Did My Professor Waste Money?:
Computational Cluster Configuration Variations and Cost-Efficiency

Christian Garcia

Abstract

Parallel and distributed computing (PDC) and high performance computing (HPC) tools and techniques are becoming increasingly common and even necessary in many disciplines, particularly as large data sets and high computational loads become commonly encountered challenges. Experimentation in this space is somewhat constrained by the cost and availability of the assets that handle even small and midsize tasks in this space; initial tests and bench-marking may be the only chance to put systems through their paces. Therefore, it has become important to provide information on optimizing these smaller but still impactful systems. This requires experiments focused on delivering more capability with commodity options and configurations.

A series of HPC tools were identified for testing against commonly used benchmarks that use parallelization techniques such as MPI and OpenMP. For the selection of each benchmark, networking speed was considered above CPU power as this may be a road to increased performance at minimal additional cost. The configuration used for this experiment consists of a computing cluster with three AMD Ryzen 9 5950 CPUs and one AMD Ryzen 9 7950X CPU, all readily available. Variations in performance were investigated using the selected benchmarks: specific tests were done with both homogeneous and heterogeneous CPU configurations, as well bonded and unique-IP configurations for 10GbE network adapters.

Background

The high performance computing cluster tested throughout this study provides a valuable resource to those who use them. These machines are used by departments around campus such as Data Science, Physics, Mathematics, Engineering, and Astronomy to name a few. Therefore, performance of applications run by these users is essential. To test configurations for performance, benchmarks were carefully selected to determine if these users could take advantage of 10GbE network adapters to increase performance and efficiency.

Methodology

Step 1: Benchmarks were carefully selected considering both stressing out the CPU but more importantly network bandwidth.
Step 2: Once the benchmarks were selected, it was important to determine how to compile and run them due to some benchmarks being quite complex.
Step 3: An initial test run of each benchmark was conducted to determine which data would be a good indicator to assist in solving our questions.
Step 4: An Excel spreadsheet was created to store all the collected data.
Step 5: Each benchmark was run twice to determine if there was any variation between using the 10GbE network adapter and not using it.
Step 6: A graphical representation of the data was created to easily parse the results of the research. (See figures 1-3 below)

