UBICOMP

Episode 12: TCP/IP & UbiComp

Hannes Frey and Peter Sturm
University of Trier

Outline

- Introduction
- Mobile IP
- TCP and Mobility
- Conclusion

References

[1] James D. Solomon, “Mobile IP: The Internet Unplugged”, Prentice Hall, 1998

[2] Hala ElAarag, “Improving TCP Performance over Mobile Networks”, ACM Computing Surveys, Vol.34, No.3, September 2002, pp. 357-374

UBICOMP

Introduction

Mobile Networking

- Multihop Ad-Hoc Networks, Sensor Networks
  - No wired infrastructure needed
- Supporting mobility for existing wired infrastructure
  - Connecting mobile nodes with Base Stations
  - Closing the gap between multihop ad-hoc networks and wired infrastructure
- Single-hop ad-hoc network architecture
  - Mobile Host (MH), Fixed Host (FH), Wired Backbone, Wireless Network, Cellular Network Infrastructure, Base Station (BS), Handoff

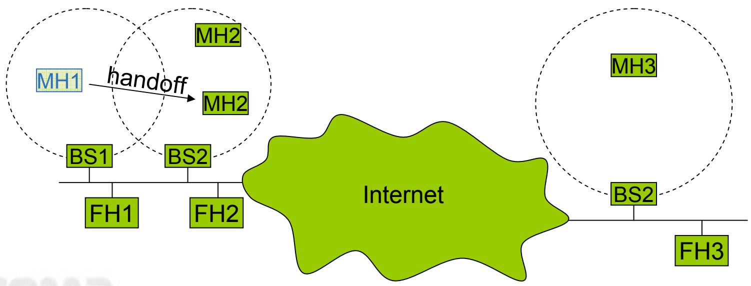

Protocol Design Strategies for Mobile Networking

- Multihop Networks
  - No legacy applications/protocols
  - Design from scratch possible
- Single Hop Networks
  - Redesign of the global Internet?
  - Mobility support as an extended feature for the Internet
    - Should not affect fixed hosts and wired backbone
  - Modify existing solutions
    - Mobile IP
    - Adapt TCP for mobility

UBICOMP

Mobile IP

IP Routing Revised

- The network protocols IP and ICMP
- IP packet: Header and Payload
- IP Addresses
  - 32-Bit number assigned to each network interface of a node
  - Dotted Decimal, Network Prefix, and Host Portion
    - Shorthand-notation subsequently used: address/prefix-length
- Routing Tables
  - Mapping of Target/Prefix-Length to Next Hop and Interface
    - All nodes on a link must have identical network prefixes
  - Host-specific, network-prefix, and default routes
- Routing decision based upon IP header and table entries
  - What if there is more than one matching entry?
- (How are those tables created? Statically, ICMP redirect, OSPF, BGP, RIP)

IP Packet: Header and Payload

0-3	4-7	8-15	16-18	19-31
Vers=4	IHL	Type of Service	Total Length	
Identification			Flags	Fragment Offset
Time To Live		Protocol	Header Checksum	
Source Address				
Destination Address				
IP Options …				
Transport Layer (e.g. TCP) …				
Application Layer (e.g. HTTP) …				
Application Layer Data (e.g. Web Page) …				

IP Header
IP Payload

Some more Internet Terminology revised

- Address resolution protocol
  - ARP Request, ARP Reply, and ARP cache
  - Proxy ARP
  - Gratuitous ARP
- Tunneling

IPsrc = Original Sender, IPdst = Ultimate Destination

| Header | Payload |

IPsrc = Tunnel Entry Point, IPdst = Tunnel Exit Point

| Outer Header | Header | Payload |

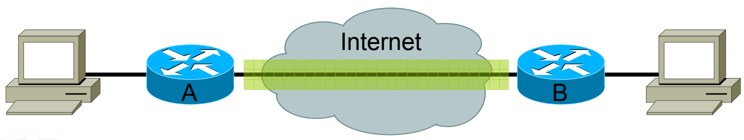

Traditional IP Routing

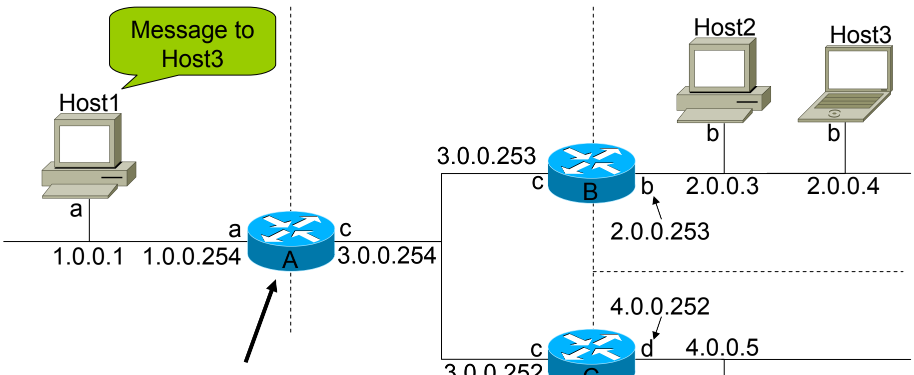

1.0.0.0/24	direct	a
3.0.0.0/24	direct	c
2.0.0.0/24	3.0.0.253	c
4.0.0.0/24	3.0.0.252	c

What Happens When a Host Moves?

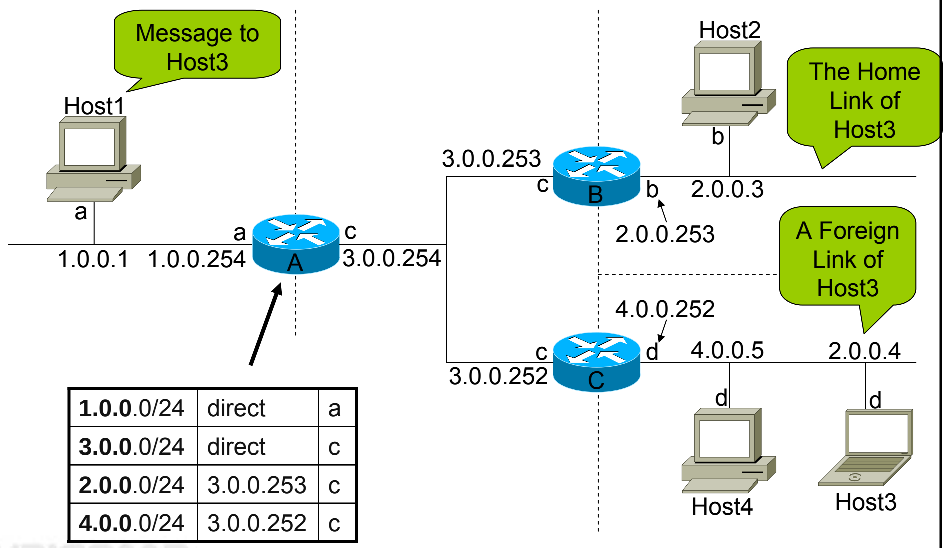

1.0.0.0/24	direct	a
3.0.0.0/24	direct	c
2.0.0.0/24	3.0.0.253	c
4.0.0.0/24	3.0.0.252	c

