A Transport Layer Approach to Host Mobility

Final Exam

PhD Candidate:
Luiz Magalhães

Committee:
Robin Kravets
Roy Campbell
Klara Nahrstedt
Benjamin Wah

Motivation: Wireless

- Wireless links allows host to roam but
  - No single wireless solution will dominate
    - tradeoffs on range, BW and number of hosts
  - Traditional IP routing does not allow for mobility
  - Changes to IP have negative impact on TCP
  - Wireless links have different characteristics from wired links

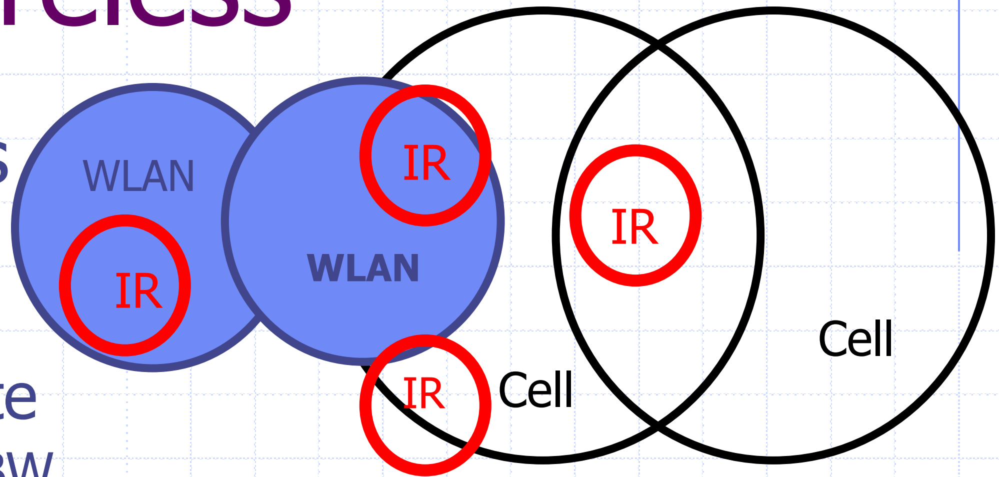

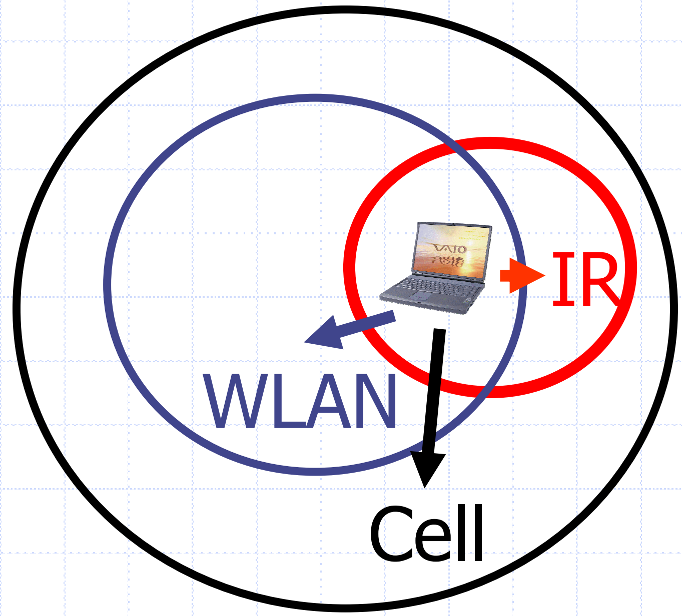

Current Mobility Solutions

- Link layer:
  - Cell phones
  - Wireless Ethernet
- Network layer:
  - Mobile IP
- Limitations:
  - Neither supports multihomed devices
  - Localized decisions

Transport Layer Approach

- Support for multihoming
- Access to
  - end-to-end information
    - Path bandwidth
    - Latency
    - Error characteristics
  - local information
    - Link-layer aware
  - application requirements

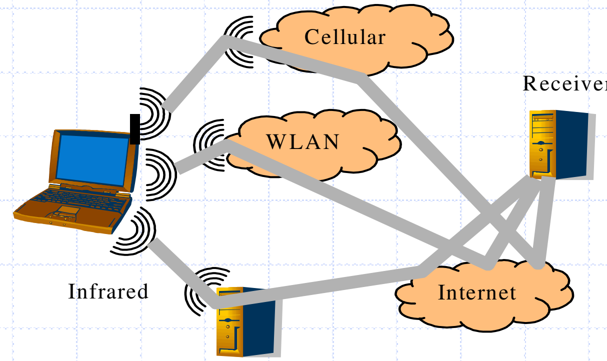

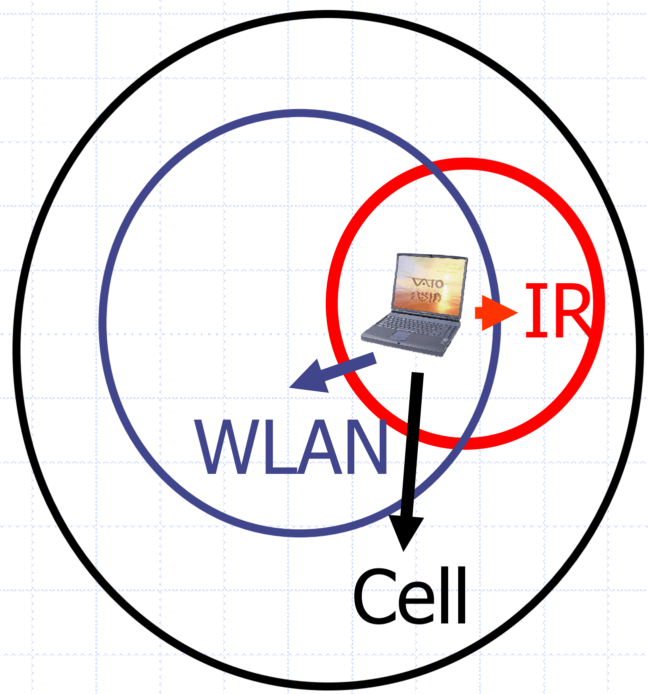

Contribution: Paradigm Change

- Move mobility support to the transport level using existing IP infrastructure
- Thesis work
  - Mobility architecture
  - Framework for multiplexing transport protocols
  - Protocol examples
  - In-depth study of protocol elements
    - Multiplexing
    - Congestion control for rate-based protocols
    - Transport protocols for wireless host
      - Loss discrimination

Outline

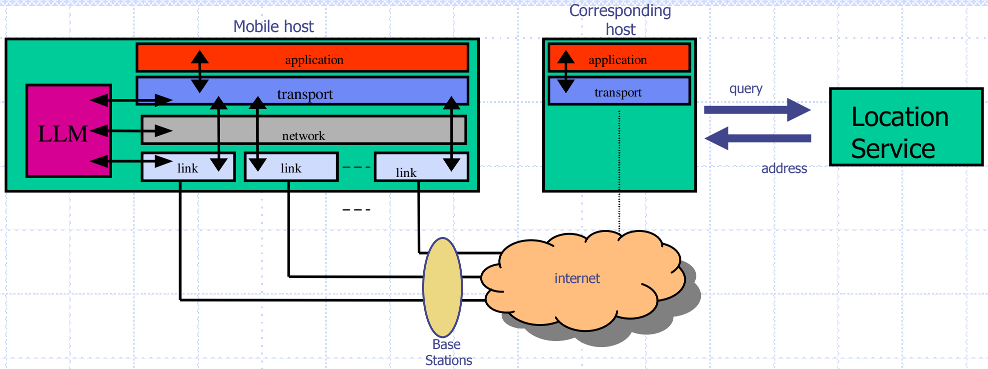

- Approach
- Architecture for link-layer aware, inverse multiplexing transport protocols
- Protocol suite
  - Multimedia Multiplexing Transport Protocol
  - Reliable Multiplexing Transport Protocol
- Protocol characteristics
- Protocol mechanisms
  - Inverse Multiplexing
  - Congestion Control
  - Loss Discrimination

Approach

Design Constraints

- Environmental limitations:
  - Wireless:
    - Low bandwidth
    - High error rate
    - High latency
  - Mobility:
    - Low power
    - Changing attachment points
- But: possible multiplicity of access points of different technologies in any area

Transport Layer Approach

- Inverse Multiplexing by B/W measurement
  - Sending one flow through multiple interfaces
- Transport protocols
  - Link layer-aware
  - Network layer-independent
- Benefits
  - Built-in mobility
  - Seamlessness
  - Adaptability
  - Bandwidth aggregation
  - Informed choice on which link-layer to use

