

A Data Communication Reliability and Trustability Study for Cluster Computing

Speaker: Eduardo Colmenares
Midwestern State University
Wichita Falls, TX

Introduction

HPC

- Relevant to a variety of sciences, not only CS
- More and more researchers considering new technologies capable of considerable computational power
 - Distributed Shared Memory Systems (Clusters)
 - Combine computational capabilities of multiple nodes via a high speed communication network.
 - Clusters, if public
 - Multiple users
 - » Need policies to schedule jobs and manage resources
- We acknowledge to potential characteristics of interest
 - Performance
 - Consistent time accuracy

Ping-Pong Test

Purpose:

- To measure the end-to-end delay time associated with sending a message back and forth between processes in a cluster of workstations (or any other parallel system)
- Two variants

Both Ping-Pong Tests

Testing Environments

Community-Cluster

This cluster has 12TB of public shared Lustre storage and three groups of public and private nodes, all connected by SDR Infiniband and Gigabit Ethernet. The quad-core nodes have Infinihost III Lx (PCI-e) cards, and the older nodes have Infinihost (PCI-X) cards.

Public quad-core (512 cpu, 4.77 TF).

64 nodes with dual quad-core Intel 5345 processors (2.33 GHz) and 12GB of memory each. Designated compute-1-x, 2-x.

Public single-core (128 cpu, 0.82 TF).

64 nodes with dual single-core Intel "Irwindale" processors (3.2 GHz) and 4 GB of memory each. Designated compute-3-x, 4-x, 5-x.

Public AMD dual-core (8 cpu, .04 TF).

1 node with quad dual-core AMD 8218 processors (2.60 GHz) and 64GB of memory. Designated compute-8-1.

Multi-user and 3 queues {2WKpar, 48Hquadpar, 2WKpar}

My-Cluster

Each one of the nodes in this cluster has the following characteristics: one Intel(R) Pentium(R) 4 CPU at 1.70GHz, one 3Com PCI 3c905C Tornado network card. All the nodes in this cluster are interconnected via a 3Com® Super Stack® 3 Switch 3300 12-Port. Table 1 summarizes the major hardware differences between nodes.

NODE	MEMORY (MB)	HARD DISK
1	511.46	40020 MB-T340016A, ATA
2	1023.4	40020 MB-T340016A, ATA
3	1023.4	20547 MB-MAXTOR 6L020J1, ATA
4	511.46	40020 MB-T340016A, ATA
5	1023.4	20547 MB-MAXTOR 6L020J1, ATA
6	1023.4	20547 MB-MAXTOR 6L020J1, ATA
7	1023.4	40020 MB-WDC WD400BB-75DEA0, ATA
8	1023.4	40027 MB-MAXTOR 6L040J2, ATA
9	1023.4	40027 MB-MAXTOR 6L040J2, ATA

1 user

Experimental results

Ping-Pong-A


Community-Cluster (queue: 2WKpar)

Community-Cluster - 2WKpar - 9P - 1st Sample Included				
Proc-Pair	Average (usecs)	Stdv (usecs)	Min (usecs)	Max (usecs)
P0 - P1	1218.29	12091.28	8	120922
P0 - P2	751.8	7424.667	8	74256
P0 - P3	598.89	5895.264	8	58962
P0 - P4	496.8	4885.778	7	48866
P0 - P5	599.54	5911.764	7	59126
P0 - P6	600.08	5905.447	8	59064
P0 - P7	597.5	5893.182	7	58940
P0 - P8	599.12	5908.675	7	59095

Experimental results

Ping-Pong-A

Community-Cluster (queue: 2WKpar)



Experimental results

Ping-Pong-A


Community-Cluster (queue: 48Hpar)

Community-Cluster - 48Hpar - 9P - 1st Sample Removed				
Proc-Pair	Average (usecs)	Stdv (usecs)	Min (usecs)	Max (usecs)
P0 - P1	12.80808	2.961235	10	27
P0 - P2	13.86869	2.448017	12	22
P0 - P3	13.49495	2.800779	11	26
P0 - P4	12.31313	2.648243	10	24
P0 - P5	13.88889	2.754917	11	27
P0 - P6	13.48485	2.708127	11	26
P0 - P7	13.52525	2.749187	11	25
P0 - P8	13.50505	2.588726	11	24

Experimental results

Ping-Pong-A

Community-Cluster (queue: 48Hpar)



Experimental results

Ping-Pong-A


Community-Cluster (queue: 48Hquadpar)

Community-Cluster - 48Hquadpar-9P-1 st Sample Removed				
Proc-Pair	Average (usecs)	Stdv (usecs)	Min (usecs)	Max (usecs)
P0 - P1	10.38384	2.141462	8	19
P0 - P2	10.25253	1.745802	8	21
P0 - P3	10.54545	1.540426	9	16
P0 - P4	9.373737	1.7237	7	18
P0 - P5	16.66667	53.26733	8	526
P0 - P6	10.14141	1.51866	8	18
P0 - P7	9.282828	1.450126	8	15
P0 - P8	9.363636	1.548235	8	16