Host-Specific Routing is a Possible Solution

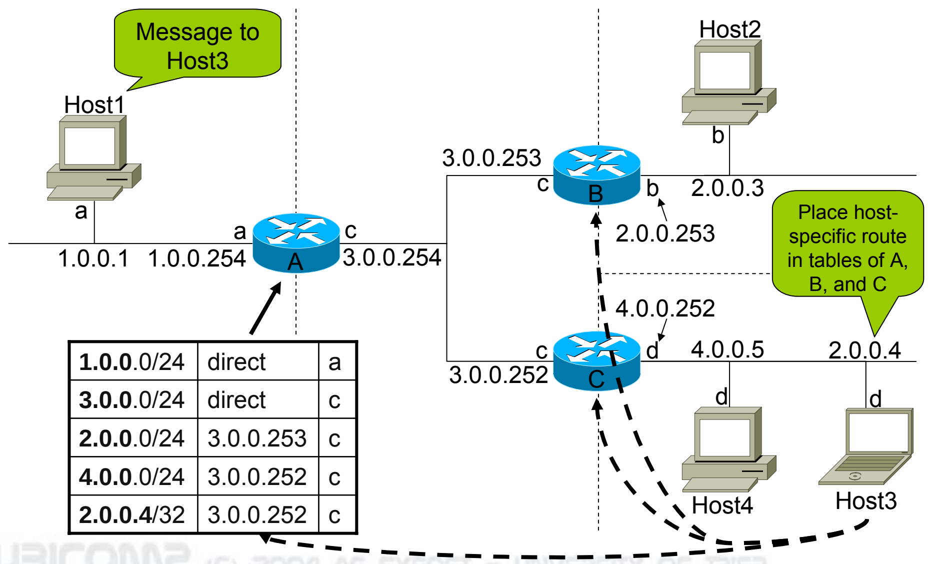

1.0.0.0/24	direct	a
3.0.0.0/24	direct	c
2.0.0.0/24	3.0.0.253	c
4.0.0.0/24	3.0.0.252	c
2.0.0.4/32	3.0.0.252	c

A Lower Bound for Host-Specific Routes

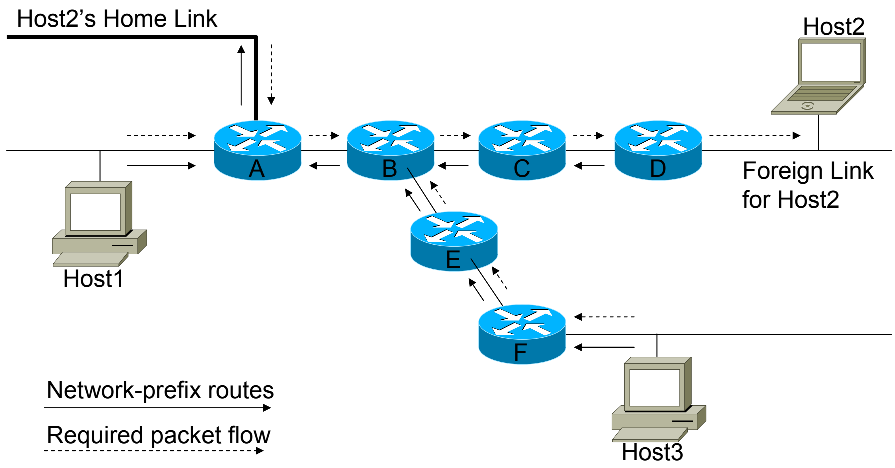

A Lower Bound for Necessary Updates

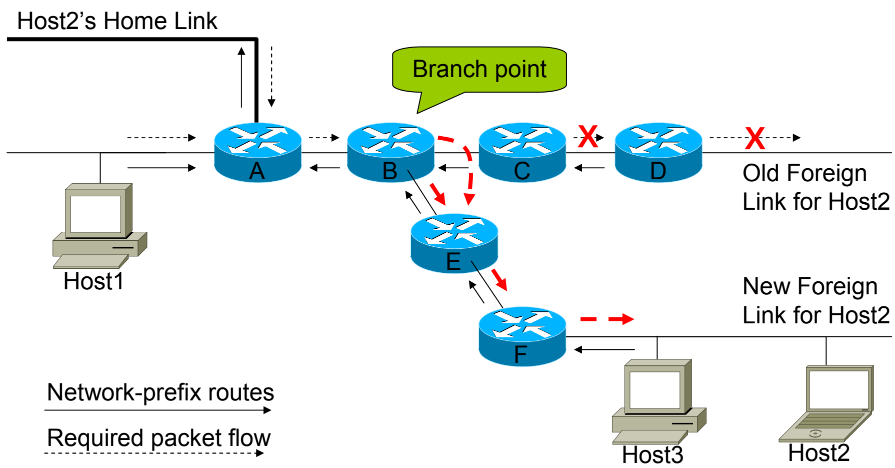

Summarizing Host-Specific Routing

- We have already millions of mobile devices
- Host-specific routes needed at least for all routers along the path from home link to foreign link
  - This may be hundreds of routers
  - Each of those routers becomes a SPOF
- Frequently changing foreign links?
  - From branch point to old link: delete host-specific routes
  - From branch point to new link: add host-specific routes
  - In branch point: modify host-specific route
- A possible attack to this solution
  - Bad guy sending wrong location updates

What About Assigning a new IP Address?

