

RSC


Part II: Network Layer

3. IP addressing (2nd part)

Redes y Servicios de Comunicaciones Universidad Carlos III de Madrid

These slides are, mainly, part of the companion slides to the book “Computer Networking: A Top Down Approach” generously made available by their authors (see copyright below). The slides have been adapted, where required, to the teaching needs of the subject above.

All material copyright 1996-2009
J.F Kurose and K.W. Ross, All Rights Reserved



*Computer Networking:
A Top Down Approach
5th edition.*

*Jim Kurose, Keith Ross
Addison-Wesley, April
2009.*

RSC Part II: Network Layer

- ❑ II. 1 Basic Network layer concepts
- ❑ II.2 Introduction to IP
 - Datagram format
 - ICMP
- ❑ II.3 IP addressing
 - Obtaining addresses, DHCP, NAT
- ❑ II.4 IP in operation
 - ARP
- ❑ II.5 Network routing
 - Link state
 - Distance Vector
 - Hierarchical routing
- ❑ II.6 Routing in the Internet
 - RIP
 - OSPF
 - BGP
 - Broadcast and multicast

IP addresses: how to get one?

Q: How does a *host* get IP address?

- ❑ hard-coded by system admin in a file
 - Windows: control-panel->network->configuration->tcp/ip->properties
 - UNIX: /etc/rc.config
- ❑ **DHCP: Dynamic Host Configuration Protocol:**
dynamically get address from as server
 - "plug-and-play"

IP addresses: how to get one?

Q: How does *network* get subnet part of IP addr?

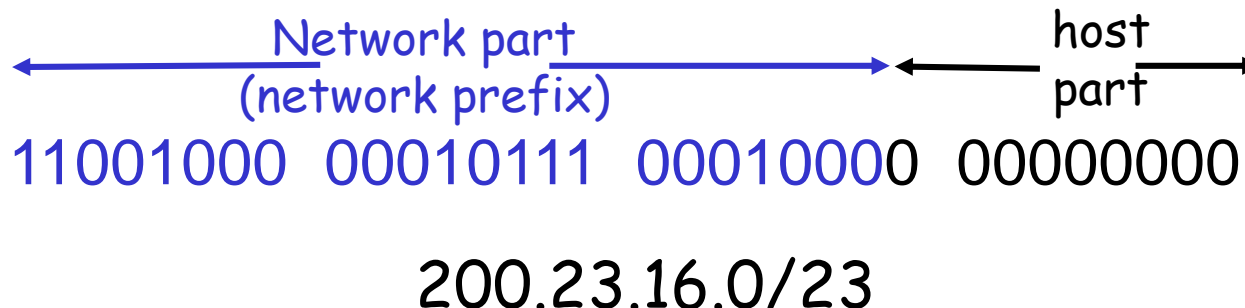
A: gets allocated portion of its provider ISP's address space

ISP's block	<u>11001000 00010111 00010000</u> 00000000	200.23.16.0/20
Organization 0	<u>11001000 00010111 00010000</u> 00000000	200.23.16.0/23
Organization 1	<u>11001000 00010111 00010010</u> 00000000	200.23.18.0/23
Organization 2	<u>11001000 00010111 00010100</u> 00000000	200.23.20.0/23
...
Organization 7	<u>11001000 00010111 00011110</u> 00000000	200.23.30.0/23

