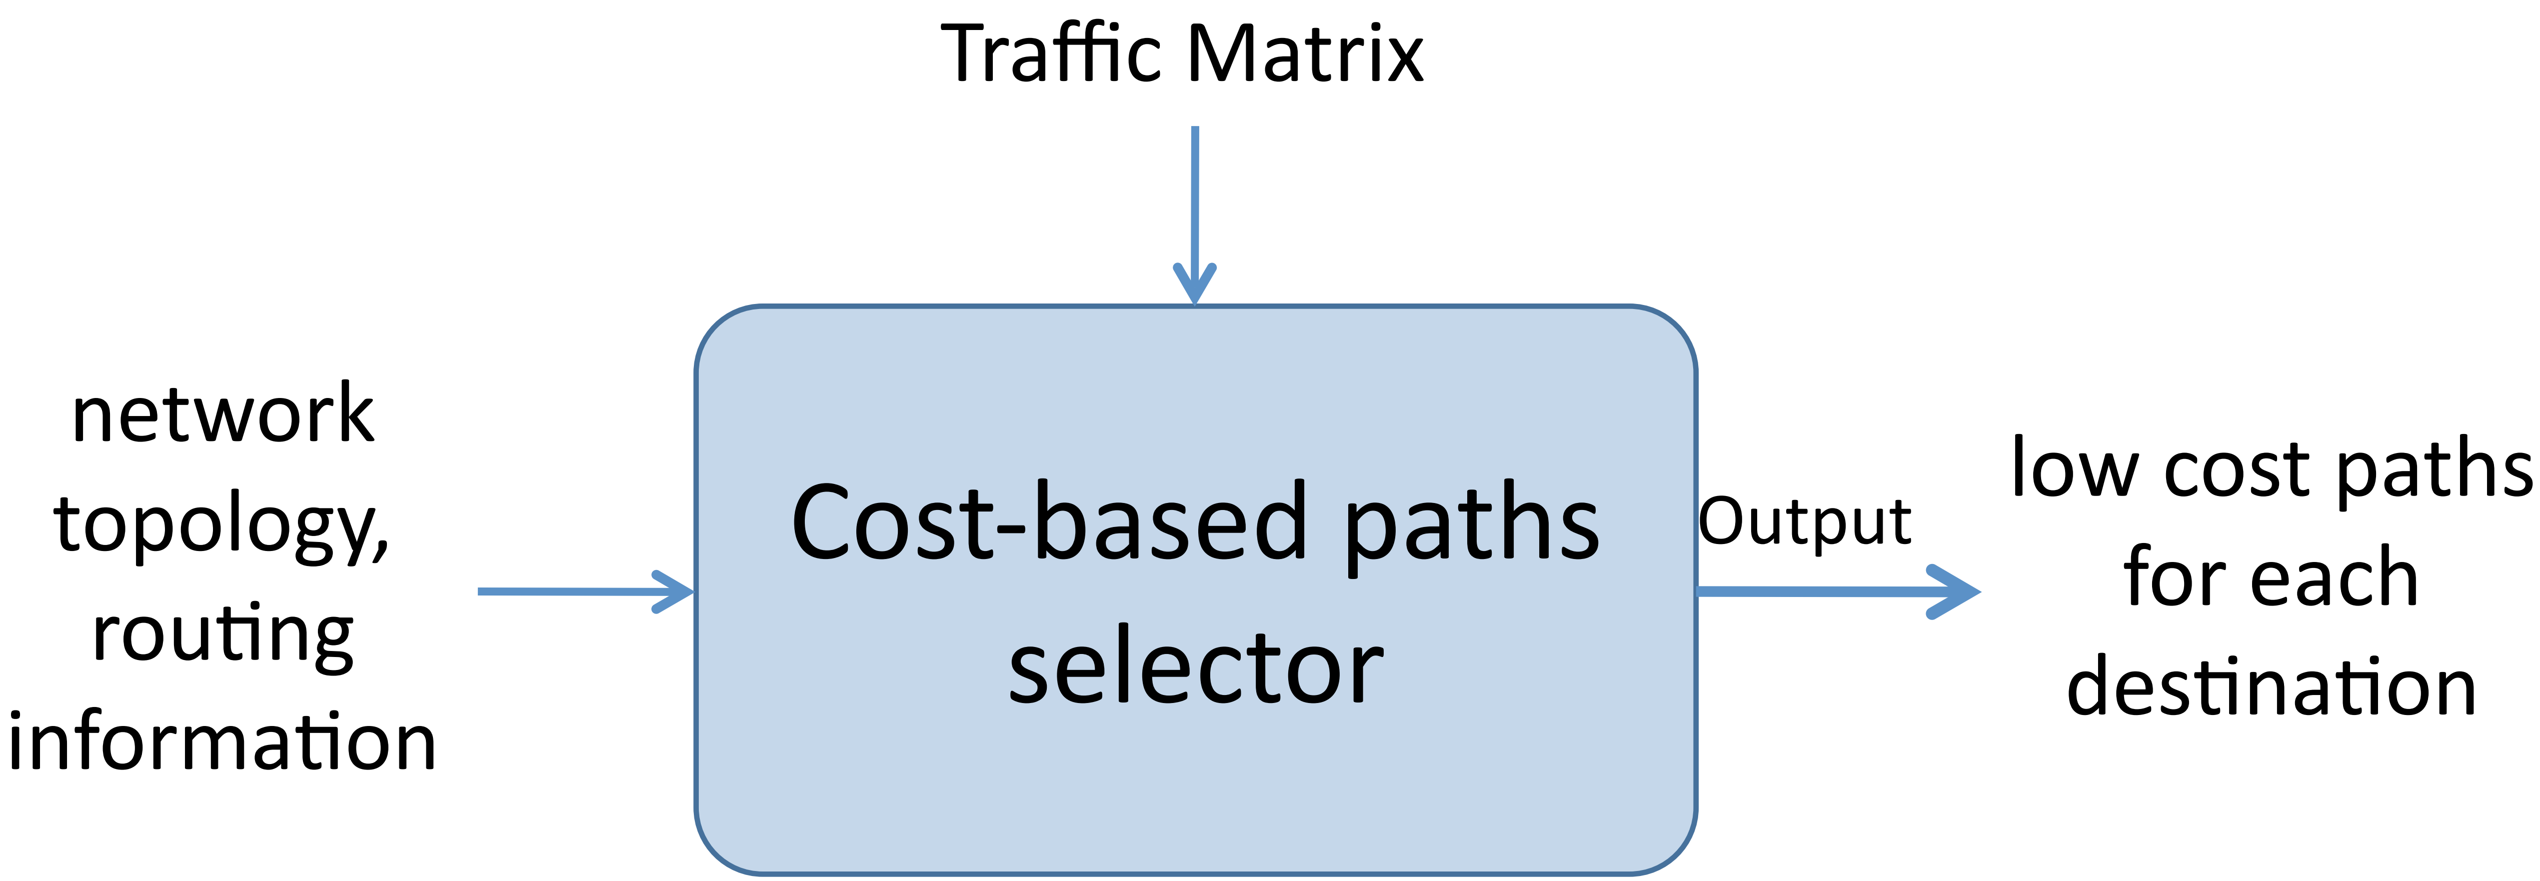

Need for a Traffic Cost Model

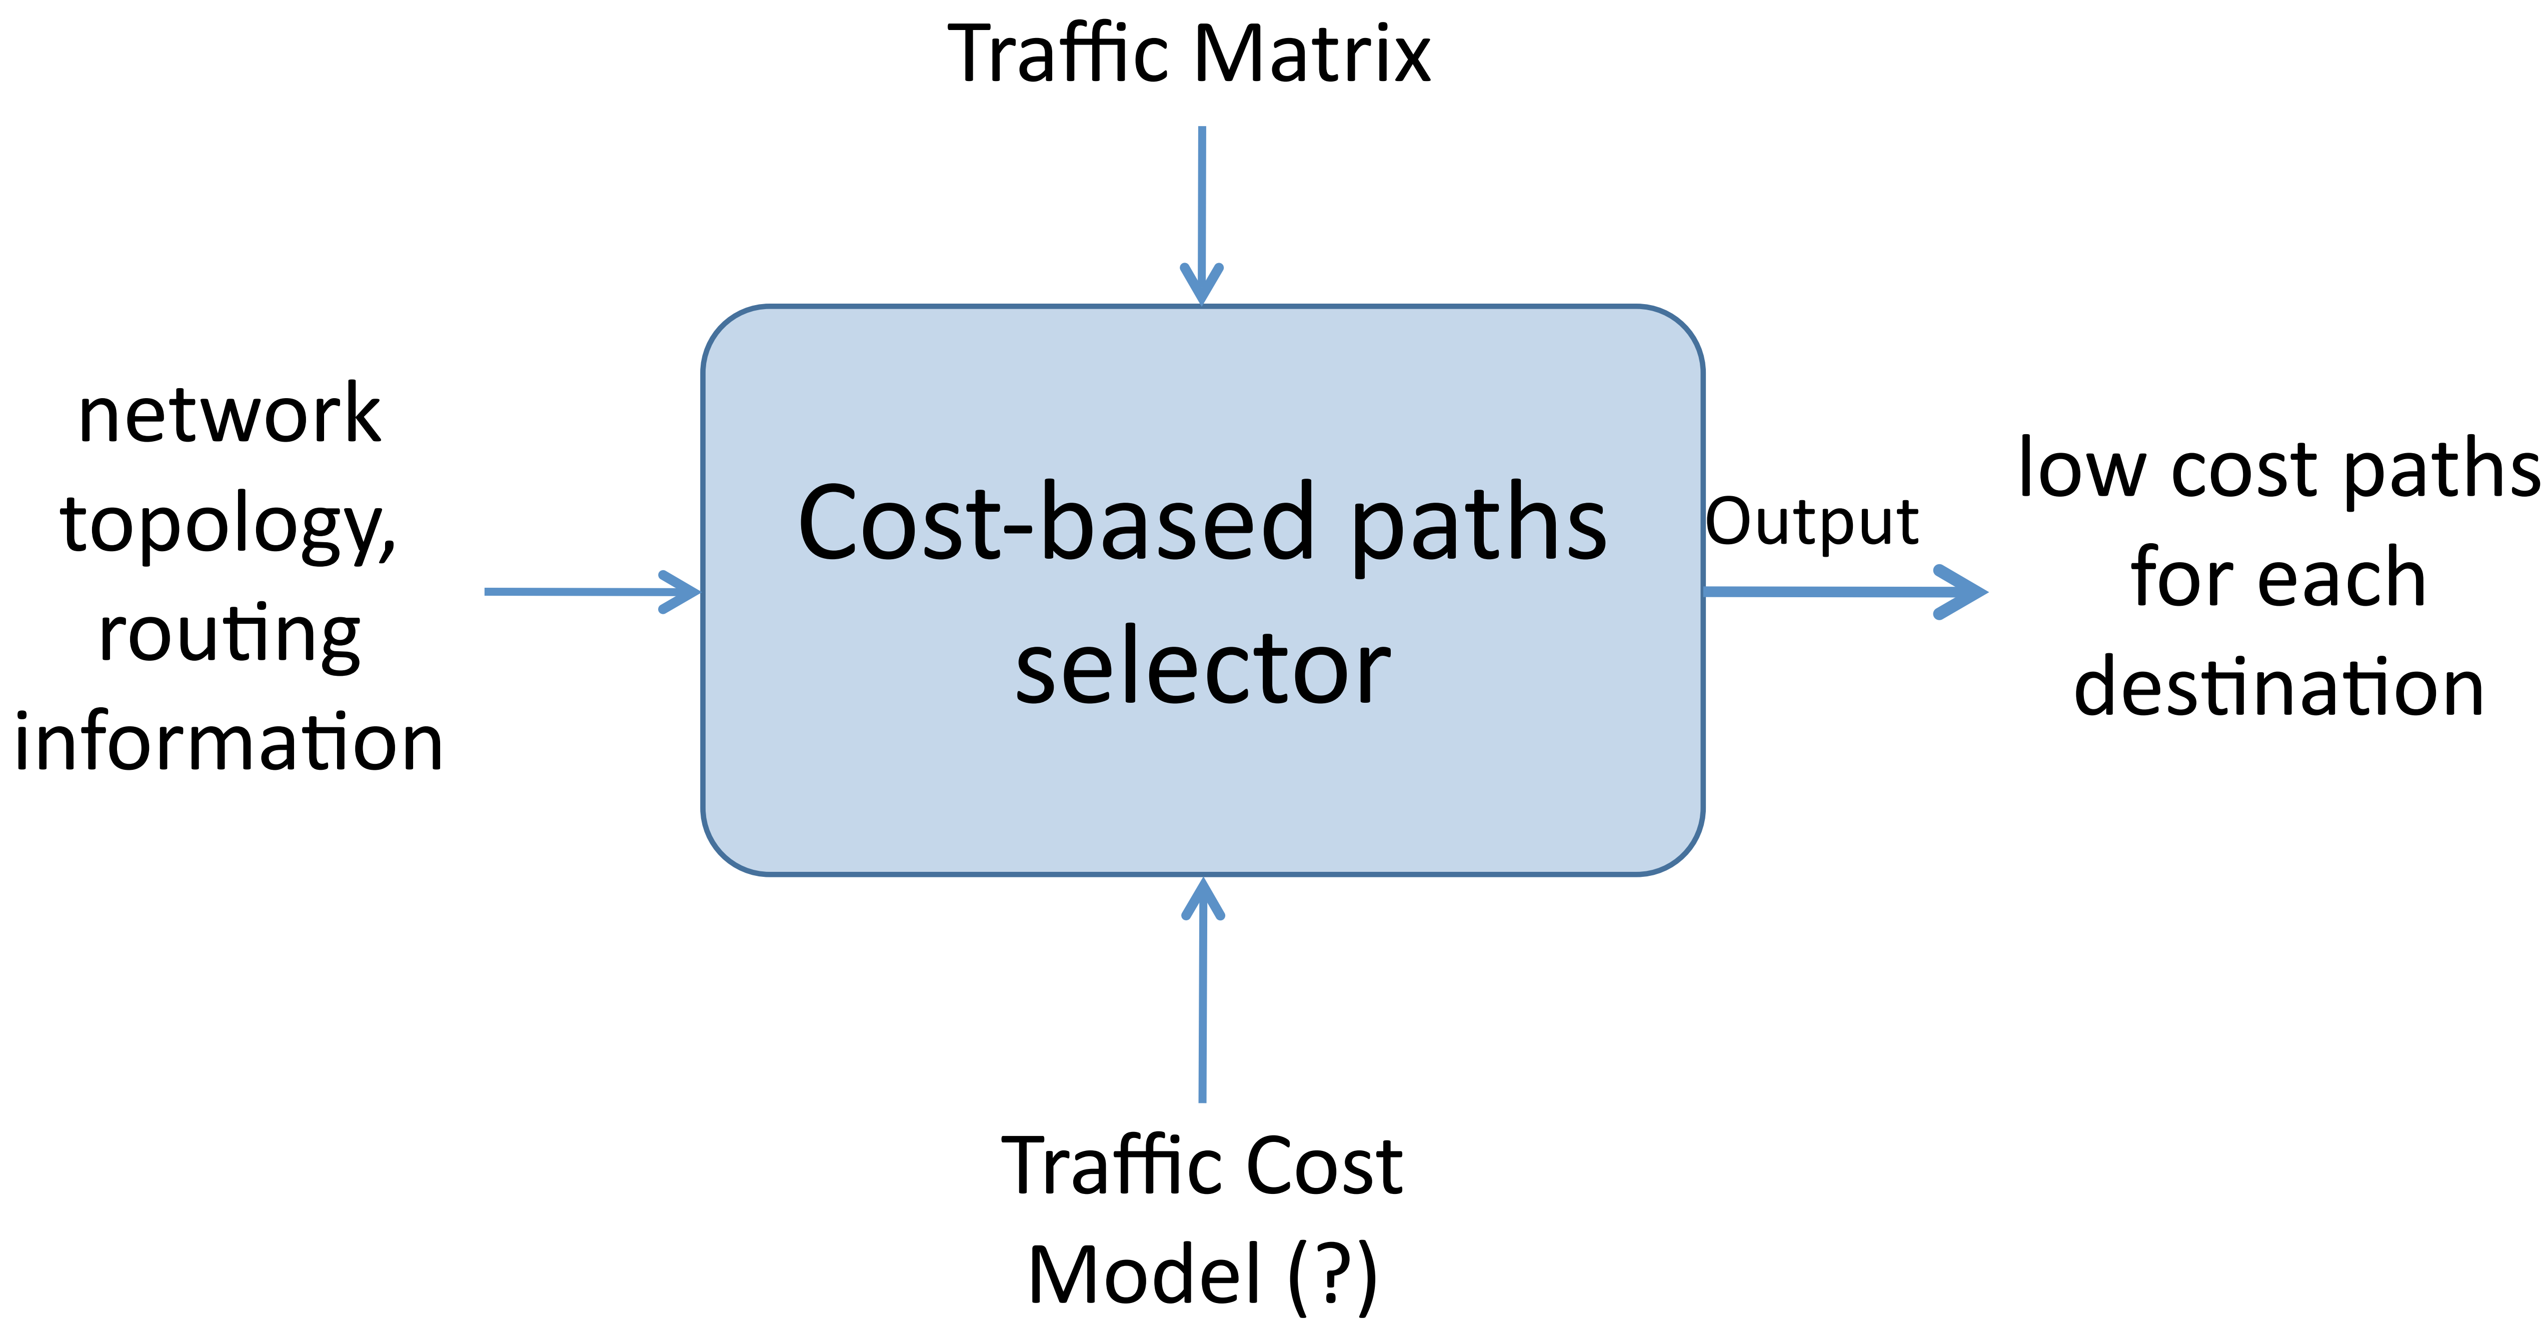

Need for a Traffic Cost Model

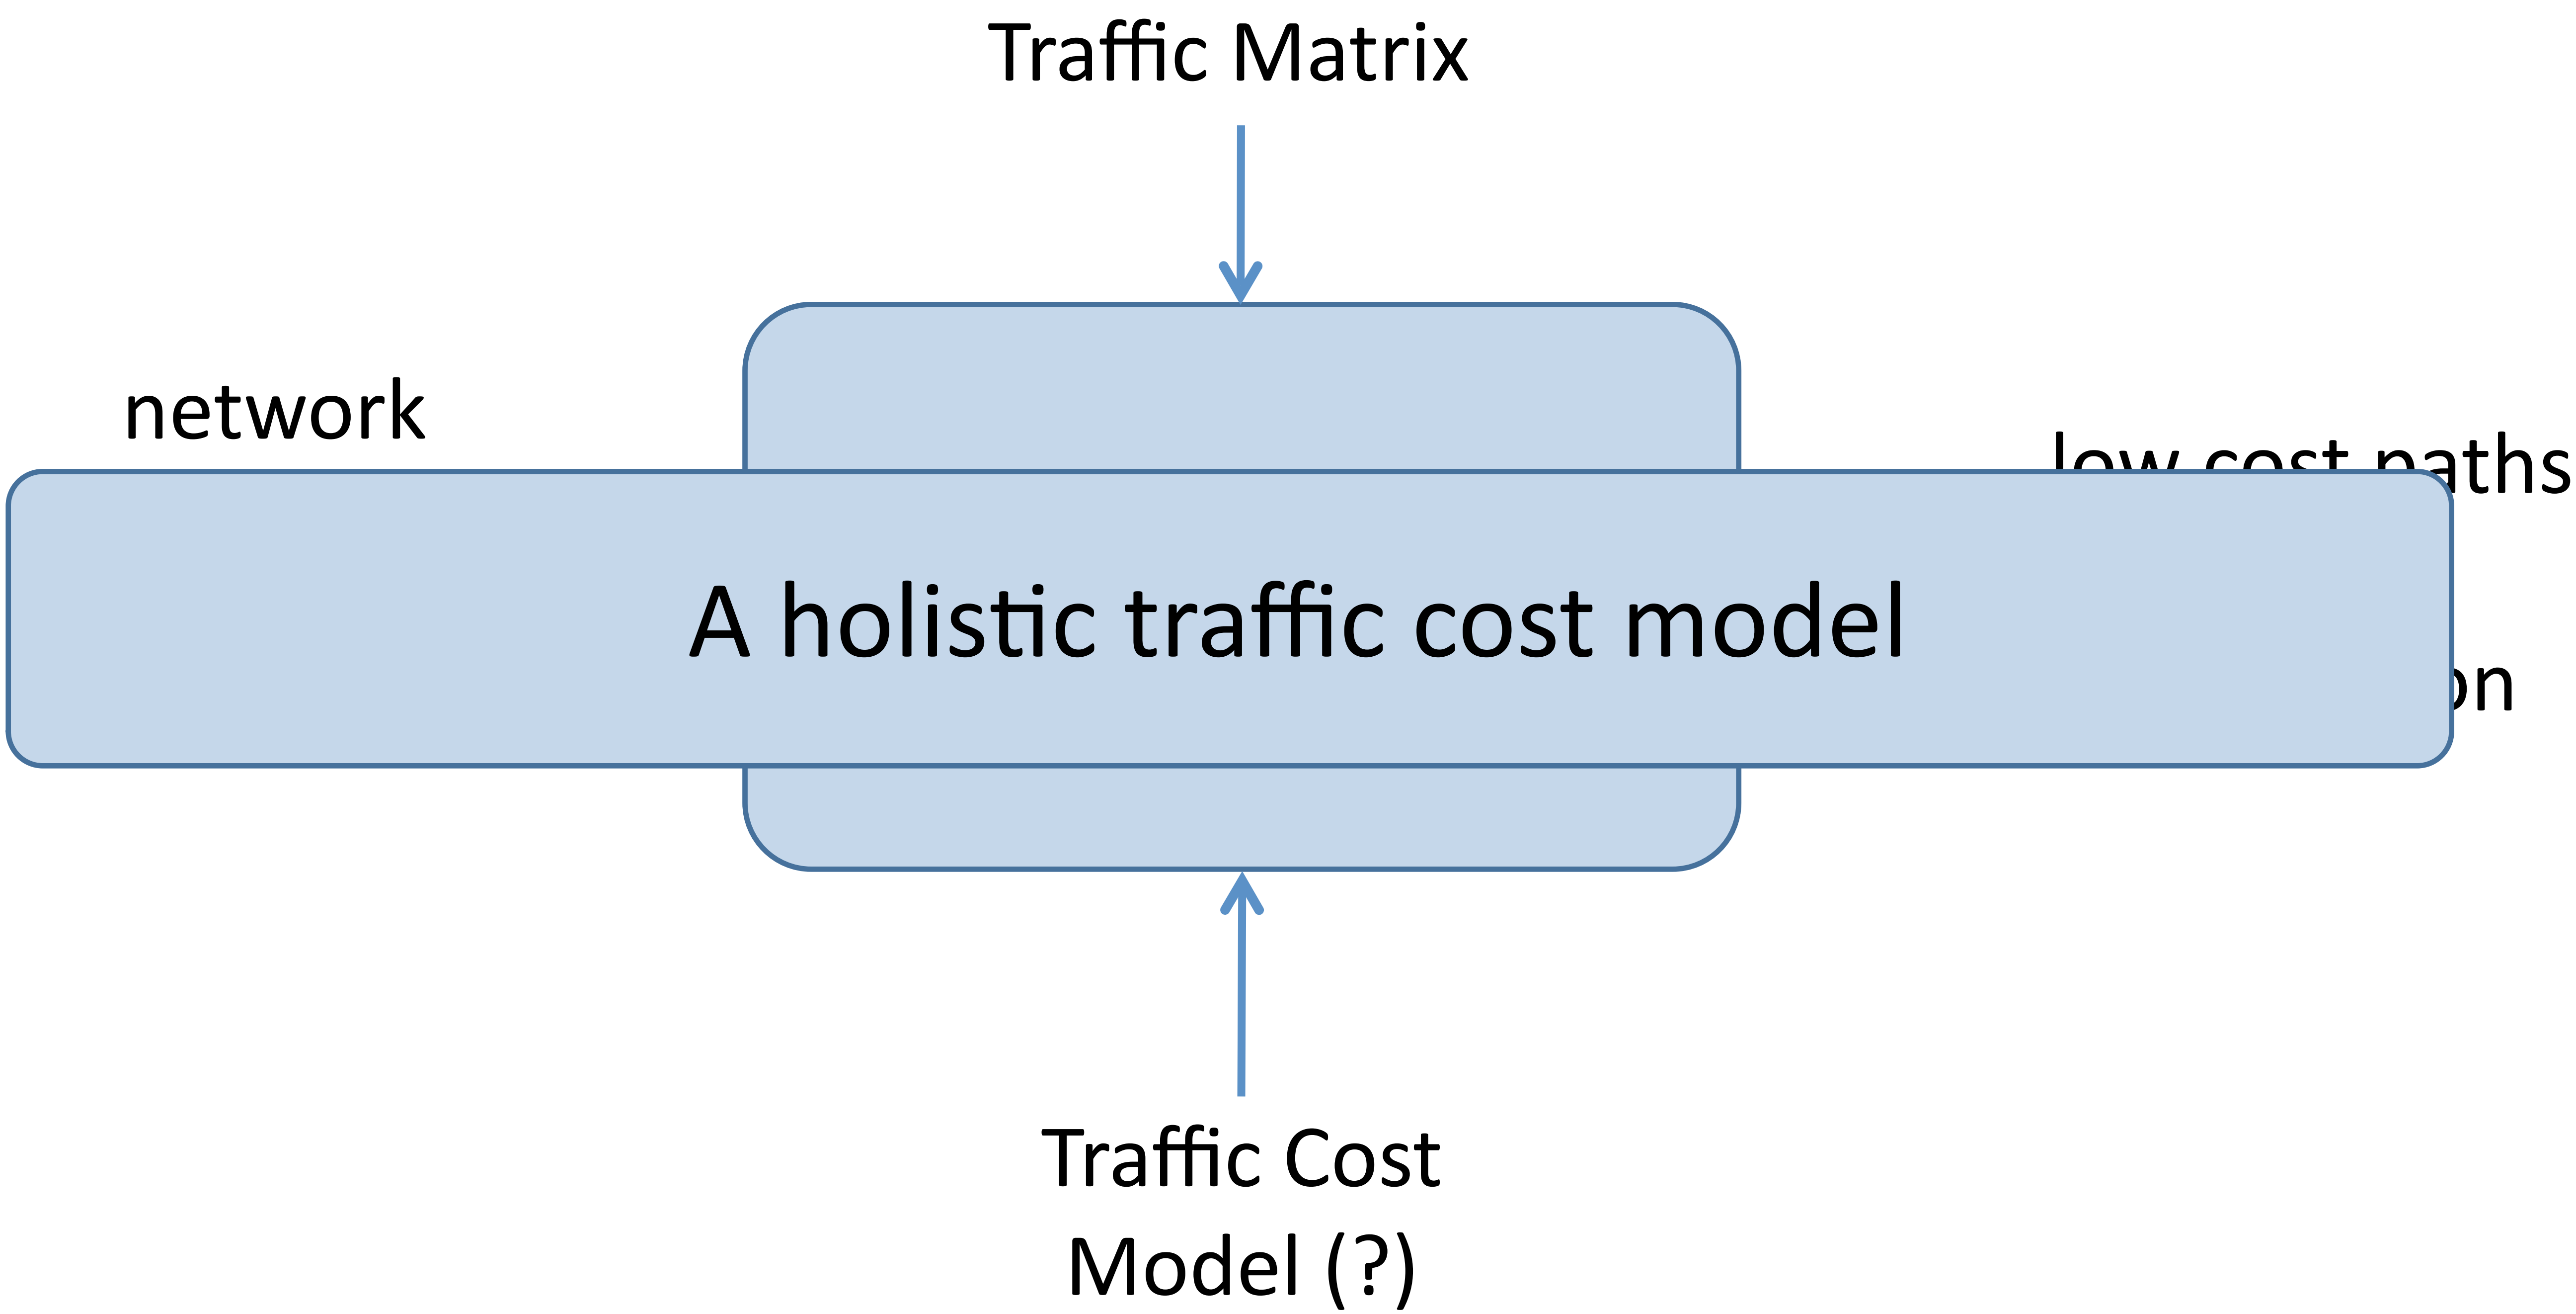

Costs for operating a network

- Traffic costs
  - Paying for transit, port costs, cost for laying fiber
- Operational costs
  - Paying salaries to employees
- Equipment and maintainenance costs
  - Buying networking gear, service fees to vendors
- Miscellaneous costs
  - IT related, real-estate, etc

Costs for operating a network

- Traffic costs
  – Paying for transit, port costs, cost for laying fiber

Goal: A simple but still useful cost model

  – Buying networking gear, service fees to vendors
- Miscellaneous costs
  – IT related, real-estate, etc

Applications

- Min-cost routing: Optimal routing of ingress-egress flows to minimize cost
- Peering Location selection: Which location to establish peering with a neighbor?
- Peering evaluation: What is the “value” of a peering link?

