# *Block-switched Networks: A New Paradigm for Wireless Transport*

**Ming Li, Devesh Agrawal, Deepak Ganesan, and Arun Venkataramani**

***University of Massachusetts Amherst***

# What You Buy vs. What You Get

- **TCP performs poorly over wireless links**

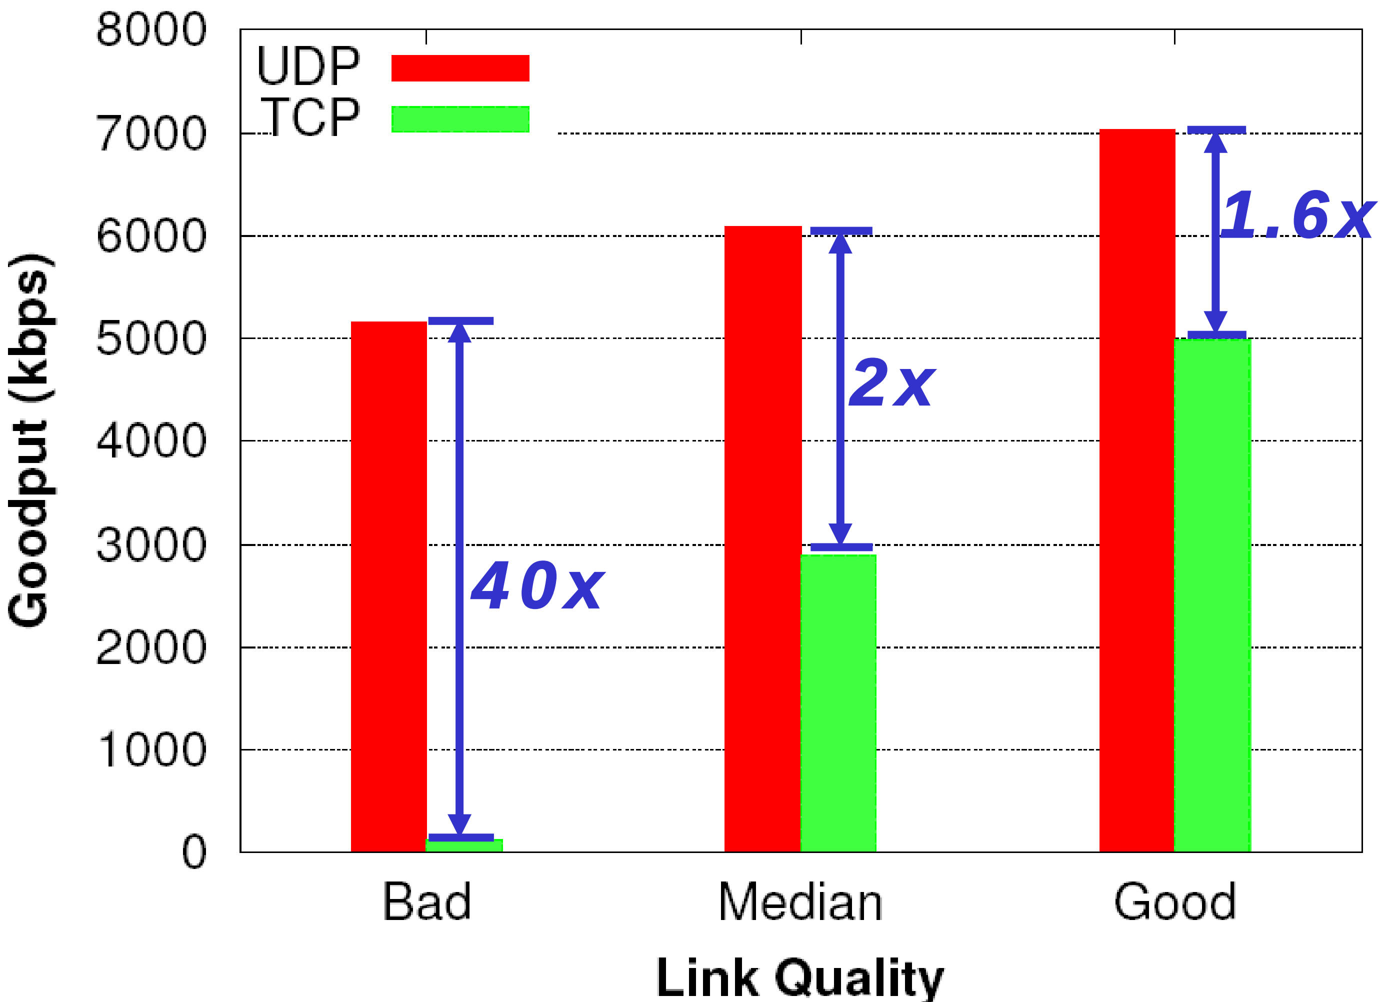


**Advertised capacity: 11Mbps**

# 1. E2E Transport

- ***E2E rate control is error-prone***
- *E2E retransmissions are wasteful*
- *E2E route disruptions cause unavailability*

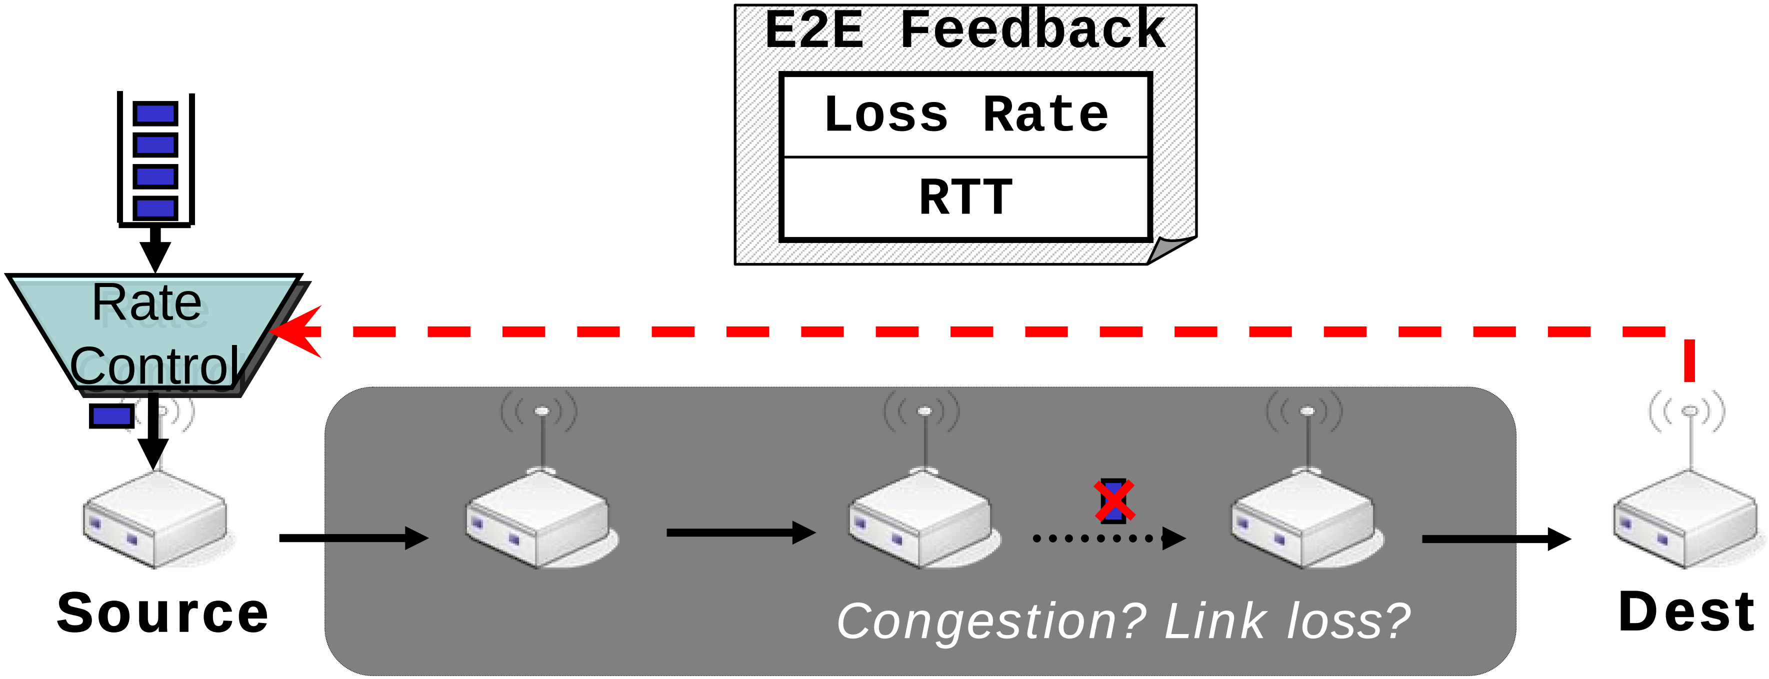

# 1. E2E Transport

- *E2E rate control is error-prone*
- ***E2E retransmissions are wasteful***
- *E2E route disruptions cause unavailability*

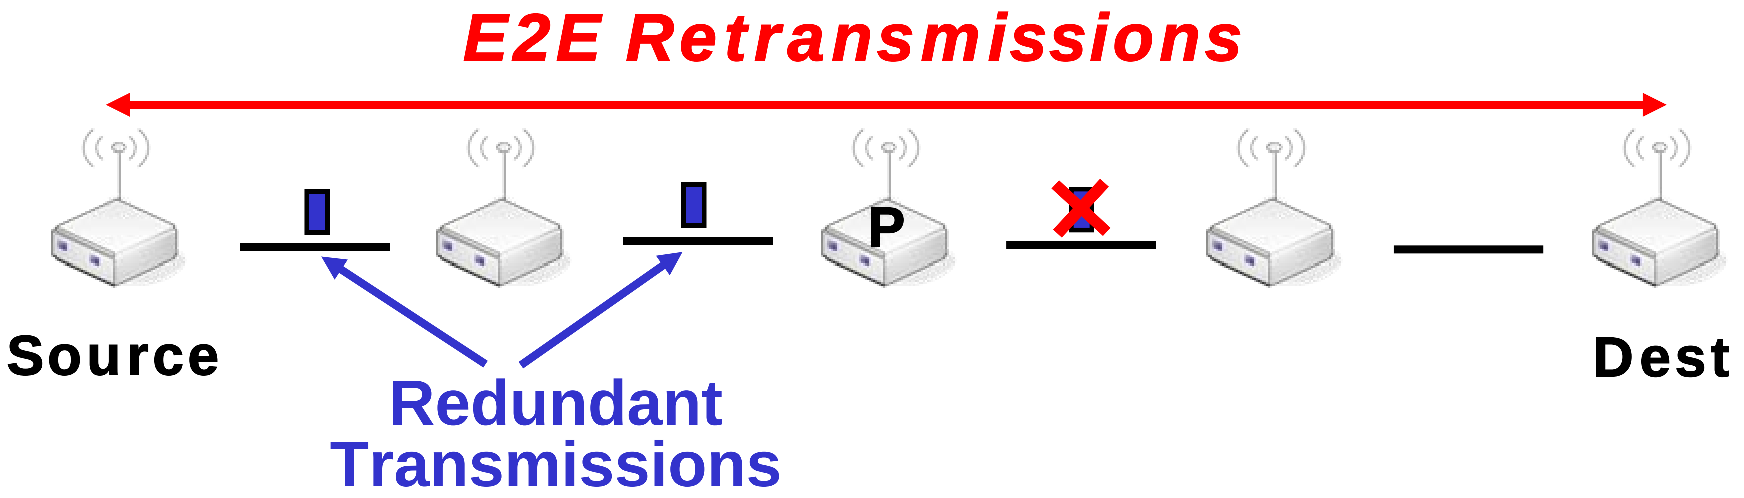

# 1. E2E Transport

- *E2E rate control* is error-prone
- *E2E retransmissions* are wasteful
- ***E2E route disruptions* cause unavailability**

