E6998-02: Internet Routing

Lecture 5
Fundamentals of Routing Protocols

John Ioannidis

AT&T Labs – Research

ji+ir@cs.columbia.edu

Announcements

Lectures 1-5 are available.

Guest Lecturer on 9/19: Noel Chiappa.

BE THERE!

In a different room (check the bboard right before class).

Lecture on 9/26 will be at 4:10pm in 1024.

W4180 on 9/26 will be at 2:40pm in 1127.

Homeworks submitted incorrectly. Don't do it again.

Homework 2 will be out later tonight.

I assume that only people who submitted homeworks are in the class.

Routing

- Internet is a collection of connected nodes.
  - End nodes (“end systems”): hosts.
  - Connecting nodes (“intermediate systems”): routers.
- Internet is organized as a collection of networks.
  - Customer networks.
  - ISPs.
    - Of varying sizes and purposes.
- Within a customer network or ISP:
  - Cohesive routing policy.
  - Performance and optimal internal routing is top priority.
- Between networks or ISPs:
  - Different routing policies.
  - Connectivity and abiding to policies is top priority.

Routing Protocols

- Manual configuration (static routes).
  - Too hard.
  - Does not scale.
  - Does not respond quickly to changing topologies.
- Source routing.
  - Source decides path of packets (LSRR/SSRR options).
  - Has been considered in the past (SDRP).
  - Still question about how routes are computed.
- Dynamic routing.
  - Distance-vector protocols.
  - Link-state protocols.
  - Other.

