LECTURE 15: DATACENTER NETWORK: TOPOLOGY AND ROUTING

Xiaowei Yang

OVERVIEW

- Portland: how to use the topology feature of the datacenter network to scale routing and forwarding

- ElasticTree: topology control to save energy
  - Briefly

BACKGROUND

- Link layer (layer 2) routing and forwarding
- Network layer (layer 3) routing and forwarding

- The FatTree topology

LINK LAYER ADDRESSING

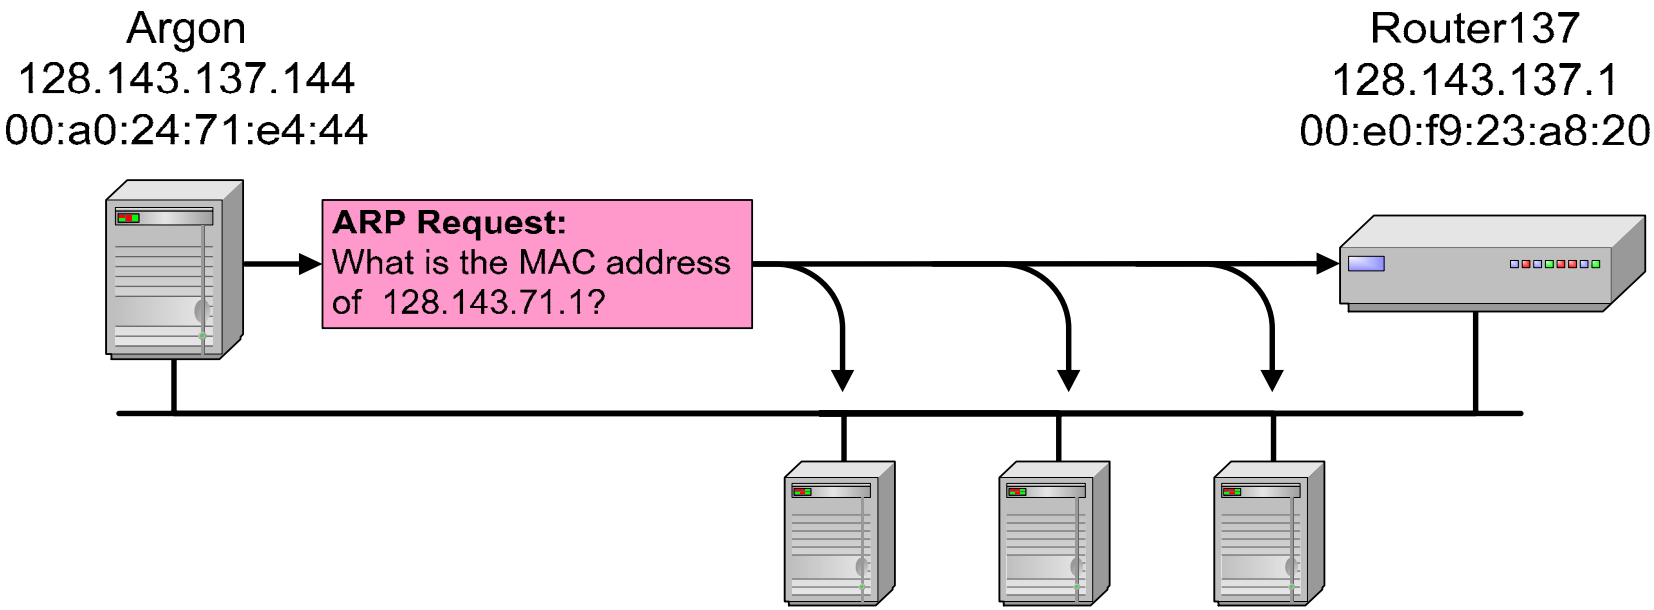

- To send to a host with an IP address p, a sender broadcasts an ARP request within its IP subnet
- The destination with the IP address p will reply
- The sender caches the result

LINK LAYER FORWARDING

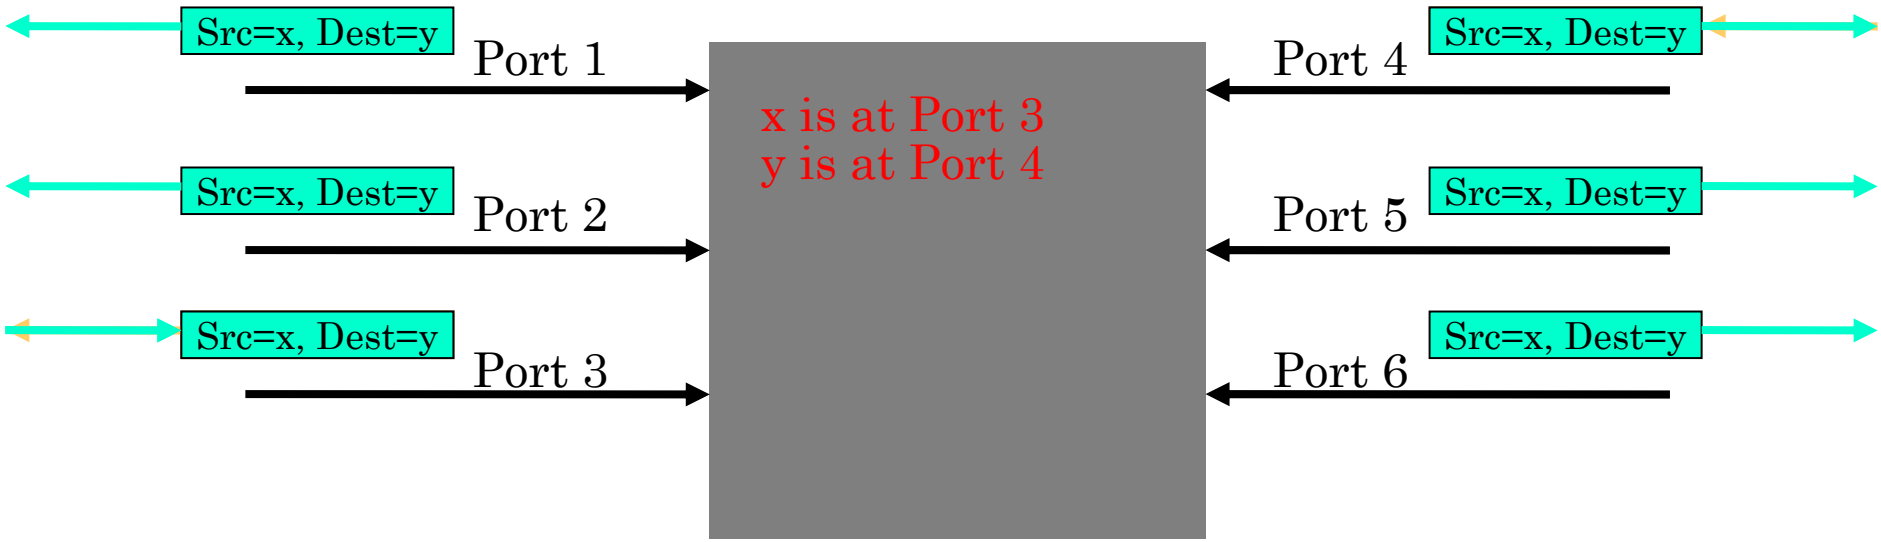

- Done via learning bridges
- Bridges run a spanning tree protocol to set up a tree topology
- First packet from a sender to a destination is broadcasted to all destinations in the IP subnet along the spanning tree
- Bridges on the path learn the sender's MAC address and incoming port
- Return packets from a destination to a sender are unicast along the learned path