

Modeling of NAMD's Network Input/Output on Large PC Clusters*

Nancy Tran¹ and Daniel A. Reed²

¹University of North Carolina
Dept. of Pathology & Laboratory Medicine
Chapel Hill, NC 27599 USA
nancytran@unc.edu

²University of North Carolina
Renaissance Computing Institute
Chapel Hill, NC 27599 USA
dan_reed@unc.edu

Abstract

This study examined the interplay among processor speed, cluster interconnect and file I/O, using parallel applications to quantify interactions. We focused on a common case where multiple compute nodes communicate with a single master node for file accesses. We constructed a predictive model that used time characteristics critical for application performance to estimate the number of nodes beyond which further performance improvement became unattainable. Predictions were experimentally validated with NAMD [12, 14], a representative parallel application designed for molecular dynamics simulation. Such predictions can help guide decision making to improve machine allocations for parallel codes in large clusters.

1. Introduction

Many computationally demanding scientific applications are being ported to commodity clusters to exploit the broad availability of inexpensive commodity components. However, performance rarely scales linearly with the number of processors, due to a combination of algorithmic and processor limitations, communication bandwidth constraints, and bounded I/O parallelism. I/O parallelism is a particularly pernicious constraint, as many applications rely on a "single I/O master, multiple compute nodes" paradigm where all worker nodes converge onto a single master node for file I/O. Although parallel file systems (e.g., GPFS and Lustre) are now emerging, applications designed for portability normally eschew dependence on a parallel file system in favor of application managed I/O. For these applications, cluster

architectures that distribute I/O nodes in proportion to compute nodes (e.g., one I/O server for every twenty compute nodes) cannot help improve performance.

In addition, determining node allocations that can maximize both application performance and job throughput remains an outstanding problem for large clusters. For example, when improvement in application performance from a 256-node allocation is small compared to a 128-node allocation, how many additional nodes beyond 128 but far less than 256 should be used? The remaining nodes can be released to other jobs for throughput increase.

To address the above issues, we constructed a performance model that analyzed the overlapping interplay among three critical characteristics – computation time used to solve a target problem, end-to-end data transmission delay, and file I/O duration. We applied the model to predict an estimate on node allocations beyond which further performance improvement could not be achieved. Finally, we validated model predictions via a case study using the NCSA (National Center for Supercomputing Application) TeraGrid cluster and NAMD [12], a representative parallel application.

The remainder of this paper is organized as follows. We begin in Section 2 by constructing a model to examine the interactions among different time parameters. Section 3 describes the characteristics of the NAMD code and experiments with the APOA1 benchmark. In Section 4, we present experimental results and model validation. Section 5 describes related work and Section 6 concludes this paper with ideas for future directions.

2. The Model

We focused on a prevalent I/O architecture in which one I/O master serves I/O requests from all compute nodes in a cluster, as shown in Figure 1. This "single I/O master, multiple compute nodes" paradigm, despite its scalability

*This work was supported in part by the National Science Foundation under grants NSF ACI 01-22296, NARA NSF 02-02 GPG, the State of Illinois, and the University of North Carolina at Chapel Hill. Experiments and analysis were conducted at NCSA.

1-4244-0910-1/07/\$20.00 ©2007 IEEE.

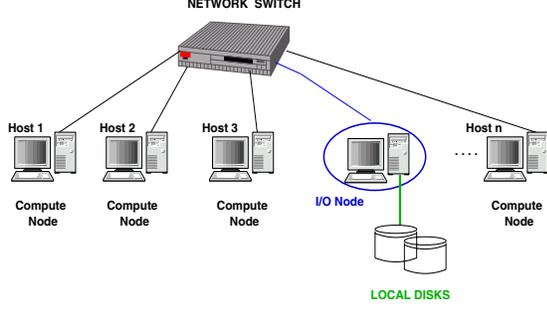


Figure 1. I/O control via a single master node.

problem, is being adopted by a large class of scientific parallel applications including NAMD – a molecular dynamics simulation code [12] at the University of Illinois, ESCAT – an electron scattering code [16] at Caltech, and CACTUS – a gravitational physics code [2] at the Max Planck Institute.

For this I/O paradigm, applications commonly use a master task to receive and assemble data packets sent by various compute nodes. The final assembled data are written to local disks without further network intervention. Parallel computations are overlapped with parallel data transfers to reduce execution times. Below, we present a model that captures and examines this overlapping interplay.

