Performance Modelling and Dynamic Scheduling on Heterogeneous-ISA multi-core Architectures

Nirmal Kumar Boran, Dinesh Kumar Yadav, Rishabh Iyer
CADSL Lab, IIT Bombay

Introduction

- **Unicore** : Single threaded programs

  Uniprocessor e.g. Pentium 4

- **Multicore** : For multithreading

  Homogeneous-CMP: e.g. intel i7, HyperX Processors

  Heterogeneous-CMP: To optimize both serial and parallel code e.g. big.LITTLE

  Heterogeneous-ISA CMP: ISA affinity is exploited

ISA Diversity

- Code Density
- Dynamic Instruction Count
- Register Pressure
- Floating point and SIMD support

[1] Venkat, Ashish, and Dean M. Tullsen. ”Harnessing ISA diversity: Design of a heterogeneous-ISA chip multiprocessor.” 2014 ACM/IEEE 41st International Symposium on Computer Architecture (ISCA). IEEE, 2014.

Motivation

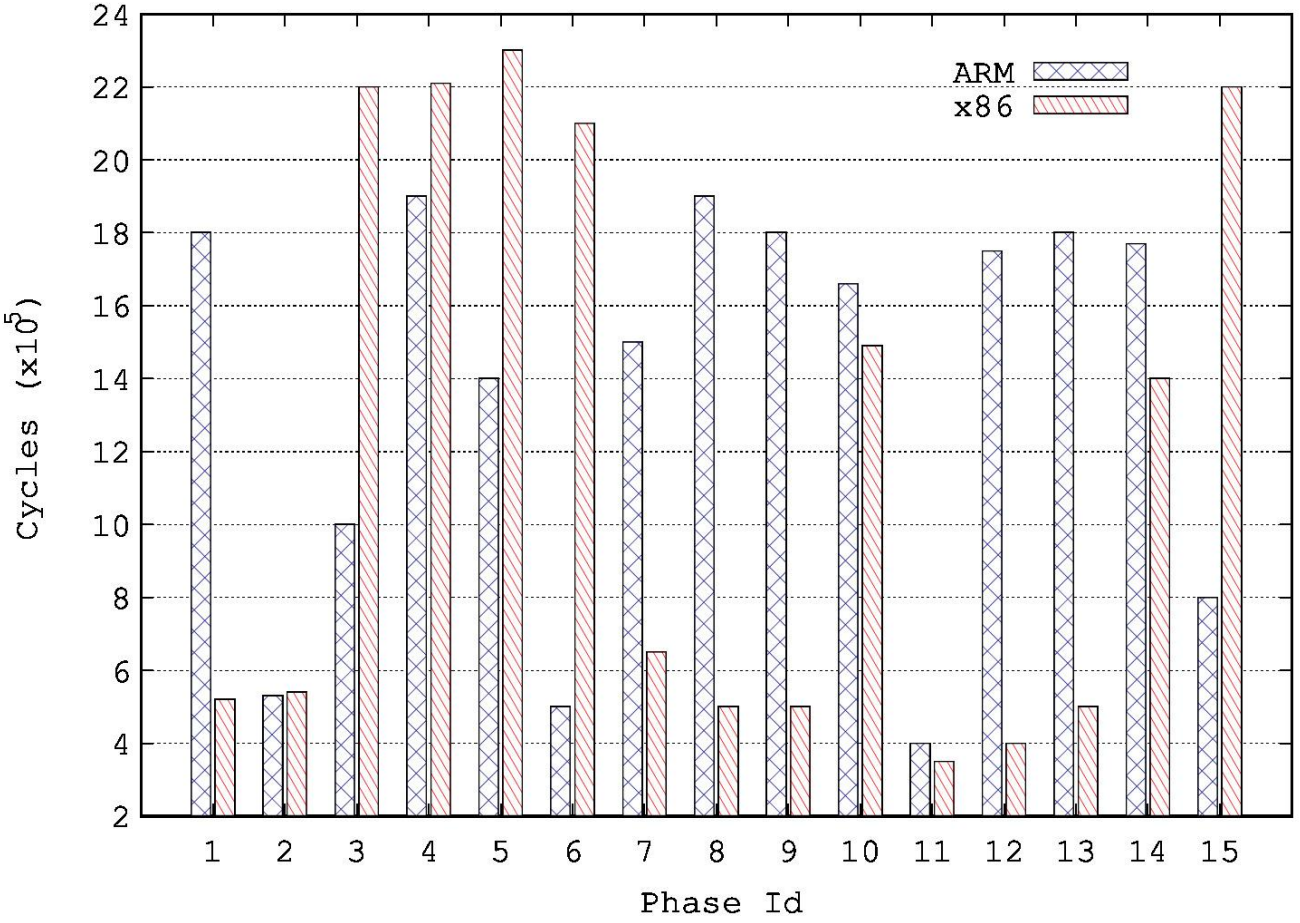

Execution time for different phases of ‘astar’ benchmark

Motivation

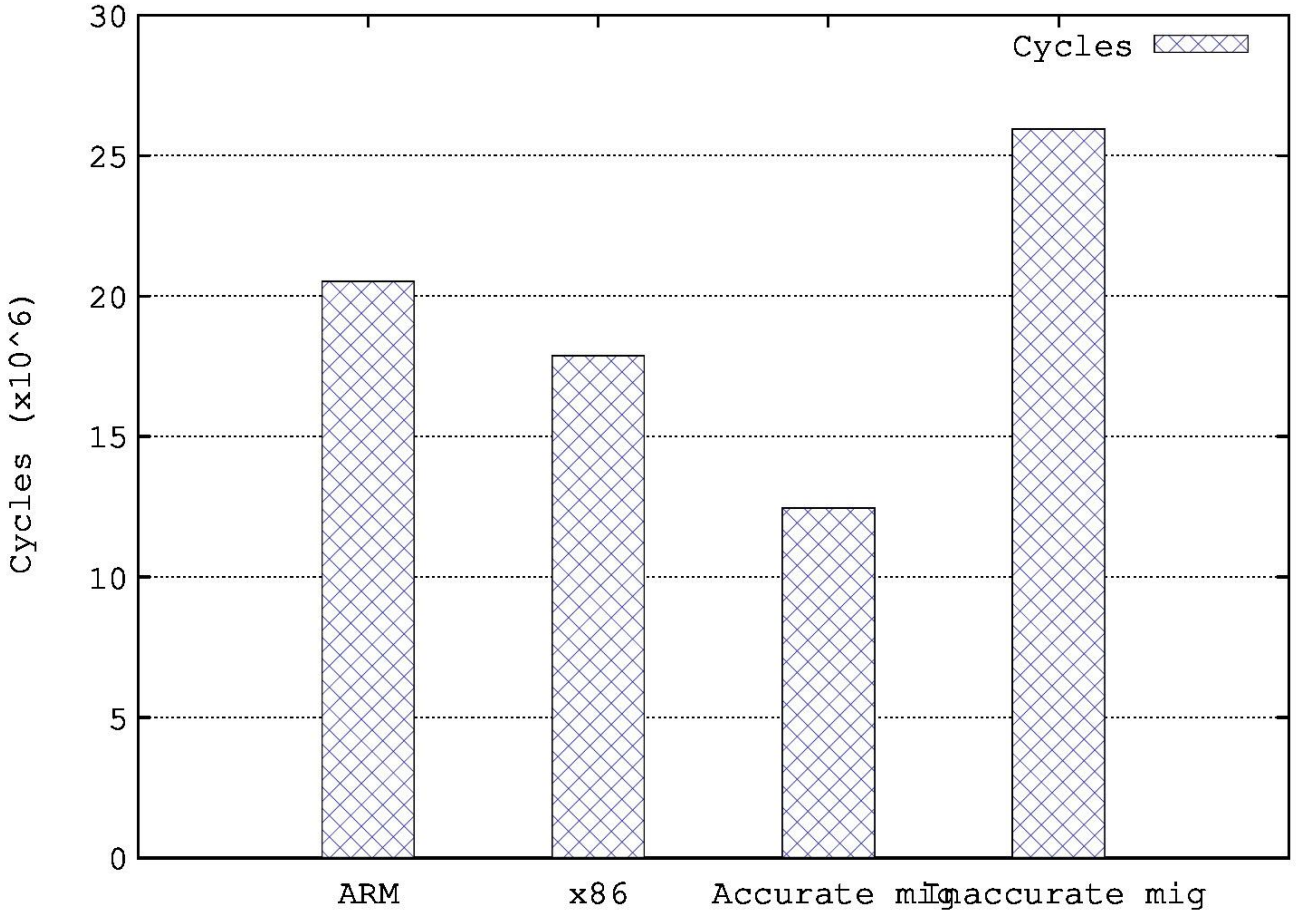
Cycles
Cycles (x10^6)
30
25
20
15
10
5
0
ARM
x86
Accurate mig
Inaccurate mig

