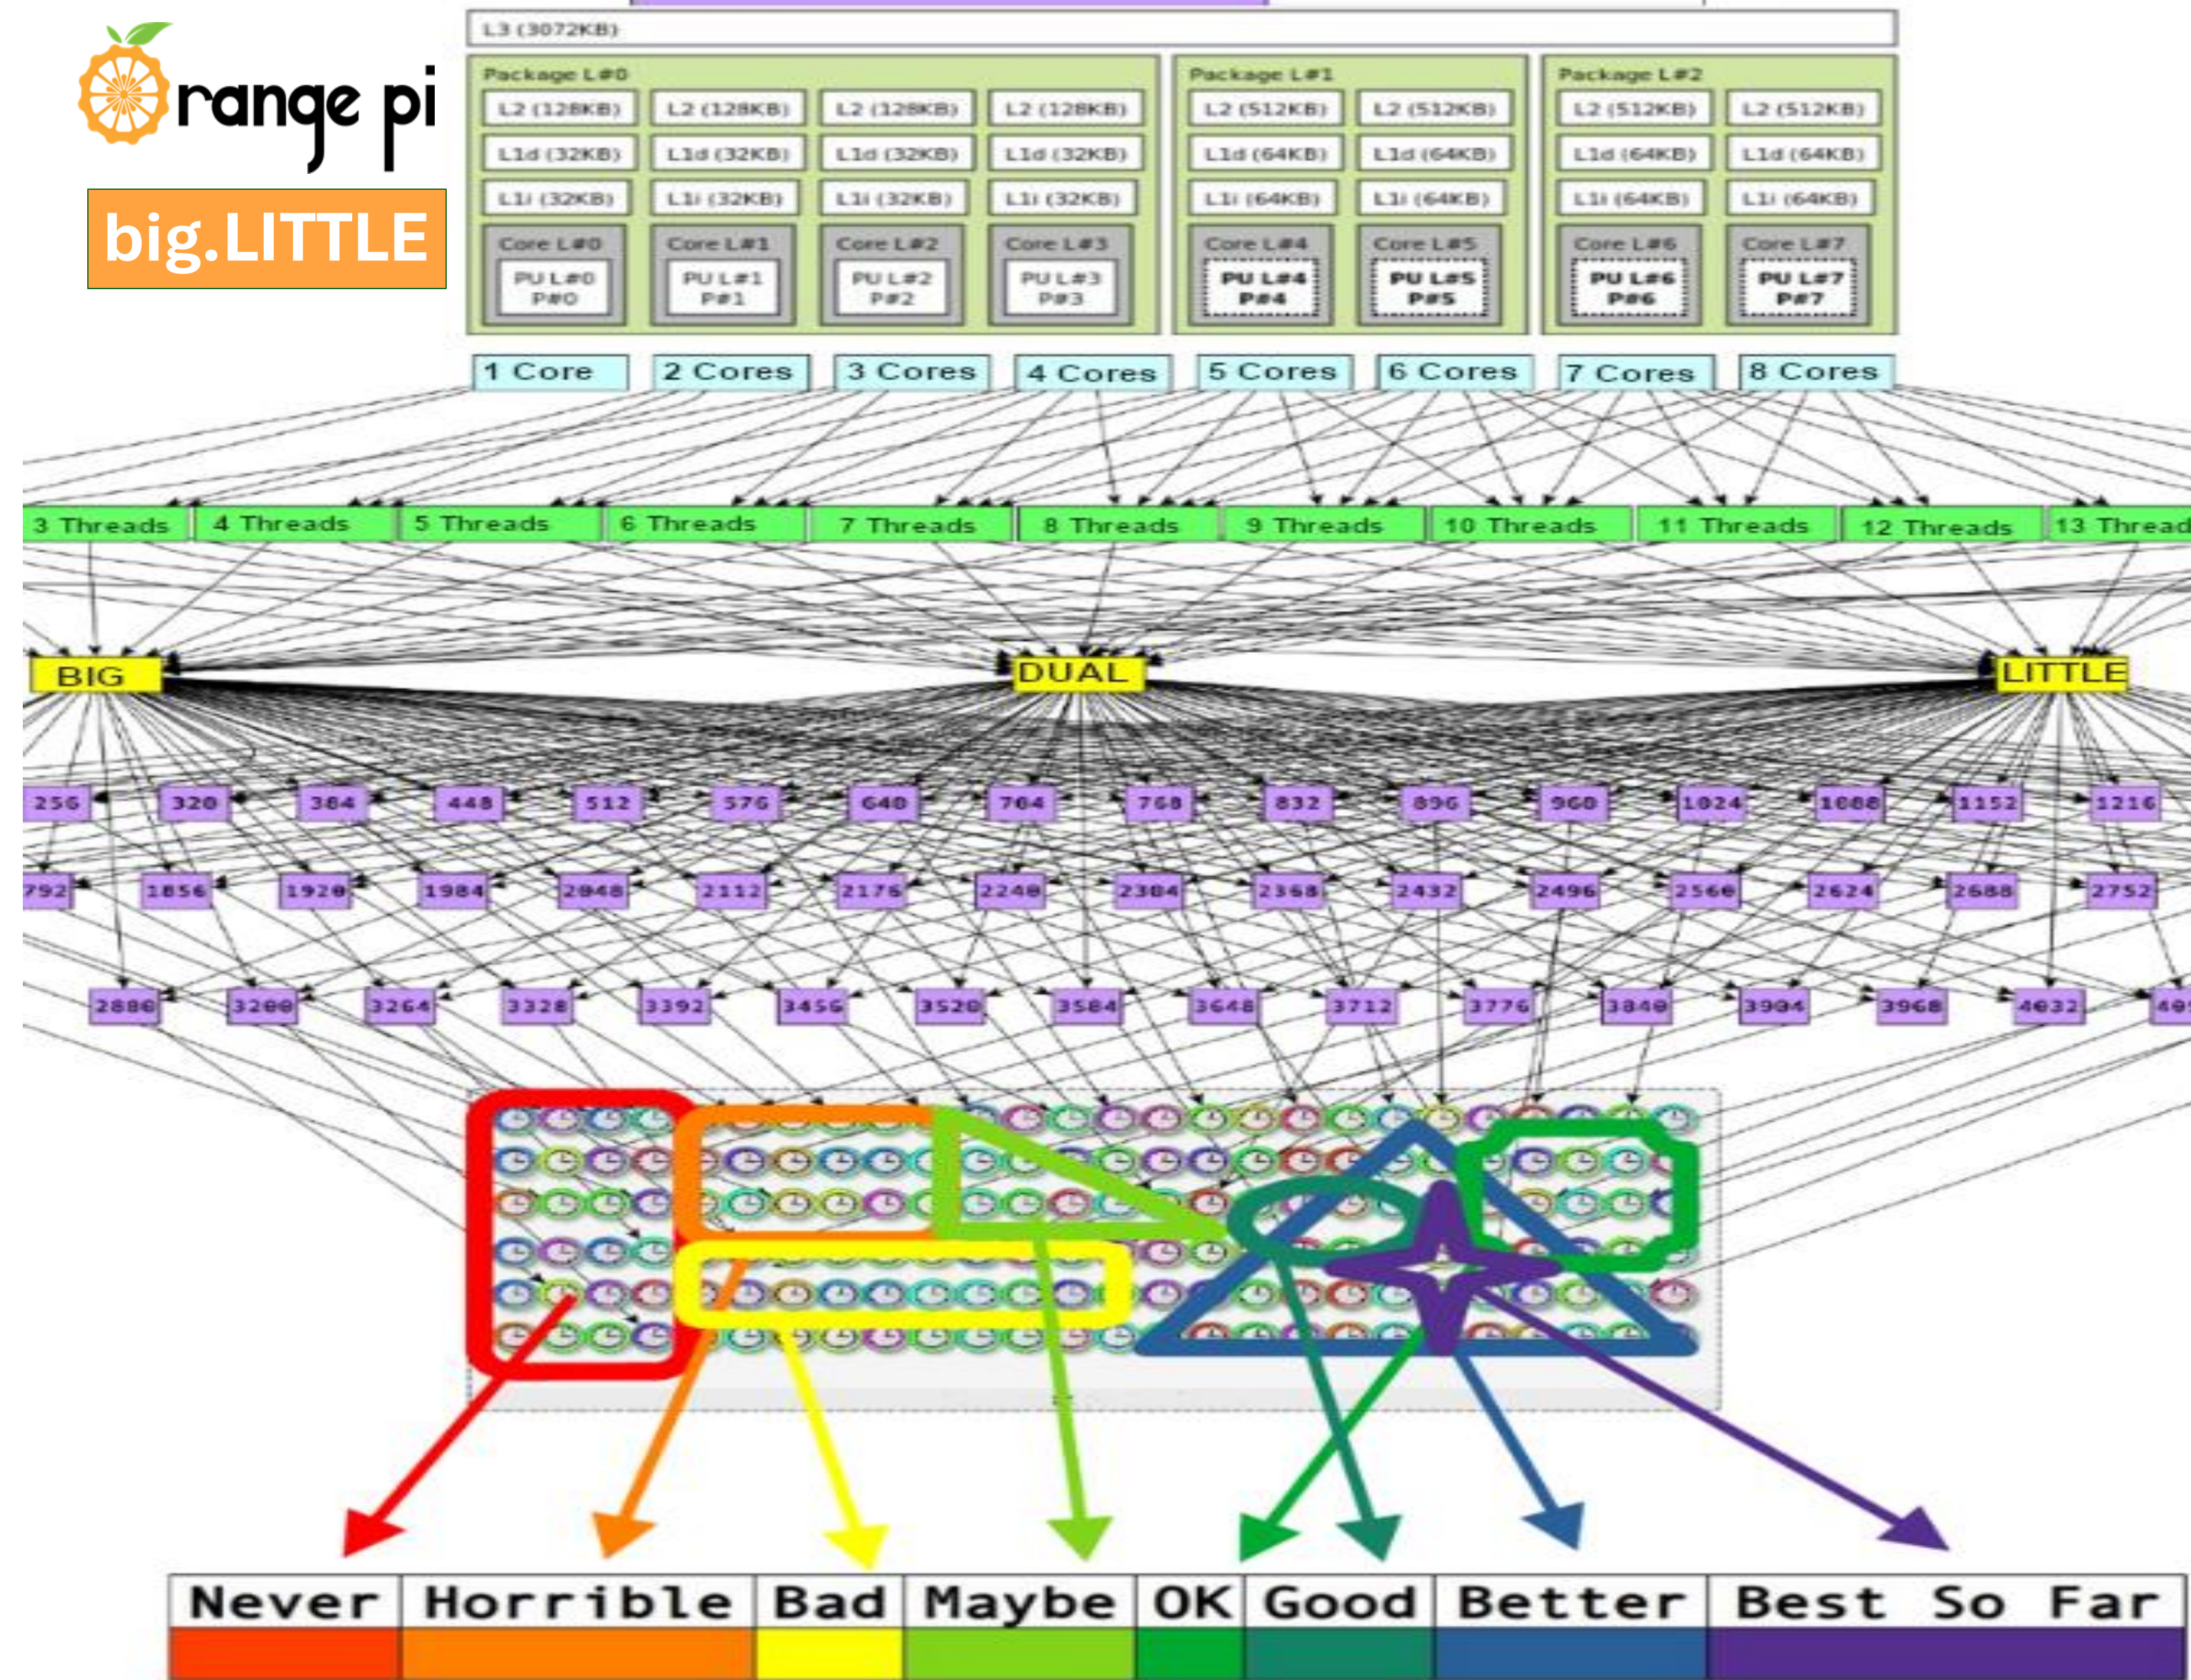
