Inter-Domain Routing: Stability, Policies, and Incentives

Vijay Ramachandran
http://www.icsi.berkeley.edu/~vijayr

Supported by the DoD URI program
under ONR grant N00014-01-1-0795

Overview of Results: Two SPYCE-Sponsored Collaborations

- Aaron Jaggard
  - ✓ Framework for design and analysis of inter-domain routing protocols
  - NEW Expanded to include more complex policies

- Joan Feigenbaum and Michael Schapira
  - NEW New class of policies for BGP-based, incentive-compatible routing; partially answers goal for option period

Inter-Domain Routing

Establish routes between autonomous systems (ASes).

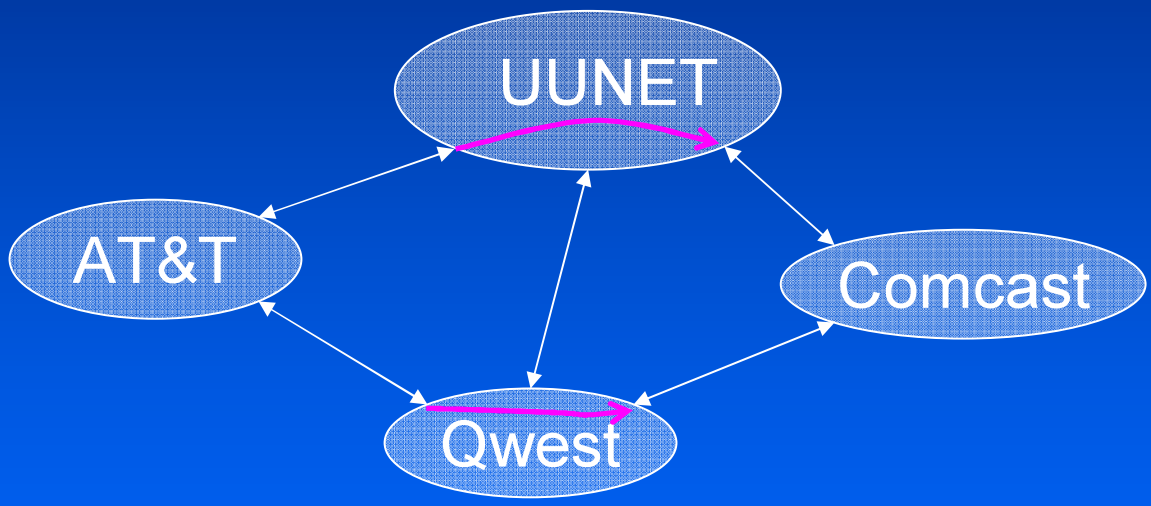

Currently done with the Border Gateway Protocol (BGP).

Why is Inter-Domain Routing Hard?

- Route choices are based on *local policies.*
- Autonomy: Policies are uncoordinated.
- Expressiveness: Policies are complex.