2.1. Model Parameters and Construction

Figure 2 illustrates key components used to model overlapping interactions between the I/O master and a representative compute node, assuming homogeneous nodes. Defining a transaction as a set of operations used to produce a set of results, our model identifies four metrics per transaction:

1. *Total computation time* which includes result computation time and communication software overhead for transmitting results in messages to the I/O master.
2. *Message transmission delay* which is composed of all the delays incurred in transmitting a message across the network (i.e., propagation latency + packet delay + router delay + any other delay) and processing overhead incurred at the I/O master for message reception and assembly. In large clusters, such overhead increases with the number of messages sent by compute nodes.
3. *Duration of file I/O* at the I/O master’s local disks.
4. *Computation improvement ratio* – if an application distributes its workload uniformly and effectively, we can assume that computation time of a representative compute node decreases exponentially when the number of nodes doubles [4]. However, in general, this improvement ratio is below 2:1 depending on applications’ implementations.

Our model is based on a known observation that *when the total computation time is smaller than message trans-*

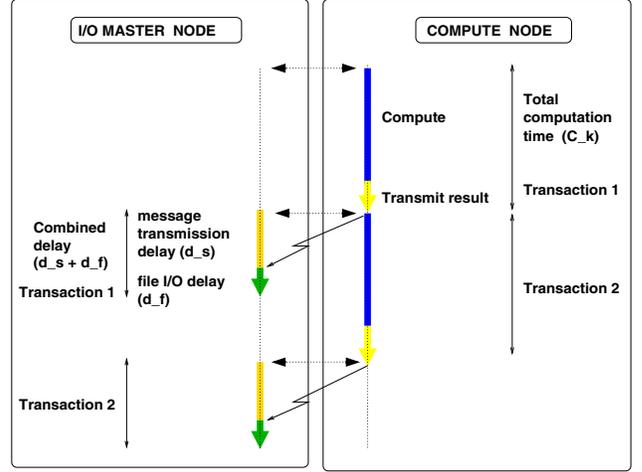


Figure 2. Interactions between I/O master and a compute node in a k -node configuration.

mission delay, it is no longer beneficial to further divide a workload. Denoting C_k as the total computation time at a compute node in a k -node configuration, and R as the associated improvement ratio, the computation time C_n for a n -node configuration can be derived as:

$$C_n = \frac{C_1}{R^{\log_2(n)}} = \frac{C_k \times R^{\log_2(k)}}{R^{\log_2(n)}} \quad (1)$$

$$C_n = C_k \times R^{\log_2(k/n)}$$

Similarly, if message transmission delay also improves exponentially when doubling the number of nodes, but at a smaller ratio (r), the delay for a n -node configuration is:

$$D_n = \frac{D_1}{r^{\log_2(n)}}$$

Comparing computation time with message transmission delay and assuming $R > r$, we have:

$$\frac{C_1}{R^{\log_2(n)}} \leq \frac{D_1}{r^{\log_2(n)}}$$

$$n \geq 2^{\frac{\log_e(C_1/D_1)}{\log_e(R/r)}} \quad (2)$$

Equation 2 provides an estimate on the number of nodes needed if the improvement ratio r for transmission delay is known. However, depending on the job mix and network traffic in a cluster, transmission delays can be highly variable, and r often does not scale with the number of nodes. Instead of r , we propose using the actual message transmission delay d_s and file I/O duration d_f to predict the number of nodes n :

$$\frac{C_1}{R^{\log_2(n)}} \leq d_s + d_f$$

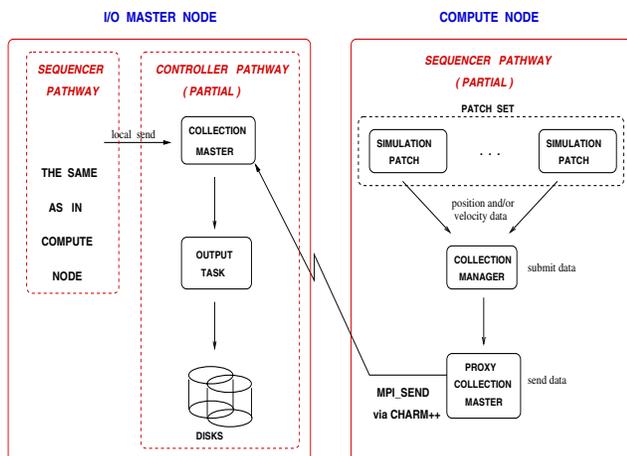


Figure 3. Data flow from a compute node to the I/O master.

