A Receiver-Centric Transport Protocol for Mobile Hosts with Heterogeneous Wireless Interfaces

Hung-Yun Hsieh,
Kyu-Han Kim, Yujie Zhu,
and Raghupathy Sivakumar
GNAN Research Group
School of Electrical and Computer Engineering
Georgia Institute of Technology
September 16, 2003

Untethered Communication

- ✓ Loss Recovery
- ✓ Congestion Control
- ✓ Power Management
- ✓ Seamless Handoffs
- ✓ Server Migration
- ✓ Bandwidth Aggregation

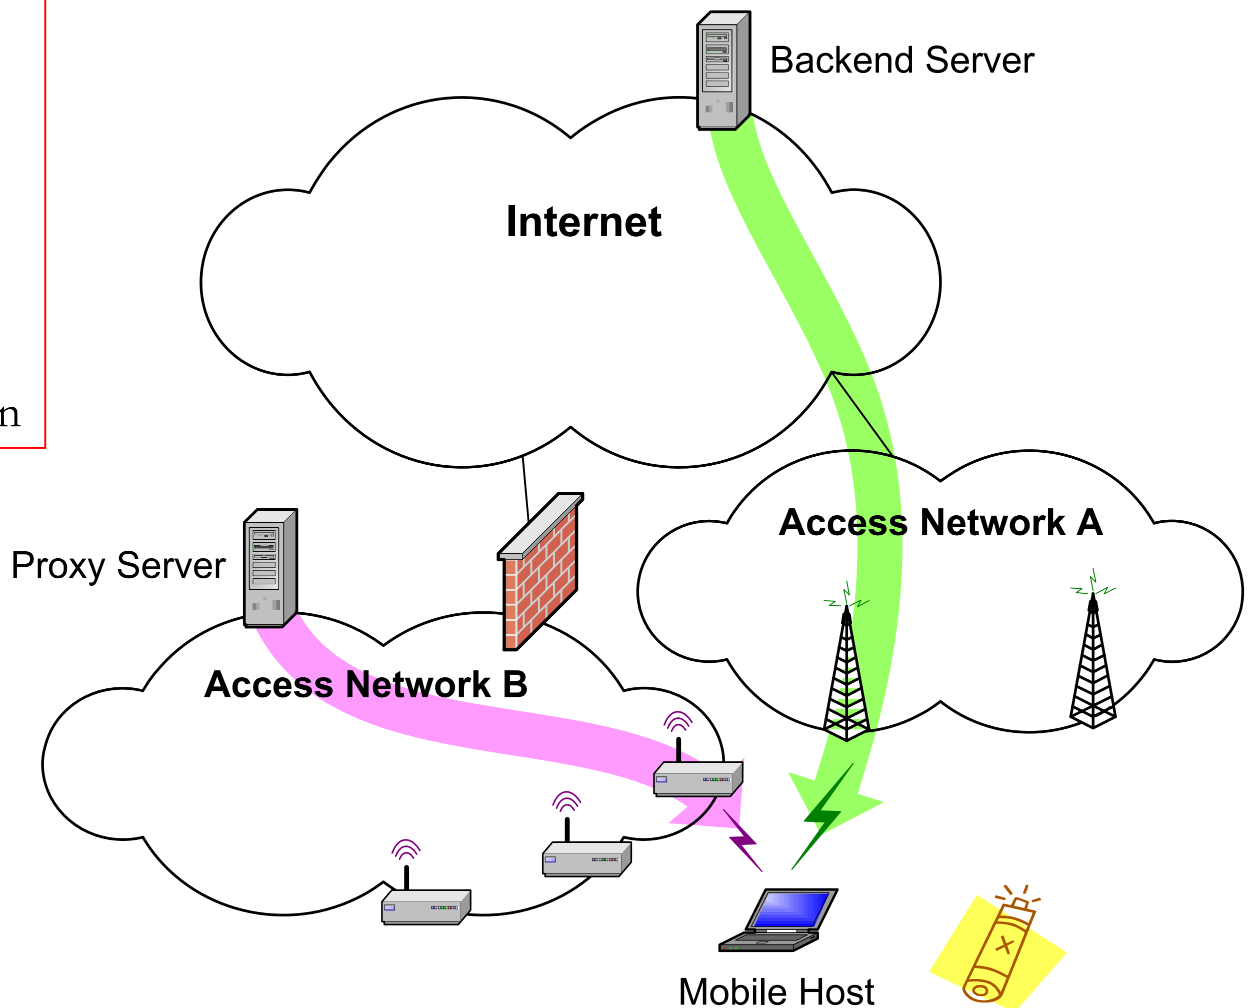

GNAN Research Group

Role of the Mobile Host

Focus of this work

- Systematic motivation for mobile-host-centric approaches in transport layer protocol design
- Propose and implement receiver-centric protocols for mobile hosts with one or multiple interfaces