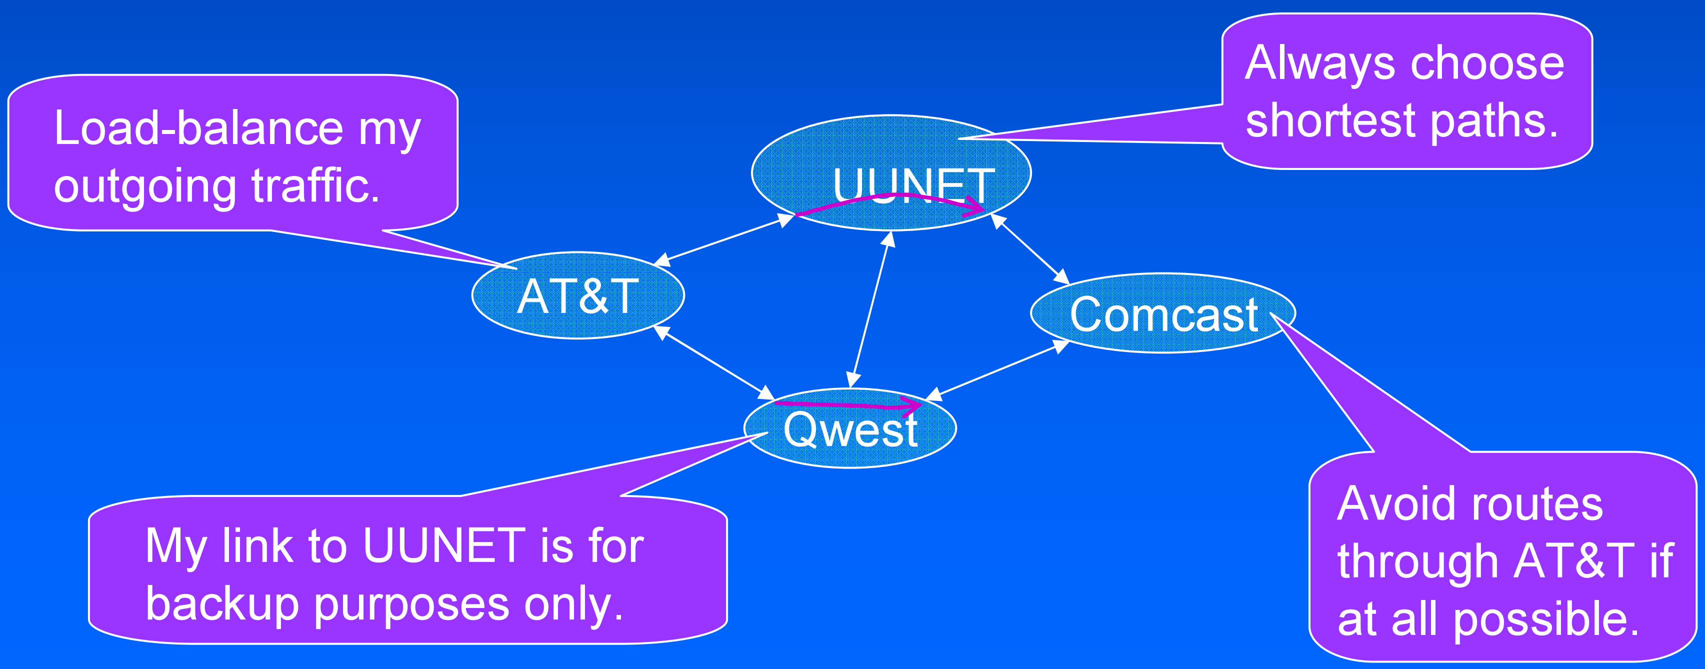

Three Desiderata

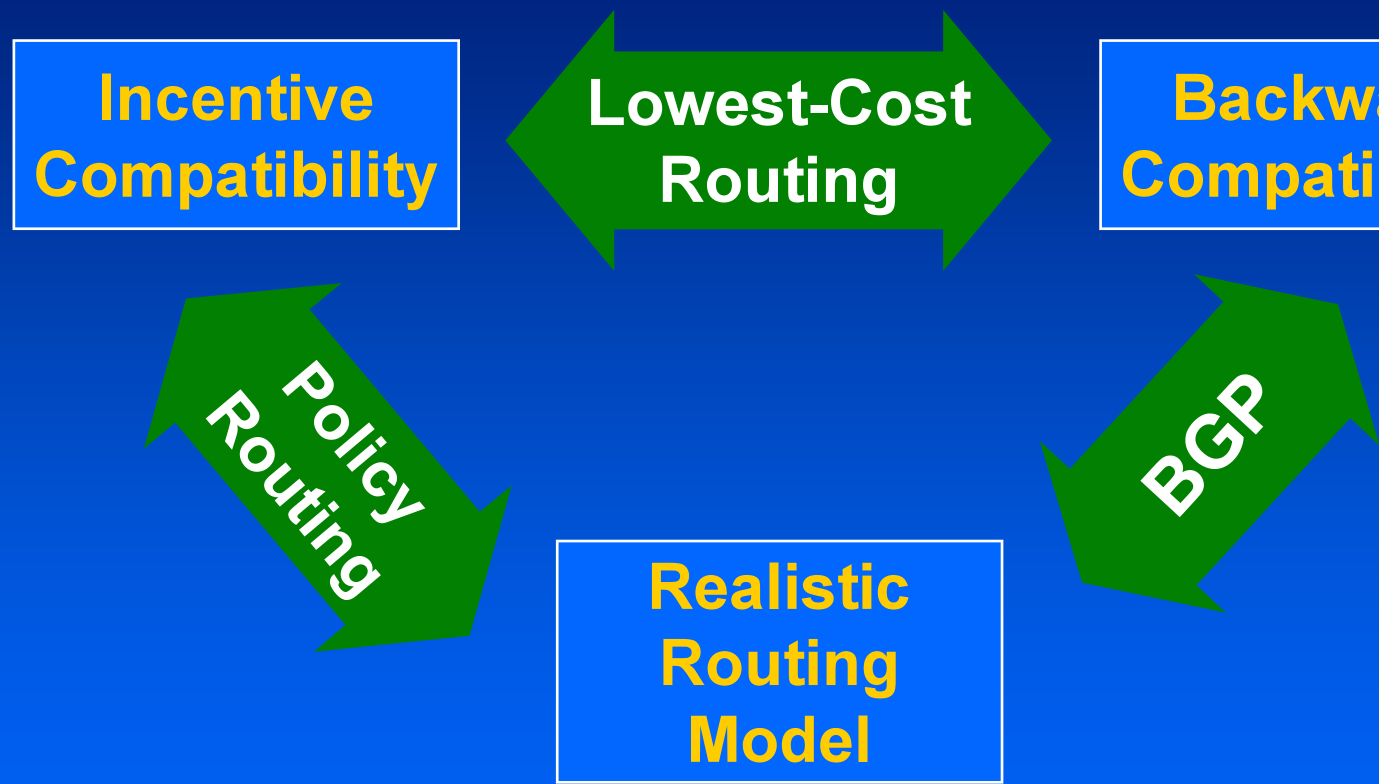

Can get 2 out of 3
2 are SPYCE achievements

Open question for option:
Can we satisfy all 3?

Yes!

Realistic Routing Model

- Next-hop routing with today's Internet business relationships.

Incentive Compatibility

- Optimal routing tree given nodes' valuations; incentivize nodes to report truthful valuations using VCG payments.

Backward Compatibility