A Traffic Cost Model

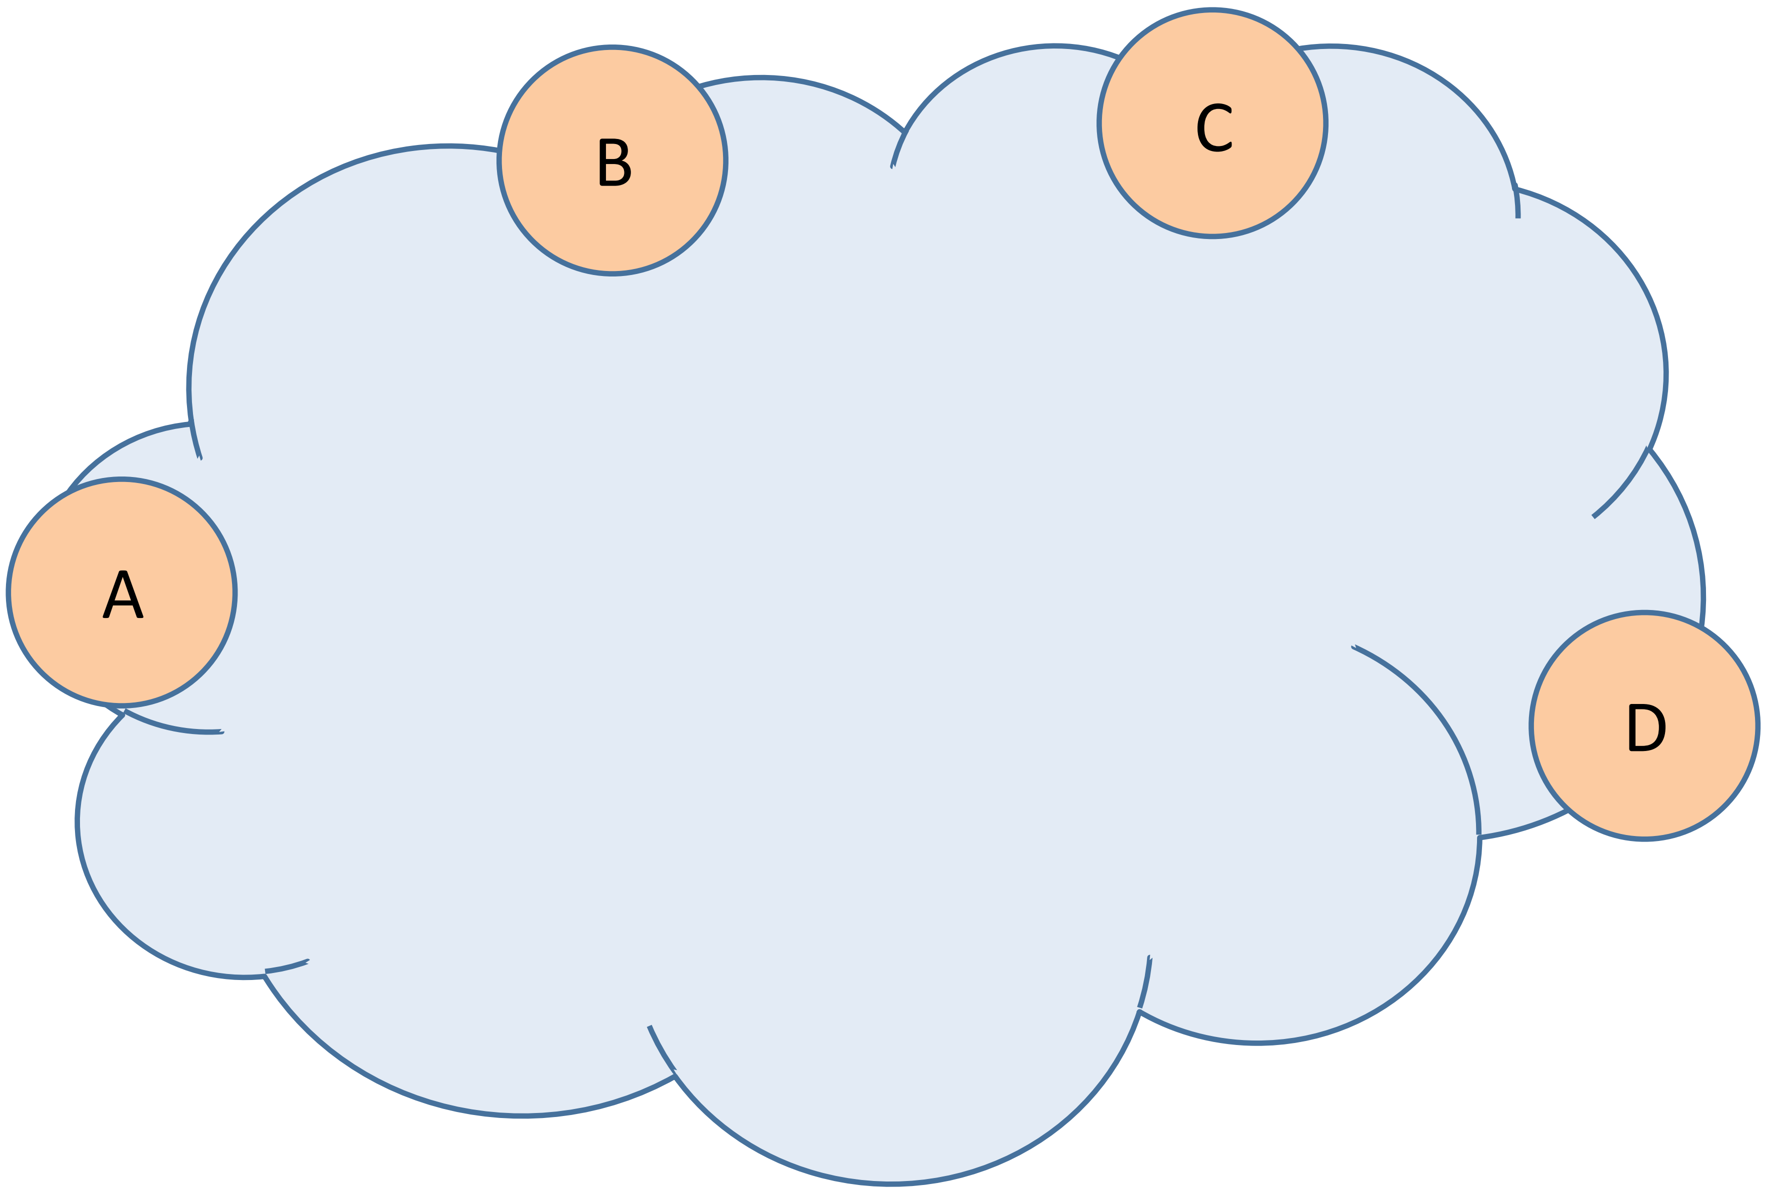

A Traffic Cost Model

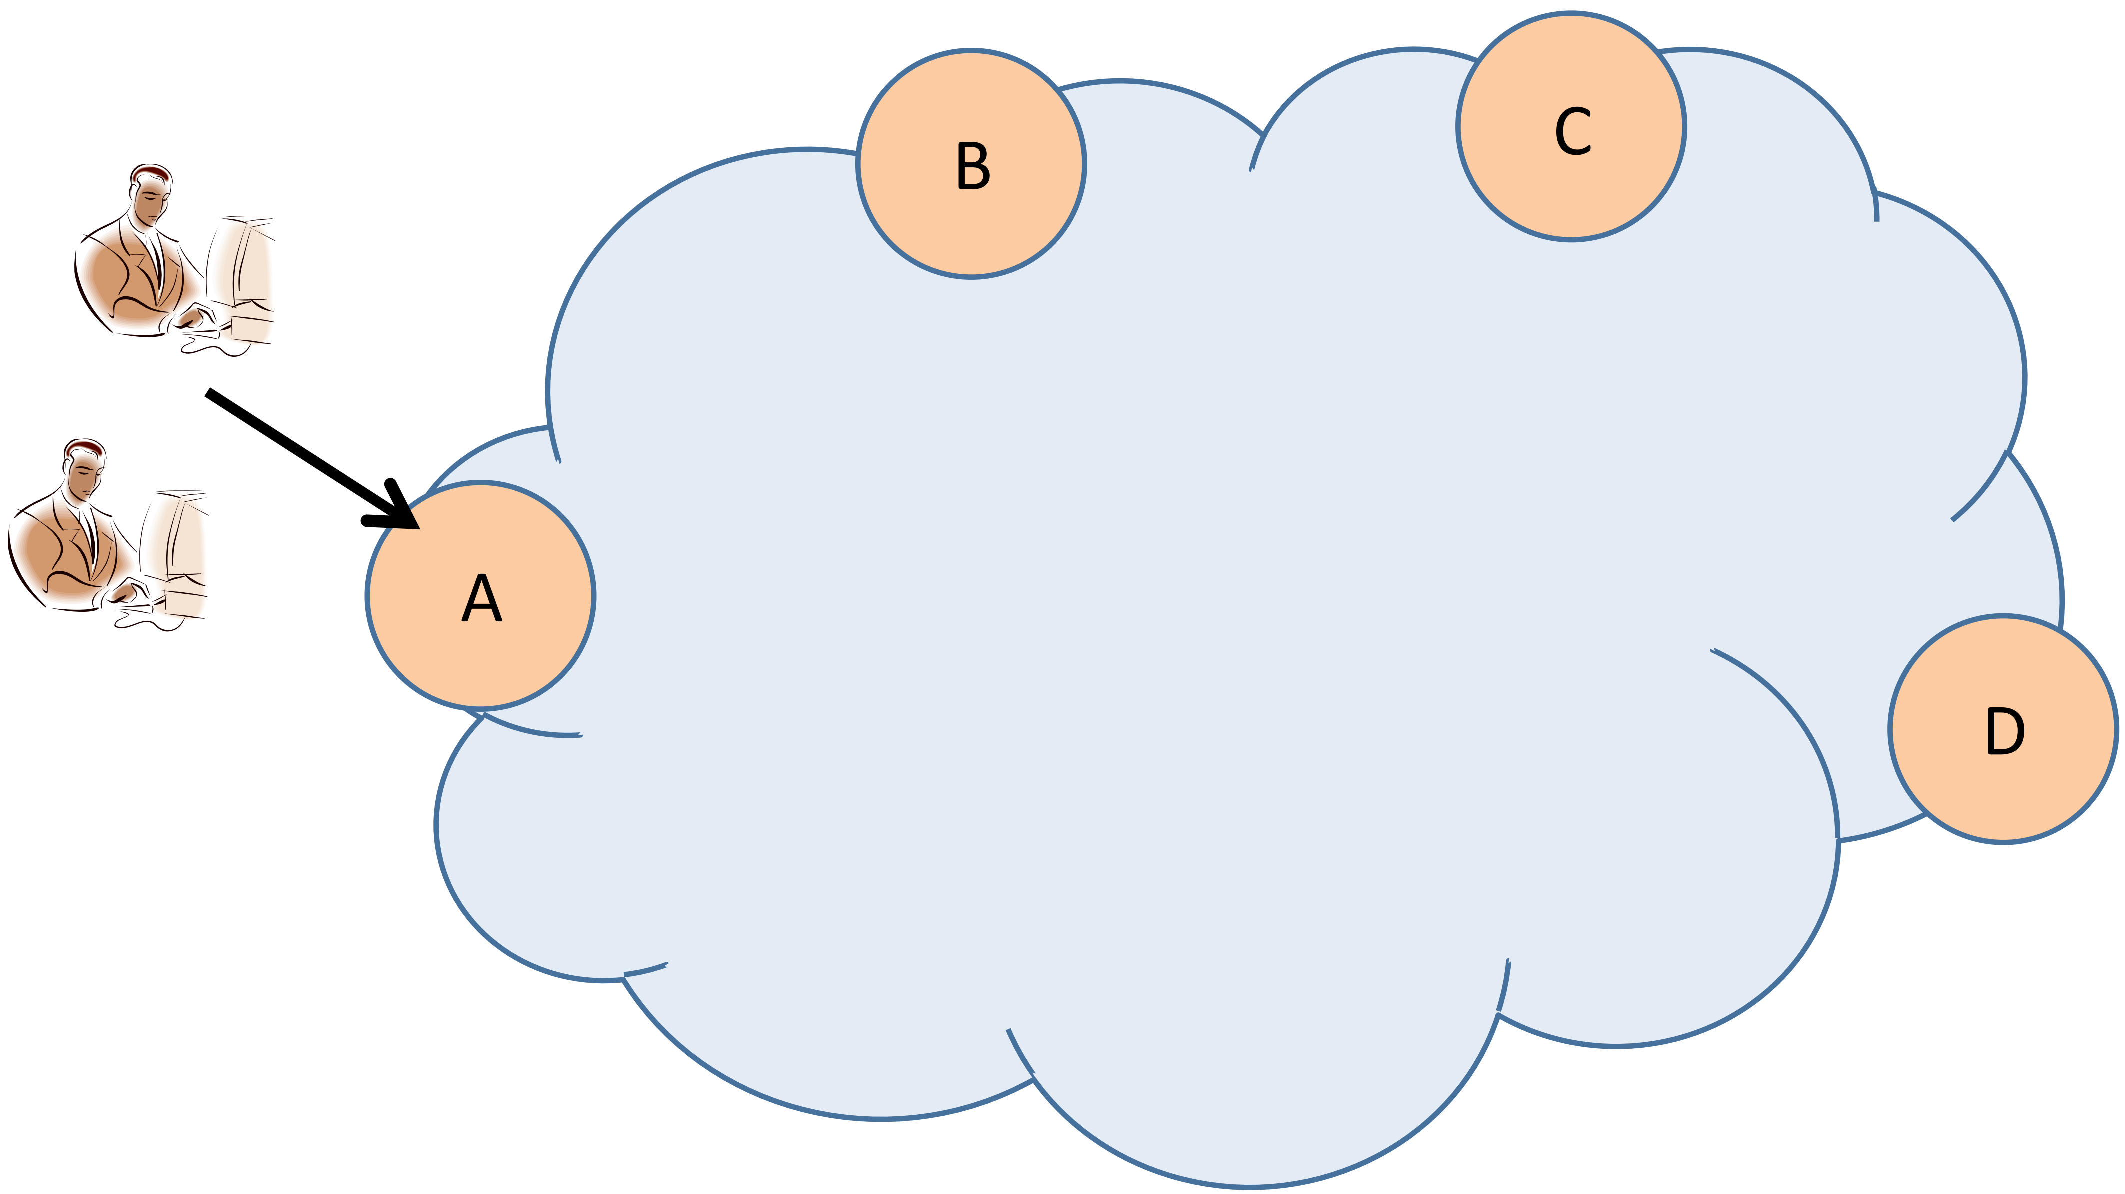

A Traffic Cost Model

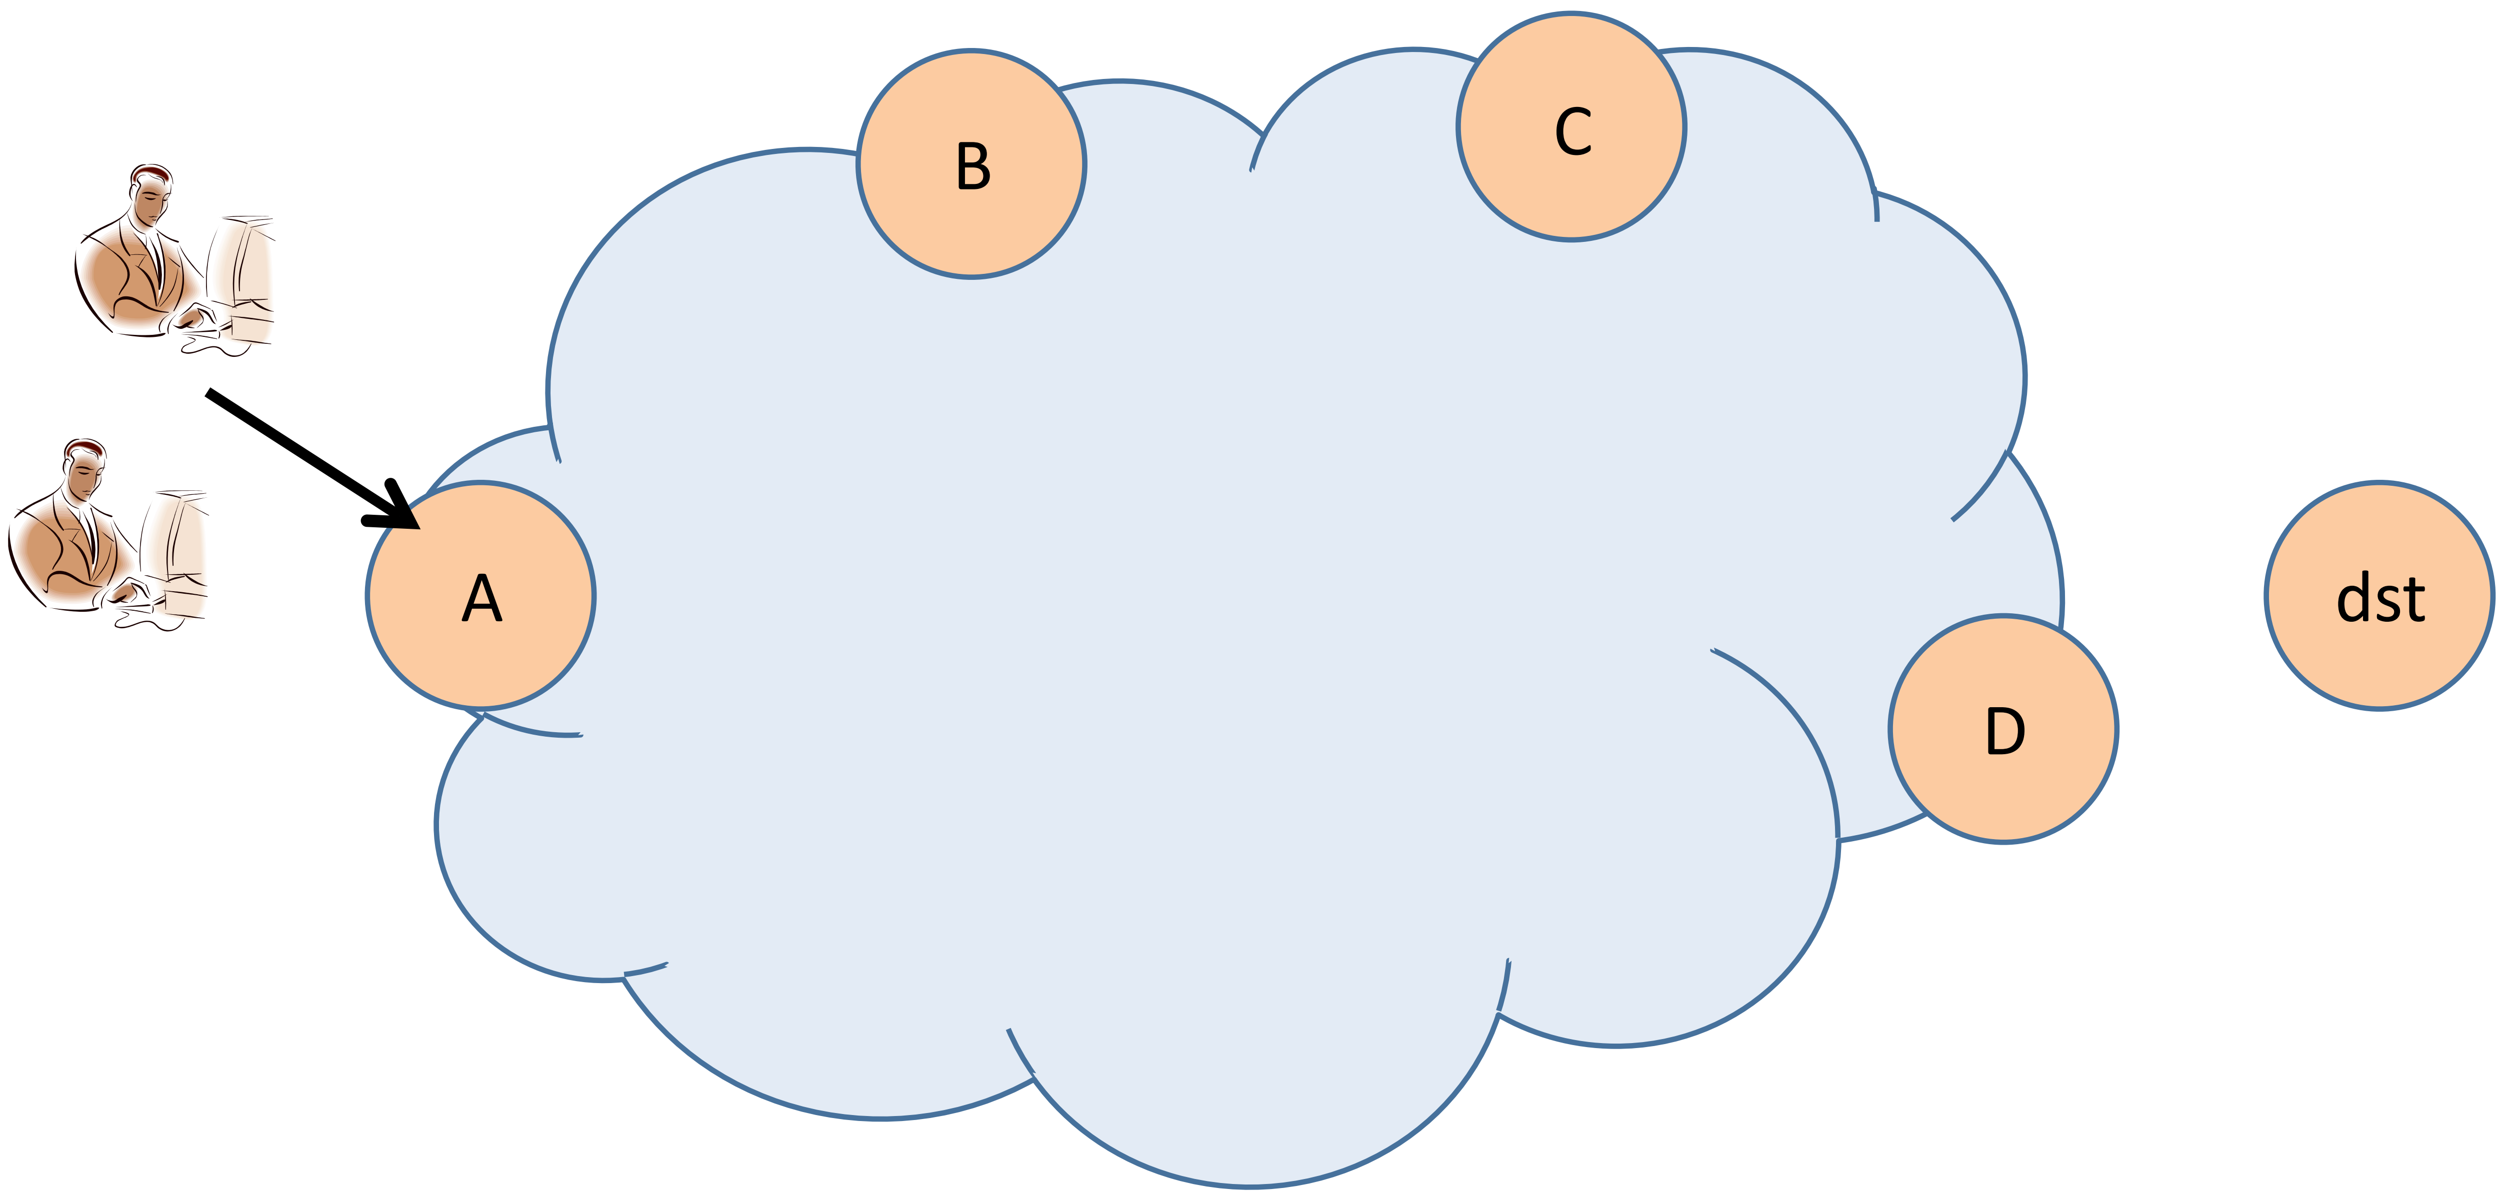

A Traffic Cost Model

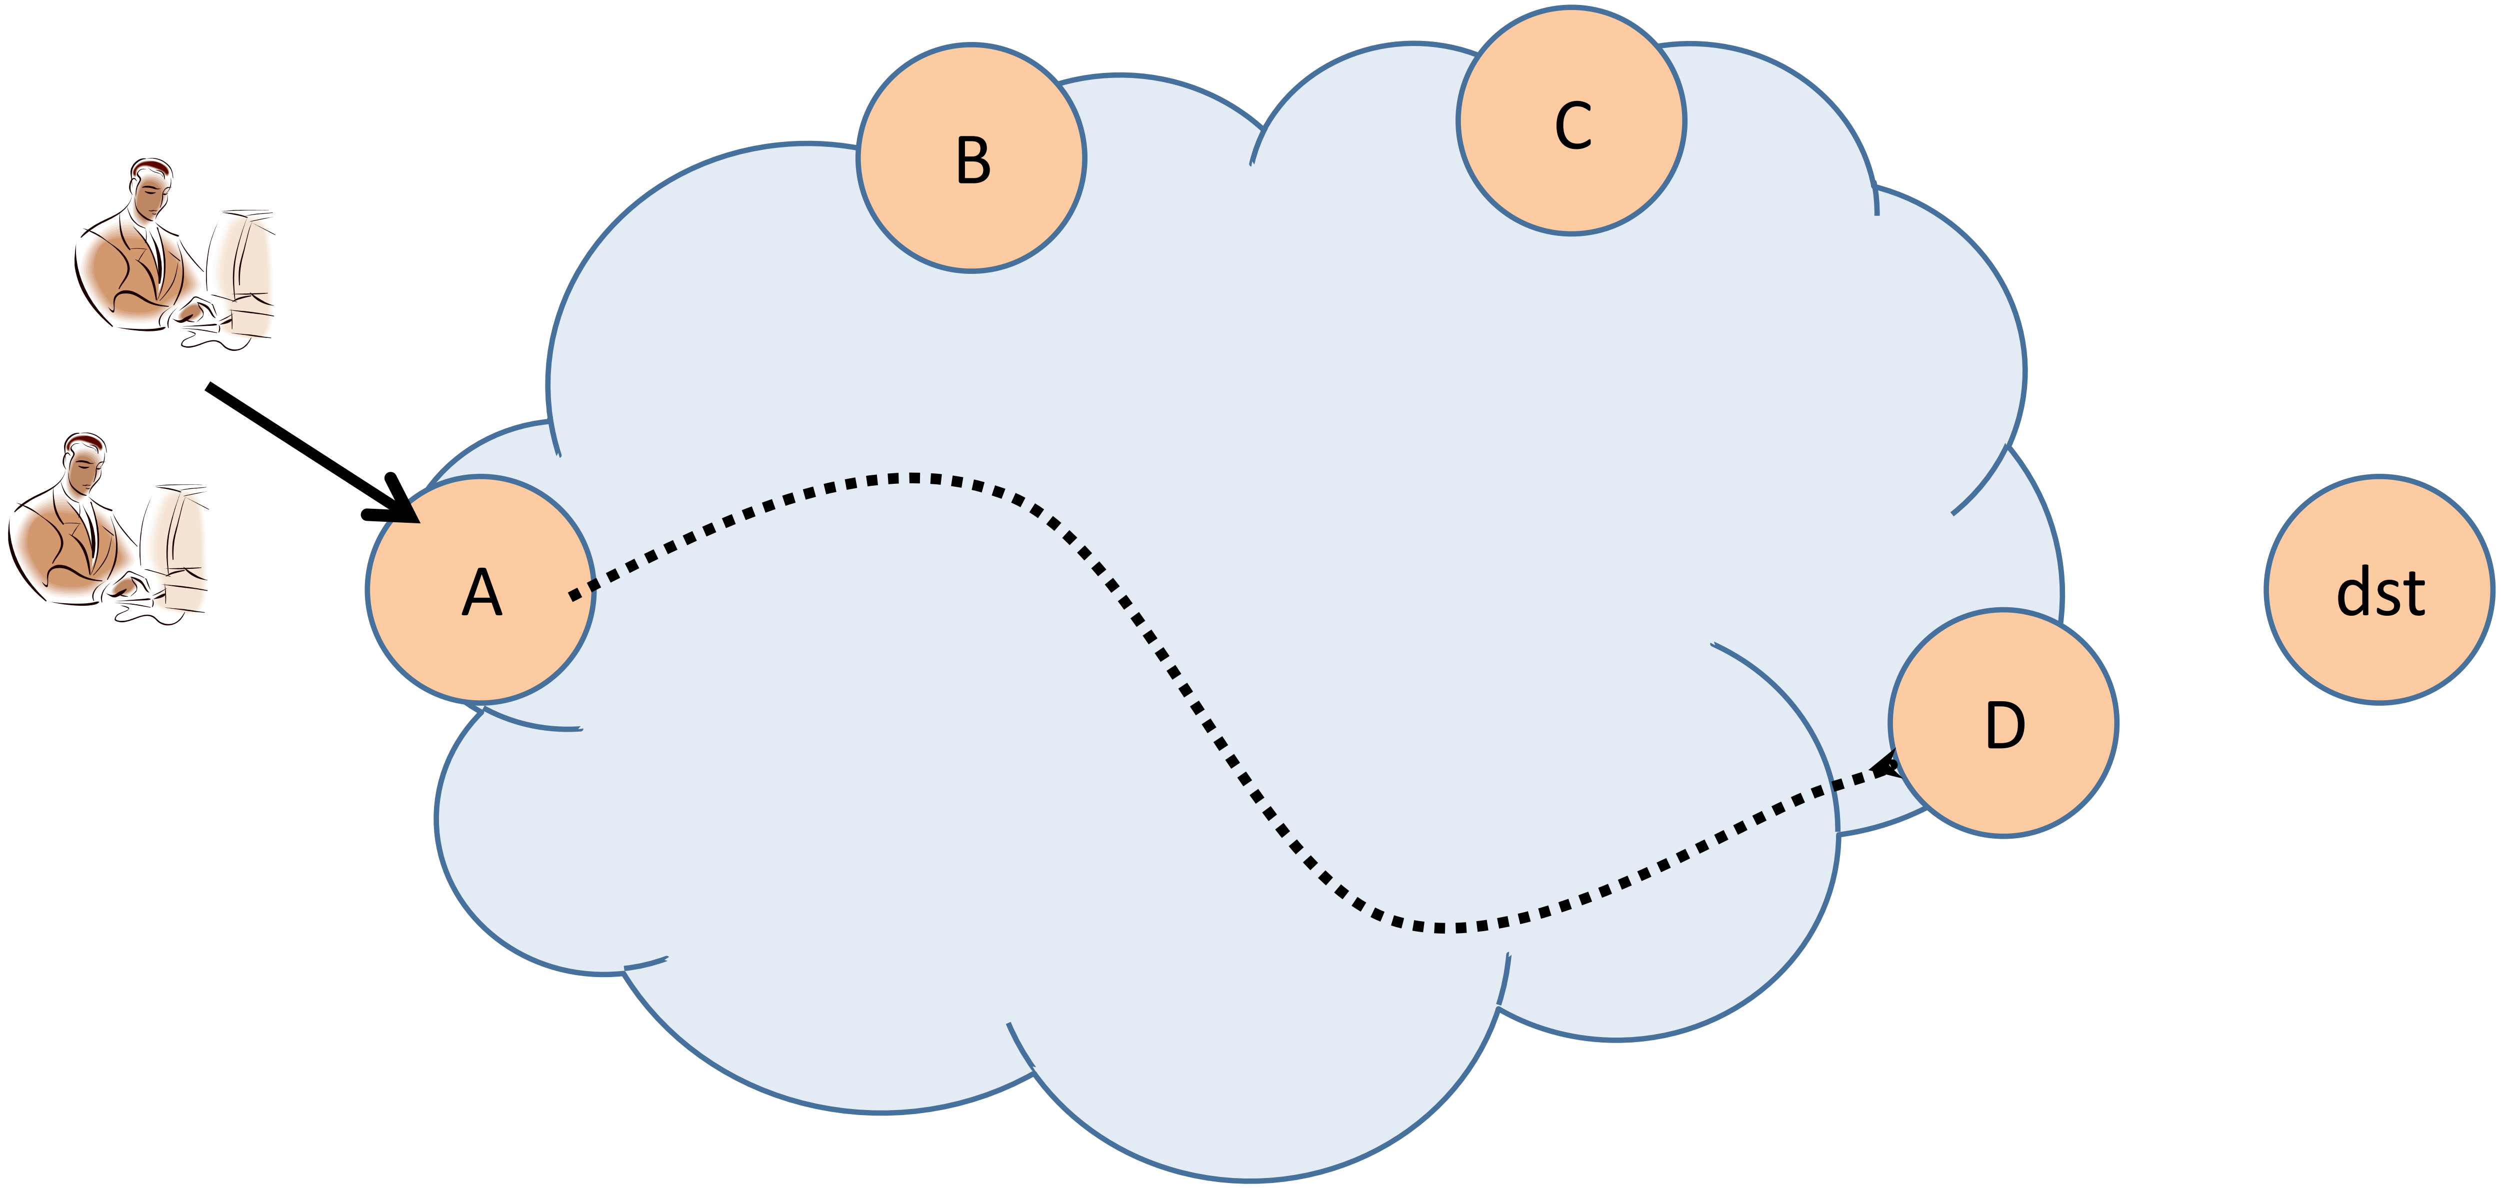

A Traffic Cost Model

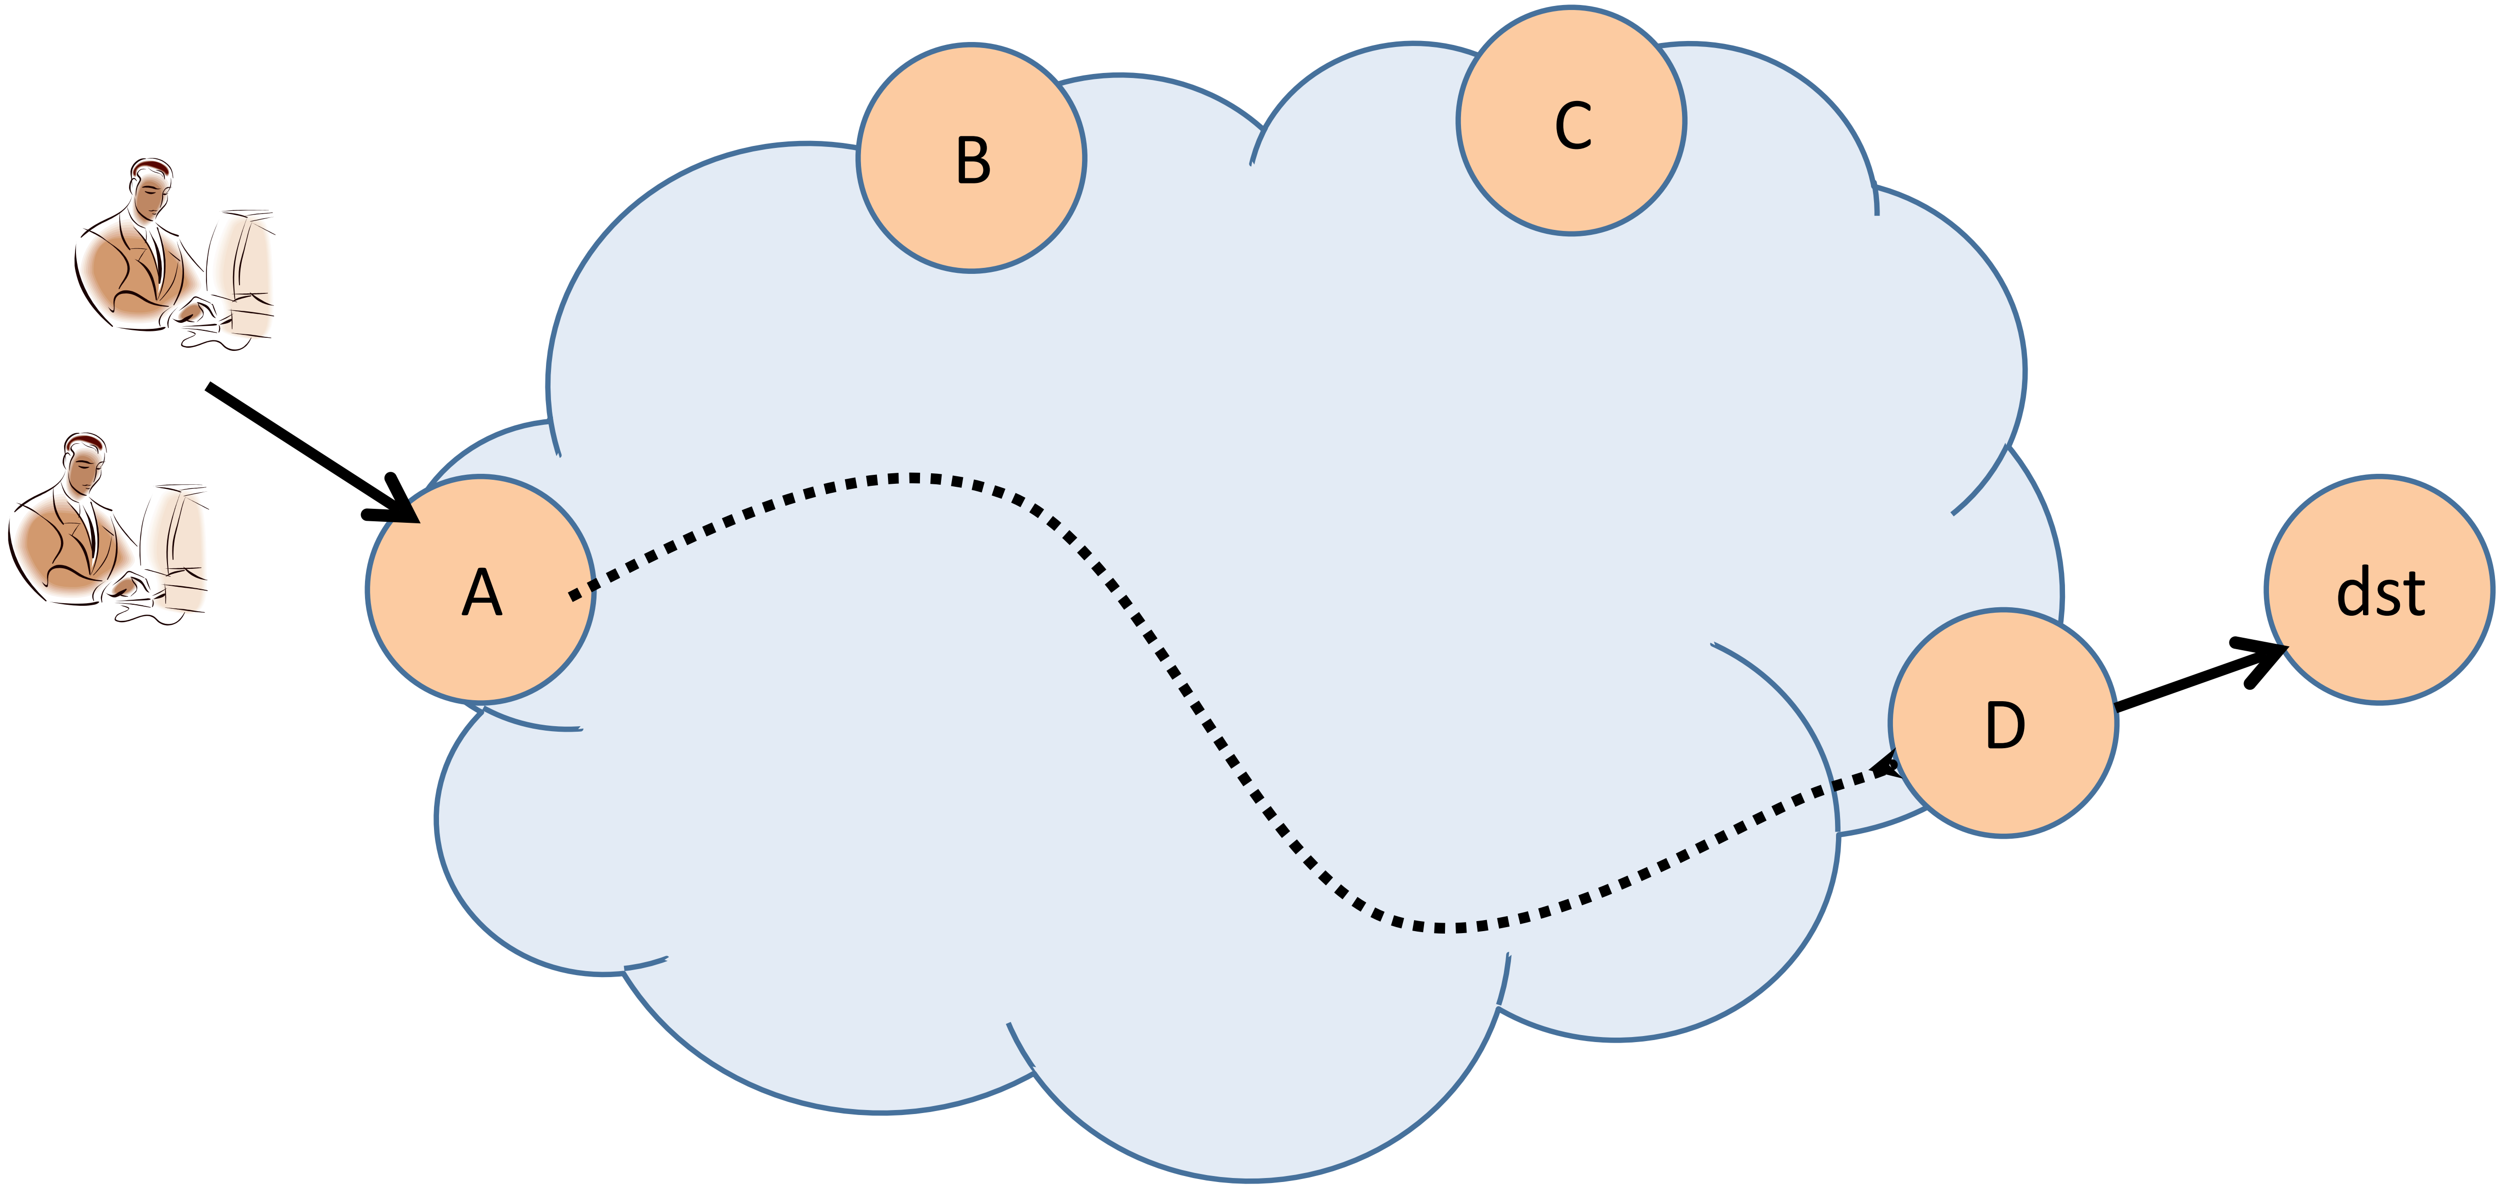

A Traffic Cost Model

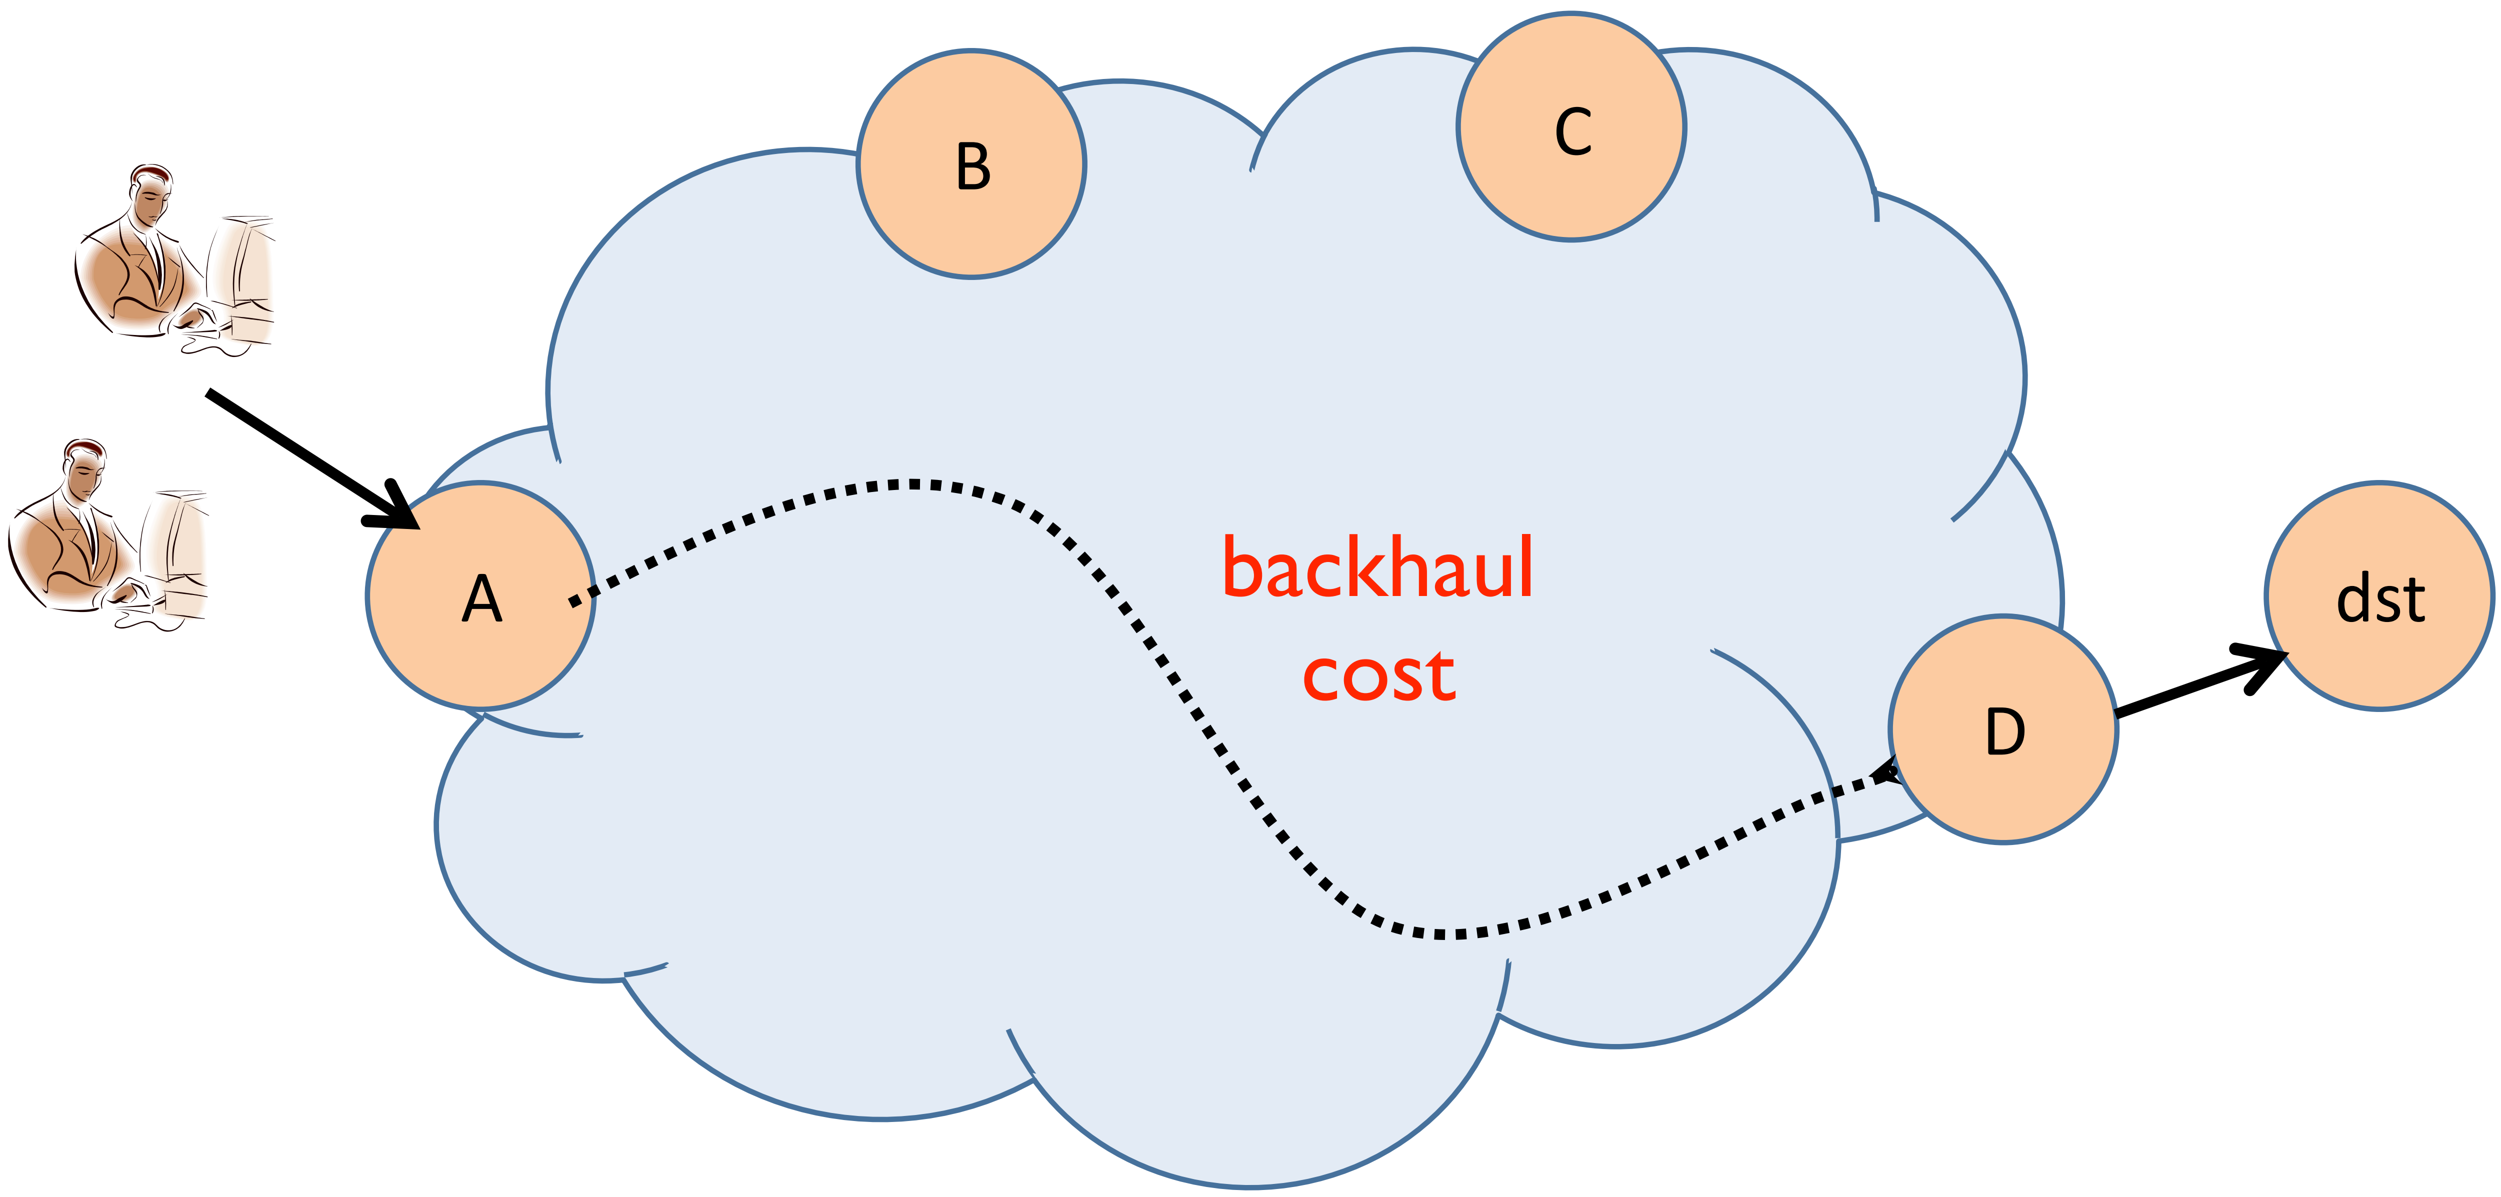

A Traffic Cost Model

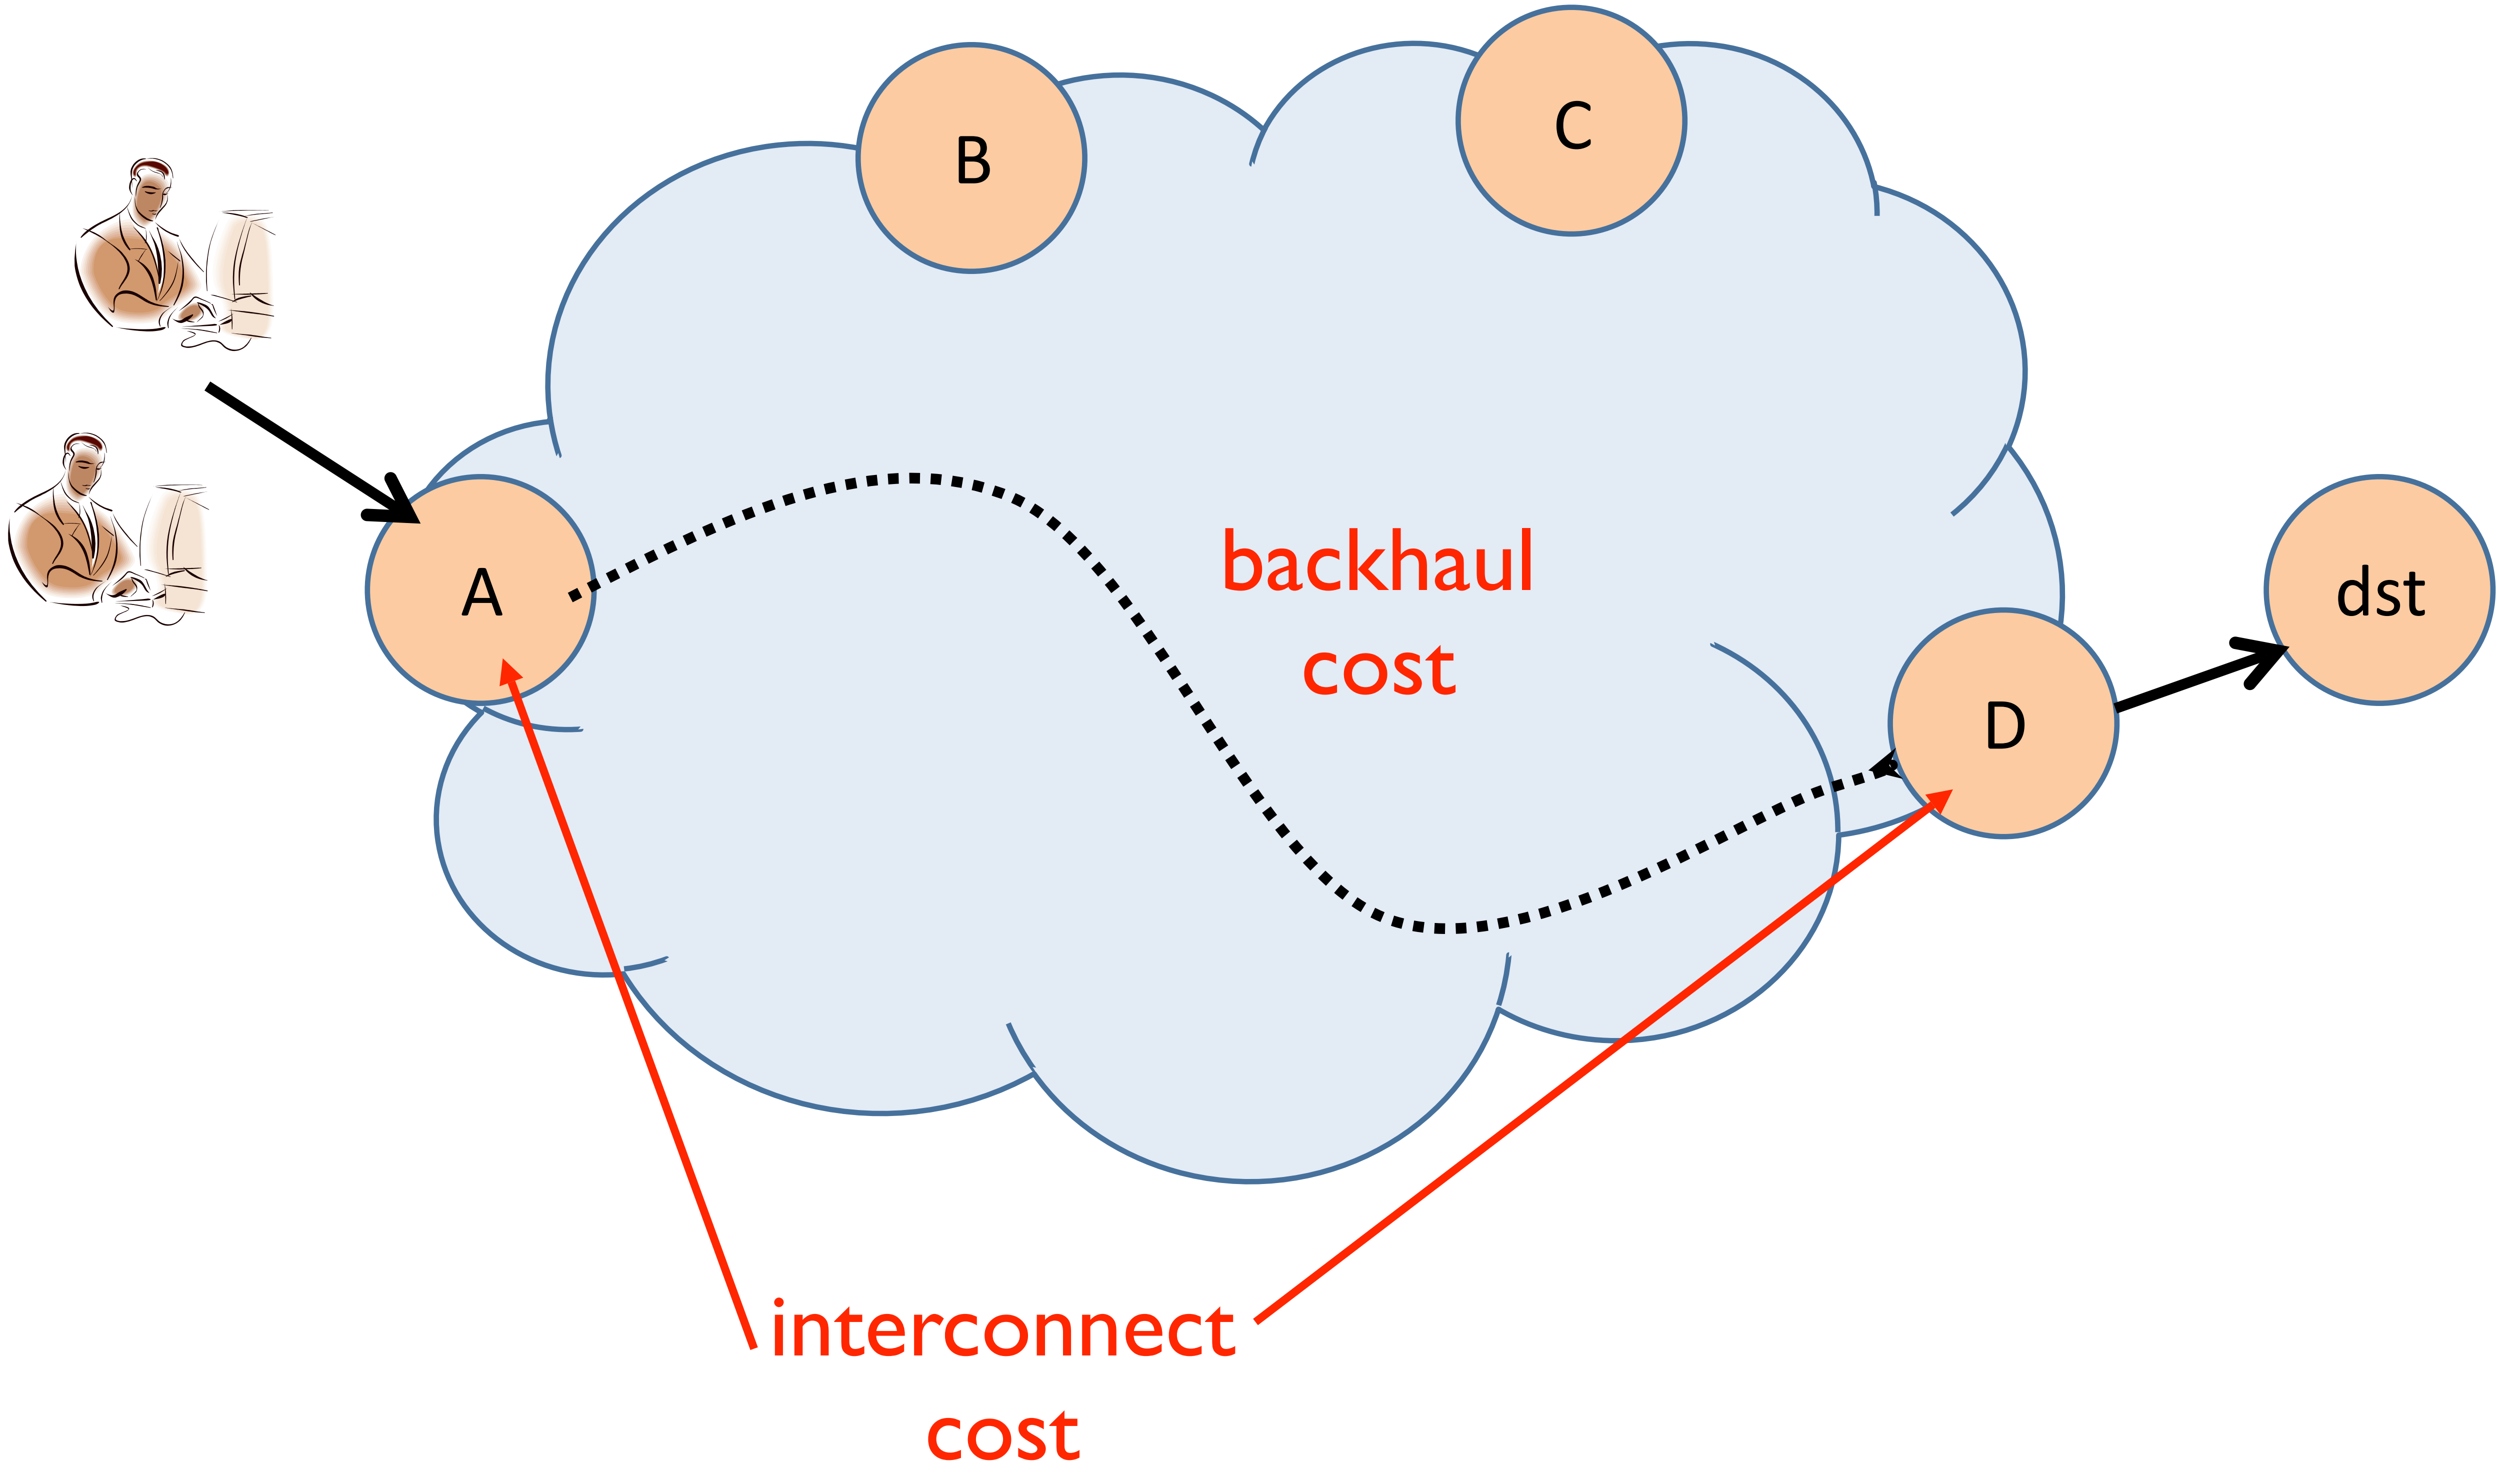

A Traffic Cost Model

A Traffic Cost Model

$C_N = C_F + C_U$

Total Cost = Fixed Cost + Usage-based Cost

A Traffic Cost Model

$C_N = C_F + C_U$

Total Cost = Fixed Cost + Usage-based Cost

$$C_F = \sum_{p1,p2} c_b(p1,p2) \quad +$$

Fixed Backhaul cost for all (PoP, PoP) pairs

A Traffic Cost Model

$C_N = C_F + C_U$

Total Cost = Fixed Cost + Usage-based Cost

$$C_F = \sum_{p1,p2} c_b(p1,p2) \quad + \quad \sum_{a,p} c_i\,(a,p)$$

Fixed Backhaul cost for all (PoP, PoP) pairs

Fixed Interconnect cost for all (AS, PoP) pairs

A Traffic Cost Model

$C_N = C_F + C_U$

Total Cost = Fixed Cost + Usage-based Cost

$$C_F = \sum_{p1,p2} c_b(p1,p2) \quad + \quad \sum_{a,p} c_i\,(a,p)$$

Fixed Backhaul cost for all (PoP, PoP) pairs

Fixed Interconnect cost for all (AS, PoP) pairs

A Traffic Cost Model

$C_N = C_F + C_U$

Total Cost = Fixed Cost + Usage-based Cost

$C_F = \sum_{p1,p2} c_b(p1,p2) \quad + \quad \sum_{a,p} c_i(a,p)$

Fixed Backhaul cost for all (PoP, PoP) pairs

Fixed Interconnect cost for all (AS, PoP) pairs

A Traffic Cost Model

$C_N = C_F + C_U$

Total Cost = Fixed Cost + Usage-based Cost

$C_F = \sum_{p1,p2} c_b(p1,p2) \quad + \quad \sum_{a,p} c_i\,(a,p)$

Fixed Backhaul cost for all (PoP, PoP) pairs

Fixed Interconnect cost for all (AS, PoP) pairs

$C_U = \sum_{f} \; (c_i(f) \quad +$

Ingress Interconnect Cost

A Traffic Cost Model

$C_N = C_F + C_U$

Total Cost = Fixed Cost + Usage-based Cost

$C_F = \sum_{p1,p2} c_b(p1,p2) \quad + \quad \sum_{a,p} c_i(a,p)$

Fixed Backhaul cost for all (PoP, PoP) pairs

Fixed Interconnect cost for all (AS, PoP) pairs

$C_U = \sum_{f} (c_i(f) \quad + \quad c_b(f) \quad +$

Ingress Interconnect Cost

Backhaul Cost

A Traffic Cost Model

$C_N = C_F + C_U$

Total Cost = Fixed Cost + Usage-based Cost

$C_F = \sum_{p1,p2} c_b(p1,p2) \quad + \quad \sum_{a,p} c_i(a,p)$

Fixed Backhaul cost for all (PoP, PoP) pairs

Fixed Interconnect cost for all (AS, PoP) pairs

$C_U = \sum_f \; (c_i(f) \quad + \quad c_b(f) \quad + \quad c_e(f))$

Ingress Interconnect Cost

Backhaul Cost

Egress Interconnect Cost

A Traffic Cost Model

$C_N = C_F + C_U$

Total Cost = Fixed Cost + Usage-based Cost

$C_F = \sum_{p1,p2} c_b(p1,p2) \quad + \quad \sum_{a,p} c_i(a,p)$

