Can Realistic BitTorrent Experiments Be Performed on Clusters?

Ashwin Rao, Arnaud Legout, and Walid Dabbous

INRIA, Project Planète
(ashwin.rao,arnaud.legout,walid.dabbous)@inria.fr

Overview of BitTorrent

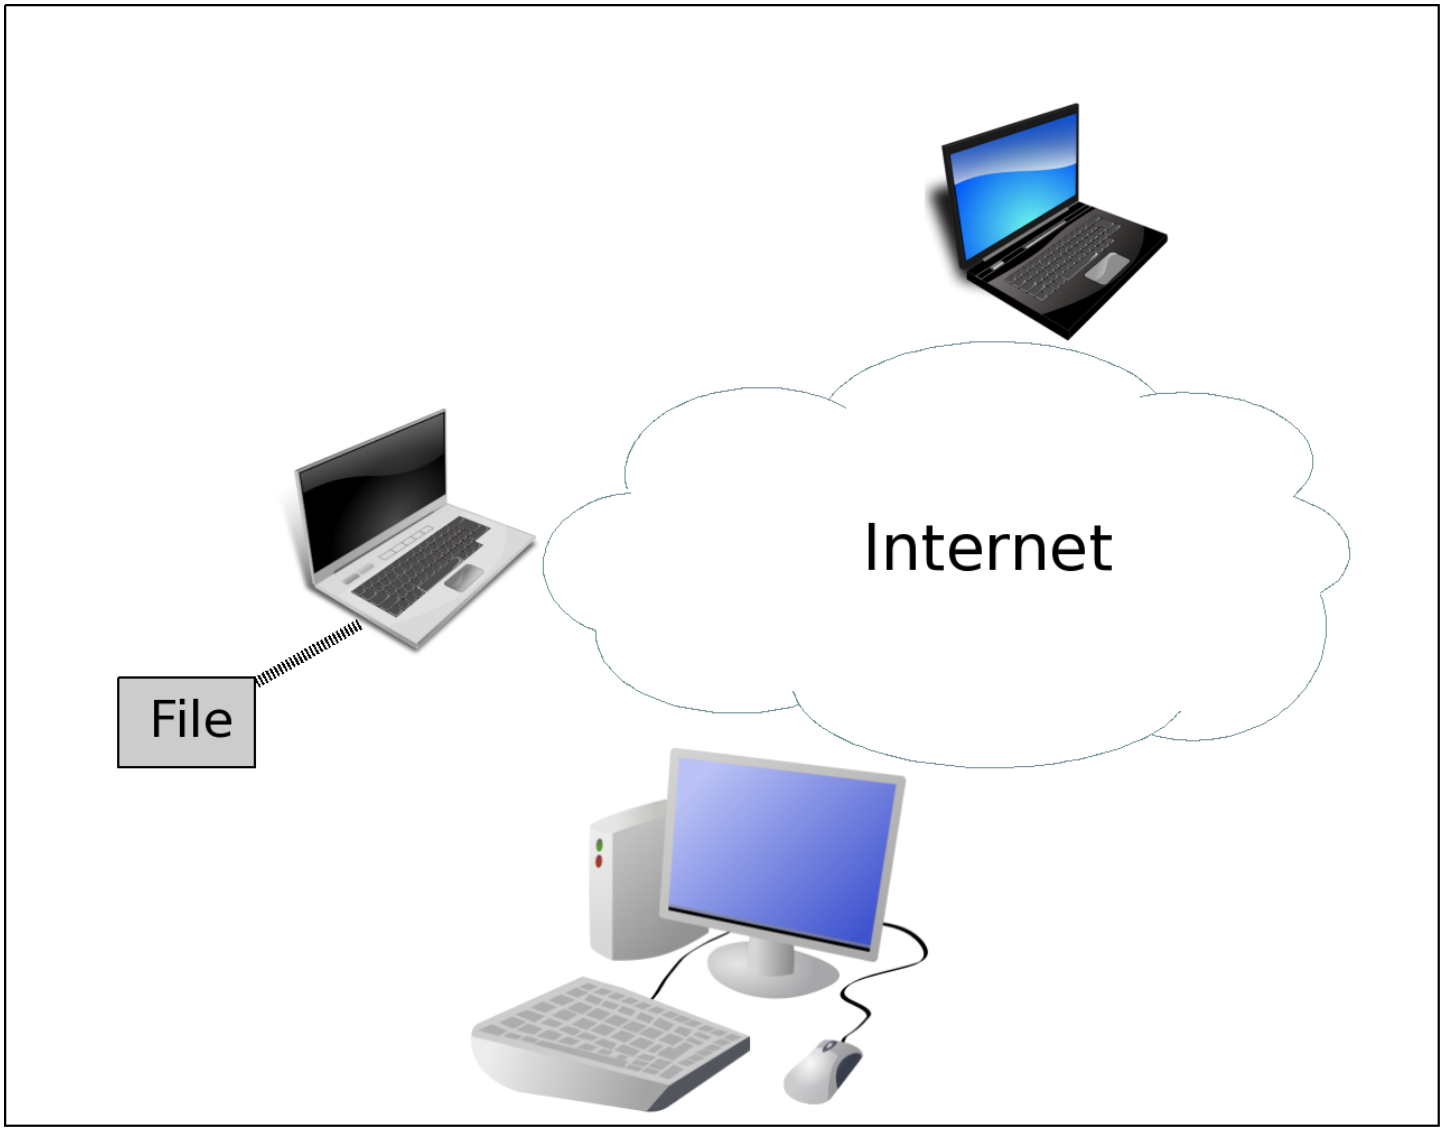

Overview of BitTorrent

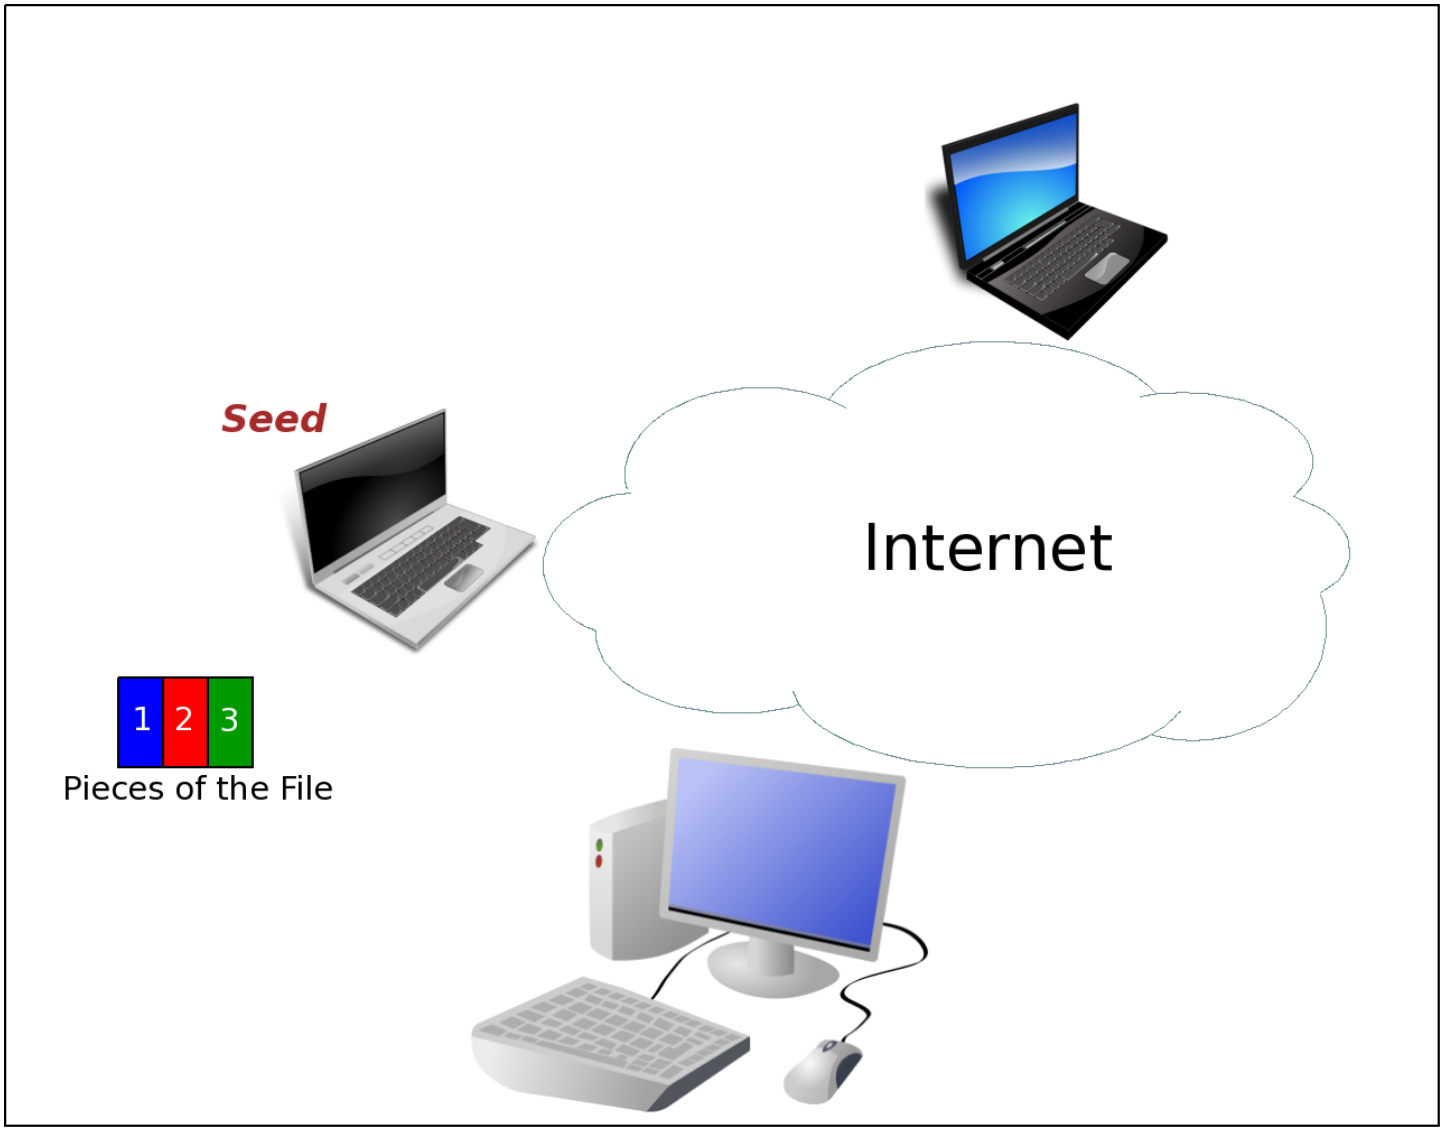

Overview of BitTorrent

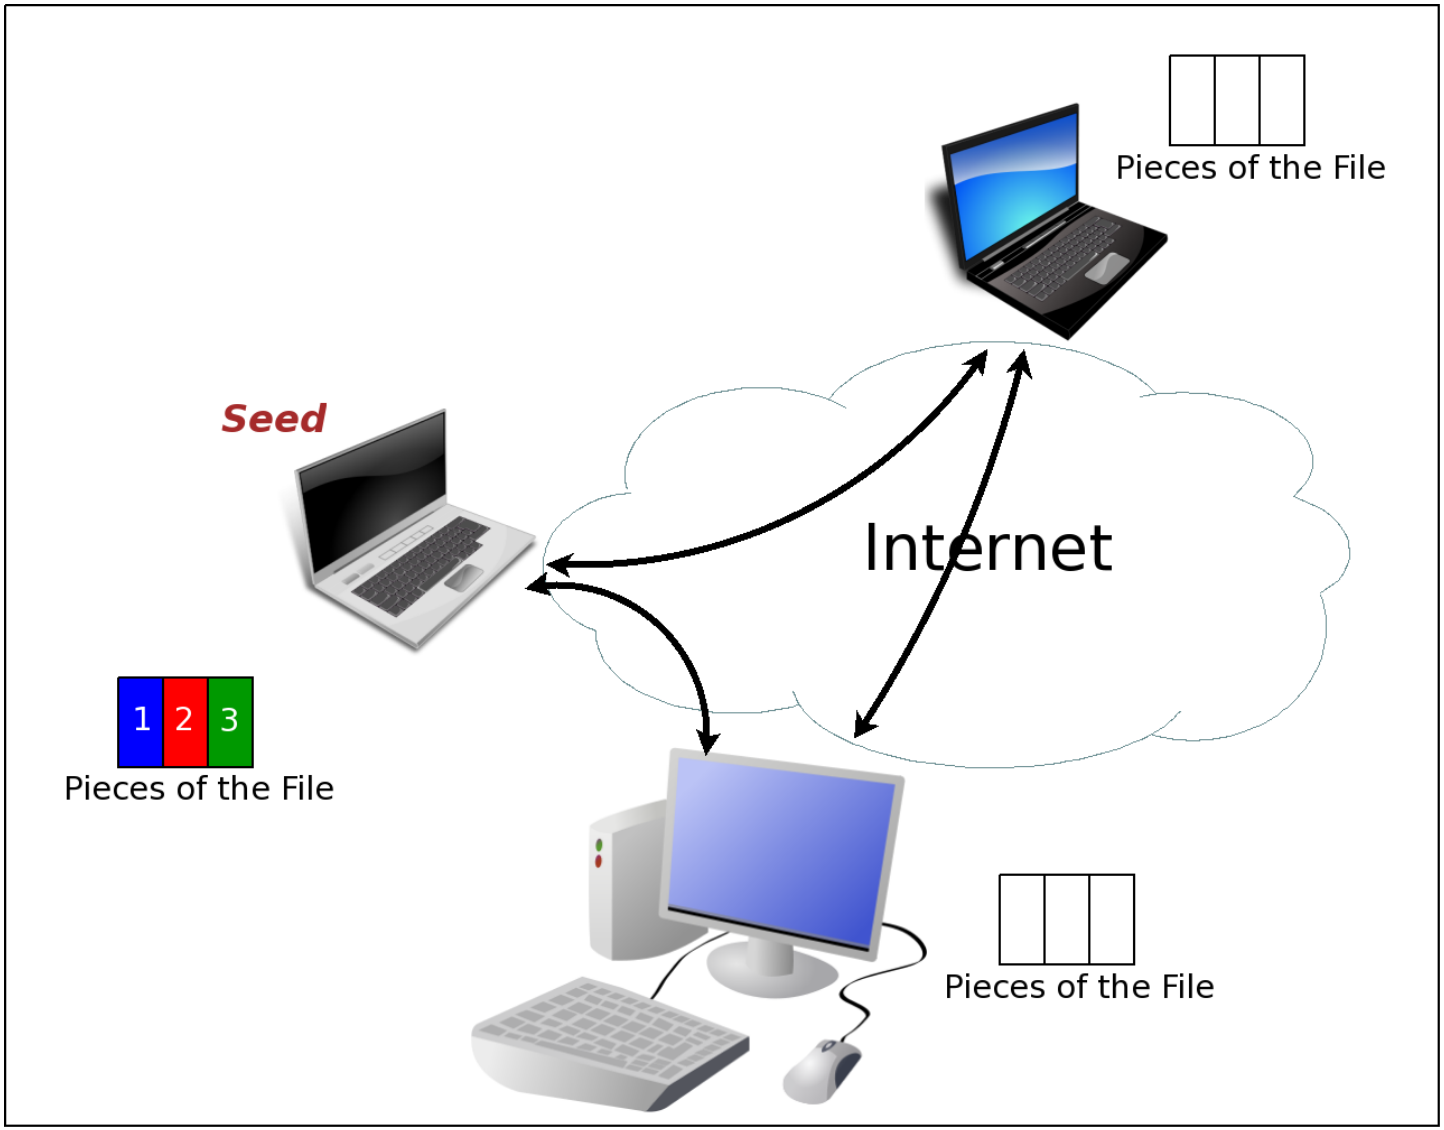

Overview of BitTorrent

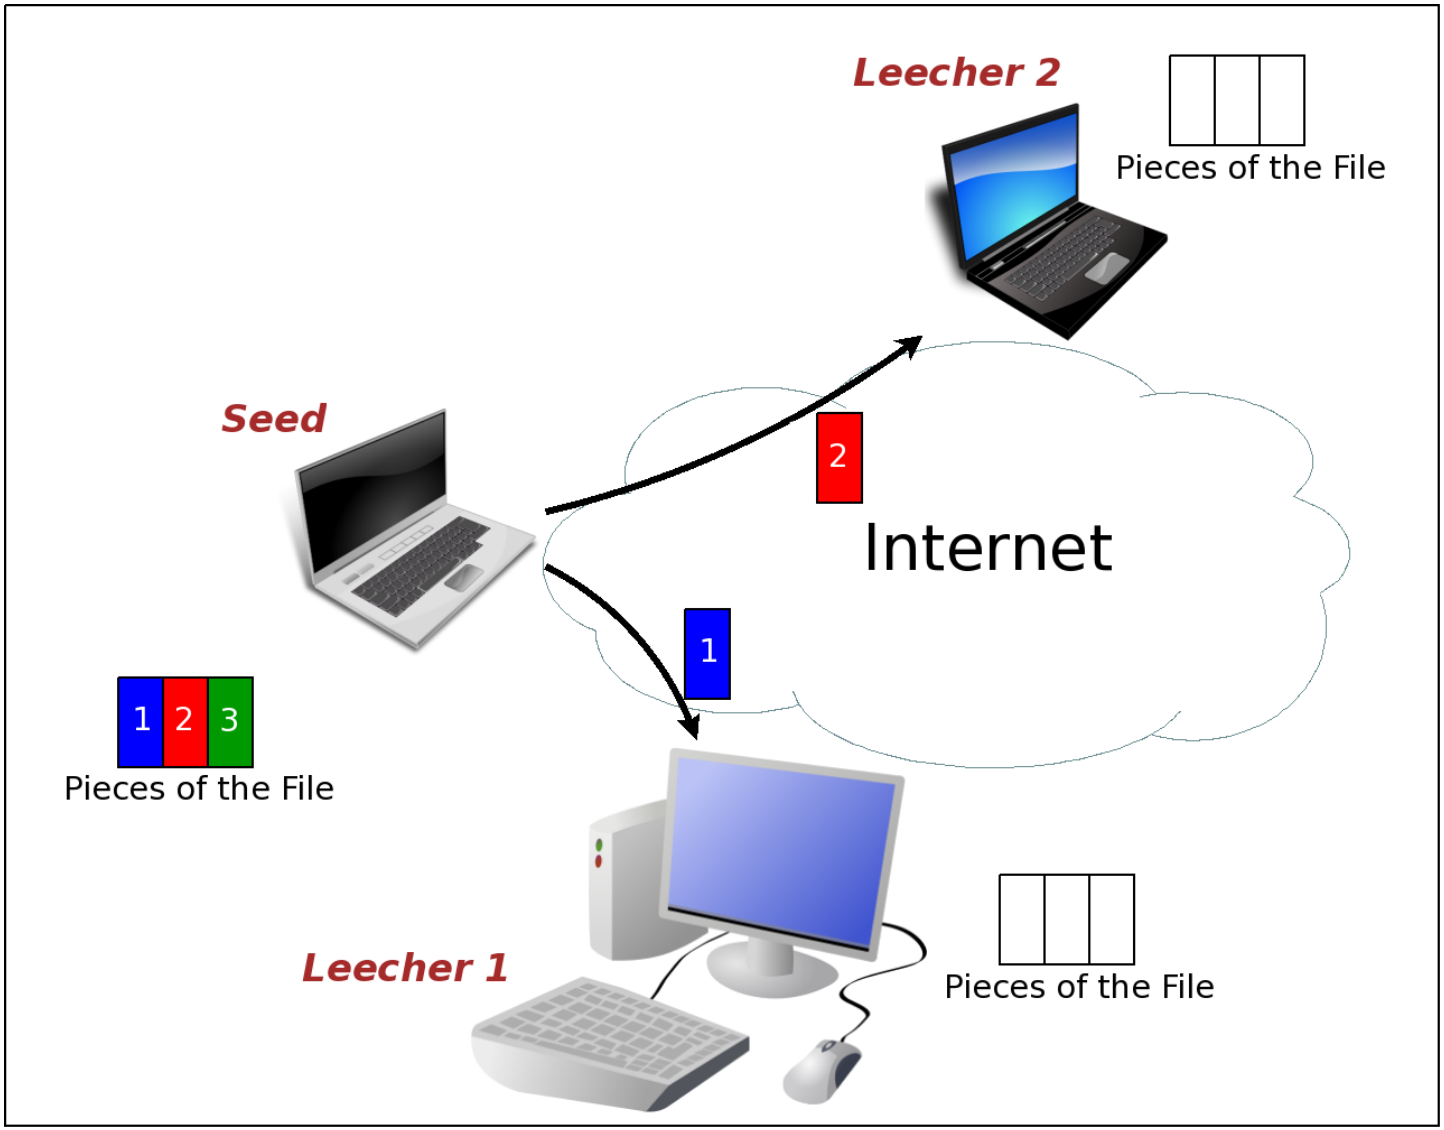

Overview of BitTorrent

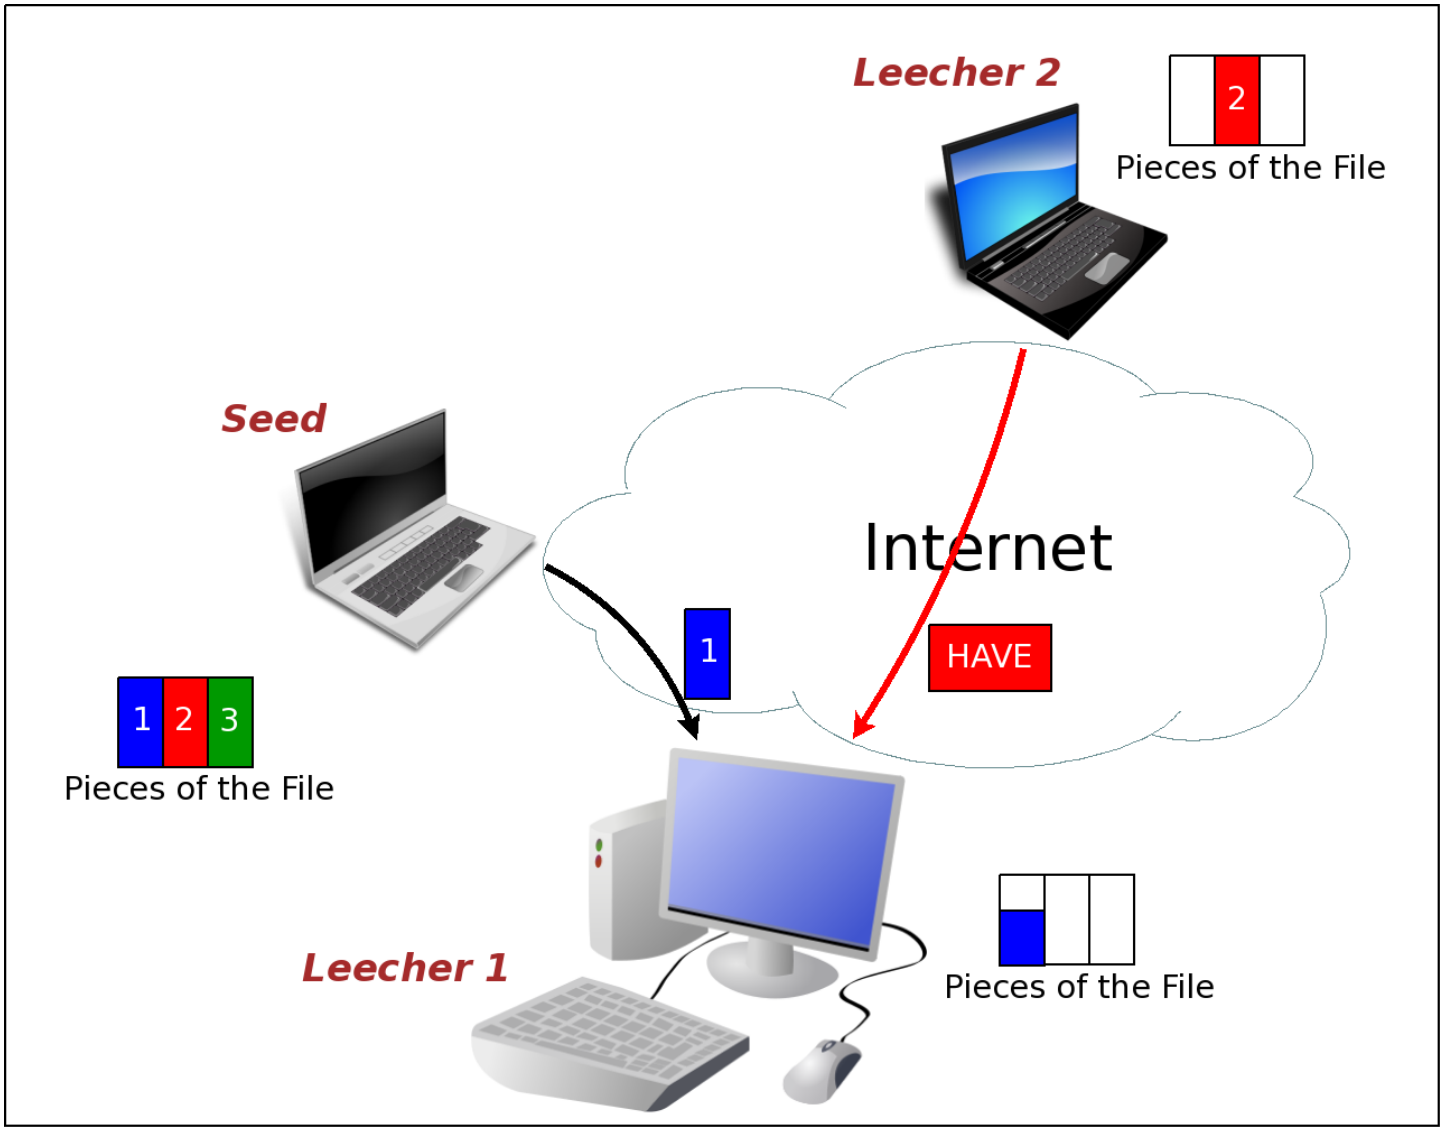

Overview of BitTorrent

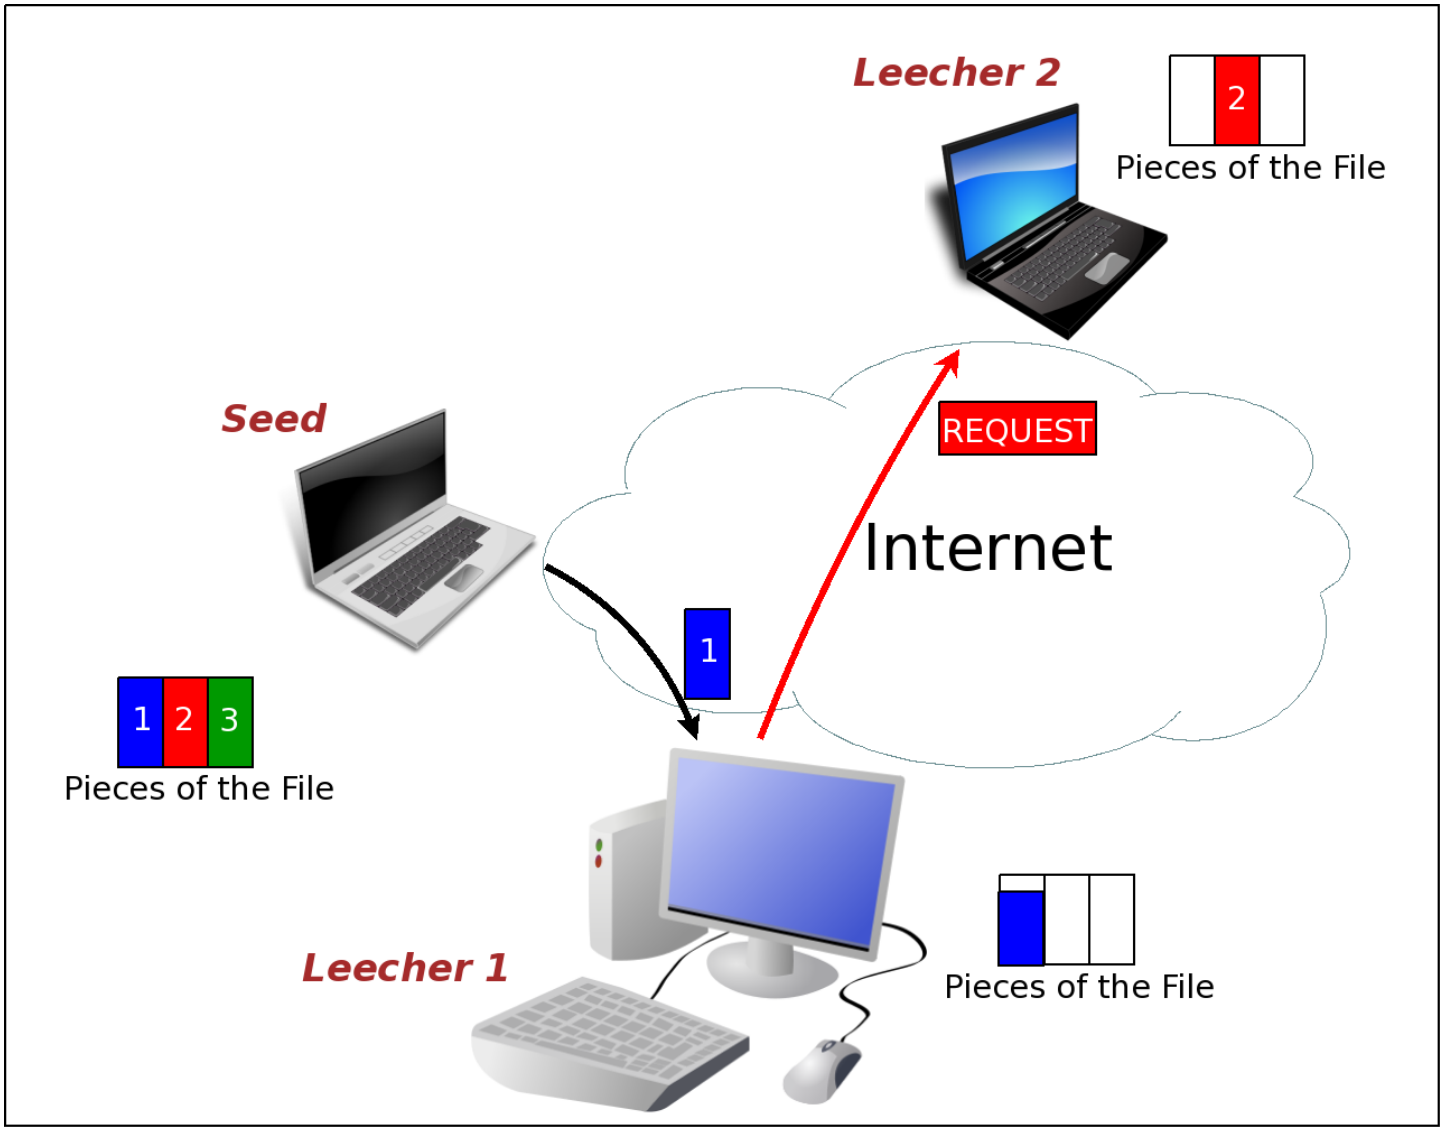

Overview of BitTorrent

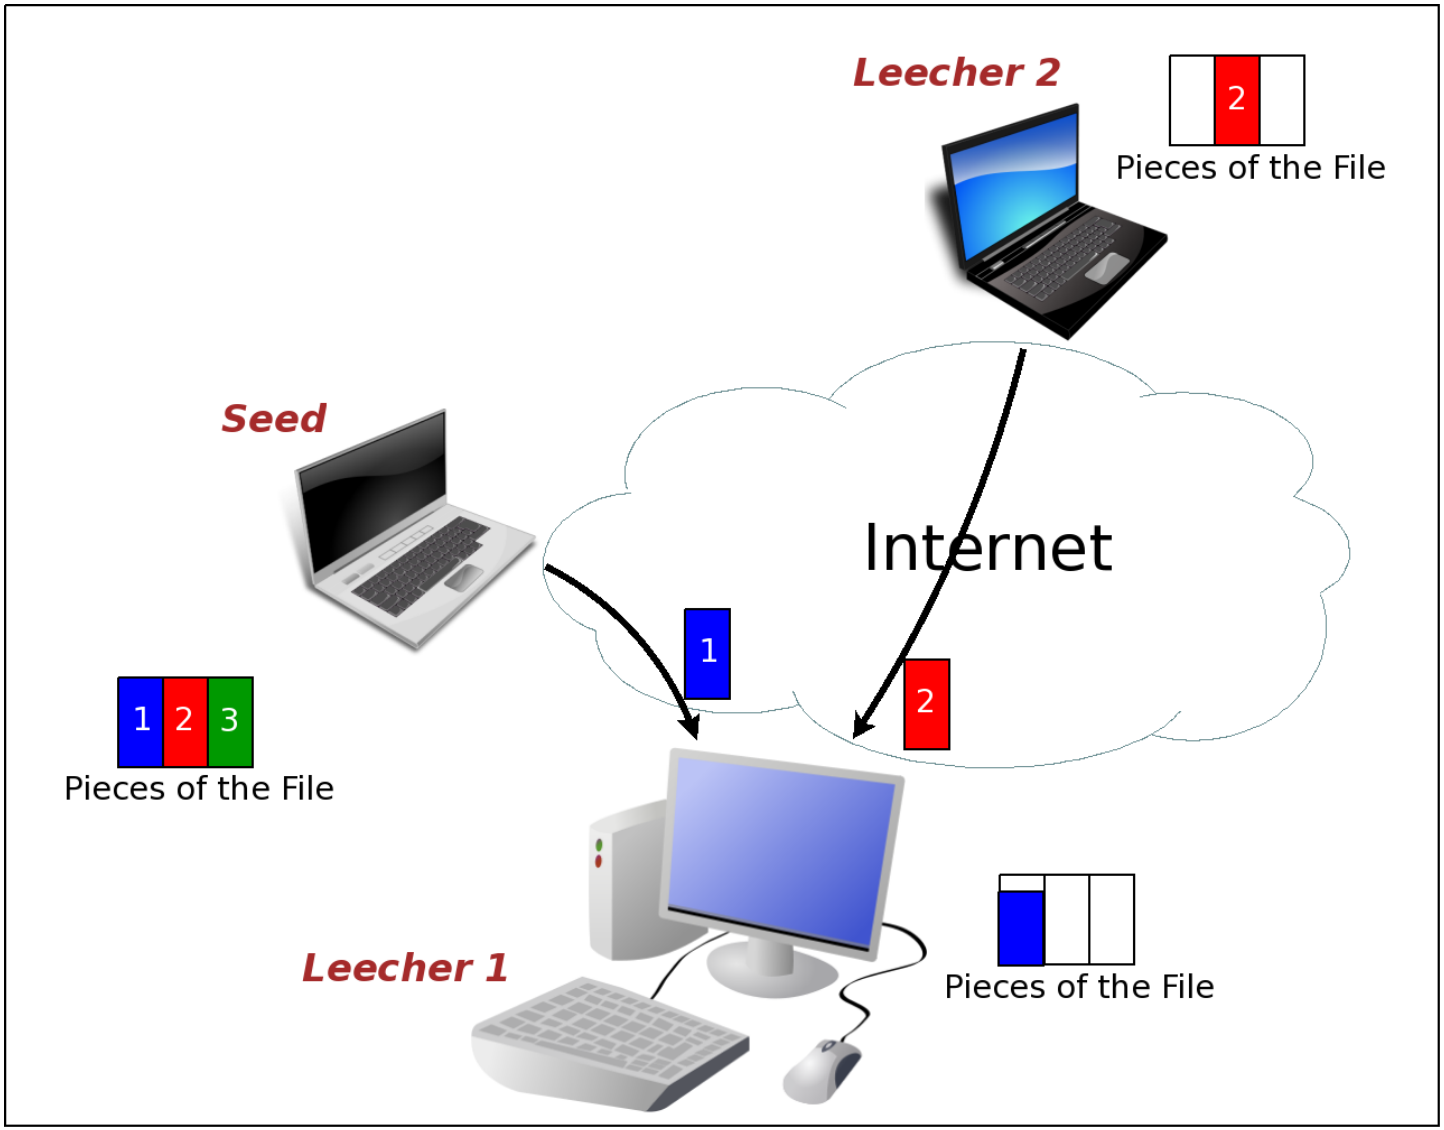

Platforms for BitTorrent Experiments