- Compute routes and payments in a BGP-compatible way.

BGP Route Processing

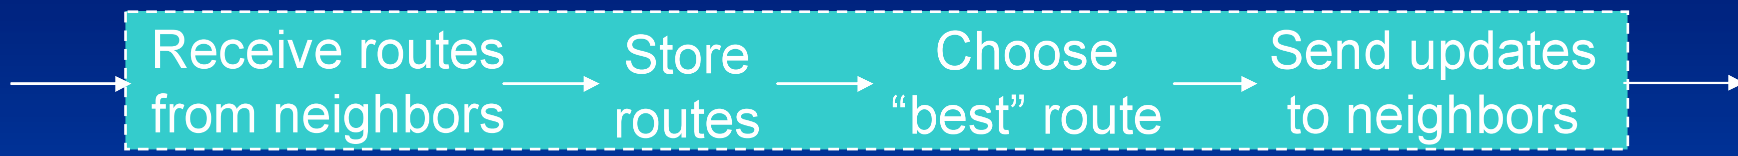

- Pros:
  - Route choices depend on neighbors' choices.
    => enforces consistency
  - Best-route choices are made locally.
    => allows autonomy & expressiveness.
  - Routes are loop free and can change with topology, without any node's knowing the whole network.
- Cons:
  - Policy-induced routing anomalies
    => Routes may not be stable.