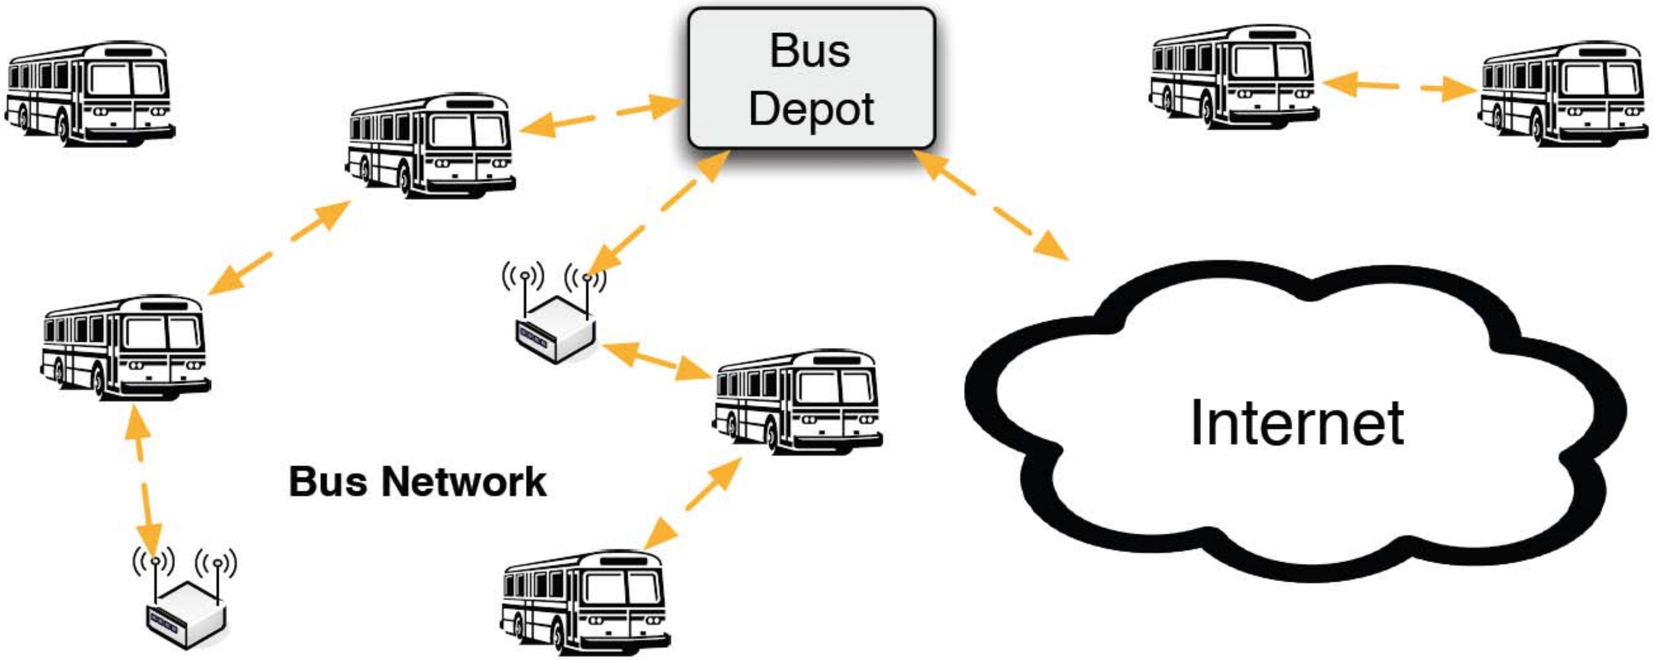

# 2. Packet as Unit of Control

- **Channel access**
- Link layer ARQ

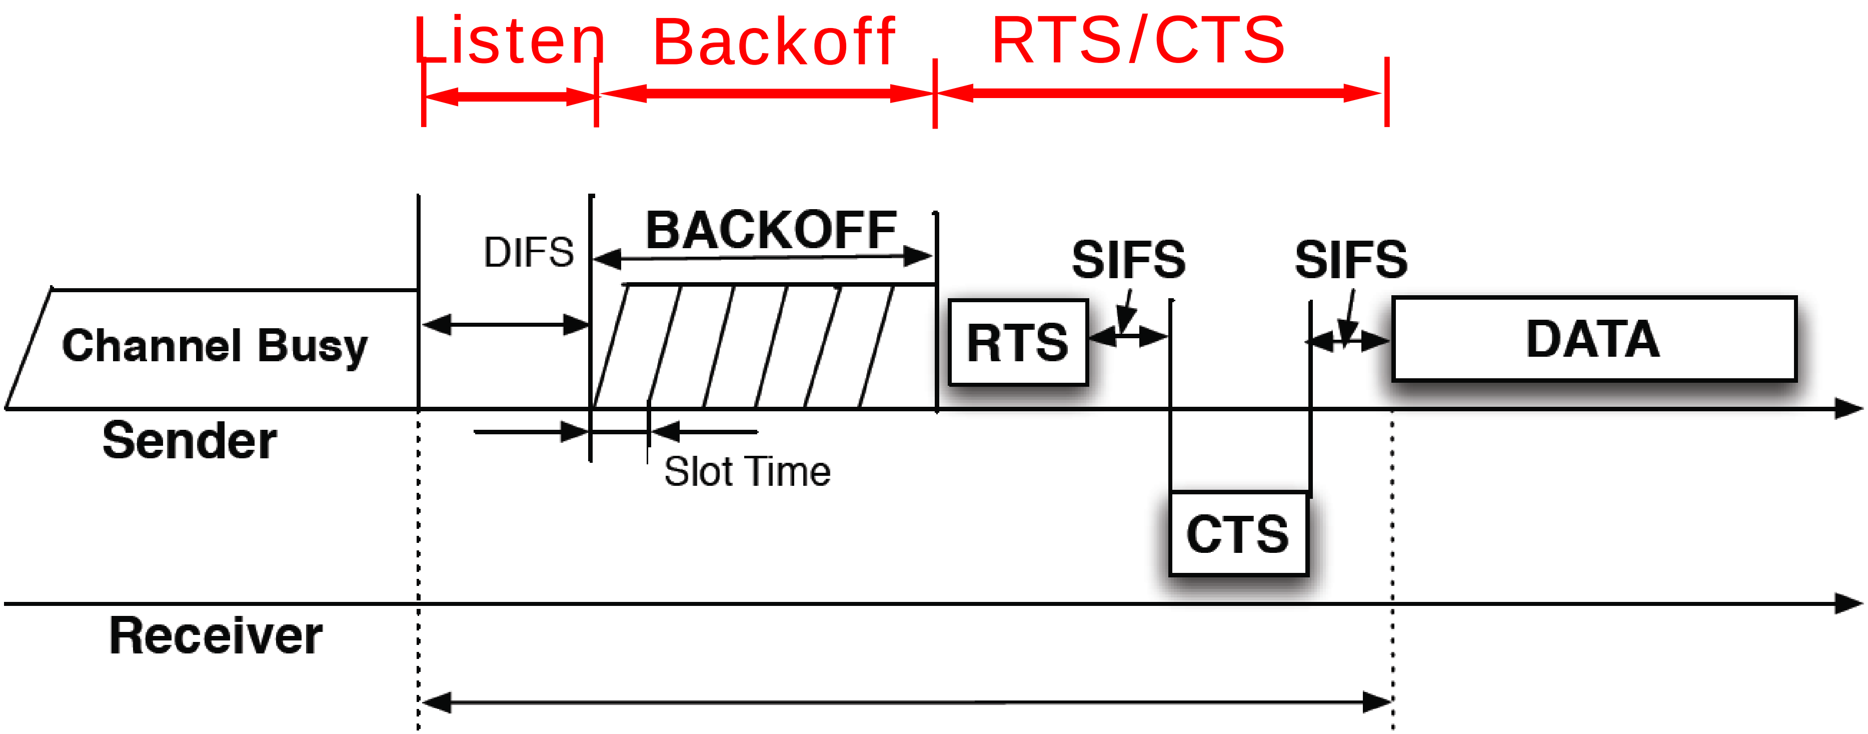


Channel Access Overhead for Each Frame

# 2. Packet as Unit of Control

- Channel access
- **Link layer ARQ**

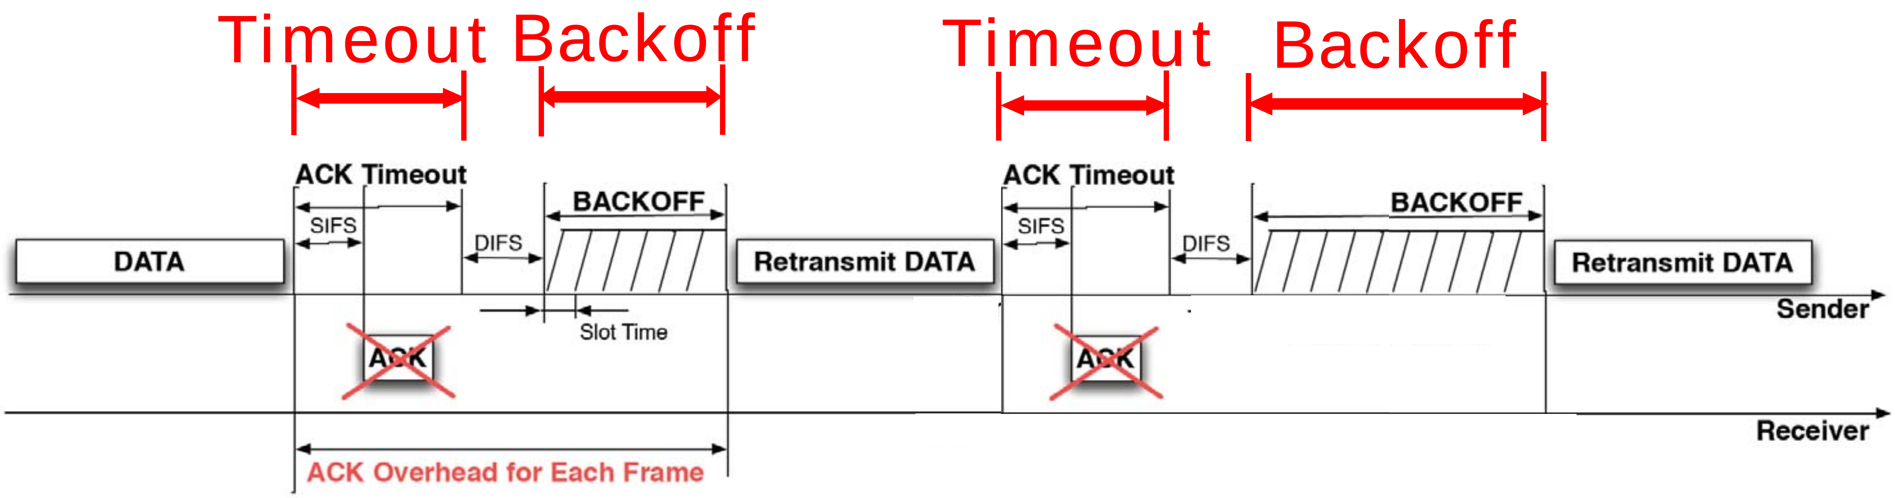

# 3. Complex Cross-Layer Interaction

- **Link-layer ARQs/backoffs hurt TCP rate control**

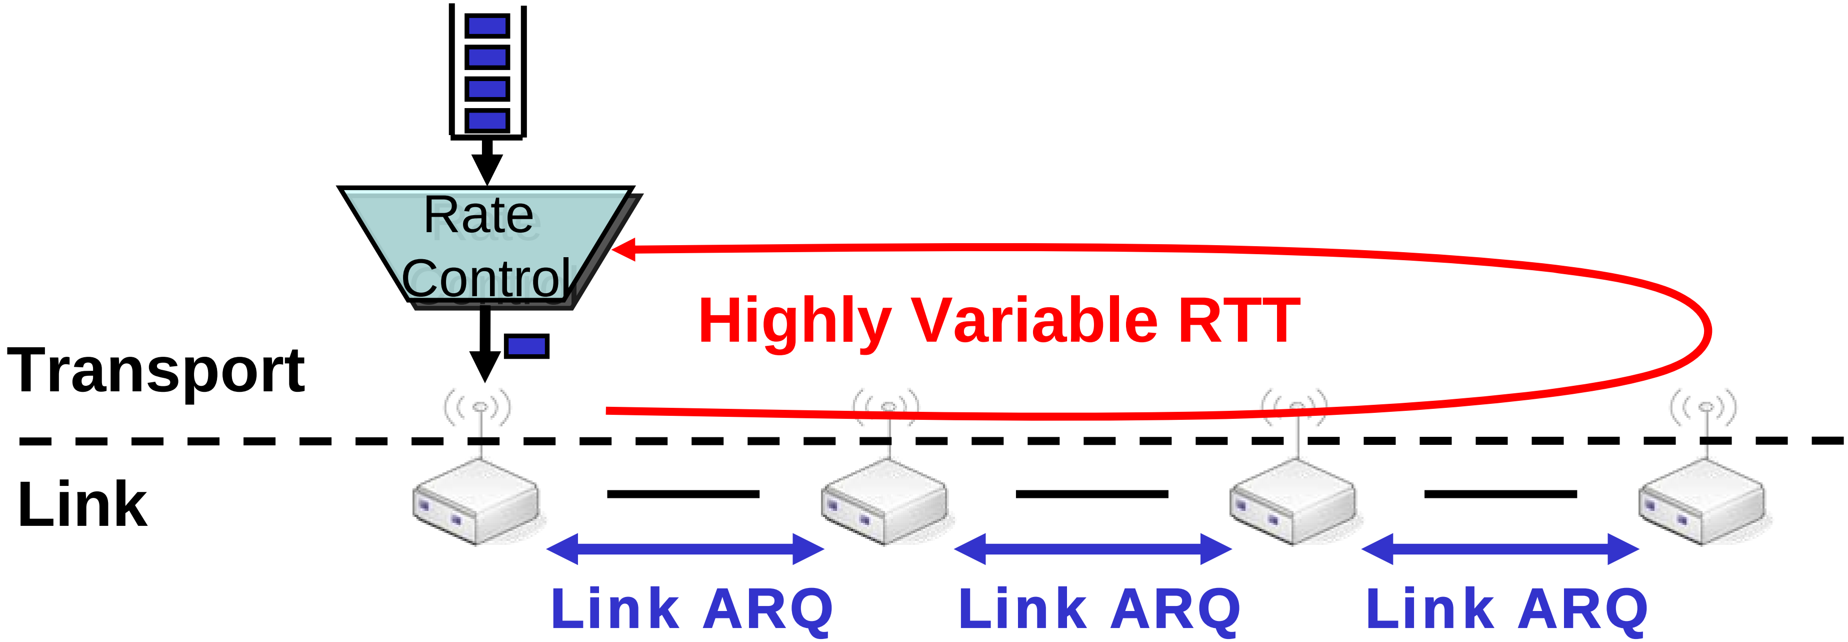

# *Hop*: A Clean Slate Re-design

- End-To-End ➡ Hop-by-Hop
- Packets ➡ Blocks
- Complexity ➡ Minimalism

# *Hop* Design

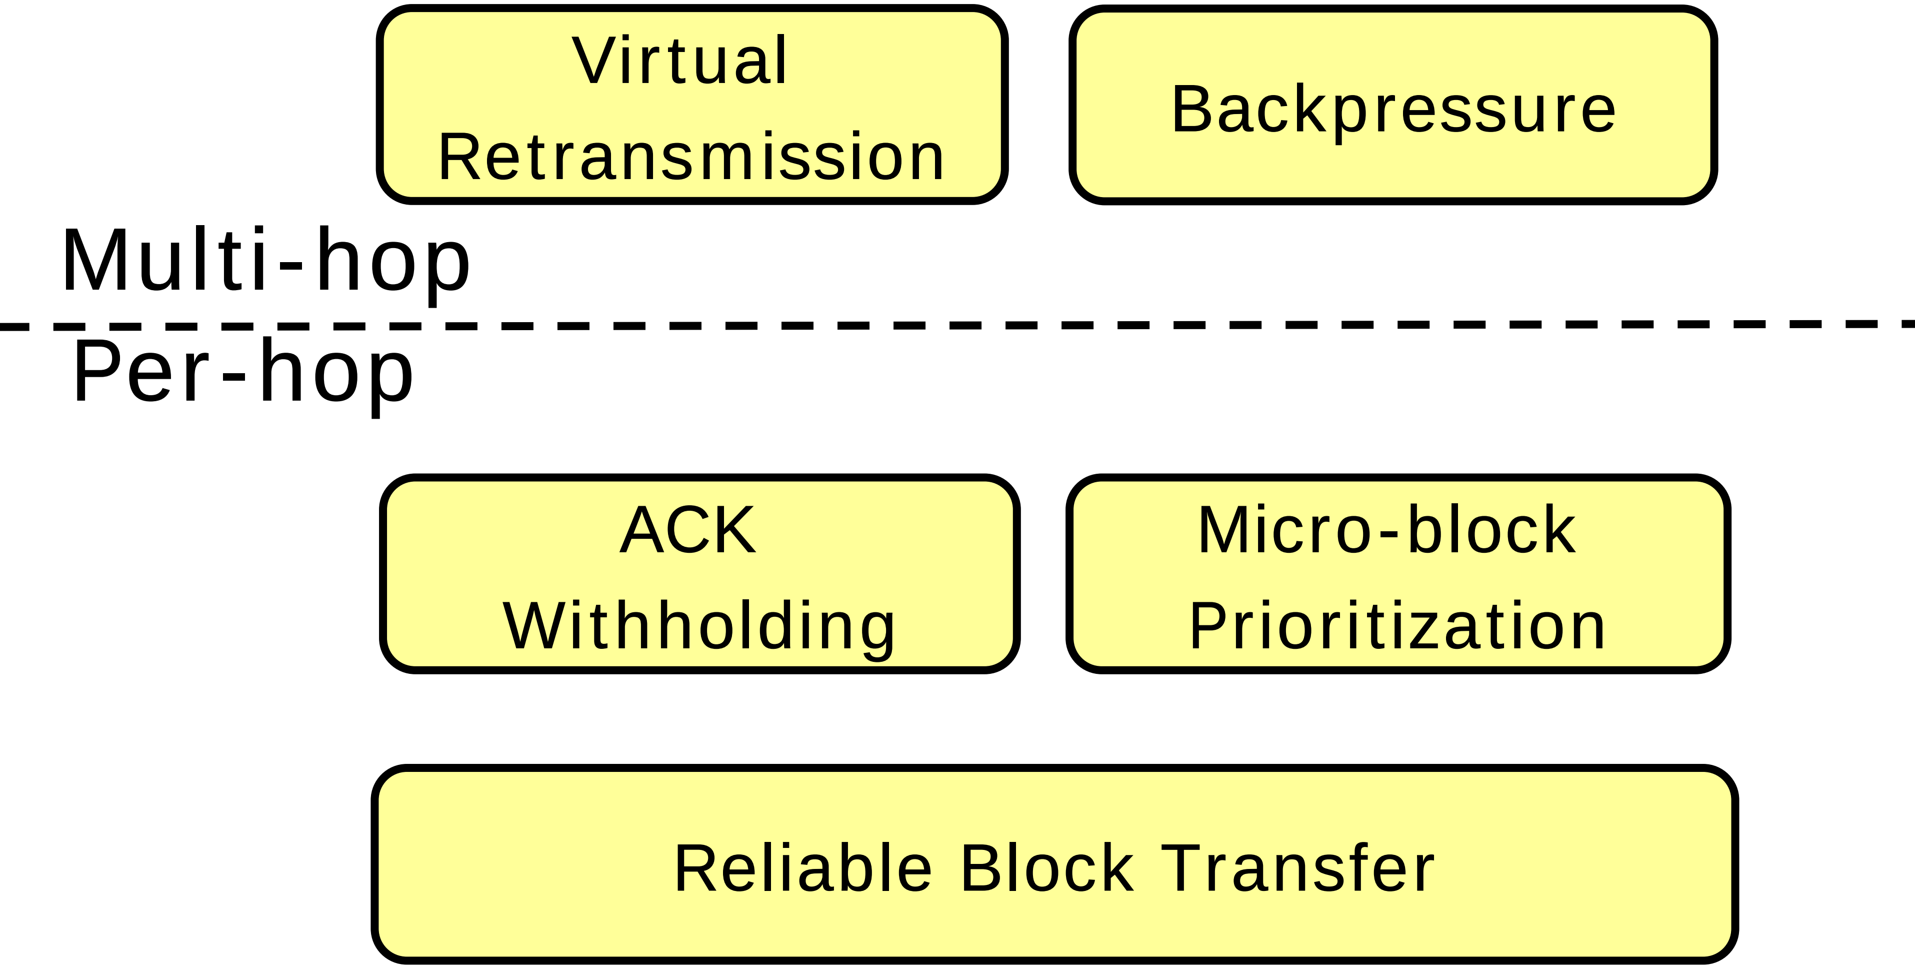
Virtual Retransmission
Backpressure
Multi-hop
Per-hop
ACK Withholding
Micro-block Prioritization
Reliable Block Transfer

