Online Learning for fine-grain CPU frequency control with limited feedback

1 Keywords:*

Online Learning, High Performance Computing, Power-aware Computing, Frequency Scaling, Bandit Algorithms.

2 Description

Optimizing multi-threaded applications for multi-core processors is technically demanding. Such applications ideally maximize the use of all allocated cores, or offer enough legibility on their workload for underutilized resources to be de-allocated or throttled. One of the challenges for power-aware High Performance Computing is that is often just an ideal case. Accordingly, we want to explore optimization of the energy to computation ratio via external resource management approaches. More precisely, in this internship we'd like to explore the downscaling of core frequency scaling on throughput-instrumented applications. While hardware frequency control and operation counter reads are both possible at the core level, power measurements are only feasible at the socket level. This is not prohibitive in practice, since some limited-feedback frameworks from the online learning literature can fit this problem. In particular, some multi-armed bandit settings may be used. We'd like the intern to explore such limited feeback online decision making algorithmic approaches.

3 Expectations

This work is rather open in nature, so it will be up to the intern to decide whether they want to spend more time on technical aspects, experimental work or algorithmic development. We do expect however that the intern will work in collaboration with the team to run experiments on real platforms.

A minima, this means extracting relevant data for demonstrating a proof-of-concept black-box controller. Time permitting, we'd like to evaluate the policy experimentally and implement it in a local software effort (python).

4 Skills

This sort of project challenges people broadly, so a simple predictor of your success is how you far you can engage with the research ideas above. The only technical requirement is that you should be able to write Python and play around with git and Linux machines - that will come in handy to extract data, and tools you use are otherwise your prerogative. A simple requirement on your mathematical/CS culture is that you should arrive on day one understanding what the multi-armed bandits, framework is. This is applied work, algorithms exist and we won't ask you to write proofs.

5 Environment

This internship is part of a collaboration between Grenoble University and Argonne National Laboratory. A successful candidate will have the opportunity of visiting Argonne during the summer to further this work.

6 Contact

- Guillaume Huard <Guillaume.Huard@imag.fr>
- Valentin Reis <vreis@anl.gov>
- Swann Perarnau <swann@anl.gov>