- TCP and UDP use IP addresses as end-point identifiers
- Most IPv4 nodes assume source/destination address/port to be constant
  - Drop connection if destination IP address is changed
- Thus, changing a nodes IP address means
  - Terminate all ongoing communications
  - Initiate new connections with the new address
- Technique is referred as *nomadic computing*

Why is Mobility Preferable to Nomadicity?

- DNS queries and updates will increase significantly
  - Update DNS entry after each link change
  - Retrieve IP address periodically
- DNS updates must additionally be secured!
- Legacy applications
  - Configuration databases depending on IP addresses
- Network-licensing which restricts access to nodes possessing specific ranges of IP addresses
- Security mechanisms providing access-privileges to nodes based upon their IP addresses
- Maintaining a pool of IP addresses for assignment to nomadic nodes

Solving Mobility Problem at the Link Layer

- Existing link layer solutions
  - CDPD (Cellular Digital Packet Data)
    - Users are sending IP packets over unoccupied radio channels
    - Each user provided with an IP address for use throughout the CDPD network
    - Different to dial-up access
  - IEEE 802.11
- Link layer solutions are insufficient for a global internet
  - Provides node mobility only in a single type of medium
  - Limited regarding global coverage
  - Switching from wireless to wired networking not possible
  - Single solution for mobility problem preferable to a plethora of medium-specific solutions

The Scope of Mobile IP

- Provide mechanisms necessary to route packets to mobile nodes at the network layer
- Complete mobility solution would involve enhancements of other protocol layers as well!
- Mobile IP characteristics
  - Maintain communication despite of link changes
  - Retain one IP address regardless the current link
  - Enable communication with others not implementing mobile IP
  - Do not require any changes to existing hosts and routers
  - Avoid new security threads
  - Allow heterogeneous mobility
  - Keep size and frequency of routing updates as small as possible
  - Make it as simple as possible to implement even on small devices
  - Avoid waste of limited IPv4 addresses

Mobile IP: Functional Entities

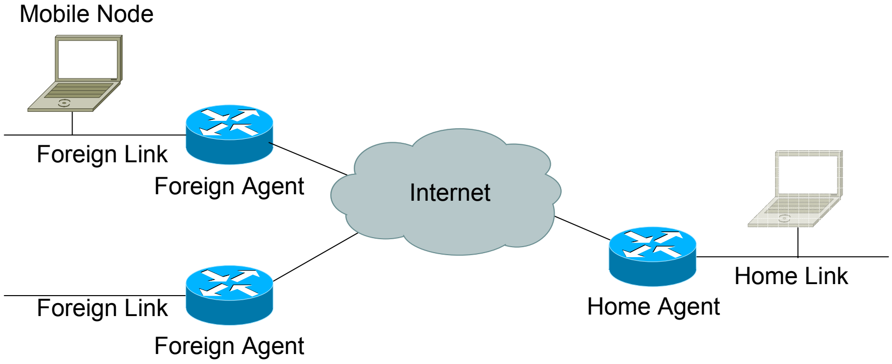

- Mobile IP protocol resides on MN, FA, and HA
- Agents are not necessarily dedicated routers

Mobile IP: Agent Discovery

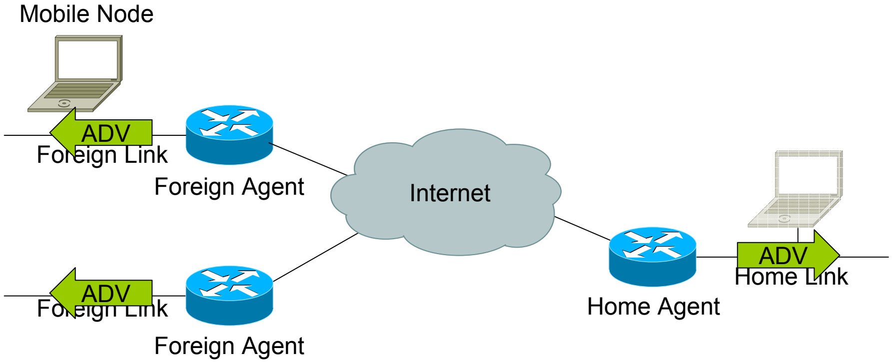

- Periodically broadcast Agent Advertisements received by all nodes on the link

Mobile IP: Agent Discovery

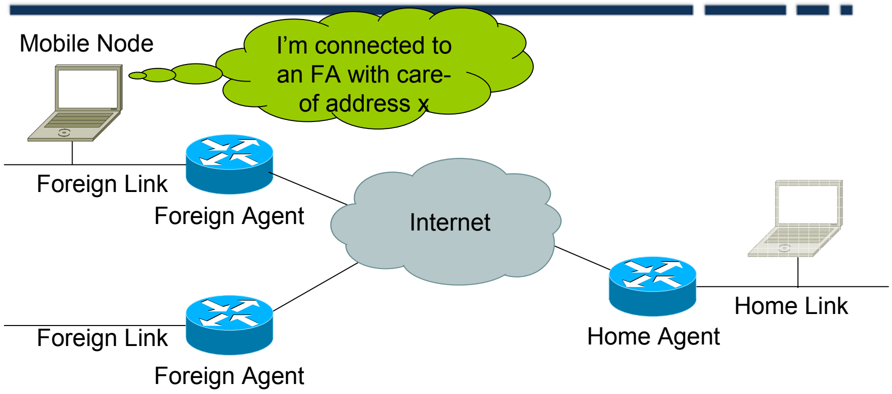

- Use Agent Advertisement to determine whether connected to home or foreign link
  - Home link: make no use of mobile IP (act as stationary node)
  - Foreign link: obtain *care-of address* from Agent Advertisement

