

Comparing the Performance of IP Over Ethernet and IEEE-1394 on a Java Platform

R.C. Norris and D.M. Miller
VLSI Design and Test Group
Department of Computer Science
University of Victoria
cnorris,mmiller@csc.uvic.ca

Abstract

Multimedia communication requires predictable network performance to operate smoothly and correctly. Our work quantifies inter-computer communication performance of Java using TCP/IP and UDP/IP over Ethernet and IEEE-1394 isochronous streams.

1. Introduction

Java [1] is an effective platform for distributed multimedia. Java is a modern, robust, object-oriented language and runtime system with a rich set of libraries. The Java platform provides a virtual machine (JVM) and services to decrease application complexity such as automated garbage collection and transparent scheduling of multiple threads. The Java platform also provides libraries to support time-based media and network communication.

Ethernet [2,3] is the dominant network technology for small to medium-sized networks. Some Ethernet technologies can transmit data at speeds of 1000 Mbps, but the most common Ethernet variants transmit data at 10Mbps or 100Mbps. Ethernet's only transmission mode, Collision Sense Multiple Access Collision Detect (CSMA/CD), does not provide deterministic network access. The time an application must wait before writing to the Ethernet media is unbounded.

The Internet Protocol (IP) [4,7] is a common networking protocol for inter-computer communication. The IP works on top of a physical network such as Ethernet. Two common protocols that provide additional services on top of IP are the Transmission Control Protocol (TCP/IP) [5,7] and the User Datagram Protocol (UDP/IP) [6,7]. The TCP/IP is a stream-based protocol with error detection, error correction and guaranteed in-order delivery of data. The UDP/IP is a datagram-based protocol with error

detection, no error correction and out of order delivery of data.

IEEE-1394 [8] (also known as FireWire and iLink) is a high performance serial bus designed for multimedia data transmission. The current version of IEEE-1394 transmits data at 400 Mbps in a local environment to up to 64 nodes. IEEE-1394 provides two transmission modes: asynchronous and isochronous streams. The asynchronous transmission of IEEE-1394 provides bounded access to the network with error correction and in-order delivery of data. The isochronous streams of IEEE-1394 provide guaranteed constant bit rate bandwidth with no error detection and in-order delivery of data. IEEE-1394 provides additional services reducing the need for higher-level protocols.

2. Experiment Configuration

In this research we compare the network performance of the TCP/IP and the UDP/IP over Ethernet and IEEE-1394 isochronous streams. All experiments were conducted with Sun's Java virtual machine version 1.3 for Linux. Our primary performance measures are the first and second statistical moments of inter-packet delay and transmit lag.

Our experiments use two Java applications executing on two identical machines connected by 100 Mbps Ethernet and 400 Mbps IEEE-1394. The sender Java application transmits 50,000 fixed sized packets to the receiver Java application. We recorded the time of each packet as it was sent and received. The packet transmission occurs across a private network with no other activity. All timestamps were recorded with microsecond resolution. A network time protocol daemon (ntpd) [9] kept both machines' clocks synchronized with a maximum estimated error of 6 microseconds. In order to minimize the effects of garbage collection not related to the network communication we forced a complete garbage collection prior to performing each experiment.

Each experiment was repeated five times for each of the three packet sizes and each of the three network types. We conducted a total of 45 experiments to produce the raw data used to derive inter-packet delay and transmit lag. We derived inter-packet delay from the difference from two consecutive packet timestamps on a single machine. We derived the transmission lag from the difference from a single packet's timestamps from both machines. We present the summary statistics from our experiments with the 99% and 95% quantiles removed. Removing upper quantiles eliminates large extreme values not directly associated with network access.

3. Experimental Results

The characteristics of data transmitted across the two physical networks differ in several areas. The first difference is the minimum time between two packets sent from the same host. An IEEE-1394 host must wait 125 microseconds after isochronous transmission before transmitting another isochronous packet. An Ethernet host must only wait long enough to gain exclusive access to the physical media (approximately 0.096 microseconds on a quiet 100baseTX network). The second difference is the factors that determine maximum transmit rate. The maximum Ethernet rate is a function of electrical signaling rate, higher-level protocol overhead, and packet. The maximum IEEE-1394 isochronous data rate is a function of packet size only.

