ADDITIVITY: A CRITERIA TO SELECT PERFORMANCE EVENTS FOR RELIABLE ENERGY PREDICTIVE ANALYTICAL MODELING

Arsalan Shahid, Muhammad Fahad, Ravi Reddy, Alexey Lastovetsky

Heterogeneous Computing Laboratory,
School Of Computer Science And Informatics,
University College Dublin,
Belfield, Dublin 4, Ireland

arsalan.shahid@ucdconnect.ie

ENERGY

A First-Class Design Constraint in All Computing Platforms

ENERGY MEASUREMENT/ESTIMATION

Accurately measuring 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 and on-chip sensors

Software Models

PMC BASED MODELS

Majority of energy predictive software models are linear predominantly use performance events or performance monitoring counters (PMCs) to predict energy consumption

Issues with PMC based Models

- Large number of PMCs to consider
- Require tremendous programming effort and time to collect PMCs
- Pure PMC based model lacks portability

PMC BASED MODELS (CONT.)

Our focus is mainly on techniques employed to extract a subset of PMCs suitable for linear energy predictive modelling

Categorizing existing techniques to select a subset of PMCs

- Techniques that consider all PMCs to capture all possible contributors to energy consumption
- Techniques that are based on a statistical methodology
- Techniques that use expert advice or intuition to pick a subset

What is missing?

ADDITIVITY

A selection criteria to determine a set of PMCs that can potentially be used for reliable energy predictive modeling

Based on an experimental observation: Energy consumption of a serial execution of two applications is the sum of energy consumptions observed for the individual execution of each application

The ***additivity*** is based on simple and intuitive rule that the value of a PMC for a compound application is equal to the sum of its values for the executions of the base applications constituting the compound application

- If the experimentally observed PMCs (sample means) of two base applications are $\overline{e1}$ and $\overline{e2}$ respectively, then a non-additive PMC of the compound application will exhibit a count experimentally that does not lie between $(\overline{e1} + \overline{e2})(1 - e)$ and $(\overline{e1} + \overline{e2})\ (1 + e)$, where the tolerance, $e = 0.05$

ADDITIVITY (CONT.)

We brand a PMC ***non-additive*** on a platform if there exists a compound application for which the calculated value significantly differs from the value observed for the application execution on the platform (within a tolerance of 5.0%).

If we fail to find a compound application (typically from a sufficiently large suite of compound applications) for which the additivity criterion is not satisfied, we term the PMC as ***potentially additive***, which means that it can potentially be used for reliable energy predictive modeling.

By definition, a potentially ***additive*** PMC must be deterministic and reproducible, that is, it must exhibit the same value (within a tolerance of 5.0%) for different executions of the same application with same runtime configuration on the same platform.

The use of a ***non-additive*** PMC as a predictor variable in a model renders it inconsistent and therefore unreliable.

EXPERIMENTAL SETUP

We study the ***additivity*** of PMCs offered by two popular tools, *Likwid* and PAPI, by employing a detailed statistical experimental methodology on a modern Intel Haswell multicore server CPU