Distance-Vector

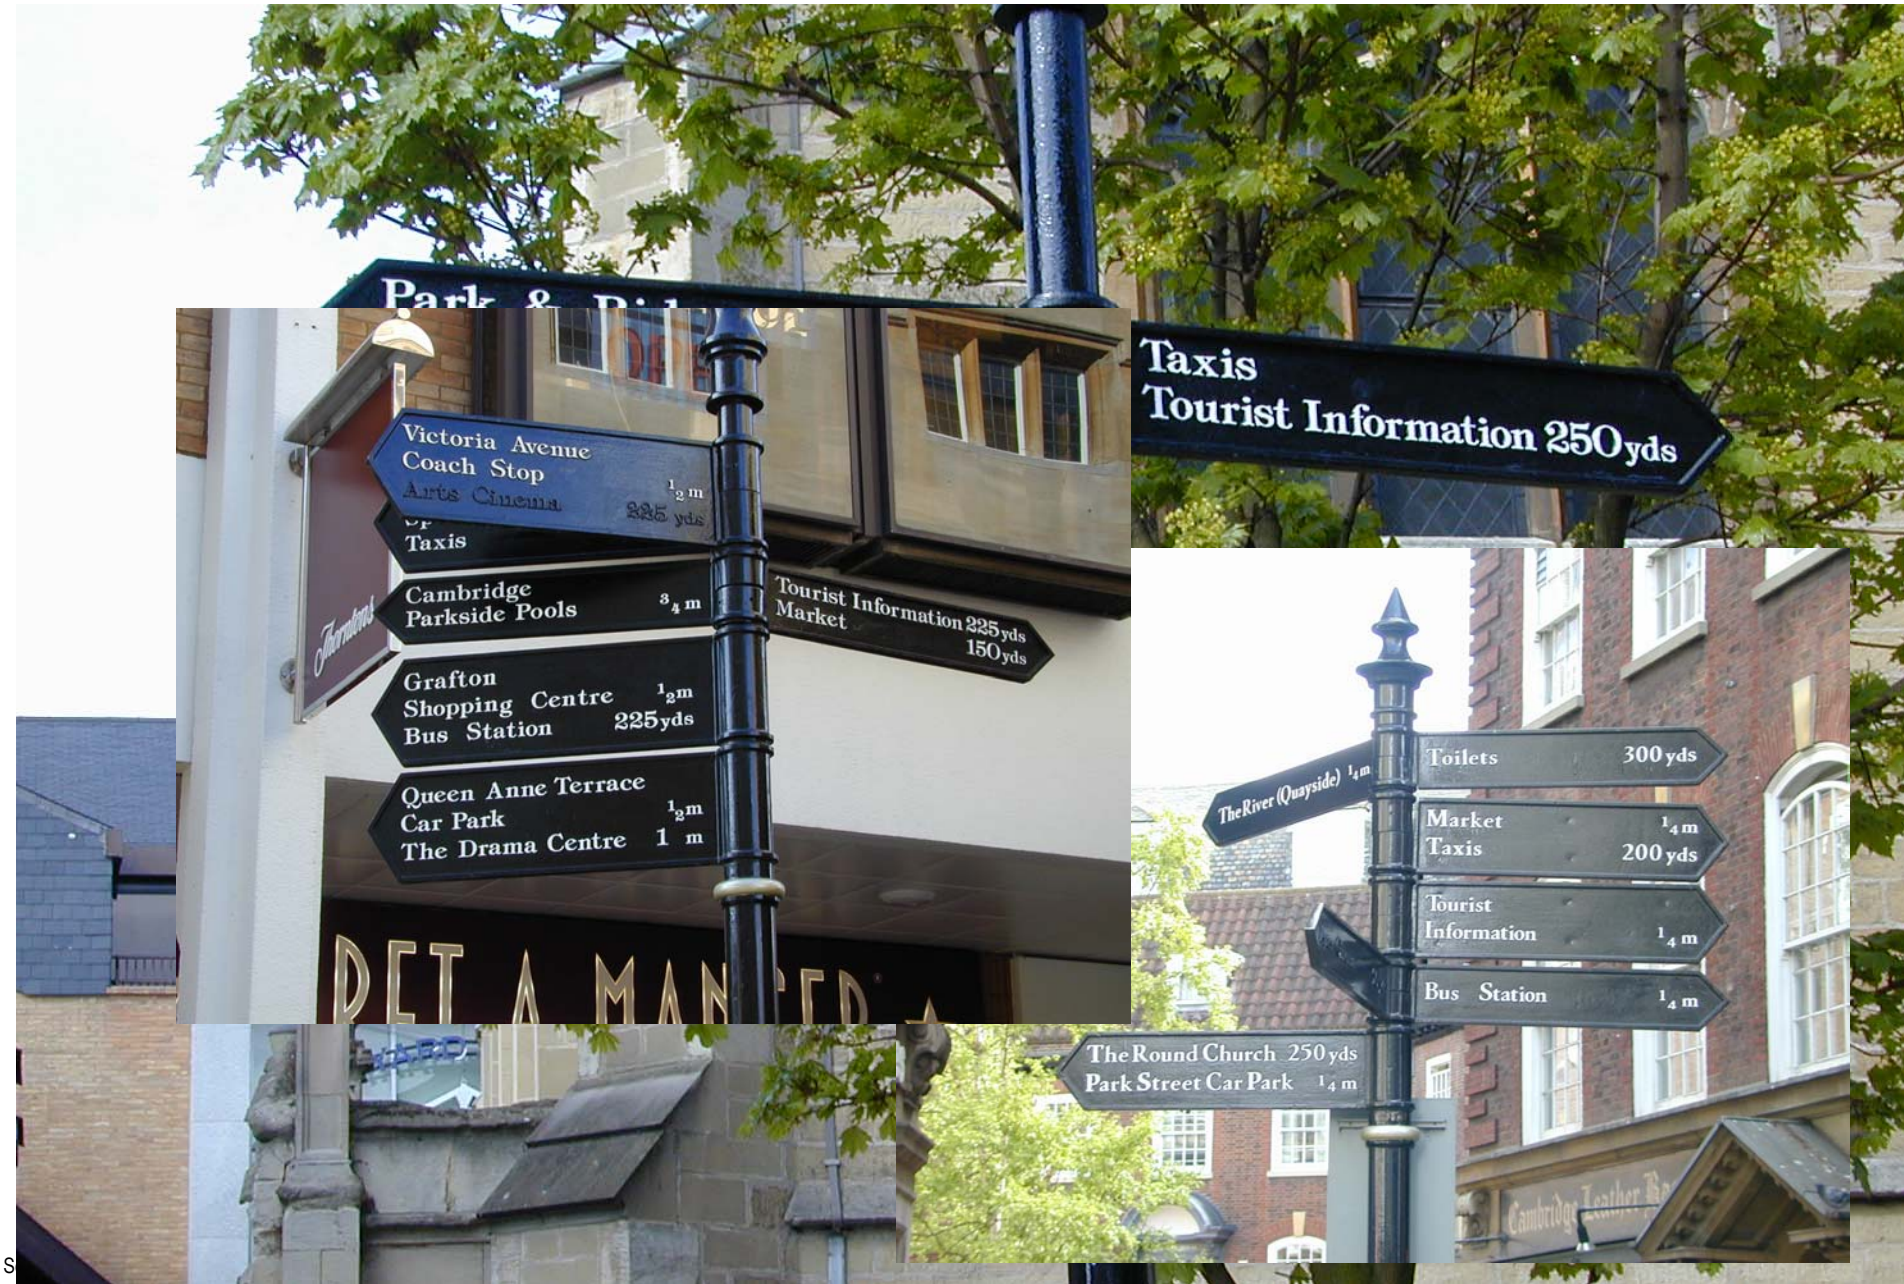

Distance-Vector

- Variations of Bellman-Ford algorithm.
- Each router starts by knowing:
  - Prefixes of its attached networks (“zero” distance).
  - Its next hop routers (how to find them?)
- Each router advertises only to its neighbors:
  - All prefixes it knows about.
  - Its distance from them.
- Each router learns:
  - All prefixes its neighbors know about.
  - Their distance from them.
- Each router figures out, for each destination prefix:
  - The “distance” (how far away it is).
  - The “vector” (the next hop router).

Link-State

- Based on Dijkstra's Shortest-Path-First algorithm.
- Each router starts by knowing:
  - Prefixes of its attached networks.
  - Links to its neighbors.
- Each router advertises to the entire network (flooding):
  - Prefixes of its directly connected networks.
  - Active links to its neighbors.
- Each router learns:
  - A complete topology of the network (routers, links).
- Each router computes shortest path to each destination.
- In a stable situation, all routers have the same graph, and compute the same paths.