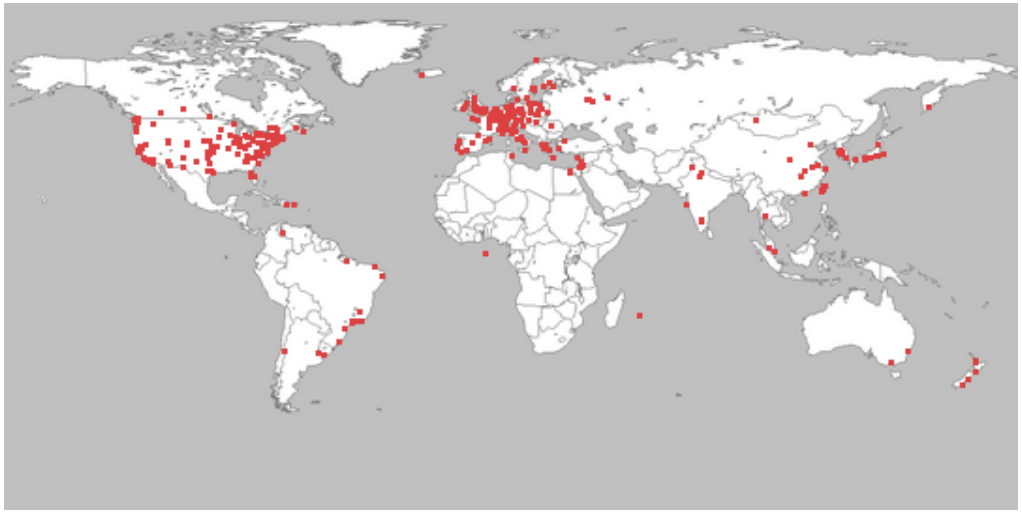

PlanetLab sites [planet-lab.org]

Platforms for BitTorrent Experiments

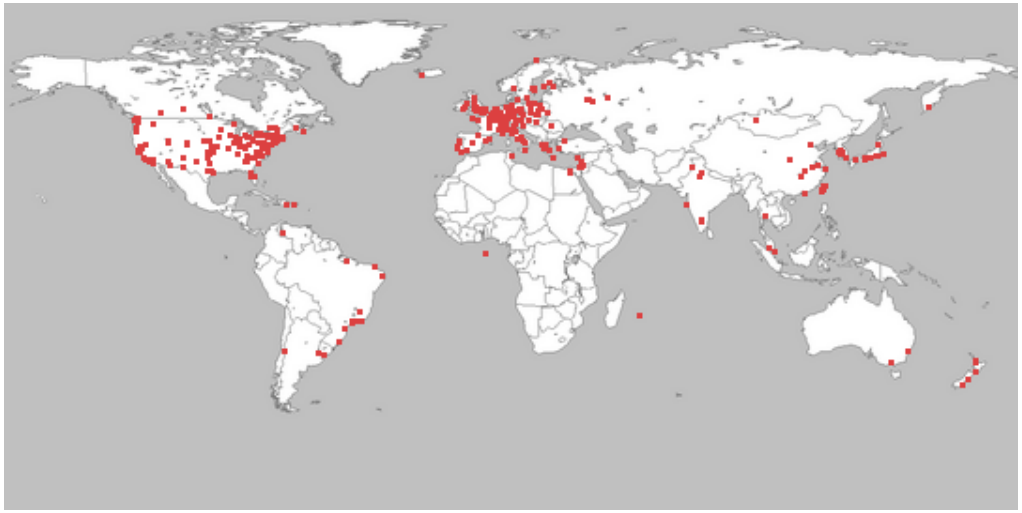

PlanetLab sites [planet-lab.org]

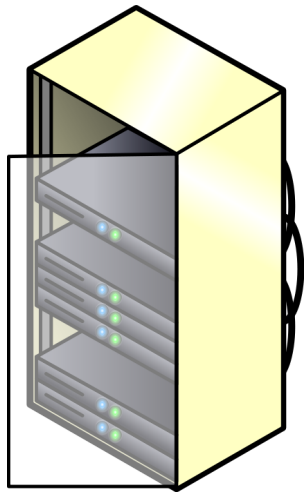

Clusters

Platforms for BitTorrent Experiments

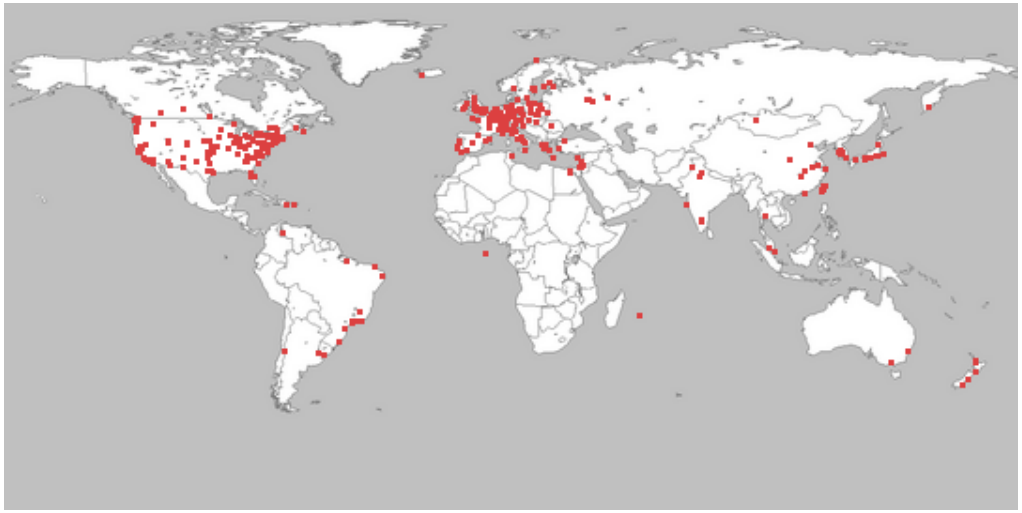

PlanetLab sites [planet-lab.org]

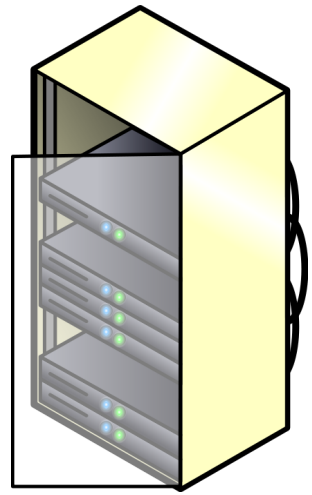

Clusters

	PlanetLab	Clusters
Reproducibility	✗	✓
Scalability	✗	✓

Platforms for BitTorrent Experiments

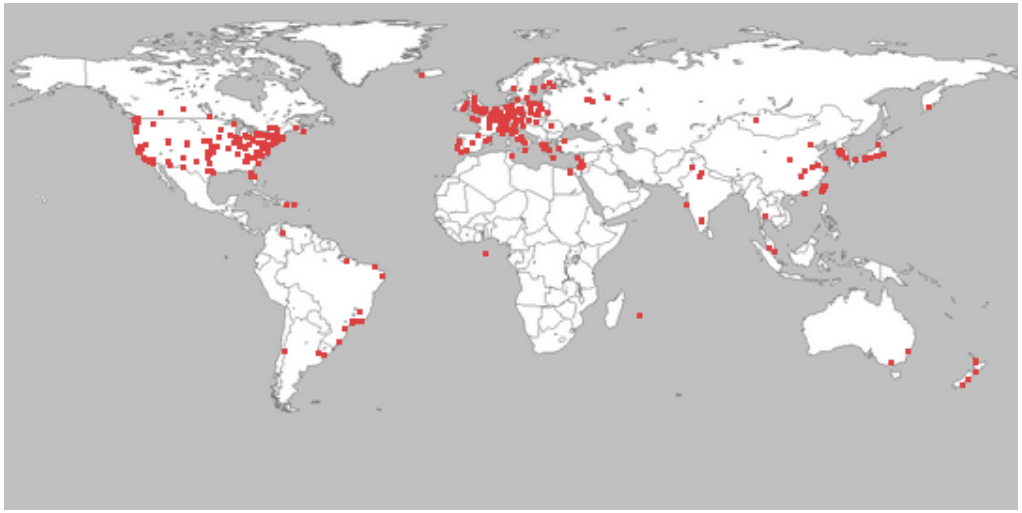

PlanetLab sites [planet-lab.org]

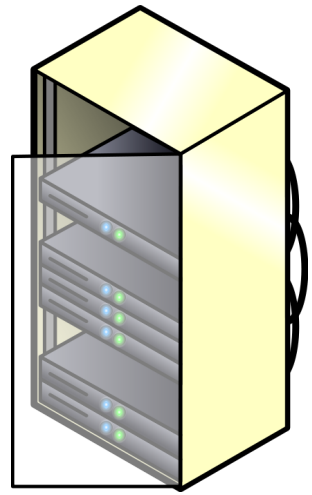

Clusters

	PlanetLab	Clusters
Reproducibility	✗	✓
Scalability	✗	✓
RTT & Packet Loss	✓	?