UBICOMP (C) 2004 AG SYSOFT – UNIVERSITY OF TRIER

Mobile IP: Registration

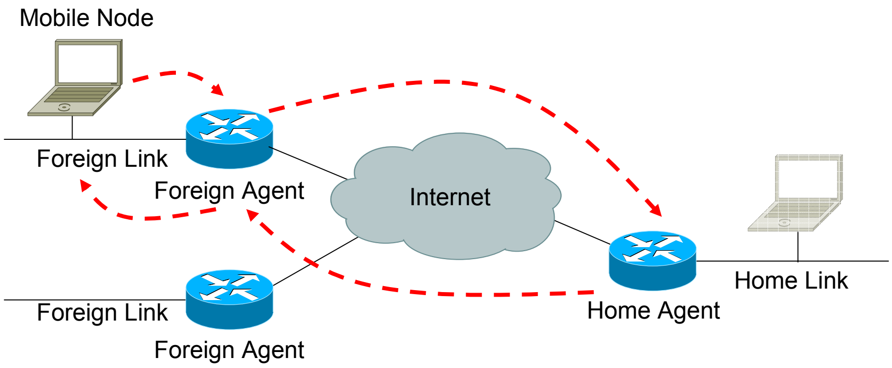

- Ask for service from the foreign agent
- Register care-of address with home agent
- Authentication to prevent remote denial-of-service attack

Mobile IP: Routing to the Mobile Node

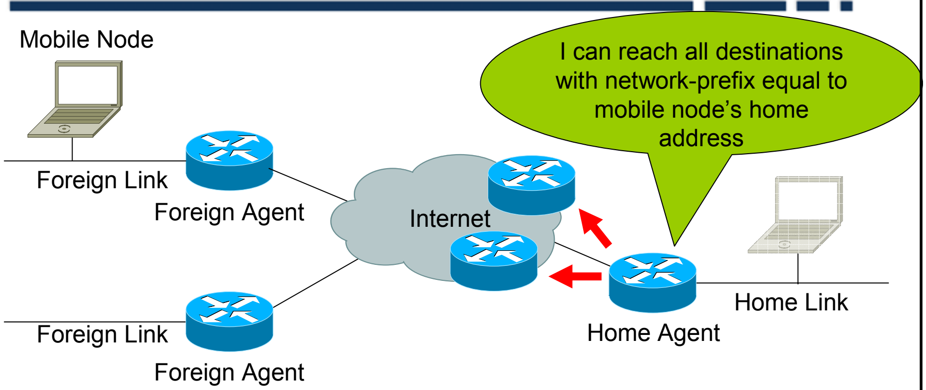

- Home node advertises reachability to mobile node's home address
- Thus attracting all packets destined to mobile node's home address

Mobile IP: Routing to the Mobile Node

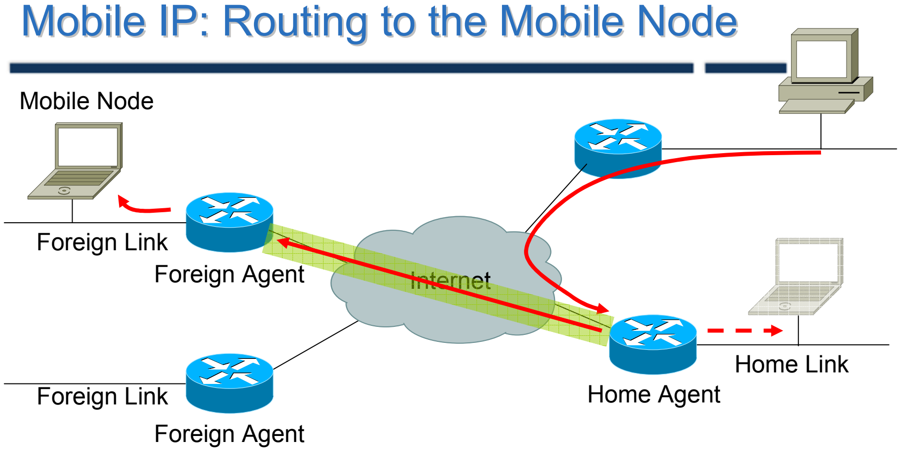

- Home agent intercepts packets destined to the home address and tunnels them to the care-of address
- Foreign Agent extracts packet from the tunnel and delivers it to the mobile node

Mobile IP: Routing from the Mobile Node

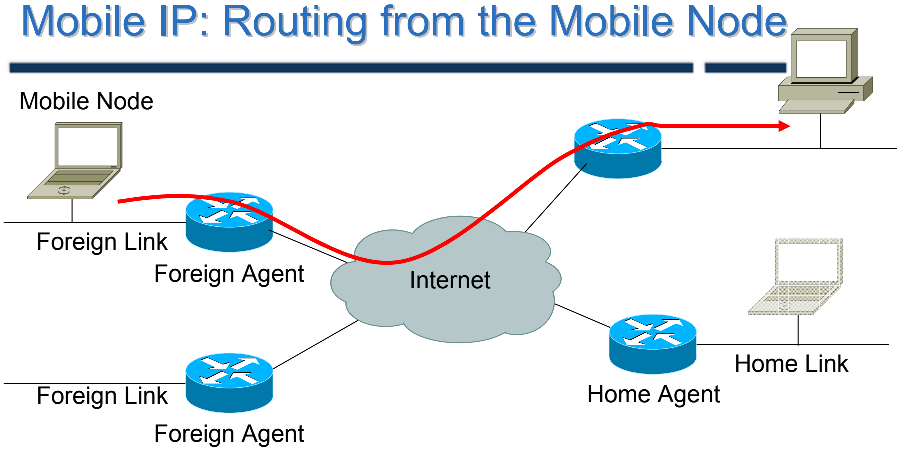

- Packets sent by mobile node are routed directly
- Foreign agent serves as router for visiting mobile node

Mobile IP: Missing Foreign Agent

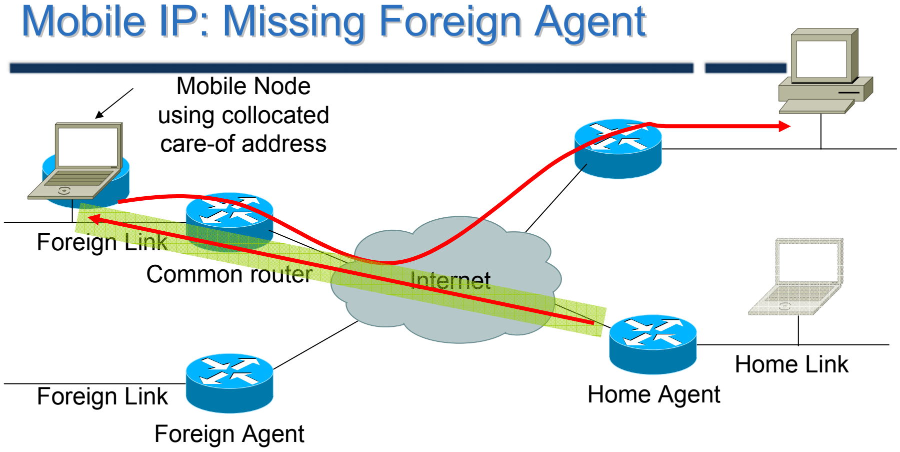

- Implement functionality of foreign agent in mobile node (tunneling)
- Use an acquired IP address (DHCP, PPP, manually) as a collocated care-of-address
  - Address is used only once (opposed to foreign agent care-of-address)
  - Address has the same network prefix