$$\log_2(n) \geq \frac{\log_e(C_1/d_s + d_f)}{\log_e(R)}$$

$$n \geq 2^{\frac{\log_e(C_1/d_s + d_f)}{\log_e(R)}} \quad (3)$$

3. Experiments with NAMD

NAMD is a molecular dynamics code designed for the simulation of molecules on high performance computing platforms. NAMD allows researchers to investigate various aspects associated with bio-macromolecules – effects of temperatures and forces, interactions among proteins, lipids, water molecules, and nucleic acids. It achieves good performance and scalability via several mechanisms, the most notable of which are:

- Spatial decomposition and parallelization of a simulation problem into patches (representing some units of work) which are distributed to different processors for calculations. To improve communication locality, neighboring patches are assigned to the same processor. To achieve fast calculations of full electrostatics interactions, NAMD uses a parallel version of the particle mesh Ewald method [7].
- Overlapping of computation with network communication. NAMD uses the Charm++/Converse object-based parallel libraries and runtime systems [11] to support non-blocking message communication and load balancing [5] among processors.

3.1. Target Components

This study focused on a portion of NAMD 2.5’s architecture that supports collection of simulation results, trans-

mission and output of results to disks for future visualization. Figure 3 highlights the architectural components we targeted and the flow of simulation results between these components, from a compute node to the I/O master. Each node may be equipped with a single or multiple processors.

As noted earlier, simulation is divided into patches and distributed to various compute nodes for execution. Depending on the grid size mapped to a simulation problem, there can be more than one patch assigned to a processor, if the number of patches exceeds the number of processors.

Activities in a *compute node* are coordinated by a sequencer object thread. In the sequencer pathway, partially shown in Figure 3, a collection manager object gathers and integrates simulation results on atom coordinate positions and/or velocities calculated by various patches. Results are then packaged into two separate messages (one for positions, the other for velocities) to be MPI-sent asynchronously by a proxy collection master to the I/O master.

In the *I/O master*, NAMD uses a controller object thread to manage data flow. The collection master of the controller receives, queues, and assembles messages into buffers which are finally written to two separate files – position and velocity. If there are more patches than processors, the I/O master will also participate in simulation computation via a sequencer object. Simulation results are sent locally to the controller, trading network bandwidth for processor cycles.

3.2. Experiment Benchmark

We experimented with the APOA1 benchmark provided by NAMD. APOA1 models the dynamics of a high density lipoprotein, known to protect against the accumulation of platelets in blood vessels [1]. The benchmark simulates over 92K atoms of proteins, lipids and water. During simulation start-up, about 7 MB of position data, 13 MB of atom structure information (bonds, angles, hydrogen donors and acceptors, etc.) and three small parameter files are sequentially fetched into memory by the I/O master.

Based on the above inputs, NAMD spatially decomposes APOA1 into a $6 \times 6 \times 4$ grid, creating 144 patches to be uniformly distributed among cluster nodes. For a n -node cluster, each node receives $144/n$ neighboring patches. Each patch has ≈ 640 atoms ($92K/144$) on average.

We ran APOA1 for 5000 time steps, varying node configurations from 4 to 255 nodes, doubling the number of nodes each time. For each configuration, the benchmark was executed three times and measurements were averaged over the three experiments. Position and velocity data were written to files every four time steps, resulting in 1250 transactions. We focused on file writes that occurred after startup because they represented the major bulk of I/O in APOA1.

Nodes	Execution Time (sec)	Improvement Ratio
4	3661	–
8	1994	1.84
16	1017	1.96
32	586	1.74
64	313	1.87
128	192	1.63
255	146	1.32

Table 1. APOA1’s overall execution times.

3.3. Computing Platform

The APOA1 benchmark was run on the NCSA Teragrid cluster (phase I). At the time of our experiments, the cluster had 256 nodes of Itanium-2, each equipped with 1.3 GHz dual processors of 64-bit addressing and 3 MB L3 cache/processor. Half of the nodes had 14 GB of main memory; the other half 12 GB. Jobs were submitted via the PBS scheduling system [10] with machine allocations pre-specified by users.

The cluster was interconnected via Myrinet and Gigabit Ethernet. Each node was installed with a Myrinet 2000-fiber/PCI interface card, linked to a Myrinet switch via 2Gb/sec fiber cables. It also had two internal SCSI disks providing 18 and 73 GB of local storage respectively. Each disk, with an 8 MB buffer cache, transferred at a peak external burst data rate of 320 MB/sec and a maximum sustained rate (disk surface transfer rate) of 98 MB/sec.