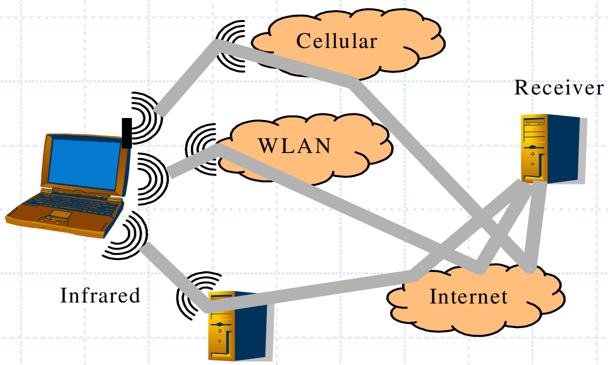

Full Mobility Architecture

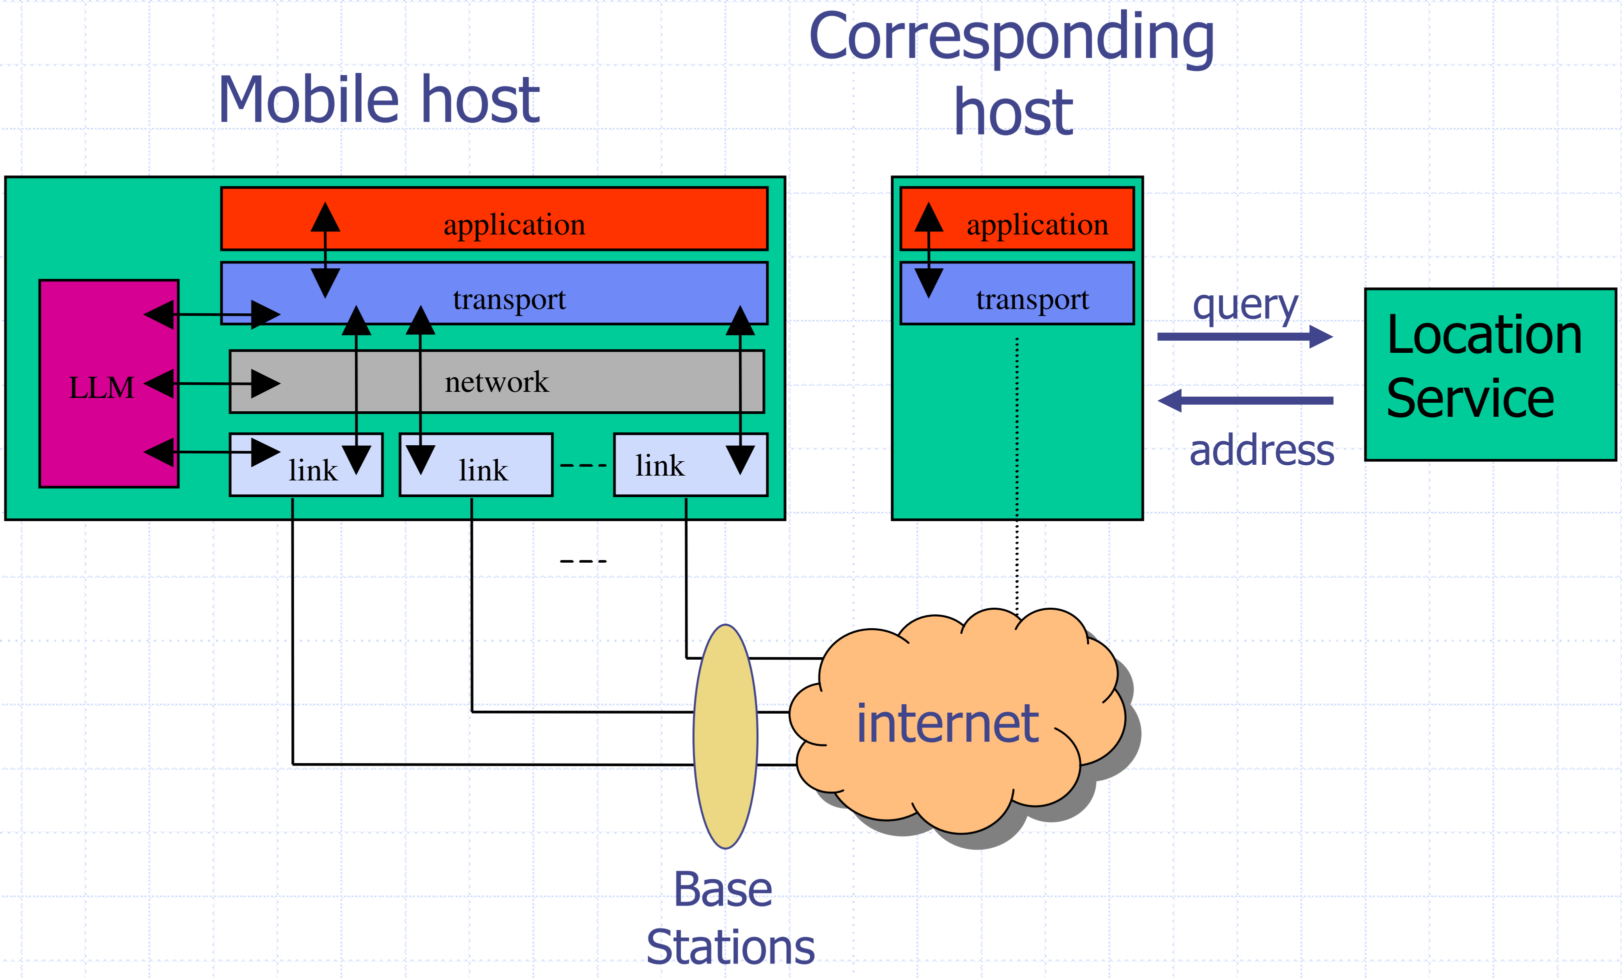

Protocols

MMTP

- Service
  - Best effort delivery of multimedia
- Protocol Details
  - If excess BW is available, channel with highest latency is filled first
  - If not enough BW, frames are dropped at sender

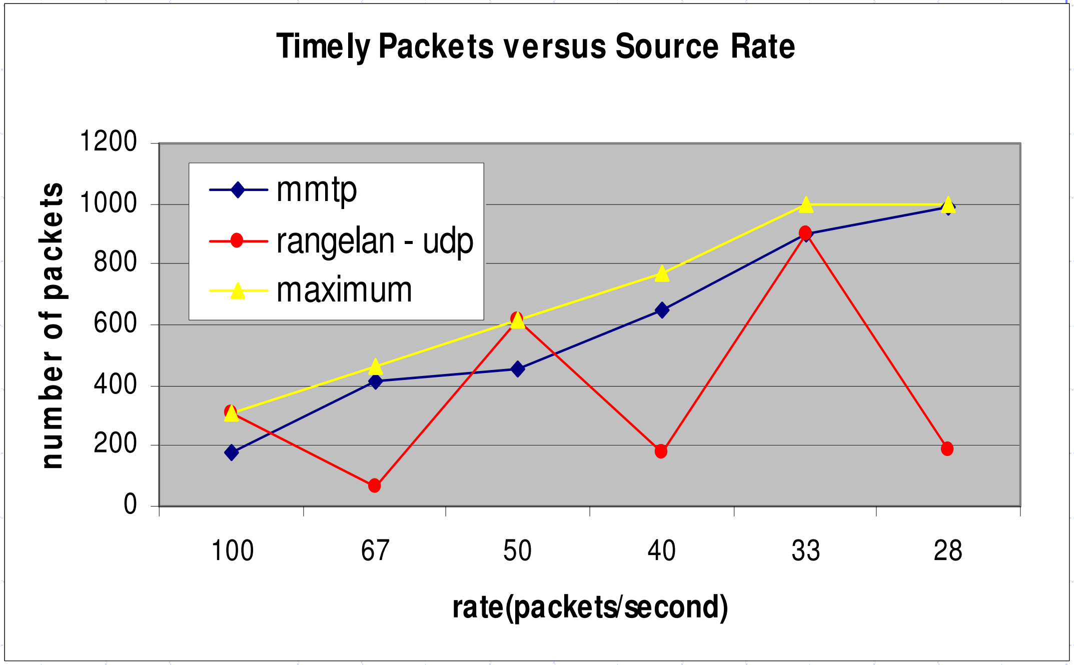

- In SIGCOMM-LA "MMTP - Multimedia Multiplexing Transport Protocol"

R-MTP

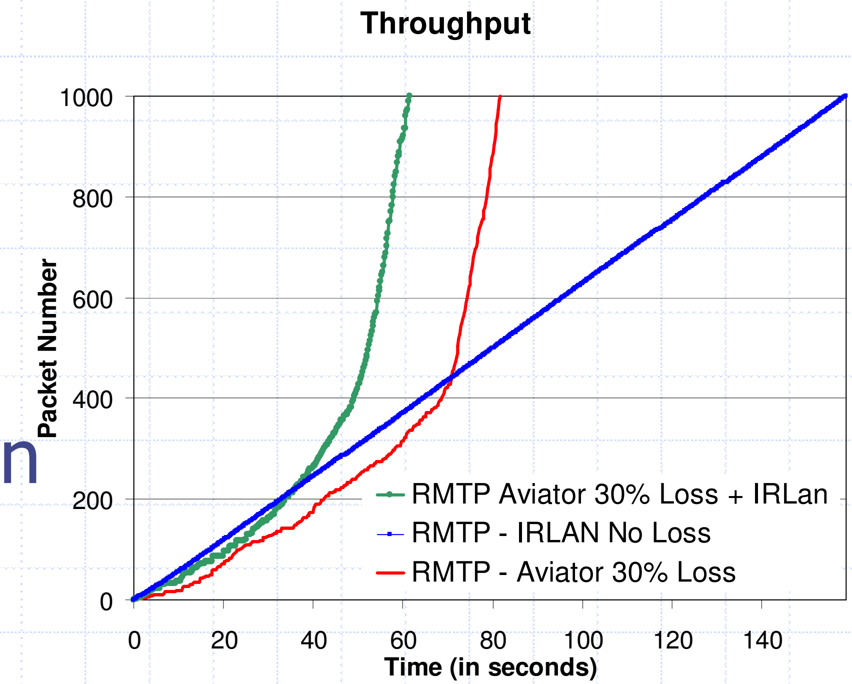

- Service
  - Reliable transmission of bulk data
- Protocol details
  - Multiple channel, rate-based
  - Selective acknowledgements for reliability
  - Bandwidth estimation for flow and congestion control

- ICNP 2001 "**Transport Level Mechanisms for Bandwidth Aggregation on Mobile Hosts**"

Characteristics

Multiplexing

- Load balancing of data transmission
  - Base individual channel load on corresponding share of total bandwidth

