

Can high-speed transport protocols be deployed on the Internet ? : Evaluation through experiments on JGNII

*Kazumi Kumazoe, Katsushi Kouyama,
Yoshiaki Hori, Masato Tsuru, Yuji Oie
Feb. 3 2006*

Outline

- Background
- Configuration of our setup
- Experimental Result
- Summary

Background

- Most high-speed transport protocols were originally proposed to meet the requirement of high throughput data transfer over fast long-distance network, such as Grid applications.
- The Internet
 - Bandwidth : increase (in both access and core networks, e.g. 1Gbps FTTH)
 - Users including ASPs : may be interested in transferring larger amount of data using high-speed transport protocols. (regardless of the intension of the original developers)
 - Little attention has been paid to the problems when those transport protocols are deployed on shared and heterogeneous network such as the Internet
- What will happen if high-speed transport protocols are running in the Internet ?

objective

■ Basic Characteristics

- A single flow
- various receiver side OS

■ How rapid changes in network conditions affected the performance of high-speed transport protocol?

- delay change
- bandwidth

■ How the high-speed transport protocol flows affected the performance of coexisting multiple heterogeneous flows and vice versa?

- flows coexisting different RTT values
- heterogeneous high-speed flows
- short-lived TCP flows
- UDP CBR flows

- Background
- Configuration of our experiment
- Experimental Result
- Summary

■ Target Protocols

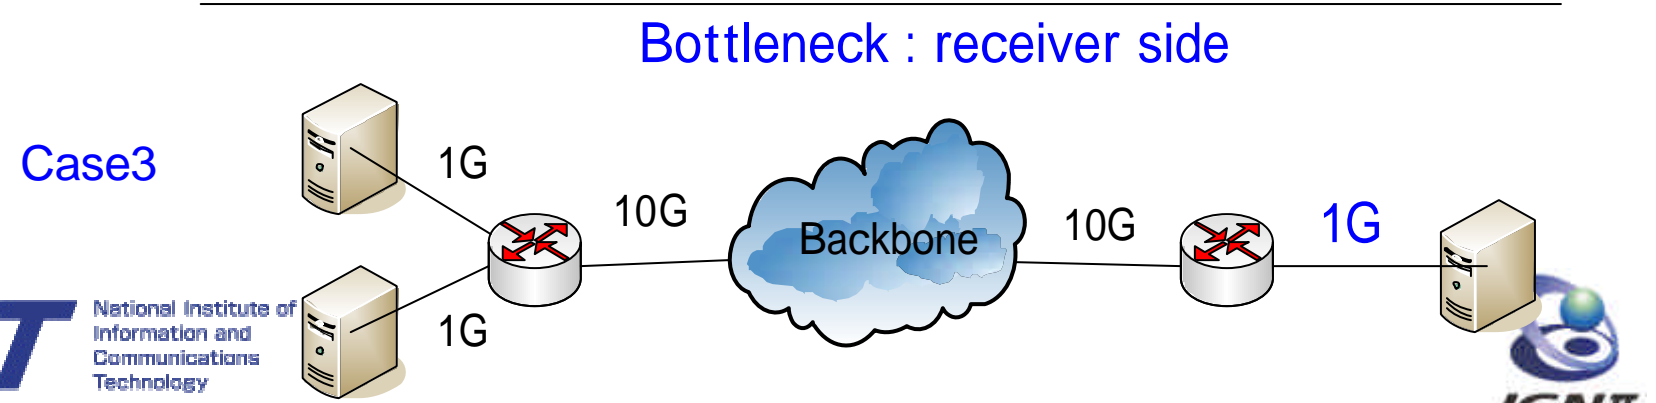
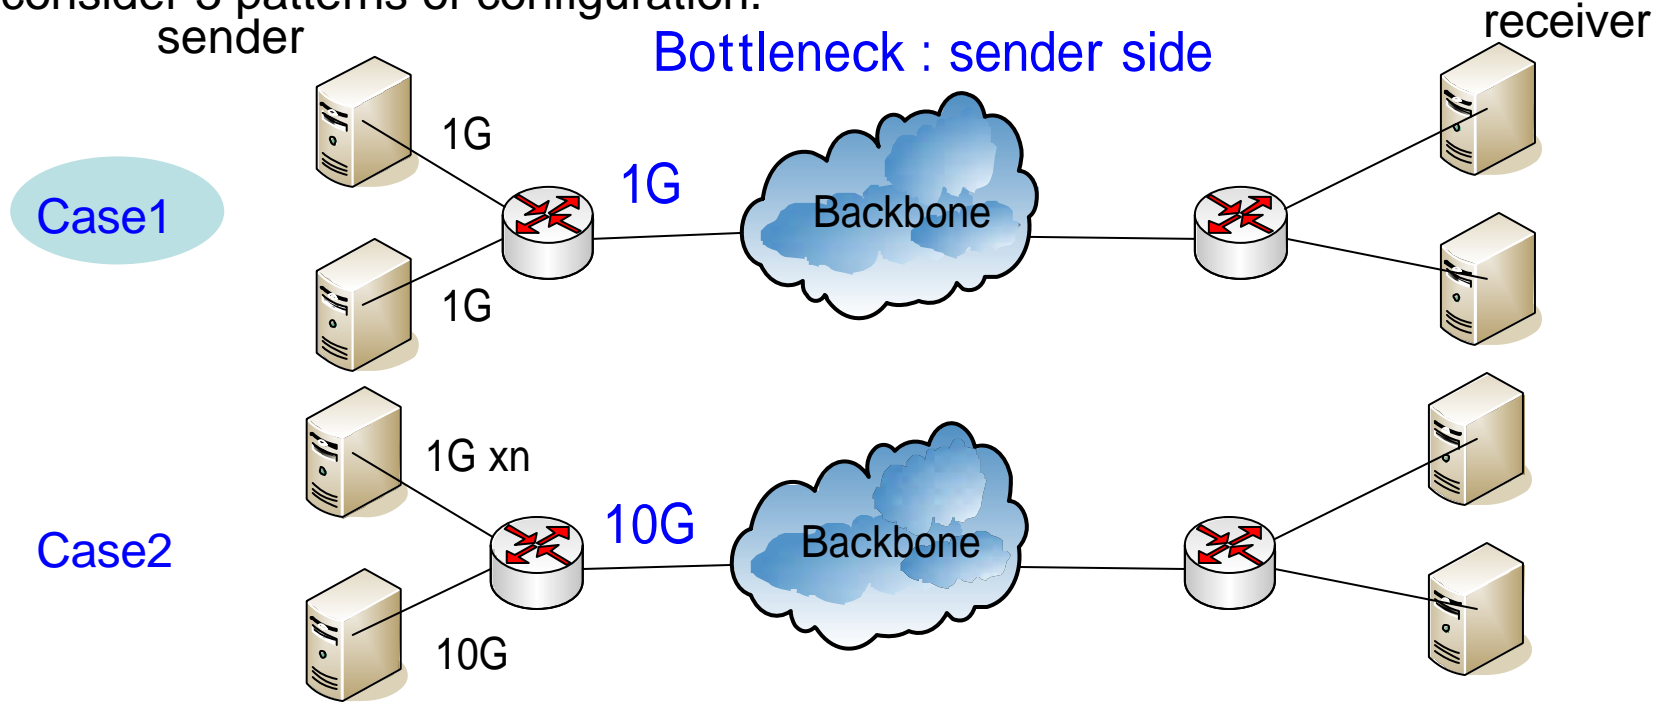
Protocol	Version
Standard TCP	Linux 2.4.20(NewReno)
HSTCP	patch for Linux 2.4.20
Scalable TCP	patch for Linux 2.4.20
FAST	patch for Linux 2.4.22
CUBIC	patch for Linux 2.4.25
HTCP	patch for Linux 2.4.23
UDT	version 2.0 on Linux 2.4.20