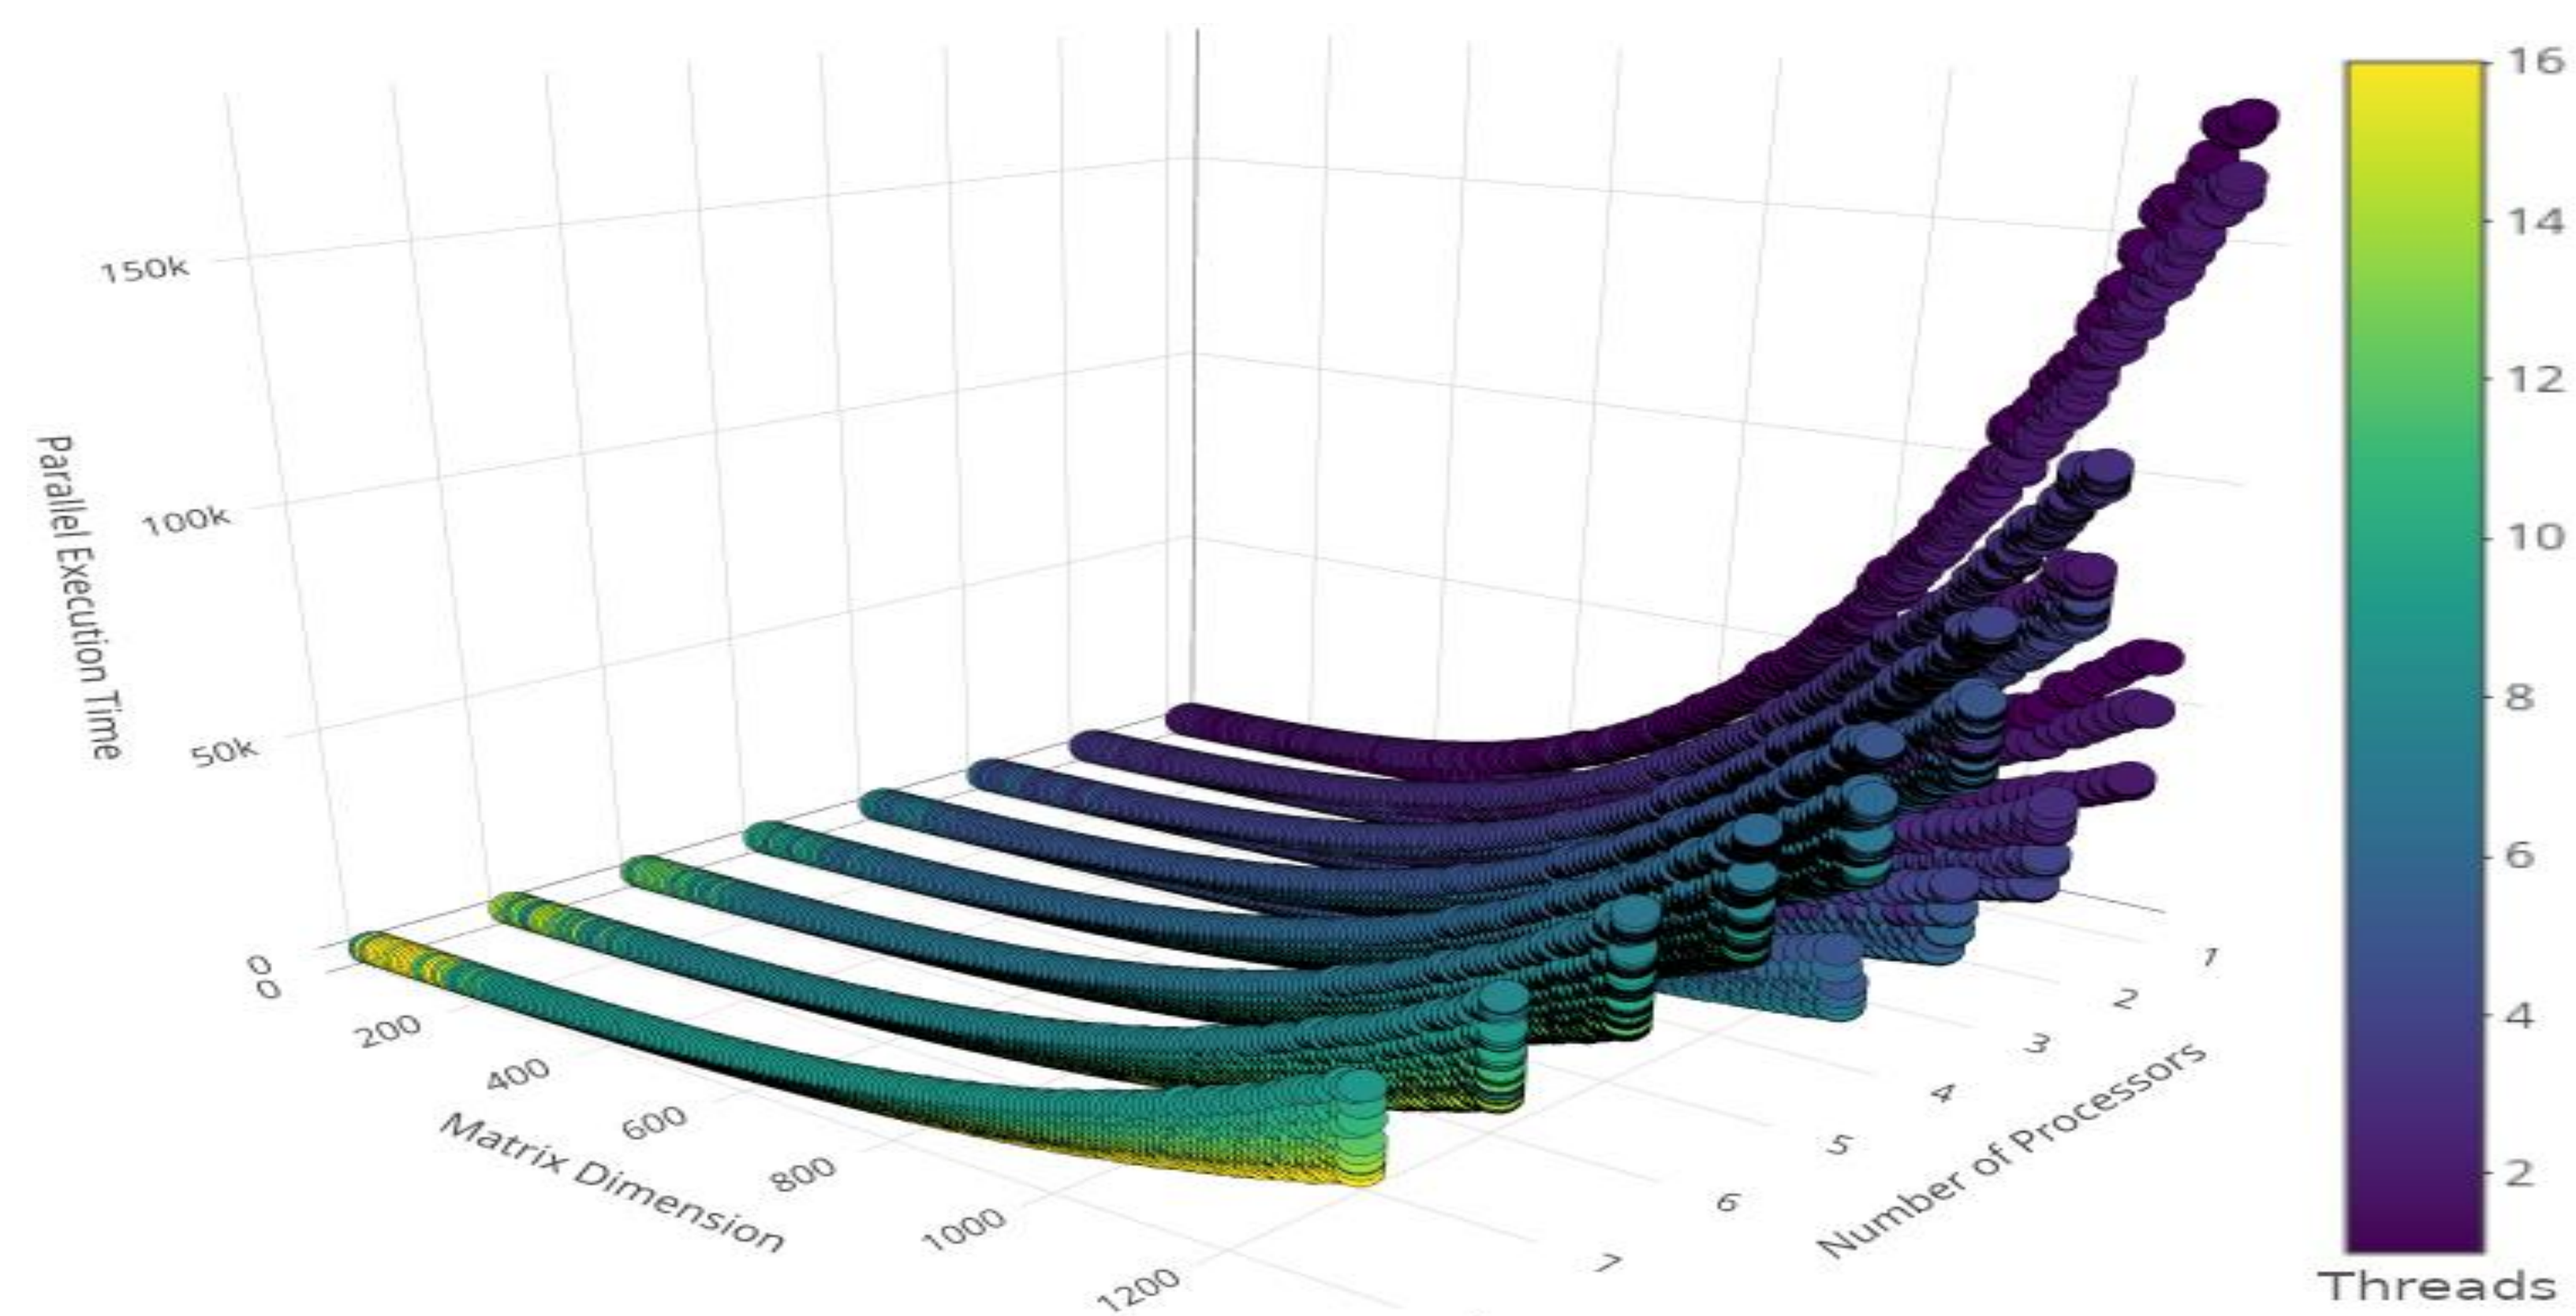


