

HTPPR: Tackling Wireless Losses by Exploiting Wired Reliability

Anirudh Badam

Dongsu Han

David G. Andersen

Michael Kaminsky

Dina Papagiannaki

Srinivasan Seshan

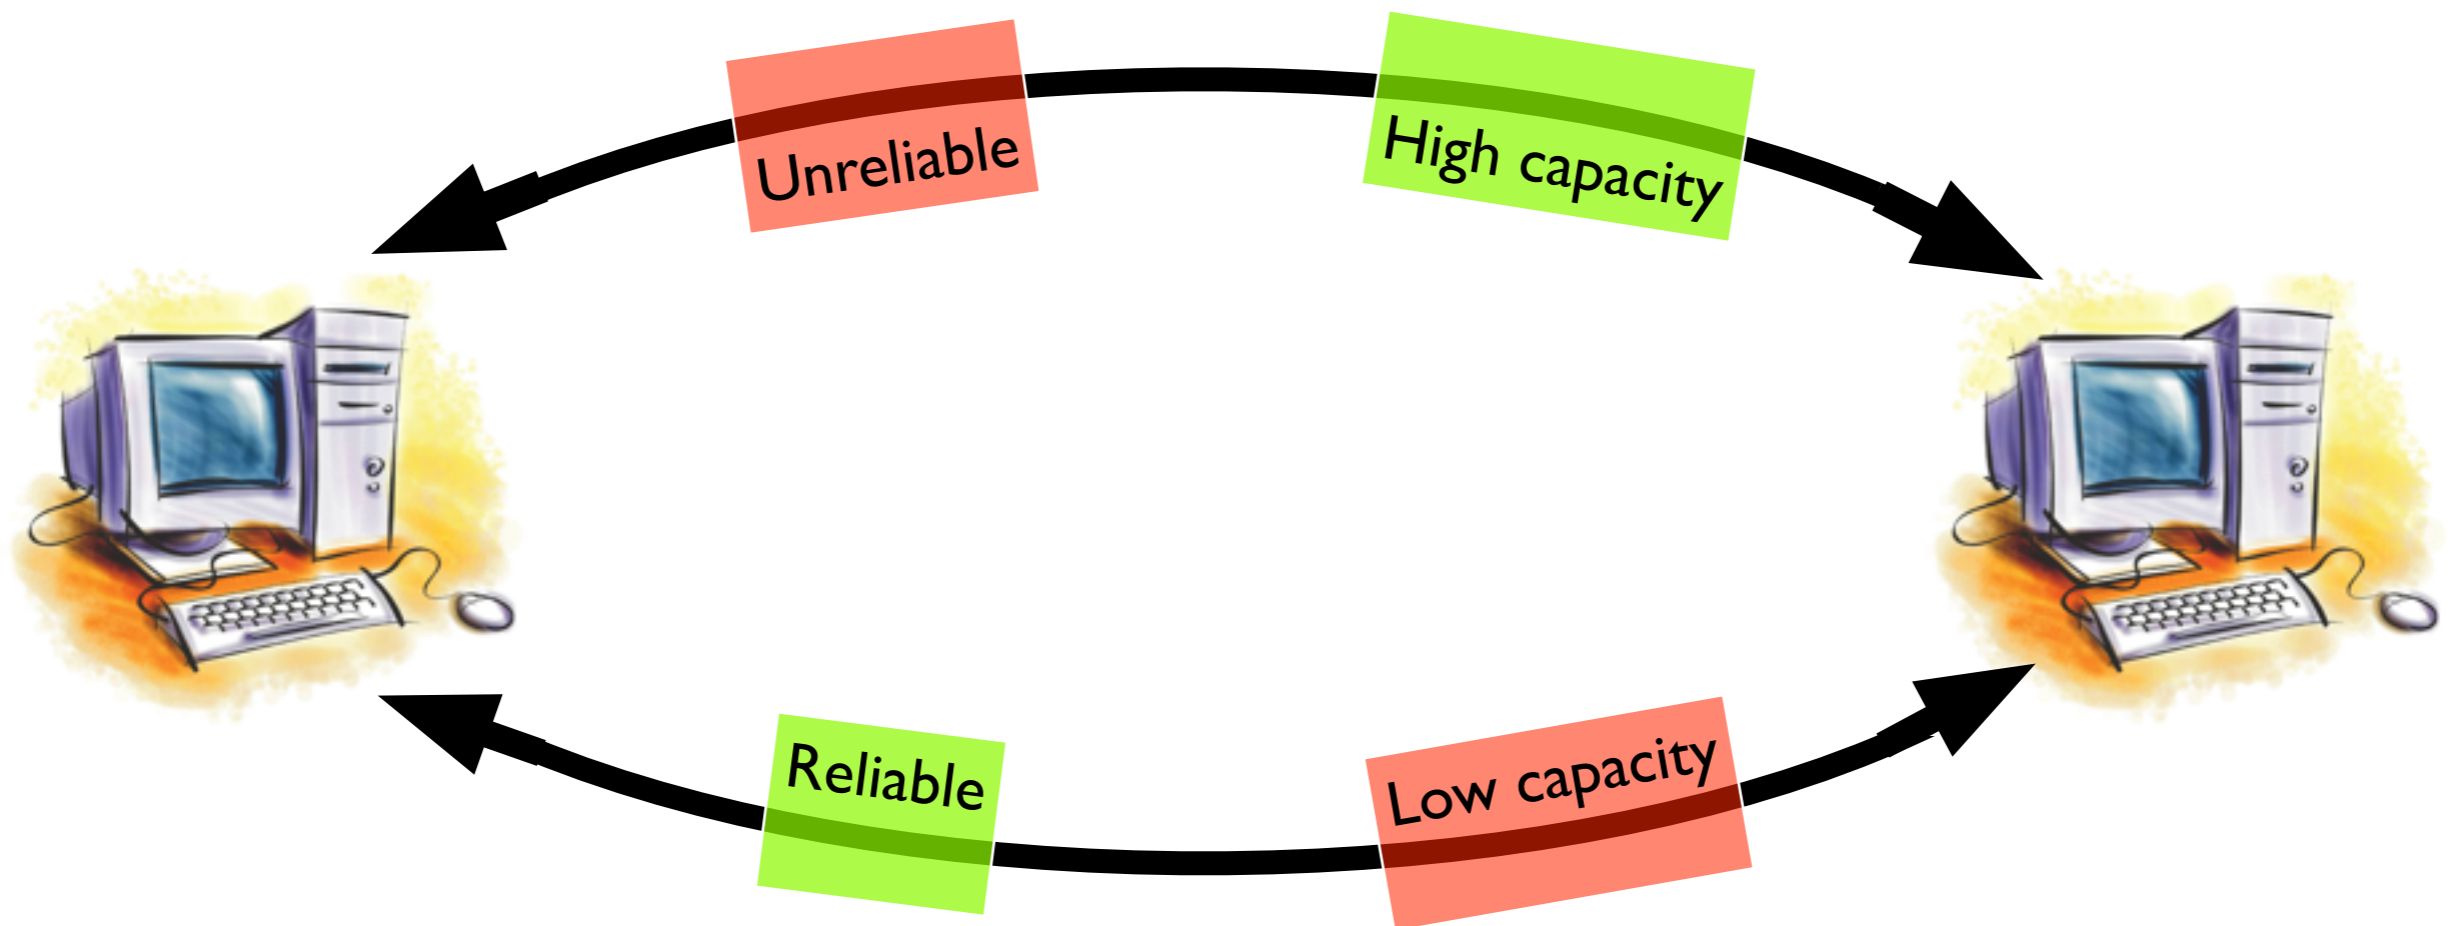


Carnegie Mellon




Goal Overview

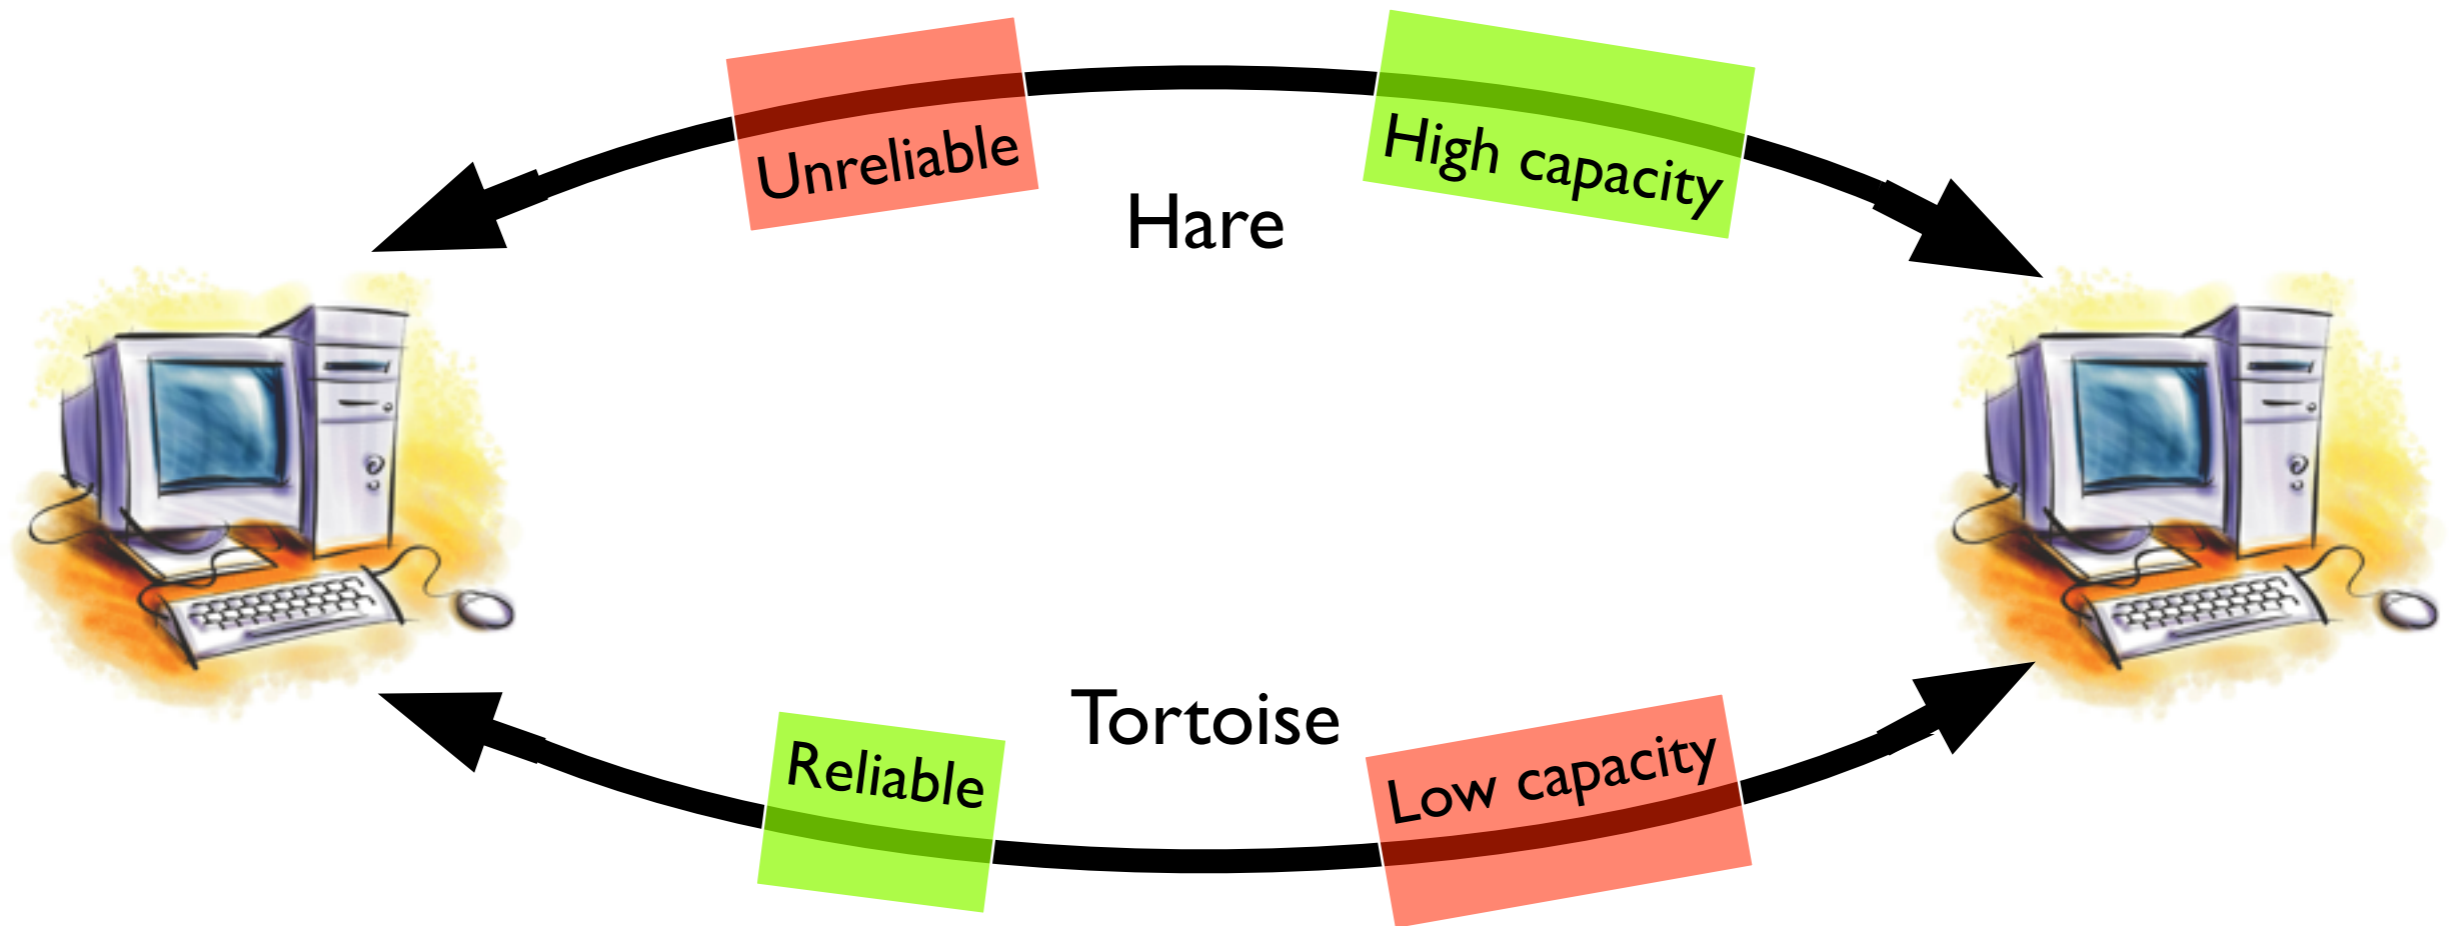
Goal Overview



Goal Overview

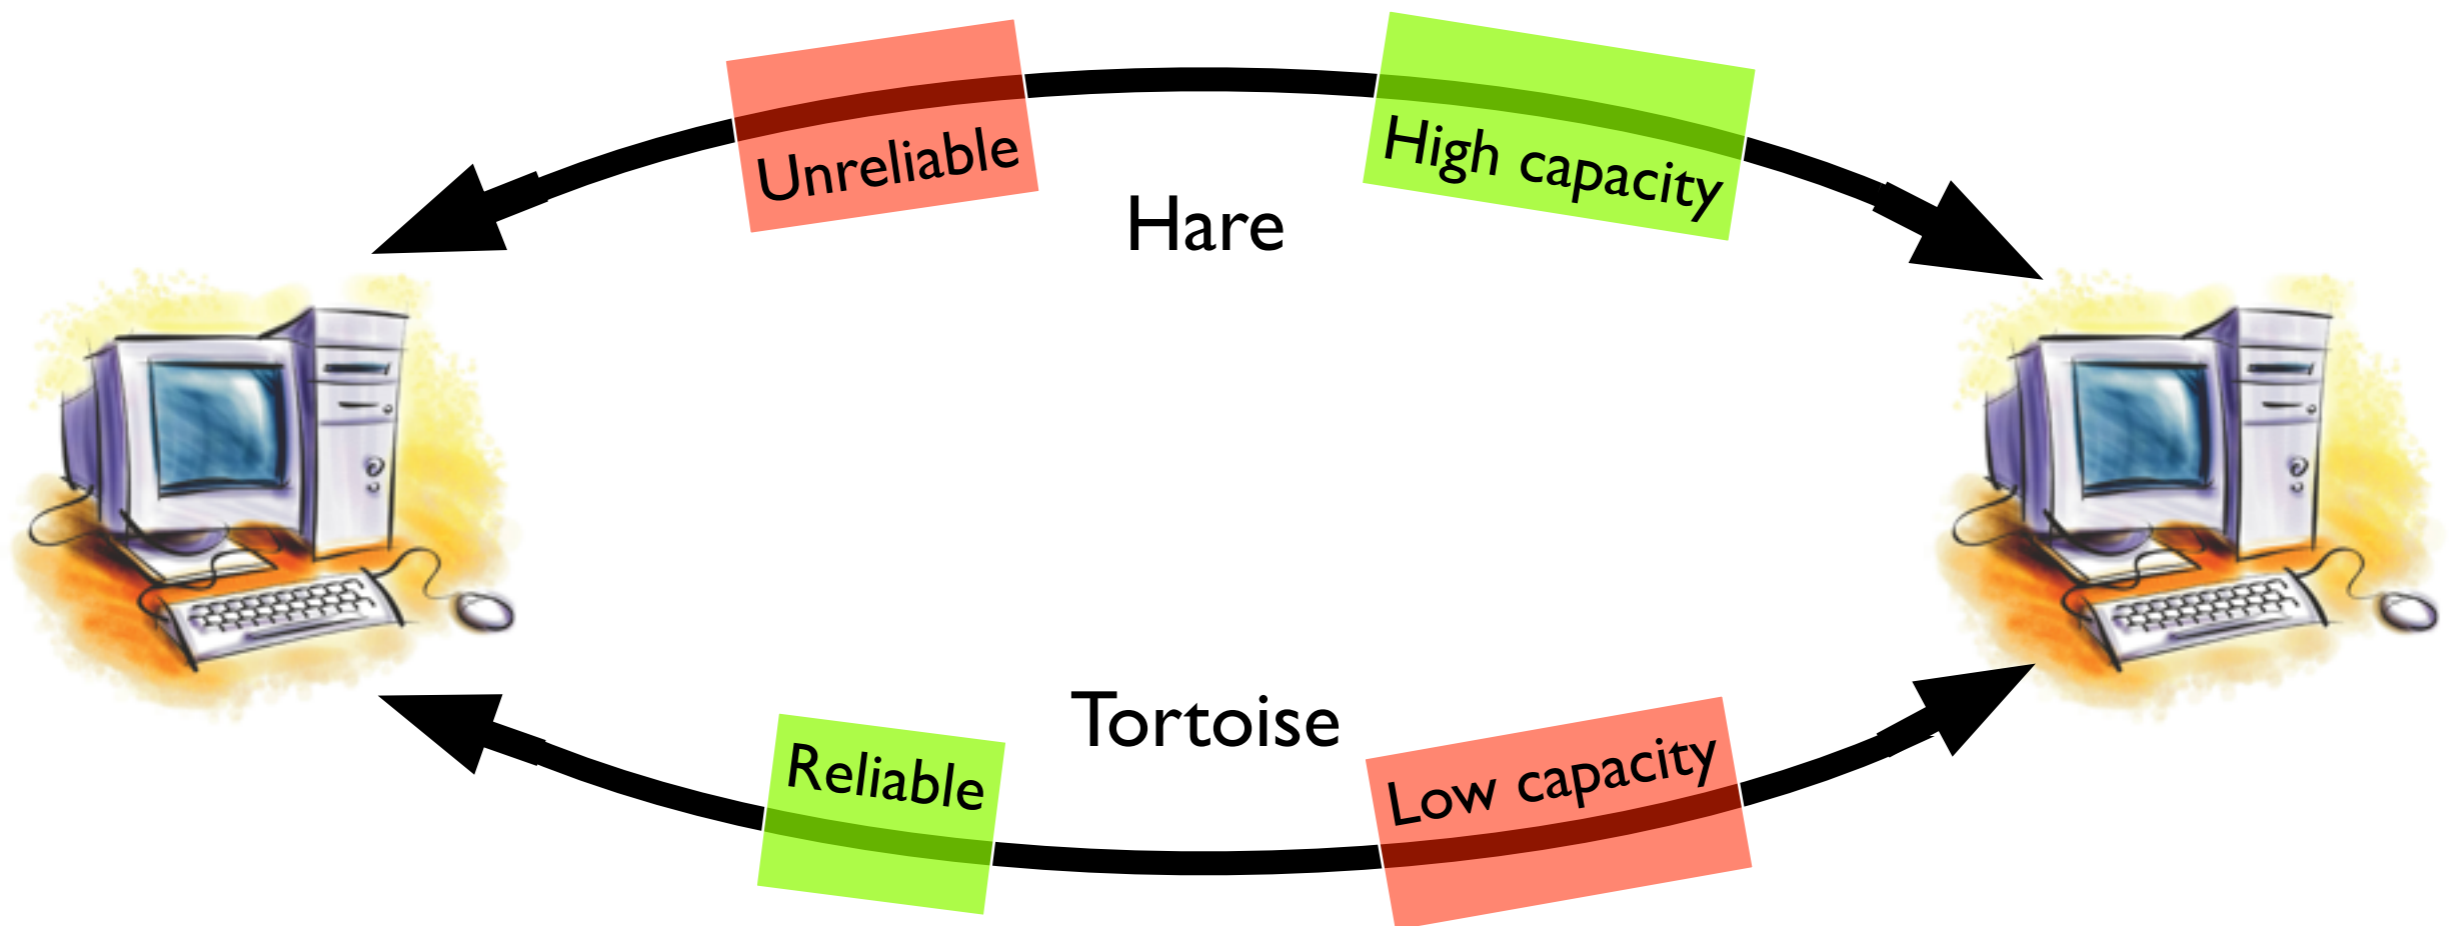


Goal Overview



$$Cap_{Total} = Cap_{Hare} * (1 - Loss_{hare}) + Cap_{Tortoise} * (1 - Loss_{Tortoise})$$

Goal Overview

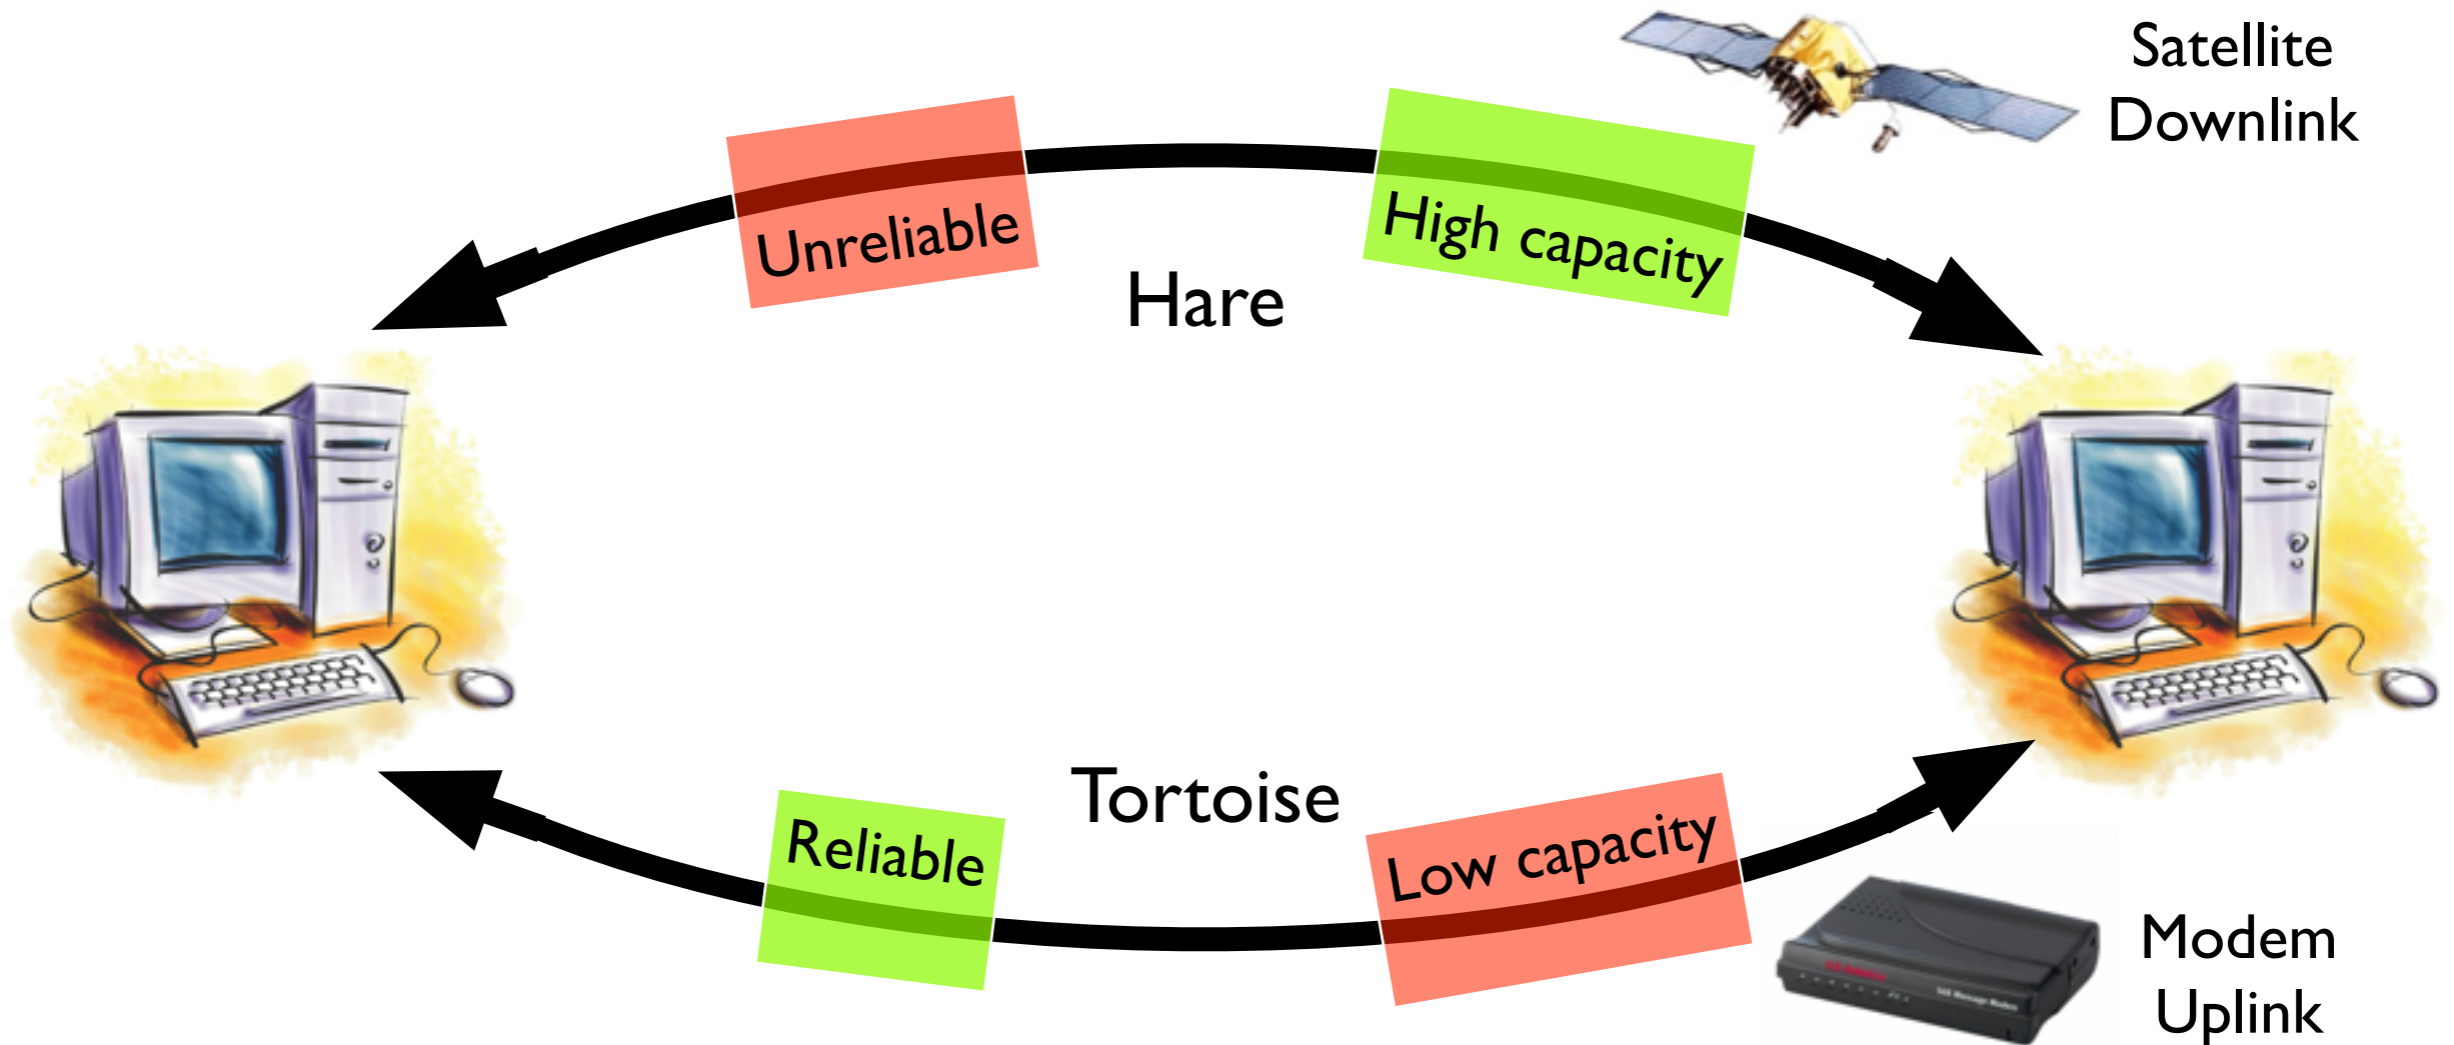


$$\text{Cap}_{\text{Total}} = \text{Cap}_{\text{Hare}} * (1 - \text{Loss}_{\text{hare}}) + \text{Cap}_{\text{Tortoise}} * (1 - \text{Loss}_{\text{Tortoise}})$$