# Reliable Per-Hop Block Transfer

- **Mechanisms**
  - Burst mode (TXOP)
  - Block ACK based ARQ
- **Benefits**

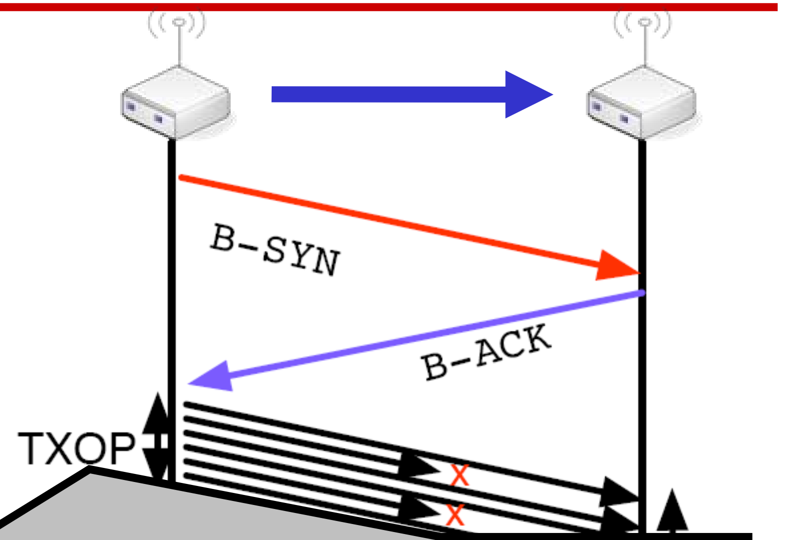


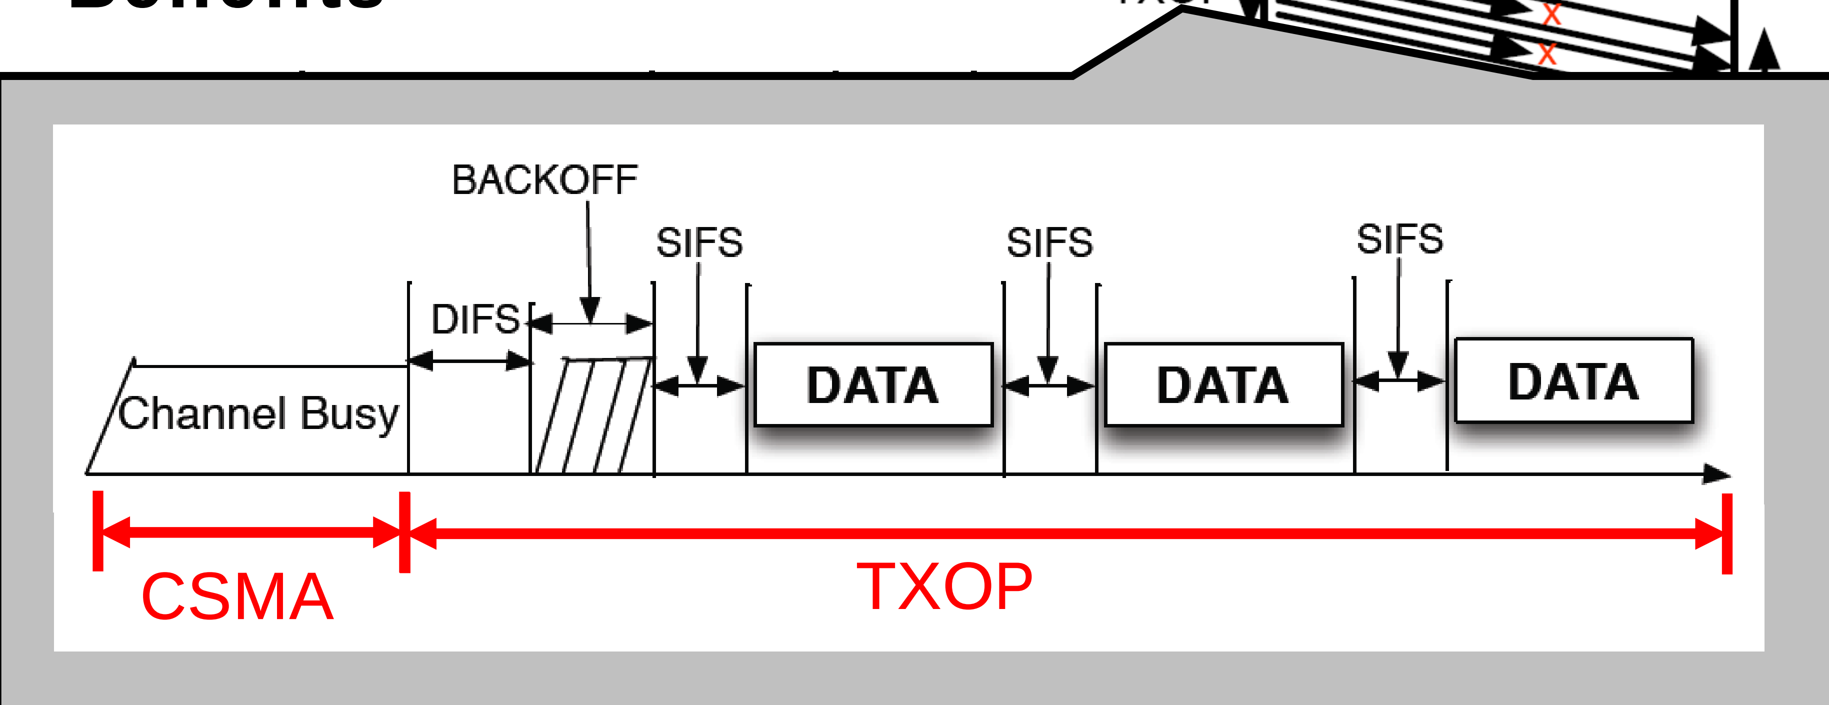

# *Hop* Design

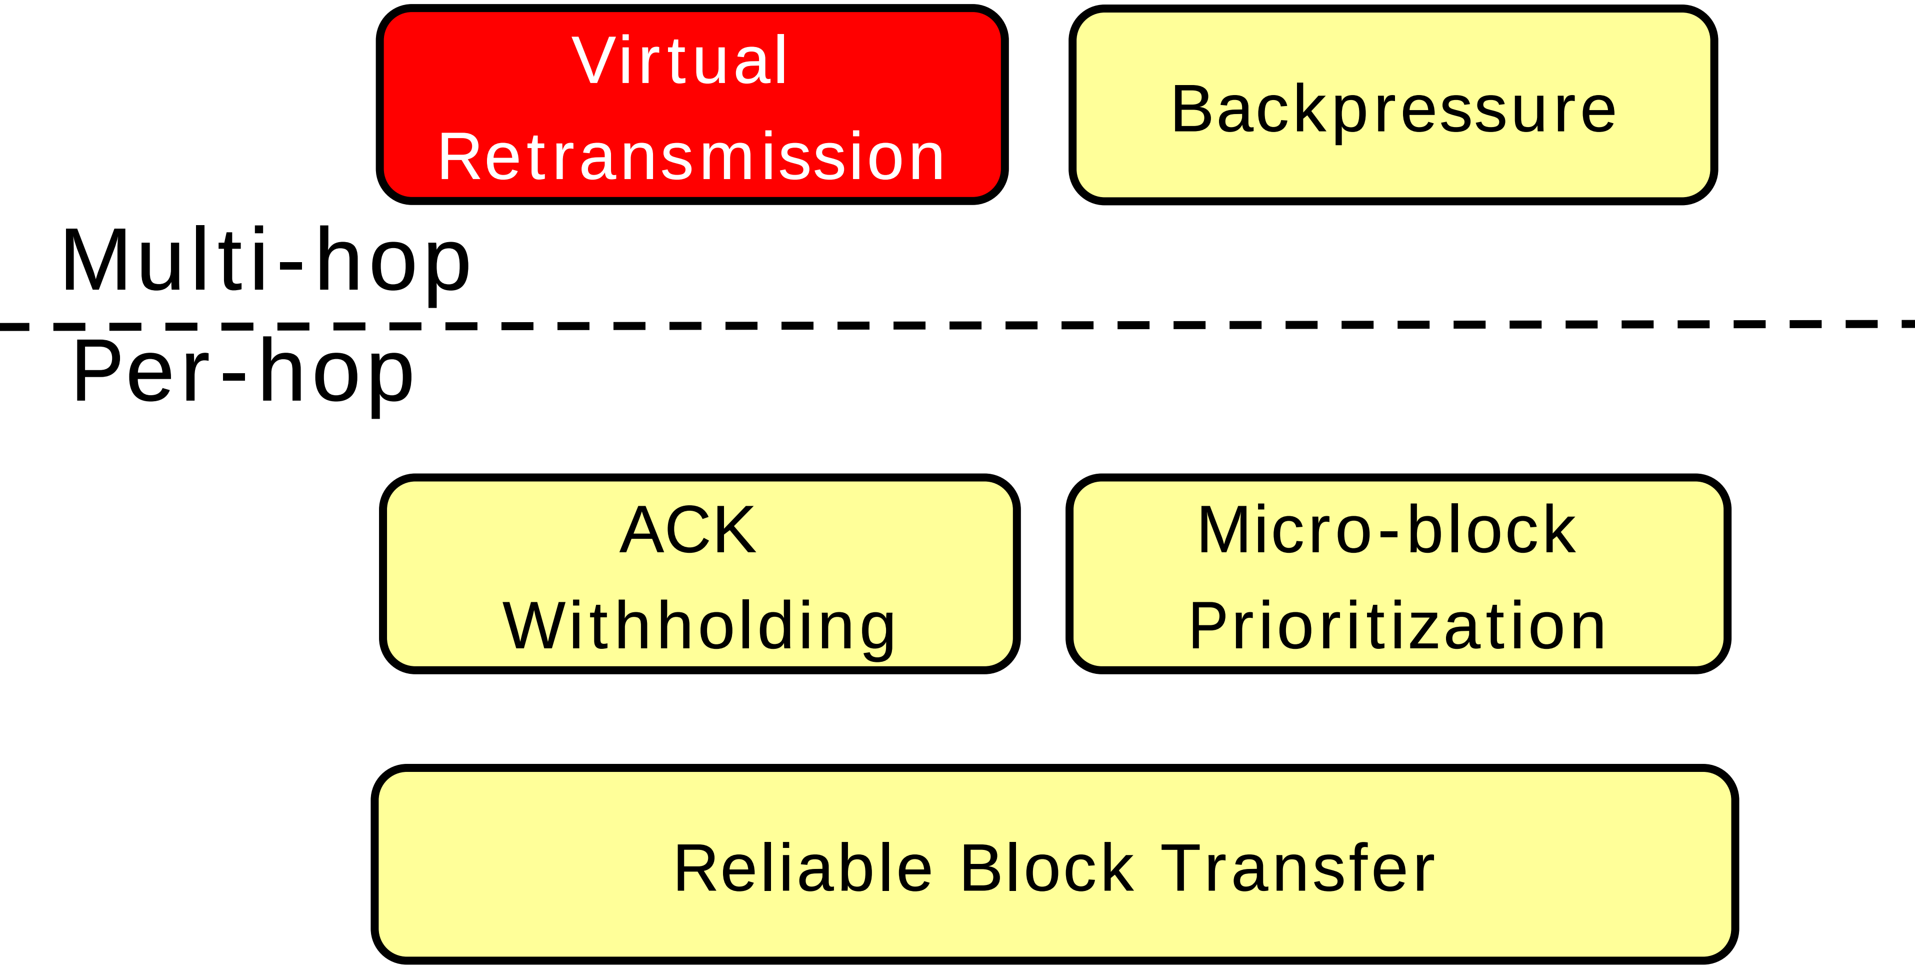
Virtual Retransmission
Backpressure
Multi-hop
Per-hop
ACK Withholding
Micro-block Prioritization
Reliable Block Transfer

# Virtual Retransmission (VTX)

- **Mechanism**
  - Leverages in-network caching
  - Re-transmits blocks only when unavailable in cache
- **Benefits**
  - Fewer transmissions
  - Low overhead
  - Simple