4. Experimental Results and Analysis

4.1. Overall Execution Time

NAMD achieved significant performance improvement for applications executing on clusters via workload partitioning, overlapping computation with network I/O and file I/O. Table 1 summarizes APOA1’s total execution times for various node configurations. Improvement ratios between consecutive execution times averaged $\approx 1.8:1$ when the number of nodes was doubled from 4 to 64. However, the speedup was smaller for configurations of 128 nodes and beyond.

In the next section, we will present measurements of APOA1’s time parameters – computation time, message transmission delay, and file write duration – followed by analysis and validation of model prediction.

4.2. Computation Time of Simulation

We timestamped messages containing simulation results before they were sent by sequencers to the I/O master. Timestamp differences between two consecutive messages

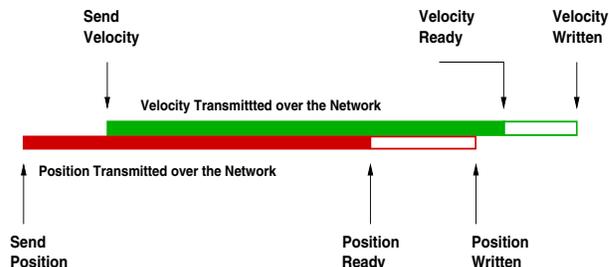


Figure 4. Overlap of data transmissions and file writes.

represented a transaction’s (four time steps) simulation time and message transmission overhead. Mean computation times, shown in Column 2 of Table 2, were averaged over three runs. Computation improvement ratios, displayed in Column 3, were comparable to those obtained for APOA1’s overall execution times (Table 1).

Because NAMD partitioned APOA1’s simulation into 144 patches, there were more patches than nodes for configurations ranging from 4 to 128 nodes. For example, 36 patches were assigned to each node in a 4-node configuration, including the I/O master. In contrast, any configuration beyond 145 nodes (144 compute nodes + 1 I/O master) has more machines than needed with one patch per machine. In a 255-node configuration, only 145 nodes were used, explaining the smaller improvement ratios observed.

4.3. Message Size and Transmission Delay

Table 3 shows that, for configurations of 4 to 128 nodes, doubling the number of nodes reduced message sizes for both position and velocity by a ratio of $\approx 2:1$. However, the 255-node configuration had a ratio of $\approx 1:1$ due to node over-allocation.

Message transmission delay measured the interval between the time a message was sent and the time when it was ready for output to disks. Figure 4 illustrates the over-

Nodes	Position Mean Msg Size (KB)	Ratio	Velocity Mean Msg Size (KB)	Ratio	Combined Mean Msg Size (KB)
4	360.9		631.5		992.4
8	180.4	2	315.7	2	496.1
16	80.7	2.2	141.1	2.2	221.8
32	40.3	2	70.5	2	110.8
64	19.8	2	34.6	2	54.4
128	10.1	1.9	17.7	1.9	27.8
255	9.7	1	17.0	1	26.7

Table 3. Message sizes.

Nodes	Mean Computation Time (sec)	Computation Improvement Ratio	Mean Aggregate Message Transmission Delay (sec)	Mean Aggregate Write Duration (sec)	Mean Combined Delay (sec)	Combined Delay / Computation Time (%)
k	C_k	R				
4	2.918	–	0.092	0.029	0.103	3.5
8	1.585	1.84	0.187	0.024	0.198	12.5
16	0.804	1.97	0.132	0.025	0.143	17.8
32	0.460	1.75	0.074	0.031	0.088	19.2
64	0.240	1.92	0.060	0.022	0.072	30.0
128	0.141	1.70	0.048	0.023	0.059	41.8
255	0.104	1.36	0.061	0.022	0.073	70.2

Table 2. Time characteristics of APOA1’s transactions.

lapping transmissions of position and velocity messages for which two files were simultaneously served by the I/O master. Based on this observation, we computed *mean aggregate message transmission delays* using: $velocity\text{-ready timestamp} - send\text{-position timestamp}$. Column 4 of Table 2 displays these delays for APOA1, averaged over three runs.

Figure 7 presents a 3D plot associating mean aggregate delays with node configurations and the combined sizes of position and velocity messages. Overall, these delays did not reduce commensurately with message sizes. To better understand this unexpected behavior, we examined the actual transmission delays and their distributions for entire executions of APOA1, shown in Figures 5 and 6 respectively.

