Annexure A

Data Distribution over Multipath

A.1 INTRODUCTION

Now a days data centers, smart phones and Internet devices are becoming multi-homed. To effectively utilize multiple links between two communicating end points, it is needed to use all the available interfaces for data transmission. MultiPath Transmission Control Protocol (MPTCP) is an effort that enables end points to distribute data over all the available interfaces simultaneously. However, the mechanism proposed for distributing data over available multiple links can not fully utilize the network resources because it does not consider the end-to-end delay and congestion state of subflows. As the conditions in each path are different, the data transmitted by different subflows arrives out of order at the receiver. The problem is more severe in an environment where the characteristics of each available link are drastically different. This causes the receiver to send duplicate acknowledgements which may be wrongly inferred as packet loss by the sender. This degrades the performance of MPTCP significantly because re-ordering is interpreted as a sign of congestion. This Annexure proposes a scheduling algorithm to reduce the packet re-ordering by predicting the data segments to be transmitted by each of the subflows of MPTCP. The proposed algorithm works based on the RTT and congestion state of subflows. The proposed scheduling algorithm tries to cope up with the asymmetry of the links.

A.2 MPTCP

Multipath TCP, as proposed by the Internet Engineering Task Force (IETF) [Raiciu et al., 2011], allows a single data stream to split across multiple paths. This has an obvious benefit for reliability, and it can also lead to more efficient use of network resources. Since 1980's TCP/IP communication is using single path per TCP connection, while today most of the networking devices are equipped with multiple interfaces. TCP by itself is not capable of efficiently and transparently using the interfaces available on a multi-homed host. Therefore, in order to fully utilize all the network resources it is needed to move from single path protocols to multipath protocols.

In MPTCP, available resources are pooled in such a way that they appear as a single connection to application users. Multipath transmission capability uses a resource pooling proposed by [Wischik et al., 2008] to increase bandwidth, by simultaneously making use of multiple disjoint (or partially disjoint) paths across a network. The initial goal of multipath TCP is to use multiple paths between the end hosts when one or both are multi-homed. MPTCP is more attractive because it is backward compatible with existing TCP applications. Each subflow of an MPTCP connection is treated as a single TCP connection. In MPTCP, every subflow has its own sequence number and congestion window.

A.3 MOTIVATION

The sender generates a data stream with an in-order sequence of data packets. Due to dynamics of the network, the sequence of data packets may change. Then when an out of order packet is received, the receiver responds with duplicated acknowledgements (DupACK) indicating

the sender to infer wrongly a packet loss and enter into congestion control stage unnecessarily. This results in lower overall end-to-end performance. In multipath TCP, segment re-ordering would happen if paths have different characteristics such as end-to-end delay, congestion state, and data rates etc. as described in [Barré et al., 2011]. Figure A.1 shows the impact of re-ordering at the receiver side. There are two paths, subflow 1 is low delay path (fast path) and subflow 2 is on high delay path (slow path). Segments are numbered as 1, 2, 3....etc. and transmitted on both subflows. Segments 1 and 2 have been transmitted by fast path and received in-order and directly went to the receive buffer at the receiver side. However the segment 3 is scheduled on slow path and it is still in-transit before segments 4, 5, 6, and 7 have been reached to the receiver side by the fast path and waiting for segment numbered 3 into out of order buffer. Segments waiting in out of order buffer will go in receive buffer when segment number 3 is received by the receiver.

In general, when a data segment arrives out of order at the receiver, it goes into out of order buffer and waits until receiving a segment that resolves the order. If a packet arrives in-order then it will go directly into receiver buffer from where data is transmitted to the application layer.

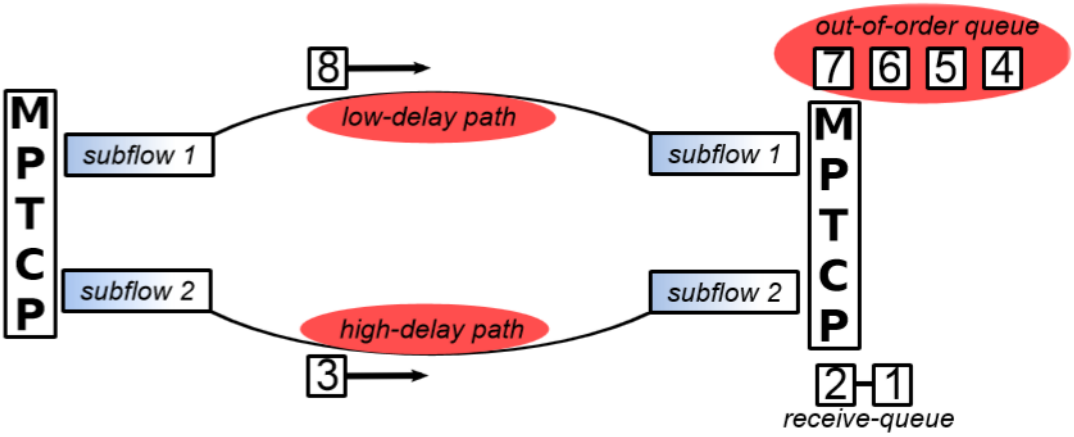

Figure A.1. : Re-ordering problem

A.3.1 Problem with Current Scheduler

Current multipath TCP scheduler takes the next available segment from the shared send buffer, whenever it is called. To illustrate the problem, consider the following scenario which is also illustrated by Figure A.2. Suppose two subflows are used, subflow 1 has an estimated RTT of 100ms, while subflow 2 has an estimated RTT of 10ms and both have a current congestion window of 3000 bytes. Subflow 1 requests the scheduler for new data, because its congestion window has been fully acknowledged. On the other hand, the faster subflow 2 is not available currently. In such a situation, current scheduler will allocate segments 3 and 4 (assuming an MSS of 1460 bytes) to subflow 1, and will receive the corresponding acknowledgements 100ms later. This is clearly sub-optimal, as by waiting a maximum of 10ms, subflow 2 would have been able to transmit, allowing the data to be acknowledged within 20ms instead of 100ms. Continuing this reasoning, one can observe that as many as 18 segments (with a MSS of 1460 bytes) would have been be acknowledged faster if they had been sent over subflow 2.