■ Performance measures

□ Throughput Characteristics

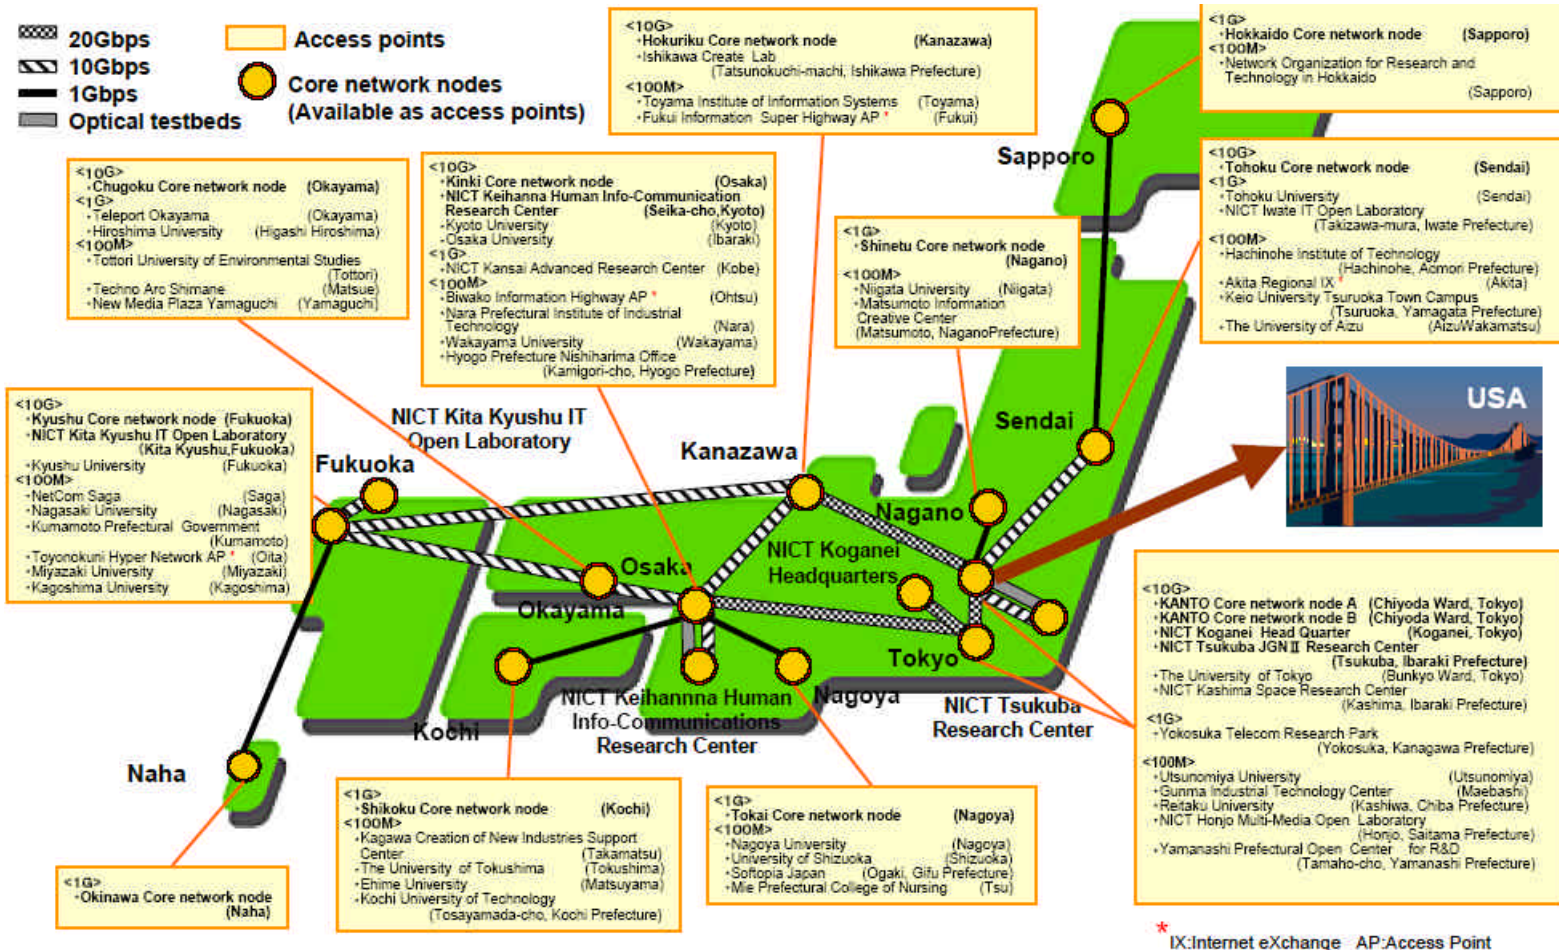
Location of Bottleneck for TCP flows in the Internet

- Bottleneck for TCP in the Internet locate at sender side and receiver side
- We consider 3 patterns of configuration.




Experimental Setup(1)

■ Backbone Network : JGN(Japan Gigabit Network) II



Experimental Setup(2)

- traffic on *JGN(Japan Gigaabit Network) II*
`Weekly' Graph (30 Minute Average)



- We will show the results of Bottleneck is 1[Gbps].
- When not congested, no losses observed in 904[Mbps] UDP flow by Iperf test.


