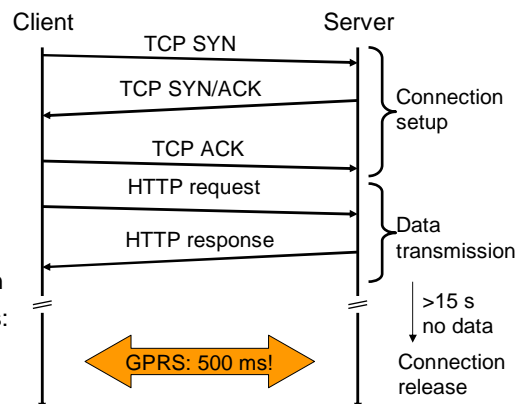


9. Mobile Transport Layer

Motivation
TCP Mechanisms
Classical Approaches
Optimizations
TCP for 2.5/3G Wireless

Transport Layer

- *E.g.*, HTTP (used by web services) typically uses TCP:
 - Reliable transport between client and server required
- TCP:
 - Stream-oriented, not transaction-oriented
 - Network friendly: time-out
 - Congestion
 - Slow down transmission
- Well known – TCP guesses quite often wrong in wireless and mobile networks:
 - Packet loss due to transmission errors
 - Packet loss due to change of network
- Result:
 - Severe performance degradation



Motivation (1)


- Transport protocols typically designed for:
 - Fixed end-systems
 - Fixed, wired networks
- Research activities:
 - Performance
 - Congestion control
 - Efficient retransmissions
- TCP congestion control:
 - Packet loss in fixed networks typically due to (temporary) overload situations
 - Router have to discard packets as soon as the buffers are full
 - TCP recognizes congestion only indirect via missing acknowledgements, retransmissions unwise, they would only contribute to the congestion and make it even worse
 - Slow-start algorithm as reaction

Influences of Mobility on TCP Mechanisms

- TCP assumes congestion, if packets are dropped:
 - Typically wrong in wireless networks, here we often have packet loss due to *transmission errors*
 - Furthermore, *mobility* itself can cause packet loss, if, e.g., a mobile node roams from one access point (e.g., foreign agent in Mobile IP) to another, while there are still packets in transit to the wrong access point and forwarding is not possible
- The performance of an unchanged TCP degrades severely:
 - However, TCP cannot be changed fundamentally due to the large base of installation in the fixed network, TCP for mobility has to remain compatible
 - The basic TCP mechanisms keep the whole Internet together

Original TCP Congestion Control (1)

- Congestion control:
 - Addresses the queuing situation in routers
 - Congestion leads to retransmissions in the transport layer
 - This leads to an increase of the congestion
 - Congestion control needs to adapt sending rate(s)
 - Queue management within routers need to drop packets




© 2005 Burkhard Stiller and Jochen Schiller FU Berlin

M9 – 6



Original TCP Congestion Control (2)



- Approaches for TCP:
 - Implicit mechanisms:
 - Increase of sending credit after receipt of an ACK
 - Decrease of sending credit upon lack of an ACK
 - Explicit mechanisms:
 - Explicit Congestion Notification (RFC 2481)
- AQM: Alternate Queue Management

© 2005 Burkhard Stiller and Jochen Schiller FU Berlin

M9 – 7



Window-based Congestion Control (1)

- Window W defines the number of packets allowed to be sent upon receipt of an ACK

- Bursty Window Control:


- ACK step = W

- Advantage:

- Small overhead due to ACKs

- Drawbacks:

- Bursty nature
- Loss of ACKs



© 2005 Burkhard Stiller and Jochen Schiller FU Berlin


M9 - 8



Window-based Congestion Control (2)

- Sliding Window Control:

- ACK step = 1



- Advantages:

- Not bursty
- Loss of ACKs bearable

- Drawbacks:

- Large overhead due to many ACKs
- Interdependency between error control and window size

© 2005 Burkhard Stiller and Jochen Schiller FU Berlin

M9 - 9



Window-based Congestion Control (3)


- Control of transmission volume, bytes or packets

- Effective rate = $W(t) / RTT(t)$;
 - $W(t) = f(p(t))$
 - $RTT(t) = q(t) + 2*d$

b: Bottleneck Bandwidth
 q: Queue Size
 p: Loss Probability
 d: Propagation Delay

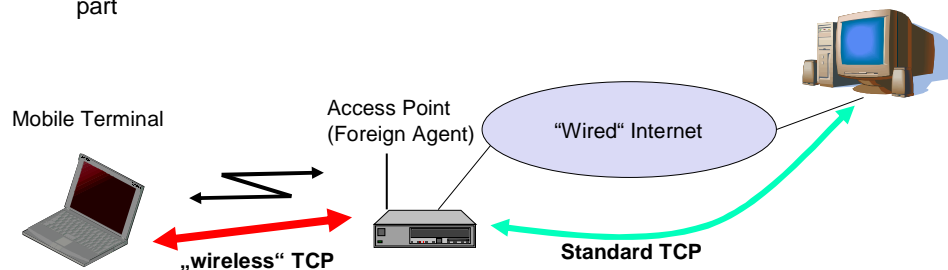
- Optimal Window: $W = b * 2d$
 - Optimal link utilization:
 - Effective rate = $b*(2d/2d) = b$
 - No queue at bottleneck:
 - Optimal network performance (power)

