

TCP over Wireless Networks

Raj Jain
Professor of Computer Science and Engineering
Washington University in Saint Louis
Saint Louis, MO 63130

Audio/Video recordings of this lecture are available at:

<http://www.cse.wustl.edu/~jain/cse574-08/>



- ❑ TCP: Key Features
- ❑ TCP Congestion Mechanisms
- ❑ Our Initial Research on TCP Congestion
- ❑ TCP Over Wireless: Issues and Solutions
- ❑ TCP over Satellite
- ❑ Our research on TCP over Satellite and Wireless

TCP: Key Features

- 1. Stream-Oriented Transmission:** Multiple application packets may be send in one TCP “**Segment.**”
Maximum Segment Size (**MSS**).
All acks are byte numbers. Segment # used in all discussions.
- 2. Reliable Delivery:** Segments are buffered at the source until acked. Retransmitted if not acked.
- 3. In-Order Delivery:** Destination delivers segments to application only when all previous segments received.
- 4. End-to-End Semantics:** Ack \Rightarrow Data received at destination
- 5. Congestion Control:**
Increases load slowly from a low initial start.
Reduces load if network congested (based on segment timeout, duplicate acks)
- 6. Congestion Avoidance:** Explicit Congestion Notification (ECN) bits in TCP/IP headers based on DECbit research

Washington University in St. Louis

CSE574s

©2008 Raj Jain

18-3

TCP Flow Control

- ❑ **Cumulative Acks:** Acks all bytes up to the ack
- ❑ **Piggybacked Acks:** Acks are sent in the reverse packets if possible.
- ❑ **Delayed Acks:** Ack delayed in case another segment is received or segment needs to be sent. Typically 200 ms
- ❑ **Duplicate Acks:** If an out of order packet is received, the previous ack is resent. Duplicate acks are not delayed.
- ❑ **Window Flow Control:**
$$\text{Throughput} = \text{Window} / \text{Round Trip Time}$$
- ❑ Ideal Window Size = Round Trip Time * Link Capacity
= Delay-bandwidth product
- ❑ TCP sets retransmission timer for only one packet. If the ack is not received and the timer expires, the packet is assumed lost.