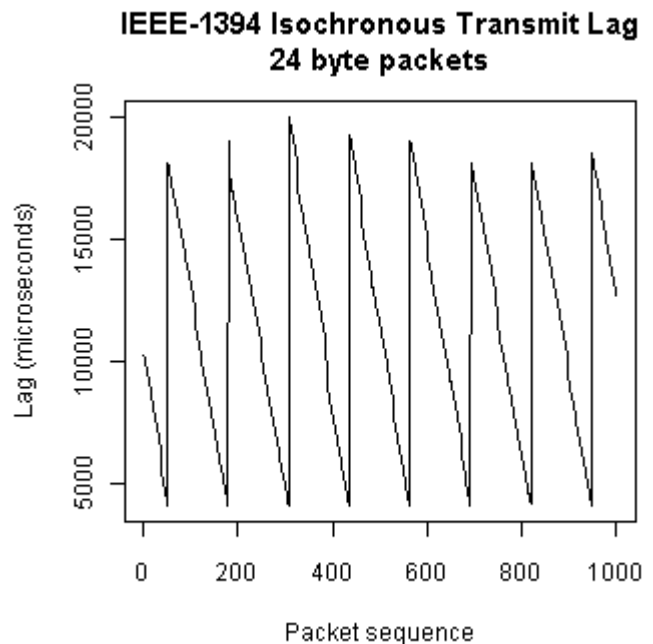
We must consider a special aspect of the IP network access to evaluate the collected data. Both the TCP/IP and the UDP/IP implementations buffer packets before sending them on the physical network and after receiving them from the physical network. Most values of inter-packet delay from the IP networks record buffer copying operations not network access operations.

The ideal IEEE-1394 isochronous transmission rate is 125 microseconds per packet regardless of packet size. Our Java implementation transmits data at 145 microseconds per packet with a standard deviation of 244. It may appear that isochronous transmission is bursty. Further investigation reveals that the 10% and 90% quantile values from inter-packet delay across all IEEE-1394 isochronous experiments are 123 to 136 microseconds per packet with standard deviations of 0

and 1.06 respectively. When the 99% quantile is removed from the data set, the mean transmit rate is 130 microseconds per packet with a standard deviation of 37.9. The observed isochronous inter-packet delay receive rate was similar to the inter-packet delay transmit rate.

IEEE-1394 isochronous streams transmit lag should be independent of packet size, but input buffering by the receiver increases delay inversely to packet size. The Linux implementation of the IEEE-1394 driver buffers incoming isochronous packets until 4096 bytes accumulate before forwarding isochronous packets to user applications. The effects of input buffering can clearly be seen in Figure 1. Input buffering also contributes to transmit lag variability. IEEE-1394 isochronous packets waiting in the input buffer are delayed until the buffer is filled. The last packet arrives to arrive may be an extreme value from the 99% quantile. Since the entire buffer is delayed, larger lag values are recorded for all packets in the receiver's input buffer.