Platforms for BitTorrent Experiments

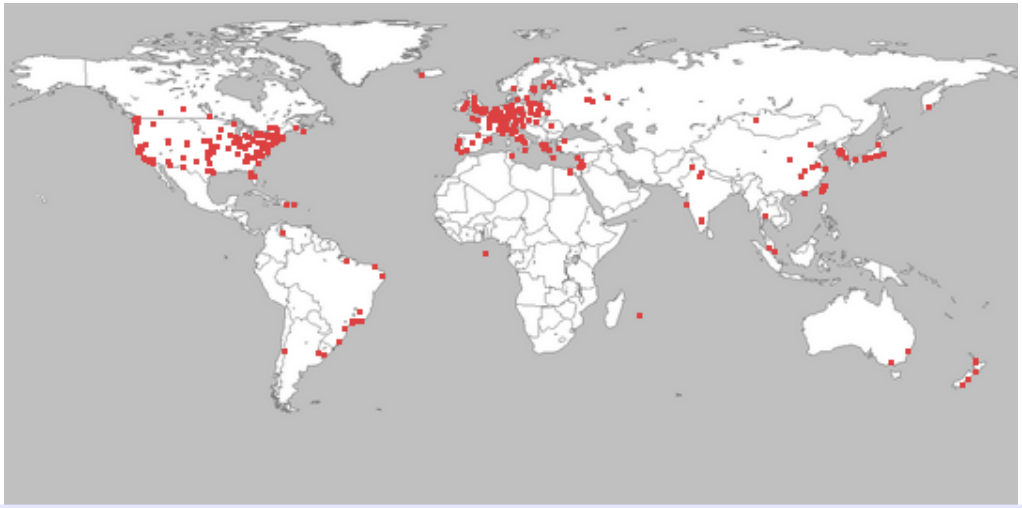

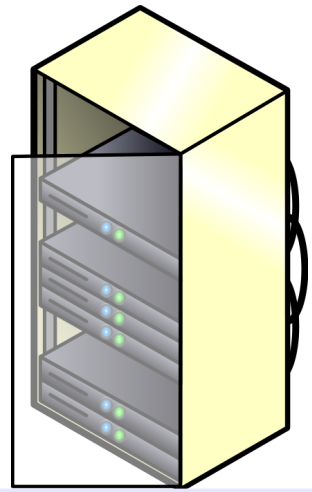

Experiments on Clusters Without Emulating RTT & Packet Loss?

Reproducibility	✗	✓
Scalability	✗	✓
RTT & Packet Loss	✓	?

Impact of RTT and Packet Loss

- BitTorrent Protocol Dynamics

Impact of RTT and Packet Loss

- BitTorrent Protocol Dynamics

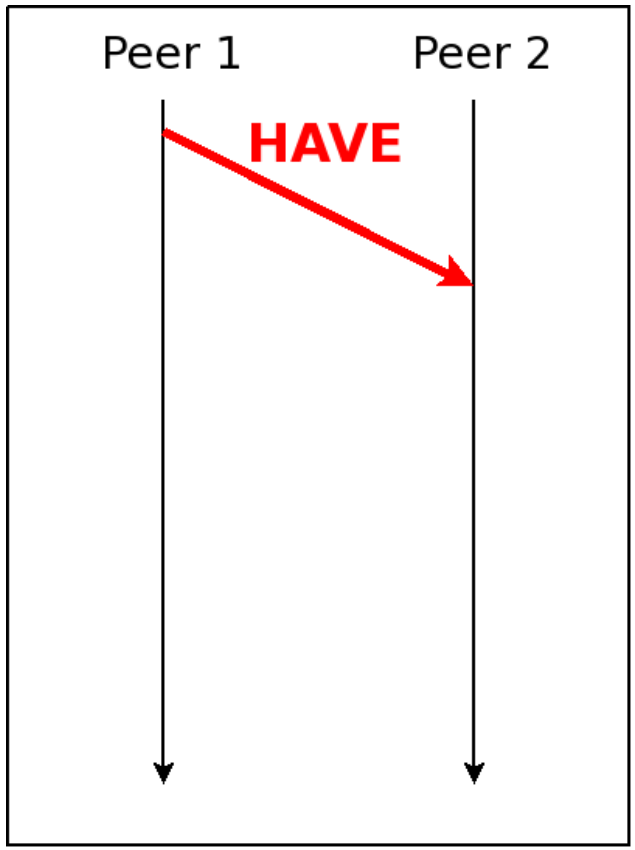

Impact of RTT and Packet Loss

- BitTorrent Protocol Dynamics

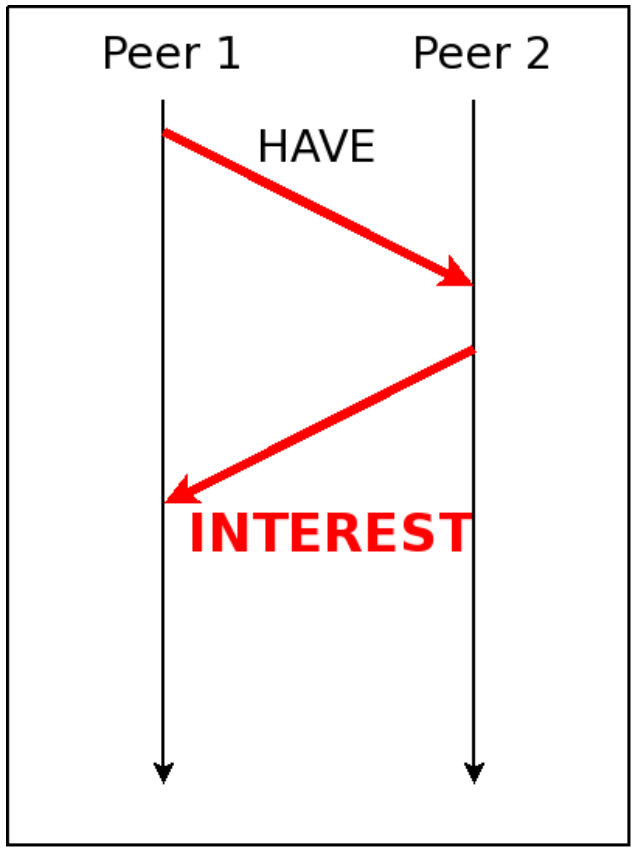

Impact of RTT and Packet Loss

- BitTorrent Protocol Dynamics

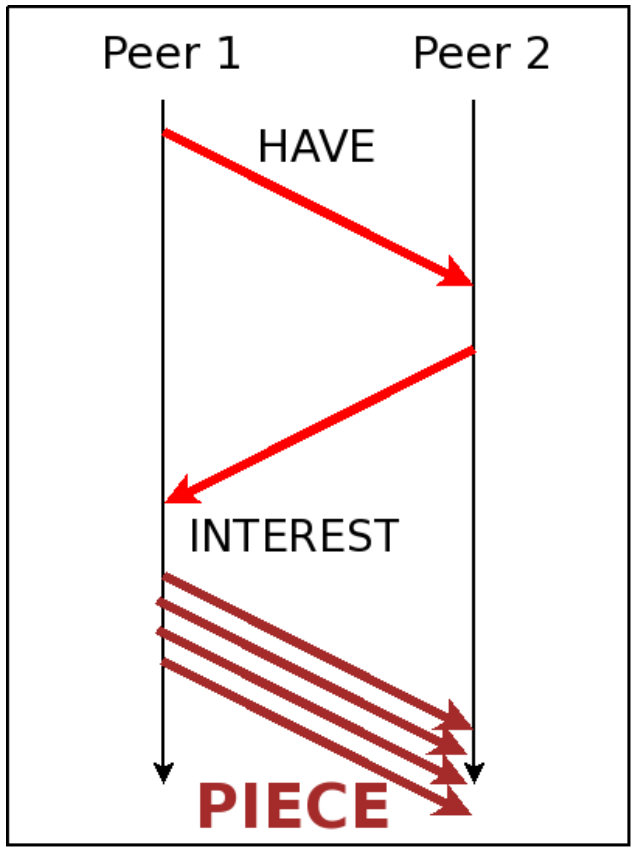

Impact of RTT and Packet Loss

- BitTorrent Protocol Dynamics
  - Pipeline Requests
  - Time between decisions $\approx$ 10 seconds

Impact of RTT and Packet Loss

- BitTorrent Protocol Dynamics
  - Pipeline Requests
  - Time between decisions $\approx$ 10 seconds
- Throughput

Impact of RTT and Packet Loss

- BitTorrent Protocol Dynamics
  - Pipeline Requests
  - Time between decisions $\approx$ 10 seconds
- Throughput

$$Throughput \propto \frac{1}{RTT\sqrt{Loss\ Rate}}$$

Impact of RTT and Packet Loss

- BitTorrent Protocol Dynamics
  - Pipeline Requests
  - Time between decisions $\approx$ 10 seconds
- Throughput
  - Users limit Upload Rate

Impact of RTT and Packet Loss

- BitTorrent Protocol Dynamics
  - Pipeline Requests
  - Time between decisions $\approx$ 10 seconds
- Throughput
  - Users limit Upload Rate