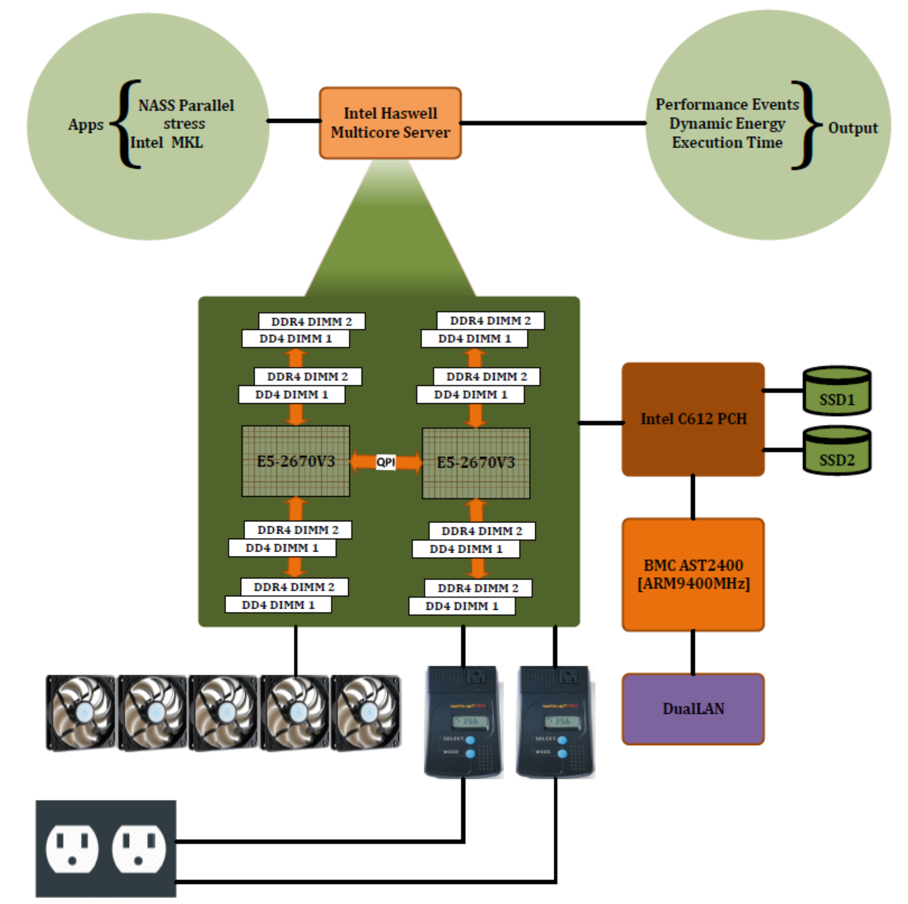

EXPERIMENTAL SETUP (CONT.)

SLOPE-PMC: Towards the Automation of PMCs Collection for Intel Based Multicore Platforms