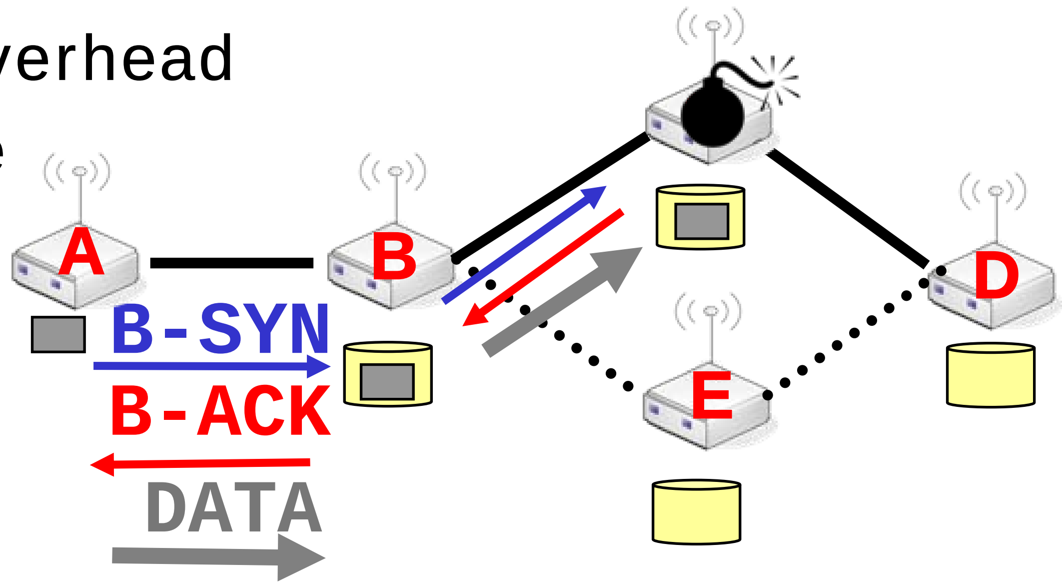

# Virtual Retransmission (VTX)

- **Mechanism**
  - Leverages in-network caching
  - Re-transmits blocks only when unavailable in cache
- **Benefits**
  - Fewer transmissions
  - Low overhead
  - Simple

| | |
|---|---|
| **B-SYN$_{VTX}$** | **B-SYN** with VTX flag set |

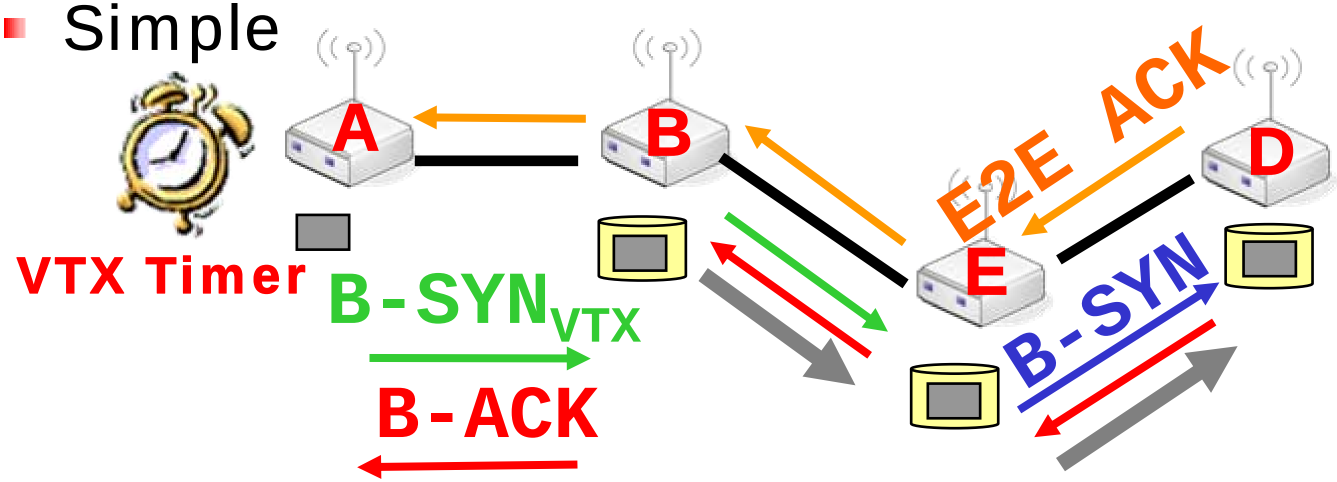

# *Hop* Design

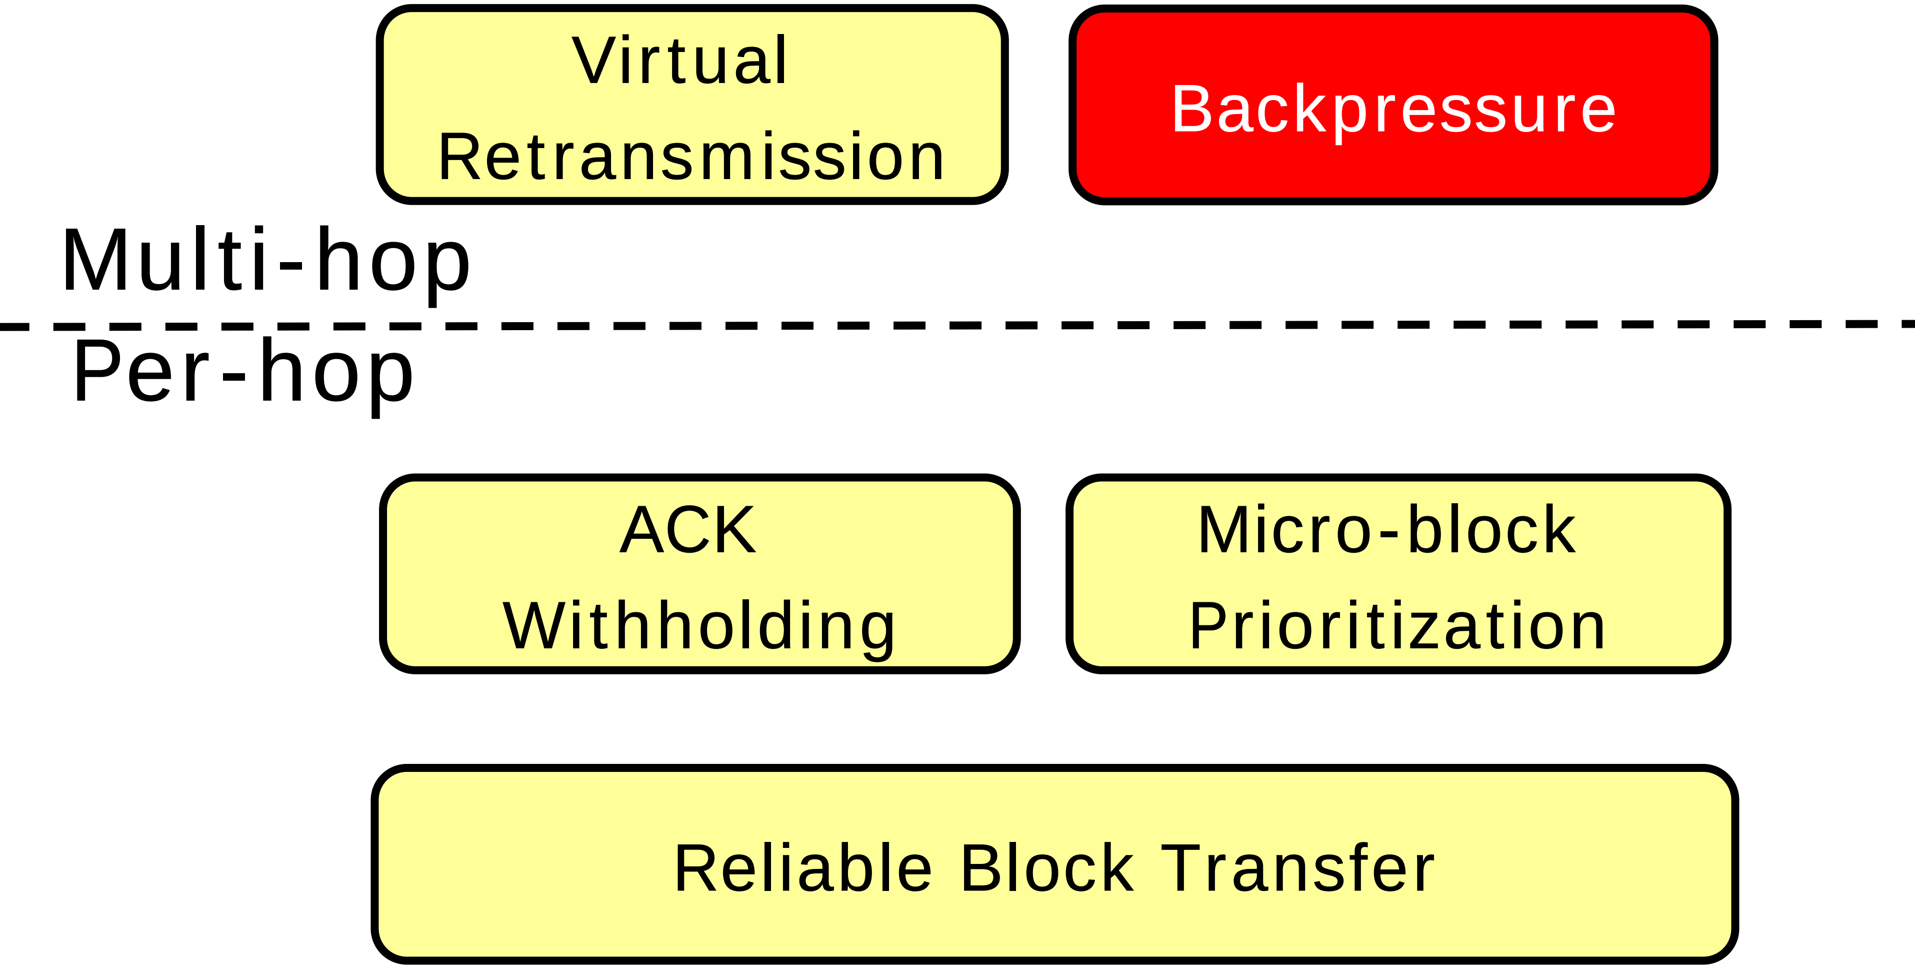
Virtual Retransmission
Backpressure
Multi-hop
Per-hop
ACK Withholding
Micro-block Prioritization
Reliable Block Transfer

# Backpressure

- **Mechanism**
  - Limits #outstanding_blocks per-flow at forwarder

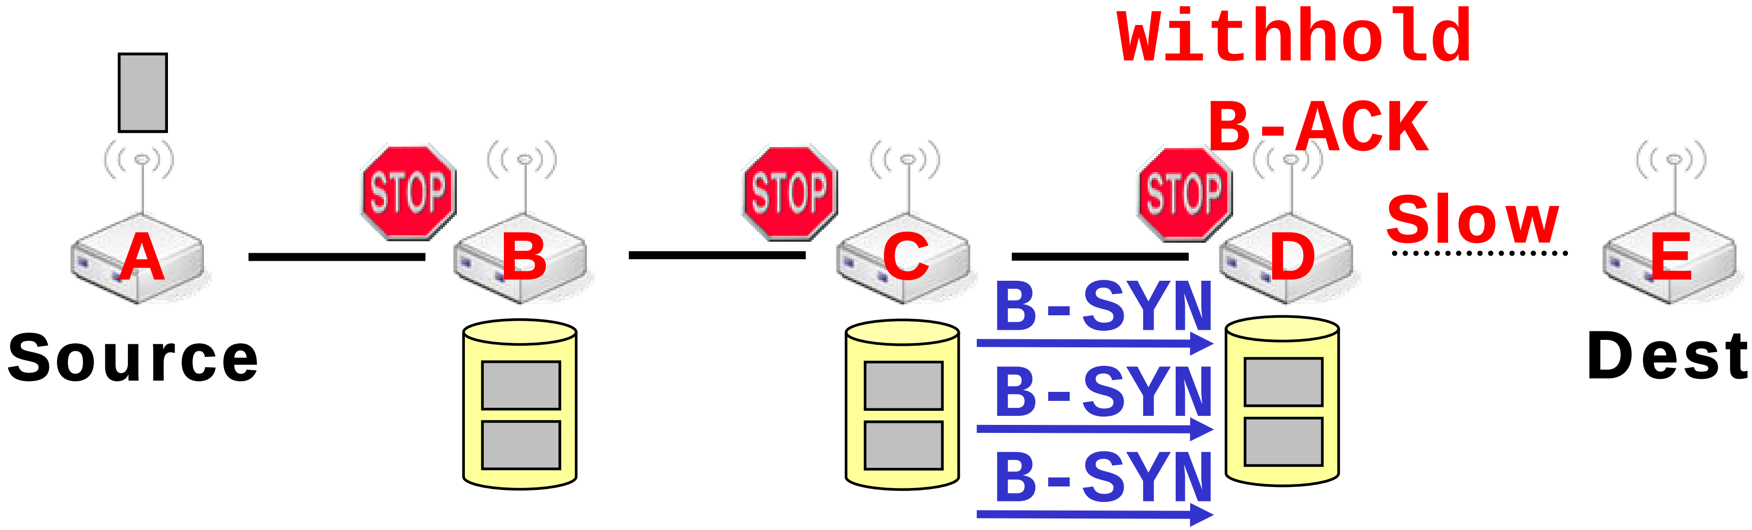


Limit of Outstanding Blocks=2

# Backpressure

- **Mechanism**
  - Limits #outstanding blocks per-flow at forwarder
- **Benefits**
  - Improves network utilization