Reduce effective $\text{Loss}_{\text{Hare}}$ by exploiting the fact that $\text{Loss}_{\text{Tortoise}}$ is low

Goal Overview

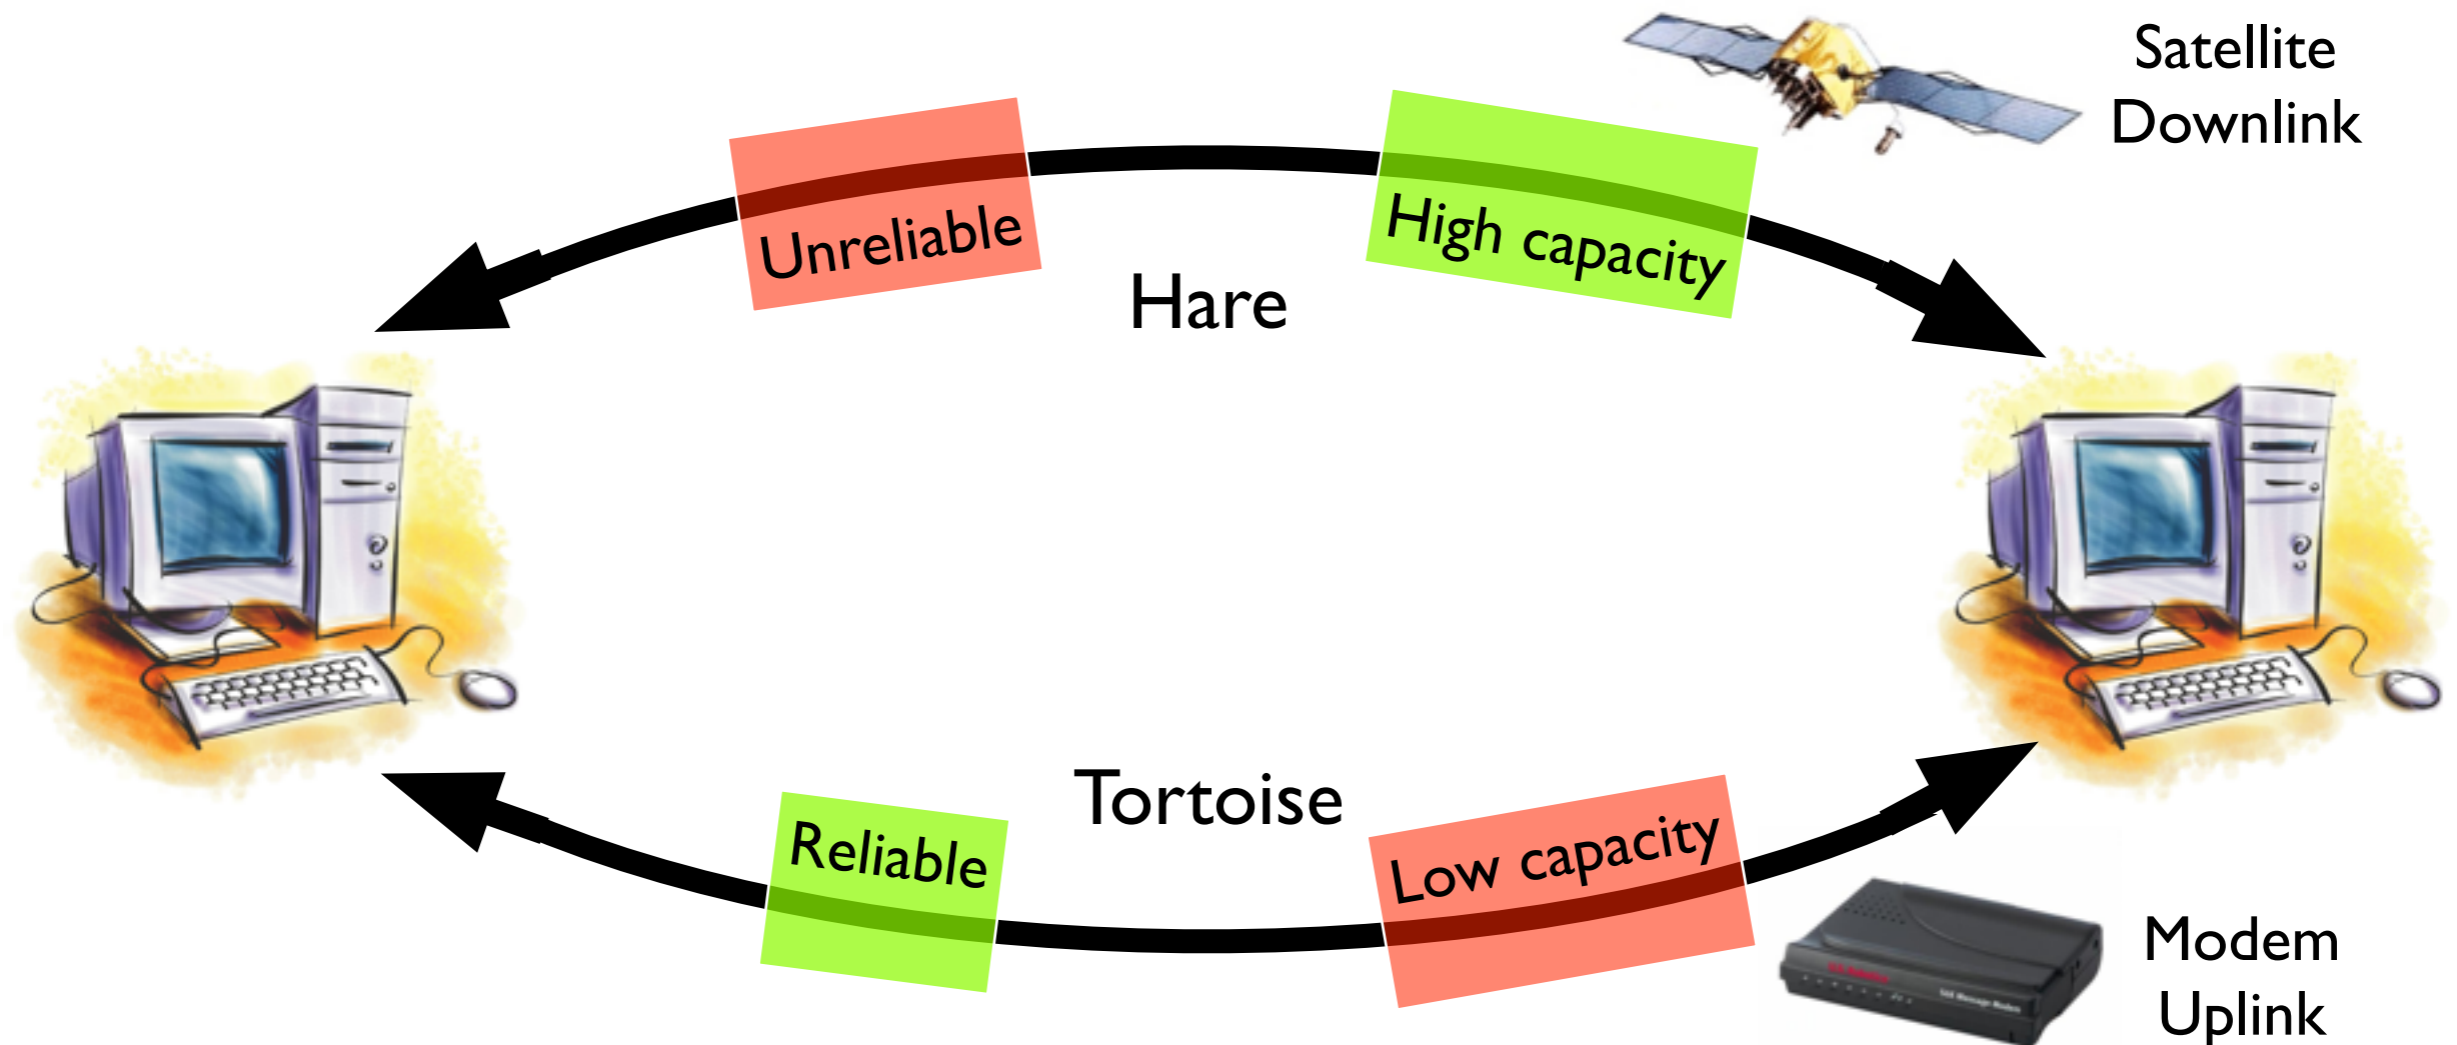


$$Cap_{Total} = Cap_{Hare} * (1 - Loss_{hare}) + Cap_{Tortoise} * (1 - Loss_{Tortoise})$$



Reduce effective $Loss_{Hare}$ by exploiting the fact that $Loss_{Tortoise}$ is low

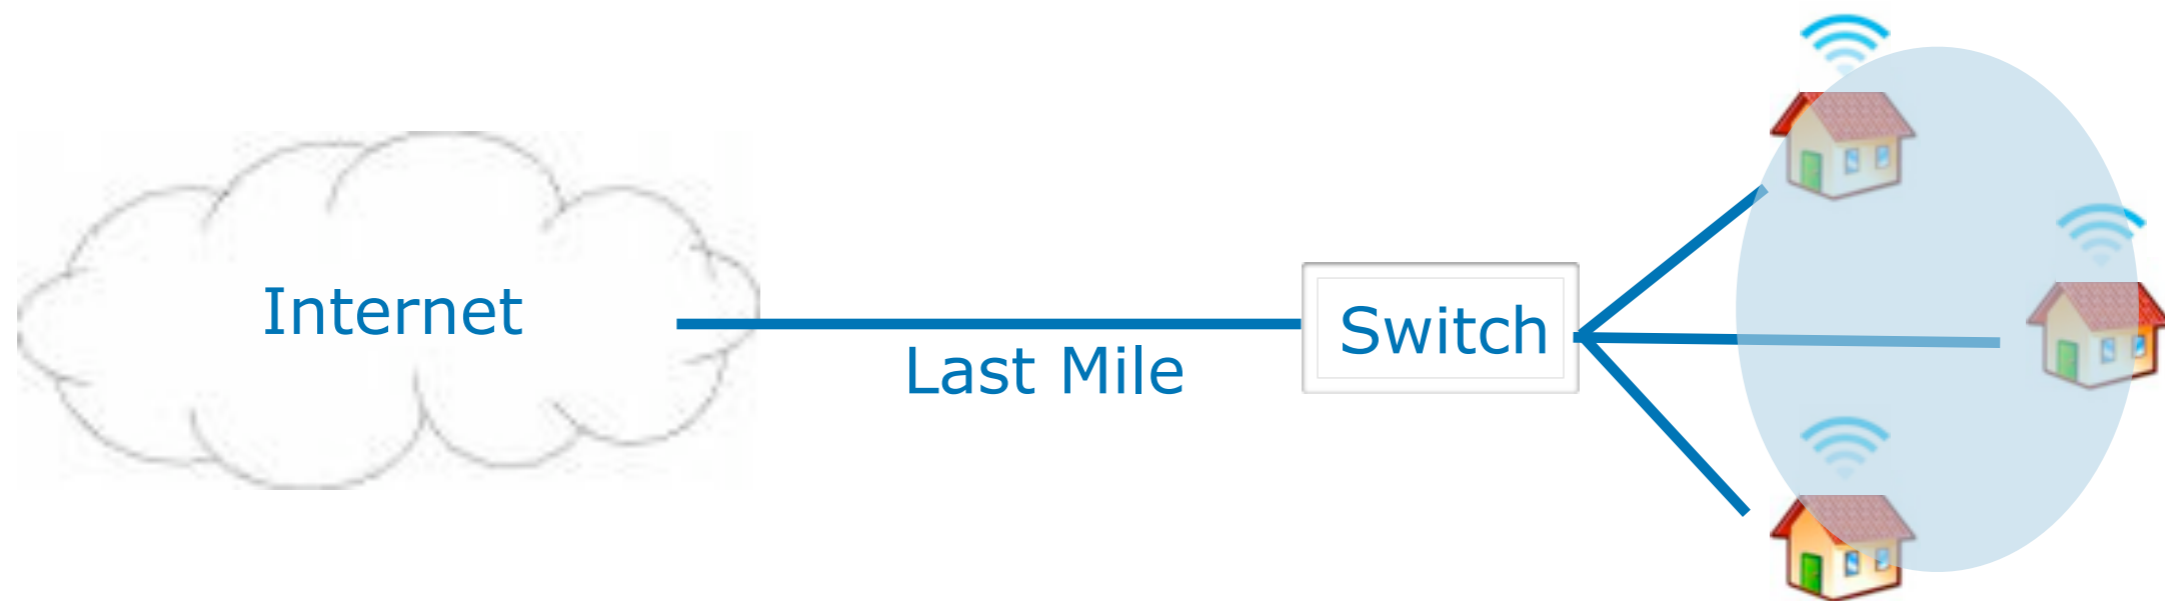
Goal Overview



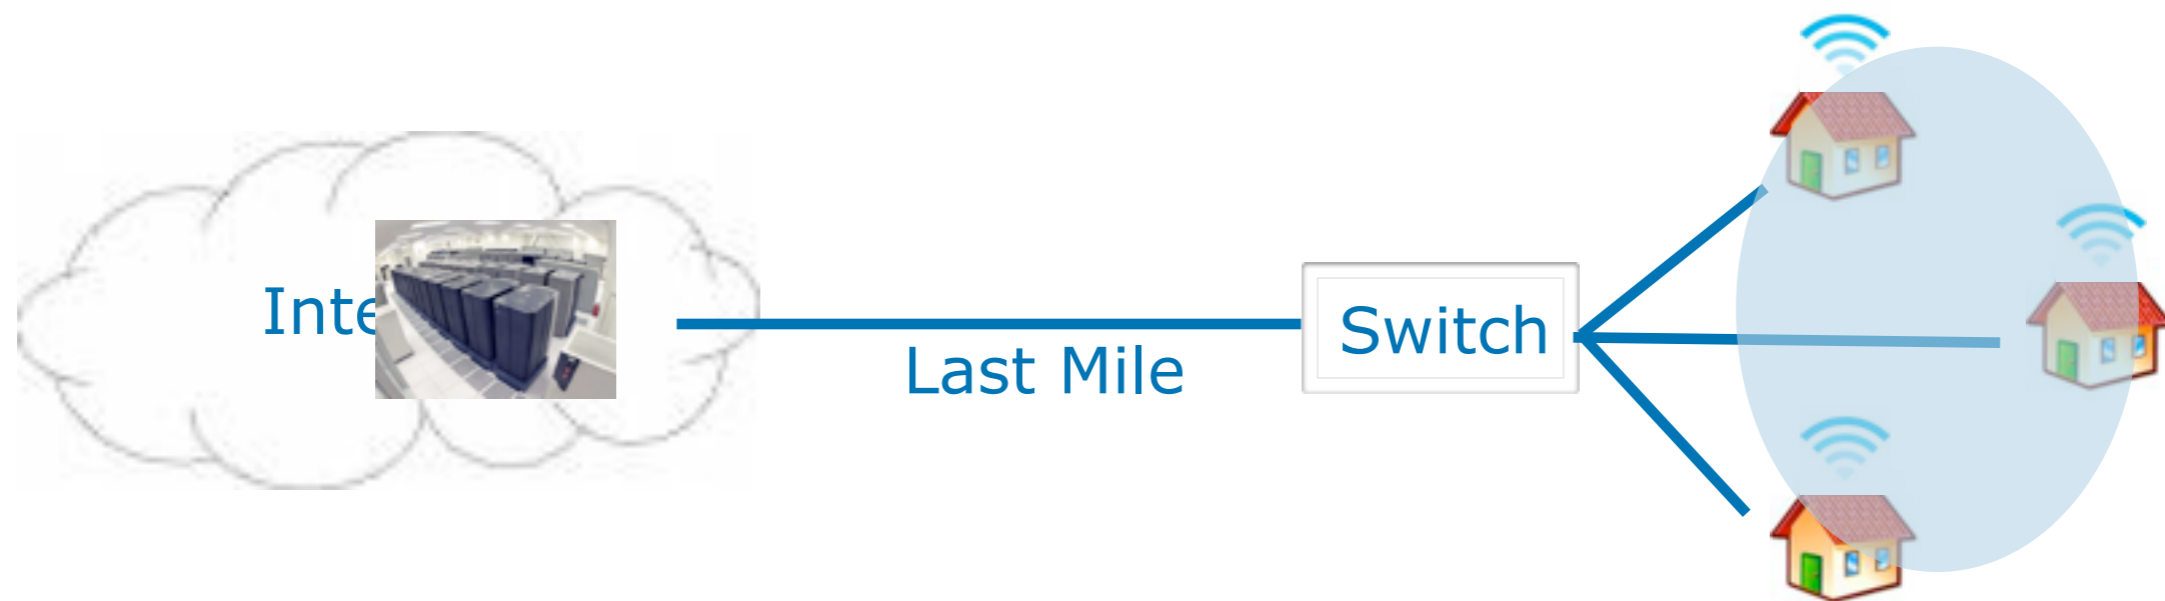
- Long distance WiFi + PPPoE to remote places
- Wireless in predominantly ethernet based datacenters
- Neighborhood wireless + broadband

