On Transport Layer Support for Peer-to-Peer Networks

Hung-Yun Hsieh
and Raghupathy Sivakumar

GNAN Research Group
School of Electrical and Computer Engineering
Georgia Institute of Technology
February 26, 2004

Transport Layer Support

- Existing peer-to-peer systems use transport layer protocols designed for client-server networks
  - TCP is the predominant transport protocol used in peer-to-peer networks
- The unique characteristics of peer-to-peer networks render existing transport protocols inefficient
  - Multiple sources with replicated content
  - Sources with non-server-like behaviors
- ☞ Transport layer support for peer-to-peer networks
  - Existing transport protocols support only point-to-point (unicast) and/or point-to-multipoint (multicast) connections
  - We argue for designing transport protocols that can support multipoint-to-point connections (MP2P) for effective data transport in peer-to-peer networks

GNAN Research Group

Role of the Transport Layer

- The transport layer is the heart of the whole protocol hierarchy [Tanenbaum 96]
  - Provides effective data transport from the source to the destination
  - Translates the service provided by the network layer for use by the application, independent of the characteristics of the network in use
- We do not differentiate between the session layer and the transport layer in this talk
  - Arguments and solutions can apply to the session layer, provided there is sufficient support from the transport layer
- ☞ Application layer vs. transport layer
  - Application layer striping / parallel downloads
  - Maintains multiple (unicast) connections without requiring any change at the transport layer

MP2P vis-à-vis Destination

- Better tackling of the artifacts in peer-to-peer networks
  - Limited capacity and transient availability at the source
- Limited capacity of the source
  - Peers are typically located at the edge of the Internet with asymmetric links
  - Difficulties involved in performing peer selection
- Transient availability of the source
  - Peer departures and/or link failures
  - Size of the content being shared
- ☞ Application layer striping
  - Resequencing requirement: buffer & delay
  - Communication overheads

GNAN Research Group

MP2P vis-à-vis Source

- Better utilization of the resources in peer-to-peer networks
  - Hosts with heterogeneous resources
- Low-profile hosts
  - To use or not to use
  - Not to become the bottleneck of the network
- High-profile hosts
  - Disincentives to share due to resource contention
  - Not to become the hotspot of the network
- ☞ Application layer striping
  - Data partitioning: content coding or source coordination
  - Content coding: efficiency and feasibility
  - Source coordination: granularity vs. overheads

GNAN Research Group

MP2P vis-à-vis Content

- Better preservation of the content in peer-to-peer networks
  - The integrity / fidelity of the content as it is propagated
- Content replication
  - Reliability vs. timeliness
  - Requesting peers later become content suppliers
  - Out-of-band loss recovery
  - Reliable delivery can be achieved without impacting the timely delivery of the data from the (primary) source
- ☞ Application layer striping
  - Repetitive implementations of transport layer functionalities (loss detection and loss recovery)
  - Inefficient implementations

GNAN Research Group

Transport Protocol Design (1)

- A multipoint-to-point transport protocol should support the API that a point-to-point one supports
  - In-sequence data delivery semantics
- Multiple states
  - TCB state: congestion window, round-trip time, timers, etc
  - Peer heterogeneity
  - Packet reordering due to multiple paths
- Decoupling of functionalities
  - A multipoint-to-point connection with only one source should incur the same overheads as a point-to-point one
  - Per-connection vs. per-path functionalities
  - Buffer and socket management: per connection
  - Congestion control: per path
  - Reliability: per-connection or per-path?

GNAN Research Group

Transport Protocol Design (2)

- Packet scheduling
  - Packets from multiple sources should be received in-sequence to avoid head-of-line blocking and/or minimize resequencing delay
  - Peer heterogeneity (bandwidth/latency mismatches and fluctuations) and peer dynamics (departures and/or arrivals)
  - Scheduling granularity vs. synchronization overheads
- Receiver-centric operations
  - The (data) receiver is the common point of all paths and the invariant of the connection
  - Packet scheduling: to minimize communication overheads (synchronization) and the impact of peer transience
  - Congestion control and reliability

GNAN Research Group

Case Study: R^{2}CP

- R^{2}CP: Radial reception control protocol
  - Proposed in [MOBICOM 03] for multi-homed mobile hosts with heterogeneous wireless interfaces
  - R^{2}CP supports the API that TCP supports
  - R^{2}CP with one source = RCP = TCP in terms of congestion control and reliability semantics