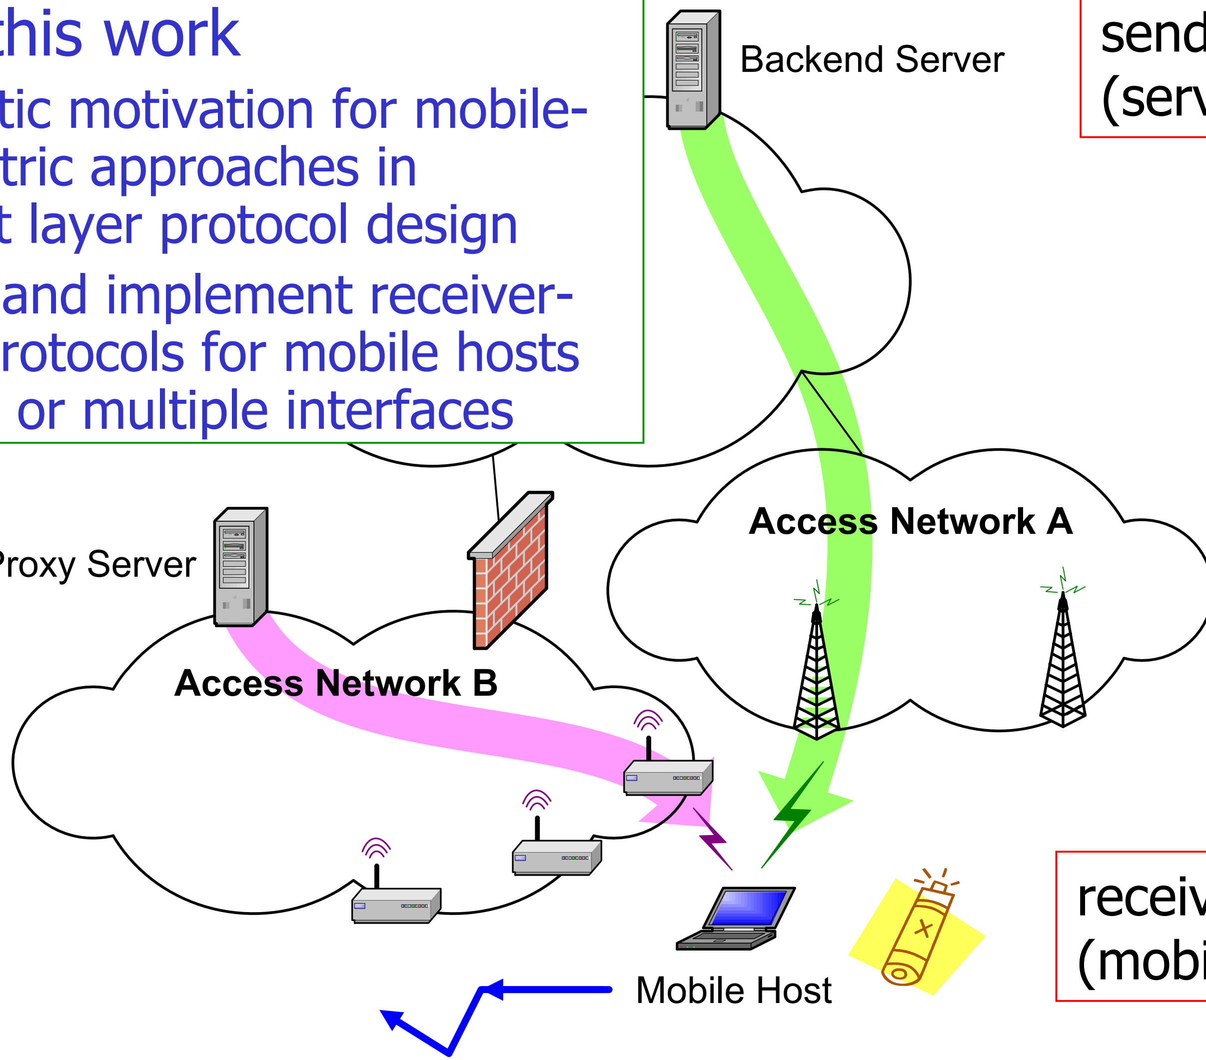

sender-centric
(server)

receiver-centric
(mobile host)

Outline

- Tackling the wireless last-hop
  - Motivation
    - Loss recovery
    - Congestion control
    - Power management
  - RCP: reception control protocol
- Supporting multi-homed mobile hosts
  - Motivation
    - Seamless handoffs
    - Server migration
    - Bandwidth aggregation
  - R^2CP: radial RCP
- Summary

GNAN Research Group

Loss Recovery

- Loss recovery
  - Detect the occurrence of losses, and react according to the nature of losses
  - Loss discrimination [Balakrishnan 96, Biaz 99]
- Sender-centric approaches
  - The sender needs feedback from the receiver to identify the nature of losses over the wireless link
  - ☞ Feedback information incurs overheads, latency, and can potentially be incomplete or unreachable
- Receiver-centric approaches
  - The receiver is fully aware of the loss occurrences in the receive buffer
  - The receiver can obtain first-hand information about the nature of losses over the wireless link

GNAN Research Group

Congestion Control

- Congestion control
  - The wireless last-hop has a defining role in the congestion control scheme used by the transport protocol
  - Optimality of interface-specific congestion control: STP [Henderson 97], WTCP [Sinha 99], TCP-Peach [Akyildiz 01]
- Sender-centric approaches
  - Optimal congestion control scheme cannot be used without the support from the sender
  - ☞ The sender needs to be configured with all possible schemes and be updated for any new access technology
- Receiver-centric approaches
  - The receiver needs to implement only the congestion control schemes pertaining to the interfaces it possesses
  - The sender does not need to be involved in the interface-specific schemes used by the receiver