Neighborhood Aware Networks

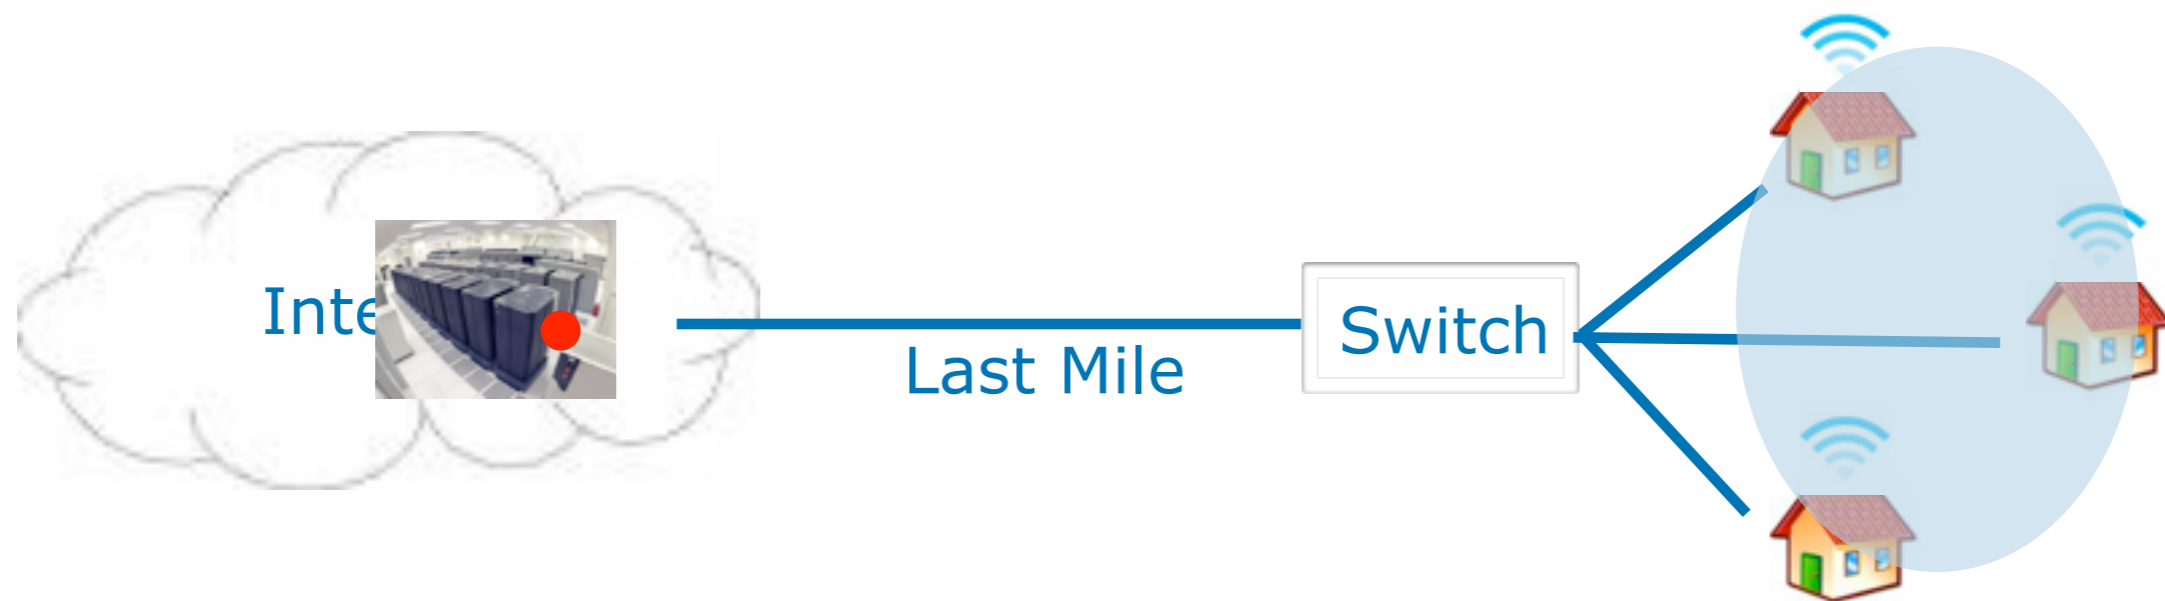
Neighborhood Aware Networks



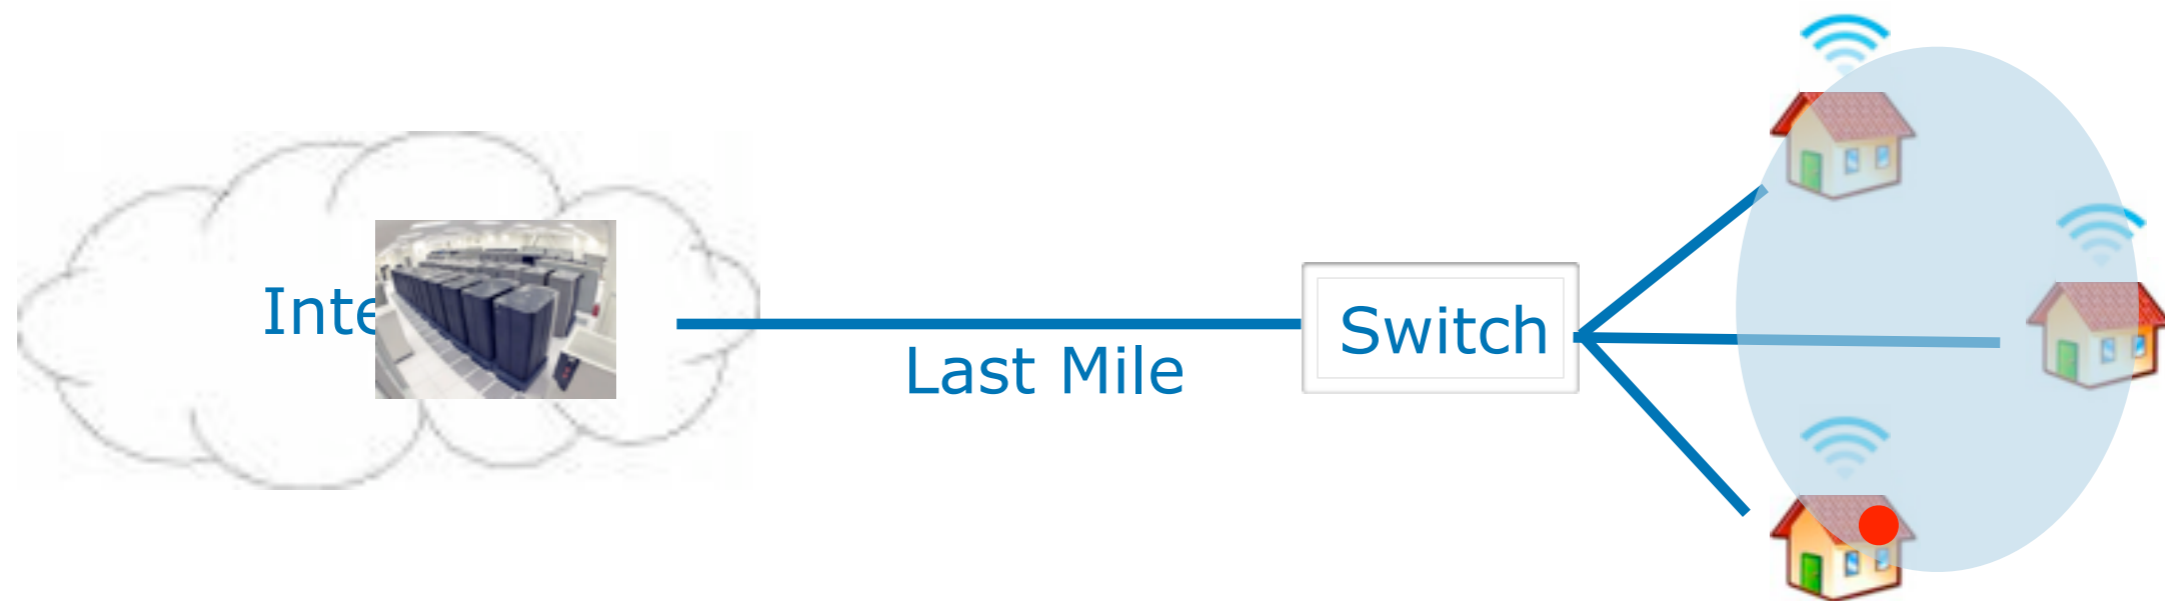
Neighborhood Aware Networks




Neighborhood Aware Networks



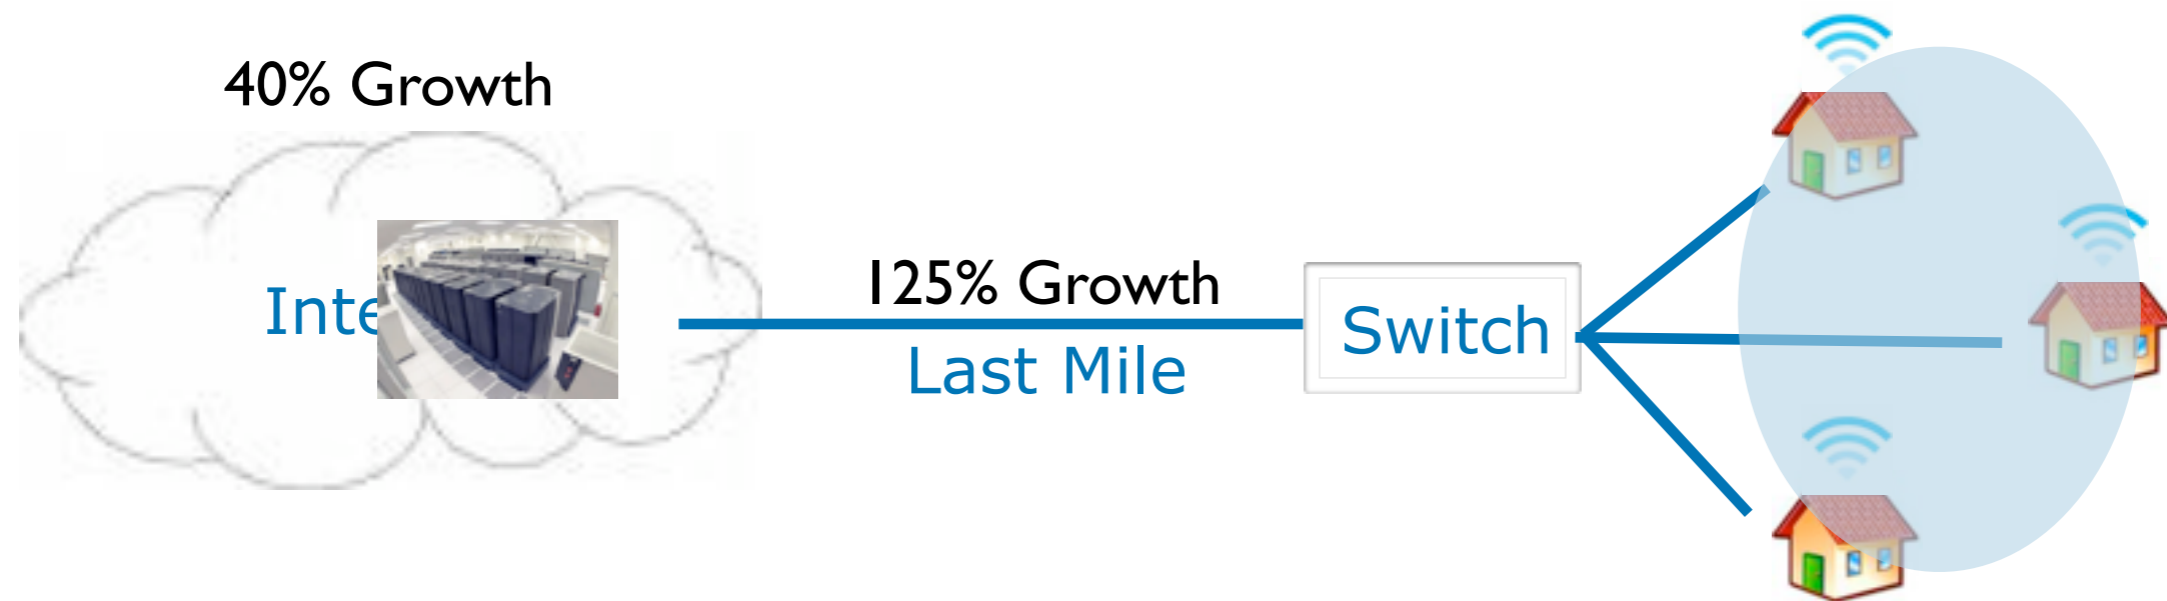
Neighborhood Aware Networks




Neighborhood Aware Networks



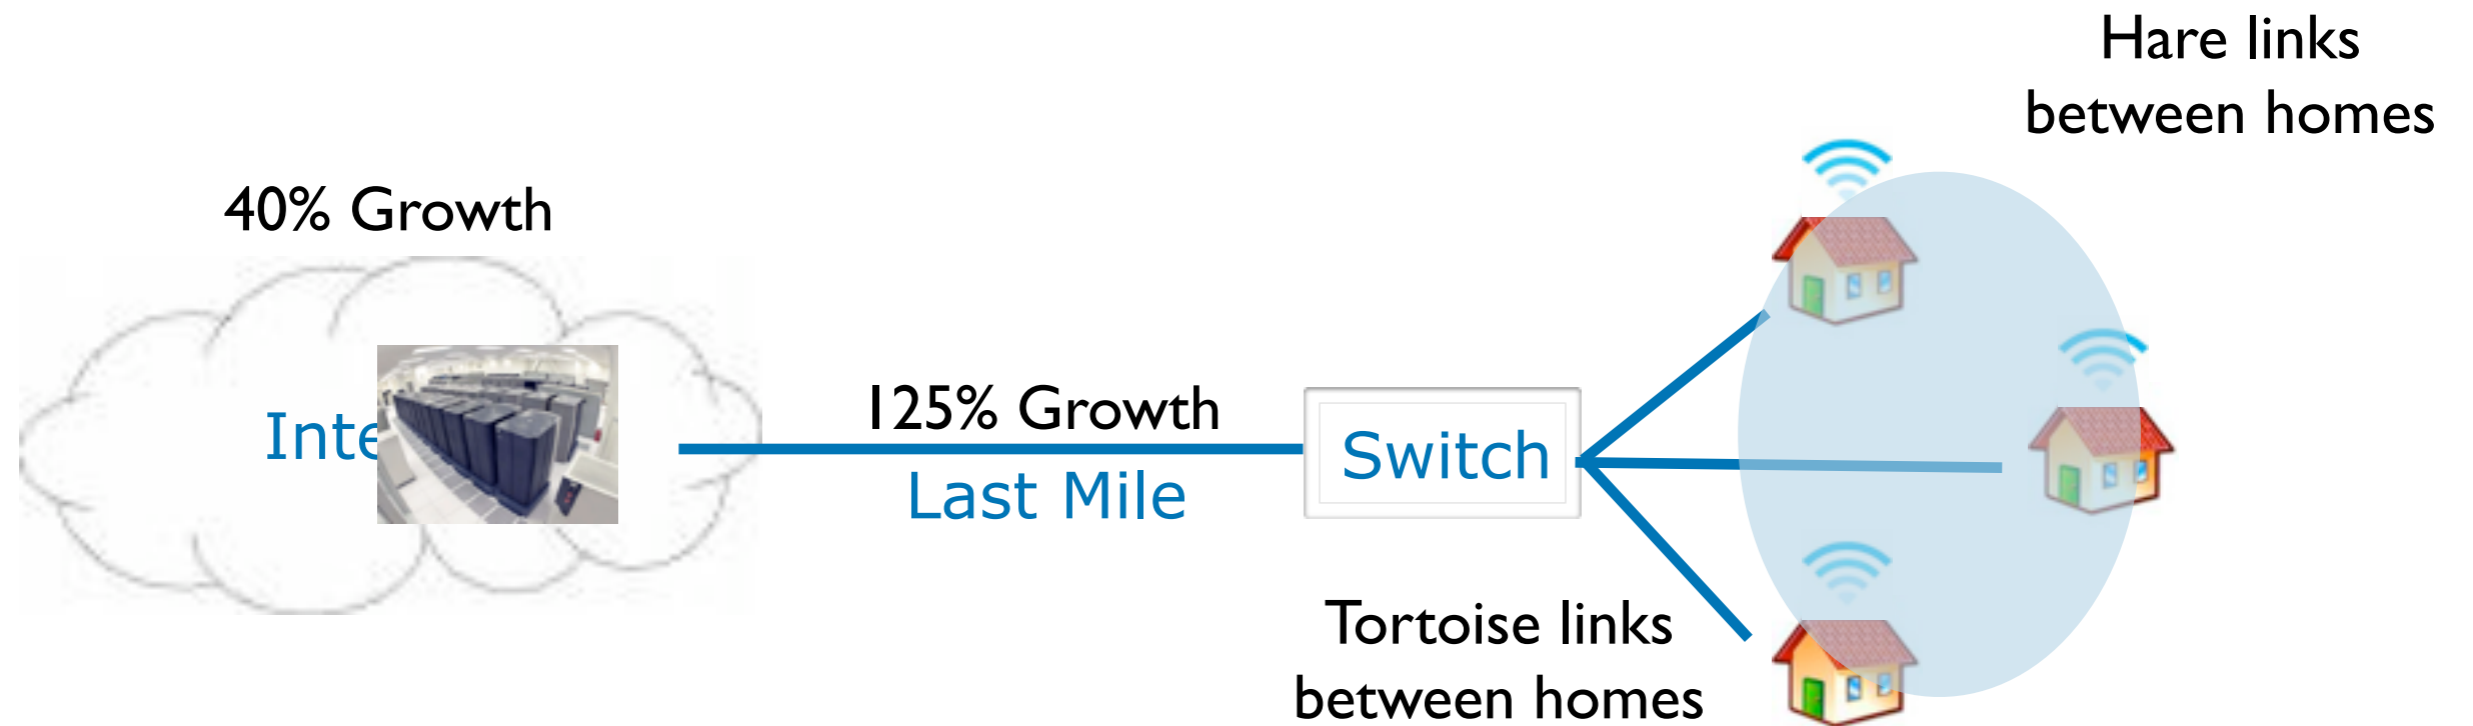
Neighborhood Aware Networks



Neighborhood Aware Networks



Neighborhood Aware Networks



Our Goal

Our Goal

- To decrease the effective loss rate on the hare channel by exploiting the reliability of the tortoise channel

Our Goal

- To decrease the effective loss rate on the hare channel by exploiting the reliability of the tortoise channel
- To produce a link whose bandwidth is more than the simple sum of the two links

Our Goal

- To decrease the effective loss rate on the hare channel by exploiting the reliability of the tortoise channel
- To produce a link whose bandwidth is more than the simple sum of the two links
- Focus on boosting hare link bandwidth by minimizing the overhead on the tortoise

Tackling Loss

Tackling Loss

- Partial packet recovery (PPR)
 - Send additional information to harness correctly delivered parts of a corrupted packet

Tackling Loss

- Partial packet recovery (PPR)
 - Send additional information to harness correctly delivered parts of a corrupted packet
- Bit-error rate based rate selection scheme
 - Send error-rate information at a granularity much smaller than a packet

Tackling Loss

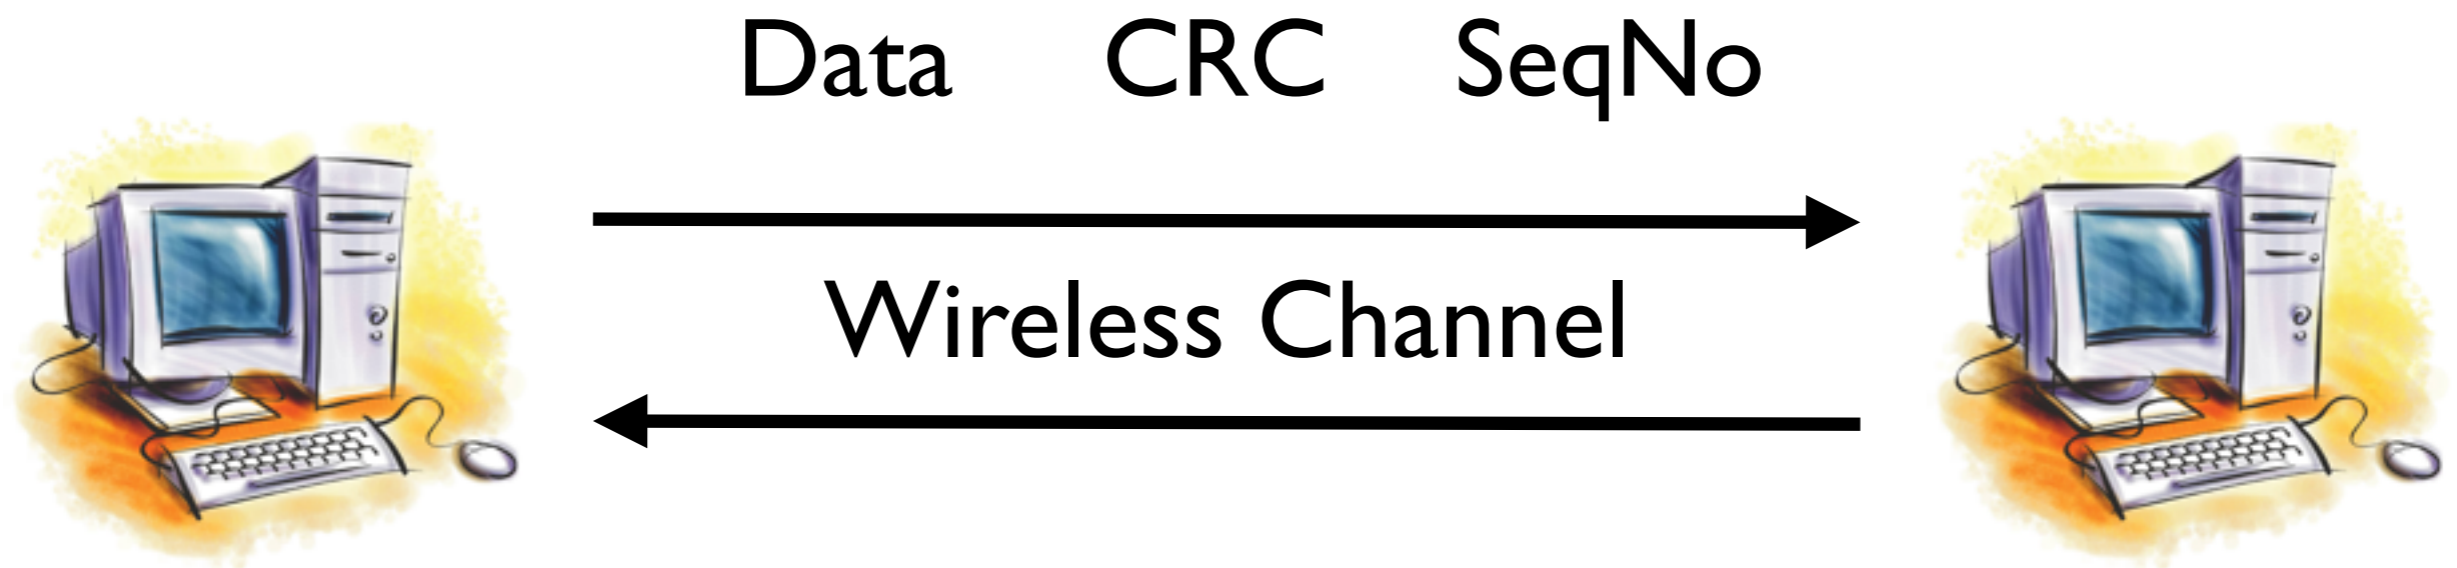
- Partial packet recovery (PPR)
 - Send additional information to harness correctly delivered parts of a corrupted packet
- Bit-error rate based rate selection scheme
 - Send error-rate information at a granularity much smaller than a packet
- Existing schemes for tackling wireless loss rely on:
 - Reliable delivery of metadata and control data
 - Timely delivery of receiver feedback to the sender