Migration Decision: Equivalence Point

- Function calls are potential Equivalence states[1][4]

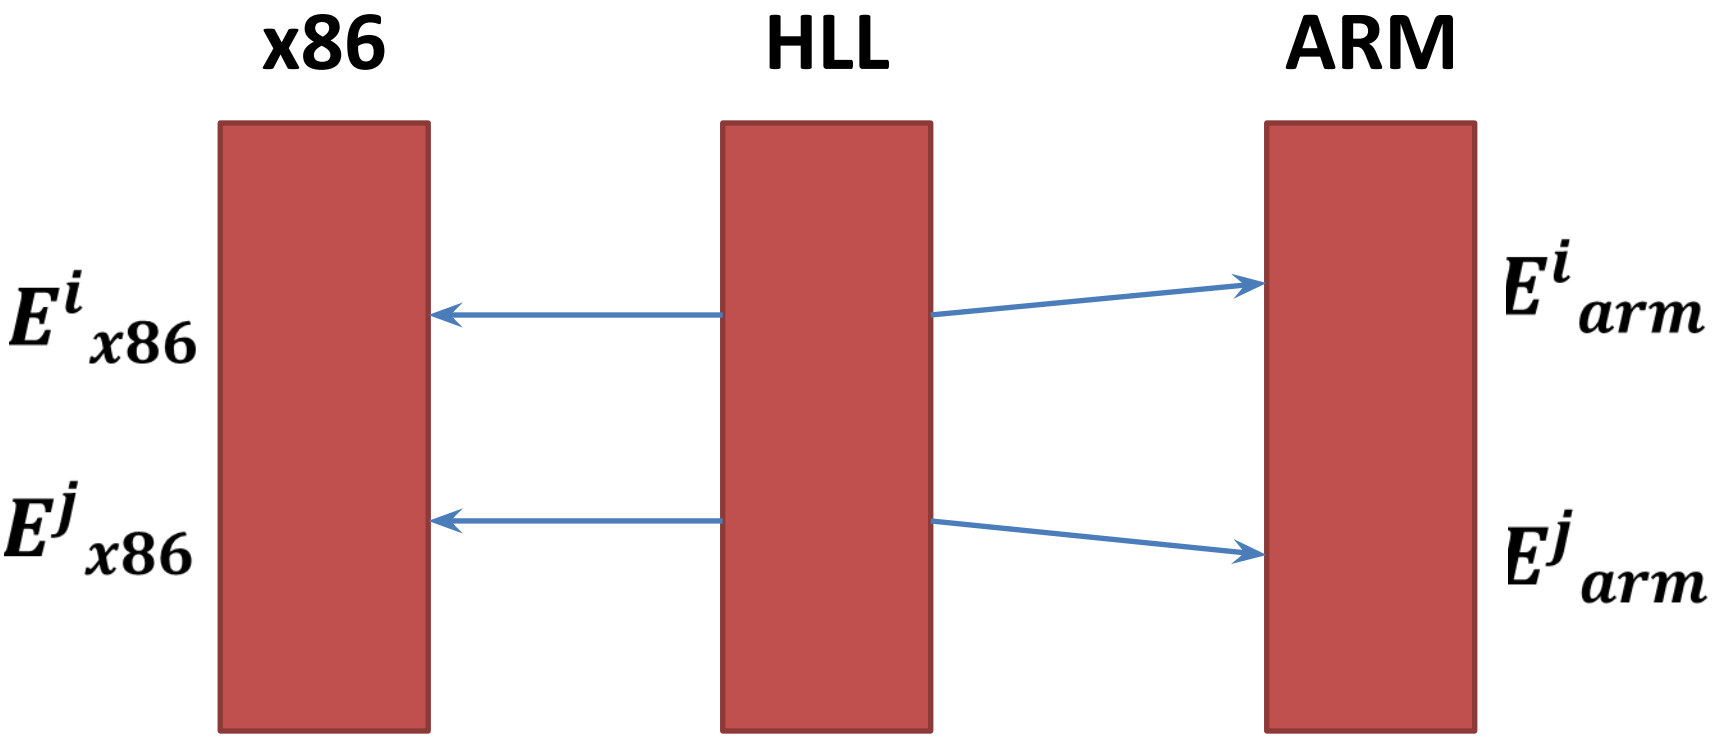

Problem Statement

- "Building a performance model and scheduler for heterogeneous ISA asymmetric multi-core architectures“