In configurations of 128 and 255 nodes, message transmission delays, aggregated for position and velocity, were mostly uniform with the exception of a few outliers. This uniformity could be attributed to the absence of perturbations from other applications as the entire cluster was dedicated to APOA1. In contrast, transmission delays in smaller configurations were highly variable and bi-modally distributed. Two possible causes may have contributed to this fluctuating behavior: a) perturbations introduced by other jobs that competed for network resources causing routing delays and b) context switchings between computation of simulation and message processing at the I/O master.

4.4. File Write Duration

During a transaction, APOA1 wrote 1.05 MB of simulation data (3 coordinates for x,y,z \times 4 bytes/coordinate \times 92224 atoms) to each of the position and velocity files, producing \approx 1.3 GB per file (1250 transactions \times 1.05 MB). According to Figure 4 which depicts overlapping interactions between file I/O and data transmissions, APOA1’s combined write durations were calculated as:

$$\begin{aligned} &(\text{position-written timestamp} - \text{position-ready timestamp}) + \\ &(\text{velocity-written timestamp} - \text{velocity-ready timestamp}) \quad (4) \end{aligned}$$

Column 5 of Table 2 presents APOA1’s mean aggregate write durations which varied between 22 to 31 msec for 2.1 MB (2×1.05) of data. As result, write throughput was \approx 70-85 MB/sec for configurations of 4 to 32 nodes and 95 MB/sec for others. The higher throughput observed with large configurations was due to fewer simulation patches assigned to the I/O master, releasing processor cycles for file writes.

In a n-node configuration where only (n-1) nodes sent messages over the network to the I/O master,

$$\text{write efficiency} = \frac{1.05 \text{ (MB)}}{(n - 1) \times \text{message size (MB)}} \times 100$$

Using message sizes compiled in Table 3, the average write efficiency was \approx 85% for position data, but only 50% for velocity data. The highest efficiency, 99% for position, 57% for velocity, was associated with the 4-node configuration which had the largest message sizes. Inevitably, smaller messages generated by larger configurations required more overhead for message packaging and metadata processing.

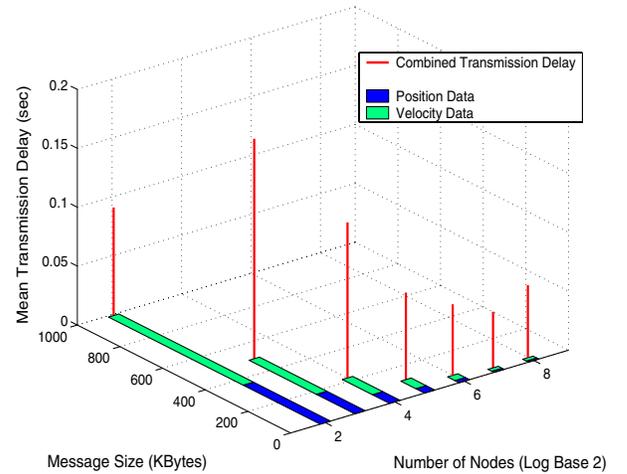


Figure 7. Mean aggregate transmission delays.