Washington University in St. Louis

CSE574s

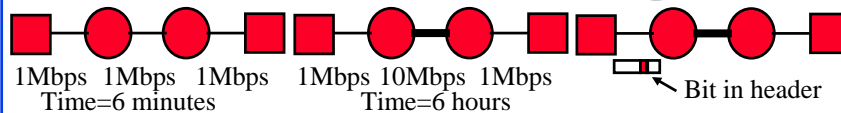
©2008 Raj Jain

18-4

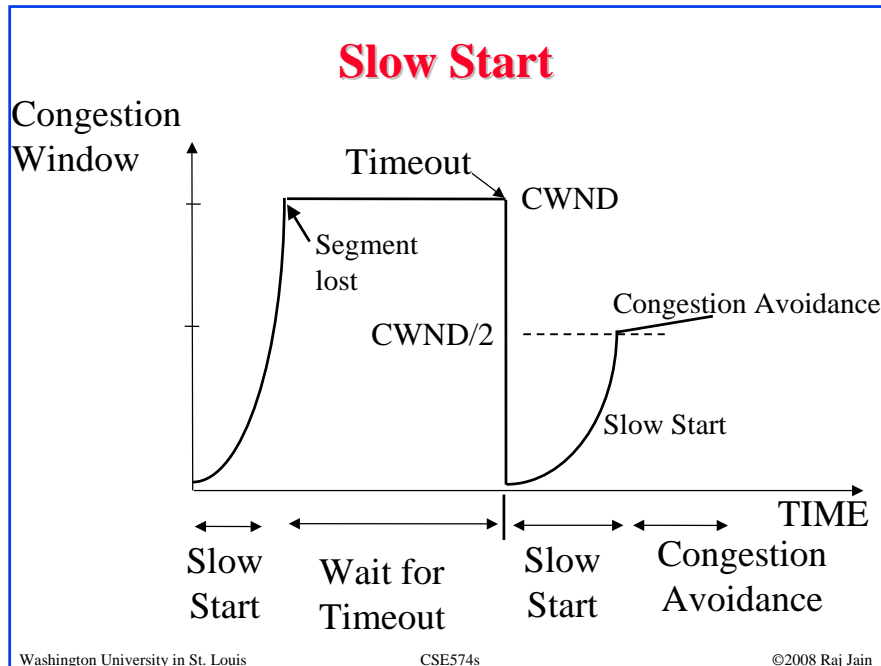
TCP Congestion Mechanisms

- ❑ Slow Start
- ❑ Fast retransmit and recovery
- ❑ New Reno
- ❑ Selective Acknowledgement
- ❑ Explicit Congestion Notification

Our Research on TCP Congestion



- ❑ Early 1980s Digital Equipment Corporation (DEC) introduced Ethernet products
- ❑ Noticed that throughput goes down with a higher-speed link in middle (because no congestion mechanisms in TCP)
- ❑ Results:
 1. Timeout \Rightarrow Congestion
 \Rightarrow Reduce the TCP window to one on a timeout [Jain 1986]
 2. Routers should set a bit when congested (DECbit). [Jain, Ramakrishnan, Chiu 1988]
 3. Introduced the term “Congestion Avoidance”
 4. Additive increase and multiplicative decrease (AIMD principle) [Chiu and Jain 1989]
- ❑ There were presented to IETF in 1986.
 \Rightarrow Slow-start based on Timeout and AIMD [Van Jacobson 1988]



18-7

- ## Slow Start (Cont)
- ❑ Receiver sends “Receive window” (for flow control)
 - ❑ Sender maintains a **Congestion Window** (CWND)
 $CWND \leq \text{Receiver Window}$
 - ❑ Set “Slow Start Threshold”
 $SSThresh = 64 \text{ kB}$ initially
 - ❑ Start with a CWND W of 1
 - ❑ Increase W by 1 after every ack until $SSThresh$
 \Rightarrow Exponential increase (**Slow Start**. W doubles every RTT)
 - ❑ Increase W by $1/W$ after every Acks (W increases by per RTT)
 \Rightarrow Linear increase (**Congestion Avoidance**)
 - ❑ On a timeout, Set $SSThresh$ to half the current window and set window to 1.
 $SSThresh \leftarrow \text{Max}\{2, 0.5W\}, W \leftarrow 1$
- Washington University in St. Louis CSE574s ©2008 Raj Jain

18-8

Fast Retransmit and Recovery (FRR)

- ❑ Also known as TCP-Reno
- ❑ Ideas:
 - Don't have to wait for timeout on a loss
 - Don't reduce to one on single loss
 - Duplicate acks \Rightarrow Loss
- ❑ On three duplicate acks:
 - Retransmit the lost segment (Fast Retransmit)
 - Set SS_{Thresh} to $\text{Max}\{2, 0.5 \times CWND\}$
 - Reduce $CWND$ to $0.5 \times CWND + \# \text{ of dupacks}$
 - New ack $\Rightarrow CWND > SS_{Thresh} \Rightarrow$ Linear increase
 - Duplicate ack \Rightarrow inflate $CWND$ by 1. Send a pkt if allowed
- ❑ **Advantage:** Recovers from loss without a timeout
- ❑ **Problem:** Cannot recover from bursty (3+) losses. Dupacks are also generated if pkts out-of-order (no loss).