Experimental results

Ping-Pong-A

Community-Cluster (queue: 48Hquadpar)



Experimental results

Ping-Pong-A


My-Cluster

My-Cluster - 9P - 1 st SAMPLE INCLUDED				
Proc-Pair	Average (usecs)	Stdv (usecs)	Min (usecs)	Max (usecs)
P0 - P1	203.41	693.2721	127	7064
P0 - P2	211.21	671.9197	126	6845
P0 - P3	200.24	617.0202	114	6074
P0 - P4	196.65	594.7444	115	6001
P0 - P5	222.55	626.2541	118	6146
P0 - P6	203.47	605.3599	117	5977
P0 - P7	213.4	604.9634	117	6010
P0 - P8	200.64	638.3008	116	6401

Experimental results

Ping-Pong-A


My-Cluster



Experimental results

Ping-Pong-A

Let's compare



Experimental results

Ping-Pong-B


Community-Cluster (queue: 2WKpar)

Community-Cluster- 2WKpar - 9P - 1 st SAMPLE INCLUDED				
Proc-Pair	Average (usecs)	Stdv (usecs)	Min (usecs)	Max (usecs)
P0 - P1	565.79	5562.042	8	55630
P0 - P2	508.85	4995.571	8	49965
P0 - P3	511.28	4997.346	10	49985
P0 - P4	509.17	4989.175	9	49902
P0 - P5	513.83	5007.899	9	50092
P0 - P6	509.03	4997.472	8	49984
P0 - P7	524.87	4995.444	9	49972
P0 - P8	510.24	4999.774	9	50008

Experimental results

Ping-Pong-B

Community-Cluster (queue: 2WKpar)



Experimental results

Ping-Pong-B

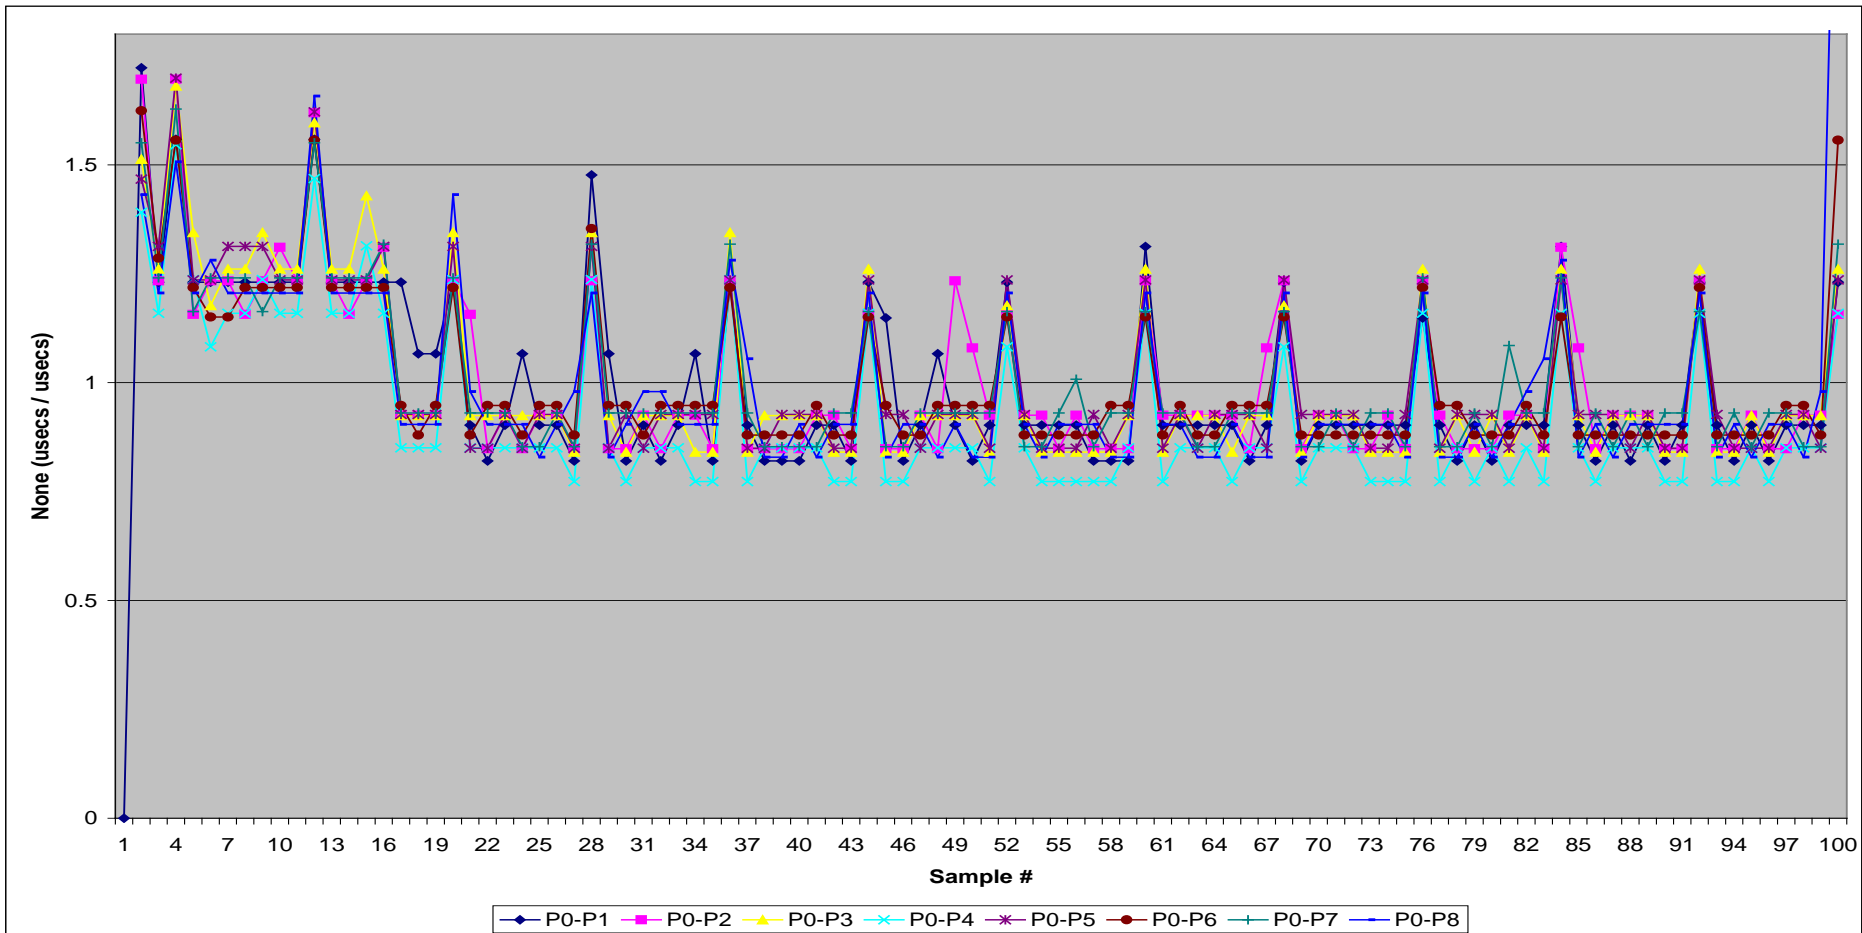
Community-Cluster (queue: 48Hpar)