Results

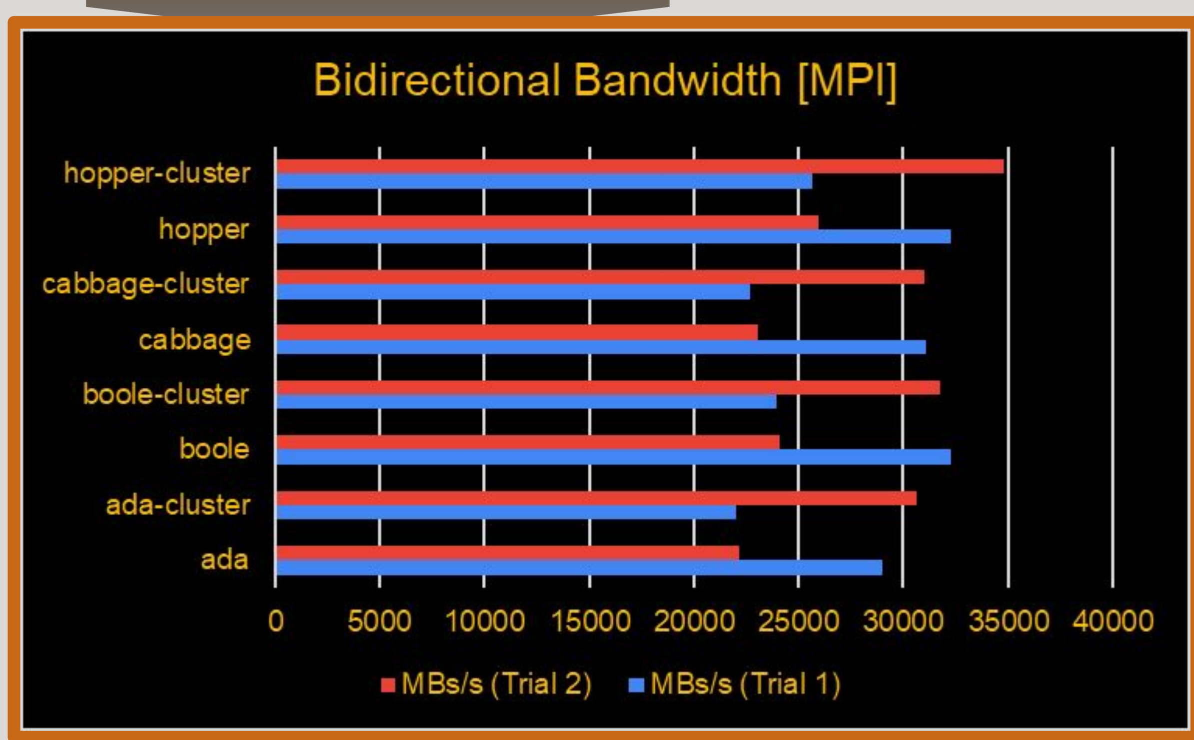

Figure 1:
Bidirectional Bandwidth Benchmark

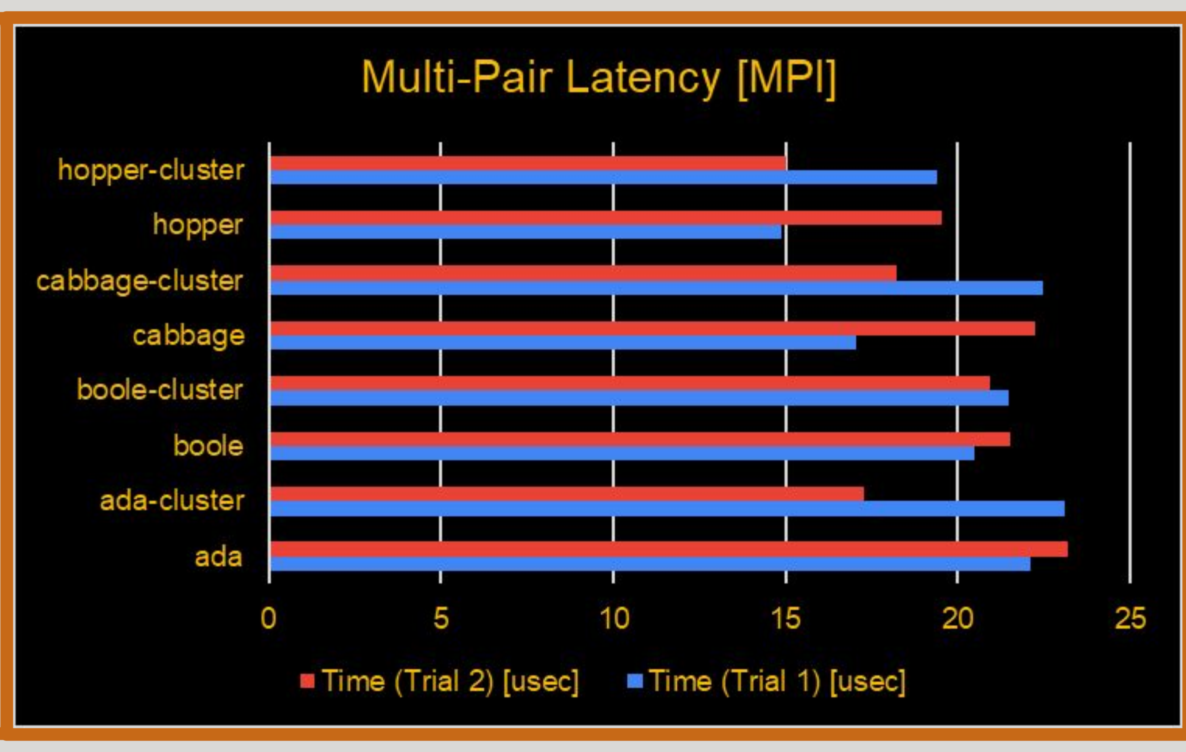

Figure 2:
Multi-Pair Latency Benchmark

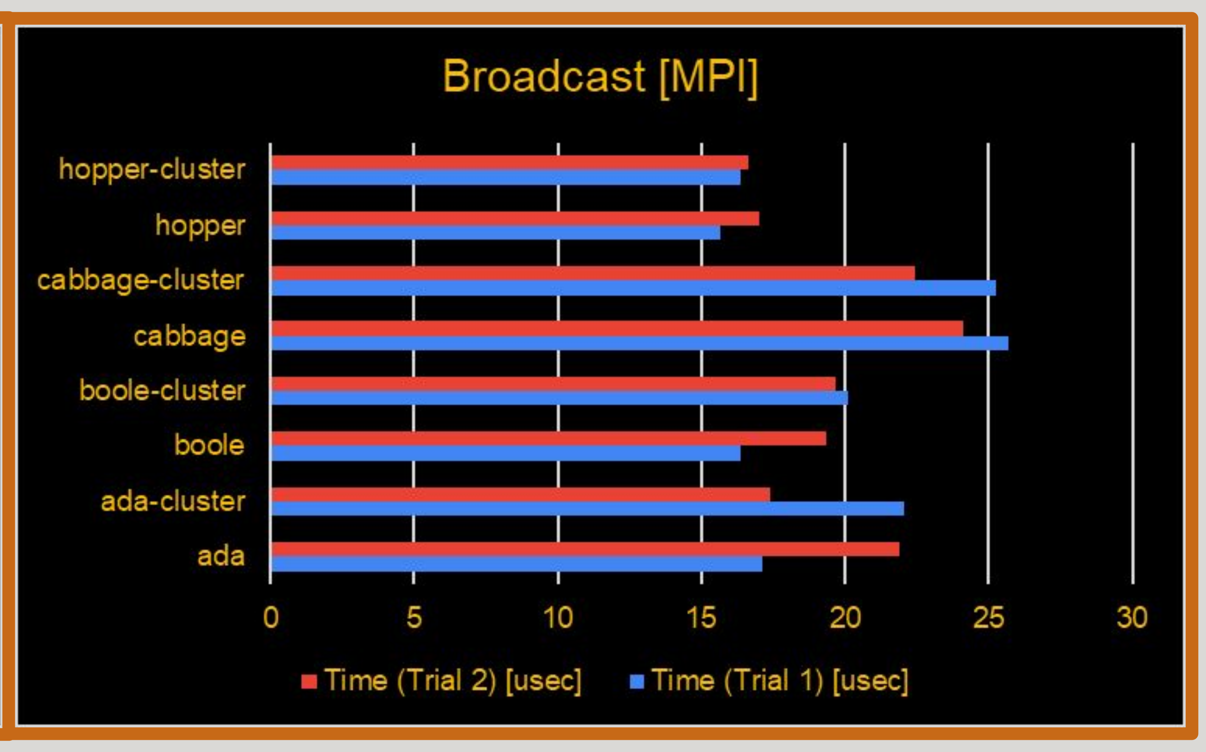

Figure 3:
Broadcast Benchmark

Conclusion

Yes, my professor has wasted money on excess network capacity. While the data above shows a lack of consistency between using the 1GbE and 10GbE network adapters to improve performance on each benchmark. A closer look at figure 1 shows a fluctuation in megabytes per second (mb/s) when running the benchmark with the 10GbE network adapter enabled and disabled over the two trials. Although this does not invalidate the results, on a job of this type there is no clear advantage. However, more consistency in the data needs to be established before coming to a more general conclusion.

It is also of note that there the potential for other benchmarks to be more appropriate for the question at hand, as there is no general purpose instrument.. Benchmarks fit in the high performance computing space is highly variable. More testing is necessarily, especially given the constraints that existed for this project. One of the major limitations was the lack of extensive replication of these benchmarks over an extended period of time. Accounting for these factors would allow for more consistent and valid data for this question. This leaves significant room for further research on the topic accounting for these limitations.

However, the research project resulted in a valid model to continue to further examine performance impacts with the 10GbE adapter - especially bonded configurations. Future research options also include testing the utility of the additional components on jobs that require access to high volume storage systems.

Resources

OpenMP
AMD
MPI
UNIVERSITÉ DU LUXEMBOURG
intel
CSinParallel
Parallel Computing in the Computer Science curriculum

Acknowledgements

Professor Nicholas Rosasco, DSc
Professor Elizabeth Shoop, PhD, Macalester College
Valparaiso University Guild (*Systems acquisition funding*)
Department of Computing and Information Sciences

VALPARAISO UNIVERSITY | Computing and Information Sciences