Automatially Generating Test Data for MPI Programs: An Annotated Bibliography

Dan Bennett September 16, 2021

References

[1] Sudheer Chunduri, Scott Parker, Pavan Balaji, Kevin Harms, and Kalyan Kumaran. Characterization of mpi usage on a production supercomputer. In *SC18: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 386–400. IEEE, 2018.

Overview: This article presents a tool, Autoperf, which can be used to monitor the execution of an MPI program. This tool is lightweight and can be used with minimal impact on the overall execution of the simulation. Using this tool they discovered that we do not use MPI as expected. They provide several suggestions for improved program performance when MPI is involved.
Evaluation: This is a good article. A limitation is that they only studied jobs on two different supercomupters. In addition, they mainly focused on long running processes. Finally, they only collected data from "production scale" projects.
Comments: This tool is available. It might be interesting to test this on a number of the "training" examples commonly used to see if they mimic the behavior of the production code.

[2] Xiangying Dang, Xiangjuan Yao, Dunwei Gong, and Tian Tian. Efficiently generating test data to kill stubborn mutants by dynamically reducing the search domain. *IEEE Transactions on Reliability*, 69(1):334–348, 2019.

Overview: Provide an overview of the article.
Evaluation: Provide an evaluation of the article.
Comments: Provide additional comments on the article.

[3] Tian Tian, Dunwei Gong, Fei-Ching Kuo, and Huai Liu. Genetic algorithm based test data generation for mpi parallel programs with blocking communication. *Journal of Systems and Software*, 155:130–144, 2019.

Overview: Provide an overview of the article.
Evaluation: Provide an evaluation of the article.
Comments: Provide additional comments on the article.