■ Path Characteristics

	RTT[ms]	Max. Bandwidth[Gbps]
Network Emulator	0.1 - 10000	1
JGNII Domestic Line	38	10
JGNII International Line	180	10

Experimental Setup(3)

■ JGNII domestic line, Network Emulator

- JGNII domestic Line
- Kitakyushu –Tokyo – Kitakyushu (RTT=38[ms])
- Bottleneck Bandwidth (\approx 1Gbps)




- Network Emulator (Packet Storm)
 - Configurable : RTT(0.1-10000[ms])
- Bottleneck Bandwidth (\sim 1Gbps)

Experimental Setup(4)

■ JGNII International Line

- Kitakyushu – US (RTT=180 [ms])
- Bottleneck Bandwidth (\approx 1Gbps)




Outline

- Background
- Configuration of our setup
- Experimental Result
 - Basic Characteristics
 - single flow
 - receiver side OS
 - **Impact on the co-existing multiple flows**
 - heterogeneous high-speed flows
 - short-lived TCP flows,
 - UDP CBR flows
- Summary


Throughput of a single flow (International line)

- JGNII International Line (RTT = 180[ms]), 100[s]
- Output Buffer size at edge router = (63[packet], 255[packet]

BUF= 63[packet]



BUF= 255[packet]



Scalable —


FAST —

CUBIC —

HTCP —

Change Receiver side OS (Emulator) ¹⁴

- Network Emulator (RTT=180[ms]) ,
- Edge Router Buffer Size = 63[packet], duration =100[s]
- Receiver side OS (FreeBSD 5.3 and XP)
- Parameter tuning (expand socket buffer size, set RFC1323 etc.)



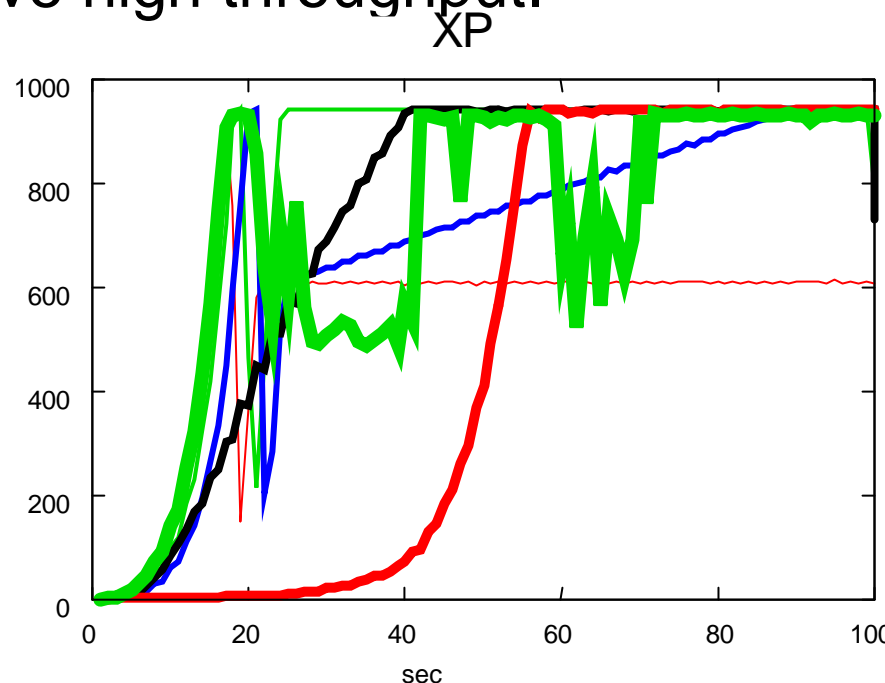
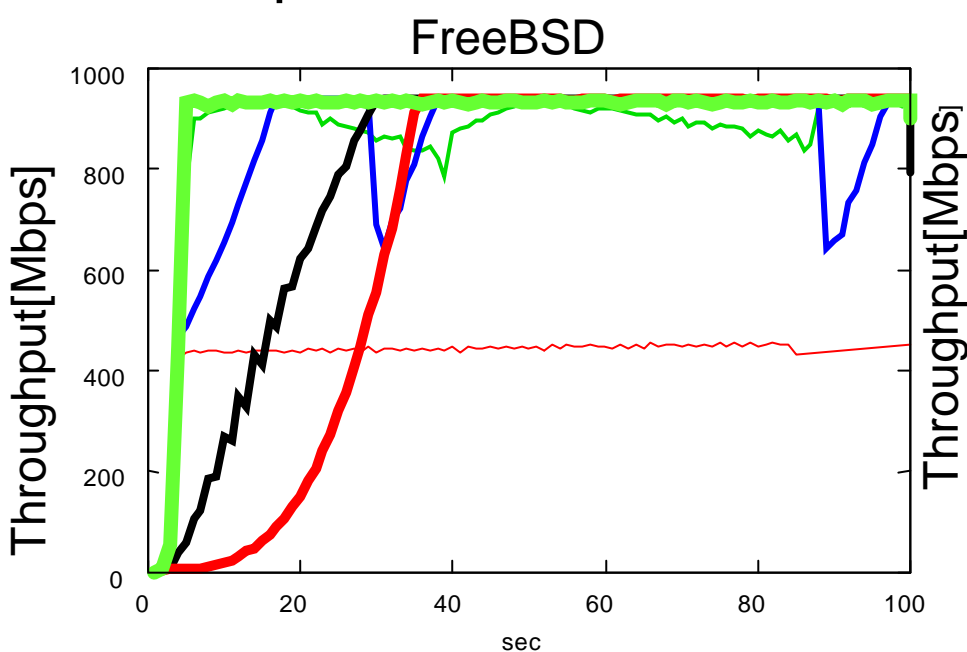
Linux