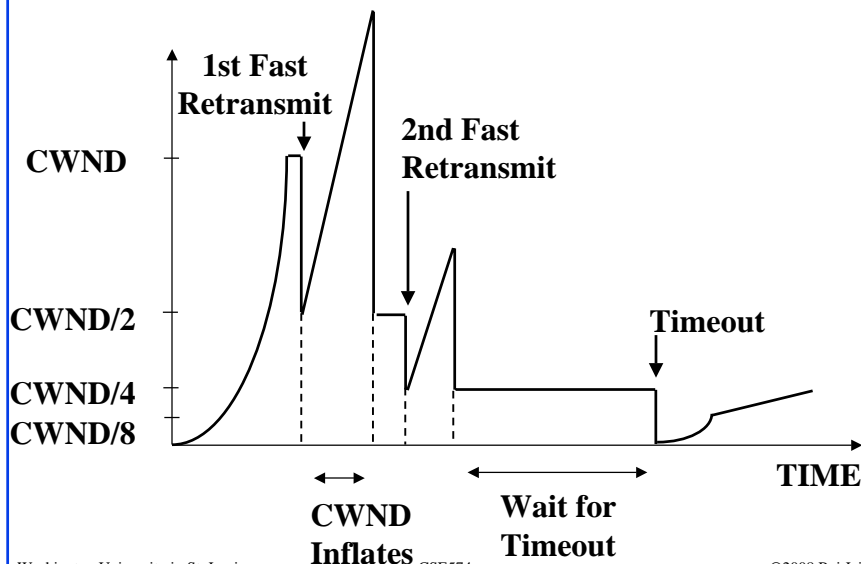
Washington University in St. Louis

CSE574s

©2008 Raj Jain

18-9

FRR (Cont)



Washington University in St. Louis

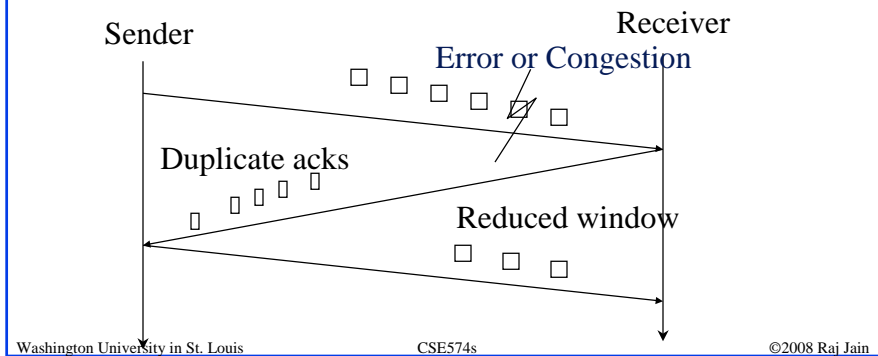
CSE574s

©2008 Raj Jain

18-10

Problems of Current TCP

- ❑ TCP cannot distinguish wireless errors from congestion.
- ❑ Frequent errors \Rightarrow Frequent window reductions
 \Rightarrow Low throughput
- ❑ On CDMA, Overload \Rightarrow Errors. Otherwise no relationship.



18-11

TCP Over Wireless

- ❑ Link Layer Mechanisms
- ❑ Split TCP Solutions
- ❑ TCP Aware Link Layer Protocols
- ❑ Explicit Notification Schemes
- ❑ TCP Over Satellite
- ❑ Our Results for Satellite and Wireless Networks


Washington University in St. Louis

CSE574s

©2008 Raj Jain

18-12

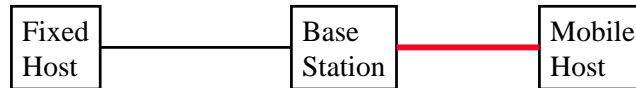
Link Layer Mechanisms

- ❑ Forward Error Correction (FEC):
 - Reduces loss due to errors.
 - Reduced link throughput even if no errors.
- ❑ Automatic Repeat Request (ARQ):
 - Link layer retransmission and acknowledgement
 - No reduction in throughput if no errors
 - Reduced throughput and increased delay at link layer
 - May cause congestion
 - May increase variance of RTT \Rightarrow Increased RTO
 - May cause head-of-line blocking 
- ❑ Adaptive Link layer strategies:
 - Dynamically vary FEC code, retransmission limit, frame size

