Date: Mar 27, 2000

To: T10 Committee (SCSI)

From: George Penokie (IBM)

Subject: Training Timings Modeling Results

Overview

The following charts show the results of adding in training at the start of each physical connection. The charts compare training lengths of 500 nsec, 1000 nsec, 1500 nsec, and 2000 nsec to no training pattern. The work load is set to queue 6 read and 4 write commands with a 'normal' number of disconnects and re-connects.

The charts contain, in order:

a) Parallel running at fast-160
b) Parallel running at fast-320
c) Packetized running at fast-160
d) Packetized running at fast-320

Summary

The results indicate the performance impact of training at the beginning of each connection to be minimal. I recommend that SPI-4 define the training pattern to always occur at each connection. The target would train on only the first occurrence of a DT DATA IN phase and a DT DATA OUT phase during a connection.

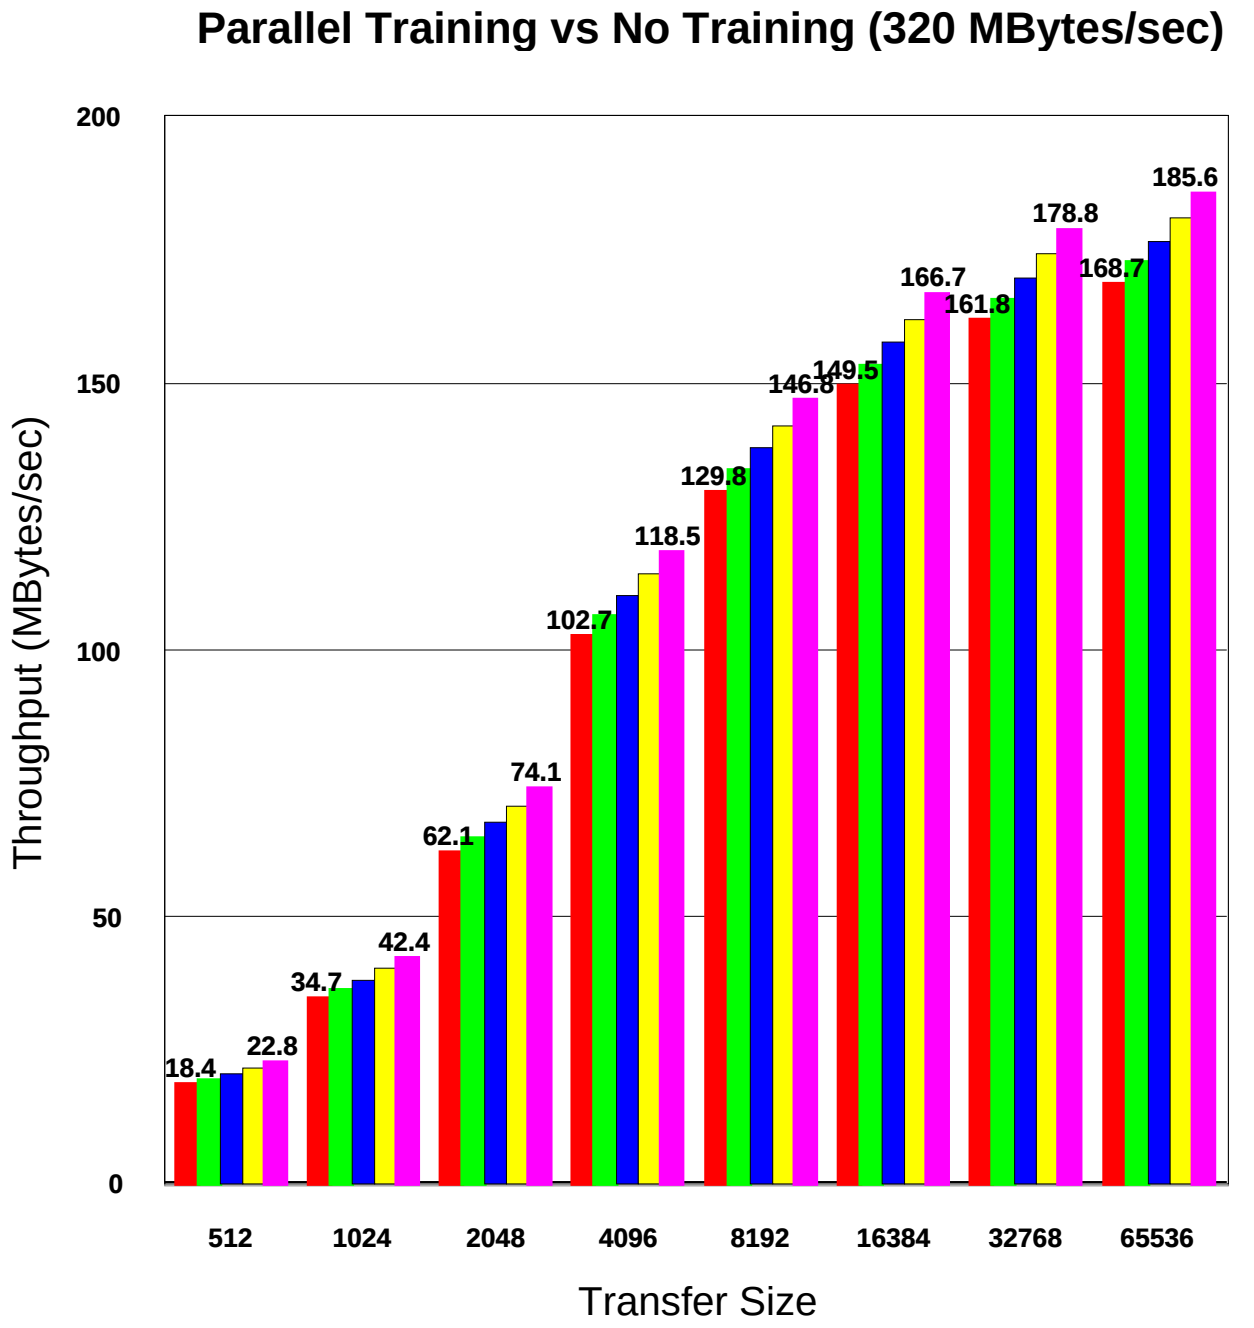

Figure 1 - Parallel training vs. no training at 320 Megabytes/sec

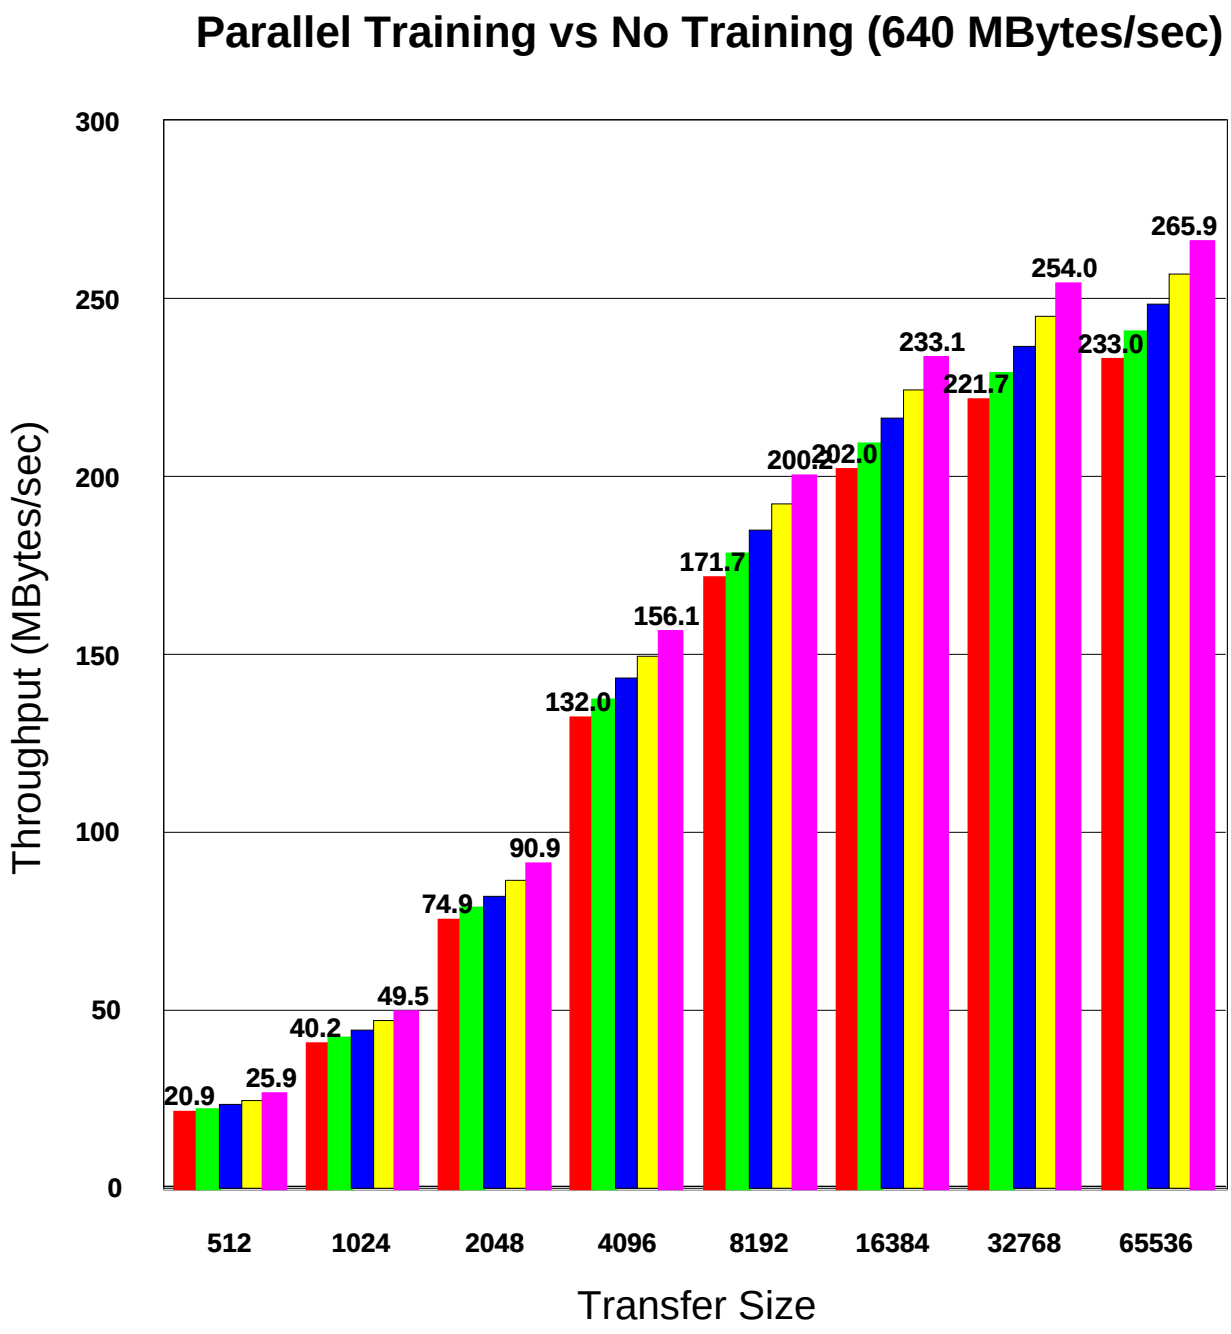