IP addressing: CIDR

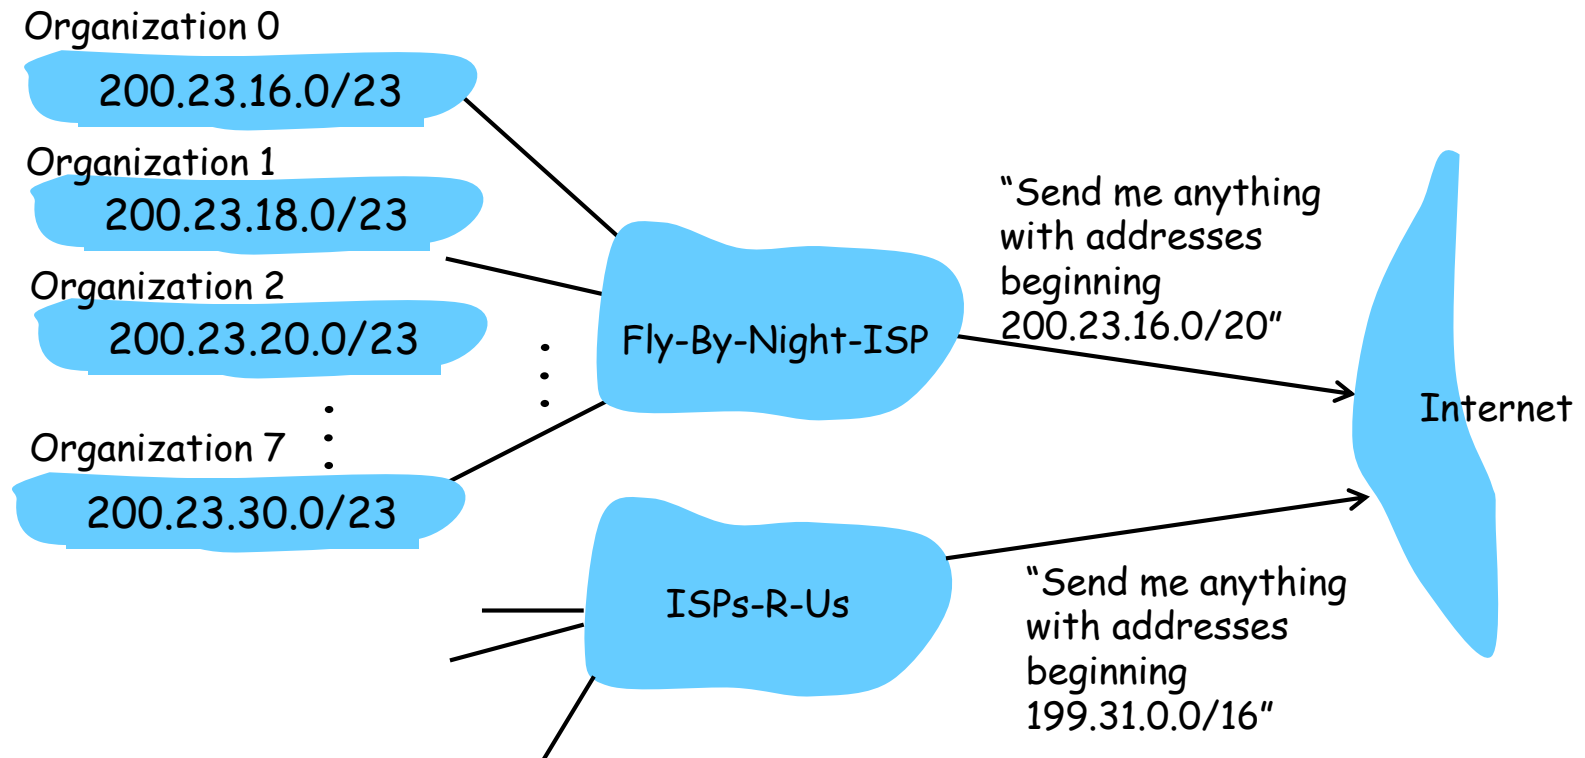
CIDR: Classless InterDomain Routing

- network portion of address of arbitrary length
- address format: $a.b.c.d/x$, where x is number of bits in the network portion of the address