SPYCE Contributions: Policy Constraints for Stability

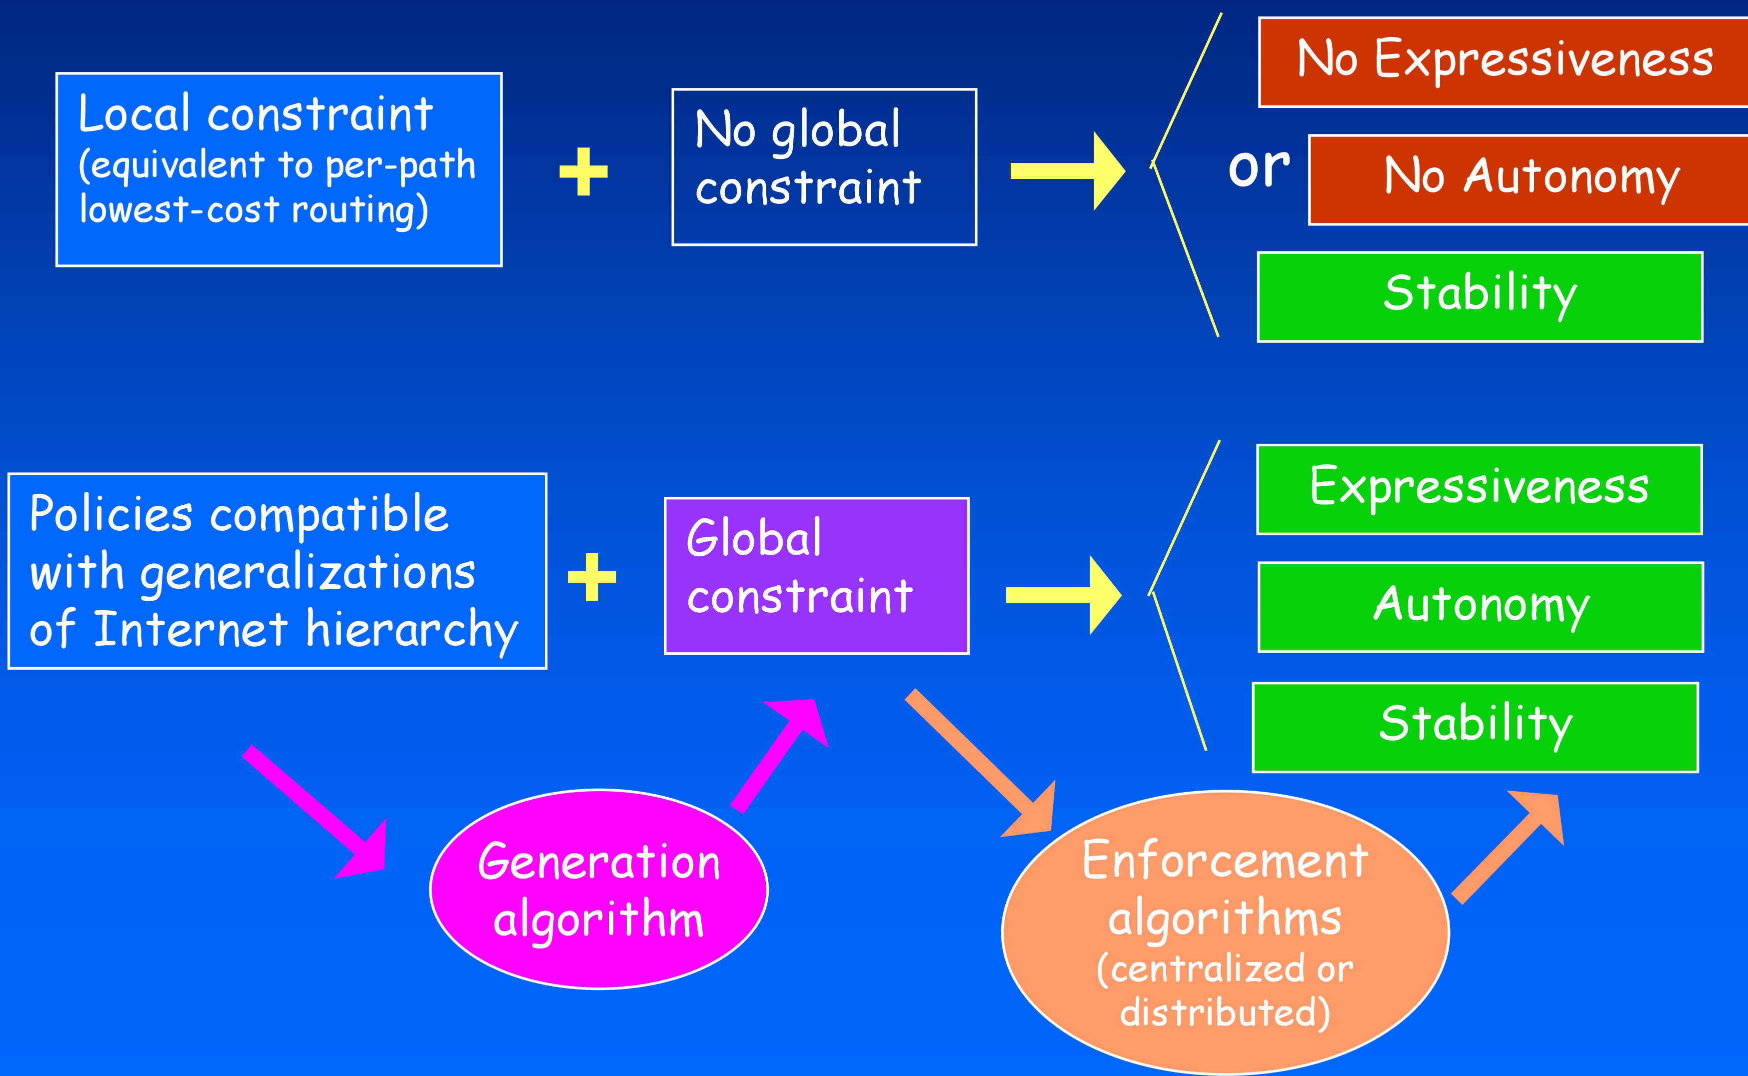

Expanding the Model

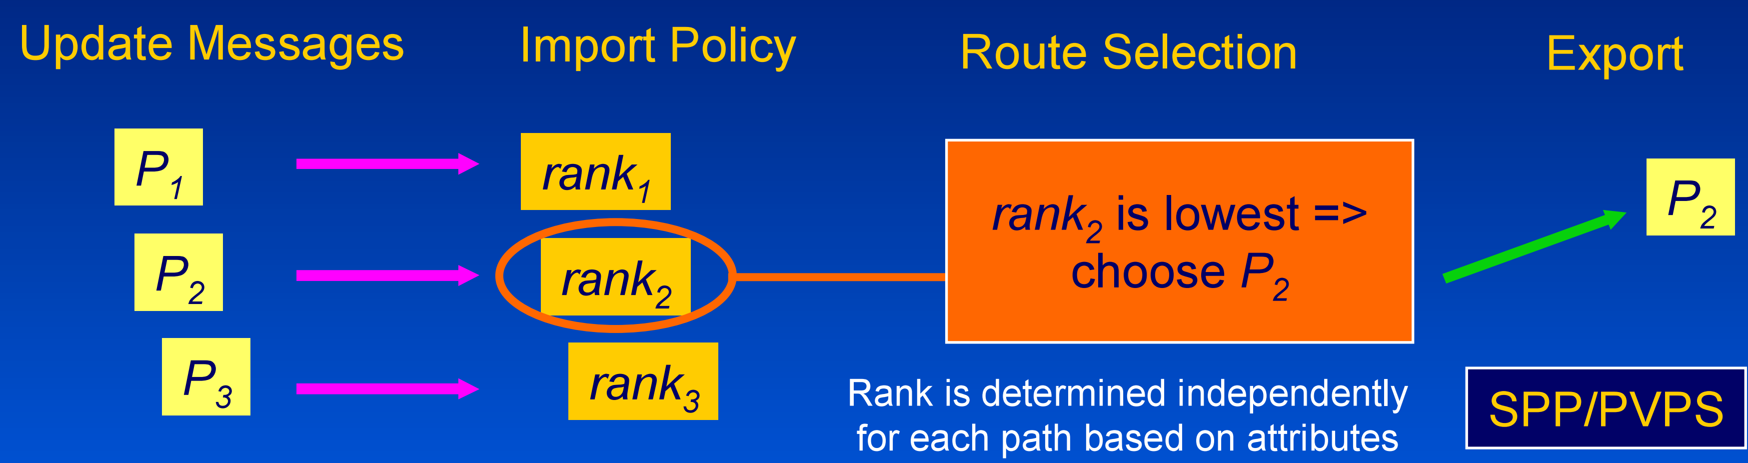