Fixed Backhaul cost for all (PoP, PoP) pairs

Fixed Interconnect cost for all (AS, PoP) pairs

$C_U = \sum_{f} (c_i(f) \quad + \quad c_b(f) \quad + \quad c_e(f))$

Ingress Interconnect Cost

Backhaul Cost

Egress Interconnect Cost

A Traffic Cost Model

$C_N = C_F + C_U$

Total Cost = Fixed Cost + Usage-based Cost

$C_F = \sum_{p1,p2} c_b(p1,p2) \quad + \quad \sum_{a,p} c_i\,(a,p)$

Fixed Backhaul cost for all (PoP, PoP) pairs

Fixed Interconnect cost for all (AS, PoP) pairs

$C_U = \sum_{f} \; (c_i(f) \quad + \quad c_b(f) \quad + \quad c_e(f))$

Ingress-egress flows

Ingress Interconnect Cost

Backhaul Cost

Egress Interconnect Cost

Cost Optimization

- We focus on optimizing traffic-dependent costs
- Requires an operator to determine the cost associated with each ingress-egress flow
- Interconnect costs based on 95th percentile of total volume at that interconnect: cost = per_bit_price * V_95
- Approach 1: Assume V_95 is linear function of average rate
  - Flow's contribution = per_bit_price * constant * flow_rate
- Approach 2: Use Shapley Value
  - Flow's contribution = Shapley value across all flows at that interconnect

Cost Optimization

- We focus on optimizing traffic-dependent costs
- Requires an operator to determine the cost associated with each ingress-egress flow
- Interconnect costs based on 95th percentile of total volume at that interconnect: cost = per_bit_price * V_95
- Approach 1: Assume V_95 is linear function of average rate
  - Flow's contribution = per_bit_price * constant * flow_rate

Approximation!

- Approach 2: Use Shapley Value
  - Flow's contribution = Shapley value across all flows at that interconnect

Cost Optimization

- We focus on optimizing traffic-dependent costs
- Requires an operator to determine the cost associated with each ingress-egress flow
- Interconnect costs based on 95th percentile of total volume at that interconnect: cost = per_bit_price * V_95
- Approach 1: Assume V_95 is linear function of average rate
  - Flow's contribution = per_bit_price * constant * flow_rate

Approximation!

- Approach 2: Use Shapley Value
  - Flow's contribution = Shapley value across all flows at that interconnect

Computationally expensive!

Application 1: Minimize Cost of Routing Traffic

- Objective: Minimize total cost of all ingress-egress flows
- Constraints: Backhaul, interconnect link capacities
- Knobs: egress (pop, AS) for each flow
- NP-hard to determine optimal routing! Use a greedy algorithm to approximate the optimal solution
- Iteratively, for each flow f in descending order of flow costs
  - For each PoP p, find the cheapest AS at p which has a route to f 's destination
  - Assign f to the cheapest egress (PoP, AS) pair