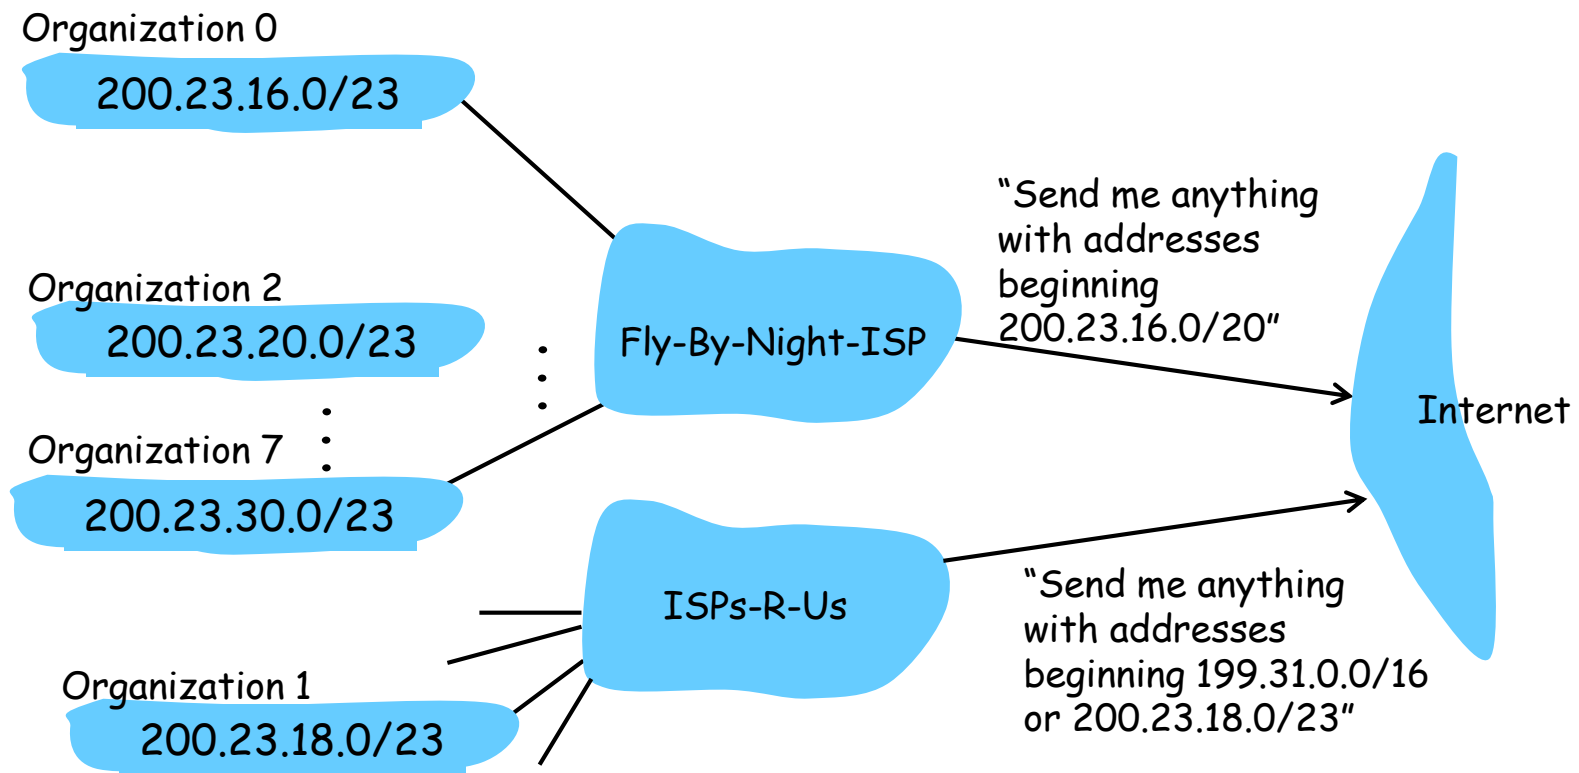
Hierarchical addressing: route aggregation

Hierarchical addressing allows efficient advertisement of routing information:



Hierarchical addressing: more specific routes

ISPs-R-Us has a more specific route to Organization 1



IP addressing: the last word...

Q: How does an ISP get block of addresses?

A: **ICANN**: Internet Corporation for Assigned Names and Numbers

- Allocates and registers addresses
- manages DNS
- assigns domain names, resolves disputes
- The Internet Assigned Numbers Authority (**IANA**), operated by the ICANN (Internet Corporation for Assigned Names and Numbers)
 - Manages IP address allocation (and parameters of Internet protocols)
 - Delegates address assignment to RIRs (Regional Internet Registry):
 - American Registry for Internet Numbers (ARIN) for North America and parts of the Caribbean
 - **RIPE Network Coordination Centre (RIPE NCC) for Europe, the Middle East and Central Asia**
 - Asia-Pacific Network Information Centre (APNIC) for Asia and the Pacific region
 - Latin American and Caribbean Internet Addresses Registry (LACNIC) for Latin America and parts of the Caribbean region
 - African Network Information Centre (AfriNIC) for Africa
 - ISPs get IP address blocks from RIRs or Local Internet Registry (LIR)

DHCP: Dynamic Host Configuration Protocol

Goal: allow host to *dynamically* obtain its IP address from network server when it joins network

Can renew its lease on address in use


Allows reuse of addresses (only hold address while connected an "on")

Support for mobile users who want to join network (more shortly)

DHCP overview:

- host broadcasts "DHCP discover" msg
- DHCP server responds with "DHCP offer" msg
- host requests IP address: "DHCP request" msg
- DHCP server sends/commits address: "DHCP ack" msg