https://git.ucd.ie/hcl/SLOPE/tree/master/SLOPE-PMC

SLOPE-PMC-LIKWID

SLOPE-PMC-PAPI

ADDITIVITY TEST FOR LIKWID PMCs

List of *potentially additive* PMCs

BR_INST_EXEC_ALL_BRANCHES	IDQ_UOPS_NOT_DELIVERED_CYCLES_0 _UOPS_DELIV_CORE
BR_MISP_EXEC_ALL_BRANCHES	IDQ_UOPS_NOT_DELIVERED_CYCLES_FE_WAS_OK
BR_INST_RETIRED_ALL_BRANCHES	UOPS_EXECUTED_PORT_PORT_0
BR_MISP_RETIRED_ALL_BRANCHES	UOPS_EXECUTED_PORT_PORT_1
DRAM_CLOCKTICKS	UOPS_EXECUTED_PORT_PORT_2
SNOOPS_RSP_AFTER_DATA_LOCAL	UOPS_EXECUTED_PORT_PORT_3
SNOOPS_RSP_AFTER_DATA_REMOTE	UOPS_EXECUTED_PORT_PORT_4
RXL_FLITS_G1_DRS_NONDATA	UOPS_EXECUTED_PORT_PORT_5
RXL_FLITS_G0_NON_DATA	UOPS_EXECUTED_PORT_PORT_6
TXL_FLITS_G0_NON_DATA	UOPS_EXECUTED_PORT_PORT_7
CPU_CLK_UNHALTED_ANY	UOPS_EXECUTED_PORT_PORT_0_CORE
CPU_CLOCK_UNHALTED_THREAD_P	UOPS_EXECUTED_PORT_PORT_1_CORE
CPU_CLOCK_UNHALTED_THREAD_P_ANY	UOPS_EXECUTED_PORT_PORT_2_CORE
CPU_CLOCK_UNHALTED_REF_XCLK	UOPS_EXECUTED_PORT_DATA_PORTS
CPU_CLOCK_UNHALTED_REF_XCLK_ANY	L2_RQSTS_ALL_DEMAND_REFERENCES
HA_R2_BL_CREDITS_EMPTY_LO_HA0	L2_RQSTS_L2_PF_MISS
HA_R2_BL_CREDITS_EMPTY_LO_HA1	MEM_UOPS_RETIRED_ALL
CPU_CLOCK_THREAD_UNHALTED _ONE_THREAD_ACTIVE	UOPS_EXECUTED_PORT_PORT_3_CORE
CPU_CLOCK_UNHALTED_TOTAL_CYCLES	UOPS_EXECUTED_PORT_PORT_4_CORE
OFFCORE_REQUESTS_OUTSTANDING _DEMAND_DATA_RD	UOPS_EXECUTED_PORT_PORT_5_CORE
OFFCORE_REQUESTS_OUTSTANDING _CYCLES_WITH_DATA_RD	UOPS_EXECUTED_PORT_PORT_6_CORE
OFFCORE_REQUESTS_OUTSTANDING _DEMAND_DATA_RD_C6	UOPS_EXECUTED_PORT_PORT_7_CORE
UOPS_EXECUTED_PORT_DATA_PORTS	UOPS_EXECUTED_PORT_ARITH_PORTS
OFFCORE_REQUESTS_DEMAND_DATA_RD	UOPS_EXECUTED_PORT_ARITH_PORTS_CORE
HA_R2_BL_CREDITS_EMPTY_HI_R2_NCB	UOPS_EXECUTED_PORT_DATA_PORTS
CPU_CLOCK_UNHALTED_THREAD_P	UOPS_RETIRED_CORE_TOTAL_CYCLES
CPU_CLOCK_UNHALTED_THREAD_P_ANY	LSD_CYCLES_4_UOPS
CPU_CLOCK_UNHALTED_REF_XCLK	UOPS_EXECUTED_THREAD
CPU_CLOCK_UNHALTED_REF_XCLK_ANY	UOPS_EXECUTED_USED_CYCLES
CPU_CLOCK_THREAD_UNHALTED _ONE_THREAD_ACTIVE	UOPS_EXECUTED_STALL_CYCLES
CPU_CLOCK_UNHALTED_TOTAL_CYCLES	UOPS_EXECUTED_TOTAL_CYCLES
ICACHE_MISSES	UOPS_EXECUTED_CYCLES_GE_1_UOPS_EXEC
L2_RQSTS_RFO_MISS	UOPS_EXECUTED_CYCLES_GE_2_UOPS_EXEC
L2_RQSTS_ALL_RFO	UOPS_EXECUTED_CYCLES_GE_3_UOPS_EXEC
L2_RQSTS_CODE_RD_HIT	UOPS_EXECUTED_CYCLES_GE_4_UOPS_EXEC
L2_RQSTS_CODE_RD_MISS	UOPS_EXECUTED_CORE
UOPS_EXECUTED_PORT_DATA_PORTS	UOPS_EXECUTED_CORE_USED_CYCLES
MEM_LOAD_UOPS_RETIRED_ALL_ALL	UOPS_EXECUTED_CORE_STALL_CYCLES
UOPS_ISSUED_ANY	UOPS_EXECUTED_CORE_TOTAL_CYCLES
UOPS_ISSUED_USED_CYCLES	UOPS_EXECUTED_CORE_CYCLES_GE_1_UOPS_EXEC
UOPS_ISSUED_STALL_CYCLES	UOPS_EXECUTED_CORE_CYCLES_GE_2_UOPS_EXEC
UOPS_ISSUED_TOTAL_CYCLES	UOPS_EXECUTED_CORE_CYCLES_GE_3_UOPS_EXEC
UOPS_ISSUED_CORE_USED_CYCLES	UOPS_EXECUTED_CORE_CYCLES_GE_4_UOPS_EXEC
UOPS_ISSUED_CORE_STALL_CYCLES	UOPS_RETIRED_ALL
UOPS_ISSUED_CORE_TOTAL_CYCLES	UOPS_RETIRED_CORE_ALL
IDQ_MITE_ALL_UOPS	UOPS_RETIRED_RETIRE_SLOTS
IDQ_DSB_UOPS	UOPS_RETIRED_CORE_RETIRE_SLOTS
IDQ_MS_UOPS	UOPS_RETIRED_USED_CYCLES
IDQ_ALL_DSB_CYCLES_ANY_UOPS	UOPS_RETIRED_STALL_CYCLES
IDQ_ALL_DSB_CYCLES_4_UOPS	UOPS_RETIRED_TOTAL_CYCLES
IDQ_ALL_MITE_CYCLES_ANY_UOPS	UOPS_RETIRED_CORE_USED_CYCLES
IDQ_UOPS_NOT_DELIVERED_CORE	UOPS_RETIRED_CORE_STALL_CYCLES
CAS_COUNT_RD	CAS_COUNT_WR
CAS_COUNT_ALL	