Need to Validate Impact of RTT & Packet Loss

Outline

Motivation

- Experiment Setup
- Impact of RTT
- Impact of Packet Loss

Testbed Setup

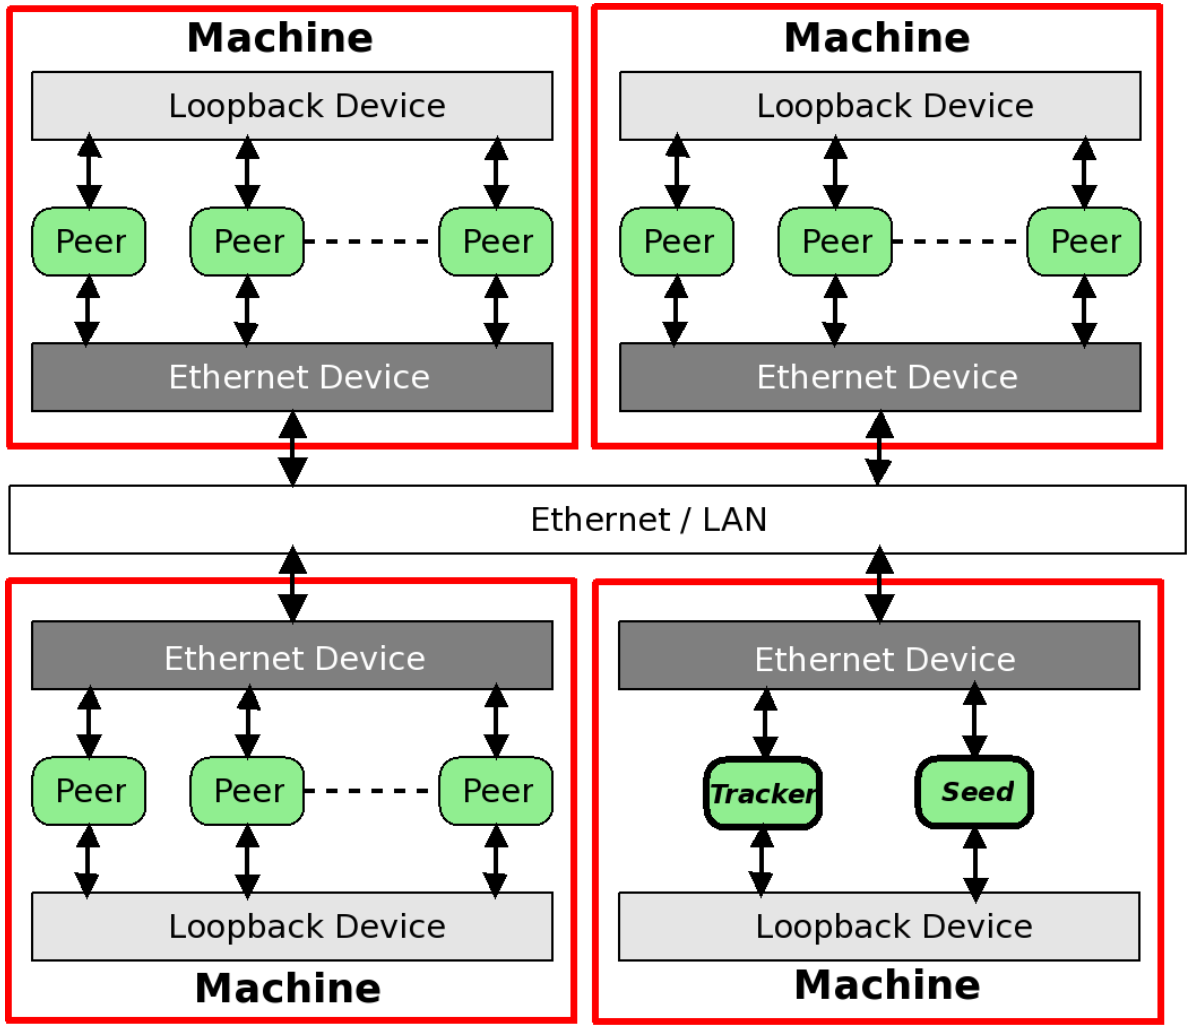

4 Machines of a Cluster

Testbed Setup

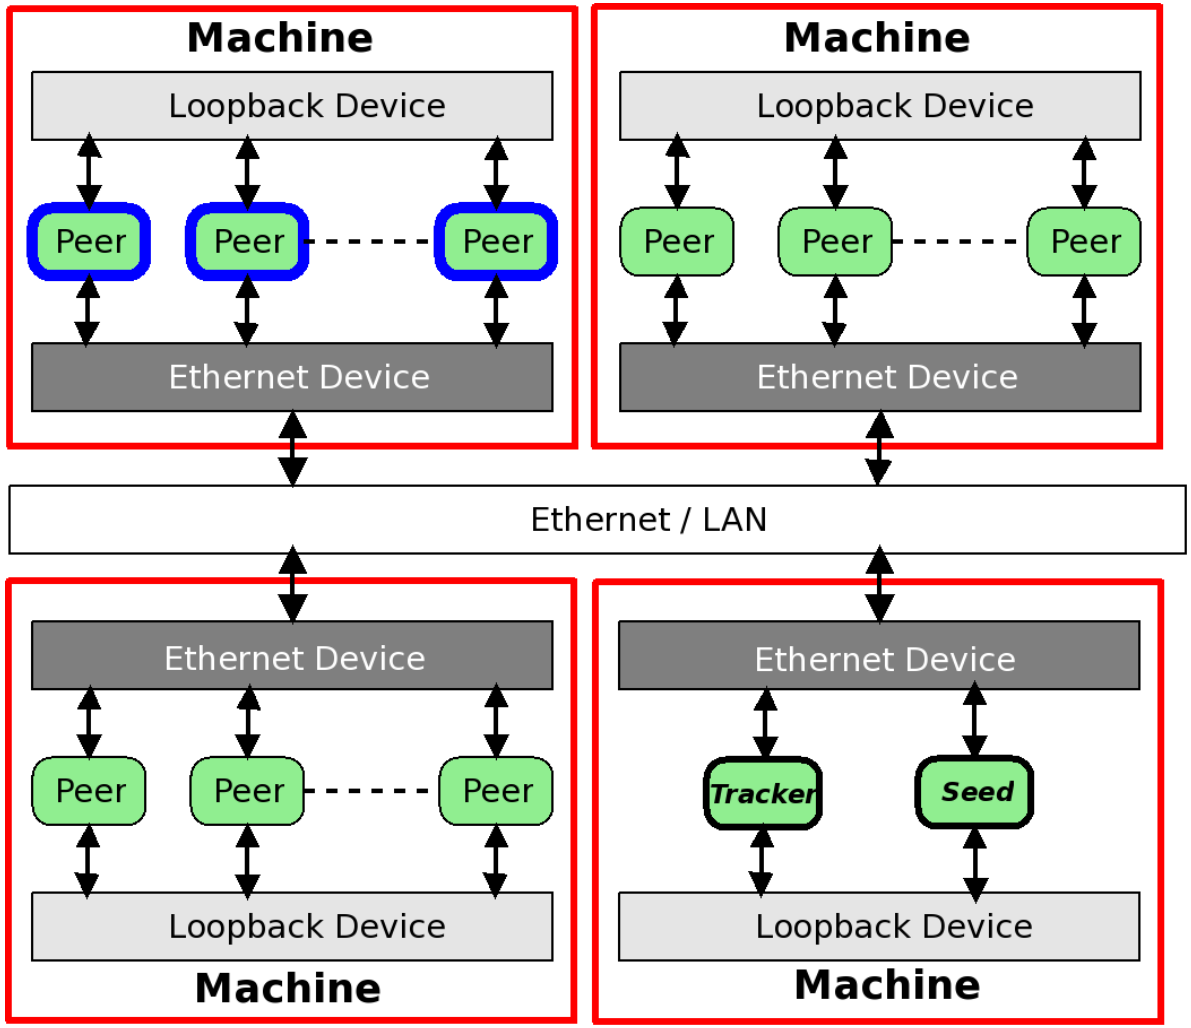

A machine can have 100 instances of Real BitTorrent Client

1 tracker, 1 seed, 300 leechers

Testbed Setup

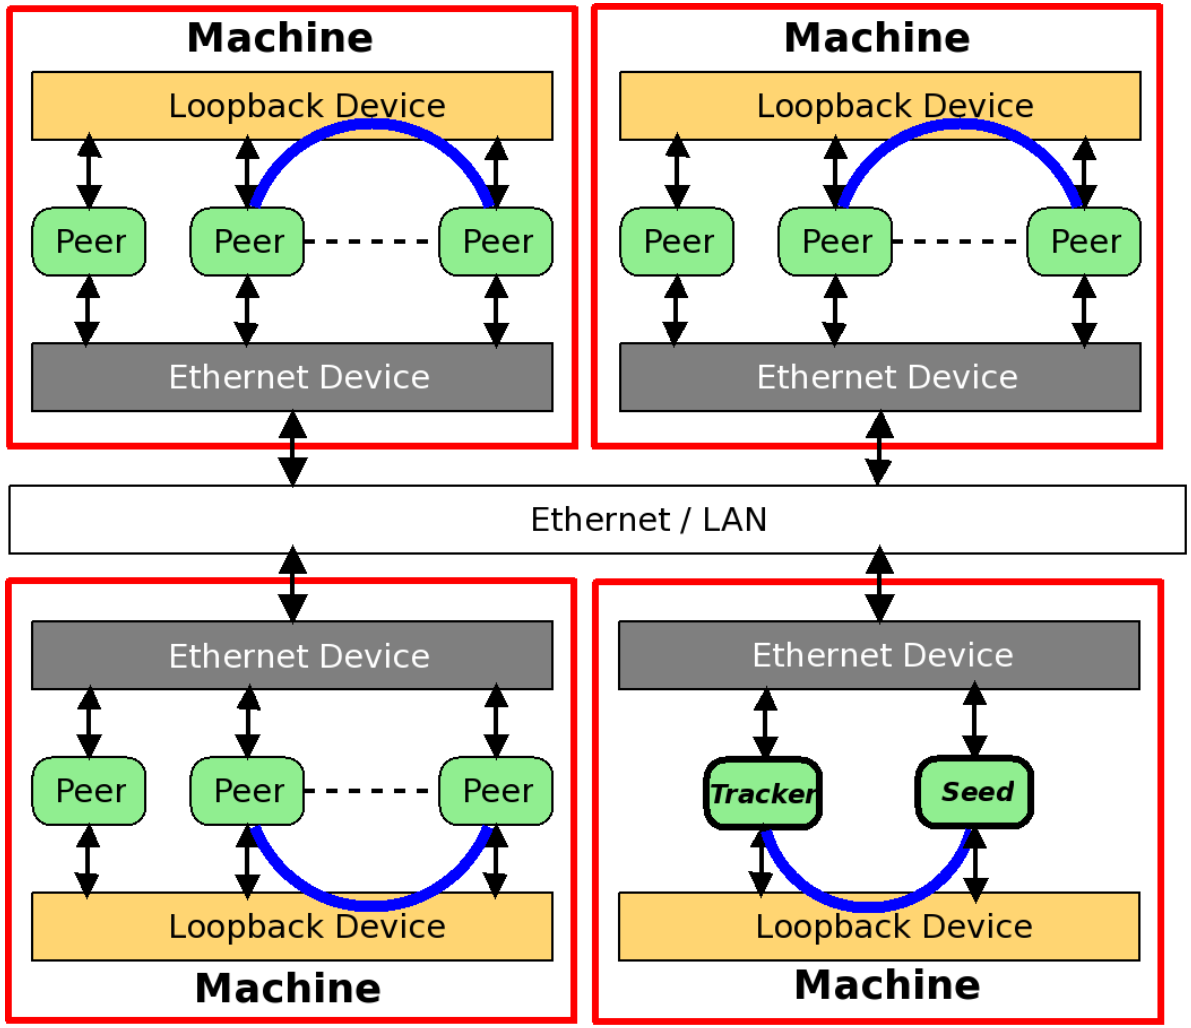

Loopback device: Peers on the same machine

Testbed Setup

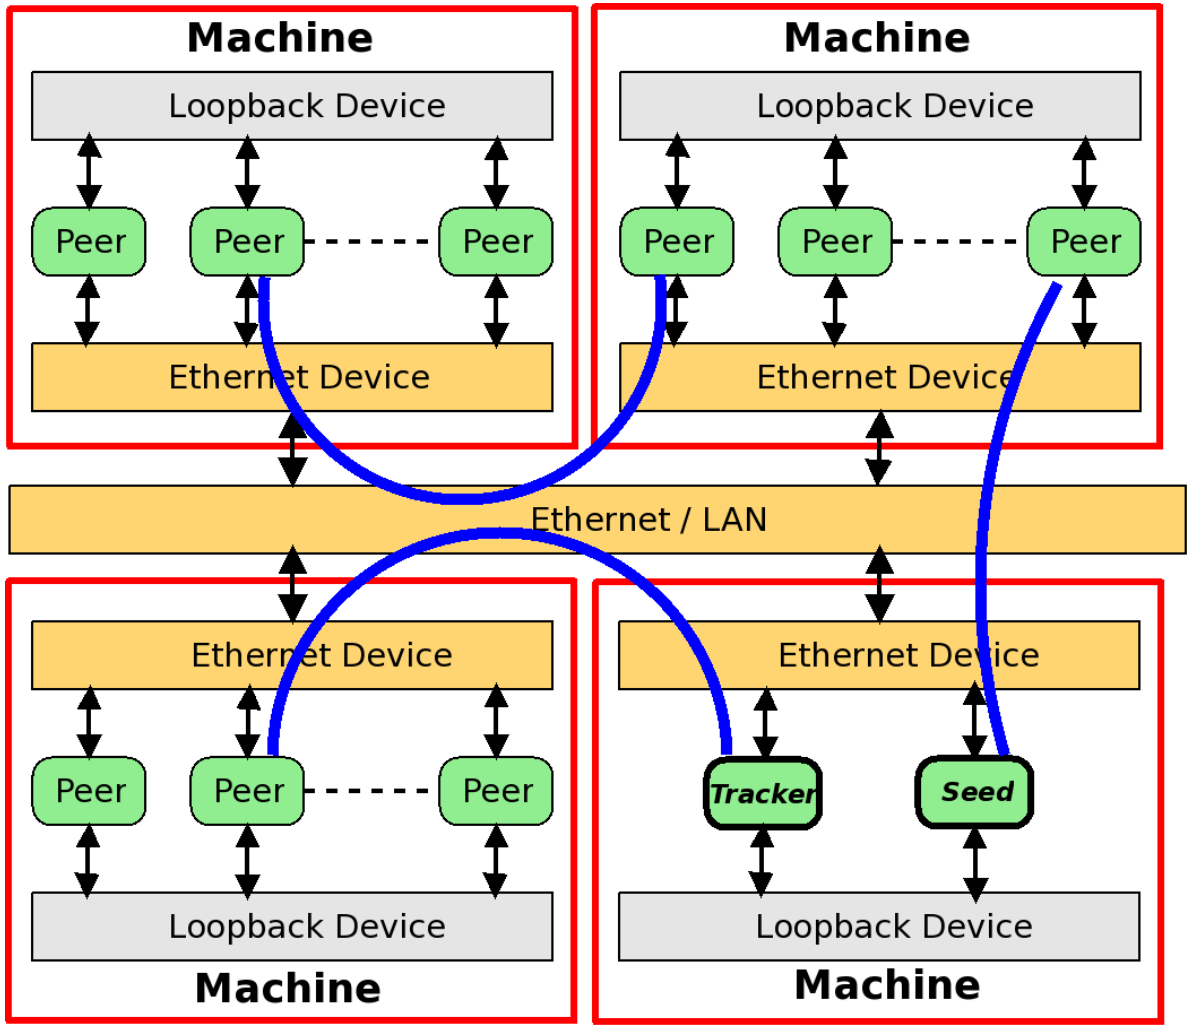

Ethernet device: Peers on adjacent machines

Testbed Setup

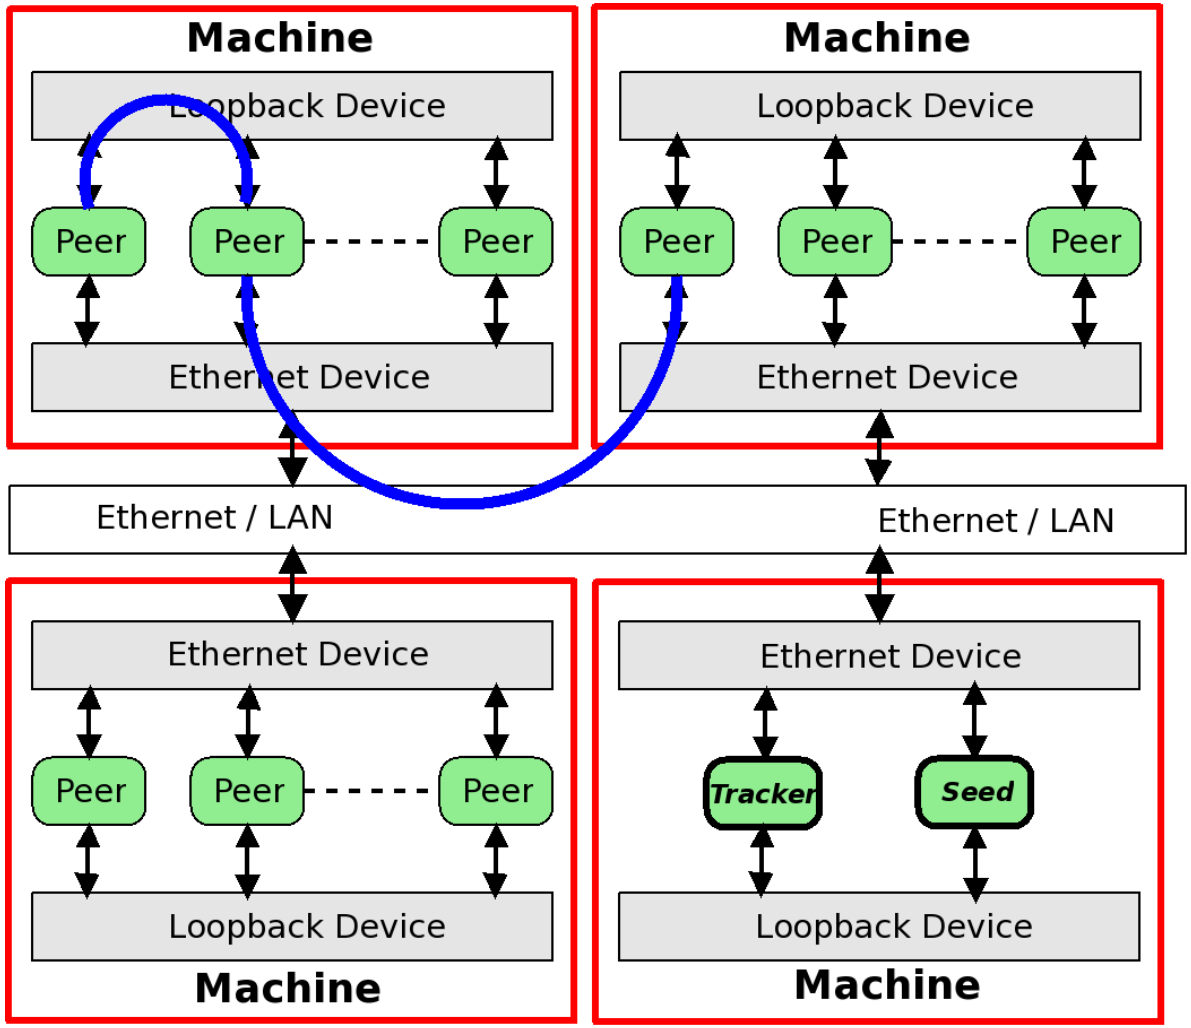

Emulate Same RTT and Packet Loss Rate

Experiment Parameters and Metric

iPlane Measurement: 95% Paths

- RTT $\leq$ 400 ms
- Packet Loss Rate $\leq$ 5%

iPlane: An Information Plane for Distributed Services.
H. Madhyastha *et al.*, OSDI 2006.
http://iplane.cs.washington.edu/