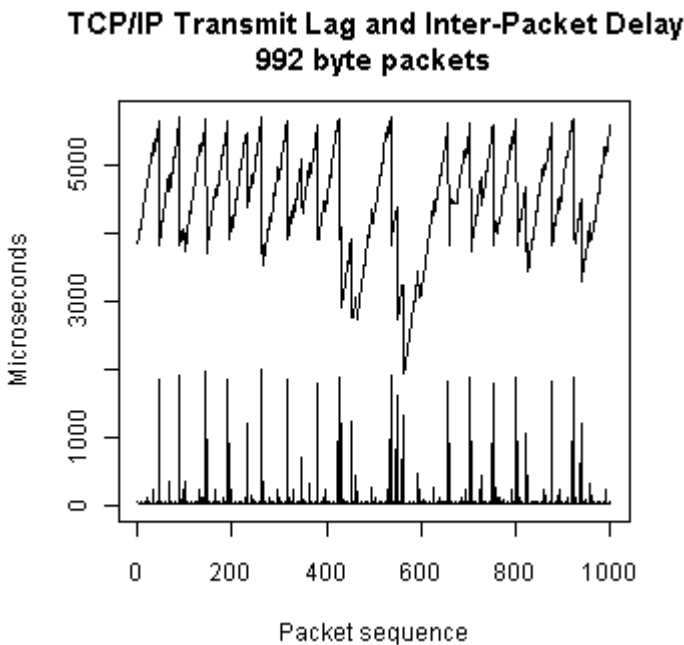
Figure 1



The transmit lag experienced by 992 byte TCP/IP packets is surprisingly high at more than 4 milliseconds. The cause of the large transmit lag value can be seen in Figure 2. The lower sender inter-packet delay plot has spikes that coincide in time and magnitude with the drops in the upper transmit lag

plot. The steep positive slope of transmit lag indicates that a buffer is filling faster than it can be emptied. The bottleneck may be between the sender's output buffer and the physical network, or between the receiver's input buffer and the TCP/IP protocol processing. The former scenario appears more likely because the spikes comprise 45% of the inter-packet delay in Figure 2 and the Ethernet utilization is 99%. If the bottleneck were the receiver's input buffer, the network would be idle during sender inter-packet delay spikes. The sender's output buffer delays packets several milliseconds before transmission, significantly contributing to TCP/IP transmit lag. Transmission of other TCP/IP packet sizes does not saturate the Ethernet and therefore is not subject to input buffer induced transmit lag.

Figure 2

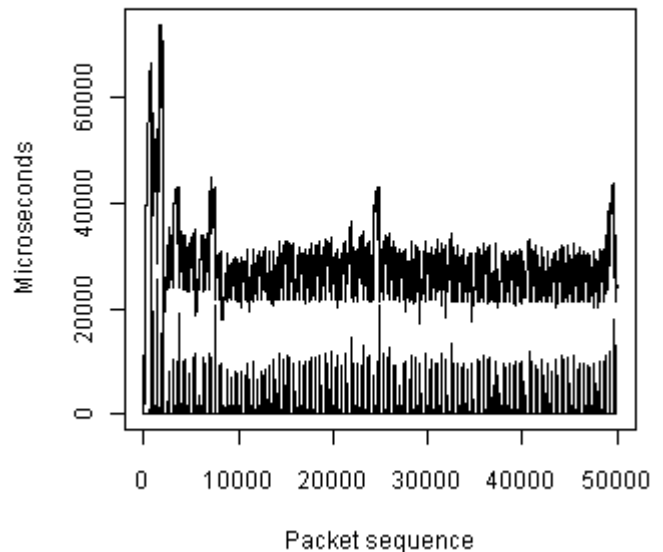


The transmit lag experienced by UDP/IP packets increases as packet size decreases. For each of the transmit lag values for the three packet sizes, the mean and median values are similar. Spikes in UDP/IP transmit lag exist but their contribution to mean is smaller than the other network transmissions. Removing the spikes contained in the 95% quantile does not reduce the transmit lag mean as significantly as the TCP/IP or IEEE-1394 isochronous streams. Removing the 95% quantile does significantly reduce variance of UDP/IP transmit lag. Buffering contributes to transmit lag values for all packet sizes, but increases

as packet size decreases. Transmit lag experienced by UDP/IP packets of 24 bytes is much larger packets of TCP/IP and IEEE-1394 isochronous streams. The upper plot in Figure 3 is transmit lag and the lower plot is inter-packet delay.

Figure 3

UDP/IP Transmit Lag and Inter-Packet Delay
24 byte packets



4. Conclusion

From the statistical summary of IEEE-1394 isochronous transmit lag in Table 1, we observe that removing the 99% quantile has a large effect on transmit lag and variance for packet sizes 992 bytes and 396 bytes, but not 24 bytes. Input buffering of small packets overwhelms all other factors contributing to IEEE-1394 isochronous transmit lag.