ADDITIVITY TEST FOR LIKWID PMCs (CONT.)

List of *non-additive* PMCs

Event Name	Maximum Percentage Error (%)
UNCORE_CLOCK	16.98
CBOX_CLOCKTICKS	16.98
SBOX_CLOCKTICKS	17.08
WBOX_CLOCKTICKS	17.57
BBOX_CLOCKTICKS	16.98
PBOX_CLOCKTICKS	16.98
RBOX_CLOCKTICKS	16.98
QBOX_CLOCKTICKS	17.57
HA_R2_BL_CREDITS_EMPTY_LO_R2_NCB	45.27
HA_R2_BL_CREDITS_EMPTY_LO_R2_NCS	48.28
HA_R2_BL_CREDITS_EMPTY_HI_HA0	203.15
HA_R2_BL_CREDITS_EMPTY_HI_HA1	213.15
HA_R2_BL_CREDITS_EMPTY_HI_R2_NCS	250.56
OFFCORE_RESPONSE_0_DMND_DATA_RD_ANY	47.50
ICACHE_IFETCH_STALL	86.60
L2_RQSTS_RFO_HIT	27.44
ARITH_DIVIDER_UOPS	3075.23
IDQ_UOPS_NOT_DELIVERED_CYCLES_LE_1_ UOP_DELIV_CORE	163.64
IDQ_UOPS_NOT_DELIVERED_CYCLES_LE_2_ UOP_DELIV_CORE	89.16
L2_RQSTS_L2_PF_HIT	39.41
ICACHE_HIT	105.45
RXL_FLITS_G0_DATA	176.62
OFFCORE_REQUESTS_OUTSTANDING_ ALL_DATA_RD	33.76
OFFCORE_REQUESTS_ALL_DATA_RD	42.45
IDQ_MITE_UOPS	42.06
L2_RQSTS_ALL_DEMAND_DATA_RD	52.76
L2_TRANS_DEMAND_DATA_RD	24.29
L2_RQSTS_ALL_DEMAND_DATA_RD_MISS	29.14
L2_RQSTS_ALL_DEMAND_DATA_RD_HIT	35.09
L2_RQSTS_ALL_DEMAND_DATA_RD	39.43
L2_TRANS_DEMAND_DATA_RD	52.43
L2_RQSTS_ALL_DEMAND_DATA_RD_MISS	56.23
L2_RQSTS_ALL_DEMAND_DATA_RD_HIT	72.32
L2_RQSTS_ALL_DEMAND_DATA_RD	35.03
L2_TRANS_DEMAND_DATA_RD	75.24
L2_RQSTS_ALL_DEMAND_DATA_RD	80.33
RXL_FLITS_G2_NCB_DATA	100
RXL_FLITS_G2_NCB_NONDATA	100
TXL_FLITS_G0_DATA	100
TXL_FLITS_G1_DRS_DATA	100
TXL_FLITS_G1_DRS_NONDATA	100
TXL_FLITS_G2_NCB_DATA	100
LSD_UOPS	42

