

Lossless Ethernet? Performance Evaluation of Paused-Based Networks

Chia-Hui Tai

Kok Kiong Yap

Email:chtai@stanford.edu

Email:yapkke@stanford.edu

April 26, 2008

1 Introduction

Pause-based flow control avoids data loss by temporarily stopping data transmission before buffer overflow occurs. This mechanism has been proposed as an optional standard in 802.3x. The purpose was to exploit commodity Ethernet for lossless networks such as SONET and Fiber Channel. For example, this standard benefits Fibre Channel over Ethernet (FCoE), which gains increasing popularity in storage area network because of its economical cost while having comparable performances in line rate as the Fiber Channel. Literature also shows that this standard was brought up for Instant Messaging and Metro Ethernet.

Despite the appealing concept of a lossless Ethernet, many people have raised concerns about the pause-based flow control. For instance, pauses in a fully connected network are likely to cause deadlocks and freeze the network connections. As another example, the pause-based flow control, defined as a link-layer protocol, duplicates the functionality that already resides in TCP and other transport-layer protocols. The performance of TCP over a pause-enabled Ethernet would become an issue.

While many people have concerns about pause-based networks, there is not yet any detailed and enlightening evaluation of those raised problems. In this project, we will address those issues systematically, in a bid to provide clear insights into the pause-based flow control for those who are interested in adopting this mechanism.

In this project, we will first concentrate on: (1) the possible deadlocks in a pause-based network, and (2) the interactions between the pause-based flow control and TCP. As we progress along, we hope to provide a clear crisp view of the current affairs in pause-based networks. We would target at delivering relevant performance evaluations to fulfill this purpose.

2 Objectives

In this project, we expect to have the following deliverables:

1. Deadlock-free pauses
 - (a) A survey of the current literature on deadlock free networks.
 - (b) The necessary and sufficient conditions for deadlocks in a pause-based network. Network topologies and routing algorithms are likely to be included in this characterization.
 - (c) Discussions of existing solutions to deadlocks in a pause-based network.
 - (d) A review and evaluations of the performance of deadlock free topologies and routing algorithms.
2. Interactions with TCP
 - (a) Performance evaluations of TCP running on top of a pause-enabled Ethernet.
 - (b) Characterizations of an appropriate congestion control mechanism for pause-based networks.

3 Timeline

- 4/21 - 5/04 Collect relate work from publications, forums, and IT news
- 5/05 - 5/11 Formulate tractable problems; Obtain initial results
- 5/12 - 5/25 Derive final results; Verify outcomes with simulation data
- 5/26 - 6/02 Wrap up; Prepare for the presentation

References

- [1] F. C. I. Assocation. Fibre channel over ethernet in the data center, April 2008.
- [2] L. Cherkasova, V. Kotov, and T. Rokicki. Fibre channel fabrics: Evaluation and design. In *Proceedings of the 29th Annual Hawaii international Conference on System Sciences*, 1996.
- [3] K. D. Gunther. Prevention of deadlocks in packet-switched data transport systems. *IEEE Transactions on Communications*, 1981.
- [4] QLogic. Ethernet: The drive for convergence.