

CS 655 – System and Network Architectures and Implementation


Module 4 – Naming, Mobility, Messaging

Martin Karsten

mkarsten@uwaterloo.ca

Notice

Some slides and elements of slides are taken from third-party slide sets. In this module, parts are taken from the Kurose/Ross slide set. See detailed statement on next slide...



A note on the use of these ppt slides:

We're making these slides freely available to all (faculty, students, readers). They're in PowerPoint form so you can add, modify, and delete slides (including this one) and slide content to suit your needs. They obviously represent a *lot* of work on our part. In return for use, we only ask the following:

- If you use these slides (e.g., in a class) in substantially unaltered form, that you mention their source (after all, we'd like people to use our book!)
- If you post any slides in substantially unaltered form on a www site, that you note that they are adapted from (or perhaps identical to) our slides, and note our copyright of this material.

Thanks and enjoy! JFK/KWR

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.

Overview

- naming
 - principles
 - examples: DNS, DHT
- mobility
 - principles
 - example: Mobile IP
- messaging
 - principles
 - example: Email

Transport – Review

- end-to-end session characteristics
 - reliability
 - performance – sender rate control
- assume both end points can rendezvous
 - can find each other
 - both are online

Naming, Mobility, Messaging

- manage and find entities and services
 - mobile systems
 - content-based naming
- maintain connectivity
 - mobile systems
- communicate at different times
 - asynchronous messaging

Naming

“Of what one cannot speak,
one must pass over in silence.”

Ludwig Wittgenstein, *Tractatus*

Naming

- names in (distributed) computer systems
 - identification, (permanent) uniqueness – *scope*
 - facilitate communication / access – *resolution*
 - description of entity – *context*
- static vs. dynamic naming
 - offline vs. offline agreement
 - online: distributed vs. centralized
 - manual configuration

Name Resolution

- access named object / entity
- direct access
 - forward message to destination
- indirect access
 - map name to other name using database (*lookup*)
 - might be a distributed database
 - > forwarding within distributed database?

Definition Attempt

- a *name* is a handle/reference, valid and unique in a *scope*, that can be used to access a (group of) object(s) via a *resolution* mechanism
- try your own...

Everything is a Name

- memory address
- file system inode / name, socket handle
- MAC address, IP address, port, DNS name
- URL
- email address, Skype ID, Twitter ID
- phone number
- service identifier: WSDL, WSIL, UDDI
- “plumbing contractor kitchener ontario”

Conflicting Goals

- uniqueness, permanence - *identifier*
 - numbering scheme, large range
- access / communication - *locator*
 - location-dependent name (distance metric?)
 - efficient processing for forwarding
- description - *descriptor*
 - precision
 - processing overhead

Discovery

- human-oriented description
 - expressive vs. concise
- vs. machine processing, efficiency
 - basic semantic schism: human vs. machine
- essential mechanism: search
 - range queries
 - typos, unclear intent
 - multiple results

Name Processing

- At any point during name processing:
scope, resolution, and context are implicit!
- if explicit, they are described by a label
... which is mapped into something else
... therefore the label is a name
... which needs scope, resolution, and context
=> recursive contradiction!

Abstract Names

- role name
 - 'mkarsten' vs. 'root'
 - 'joe@email.com' vs. 'admin@isp.com'
- service name
 - print.cs.uwaterloo.ca
 - ipp://print.cs.uwaterloo.ca:631/printers/color


Distributed Naming System

- name assignment
 - alias support
- name resolution
 - resolution overhead
 - storage overhead
 - scalability, caching
 - size of name space, number of managed entities
 - frequency of updates and lookups
 - relative vs. absolute names

Example: DNS

- Domain Name System
- hierarchical host names in the Internet
 - example: `cpu08.student.cs.uwaterloo.ca`
 - sequence of labels, written with separator
 - cf. file system name: `/home/mkarsten/cs655/network.pdf`
 - naming conventions
 - top level domain: country code or orga-type code
 - 2nd level: organizational name
 - etc. - local conventions (org units, depts, ...)
 - lowest level: role-based name (www, print, cpu, ...)

DNS Database




- distributed database - entry point: root server(s)
 - hierarchy (mostly) follows naming hierarchy
 - resolution requests traverse hierarchy
- local caching

Caching

- local DNS server
 - not part of authoritative hierarchy
 - caching proxy for DNS requests
 - often co-located with regular DNS server
 - “default name server”
- hierarchy of caching servers
- DNS updates are not frequent
 - and not expected to propagate fast


DNS Lookup

- iterative resolution
 - server replies with name/address of next server to contact
 - resolver can cache multiple intermediate results



DNS Lookup

- recursive resolution
 - server forwards request and replies with ultimate response
 - high-level servers can cache lots of results
 - load on high-level servers?
- recursive vs. iterative?
=> mix & match!



DNS Records

- DNS resource records (RR)
name, value, type, ttl
- Type A: hostname -> IP address
- Type NS: domain name -> name server name
- Type CNAME: hostname -> hostname (alias)
- Type MX: email domain -> mail server name
- etc.

DNS – Notes

- simple request/reply protocol using UDP
 - retransmission, message identifiers
 - recursion optional
- ownership & regulation
 - IP address has technical meaning
 - DNS name has business value
- security & authentication
 - 'www.personalbank.com' redirected to bad party?

DNS – Notes

- replicated root servers
- relative names?
 - difficult, because resolution logically starts at root
- special functionality
 - multiple entries for name
 - return random value -> basic load balancing
 - server location -> geographical load balancing
- relevance in the age of Google?
 - how important is it to own `www.mycompany.com`?


Example: Distributed Hash Table

- flat identifier space for peer-to-peer applications
- DHT: distributed database for (key,value) pairs
- peers can insert (key,value) pairs
- peers query with key
 - DB returns with value that matches key
- identifier for each peer in $[0 \dots 2^n - 1]$
- keys are taken from the same range
 - e.g., use hashing

Location of Values


- assign (key,value) to immediate successor peer of key, i.e., peer with smallest larger ID
 - example: $n = 4$, peers: 1, 3, 4, 5, 8, 10, 12, 14
 - key = 14, successor peer = 14
 - key = 15, successor peer = 1
- => logical ring

Logical Ring Structure



- virtual network (“overlay”)
- linear search?

Circular DHT with Shortcuts



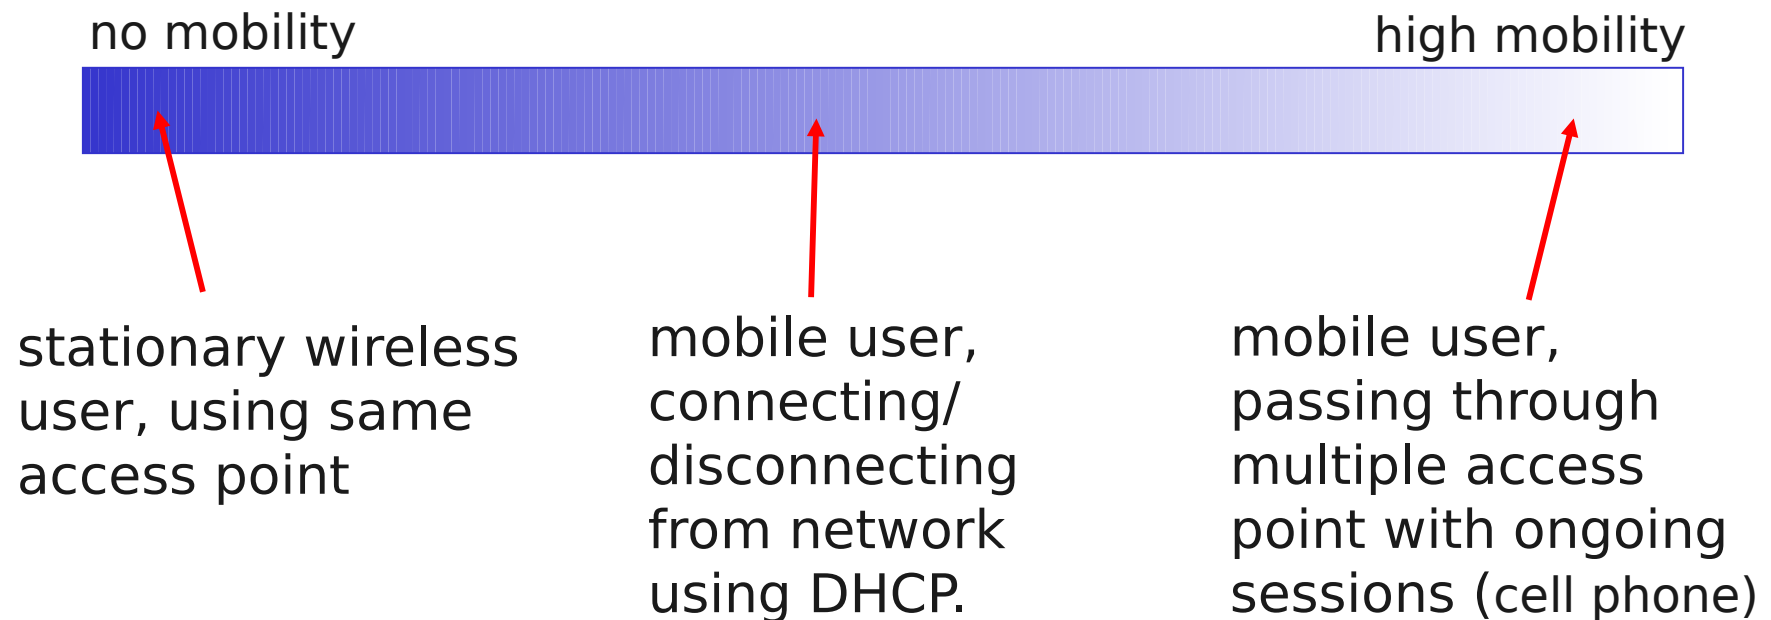
- each peer has shortcuts to $O(\log N)$ peers
 - average distance 1, 2, 4, 8, 16, etc. (randomized)
- expected lookup in $O(\log N)$ steps