Cost Optimization

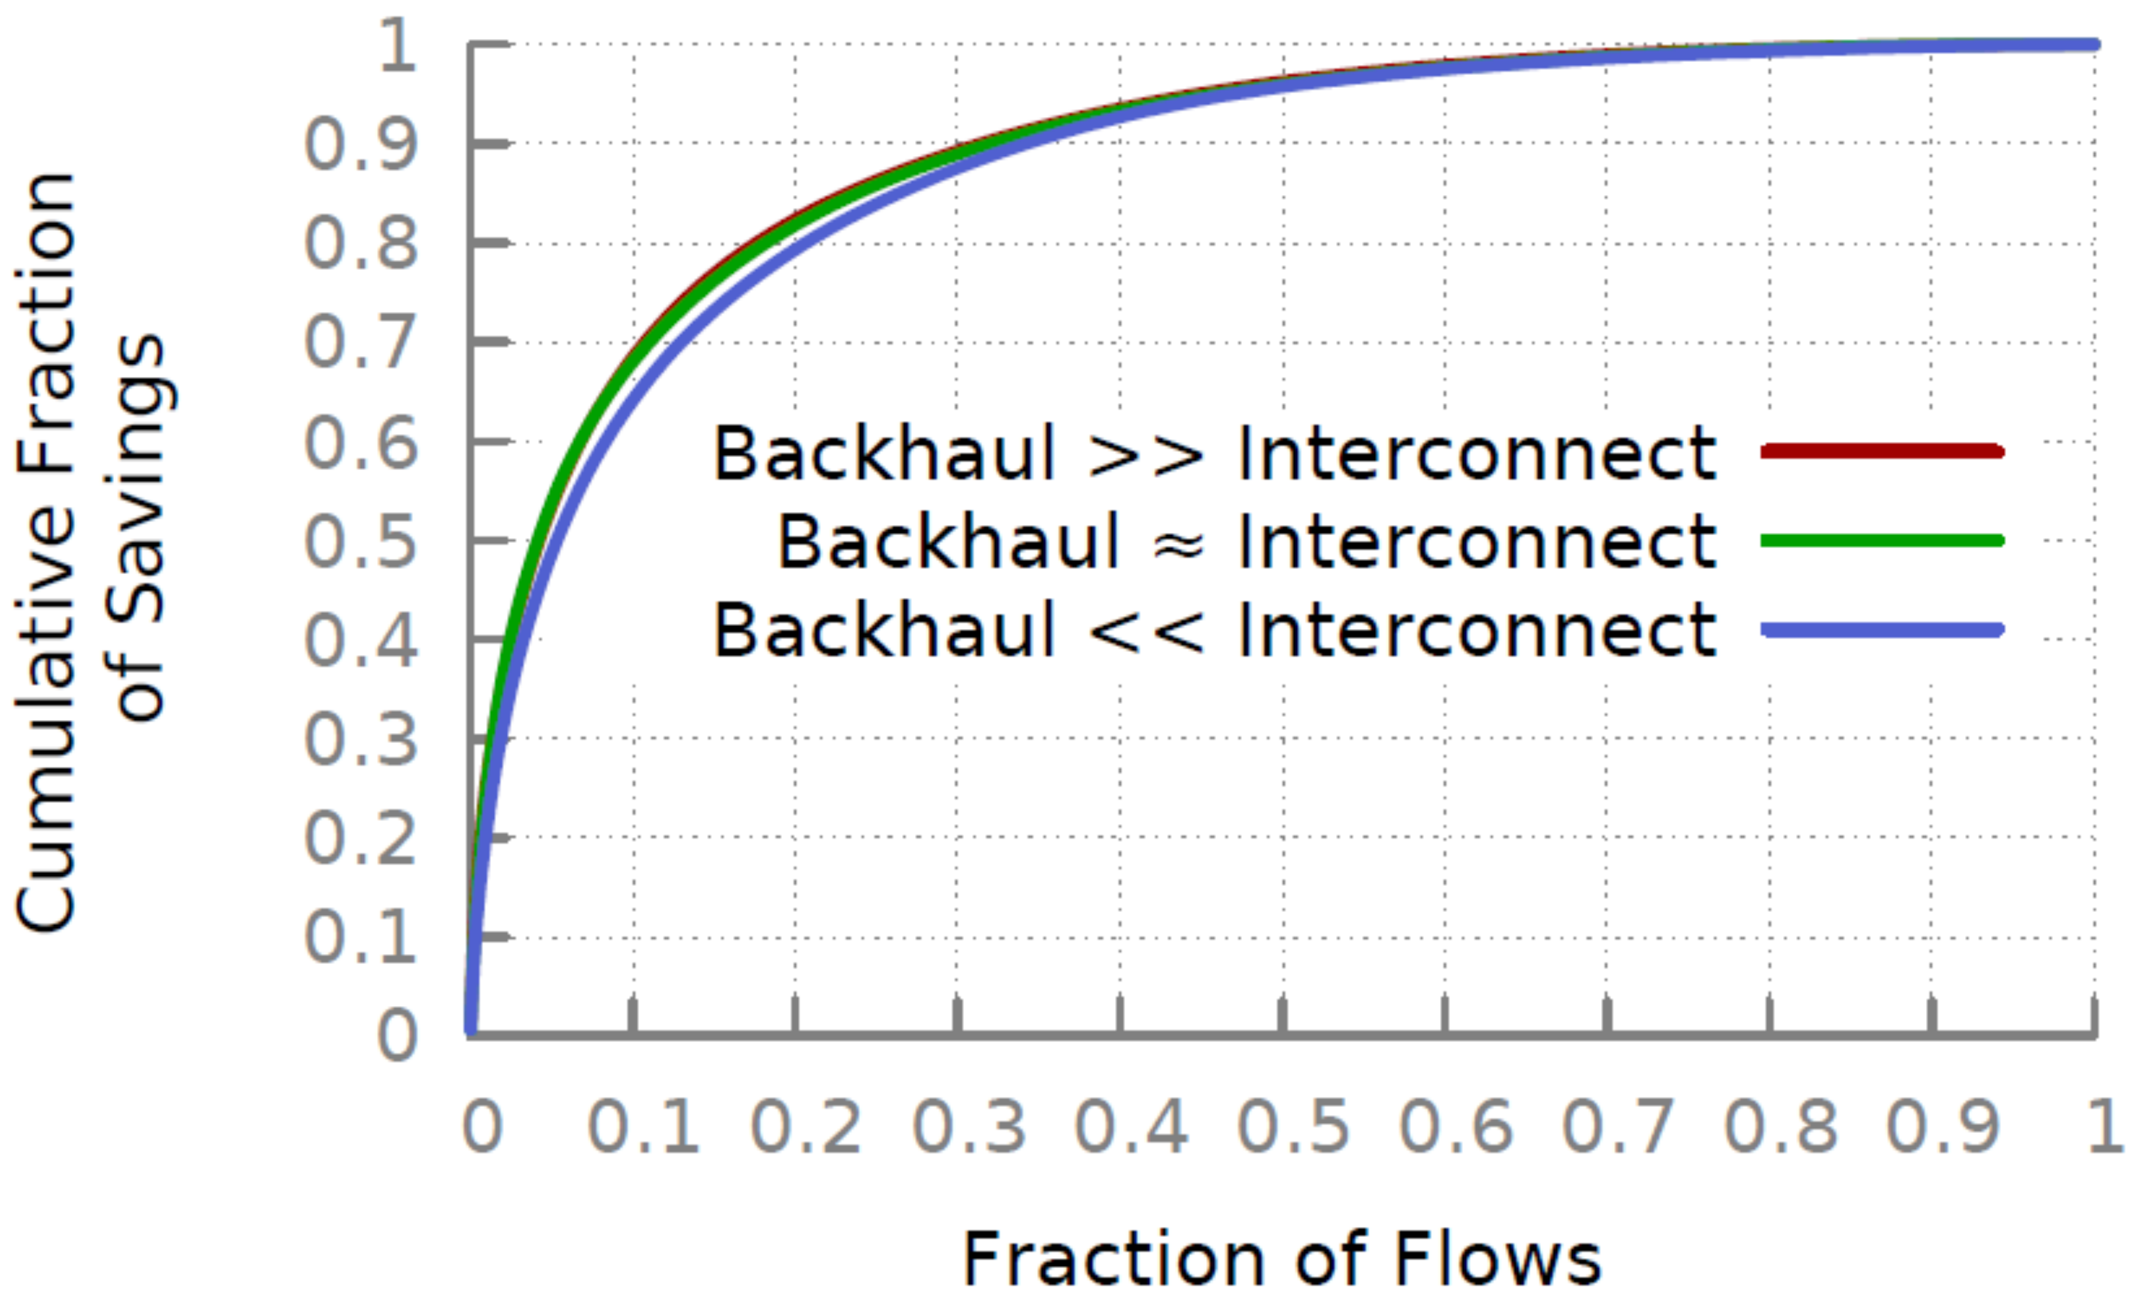

Cost Optimization

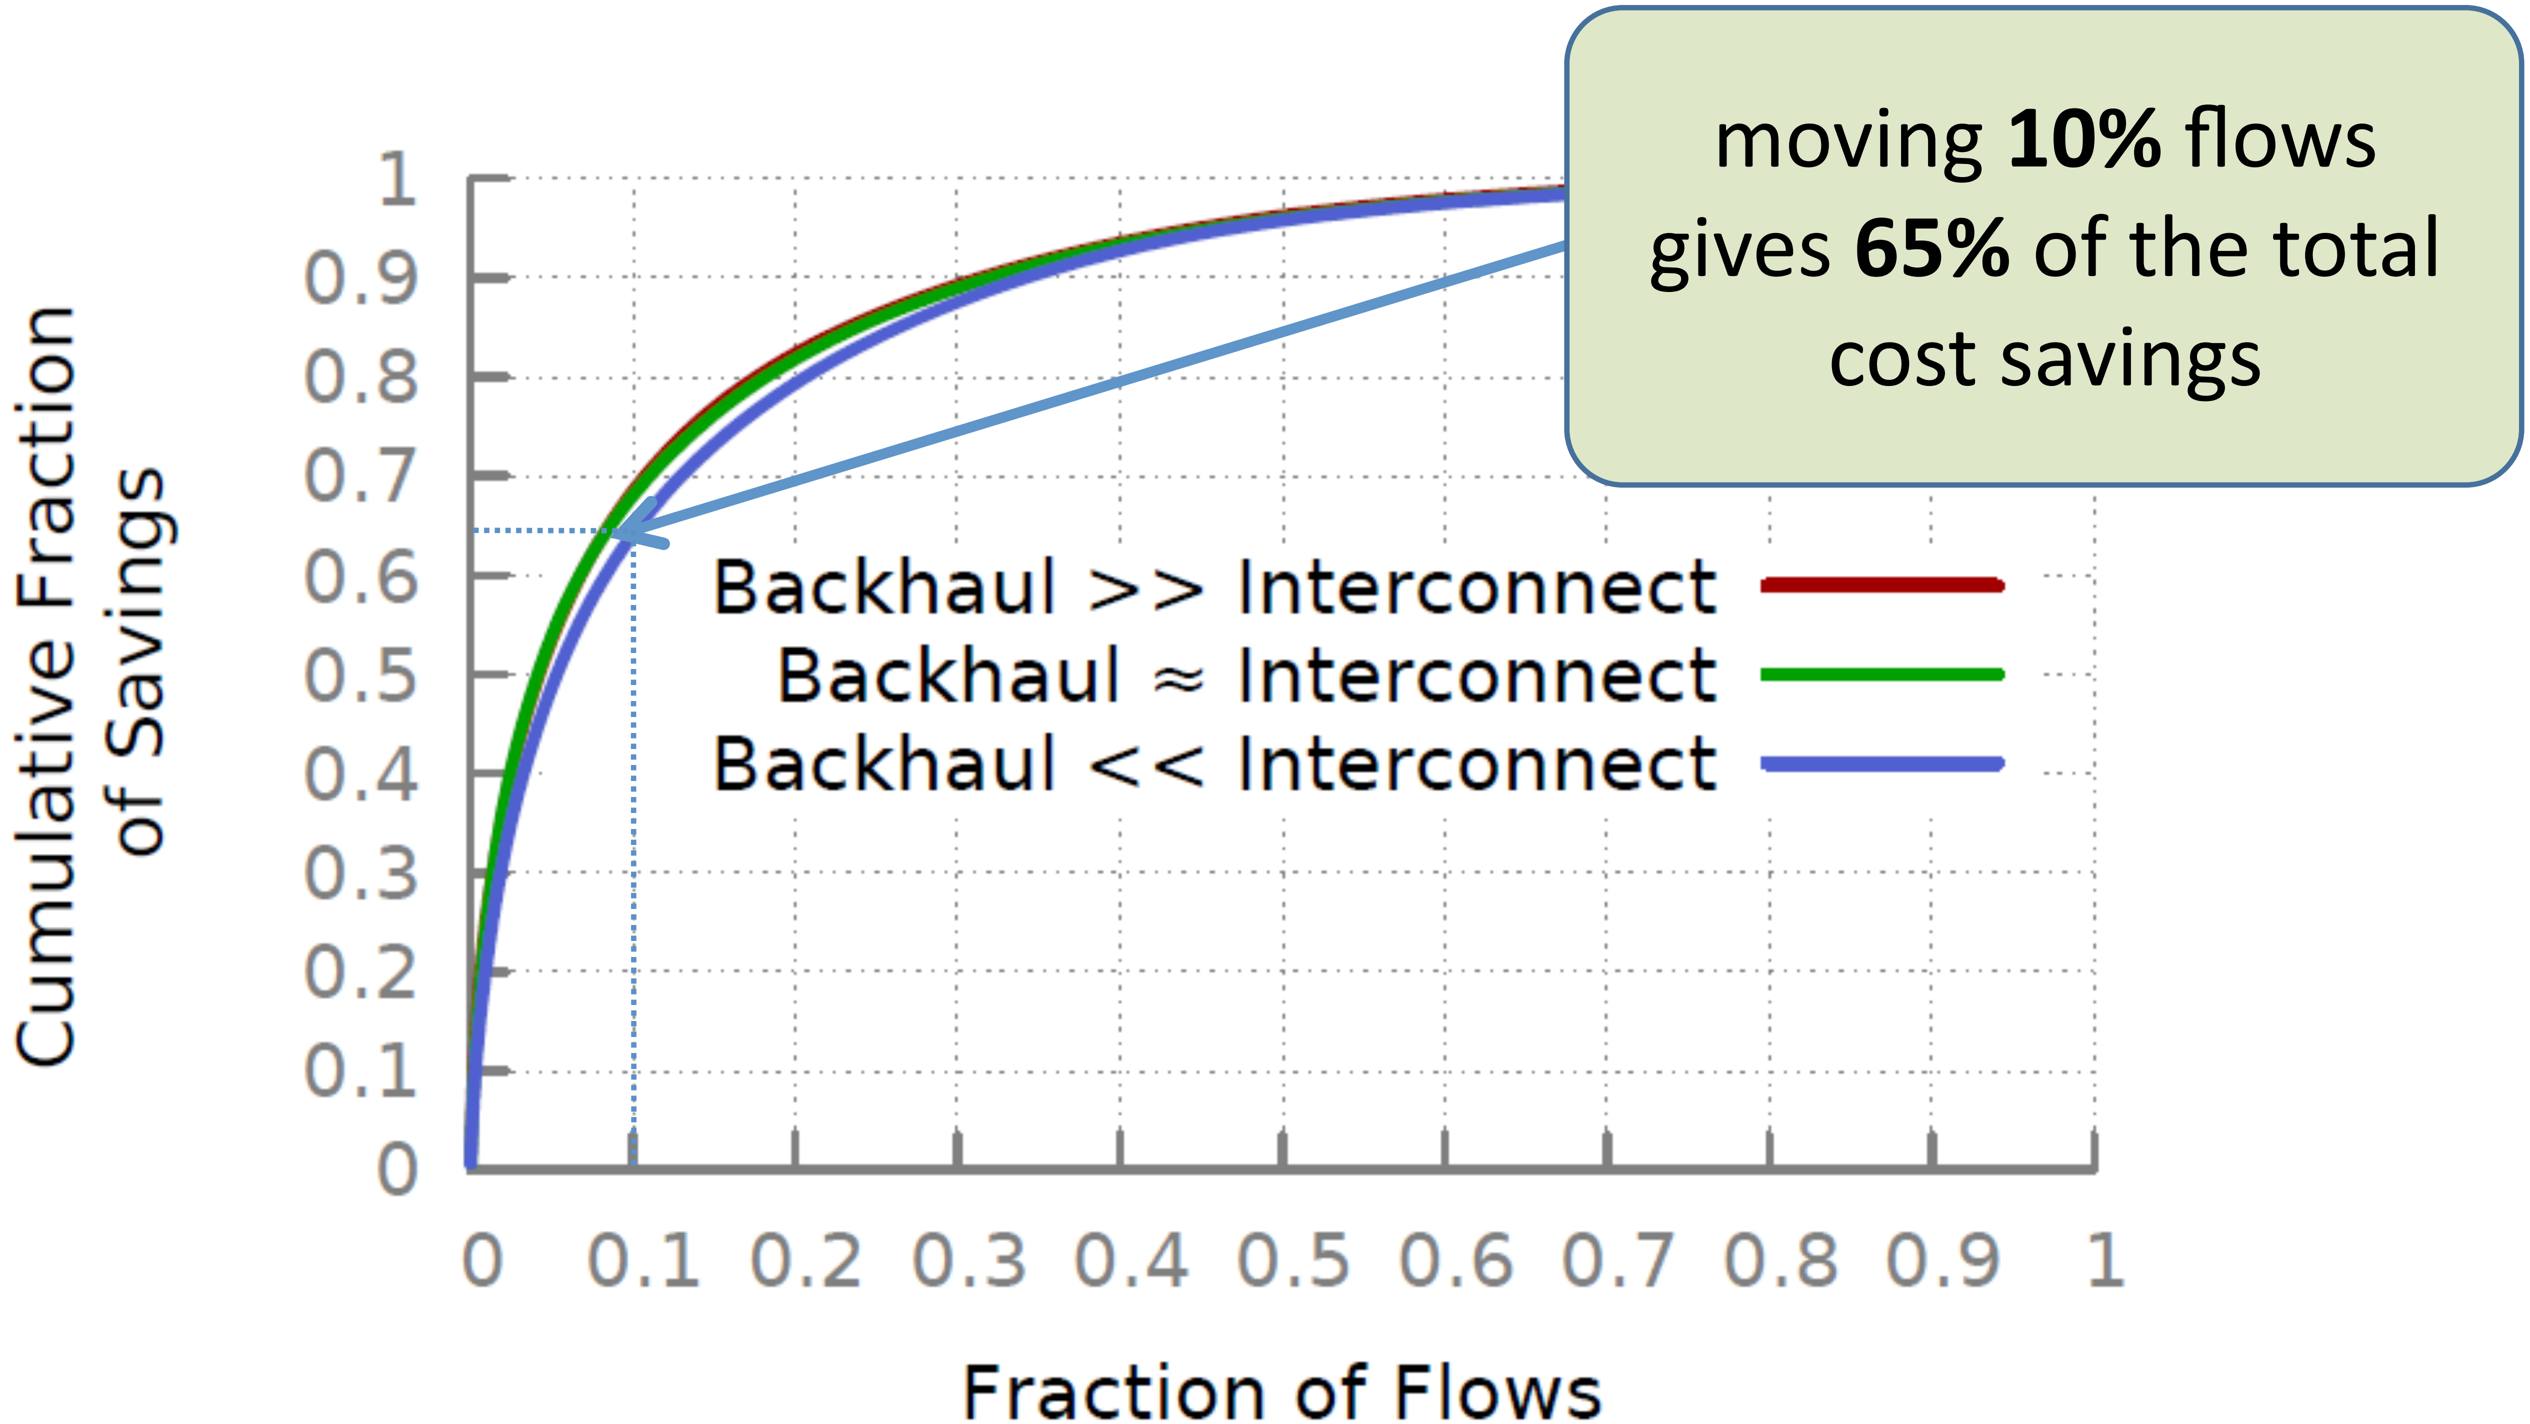

Part 2: A Value-based Framework for Peering Decisions

Evaluating Potential Peers

- To peer or not to peer...

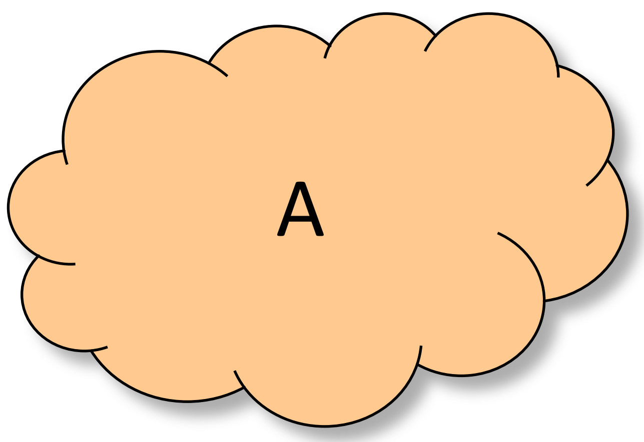

Evaluating Potential Peers

- To peer or not to peer...

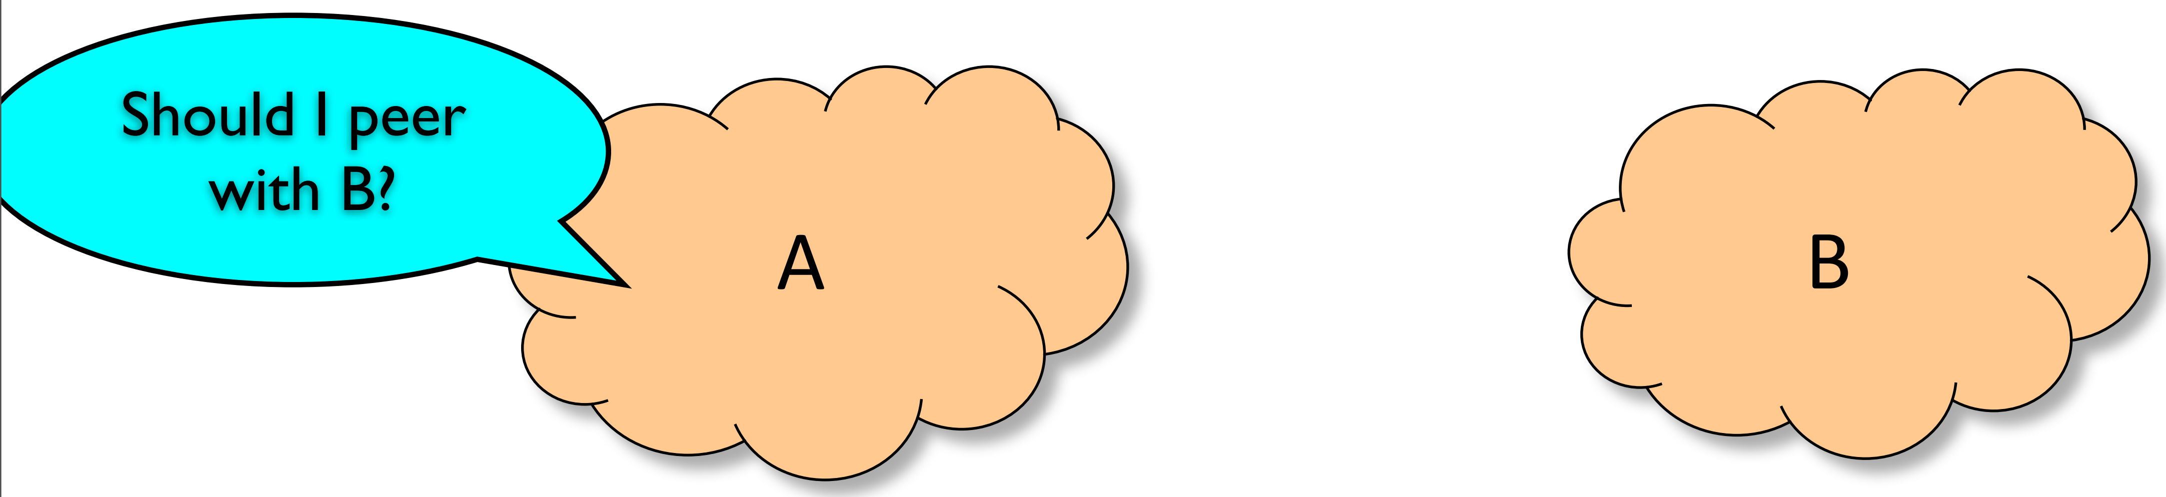

Evaluating Potential Peers

- To peer or not to peer...

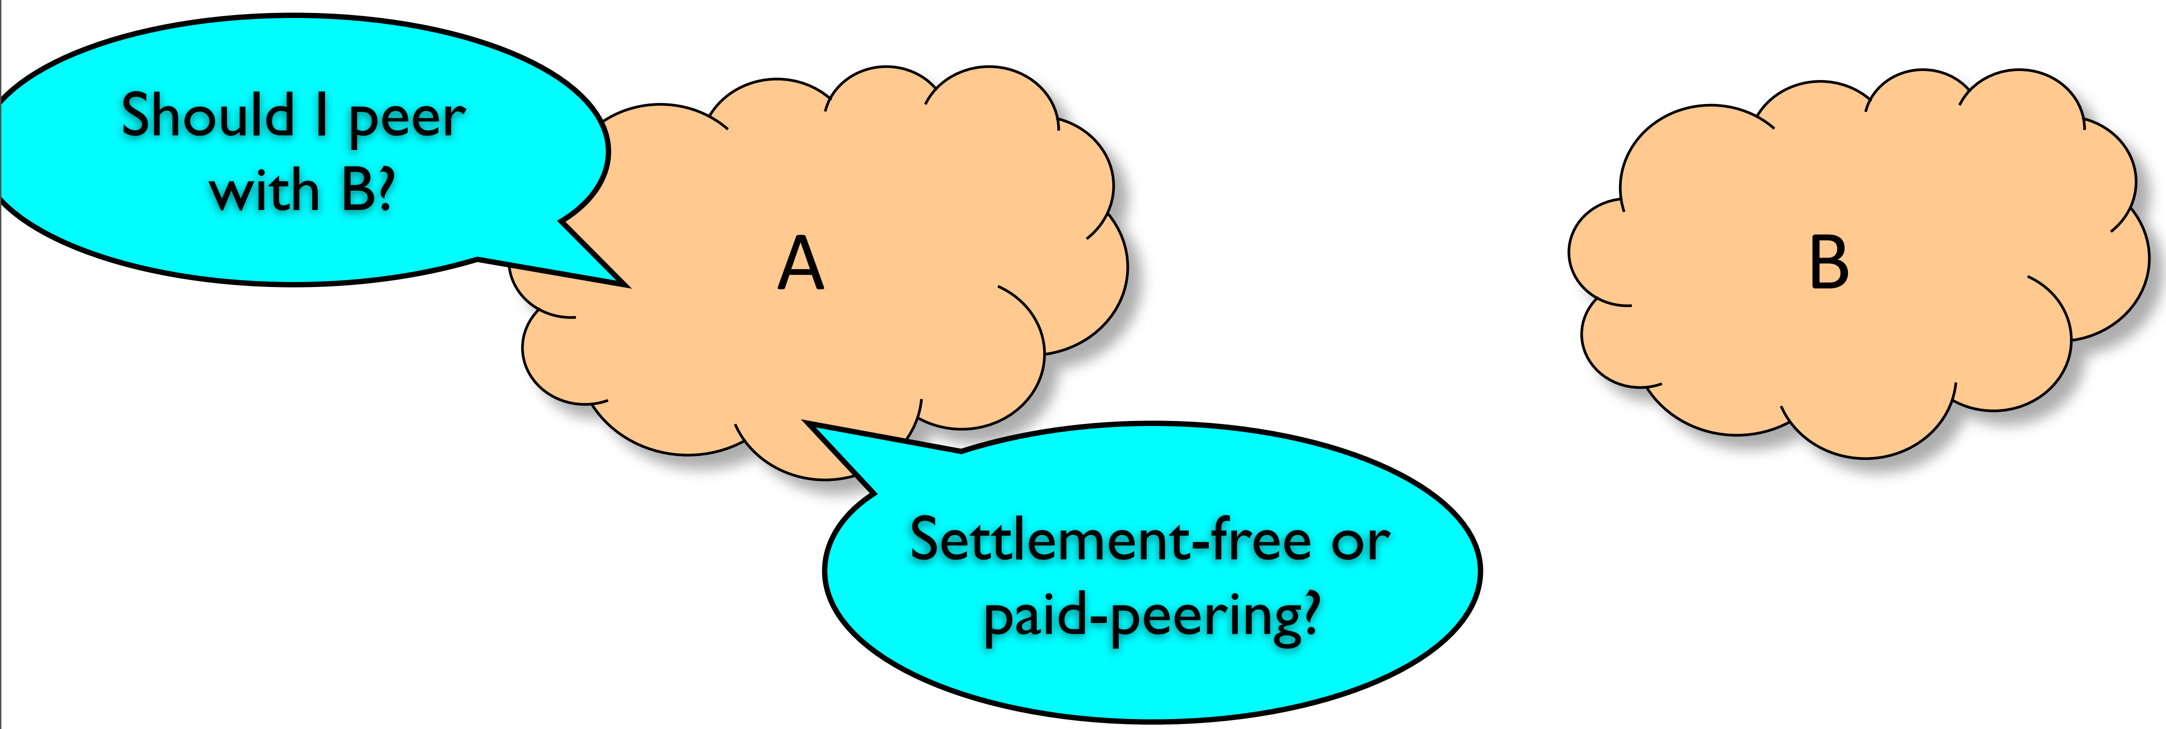

Evaluating Potential Peers

- To peer or not to peer…

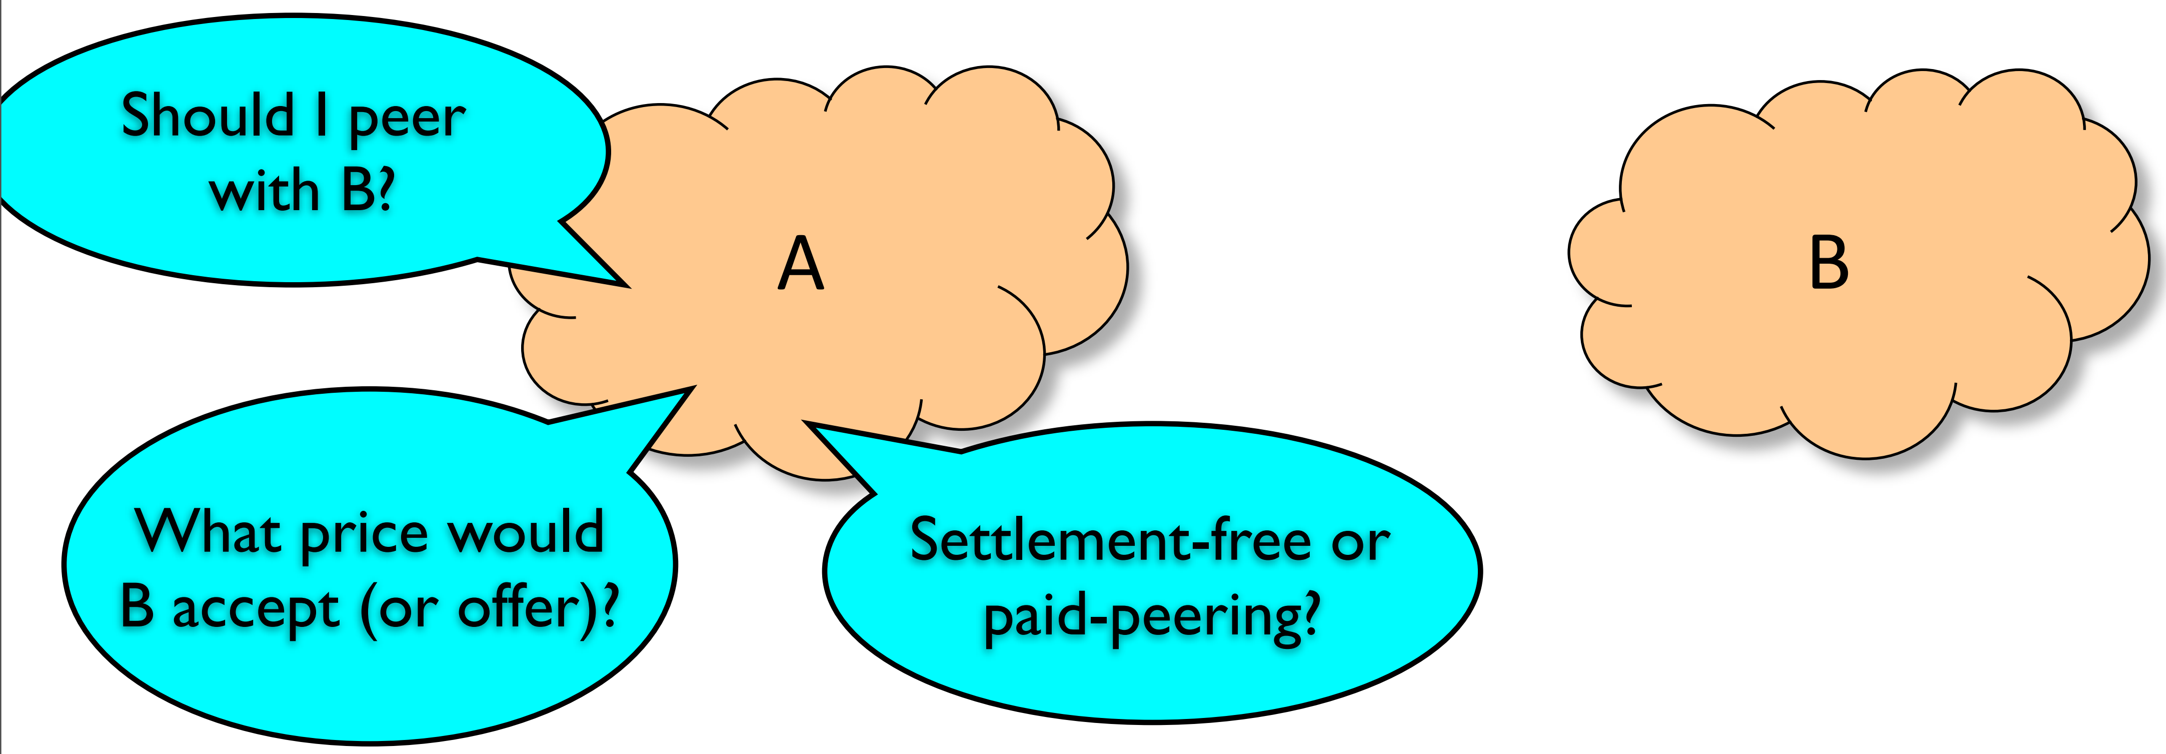

Evaluating Potential Peers

- To peer or not to peer...

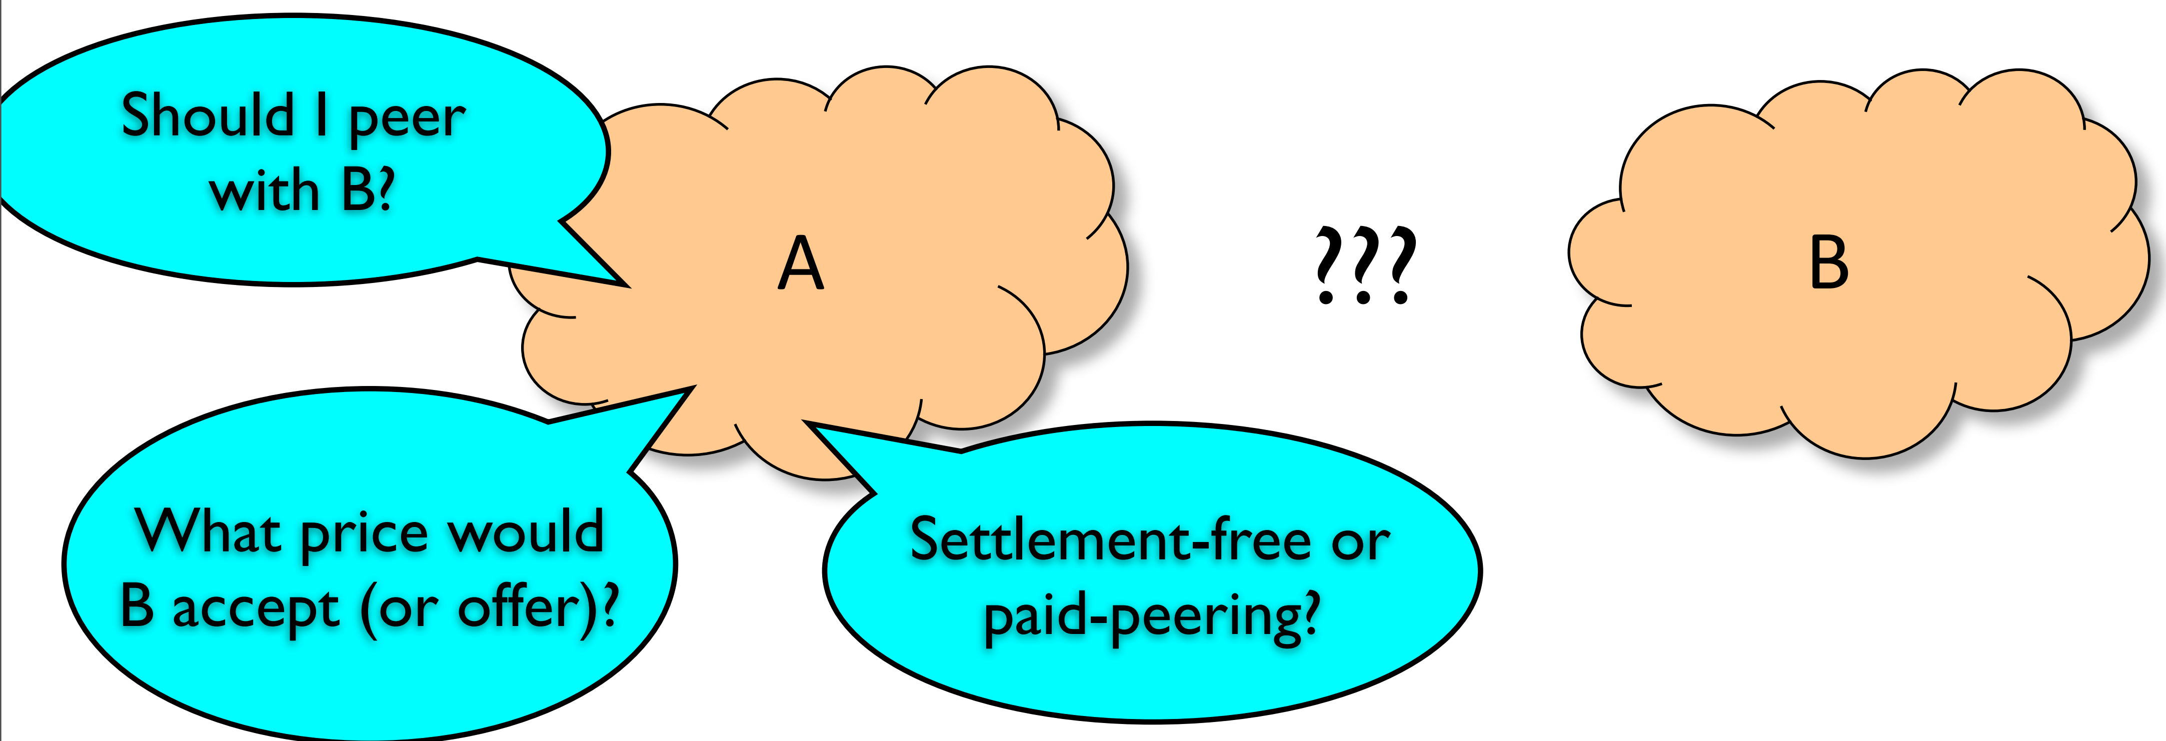

Value Based Peering

- Price based on the “value” of the link

- For a network, define the notion of economic “fitness”
- fitness = revenue – interconnect costs – backhaul cost

- Value of a peering link is the difference in fitness with and without the link
- Value = f_{with} - $f_{without}$

- Revenue and costs could change on peering/depeering

What Affects Peering Value?