DHCP client-server scenario



DHCP client-server scenario

DHCP server: 223.1.2.5

DHCP discover

src : 0.0.0.0, 68
dest.: 255.255.255.255,67
yiaddr: 0.0.0.0
transaction ID: 654

arriving
client



DHCP offer

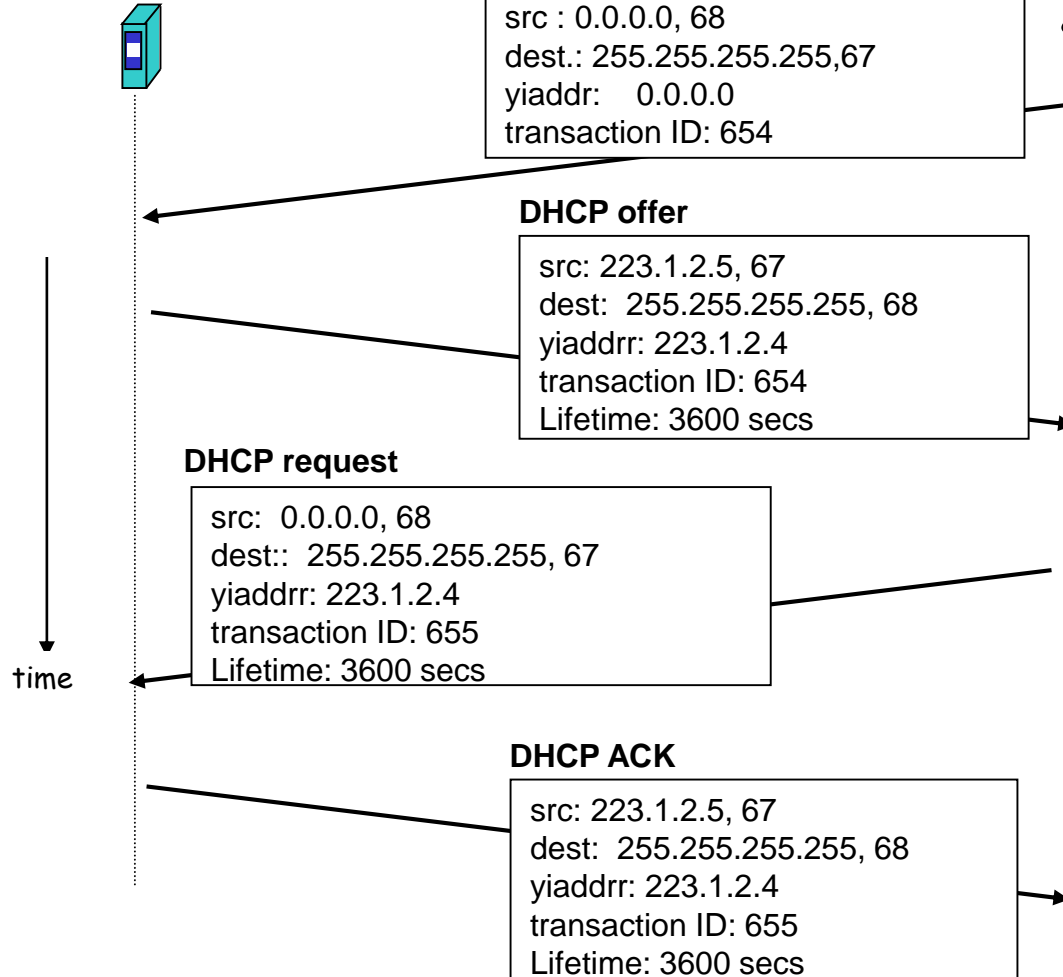
src: 223.1.2.5, 67
dest: 255.255.255.255, 68
yiaddr: 223.1.2.4
transaction ID: 654
Lifetime: 3600 secs