HTPPR Overview

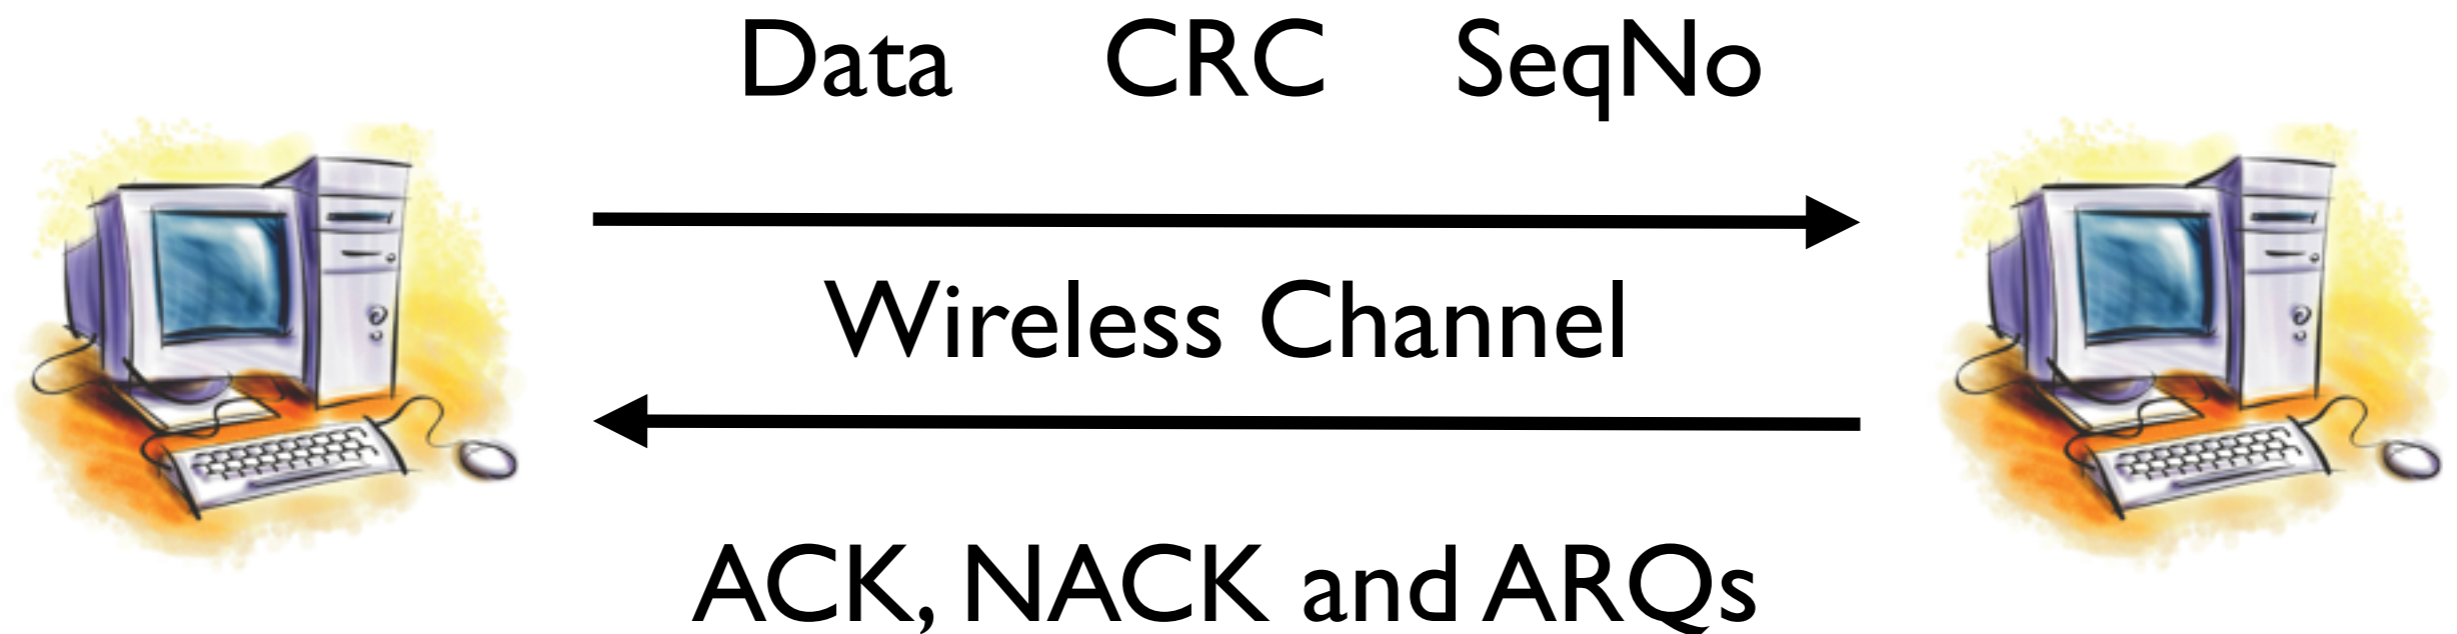
HTPPR Overview



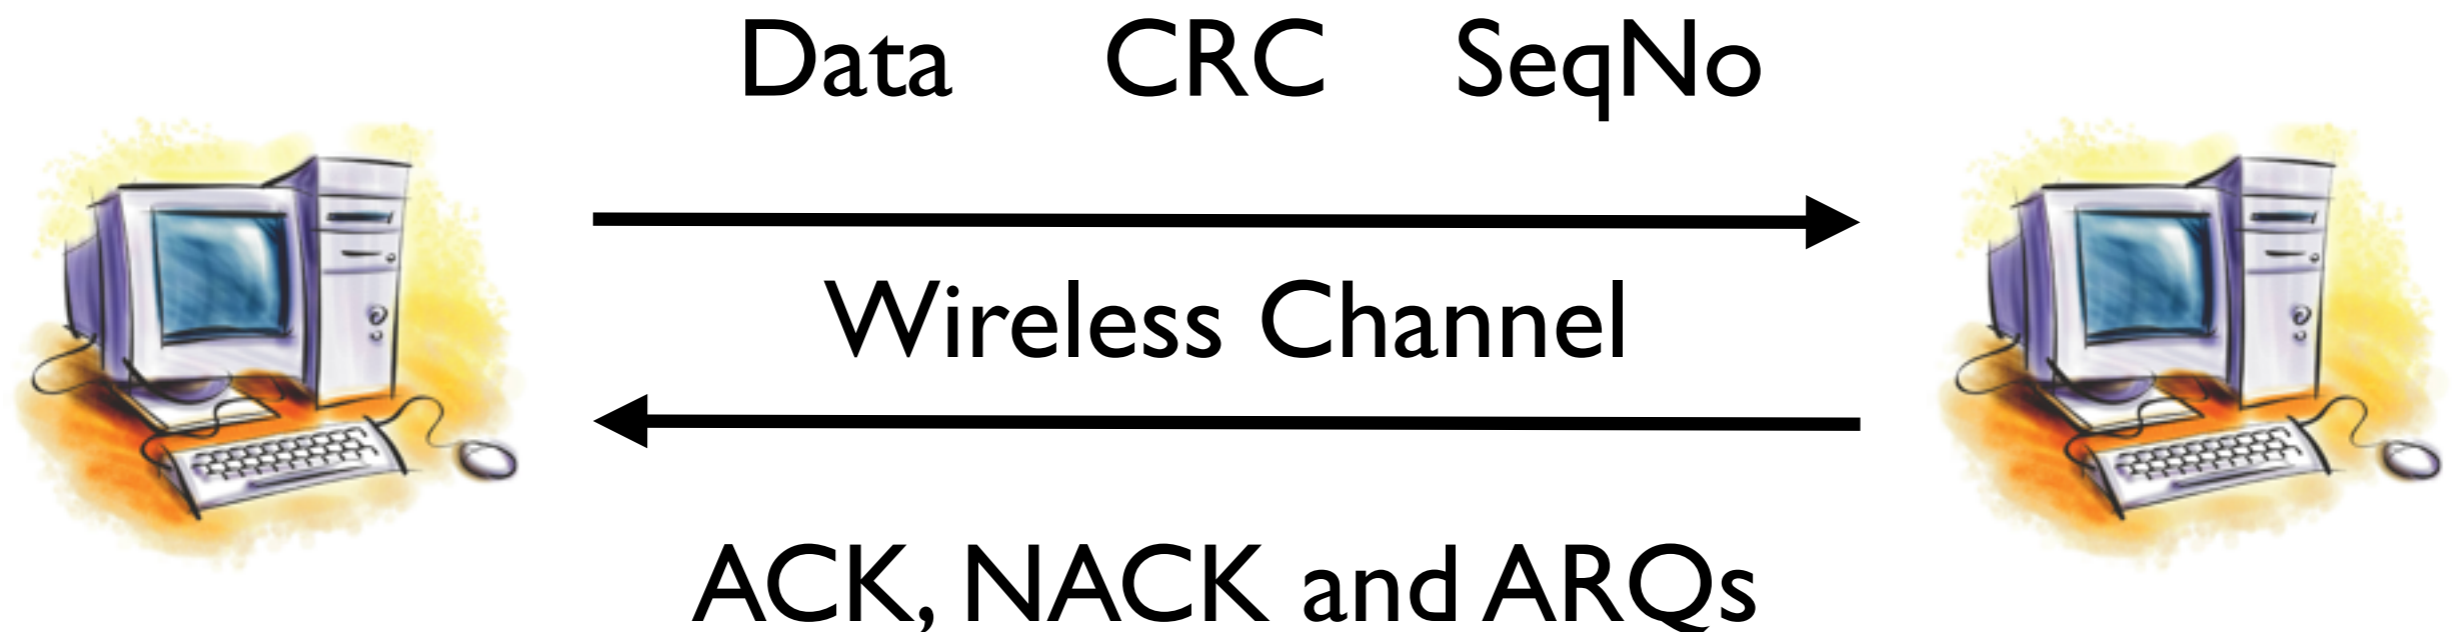
HTPPR Overview



HTPPR Overview



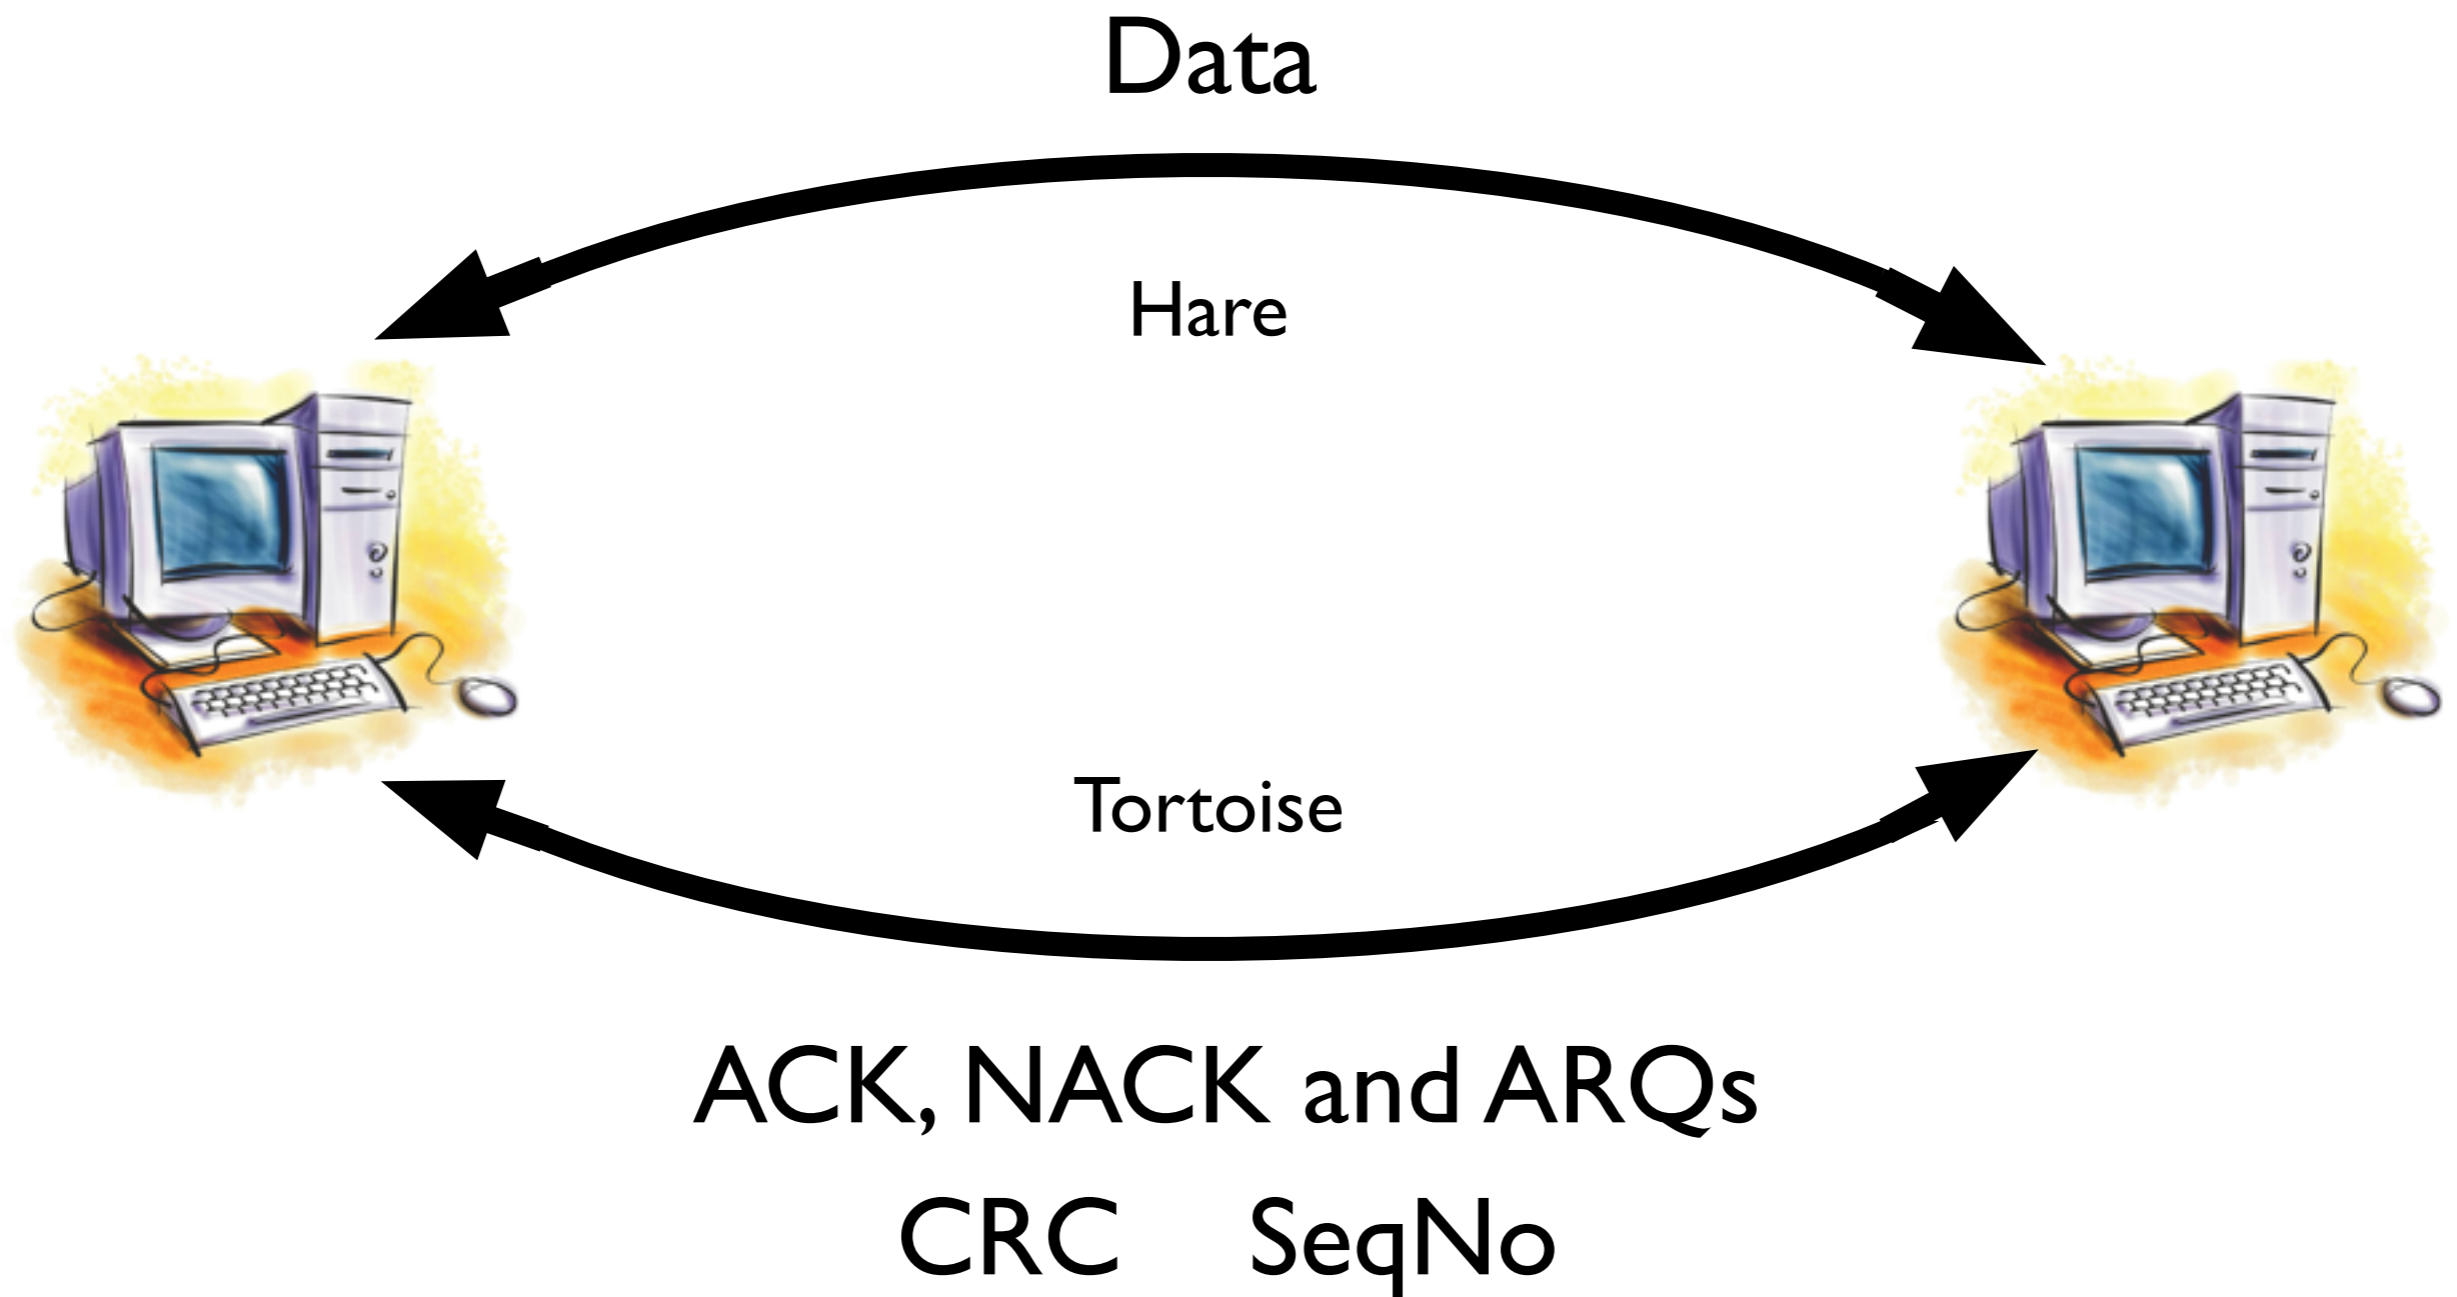
HTPPR Overview



Metadata is important for the PPR to work well


HTPPR Overview

HTPPR Overview





Partial Packet Recovery



Partial Packet Recovery



Partial Packet Recovery



Partial Packet Recovery



Wireless

Wired

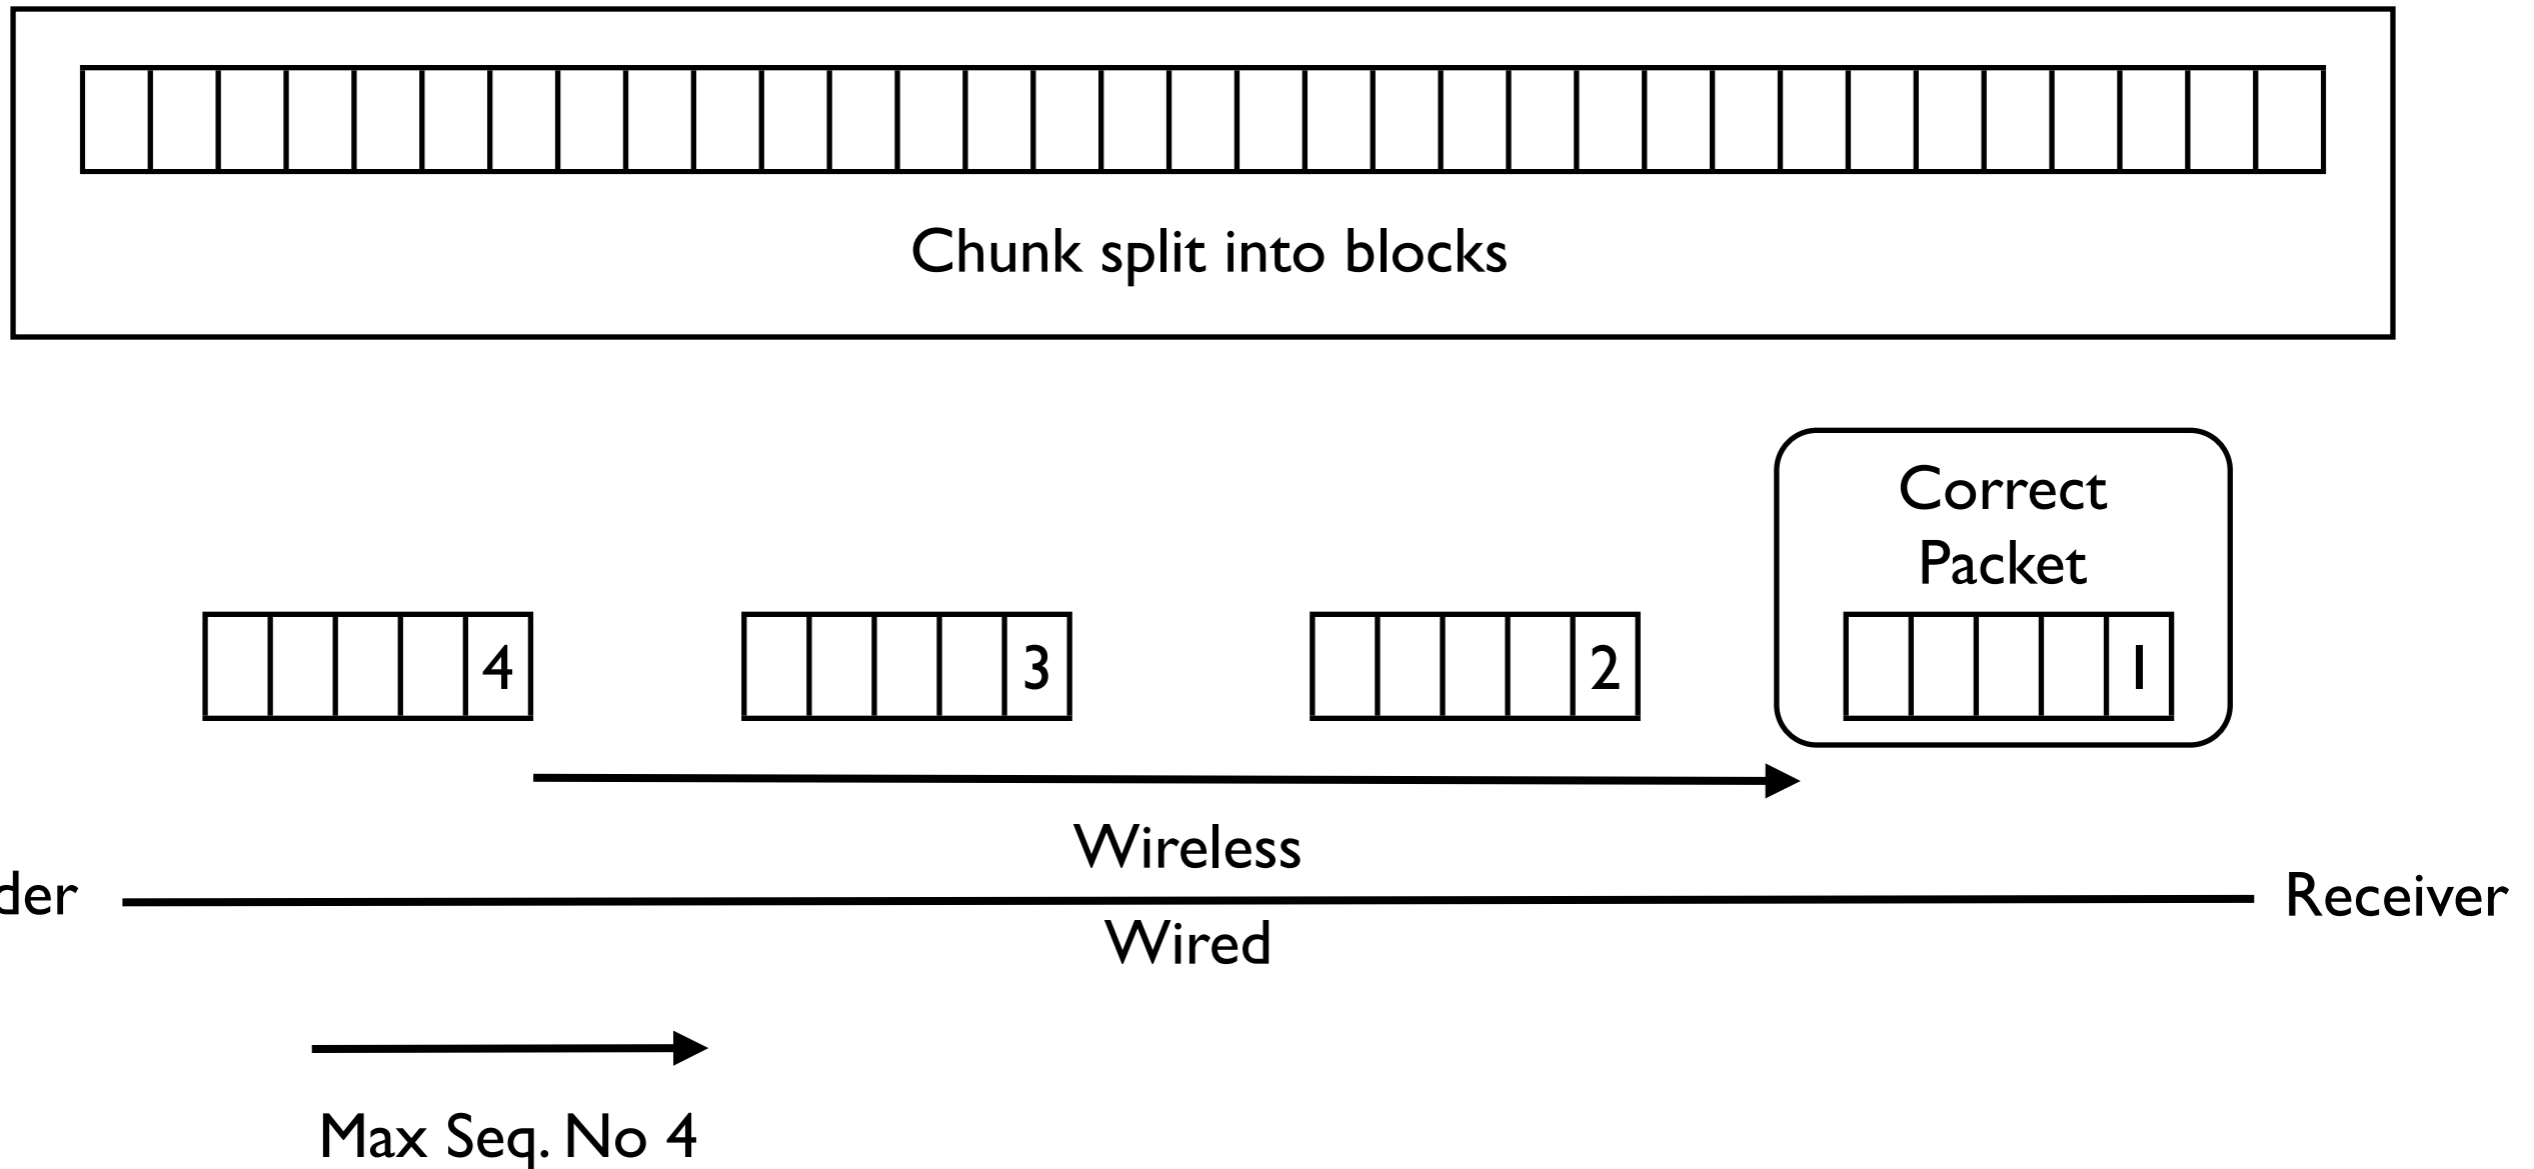
Sender

Receiver

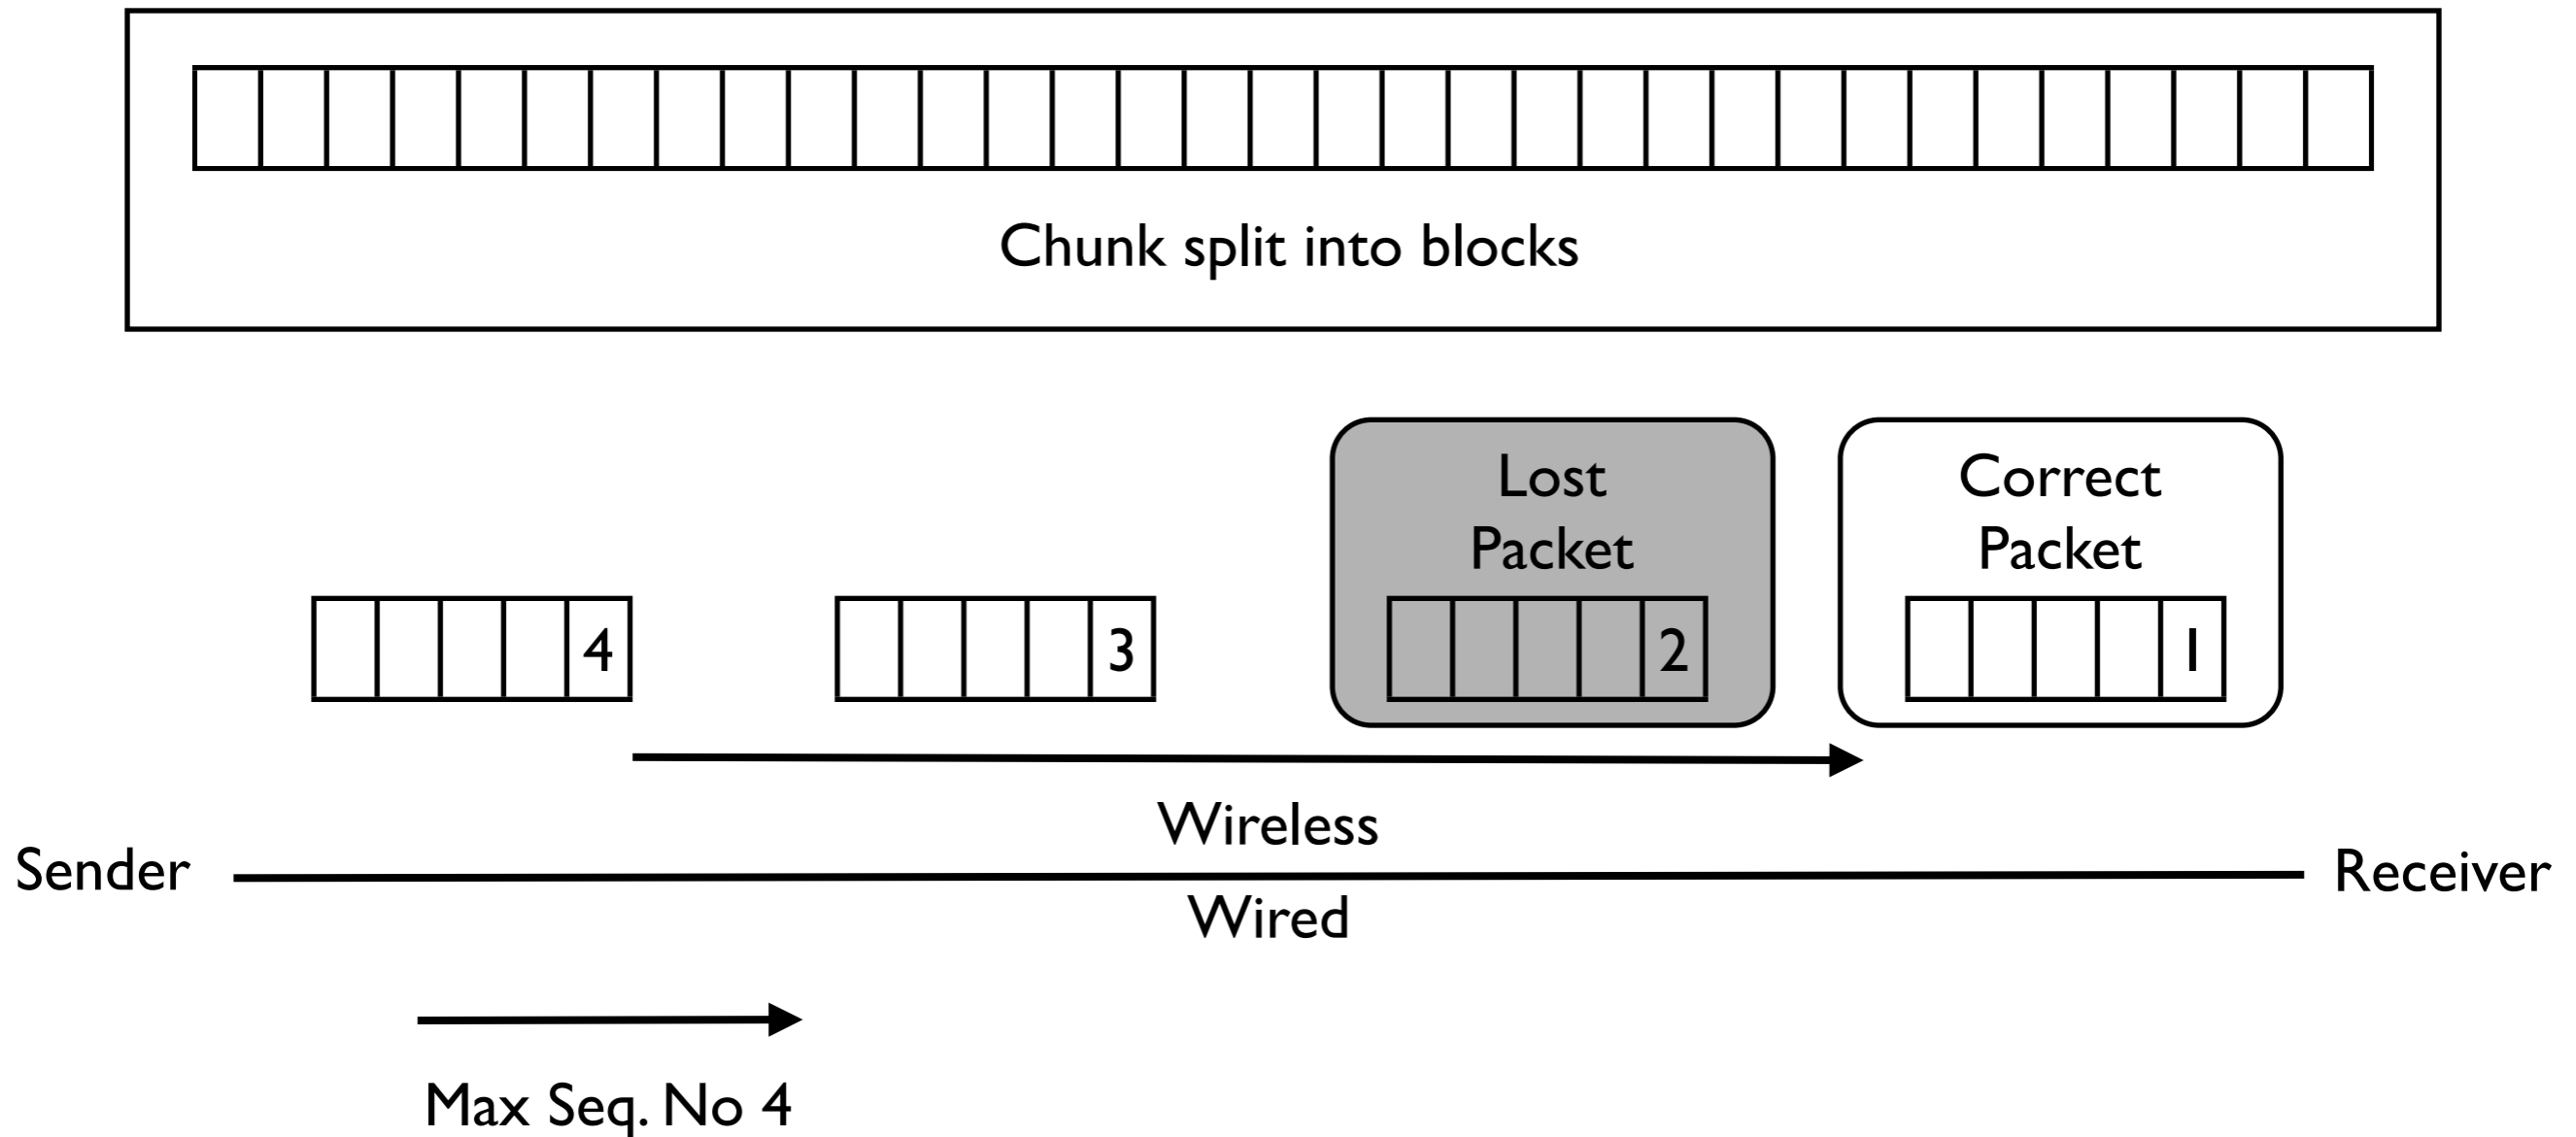


Max Seq. No 4

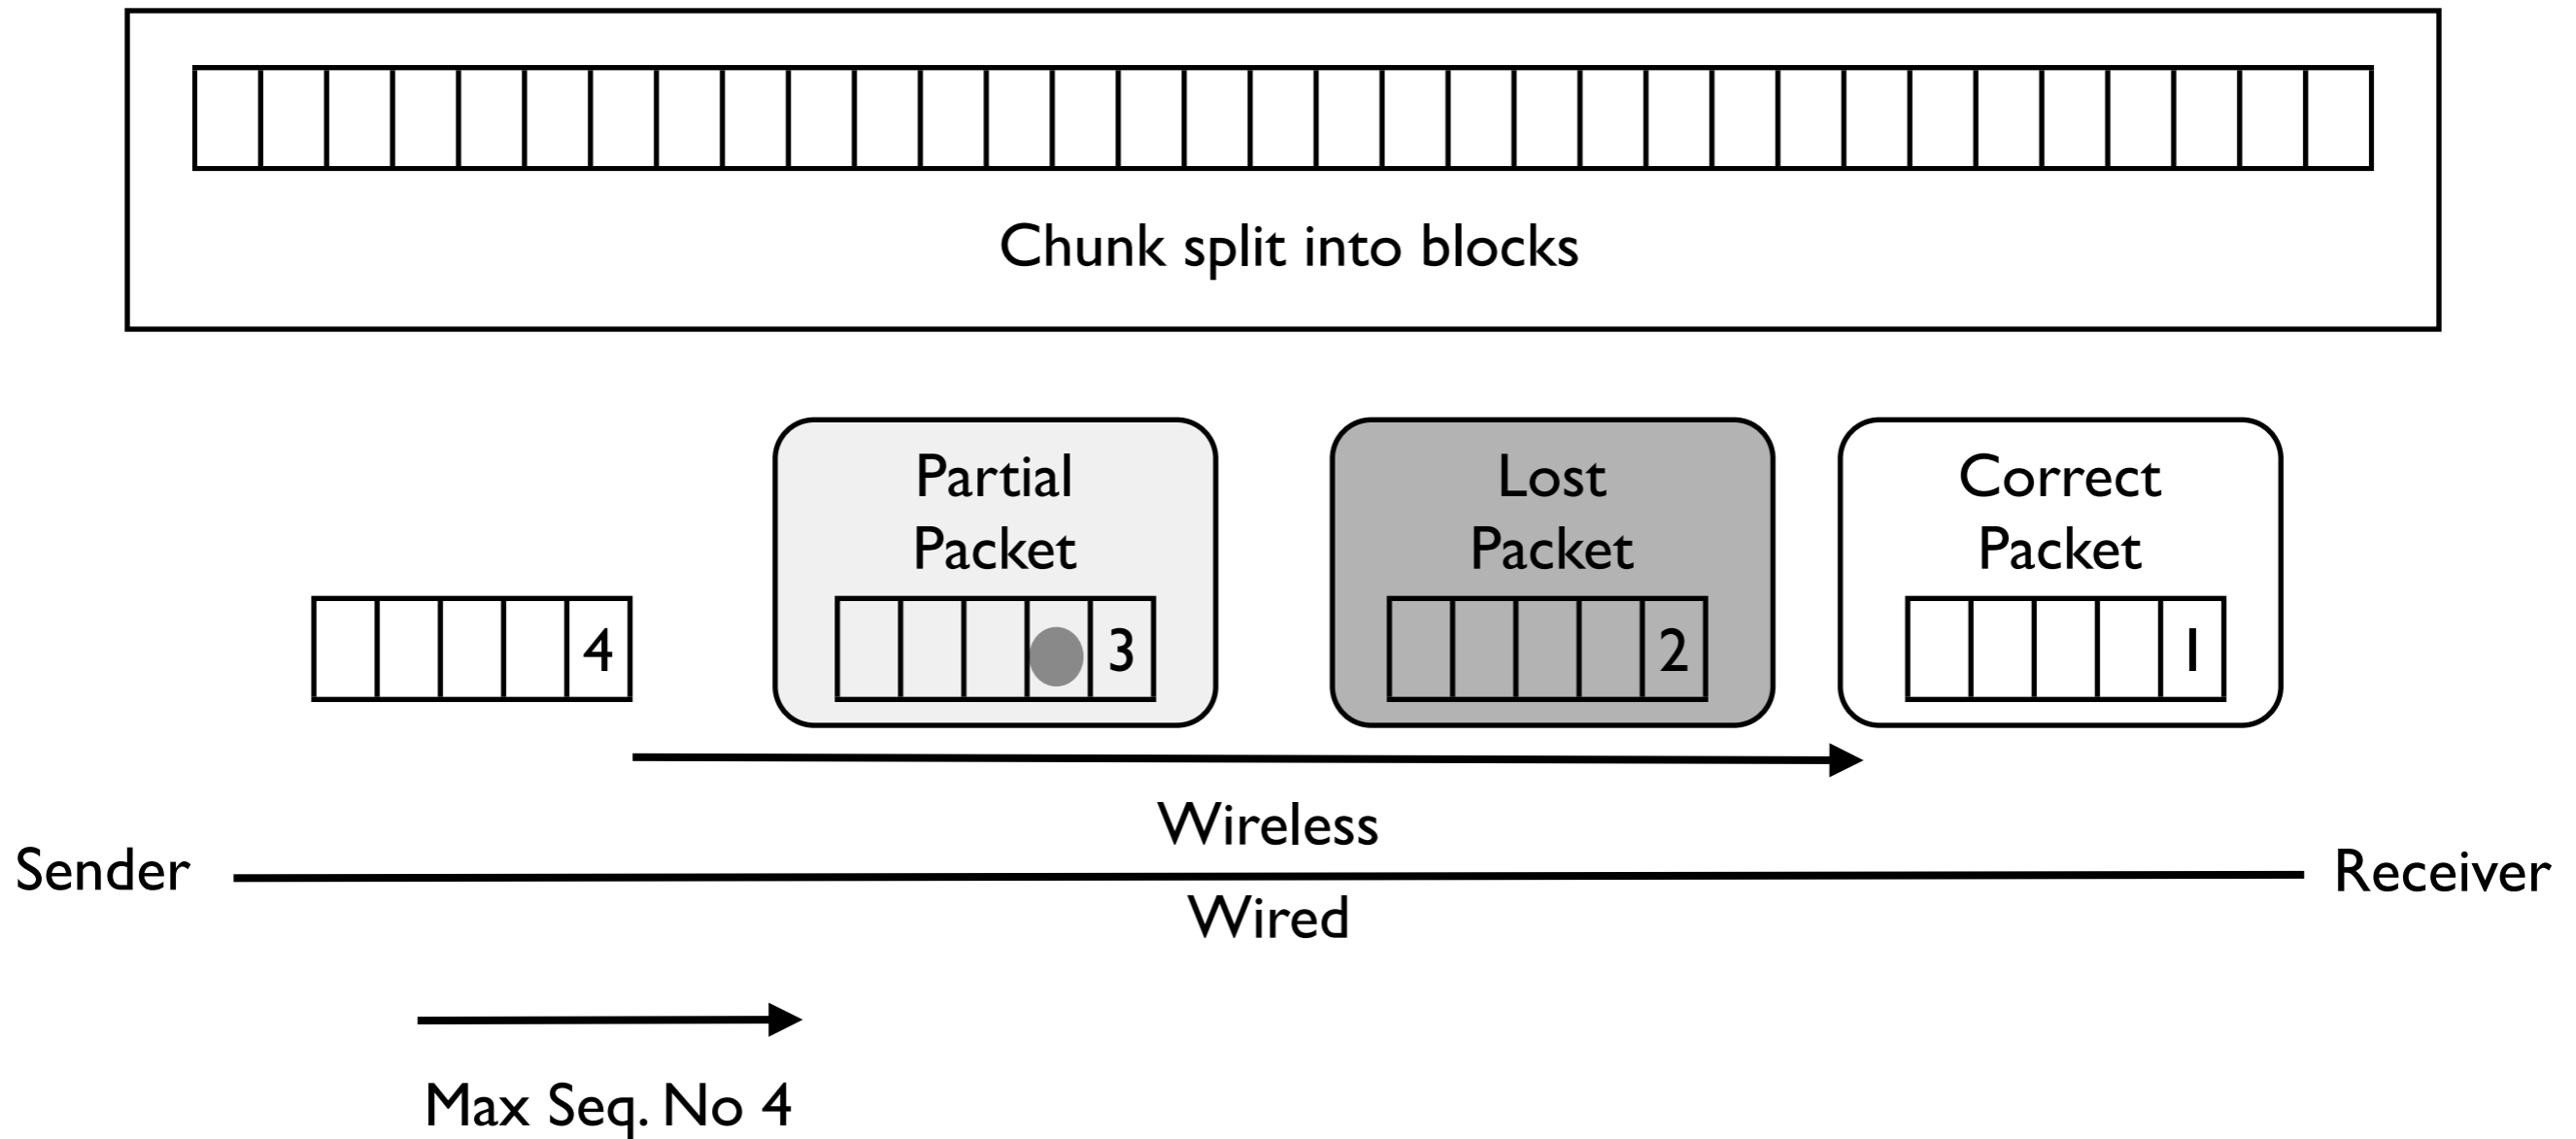
Partial Packet Recovery



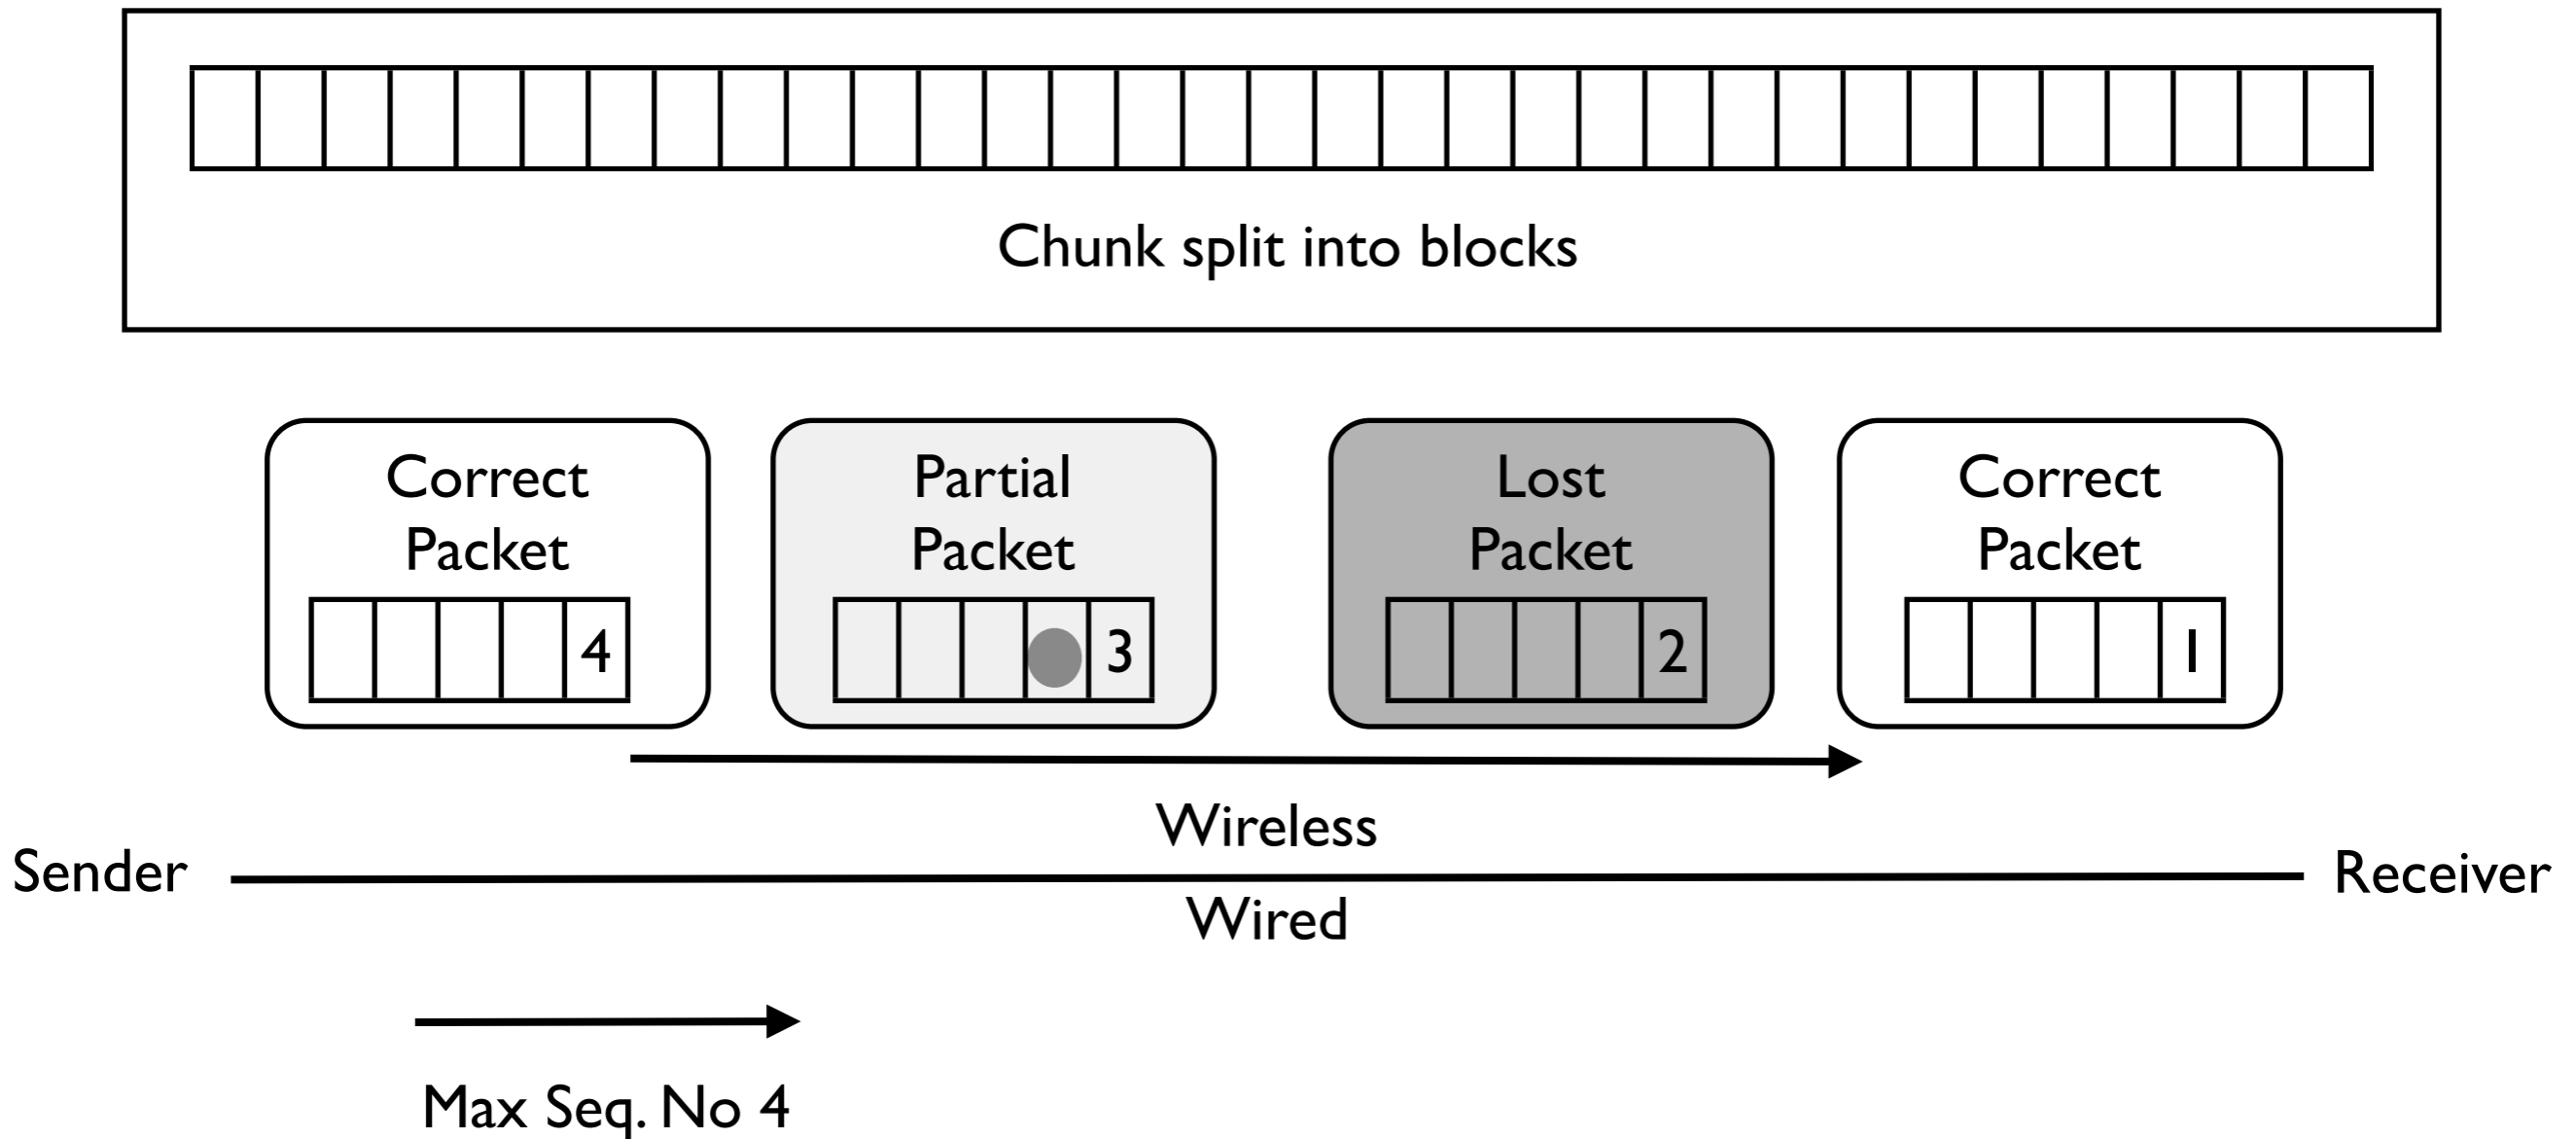
Partial Packet Recovery



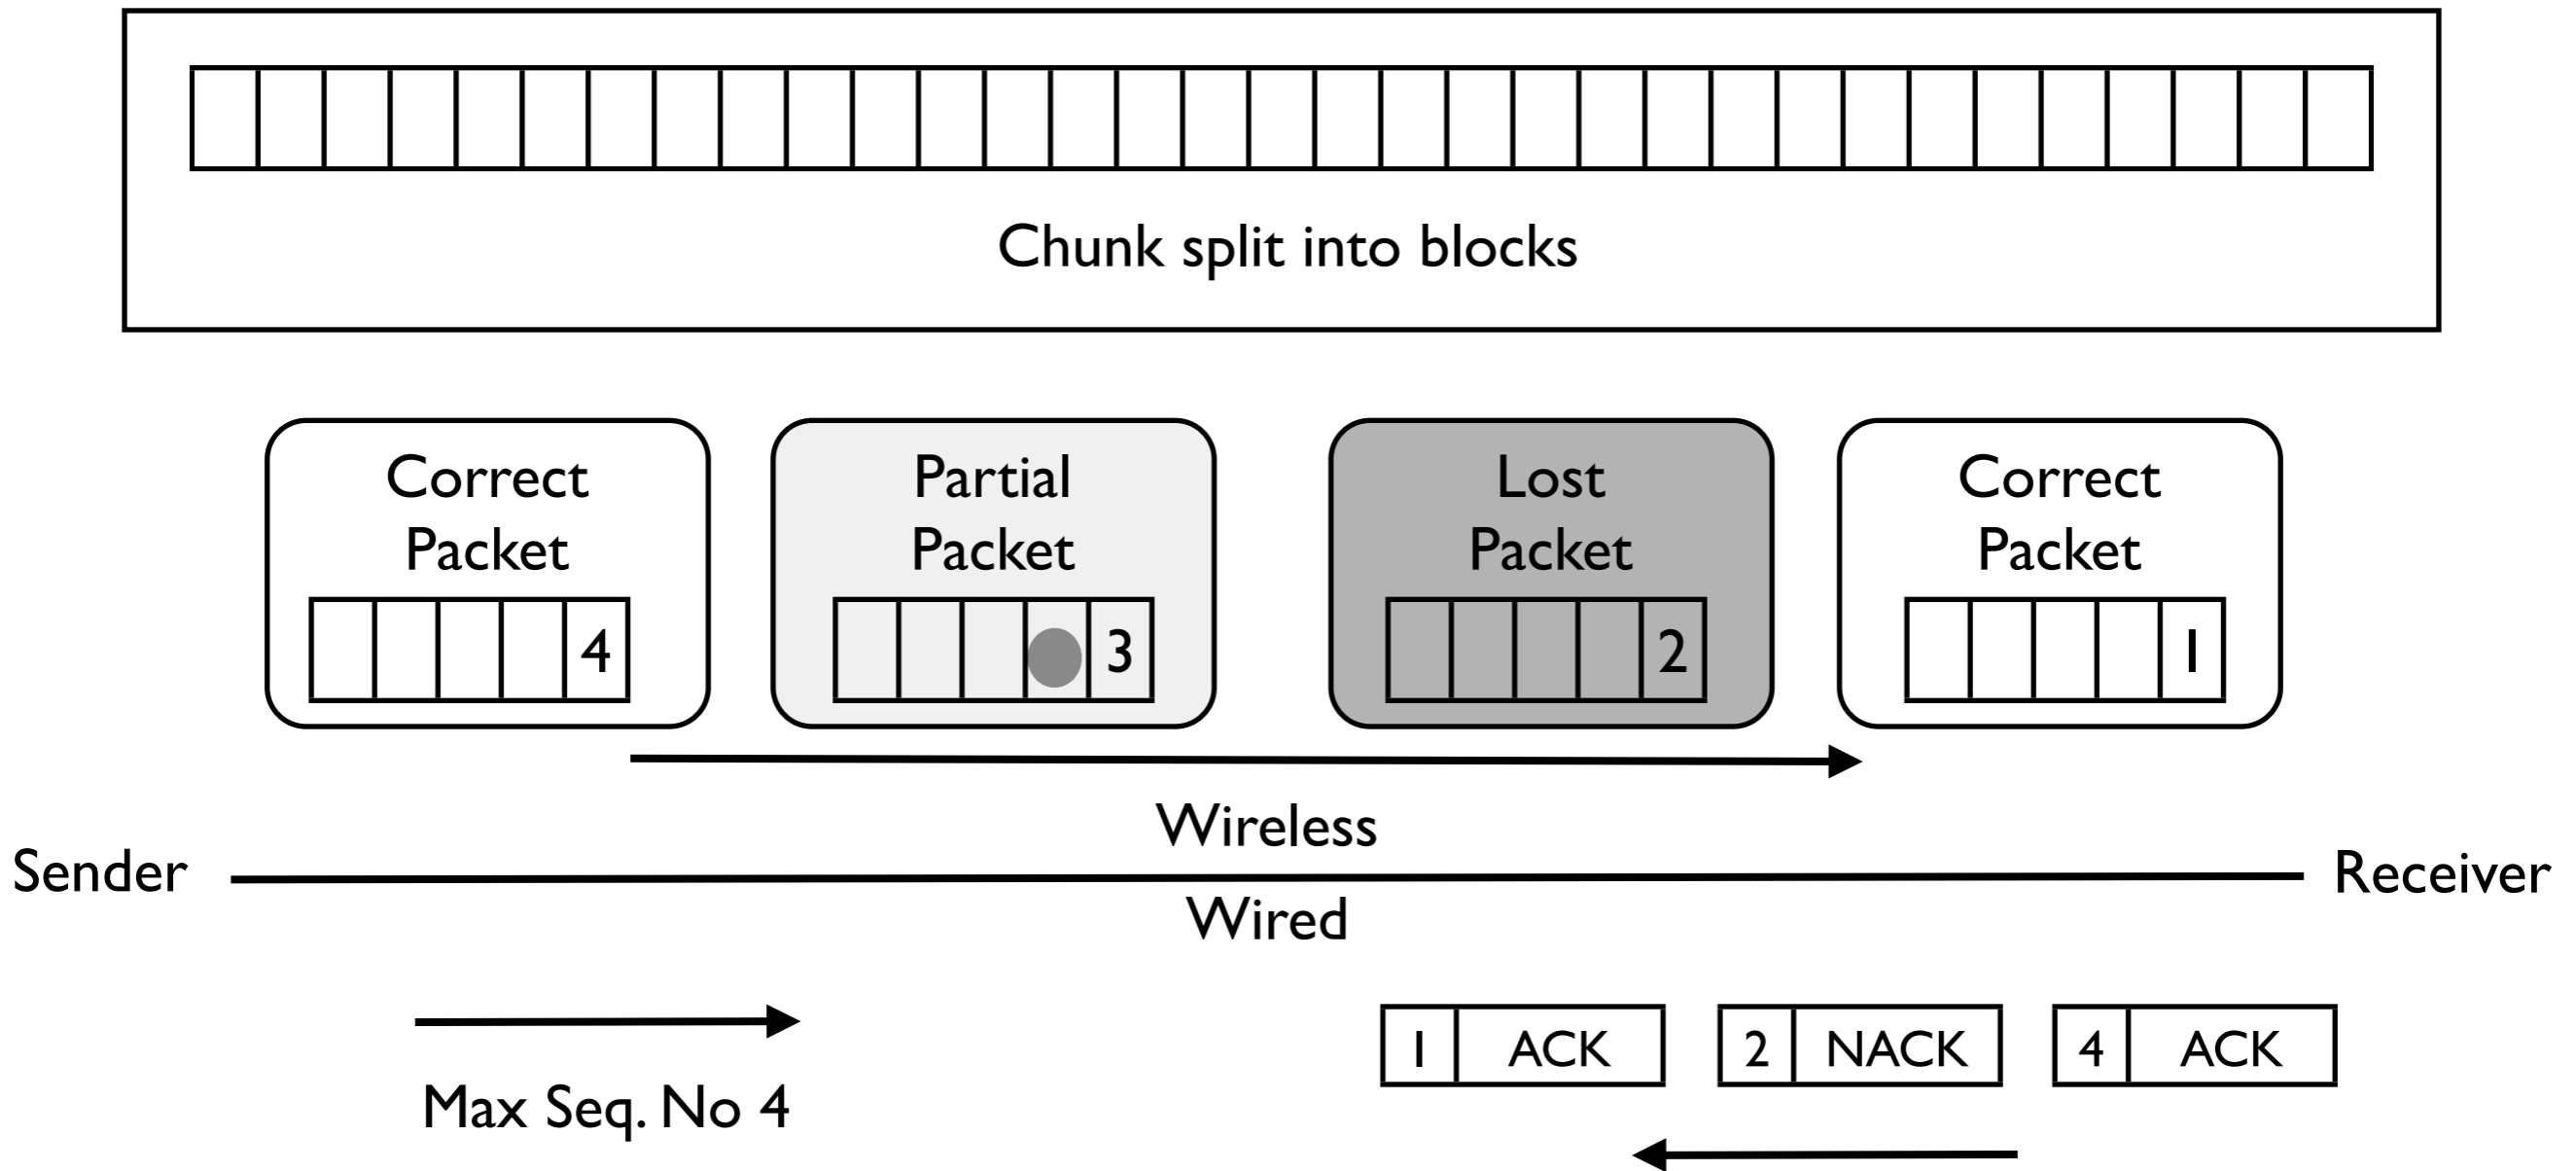
Partial Packet Recovery



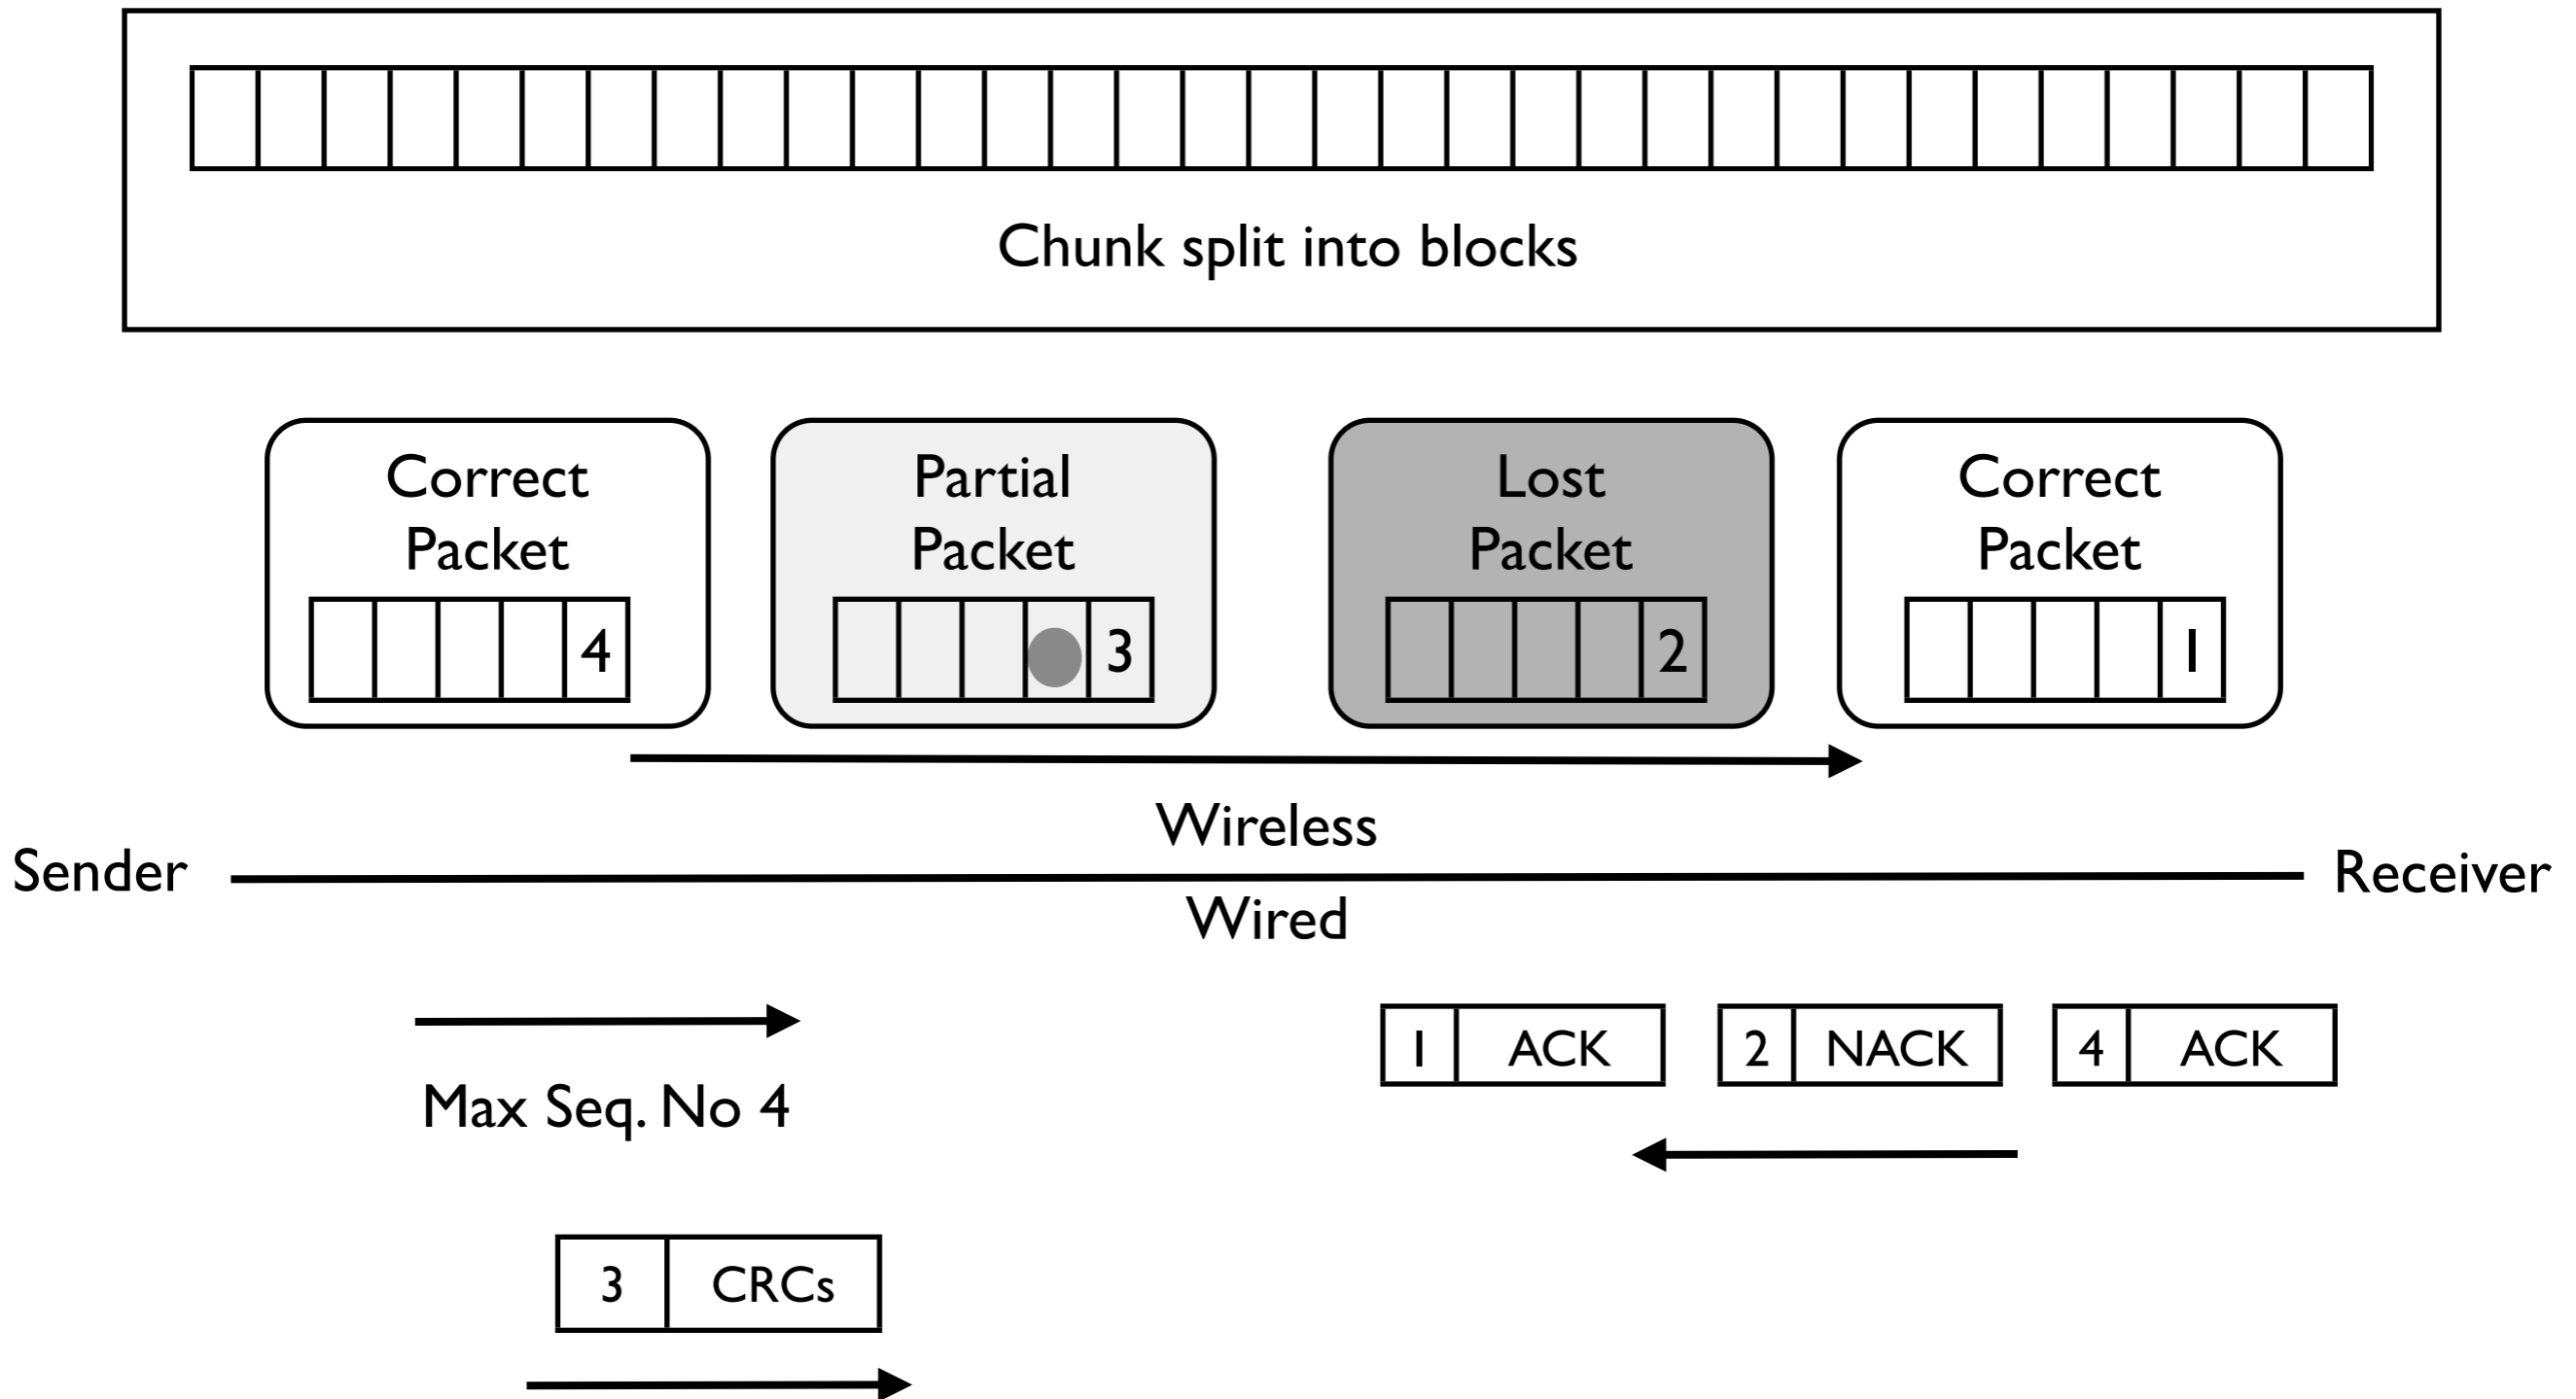
Partial Packet Recovery



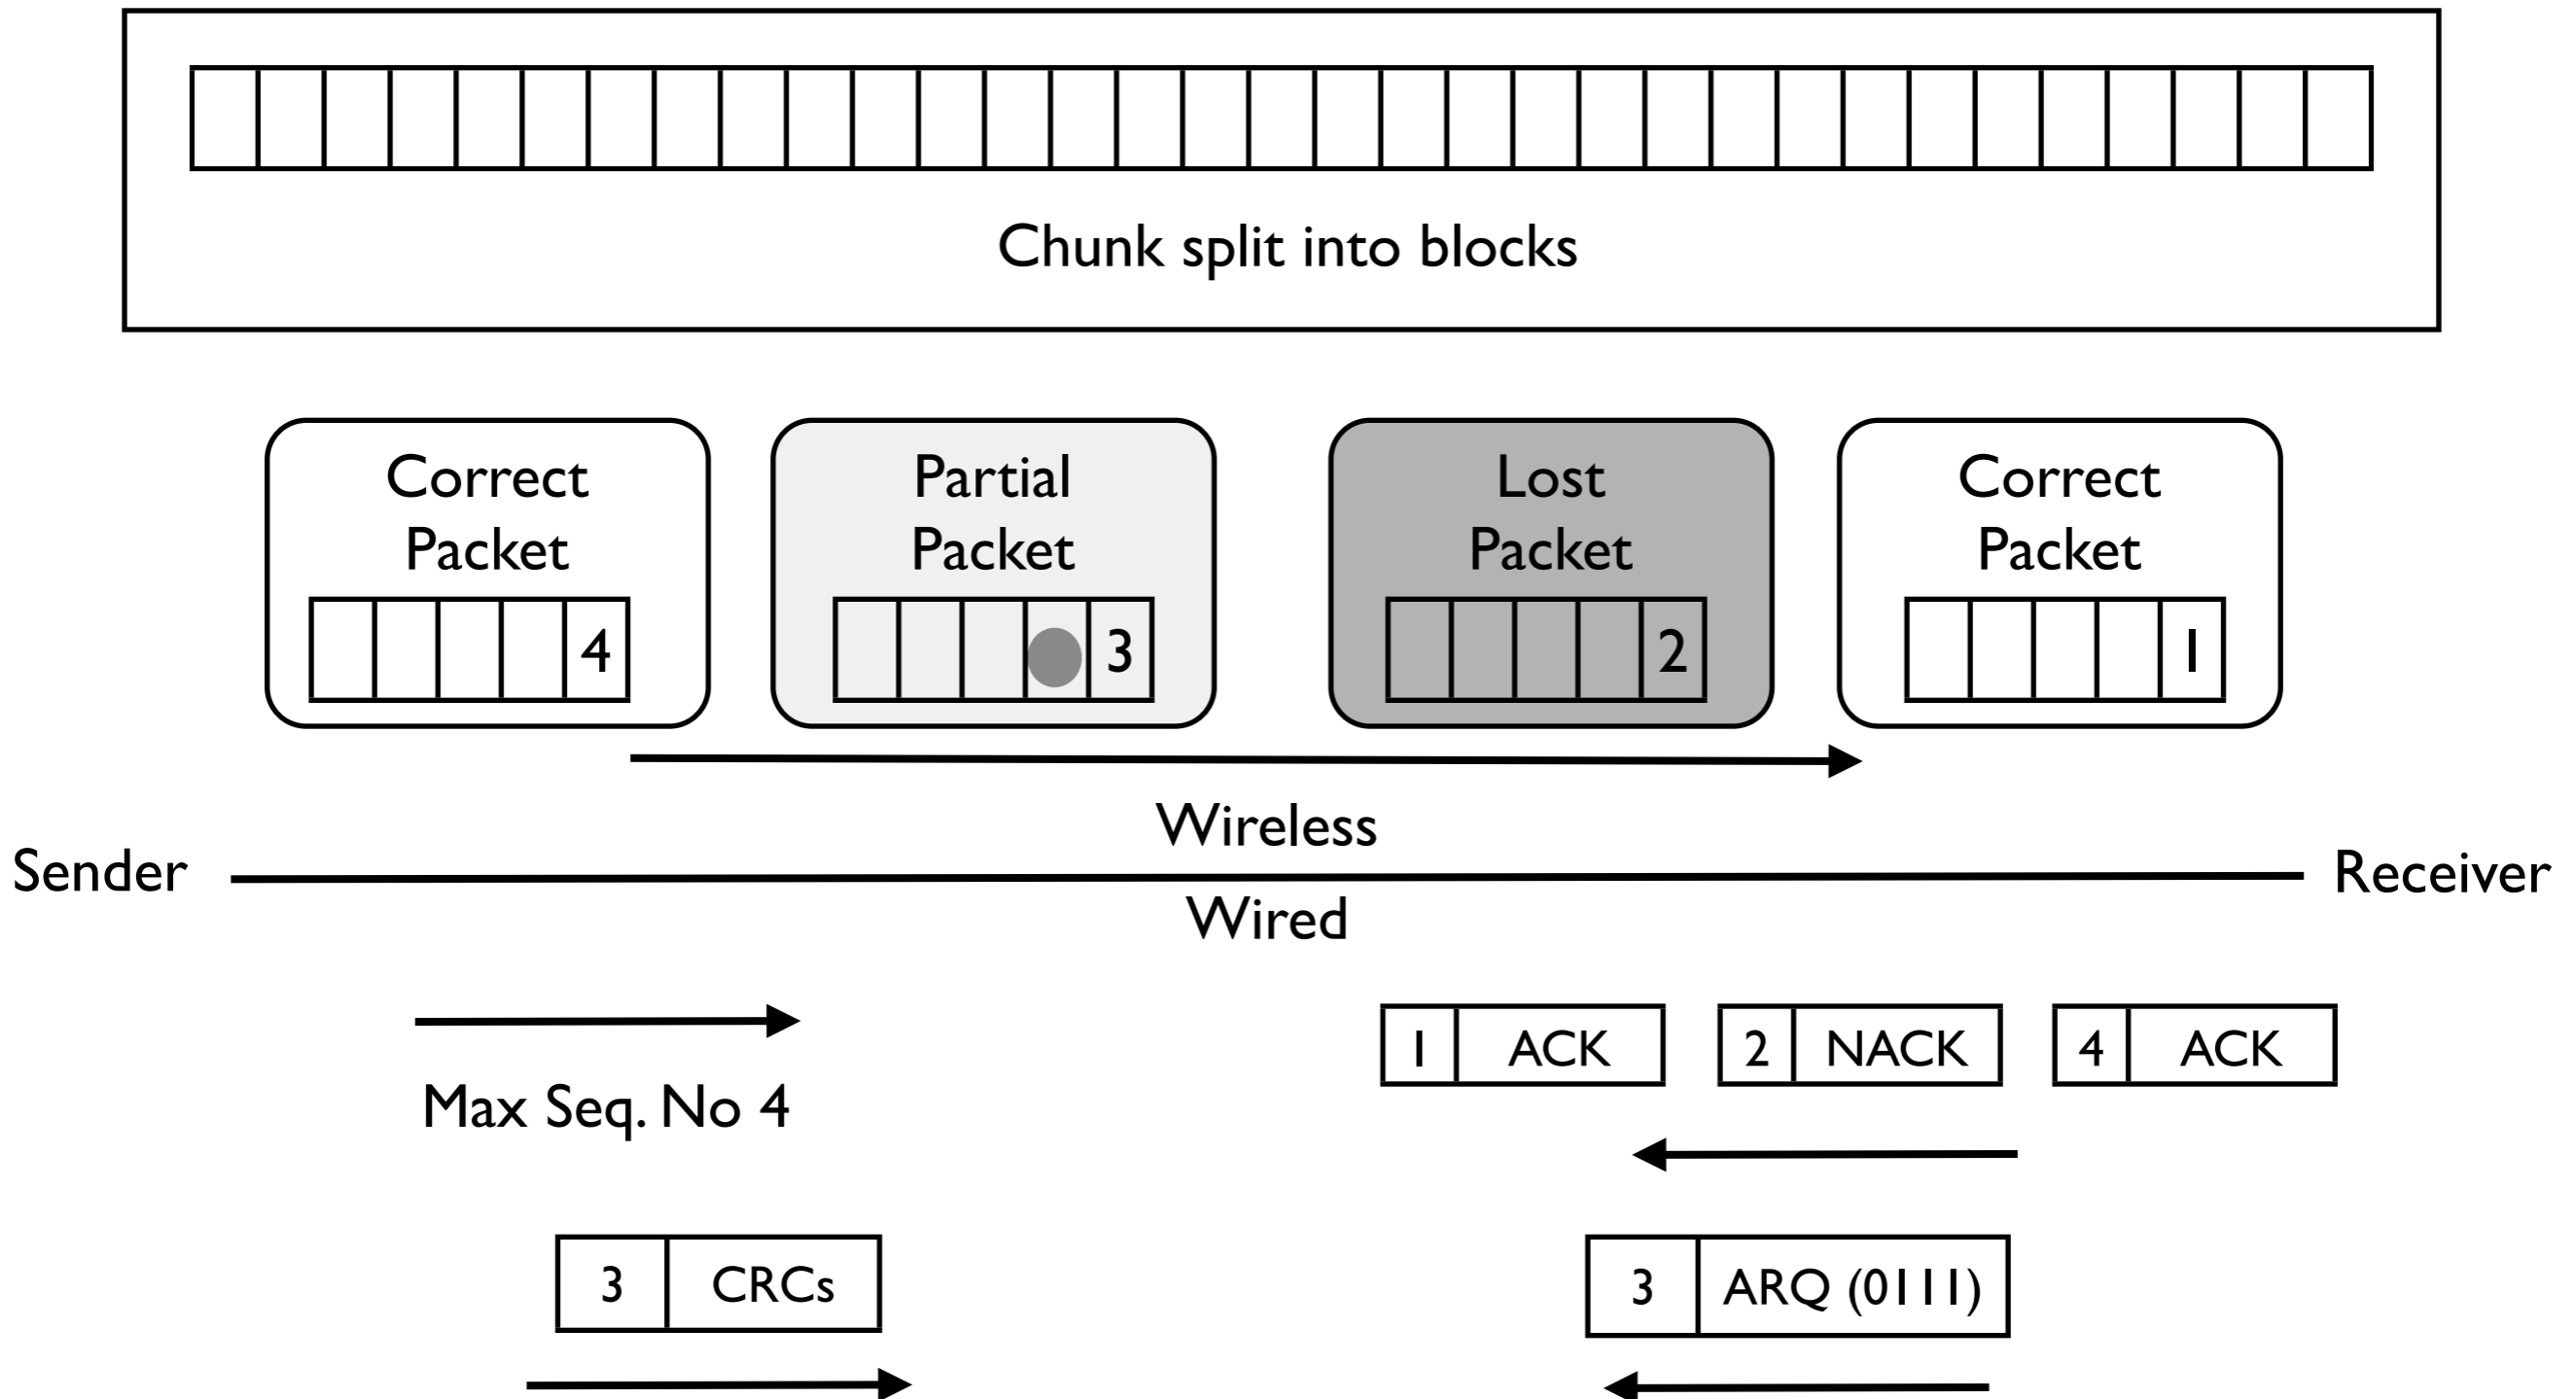
Partial Packet Recovery



Partial Packet Recovery



Partial Packet Recovery



Bitrate Adaptation

Bitrate Adaptation


- PPR beneficial when rate can be adapted to harness partial packets even at “higher than normal” bitrates
- Block-level error information can be used to adapt the bitrate (delivered reliably over the tortoise)

Bitrate Adaptation

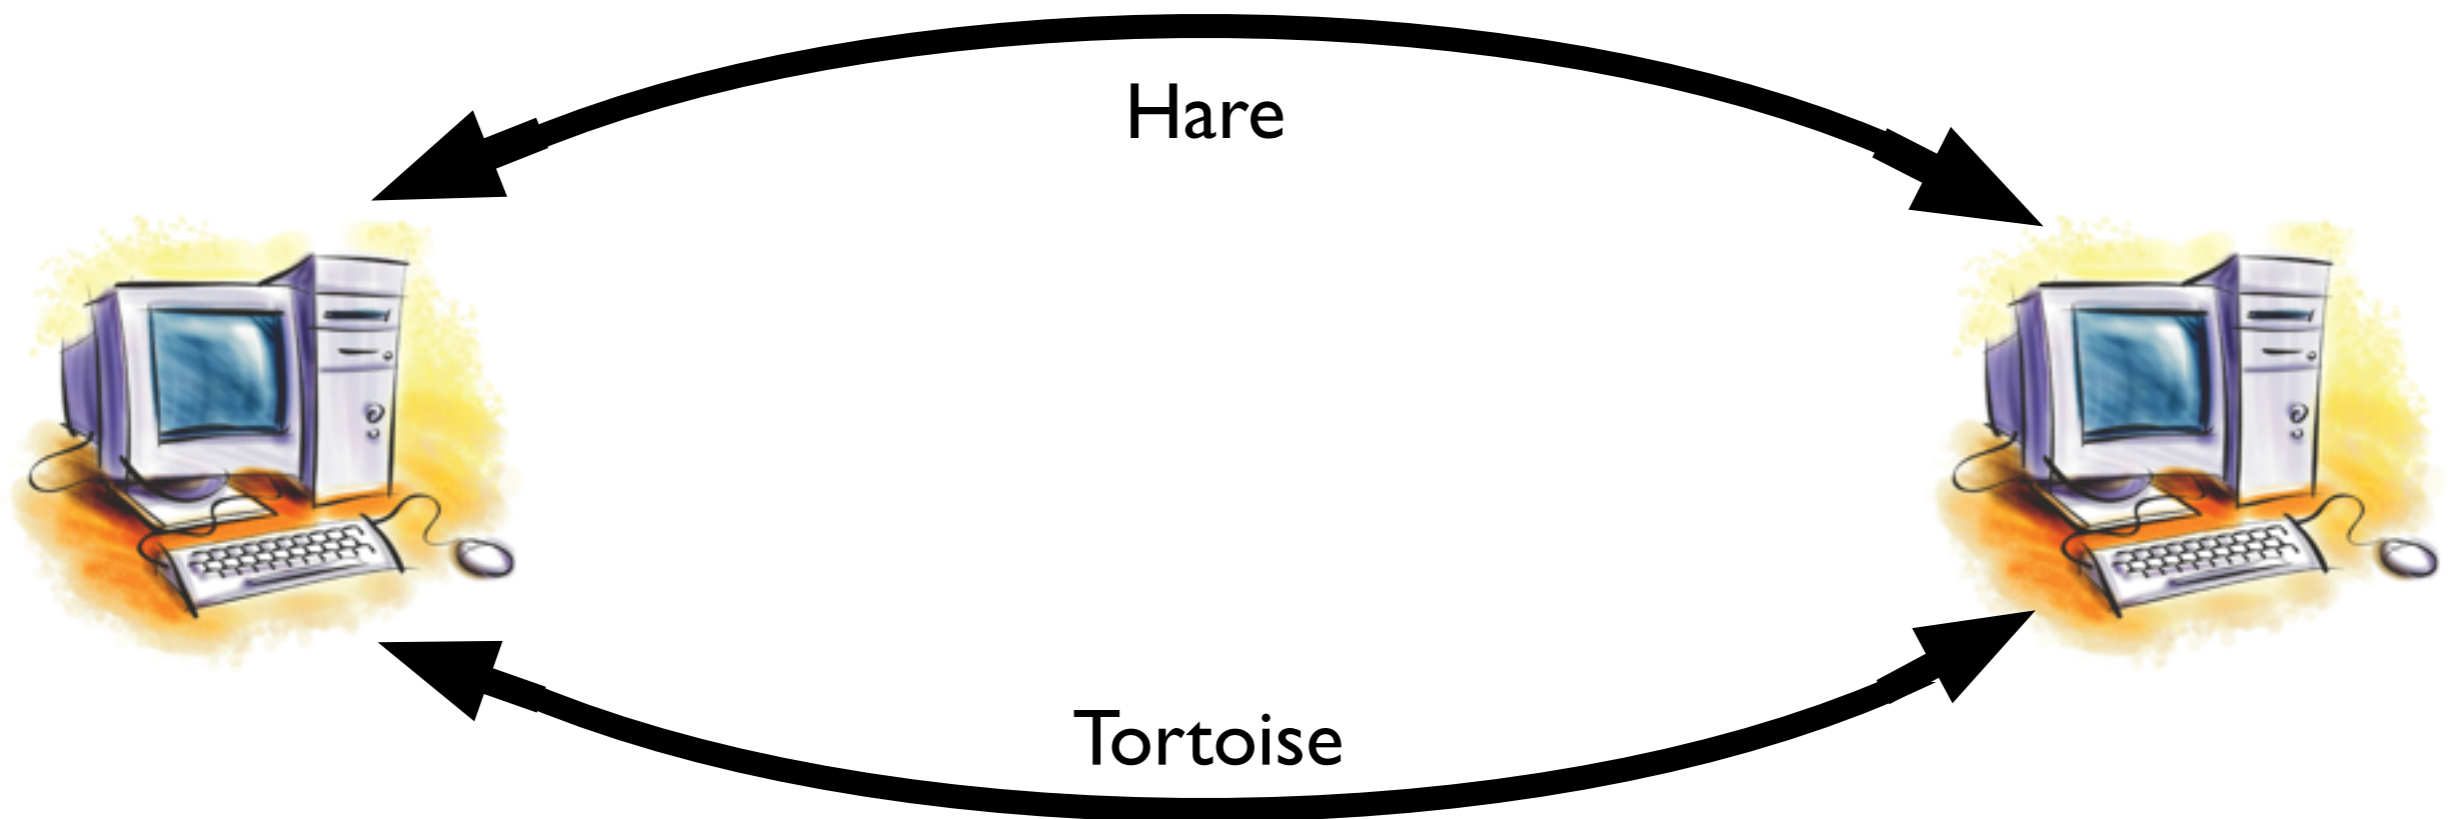