Figures A.1 and A.2 describe the problems in concurrent transmission due to an asymmetric link between two ends. Figure A.1 describes the problem at the receiver side, due which the overall application latency increases. Figure A.2 describes the problem of blocking of fast path due to slow path at sender side when bottommost segment (i.e., next available segment in send buffer) is scheduled and links are asymmetric in nature. The problem with the data scheduler of MPTCP is

how to distribute data among subflows with different characteristics?

The solution of the above stated problem is to choose an appropriate segment for subflow in order to minimize the re-ordering at receiver side. The new scheduler would then take directly segments 19, 20 and feed them to subflow 1. Doing so will reduce the connection level re-ordering at the receiver and avoid the transmission of DupACKs by the receiver and thus reducing the total latency of application.

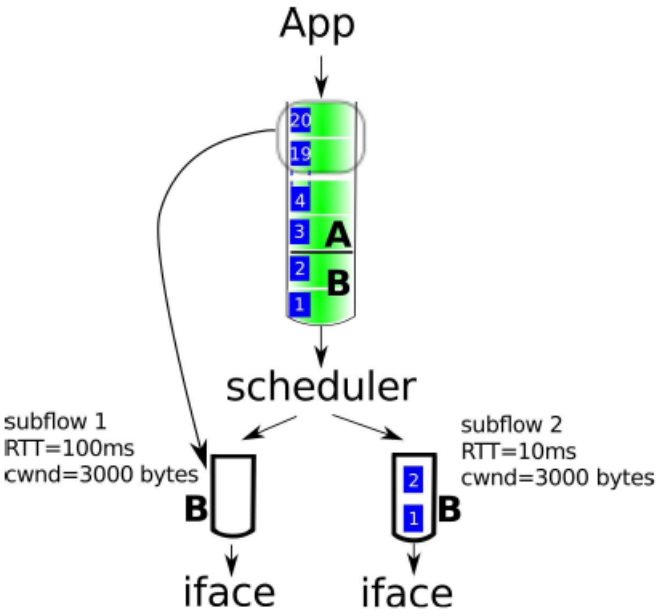

Figure A.2. : Proposed algorithm for subflow segment estimation

A.4 PROPOSED ALGORITHM

The proposed algorithm assumes that for all subflows the size of the segment is same. In this algorithm each chunk of size 1400 bytes in the sender's send buffer is assigned a number called segment number. At the receiver side, there are two buffers, one is receiver buffer (which receives the data segments that are in-order and delivers them to application layer) and other is out of order buffer (which holds the data segments that arrive out of order). The proposed algorithm tries to predict the segment number for a given subflow based on the RTT and congestion state information by assuming that the remaining segments with lower number will be sent by other subflows which have lower RTT. The algorithm predicts the segment number by integrating the round trip time and the congestion state of the subflow. Each subflow has a unique identifier, subflow id. A segment number is assigned to a chunk of bytes coming from application layer to the send buffer at transport layer. The proposed algorithm further assumes that the lost segments are re-transmitted on the same subflow.

The following notations are used in the proposed algorithm to estimate the segment number for a subflow c with round trip time rtt_c and congestion window $cwnd_c$, This approach considers all the subflows i with round trip time rtt_i and congestion window $cwnd_i$ that satisfy the following conditions.

$$\forall \text{ subflow}_\text{i} \in \text{subflow}_\text{list}$$

Algorithm 5 Segment Estimation Algorithm for Subflow c

```
Estimate Seg Number(subflow id)
rtt_c ← subflow_c(RTT) and
cwnd_c ← subflow_c(cwnd)
segment#_estimated ← next segment in send buffer
if (rtt_c = 0) then
   return (segment#_estimated) /*if this is the first packet to be transmitted by subflow c*/
else
   for each subflow_i ∈ (subflow list− subflow_c) do
      rtt_i ← subflow_i(RTT) and
      cwnd_i ← subflow_i(cwnd)
      ratio = rtt_c / rtt_i
      if (ratio > 1) then
         Estimate cwnd_i = CalCwnd(ratio, subflow_i) /* CalCwnd() returns the number of segments
         that will be sent by subflow_i before the estimated segment */
         segment#_estimated += cwnd_i
      end if
   end for
   return (segment#_estimated)
end if
```

$$\text{rtt}_i < \text{rtt}_c \text{ and } \text{subflow}_i \neq \text{subflow}_c$$

The ratio of round trip times of subflows c and i (i.e., $\frac{RTT_c}{RTT_i}$) indicates that the subflow i will get $\frac{RTT_c}{RTT_i}$ times higher number of chances (or transmission opportunities) to transmit the data before the subflow c gets a chance to transmit data. The algorithm also considers congestion window growth of the subflow i. That is on an average how much data it will be able to transmit in a iterations. To estimate the congestion window growth of subflow i, our proposed algorithm considers the congestion phases for subflow (i.e. slow start and congestion avoidance). Based on the congestion phase of subflow, the algorithm tries to calculate the growth of congestion window. To calculate the offset, the procedure CalCwnd (), given in Algorithm 6, uses the congestion control mechanism used by MPTCP.

