SLOPE: Towards Accurate and Reliable Energy Predictive Modelling using Performance Events on Modern Computing Platforms

Heterogeneous Computing Laboratory
UCD School of Computer Science

Arsalan Shahid *(www.arsalanshahid.info)*
Supervisor: Dr Alexey Lastovetsky

Energy Consumption Measurements in HPC

HPC Platforms and Key Challenges

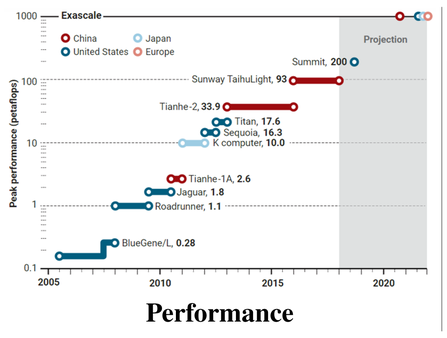

Performance

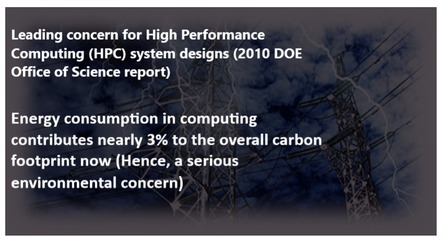

Energy

*www.sciencemag.org/news/2018/02/racing-match-chinas-growing-computer-power-us-outlines-design-exascale-computer
*science.energy.gov/ ∼ /media/ascr//pdf/reports/Exascale_subcommittee_report.pdf
*Project GreenLight: Optimizing Cyber-infrastructure for a Carbon-Constrained World, 2010

Energy Measurement / Estimation

Accurately measuring the dynamic energy consumption of an application during its execution is a key to application-level energy optimization techniques

Dominant Approaches

Physical measurements using external power-meters
* Accurate at *system-level*
* Do not provide fine-grained component-level decomposition of the energy
* Optimization of application for energy becomes difficult

Software Models
* Emerged as the pre-eminent alternative
* Ability to provide energy decomposition at finer granularity
* Pre-dominantly use performance monitoring counters (PMCs) or performance events
* Majority of PMC-based models are linear

Additivity of PMCs

Performance Monitoring Counters (PMCs) and PMC-based linear models

PMC Features
* Specific-purpose registers
* Aid low-level performance analysis/tuning
* Large in number
* Can not be collected all together
* Architecture specific

Dominant PMC groups
* Cache misses
* Floating point operations
* Page faults
* Memory Accesses
* Branch Instructions

A Simple Linear PMC-based Energy Model

$E = \sum_{j=0}^{M} \beta_j X_j + \epsilon$
where, $j = \{1, 2, ..., M\}$ are PMCs
and, ϵ is the error term

Existing issues for PMC-based Models
* Large number to consider
* Tremendous programming effort/time
* Pure PMC-based model lacks portability

Existing techniques to select PMCs subset
* Consider all PMCs
* Based on a statistical methodology
* Use of expert advice or intuition

Accuracy of PMC-based Energy Models

- PMC-based linear energy predictive models are inaccurate
  - K. O'Brien et al, *ACM Computing Surveys* **50(3)**, Article No. 37, 2017

Research Question: What are the underlying causes of their inaccuracy?
- Can we make them reliable and accurate?

Energy *Additivity*: An Experimental Observation

Let A and B be two applications with energy consumptions E_A and E_B respectively, and E_{AB} be their energy consumption when they are run one after the other serially, then

$$E_{AB} = E_A + E_B$$

Intuition: For a reliable linear energy predictive model, PMCs must follow the rule of *additivity*

Experimental Configurations, Results and Analysis

Additivity – A Selection Criterion

A selection criterion for Performance Events for reliable energy predictive linear modeling

***Additivity* Test**

Step 1: A PMC must be deterministic and reproducible
Step 2: If an application C with PMCs $\vec{e}_C$ composed of serial execution of two applications A and B with PMCs $\vec{e}_A$ and $\vec{e}_B$ respectively, then

$$\vec{e}_C = \vec{e}_A + \vec{e}_B \quad \pm\epsilon$$
Where, ϵ is user-specified tolerance

SLOPE Tools

SLOPE-PMC: Towards the automation of PMC collection for modern computing platforms
AdditivityChecker: Check PMCs for *Additivity*

* A Shahid and M Fahad, et al. in Supercomputing Frontiers and Innovations, Volume 4, Issue 4, 2017, DOI: 10.14529/jsfi170404
* https://git.ucd.ie/hcl/SLOPE/tree/master/SLOPE-PMC
* https://git.ucd.ie/hcl/SLOPE/tree/master/AdditivityChecker

Additivity Test Results

	Likwid PMCs		PAPI PMCs	
Tolerance (%)	Additive	Non-additive	Additive	Non-additive
5	108	43	36	17
20	116	35	38	15
30	119	32	43	10

Effect of tolerance on *additivity* of PMCs. These results have been obtained from experiments carried out at Intel Haswell Server. The test-suite contains applications from Intel MKL routines, NAS Parallel benchmarks and *Stress*. Further details can be found in [1].

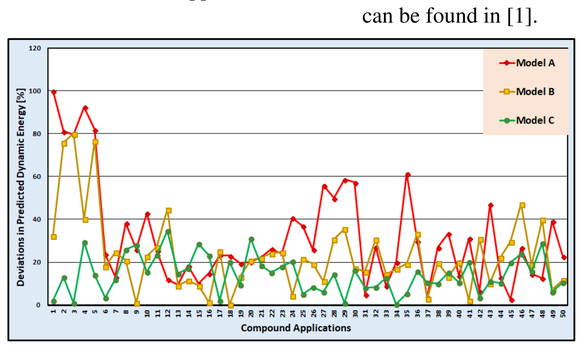

A

	Prediction Errors (%)		
Models	Max	Min	Avg
A	99.9	2.7	32
B	80	0.2	23
C	34.6	0.4	14.3

B

(A) *Non-additivity* of PMCs effecting the prediction accuracy of models. Model A uses 6 well-known PMCs, Model B using highly correlated PMCs and Model C using one highly *additive* PMC. (**B**) Models' prediction errors. Model training data set = 277 points. Model testing data set = 50 points

Discussions and Future Directions

Summary of Results and Recommendations

* Many PMCs on modern multicore machines are not *additive* [1]
* For a linear energy predictive model, all predictor variables must be *additive*.
* A PMC can be *non-additive* with error as high as 3075%
* Using *additivity* test on PMC-based models can significantly increase their prediction accuracy

On-going work and Future Research

* Study the presented linear energy predictive models in terms of *additivity* of PMCs
* Study the impact of multicores on the *additivity* of PMCs
* Generalize the assumptions behind existing linear energy predictive models
* Explore the suitability of PMCs for non-linear energy predictive modelling

References

[1] A Shahid, M Fahad, R Reddy, and A Lastovetsky, "*Additivity*: A Selection Criterion for Performance Events for Reliable Energy Predictive Modeling" in Supercomputing Frontiers and Innovations, Volume 4, Issue 4, 2017, DOI: 10.14529/jsfi170404