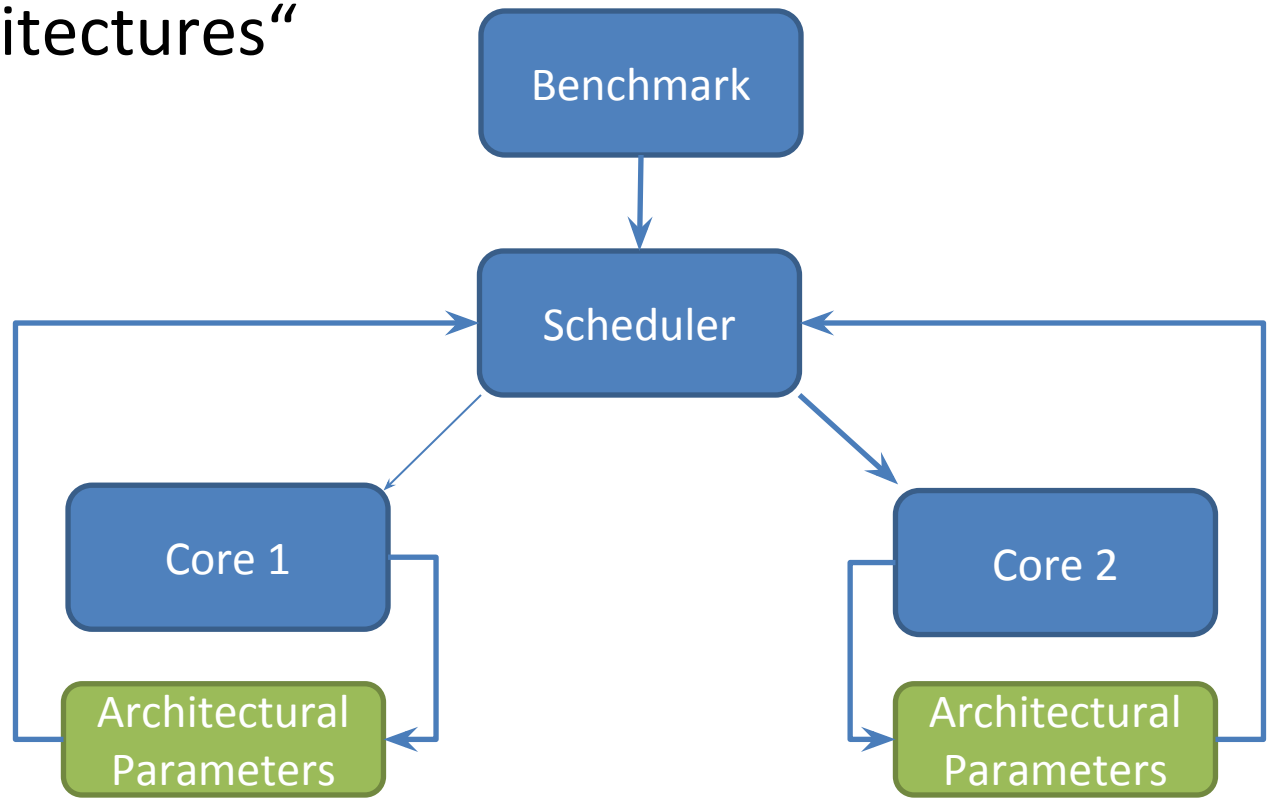

Previous Work – Training Phase

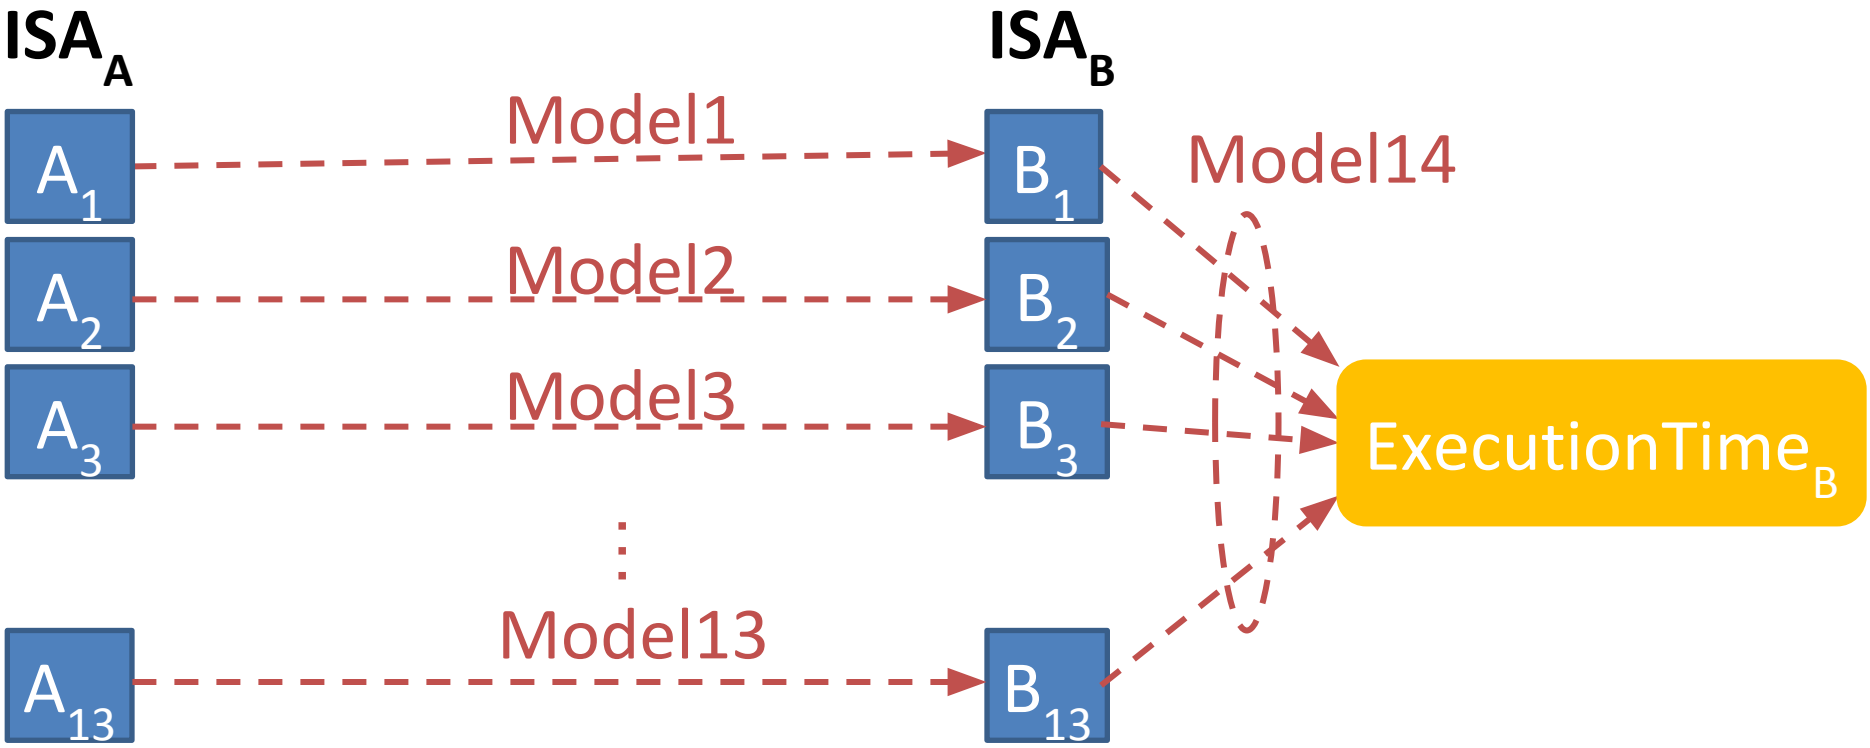

[4]: Boran, N.K., Meghwal, R.P., Sharma, K., Kumar, B., Singh, V.: Performance modelling of heterogeneous ISA multicore architectures. In: East-West Design & Test Symposium (EWDTS), 2016 IEEE. pp. 1{4. IEEE (2016)

