Overlay Networks

Definition

- Network
  - defines addressing, routing, and service model for communication between hosts
- Overlay network
  - A network built on top of one or more existing networks
  - adds an additional layer of indirection/virtualization
  - changes properties in one or more areas of underlying network
- Alternative
  - change an existing network layer

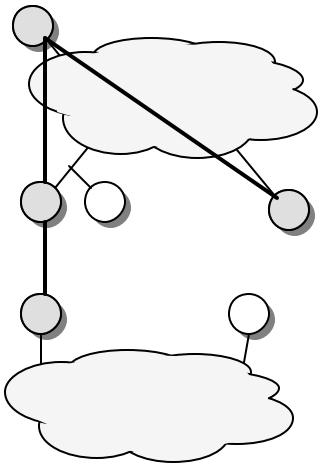

istoica@cs.berkeley.edu

Definition

- An overlay network is a virtual network of nodes and logical links that is built on top of an existing network with the purpose to implement a network service that is not available in the existing network.

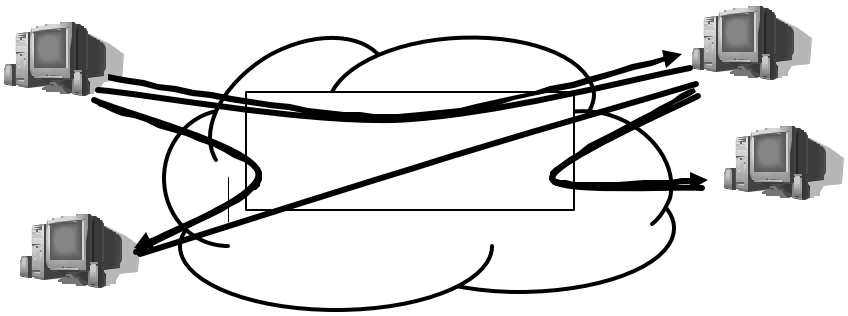

3

A Historical Example

- **The Internet is an overlay network**
  - goal: connect local area networks
  - built on local area networks (e.g., Ethernet), phone lines
  - add an Internet Protocol header to all packets

istoica@cs.berkeley.edu

4

Delivery of an IP datagram

- At data link layer :
  each data link network can transfer frames, but not delivery across multiple LANs

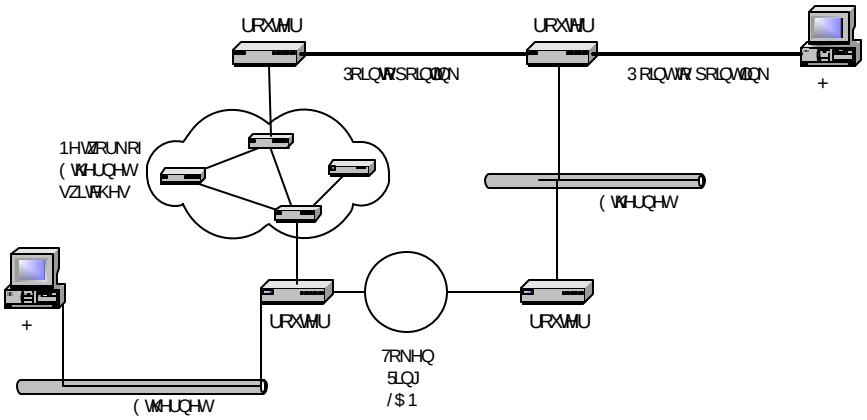

5

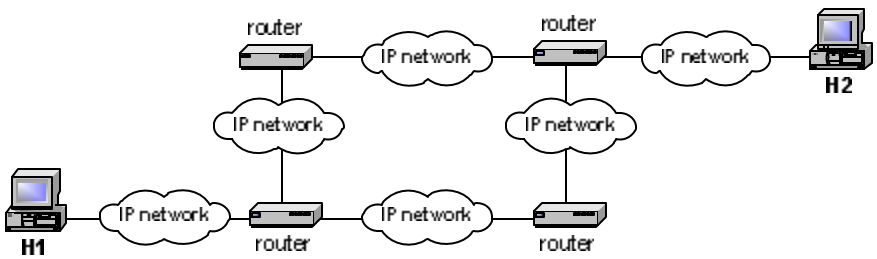

6

IP network as an overlay network

A set of routers that are all connected to a network of Layer 2 switches (ATM, Ethernet) ….