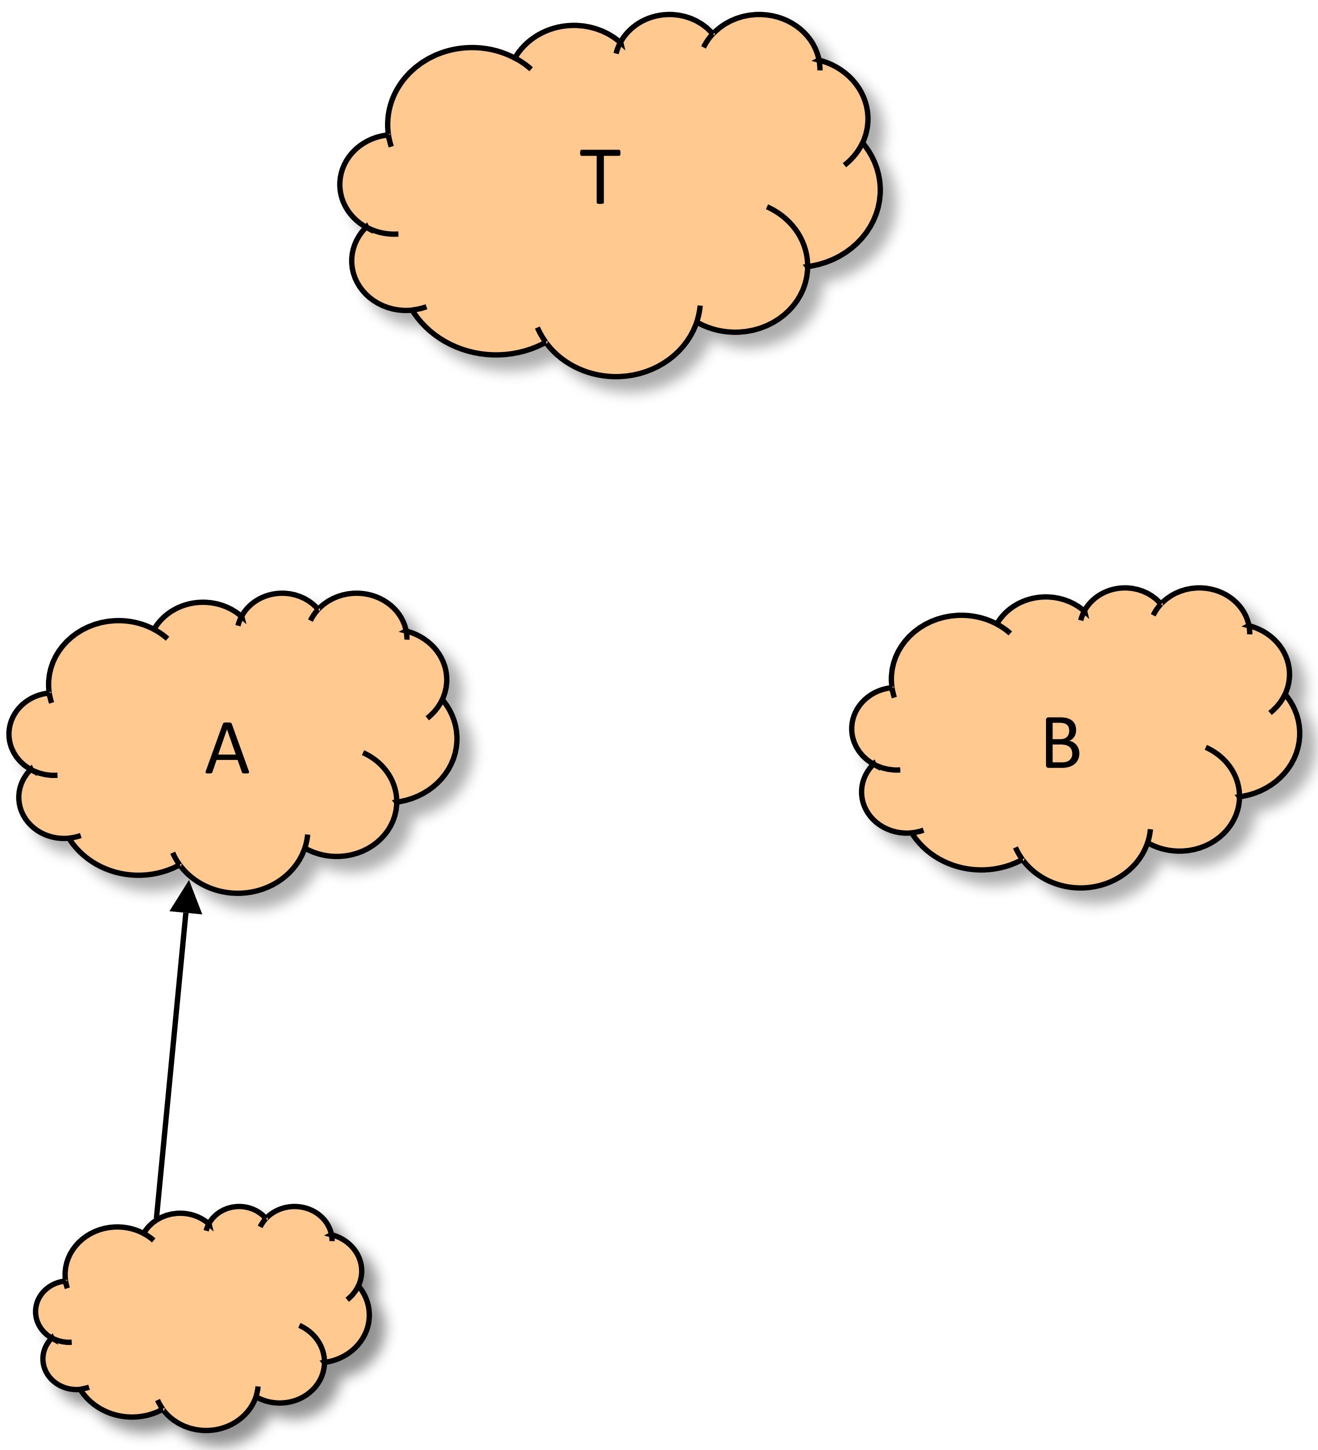

What Affects Peering Value?

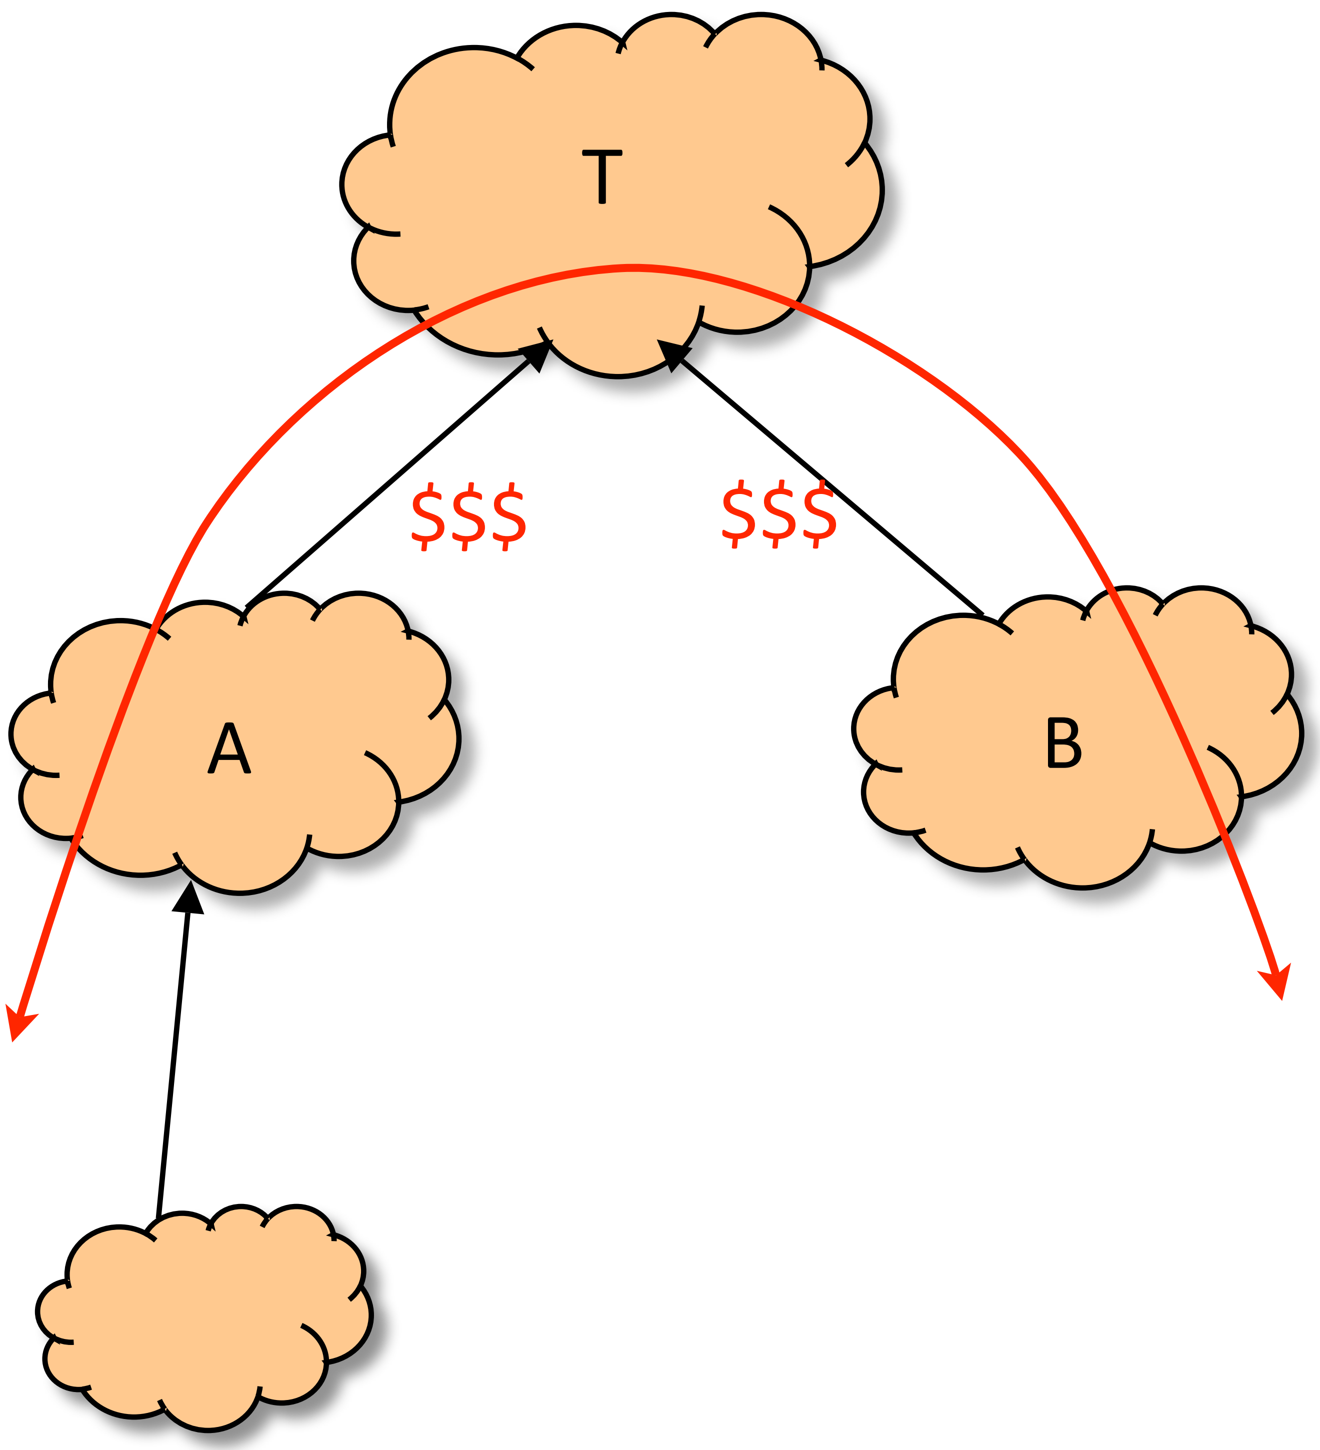

What Affects Peering Value?

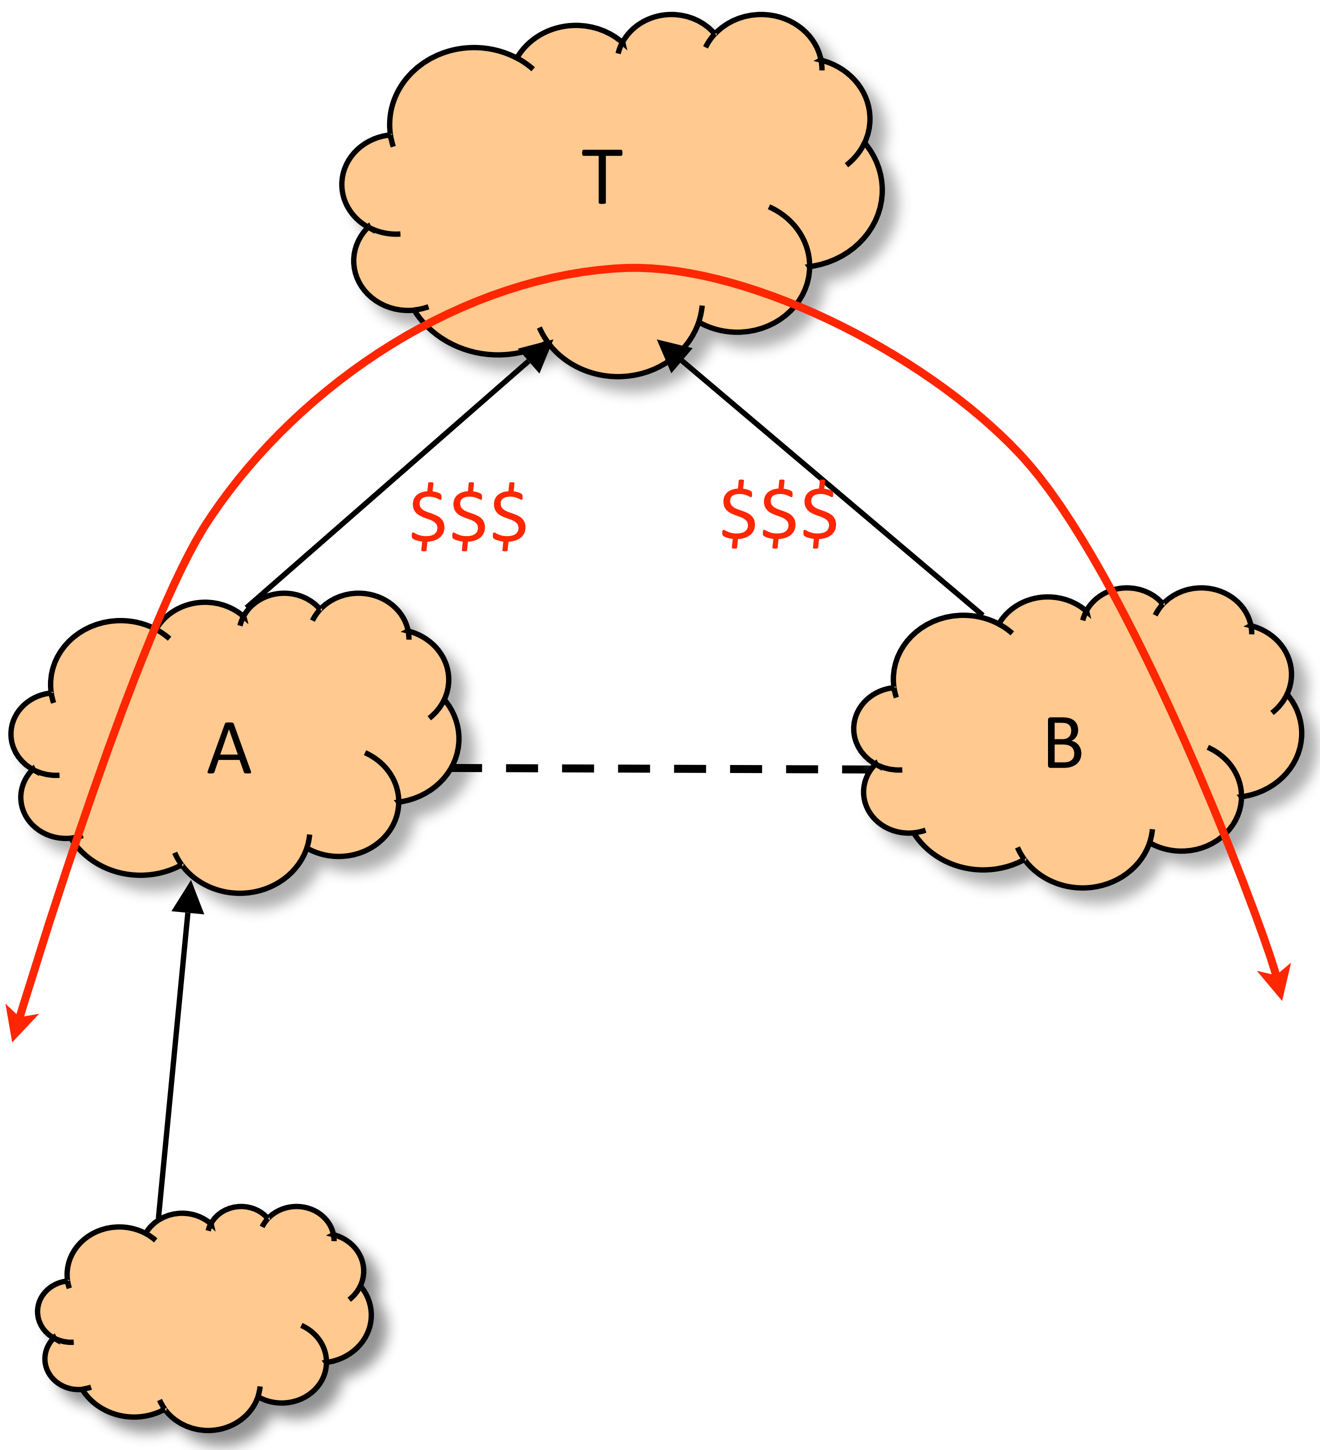

What Affects Peering Value?

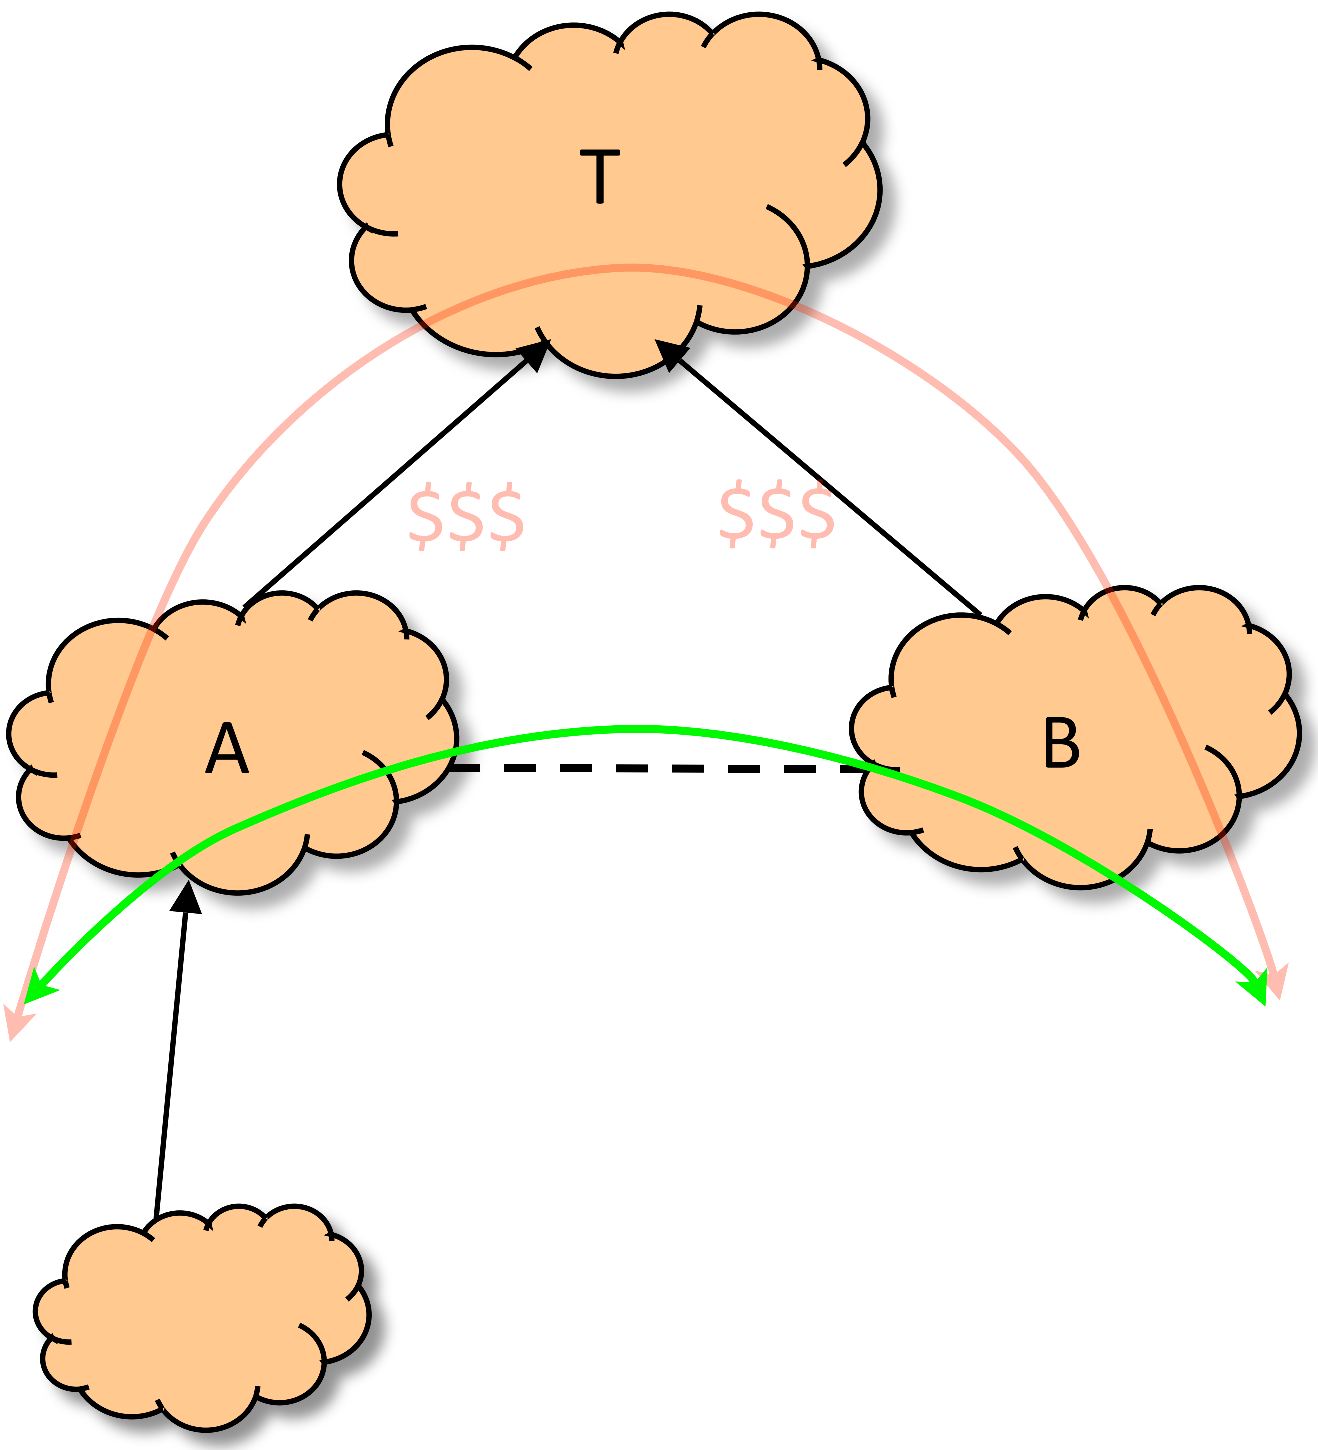

- Interconnect cost changes: Avoid a transit provider

What Affects Peering Value?

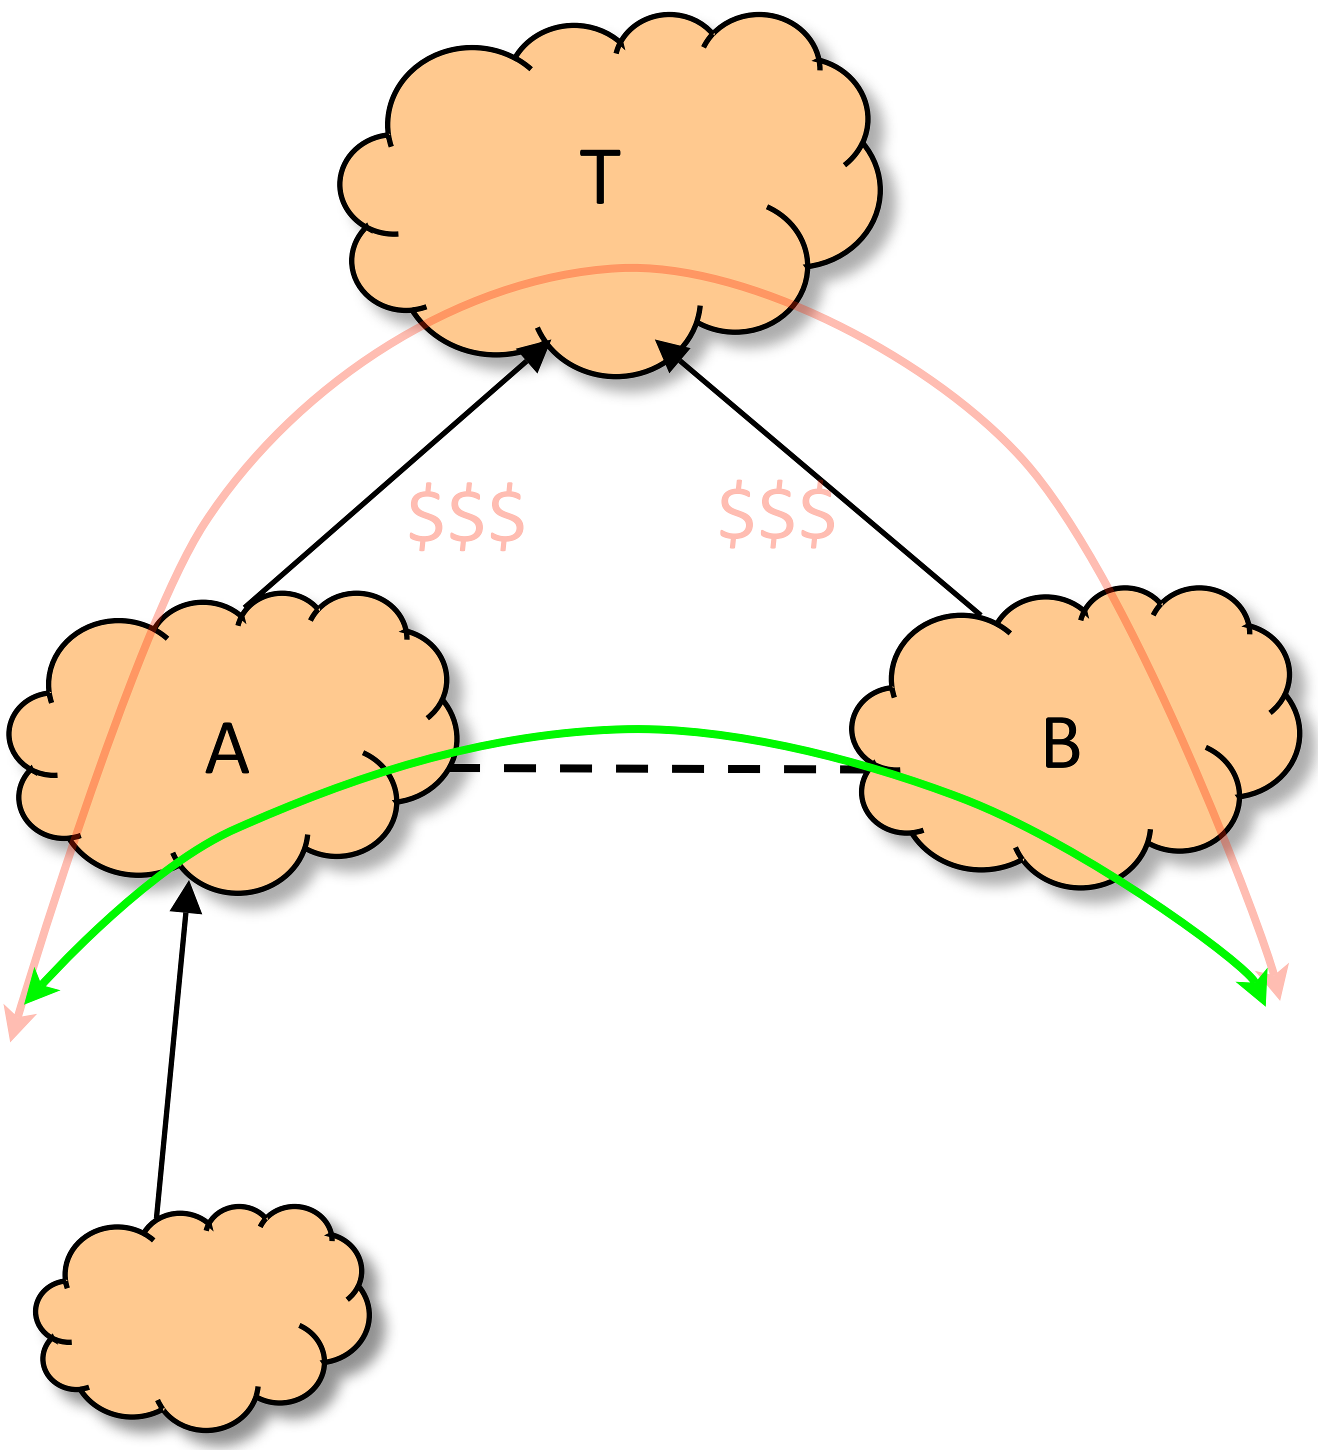

- Interconnect cost changes: Avoid a transit provider
- Backhaul cost changes: Peering link changes how traffic is routed in a network

What Affects Peering Value?

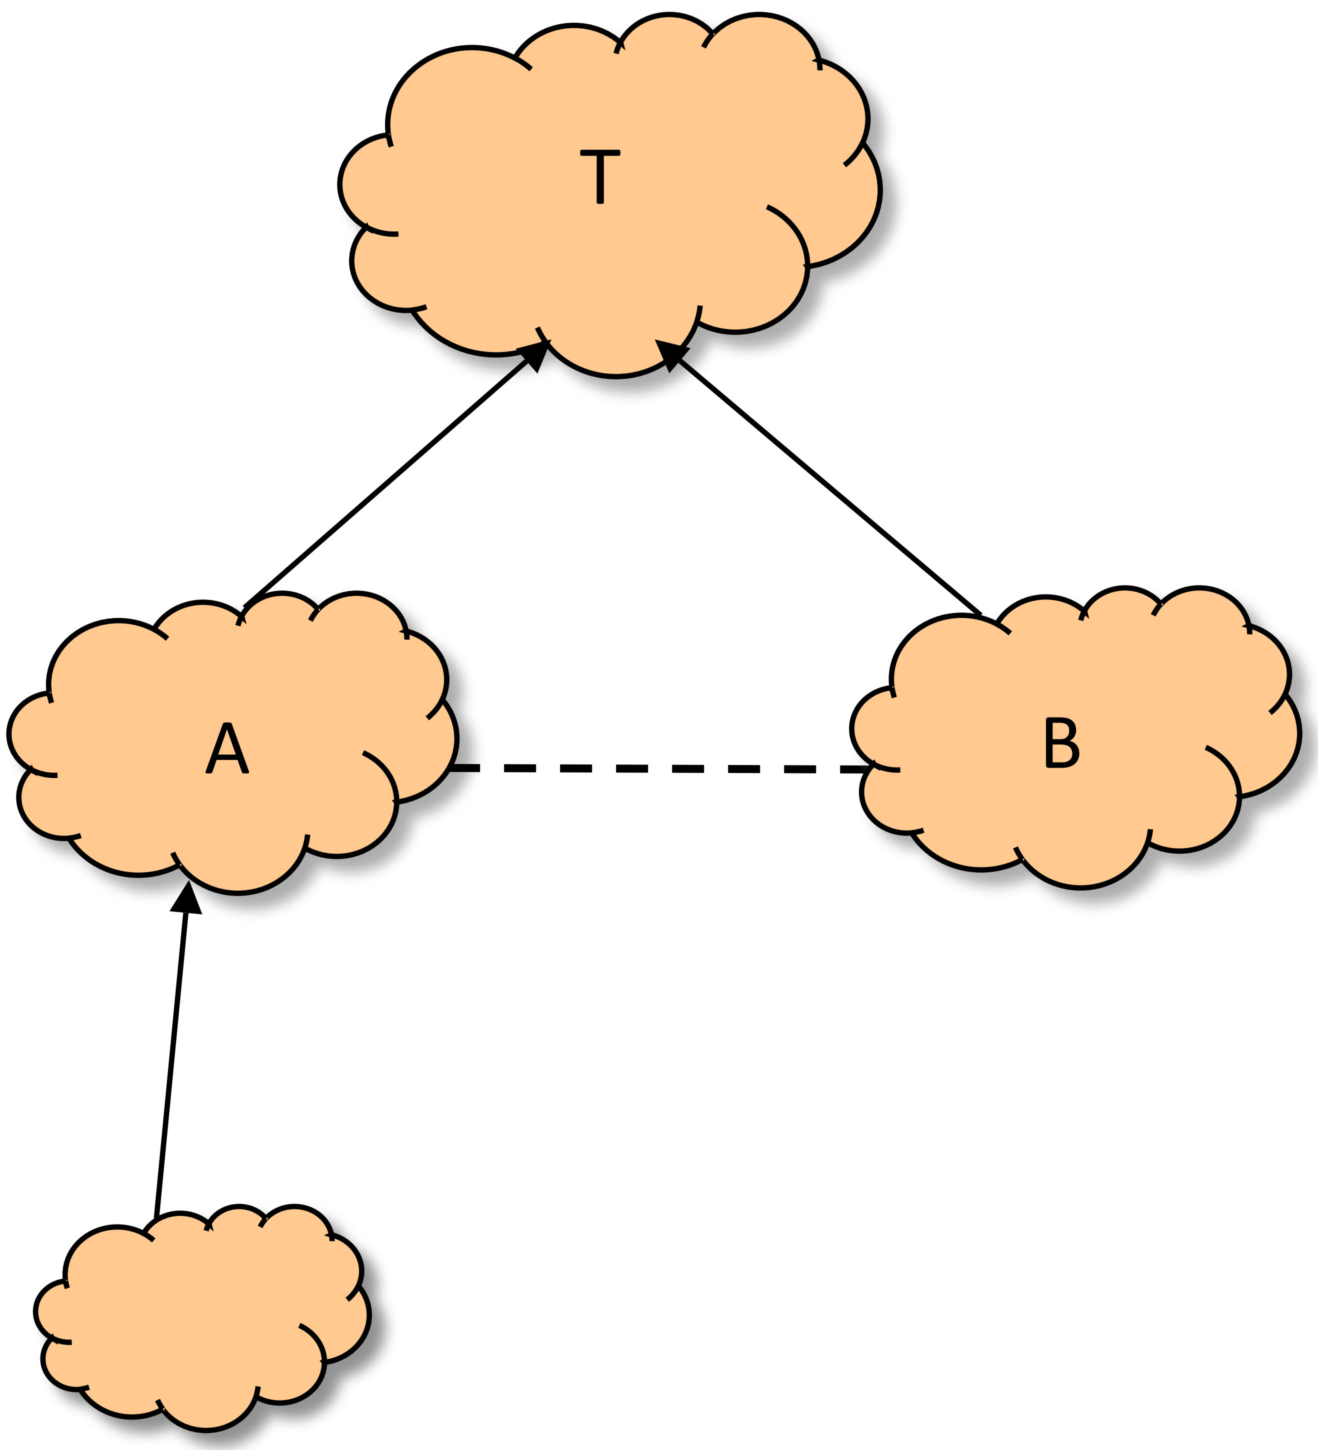

- Interconnect cost changes: Avoid a transit provider
- Backhaul cost changes: Peering link changes how traffic is routed in a network

What Affects Peering Value?