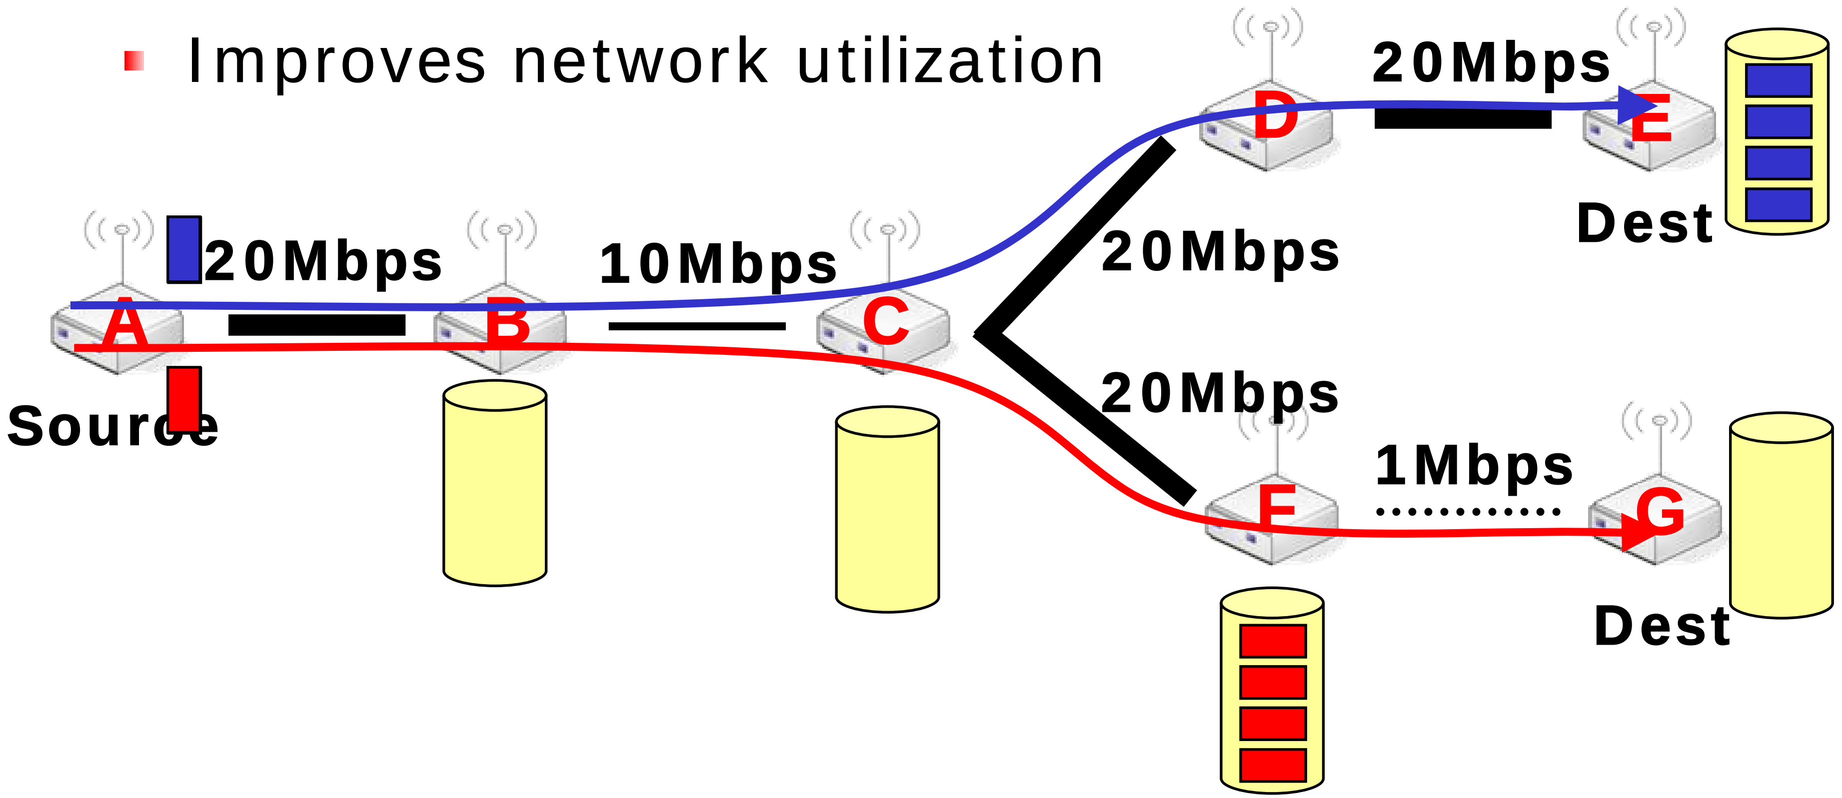


Aggregate goodput without backpressure: ***6Mbps***

# Backpressure

- **Mechanism**
  - Limits #outstanding blocks per-flow at forwarder
- **Benefits**
  - Improves network utilization

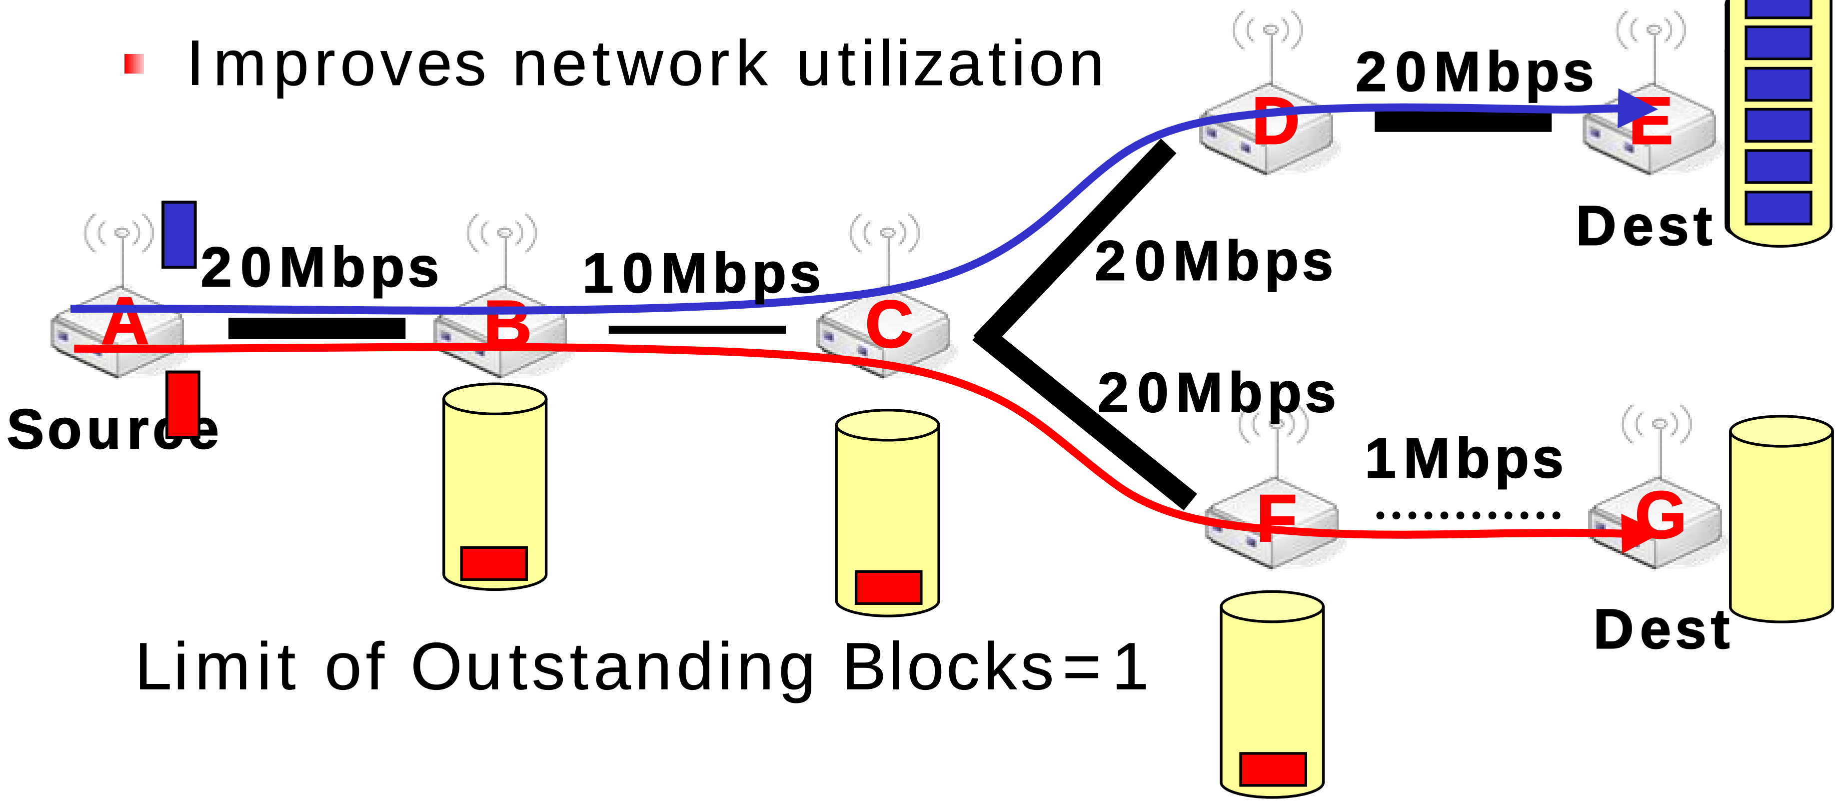


Limit of Outstanding Blocks=1

Aggregate goodput with backpressure: ***10Mbps***

# *Hop* Design

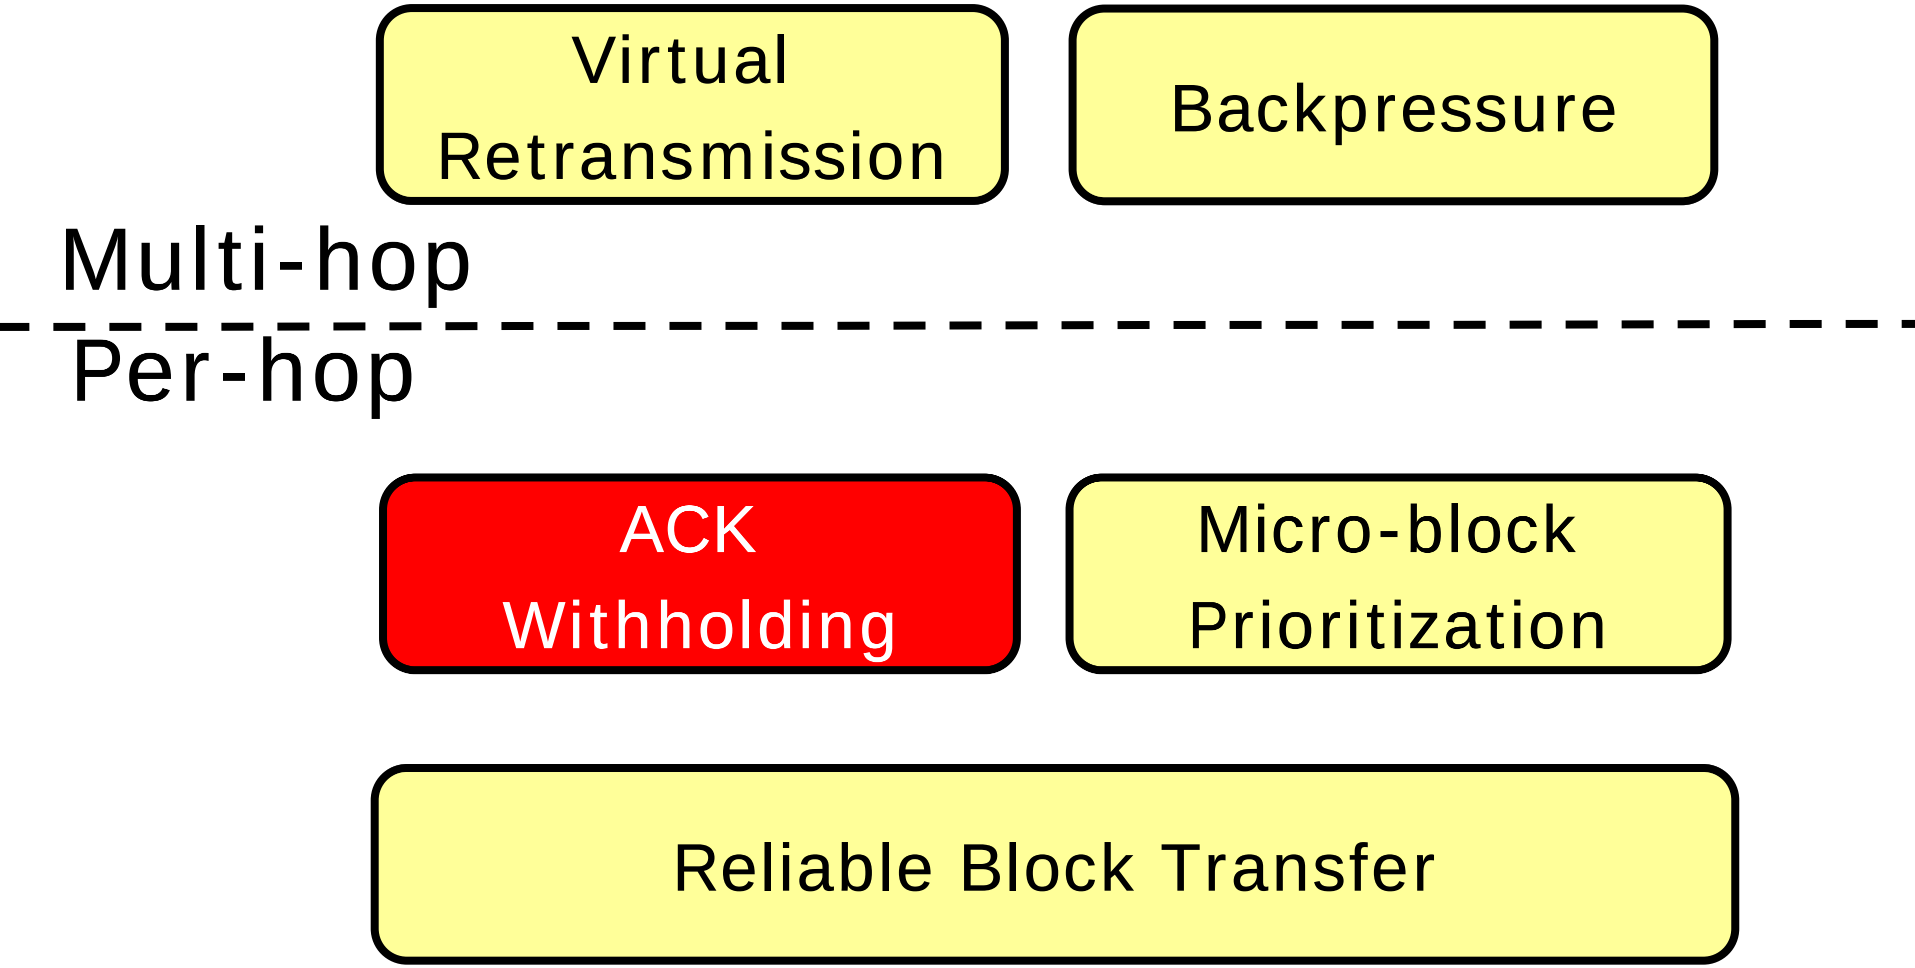


***RTS/CTS*** is overly conservative and incurs high overhead.

# Ack Withholding

- **Mechanism:**
  - Receiver withholds all but one **B-ACK**
- **Benefit:**
  - Low overhead
  - Less conservative
  - Simple