Platform Complexity

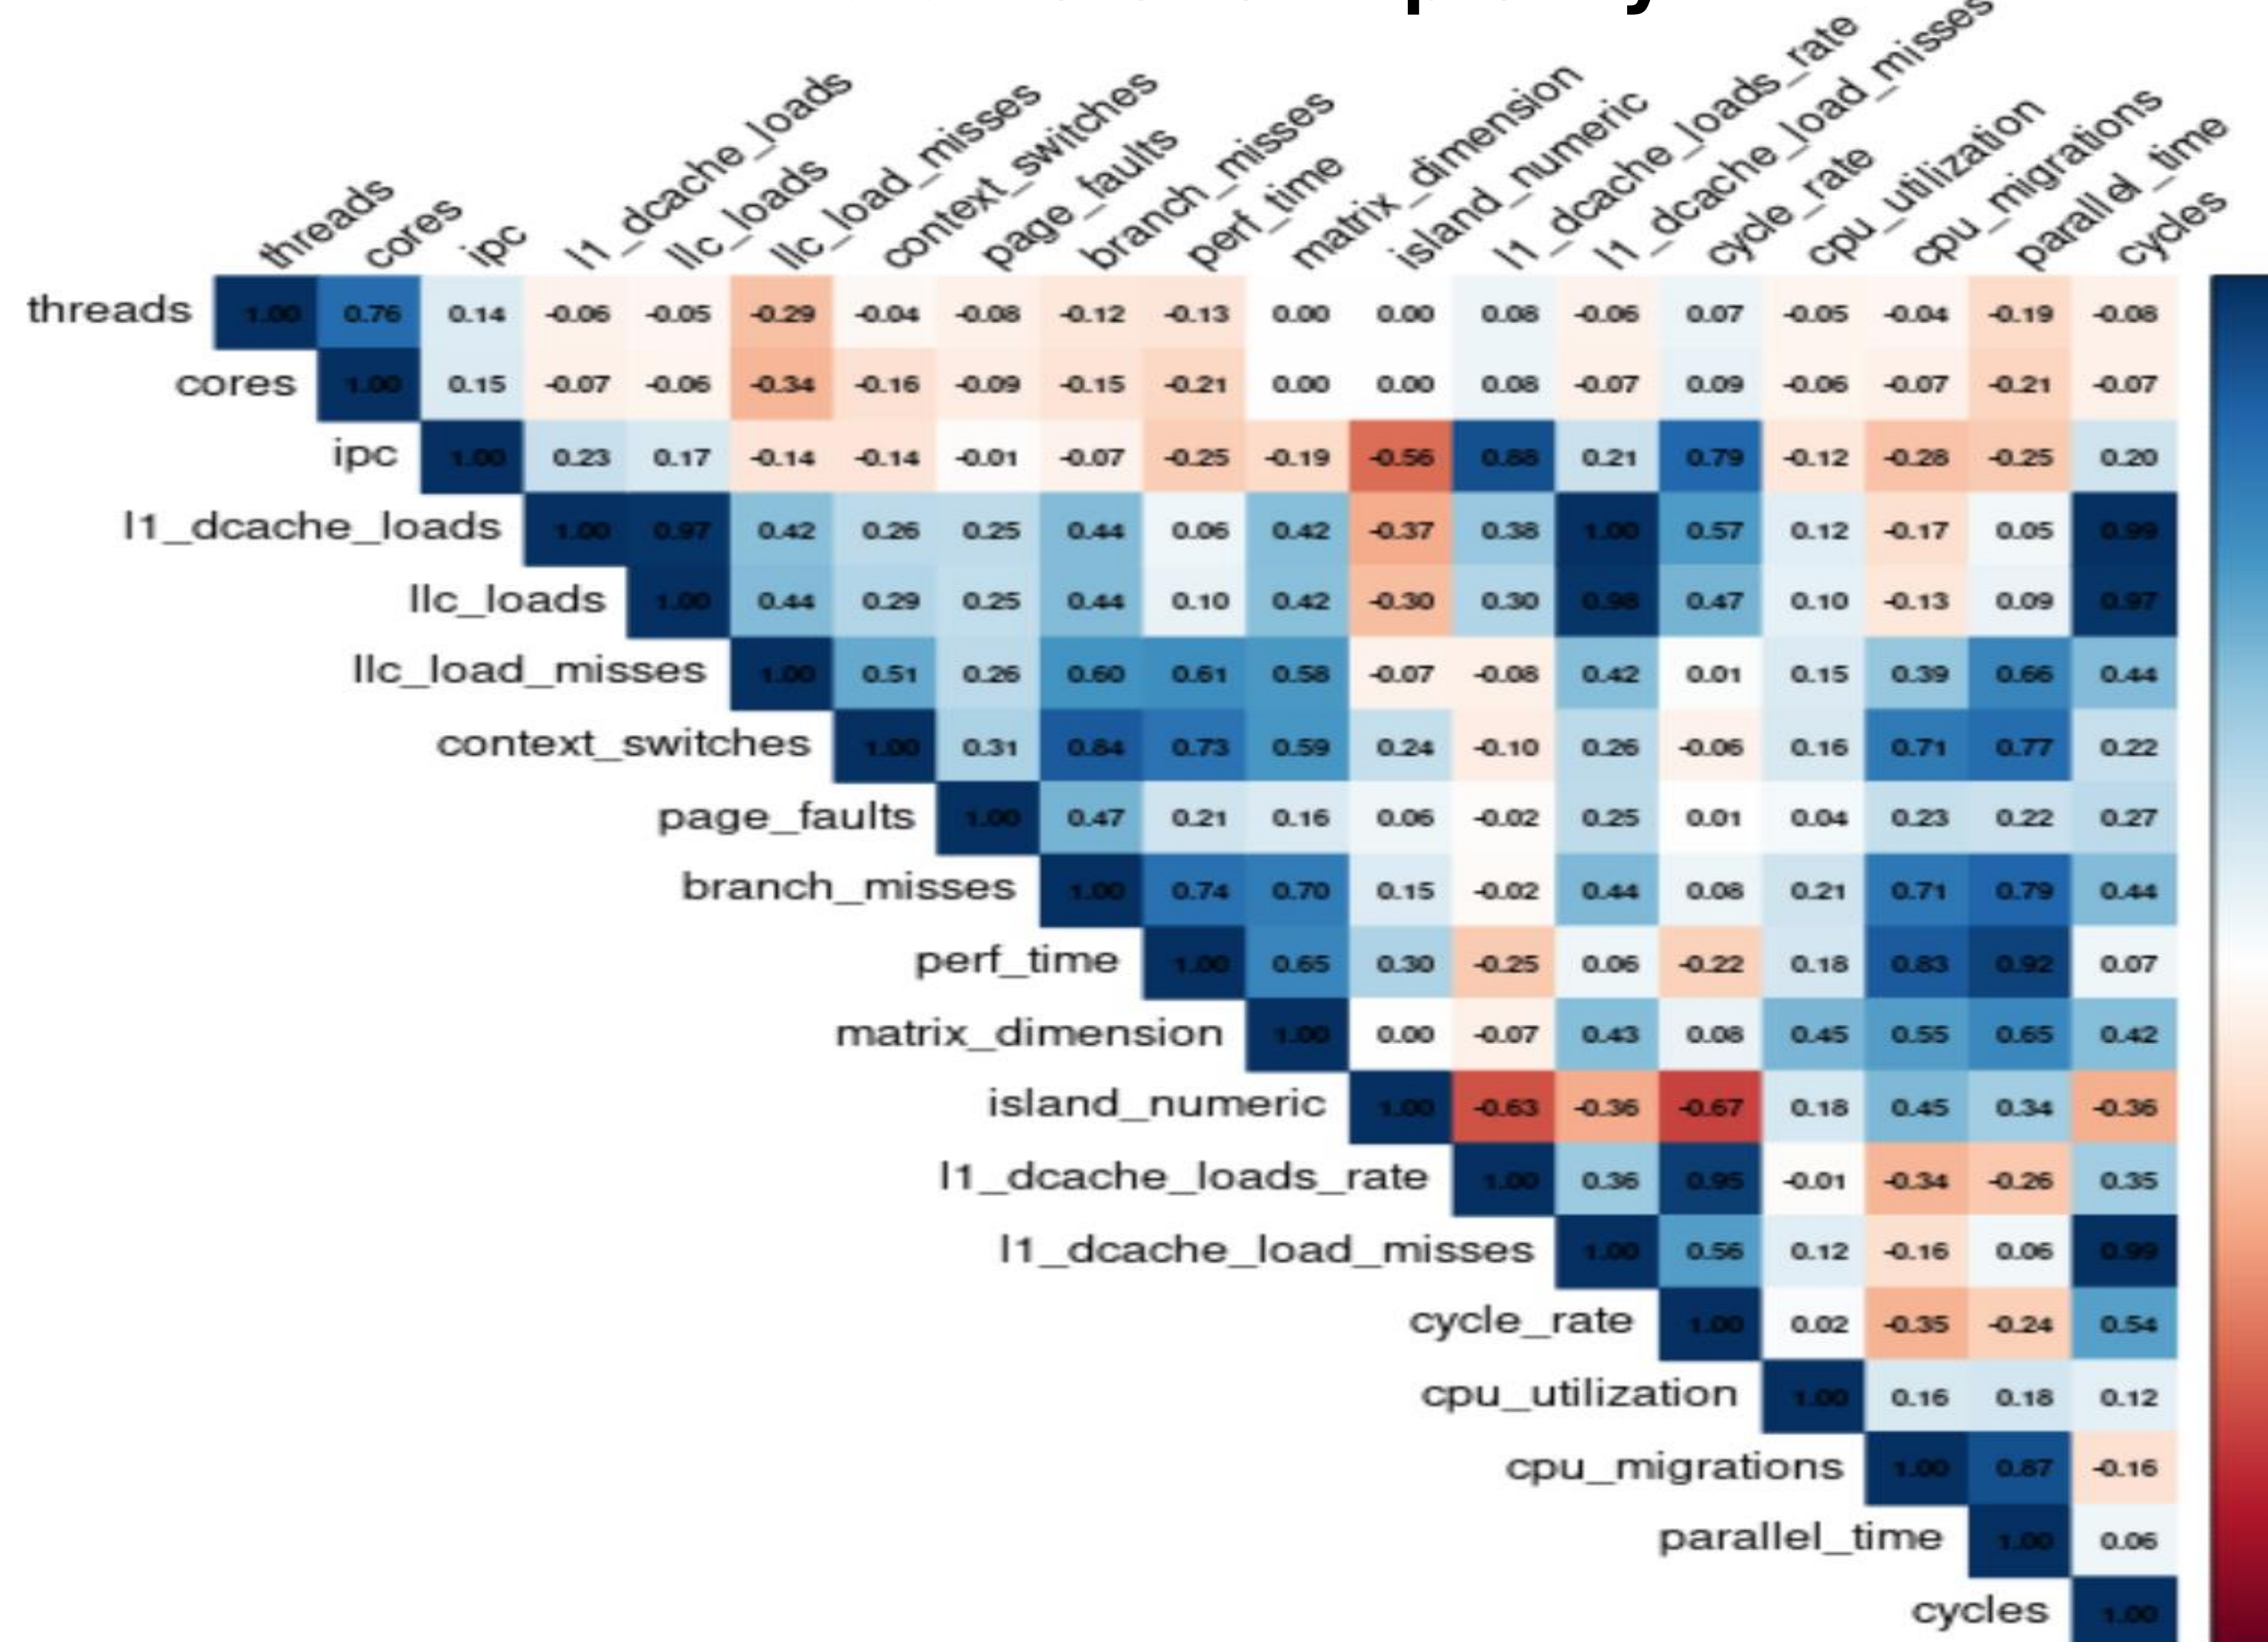
PROCESSING UNITS	8 PROCESSORS
THREADS	16 THREADS
ISLANDS	3 REGIONS
WORKLOAD PARAMETERS	128 OPTIONS