To obtain the data segment to be transmitted by sublows c, this scheme adds the offset for each subflow i that satisfies the following condition.

$$\frac{RTT_c}{RTT_i} > 1$$

Initially, when a subflow requests data, the algorithm returns the bottommost segment (i.e., next available segment) from send buffer as the scheduler has no idea about the characteristics of subflows (or paths). From the next time step, the segment is to be transmitted by a subflow is estimated by considering the RTTs and the state of the congestion information of all subflows. The algorithm estimates the segments to be transmitted by subflow c when the ratio of the RTTs of subflow c and subflow i is greater than 1. If the ratio is less than 1, the proposed algorithm simply schedules the next available segment in send buffer for subflow c. In other words, all subflows having larger RTTs compared to that of the subflow c will not get a chance to transmit in the current time instance.

Algorithm 6 Segment Estimation Algorithm

```
CalCwnd(ratio, subflow_i)
i_cwnd ← subflow_i(cwnd)
Increment_ss ← 1 /* for slow start */
Increment_ca ← 0 /* for congestion avoidance */
if (subflow_i is in slow start phase) then
  for i = 1 to ratio do
    Increment_ss ← (2*Increment_ss)
  end for
  for j = 1 to (Increment_ss*cwnd_i) do
    cwnd_estimated = INCR_CWND(subflow_i, cwnd_i)
    cwnd_i += cwnd_estimated
  end for
  return (cwnd_i)
else
  /* subflow is in congestion avoidance phase */
  for i = 1 to ratio do
    Increment_ca ← (1+Increment_ca)
  end for
  for j = 1 to (Increment_ca + cwnd_i) do
    cwnd_estimated = INCR_CWND(subflow_i, cwnd_i) /* INCR_CWND() is used by MPTCP to
    increase the congestion window when it receives acknowledge */
    cwnd_i += cwnd_estimated
  end for
  return(cwnd_i)
end if
```

During the estimation of segment number for a subflow c, for all subflows i which satisfy the above stated condition (ratio>1), the algorithm checks the congestion state of subflow i. Depending on the state of the congestion control mechanism, the function CalCwnd() calculates on an average how many segments will be transmitted by subflow i after $\frac{RTT_c}{RTT_i}$ iterations. As followed in previous versions of TCP, the congestion window of subflow i would become double after every RTT time when it is in slow start phase. The congestion window increases by one per RTT when it is in congestion avoidance phase. In order to avoid over estimation of segments the proposed algorithm considers only the value of congestion window after $\frac{RTT_c}{RTT_i}$ iterations for the offset.

The advantages of the proposed algorithm are:

- Less re-ordering at receiver side,
- Small receive buffer at connection level at the receiver,
- Reduced average waiting time of segment in order to deliver to the receiver application,
- Less number of re-transmissions and out of order segments received at receiver side, and
- Fast data delivery to the application.

A.5 EVALUATION

The proposed algorithm is implemented in NS-3 simulator. The segment size for each subflow is taken as 1400 bytes in simulations. The performance of the proposed algorithm is compared with that of simple MPTCP, in which data scheduler schedules the next available segments from the sender's send buffer to the available subflow whenever it is called.

A.5.1 Simulation Setup

The simulation studies are performed for two scenarios, two subflows and three subflows.

Two Subflows

The simulation topology for two subflows is shown in Figure A.3. The characteristics of subflows for different scenarios are presented in Table A.1. To simulate two subflows scenario, 6 configurations have been taken to compare the performance of the proposed algorithm with that of MPTCP.

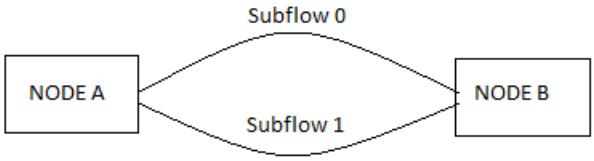

Figure A.3. : Simulation Topology for 2 subflows

Delay Ratio	$Subflow_0$	$Subflow_1$
1 : 1	10 ms,10 Mbps	10 ms,10 Mbps
1 : 5	10 ms,10 Mbps	50 ms,10 Mbps
1 : 10	10 ms,10 Mbps	100 ms,10 Mbps
1 : 15	10 ms,10 Mbps	150 ms,10 Mbps
1 : 20	10 ms,10 Mbps	200 ms,10 Mbps
1 : 25	10 ms,10 Mbps	250 ms,10 Mbps

Table A.1. : Characteristics of subflows in two subflow scenario

Three Subflows

The simulation topology for three subflows is shown in Figure A.4. The characteristics of subflows for different scenarios are presented in Table A.2 for three path scenario. For the three subflows, three configurations namely Type1, Type2 and Type3 are considered.

A.5.2 Simulation Results

To compare the results, the performance metrics used in this work are the number of retransmissions, the number of re-ordered segments received at receiver, and the average waiting time. In case of link asymmetry, the overall performance of multipath TCP suffers as re-ordering at the receiver causes re-transmission of segments and also increases the waiting time of segments at out of order buffer. One can notice that MPTCP shows improved performance when it is augmented with the proposed algorithm which predicts the segment to be transmitted in order to minimize the re-ordering at the receiver.