Split TCP Solutions

- ❑ Indirect TCP
- ❑ Selective Repeat Protocol (SRP)
- ❑ Mobile TCP
- ❑ Mobile-End Transport Protocol

Indirect TCP



- ❑ Two TCP connections:
 - Fixed host to Base
 - Base to Mobile
- ❑ Independent flow control on two connections
- ❑ Packets buffered in the base
- ❑ Ack at sender \neq MH has received
 - Violates TCP's end-to-end semantics
 - BS retains hard state. BS failure \Rightarrow loss of data
 - On handoff, stored packets must be sent to new BS
 - Does not work if connection not bi-directional. E.g., satellites
- ❑ Reference: Bakre95, Bakre 97

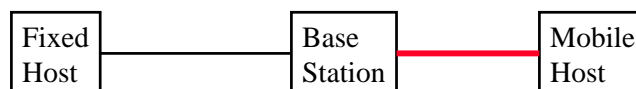
Washington University in St. Louis

CSE574s

©2008 Raj Jain

18-15

Selective Repeat Protocol (SRP)



- ❑ Two connections: Similar to Indirect TCP
- ❑ FH to BS: Standard TCP
- ❑ BS to MH: Selective repeat protocol on UDP
- ❑ Reference: Yavatkar94

Washington University in St. Louis

CSE574s

©2008 Raj Jain

18-16

Mobile TCP

- ❑ Asymmetric split connection
- ❑ Simpler protocol at mobile host
- ❑ Mobile does error detection only
- ❑ Base does error correction and error detection
- ❑ Header compression on wireless hop
- ❑ On/off flow control on wireless hop
- ❑ Ref: Haas97

Mobile-End Transport Protocol

- ❑ TCP runs only between fixed host and BS
- ❑ BS guarantees reliable ordered delivery to mobile
- ❑ Ref: Wang98

TCP Aware Link Layer Protocols

- ❑ Snoop Protocol
- ❑ WTCP
- ❑ Delayed DupAcks Protocol
- ❑ SCPS-TP

Snoop Protocol

- ❑ Split connection and link level retransmission
- ❑ Base monitors returning acks. Retransmits on duplicate acks and drops the duplicate ack
- ❑ Advantages: Only soft state at BS. Only BS modified. No changes to FH or MH.
- ❑ If wireless link delay is less than 4 packets, 3 duplicate acks will not happen and a simple link-level retransmission without dropping duplicate ack will also work.
- ❑ Disadvantages: Does not work with encrypted packets
- ❑ Does not work on asymmetric paths
- ❑ Ref: Balakrishnan95

WTCP

- ❑ Similar to Snoop
- ❑ Snoop can cause increased RTT
- ❑ WTCP corrects RTT by modifying the timestamp in returning acks
- ❑ Disadvantages:
 - Useful only if retransmission times are large (>1 tick)
 - Does not work on shared LANs, where overload => Increased delay
- ❑ Ref: Ratnam98

Delayed DupAcks Protocol

- ❑ Similar to Snoop. But no duplicate ack dropping at BS
- ❑ Link layer retransmission on wireless hop
- ❑ Third duplicate acks delayed at MH. BS does not need to look into TCP headers.
- ❑ Out-of-order packet delivery from link layer to TCP allowed at MH to avoid head-of-line blocking at MH
- ❑ Advantage: BS is not TCP aware. Can be used even if headers are encrypted.
- ❑ Disadvantages: Congestion losses are recovered later since dupacks delayed.
- ❑ Ref: Mehta98, Vaidya99

SCPS-TP

The diagram illustrates the SCPS-TP (Space Communications Transport Protocol) architecture. It shows a satellite in space (left) connected via a dashed line to a ground station on Earth (center). The ground station is connected via solid lines to two destinations, D1 and D2 (right).