GNAN Research Group

Power Management

- Power management
  - Operations of the transport protocol can have significant impact on the power consumption
  - Retransmission strategy based on the channel state [Zorzi 99, Tsaoussidis 00, Singh 02]
- Sender-centric approaches
  - The sender needs to be aware of the power conserving decisions taken at the receiver
  - ☞ Feedback consumes power, and can potentially limit the time granularity of the power conserving decision
- Receiver-centric approaches
  - The receiver drives the protocol operations, while the sender merely responds based on the receiver's direction
  - The receiver can take power-conserving decisions independent of the sender

GNAN Research Group

RCP: Reception Control Protocol

- A TCP clone in its behavior
  - TCP semantics: reliable, in-sequence data delivery
  - TCP-equivalent operations [Bansal 01]
  - Transposition of protocol functionalities
  - Receiver-centric congestion control, reliability, and flow control

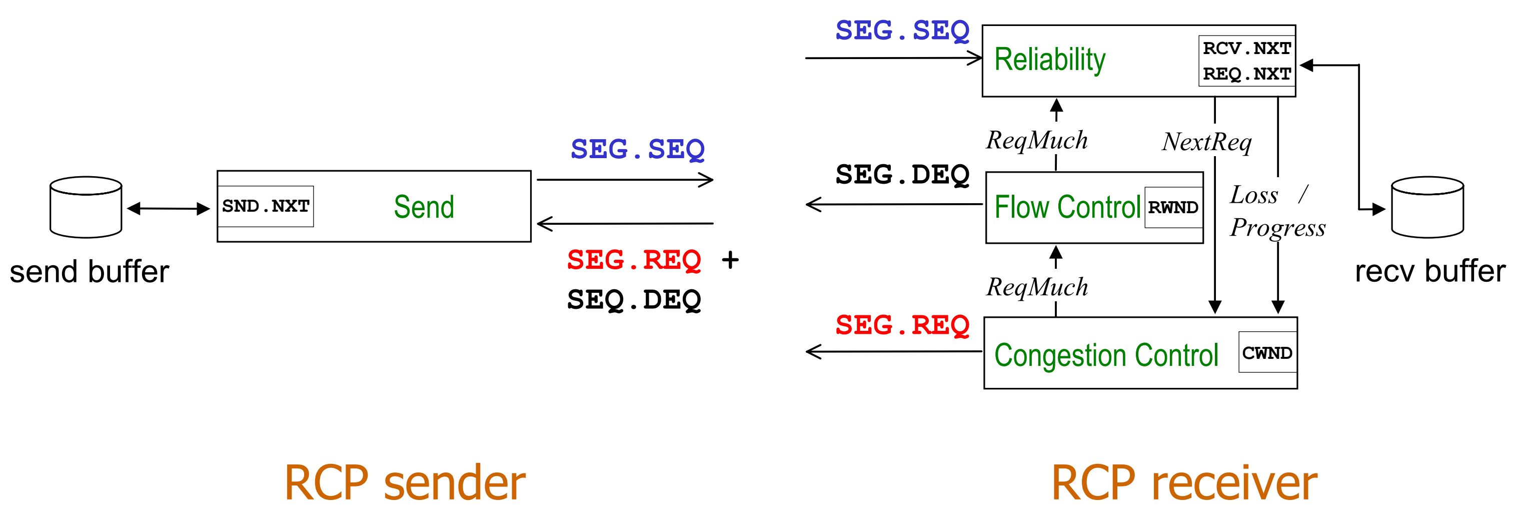

RCP Operations

- REQ – DATA handshake
  - Retain the self clocking mechanism in TCP that uses the DATA – ACK handshake
  - Allow the receiver to control which data the sender should send on a per-packet granularity
- REQ
  - Generated based on the available buffer space at the receiver (flow control) and the congestion window size
  - The window size controls the amount of pending REQs, while the return of DATA clocks the progression of the window (congestion control)
  - Loss is inferred through out-of-order DATA arrivals or timeouts – the receiver has complete knowledge of the state of the received buffer (reliability)
  - Improved robustness using two modes: cumulative mode and pull mode

GNAN Research Group

RCP Performance

- TCP friendliness
- REQ is robust to losses
  - Introduce on/off traffic in the reverse path (drop rate ~ 40%)
- Intelligent loss recovery
  - RCP-ELN achieves better performance than TCP-ELN (56% improvement for 1% packet loss rate)
- Scalable congestion control
  - Implement a receiver-centric version of STP called RCP-STP, which shares the same sender as RCP (RCP-NewReno)
  - RCP-STP achieves the desired performance as STP in a satellite environment (RTT ~ 500ms)
- Efficient power management
  - Use the WiFi device and two-state Markov error model
  - RCP achieves power savings without suffering from the performance degradation in TCP

GNAN Research Group

Outline

- Tackling the wireless last-hop
  - Motivation
    - Loss recovery
    - Congestion control
    - Power management
  - RCP: reception control protocol
- Supporting multi-homed mobile hosts
  - Motivation
    - Seamless handoffs
    - Server migration
    - Bandwidth aggregation
  - R^2CP: radial RCP
- Summary

Multi-homed Mobile Hosts

- Performance tradeoffs of different access technologies
  - Network capacity (data rate), coverage area, mobility support, and transmission power
- Challenges
  - Different autonomous domains
  - Heterogeneous networks with very different characteristics in terms of bandwidth, delay, loss rate, and fluctuation
  - ☞ seamless handoffs
  - ☞ server migration
  - ☞ bandwidth aggregation

