

Noise Resilient Empirical Performance Modeling with Deep Neural Networks



Project Manager
Marcus Ritter

Researchers
Benedikt Naumann and Alexander
Geiß

Principal Investigator
Prof. Dr. Felix Wolf

Project Term
2020 - 2021

Clusters
Lichtenberg Cluster Darmstadt

Software
FASTEST, TensorFlow

Institute
Department of Computer Science

University
Technische Universität Darmstadt

Introduction

Empirical performance modeling is a proven instrument to analyze the scaling behavior of HPC applications. Using a set of smaller-scale experiments, it can provide important insights into application behavior at larger scales. Extra-P is an empirical modeling tool that applies linear regression to automatically generate human-readable performance models. Similar to other regression-based modeling techniques, the accuracy of the models created by Extra-P decreases as the amount of noise in the underlying data increases. This is why the performance variability observed in many contemporary systems can become a serious challenge. In this project, we investigate novel adaptive modeling approaches that can make Extra-P more noise resilient.

Methods

We use a noise characterization heuristic to estimate the amount/level of noise on the conducted empirical performance measurements. We then train a deep neural network at the task of creating empirical performance models describing the performance of an application as a function of its configurations parameters (e.g. the number of processes or the problem size). Using the estimated noise level we use transfer learning to further improve the trained network for modeling the performance based on noisy measurement for specific applications. We use a combination of synthetically generated performance functions adding various levels of random noise to

them, and different applications case studies to evaluate our new approach.

Results

Using the synthetic data analysis and data from three different case studies conducted on the Lichtenberg cluster, we were able to improve the model accuracy of Extra-P at high noise levels by up to 25% while increasing the predictive power of the models by about 15%.

Discussion

The results of the project show that deep neural networks can be successfully used to create accurate performance models with a high predictive power based on noisy measurements. This means that we can employ Extra-P to model the performance of HPC applications even on systems with high noise levels, due to network communications or other causes. This increases the general applicability of our tool.

Outlook

For future work, we want to analyze if we can characterize the type of noise and its behavior found in the measurements even further.

Publications

Ritter, M.; Geiß, A.; Wehrstein, J.; Calotoiu, A.; Reimann, T.; Hoefler, T.; Wolf, F. : "Noise-Resilient Empirical Performance Modeling with Deep Neural Networks." In Proc. of the 35th IEEE International Parallel and Distributed Processing Symposium (IPDPS), Portland, Oregon, USA, pages 23-34, IEEE, May 2021
<https://doi.org/10.1109/IPDPS49936.2021.00012>

Last Update: 2022-03-27 12:56