… can be configured to be a full mesh at the IP layer.

How?

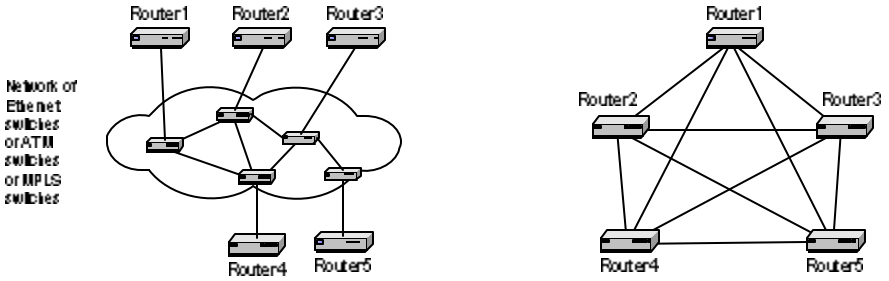

7

Applications of Overlays

- Routing
- Addressing
- Security
- Multicast
- Mobility

istoica@cs.berkeley.edu

8

Benefits

- Do not have to deploy new equipment, or modify existing software/protocols
  - probably have to deploy new software on top of existing software
  - e.g., adding IP on top of Ethernet does not require modifying Ethernet protocol or driver
  - allows bootstrapping
    - expensive to develop entirely new networking hardware/software
    - all networks after the telephone have begun as overlay networks

istoica@cs.berkeley.edu

Benefits

- Do not have to deploy at every node
  - Not every node needs/wants overlay network service all the time
    - e.g., QoS guarantees for best-effort traffic
  - Overlay network may be too heavyweight for some nodes
    - e.g., consumes too much memory, cycles, or bandwidth
  - Overlay network may have unclear security properties
    - e.g., may be used for service denial attack
  - Overlay network may not scale (not exactly a benefit)
    - e.g. may require n^2 state or communication

istoica@cs.berkeley.edu

Costs

- Adds overhead
  - Adds a layer in networking stack
    - Additional packet headers, processing
  - Sometimes, additional work is redundant
    - E.g., an IP packet contains both Ethernet (48 + 48 bits) and IP addresses (32 + 32 bits)
    - Eliminate Ethernet addresses from Ethernet header and assume IP header(?)
- Adds complexity
  - Layering does not eliminate complexity, it only manages it
  - More layers of functionality → more possible unintended interaction between layers
  - E.g., corruption drops on wireless interpreted as congestion drops by TCP

istoica@cs.berkeley.edu

Applications: Increasing Routing Robustness

- **Resilient Overlay Networks (RONs)**
  [Anderson et al 2001]
  - Overlay nodes form a complete graph
  - Nodes probe other nodes for lowest latency
  - Knowledge of complete graph → lower latency routing than IP, faster recovery from faults

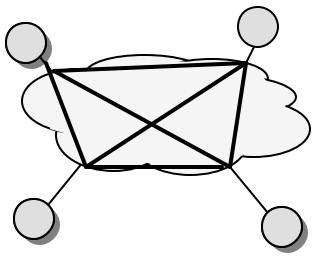

istoica@cs.berkeley.edu

Case Study: Multicasting

- Multicast communications refers to one-to-many or many-to-many communications.
- Multicast implements a one-to-many send operation:

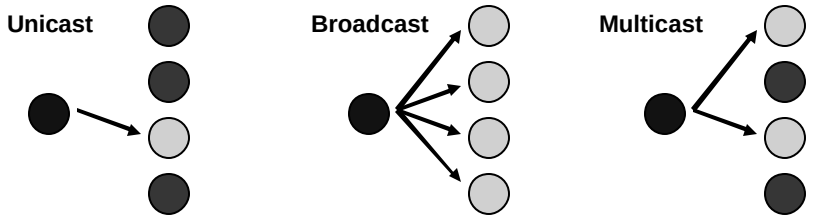

13

Case Study: Need for multicast

- Without support for multicast at the network layer:

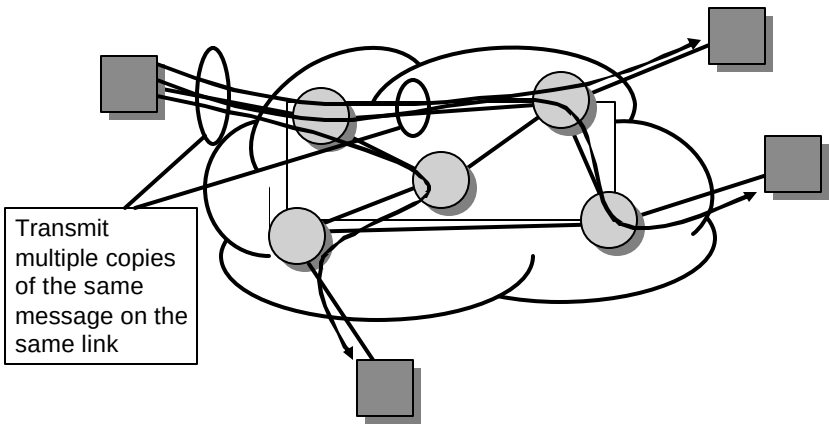

14

Multicasting at the network layer

- With support for multicast at the network layer:

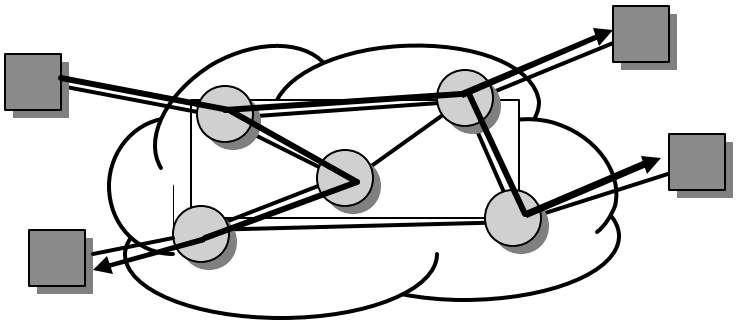

Requires a set of mechanisms at the network layer:

(1) Routers must be able to send multiple copies of same packet
(2) Multicast routing algorithms needed to build up a dissemination tree

15

How does multicast work with IP?

- IP multicast addresses are allocated a certain range:

Class D

Class	From	To
D	224.0.0.0	239.255.255.255

- Each multicast group designates a “multicast group”.
- Hosts can “join” a multicast group.
- An IP datagram sent to a multicast address is forwarded to everyone who has joined the multicast group.

16

MBONE: An overlay network for incremental IP multicast deployment

- IP multicast deployment in the Internet began in early 1990s with the creation of the Multicast Backbone (MBONE)
- MBONE solved the problem of wide-area IP multicast routing on the Internet where only few routers were capable of IP multicast routing, by setting up a virtual network of multicast routers that are connected by unicast path.
- MBONE uses the concept of IP tunneling (IP-in-IP encapsulation)

17

MBONE: Overlay Network of multicast routers

- MBONE is the set of all of multicast routers
- MBONE is an overlay network at the IP layer that provides multicast service