FreeBSD5.3

XP

Change Receiver side OS (Emulator) ¹⁵

- Network Emulator (RTT=180[ms]),
- Edge router buffer size = 63
- FreeBSD is similar in their behavior to Linux flows
- XP receiver increased more slowly.
- All protocol flows could achieve high throughput.



Standard — Scalable — CUBIC —
HSTCP — FAST — HTCP —


Outline

- Background
- Configuration of our setup
- Experimental Result
 - Basic Characteristics
 - single flow
 - receiver side OS
 - **Impact on the co-existing multiple flows**
 - heterogeneous high-speed flows
 - short-lived TCP flows,
 - UDP CBR flows
- Summary


Coexistence of long-lived TCP flows (2)

- JGNII international line(RTT=180[ms]), Edge Router Buffer Size = 63, 300[s]
- Two different high-speed flows share a high-speed and long distance path.
- link utilization degrades when different kinds of high-speed transport protocol flows coexist
- Unfairness between heterogeneous protocol flows

(1) HTCP



(1) CUBIC

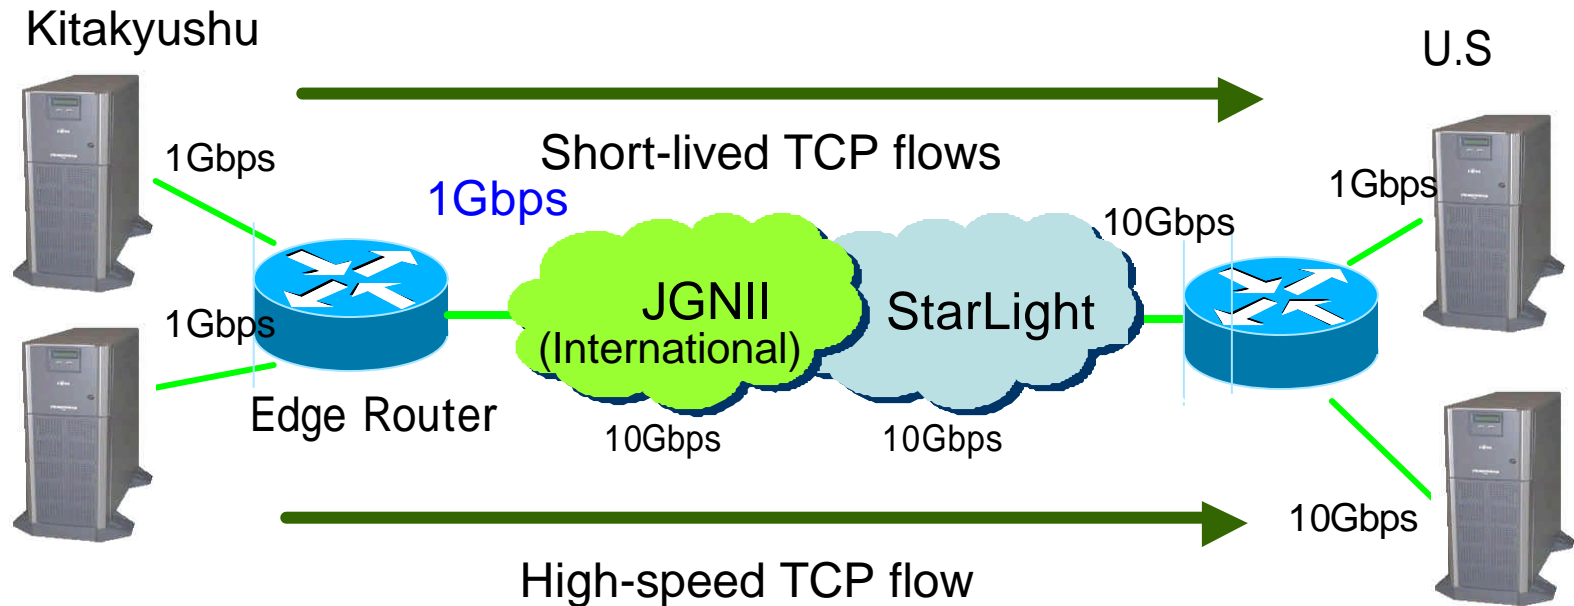


Outline

- Background
- Configuration of our setup
- Experimental Result
 - Basic Characteristics
 - single flow
 - receiver side OS
 - **Impact on the co-existing multiple flows**
 - heterogeneous high-speed flows
 - short-lived TCP flows,
 - UDP CBR flows
- Summary

Coexistence of Short-Lived TCP Flows(1)

- a single high-speed flow + short-lived TCP flows
- JGNII international line(RTT=180[ms]), 300[s]
- Edge Router Buffer Size = 63[packet]
- Short-lived TCP flow (web traffic)
 - 3000 files / 300 seconds
 - file size : Pareto distribution with the mean of 100,300,500[Kbyte]
($\alpha = 1.3$)



Coexisting of short-lived TCP flows(2)


■ throughput of short-lived TCP flows

□ Performance measure

averaged throughput of short-lived TCP = $\sum_i S_i / \sum_i T_i$


S_i : file size of flow i , T_i : transfer time of flow i

- short-lived TCP flow slightly affected by coexisting high-speed transport protocol flow (the larger files were transferred, the more degradation were experienced.)




Coexisting of short-lived TCP flows(3)

- throughput of high-speed transport protocol
 - averaged throughput over 300[s]
 - high-speed transport protocol flow was strongly affected by coexisting short-lived TCP flows. (except for UDT.)
 - influence on TCP-based Transport protocols
average file size 100 KB < average file size 500KB