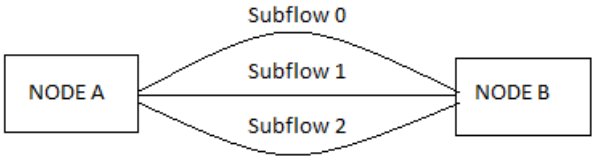

Figure A.4. : Simulation Topology for 3 subflows

(a) Type1

$\textbf{Subflow}_0$	$\textbf{Subflow}_1$	$\textbf{Subflow}_2$
20 ms	50 ms	100 ms
25 Mbps	10 Mbps	10 Mbps

(b) Type2

$\textbf{Subflow}_0$	$\textbf{Subflow}_1$	$\textbf{Subflow}_2$
20 ms	50 ms	100 ms
20 Mbps	10 Mbps	25 Mbps

(c) Type3

$\textbf{Subflow}_0$	$\textbf{Subflow}_1$	$\textbf{Subflow}_2$
10 ms	50 ms	250 ms
10 Mbps	50 Mbps	25 Mbps

Table A.2. : Characteristics of subflows in three subflow scenario

Two Subflows

Figures A.5, A.6, and A.7 show the performance of MPTCP and the proposed algorithm for two subflows. It is clear from these results that the proposed algorithm augmented with MPTCP is able to reduce the number of re-transmitted segments, total number of re-ordered segments received at receiver and average waiting time per segments as compared to MPTCP.

In case of symmetric links, the performance (in terms of total re-transmissions, total re-ordered segments and average waiting time per segments) of the proposed algorithm augmented with MPTCP is same as that of original MPTCP as shown in Figures A.5, A.6 and A.7 for two subflows when delay ratio is 1: 1. In asymmetric scenarios when delay ratio between subflows are 1: 5, 1: 10, 1: 15, 1: 20 and 1: 25 with two subflows, the proposed data distribution algorithm estimates segments for subflows by considering their characteristics and performs much better than original MPTCP.

Three Subflows

Figures A.8, A.9, and A.10 show the performance of MPTCP and the proposed algorithm for three subflows with different flows characteristics. In the case of three subflows, Figures A.8, A.9, and A.10 show the performance improvement of the proposed algorithm over MPTCP.

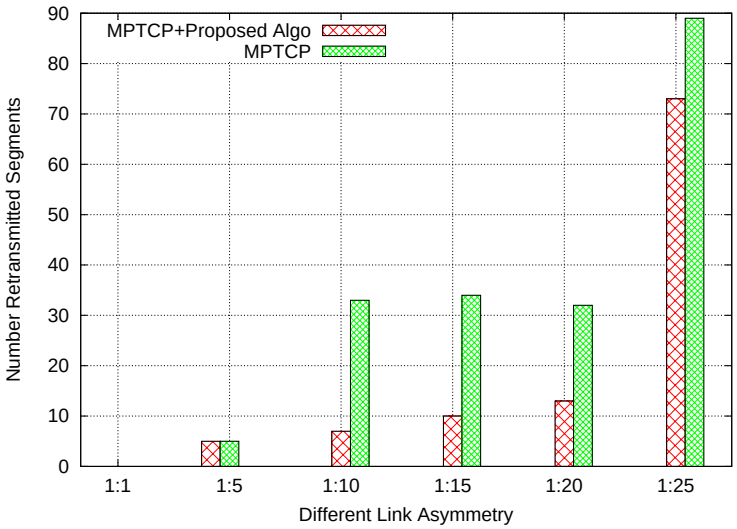

Figure A.5. : Total number of re-transmissions

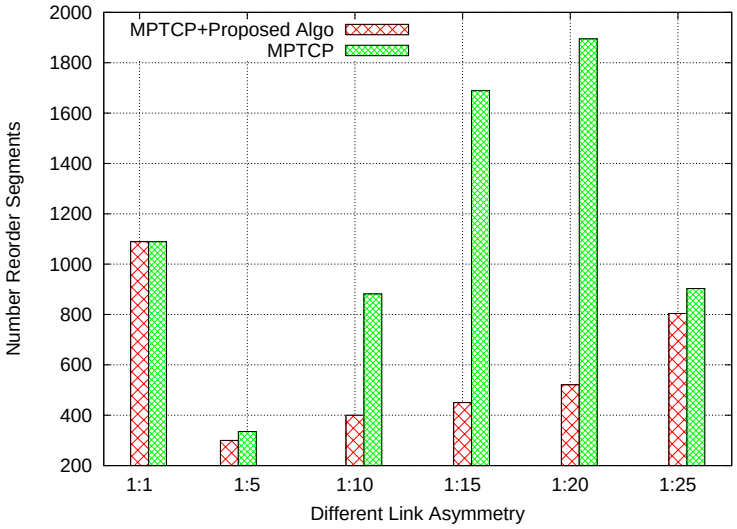

Figure A.6. : Total number of re-ordered segments

A.5.3 Our Contribution

To meet the requirements of data distribution mechanism in multipath TCP, this algorithm proposes a scheduler at transport layer that predicts the data segment to be transmitted for a subflow that is currently requesting segments for transmission. The segment to be transmitted by a subflow is estimated based on round trip time and current congestion window (CWND). The scheduler minimizes the total re-ordering at the destination and reduces the segment waiting time of segment in the receive buffer at the receiver.