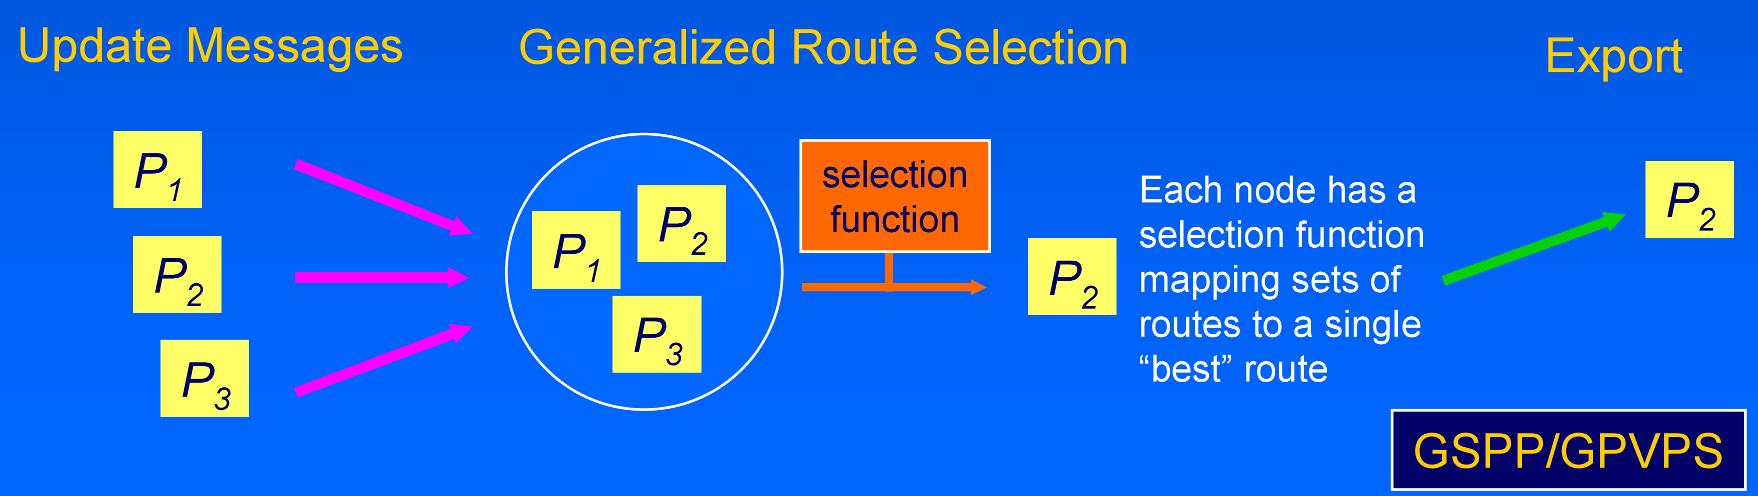

Results about GSPP / GPVPS

- Ability to model iBGP and complex BGP features, *e.g.*, MED
- Analogous sufficient conditions for stability in the general case, *e.g.*, the absence of *generalized dispute wheels*
- Improved space efficiency for methods to fix oscillations in BGP

