Reading and Thinking

MULTITASK AND TRANSFER LEARNING FOR AUTOTUNING EXASCALE APPLICATIONS

WISSAM M. SID-LAKHDAR, MOHSEN MAHMOUDI AZNAVEH, XIAOYE S. LI, AND JAMES W. DEMMEL

ABSTRACT. Multitask learning and transfer learning have proven to be useful in the field of machine learning when additional knowledge is available to help a prediction task. We aim at deriving methods following these paradigms for use in autotuning, where the goal is to find the optimal performance parameters of an application treated as a black-box function. We show comparative results with state-of-the-art autotuning techniques. For instance, we observe an average $1.5x$ improvement of the application runtime compared to the OpenTuner and HpBandSter autotuners. We explain how our approaches can be more suitable than some state-of-the-art autotuners for the tuning of any application in general and of expensive exascale applications in particular.

https://arxiv.org/pdf/1908.05792.pdf

- Are the algebraic solvers robust in your application?
- What solvers do you use? Do you think they are good enough?
- What strategies for improving robustness will fit your application the best? Why?
- How do you want to improve robustness?