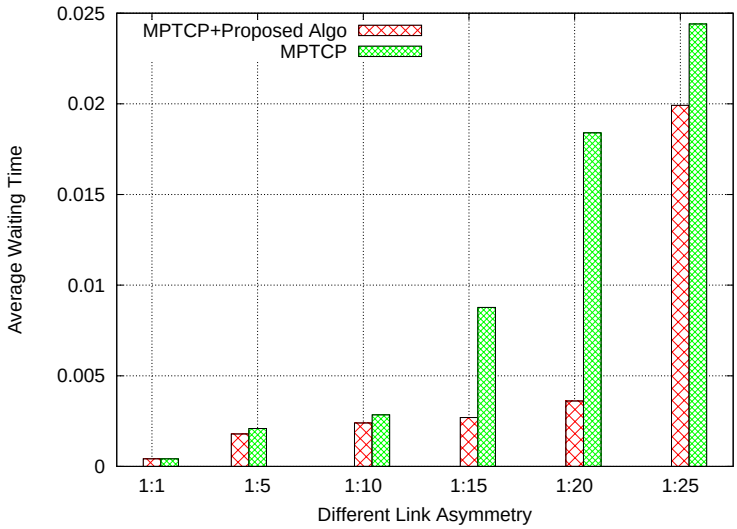

Figure A.7. : Average waiting time per segment

A.6 SUMMARY

This work examined that how to improve the overall performance of multipath TCP by estimating the segment for a given subflow in order to minimize the re-ordering at the receiver. With the proposed algorithm, the number of re-transmissions, average waiting time of segments and out of order segments received at receiver side are reduced. This work found that in case of asymmetric link, the overall performance of simple multipath TCP suffers due to re-ordering at the receiver. The performance of the proposed algorithm with MPTCP is same as that of simple MPTCP in the case of symmetric links between end nodes.

...

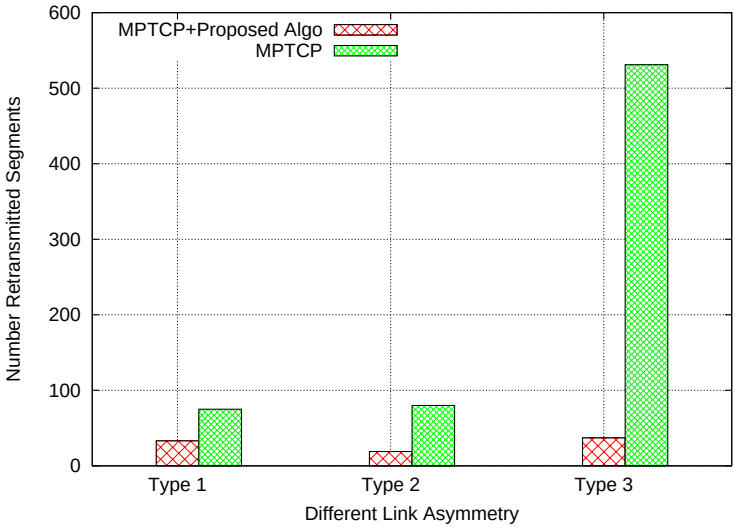

Figure A.8. : Total number of re-transmissions

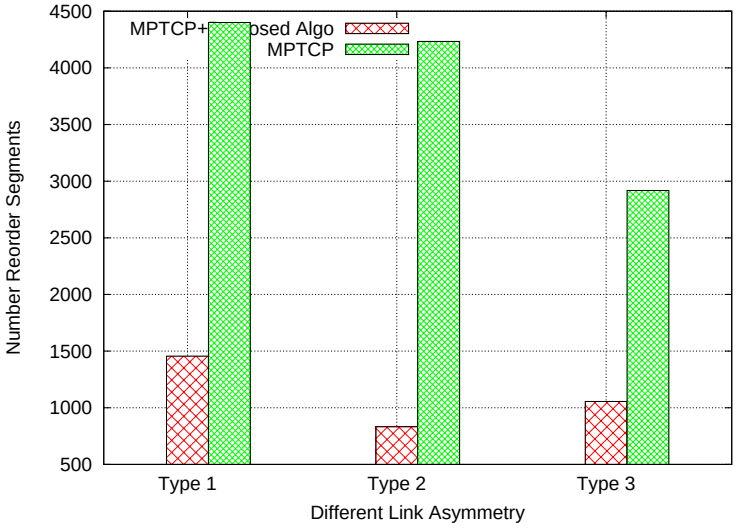

Figure A.9. : Total number of re-ordered segments

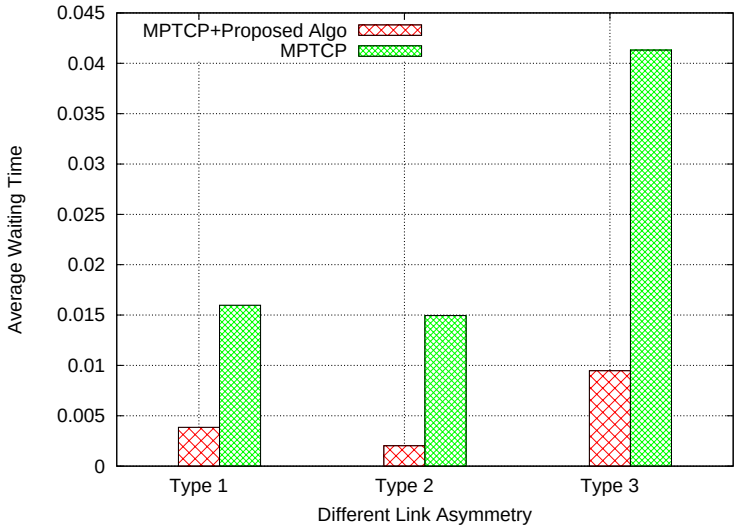

Figure A.10. : Average waiting time per segment

References

Abts, D. and Felderman, B., (2012), "A guided tour through data-center networking", Queue, Vol.10, No.5, pp.10, 2012