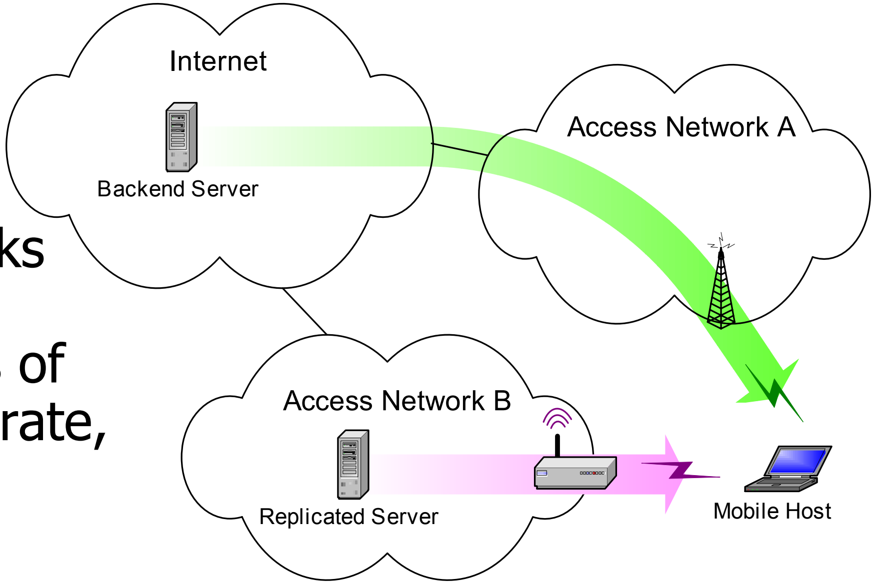

☞ An end-to-end approach without network support (cf. interworking architecture [BRAN 01])

GNAN Research Group

Seamless Handoffs

- Seamless handoffs
  - A transport layer solution does not incur latency exhibited in using Mobile IP [Snoeren 00 (TCP-Migrate)]
  - A transport protocol with multiple states can facilitate seamless handoffs [Riegel 03 (SCTP), Hsieh 03 (pTCP)]
- Sender-centric approaches
  - The sender needs to be informed of the handoff decision and status at the receiver
  - Interface-specific congestion control schemes cannot be used in a live connection without the sender support
- Receiver-centric approaches
  - The receiver is fully aware of the handoff decision
  - The receiver can use interface-specific congestion control schemes required during handoffs without involving the sender

GNAN Research Group

Server Migration

- Server migration
  - Server migration is necessary or desirable for achieving service continuity when the mobile host moves across different autonomous domains
  - A transport layer protocol can facilitate state synchronization required for server migration [Snoeren 01]
- Sender-centric approaches
  - Overheads in synchronizing the (complex) sender states
  - The receiver needs to purge the receiver buffer to avoid adverse reaction from the new sender
- Receiver-centric approaches
  - The sender maintains the minimal state, and hence the overhead is minimized
  - The receiver can request only the un-received data from the new sender after server migration

GNAN Research Group

Bandwidth Aggregation

- Bandwidth aggregation (striping)
  - Best-effort vs. policy-based bandwidth aggregation [Hsieh 03]
  - Point-to-point (one source) vs. multipoint-to-point (multiple sources) bandwidth aggregation
- Sender-centric approaches
  - Multipoint-to-point bandwidth aggregation cannot be effectively supported due to the required sender coordination
  - Interpreting and conveying the policy supplied by the user can be unwieldy
- Receiver-centric approaches
  - The receiver can easily control and coordinate transmissions from different senders for multipoint-to-point striping
  - The receiver can easily access the policy, and follow the policy that involves monitoring the dynamics of the wireless link

GNAN Research Group

R^2CP: Radial RCP

- RCP for multi-homed mobile hosts
  - A receiver-only extension which can communicate with one or multiple RCP senders
  - Multi-state transport protocol built atop RCP
  - Multipoint-to-point communication (as well as unicast)

