


## Inter-Domain Routing


Arvind Krishnamurthy  
Fall 2003

## A Logical View of the Internet

- Hierarchical ordering of autonomous systems
  - Area routing
    - Details of Internals not necessary for routing
    - Provides scalability but might degrade performance
- Relationships between autonomous systems:
  - Transit relationship: paying relationship
  - Peering relationship: non-paying one
    - Avoid traffic through a transit relationship
    - Find lower latency paths




## Economics of the Internet




- Transit relationship:
  - Reveal customers to everyone
  - Reveal all known paths to customers
  - ISP Y will advertise paths to C2 and C3

## Peering Relationships



- ISP Y peers with ISP Z
  - Y informs Z of C2 and C3
  - Z informs Y of C4

## Peering Relationships



- ISP Y peers with ISP X
  - Y informs X of C2 and C3
  - X informs Y of C1
  - Y does not tell X of the path to C4
  - Y needs to have transit relationship with P to route to C4

## Multi-Exit Discriminator

- C has a transit relationship with P
- C needs to send a packet into P's domain
  - Packet origin: Boston, packet destination: SF
  - It has two choices: transit into P earlier (near Boston) or transit later (near SF)
  - Prefer early transit



## Routing Preferences

- In general, customer > peer > provider
  - Use LOCAL PREF to ensure this
  - Routing through customer and peer could lower latency and lower traffic costs
- Processing order of path attributes:
  - Select route with highest LOCAL-PREF
  - Select route with shortest AS-PATH
  - For routes learned from same neighbor:
    - Apply multi-exit discriminator


## Routing algorithms for the internet

- Link state or distance vector?
- Problems with link state:
  - LS database too large – entire Internet
  - Metric used by routers not the same
  - May expose policies to other AS's
- Can we use distance-vector algorithms for policy routing?


## Announcements

- Reminder:
  - First quiz will be held next Monday
- Project dynamics will be announced later this week


## The Bouncing Effect




## C Sends Routes to B




## B Updates Distance to A



### B Sends Routes to C




### C Sends Routes to B



### Solutions

- Problems arise:
  - When metric increases
  - Implicit path has loops
- Solution 1: If metric increases, delay propagating information
  - In our example, B delays advertising route
  - C eventually thinks B's route is gone, picks its own route
  - B then selects C as next hop
- Adversely affects convergence
- Other "Solutions":
  - Split horizon: C does not advertise route to B
  - Poisoned reverse: C advertises route to B with infinite distance
- Works for two node loops
  - Does not work for loops with more nodes

### Example Where Split Horizon Fails



- When link breaks, C marks D as unreachable and reports that to A and B
- Suppose A learns it first
  - A now thinks best path to D is through B
  - A reports D unreachable to B and a route of cost=3 to C
- C thinks D is reachable through A at cost 4 and reports that to B
- B reports a cost 5 to A who reports new cost to C
- etc...

### Solution: Distance Vector with Path

- Each routing update carries the entire path
- Loops are detected as follows:
  - When AS gets route check if AS already in path
    - If yes, reject route
    - If no, add self and (possibly) advertise route further
- Advantage:
  - Metrics are local - AS chooses path, protocol ensures no loops
- Hop-by-hop Model
  - BGP advertises to neighbors only those routes that it uses
    - Consistent with the hop-by-hop Internet paradigm
    - e.g., AS1 cannot tell AS2 to route to other AS's in a manner different than what AS2 has chosen (need source routing for that)

### What problem is BGP solving?

Underlying problem	Distributed means of computing a solution.
Shortest Paths	BF: RIP, OSPF ...
X?	BGP

#### Having an X can

- aid in the design of policy analysis algorithms and heuristics,
- aid in the analysis and design of BGP and extensions,
- help explain some BGP routing anomalies,
- provide a fun way of thinking about the protocol

### Stable-Paths Problem

Input:

- A graph of nodes and edges,
- Node 0, called *the origin*,
- For each non-zero node, a set or permitted paths to the origin. This set always contains the "null path".
- A ranking of permitted paths at each node. Null path is always least preferred. (Not shown in diagram)

most preferred  
...  
least preferred (not null)

### Solution to SPP

A **solution** is an assignment of permitted paths to each node such that

- node  $u$ 's assigned path is either the null path or is a path  $uwP$ , where  $wP$  is assigned to node  $w$  and  $(u,w)$  is an edge in the graph,
- each node is assigned the highest ranked path among those consistent with the paths assigned to its neighbors.