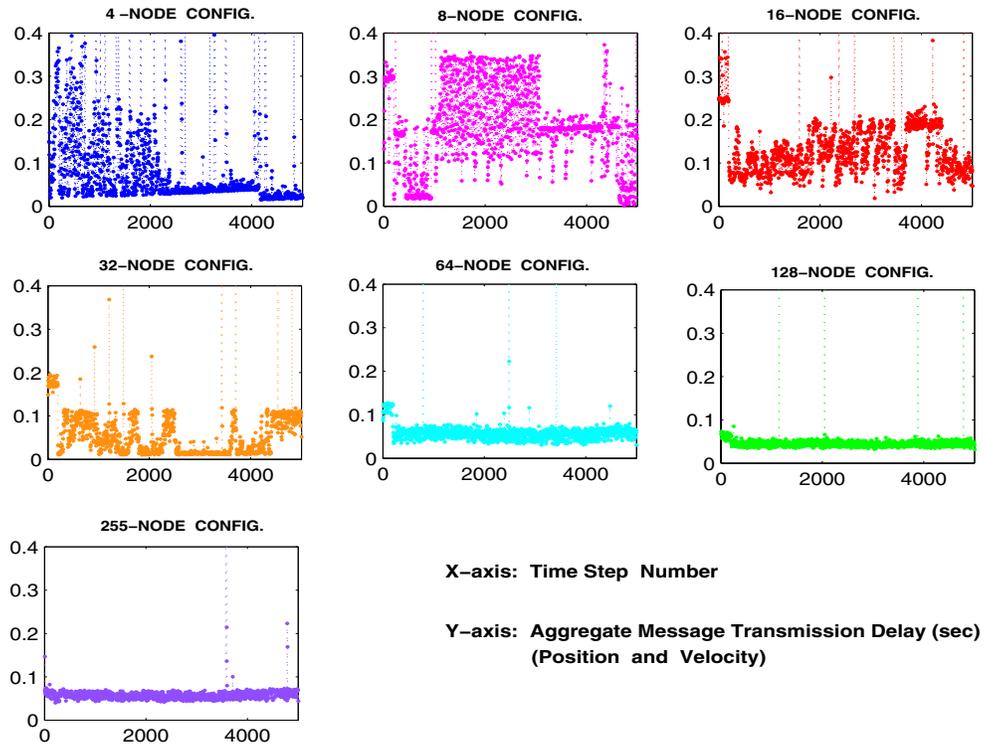


Figure 5. Aggregate message transmission delay behaviors.

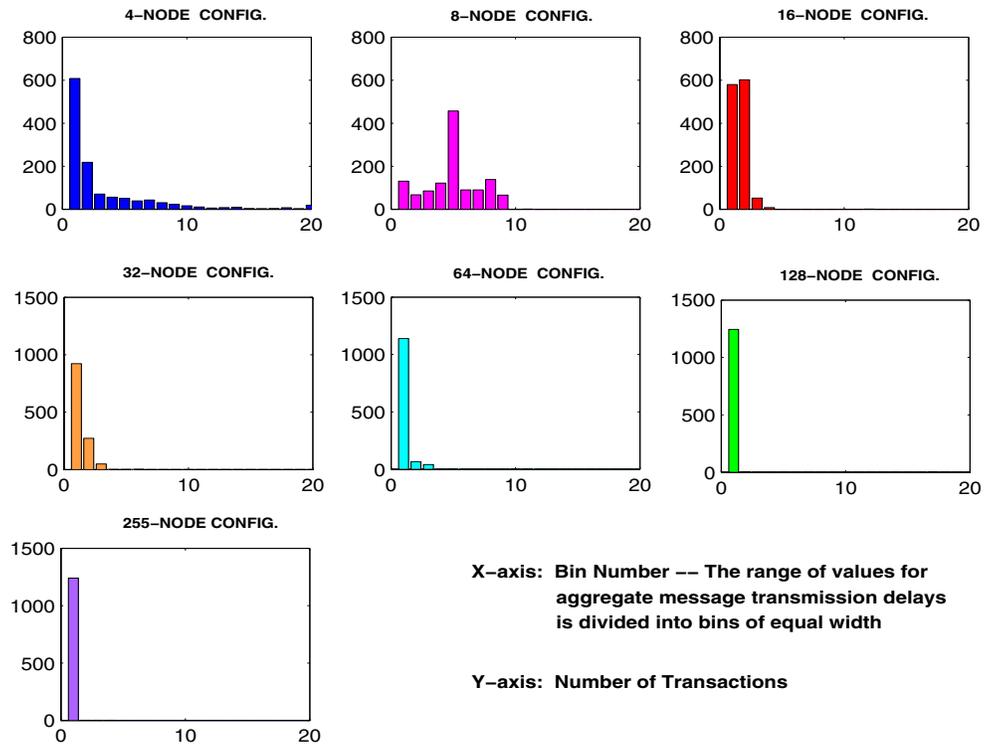


Figure 6. Distributions (histograms) of aggregate message transmission delays.

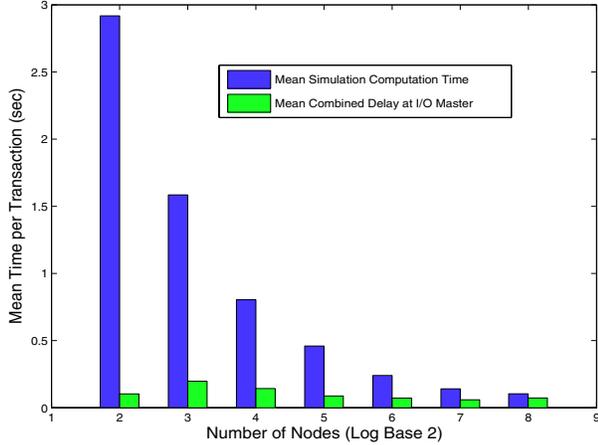


Figure 8. Simulation computation times compared with combined delays.