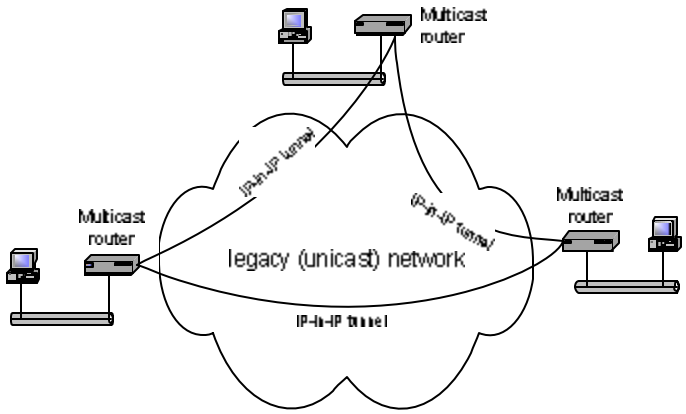

18

IP Tunneling

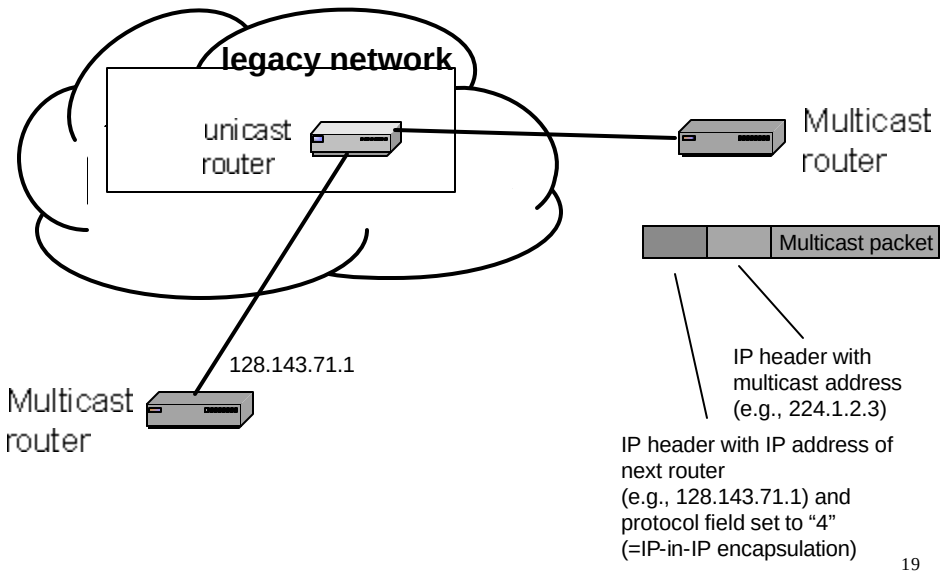

19

Incremental deployment with IP tunneling

- IP tunneling in MBONE permits gradual deployment of a service (here: multicast service)
- **Phase 1: multicast enabled routers only at edges**

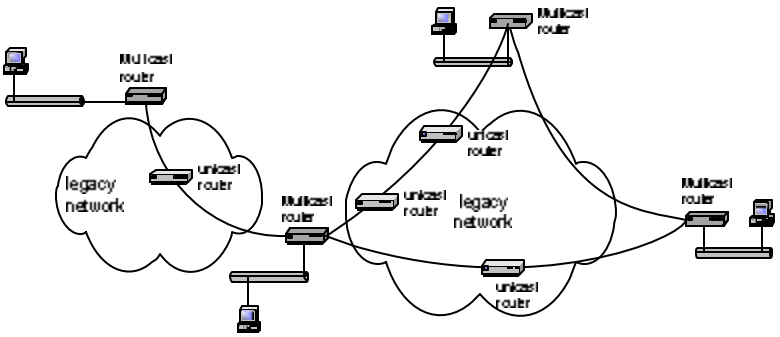

20

Incremental deployment with IP tunneling

- **Phase 2: some routers in backbone are multicast capable**

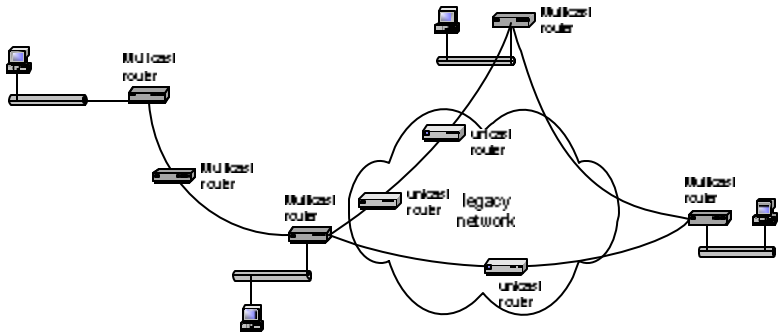

21

Incremental deployment with IP tunneling

- **Phase 3: more routers in backbone are multicast capable**

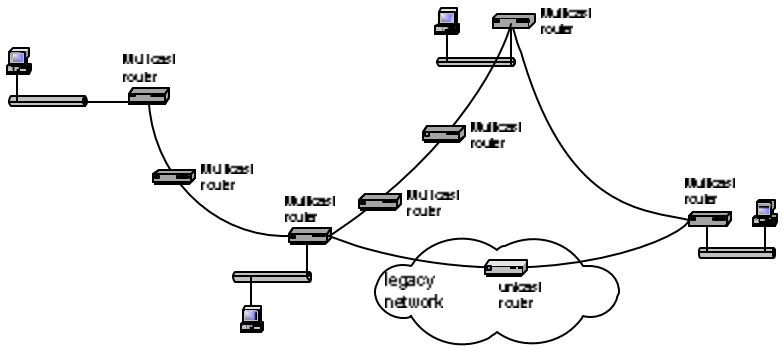

22

Incremental deployment with IP tunneling

- Finally, when all routers are multicast enabled, the MBONE overlay has been replaced by an infrastructure
- **Phase 4: entire network is multicast enabled**