Previous Work – Prediction Phase

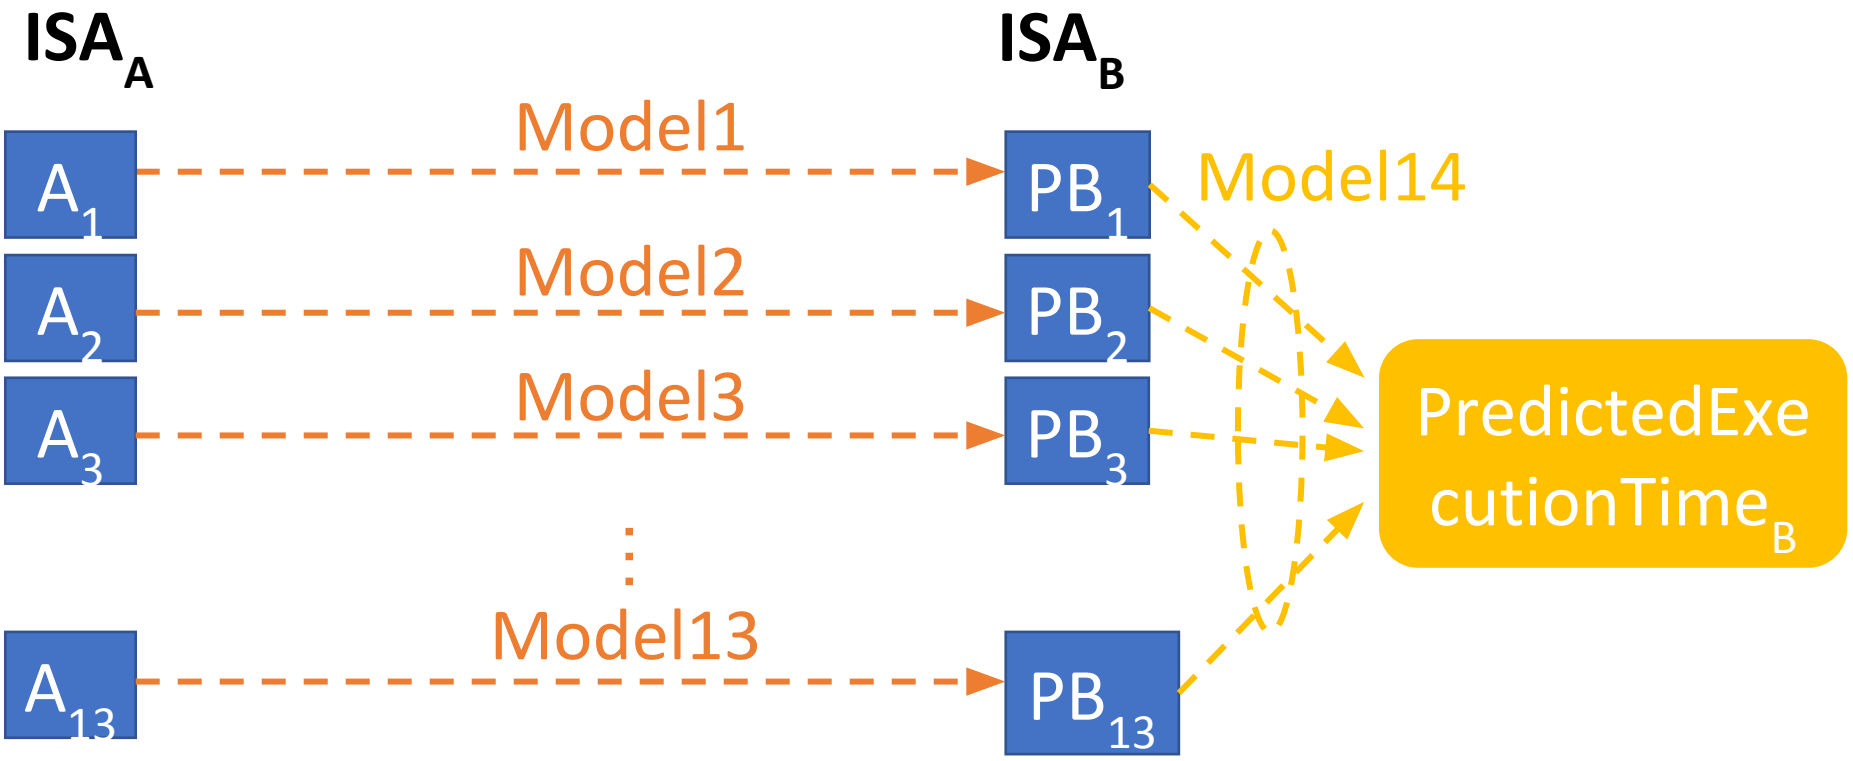

[4]: Boran, N.K., Meghwal, R.P., Sharma, K., Kumar, B., Singh, V.: Performance modelling of heterogeneous ISA multicore architectures. In: East-West Design & Test Symposium (EWDTS), 2016 IEEE. pp. 1{4. IEEE (2016)