ADDITIVITY TEST FOR PAPI PMCs

List of *potentially additive* PMCs

PAPI_L1_DCM	PAPI_FUL_CCY	PAPI_L2_DCW
PAPI_L2_DCM	PAPI_BR_UCN	PAPI_L3_DCW
PAPI_CA_SHR	PAPI_BR_CN	PAPI_L3_TCR
PAPI_CA_CLN	PAPI_BR_TKN	PAPI_L2_TCW
PAPI_CA_INV	PAPI_BR_NTK	PAPI_L3_TCW
PAPI_CA_ITV	PAPI_BR_MSP	PAPI_REF_CYC
PAPI_L1_STM	PAPI_BR_PRC	PAPI_L1_TCM
PAPI_L2_LDM	PAPI_TOT_INS	PAPI_L2_TCM
PAPI_L2_STM	PAPI_L2_DCR	PAPI_BR_INS
PAPI_PRF_DM	PAPI_L3_DCR	PAPI_RES_STL
PAPI_TOT_CYC	PAPI_L2_DCA	PAPI_L3_DCA
PAPI_L2_TCA	PAPI_L2_TCR	PAPI_L3_TCA

ADDITIVITY TEST FOR PAPI PMCs (CONT.)

List of *non-additive* PMCs

Event Name	Maximum Percentage Error (%)
PAPI_CA_SNP	40.23
PAPI_TLB_DM	31.54
PAPI_TLB_IM	23.70
PAPI_STL_CCY	31.43
PAPI_LD_INS	32.06
PAPI_SR_INS	21.98
PAPI_LST_INS	45.87
PAPI_L1_ICM	37.28
PAPI_L2_ICM	37.50
PAPI_L2_ICH	107.12
PAPI_L2_ICA	30.65
PAPI_L3_ICA	30.2
PAPI_L2_ICR	30.65
PAPI_L3_TCM	14.54
PAPI_L3_LDM	74.68
PAPI_L1_LDM	200.82
PAPI_L3_ICR	19.48

DISCUSSIONS

- For Likwid PMCs, out of a total of 151 PMCs, 43 PMCs are ***non-additive***
- For PAPI PMCs,17 PMCs out of a total of 53 PMCs are ***non-additive***
- If we increase the tolerance to about 20%, then only 8 non-additive Likwid PMCs will become potentially additive. For PAPI, only two non-additive PMCs will become *potentially additive*. Increasing the tolerance to about 30% results in other 3 non-additive Likwid PMCs and 5 ***non-additive*** PAPI PMCs becoming potentially additive
- Some of these PMCs have been used as key predictor variables in energy predictive models
  - To summarize, the non-additive PMCs that exceed a specified tolerance must be excluded from the list of PMCs to be considered as predictor variables for energy predictive modeling, because they can potentially damage the prediction accuracy of these models due to their highly non-deterministic nature. Also the list of potentially additive PMCs must be further tested exhaustively for more diverse applications and platforms to secure more confidence in their additivity.

PUBLICATION

Arsalan Shahid, Muhammad Fahad, Ravi Reddy, Alexey Lastovetsky, ”Additivity: A Selection Criterion for Performance Events for Reliable Energy Predictive Modeling” in Supercomputing Frontiers and Innovations: An International Journal, Volume 4, Issue 4, DOI: 10.14529/jsfi170404 *(Published: December 29, 2017)*

QUESTIONS/COMMENTS

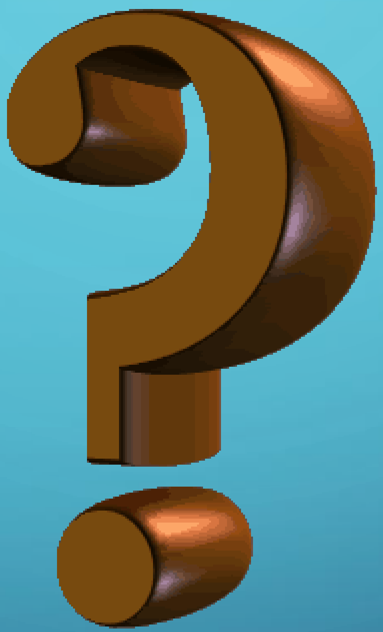

THANK YOU!

Arsalan Shahid

Ph.D. Researcher,

HCL, UCD, Ireland

https://sites.google.com/site/arsalanshahidspace/

3rd NESUS PhD Symposium, Zagreb, Croatia
January 25, 2018

arsalan.shahid@ucdconnect.ie