Al-Fares, M., Loukissas, A., and Vahdat, A., (2008), "A scalable, commodity data center network architecture", ACM SIGCOMM Computer Communication Review, Vol.38, No.4, pp.63–74, 2008

Al-fares, M., Radhakrishnan, S., Raghavan, B., Huang, N., and Vahdat, A., (2010), "Hedera: Dynamic flow scheduling for data center networks", In In Proc. of Networked Systems Design and Implementation (NSDI) Symposium

Alizadeh, M., Edsall, T., Dharmapurikar, S., Vaidyanathan, R., Chu, K., Fingerhut, A., Matus, F., Pan, R., Yadav, N., Varghese, G., et al., (2014), "CONGA: Distributed congestion-aware load balancing for datacenters", In Proceedings of the 2014 ACM conference on SIGCOMM, pp. 503–514, ACM

Alizadeh, M., Greenberg, A., Maltz, D. A., Padhye, J., Patel, P., Prabhakar, B., Sengupta, S., and Sridharan, M., (2010), "Data Center TCP (DCTCP)", In Proceedings of the ACM SIGCOMM 2010 Conference, SIGCOMM '10, pp. 63–74, New York, NY, USA, ACM

Alizadeh, M., Javanmard, A., and Prabhakar, B., (2011), "Analysis of DCTCP: stability, convergence, and fairness", In Proceedings of the ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems, pp. 73–84, ACM

Alizadeh, M., Kabbani, A., Edsall, T., Prabhakar, B., Vahdat, A., and Yasuda, M., (2012), "Less is More: Trading a Little Bandwidth for Ultra-low Latency in the Data Center", In Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation, NSDI'12, Berkeley, CA, USA, USENIX Association

Alizadeh, M., Yang, S., Sharif, M., Katti, S., McKeown, N., Prabhakar, B., and Shenker, S., (2013), "pFabric: Minimal Near-optimal Datacenter Transport", In Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM, SIGCOMM '13, pp. 435–446, New York, NY, USA, ACM

Barré, S., Paasch, C., and Bonaventure, O., (2011), "Multipath TCP: from theory to practice", NETWORKING 2011,, pp.444–457, 2011

Chen, K., Hu, C., Zhang, X., Zheng, K., Chen, Y., and Vasilakos, A. V., (2011), "Survey on routing in data centers: insights and future directions", IEEE network, Vol.25, No.4, 2011

Chen, L., Li, B., and Li, B., (2013), "On meeting deadlines in datacenter networks", Tsinghua Science and Technology, Vol.18, No.3, pp.273–285, 2013

Chowdhury, M., Zaharia, M., Ma, J., Jordan, M. I., and Stoica, I., (2011), "Managing Data Transfers in Computer Clusters with Orchestra", In Proceedings of the ACM SIGCOMM 2011 Conference, SIGCOMM '11, pp. 98–109, New York, NY, USA, ACM

Cisco, (2015), "Cisco Global Cloud Index: Forecast and Methodology, 2015–2020", 2015

Cui, W. and Qian, C., (2014), "Difs: Distributed flow scheduling for adaptive routing in hierarchical data center networks", In Proceedings of the tenth ACM/IEEE symposium on Architectures for networking and communications systems, pp. 53–64, ACM

Divakaran, D. M., Altman, E., Post, G., Noirie, L., and Primet, P. V.-B., (2009), "From packets to XLFrames: sand and rocks for transfer of mice and elephants", In INFOCOM Workshops 2009, IEEE, pp. 1–6, IEEE

Dixit, A., Prakash, P., Hu, Y. C., and Kompella, R. R., (2013), "On the impact of packet spraying in data center networks", In INFOCOM, 2013 Proceedings IEEE, pp. 2130–2138, IEEE

Dukkipati, N. and McKeown, N., (2006), "Why flow-completion time is the right metric for con-

gestion control", ACM SIGCOMM Computer Communication Review, Vol.36, No.1, pp.59–62, 2006

Farrington, N., Porter, G., Radhakrishnan, S., Bazzaz, H. H., Subramanya, V., Fainman, Y., Papen, G., and Vahdat, A., (2010), "Helios: a hybrid electrical/optical switch architecture for modular data centers", ACM SIGCOMM Computer Communication Review, Vol.40, No.4, pp.339–350, 2010

Greenberg, A., Hamilton, J., Maltz, D. A., and Patel, P., (2008), "The Cost of a Cloud: Research Problems in Data Center Networks", SIGCOMM Comput. Commun. Rev., Vol.39, No.1, pp.68–73, December 2008

Greenberg, A., Hamilton, J. R., Jain, N., Kandula, S., Kim, C., Lahiri, P., Maltz, D. A., Patel, P., and Sengupta, S., (2009), "VL2: a scalable and flexible data center network", In ACM SIGCOMM computer communication review, volume 39, pp. 51–62, ACM