DHCP request

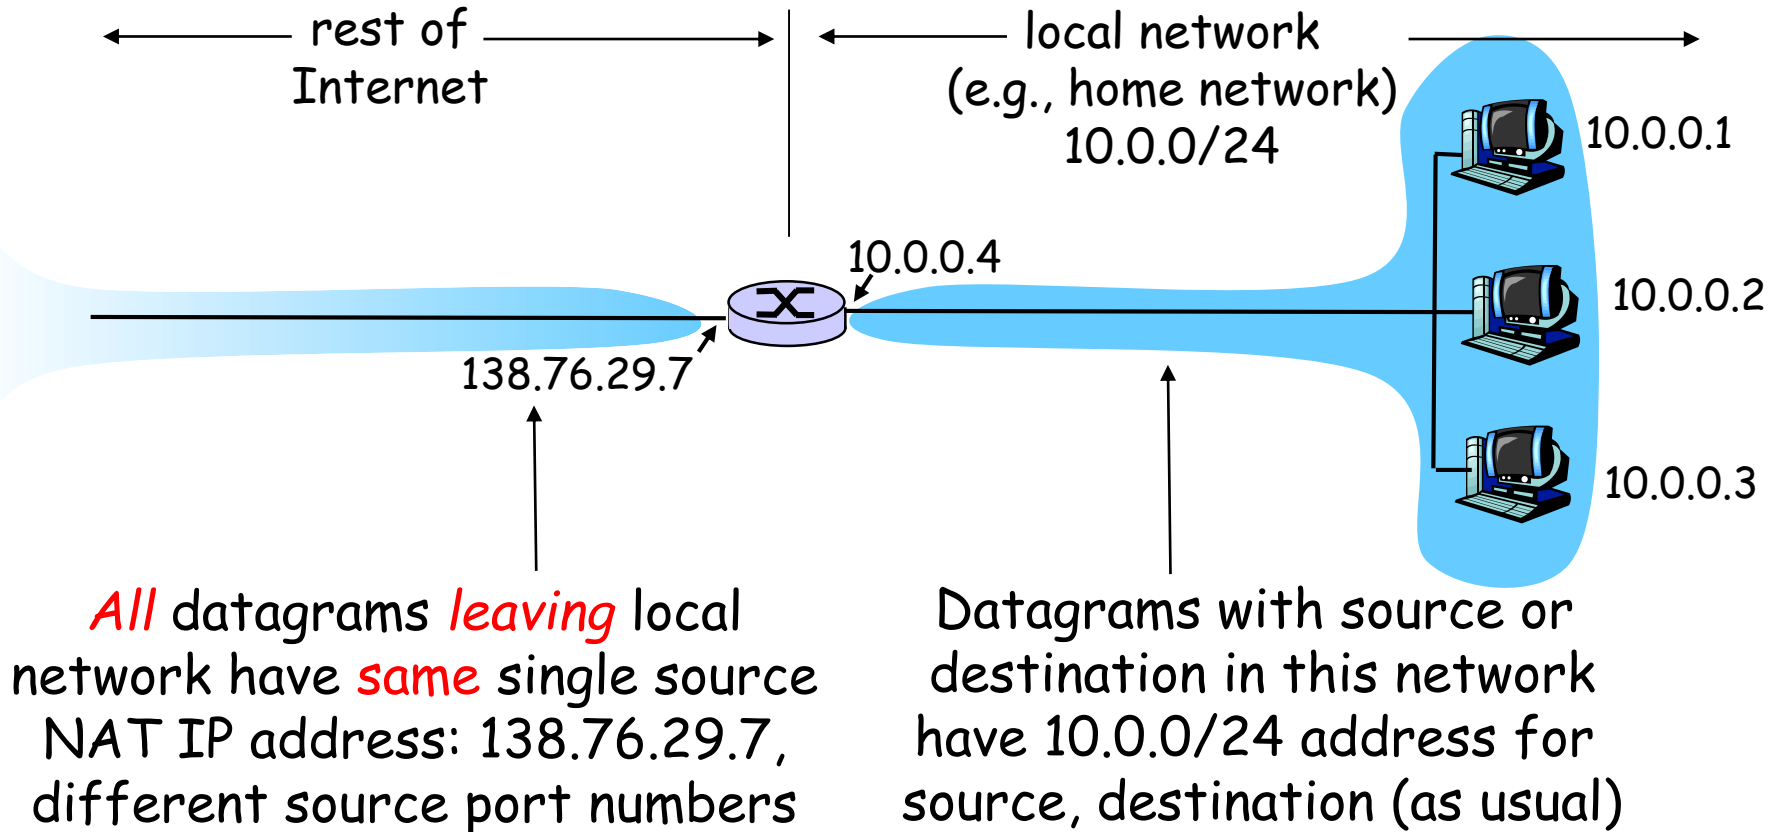
src: 0.0.0.0, 68
dest.: 255.255.255.255, 67
yiaddr: 223.1.2.4
transaction ID: 655
Lifetime: 3600 secs