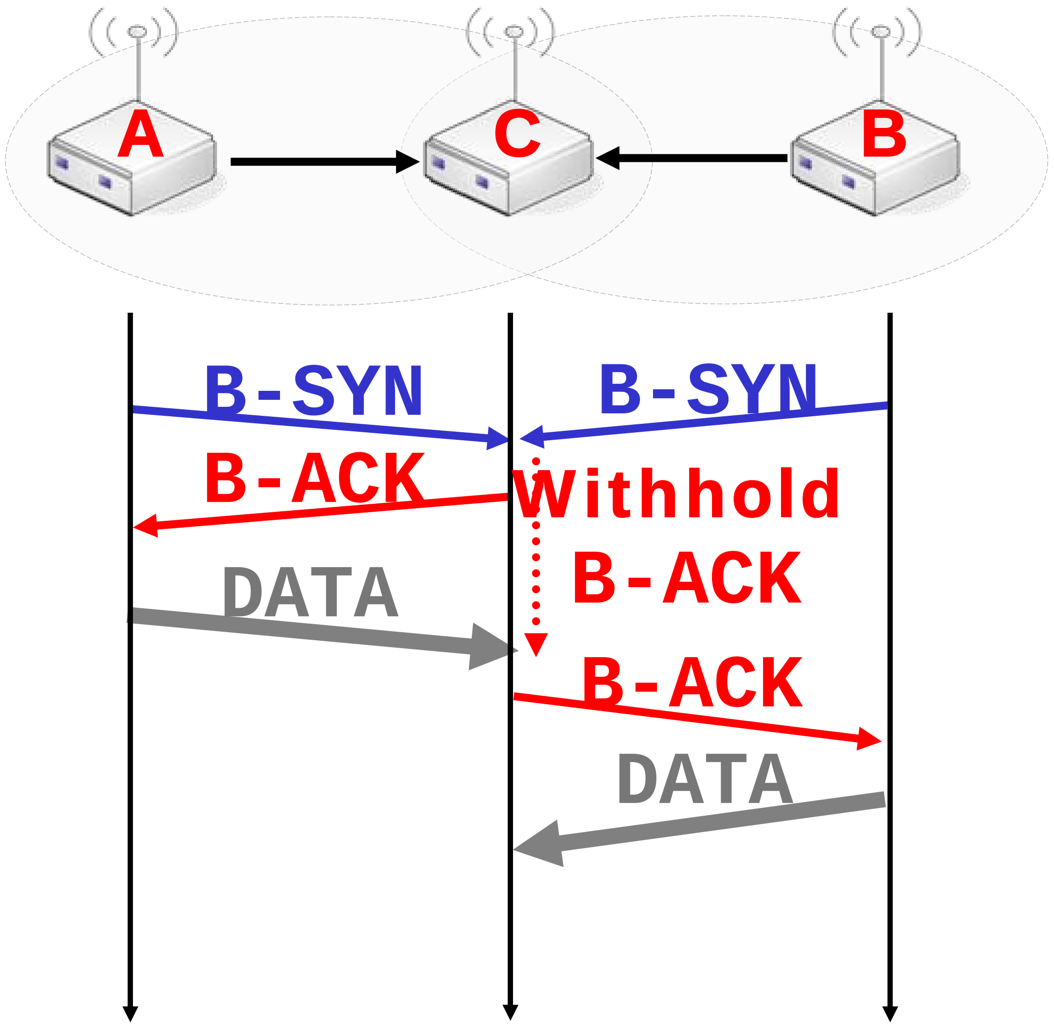

# *Hop* Design

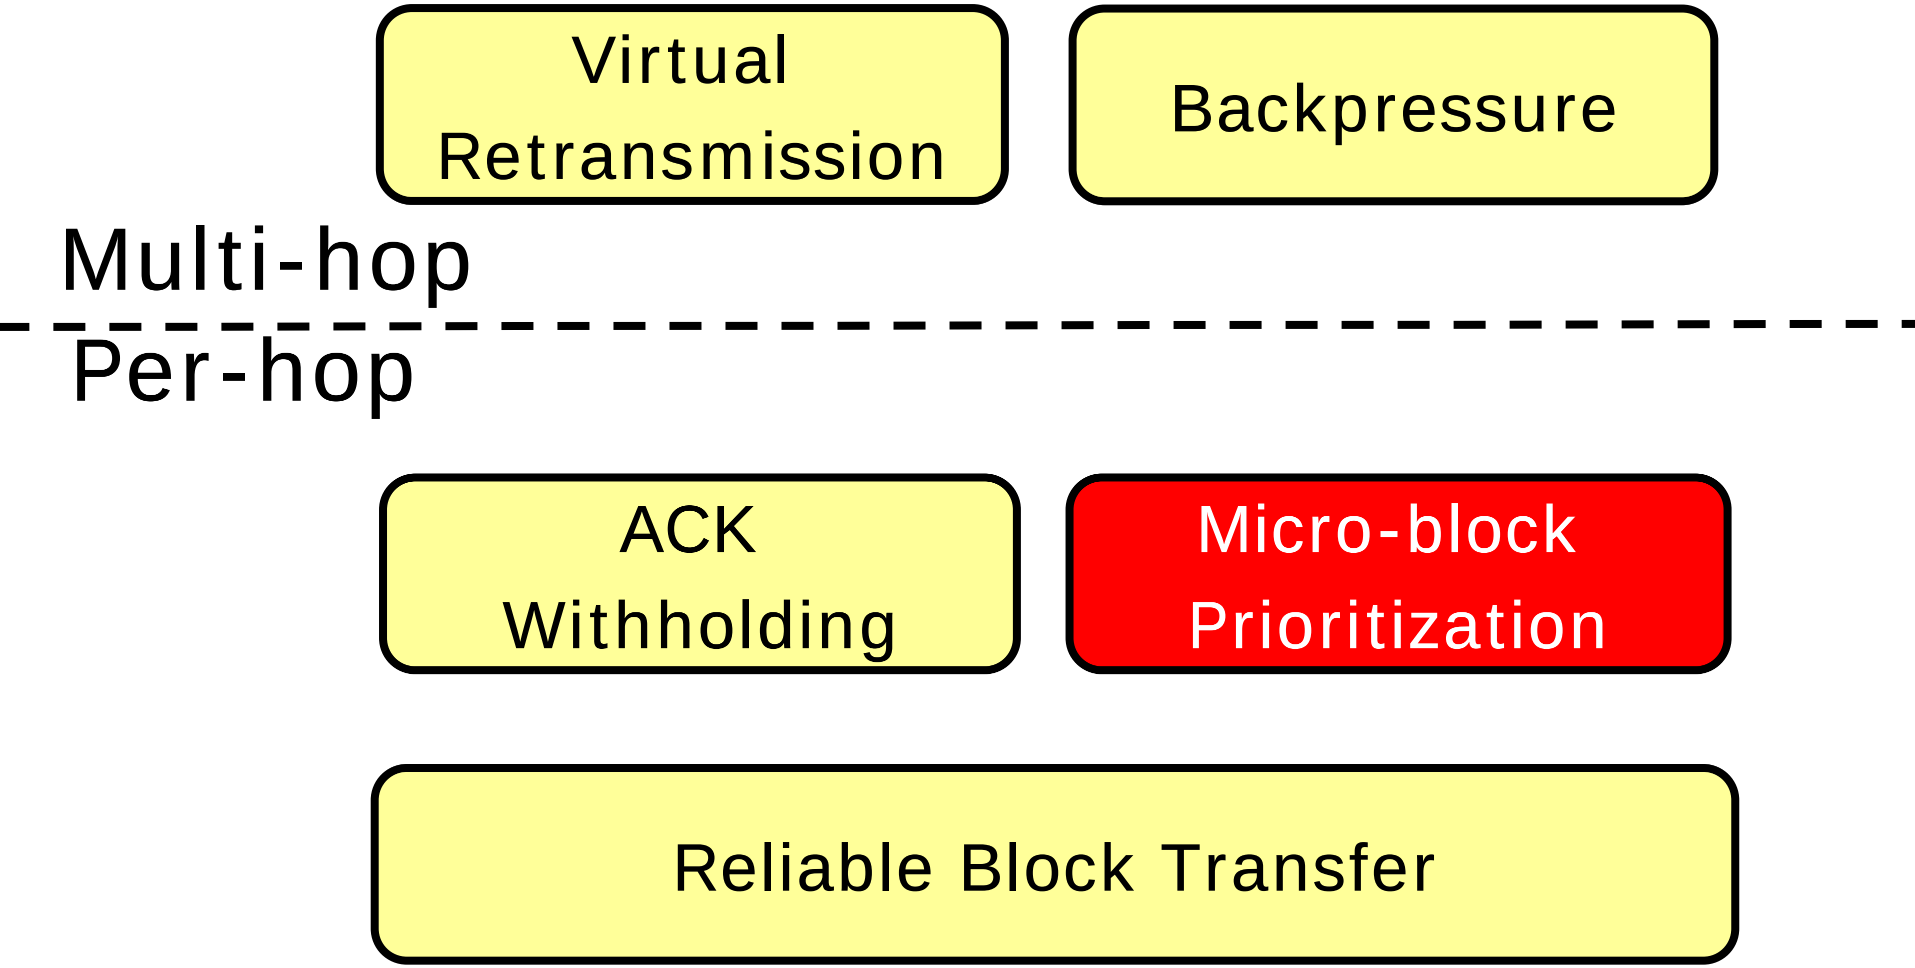

Virtual Retransmission
Backpressure
Multi-hop
Per-hop
ACK Withholding
Micro-block Prioritization
Reliable Block Transfer

# Micro-block Prioritization

- **Mechanisms**
  - Sender piggybacks small blocks to B-SYN
  - Receiver prioritizes small block's B-ACK
- **Benefits**
  - Low delay for small blocks