Grosvenor, M. P., Schwarzkopf, M., Gog, I., Watson, R. N., Moore, A. W., Hand, S., and Crowcroft, J., (2015), "Queues don't matter when you can jump them!", In NSDI, pp. 1–14

Guo, C., Lu, G., Li, D., Wu, H., Zhang, X., Shi, Y., Tian, C., Zhang, Y., and Lu, S., (2009), "BCube: a high performance, server-centric network architecture for modular data centers", ACM SIGCOMM Computer Communication Review, Vol.39, No.4, pp.63–74, 2009

Guo, C., Wu, H., Tan, K., Shi, L., Zhang, Y., and Lu, S., (2008), "Dcell: a scalable and fault-tolerant network structure for data centers", In ACM SIGCOMM Computer Communication Review, volume 38, pp. 75–86, ACM

Handley, M., Raiciu, C., Agache, A., Voinescu, A., Moore, A. W., Antichi, G., and Wójcik, M., (2017), "Re-architecting datacenter networks and stacks for low latency and high performance", In Proceedings of the Conference of the ACM Special Interest Group on Data Communication, pp. 29–42, ACM

He, K., Rozner, E., Agarwal, K., Felter, W., Carter, J., and Akella, A., (2015), "Presto: Edge-based load balancing for fast datacenter networks", In Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication, pp. 465–478, ACM

T. Hoff,Latency is everywhere and it costs you sales - How to crush it,

Scott Hogg,Does your network support Jumbo Frames and should you enable it?,

Hong, C.-Y., Caesar, M., and Godfrey, P., (2012), "Finishing flows quickly with preemptive scheduling", ACM SIGCOMM Computer Communication Review, Vol.42, No.4, pp.127–138, 2012

Hopps, C. E., (2000), "Analysis of an equal-cost multi-path algorithm", 2000

Jalaparti, V., Bodik, P., Kandula, S., Menache, I., Rybalkin, M., and Yan, C., (2013), "Speeding Up Distributed Request-response Workflows", In Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM, SIGCOMM '13, pp. 219–230, New York, NY, USA, ACM

Joy, S. and Nayak, A., (2015), "Improving flow completion time for short flows in datacenter networks", In Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on, pp. 700–705, IEEE

Kabbani, A., Vamanan, B., Hasan, J., and Duchene, F., (2014), "FlowBender: Flow-level Adaptive Routing for Improved Latency and Throughput in Datacenter Networks", In Proceedings of the 10th ACM International on Conference on emerging Networking Experiments and Technologies, pp. 149–160, ACM

Kant, K., (2009), "Data center evolution: A tutorial on state of the art, issues, and challenges", Computer Networks, Vol.53, No.17, pp.2939–2965, 2009

Krevat, E., Vasudevan, V., Phanishayee, A., Andersen, D. G., Ganger, G. R., Gibson, G. A., and Seshan, S., (2007), "On application-level approaches to avoiding TCP throughput collapse in cluster-based storage systems", In Proceedings of the 2nd international workshop on Petascale data storage: held in conjunction with Supercomputing'07, pp. 1–4, ACM

Kumar, R., (2014), "The Mathematical way to Decide", 2014

Lebiednik, B., Mangal, A., and Tiwari, N., (2016), "A Survey and Evaluation of Data Center Network Topologies", CoRR, Vol.abs/1605.01701„ 2016

Lee, C., Park, C., Jang, K., Moon, S. B., and Han, D., (2015), "Accurate Latency-based Congestion Feedback for Datacenters.", In USENIX Annual Technical Conference, pp. 403–415

Liu, S., Xu, H., and Cai, Z., (2013), "Low Latency Datacenter Networking: A Short Survey", CoRR, Vol.abs/1312.3455,, 2013

Mai, L., Rupprecht, L., Alim, A., Costa, P., Migliavacca, M., Pietzuch, P., and Wolf, A. L., (2014), "NetAgg: Using Middleboxes for Application-specific On-path Aggregation in Data Centres", In Proceedings of the 10th ACM International on Conference on Emerging Networking Experiments and Technologies, CoNEXT '14, pp. 249–262, New York, NY, USA, ACM

Mudigonda, J., Yalagandula, P., Al-Fares, M., and Mogul, J. C., (2010), "Spain: Cots data-center ethernet for multipathing over arbitrary topologies.", In NSDI, volume 10, pp. 18–33

Munir, A., Baig, G., Irteza, S. M., Qazi, I. A., Liu, A. X., and Dogar, F. R., (2014), "Friends, Not Foes: Synthesizing Existing Transport Strategies for Data Center Networks", In Proceedings of the 2014 ACM Conference on SIGCOMM, SIGCOMM '14, pp. 491–502, New York, NY, USA, ACM

Munir, A., Qazi, I. A., Uzmi, Z. A., Mushtaq, A., Ismail, S. N., Iqbal, M. S., and Khan, B., (2013), "Minimizing flow completion times in data centers", In INFOCOM, 2013 Proceedings IEEE, pp. 2157–2165, IEEE

Murray, D., Koziniec, T., Lee, K., and Dixon, M.,, "Large MTU's and internet performance.", In 13th IEEE Conference on High Performance Switching and Routing (HPSR 2012), pp. 82–87

Niranjan Mysore, R., Pamboris, A., Farrington, N., Huang, N., Miri, P., Radhakrishnan, S., Subramanya, V., and Vahdat, A., (2009), "Portland: a scalable fault-tolerant layer 2 data center network fabric", In ACM SIGCOMM Computer Communication Review, volume 39, pp. 39–50, ACM