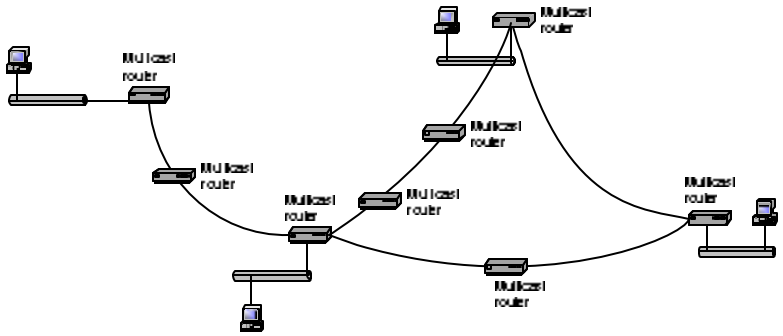

23

Other networks with tunneling: 6Bone, VPN

- 6Bone is an IPv6 overlay that can be used to transmit IPv6 packets over an IPv4 network

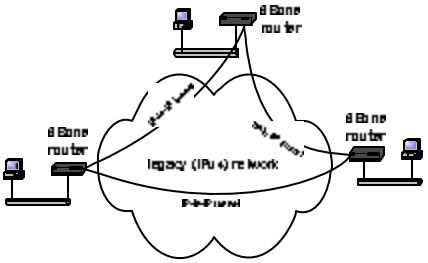

- Virtual Private Networks (VPNs) encrypt and encapsulate IPv4 traffic.
- Encryption and encapsulation is done by VPN routers

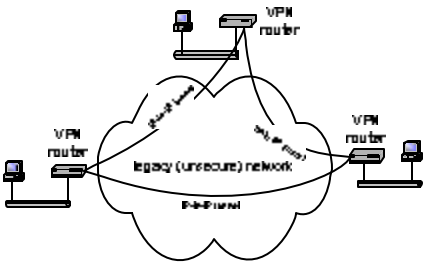

24

The evolution of IP multicast

- **Early 1990s:** MBONE routers were developed and deployed
- **Mid-1990s:**
  - MBONE has thousands of MBONE routers
  - IP routers increasingly support multicast routing
- **End-1990s:**
  - Most routers can support multicasting
  - **But IP multicast is not turned on most routers because of concerns pertaining to scalability, network management, deployment and support for error, flow and congestion control**
- **Since early 2000s:**
  - Multicast through application layer overlays

25

Application-layer multicasting

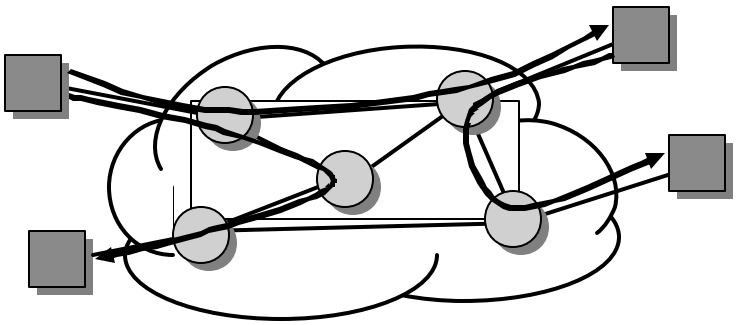

- Provide multicast functionality above IP layer
- Data is transmitted between neighbors in the overlay
- No multicast needed in overlay network

26

Potential Benefits

- Scalability
  - Routers do not maintain per-group state
  - End systems do, but they participate in very few groups
- Easier to deploy
  - Only requires adding software to end hosts
- Potentially simplifies support for higher level functionality
  - Use hop-by-hop approach, but end hosts are routers
  - Leverage computation and storage of end systems
  - E.g., packet buffering, transcoding of media streams, ACK aggregation
  - Leverage solutions for unicast congestion control and reliability

Overlay Construction Problems

- Dynamic changes in group membership
  - Members may join and leave dynamically
  - Members may die
- Dynamic changes in network conditions and topology
  - Delay between members may vary over time due to congestion, routing changes
- Knowledge of network conditions is member specific
  - Each member must determine network conditions for itself

What is the best overlay?

