Homa: A Data Center Transport Protocol

Harsh Gaurang Kapadia
harshk@bu.edu
U22820240

1. Objective of Study

Homa is a new transport protocol that is touted to be the replacement for the Transmission Control Protocol (TCP) in data centers. It is designed in a manner that eliminates the problems that TCP's connection orientation, stream orientation, in-order packet delivery and congestion control cause, thereby reducing 'datacenter tax' (a collection of low-level overheads that consume a significant fraction of all processor cycles in data centers). [4.1]

I plan to study, understand, compare and analyse the performance improvements that Homa showcases in comparison to transport protocols like Data Center TCP (DCTCP). I also plan to explore the concerns with Homa and find ideas to improve upon those issues.

2. Plan of Study

2.1. Study Structure

- 2.1.1. Understand data center topology and networking
  - 2.1.1.1. Understand buffer sizing and utilization in routers
- 2.1.2. Understand DCTCP
- 2.1.3. Understand Homa
  - 2.1.3.1. Build Homa Linux module
- 2.1.4. Compare the performance of Homa and DCTCP
  - 2.1.4.1. Decide experiments to be conducted
  - 2.1.4.2. Check Buffer Utilization for both
- 2.1.5. Explore the concerns with Homa

2.2. Deliverables

- 2.2.1. Mid-semester report
- 2.2.2. Final report
- 2.2.3. Code and reproducible experiments (GitHub)

2.3. Grade Distribution

2.3.1. Progress/discussion meetings every two weeks: 10%
2.3.2. Mid-semester report: 20%
2.3.3. Final report: 30%
2.3.4. Code and reproducible experiments (GitHub): 40%

3. Rationale behind Study

I am interested in Computer Networks, particularly the Transmission Control Protocol (TCP). I was a student in the BU GRS CS 655 Graduate Introduction to Computer Networks course taught by Professor Abraham Matta in Fall 2022. The course covered TCP extensively. I also did a Performance Comparison of TCP Versions mini-project as a part of the course. This Directed Study builds upon the course and the mini-project to go beyond and explore a newer implementation of transport protocols in the data center called 'Homa'.

4. Bibliography of Readings

4.1. J. Ousterhout. *It's Time to Replace TCP in the Datacenter.* (The second version released in January 2023.)

4.2. B. Montazeri, Y. Li, M. Alizadeh, and J. Ousterhout. *Homa: A Receiver-Driven Low-Latency Transport Protocol Using Network Priorities.* In Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication, SIGCOMM '18, pages 221 - 235, New York, NY, USA, 2018. Association for Computing Machinery.

4.3. J. Ousterhout. *A Linux Kernel Implementation of the Homa Transport Protocol.* In 2021 USENIX Annual Technical Conference (USENIX ATC 21), pages 99 - 115. USENIX Association, July 2021.

4.3.1. The official Homa Linux module source code

4.4. The official Homa Wiki.

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

4.6. M. Noormohammadpour and C. S. Raghavendra. *Datacenter Traffic Control: Understanding Techniques and Tradeoffs.* In IEEE Communications Surveys & Tutorials, vol. 20, no. 2, pp. 1492-1525, Secondquarter 2018, doi:

10.1109/COMST.2017.2782753.

4.7. Theophilus Benson, Aditya Akella, and David A. Maltz. 2010. *Network traffic characteristics of data centers in the wild.* In Proceedings of the 10th ACM SIGCOMM conference on Internet measurement (IMC '10). Association for Computing Machinery, New York, NY, USA, 267–280.

4.8. Svilen Kanev, Juan Pablo Darago, Kim Hazelwood, Parthasarathy Ranganathan, Tipp Moseley, Gu-Yeon Wei, and David Brooks. 2015. *Profiling a warehouse-scale computer.* In Proceedings of the 42nd Annual International Symposium on Computer Architecture (ISCA '15). Association for Computing Machinery, New York, NY, USA, 158–169.

4.9. G. Appenzeller, I. Keslassy, and N. McKeown. *Sizing router buffers.* In SIGCOMM, 2004.

4.10. Neda Beheshti, Yashar Ganjali, Monia Ghobadi, Nick McKeown, and Geoff Salmon. 2008. *Experimental study of router buffer sizing.* In Proceedings of the 8th ACM SIGCOMM conference on Internet measurement (IMC '08). Association for Computing Machinery, New York, NY, USA, 197–210.

4.11. Damon Wischik and Nick McKeown. 2005. *Part I: buffer sizes for core routers.* SIGCOMM Comput. Commun. Rev. 35, 3 (July 2005), 75–78.

4.12. W. Bai, S. Hu, K. Chen, K. Tan and Y. Xiong. *One More Config is Enough: Saving (DC)TCP for High-Speed Extremely Shallow-Buffered Datacenters.* In IEEE/ACM Transactions on Networking, vol. 29, no. 2, pp. 489-502, April 2021, doi: 10.1109/TNET.2020.3032999.