Power = Throughput / Delay:

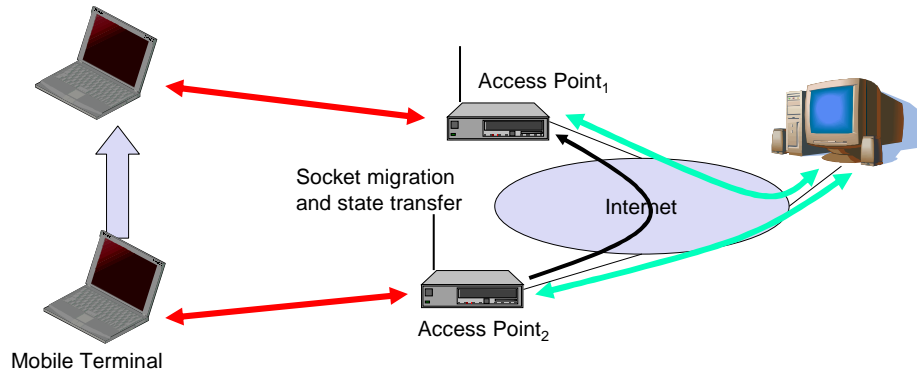


Early Approach: Indirect TCP

- Indirect TCP or I-TCP segments the connection:
 - No changes to the TCP protocol for hosts connected to the wired Internet, millions of computers use (variants of) this protocol
 - Optimized TCP protocol for mobile hosts
 - Splitting of the TCP connection at, e.g., the foreign agent into 2 TCP connections, no real end-to-end connection any longer
 - Hosts in the fixed part of the net do not notice the characteristics of the wireless part



I-TCP Socket and State Migration



© 2005 Burkhard Stiller and Jochen Schiller FU Berlin

M9 – 12



Indirect TCP Discussion

□ Advantages:

- No changes in the fixed network necessary, no changes for the hosts (TCP protocol) necessary, all current optimizations to TCP still work
- Transmission errors on the wireless link do not propagate into the fixed network
- Simple to control, mobile TCP is used only for one hop between, e.g., a foreign agent and mobile host
- Therefore, a very fast retransmission of packets is possible, the short delay on the mobile hop is known

□ Disadvantages:

- Loss of end-to-end semantics, an acknowledgement to a sender does now not any longer mean that a receiver really got a packet, foreign agents might crash
- Higher latency possible due to buffering of data within the foreign agent and forwarding to a new foreign agent


© 2005 Burkhard Stiller and Jochen Schiller FU Berlin

M9 – 13



Early Approach: Snooping TCP

- “Transparent“ extension of TCP within the foreign agent:
 - Buffering of packets sent to the mobile host
 - Lost packets on the wireless link (both directions!) will be retransmitted immediately by the mobile host or foreign agent, respectively (so called “local” retransmission)
 - The foreign agent therefore “snoops” the packet flow and recognizes acknowledgements in both directions, it also filters ACKs
 - Changes of TCP only within the foreign agent



© 2005 Burkhard Stiller and Jochen Schiller FU Berlin

M9 – 14



Snooping TCP Details

- Data transfer to the mobile host:
 - FA buffers data until it receives ACK of the MH, FA detects packet loss via duplicated ACKs or time-out
 - Fast retransmission possible, transparent for the fixed network
- Data transfer from the mobile host:
 - FA detects packet loss on the wireless link via sequence numbers, FA answers directly with a NACK to the MH
 - MH can now retransmit data with only a very short delay
- Integration of the MAC layer:
 - MAC layer often has similar mechanisms to those of TCP
 - Thus, the MAC layer can already detect duplicated packets due to retransmissions and discard them
- Problems:
 - Snooping TCP does not isolate the wireless link as good as I-TCP
 - Snooping might be useless depending on encryption schemes

© 2005 Burkhard Stiller and Jochen Schiller FU Berlin

M9 – 15



Early approach: Mobile TCP

- ❑ Special handling of lengthy and/or frequent disconnections
- ❑ M-TCP splits as I-TCP does:
 - Unmodified TCP fixed network to supervisory host (SH)
 - Optimized TCP SH to MH
- ❑ Supervisory host:
 - No caching, no retransmission
 - Monitors all packets, if disconnection detected (no ACKs for some time)
 - Set sender window size to 0
 - Sender automatically goes into persistent mode (no sending possible)
 - Old or new SH reopen the window at original size
- ❑ Advantages:
 - Maintains semantics, supports disconnection, no buffer forwarding
- ❑ Drawbacks:
 - Loss on wireless link propagated into fixed network
 - Adapted TCP on wireless link