DHT Challenges

- decentralized management
- join / leave operations
 - periodic ping of successors
 - eliminate peers that have left from structure
 - join needs starting point and hook into the structure
- security?
- overhead with high churn?

Mobility

- spectrum of mobility, from network perspective



Mobility

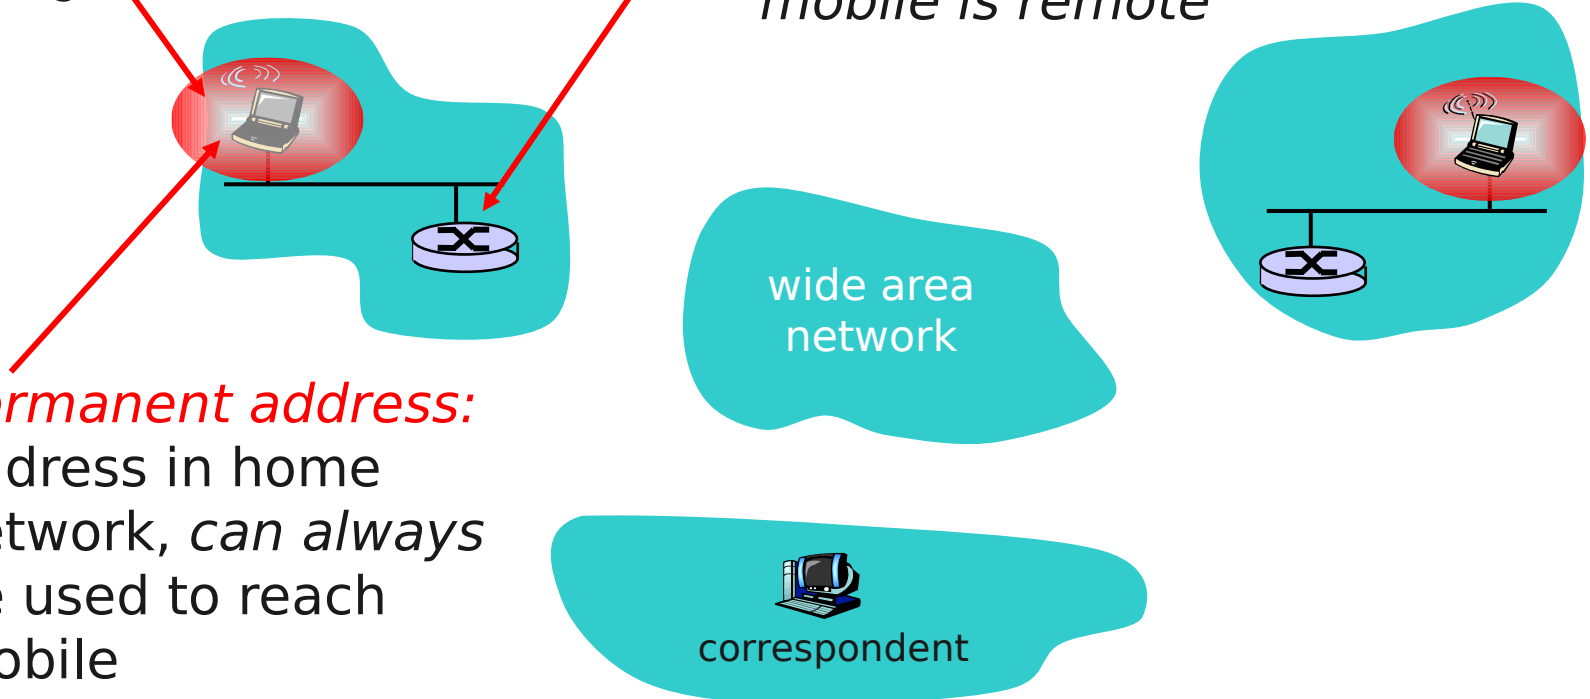
- network access point
 - nearest network-level stationary router
 - might or might not be wireless access point
- challenges
 - initial lookup of responder
 - needs some form of registry
 - ongoing connectivity with movements
 - needs some form of update/redirection
 - similar, but different: time scale – session vs. packet

Example: Mobile IP

home network:
permanent “home” of
mobile
(e.g., 128.119.40/24)

home agent: entity that will
perform mobility functions
on behalf of mobile, when
mobile is remote

Permanent address:
address in home
network, *can always*
be used to reach
mobile
e.g., 128.119.40.186

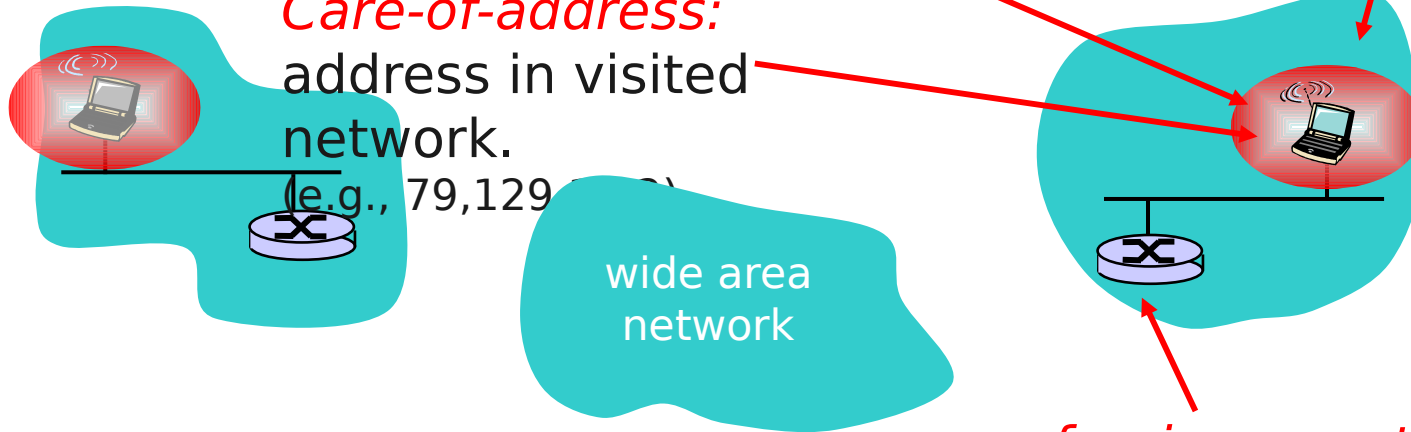


Mobile IP

Permanent address:
remains constant (e.g.,
128.119.40.186)

visited network: network
in which mobile
currently resides (e.g.,
79.129.13/24)