Preventing Mobile Node Denial-of Service

- Bad guy may register a wrong care-of-address
  - All packets get tunneled to the bad guy
  - Bad guy gets to see all packets
  - Mobile node has been cut off from all communications
- Attack is more trivial than in static networks

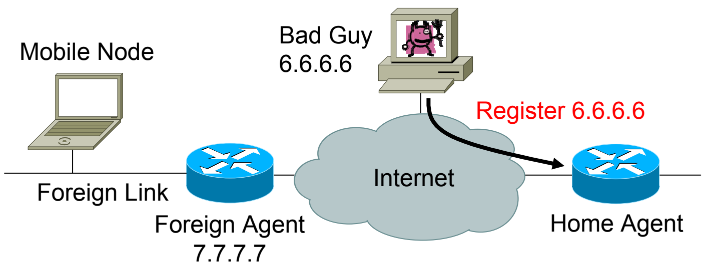

Preventing Mobile Node Denial-of Service

- Mobile IP implementations must support at least “Keyed MD5 authentication” which works as follows

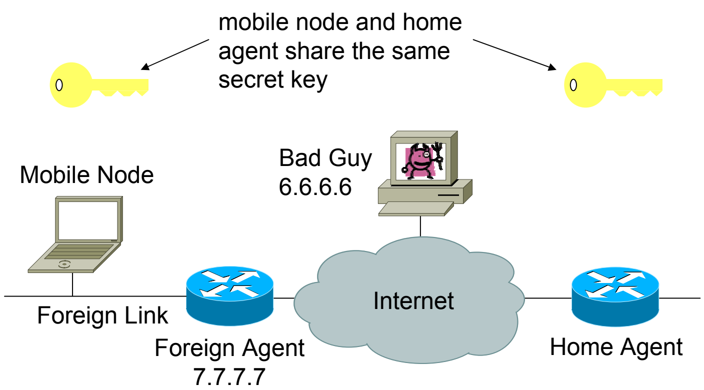

Preventing Mobile Node Denial-of Service

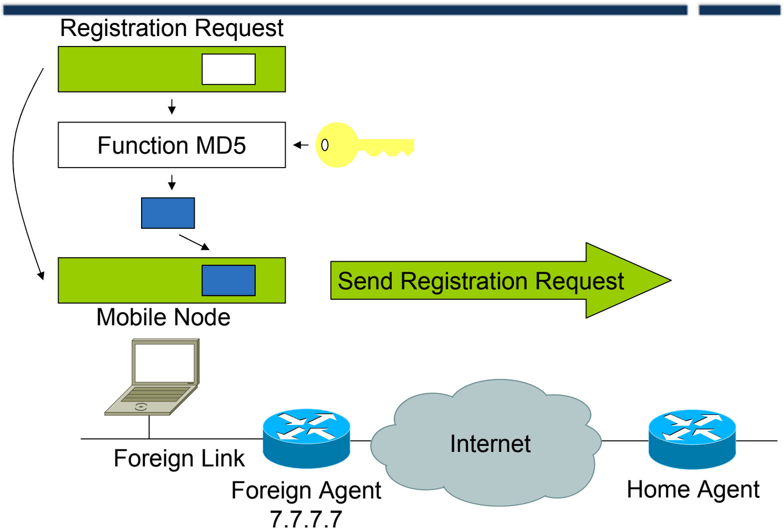

Preventing Mobile Node Denial-of Service

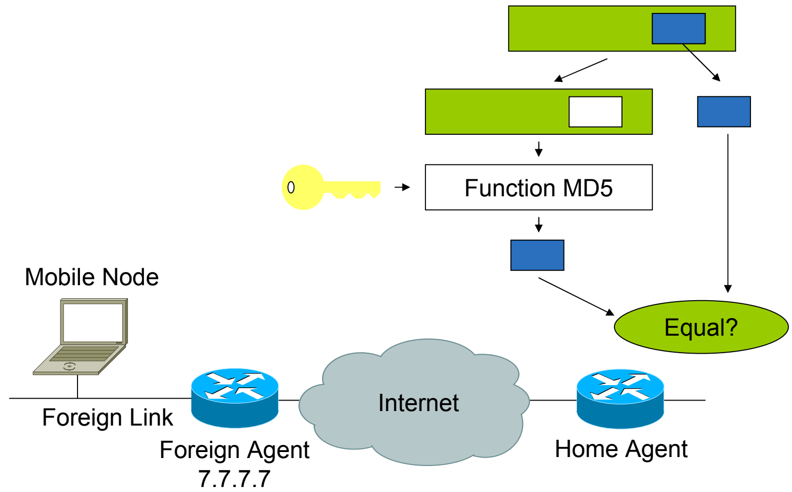

Preventing Replay Attacks

- Bad Guy may store Registration Request for later use

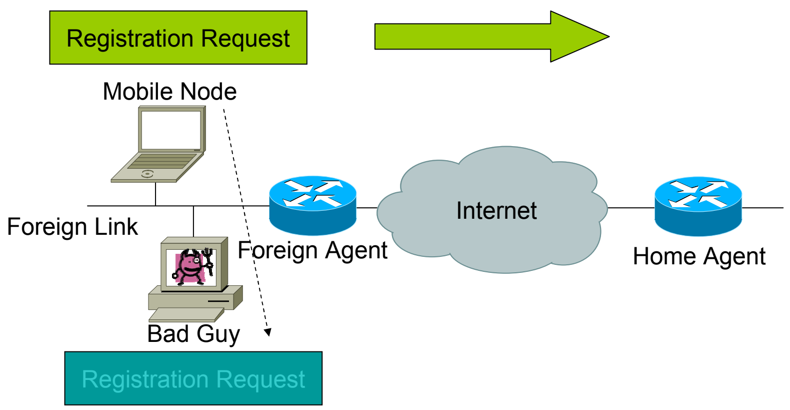

Preventing Replay Attacks

- Bad Guy registers after mobile node disappeared by using the stored registration request