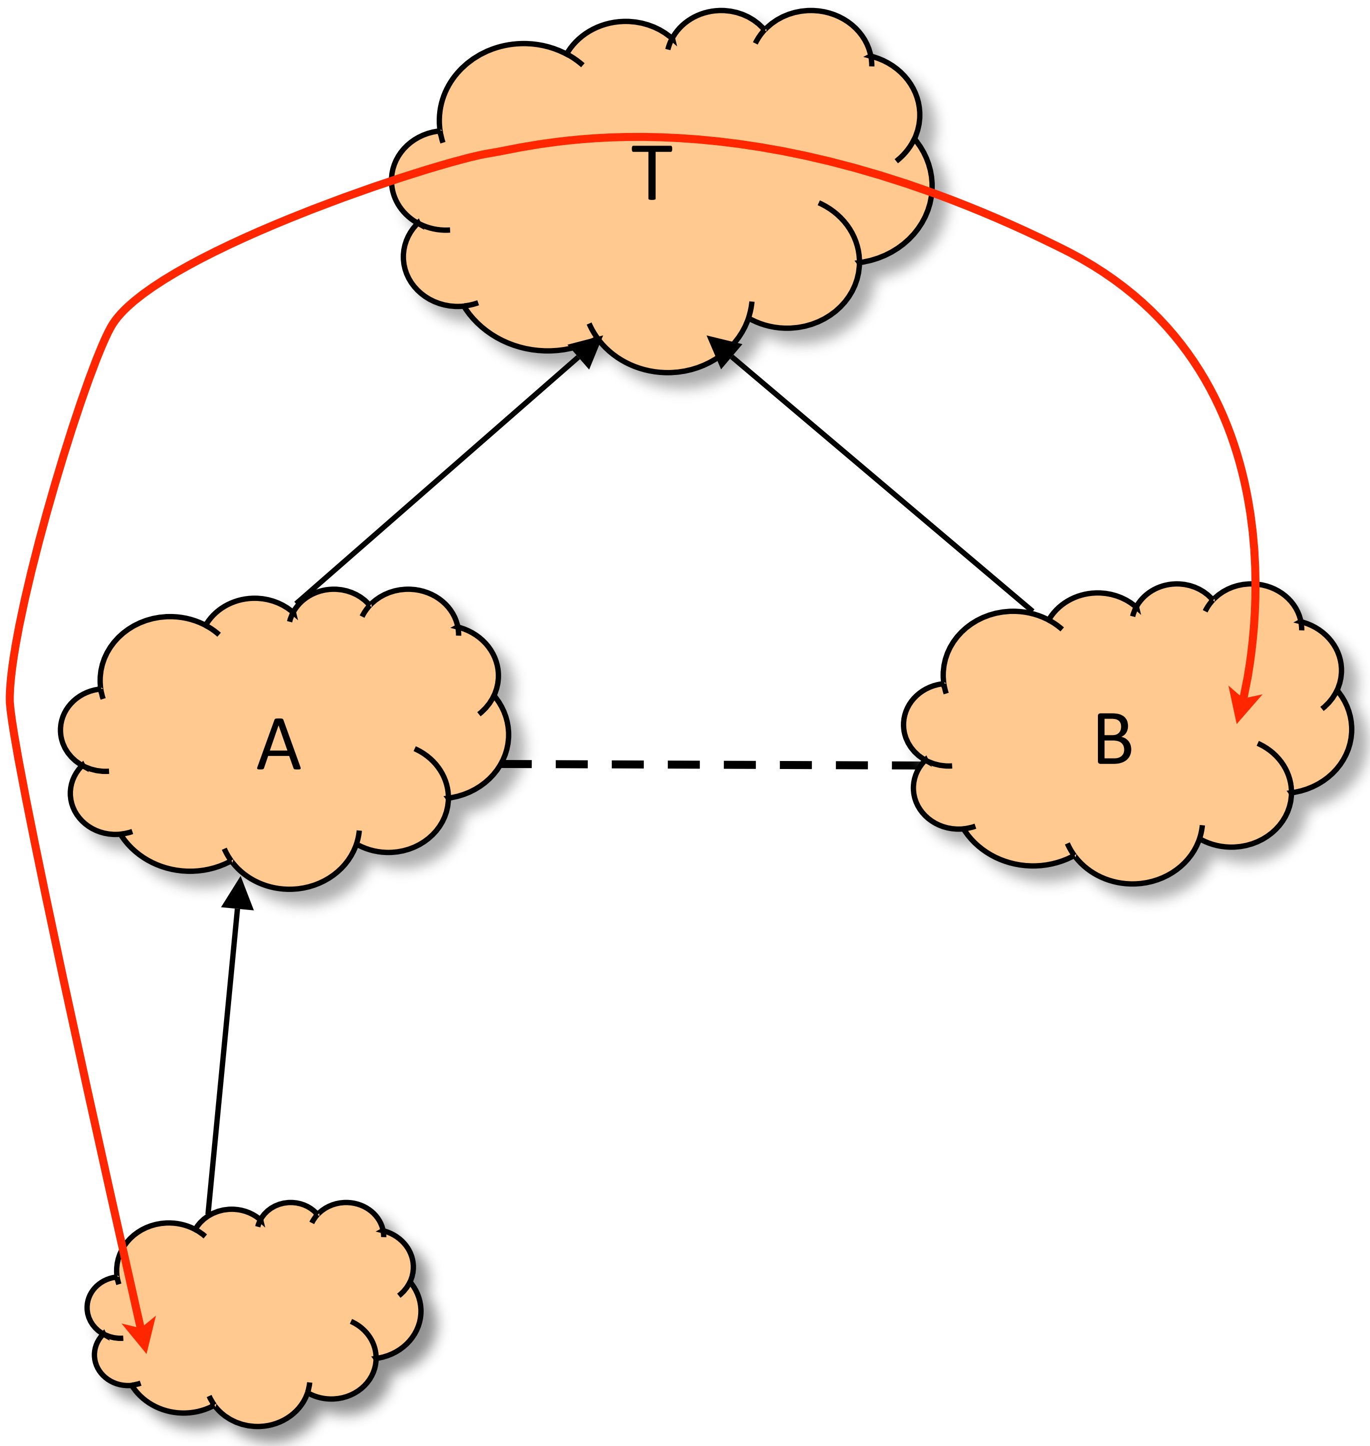

- Interconnect cost changes: Avoid a transit provider
- Backhaul cost changes: Peering link changes how traffic is routed in a network

What Affects Peering Value?

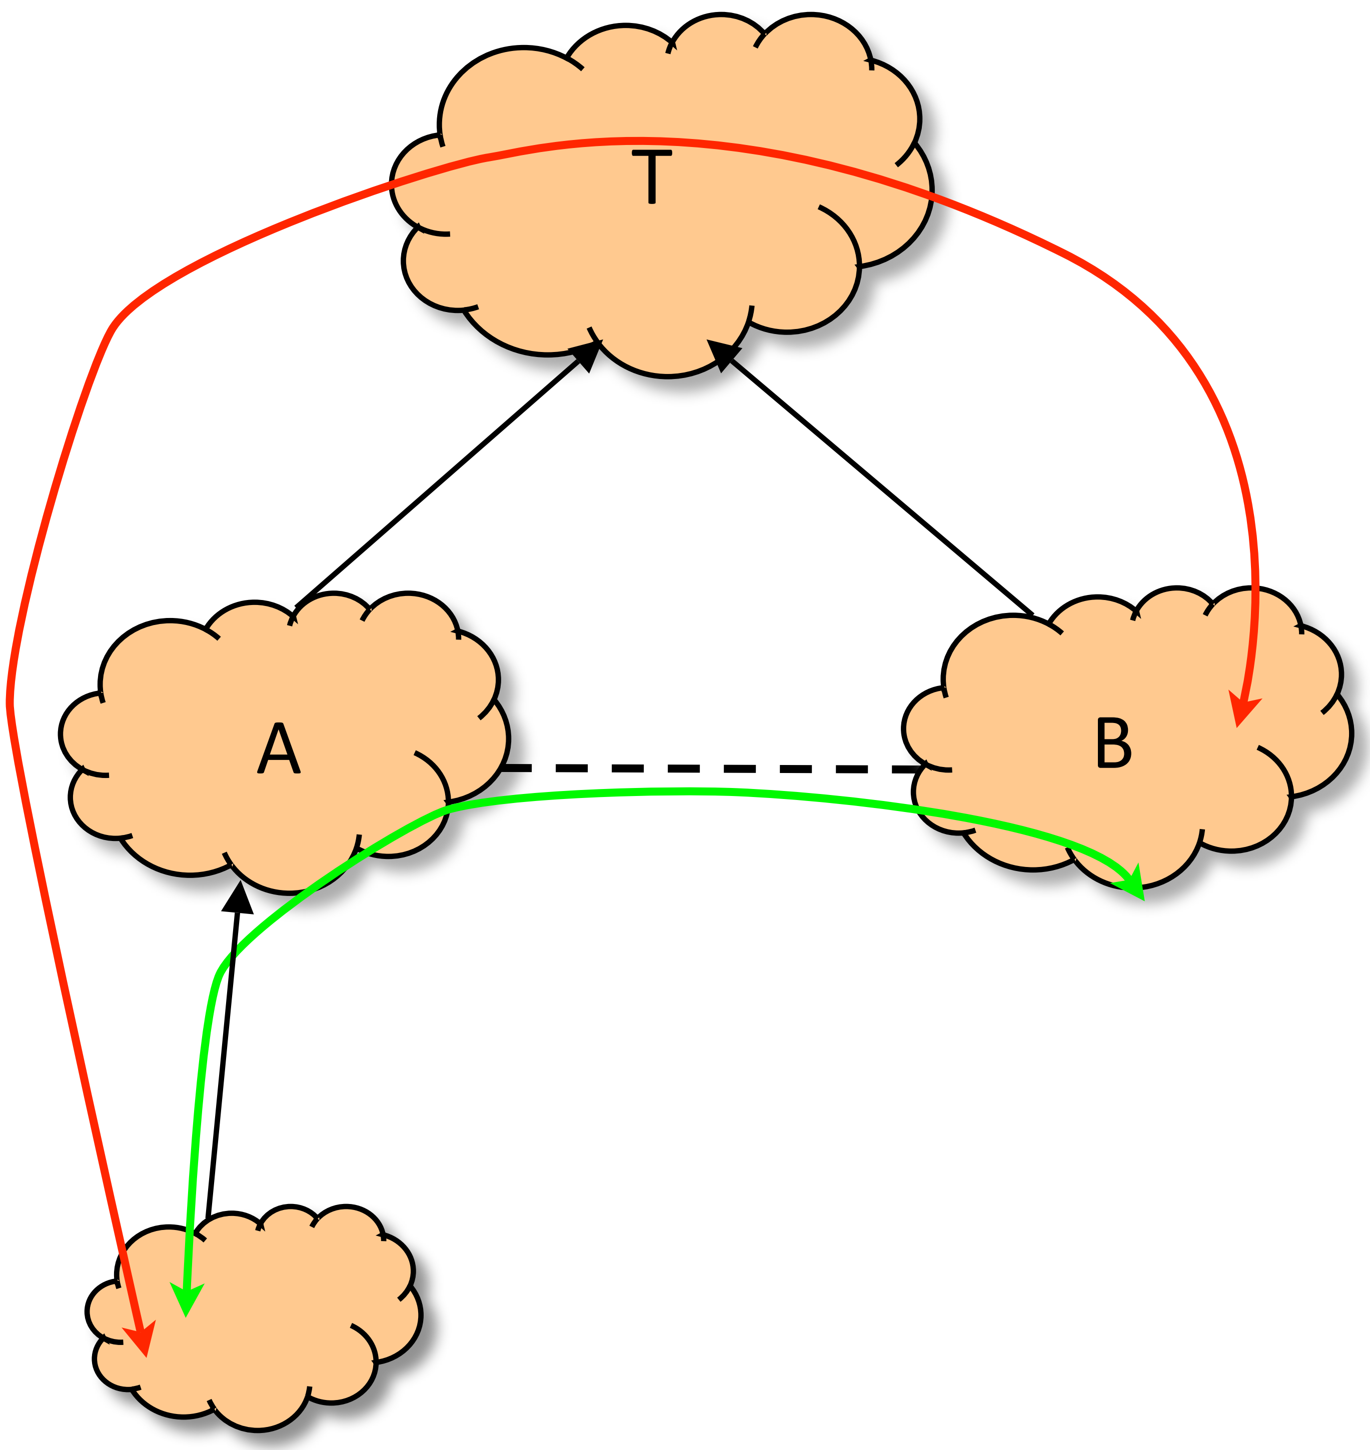

- Interconnect cost changes: Avoid a transit provider
- Backhaul cost changes: Peering link changes how traffic is routed in a network

What Affects Peering Value?

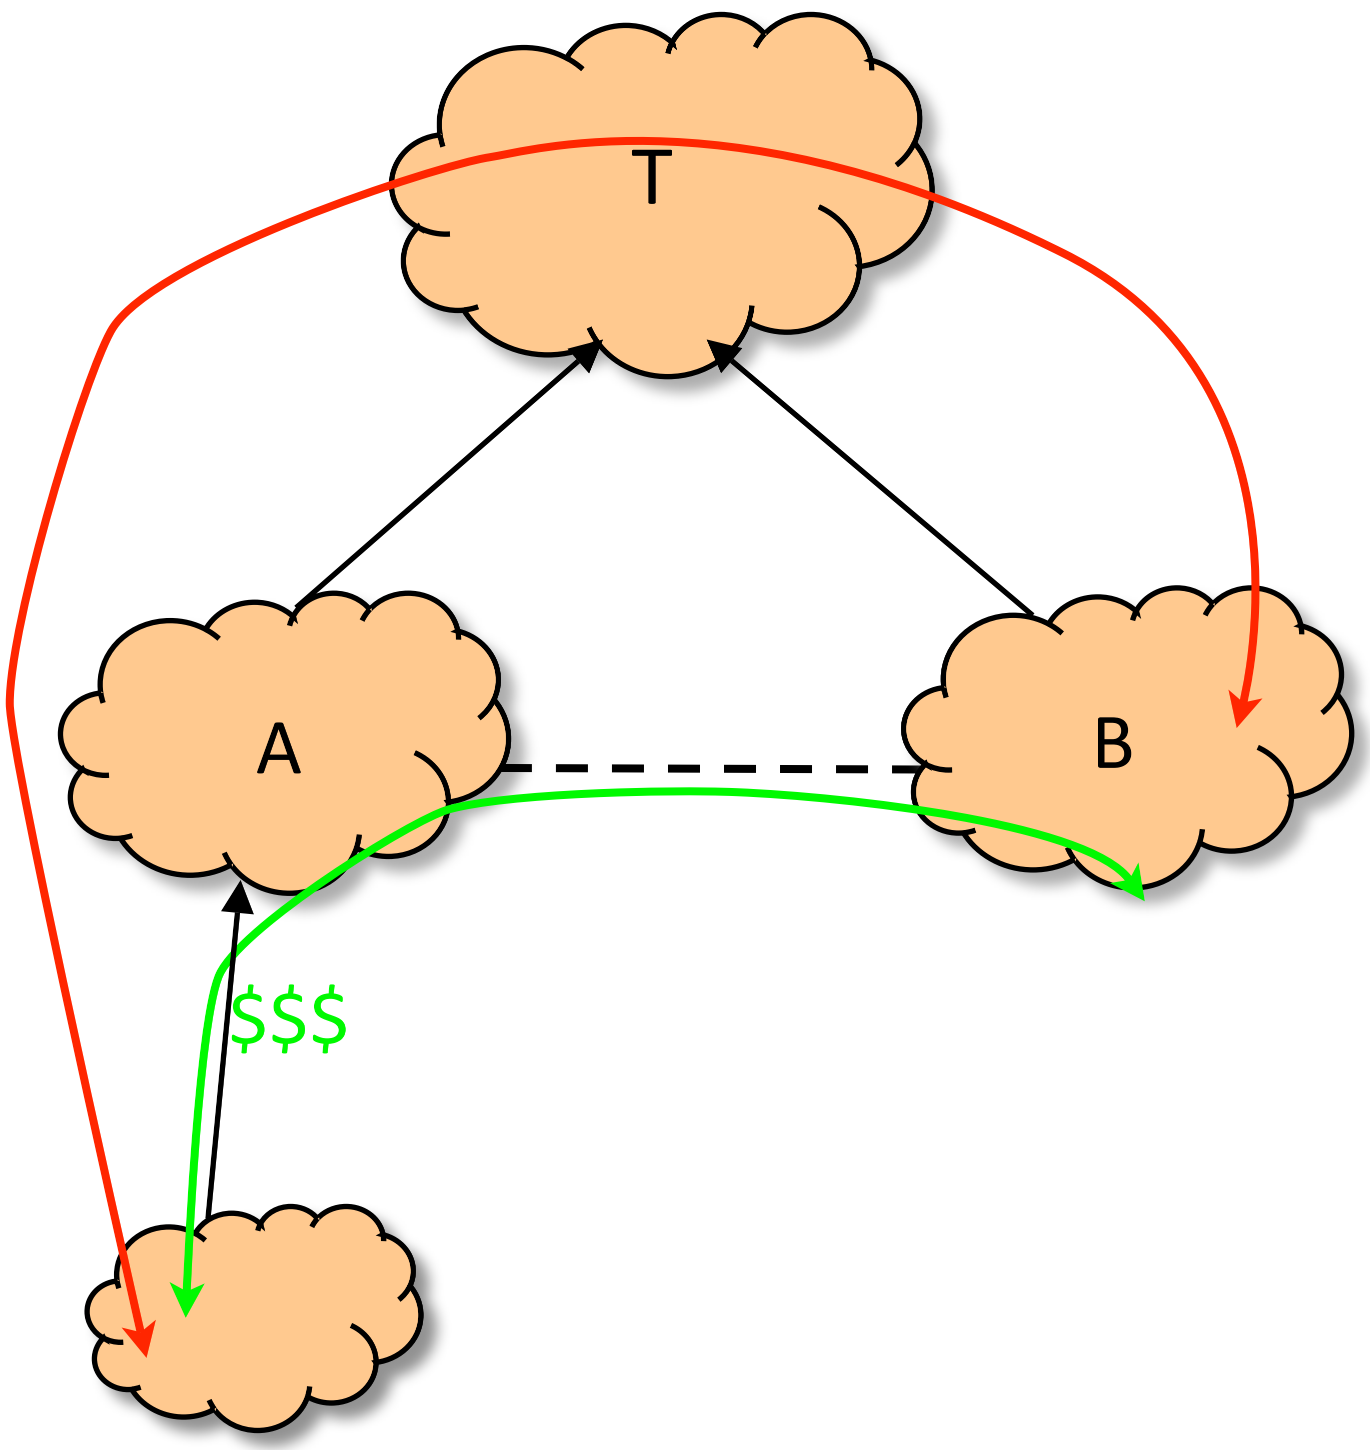

- Interconnect cost changes: Avoid a transit provider
- Backhaul cost changes: Peering link changes how traffic is routed in a network

What Affects Peering Value?

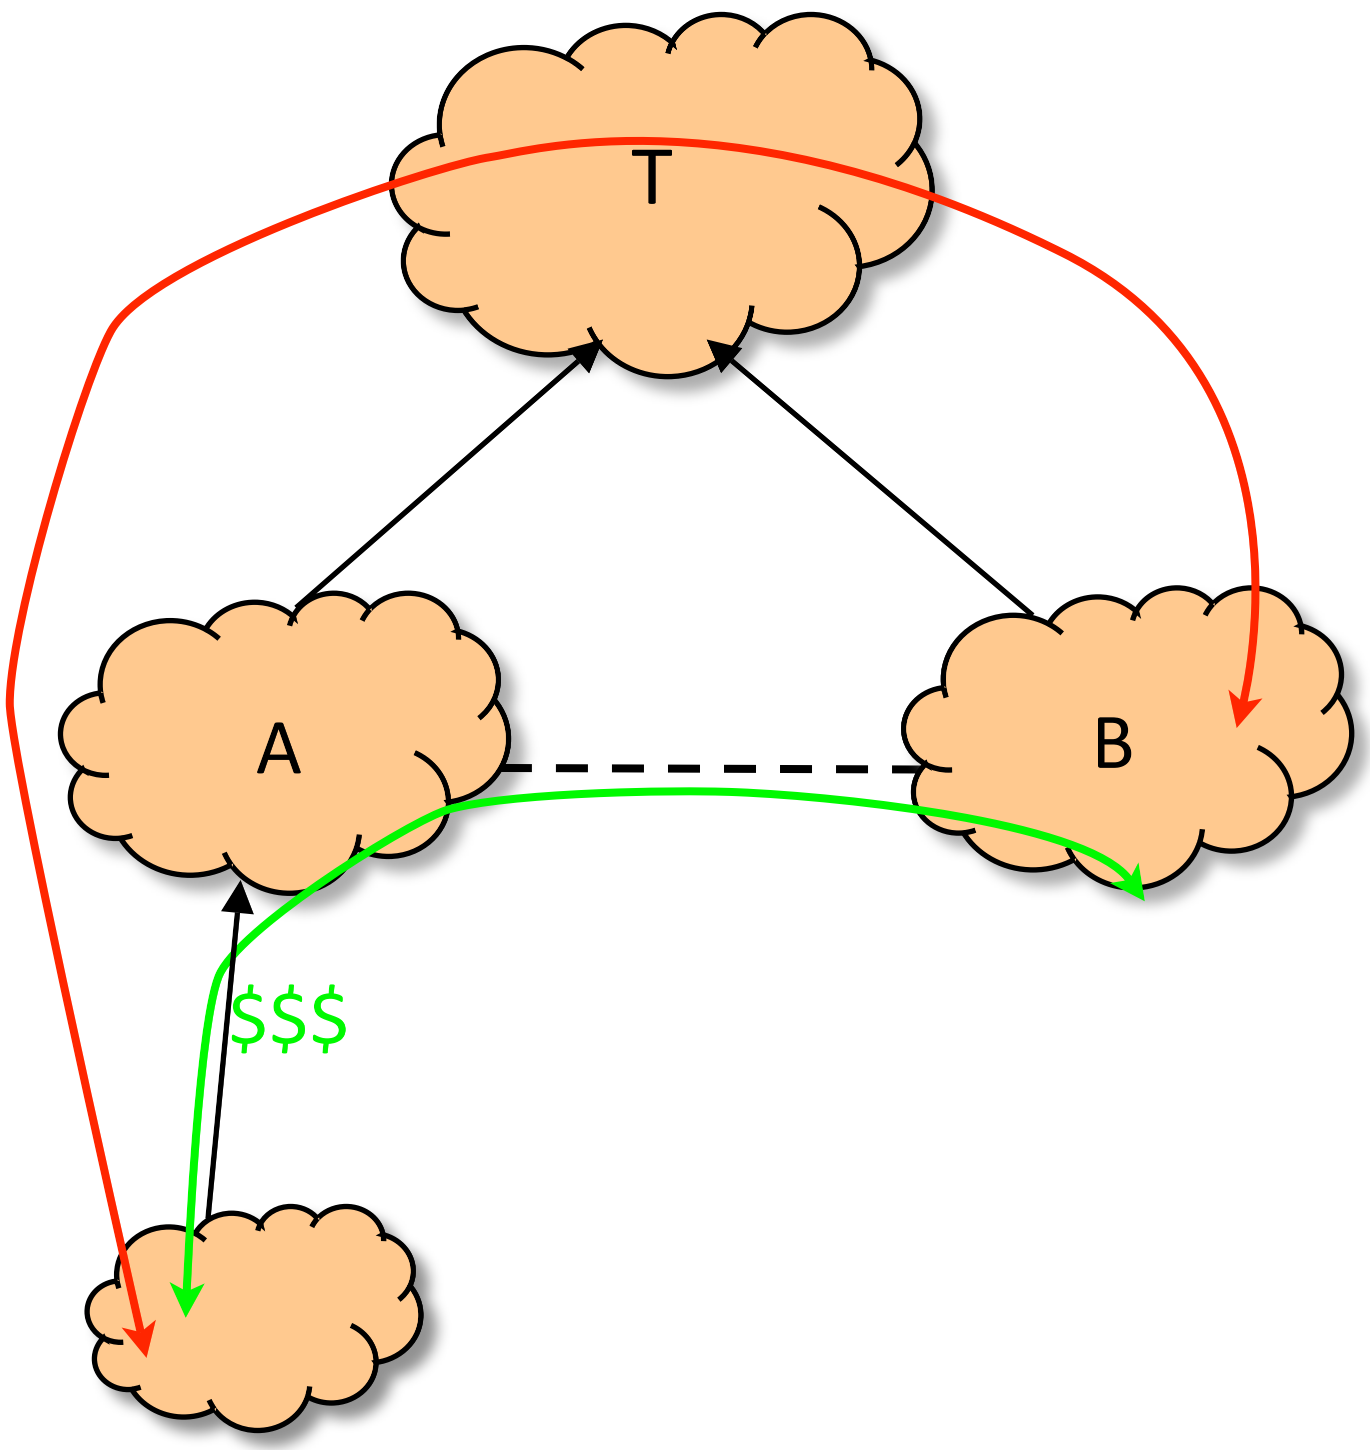

- Interconnect cost changes: Avoid a transit provider
- Backhaul cost changes: Peering link changes how traffic is routed in a network
- Revenue changes: Attract/ lose traffic due to new peering link

The Fair Peering Price

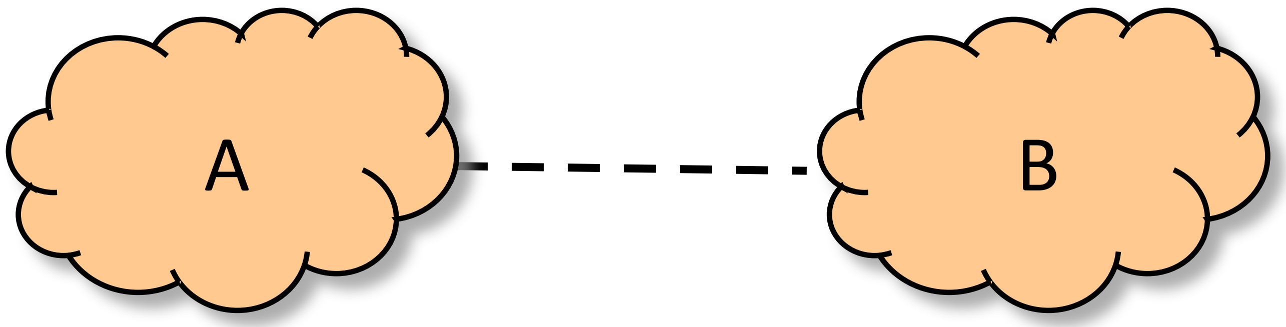

The Fair Peering Price

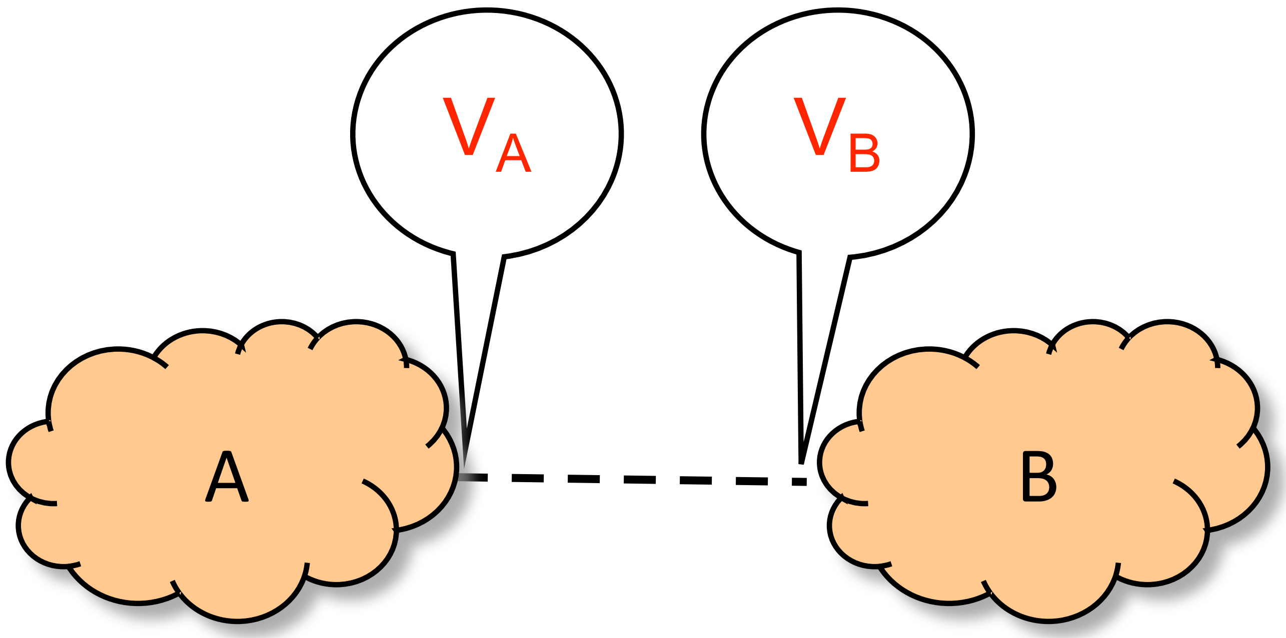

The Fair Peering Price

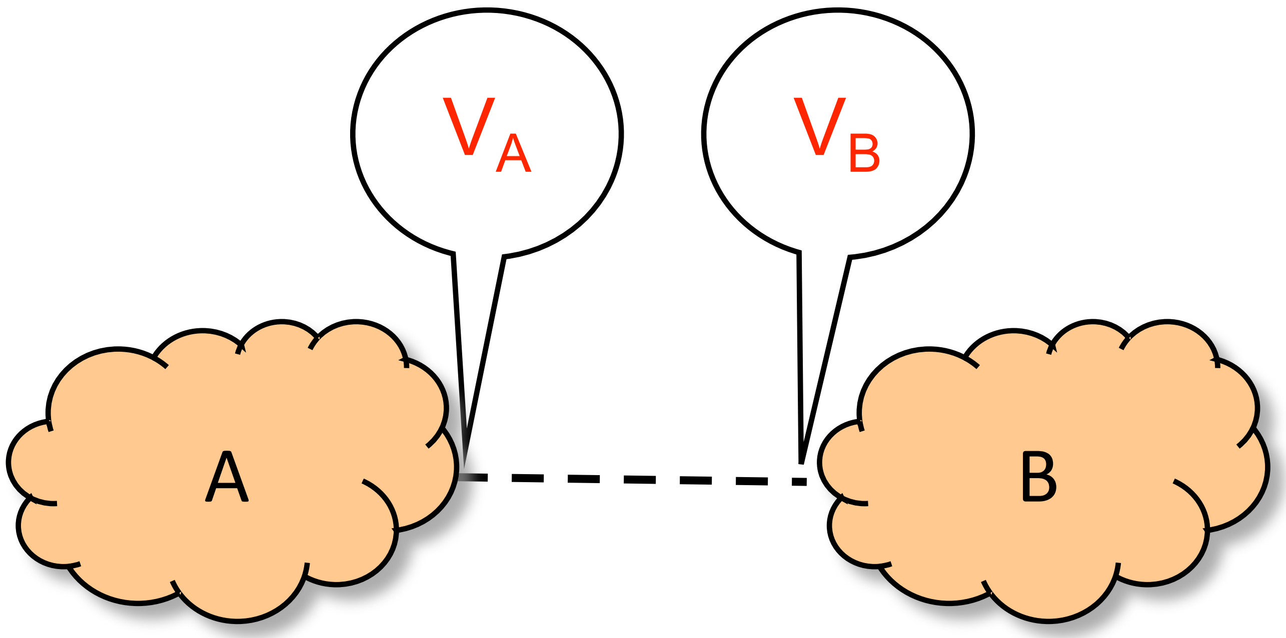

- A and B see values V_A and V_B

The Fair Peering Price

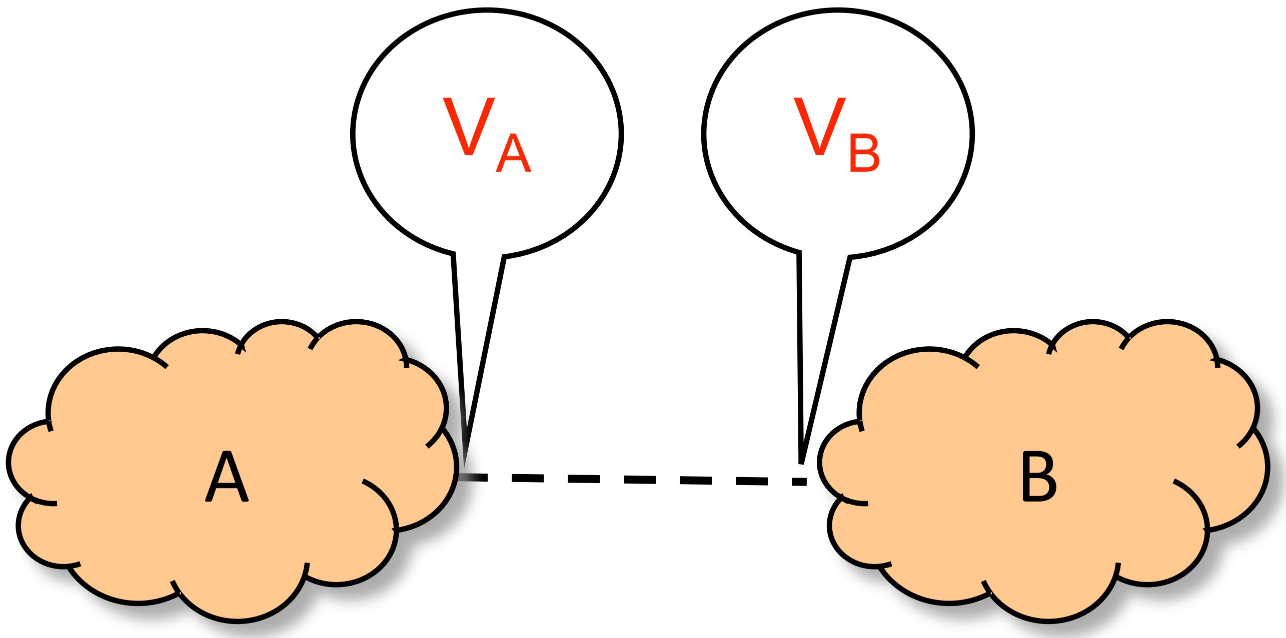

- A and B see values V_A and V_B
- What should be the paid-peering price?