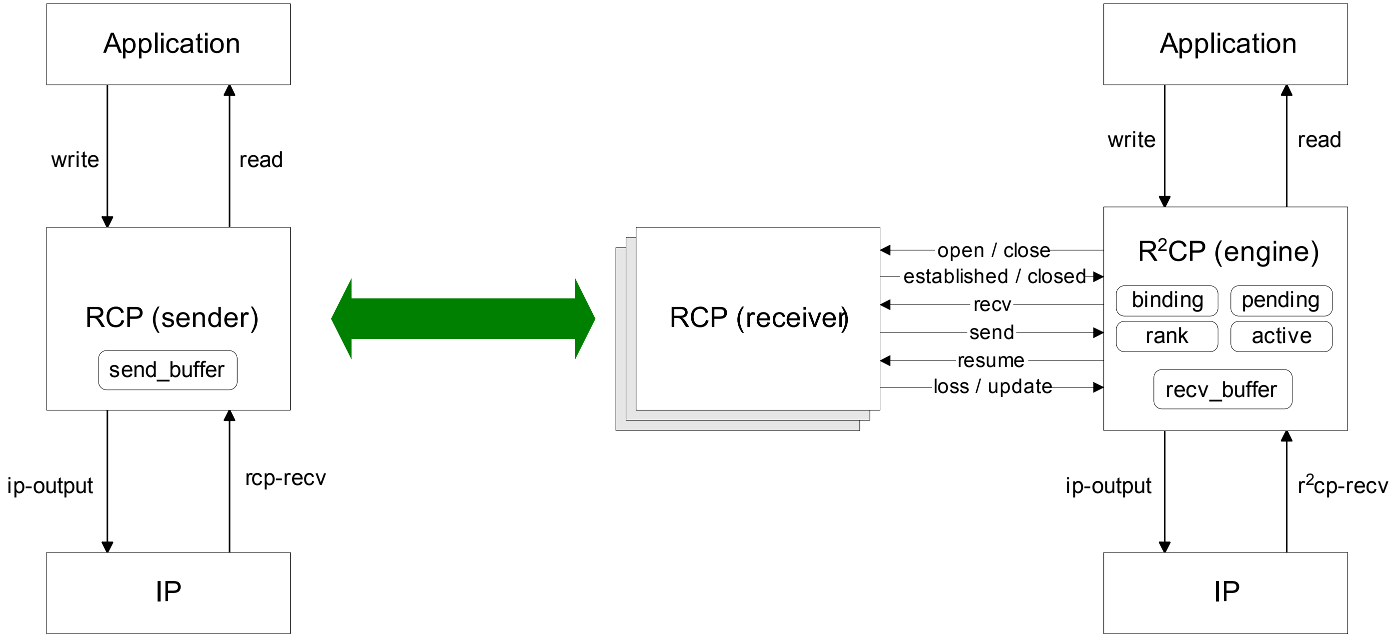

R²CP Design

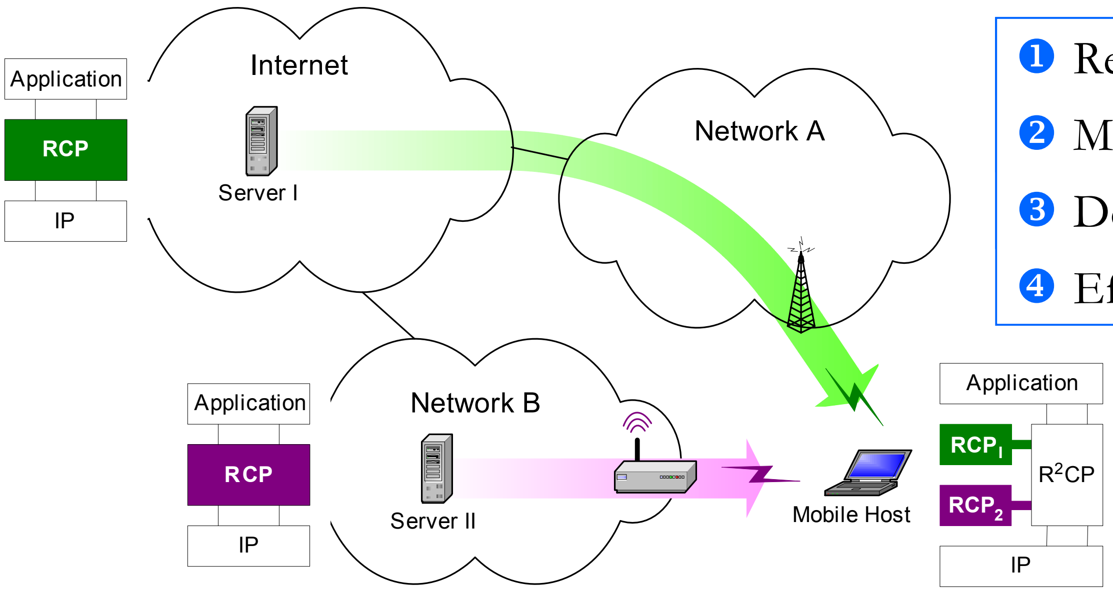

1. Receiver-centric operations
2. Maintaining multiple states
3. Decoupling of functionalities
4. Effective packet scheduling

- Effective packet scheduling
  - R²CP supports in-sequence data delivery
  - Effective packet scheduling avoids head-of-line blocking
  - R²CP schedules which data should be received through individual RCP pipes, while RCP controls when and how much data each pipe can request