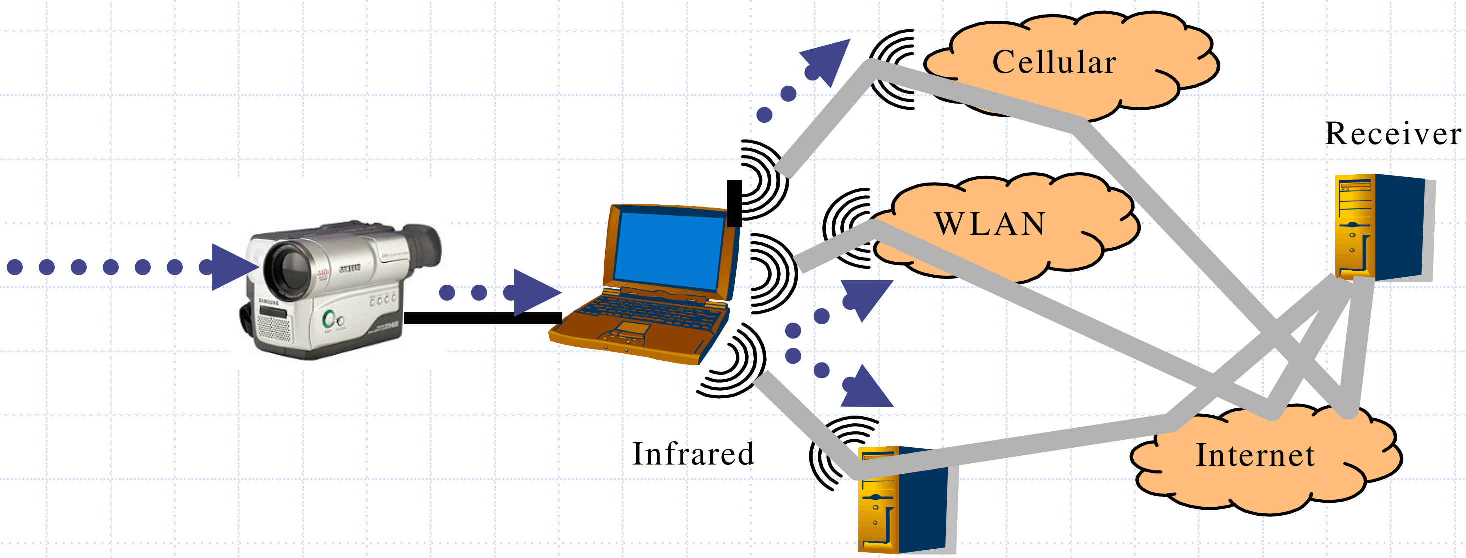

Basic Channel Mechanism

- Rate-based transmission

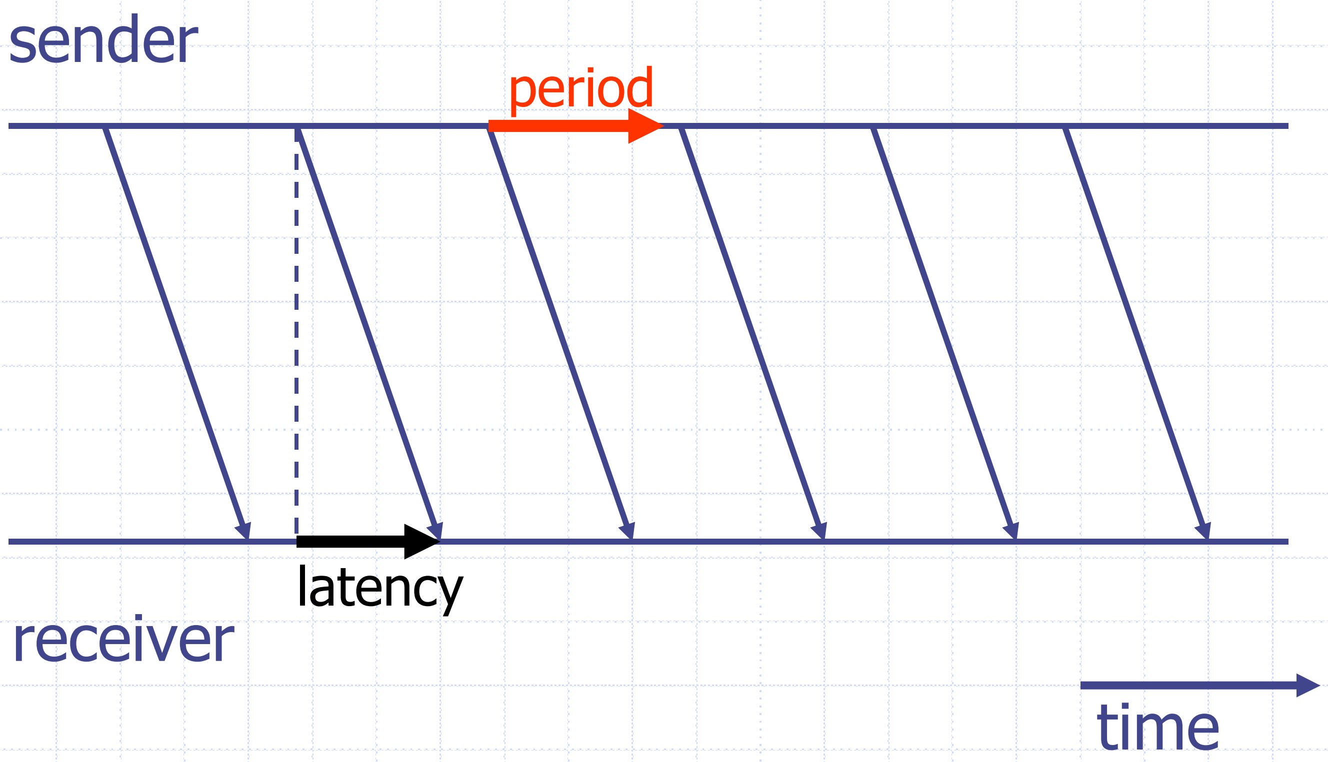

Bandwidth Tracking

- Probing using packet-pair

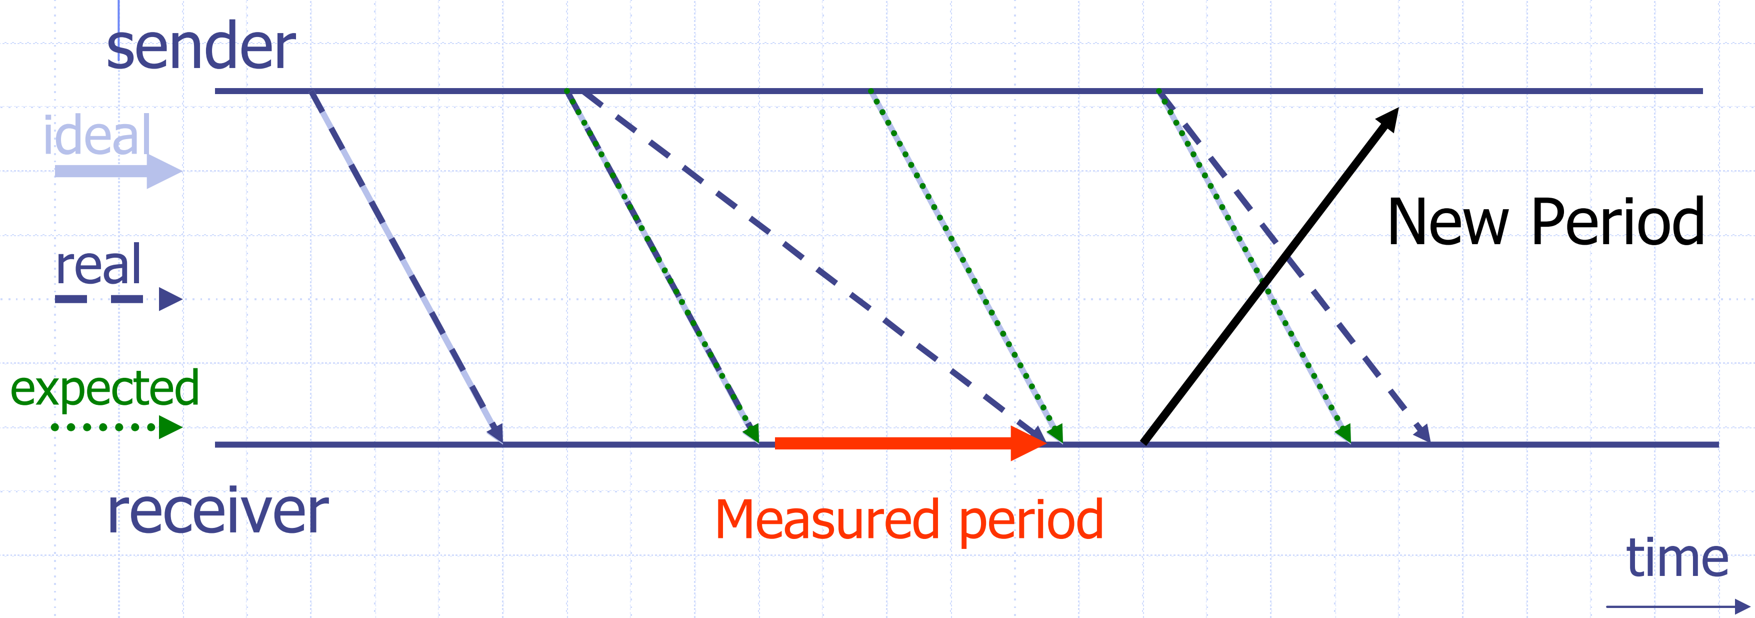

Congestion Avoidance

- Congestion avoidance
  - Reduce sending rate before causing packet loss

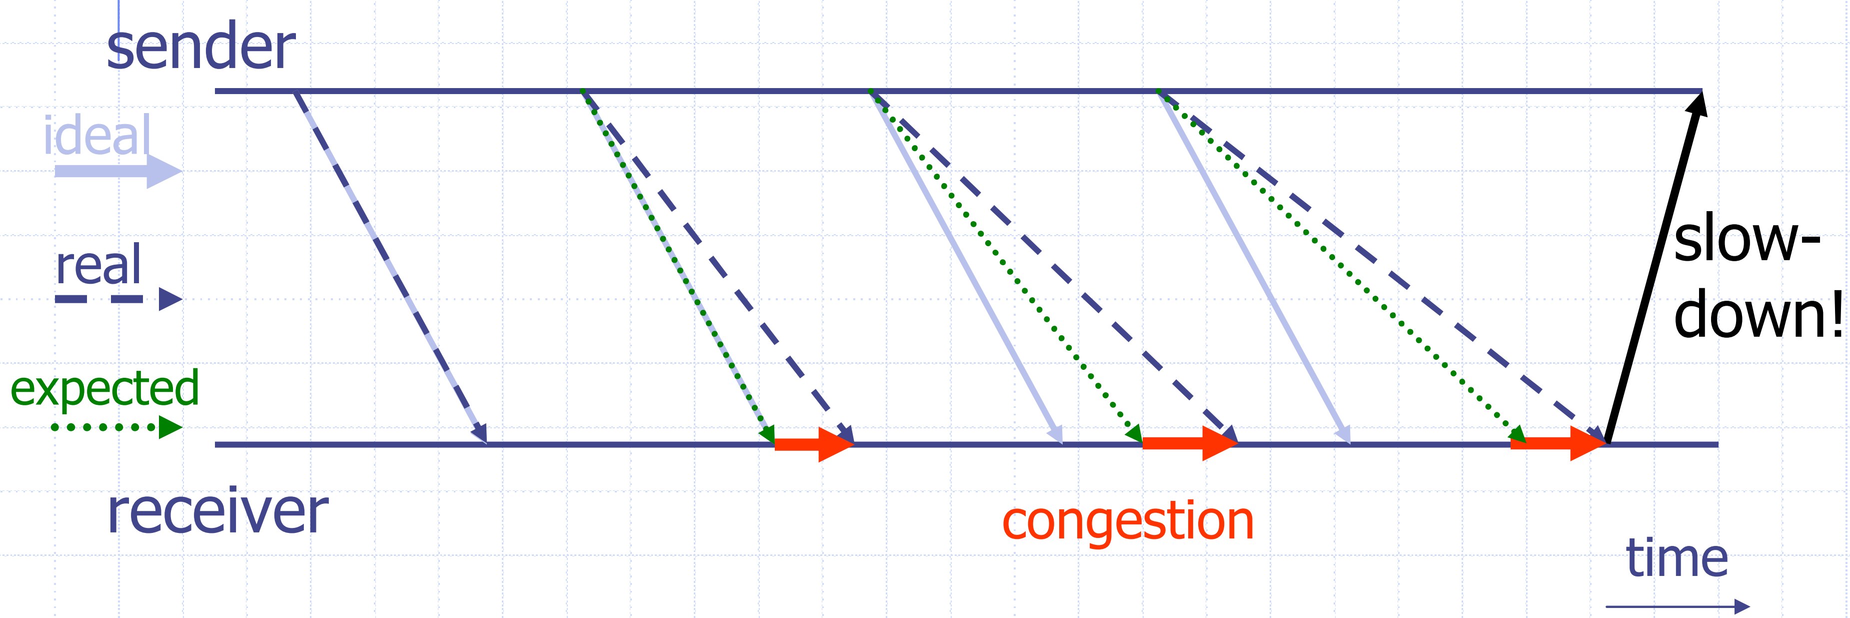

Loss Discrimination

- Distinguish congestion losses from transmission losses

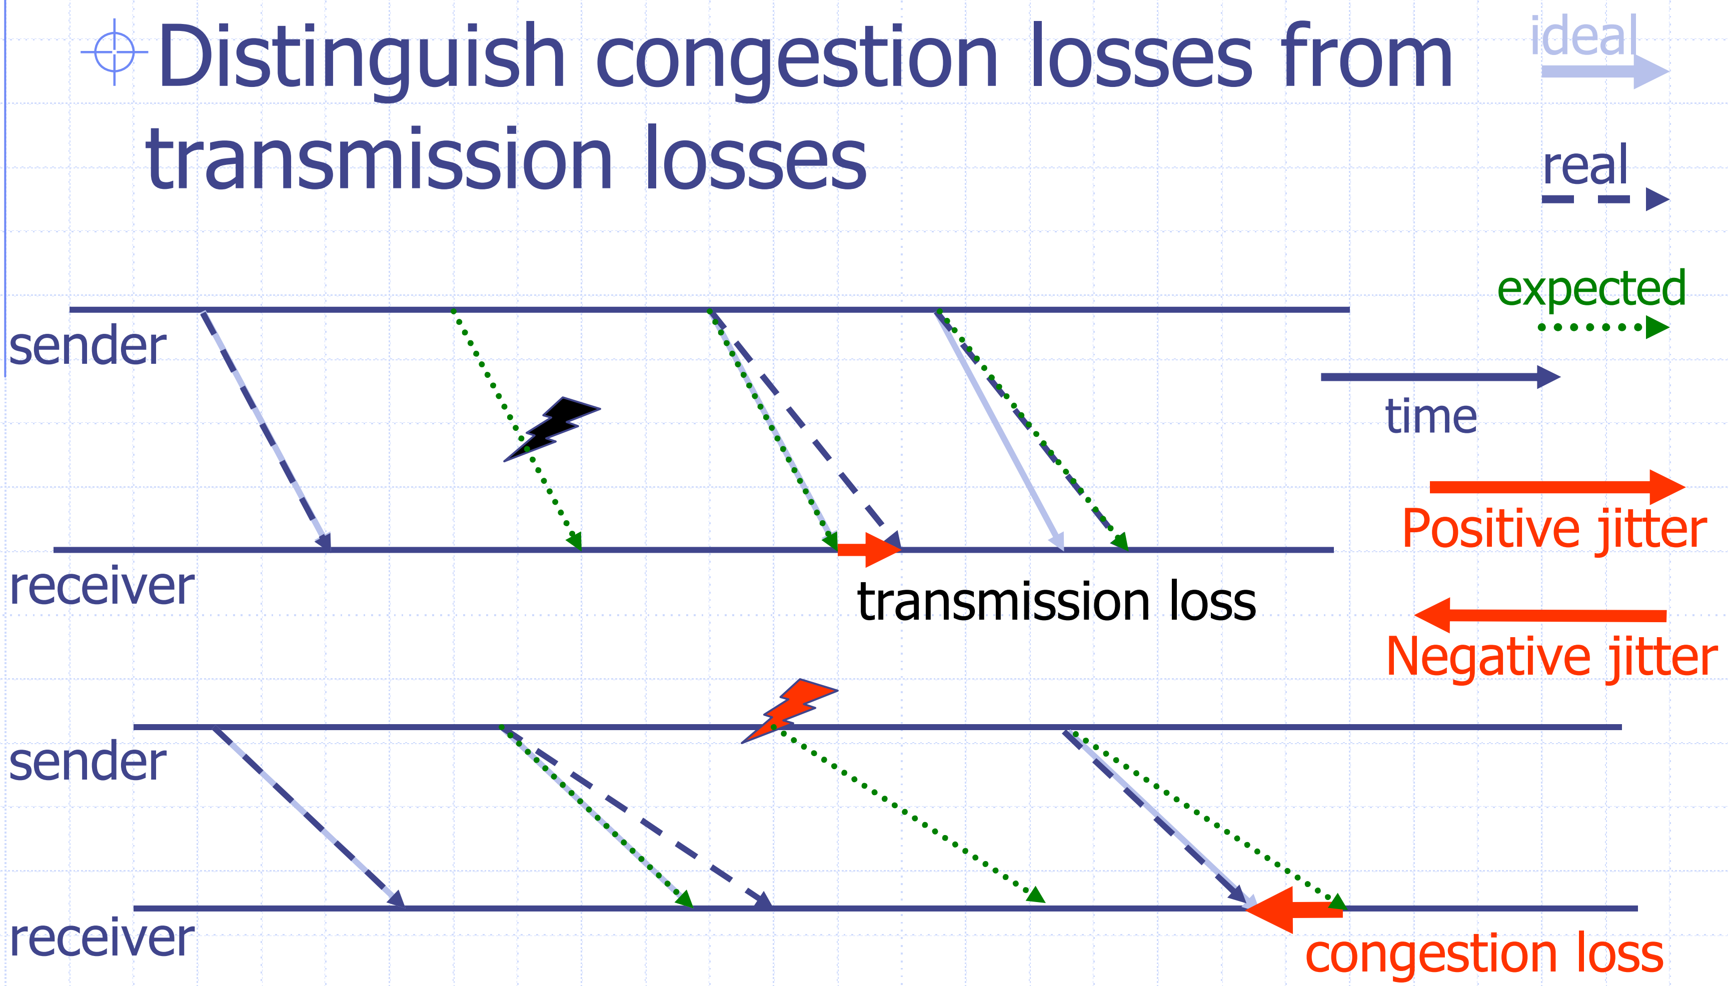

Multiplexing

Inverse Multiplexing

- Approach
  - Send data corresponding to the fair share bandwidth on each channel
- Challenges
  - Measuring fair share bandwidth
  - Data reordering
- Solutions
  - Rate-based transmission mechanism
  - Admission & monitoring of delay
    - constrain the use of channels with large delays

Advantages of Inverse Multiplexing

- Bandwidth aggregation
- Fast feedback path
- Increased performance on lossy channels
- Smooth handoffs
- Intelligent channel selection

Results: BW Aggregation

- If multiple technologies are present on the same area, RMTP can use them to increase throughput