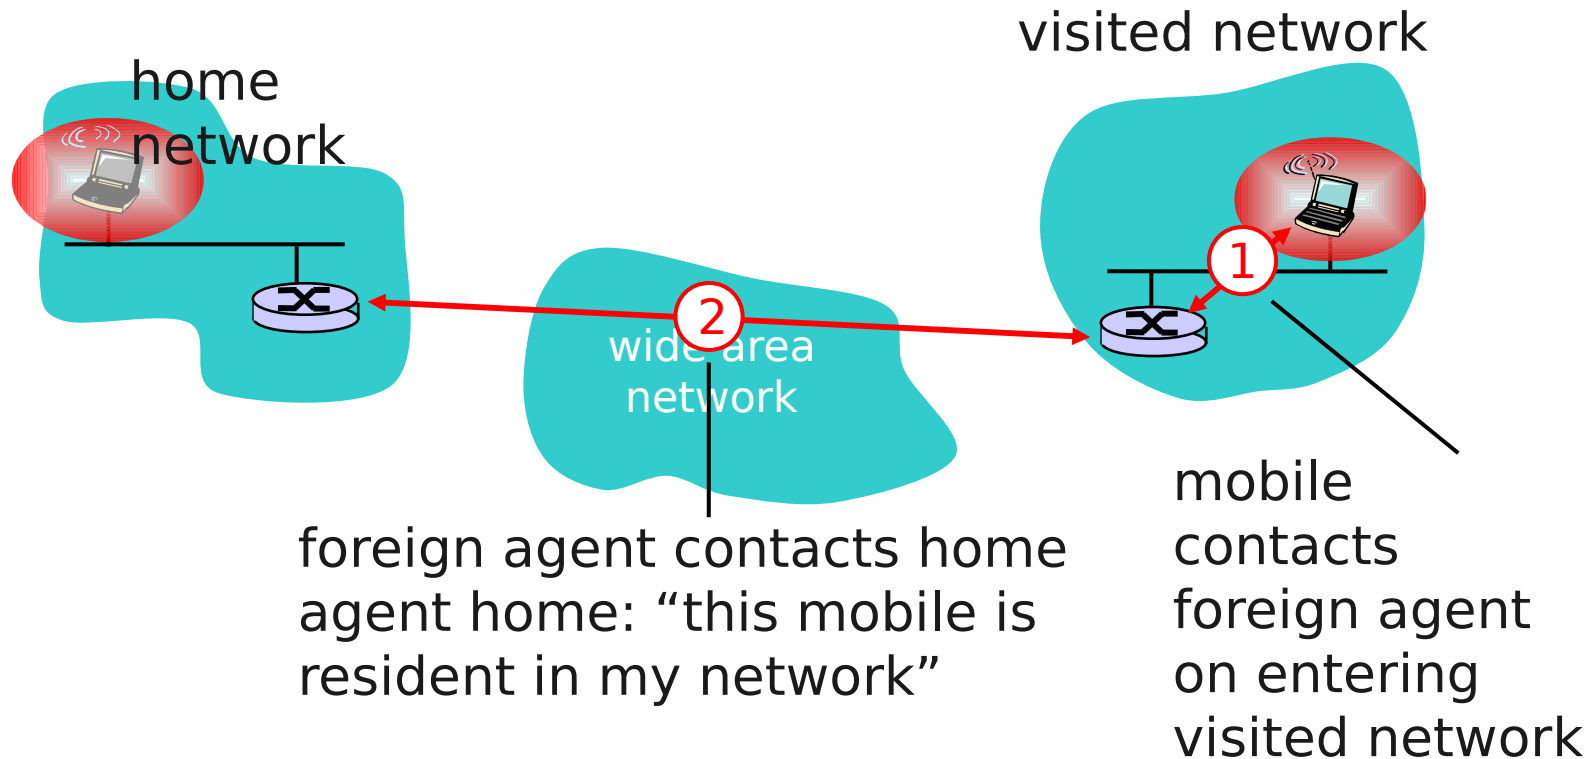
Care-of-address:
address in visited
network.
(e.g., 79.129.13.1)



correspondent:
wants to
communicate with
mobile

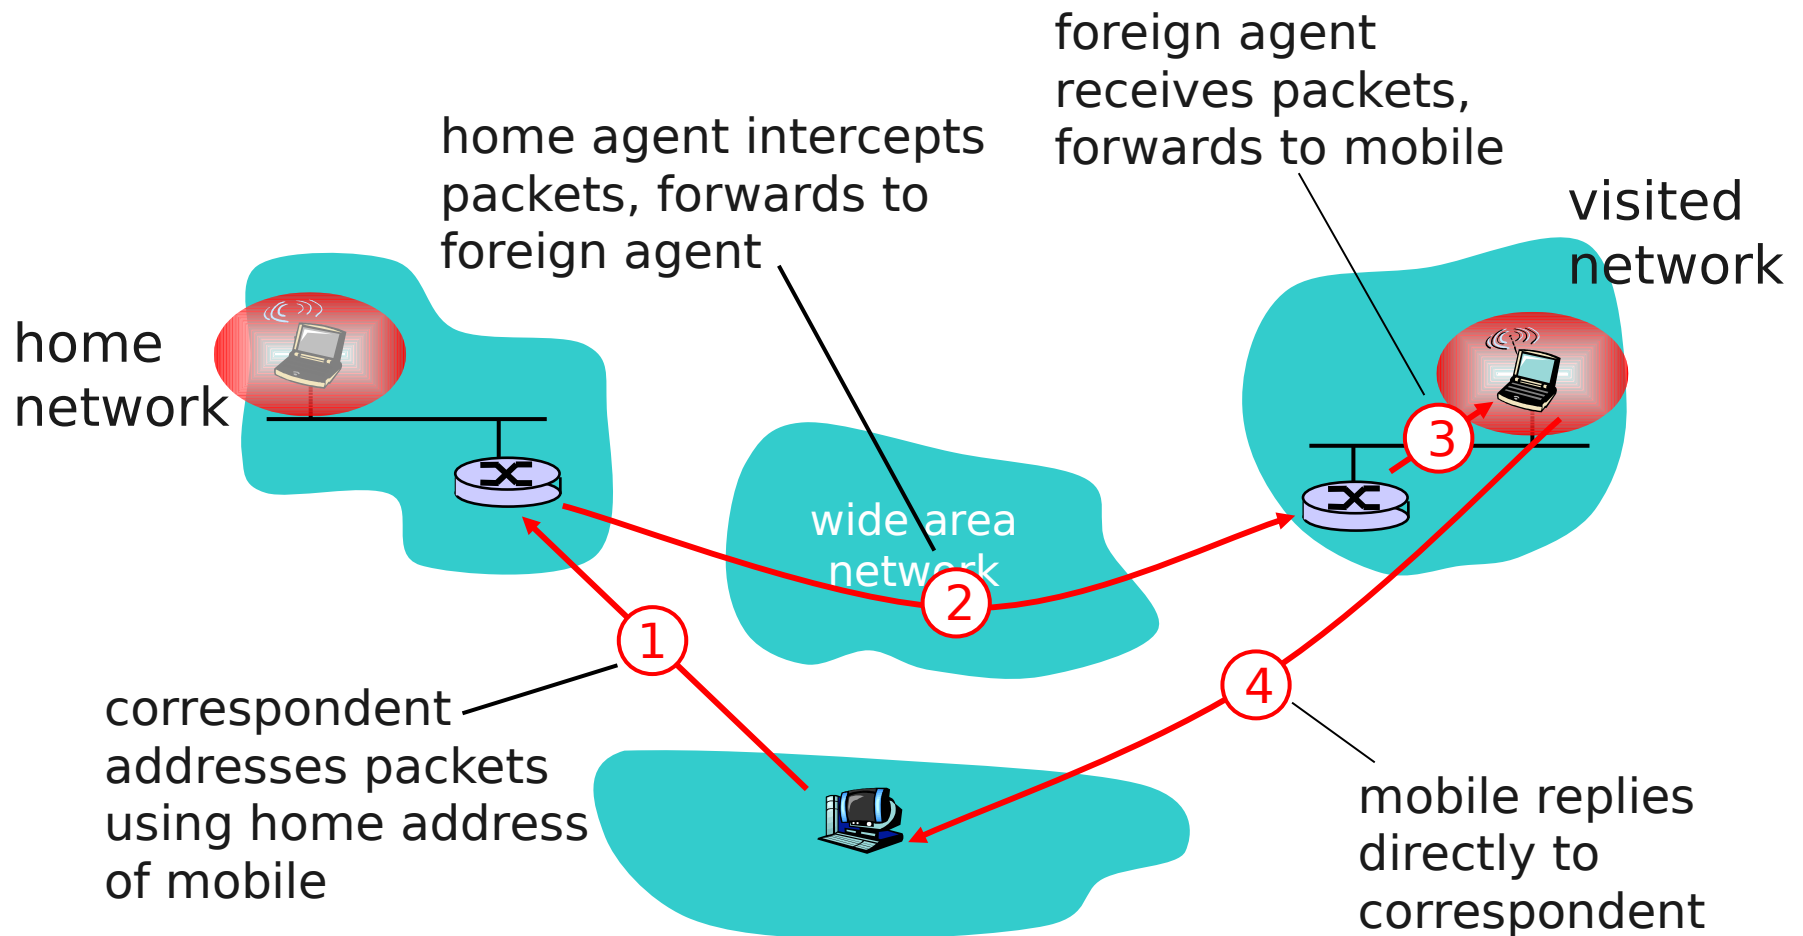
foreign agent:
entity in visited
network that
performs mobility
functions on behalf
of mobile.