The Fair Peering Price

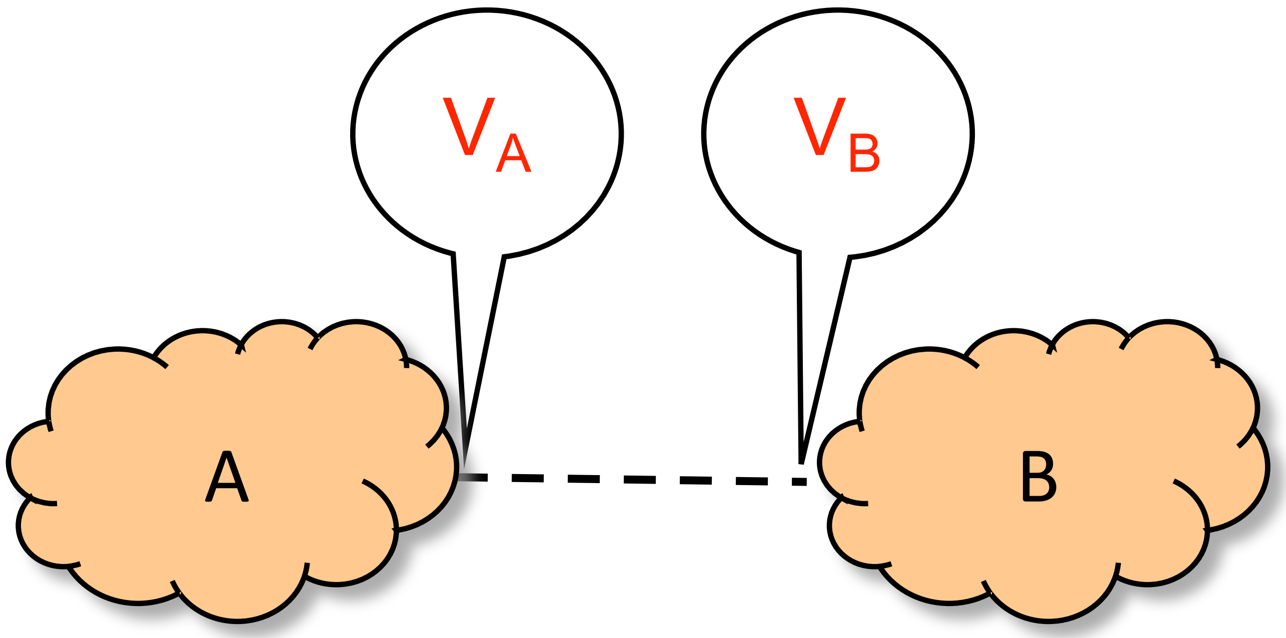

- A and B see values V_A and V_B
- What should be the paid-peering price?
- Fair price is $(V_A - V_B)/2$

The Fair Peering Price

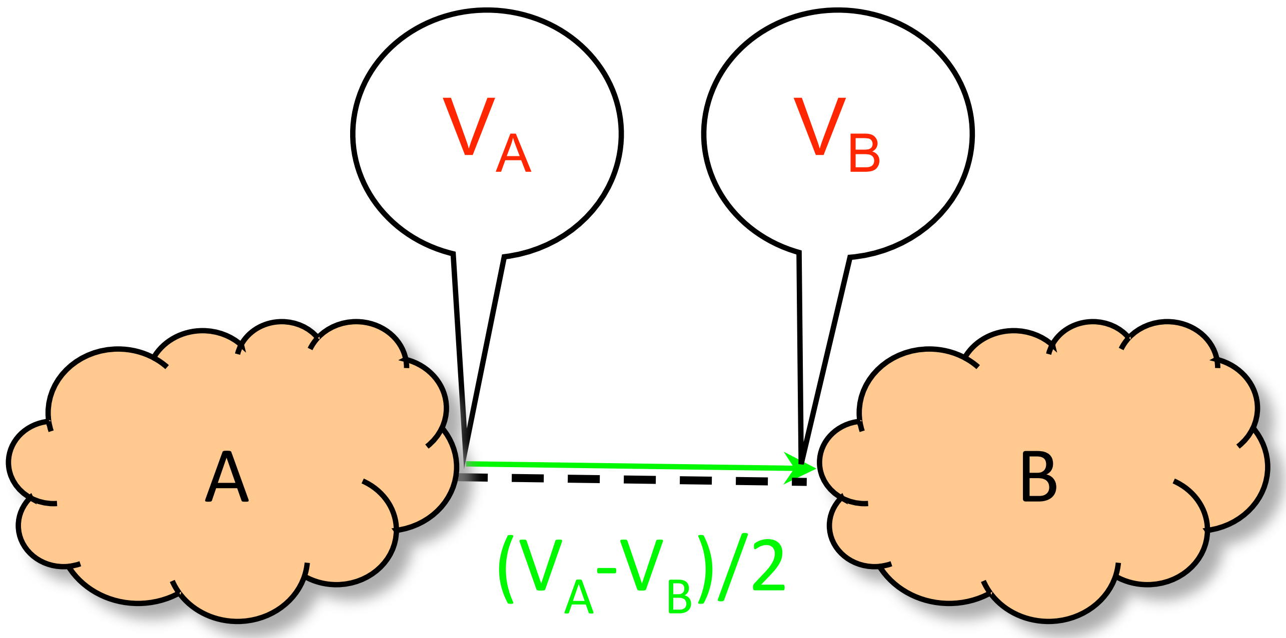

- A and B see values V_A and V_B
- What should be the paid-peering price?
- Fair price is $(V_A-V_B)/2$

The Fair Peering Price

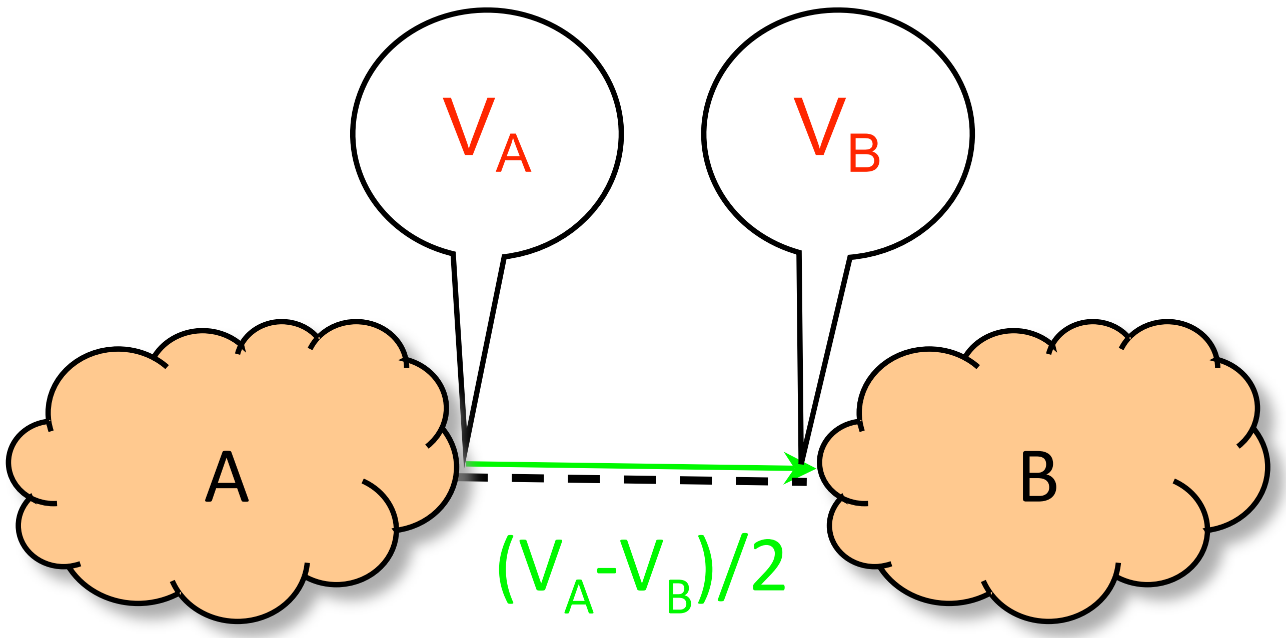

- A and B see values V_A and V_B
- What should be the paid-peering price?
- Fair price is $(V_A-V_B)/2$

The fair price equalizes the benefit that A and B see from the link

Why Peer at the Fair Price?

- Peering with the fair price is optimal
  - Both networks see better fitness by peering at the fair price
  - As compared to the case where peering link does not exist
- Peering with the fair price is stable
  - No network has the incentive to unilaterally de-peer the other
  - Unique Nash Equilibrium
- Optimal and stable as long as $V_A + V_B > 0$
  - Either V_A or V_B can be negative, as long as total is positive
  - For cost-benefit peering, both V_A and V_B must be positive

Some Hard Questions..

- Value-based peering is fair, optimal and stable.
  - But what (if any) is the right notion of fairness?
  - Equal value? Equal cost?
- How can networks estimate peering value?
  - Peering trials..
  - A cost model to estimate peering value
- What if networks lie about peering value?
- What happens if everyone uses value-based peering?
  - Density of peering links, end-to-end path lengths..

Some Hard Questions..

- Value-based peering is fair, optimal and stable.
  - But what (if any) is the right notion of fairness?
  - Equal value? Equal cost?
- How can networks estimate peering value?
  - Peering trials..
  - A cost model to estimate peering value
- What if networks lie about peering value?
- What happens if everyone uses value-based peering?
  - Density of peering links, end-to-end path lengths..

The ITER Model

- ITER: Agent-based computational model to answer "what-if" questions about Internet evolution

- Inputs:
  - Network types: transit provider, content provider, stub
  - Peer selection methods, provider selection methods
  - Geographical constraints
  - Pricing/cost parameters
  - Interdomain traffic matrix

- Output: Equilibrium internetwork topology, traffic flow, per-network fitness

The ITER Approach

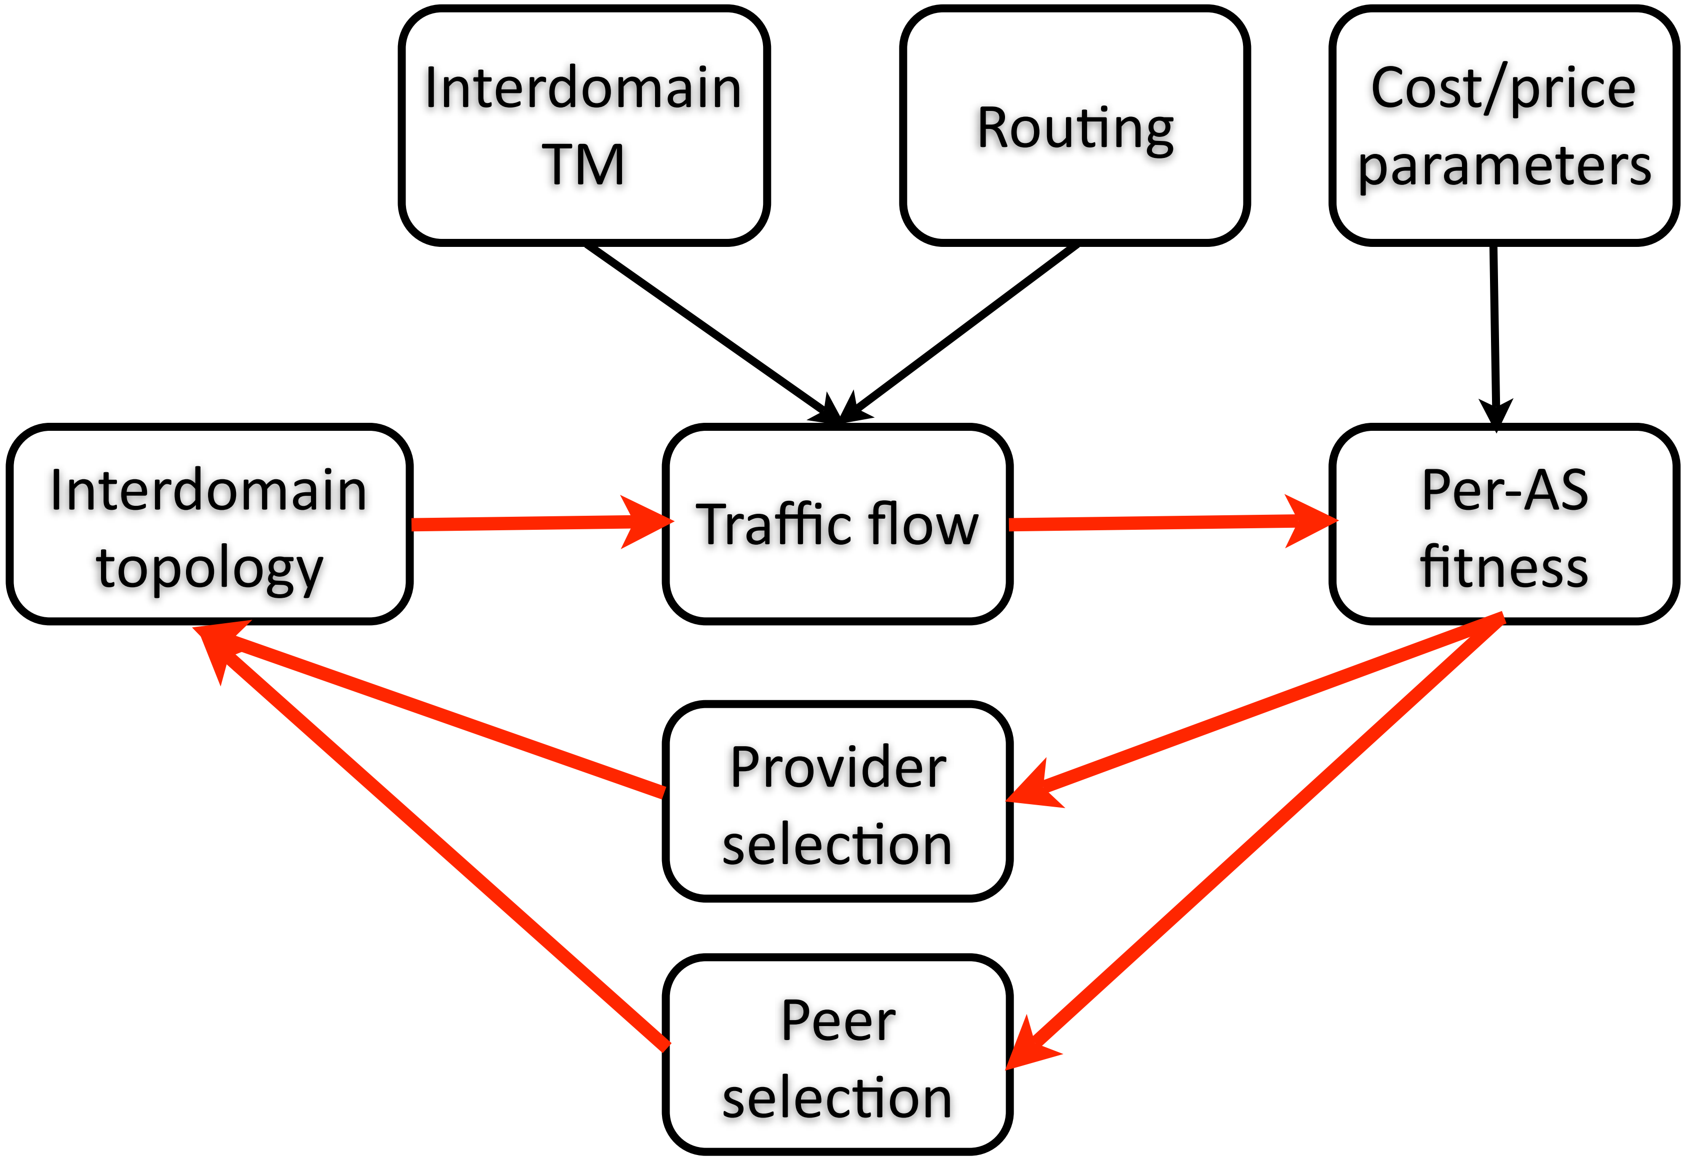

- Compute equilibrium: no network has the incentive to change its providers/peers

Using ITER to simulate value-based peering

- Small but realistic internetwork topology: transit providers, content providers, and stubs
- Interdomain traffic matrix: most traffic flows from content providers to stubs
- Provider selection: price-based
- Peer selection: value-based, cost-benefit and traffic-ratio peering

ITER Results: Value-based Peering

- Higher density of peering links with value-based peering as compared to peering by traffic ratios or peering by cost-benefit analysis
  - Peering links that cannot be formed with cost-benefit analysis are feasible with value-based peering
  - Shorter end-to-end paths
- Payment direction: The same network can end up on either side of a paid-peering relationship
  - What happens in practice?

Thanks!

The papers are online at www.caida.org/~amogh

Feedback, comments, criticism: amogh@caida.org

Evaluation

- Access to routing and traffic data from an access ISP in UK
- No access to backhaul and interconnect cost data
- Considered three cost scenarios:
  - Backhaul >> Interconnect (large ISP or cheap transit scenario)
  - Backhaul ≈ Interconnect
  - Backhaul << Interconnect (content provider or expensive transit scenario)
- Evaluated cost optimization using the greedy approach
- What-if scenarios:
  - Where to peer?
  - Which new peer to establish peering with?
  - How useful is an existing peering relation?

Evaluation

- Access to routing and traffic data from an access ISP in UK
- No access to backhaul and interconnect cost data
- Considered three cost scenarios:
  - Backhaul >> Interconnect (large ISP or cheap transit scenario)
  - Backhaul ≈ Interconnect
  - Backhaul << Interconnect (content provider or expensive transit scenario)
- Evaluated cost optimization using the greedy approach
- What-if scenarios:
  - Where to peer?
  - Which new peer to establish peering with?
  - How useful is an existing peering relation?