Proposed Work – Training Phase

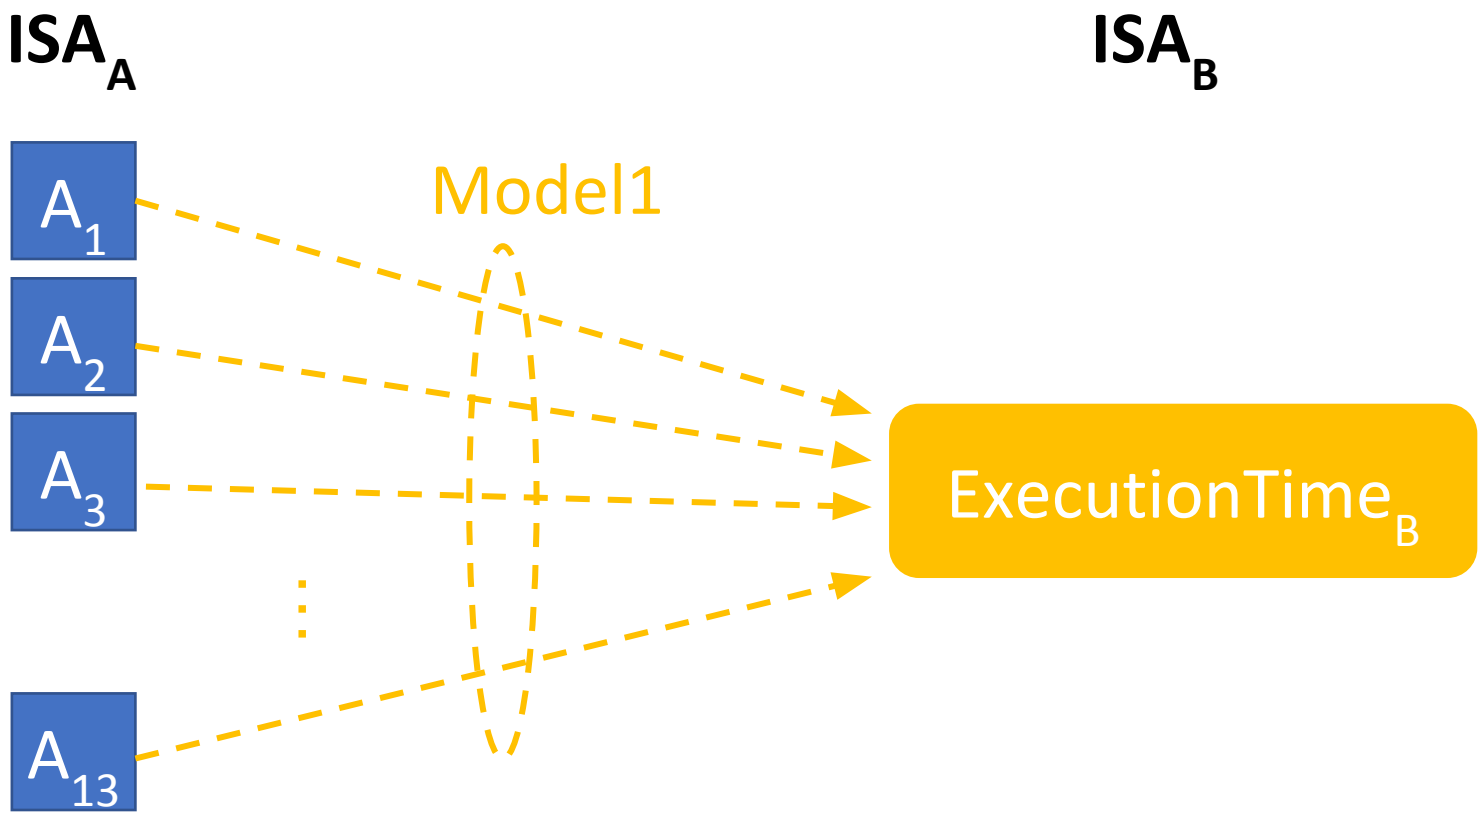

Proposed Work – Prediction Phase

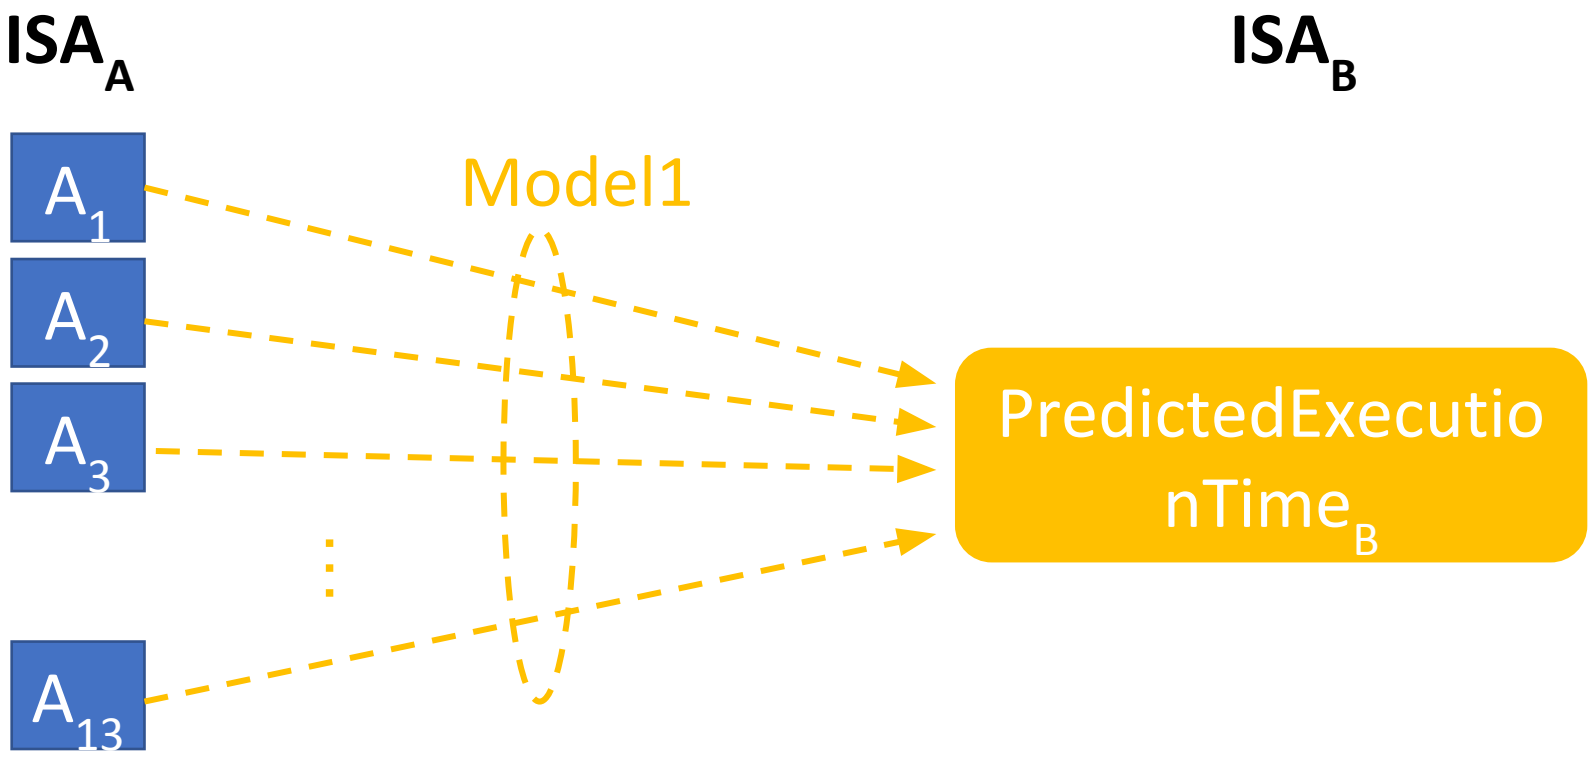

Linear Regression Equation

$$Cycle_B = K + a_1.(L1DcacheMiss_A) + a_2.(L1IcacheMiss_A) + a_3.(L2CacheMiss_A) + a_4.(IQFullEvents_A) + a_5.(SQFullEvents_A) + a_6.(ROBFullEvents_A) + a_7.(BranchMissPrediction_A) + a_8.(MLP_A) + a_9.(MSHRFullEvents_A) + a_{10}.(ILP_A) + a_{11}.(LoadCount_A) + a_{12}.(FloatInstruction_A) + a_{13}.(DynamicInstructionCount_A)$$

Configuration

Design Parameter	ARM	Alpha	x86
Architectural Registers	32 GPR	64 GPR	16 GPR
Cache line sizes (bytes)	64	64	64
LSQ size (bytes)	32	32	32
Fetch Width	4	4	4
Instruction Queue entries	64	64	64
ROB entries	192	192	192
DCache, ICache size	32KB	32KB	32KB
L2 Cache size	256KB	256KB	256KB