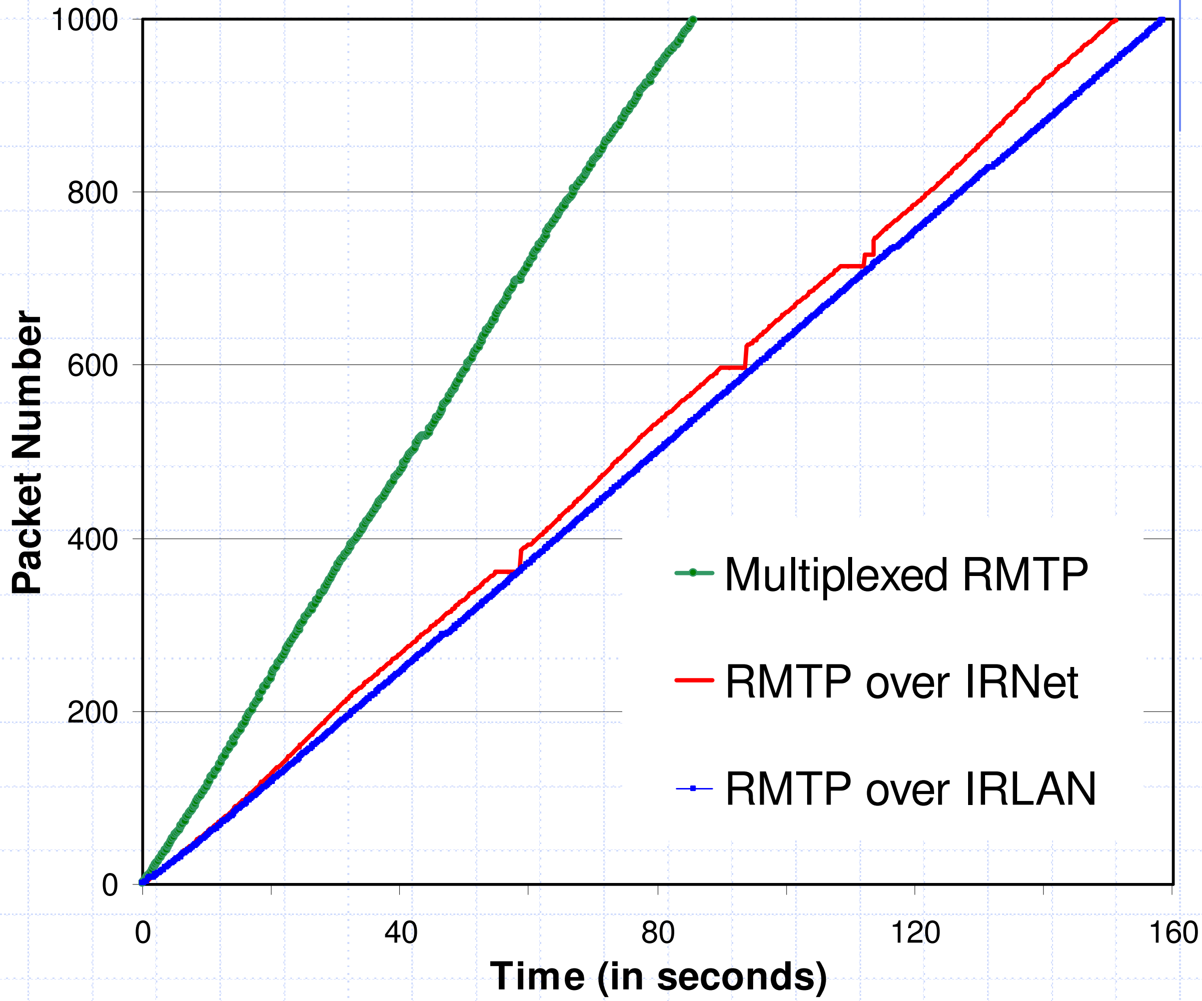

Informed Interface Use

- Measuring BW and using all channels lead to better performance
  - TCP performs better on the slower (115Kb) link with no loss than on the faster (2Mb) with losses

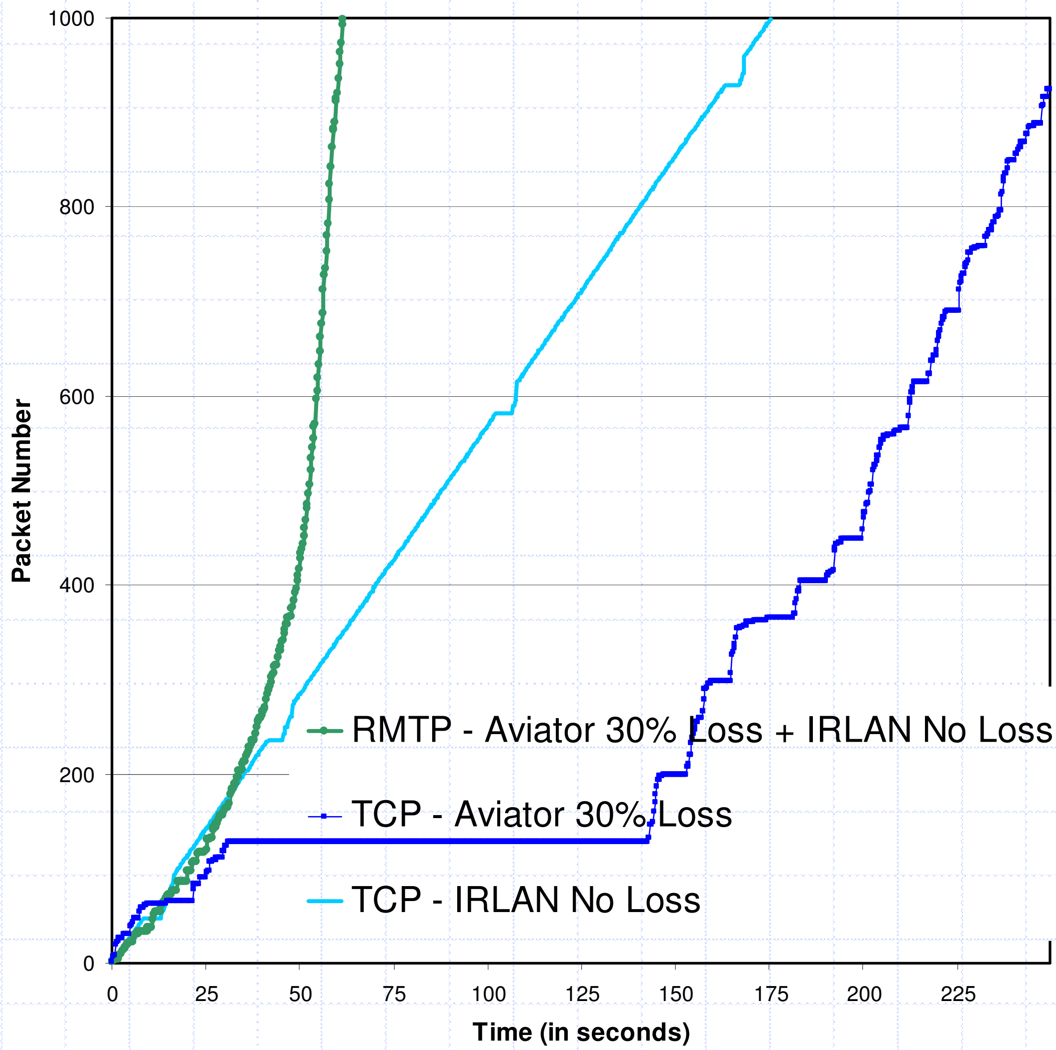

Congestion Avoidance and Control

Congestion Control for Rate-Based Protocols

- Homeostatic Congestion Control
  - Balance Point
    - Fair-share Bandwidth
  - Forces

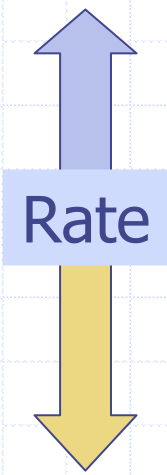

- Bandwidth Estimation
  - Probes the network to find what is the fair share of bandwidth
- Congestion Avoidance
  - Lowers the rate if it exceeds network availability
- Congestion control by loss detection

Algorithm

- Exponential increase
  - Packet pairs
    - Once every 5 packets
  - Tracking Period
    - Low pass filter
- Congestion avoidance
  - Measurement-based decrease
    - Error measurement (jitter)
- Congestion Control
  - Multiplicative decrease

$P(n+1) = (1 - \alpha) * P(n) + \alpha * MeasuredPeriod$ (1)

$\alpha \in [0,1]$

MeasuredPeriod = interarrival time

$Error = ((1 - \alpha)^n) * (P0 - POptimal)$ (2)

n - number of measurements

NewPeriod = OldPeriod + (jitter(1)+…+jitter(n))/n (3)

Where n is 2 or 3

If (current_time > time_last_loss + RTT + 2* OldPeriod)

NewPeriod = OldPeriod*2

Challenges

- Synchronization
  - All flows experience same losses
  - Lower overall network utilization
- Solution: add random quantity to period

$P(n+1) = (1 - \alpha) * P(n) + \alpha * (1 + \beta*rnd) * MeasuredPeriod$

  - In our tests, $\beta = 0.1$
- Fair share versus available bandwidth
  - Any flow can use all resources
  - Flows that only use left-over BW can starve
  - Equilibrium between flows means being aggressive
- Solution: packet pair randomness