R^2CP Packet Scheduling

- Avoid out-of-order arrival of data

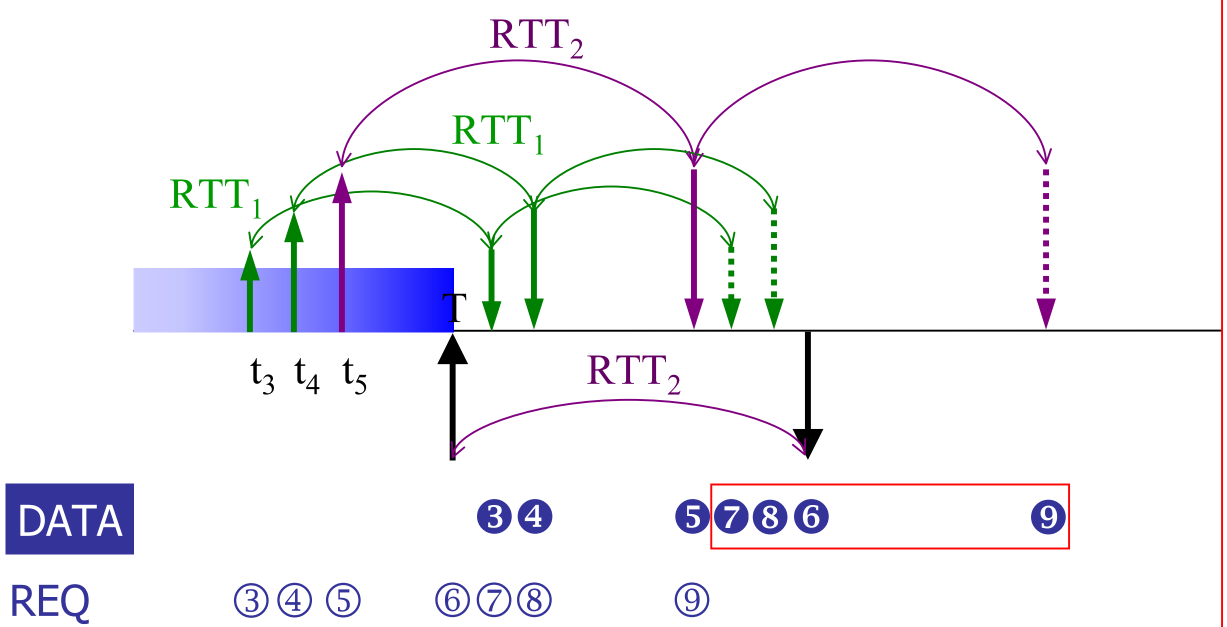

rank data structure

- Every pending REQ sent out by pipe j at time t has an entry with a timestamp of $t + 2*RTT_j$
- The rank of a new REQ issued by pipe k at time T is determined by comparing $T + RTT_k$ with entries in the **rank** data structure

The **binding** data structure maintains the mapping between the R^2CP sequence number and the RCP sequence number (along with the pipe identifier)

The **pending** data structure maintains ranges of data to be requested (includes retransmissions)

R²CP Performance

- ns-2 emulation
  - Live Internet traffic
  - Uncontrolled environment
- Testbed
  - 2 WiFi access points with different ESSIDs
  - 1 IBM laptop with two WiFi cards
  - 2 Dell desktops
  - Linux OS
  - libpcap
  - + tcpdump
  - + ns-2 trace

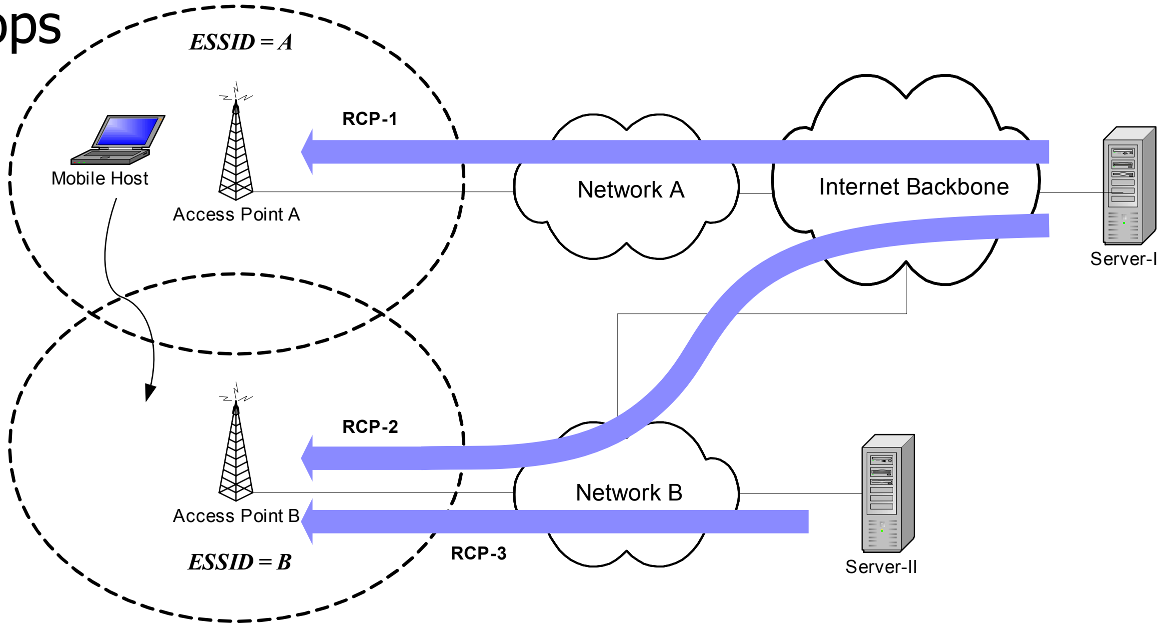

GNAN Research Group

Seamless Handoffs

- Scenario
  - The two interfaces are associated with different IP addresses in respective networks
  - The mobile host opens the RCP-2 pipe before closing the RCP-1 pipe