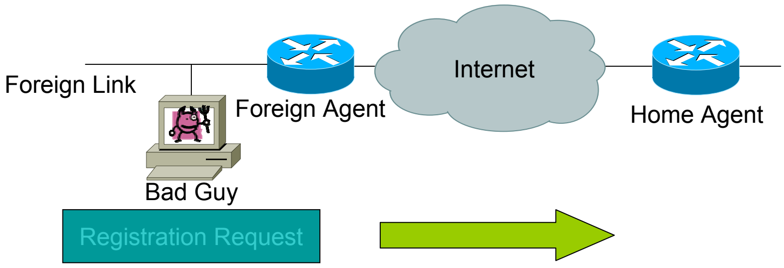

Preventing Replay Attacks

- Generate a unique identification for each successive registration request
- Home agent and mobile node agree on the same rule for successive identifications
  - Time value
  - pseudo-random-numbers (nonces)
- Home agent replies with rejection message when out of sync
  - Mobile node synchronizes its clock
  - Mobile node adjusts the pseudo number generator

Why the Triangle Route?

- Packets addressed to the mobile node are routed via the home agent
- Packets addressed to the correspondent are routed directly

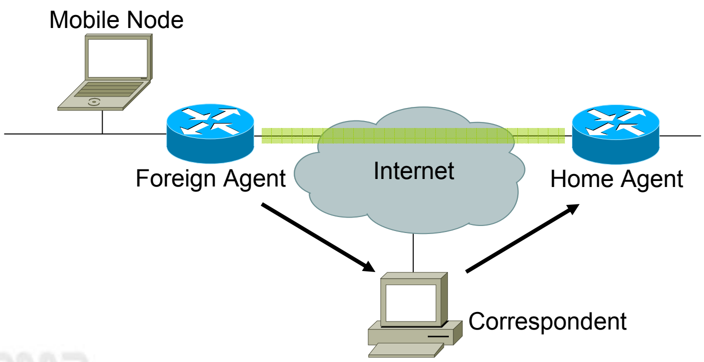

Why the Triangle Route?

- Mobile Node may establish a tunnel to the Correspondent
- This optimizes routing but is this secure?

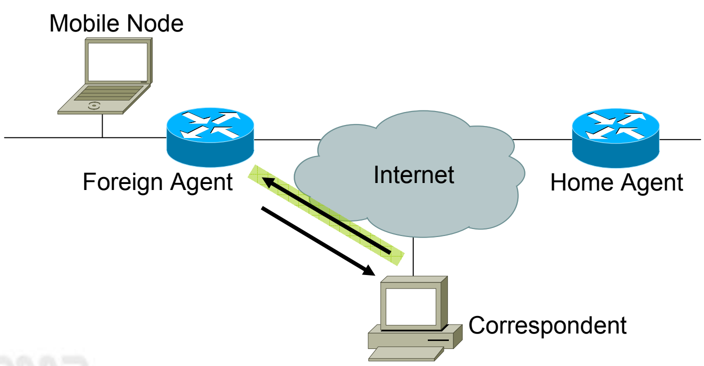

Why the Triangle Route?

- Possible attack: send a wrong care of address
  - Receive all packets addressed to mobile node
  - Cut off mobile node from communication

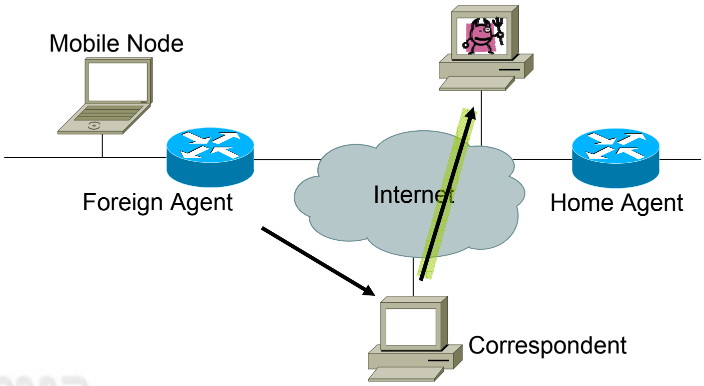

Why the Triangle Route?

- Strong authentication needed
  - Secret between mobile node and every possible correspondent needed
  - Infrastructure needed to establish secret key

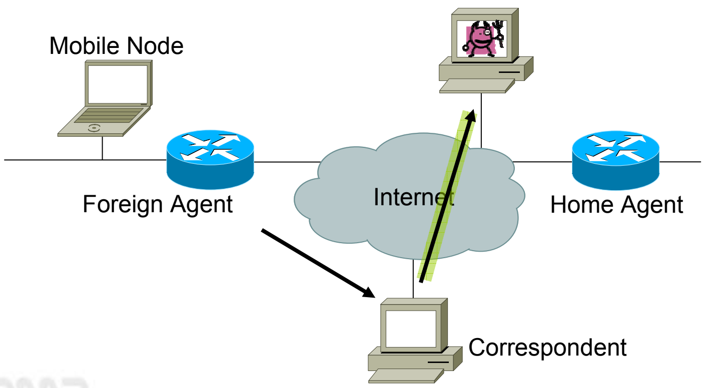

The Impact of Ingress Filtering on Mobile IP (1)

- Attacks often in combination with IP spoofing
  - E.g. Flooding Attack: Bomb the server with TCP connection setup requests
  - IP spoofing possible since unicast routing depends on destination address only
- Ingress Filtering to protect against IP spoofing
  - Router receives packet P from source A at interface X
  - Discard P if packets destined to A would not be sent over X
- How looks a Mobile IP packet from mobile host to a correspondent, when the mobile host is connected to a foreign link?
- Router with ingress filtering may drop Mobile IP packets!