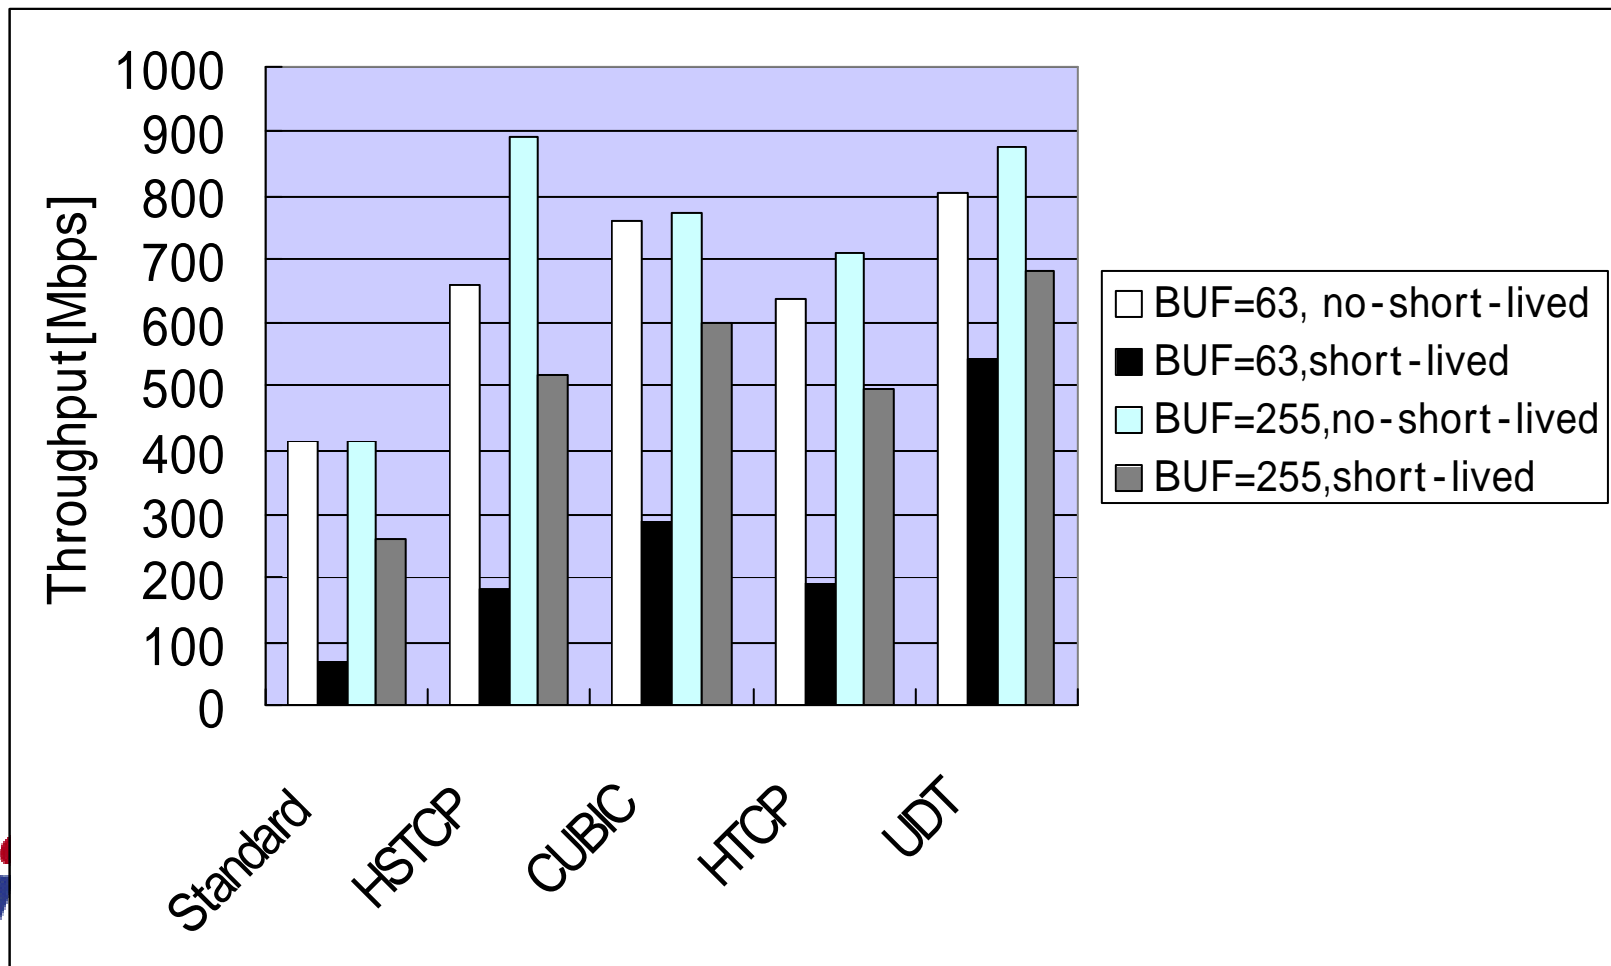
Coexisting of short-lived TCP flows(4)

- Throughput of high-speed transport protocol
- coexisting short-lived TCP flows (average file size : 500[KB]) during 50[s] and 350[s]
- high-speed transport protocol flow is strongly affected by coexisting short-lived TCP flows



Coexistence of short-lived TCP flows(5)

- Throughput of high-speed transport protocol
- buffer size at edge router : 63[packet] and 255[packet]
- averaged file size : 500[KB]
- throughput : (edge router : 255[packet]) > (edge router 63[packet])



Outline


- Background
- Configuration of our setup
- Experimental Result
 - Basic Characteristics
 - single flow
 - receiver side OS
 - **Impact on the co-existing multiple flows**
 - heterogeneous high-speed flows
 - short-lived TCP flows,
 - UDP CBR flows
- Summary

Coexistence of CBR UDP flows(1)

- a single high-speed flow + 8[Mbps] UDP x 2
- JGNII international line (RTT=180[ms])
- UDP flow : Jitter, Packet Loss Rate, TCP flow : throughput

Case1

300[s]



Case2

300[s]





Jitter of UDP flows

- variation of received packet interval measured by Iperf (RFC 1889)
- UDP flows only < coexisting high-speed transport protocol flows (slightly affected)