Perry, J., Balakrishnan, H., and Shah, D., (2017), "Flowtune: Flowlet Control for Datacenter Networks", In 14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17), pp. 421–435, USENIX Association

Perry, J., Ousterhout, A., Balakrishnan, H., Shah, D., and Fugal, H., (2014), "Fastpass: A Centralized "Zero-queue" Datacenter Network", In Proceedings of the 2014 ACM Conference on SIGCOMM, SIGCOMM '14, pp. 307–318, New York, NY, USA, ACM

Prakash, P., Dixit, A., Hu, Y. C., and Kompella, R., (2012), "The TCP Outcast Problem: Exposing Unfairness in Data Center Networks", In 9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12), pp. 413–426, San Jose, CA, USENIX

Prakash, P., Lee, M., Hu, Y. C., Kompella, R. R., et al., (2013), "Jumbo frames or not: That is the question!", 2013

Raiciu, C., Barre, S., Pluntke, C., Greenhalgh, A., Wischik, D., and Handley, M., (2011), "Improving datacenter performance and robustness with multipath tcp", ACM SIGCOMM Computer Communication Review, Vol.41, No.4, pp.266–277, 2011

Ramaboli, A. L., Falowo, O. E., and Chan, A. H., (2012), "Bandwidth aggregation in heterogeneous wireless networks: A survey of current approaches and issues", Journal of Network and Computer Applications, Vol.35, No.6, pp.1674–1690, 2012

Rojas-Cessa, R., Kaymak, Y., and Dong, Z., (2015), "Schemes for fast transmission of flows in data center networks", IEEE Communications Surveys & Tutorials, Vol.17, No.3, pp.1391–1422, 2015

Salyers, D., Jiang, Y., Striegel, A., and Poellabauer, C., (2007), "JumboGen: dynamic jumbo frame generation for network performance scalability", ACM SIGCOMM Computer Communication Review, Vol.37, No.5, pp.53–64, 2007

Singla, A., Hong, C.-Y., Popa, L., and Godfrey, P. B., (2012), "Jellyfish: Networking Data Centers, Randomly.", In NSDI, volume 12, pp. 17–17

Vamanan, B., Hasan, J., and Vijaykumar, T., (2012), "Deadline-aware datacenter tcp (d2tcp)", ACM SIGCOMM Computer Communication Review, Vol.42, No.4, pp.115–126, 2012

Vanini, E., Pan, R., Alizadeh, M., Taheri, P., and Edsall, T., (2017), "Let It Flow: Resilient Asymmetric Load Balancing with Flowlet Switching.", In NSDI, pp. 407–420

Vasudevan, V., Phanishayee, A., Shah, H., Krevat, E., Andersen, D. G., Ganger, G. R., Gibson, G. A., and Mueller, B., (2009), "Safe and effective fine-grained TCP retransmissions for datacenter communication", In ACM SIGCOMM computer communication review, volume 39, pp. 303–314, ACM

Wang, G., Andersen, D. G., Kaminsky, M., Papagiannaki, K., Ng, T., Kozuch, M., and Ryan, M., (2010), "c-Through: Part-time optics in data centers", In ACM SIGCOMM Computer Communication Review, volume 40, pp. 327–338, ACM

Wang, T., Su, Z., Xia, Y., and Hamdi, M., (2014), "Rethinking the data center networking: Architecture, network protocols, and resource sharing", IEEE access, Vol.2„ pp.1481–1496, 2014

Wilson, C., Ballani, H., Karagiannis, T., and Rowtron, A., (2011), "Better never than late: Meeting deadlines in datacenter networks", In ACM SIGCOMM Computer Communication Review, volume 41, pp. 50–61, ACM

Wischik, D., Handley, M., and Braun, M. B., (2008), "The resource pooling principle", ACM SIGCOMM Computer Communication Review, Vol.38, No.5, pp.47–52, 2008

Wu, H., Feng, Z., Guo, C., and Zhang, Y., (2013), "ICTCP: Incast congestion control for TCP in data-center networks", IEEE/ACM transactions on networking, Vol.21, No.2, pp.345–358, 2013

Wu, X. and Yang, X., (2012), "Dard: Distributed adaptive routing for datacenter networks", In Distributed Computing Systems (ICDCS), 2012 IEEE 32nd International Conference on, pp. 32–41, IEEE

Xia, W., Zhao, P., Wen, Y., and Xie, H., (2017), "A survey on data center networking (DCN): infrastructure and operations", IEEE Communications Surveys & Tutorials, Vol.19, No.1, pp.640–656, 2017

Xu, H. and Li, B., (2013), "RepFlow: Minimizing flow completion times with replicated flows in data centers", arXiv preprint arXiv:1307.7451,, 2013

Zats, D., Das, T., Mohan, P., Borthakur, D., and Katz, R., (2012), "DeTail: reducing the flow completion time tail in datacenter networks", ACM SIGCOMM Computer Communication Review, Vol.42, No.4, pp.139–150, 2012

Zhang, H., Zhang, J., Bai, W., Chen, K., and Chowdhury, M., (2017), "Resilient datacenter load balancing in the wild", In Proceedings of the Conference of the ACM Special Interest Group on Data Communication, pp. 253–266, ACM

Zhang, T., Wang, J., Huang, J., Huang, Y., Chen, J., and Pan, Y., (2016), "Adaptive marking threshold method for delay-sensitive TCP in data center network", Journal of Network and Computer Applications, Vol.61„ pp.222–234, 2016