The Impact of Ingress Filtering on Mobile IP (2)

- Foreign Agent may provide reverse tunneling
  - Announced in agent advertisement
  - If needed mobile node selects foreign agent supporting this feature → Announced in registration request
  - Foreign agent might insist on a reverse tunnel (e.g. when service provider is using ingress filtering)

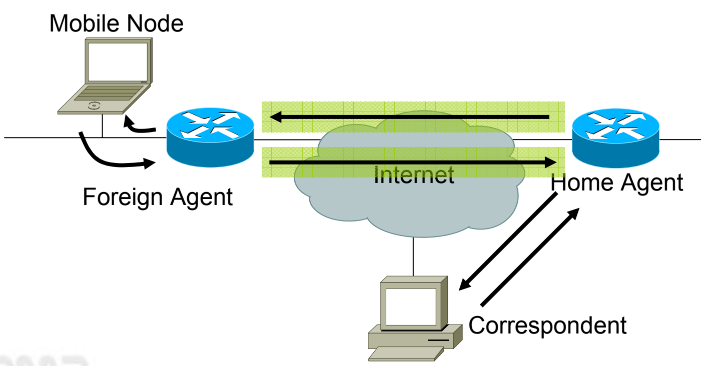

The Impact of Ingress Filtering on Mobile IP (3)

- Different approaches for reverse tunneling
  - Mobile node with collocated address simply tunnels packets to home agent
  - Foreign agent inspects sender address and tunnels if sender registered for reverse tunneling
  - A more flexible solution:

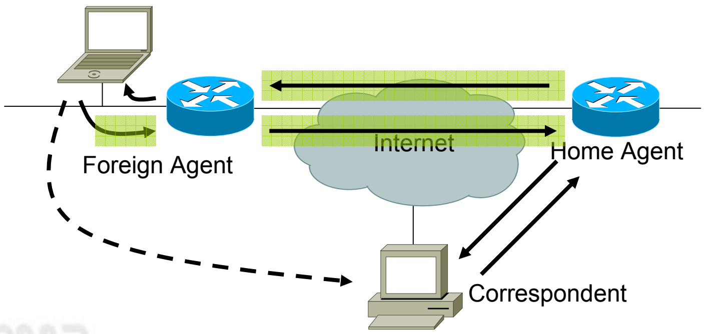

Mobile IP and Broadcasts

- Mobile node may request the home agent to forward all broadcasts (prefix.11…11)
  - Prefix-specific: 123.23.255.255
  - Link-specific: 255.255.255.255
- Receiving with a Collocated Care-of Address
  - Simply forward message to the mobile node
- Receiving with a Foreign Agent Care-of Address
  - What to do with a packet addressed to prefix.11…11?
  - Solution: nested encapsulation
- Sending a Broadcast (three cases arise)
  - Link-specific for foreign link: use link-layer broadcast address
  - Link-specific for home link: tunneling
  - prefix-specific: direct or tunneling (use the latter when intermediate routers filter prefix-specific broadcast messages)

Mobile IP and Multicasts

- Multicast address 1110.multicast-group
- To get a group member register at multicast router (IGMP Host Membership Reports)
- Multicast Routing: IP source address must be topologically correct
- Sending on a foreign link
  - Tunneling to home agent → must be a multicast router
  - Use the collocated-care-of address if available → needs multicast router on the foreign link
- Receiving on a foreign link
  - Tunnel IGMP Host Membership Reports to home agent → must be a multicast router
  - Send IGMP Host Membership Reports on the foreign link → needs multicast router on the foreign link
  - Pro and cons of both methods: message efficiency, reregistering

UBICOMP

TCP and Mobility

TCP High-Level Overview

- TCP characteristics
  - Designed for wired networks with stationary hosts
  - Reliable transport-layer protocol for the Internet
  - Deliver data sequentially and error-free
  - Full-duplex, stream oriented communication
  - Connection-oriented
- TCP is a transport-layer protocol
  - Accept data from application layer
  - Chop data in optimally sized chunks
  - Prepend TCP header
  - Send each segment in the payload of an IP packet

TCP Error Detection and Correction

- Error Detection
  - Sender: store checksum in field of the segment
  - Receiver: compare checksum with received one and discard package, if not identical
- Error Correction
  - Retransmission on missing acknowledge
    - Data or acknowledge possibly lost
  - TCP Performance depends on estimated round-trip time
    - Premature retransmission wastes network bandwidth
    - Overdue retransmission slows down data rate
  - Full-duplex property used for piggybacking acknowledgments
    - Possibly delay transmission of acknowledgment

TCP Flow Control

- The need to control the number of segments in transit
  - Waiting for acknowledgement before sending the next segment?
  - Send all segments subsequently? → Flow Control Needed!
- Sliding Window Protocol
  - Node may send a maximum number before it has to stop
  - Maximum number is defined by a *window*
  - On receipt of an acknowledge
    - Window may slide forward
    - Window size is possibly increased/decreased
    - Node may send more segments when window slides or is increased
  - Use Timeout to retransmit segments not acknowledged so far

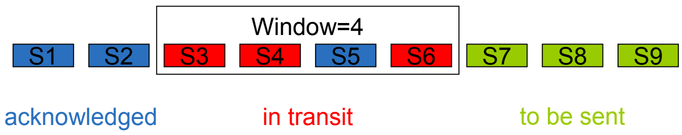

Congestion Control Scheme in Regular TCP (1)

- Slow-Start Algorithm
  - Start with cwnd=1
  - Exponentially increase cwnd for each received ACK until cwnd==ssthresh
  - If cwnd>ssthresh increase cwnd linearly for each received ACK until cwnd==receivers maximum advertised window size
- Congestion Avoidance
  - Continuously measure average of round trip delay d
    - d = (Lambda * d) + ((1-Lambda) * d_current); 0 <= Lambda <= 1
      - Lambda >> 1: Makes weighted average immune to changes lasting a short time.
      - Lambda >> 0: Makes weighted average respond quickly to delay changes.
  - Retransmit segment if current delay > 4 * d

Congestion Control Scheme in Regular TCP (2)

- Fast Retransmit Algorithm
  - On receipt of a segment acknowledge highest in-order sequence number
  - Receipt of out-of-order packets result in duplicate acknowledges
  - On receipt of three identical acknowledges immediately retransmit the segment with the least sequence number not acknowledged so far
- Example of Fast Retransmit Algorithm
  - S1, …, S9 received and acknowledged so far
  - S10 is lost during transmission
  - Receipt of S11, S12, and S13 results in three identical acknowledges of S1,…,S9
  - Sender concludes S10 is missing and retransmits S10 immediately