Mobile IP – Registration

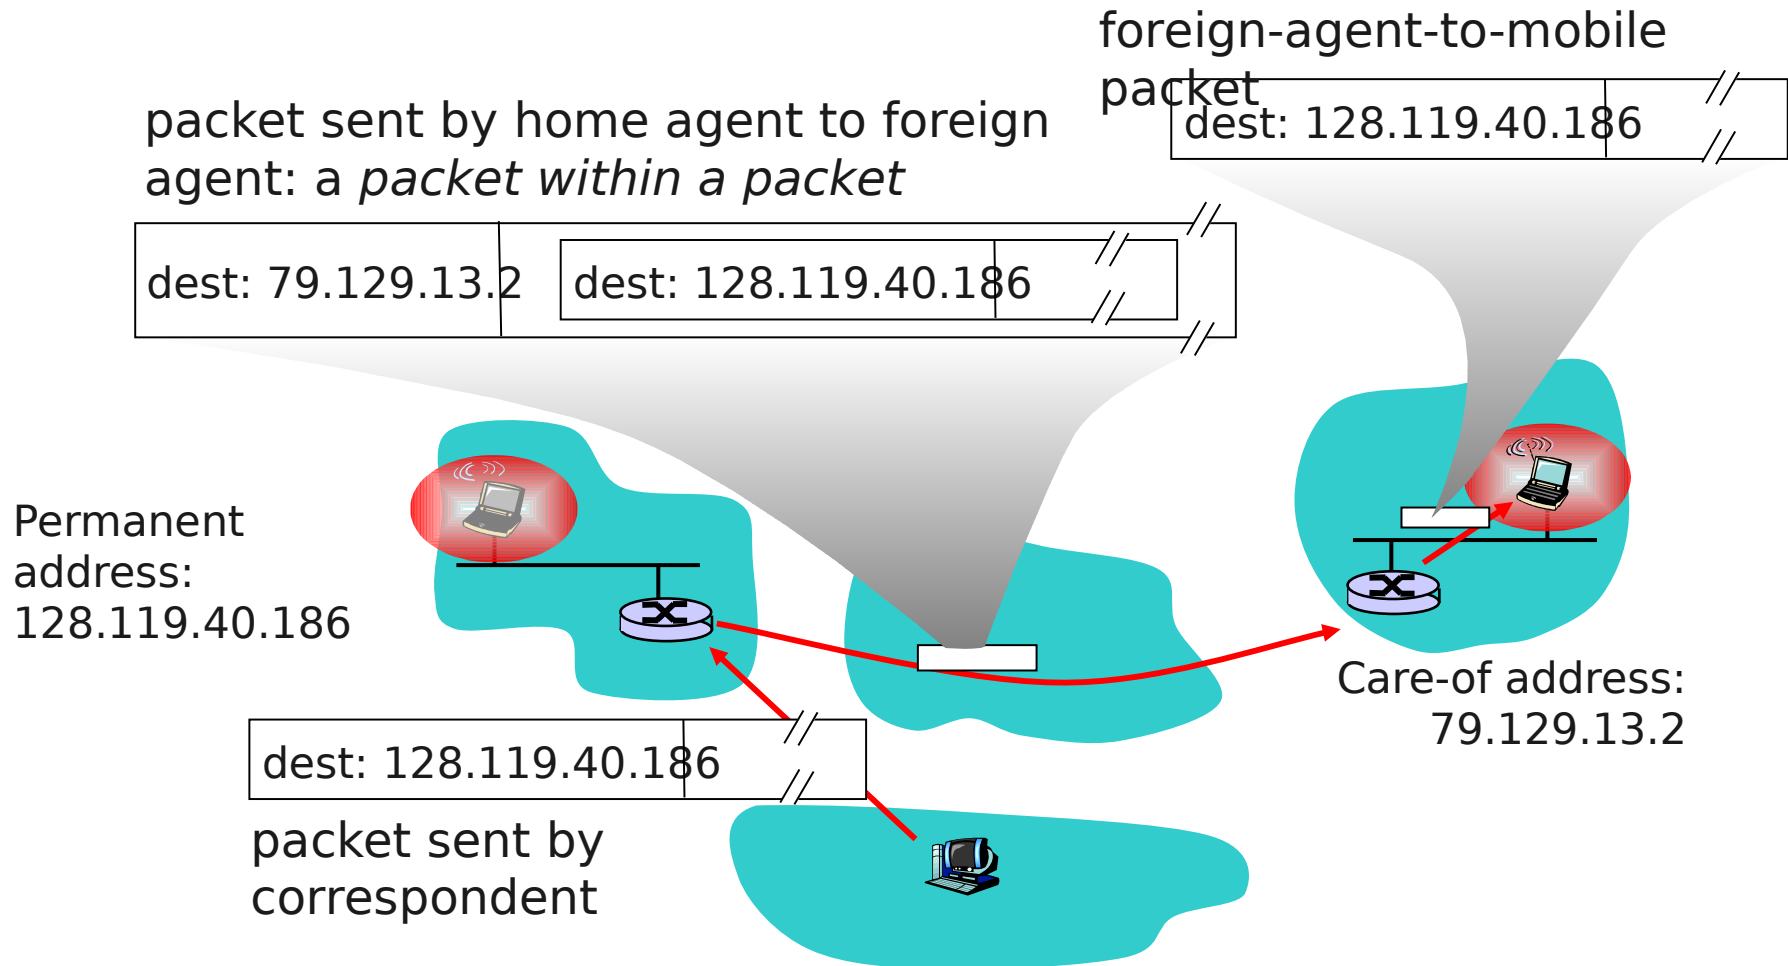


- foreign agent knows about mobile
- home agent knows location of mobile

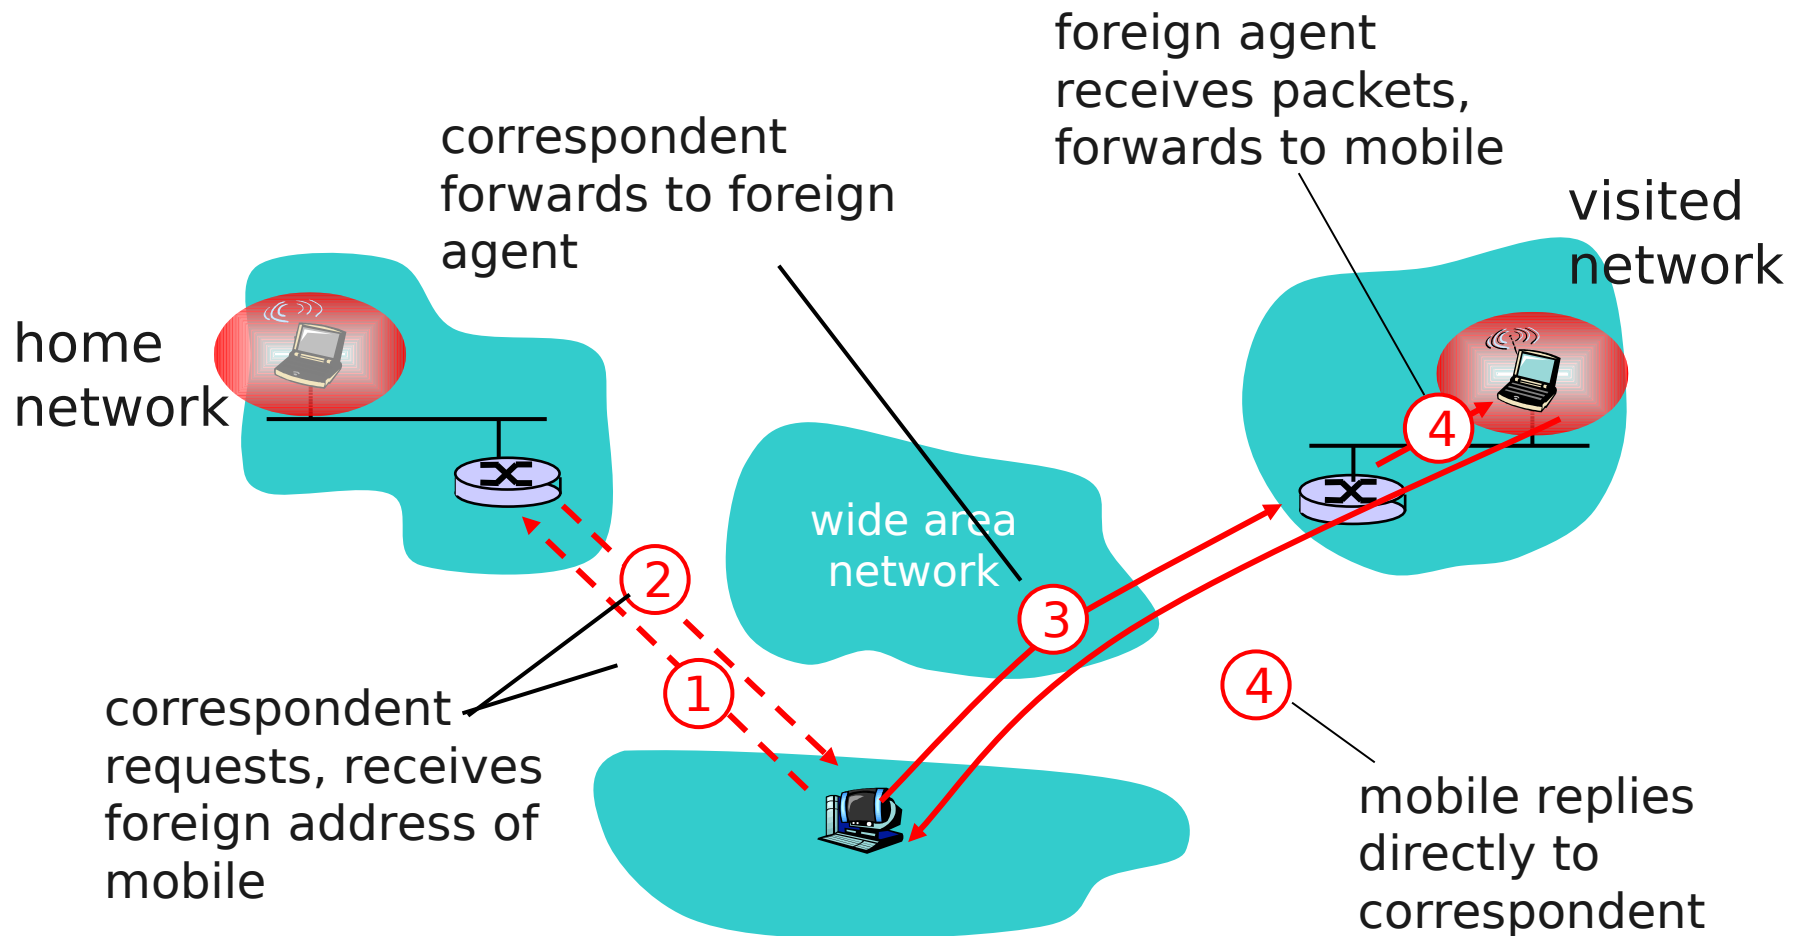
Mobile IP – Indirect Routing



Mobile IP – Indirect Routing




Mobile IP – Direct Routing



Mobile IP – Hierarchical Agents

- fast handover support – avoid end-to-end delay
- can be extended arbitrarily



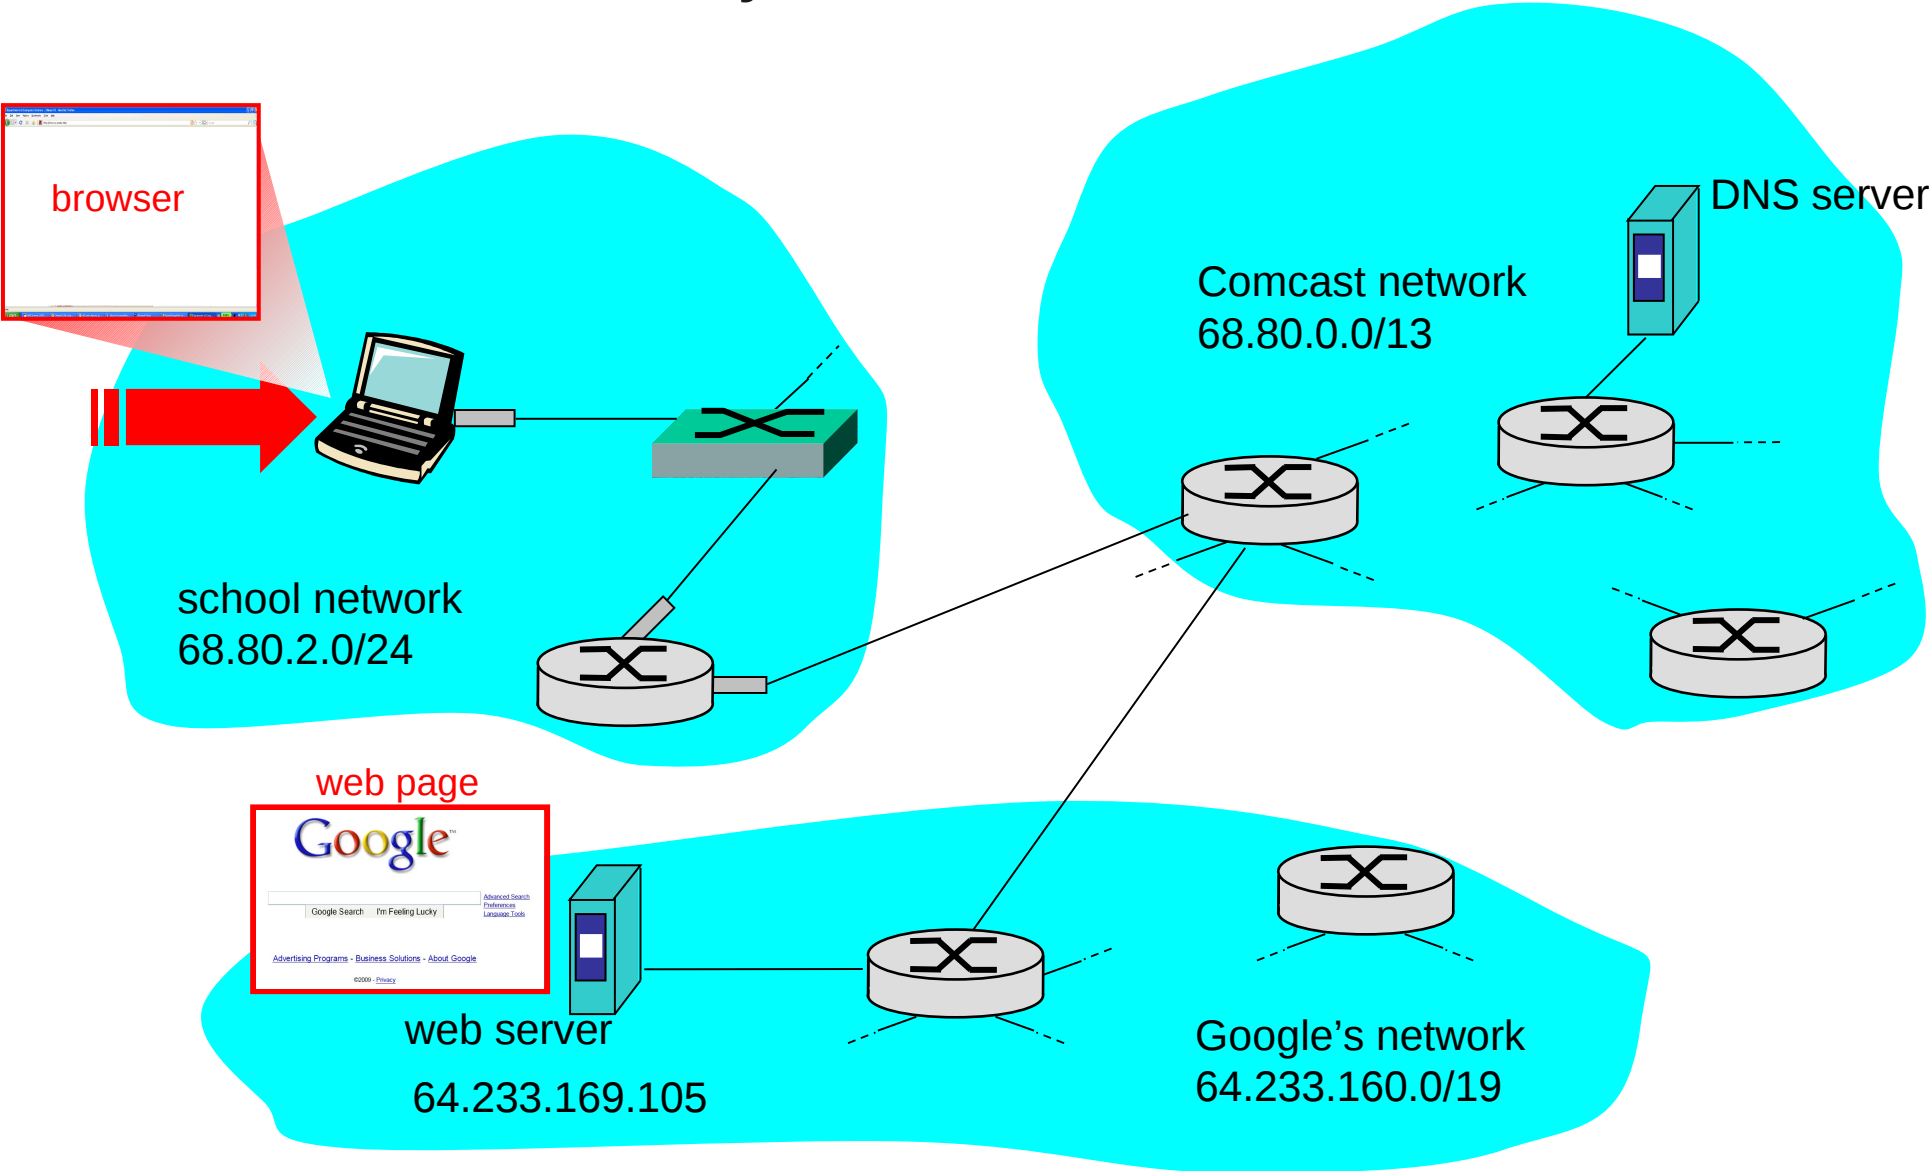
Mobile IP – Discussion

- overarching design consideration
 - transparency, backward-compatibility
 - triangle routing direct reply – source addr checks?
- foreign agent functionality can be with mobile
 - e.g., with address allocation via DHCP
- Mobile IP uses IP tunneling
 - could use direct forwarding with address rewriting
 - ... similar to reverse NAT
 - ... similar to virtual circuit -> similar to GSM


A Day in the Life of a Web Request