4.5. Combined Message Delay

Combining file write duration with message transmission delay gives the total delay. Again, Figure 4 provides the logic for computing a transaction’s total delay. Column 6 of Table 2 summarizes measurements of mean combined delays which, in general, followed a trend set by transmission delays because the latter were at least twice longer than file writes.

$$\text{combined message delay} = \text{velocity-written timestamp} - \text{send-position timestamp} \quad (5)$$

Comparing computation times with combined delays (Columns 2 and 6, Table 2) indicated that, as the number of nodes increased, total message delays constituted a larger fraction of computation times – up to 70% for 255 nodes (Column 7, Table 2). Figure 8 reaffirmed that when message delays exceeded computation times, further partitioning a workload would not improve performance.

Our measurements also indicated that the ratios (r) between consecutive mean combined delays did not improve exponentially – a characteristic that our model has considered. In fact, r was only 1.22 when the number of nodes was increased from 32 to 64, and from 64 to 128.

4.6. Model Prediction Validation

To validate node predictions generated by equation 3, values for model parameters were retrieved from Table 2. To avoid data skewing, the 255-node configuration was not selected because it did not reflect the general behavior of APOA1 relative to other configurations.

Nodes	1-Node Equivalent Computation Time (sec)
k	$C_k \times R^{\log_2 k} = C_1$
4	$2.918 \times 1.836^2 = 9.836$
8	$1.585 \times 1.836^3 = 9.809$
16	$0.804 \times 1.836^4 = 9.136$
32	$0.460 \times 1.836^5 = 9.597$
64	$0.240 \times 1.836^6 = 9.193$
128	$0.141 \times 1.836^7 = 9.916$
	$mean(C_1) = 9.581$

Table 4. Computation times equivalent to one-node configuration.

For configurations of 4 to 128 nodes, computation times had a mean improvement ratio $R = 1.836$ along with a mean value for $C_1 = 9.581$, calculated in Table 4 from times equivalent to one-node configuration via equation 1. The mean combined delay averaged across configurations was 0.110. Thus, an estimate for the number of nodes required by APOA1 was predicted as:

$$n = 2^{\frac{\log_e(9.581/0.110)}{\log_e(1.836)}} \approx 2^{7.35} \approx 163$$

As mentioned earlier, APOA1 had only 144 computation patches and the number of nodes to be allocated should be 145 (144 compute nodes + 1 node for I/O). Comparing our predicted estimate with 145 yields a prediction accuracy $\approx 87\%$ ($[1 - \frac{163-145}{145}] \times 100$). Furthermore, allocating 163 nodes instead of 255 would make available over 90 nodes for other jobs in the cluster.

5. Related Work

There have been many studies on analysis and modeling of parallel systems for performance improvement on parallel machines. Early works focused mostly on processor allocation, utilization, and performance [3, 9, 15]. Later, network communication was recognized to be a major performance limiting factor in parallel and distributed systems [8].

Cremonesi et al. [6] used single class queuing networks to model application performance on parallel machines and clusters of workstations. Similar to many extant queuing models, their studies assumed that communication delays were exponentially distributed. In practice, this assumption may not be valid for many applications, including NAMD (see Figure 6).

Sevcik examined static (non-preemptive) scheduling strategies [15] that made processor allocation decisions based on average parallelism ratio – execution time on a single processor over time on an unlimited number of processors. Allocation can be further fine-tuned by adjusting

this ratio to the overall system load and the number of processors. Our model extended this concept by examining interactions between execution time improvement ratios, network communication, and I/O delays.

Natarajan and Iyer [13] presented a measurement and simulation study to characterize performance bottlenecks and throughput limits in configurations with multiple I/O nodes and compute nodes. Their analysis compared I/O performance using remote I/O nodes to that using disks local to the compute nodes. In contrast, our study investigated I/O issued from multiple compute nodes to a single I/O master.

6. Conclusions and Future Directions

We have designed a model that analyzed the temporal interplay between computation, data communication, and network I/O for parallel applications running on large PC clusters. No assumption of exponentially distributed communication delays was needed. It has been known that there is a limit to the advantages of partitioning a computation load among cluster nodes. Our model exploited this knowledge to provide a closed form equation to predict node allocations for a large class of parallel codes that use a master node for aggregating file I/O. Model predictions were validated with NAMD resulting in a prediction accuracy $\approx 85\%$.