Results: Root mean squared error

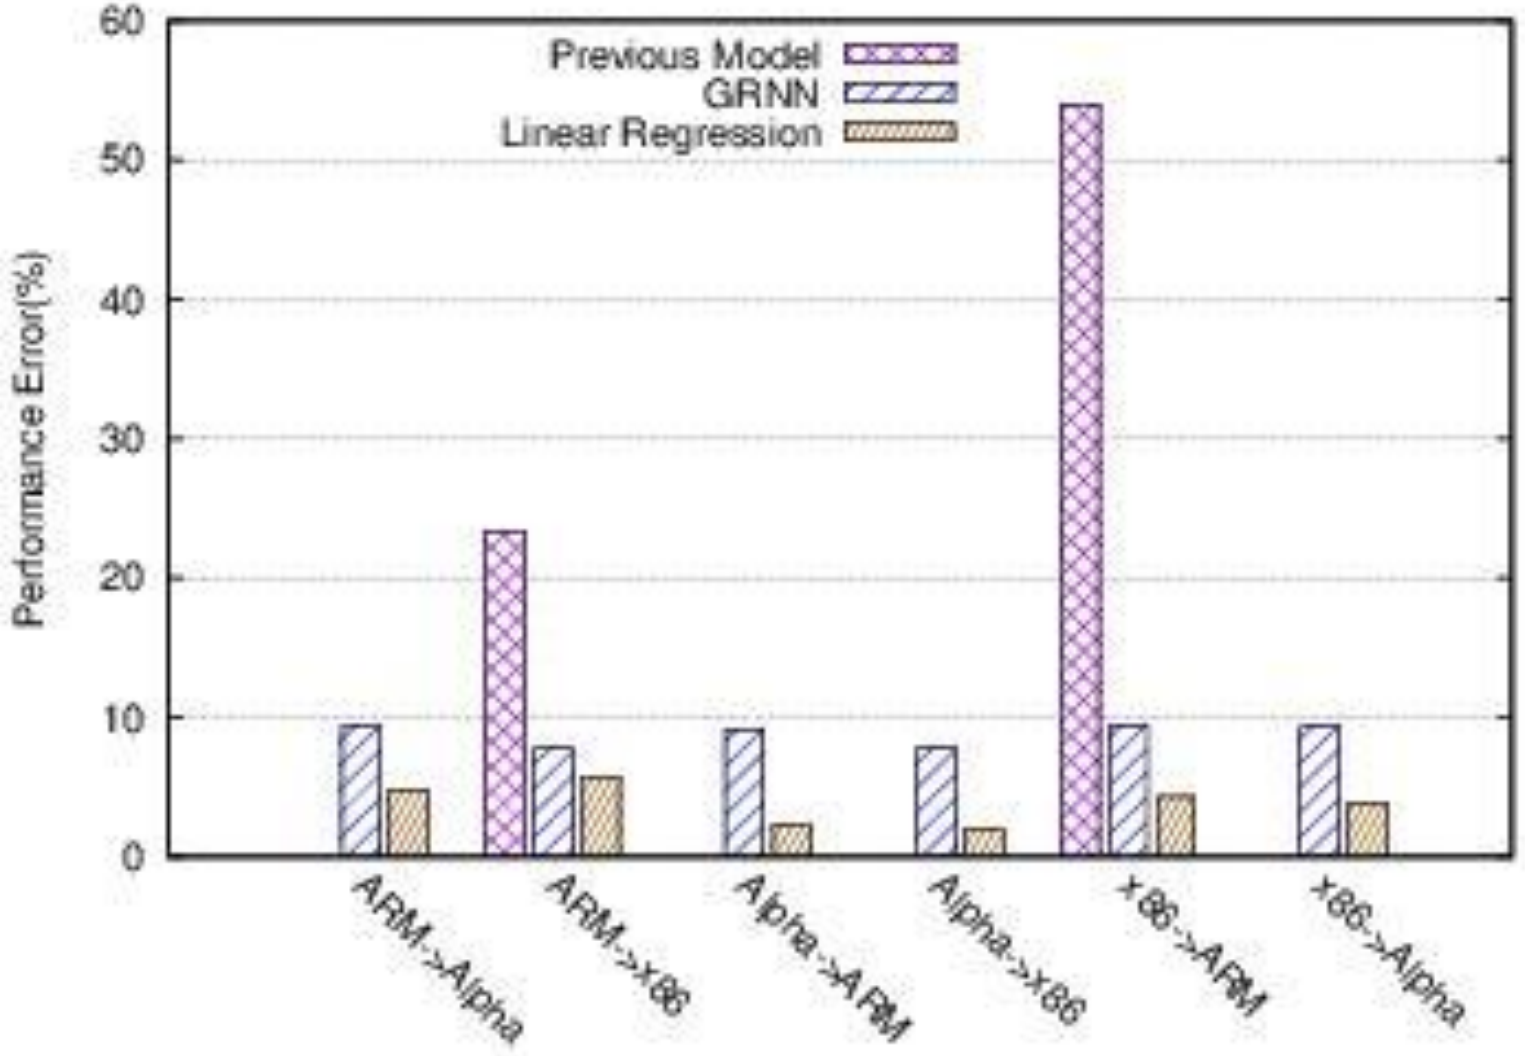

Results: Standard Deviation

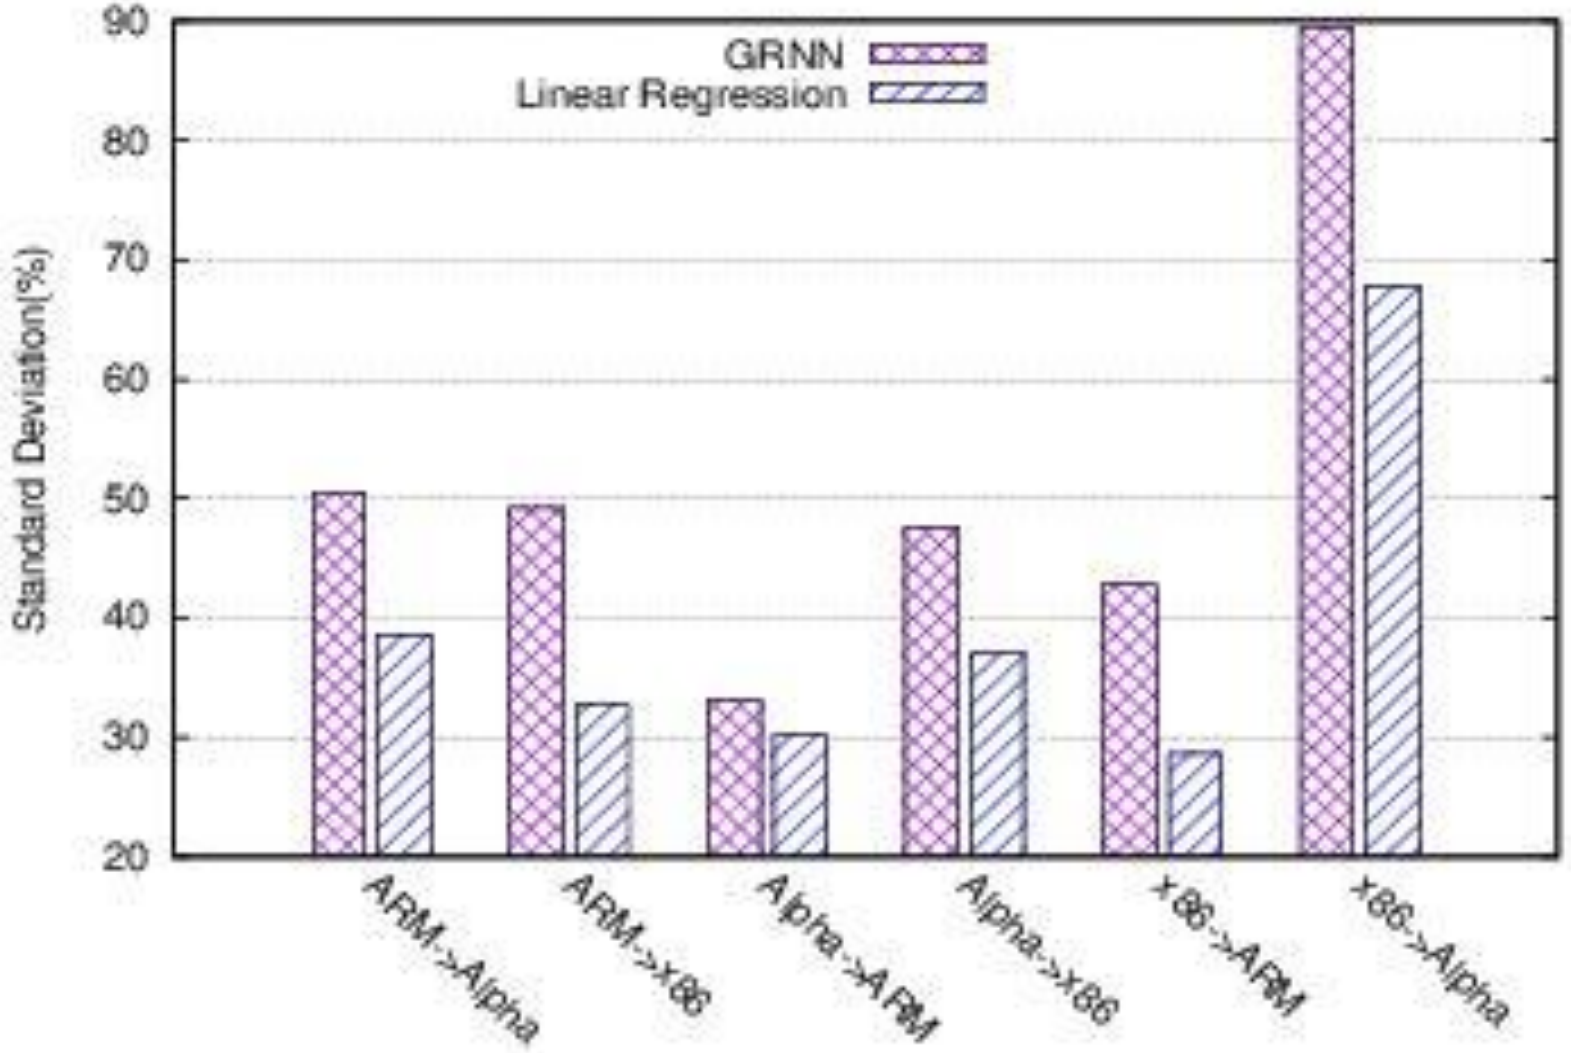

Scheduling Prediction

- If PredictedTime < (CurrentTime + migrationOverhead)