Evaluating Current Peers

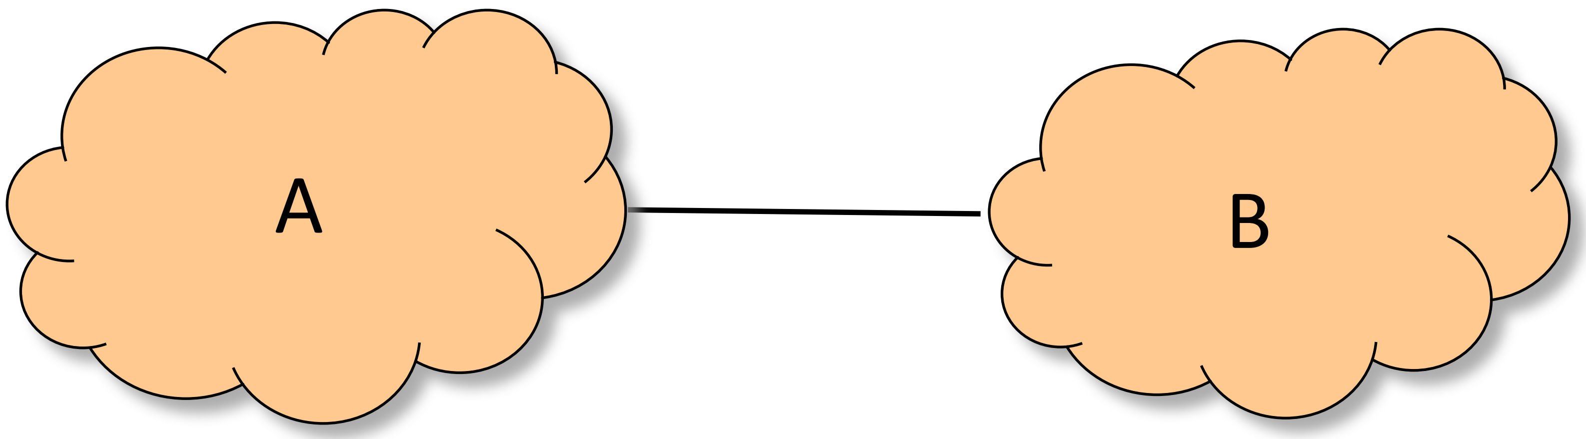

2

Evaluating Current Peers

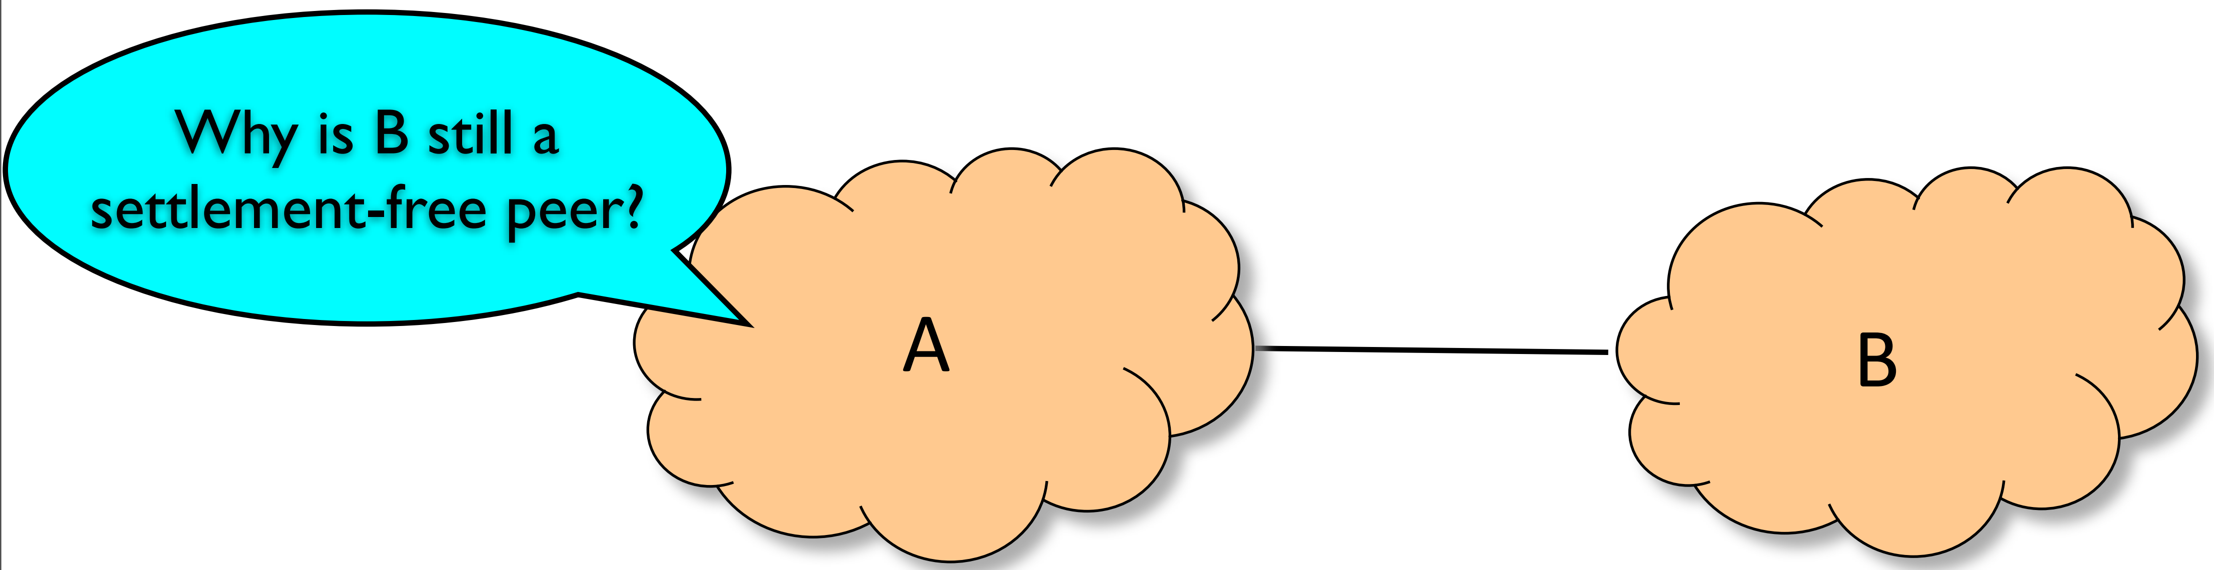

2

Evaluating Current Peers

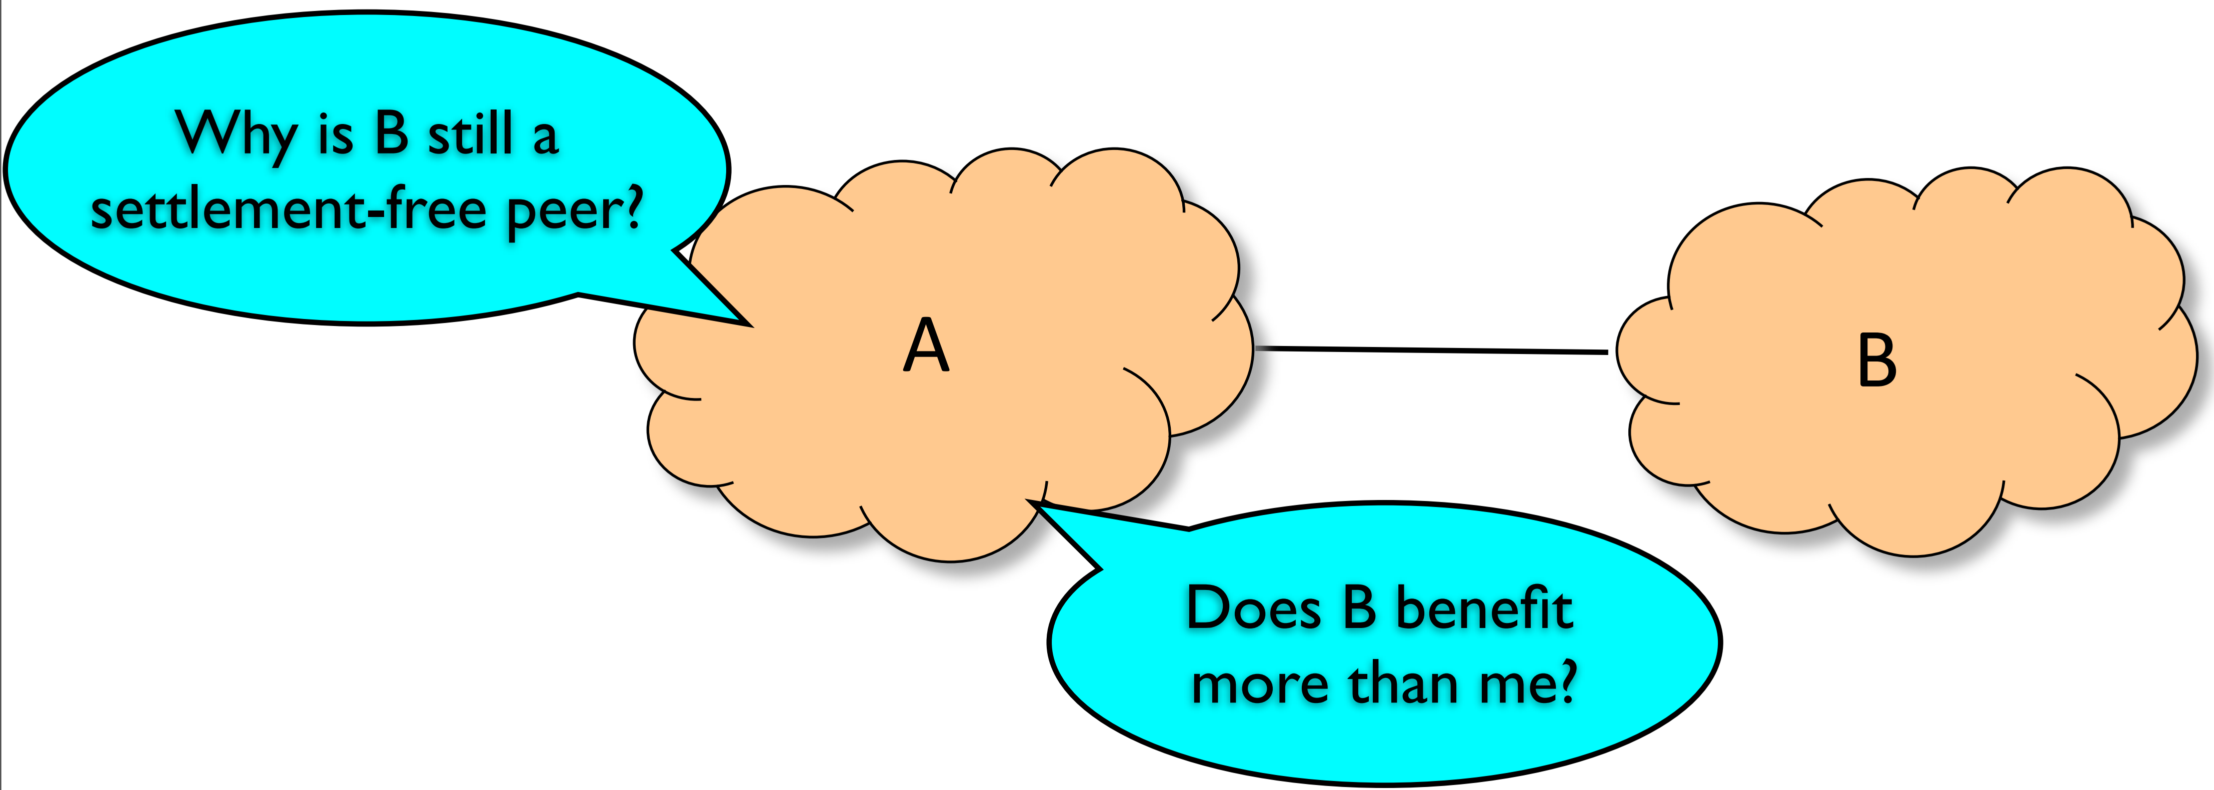

2

Evaluating Current Peers

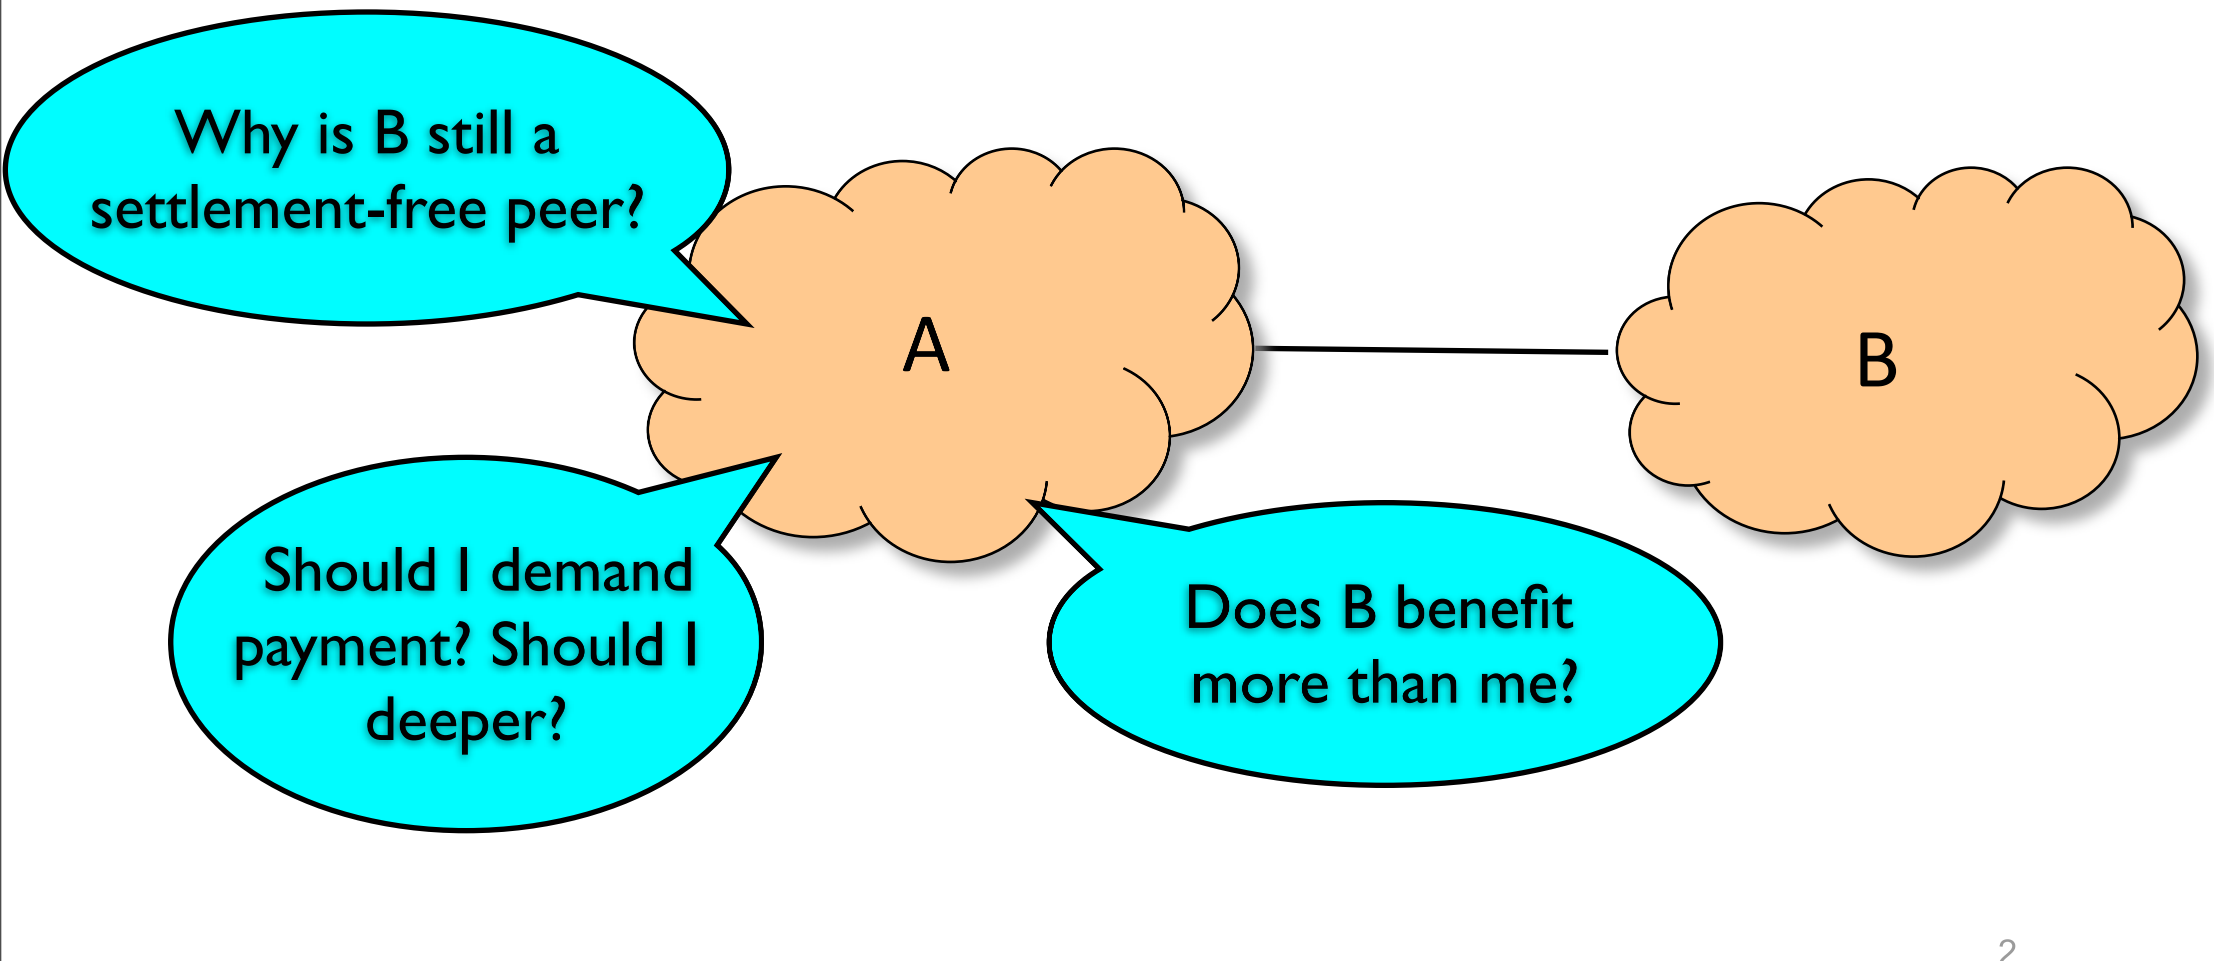

Negative Peering Value

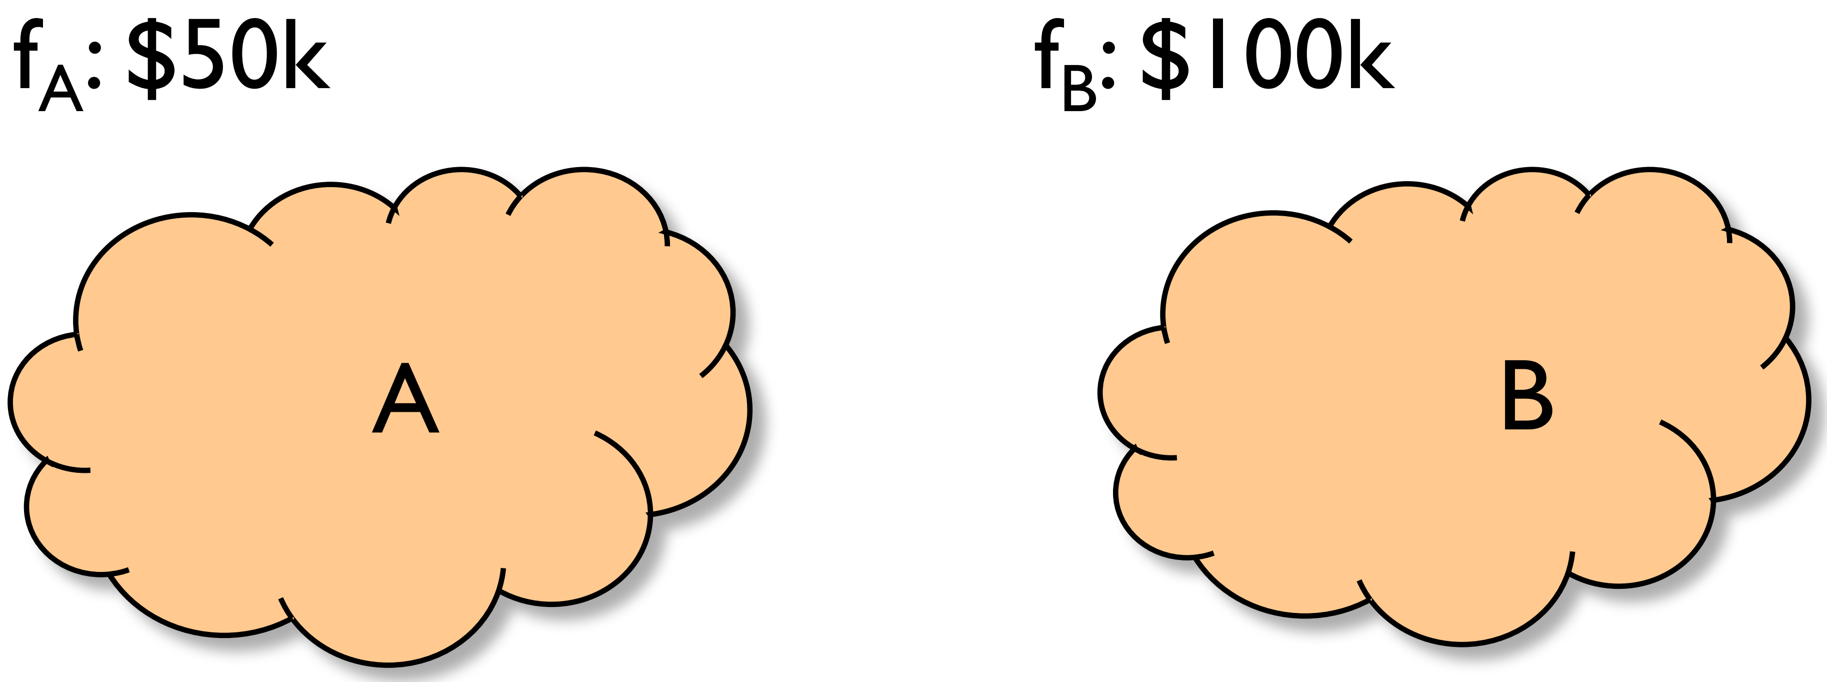

Negative Peering Value

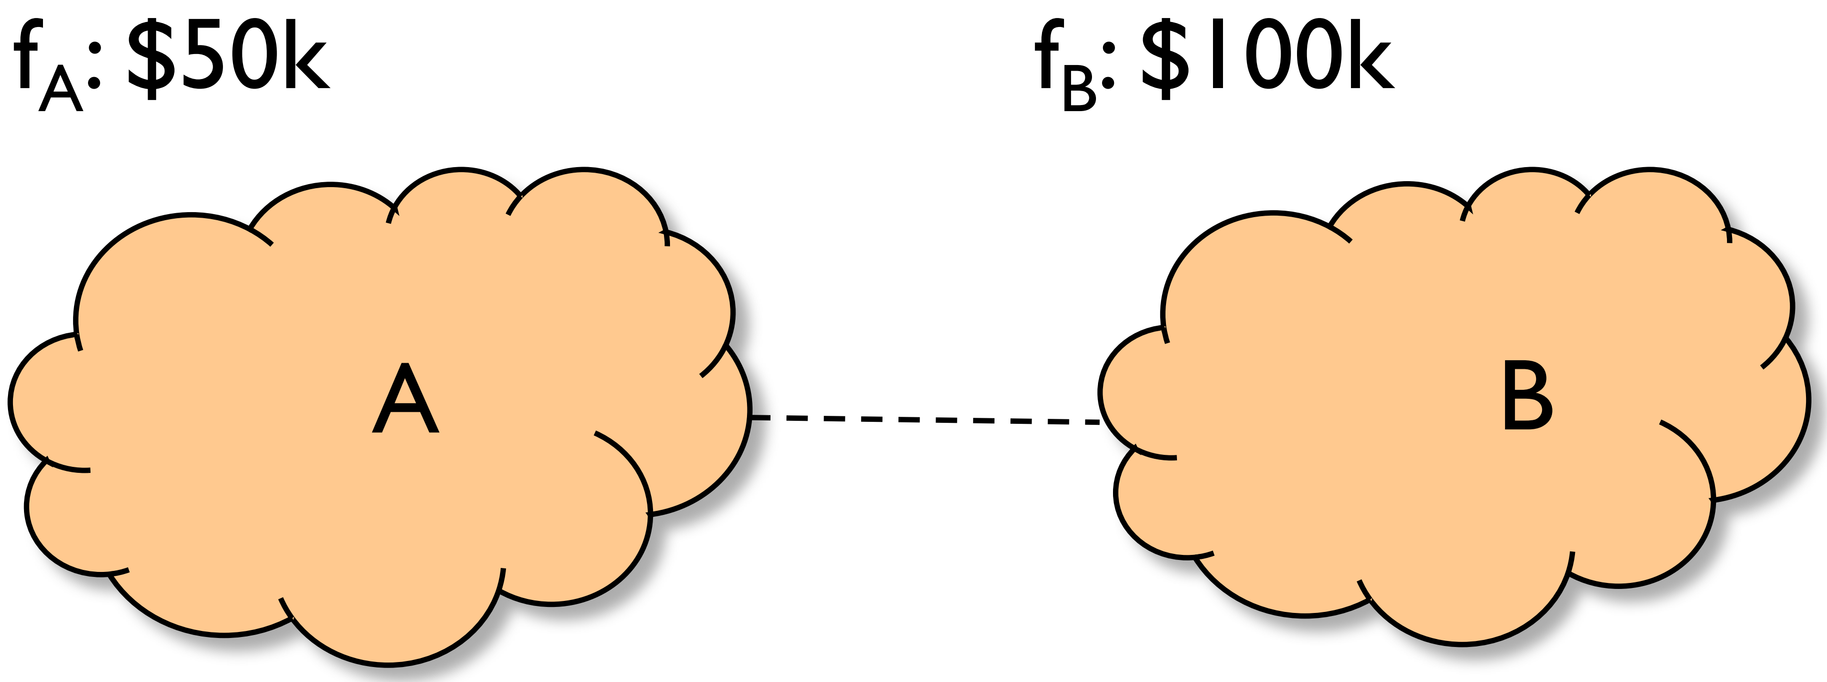

Negative Peering Value

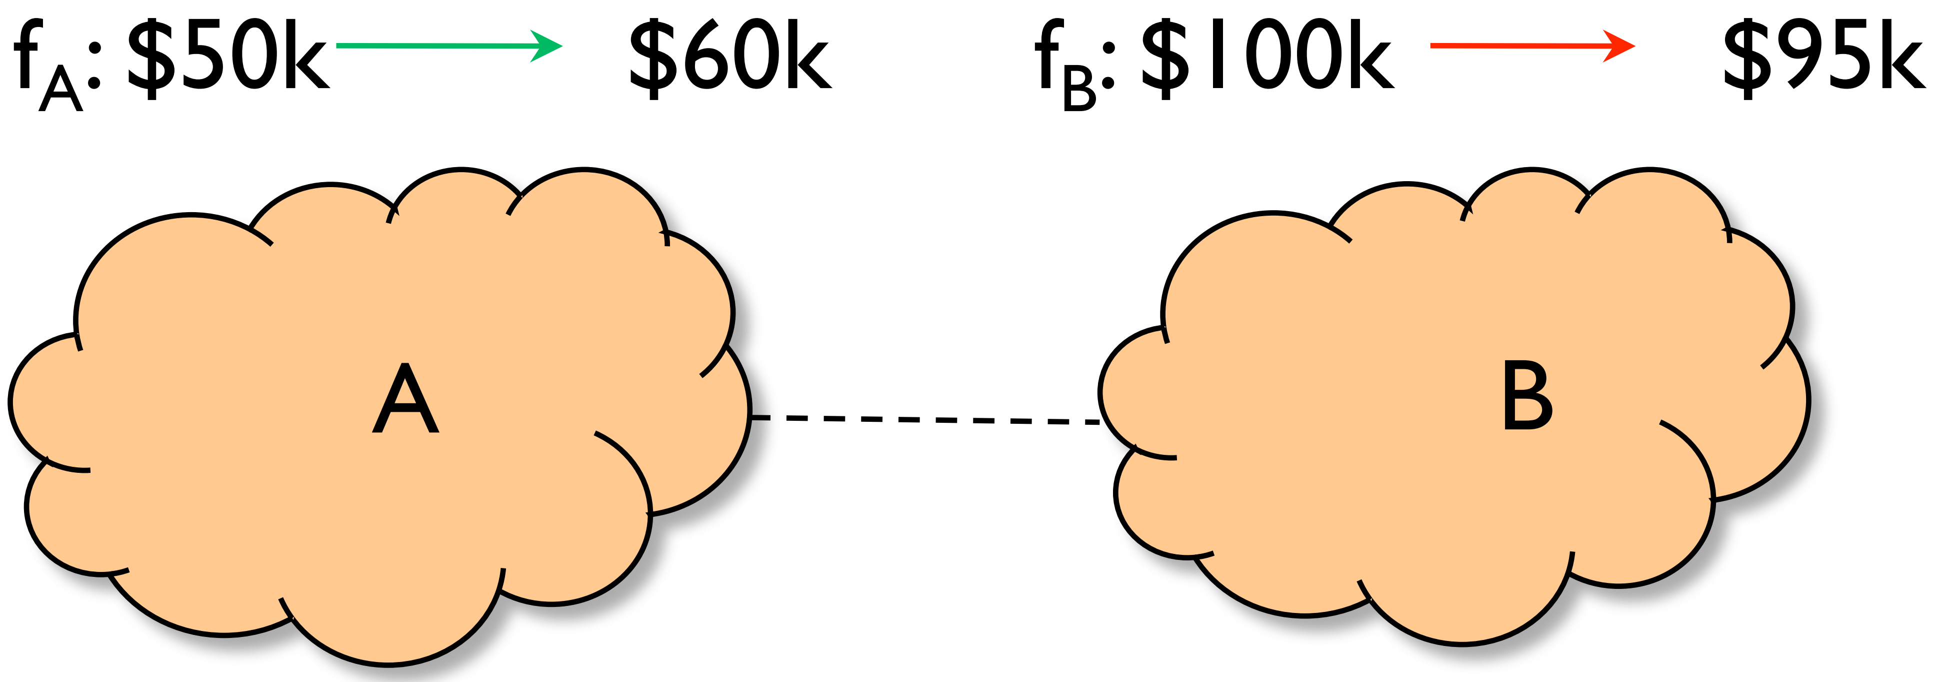

Negative Peering Value

f_A: \$50k → \$60k f_B: \$100k → \$95k

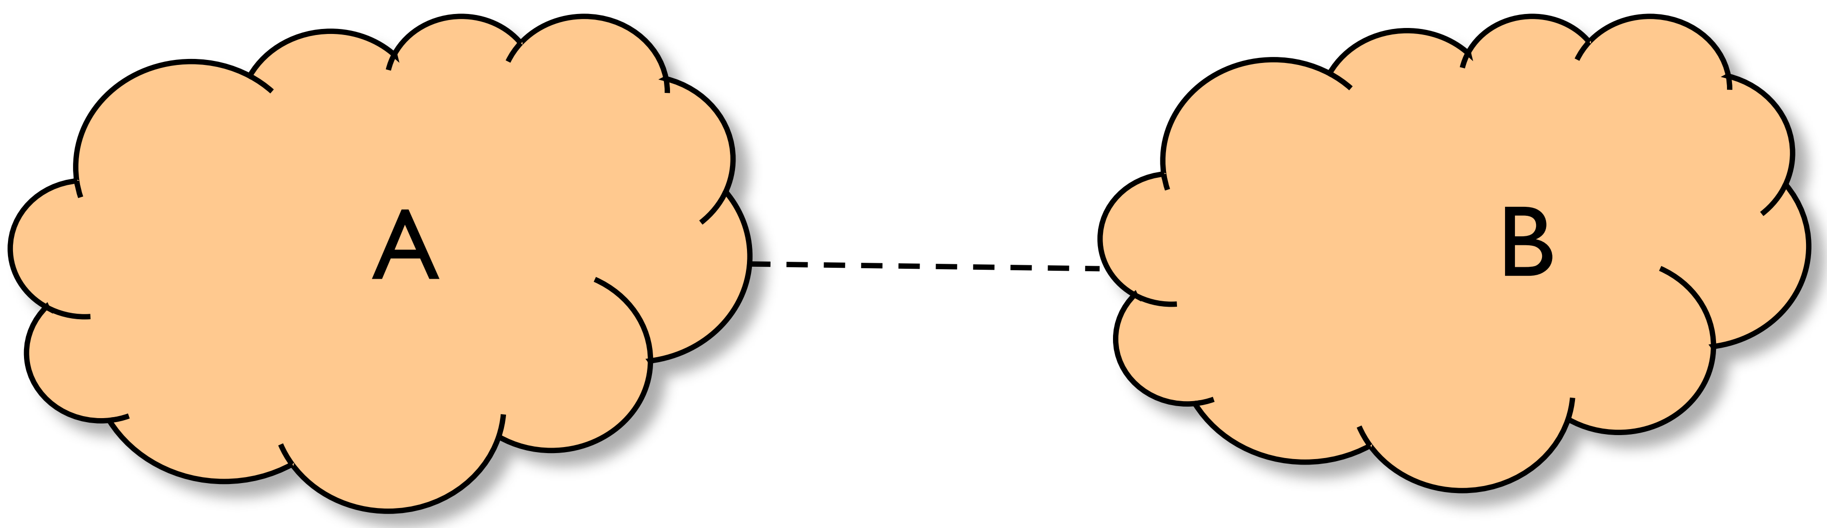

V_A=\$10k V_B=-\$5k

Negative Peering Value

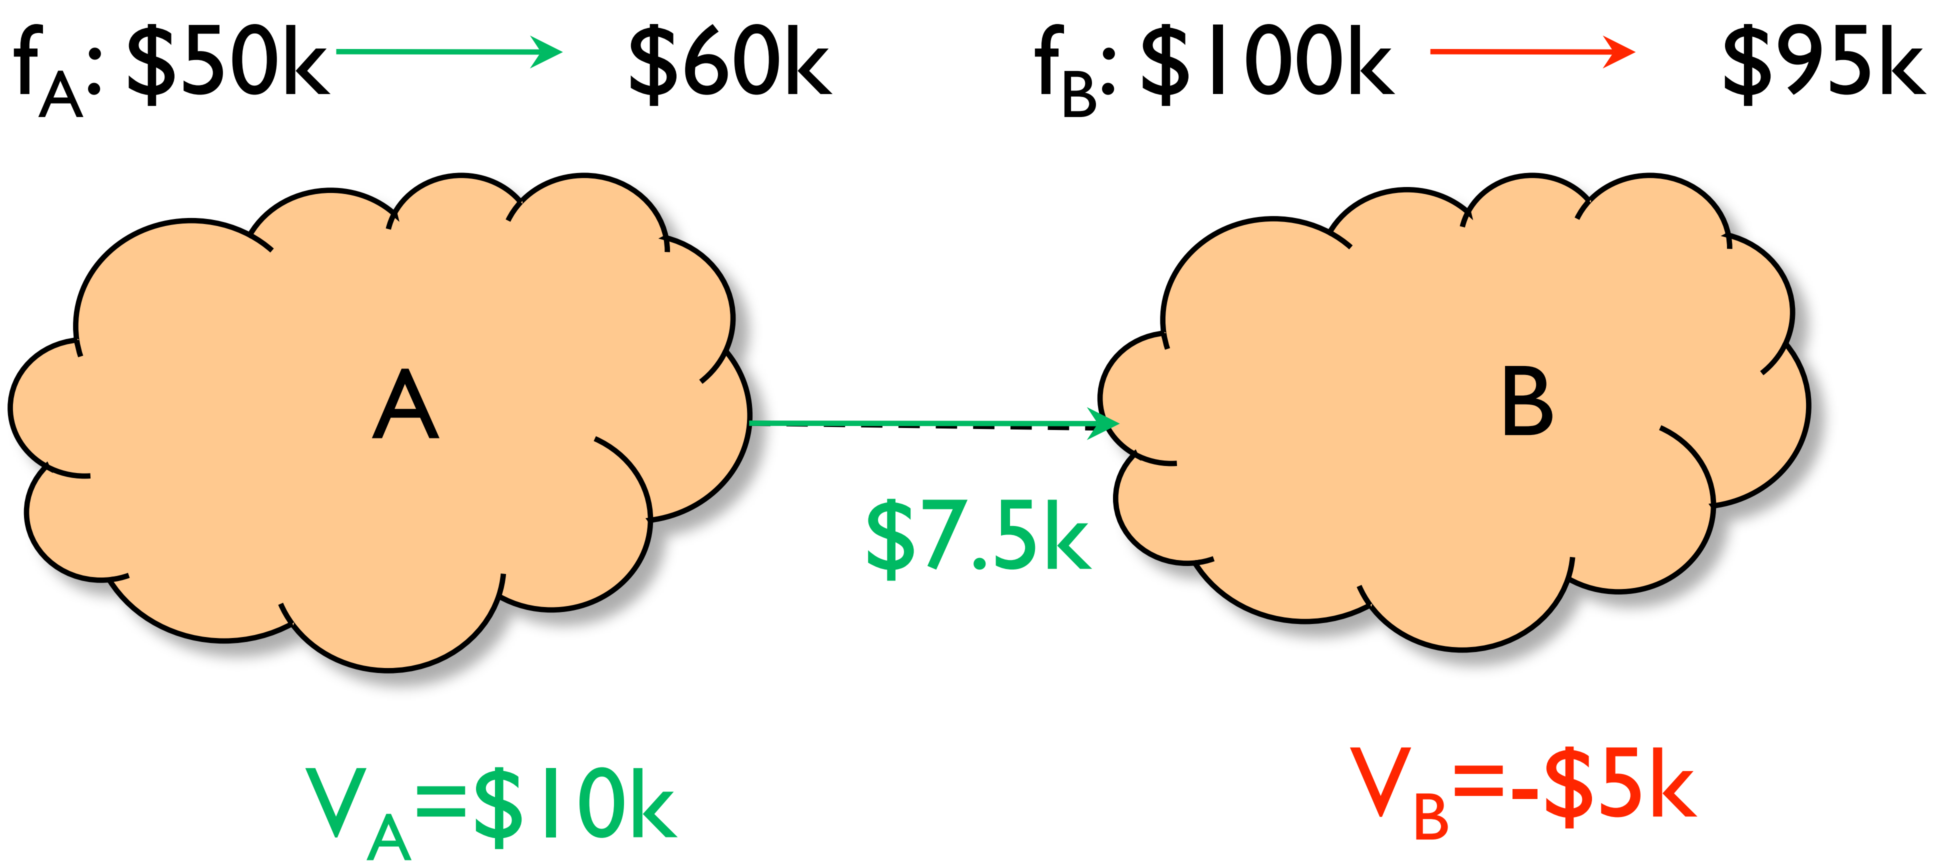

Negative Peering Value

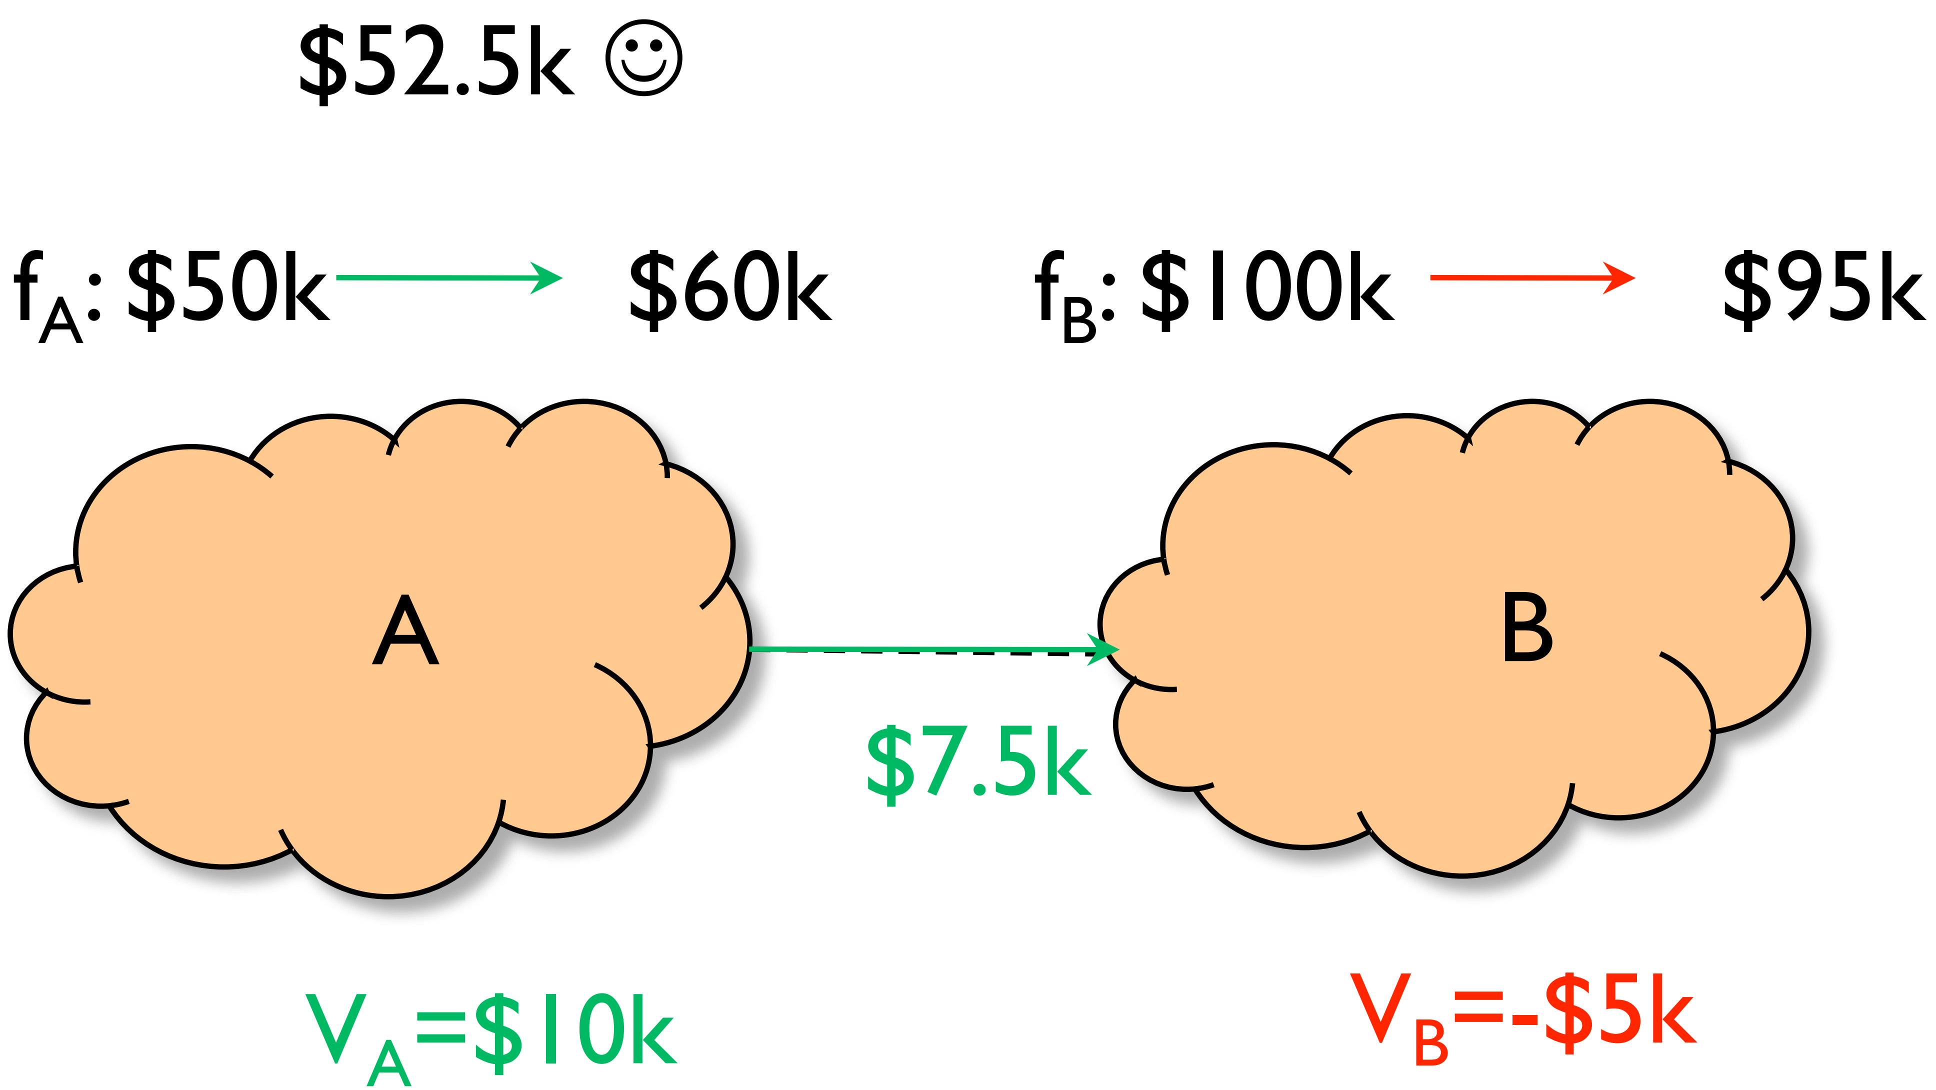

Negative Peering Value

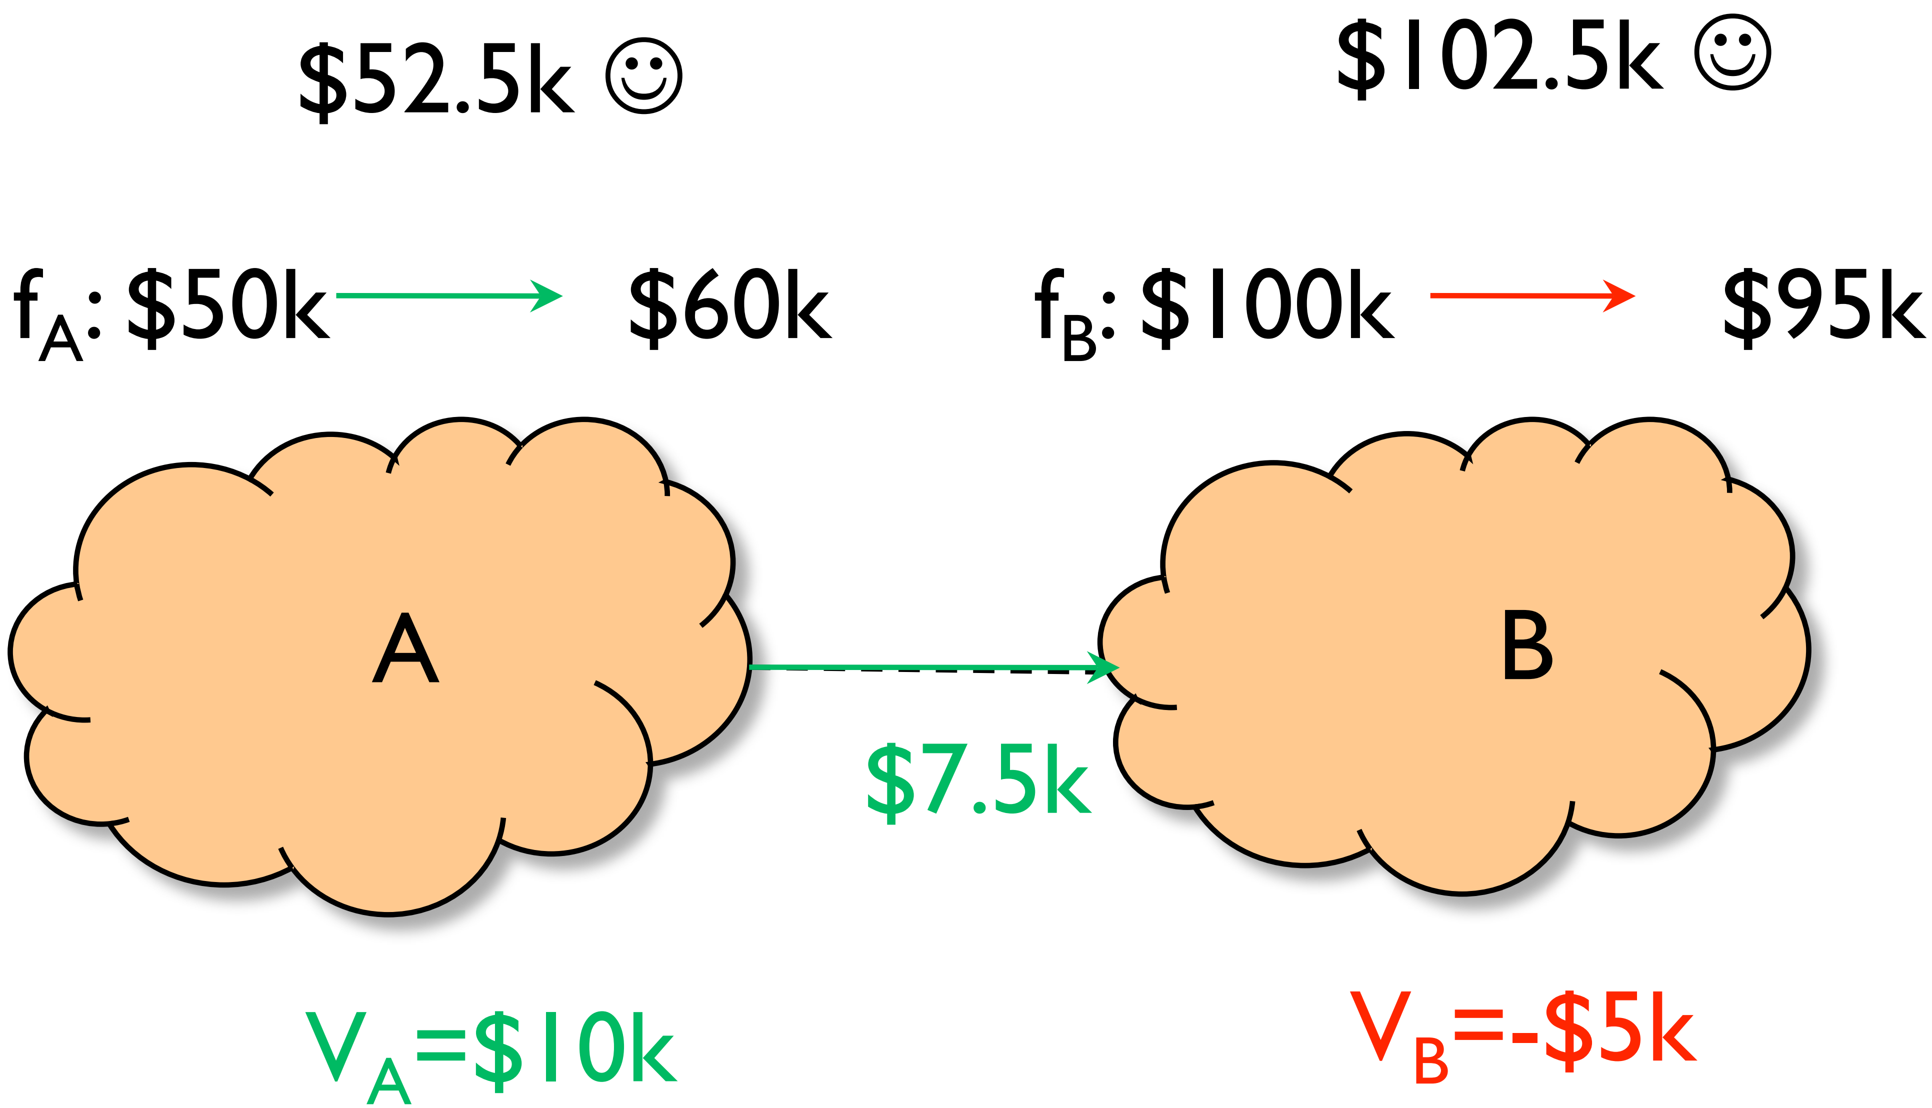

Hiding peering value

- Assume true $V_A + V_B > 0$ and $V_B > V_A$
  - A should get paid $(V_B - V_A)/2$

- If A estimates V_B correctly, and claims its peering value is V_L, where $V_L << V_A$
  - B is willing to pay more: $(V_B - V_L)/2$:)

- If A doesn't estimate V_B correctly, and $V_L + V_B < 0$, the peering link is not feasible!
  - A loses out on any payment :(