Welfare-Maximizing Routing

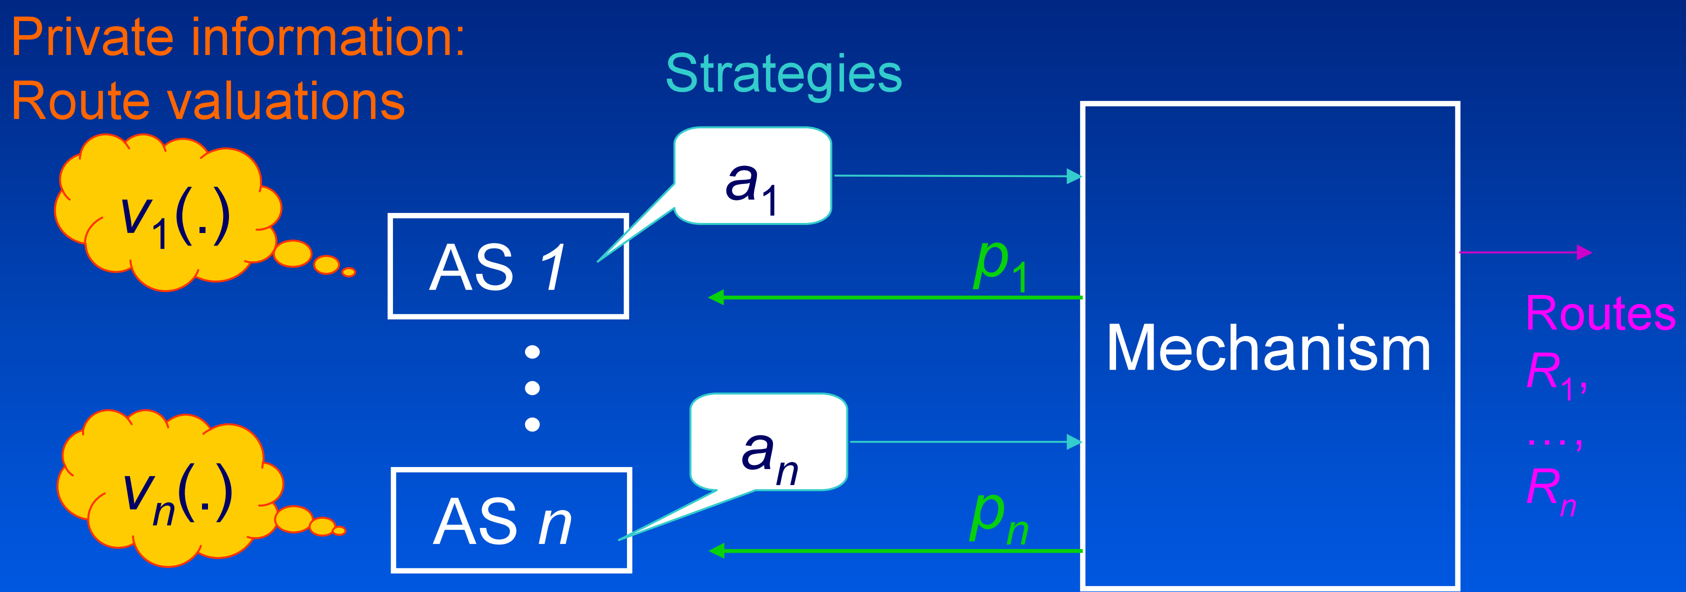

- Maximize sum of nodes' valuations = $\sum_i v_i(R_i)$.
- *A confluent routing tree and payments are computed in parallel for each destination.*
- Source nodes are paid for their contribution to the routing tree.
- We want a BGP-style algorithm that computes routes *and* payments.

SPYCE Results: Welfare Maximization and Inter-Domain Routing

Routing-Policy Class	Good Centralized Algorithm?	Good Distributed Algorithm?
LCP	✓	✓
General Policy	✗ (and hard to approximate)	✗ (and hard to approximate)
Next Hop	✓	✗
Subjective Cost	✗ (incl. some special cases)	✗ (approx. is easy if >1 tree)
Forbidden Set	✗	✗

General Approach

- Find a class of policies for which BGP converges to an optimal tree T (*i.e.*, one that maximizes the sum of the valuations of all source nodes).

- Use VCG payment formula to ensure truthfulness, *i.e.*, payment to node k is

$$p_k = \sum_{i \neq k} v_i(T) - h_k(\cdot)$$

where h_k is a function that does not depend on node k's valuation.

Gao-Rexford Framework (1)

Neighboring pairs of ASes have one of:

- a *customer-provider* relationship
  (One node is purchasing connectivity from the other node.)
- a *peering* relationship
  (Nodes have offered to carry each other's transit traffic, often to shortcut a longer route.)