DHCP ACK

src: 223.1.2.5, 67
dest: 255.255.255.255, 68
yiaddr: 223.1.2.4
transaction ID: 655
Lifetime: 3600 secs



NAT: Network Address Translation



NAT: Network Address Translation

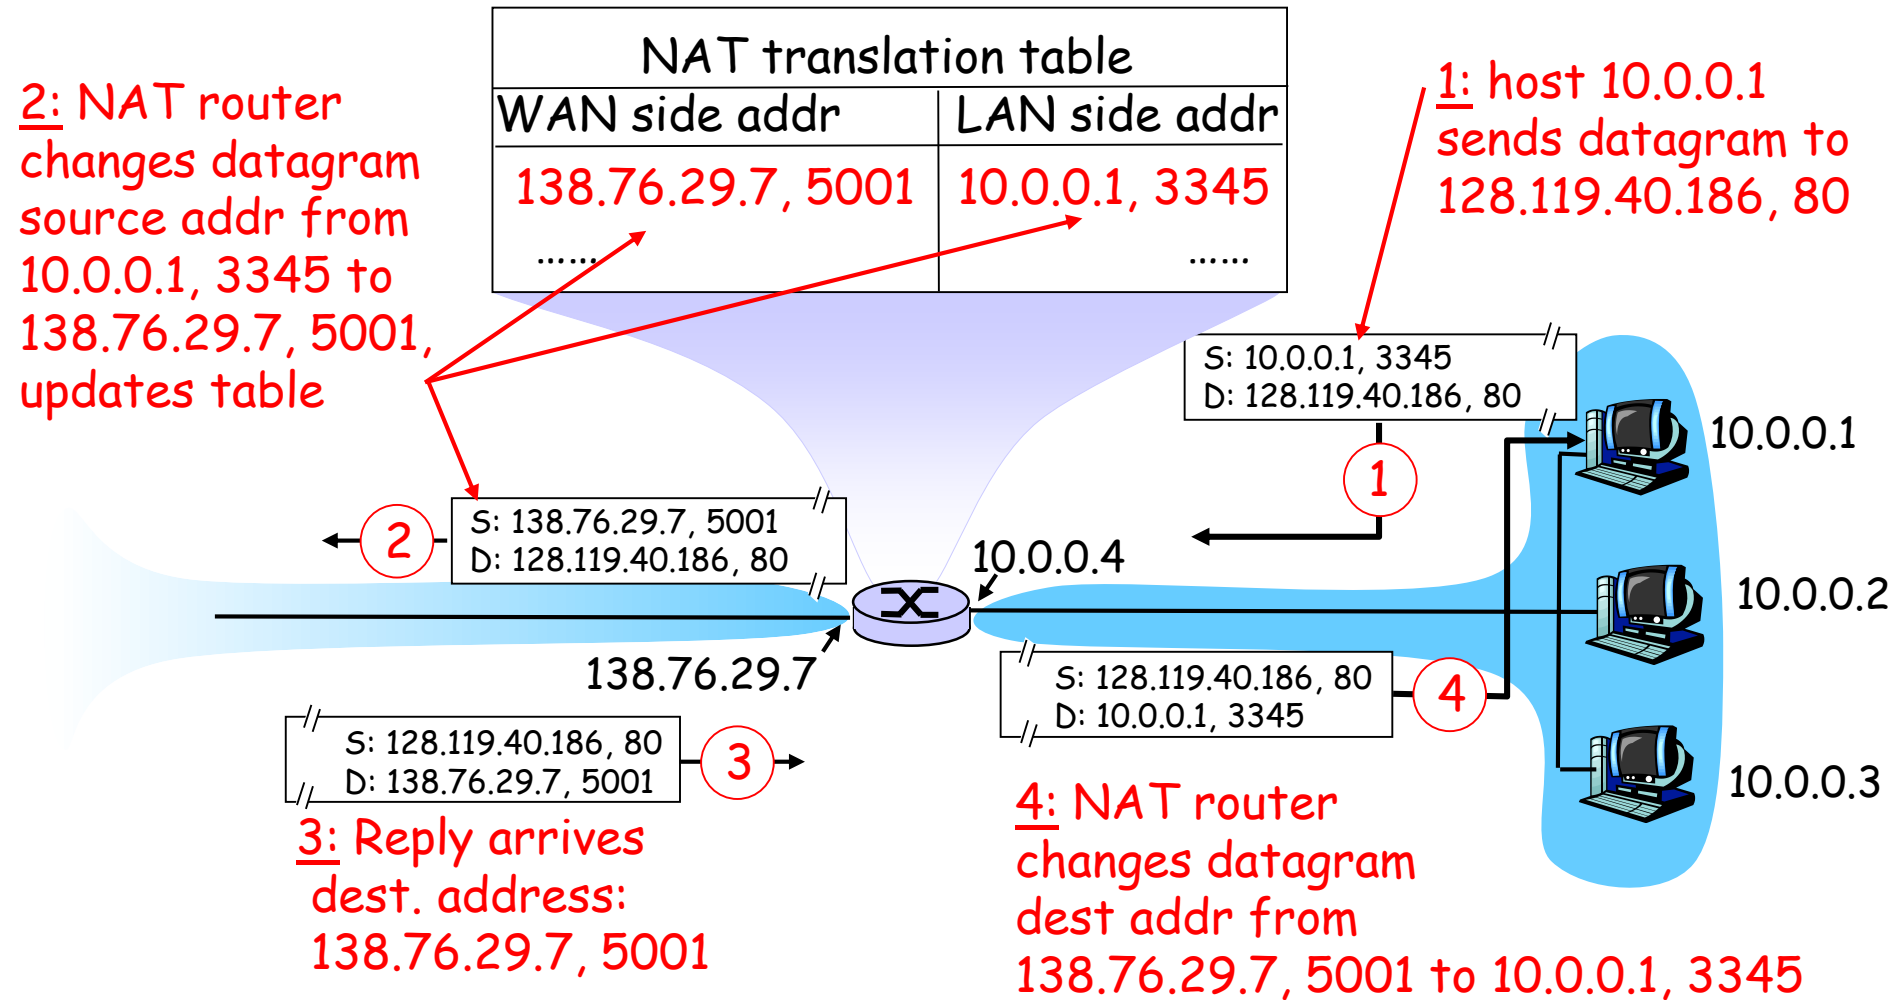
- ❑ **Motivation:** local network uses just one IP address as far as outside world is concerned:
 - range of addresses not needed from ISP: just one IP address for all devices
 - can change addresses of devices in local network without notifying outside world
 - can change ISP without changing addresses of devices in local network
 - devices inside local net not explicitly addressable, visible by outside world (a security plus).
- ❑ Private address space:
 - 10.0.0.0 - 10.255.255.255 (10/8 prefix)
 - 172.16.0.0 - 172.31.255.255 (172.16/12 prefix)
 - 192.168.0.0 - 192.168.255.255 (192.168/16 prefix)