- journey through basic functionality complete!
 - channel, network, transport, naming
- putting-it-all-together: synthesis!
 - *goal*: identify, review, understand protocols (at all layers) involved in seemingly simple scenario: requesting www page
 - *scenario*: student attaches laptop to campus network, requests/receives `www.google.com`

A day in the life: scenario




A day in the life... connecting to the Internet



- connecting laptop needs to get its own IP address, addr of first-hop router, addr of DNS server: use **DHCP**
- DHCP request **encapsulated** in **UDP**, encapsulated in **IP**, encapsulated in **802.1** Ethernet
- Ethernet frame **broadcast** (dest: FFFFFFFF) on LAN, received at router running **DHCP** server
- Ethernet **demux'ed** to IP demux'ed, UDP demux'ed to DHCP


A day in the life... connecting to the Internet



- DHCP server formulates **DHCP ACK** containing client's IP address, IP address of first-hop router for client, name & IP address of DNS server
- encapsulation at DHCP server, frame forwarded (**switch learning**) through LAN, demultiplexing at client
- DHCP client receives DHCP ACK reply


Client now has IP address, knows name & addr of DNS server, IP address of its first-hop router

A day in the life... ARP (before DNS, before HTTP)



- before sending **HTTP** request, need IP address of www.google.com: **DNS**
- DNS query created, encapsulated in UDP, encapsulated in IP, encapsulated in Eth. In order to send frame to router, need MAC address of router interface: **ARP**
- **ARP query** broadcast, received by router, which replies with **ARP reply** giving MAC address of router interface
- client now knows MAC address of first hop router, so can now send frame containing DNS query