Community-Cluster - 48Hpar - 9P - 1st Sample Removed				
Proc-Pair	Average (usecs)	Stdv (usecs)	Min (usecs)	Max (usecs)
P0 - P1	12.19192	2.414567894	10	21
P0 - P2	12.9697	2.50083474	11	22
P0 - P3	11.89899	2.384023993	10	20
P0 - P4	12.93939	2.376836269	11	22
P0 - P5	12.94949	2.480017023	11	22
P0 - P6	14.77778	2.593380725	13	24
P0 - P7	12.89899	2.296825447	11	21
P0 - P8	13.27273	3.24766221	11	35

Experimental results

Ping-Pong-B

Community-Cluster (queue: 48Hpar)



Experimental results

Ping-Pong-B


Community-Cluster (queue: 48Hquadpar)

Community-Cluster - 48Hquadpar - 9P - 1st Sample Removed				
Proc-Pair	Average (usecs)	Stdv (usecs)	Min (usecs)	Max (usecs)
P0 - P1	9.343434	1.691283875	8	19
P0 - P2	9.464646	1.547302665	8	21
P0 - P3	10.90909	2.321553443	9	27
P0 - P4	10.36364	1.438923531	9	22
P0 - P5	10.61616	1.838960835	9	22
P0 - P6	10.47475	1.053114277	9	15
P0 - P7	9.353535	0.872635362	8	13
P0 - P8	9.676768	2.668135032	8	28

Experimental results

Ping-Pong-B

Community-Cluster (queue: 48Hquadpar)



Experimental results

Ping-Pong-B


My-Cluster

My-Cluster - 9P - 1st Sample Removed				
Proc-Pair	Average (usecs)	Stdv (usecs)	Min (usecs)	Max (usecs)
P0 - P1	155.5657	54.73240312	133	583
P0 - P2	199.4545	410.5913921	138	4221
P0 - P3	214.9293	788.5313299	122	7978
P0 - P4	261.9697	1216.547821	122	12236
P0 - P5	174.8283	272.2759345	124	2767
P0 - P6	161.3636	233.9981011	122	2449
P0 - P7	202.4646	64.69146762	121	358
P0 - P8	129.4747	19.3086337	119	274

Experimental results

Ping-Pong-B


My-Cluster




Experimental results

Ping-Pong-B

Let's compare



Was all this work worth it?



Conclusions

All three queues of the Community-Cluster are prone to pulses of some nature that can negatively impact high performance applications where consecutive steady and accurate time readings are needed to statistically validate a phenomena under study.

Thank You

Questions ?