NAT: Network Address Translation

Implementation: NAT router must:

- *outgoing datagrams: replace* (source IP address, port #) of every outgoing datagram to (NAT IP address, new port #)
... remote clients/servers will respond using (NAT IP address, new port #) as destination addr.
- *remember (in NAT translation table)* every (source IP address, port #) to (NAT IP address, new port #) translation pair
- *incoming datagrams: replace* (NAT IP address, new port #) in dest fields of every incoming datagram with corresponding (source IP address, port #) stored in NAT table

NAT: Network Address Translation




NAT: Network Address Translation

- ❑ 16-bit port-number field:
 - 60,000 simultaneous connections with a single LAN-side address!
- ❑ NAT is controversial:
 - routers should only process up to layer 3
 - violates end-to-end argument
 - NAT possibility must be taken into account by app designers, eg, P2P applications
 - address shortage should instead be solved by IPv6

NAT traversal problem


- ❑ client wants to connect to server with address 10.0.0.1
 - server address 10.0.0.1 local to LAN (client can't use it as destination addr)
 - only one externally visible NATted address: 138.76.29.7
- ❑ solution 1: statically configure NAT to forward incoming connection requests at given port to server
 - e.g., (123.76.29.7, port 2500) always forwarded to 10.0.0.1 port 25000



NAT traversal problem

- solution 2: Universal Plug and Play (UPnP) Internet Gateway Device (IGD) Protocol. Allows NATted host to:
 - ❖ learn public IP address (138.76.29.7)
 - ❖ add/remove port mappings (with lease times)

i.e., automate static NAT port map configuration



NAT traversal problem

- solution 3: relaying (used in Skype)
 - NATed client establishes connection to relay
 - External client connects to relay
 - relay bridges packets between to connections