Problems with Packet Pair

- Normal flow

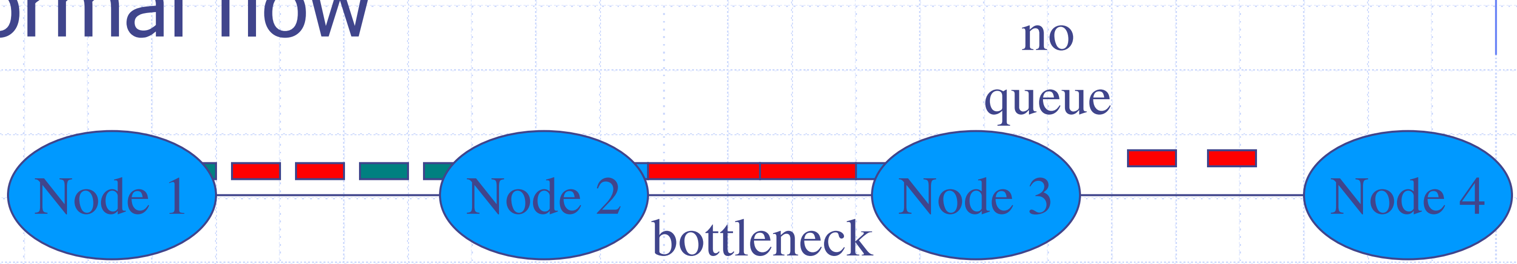

- Time compression

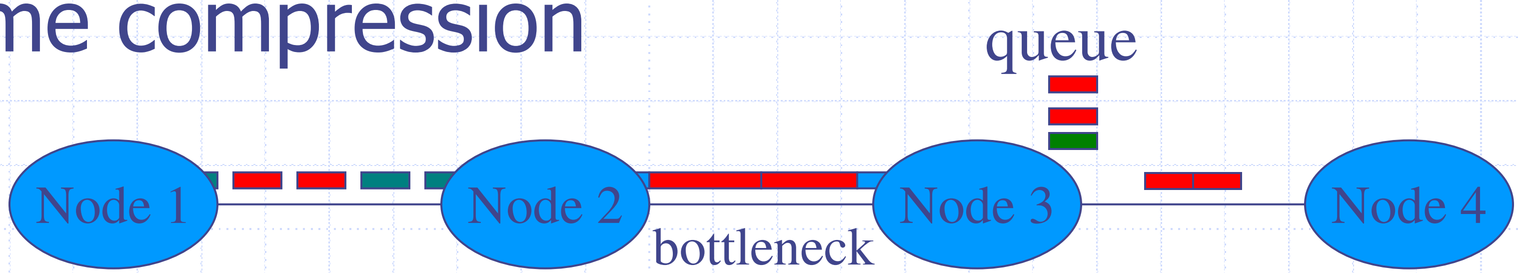

- Time expansion

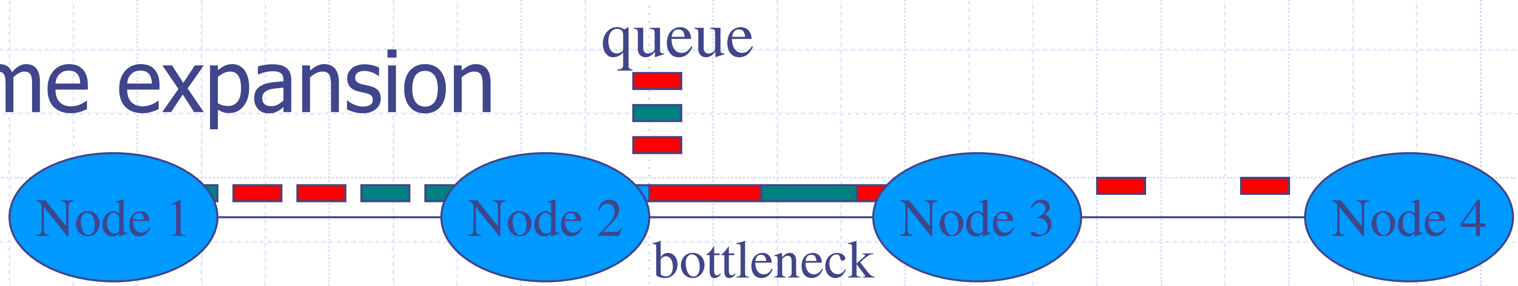

Coping with Uncertainties in the Measurements

- Let Pc be the current period, PMeasured the measured period and POptimal the true optimal period. We have 3 cases
  - Too small

    Pmeasured < (POptimal – (1–α) * Pc) / α
  - Sweet spot

    POptimal > PMeasured > (POptimal - (1-α) * Pc) / α
  - Too large

    PMeasured > POptimal

Next period:

P(n+1) = (1 – α) * P(n) + α * MeasuredPeriod

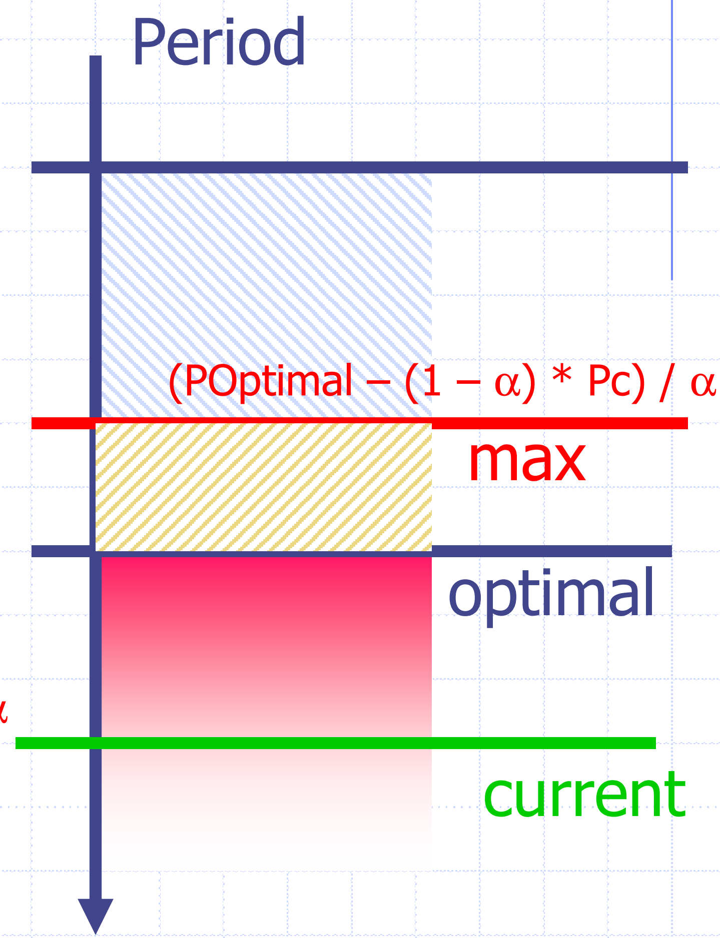

Simulation: Bandwidth Tracking

- XMTP is less sensitive to latency than TCP

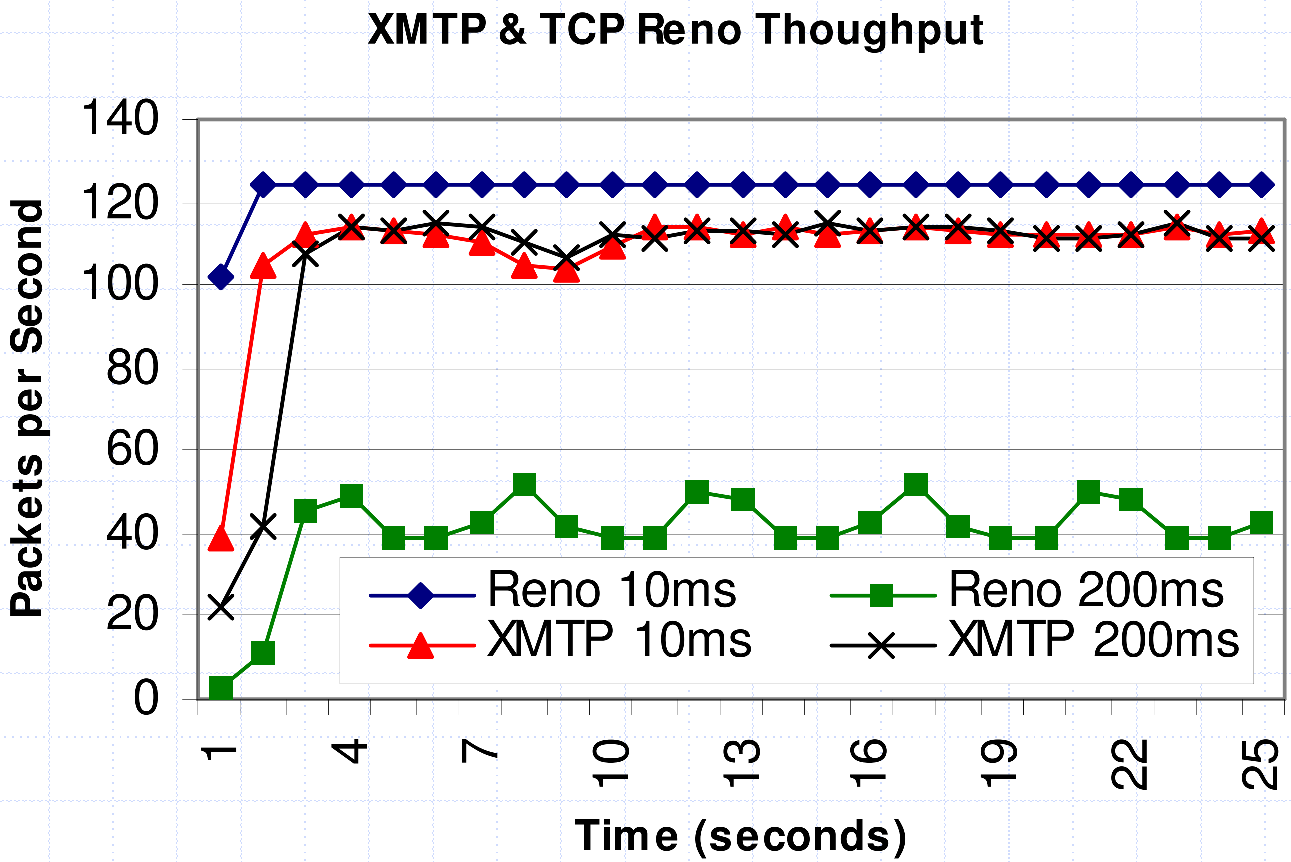

Simulation: Multiple Flows

- Bottleneck delay 50ms
- Usage goes up with mixed XMTP traffic
- Fairness does not reflect added usage.