- ❑ Space Communications Transport Protocol (SCPS-TP)
- ❑ Used in satellite communications
- ❑ Ground stations monitors packets with failed checksum and sends corruption experienced messages to destinations of recent error-free packets
- ❑ Ground stations can detect outage of incoming link and assume outage of outgoing link also.
- ❑ Destinations ack with “Corruption experienced” bit
- ❑ After receiving an ack with “Corruption Experienced” bit, sender does not back off on timeout or duplicate acks until it receives an ack without that bit.

Washington University in St. Louis CSE574s ©2008 Raj Jain

18-23

Explicit Notification Schemes

- ❑ Explicit Loss Notification
- ❑ Explicit Loss Notification 2
- ❑ Explicit Bad State Notification
- ❑ Partial Ack Protocols

Washington University in St. Louis CSE574s ©2008 Raj Jain

18-24

Explicit Loss Notification



- ❑ Works with Mobile host sources
First link on the path is wireless
- ❑ BS keeps track of missing packets from mobile
- ❑ When DupAcks is received, BS sets “ELN” bit in the DupAcks
- ❑ When mobile receives the DupAcks with ELN bit, it does not back off. Simply retransmits.

- ❑ Reference: Balakrishnan98

Explicit Loss Notification 2

- ❑ Similar to ELN. Works when receiver is mobile.
- ❑ Caches TCP sequence numbers at the base (as in snoop) but does not cache data packets
- ❑ Dupacks are tagged with ELN bit if sequence number of lost packet is cached at the base
- ❑ Ref: Biaz99

Explicit Bad State Notification

- ❑ Works when Mobile is the receiver
- ❑ Link layer retransmission on the wireless link
- ❑ If base cannot deliver the packet to Mobile, it sends a “explicit bad state notification” (EBSN) message to sender
- ❑ Ref: Bakshi97

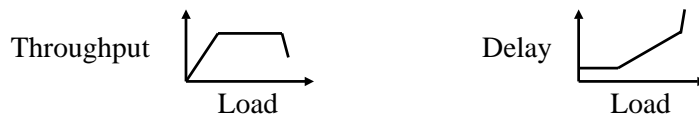
Partial Ack Protocols

- ❑ Two types of Acks:
 - Normal TCP acks
 - Partial acks: Informing source that packet was received by intermediate host, e.g., base station
- ❑ If partial ack is received but the normal ack is not received or DupAcks received, the sender does not back off, simply retransmits
- ❑ Ref: Cobb95, Biaz97

Receiver-Based Scheme

- ❑ Receiver uses inter-arrival time between packets to *guess* the cause of packet loss
- ❑ If the loss is guessed to be due to error, sender is informed via ELN bit in duplicate ack or explicit message
- ❑ Advantages:
 - Can be implemented without base modification
 - Works with encrypted packets
- ❑ Disadvantage:
 - Works only if the wireless link is the slowest link
Ensure that there is some queuing at the base
 - Queuing delays at the base for all packets should be similar
- ❑ Ref: Biaz98

Sender-Based Discrimination Scheme



- ❑ Sender uses roundtrip times, window sizes, and loss pattern to guess the cause of packet loss
- ❑ If loss is guessed to be due to errors, the sender does not back off. Simply retransmits.
- ❑ Heuristics:
 - Delay gradient $dD/dW > 0 \Rightarrow$ Congestion
 - Throughput gradient $dT/dW < 0.5 \Rightarrow$ Congestion
 - $W/RTT_{min} - W/RTT_{actual} > b \Rightarrow$ Congestion
- ❑ Disadvantages: Does not work in practice.
Delay and throughput measurements are quite noisy.
- ❑ Ref: Biaz98b, Biaz99b