Fast Retransmit and Fast Recovery

- ❑ Change of foreign agent often results in packet loss:
 - TCP reacts with slow-start although there is no congestion
- ❑ Forced fast retransmit:
 - As soon as the mobile host has registered with a new foreign agent, the MH sends duplicated acknowledgements on purpose
 - This forces the fast retransmit mode at the communication partners
 - Additionally, the TCP on the MH is forced to continue sending with the actual window size and not to go into slow-start after registration
- ❑ Advantage:
 - Simple changes result in significant higher performance
- ❑ Drawback:
 - Further mix of IP and TCP, no transparent approach

Transmission and Time-out Freezing

- ❑ Mobile hosts can be disconnected for a longer time:
 - No packet exchange possible, e.g., in a tunnel, disconnection due to overloaded cells or multiplexing with higher priority traffic
 - TCP disconnects after time-out (loss of layer 2 connectivity) completely
- ❑ TCP freezing:
 - MAC layer is often able to detect interruption in advance
 - MAC can inform TCP layer of upcoming loss of connection
 - TCP stops sending, but does not assume a congested link
 - MAC layer signals again if reconnected
- ❑ Advantage:
 - Scheme is independent of data
- ❑ Drawback:
 - TCP on mobile host has to be changed, mechanism depends on MAC layer

Selective Retransmission

- ❑ TCP acknowledgements are often cumulative:
 - ACK n acknowledges correct and in-sequence receipt of packets up to n
 - If single packets are missing quite often a whole packet sequence beginning at the gap has to be retransmitted (go-back-n), thus wasting bandwidth
- ❑ Selective retransmission as one solution:
 - RFC2018 allows for acknowledgements of single packets, not only acknowledgements of in-sequence packet streams without gaps
 - Sender can now retransmit only the missing packets
- ❑ Advantage:
 - Much higher efficiency
- ❑ Drawback:
 - More complex software in a receiver, more buffer needed at the receiver

Transaction-oriented TCP

- TCP phases:
 - Connection setup, data transmission, connection release
 - Using 3-way-handshake needs 3 packets for setup and release, respectively
 - Thus, even short messages need a minimum of 7 packets!
- Transaction-oriented TCP:
 - RFC1644, T-TCP, describes a TCP version to avoid this overhead
 - Connection setup, data transfer and connection release can be combined
 - Compare WAP Transaction Layer Protocol
 - Thus, only 2 or 3 packets are needed
- Advantage:
 - Efficiency
- Drawback:
 - Requires changed TCP
 - Mobility not longer transparent

Comparison of Approaches for “Mobile” TCP

Approach	Mechanism	Advantages	Disadvantages
Indirect TCP	splits TCP connection into two connections	isolation of wireless link, simple	loss of TCP semantics, higher latency at handover
Snooping TCP	“snoops” data and acknowledgements, local retransmission	transparent for end-to-end connection, MAC integration possible	problematic with encryption, bad isolation of wireless link
M-TCP	splits TCP connection, chokes sender via window size	Maintains end-to-end semantics, handles long term and frequent disconnections	Bad isolation of wireless link, processing overhead due to bandwidth management
Fast retransmit/ fast recovery	avoids slow-start after roaming	simple and efficient	mixed layers, not transparent
Transmission/ time-out freezing	freezes TCP state at disconnect, resumes after reconnection	independent of content or encryption, works for longer interrupts	changes in TCP required, MAC dependant
Selective retransmission	retransmit only lost data	very efficient	slightly more complex receiver software, more buffer needed
Transaction oriented TCP	combine connection setup/release and data transmission	Efficient for certain applications	changes in TCP required, not transparent

TCP Improvements (1)

- Initial research work:
 - Indirect TCP, Snoop TCP, M-TCP, T/TCP, SACK, Transmission/time-out freezing
- TCP over 2.5/3G wireless networks
 - Fine tuning today's TCP
 - Learn to live with:
 - Data rates: 64 kbit/s up, 115-384 kbit/s down; asymmetry: 3-6, but also up to 1000 (broadcast systems), periodic allocation/release of channels
 - High latency, high jitter, packet loss
 - Suggestions:
 - Large (initial) sending windows, large maximum transfer unit, selective acknowledgement, explicit congestion notification (ECN), time stamp, no header compression
 - Already in use:
 - i-mode running over FOMA
 - WAP 2.0 ("TCP with wireless profile")

$$BW \leq \frac{0.93 * MSS}{RTT * \sqrt{p}}$$

- max. TCP BandWidth
- Max. Segment Size
- Round Trip Time
- loss probability

TCP Improvements (2)

- Performance enhancing proxies (PEP, RFC 3135):
 - Transport layer:
 - Local retransmissions and acknowledgements
 - Additionally on the application layer:
 - Content filtering, compression, picture downscaling
 - E.g., Internet/WAP gateways
 - Web service gateways?
 - Big problem: breaks end-to-end semantics:
 - Disables use of IP security
 - Choose between PEP and security!
- More open issues:
 - RFC 3150 (slow links)
 - Recommends header compression, no timestamp
 - RFC 3155 (links with errors):
 - States that explicit congestion notification cannot be used
 - In contrast to 2.5G/3G recommendations!