A Solution need not represent a shortest path tree, or a spanning tree.

### SPP may have multiple solutions

DISAGREE      First solution      Second solution

### "Route Triggering"

Remove primary link      Restore primary link

### BAD GADGET: No solution

### PRECARIOUS: Has solution, but can get trapped

This part has a solution only when node 1 is assigned the direct path (1 0).

As with DISAGREE, this part has two distinct solutions

### Solving SPP

Just enumerate all path assignments  
And check stability of each....

**Exponential complexity !!**

But, in worst case you (probably)  
can't do any better...

### 3-SAT

**Variables**  $V = \{X_1, X_2, \dots, X_n\}$

**Clauses**  $C_1 = X_{17} \text{ or } \sim X_{23} \text{ or } \sim X_3,$   
 $C_2 = \sim X_2 \text{ or } X_3 \text{ or } \sim X_{12}$   
 ....  
 $C_m = X_6 \text{ or } \sim X_7 \text{ or } X_{18}$

**Question** Is there an variable assignment  
 $A : V \rightarrow \{\text{true}, \text{false}\}$  such that  
 each clause  $C_1, \dots, C_m$  is true?

**3- SAT is NP-complete**

### Assignment to variable X

$X \bar{X} 0$   
 $X 0$

$\bar{X} X 0$   
 $\bar{X} 0$

$X = \text{false}$

$X = \text{true}$

### Solvability is NP-Complete

$C X_7 0$   
 $C X_5 0$   
 $C X_3 0$

**BAD GADGET**

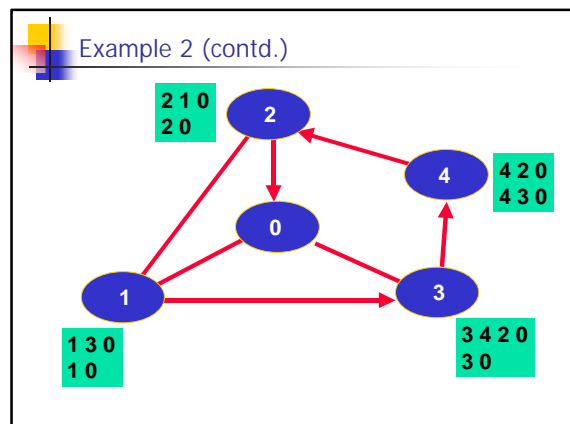
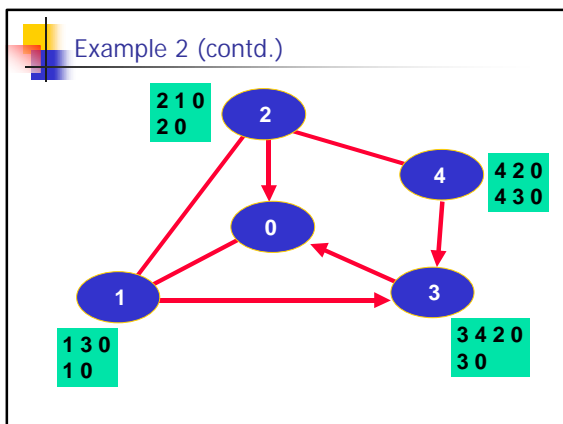
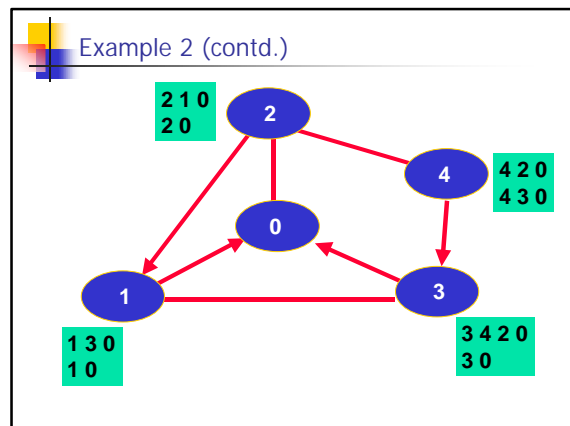
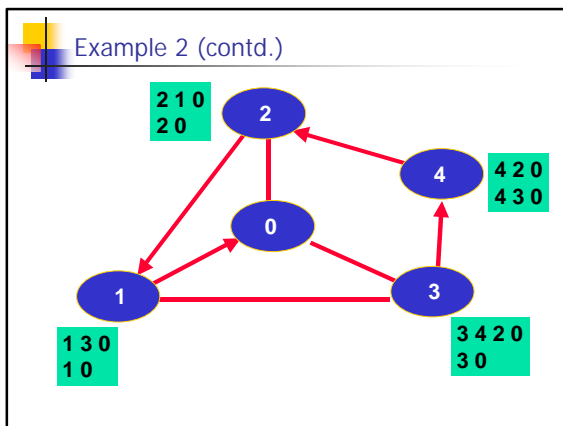
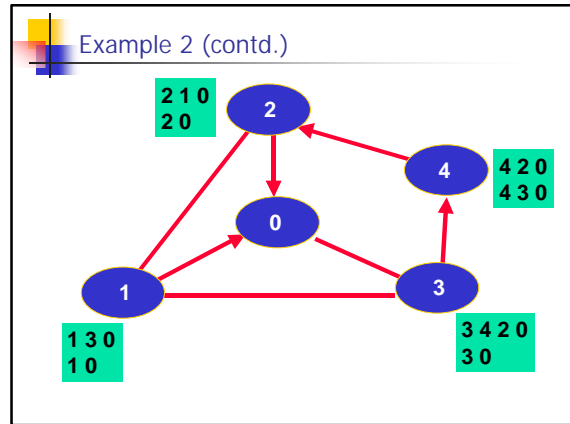
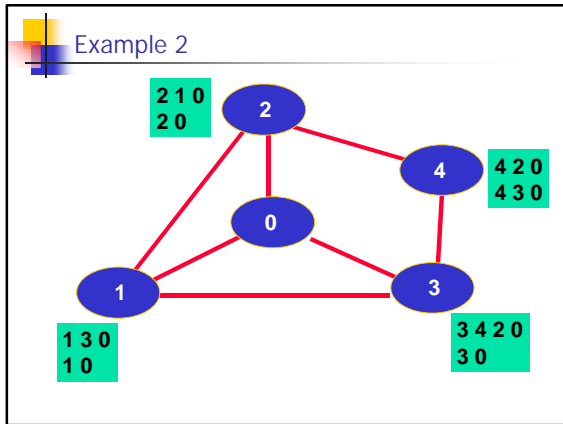
### Simple Path Vector Protocol (SPVP)

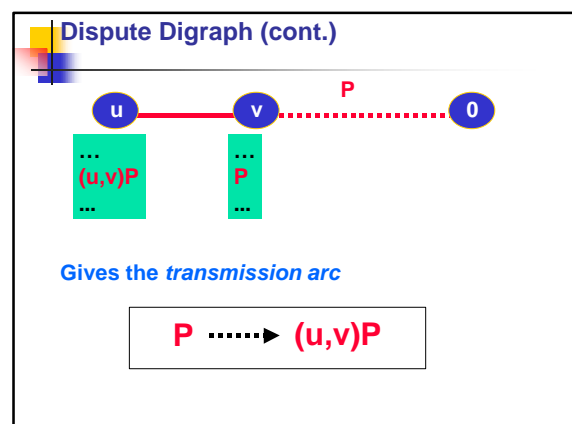
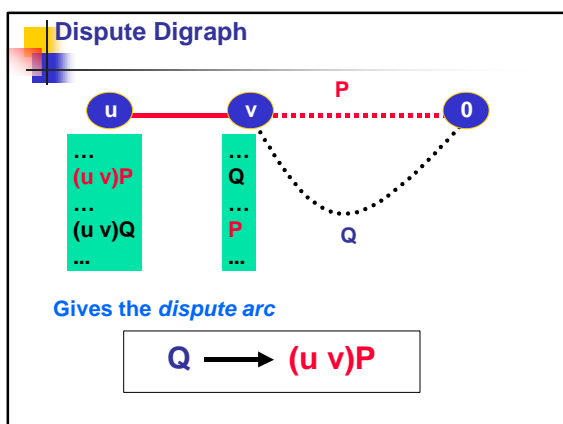