Experiment Parameters and Metric

iPlane Measurement: 95% Paths

- RTT $\leq$ 400 ms
- Packet Loss Rate $\leq$ 5%

Experiment Parameters

- RTT - up to 1000 ms
- Packet Loss Rate - up to 5%
- 50 MB file

Metric

- Download completion time

Outline

Motivation

Experiment Setup

- Impact of RTT
- Impact of Packet Loss

Impact of RTT - Experiment Results

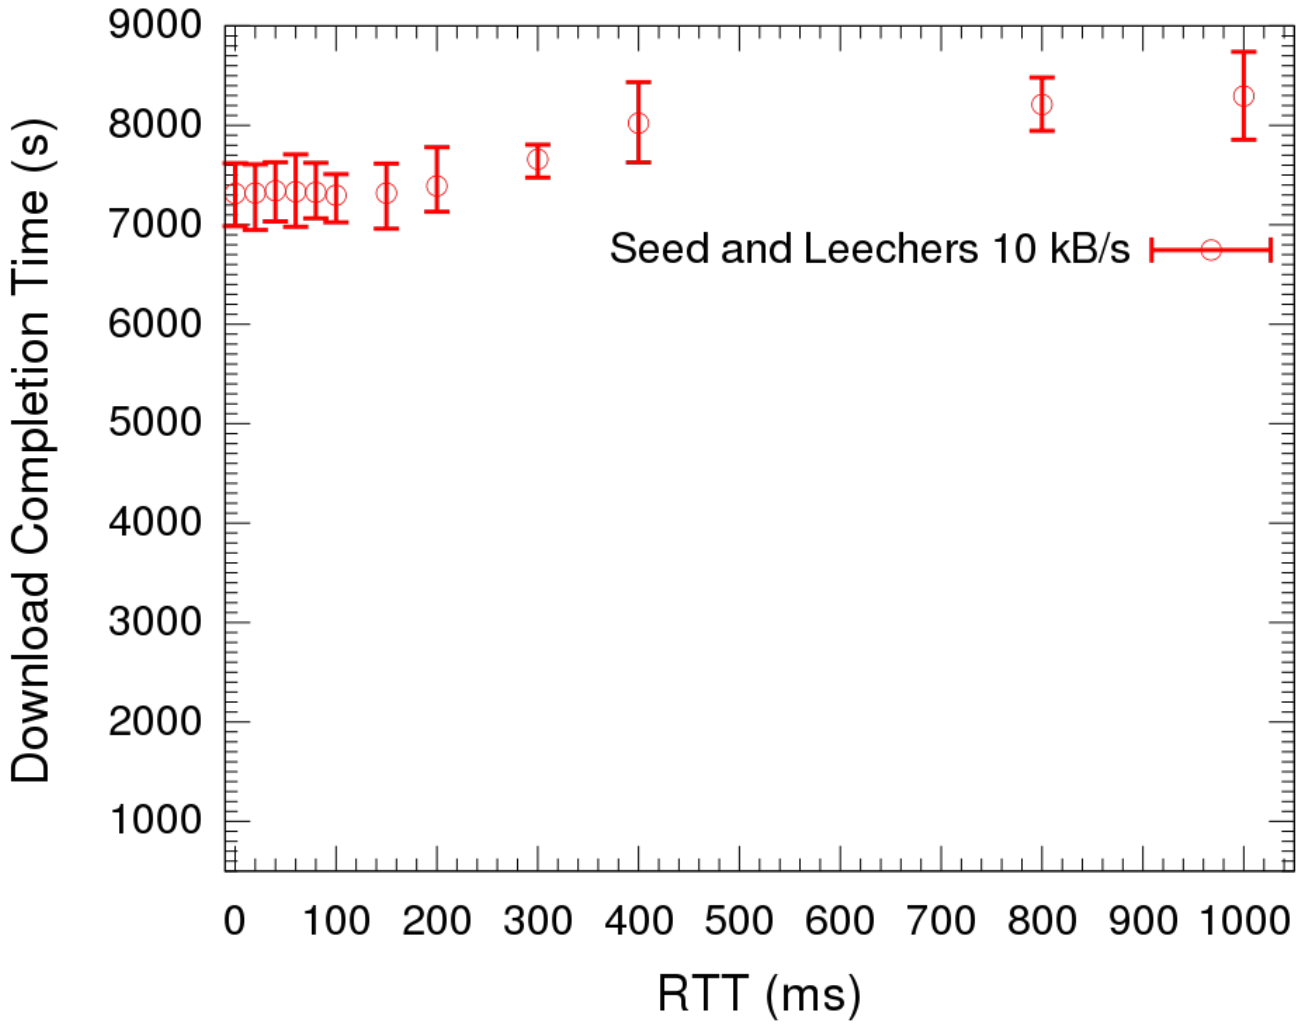

Impact of RTT - Experiment Results

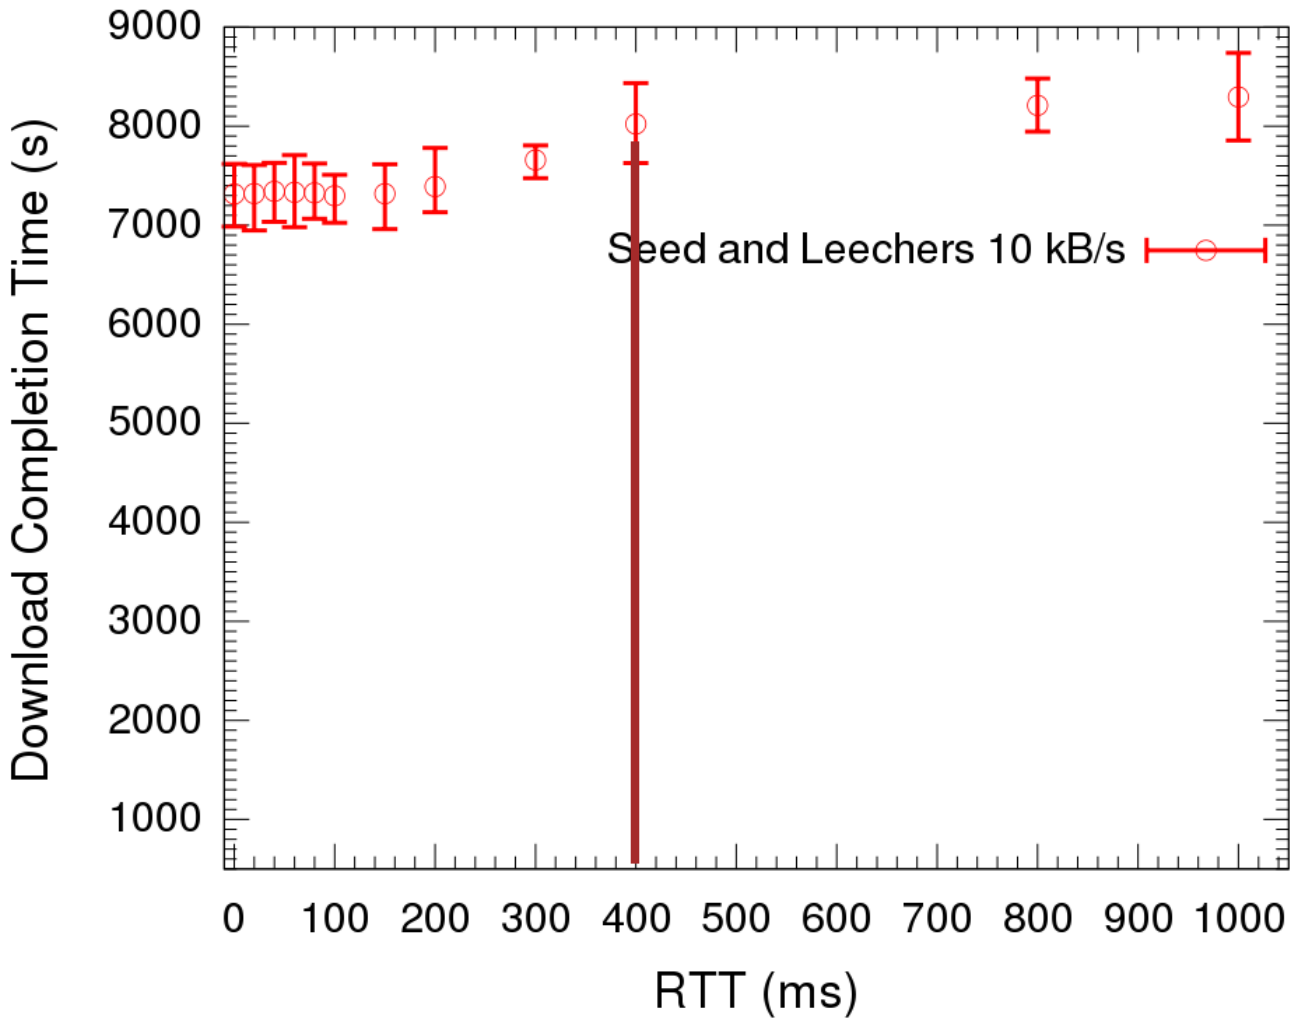

Impact of RTT - Experiment Results

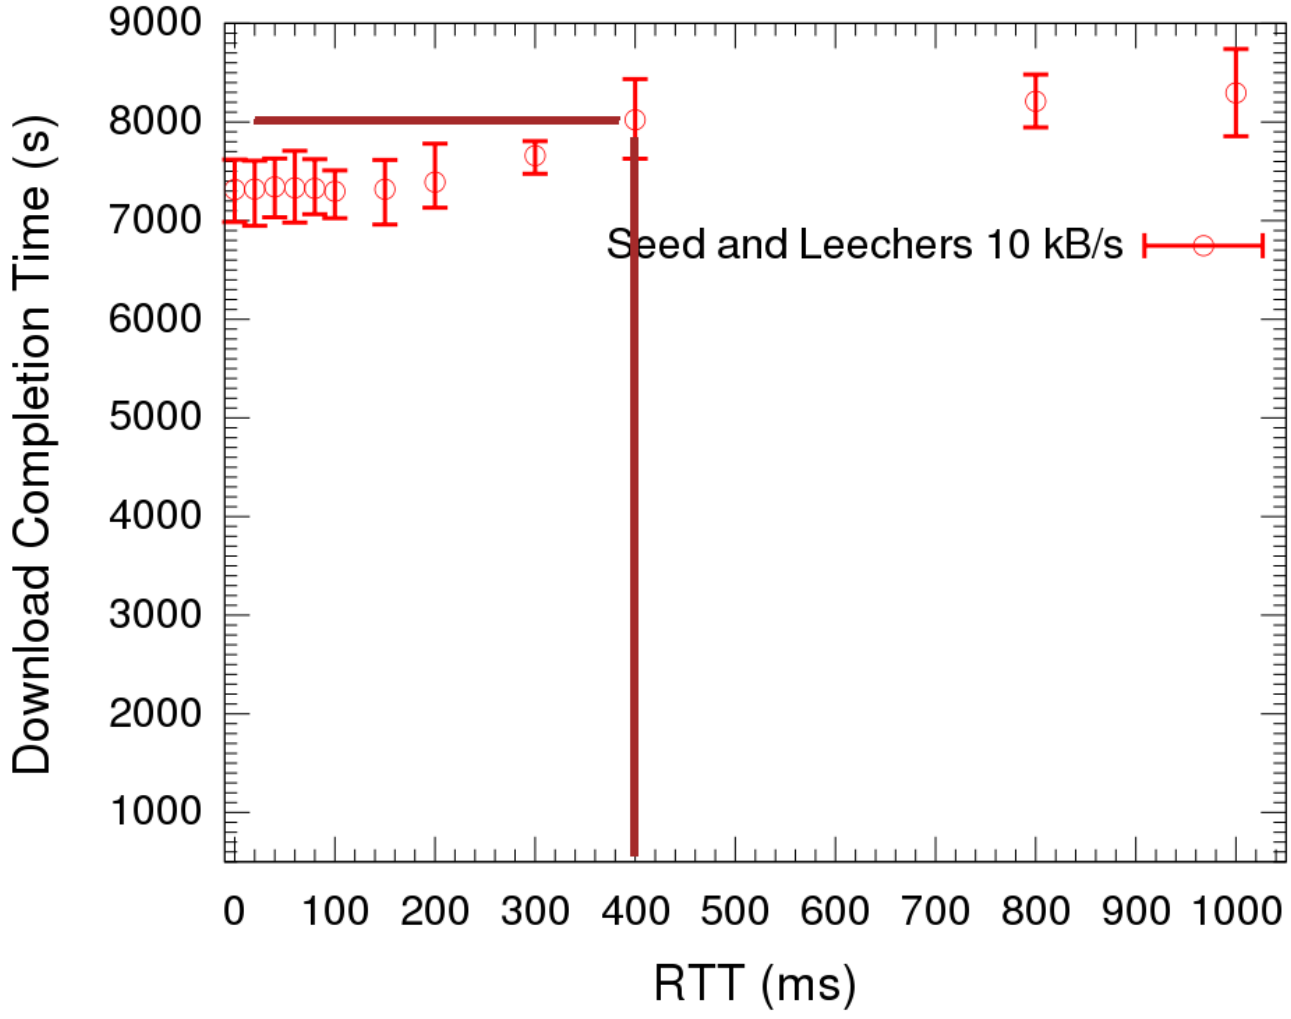

Impact of RTT - Experiment Results

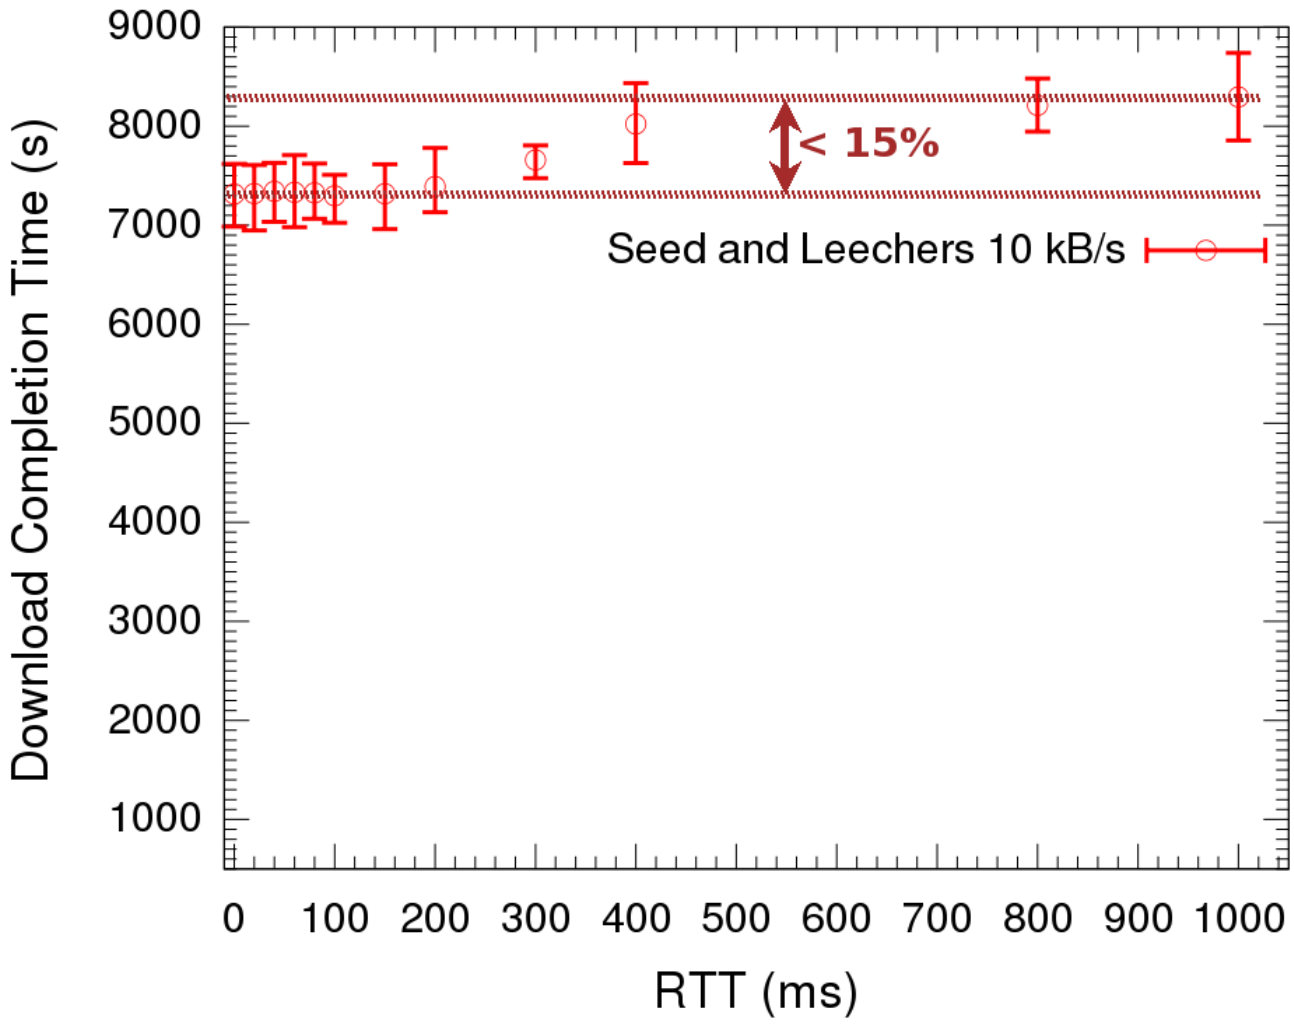

Impact of RTT - Experiment Results

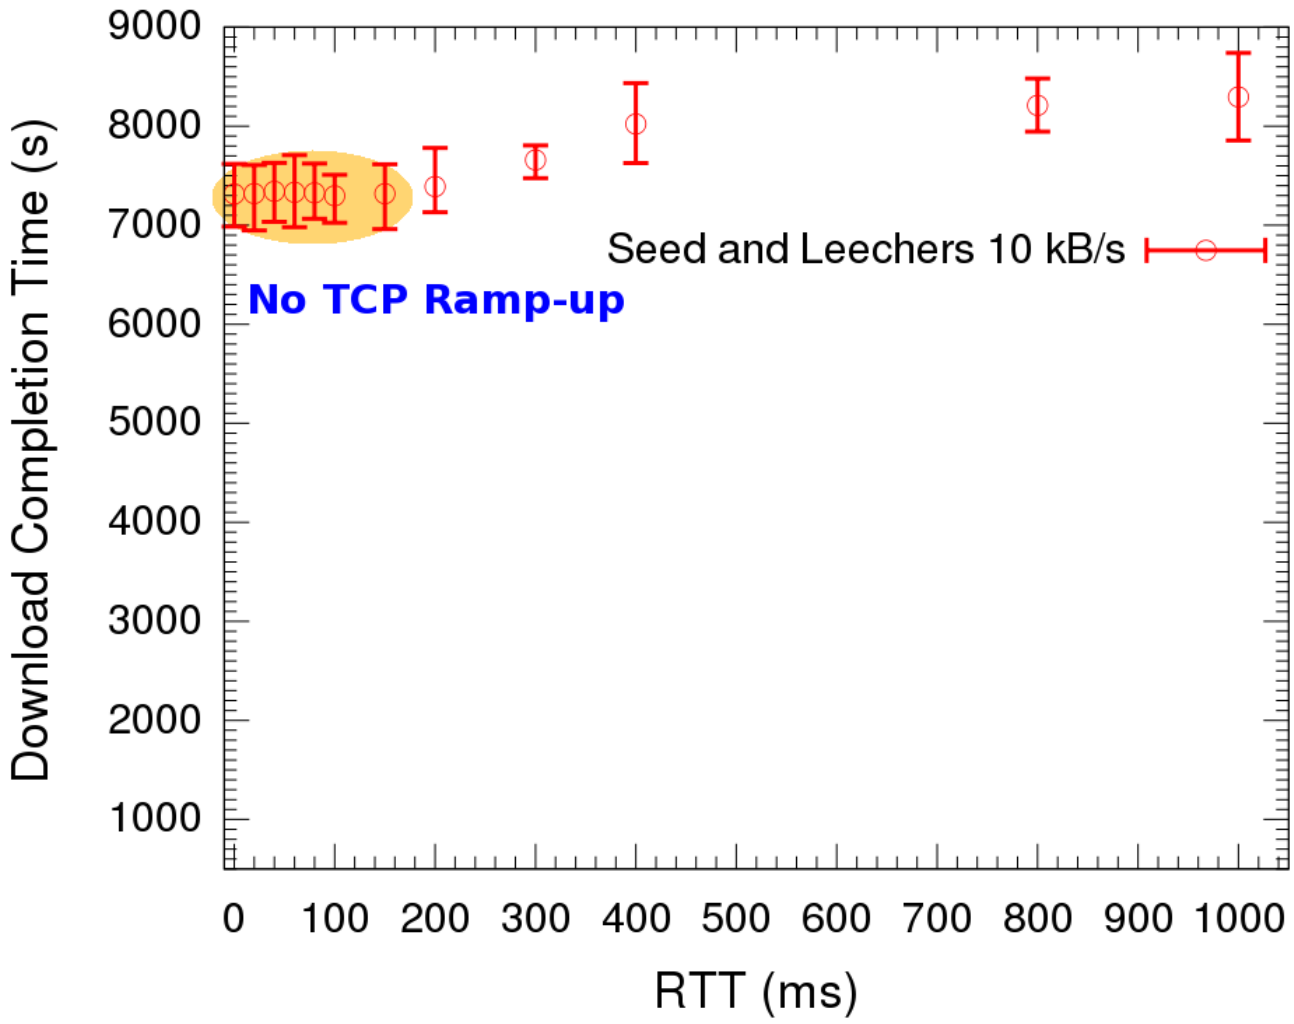

Impact of RTT - Experiment Results

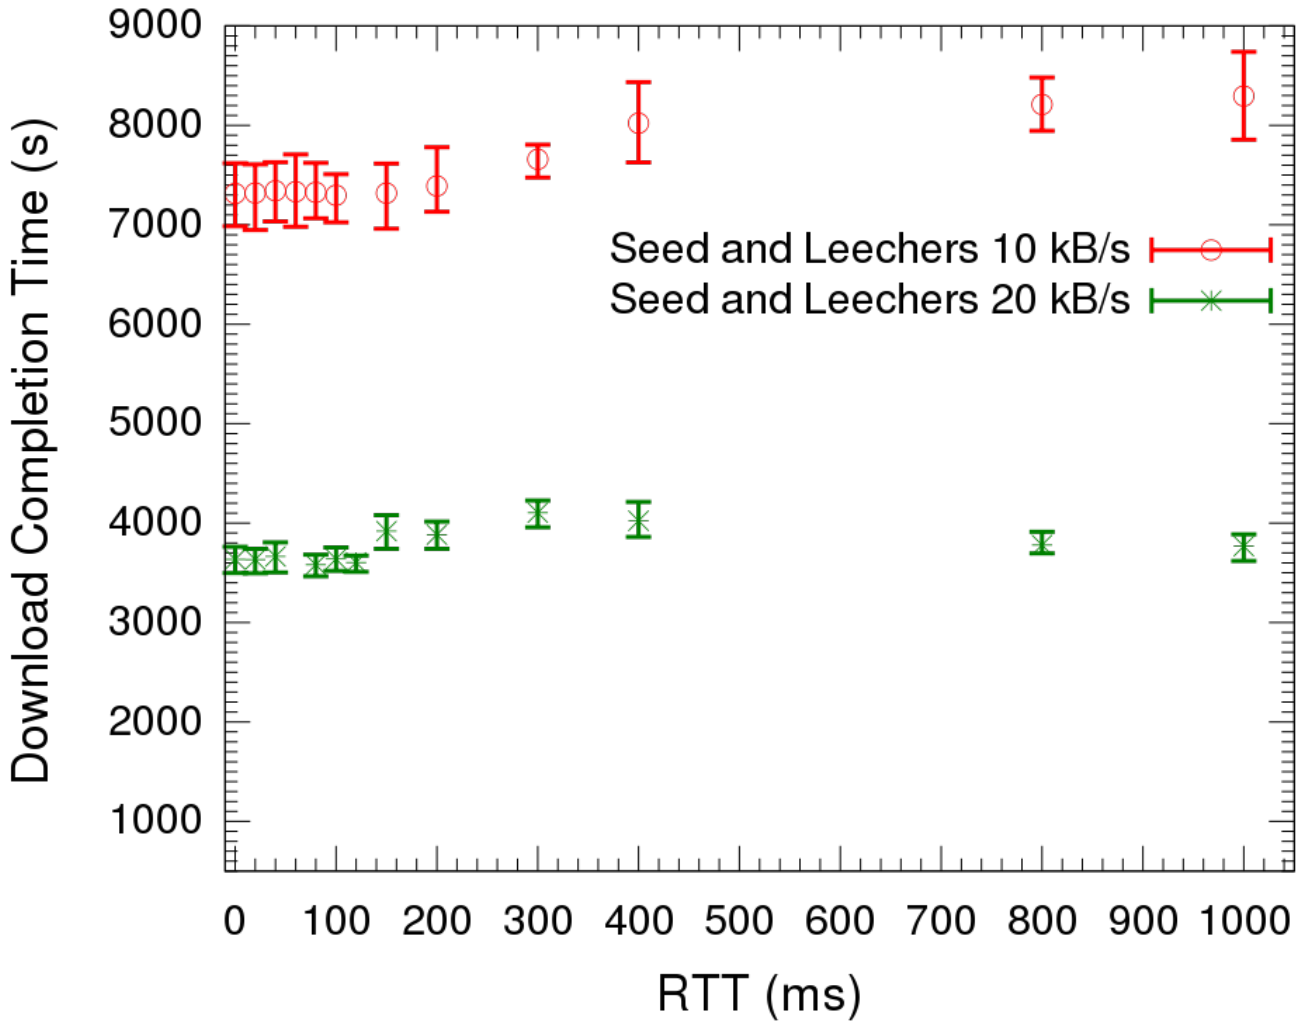