	Flow 1	Flow 2	Flow 3	Flow 4	Flow 5	Usage	Fairness
5 TCP	3450	3366	3344	3280	3380	16820	0.999733
1XMTP/4TCP	14718	3422	3420	3416	3412	28388	0.612052
2XMTP/3TCP	9947	8505	2651	2997	2759	26859	0.740284
3XMTP/2TCP	1957	7307	8493	1364	2256	21377	0.670529
4XMTP/1TCP	6904	812	5714	7289	2199	22918	0.756063
5 XMTP	3458	3389	3963	5071	4896	20777	0.971885

Simulation: TCP Friendliness

- Depending on the delay, TCP or XMTP may dominate, but they do not starve each other

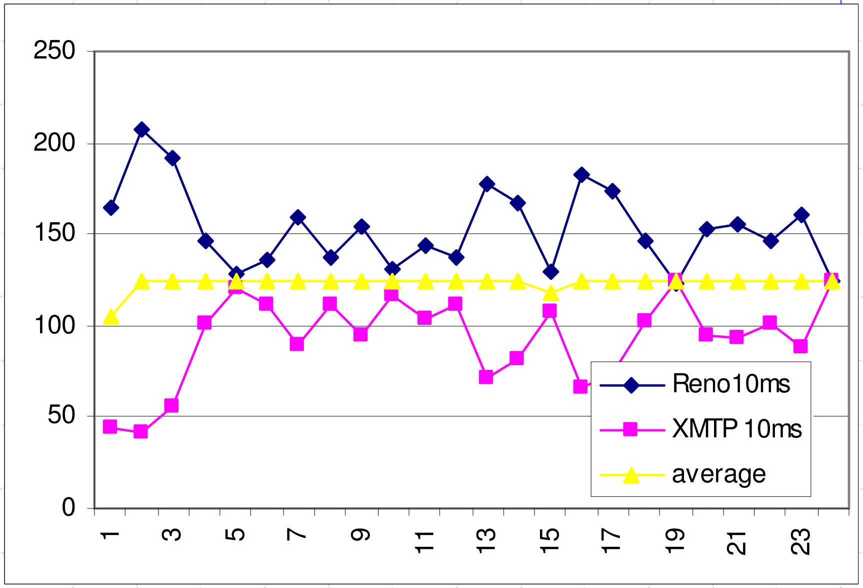

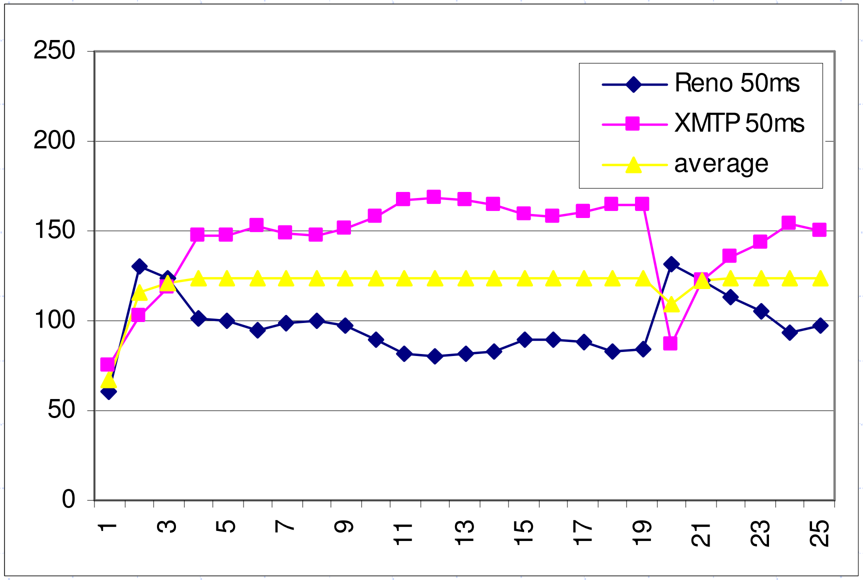

Loss Discrimination

What is Loss Discrimination?

- To decide if a loss was caused by congestion or by transmission error
- Influences the reaction to loss
  - Transmission: resend data
  - Congestion: use congestion control mechanisms

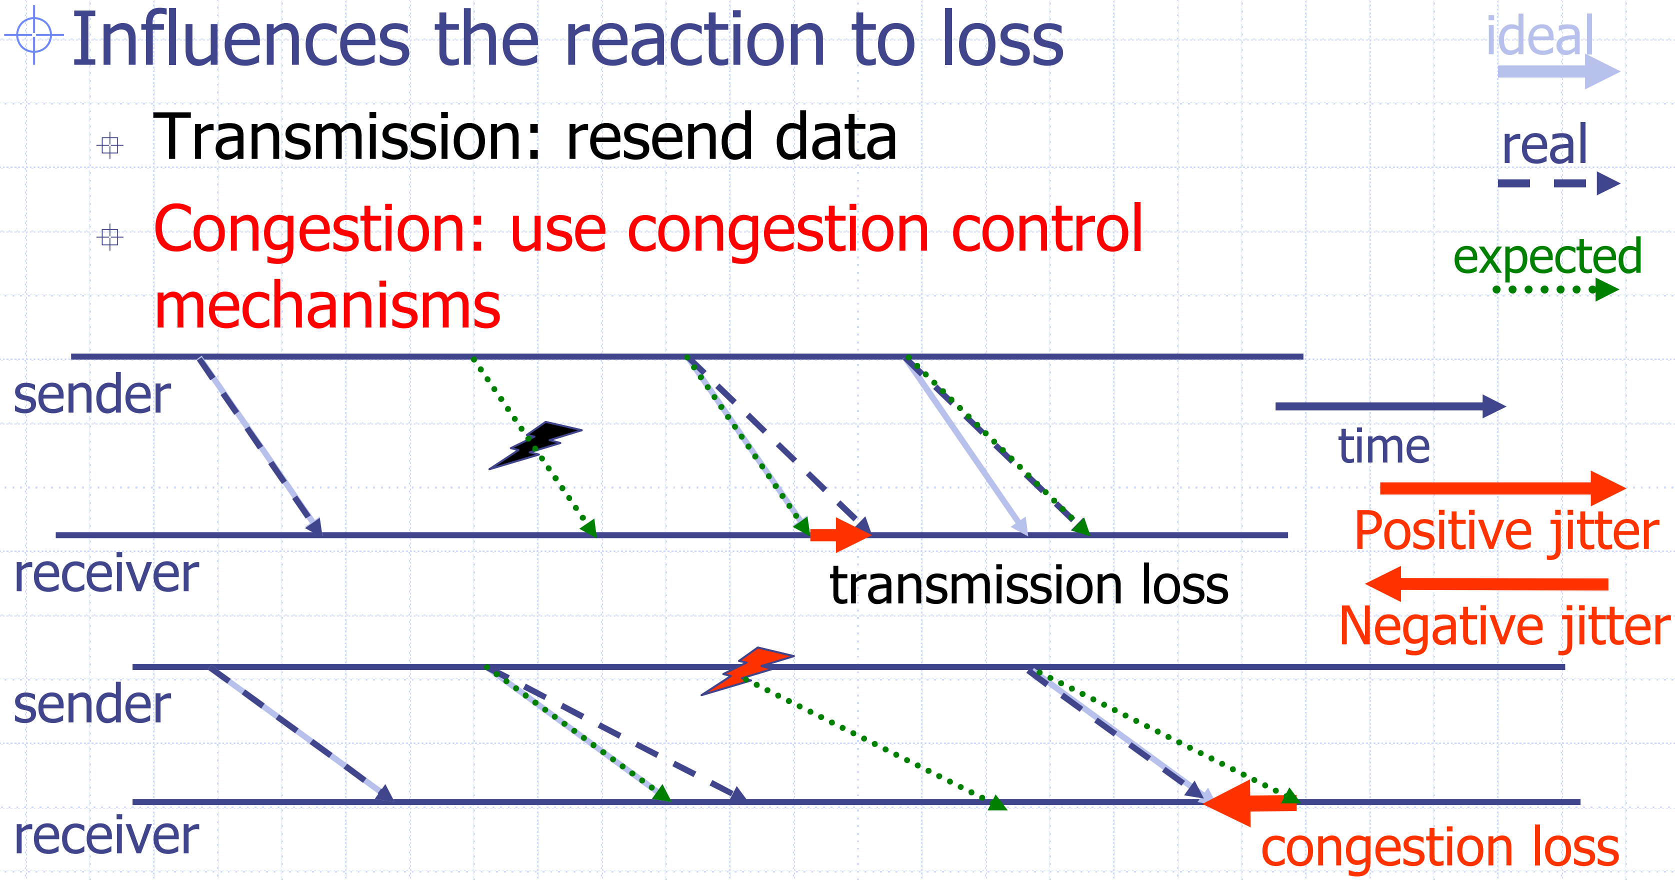

Queue Sizes and Jitter

- Interarrival time depends on the difference of sending times and transit times
- Transit time is flight time plus queue time. Flight time is invariant if
  - Routes are stable
  - Packet sizes are constant
- Jitter is caused by queue sizes seen by packets

$IAT_{i,\, i-1} = TT_i - TT_{i-1} + TS_i - TS_{i-1}$

If P is the period of a rate-based protocol

$TS_{i+1} = TS_i + P$

We expect IAT = P

Defining Jitter = IAT - P,

We get

$J = TT_i - TT_{i-1} + P - P$

Therefore

$J_i = FT_i - FT_{i-1} + Q_i - Q_{i-1}$

$J_i = Q_i - Q_{i-1}$

Heuristics

- If the jitter following a loss is negative, the loss is deemed a congestion loss.

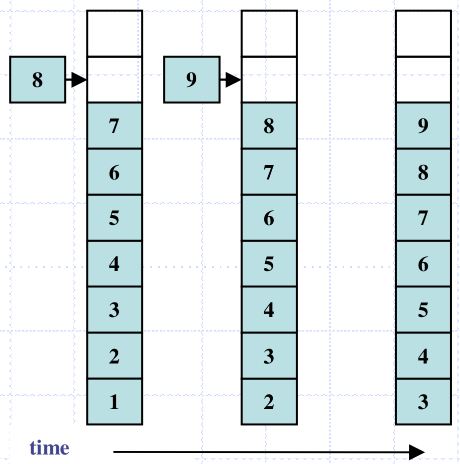

uncongested

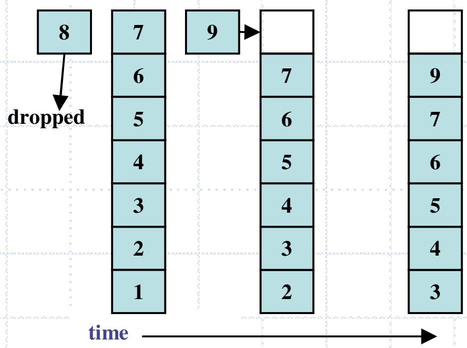

congested

Performance

- Simulation
  - 3 1MB links
  - Delay 40ms

Protocol	Number of packets		
	No Loss	0.1% Loss	1% Loss
NewReno	2694	2808	1415
Westwood	2763	2736	1812
XMTP	2855	2813	2592