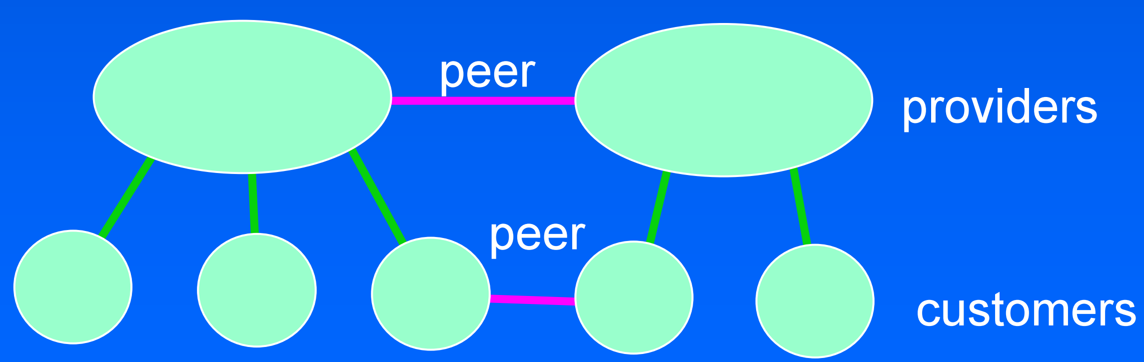

Gao-Rexford Framework (2)

- Global constraint: **no customer-provider cycles**
- Local *preference* and *scoping* constraints, which are consistent with Internet economics:

Preference Constraints

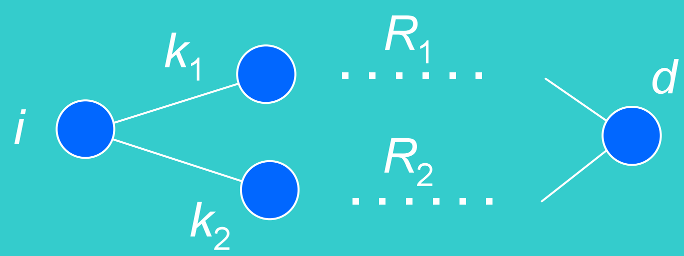

- If k_1 and k_2 are both customers, peers, or providers of *i*, then either ik_1R_1 or ik_2R_2 can be more valued at *i*.
- If one is a customer, prefer the route through it. If not, prefer the peer route.

Scoping Constraints

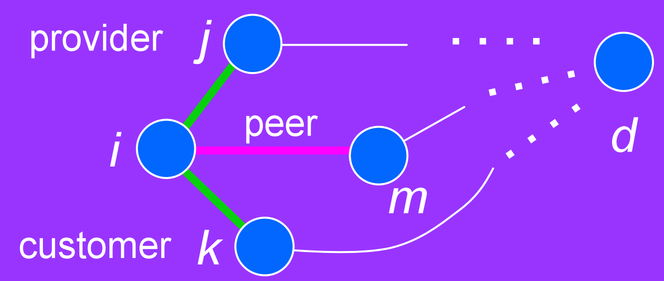

- Export customer routes to all neighbors and export all routes to customers.
- Export peer and provider routes to all customers only.

- Gao-Rexford conditions => no dispute [GR01, GGR01]

Realistic Policies for Incentive-Compatible Routing

- **Next-hop valuations:** The valuation of a route depends only on its next hop.

- **Theorem:** If ASes obey Gao-Rexford conditions and next-hop valuations:
  - BGP computes the optimal (welfare-maximizing) routing tree.
  - VCG payments can be computed along with routes, with only small increase in routing-table size (like LCP).

Summary of Sufficient Conditions

No assumptions	→	Hard (BGP may not converge)
No dispute cycle	→	Non-optimal (BGP will converge, but the solution may be arbitrarily far from optimal.)
No dispute cycle and policy consistency	→	Optimal convergence (but payment computation might be highly space-consuming)
No dispute cycle and next-hop valuations (special case of policy consistency)	→	Optimal convergence and good BGP-style algorithm (Requires $\mathcal{O}(1)$ additional space per transit node.)