Evaluation criteria:

1. **Properties of the overlay graph**
2. **Mapping of the overlay to the layer-3 network**
3. **Properties of protocol that maintains the overlay topology**

29

Overlay-based approaches for multicasting

- Build an overlay mesh network and embed trees into the mesh:
- Build a shared tree:
- Build a graph with well-known properties
  - N-dimensional torus: CAN (see reading list)
  - Hypercube inspired: Chord (see reading list)
  - Triangulation: Delaunay Triangulation (see reading list)

30

1. Properties of the overlay graph

- **Number of neighbors (routing table size)**
  - **Many DHTs, hypercubes: O(log N)** *(max.)*
  - **Triangulation graphs: O(N)** *(max.),* **6** *(avg.)*
  - **Meshes, trees:** no *a priori* bound, but bounds can be enforced
- **Path lengths in the overlay**
  - **Many DHTs, hypercubes: O(log N)** *(max.)*
  - **Triangulation graphs: O(N)** *(max.),* **O(√N)** *(best case avg.)*
  - **Meshes, trees:** no *a priori* bound

31

2. Mapping of the overlay to the layer-3 network

- **Compare overlay multicast to network-layer multicast:**
  "Stretch": Ratio of delay to shortest path delay
  "Stress": Number of duplicate transmissions over a physical link
- Overlays that provide a good mapping need to be aware of the underlying layer-3 network

32

Illustration of “Stress” and “Stretch”

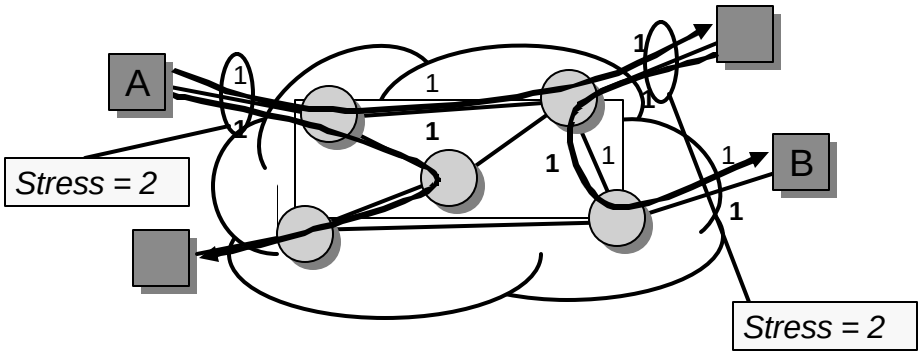

Unicast delay A→B : 4

Delay A→B in overlay: 6

Stretch for A→B: 1.5

33

Case Study: Mobile IP

- **Goal:** Provide access to the Internet to mobile IP hosts
- Reconnection occurs automatically
- Mobility transparent to applications and higher level protocols such as TCP

34

Mobile Networking problem

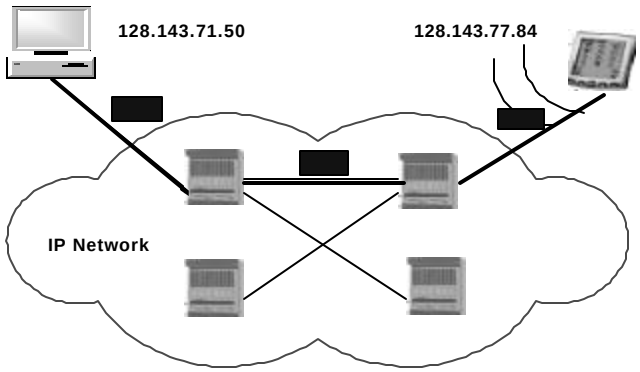

35

Mobile Networking problem

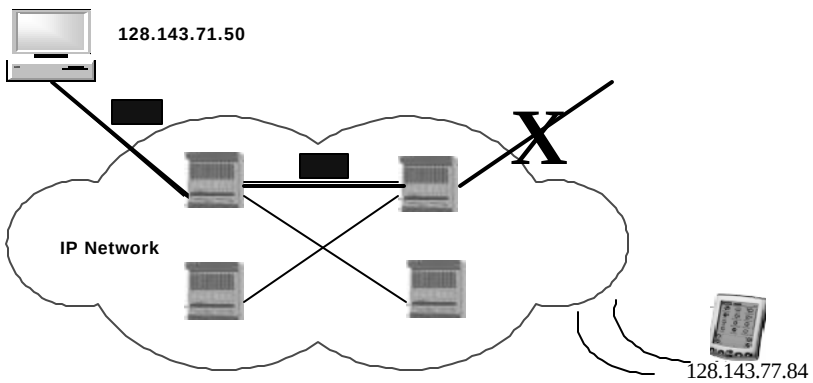

36

Mobile IP Approach

- Mobile IP uses two IP addresses :
  - **Home address:** The IP address assigned to the *mobile node*, making it logically appear attached to its *home network*.
  - **Care-of address:** An IP address at the *mobile node's* current point of attachment to the Internet, when the mobile node is not attached to the *home network*.