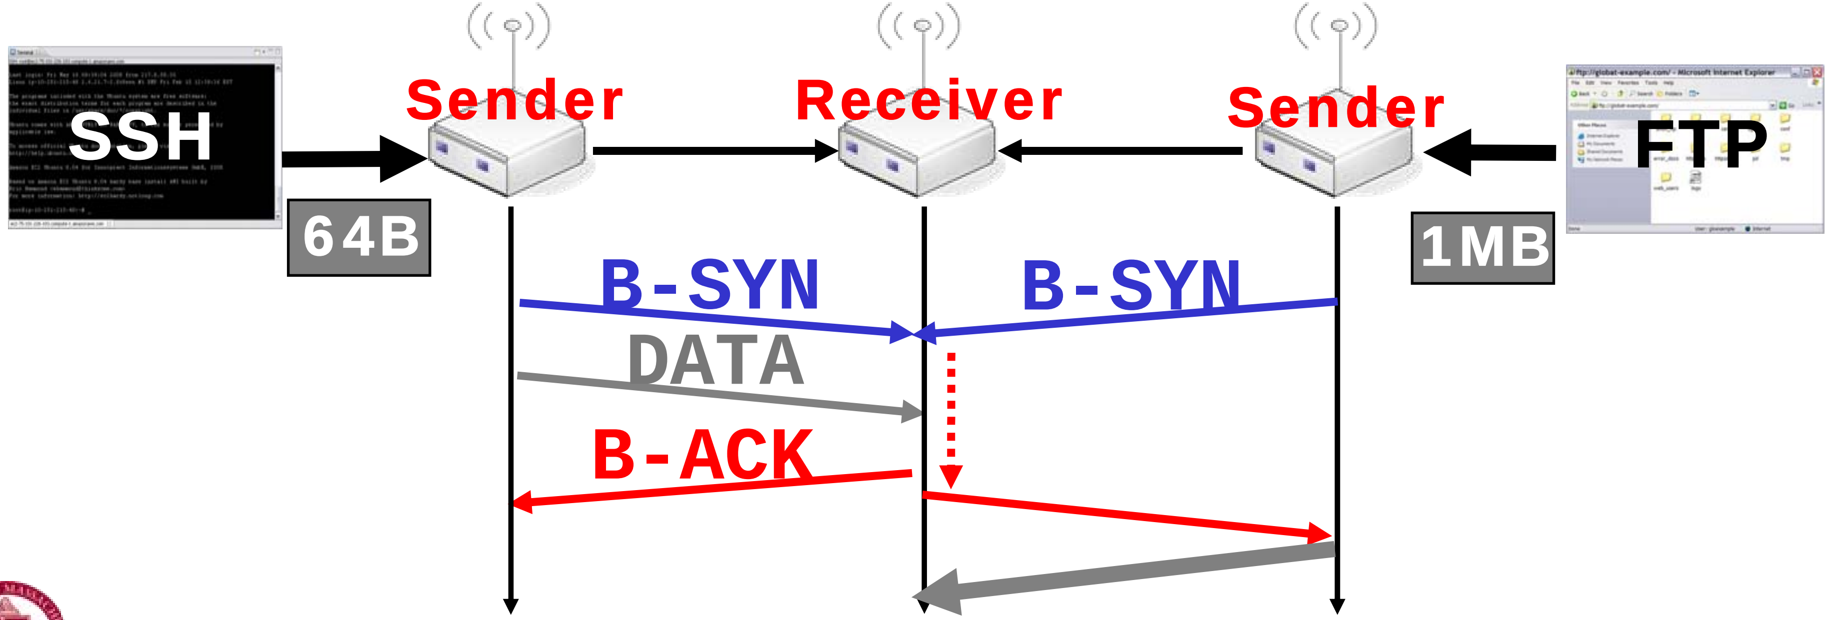

# Testbed

- **20 nodes on the 2nd floor of UMass CS building**

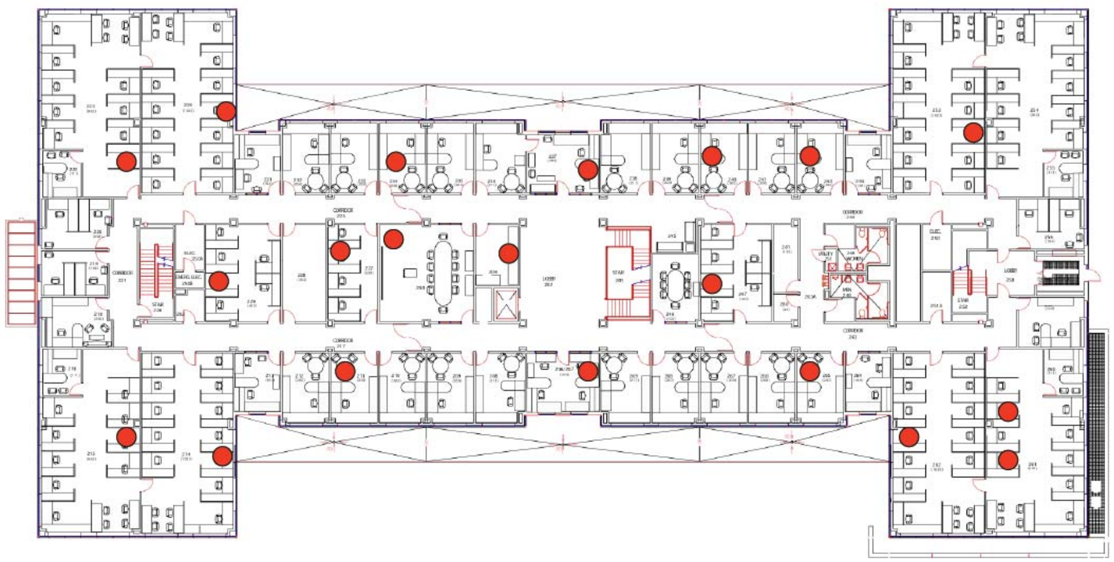

- **Apple Mac Mini**
  - Dual Core 1.8GHz, 2GB RAM, Atheros 802.11 a/b/g card

# Single-flow Single-hop Performance

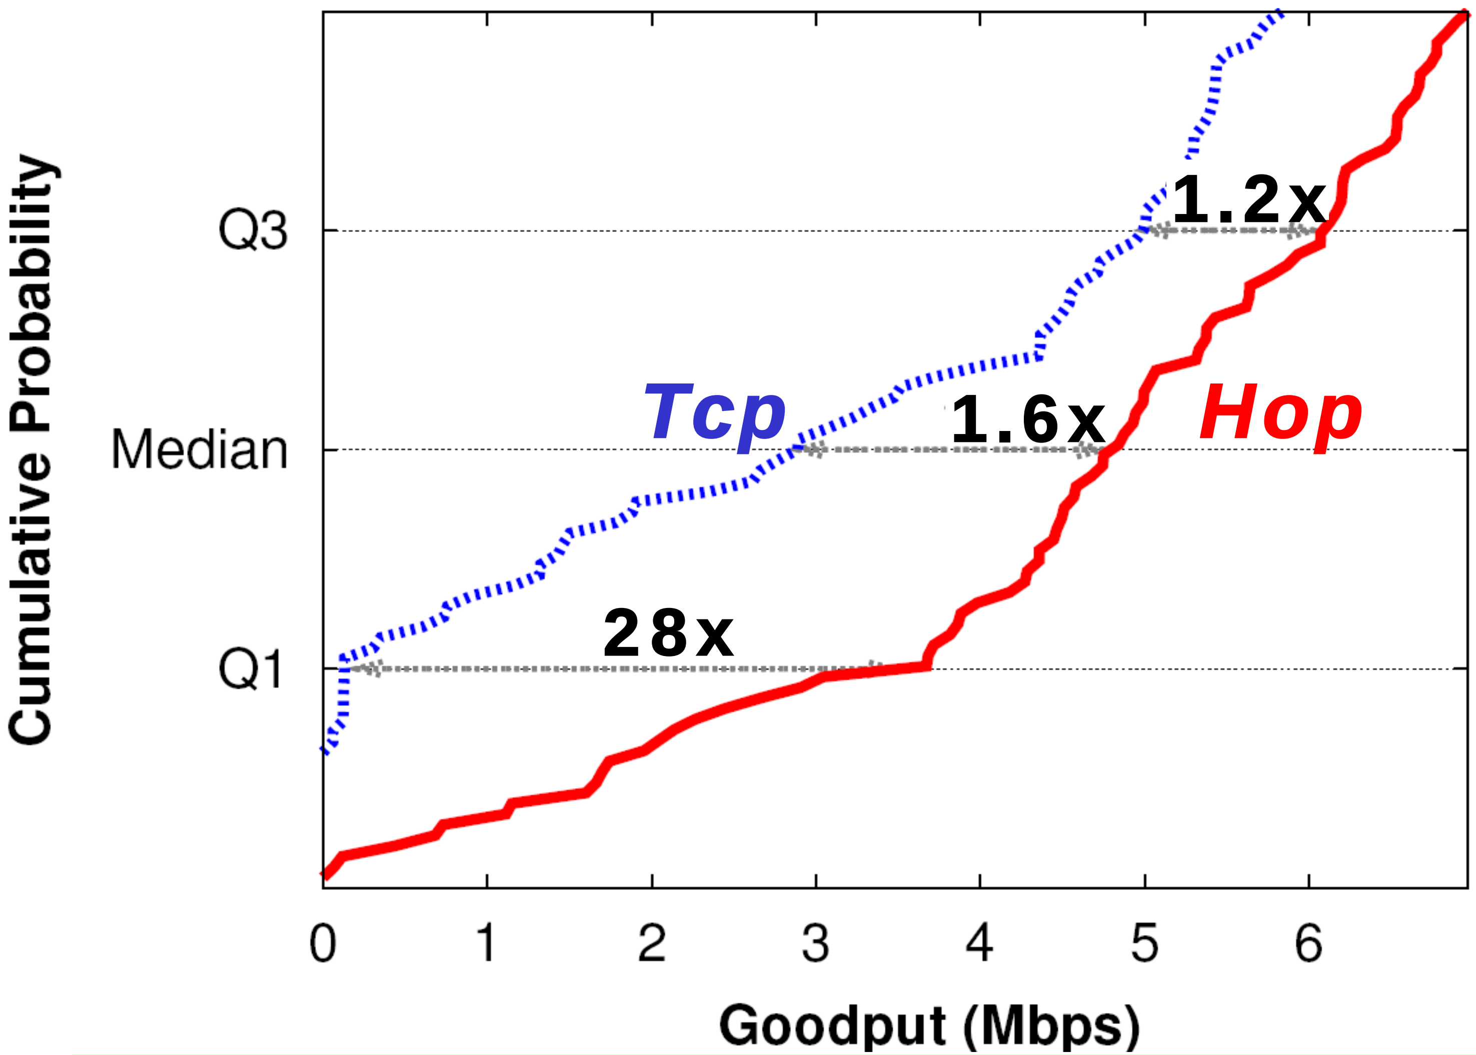


**Hop achieves significant gains over TCP**

# Single-flow Multi-hop Performance

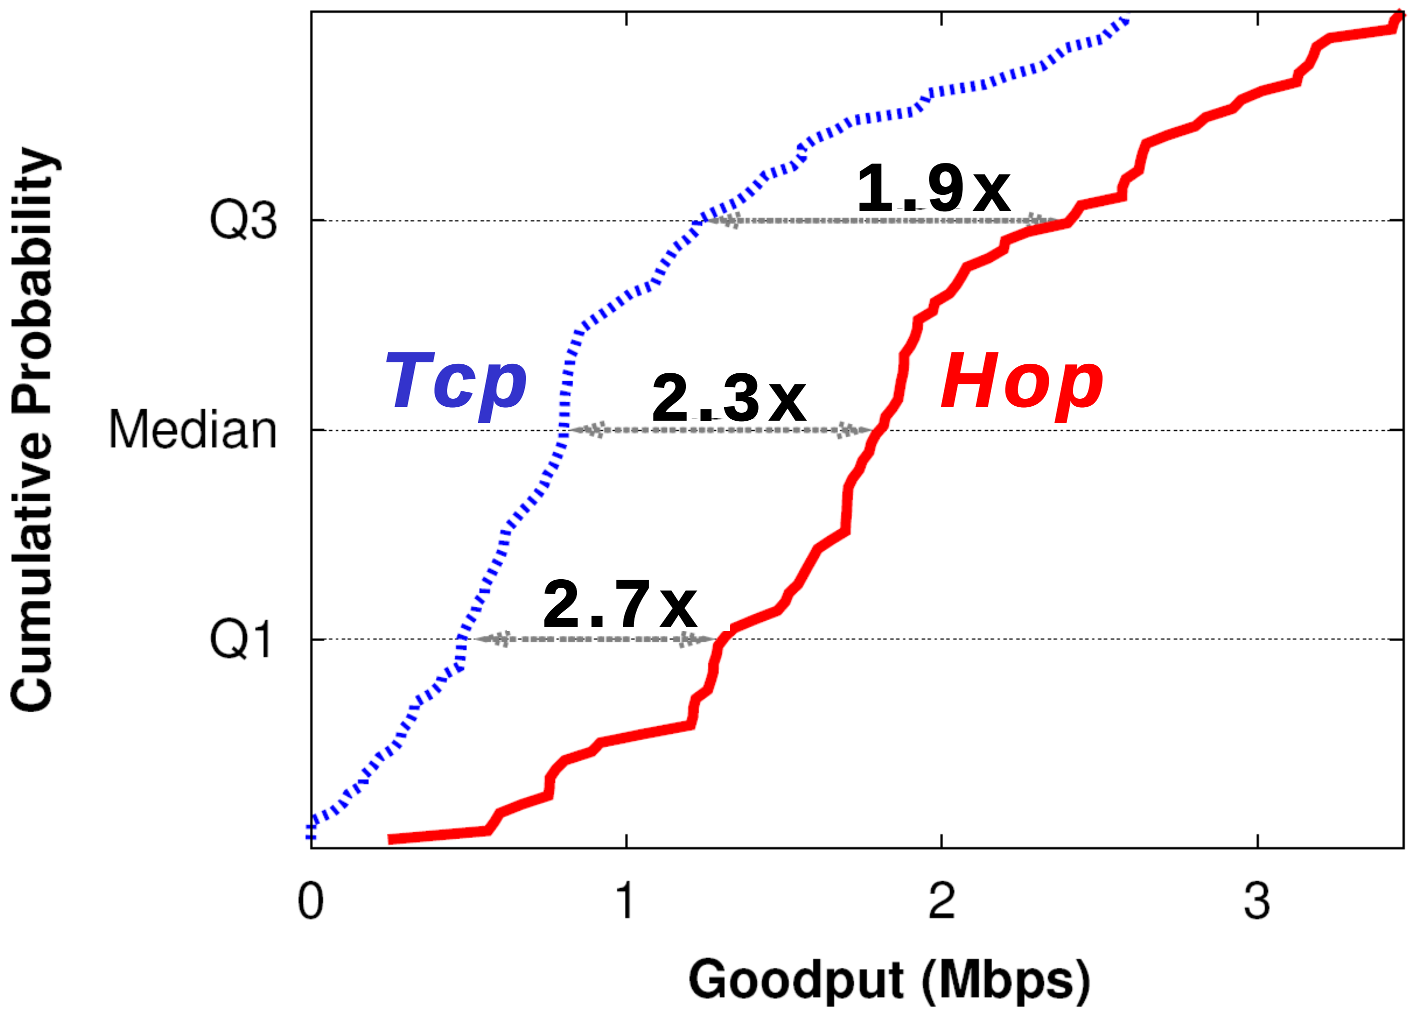


**Hop achieves significant gains over TCP**

# Graceful Degradation with Loss

- **Emulated link layer losses at the receiver**

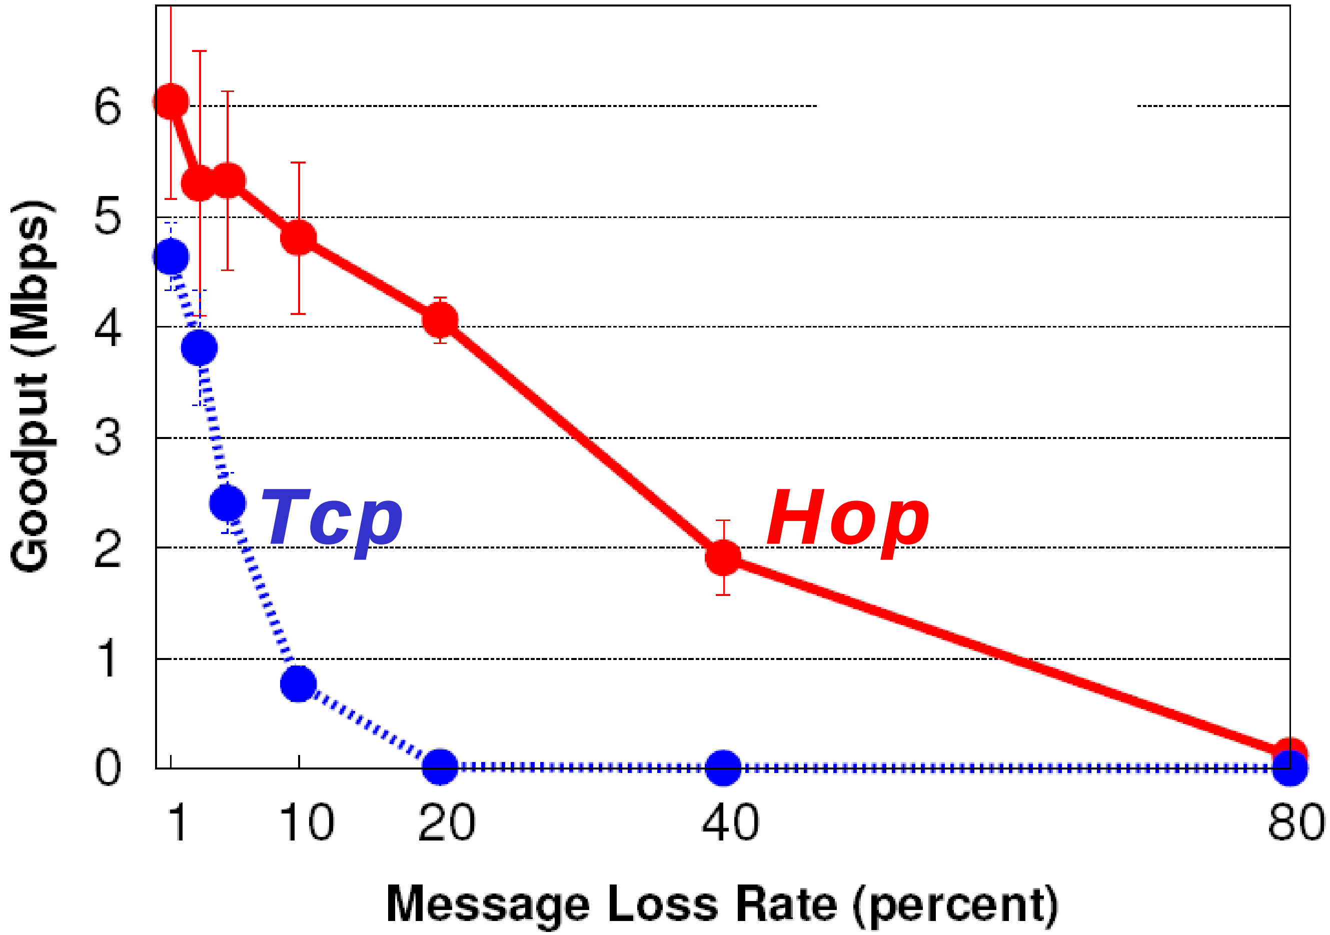


**TCP drops to zero with moderate losses**

# Scalability to High Load

- **30 concurrent flows**

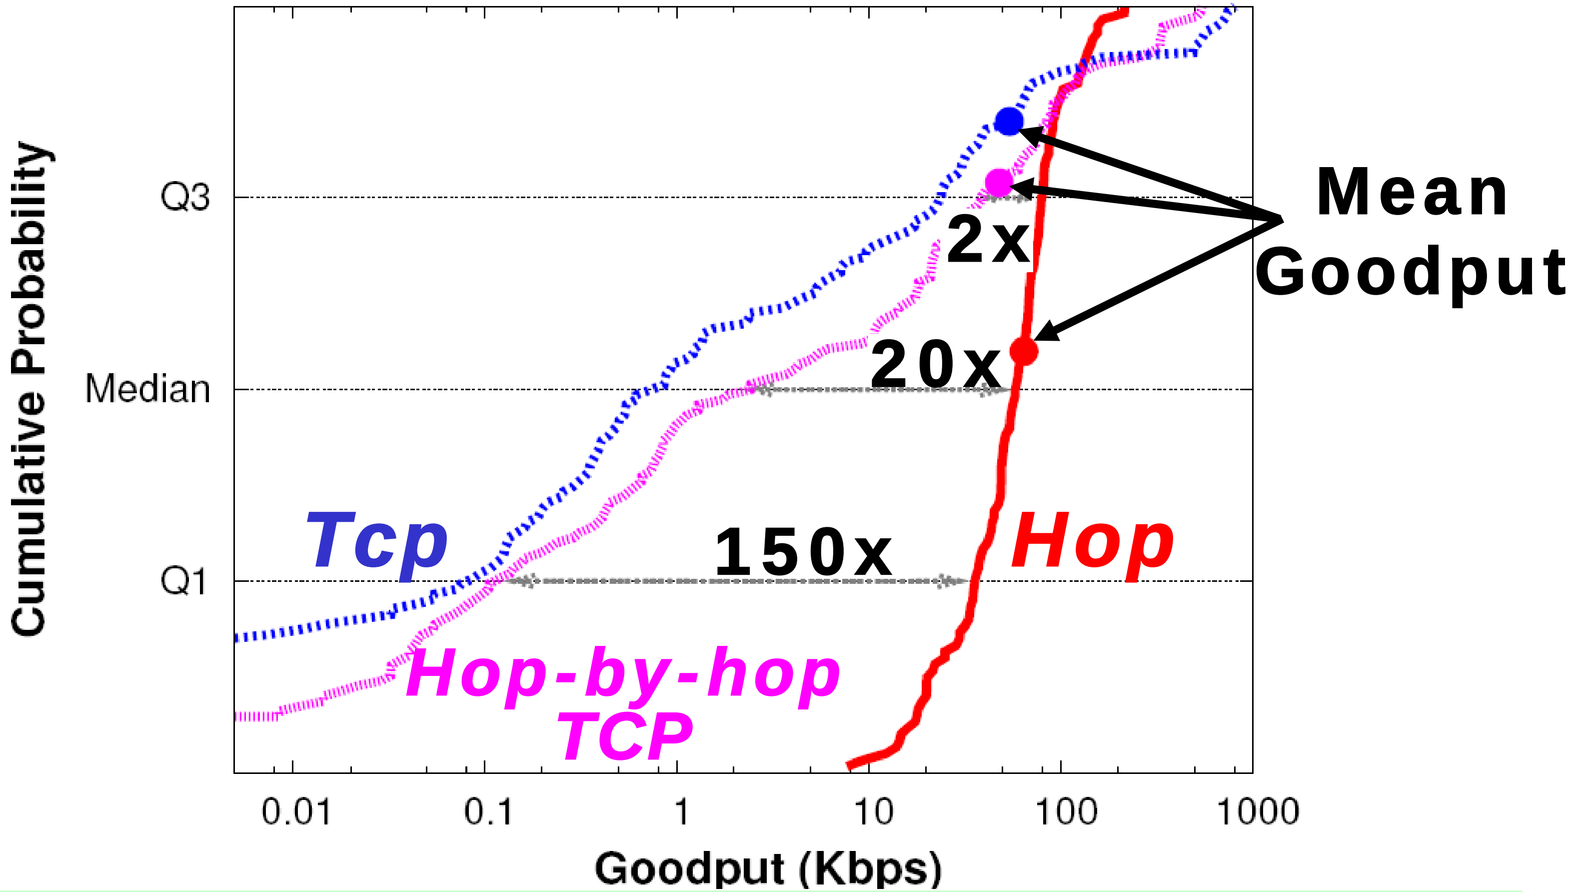


**Hop achieves massive gains over TCP and is much fairer**

# Hop over WLAN

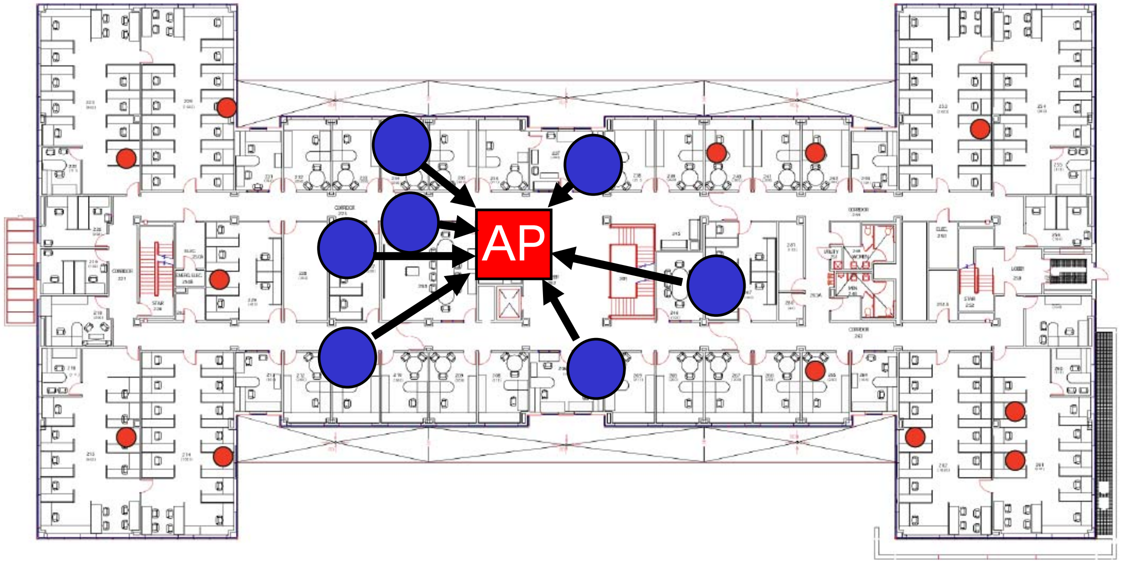


| | Mean (kbps) | Median (kbps) |
|---|---|---|
| Hop | 663 | 652 |
| TCP | 587 | 244 |
| TCP+RTS/CTS | 463 | 333 |

**Hop improves utilization over TCP+RTS/CTS**

# Low Delay for Small Transfers

- **4 nodes perform large transfers, 1 node performs small transfer**

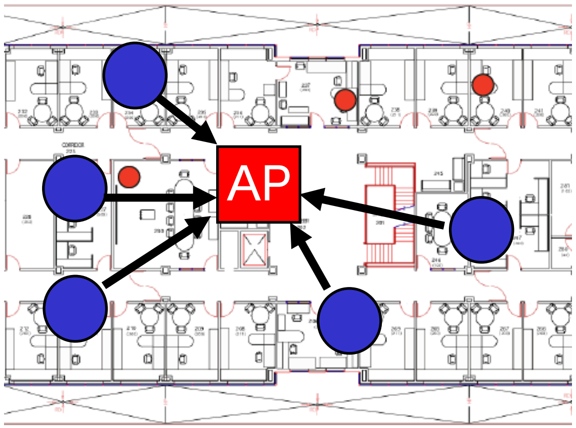


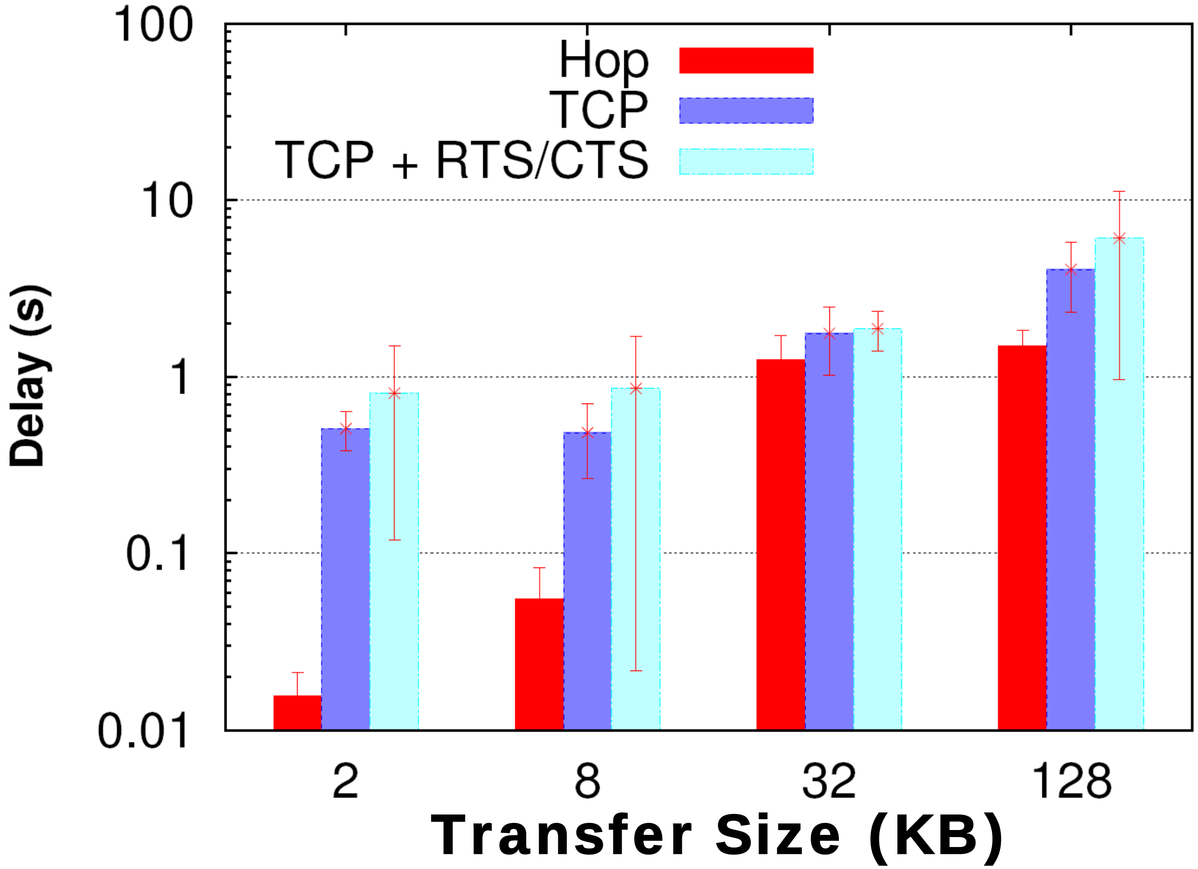


**Hop lowers delay across all file sizes**

# Summary of Other Results

- **Partitionable network**
  - TCP breaks down
  - Hop significantly outperforms (TCP-based) DTN2.5
- **Network and link layer dynamics**
  - Hop outperforms TCP under dynamic network conditions