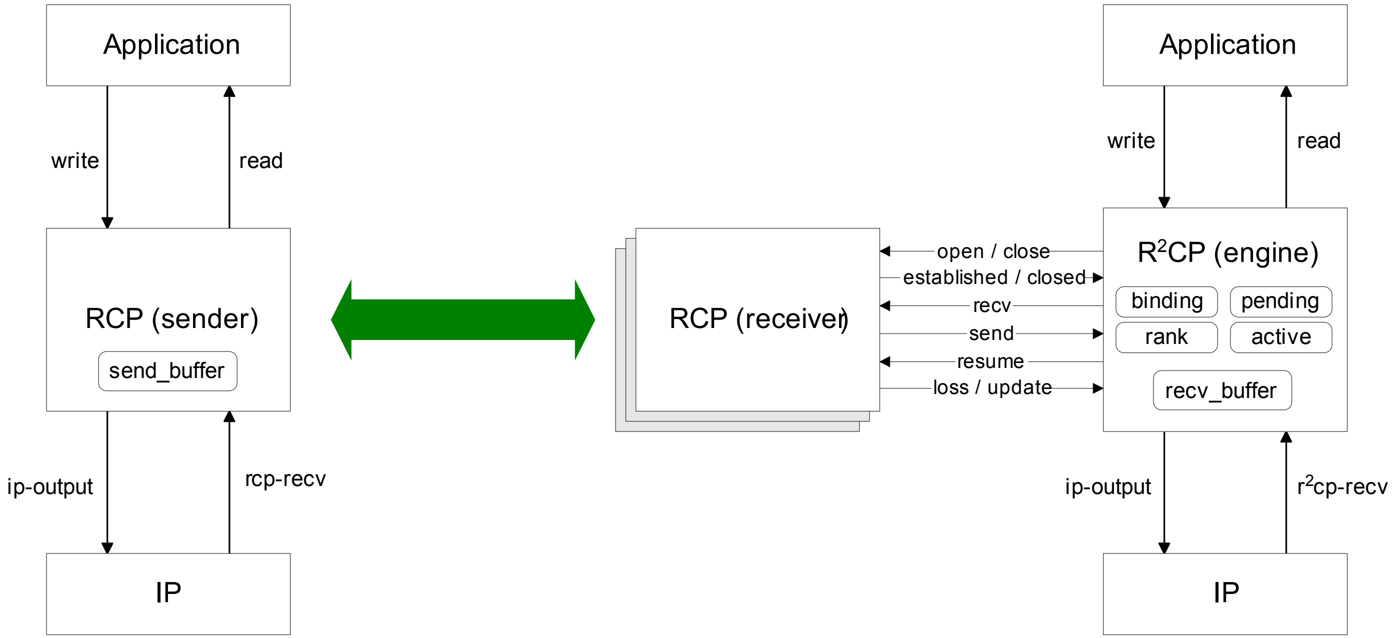

R^2CP Synopsis

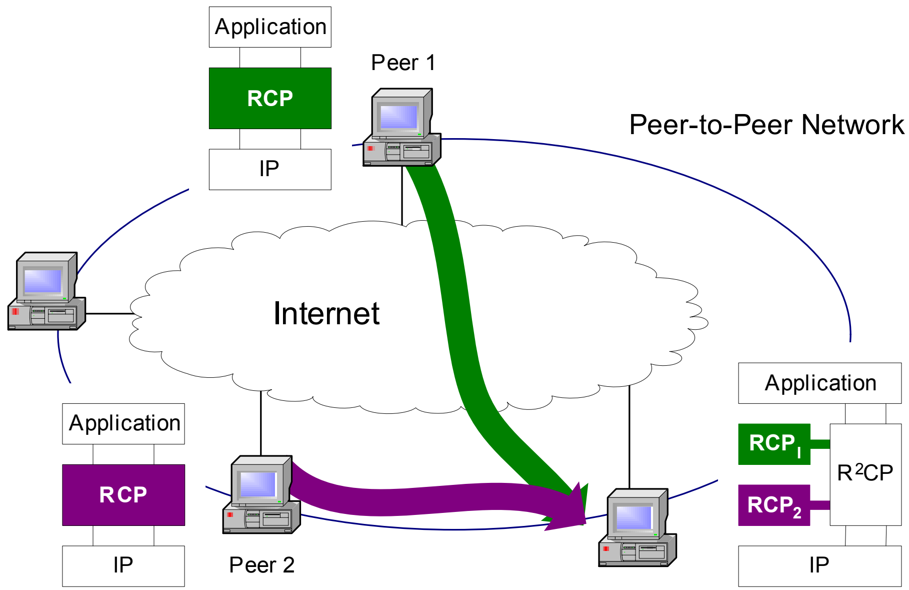

❶ Maintaining multiple states
❷ Decoupling of functionalities
❸ Effective packet scheduling
❹ Receiver-centric operations

- Receiver-centric operation
  - The receiver drives the operation of congestion control and reliability for individual paths
  - The receiver is responsible for packet scheduling from multiple sources
  - ☞ when, which, how much to request from each source

GNAN Research Group

Future Work and Summary

- Future work
  - Peer selection and load balancing
  - Peer departures and arrivals
  - Self-interest peers and untrusting peers
  - Fairness model
  - Deployment issues

- Summary
  - The transport layer plays a defining role in effective data transport between the source and the destination
  - A transport layer protocol with multipoint-to-point support can be a powerful building block for peer-to-peer applications, by masking the artifacts of the underlying network and allowing significant performance gains

GNAN Research Group