- **Home network:** The network at which the *mobile node* seems reachable, to the rest of the Internet, by virtue of its assigned IP address.
- **Foreign network:** The network to which the *mobile node* is attached when it is not attached to its *home network*, and on which the *care-of address* is reachable from the rest of the Internet.
- **Home agent:** A router on the *home network* that effectively causes the mobile node to be reachable at its home address even when the mobile node is not attached to its home network.
- **Foreign agent:** A router on the foreign network that can assist the mobile node in receiving datagrams delivered to the care-of address.

37

How Mobile IP Works

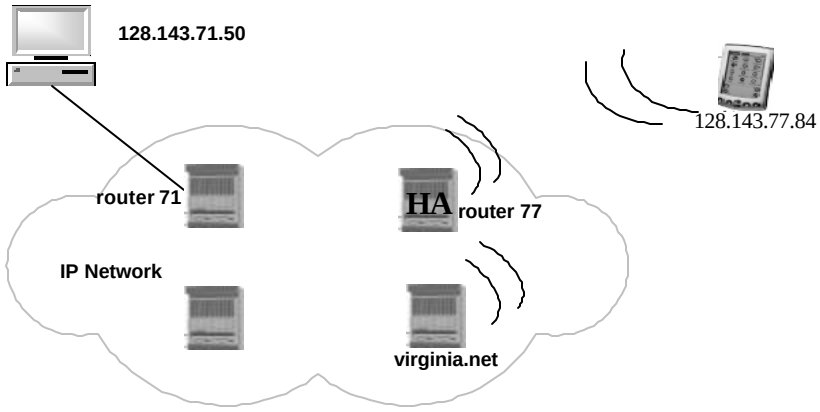

38

How Mobile IP Works

- Discovering the care-of address
- Registering the care-of address
- Tunneling to the care-of address

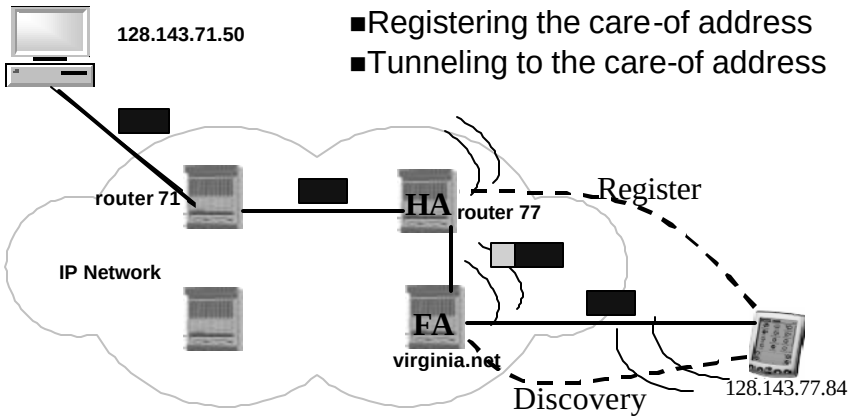

39

Tunneling to the Care of Address

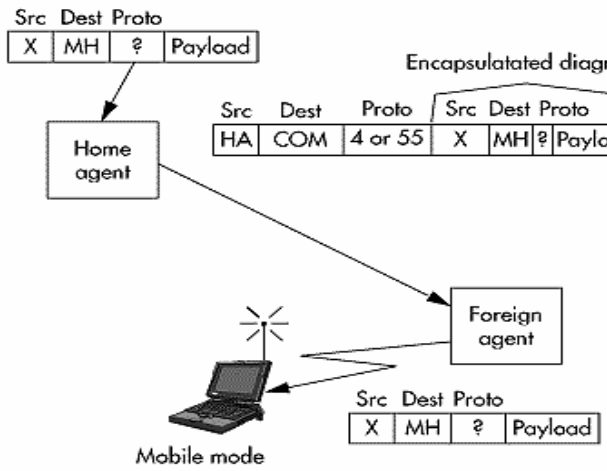

40