- PPR beneficial when rate can be adapted to harness partial packets even at “higher than normal” bitrates
- Block-level error information can be used to adapt the bitrate (delivered reliably over the tortoise)
- Exponentially weighted moving average of the block-level error rate information is used
 - More weight to recent updates
 - Does not require any hardware modifications unlike existing error-rate based rate adaptation schemes

Problem: Tortoise is Expensive

Problem: Tortoise is Expensive



Problem: Tortoise is Expensive



- Low capacity tortoise must be used effectively
- Reduce the footprint of the metadata and control data
- Main overhead source: 32 bit CRC per block of data

Solution: Multi-resolution CRC

Solution: Multi-resolution CRC

- Solution: multi-resolution CRC
 - One packet level CRC

Solution: Multi-resolution CRC

- Solution: multi-resolution CRC
 - One packet level CRC
 - Four sub-packet CRCs
 - Final block-level CRCs
- Increases latency but saves bandwidth on the tortoise
 - Our focus is mainly on static multimedia content

Further Reducing the CRC Overhead

Further Reducing the CRC Overhead

- CRCs for static content can be precomputed and preseeded
- Precomputed CRCs can then be stored uniformly across the neighborhood
- Computation and transfers can happen overnight for new and popular content

Further Reducing the CRC Overhead

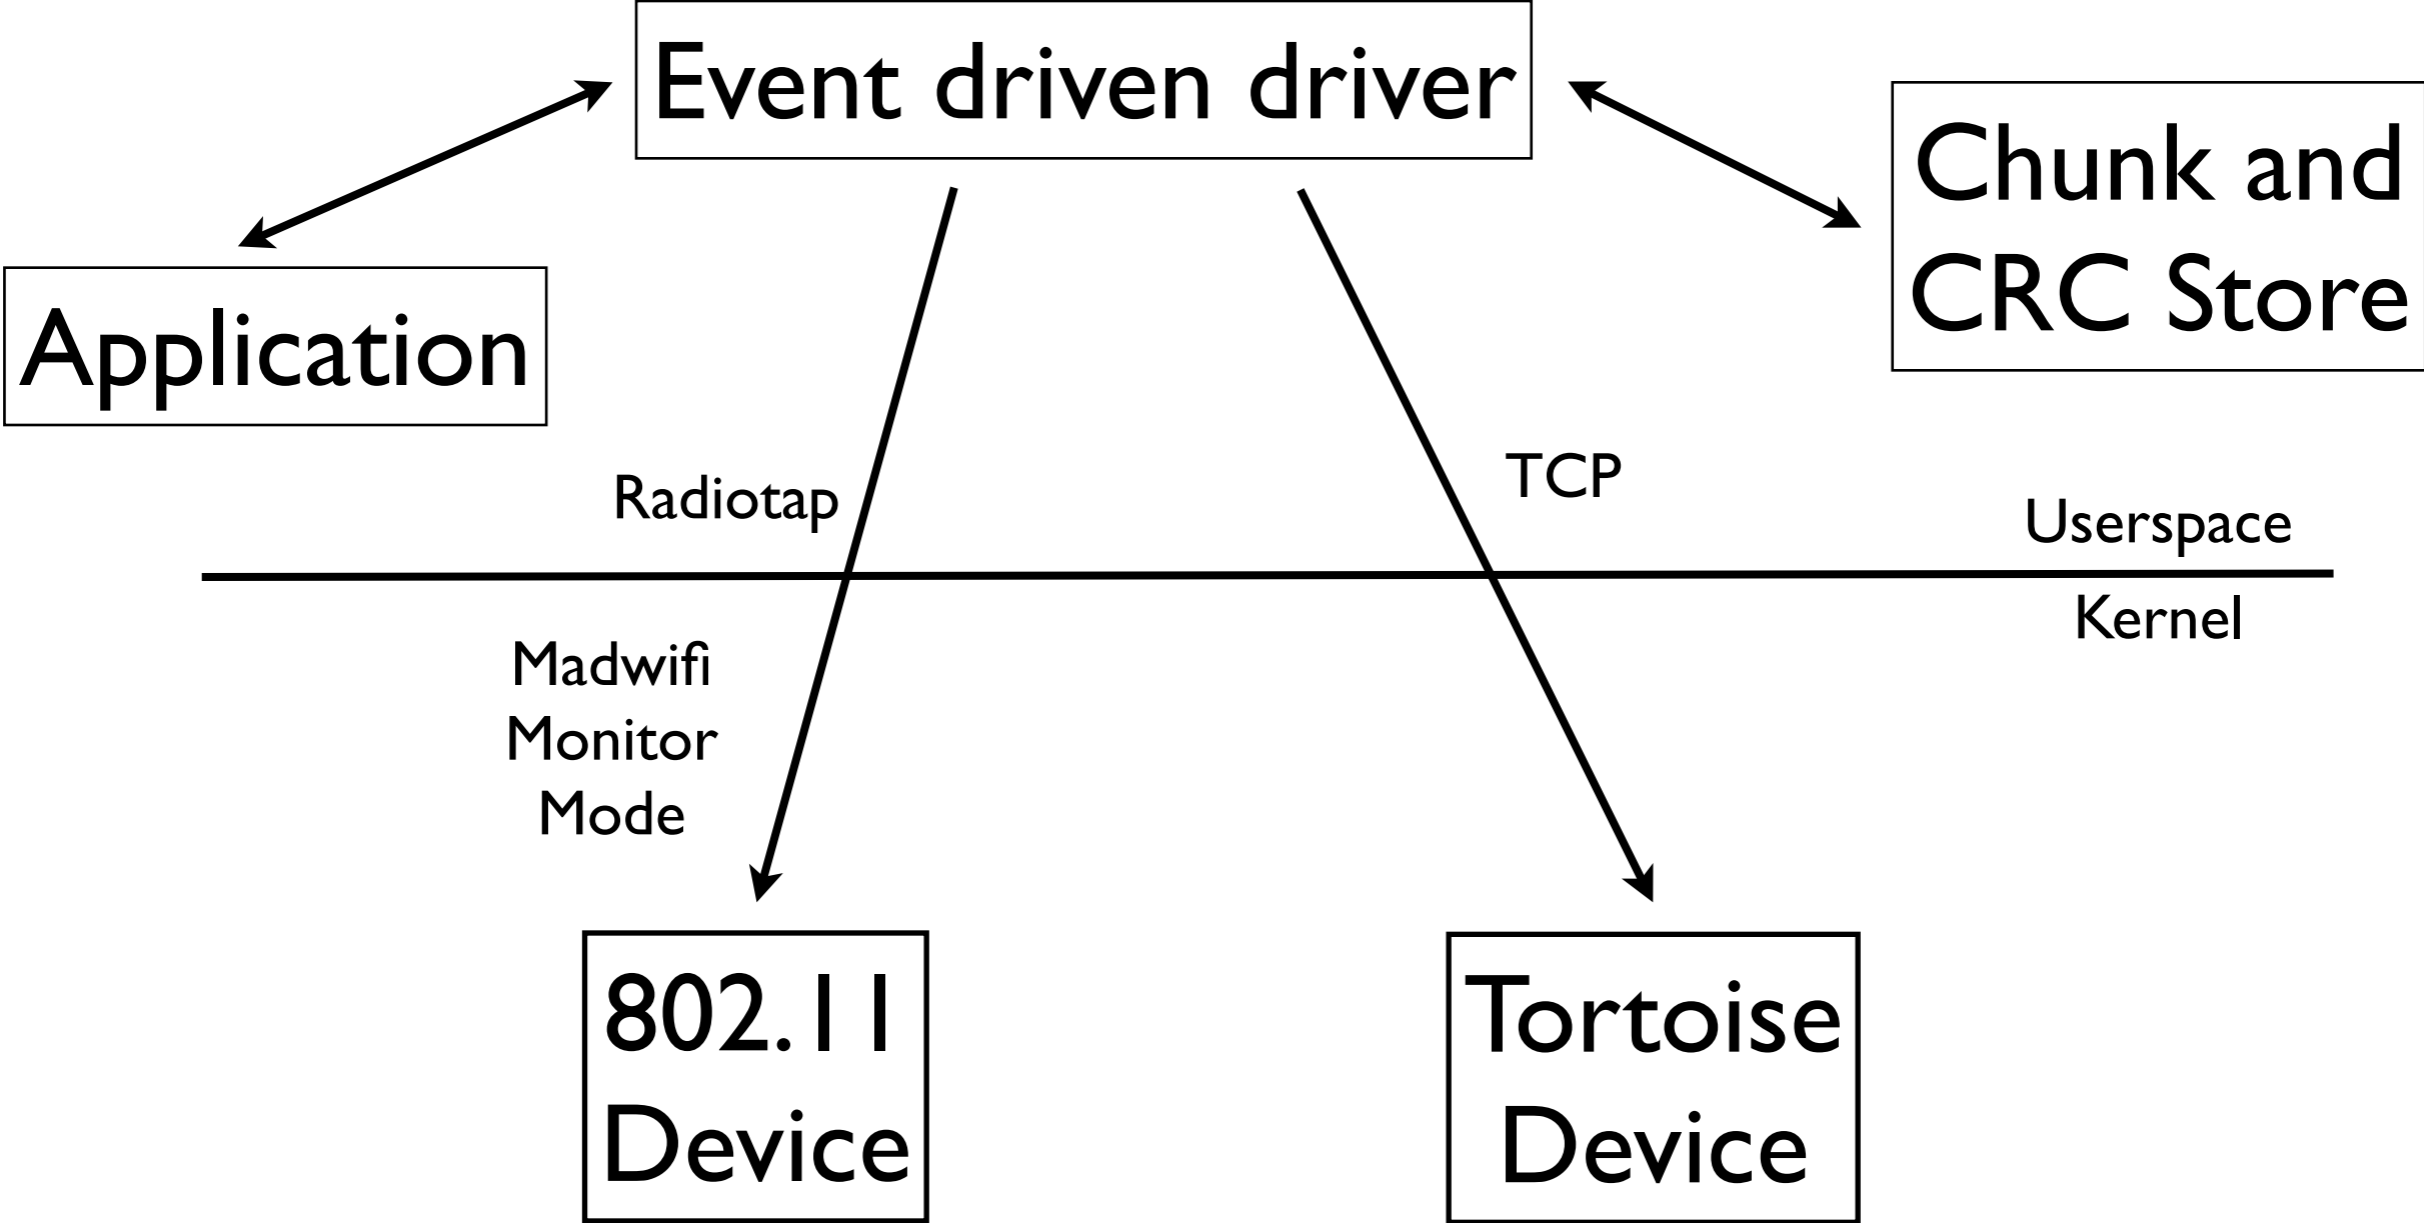
- CRCs for static content can be precomputed and preseeded
- Precomputed CRCs can then be stored uniformly across the neighborhood
- Computation and transfers can happen overnight for new and popular content
- Even one extra sender for CRCs will bring down the peak utilization of the senders by half

Further Reducing the CRC Overhead

- CRCs for static content can be precomputed and preseeded
- Precomputed CRCs can then be stored uniformly across the neighborhood
- Computation and transfers can happen overnight for new and popular content
- Even one extra sender for CRCs will bring down the peak utilization of the senders by half
- Important to reduce peak uplink utilization because tortoise links are the main internet connections for the neighborhood

Implementation

11,000+ lines of code (C++)




HTPPR vs Other Protocols

HTPPR vs Other Protocols


- HTPPR vs TCP vs Hop
 - Hop is a block based wireless transport protocol designed to reduce control traffic
 - TCP is Cubic with SampleRate
 - Tested 25 links of varying SNRs in a lab