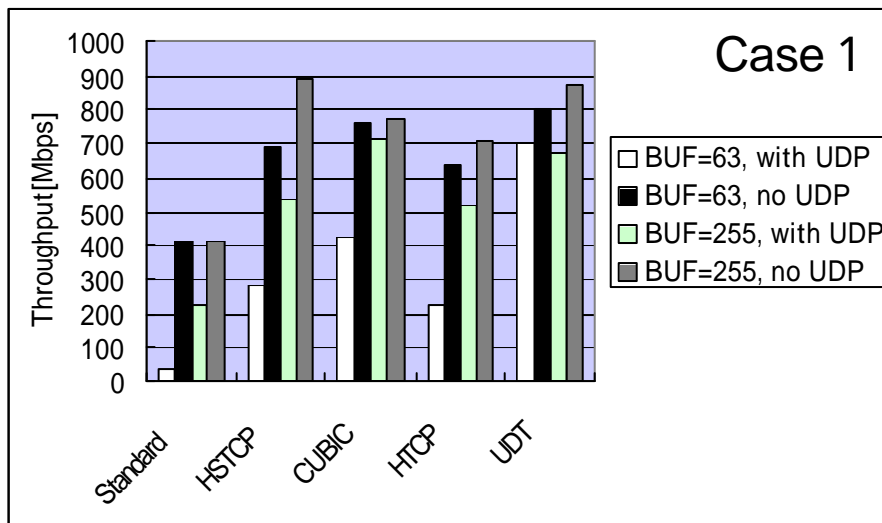
degree of degradation : depends on type of protocol and buffer size at router

Packet Loss Rate of UDP flows

- only UDP flows and coexisting Standard TCP flow : 0 %
- other case : depend on the kinds of protocol < 0.02 %

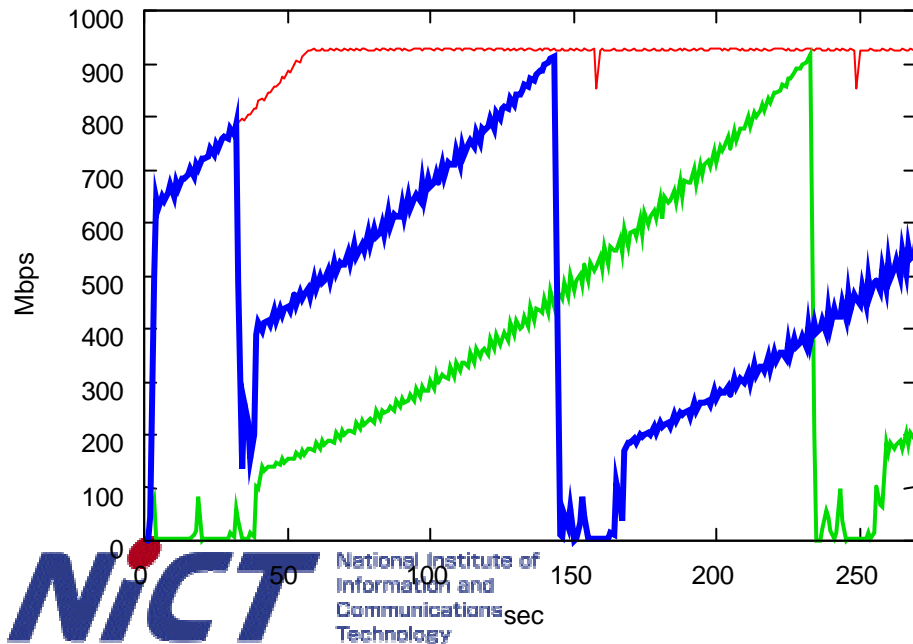


- Throughput of high-speed
- buffer size of the edge router : 63, 255[packet]
- In both case1 and case2, high-speed flow are greatly affected by coexisting UDP flows
- The averaged throughput (over 300[s]) of high-speed flow is almost same.
- The performance of Standard TCP observed in case2 is not affected by UDP flows.




- high-speed flow + two 8[Mbps] UDP flows
- compare the throughput characteristics
 - Only high-speed transport protocol flow
 - case1 (UDP flows start first)
 - case2 (high-speed TCP flow start first)
- case1 : the starting phase of the high-speed flow is strongly affected by coexisting UDP flows

HSTCP, BUF=63[packet]



HSTCP, BUF=255[packet]



Outline

- Background
- Configuration of our setup
- Experimental Result
 - Basic Characteristics
 - single flow
 - receiver side OS
 - **Impact on the co-existing multiple flows**
 - heterogeneous high-speed flows
 - short-lived TCP flows,
 - UDP CBR flows
- Summary

Summary

- We have been investigating what happens if high-speed transport protocols are used in the global Internet through experiments in the JGNII on various scenarios.
 - Throughput Behavior depends on receiver side OS, although high throughput can be achieved.
 - link utilization degrades when different kinds of high-speed transport protocol flows coexist.
 - High-speed protocols are considerably affected by coexisting short-live TCP/ small amount of UDP flows and vice versa.

- We think an improved high-speed transport protocol and/or new management mechanisms in the intermediate nodes are needed to achieve efficient and moderate bandwidth sharing between coexisting various flows in the Internet.