  Migrate to another ISA

ISA	Prediction Accuracy
ARM	92.3 %
Alpha	92.7 %
X86	83.4 %

Results: Speed Up 1

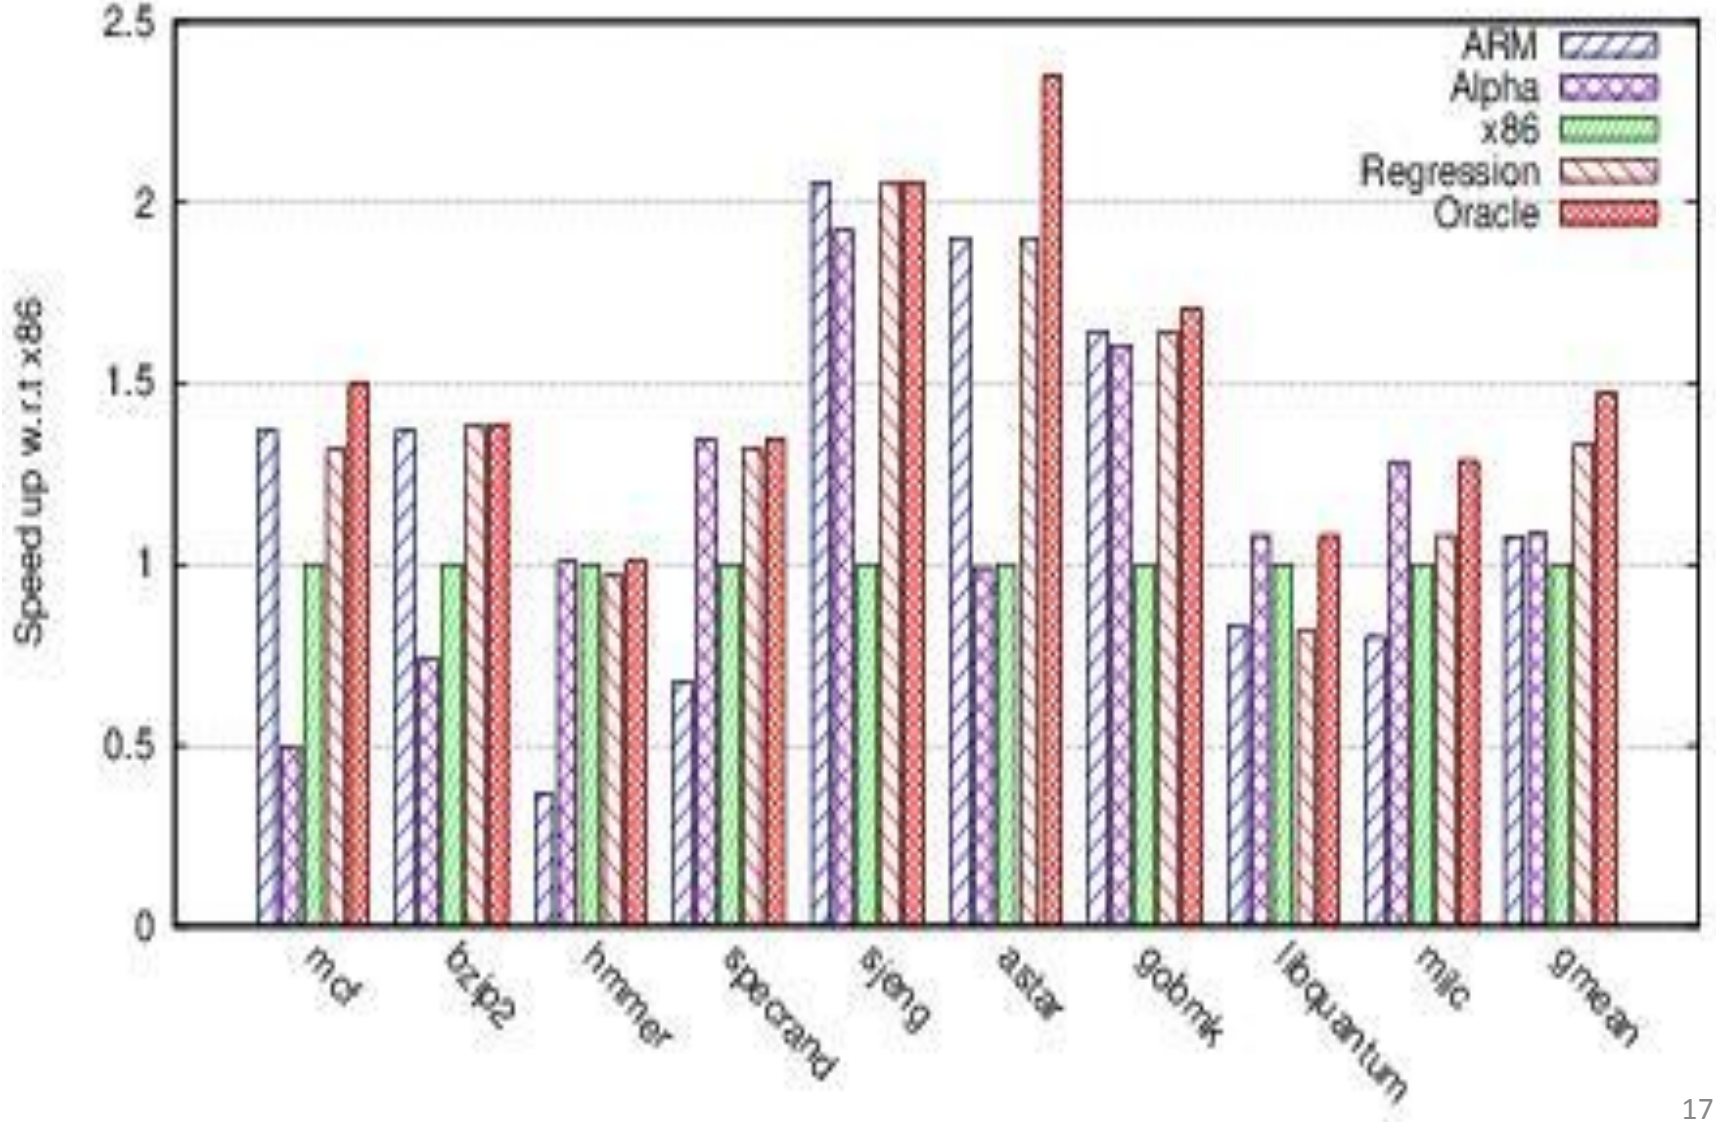

Results: Speed Up 2

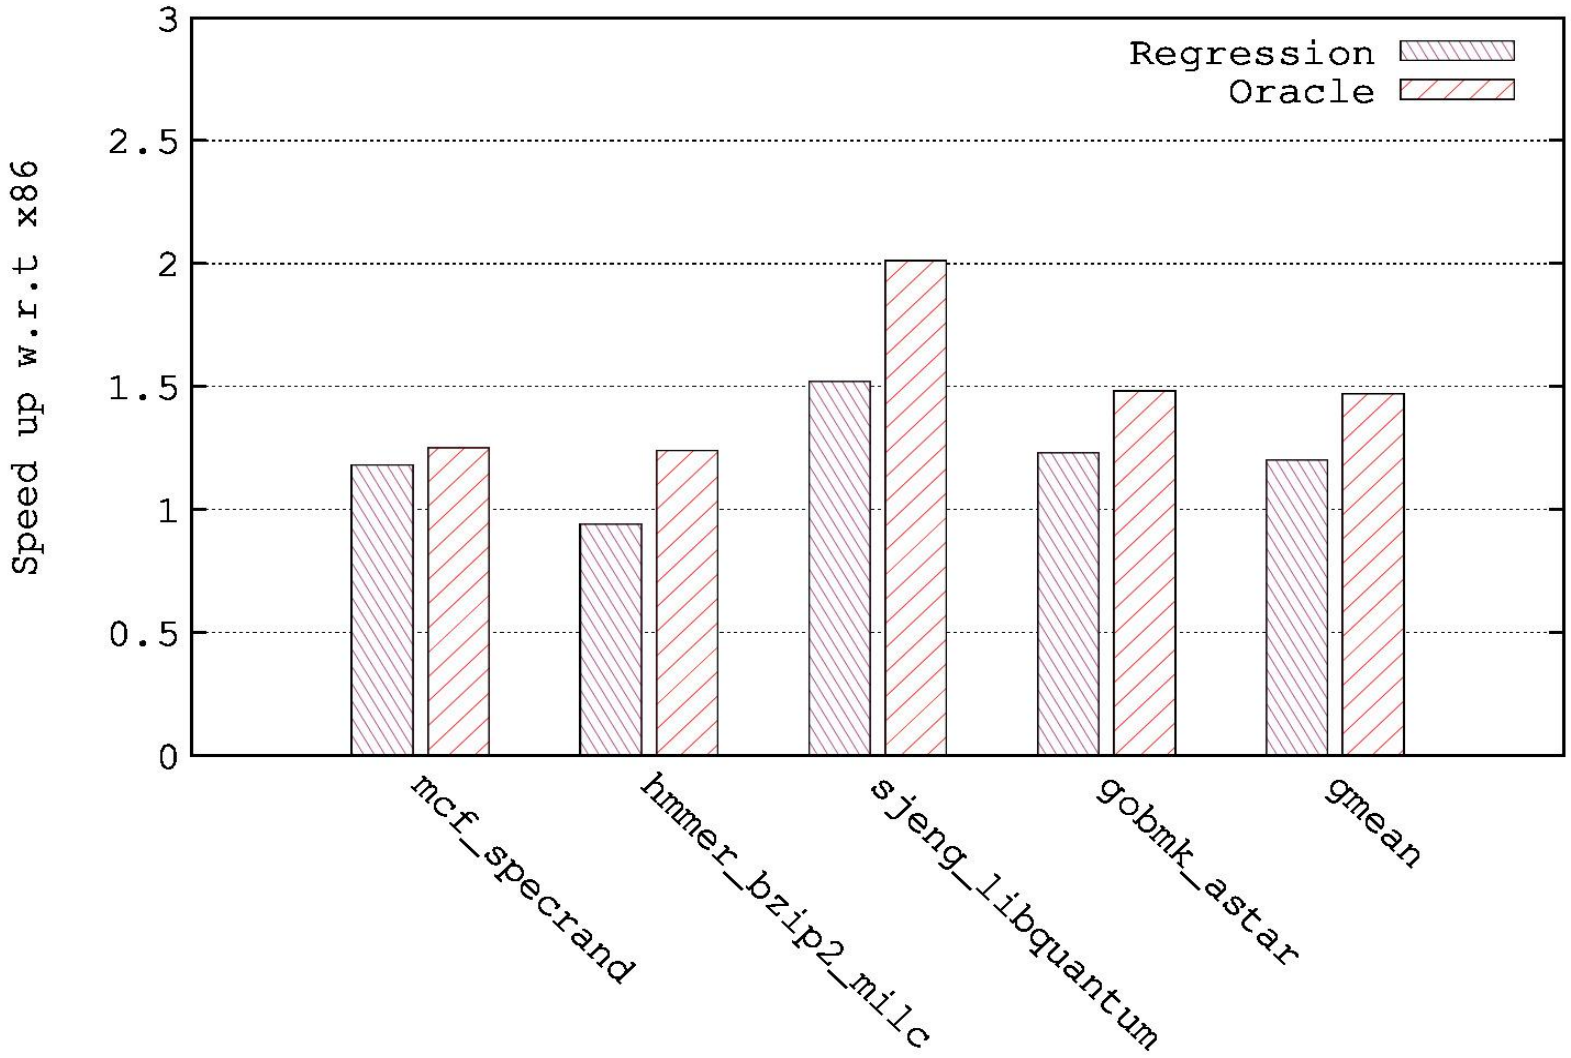

Results: Energy Comparison

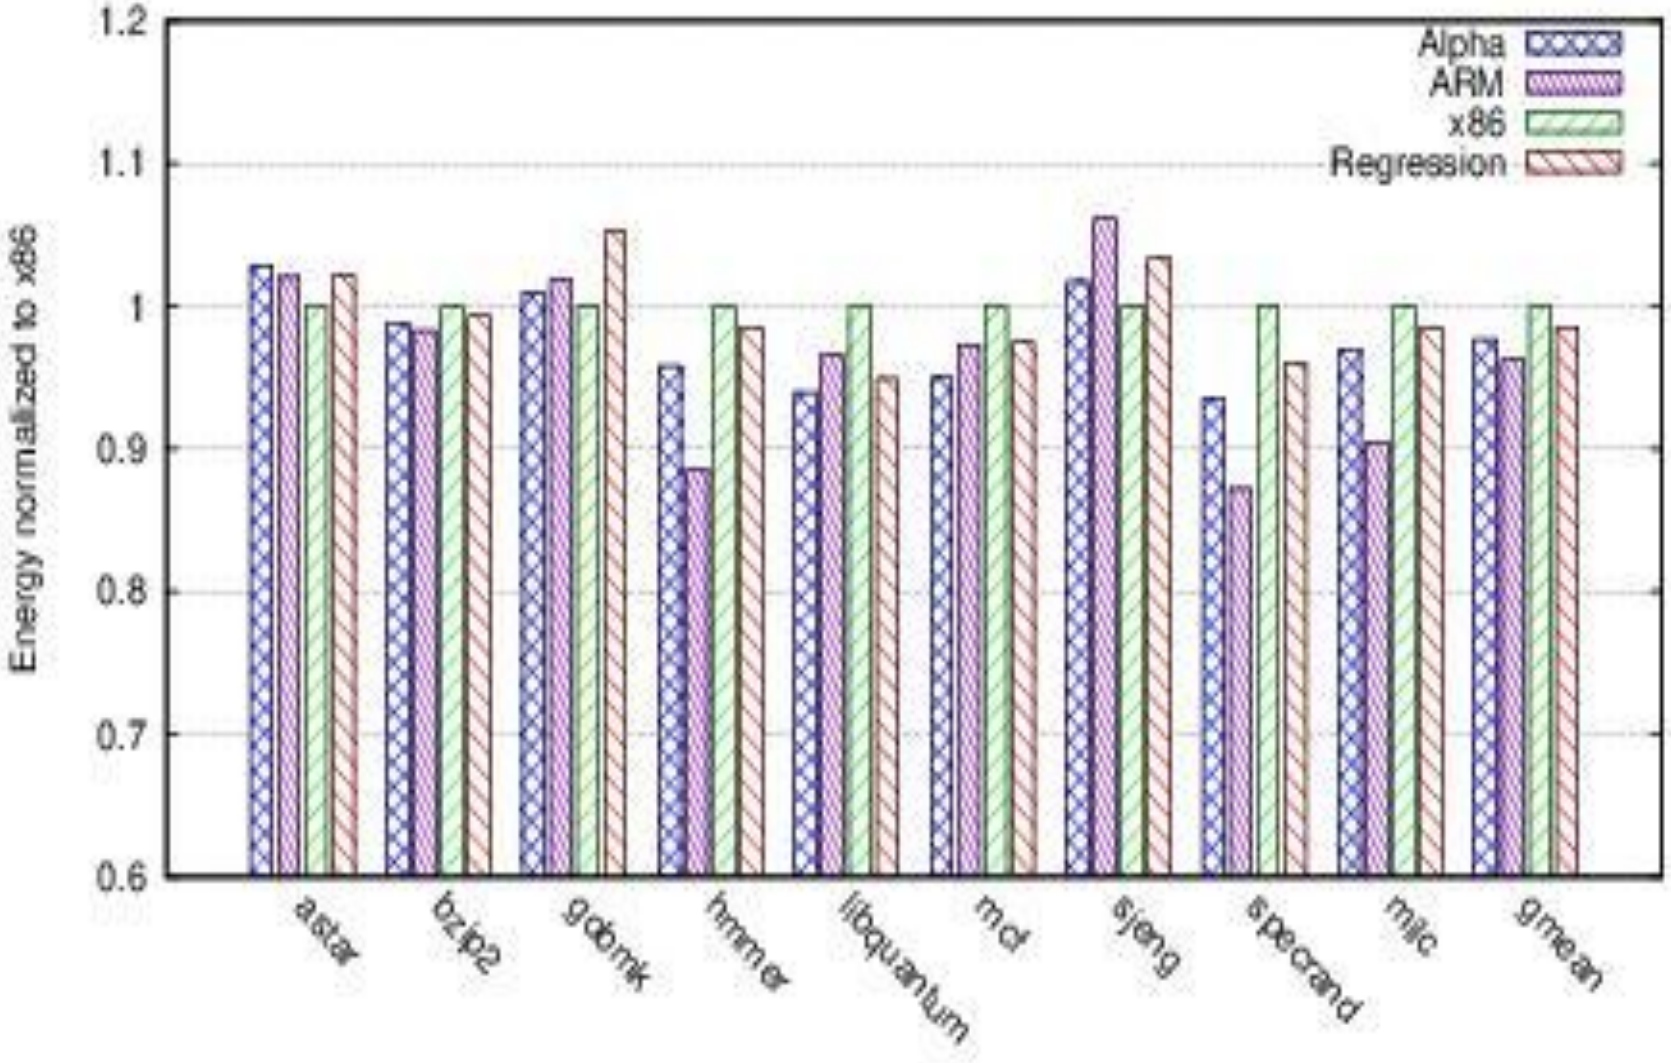

References

- [1]: Venkat, Ashish, and Dean M. Tullsen. ”Harnessing ISA diversity: Design of a heterogeneous-ISA chip multiprocessor.” 2014 ACM/IEEE 41st International Symposium on Computer Architecture (ISCA). IEEE, 2014.
- [2]: Von Bank, David G., Charles M. Shub, and Robert W. Sebesta. "A unified model of pointwise equivalence of procedural computations." *ACM Transactions on Programming Languages and Systems (TOPLAS)* 16.6 (1994)
- [3]: M. Pricopi, T. S. Muthukaruppan, V. Venkataramani, T. Mitra, and S. Vishin, “Power-performance modeling on asymmetric multi-cores,” in Proceedings of the 2013 International Conference on Compilers, Architectures and Synthesis for Embedded Systems, ser. CASES ’13.
- [4]: Boran, N.K., Meghwal, R.P., Sharma, K., Kumar, B., Singh, V.: Performance modelling of heterogeneous ISA multicore architectures. In: East-West Design & Test Symposium (EWDTS), 2016 IEEE. pp. 1{4. IEEE (2016)
- [5]: N. L. Binkert, R. G. Dreslinski, L. R. Hsu, K. T. Lim, A. G. Saidi, and S. K. Reinhardt, The m5 simulator: Modeling networked systems, IEEE Micro, vol. 26, no. 4, pp. 5260, July 2006.
- [6]: A. Lukefahr, S. Padmanabha, R. Das, F. M. Sleiman, R. Dreslinski, T. F. Wenisch, and S. Mahlke, Composite cores: Pushing heterogeneity into a core, in Proceedings of the 2012 45th Annual IEEE/ACM International Symposium on Microarchitecture, ser. MICRO-45. Washington, DC, USA: IEEE Computer Society, 2012, pp. 317328.
- [7]: Kumar, R., Farkas, K.I., Jouppi, N.P., Ranganathan, P., Tullsen, D.M.: SingleISA heterogeneous multi-core architectures: The potential for processor power reduction. In: Microarchitecture, 2003. MICRO-36. Proceedings. 36th Annual IEEE/ACM International Symposium on. pp. 81{92. IEEE (2003)

Thank you

Migration decision

- Execution of a program on a processor- sequence of states
- State change - execution of a single instruction
- There exists certain well-defined states of computation where one to one correspondence is possible with another instance of computation[2]