Impact of RTT - Experiment Results

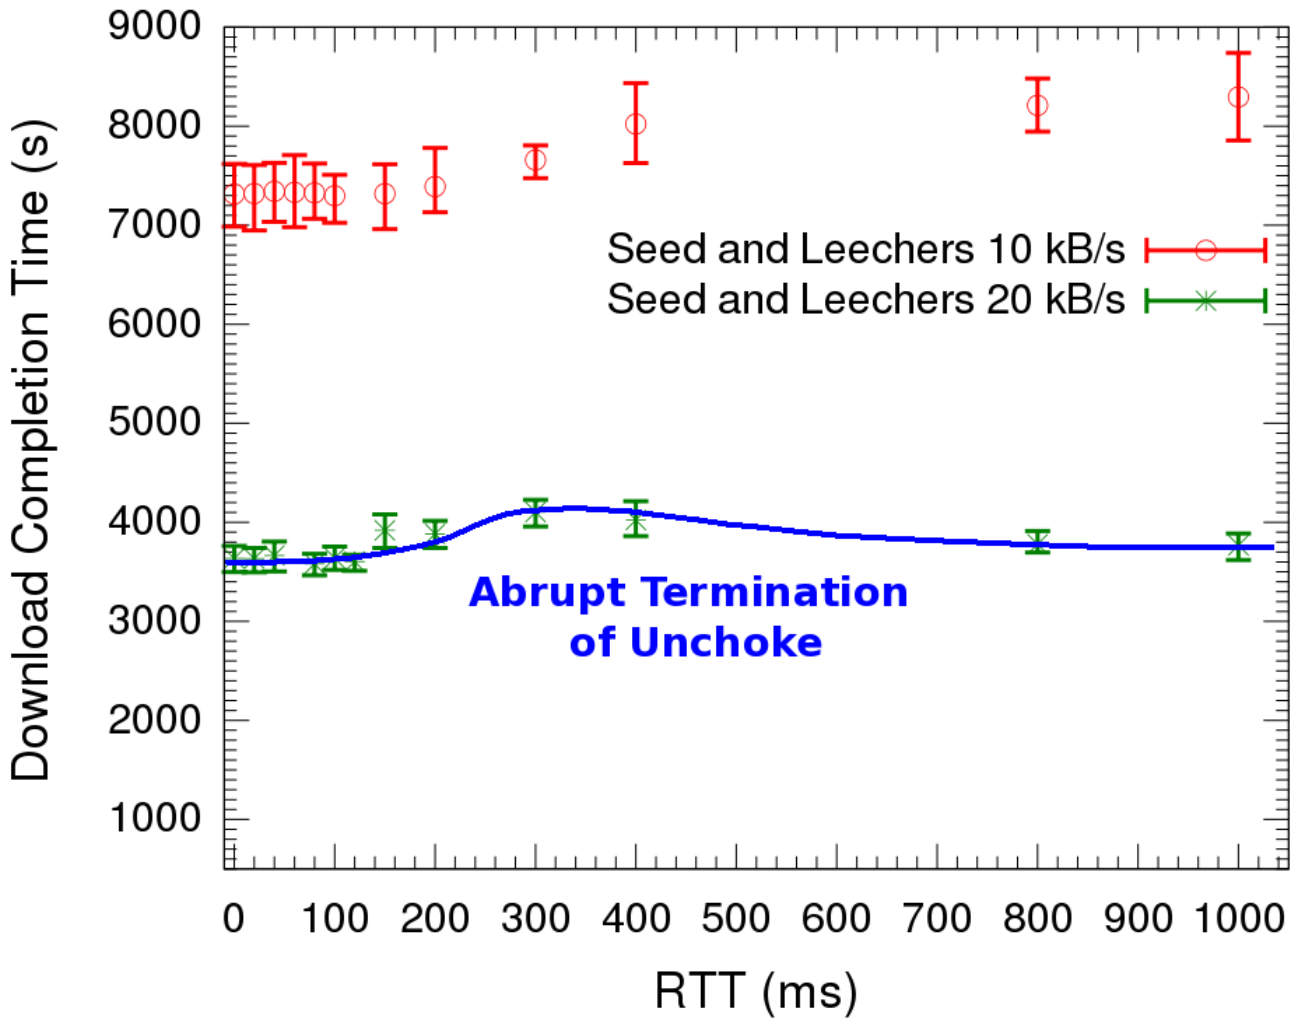

Impact of RTT - Experiment Results

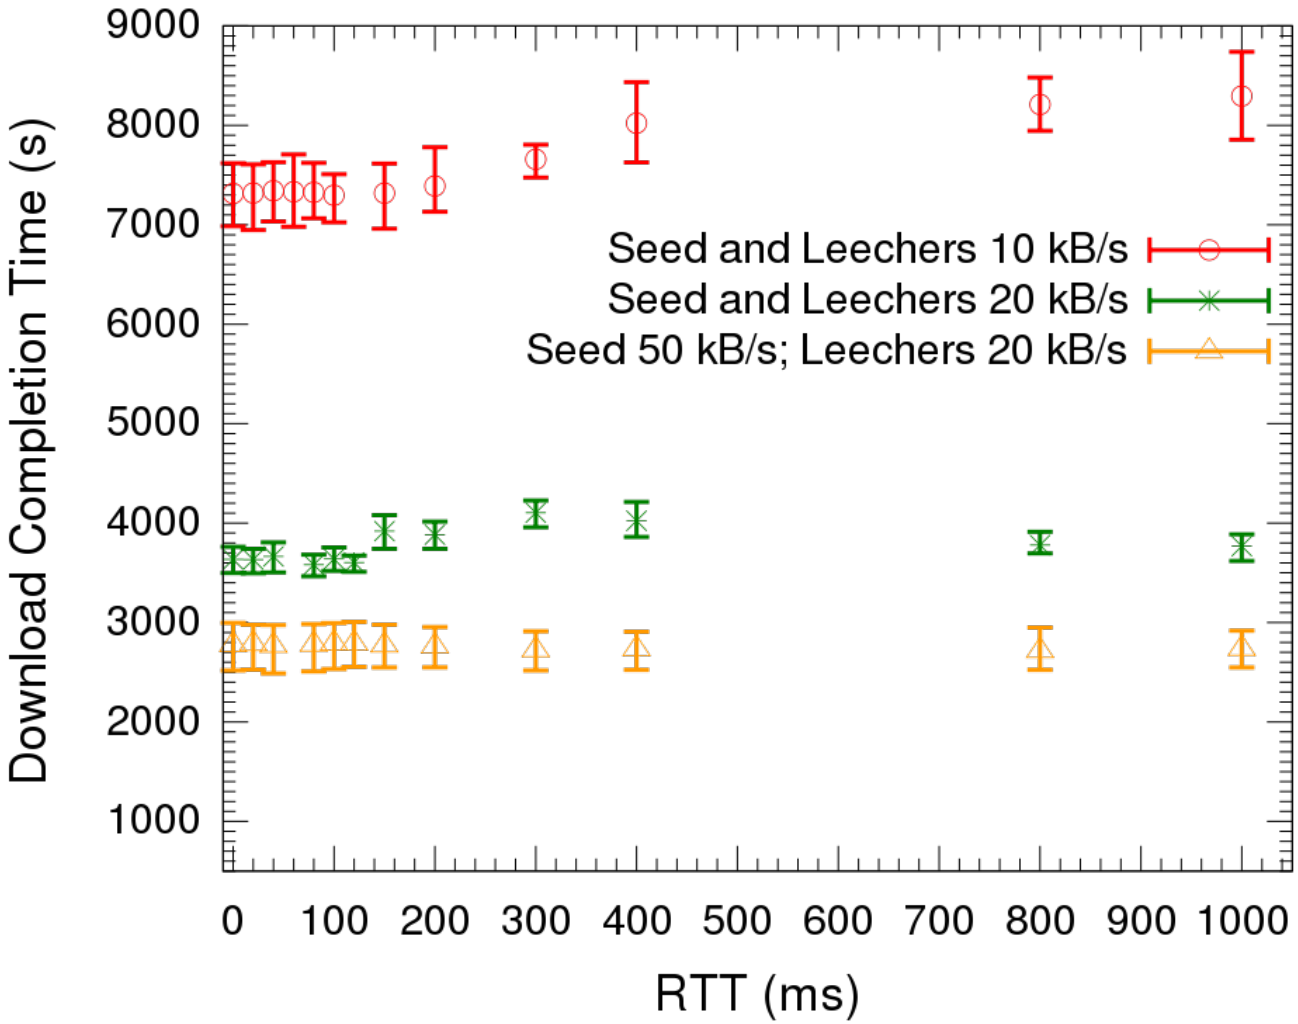

Scenario Fast Seed and Slow Leechers

Impact of RTT - Experiment Results

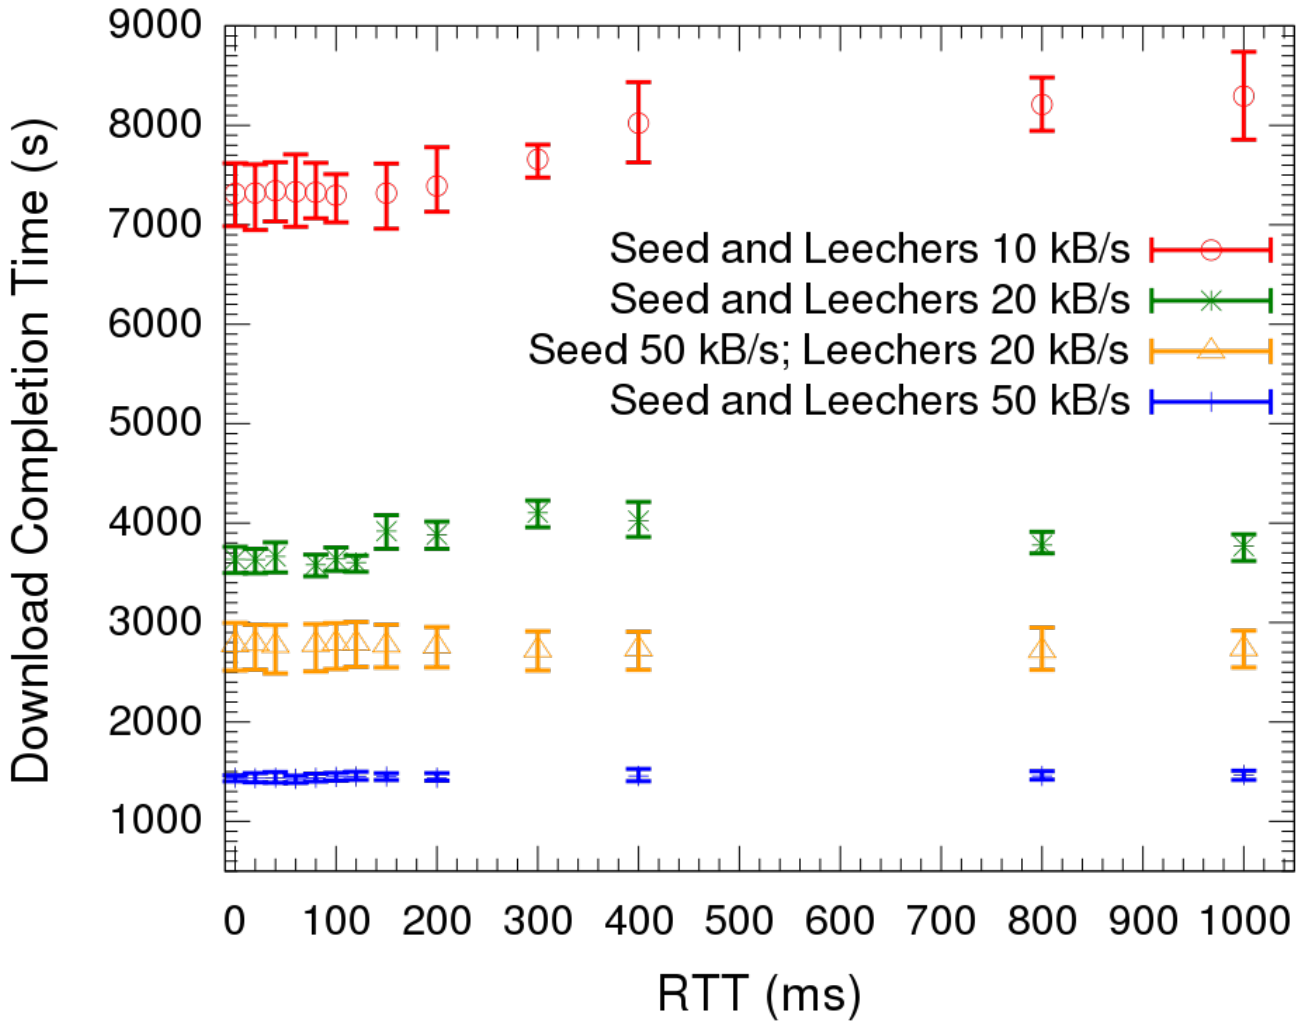

Scenario of Fast Peers

Impact of RTT - Experiment Results

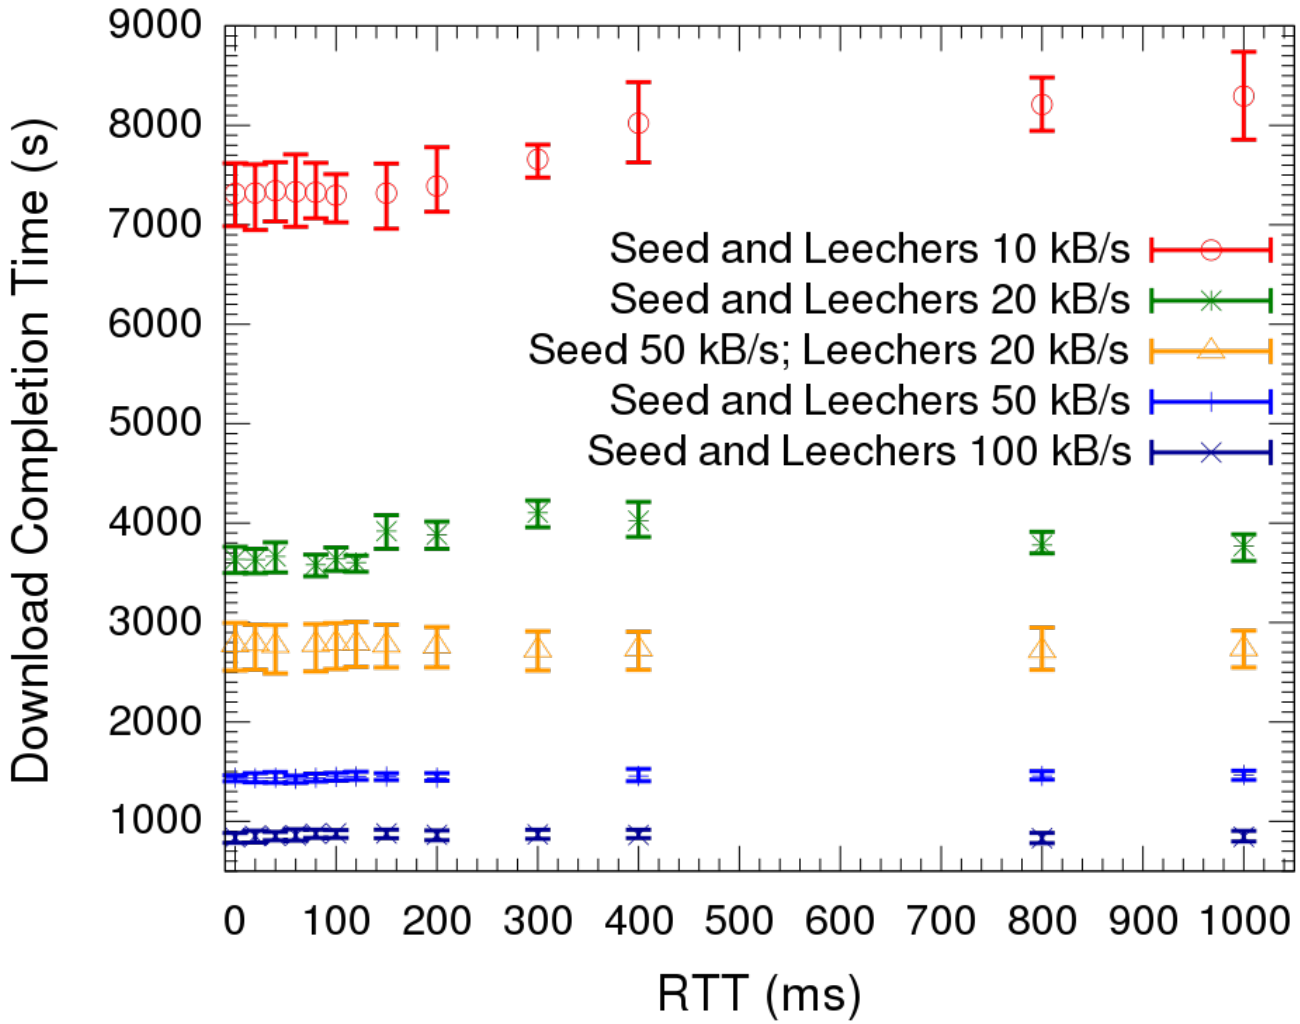

Impact of RTT - Experiment Results

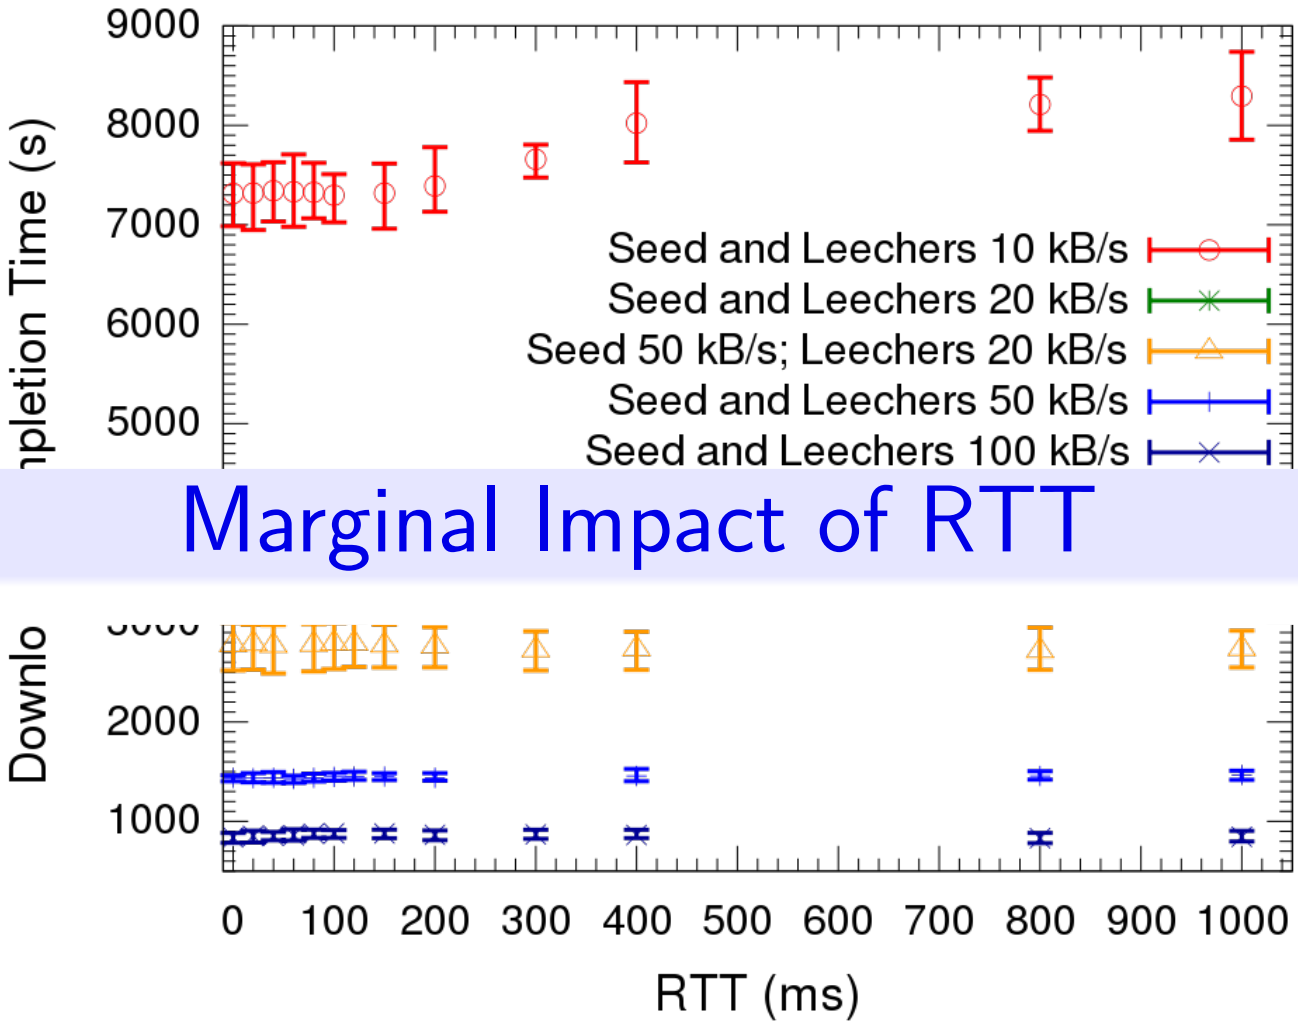

Marginal Impact of RTT

Outline

- Motivation
- Experiment Setup
- Impact of RTT
- Impact of Packet Loss

Impact of Packet Loss (20 kB/s)

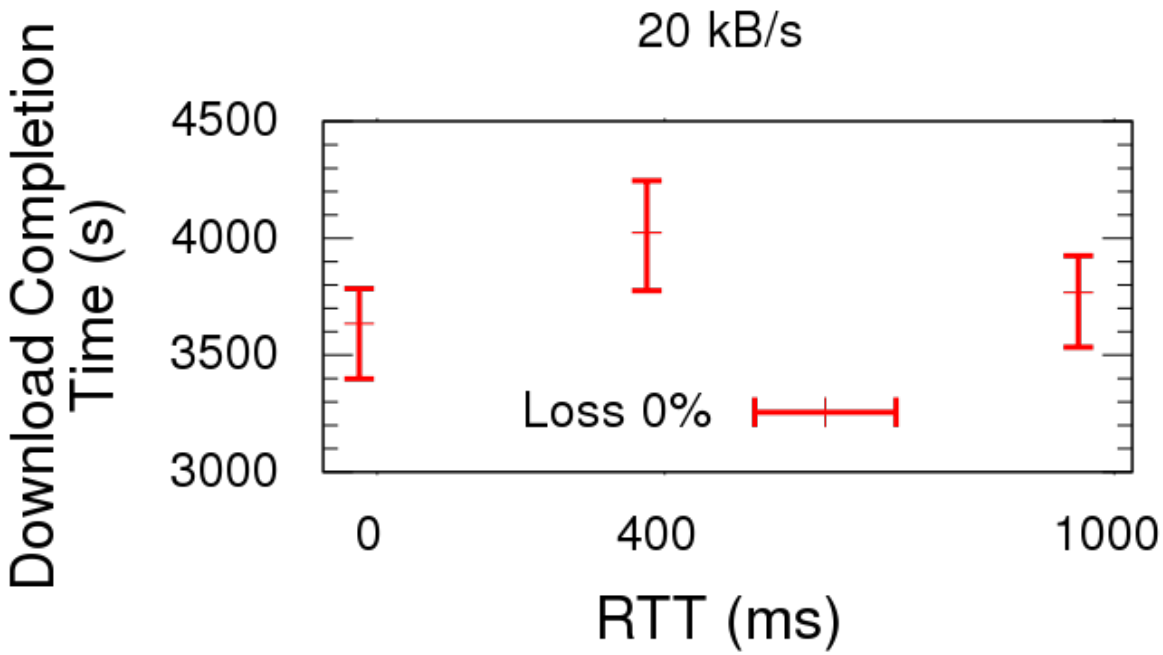

Impact of Packet Loss (20 kB/s)

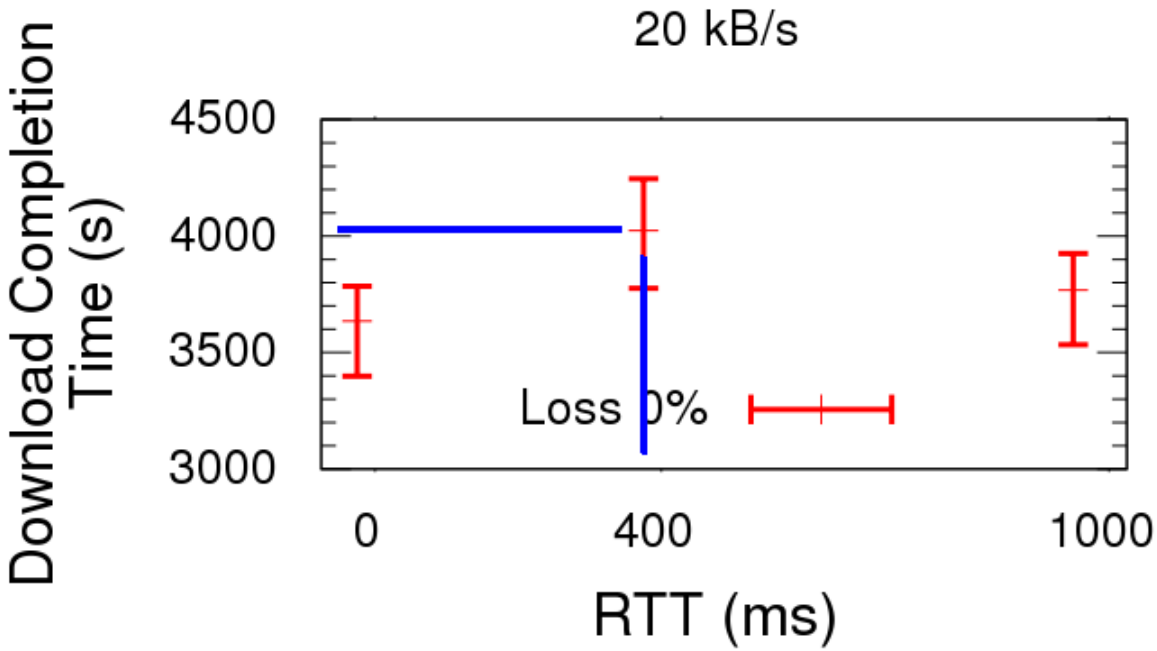

Impact of Packet Loss (20 kB/s)

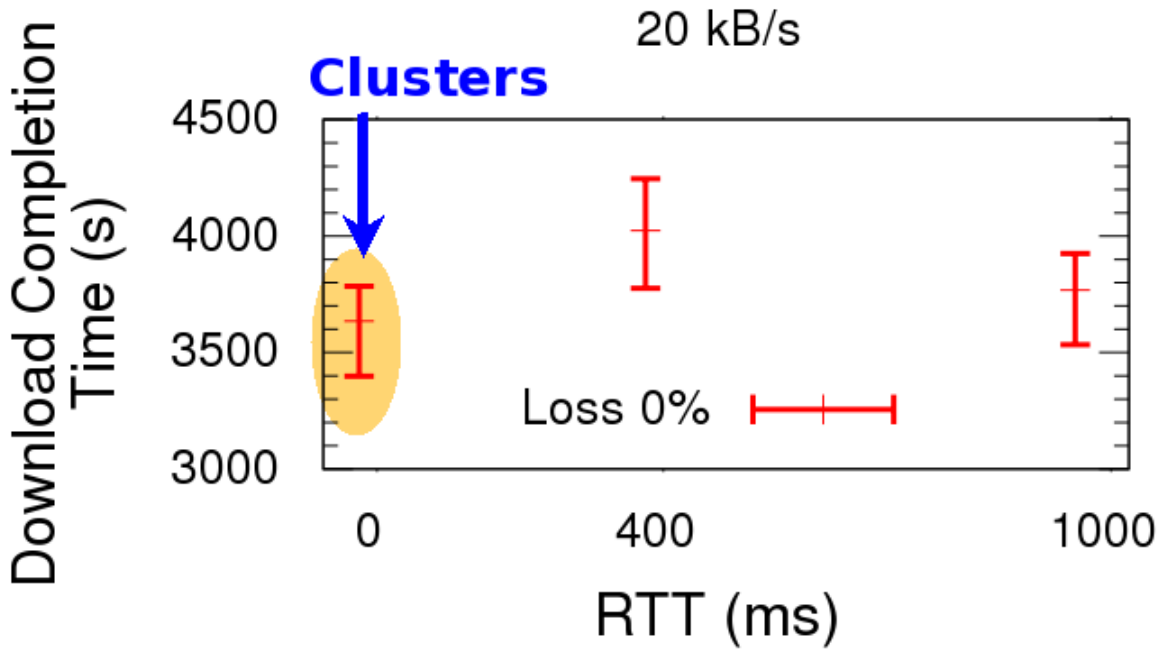

Impact of Packet Loss (20 kB/s)