HTPPR vs Other Protocols


HTPPR vs Other Protocols



HTPPR vs Other Protocols



HTPPR vs Other Protocols



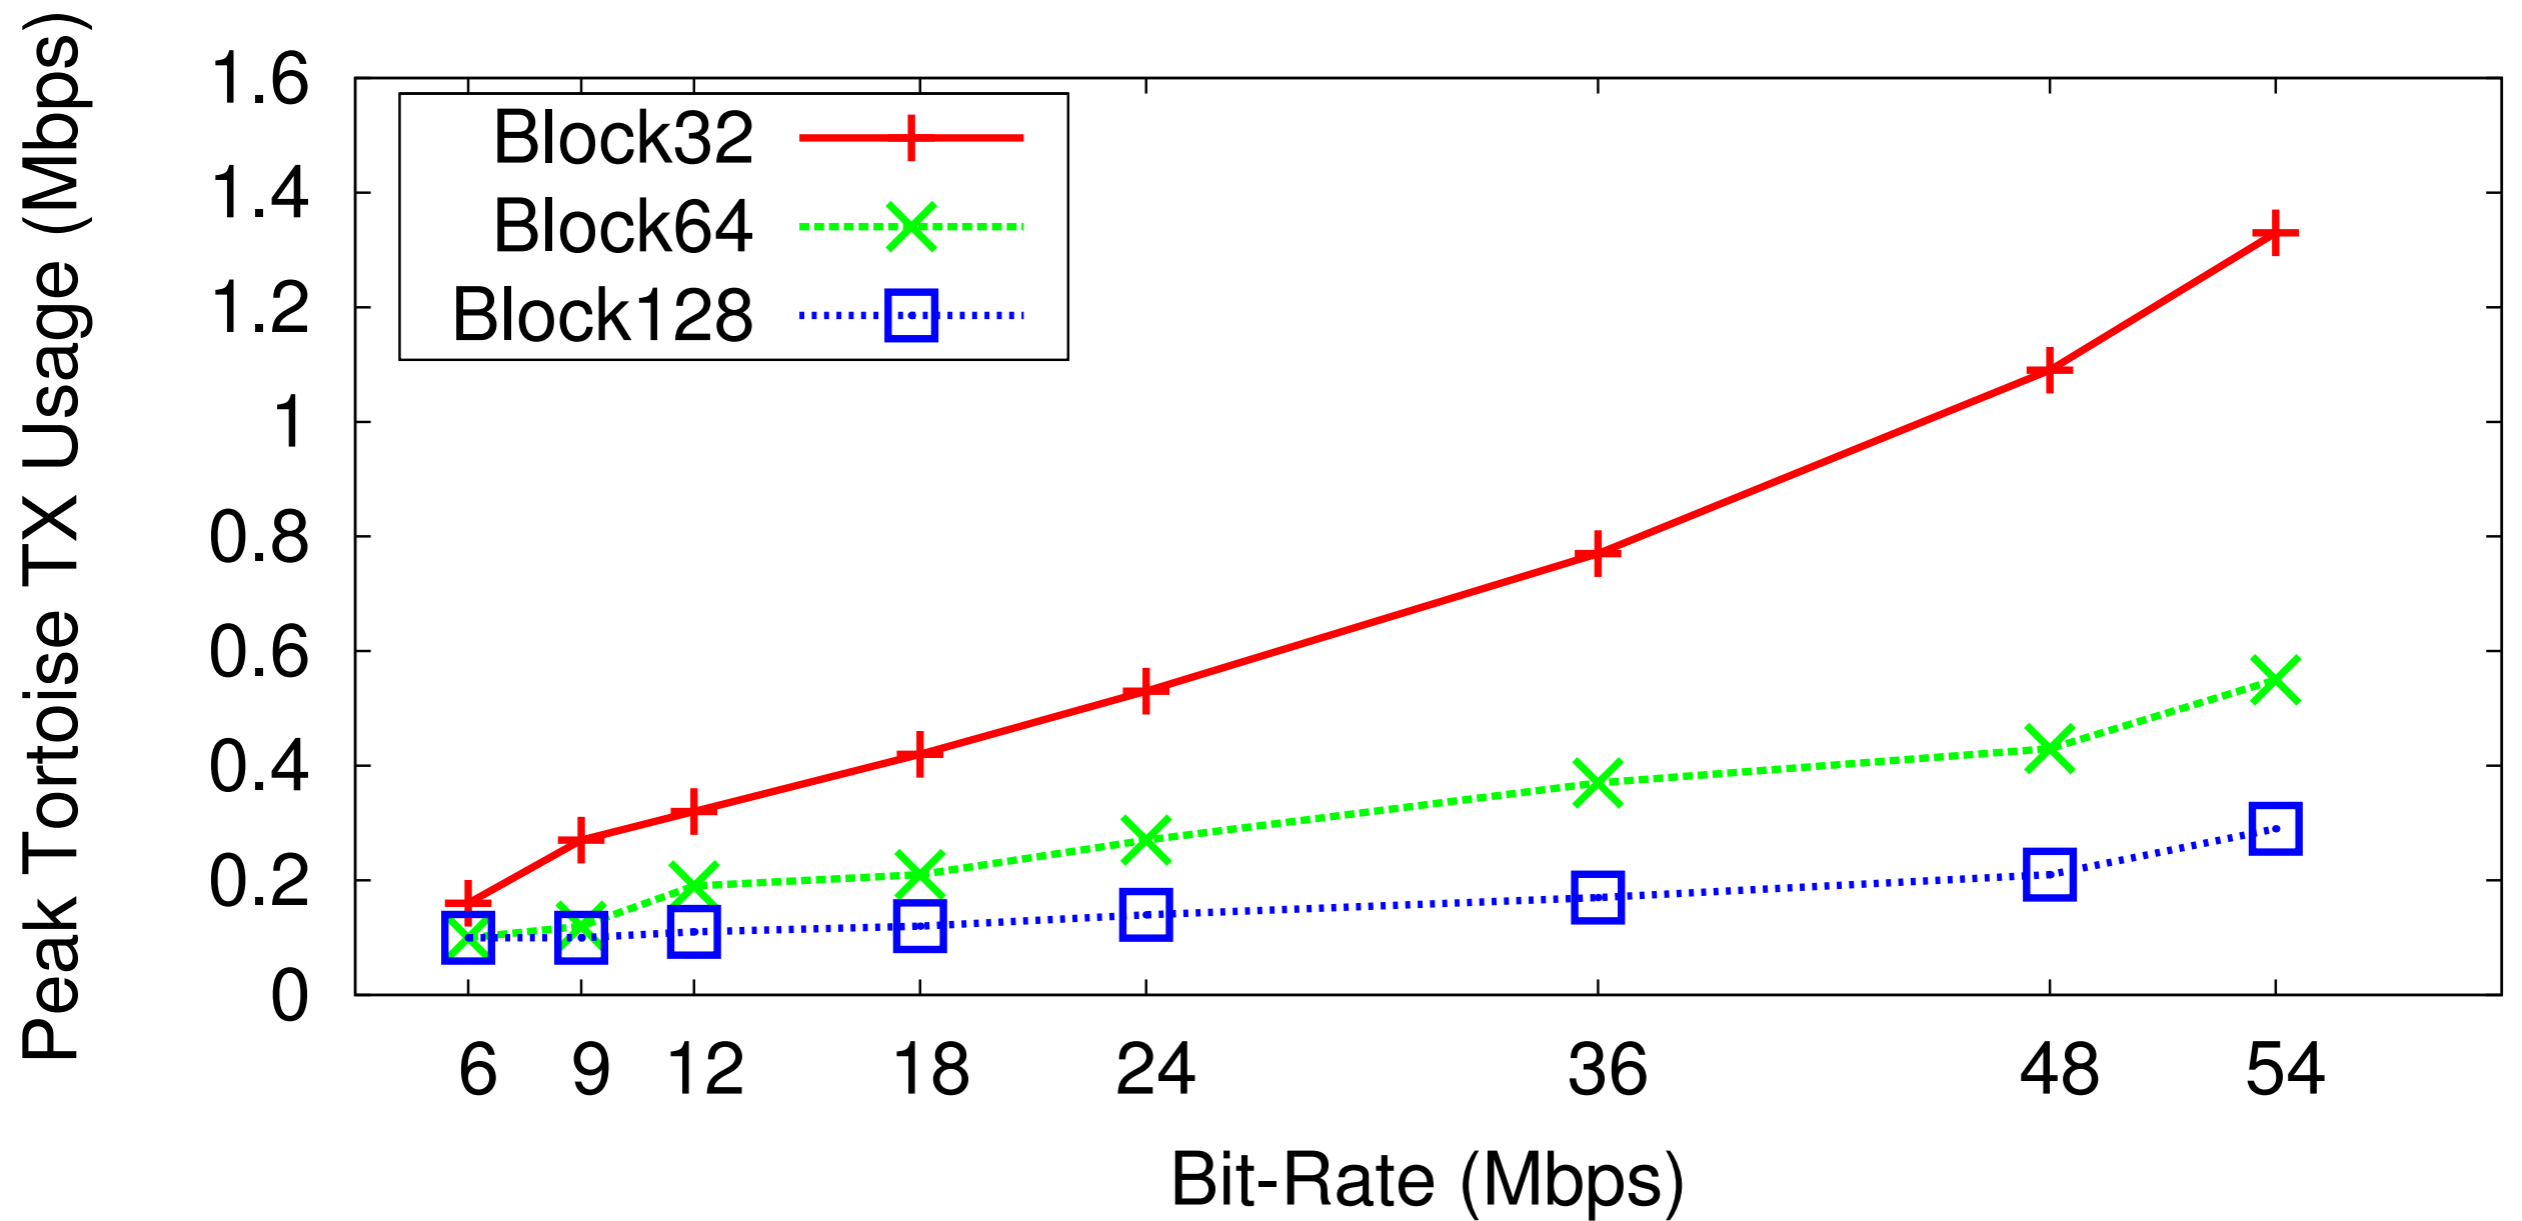
Peak Tortoise Utilization

Peak Tortoise Utilization


- Tortoise link utilization
 - Peak uplink and downlink utilizations are measured
 - Measured worst case by making every packet a partial one

Peak Tortoise Utilization

Peak Tortoise Utilization




Peak Tortoise Utilization




Peak Tortoise Utilization

Peak Tortoise Utilization



Peak Tortoise Utilization



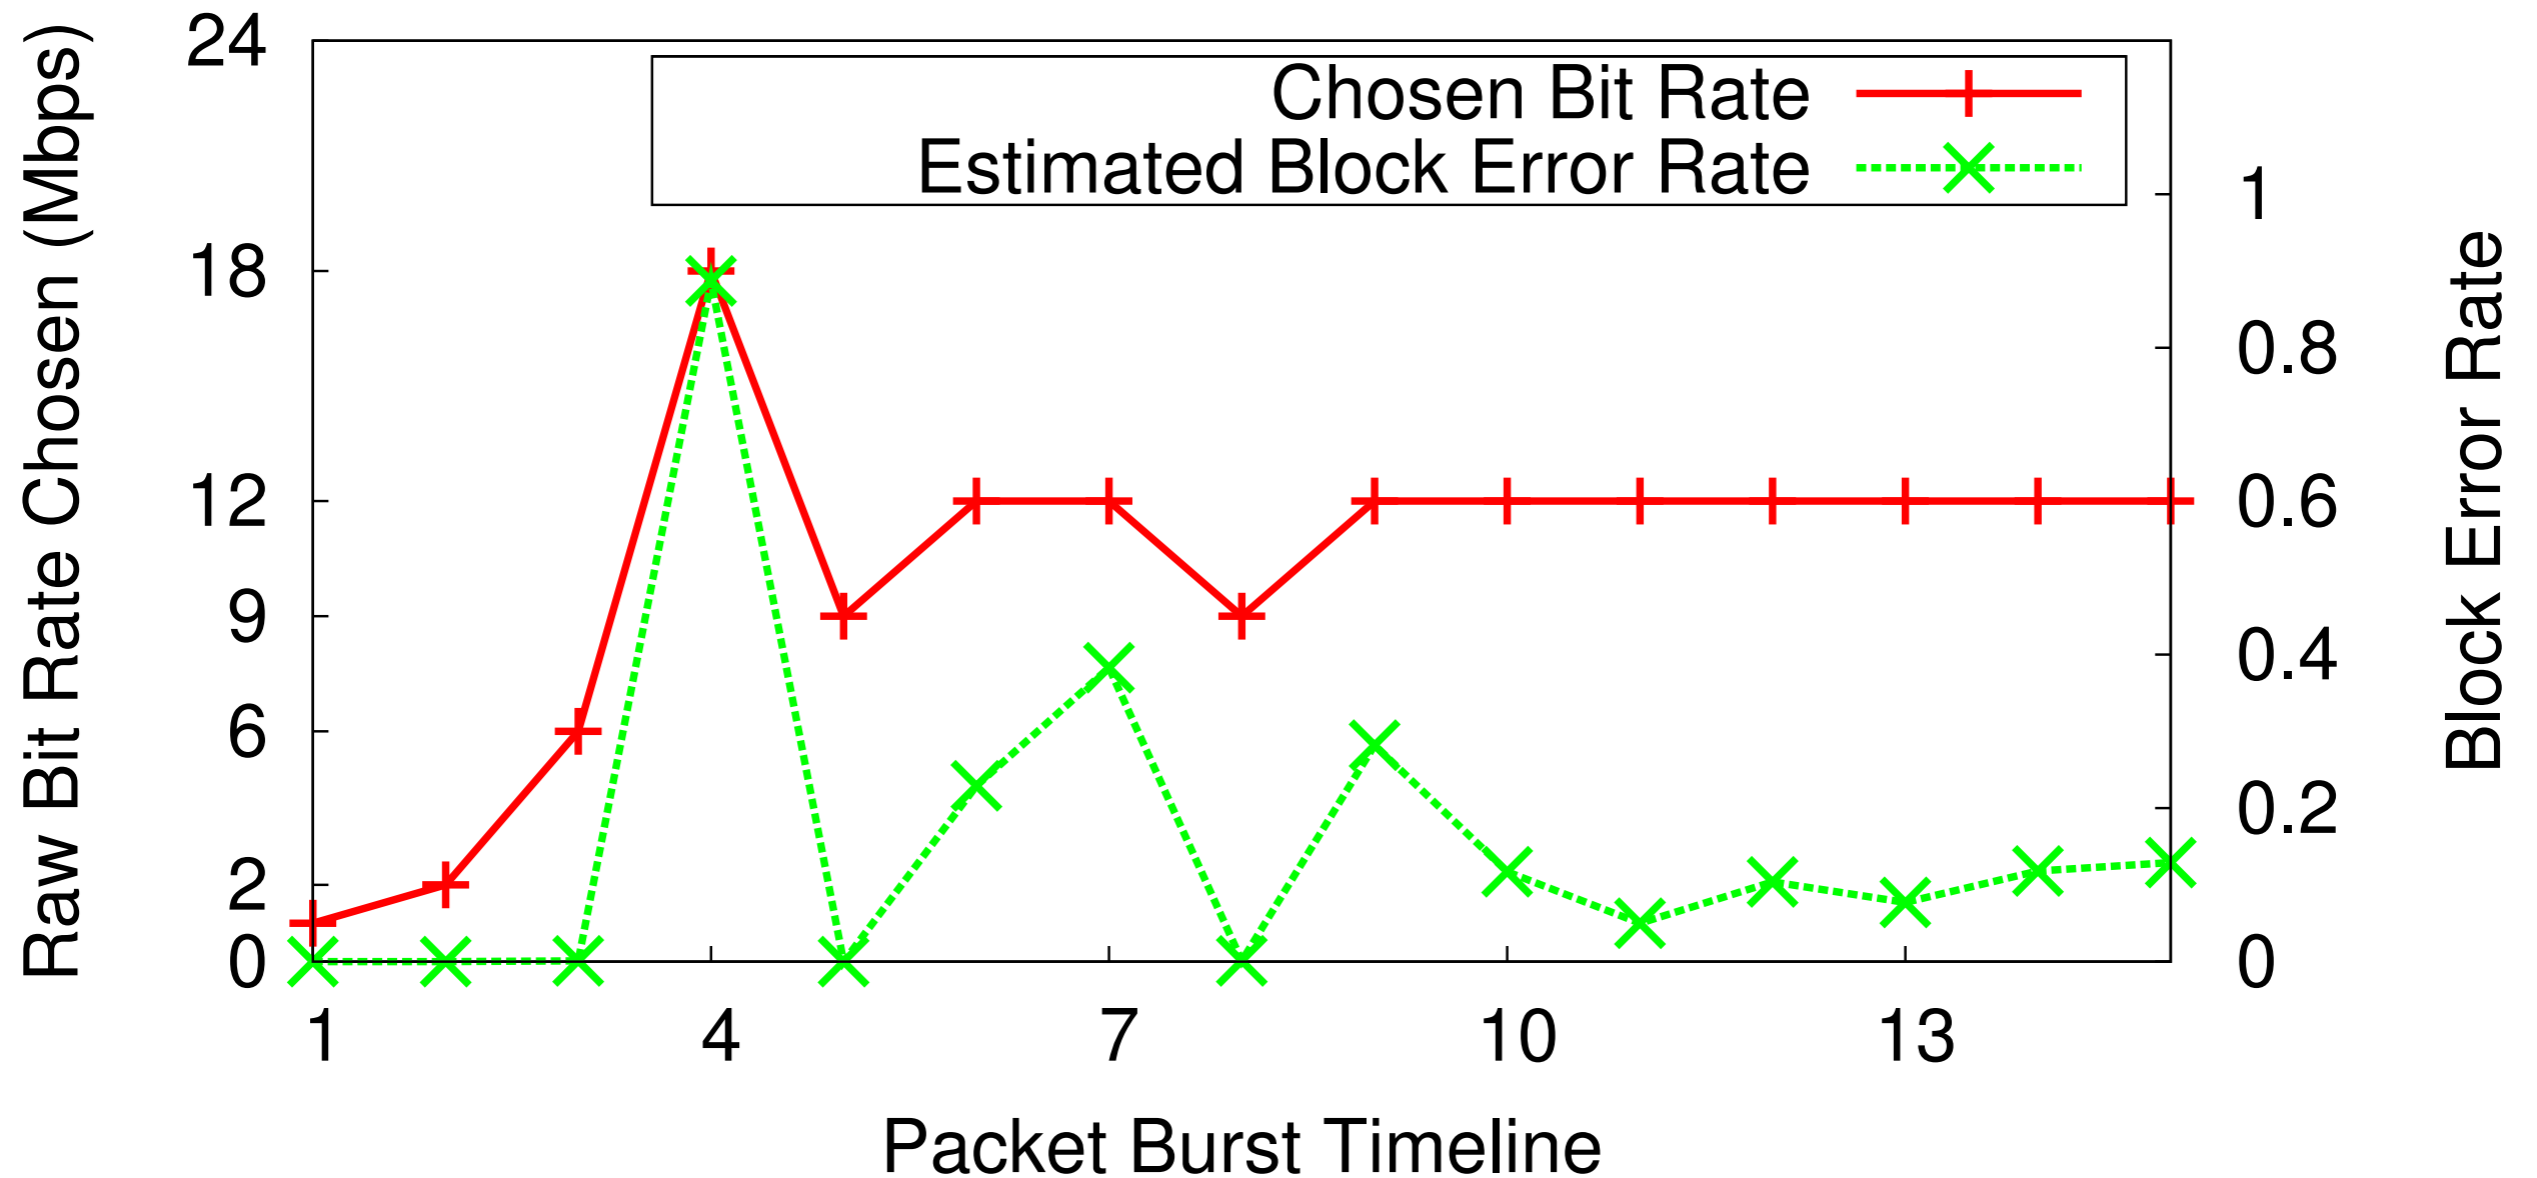
Bitrate Adaptation's Performance

Bitrate Adaptation's Performance


- Autorate performance
 - Measure obtained goodput against hand chosen one for “mostly” static links
 - Measure the number of round trips the mechanism takes to reach the ideal rate