Figure 2 - Parallel training vs. no training at 640 Megabytes/sec

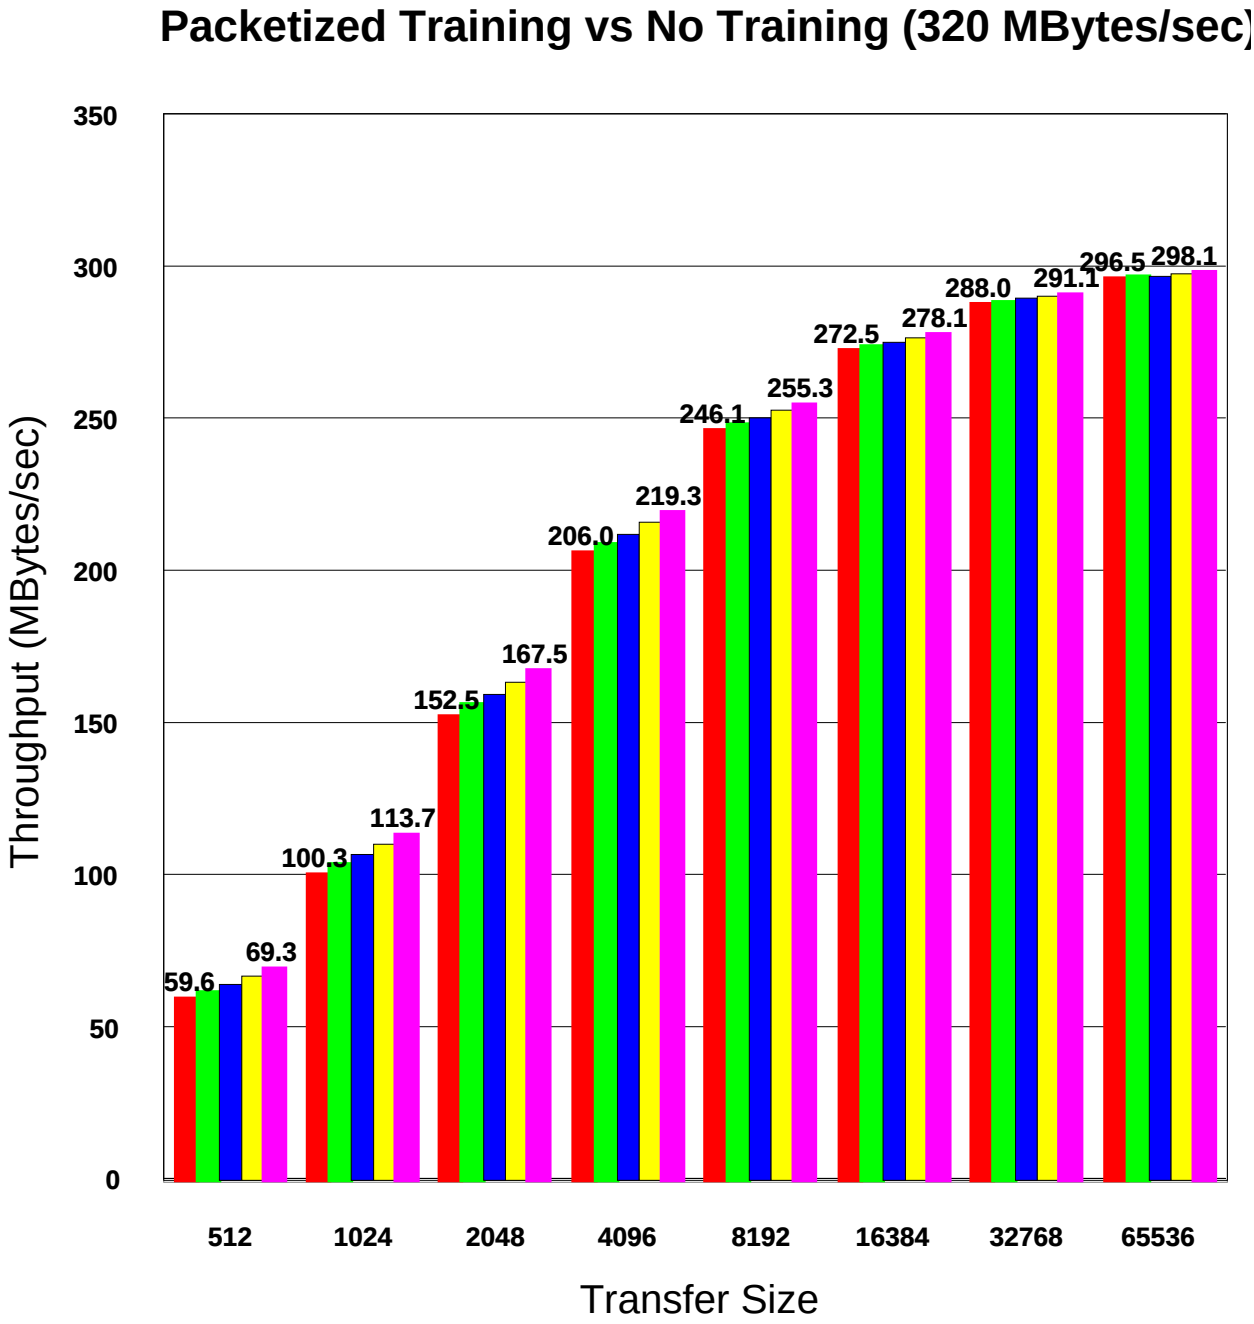

Figure 3 - Packetized training vs. no training at 320 Megabytes/sec

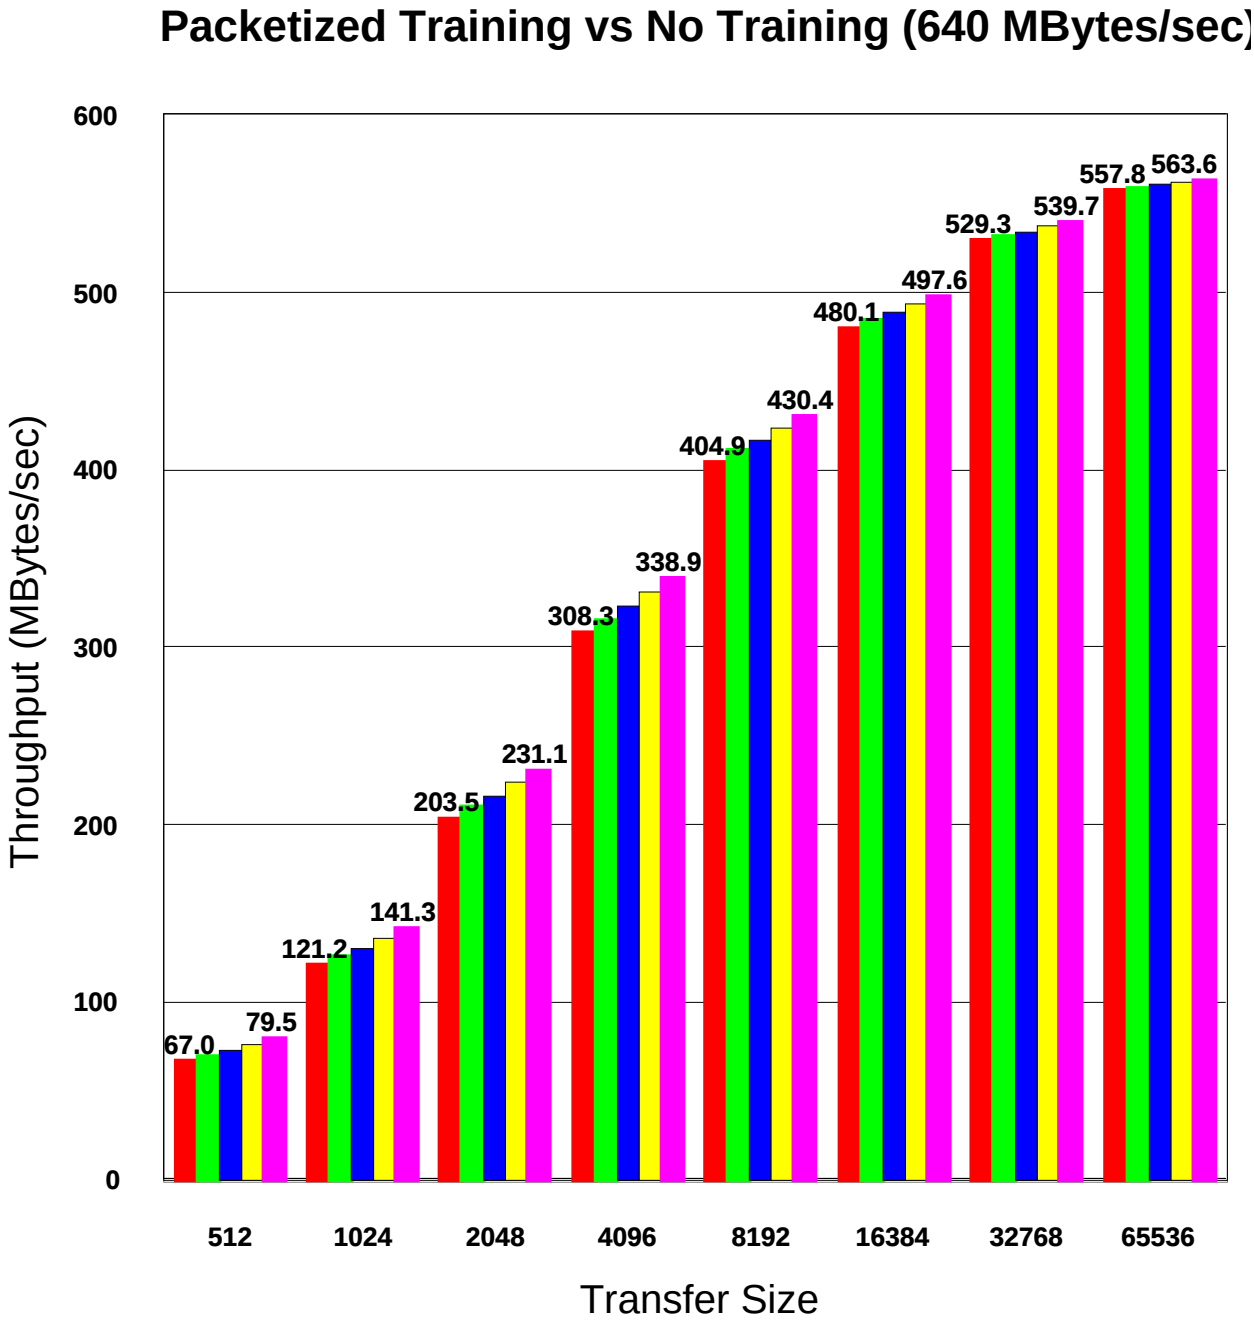

Figure 4 - Packetized training vs. no training at 640 Megabytes/sec