Our experiments with NAMD showed that, to avoid erroneous counting, special attention must be given to measuring model parameters when interactions among these parameters overlap. Experimental results have demonstrated that overly aggressive workload distribution will transform a computationally intensive task into one that is network and I/O intensive.

Future research directions include validation of the model against other parallel scientific applications and incorporation of model predictions into performance monitoring and/or instrumentation software. The goal for the latter effort is to provide guidance for automatic node allocations on large clusters.

7. Acknowledgments

Thanks to the NCSA Teragrid support staff and Andy Loftus for his efforts in providing valuable information on the Teragrid cluster configuration. We appreciate the help of Rick Kufirin, James Phillips, and Zhenbing Zhang for their technical expertise in porting NAMD to the cluster. NAMD was developed by the Theoretical and Computational Biophysics Group in the Beckman Institute for Advanced Science and Technology at the University of Illinois at Urbana-Champaign.

References

- [1] <http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?db=gene&cmd=search&term=apoa1>.
- [2] G. Allen, T. Goodale, J. Masso, and E. Seidel. The Cactus computational toolkit and using distributed computing to collide neutron stars. In *Proceedings of the Eighth IEEE International Symposium on High Performance Distributed Computing*, 1999.
- [3] G. Amdahl. Validity of the single processor approach to achieving large scale computing capabilities. In *Proceedings of the AFIPS Conference*, volume 30, pages 483–485, 1967.
- [4] S. Arvindam, V. Kumar, and V. Rao. Efficient parallel algorithms for search problems: applications in VLSI CAD. In *Frontiers of Massively Parallel Computation*, pages 166–169, 1990.
- [5] R. Brunner and L. Kale. Adapting to load on workstation clusters. In *The Seventh Symposium on the Frontiers of Massively Parallel Computation*, pages 106–112, 1999.
- [6] P. Cremonesi and C. Gennaro. Integrated performance models for SPMD applications and MIMD architectures. *IEEE Transactions on Parallel and Distributed Systems*, 13(7):745–757, July 2002.
- [7] T. Darden, D. York, and L. Pedersen. Particle mesh Ewald: An $N \cdot \log(N)$ method for Ewald sums in large systems. *Journal of Chemical Physics*, 98(12):10089–10092, June 1993.
- [8] C. Douglas, T. Mattson, and M. Schultz. Parallel programming systems for workstation clusters. In *Technical Report YALEU/DCS/TR-975, Yale University Department of Computer Science Research*, 1993.
- [9] J. Gustafson. Reevaluating Amdahl’s law. In *Communications of the ACM*, volume 31, pages 532–533, 1988.
- [10] R. Henderson. Job scheduling under the portable batch system. In *Proceedings of the Workshop on Job Scheduling Strategies for Parallel Processing*, volume 949 of LNCS, pages 279–294. Springer-Verlag, 1995.
- [11] L. Kale and S. Krishnan. Charm++, parallel programming with message-driven objects. In *Parallel Programming Using C++*, pages 175–213. MIT Press, 1996.
- [12] L. Kale, R. Skeel, M. Bhandarkar, R. Brunner, A. Gursoy, N. Krawetz, J. Phillips, A. Shinozaki, K. Varadarajan, and K. Schulten. NAMD2: Greater scalability for parallel molecular dynamics. In *Journal of Computational Physics*, volume 151, pages 283–312, 1999.
- [13] C. Natarajan and R. Iyer. Measurement and simulation based performance analysis of parallel I/O in a high-performance cluster system. In *Proceedings of the 8th IEEE Symposium on Parallel and Distributed Processing*, pages 332–339, 1996.
- [14] J. Phillips, Z. Zhang, S. Kumar, and L. Kale. NAMD: Biomolecular simulation on thousands of processors. In *Proceedings of the 2002 ACM/IEEE Conference on Supercomputing*, pages 1–18, 2002.
- [15] K. Sevcik. Characterizations of parallelism in applications and their use in scheduling. In *Proceedings of the ACM Metrics and Performance Conference*, pages 171–180, 1989.
- [16] C. Winstead and V. McCoy. Studies of electron-molecule collisions on massively parallel computers. In *Modern Electronic Structure Theory*, volume 2. World Scientific, 1994.