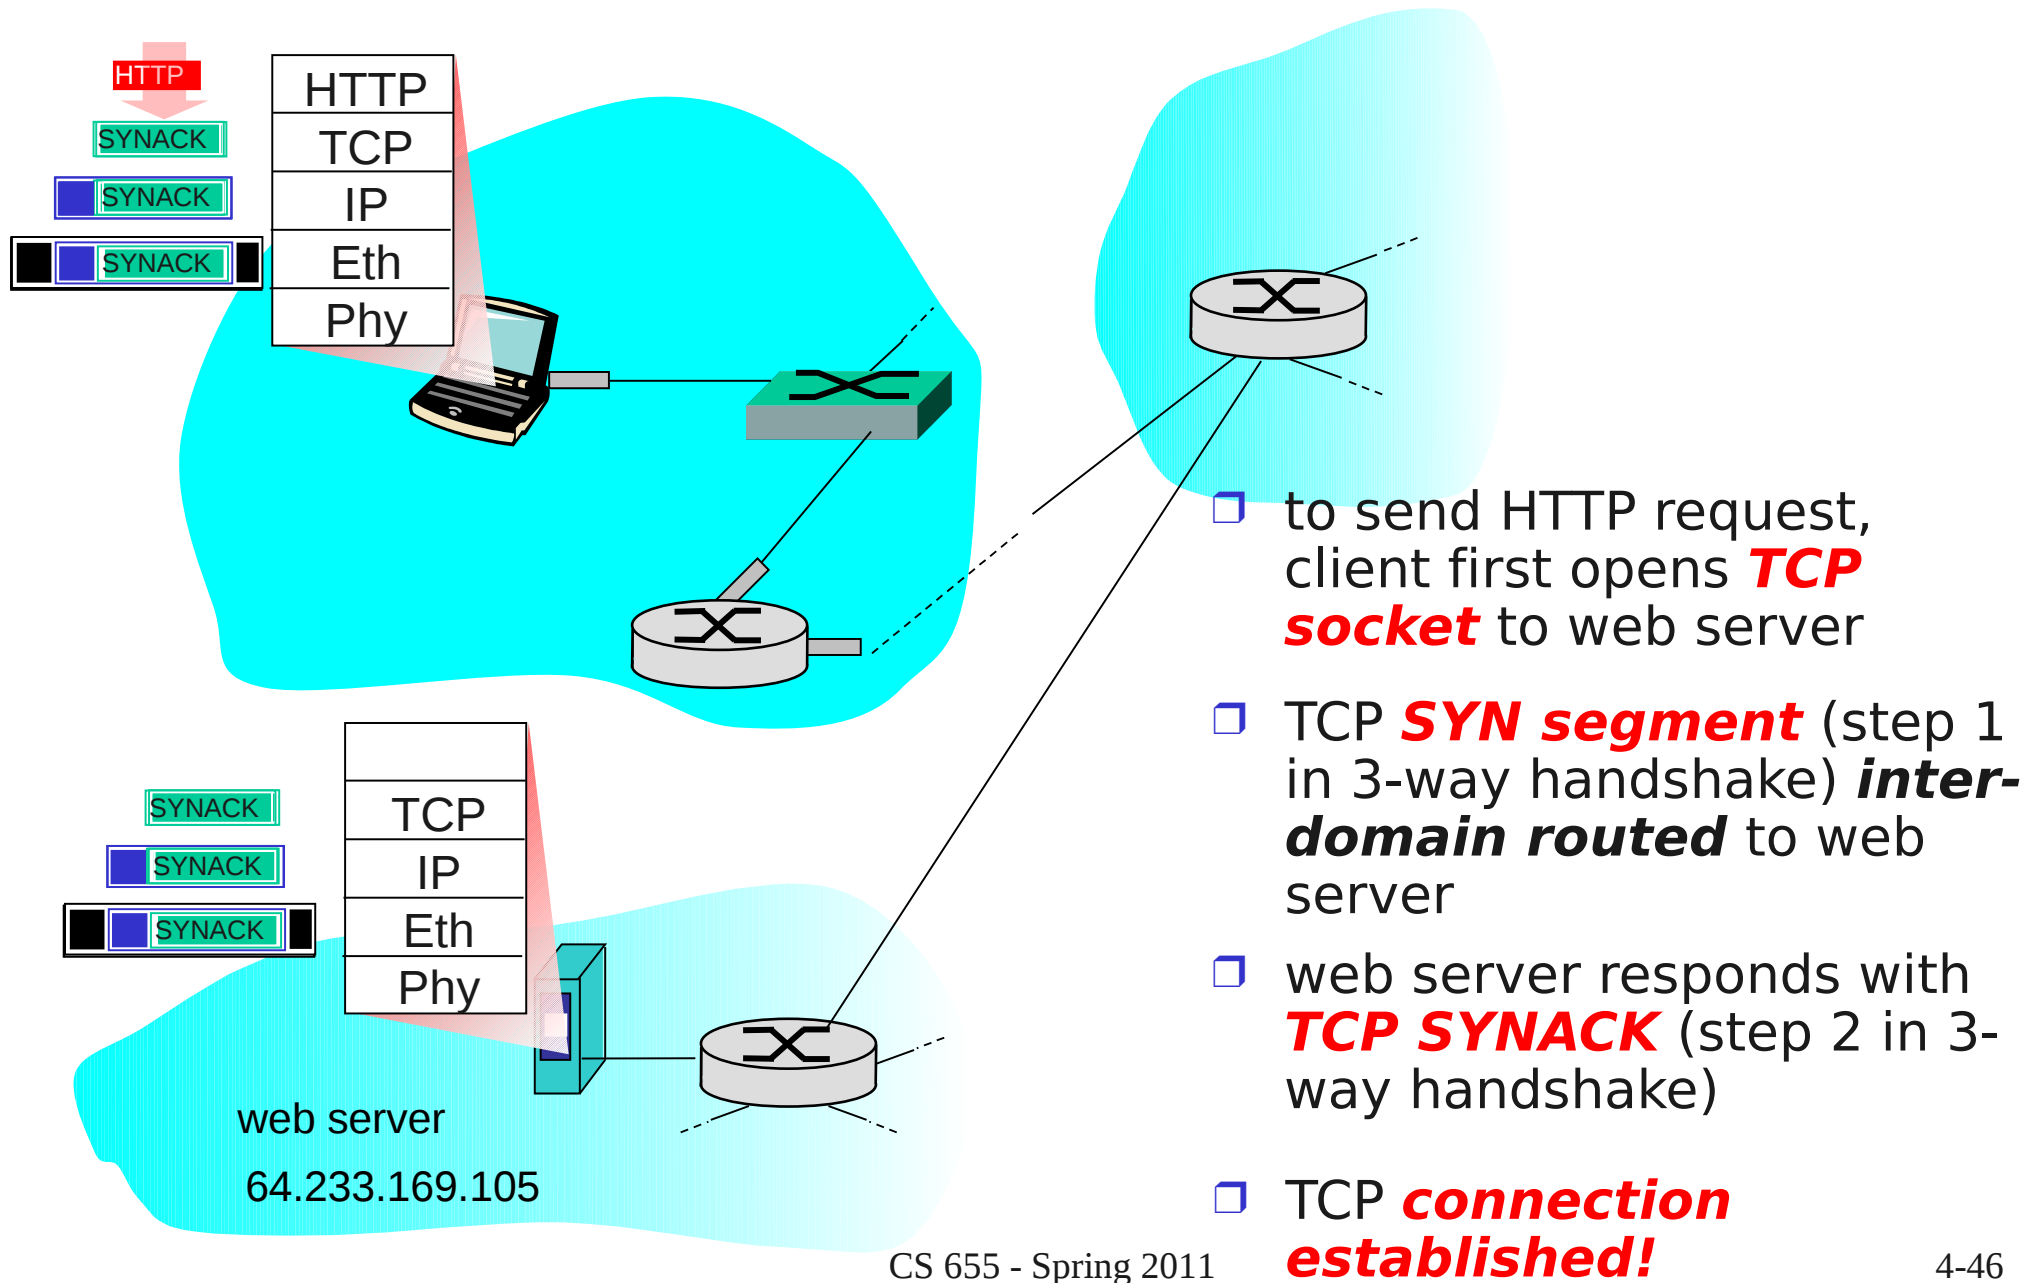
A day in the life... using DNS



- IP datagram containing DNS query forwarded via LAN switch from client to 1st hop router