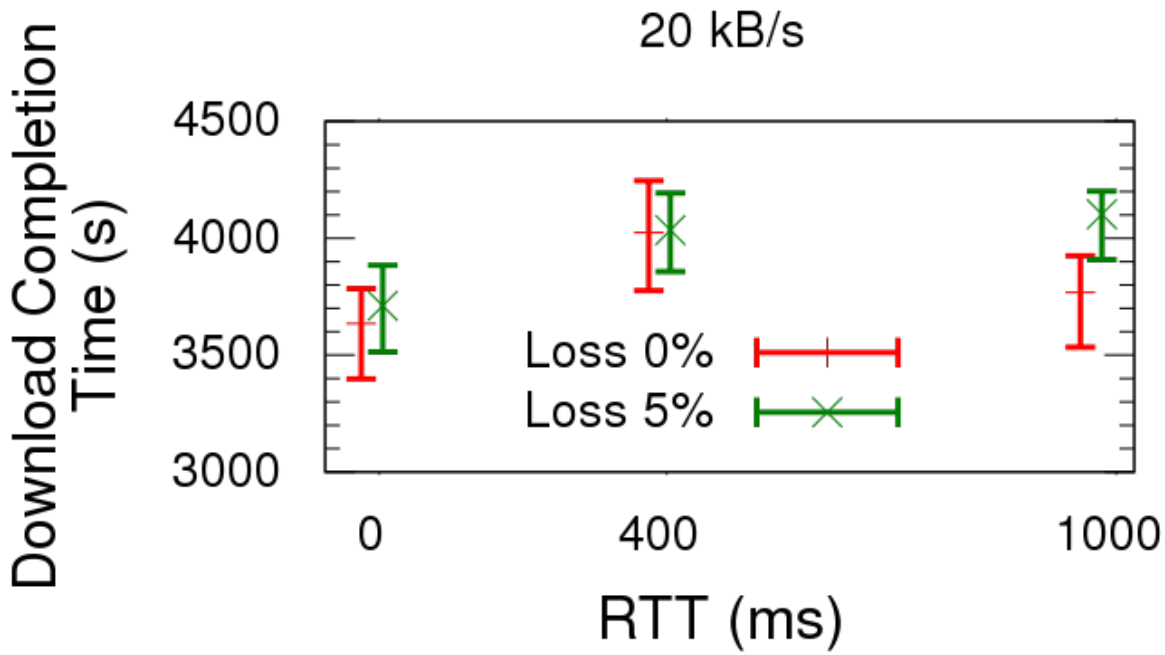

Impact of Packet Loss (20 kB/s)

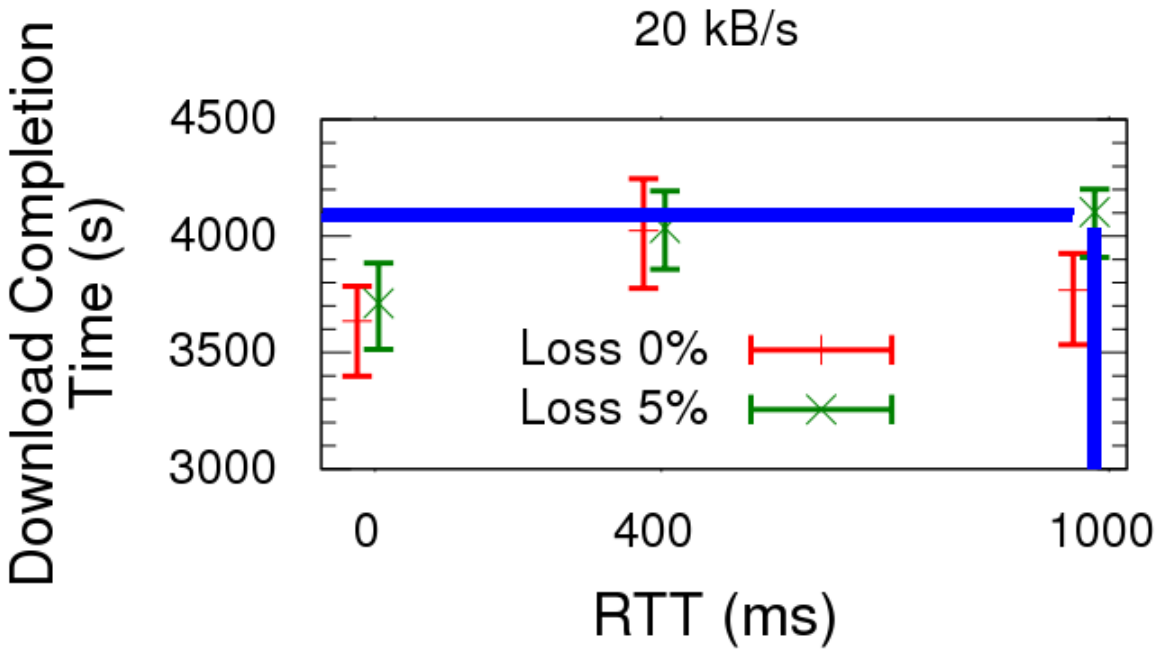

Impact of Packet Loss (20 kB/s)

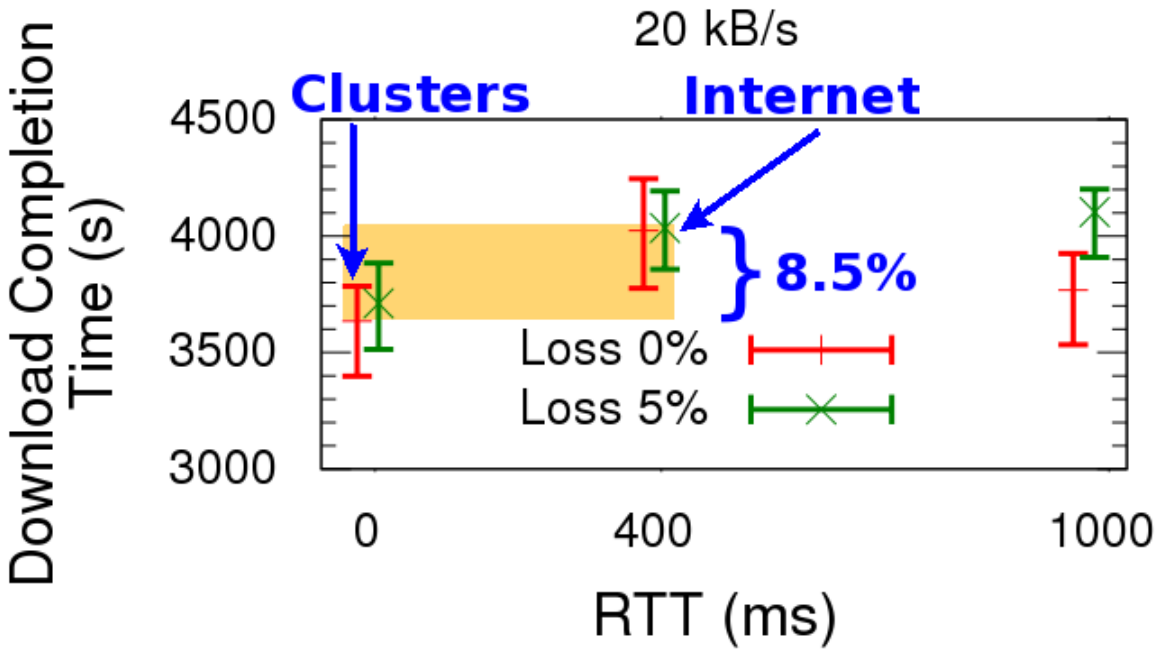

Impact of Packet Loss (20 kB/s)

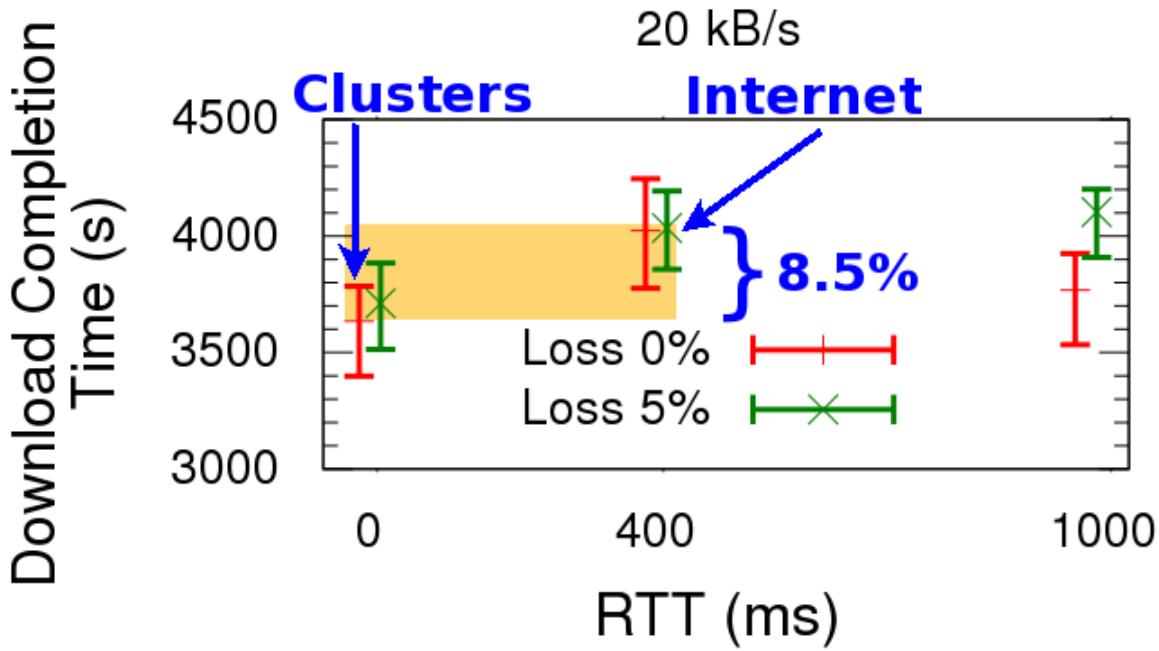

Marginal Impact of Packet Loss

Impact of Packet Loss (100 kB/s)

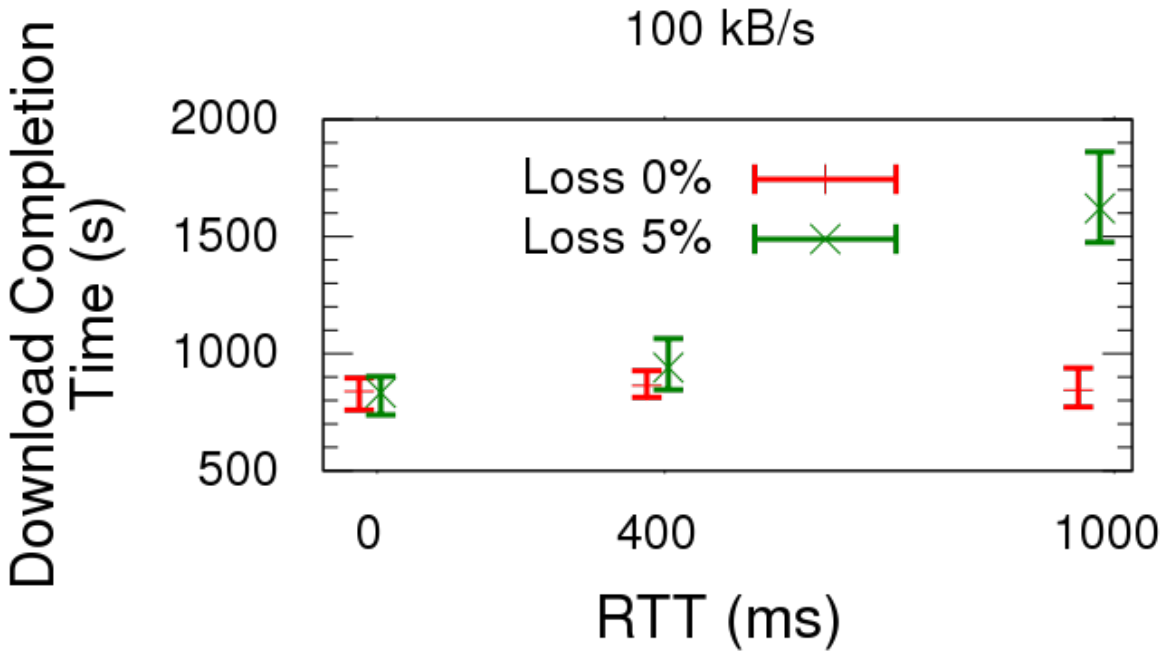

Impact of Packet Loss (100 kB/s)

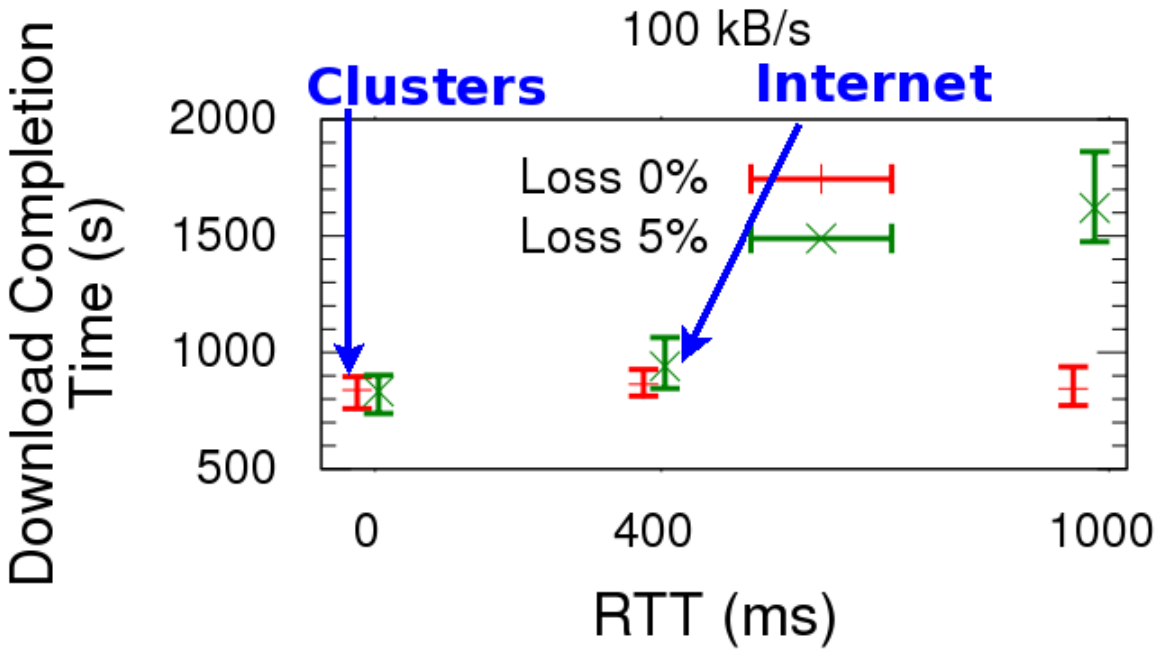

Impact of Packet Loss (100 kB/s)

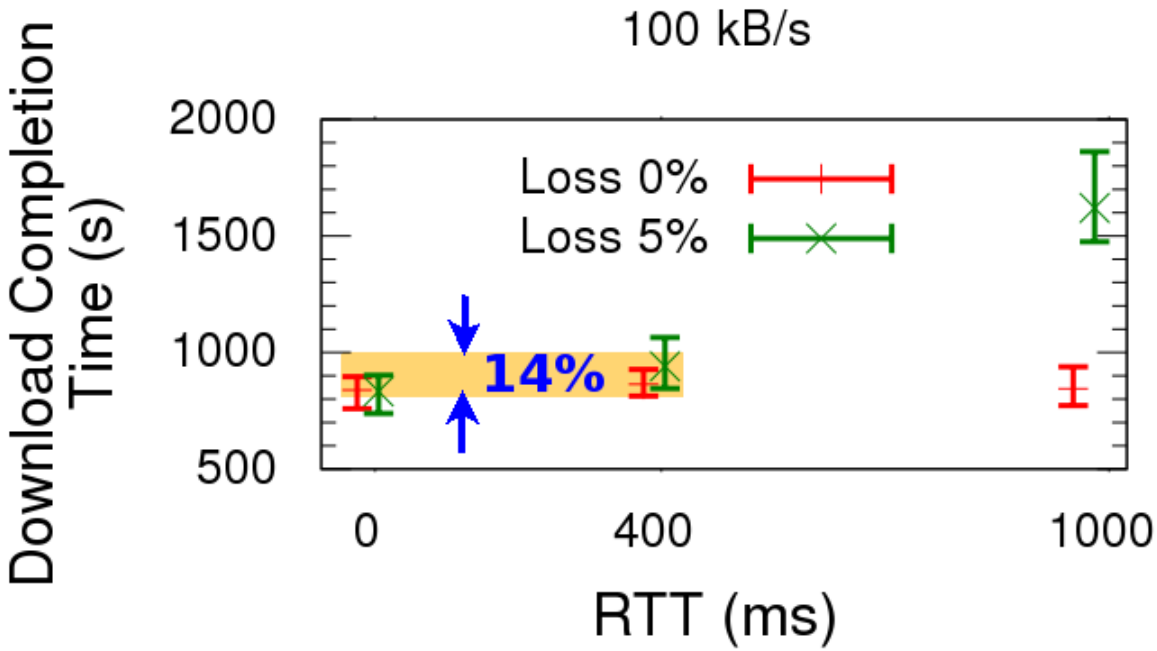

Impact of Packet Loss (100 kB/s)

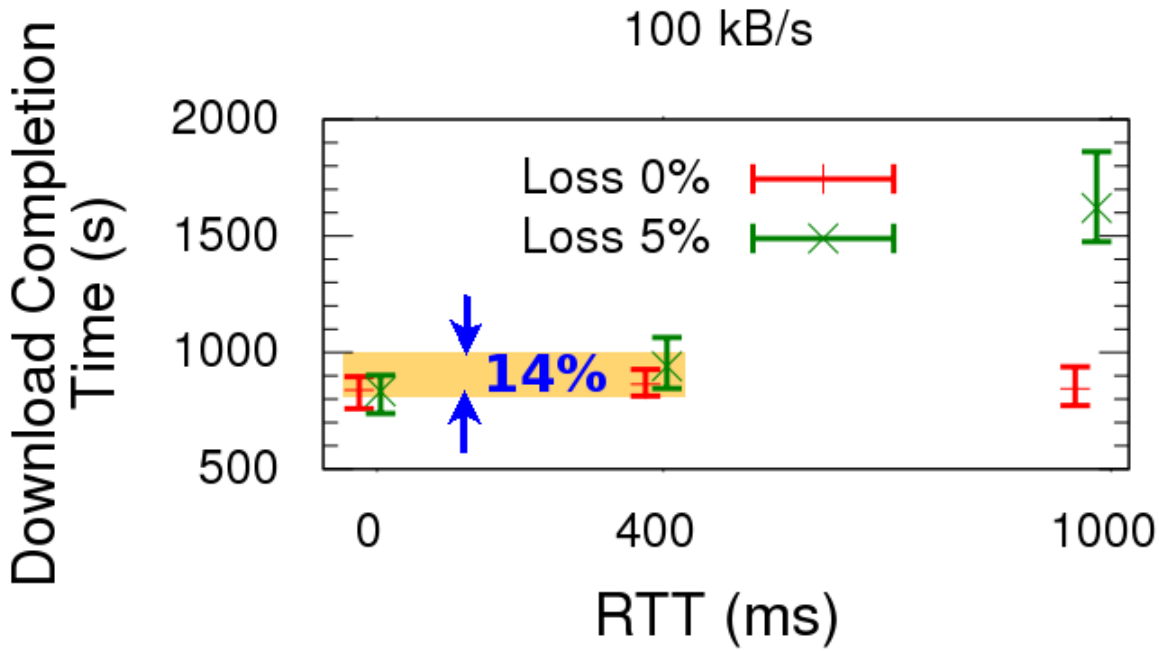

Marginal Impact of Packet Loss

Conclusions

- Marginal Impact of RTT
- Marginal Impact of Packet Loss
- Avoid Emulation of RTT and Packet Loss

Conclusions

- Marginal Impact of RTT
- Marginal Impact of Packet Loss
- Avoid Emulation of RTT and Packet Loss

Dedicated clusters can be used for BitTorrent Experiments

Can Realistic BitTorrent Experiments Be Performed on Clusters?

Questions

ashwin.rao@inria.fr

Can Realistic BitTorrent Experiments Be Performed on Clusters?