Bitrate Adaptation's Performance

Bitrate Adaptation's Performance

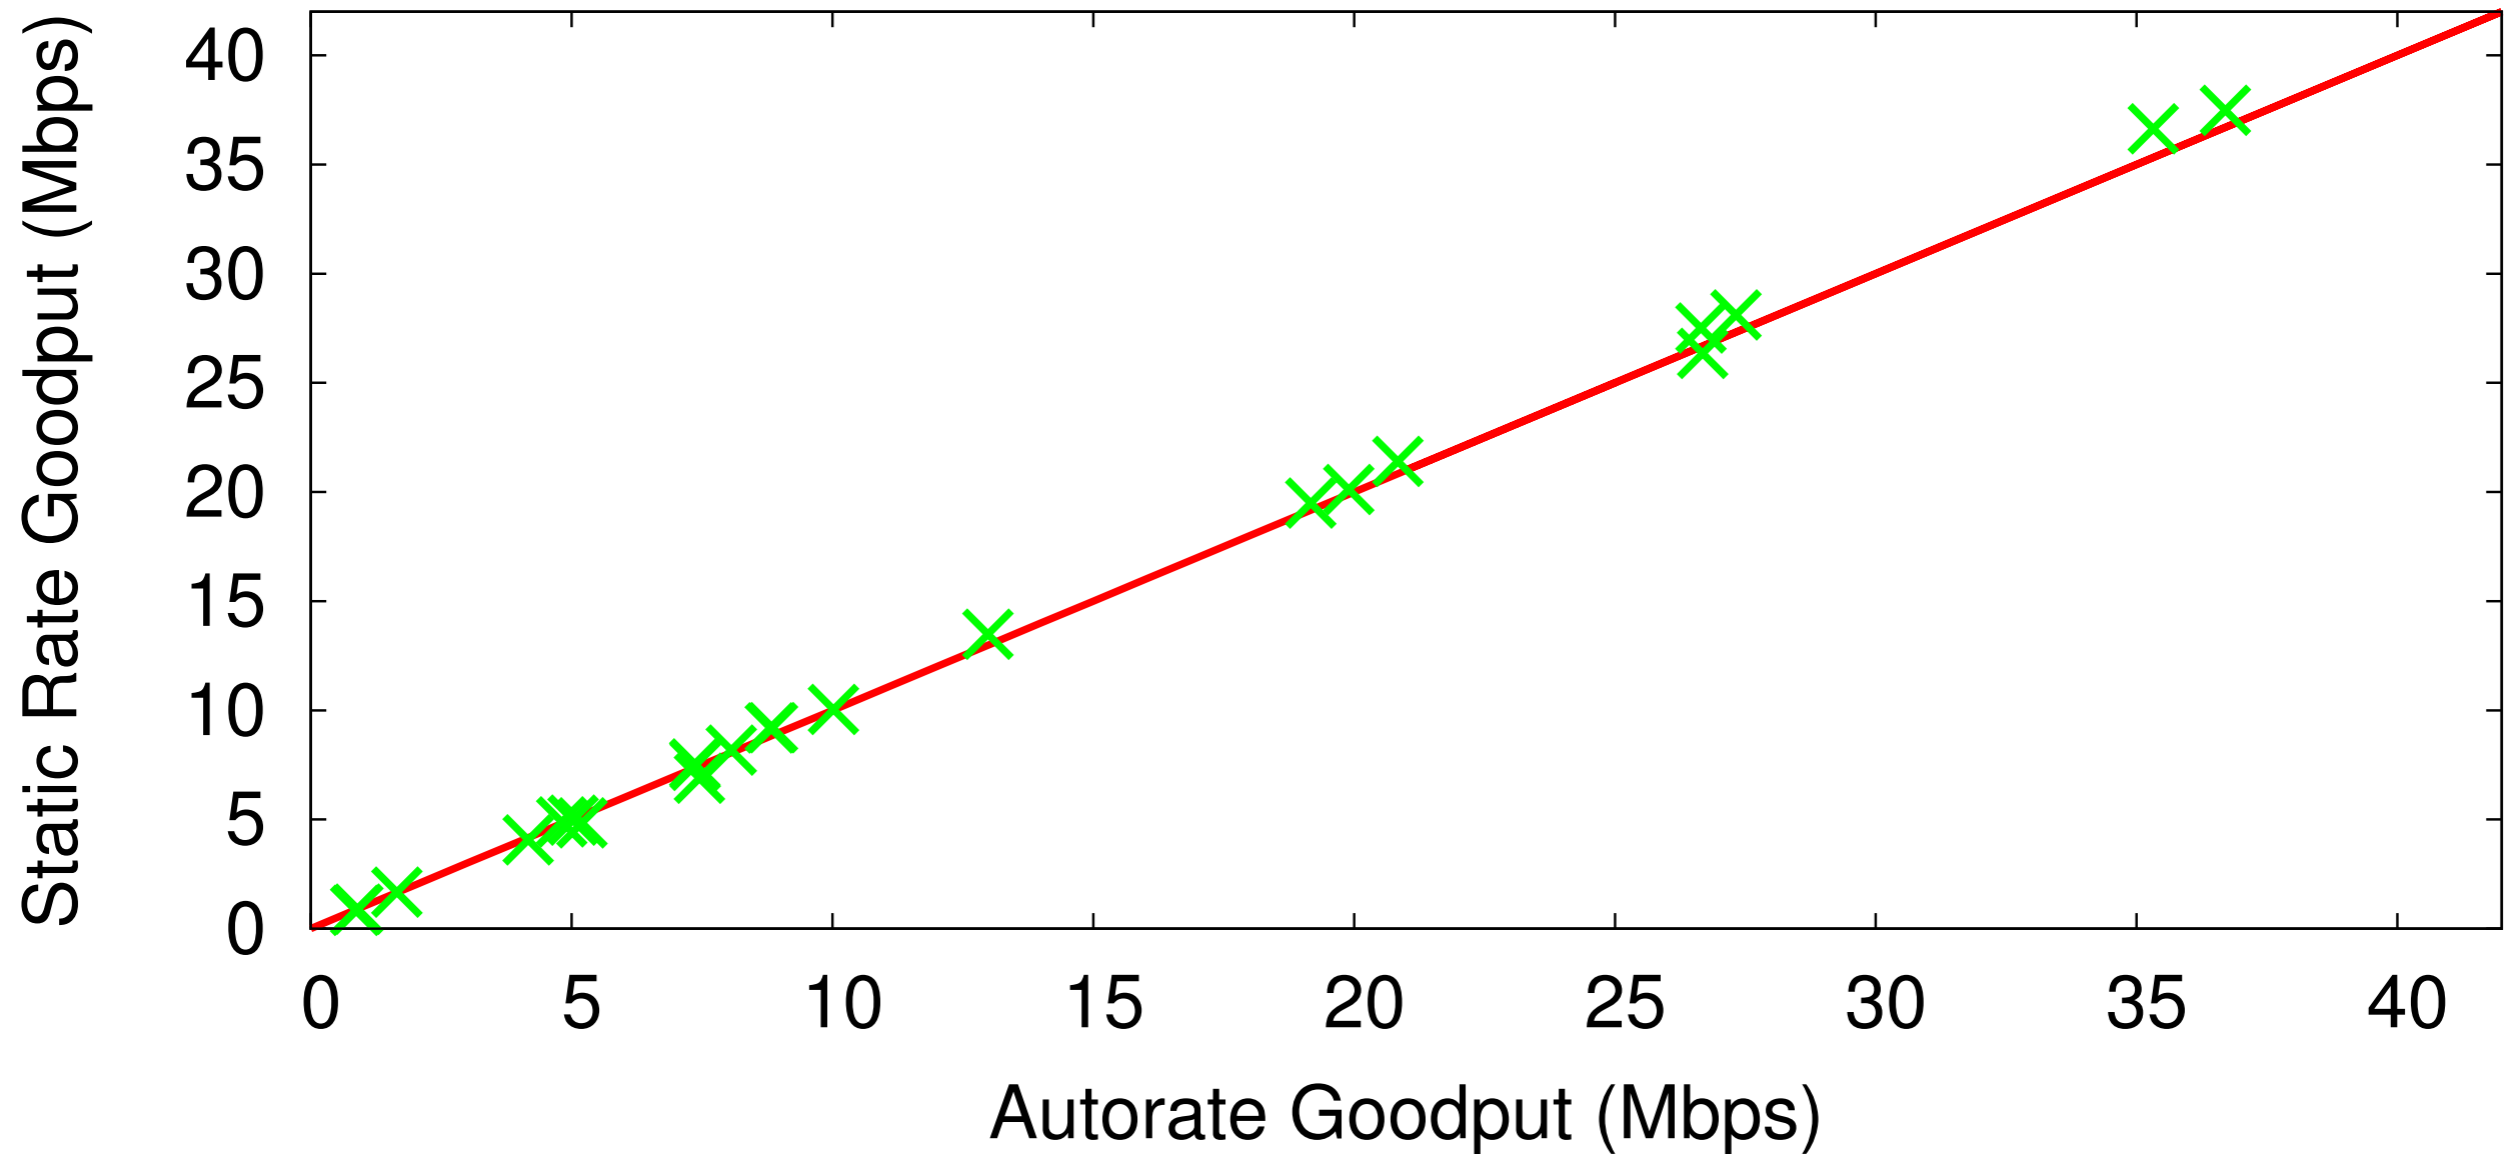


Bitrate Adaptation's Performance

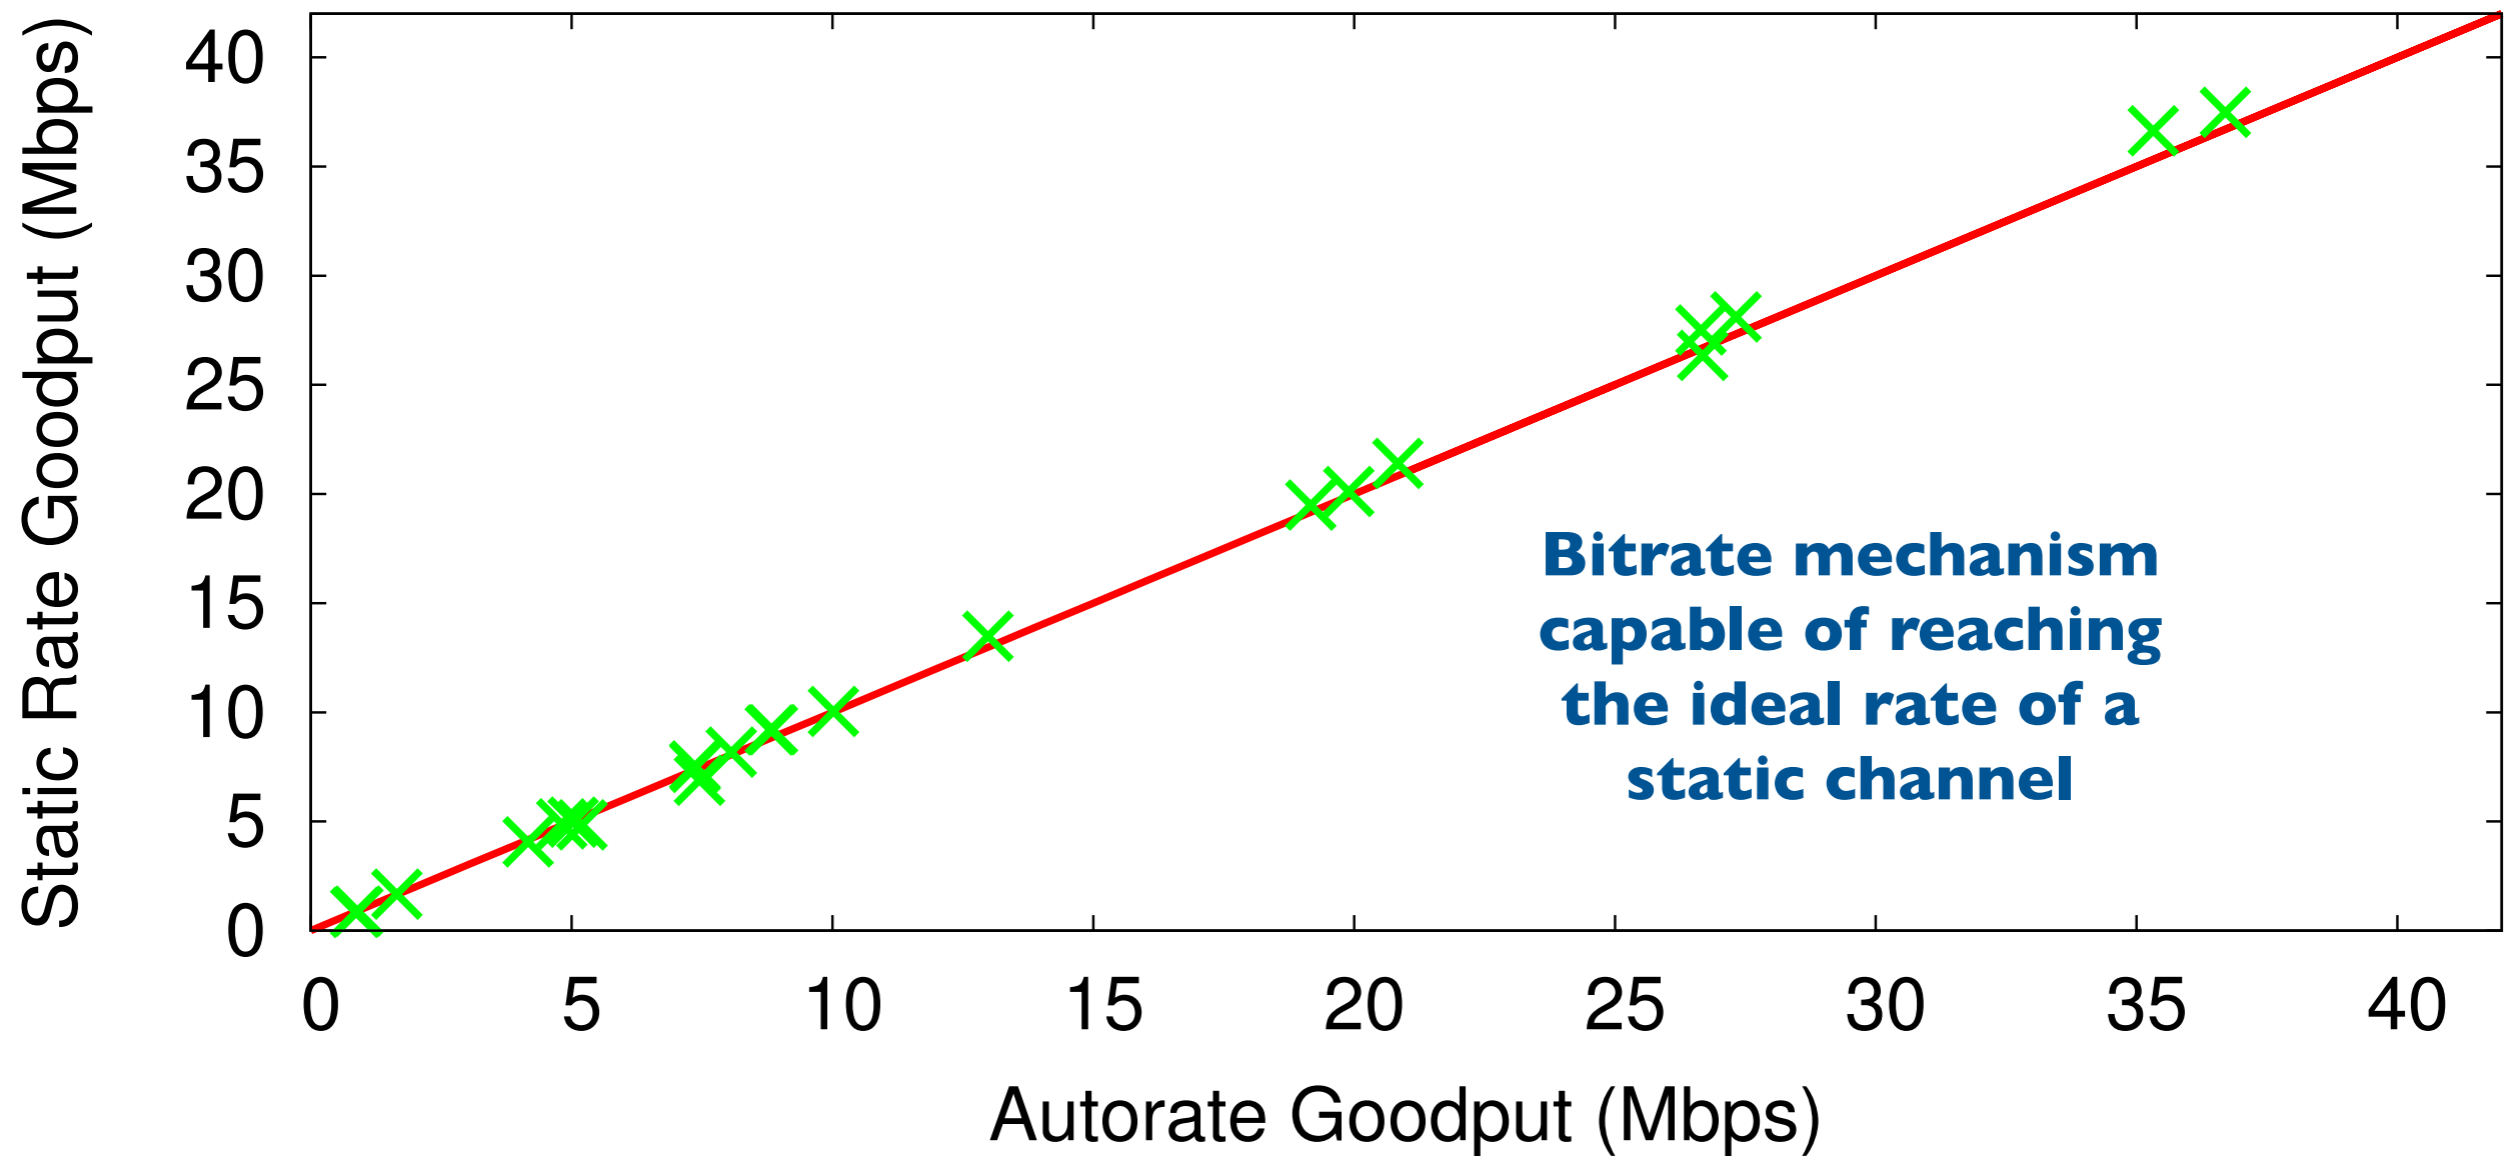


Bitrate Adaptation's Performance

Bitrate Adaptation's Performance



Bitrate Adaptation's Performance




CRC Precomputation Benefits

CRC Precomputation Benefits

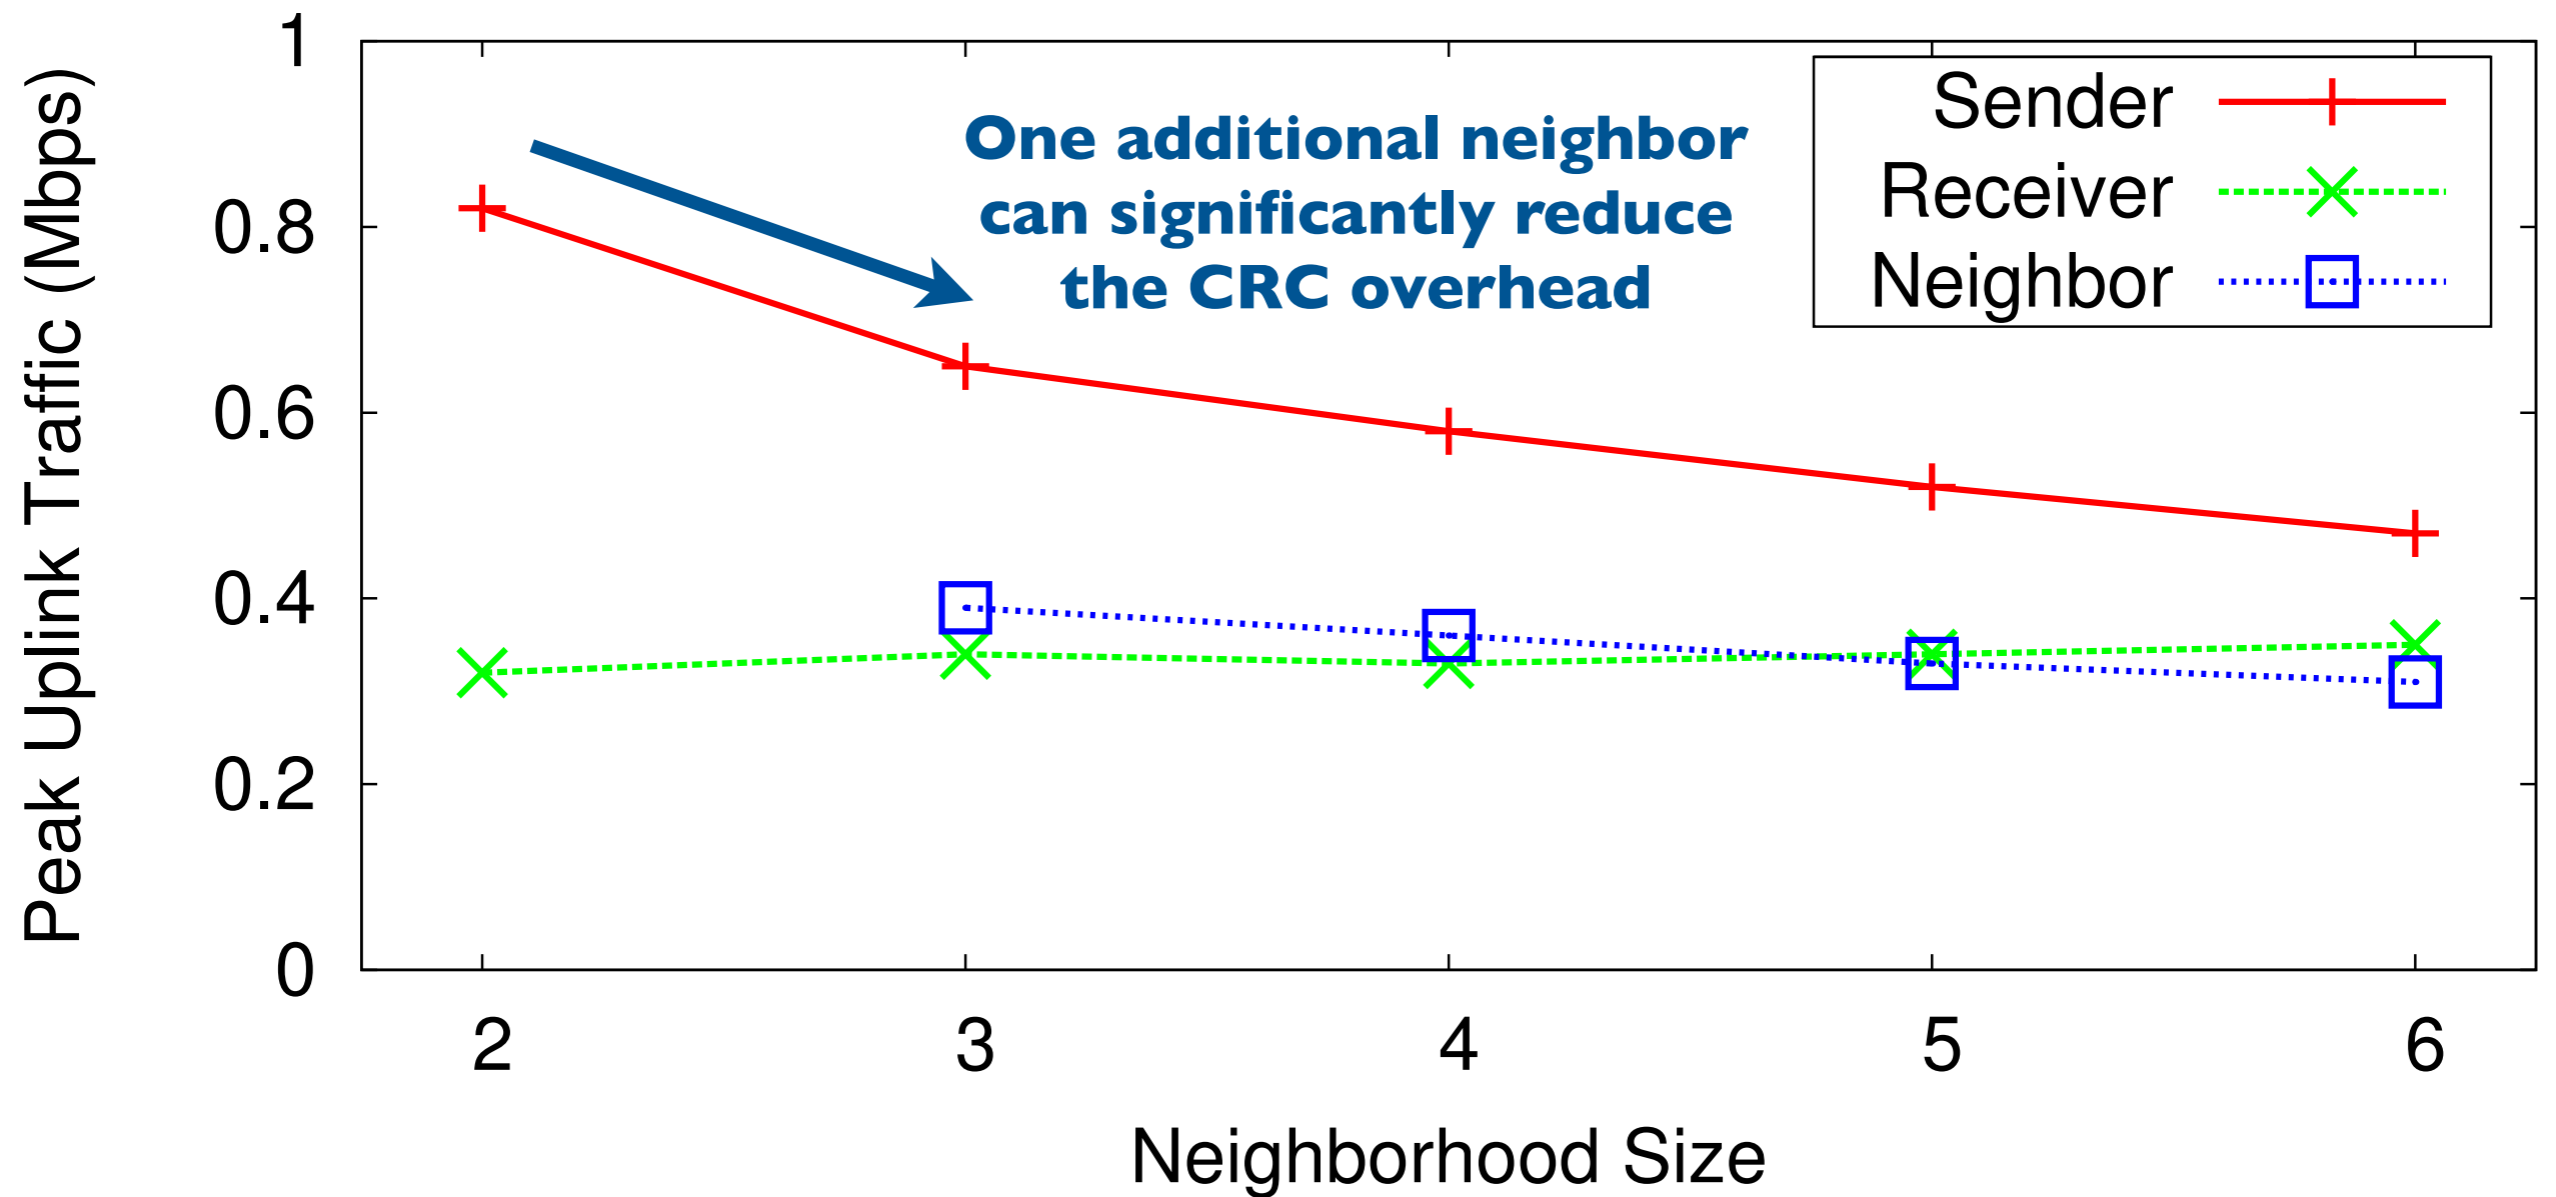
- CRC precomputation and preseeding benefits
 - Test in a neighborhood of 6 houses for static multimedia content shared between the homes
 - Measure peak-uplink tortoise link utilization

CRC Precomputation Benefits

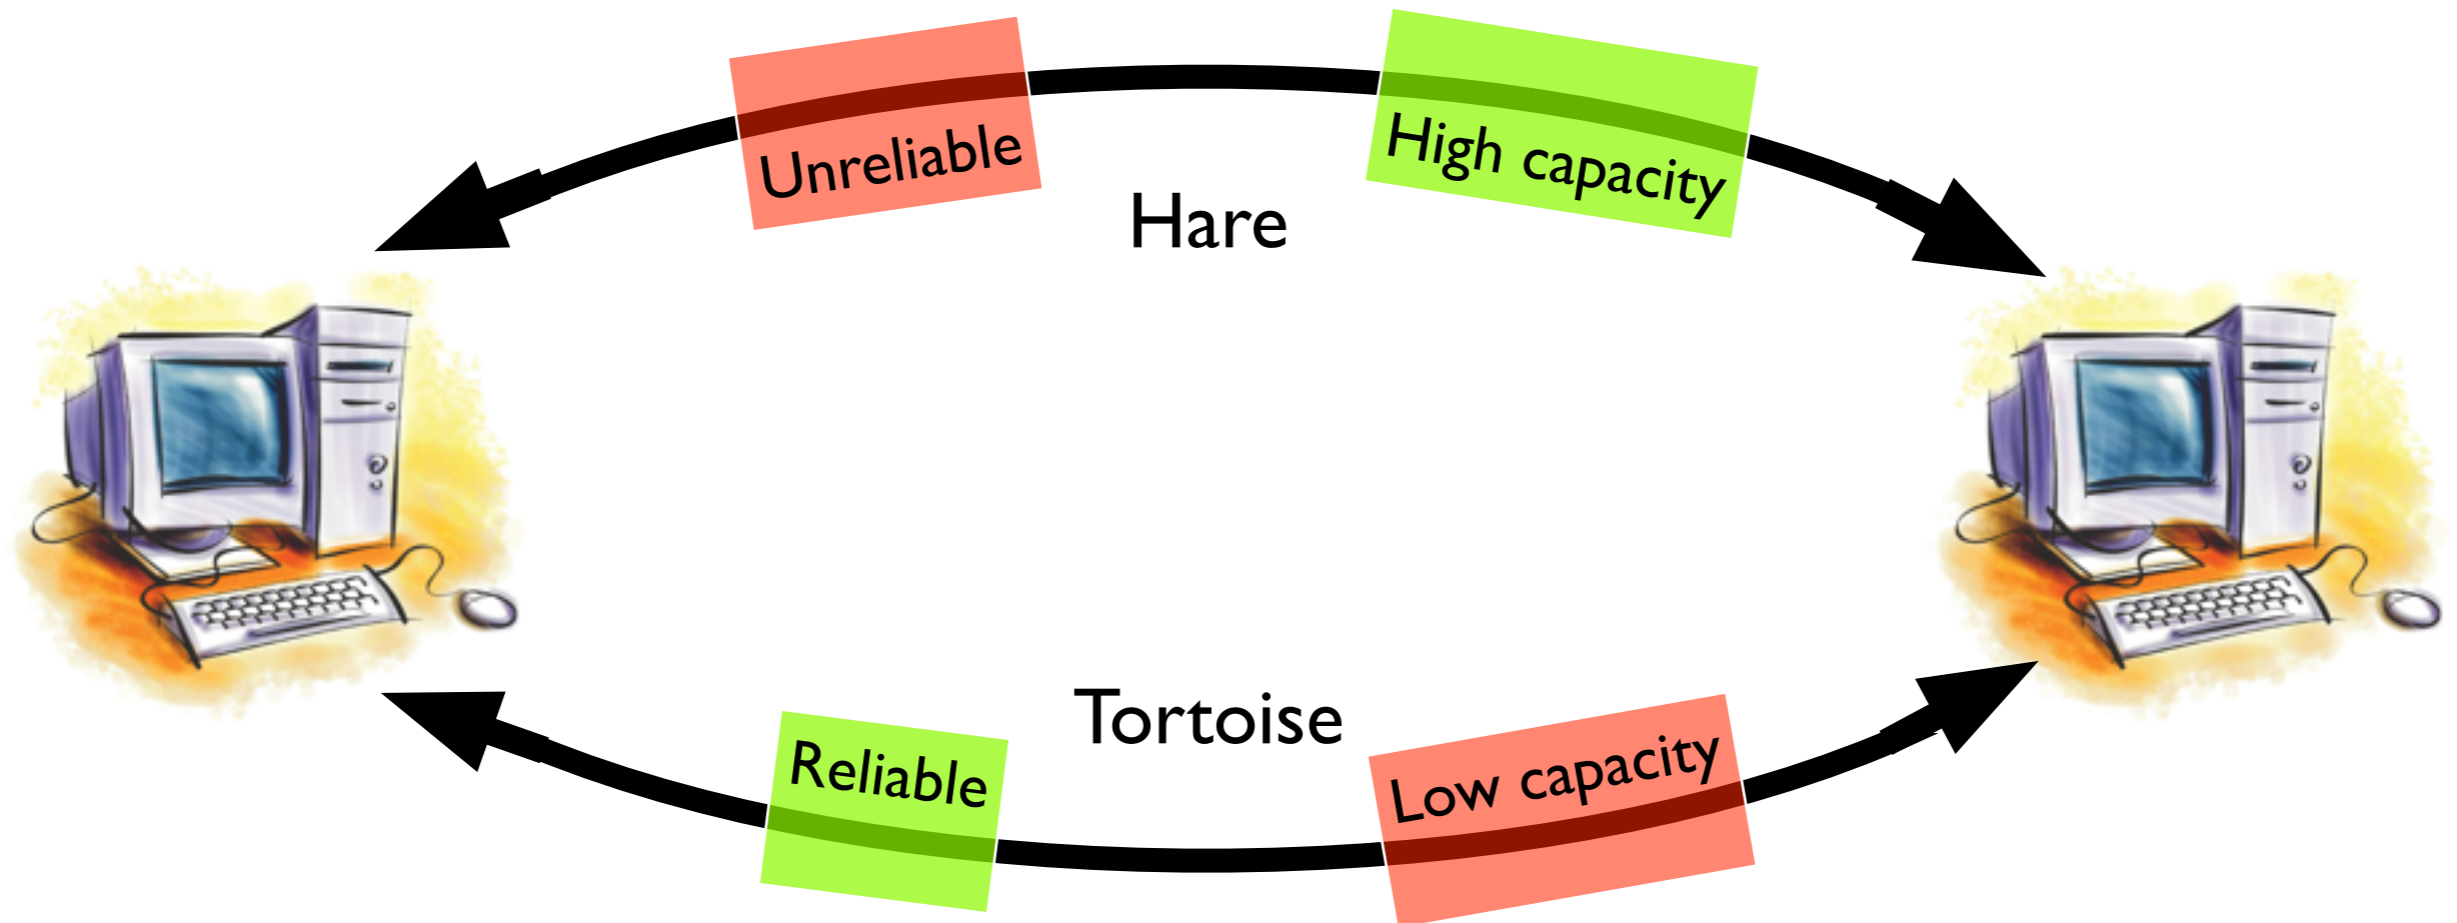
CRC Precomputation Benefits




CRC Precomputation Benefits



Conclusions



Conclusions



- Hare and the tortoise link bonding mechanism via PPR
- Low footprint metadata design for the tortoise
- Performs upto 3x better than TCP and upto 2.5x better than Hop

Thanks

Anirudh Badam

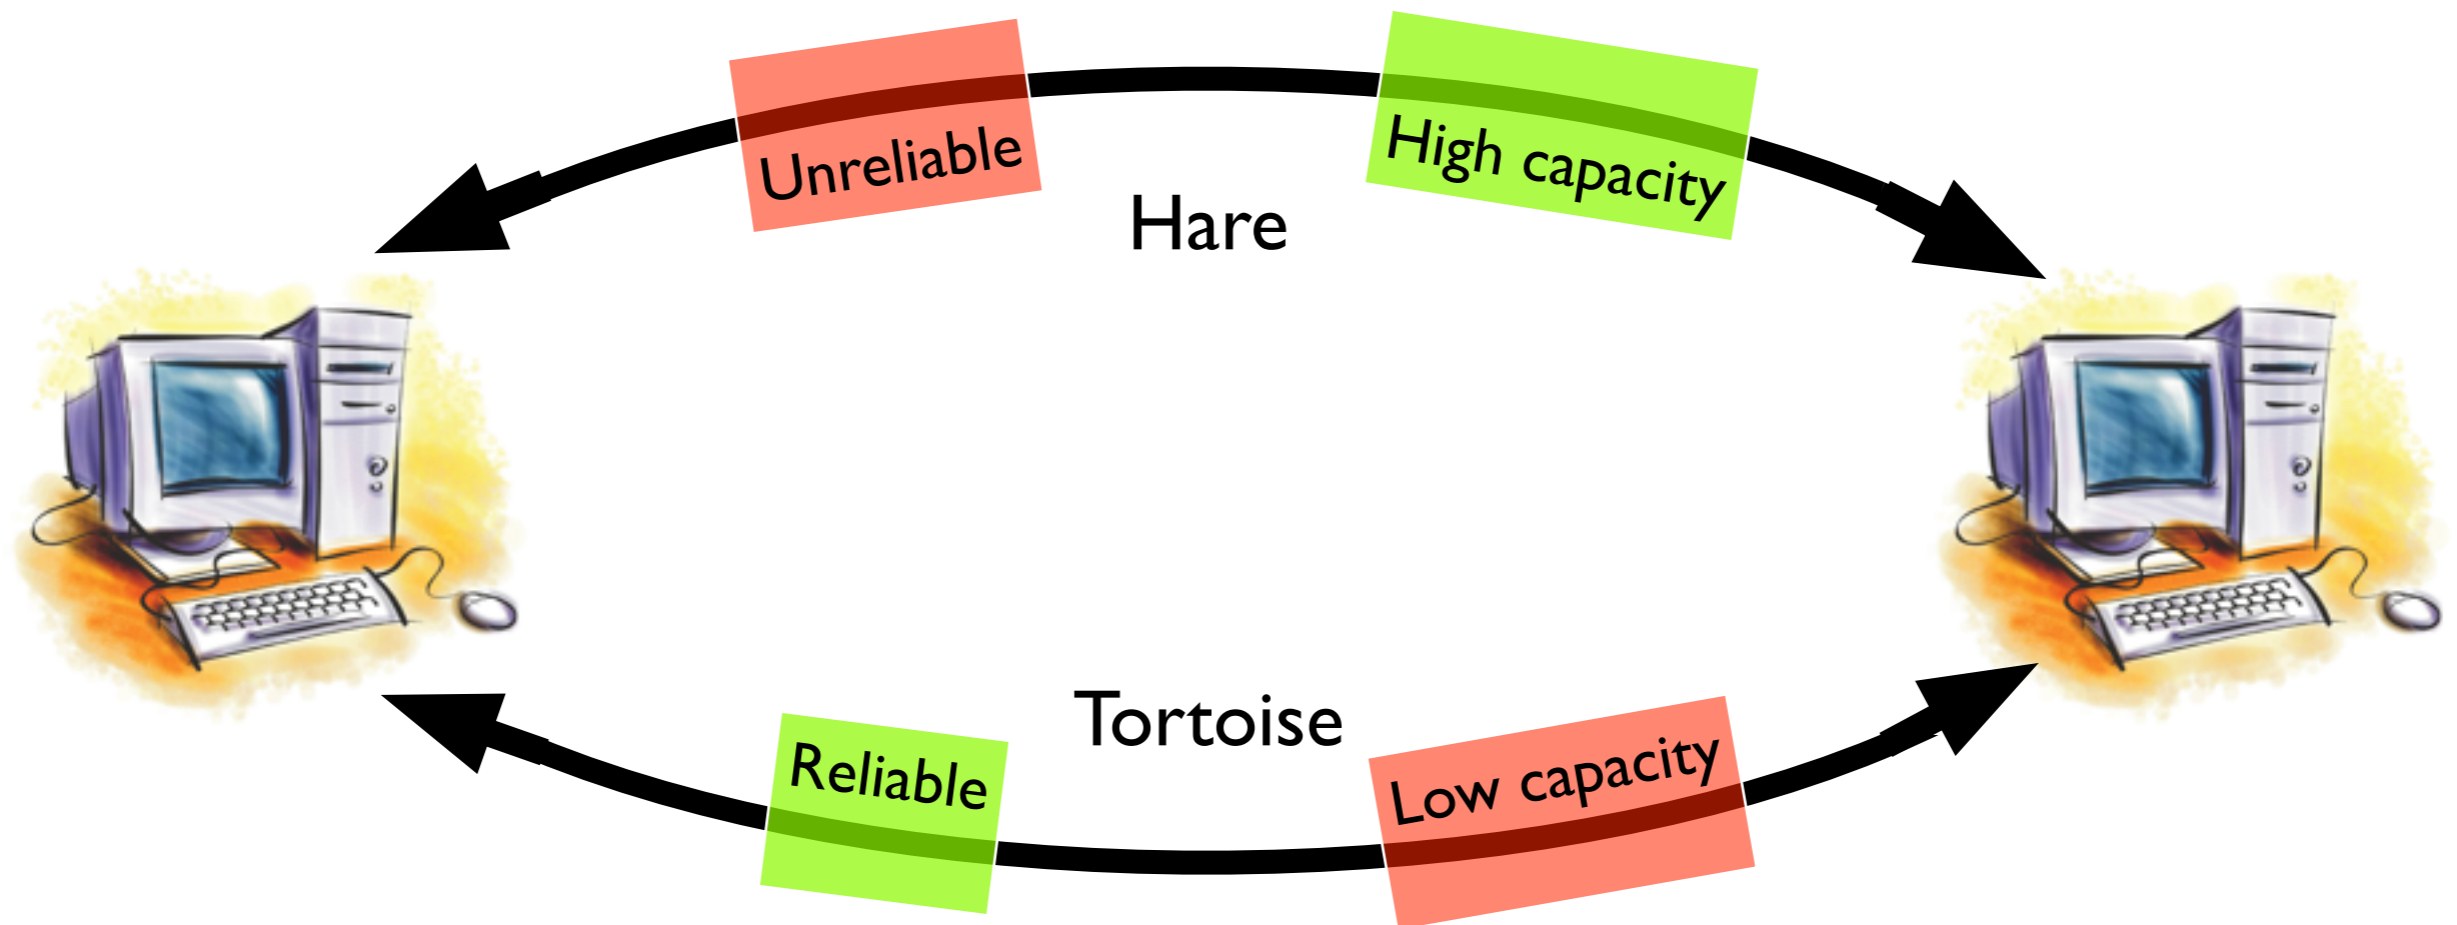
abadam@cs.princeton.edu



Carnegie Mellon



Conclusions



- Hare and the tortoise link bonding mechanism via PPR
- Low footprint metadata design for the tortoise
- Performs upto 3x better than TCP and upto 2.5x better than Hop