TCP Over Satellite

- ❑ IETF TCPSAT
- ❑ Satellite Transport Protocol (STP)
- ❑ Early Acks: ACKprime
- ❑ Our results

IETF TCPSAT

- ❑ Large propagation delays => Large bandwidth delay product
=> Large windows => Use window scaling option
$$\text{Window} = \text{window} * 2^{\text{Scaling factor}}$$
- ❑ Use Selective acknowledgements
=> Allows multiple packets to be recovered in one RTT
- ❑ Do not delay acks => Ack every packet
- ❑ Use larger initial window size (suggested 4kB)
- ❑ Byte Counting: Increase window by number of bytes acked rather than just 1 MSS per RTT
- ❑ Reduce bursts from the sender
- ❑ Ref: RFC 2488, 2760

Satellite Transport Protocol (STP)

- ❑ Non-TCP transport protocol
- ❑ No retransmission timer
- ❑ Sender periodically requests receiver to ack received packets => Save reverse bandwidth if no errors
- ❑ Receiver can also send “Selective Nacks” if packets lost
- ❑ Ref: Henderson98

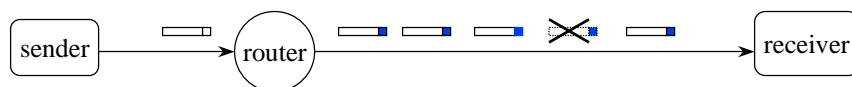
Early Acks: ACKprime

- ❑ Ground stations send partial acks => Grows congestion window
- ❑ Full acks from the receiver required for reliable delivery
- ❑ Ref: Scott98

Our Results for Satellite Networks

- ❑ End System Improvements:
 - Slow start
 - Fast Retransmit and Recovery
 - New Reno
 - SACK
- ❑ Intermediate System Improvements: Drop policies
- ❑ For satellite paths, end system improvements have more impact than intermediate-system based improvements
- ❑ SACK helps significantly
- ❑ Fairness depends upon the drop policies in the intermediate systems and not on end system policies

Wireless Networks: Our Solution



Desired Attributes of the Solution:

1. Must maintain TCP's end-to-end semantics: A packet is acked only after received by the final destination.
2. Modifications must be local: Only Base Station (BS) and Mobile Host (MH) are in the control of wireless service provider. Cannot change all locations that MH visits.
3. Must apply to two-way traffic: MH can be both a sender and a receiver.
4. Wireless links can be at the end or in the middle (satellite links)

Ref: Liu and Jain 2003

Survey of Prior Proposals

	I-TCP	Multiple Acks	Control Connection	Snoop	ELN	Delayed Dupacks
Semantic	No	Yes	Yes	Yes	Yes	Yes
Local	Yes	No	No	Yes	Yes	Yes
Encryption	No	No	Yes	No	No	Yes
Two-way	Yes	No	No	No	No	No
Intermediate links	Yes	No	Yes	No	No	Yes

Washington University in St. Louis

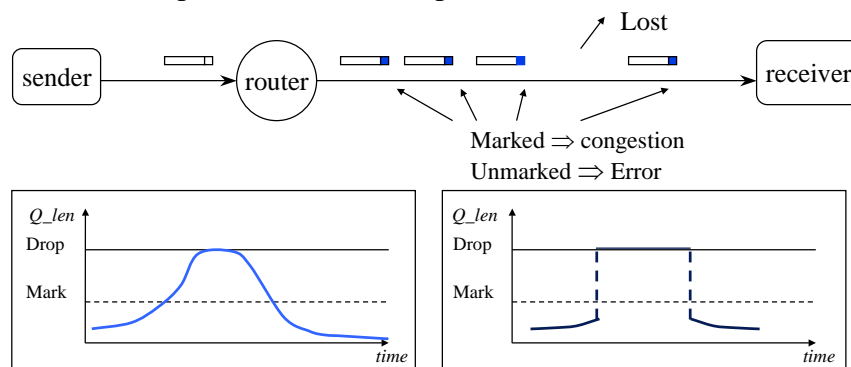
CSE574s

©2008 Raj Jain

18-37

Congestion Coherence

- ❑ Congestion does not happens nor disappear suddenly:
 - Before congestion reaches the point where a packet has to be dropped, some packets must have been marked.
 - After a packet is lost, some packets will be marked.