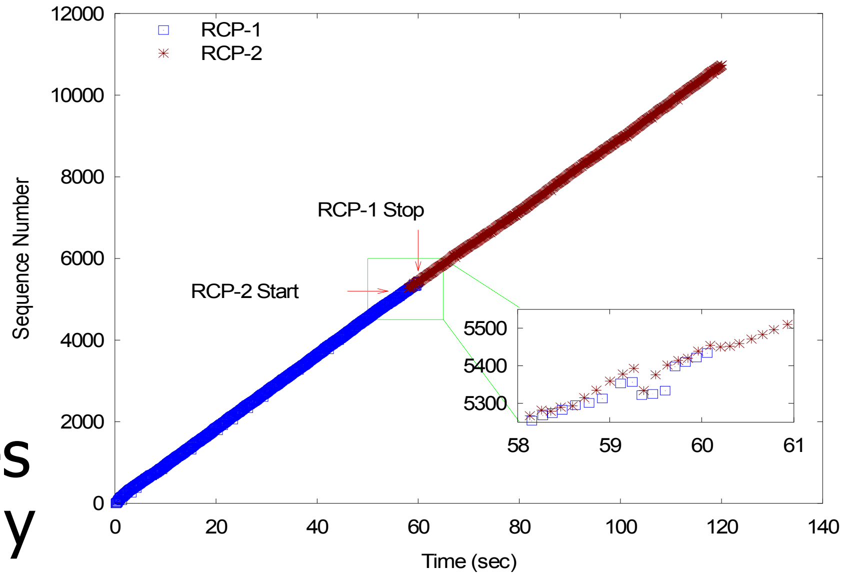

- The mobile host continues receiving data without any stall during handoffs

Server Migration

- Scenario
  - The mobile host connects to Server-II when it moves to network B
  - RCP-3 is opened between address B and Server-II
  - RCP-2 is between address B and Server-II for reference

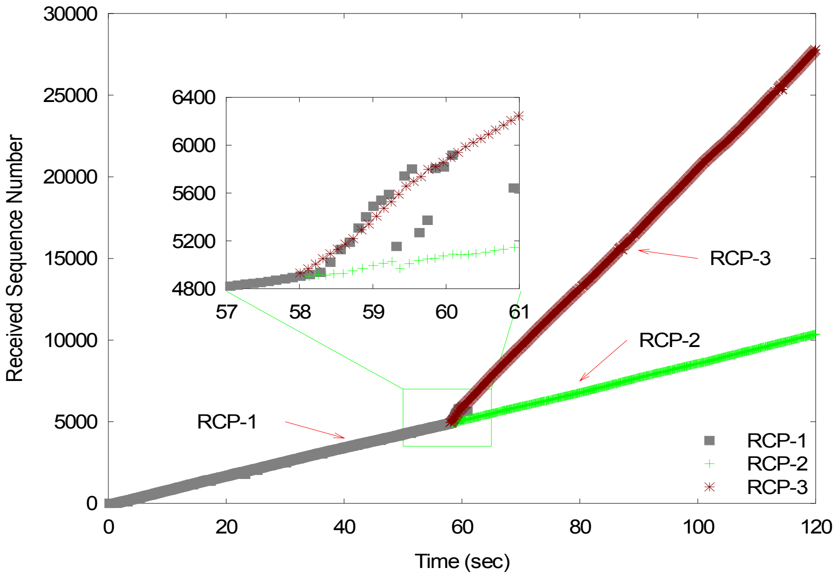

- The receiving application at the mobile host can be made unaware of the server migration

GNAN Research Group

Bandwidth Aggregation

- Scenario
  - Best-effort bandwidth aggregation
- Packet scheduling
  - Bandwidth mismatch
  - Delay mismatch
  - Buffer occupancy

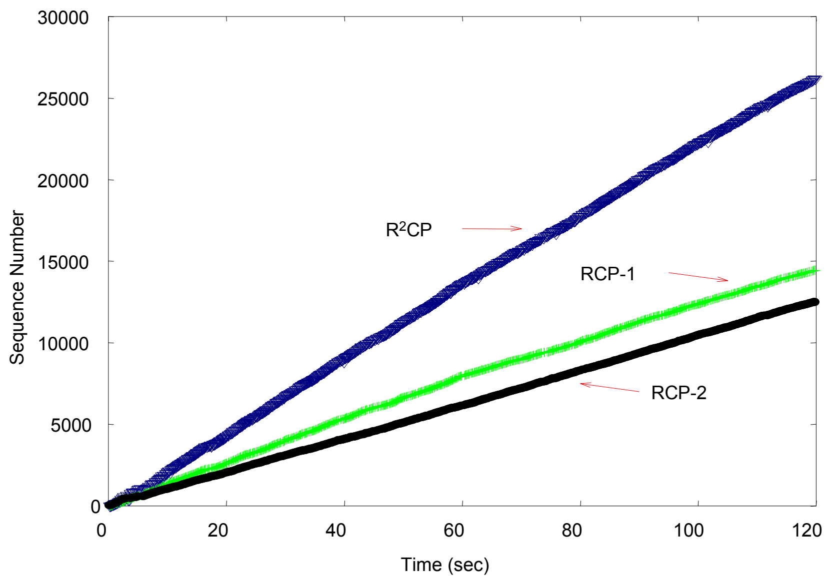

- R²CP can effective achieve bandwidth aggregation through its RTT-based packet scheduling