The Impact of Packet Loss on TCP

- Drop ssthresh into max{2, current window size / 2}
- Reset cwnd to 1→ activate Slow Start Algorithm
- Invoke Timer Backoff Strategy
  - use timers normally
  - increase time-out value if the timer expires
    - New_time-out = Alpha * time-out; where Alpha >= 2
  - Timer backoff strategy is upper-bounded
- Exception: Fast Retransmit

The Fundamental Assumption for TCP

> The majority of lost segments and acknowledgements are due to congestion

- Quite valid for most operational parts of the Internet
- Assumption prevented the Internet from melting down over the past few years!
- However assumption false for mobile networking
  - Error prone media
  - Handoff
- Mobile Computing imposes new challenges to the Internet
  - High bit error rates, Disconnections, limited and variable Bandwidth, cell size, power scarcity, dynamic network topology

The Impact of Bit Error Rate

- Difference between Congestion and BER
  - Error bursts despite sufficient communication resources
- TCP mistakenly assumes congestion
  - Drop ssthresh, start slow start, invoke timer backup strategy
  - Degrades performance of TCP significantly
- Simulation Experiment by Hala ElAarag [2]

	BER = 10^{-5}	BER = 10^{-6}
Throughput (pkts/sec)	39.439	87.455
Success Probability	0.9892	0.999
Transfer time of 5000 pkts. in secs.	123.847	58.032

- Wireless links can reach Bit Error Rate (BER) of 10^{-5}

The Impact of Disconnection

- 0.0: slow start
- 3.0: disconnect
- 3.3: connect
- 4.1: retransmit (slow start)
- 6.3: disconnect
- …

- No duplicate ACK
- No fast retransmit
- Slow start
- Serial timeout condition

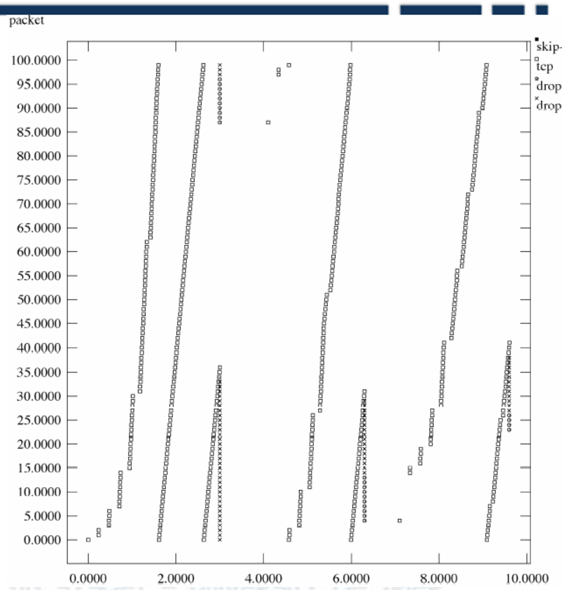

Simulation Result from [2]

Link-Layer Enhancements

- Try to increase quality of lossy wireless links
  - Forward Error Correction (FEC)
    - AIRMAIL
  - Automatic Repeat Request (ARQ)
    - AIRMAIL, RLP, SNOOP
- Problem: retransmission at the link layer
  - TCP retransmits segments after timeout interval
  - Link-layer has to process two copies of the same segment
- Solution: Enable link-layer to inspect higher-layer protocol headers → Discard duplicates
  - Higher-layer encryption?

TCP Enhancements (1)

- Fast retransmit for mobile-node movement
  - Avoid performance decrease during handoff
  - Mobile node knows when it changed its base station
    - Go into fast retransmit combined with slow start
      - Immediately retransmit unacknowledged segments without going into congestion mode
    - Force correspondents to enter fast retransmit
      - Send three duplicate acknowledgements
  - Note: mobility may require a more tight coupling of protocol layers

TCP Enhancements (2)

- Connection segmentation
  - Isolate mobility related problems form existing network protocols
  - Split connection between MH and FH in two separate connections
  - Intermediate node concatenates both connections
  - Conventional TCP need not be changed in any way
  - Specialized Protocols dealing best with mobility
  - Example protocols: MTC, I-TCP, M-TCP, WAP

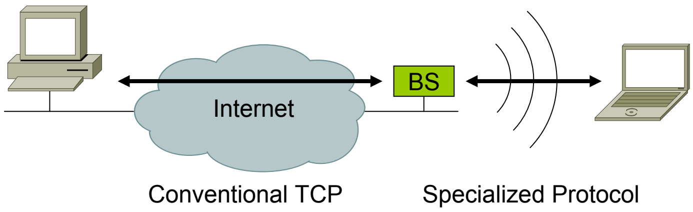

TCP Enhancements (3)

- Transmission and timeout freezing
  - Solve long, though temporary, missing data transmissions
    - E.g. mobile node in a tunnel, preempted data traffic in a mobile phone cell
  - Link-layer may possibly inform TCP of temporary disruption
  - TCP freezes any operation
    - Outstanding packets, measurements, congestion control
  - Signal TCP when media is available
  - Advantage over link-layer solution: unaffected by encryption
    - Signals from link layer based on media not TCP data content
  - Note again: mobility leads to a tighter coupling of protocol layers

Enhancements not Directly Related to Mobility

- TCP Enhancements for slow links
  - Selective Acknowledgement [RFC 2018]
    - Inform sender about all received segments
    - Prevent unnecessary retransmission when a hole appears in the data
  - Transaction TCP [RFC 1644]
    - TCP consists of three phases
      - connection establishment, data transfer, connection close
    - Reduce the amount of message exchange by combining the three phases for small application exchanges
- Application-Layer Enhancements
  - Intelligent caching, distribution of information
  - Improve performance on application layer protocols (e.g. HTTP)
    - HTTP 1.0: separate TCP connection for each web element → slow start for each connection
    - HTTP 1.1: single TCP connection → enough data to get fast beyond slow start and finally transfer with optimal transmission speed

UBICOMP

Conclusion

Conclusion

- Mobility solution for the Internet
  - Architecture: single-hop networks
  - Cope with legacy protocols and applications
- Adapting a single protocol layer does not suffice
- Paradigm shift: Transparency vs. Efficiency
  - Communication among layers sometimes needed
- Mobile IP
  - Should not affect existing static hosts and routers
  - Security concern → new threats possible
- TCP
  - TCP designed to avoid congestion control
  - Gets suboptimal for mobility and wireless networking