Protocol	Number of packets	
	0.1% Loss	1% Loss
NewReno Losses	6	18
Westwood Losses	6	22
XMTP Losses	4	37
XMTP Congestion	0	3

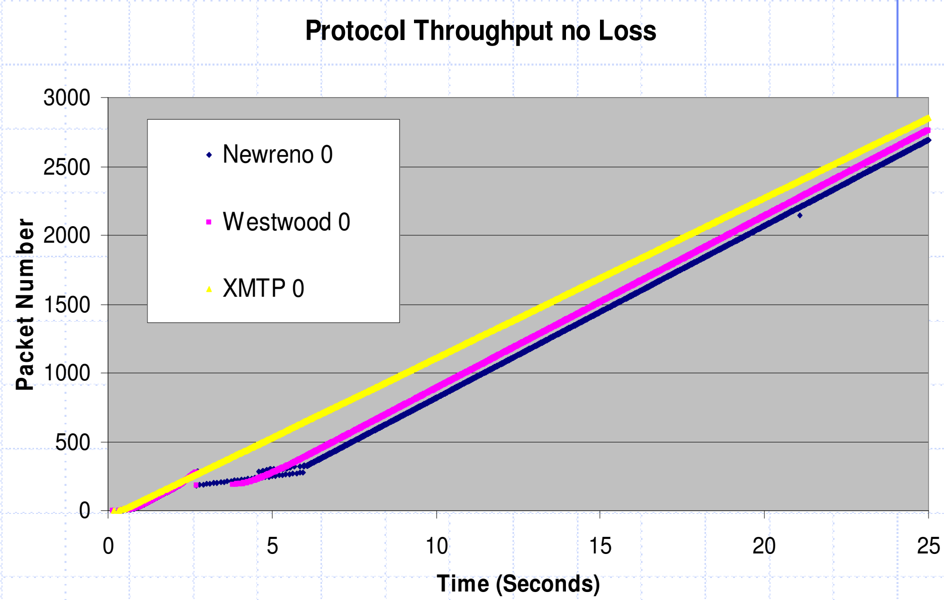

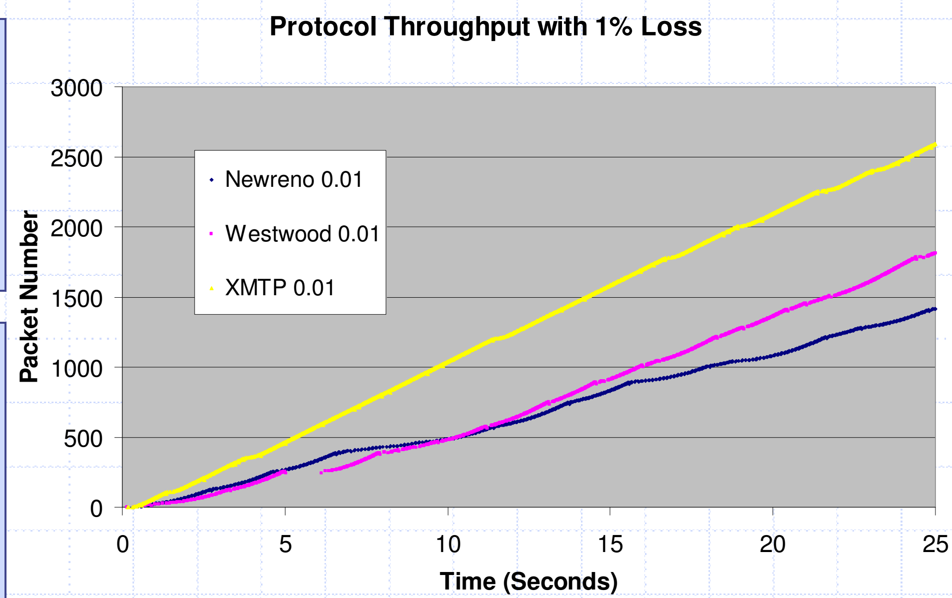

Hit Ratio

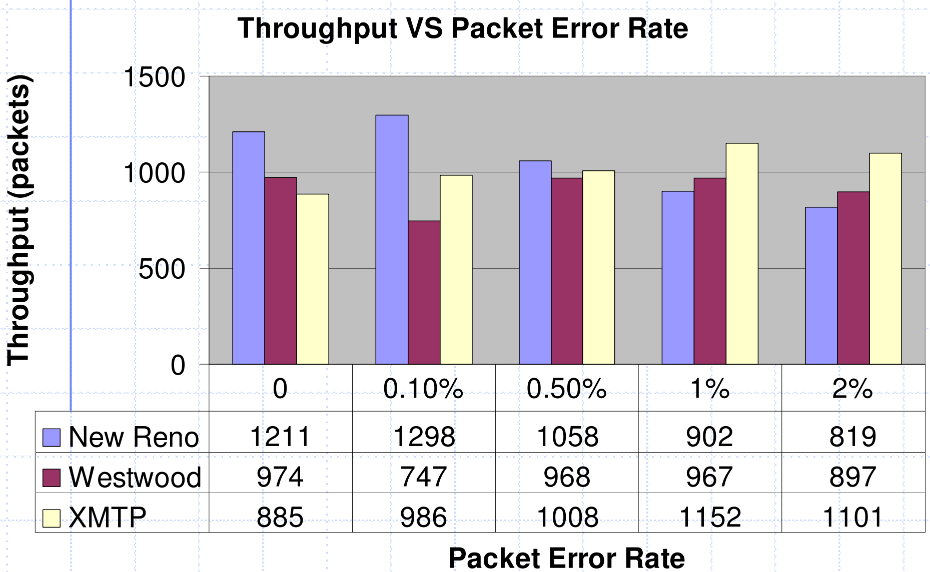

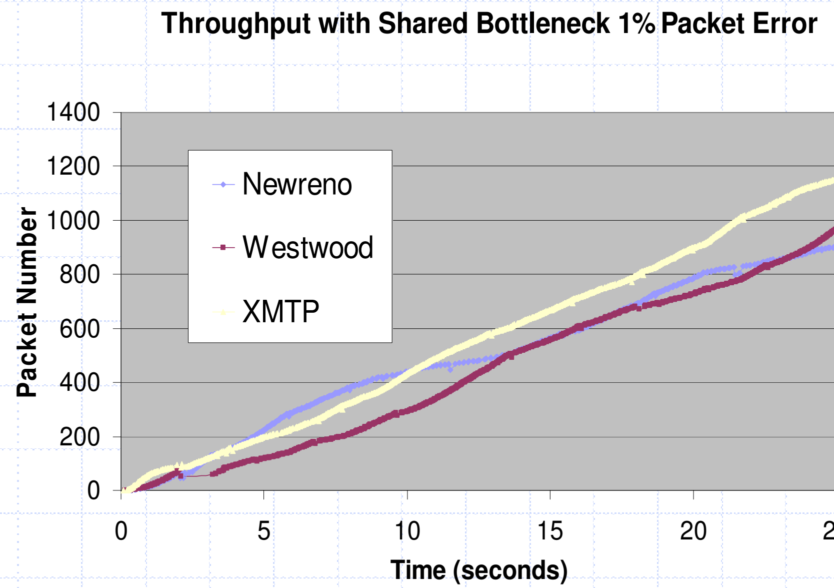

Error Rate	Total drops	Trans.	OK	Not	Cong.	OK	Not
0	32	5	0	5 16%	27	27	0
0.1%	28	8	2	6 21%	20	20	0
0.5%	23	6	3	3 13%	17	14	3 13%
1.0%	23	7	6	1 04%	16	15	1 04%
2.0%	29	9	9	0	20	0	20 68%

Packet Error Rate	Bit Error Rate (x 10^{-6})
0.1%	0.1125
0.5%	0.5625
1.0%	1.1250
2.0%	2.2500

Conclusion

Future Work

- Link Layer Manager
  - Identify available link layers
  - Establish link layer connections
  - Acquire IP addresses
- Location Service
  - Allow corresponding hosts to find current address of mobile
- Power Management
  - Efficient use of energy resources in the context of multiple channels

Contributions

- Architecture for transport level mobility
- Techniques for bandwidth aggregation
- Homeostatic congestion controller
- Techniques for detection of transmission losses
- Protocol Suite
  - MMTP
  - RMTP

Related Research

- Communication Channel Multiplexing
  - ATM, PPP-Multilink, EtherChannel, SCTP, pTCP
- Mobility
  - Barwan Project, Mobile People Project
- Reliable Transport Protocols in Wireless Environments
  - Rate-based, loss detection
- Bandwidth estimation
  - Packet pair, bandwidth measurement tools
- Loss discrimination
  - ECN, ELN, end-to-end

Publications

- **End-to-End Inverse Multiplexing for Mobile Hosts, L. Magalhaes and R. Kravets,** ***to appear in the Journal of the Brazilian Computer Society***
- **Transport Level Mechanisms for Bandwidth Aggregation on Mobile Hosts, L. Magalhaes and R. Kravets ,*****The 9th International Conference on Network Protocols (ICNP 2001), 2001***
- **End-to-End Inverse Multiplexing for Mobile Hosts, L. Magalhaes and R. Kravets,** ***The 19th Brazilian Symposium on Computer Networks, Florianopolis, Brazil, 2001***
- **MMTP: Multimedia Multiplexing Transport Protocol, L. Magalhaes and R. Kravets ,** ***The First Workshop on Data Communications in Latin America and the Caribbean (SIGCOMM-LA 2001), 2001***
- **A Cooperative Approach to User Mobility,R. Kravets, C. Carter, and L. Magalhaes,*****ACM Computer Communications Review*****, vol. 31, 2001.**
- **On-Demand TCP : Transparent peer to peer TCP/IP over IrDA, J. Tourrilhes, L. Magalhaes and C. Carter,** ***Proc. of ICC 2002***

A Transport Layer Approach to Host Mobility

Luiz Claudio Schara Magalhães
Department of Computer Science
University of Illinois at Urbana Champaign

http://mobius.cs.uiuc.edu
schara@telecom.uff.br