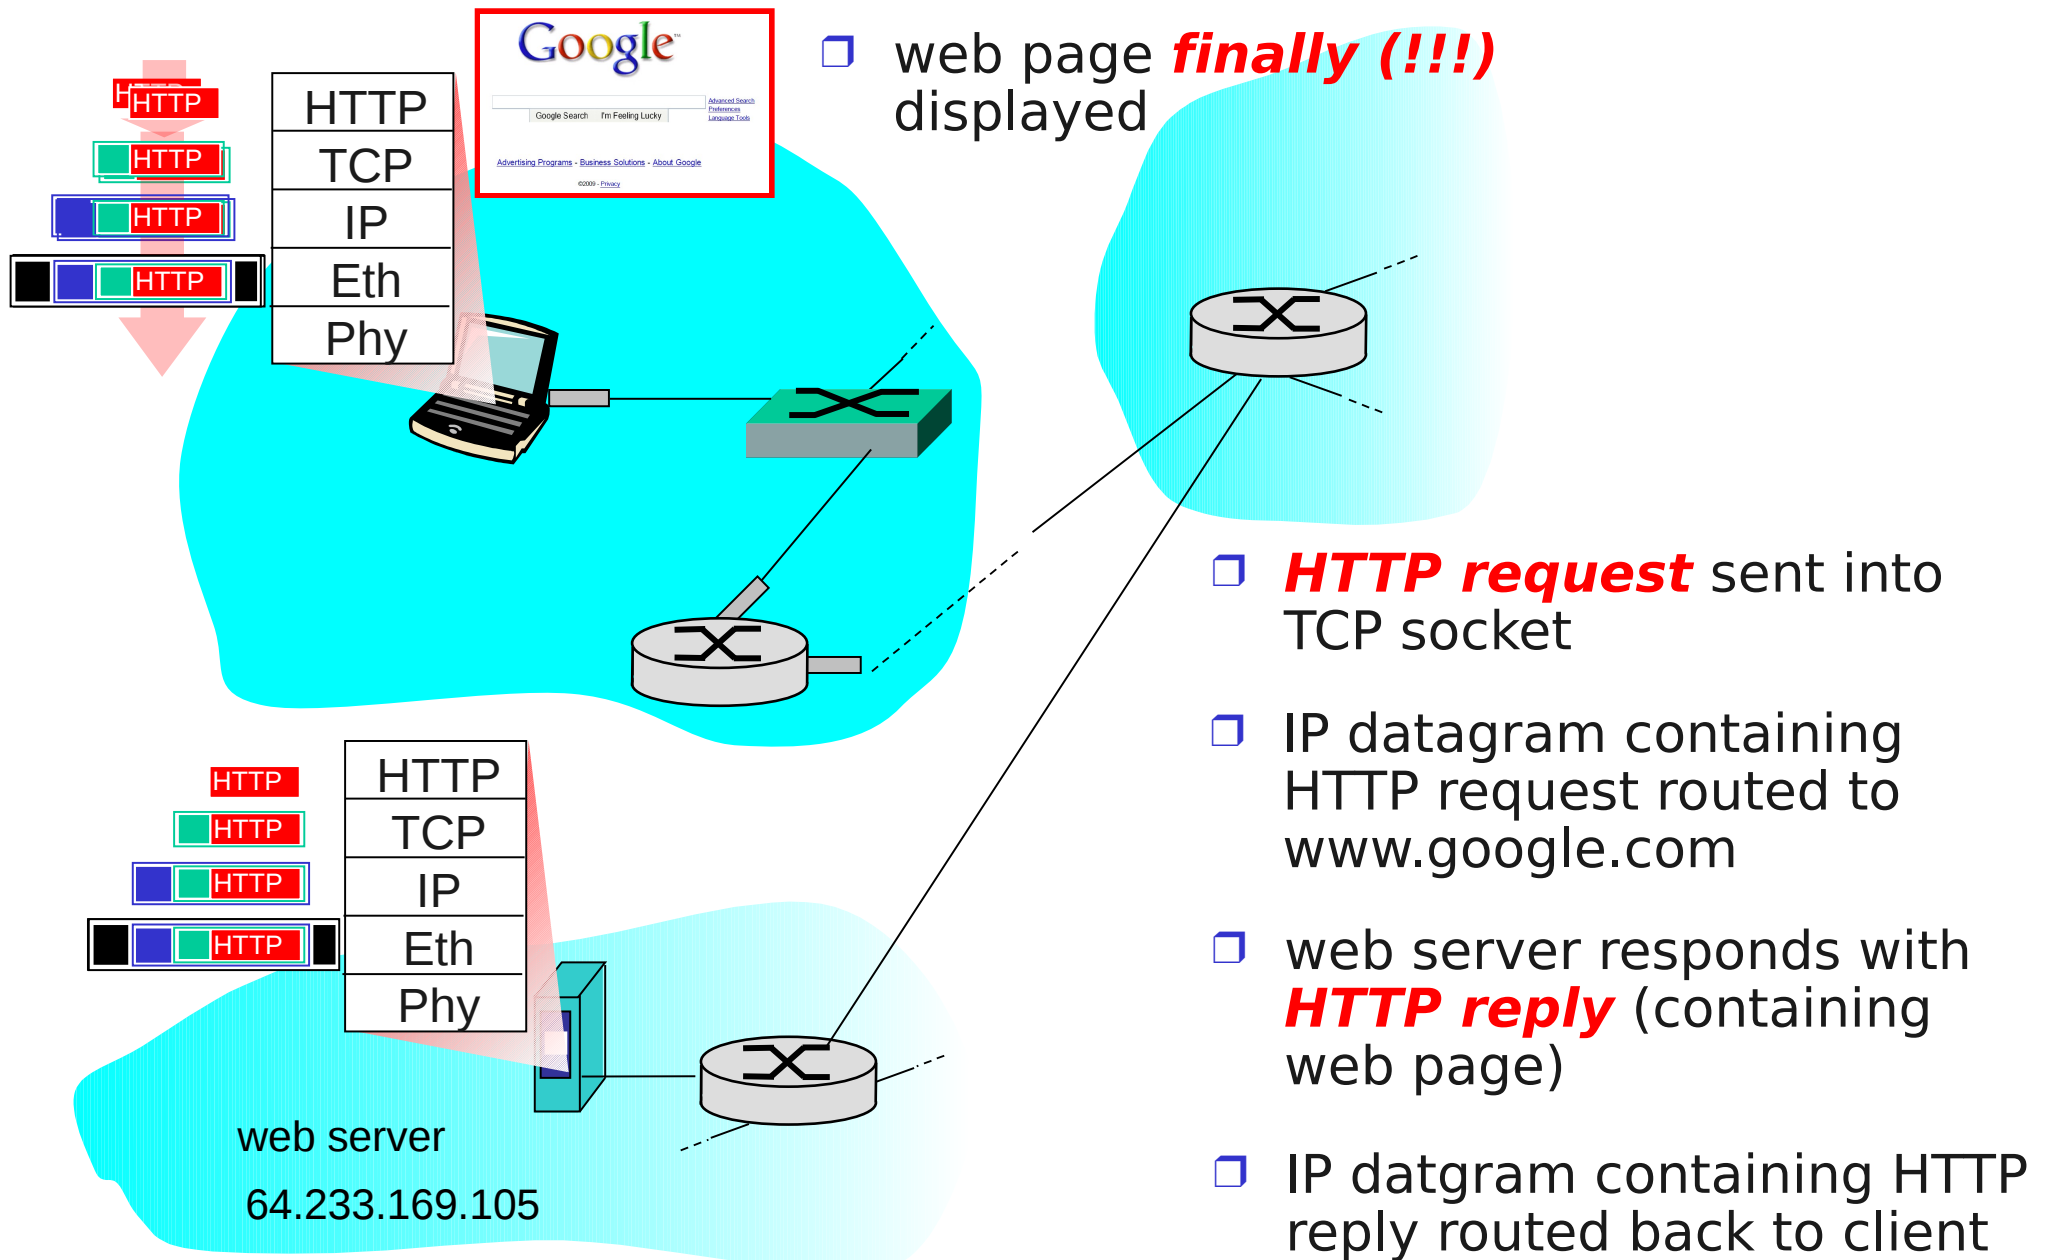
- IP datagram forwarded from campus network into comcast network, routed (tables created by **RIP, OSPF, IS-IS** and/or **BGP** routing protocols) to DNS server
- demux'ed to DNS server
- DNS server replies to client with IP address of www.google.com

A day in the life... TCP connection carrying HTTP



A day in the life... HTTP request/reply

□ web page **finally (!!!)** displayed



Messaging

- *persistent* communication
 - sender can terminate after sending message
 - receiver does not need to be online
 - *vs. transient* communication
- *asynchronous* communication
 - sender can continue other work after sending
 - *vs. sender waits for acknowledgement*
 - receiver is notified when message is available
 - *vs. receiver blocks waiting for message*

Messaging Middleware

- group communication: publish / subscribe
 - underlying distribution model: unicast vs. broadcast
 - multicast: router forwards across to multiple links
- persistence -> reliability?
- error control -> exactly-once delivery
 - more later
- flexible integration with heterogeneous systems
 - OS, network, programming language, etc.

Messaging Queueing Primitives


- Put – send message to queue
- Get – block to receive message from queue
- Poll – check (non-blocking) for message
- Notify – install asynchronous receive handler

- need buffer decoupled from sender, receiver
- relay nodes for larger networks
 - addressing, routing, forwarding, etc., as usual

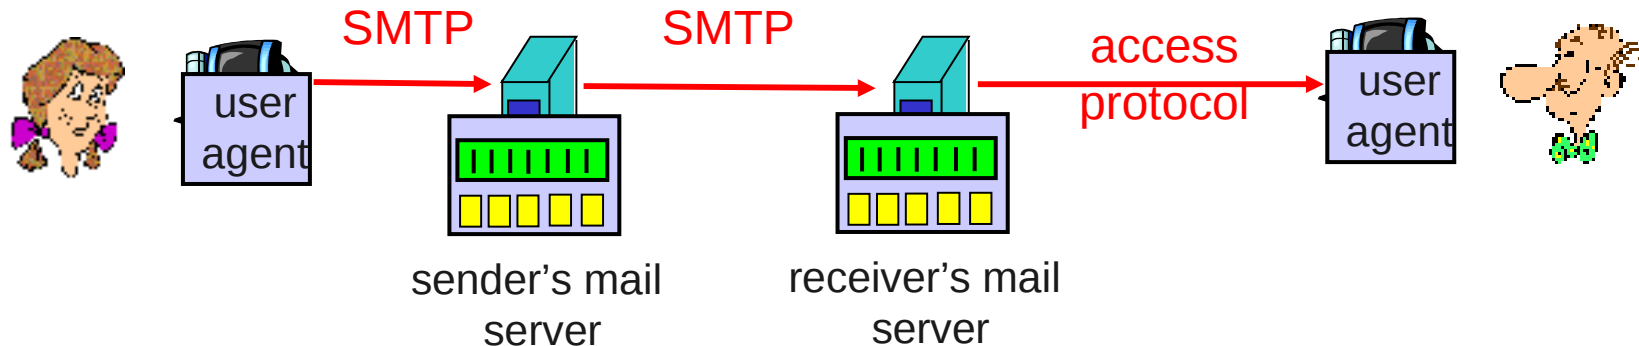
Example: Email

mail servers

- incoming messages mailbox
- outgoing message queue
- communication protocol: SMTP
 - reliable server-to-server transfer



Email Access Protocols



- sender: synchronous, transient to server
- receiver: asynchronous, persistent from server
 - Post Office Protocol (POP) – old & simple
 - Internet Mail Access Protocol (IMAP) – better
 - HTTP – POP, IMAP, etc in background
 - remote file system and file-based (elm, pine, etc.)