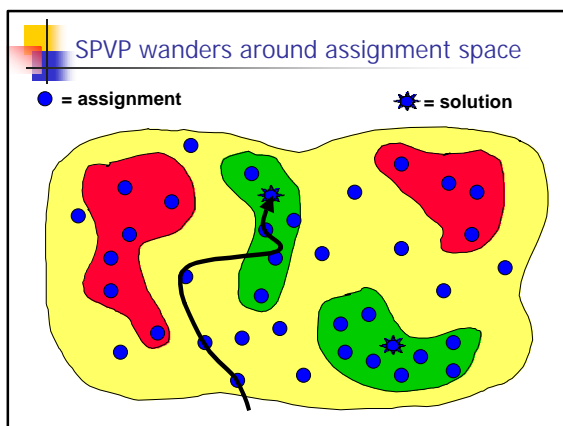
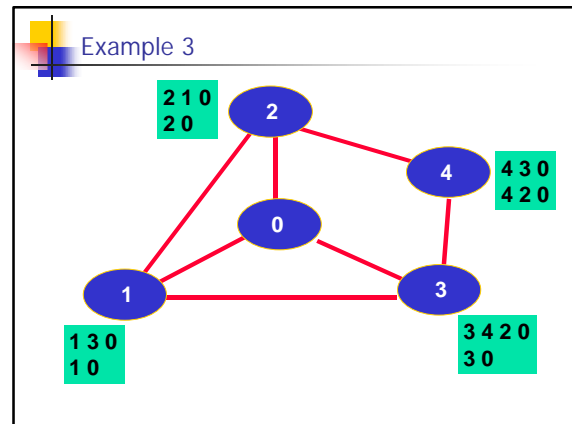
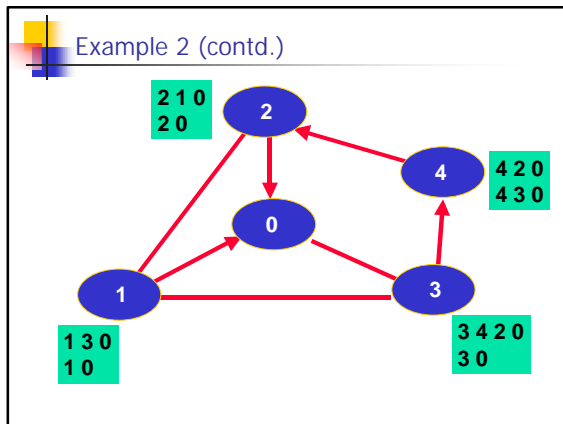
Pick the best path available at any given time...  
 $\text{rib}(u)$ : u's best path to origin  
 $\text{rib-in}(u \leftarrow w)$ : recent path sent from w


```

process spvp[u]
{
  receive P from w →
  {
    rib-in(u ← w) := u P
    if rib(u) != best(rib-in(u ← x)) {
      rib(u) := best(rib-in(u ← x))
      foreach v in peers(u) {
        send rib(u) to v
      }
    }
  }
}
  
```

### Example 1







Sufficient Condition for Sanity

If an instance of SPP has an acyclic dispute digraph, then

Static (SPP)	Dynamic (SPVP)
solvable	safe (can't diverge)
unique solution	predictable restoration
all sub-problems uniquely solvable	robust with respect to link/node failures