Washington University in St. Louis

CSE574s

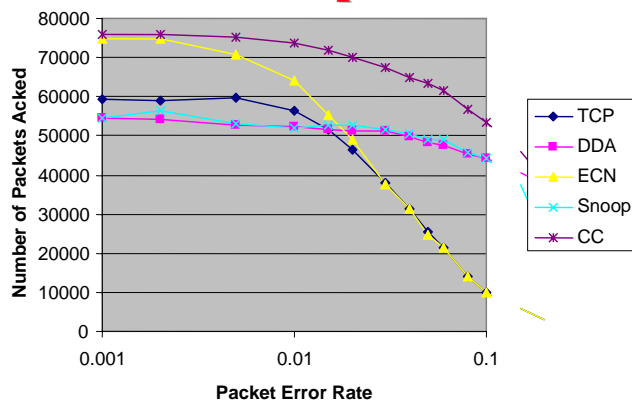
©2008 Raj Jain

18-38

Congestion Coherence Algorithm


- ❑ Link layer acks and retransmissions at all wireless nodes.
- ❑ **Receiver:**
 - Out-of-order packets received check ECN bits.
 - If any packet marked, send duplicate acks. Otherwise, defer the duplicate acks.
 - If expected packet arrives, drop deferred dupacks.
 - If the packet times out, release all deferred dupacks.
- ❑ **Sender:**
 - When the third duplicate acks arrives, MH checks the ECN-ECHO bits.
 - If any of these duplicate acks carry an ECN-ECHO, MH retransmits the lost packet and reduces the window. Otherwise, TCP defers the retransmission.
 - When the expected ack arrives, cancel the deferred retransmission.
 - If the expected ack does not arrive in certain period of time then MH starts the deferred retransmission.

Goodput



- ❑ Congestion Coherence provides the highest throughput

Summary



- ❑ Frequent errors on wireless links trigger the congestion mechanism in TCP resulting in low throughput
- ❑ Key mechanisms are link level schemes to reduce/hide error losses, split TCP, TCP modification in base, receiver, or sender
- ❑ Since congestion builds up slowly, coherence of ECN bits provides a good distinction of congestion vs. errors
- ❑ On satellite links, window scaling, large initial windows, and SACK are helpful

Reading Assignment

- ❑ Read sections 9.1 through 9.6 of Murthy and Manoj
- ❑ See also references at the end of this presentation

Homework 18

- ❑ **Exercise:** A TCP entity opens a connection and uses slow start. Approximately how many round-trip times are required before TCP can send N segments.
- ❑ **Hint:** Write down what the CWND and total segments will be after 1 round trips, 2 round trips, 3 round trips, ...

References

- ❑ C. Liu and R. Jain, "Approaches of Wireless TCP Enhancement and A New Proposal Based on Congestion Coherence", the 36th Hawaii International Conference on System Sciences, Quality of Service in Mobile and Wireless Network minitrack, Big Island, Hawaii, January 5-9, 2003, pp. 307a, <http://www.cse.wustl.edu/~jain/papers/hicss.htm>
This paper also has references to several other papers on wireless.
- ❑ Nitin Vaidya, "TCP for Mobile and Wireless Hosts," a comprehensive tutorial presentation, 291 pp., <http://research.microsoft.com/meshsummit/Presentations/DayOne/session1/nitin vaidya.ppt>
- ❑ Annika Wennström, Stefan Alfredsson, and Anna Brunstrom, "TCP over Wireless Networks," <http://www.cs.kau.se/~alfs/licmaterial/tcpreport.pdf>