From the discussion earlier in this paper and the statistical summary of TCP/IP transmit lag in Table 1, we observe that output buffering has the largest effect on transmit lag and variance for 992 byte packets. We observe the transmit lag mean and variance of the two other packet sizes are greatly affected by extreme values located in the 95% quantile.

From the statistical summary of UDP/IP transmit lag in Table 1, we observe that values in the 99% quantile have a large effect on the variance and a small effect on the mean. Buffering by the UDP/IP protocol

implementation induces large transmit lag on small packet sizes.

We observe that some services of the Java platform designed to reduce program complexity decrease the performance of multimedia network communication. The values of the 99% quantile suggest that multimedia applications must compensate for 20

millisecond processing interruptions. Buffering can reduce inter-packet delay time to almost zero, but can increase transmit lag substantially. Tuning sender and receiver buffer sizes for various network conditions will enhance network application performance, especially streaming multimedia applications.

	Packet size	Median	Mean	Standard deviation	With Q 95 removed			With Q 99 removed		
					Q 95	Mean	Standard deviation	Q 99	Mean	Standard deviation
TCP/IP	992	4435µs	4492µs	1671	5516µs	4283µs	857.6	10910µs	4367µs	974.5
	396	440µs	1135µs	2801	4982µs	573.6µs	594.7	16170µs	932.3µs	1915
	24	1910µs	5187µs	9737	30380µs	3217µs	4062	51000µs	4624µs	7993
UDP/IP	992	835µs	869.3µs	647.7	1141µs	818.3µs	175.9	1578µs	833.8µs	188.7
	396	1683µs	1716µs	1139	2168µs	1601µs	379.9	3743µs	1640µs	423.7
	24	28110µs	28780µs	8311	40920µs	27400µs	4285	67530µs	28250µs	6070
IEEE-1394	992	754µs	896.5µs	1355	1106µs	728µs	124.5	6236µs	783.9µs	372.1
	396	1258µs	1444µs	1451	2394µs	1250µs	326.6	7105µs	1333µs	553.3
	24	11860µs	11950µs	4707	18850µs	11450µs	4203	23820µs	11800µs	4474

Table 1. Statistical summary of transmit lag

5. References

- Gosling James, Joy Bill, & Steele Guy (1996). The Java Language Specification, Reading, Mass: Addison-Wesley.
- Metcalf Robert M., & Boggs David R. (1976). Ethernet: Distributed Packet Switching for Local Computer Networks. Association for Computing Machinery, Vol. 19/No 5.
- Spurgeon Charles E. (2000). Ethernet: The definitive guide, : O'Reilly & Associates.
- Information Sciences Institute (1981, September). RFC 791. University of Southern California: Author. Retrieved March 5, 2001, from World Wide Web: <ftp://ftp.isi.edu/in-notes/rfc791.txt>
- Information Sciences Institute (1981, September). RFC 793. University of Southern California: Author. Retrieved March 5, 2001, from World Wide Web: <ftp://ftp.isi.edu/in-notes/rfc793.txt>
- Information Sciences Institute (1980, August). RFC 768. University of Southern California: Author J. Postel. Retrieved March 5, 2001, from World Wide Web: <ftp://ftp.isi.edu/in-notes/rfc768.txt>
- Wright Gary & Stevens W. Richard (1995). TCP/IP Illustrated: Volume 2. The Implementation, Reading, Mass: Addison-Wesley.
- Anderson, Don (1999). FireWire system architecture (2nd ed.). Reading, Mass: Addison-Wesley.
- Mills, D.L. (1996, October). Simple network time protocol (SNTP) version 4 for IPv4, IPv6 and OSI. Network Working Group Report RFC-2030. University of Delaware: Author D.L. Mills. Retrieved May 5, 2001, from World Wide Web: <http://www.eecis.udel.edu/~mills/database/rfc/rfc2030.txt>