- **Hop under 802.11g**
  - Similar performance gains as in 802.11b
- **Impact on VoIP**
  - Hop impacts concurrent VoIP slightly more than TCP, but achieves significantly higher goodput.

# Related Work

- **Fixing E2E rate-control**
  - Separating loss/congestion [Snoop, WTCP, Westwood+, ATCP, TCP-ELFN]
  - Network-assisted rate control [ATP, NRED, IFRC, WCP]
  - **Hop circumvents rate control**
- **Backpressure**
  - ATM, theoretical work [Tassiulas,...]
  - Tree/chain sensor data aggregation [Fusion, Flush]
  - Reliable point-to-point transport [RAIN, CXCC, Horizon]
  - **Hop reduces backpressure overhead using blocks**
- **Batching**
  - Common optimization at link [802.11e/802.11n, WiLDNet, Kim08, CMAP], transport [Delayed-ACK, DTN2.5], and network [ExOR] layers
  - **Hop leverages batching across layers**

# Summary

- **Block switching > packet switching**
  - Key abstraction: Reliable per-hop block transfer
- **Hop**
  - **Fast**: Significant throughput, fairness, delay gains
  - **Robust**: Degrades gracefully to challenged networks
  - **Simple**: Minimizes complex cross-layer interaction

Can we have one transport protocol for diverse wireless networks? ***Yes, we can!***

Source code at

**http://hop.cs.umass.edu**

# *The End*

# Questions?

# Dynamic Network Conditions

- **30 concurrent flows**
- **Auto Bit-Rate Control**
- **OLSR**

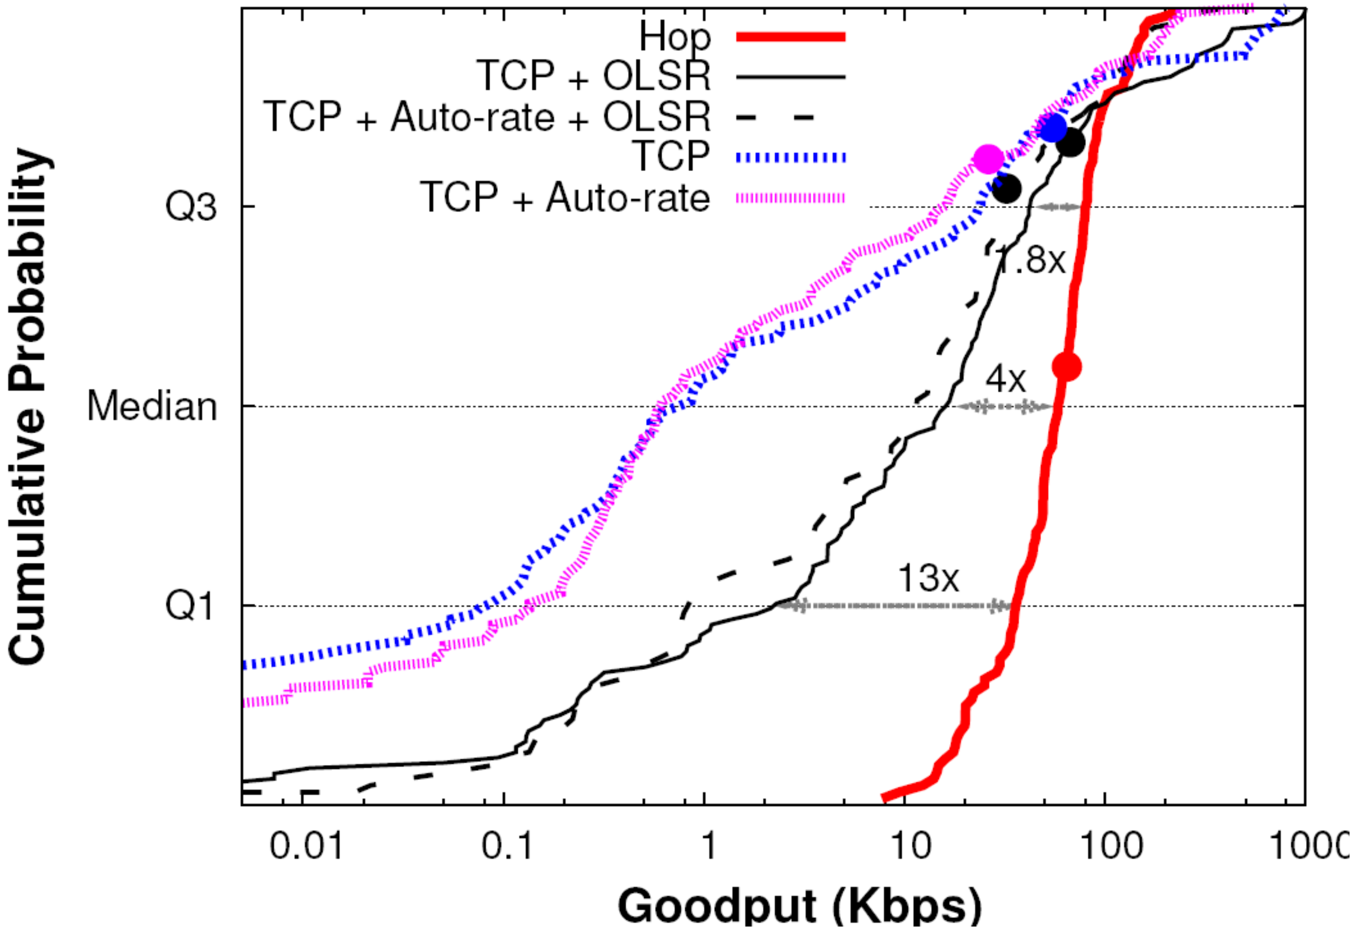


**Hop continues to significantly outperform TCP under dynamic network conditions**