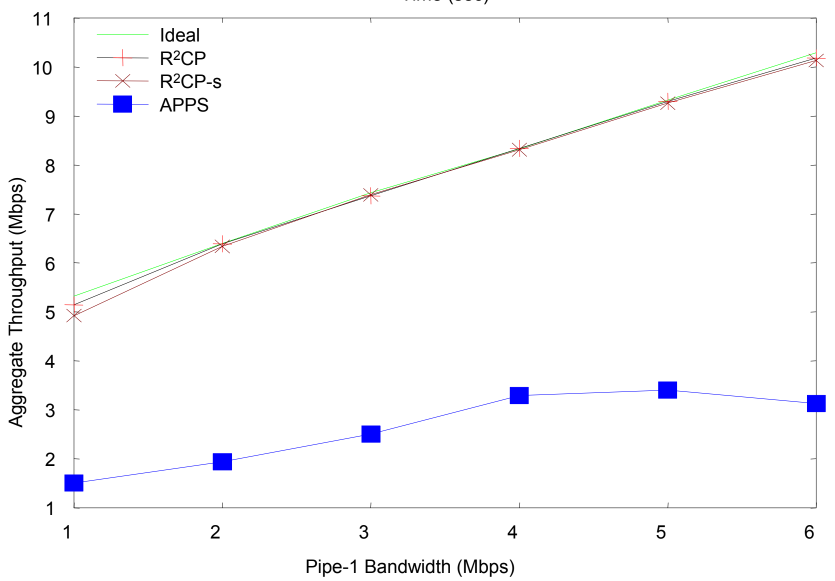

Bandwidth Aggregation

- Scenario
  - Best-effort bandwidth aggregation
- Packet scheduling
  - Bandwidth mismatch
  - Delay mismatch
  - Buffer occupancy

- R^2CP can effective achieve bandwidth aggregation through its RTT-based packet scheduling

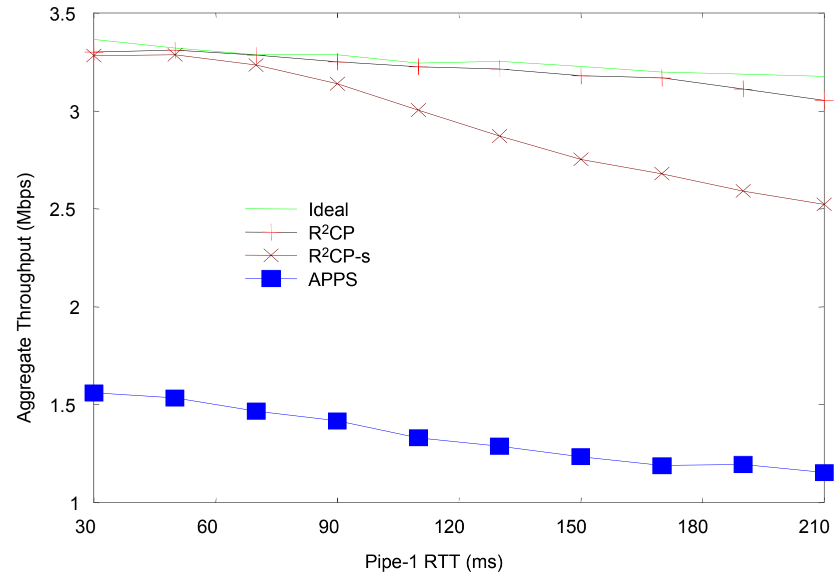

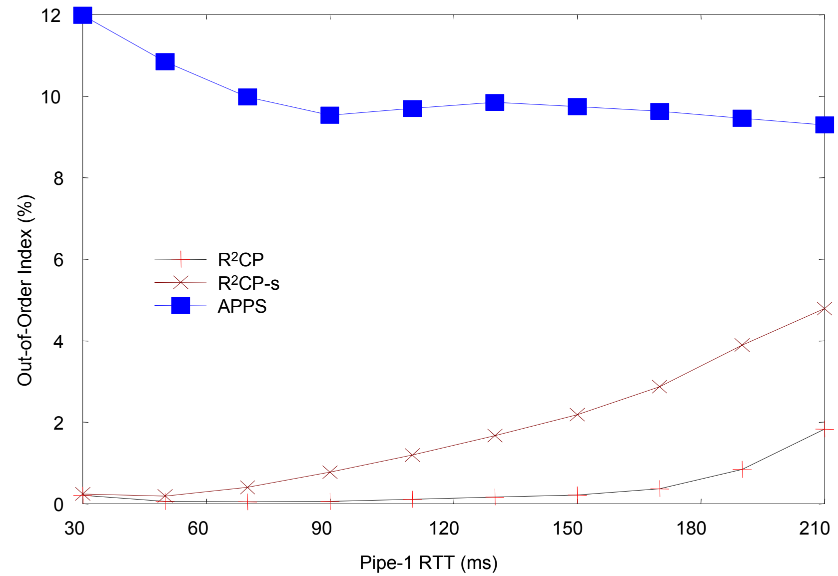

Related Work

- NETBLT [Clark 87]
  - The receiver maintains the retransmission timer
- WTCP [Sinha 99], TFRC [Handley 00], TCP-Real [Tsaoussidis 02]
  - Increased participation from the receiver, including calculating the send rate, maintaining the loss history, or tracking loss events
- WebTP [Gupta 00]
  - A receiver-centric transport protocol optimized for the Web application
  - It does not consider REQ losses, and is not TCP friendly
- Bandwidth management and sharing [Spring 00, Mehra 03]
  - RCP can allow more effective bandwidth management and sharing for mobile hosts over the wireless link than TCP

GNAN Research Group

Summary

- A receiver-centric transport protocol can achieve fundamental performance gains and functionality gains over a sender-centric one
- RCP is a TCP clone in its general behavior, but can achieve better performance in terms of loss recovery, congestion control, and power management in wireless environments
- R^2CP is a multi-state extension of RCP that enables seamless handoffs, server migration, and bandwidth aggregation for multi-homed mobile hosts with heterogeneous wireless interfaces
- Issues
  - Mobile host overheads
  - RCP extensions (upstream traffic/application-limited traffic)

GNAN Research Group

http://www.ece.gatech.edu/research/GNAN