PIRA: Performance Instrumentation Refinement Automation

Jan-Patrick Lehr
Scientific Computing, Technical University of Darmstadt

December 23, 2021

1 Abstract

PIRA [1] is a tool to automatically filter and focus Score-P's profiling to *relevant program regions*. This involves both static, i.e., source-code feature, and dynamic, i.e., runtime information, analysis. It uses the whole-program call-graph representation MetaCG [2] for its analyses and has been used for automatic (1) hot-spot detection and refinement, (2) scalability analysis, (3) kernel identification, and (4) MPI load-imbalance detection.

In this talk, we present an overview of MetaCG and PIRA together with its analyses and a focus on the most recent addition of automatic (MPI) load-imbalance detection. Our experiments on the SPEC CPU 2006 suite show that PIRA automatically constructs overview measurements with runtime overhead $< 10\%$. For the load-imbalance detection, our experiments on MPI-parallel LULESH and the Ice-sheet and Sea-level System Model (ISSM) show that PIRA keeps the runtime overhead below 15%, while correctly identifying the existing load imbalances.

PIRA and MetaCG are available under BSD 3-clause license at `https://github.com/tudasc/pira` and `https://github.com/tudasc/metacg`.

References

[1] JP Lehr, A Hück, and C Bischof. 2018. PIRA: Performance Instrumentation Refinement Automation. 5th ACM SIGPLAN International Workshop on Artificial Intelligence and Empirical Methods for Software Engineering and Parallel Computing Systems (AI-SEPS 2018). ACM, `https://doi.org/10.1145/3281070.3281071`.

[2] JP Lehr, A Hück, Y Fischler, and C Bischof. 2020. MetaCG: Annotated call-graphs to facilitate whole-program analysis. 11th ACM SIGPLAN International Workshop on Tools for Automatic Program Analysis (TAPAS 2020). ACM, `https://doi.org/10.1145/3427764.3428320`.