Thank you!

ashwin.rao@inria.fr

BACKUPS

- Variable Latency
- .torrent file
- TCP SQRT

Heterogeneous Latency - Setup

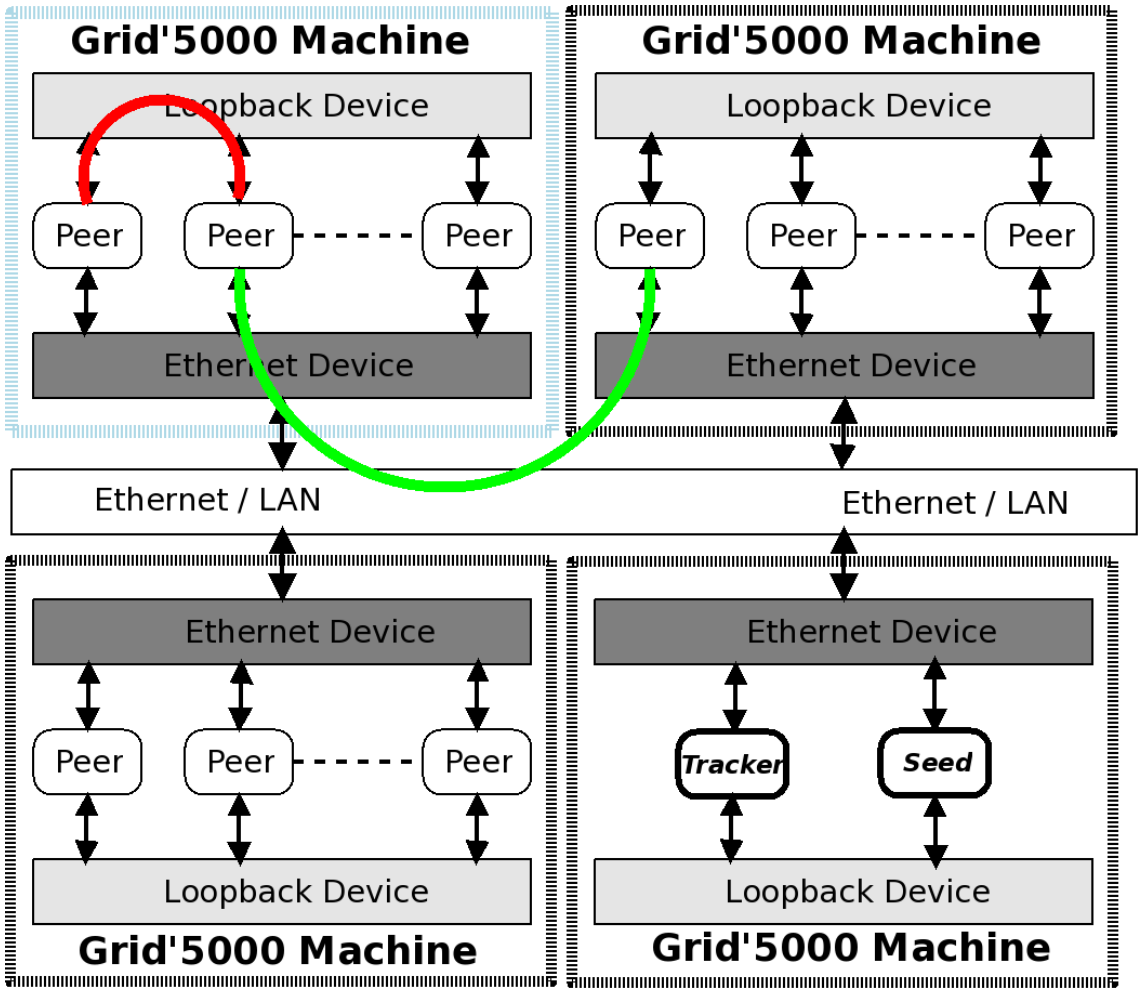

Different latency added on Loopback and Ethernet Devices

▸ BACKUP

Heterogeneous Latency - Setup

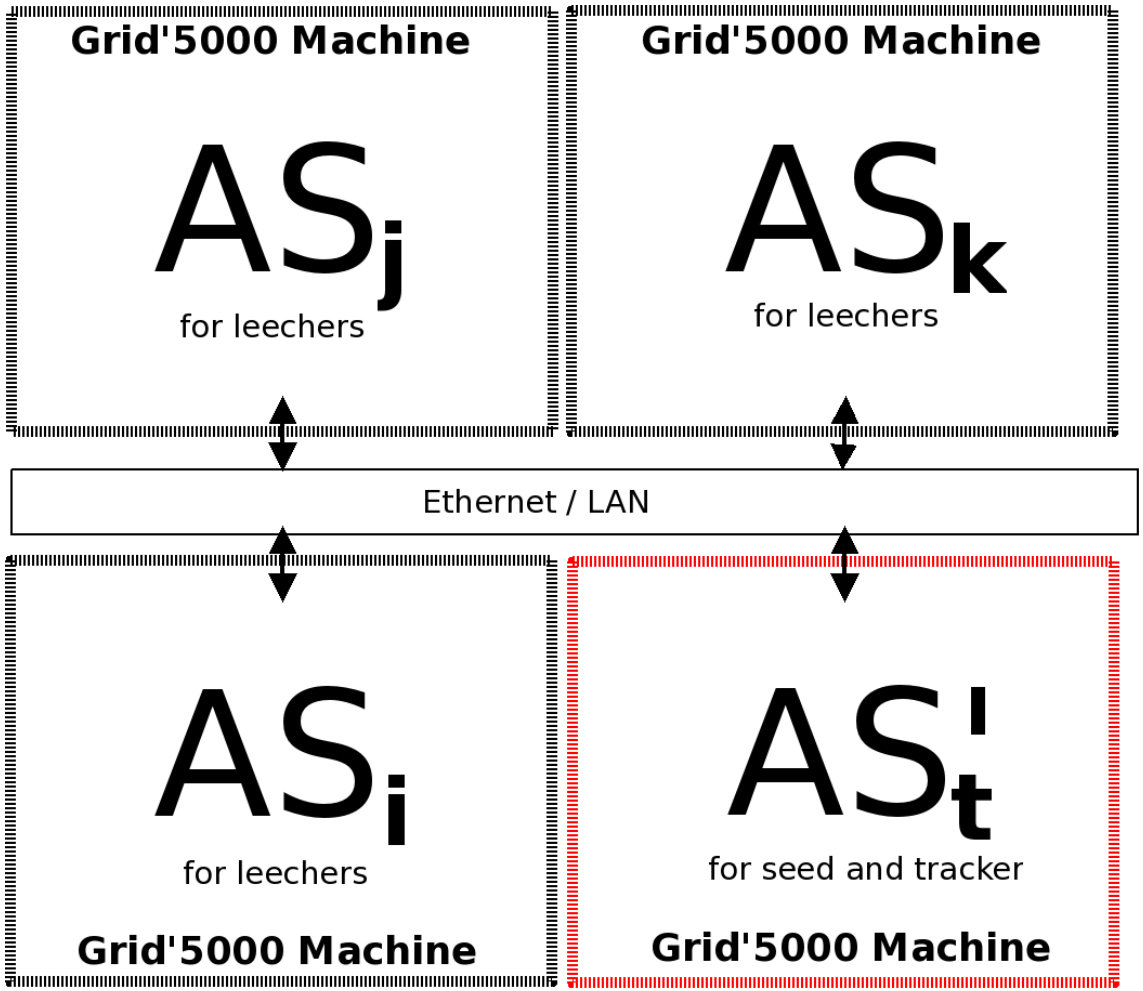

Grid'5000 Machine to Abstract an AS

▸ BACKUP

Results - Small RTT Between Peers

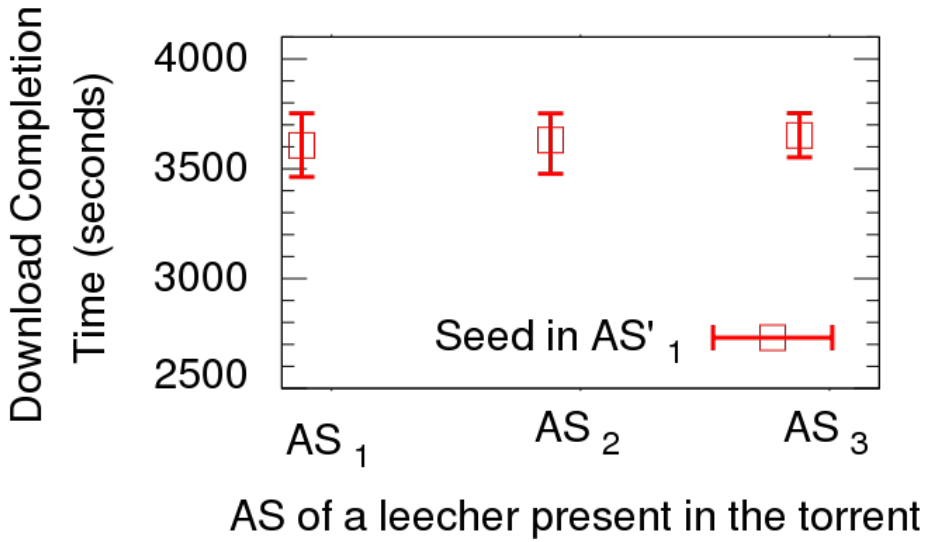

Upload Rate limited to 20 kB/s.

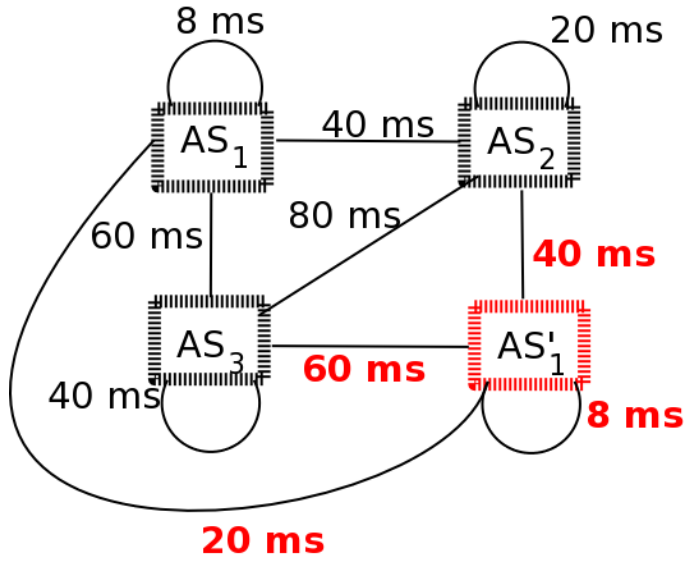

RTT between a pair of leechers.

RTT between any two peers is less than 100 ms

▸ BACKUP

Results - Small RTT Between Peers

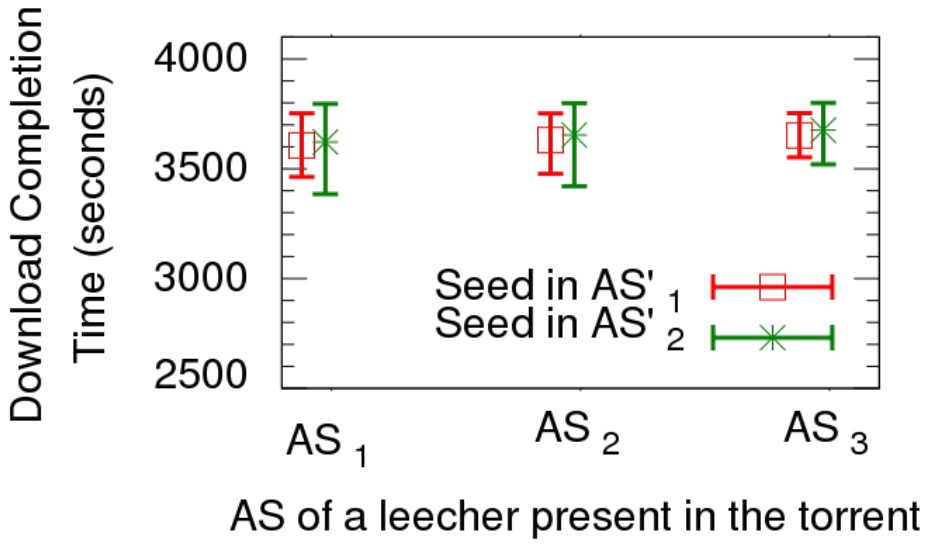

Upload Rate limited to 20 kB/s.

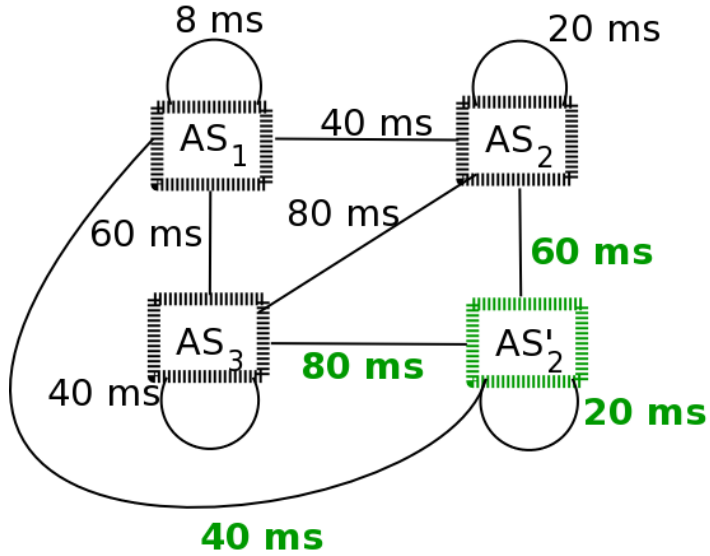

RTT between a pair of leechers.

RTT between any two peers is less than 100 ms

▸ BACKUP

Results - Small RTT Between Peers

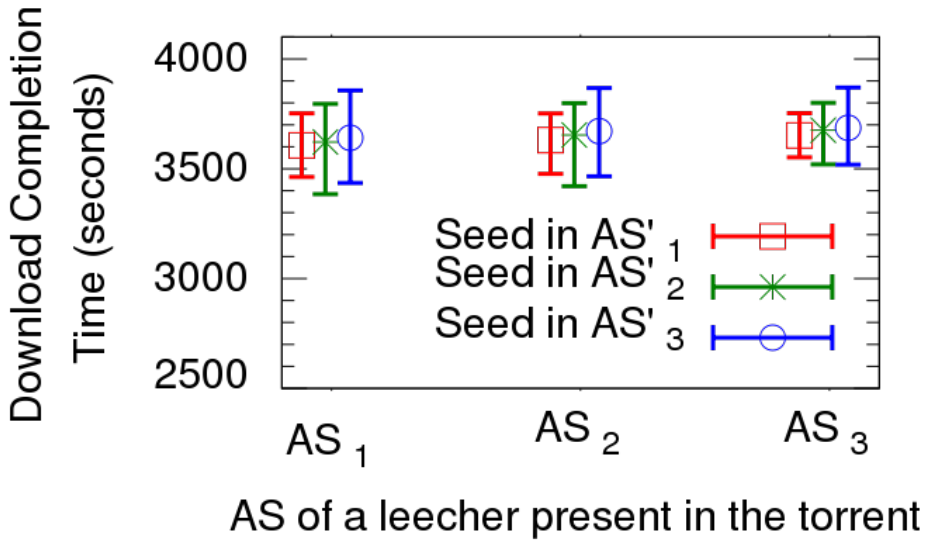

Upload Rate limited to 20 kB/s.

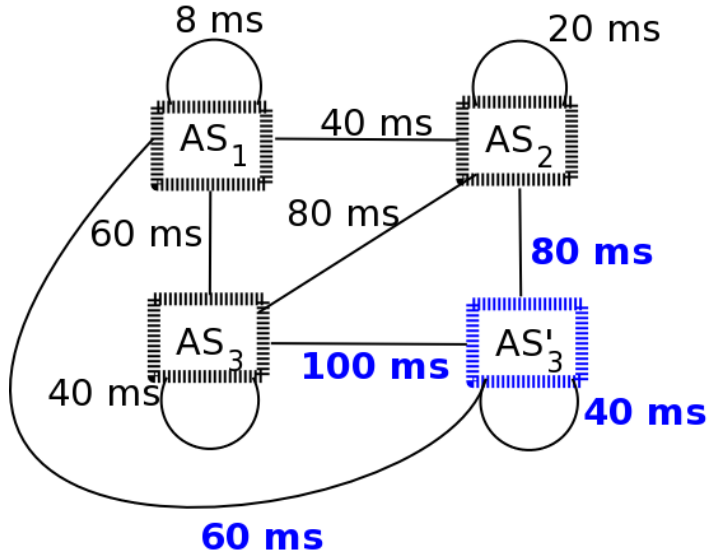

RTT between a pair of leechers.

RTT between any two peers is less than 100 ms

▸ BACKUP

Results - Large RTT Between Some Peers

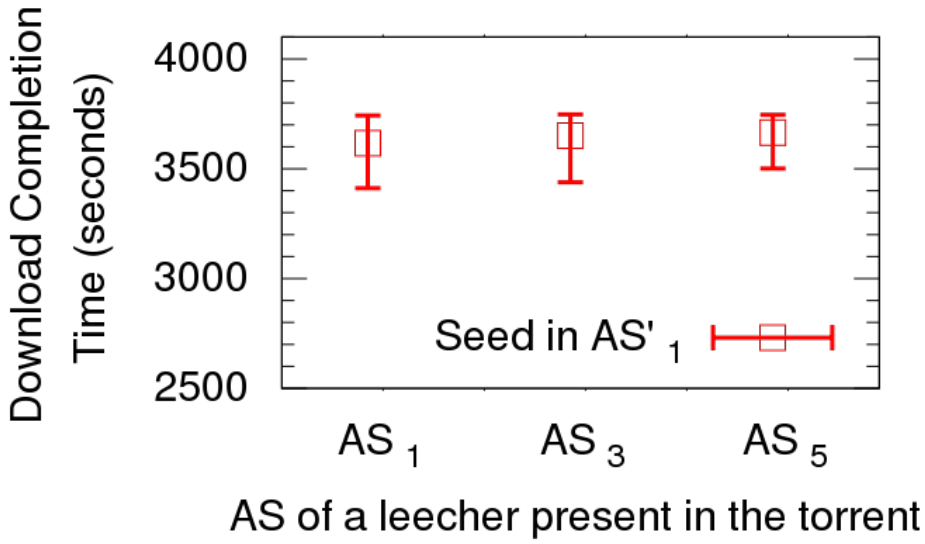

Upload Rate limited to 20 kB/s.

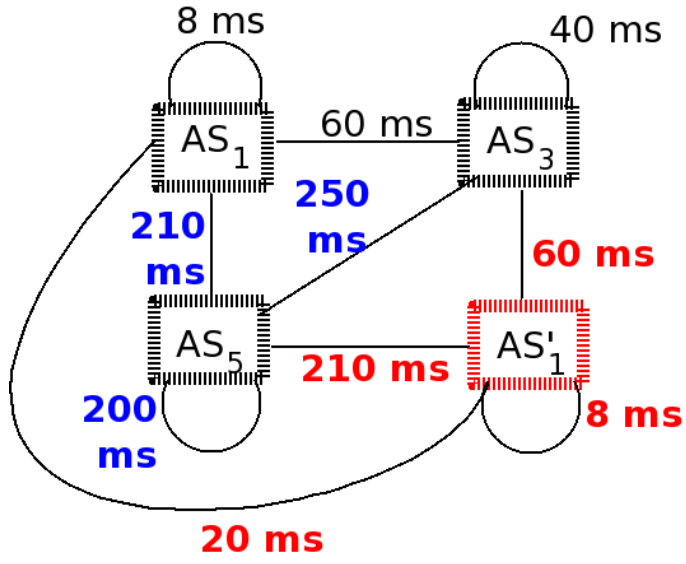

RTT between a pair of leechers.

RTT between some of the peers is greater than 200 ms

▸ BACKUP

Results - Large RTT Between Some Peers

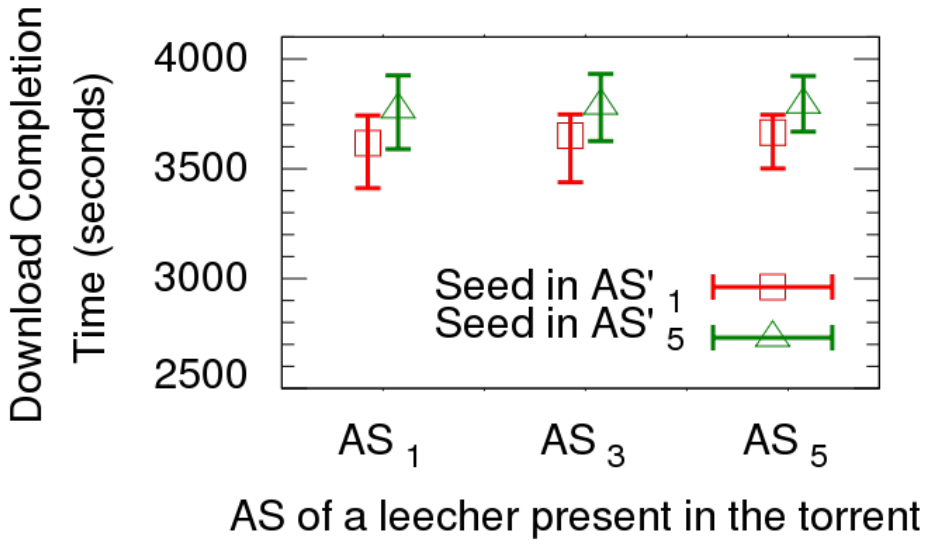

Upload Rate limited to 20 kB/s.

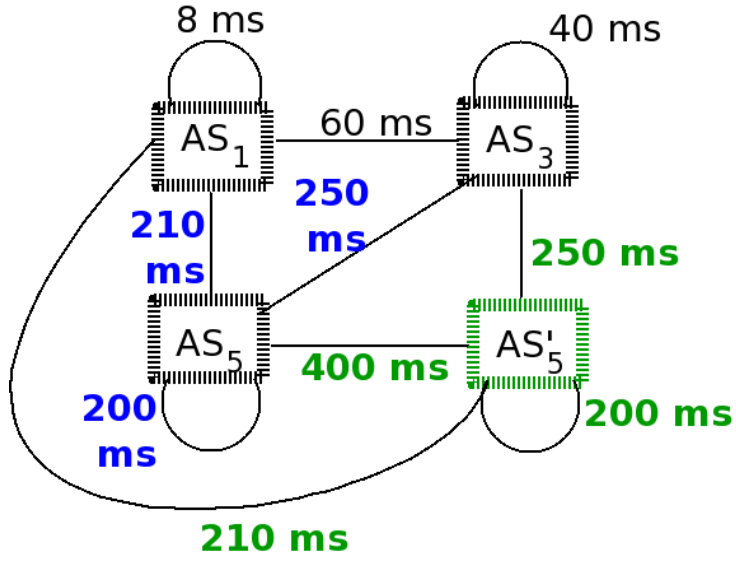

RTT between a pair of leechers.

RTT between some of the peers is greater than 200 ms

▸ BACKUP

.torrent File

info	dictionary describing files in torrent
name	File Name
length	File size in bytes
piece length	Size of a piece
pieces	20-byte hash values; one per piece
announce	The announce URL of tracker
...	

▸ BACKUP

TCP Square Root

$$Throughput = \frac{MSS}{RTT}\frac{C}{\sqrt{p}}$$

where

- MSS Maximum Segment Size
- RTT Round trip time
- p Loss rate
- C constant depending on TCP implementation

▸ BACKUP