Performance Complexity



Parameter Complexity



Overview

- Heterogeneous platforms require unique resource mappings for execution requirements.
- New analyses and empirical evaluations are required to realize platform potentials.
- Using principled analysis techniques, we model and evaluate platform performance empirically using structural equation modeling¹ (SEM) methods in the context of the Orange Pi 5 platform.
- We introduce new abstractions and coordination models to navigate the feature-rich landscape reshaping edge-computing and embedded systems architectures.

High-Level Goals

- Expressive models** that accurately capture heterogeneous platform performance.
- Systematic identification** of important performance parameters.
- Equivalence relationship determination** between performance parameters.
- Generalizable performance models** across components and platforms.

Hypotheses

- Discoverable structures exist within workload-to-platform relationships.
- Equivalence relationships exist between components and platforms.
- Generalizable models capturing heterogeneous performance characteristics exist.
- Invertible generalizable models capture parameter characteristics.

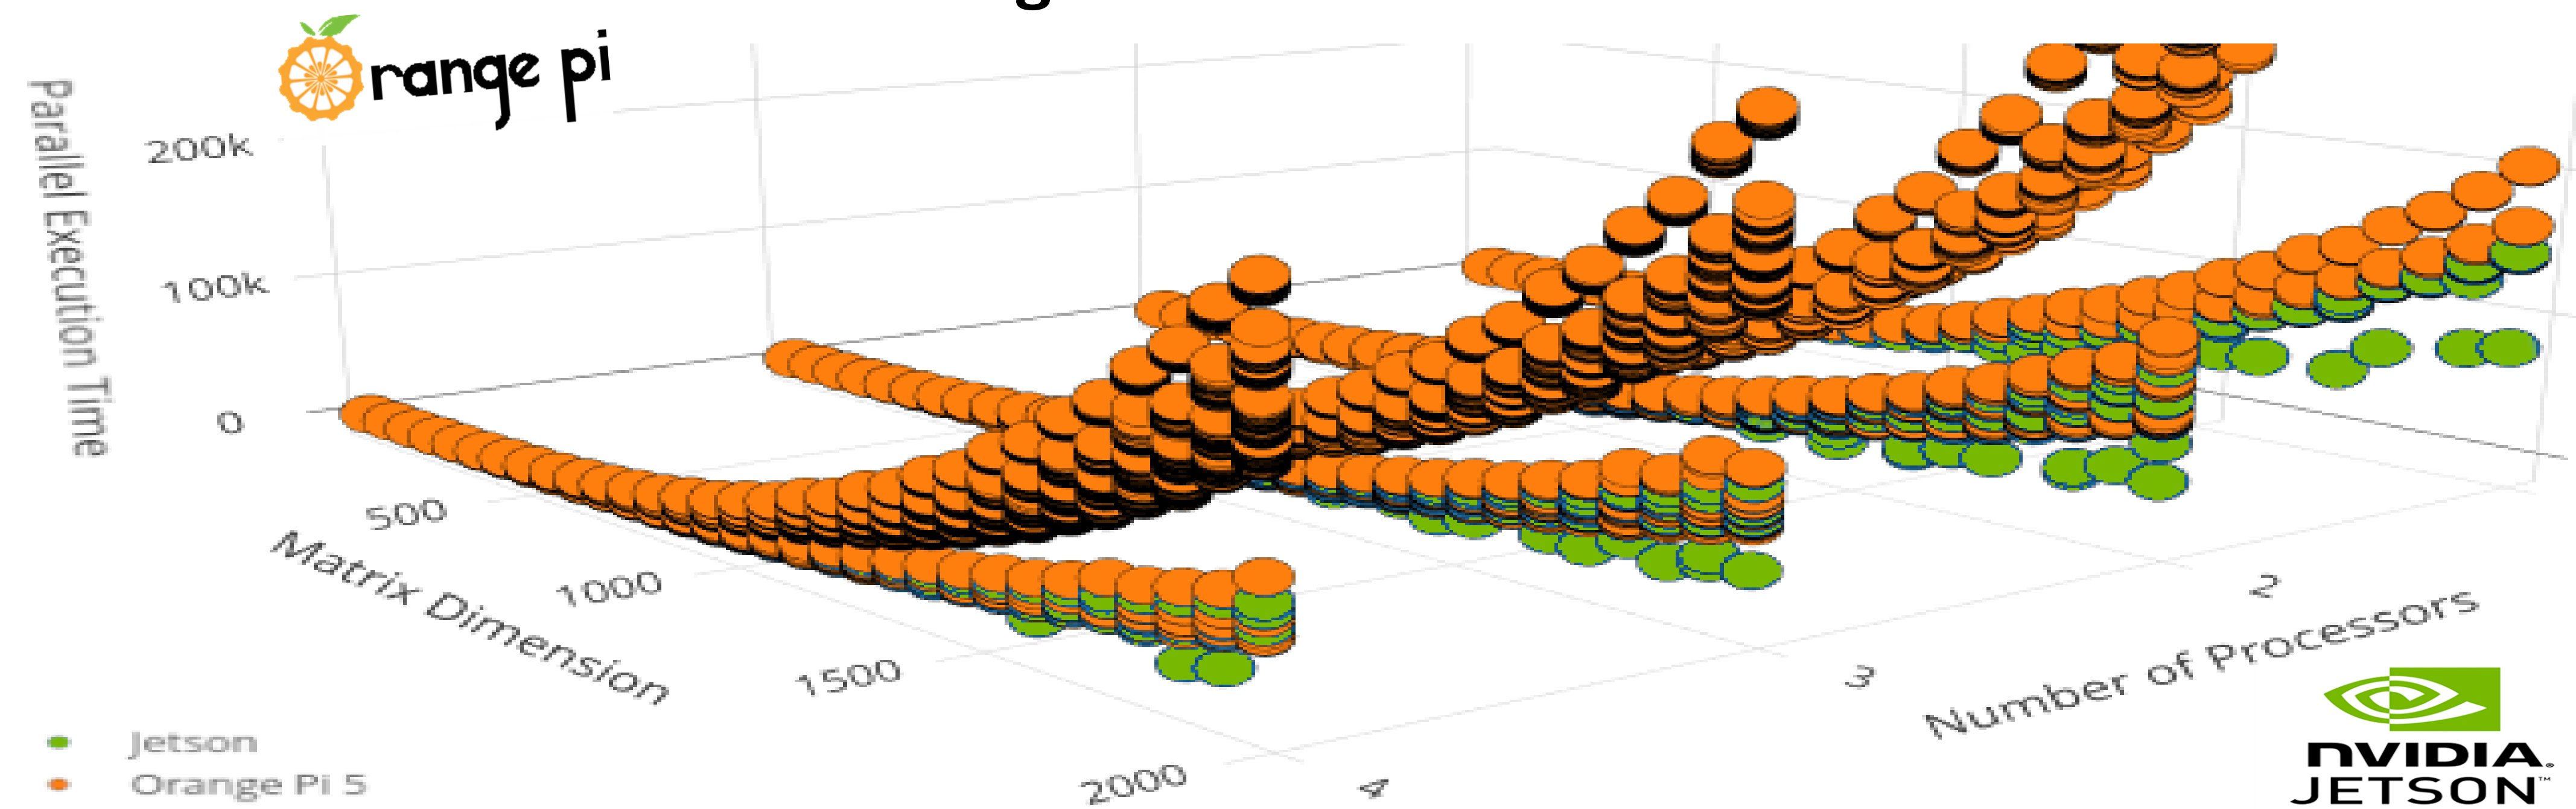
Findings

- Strongest negative impact** on execution time: Little Island + Matrix Dimension.
- Problem Complexity grows exponentially** with Island + Matrix Dimension interactions.
- Universal diminishing returns** with negative incremental gains for higher core counts.
- Optimal core count is highly dependent** on task characteristics, not simply core volume.

Conclusions

- Heterogeneous design-space cannot effectively be navigated by two-dimensional analysis.
- Platform assessments in higher dimensions incorporating parameter interactions will be necessary to determine precise platform performance.
- The findings support the need for new modeling techniques for heterogeneous platforms.
- These approaches can be extended to a range of platforms, configurations, and scenarios by leveraging insights garnered from the hardware-agnostic platform attributes.
- Startling similarities exist in the performance behavior of the NVIDIA Jetson platform.

Orange Pi 5 vs NVIDIA Jetson



Platform Modeling

Symbol	Description
$a_i, 1 \leq i \leq 3$	Cubic coefficients
M	Matrix dimension
N_B	Number of Big cores
N_L	Number of Little cores
R_{BL}	Big-to-Little ratio
S_{WS}	Working set size
T_e	Execution time

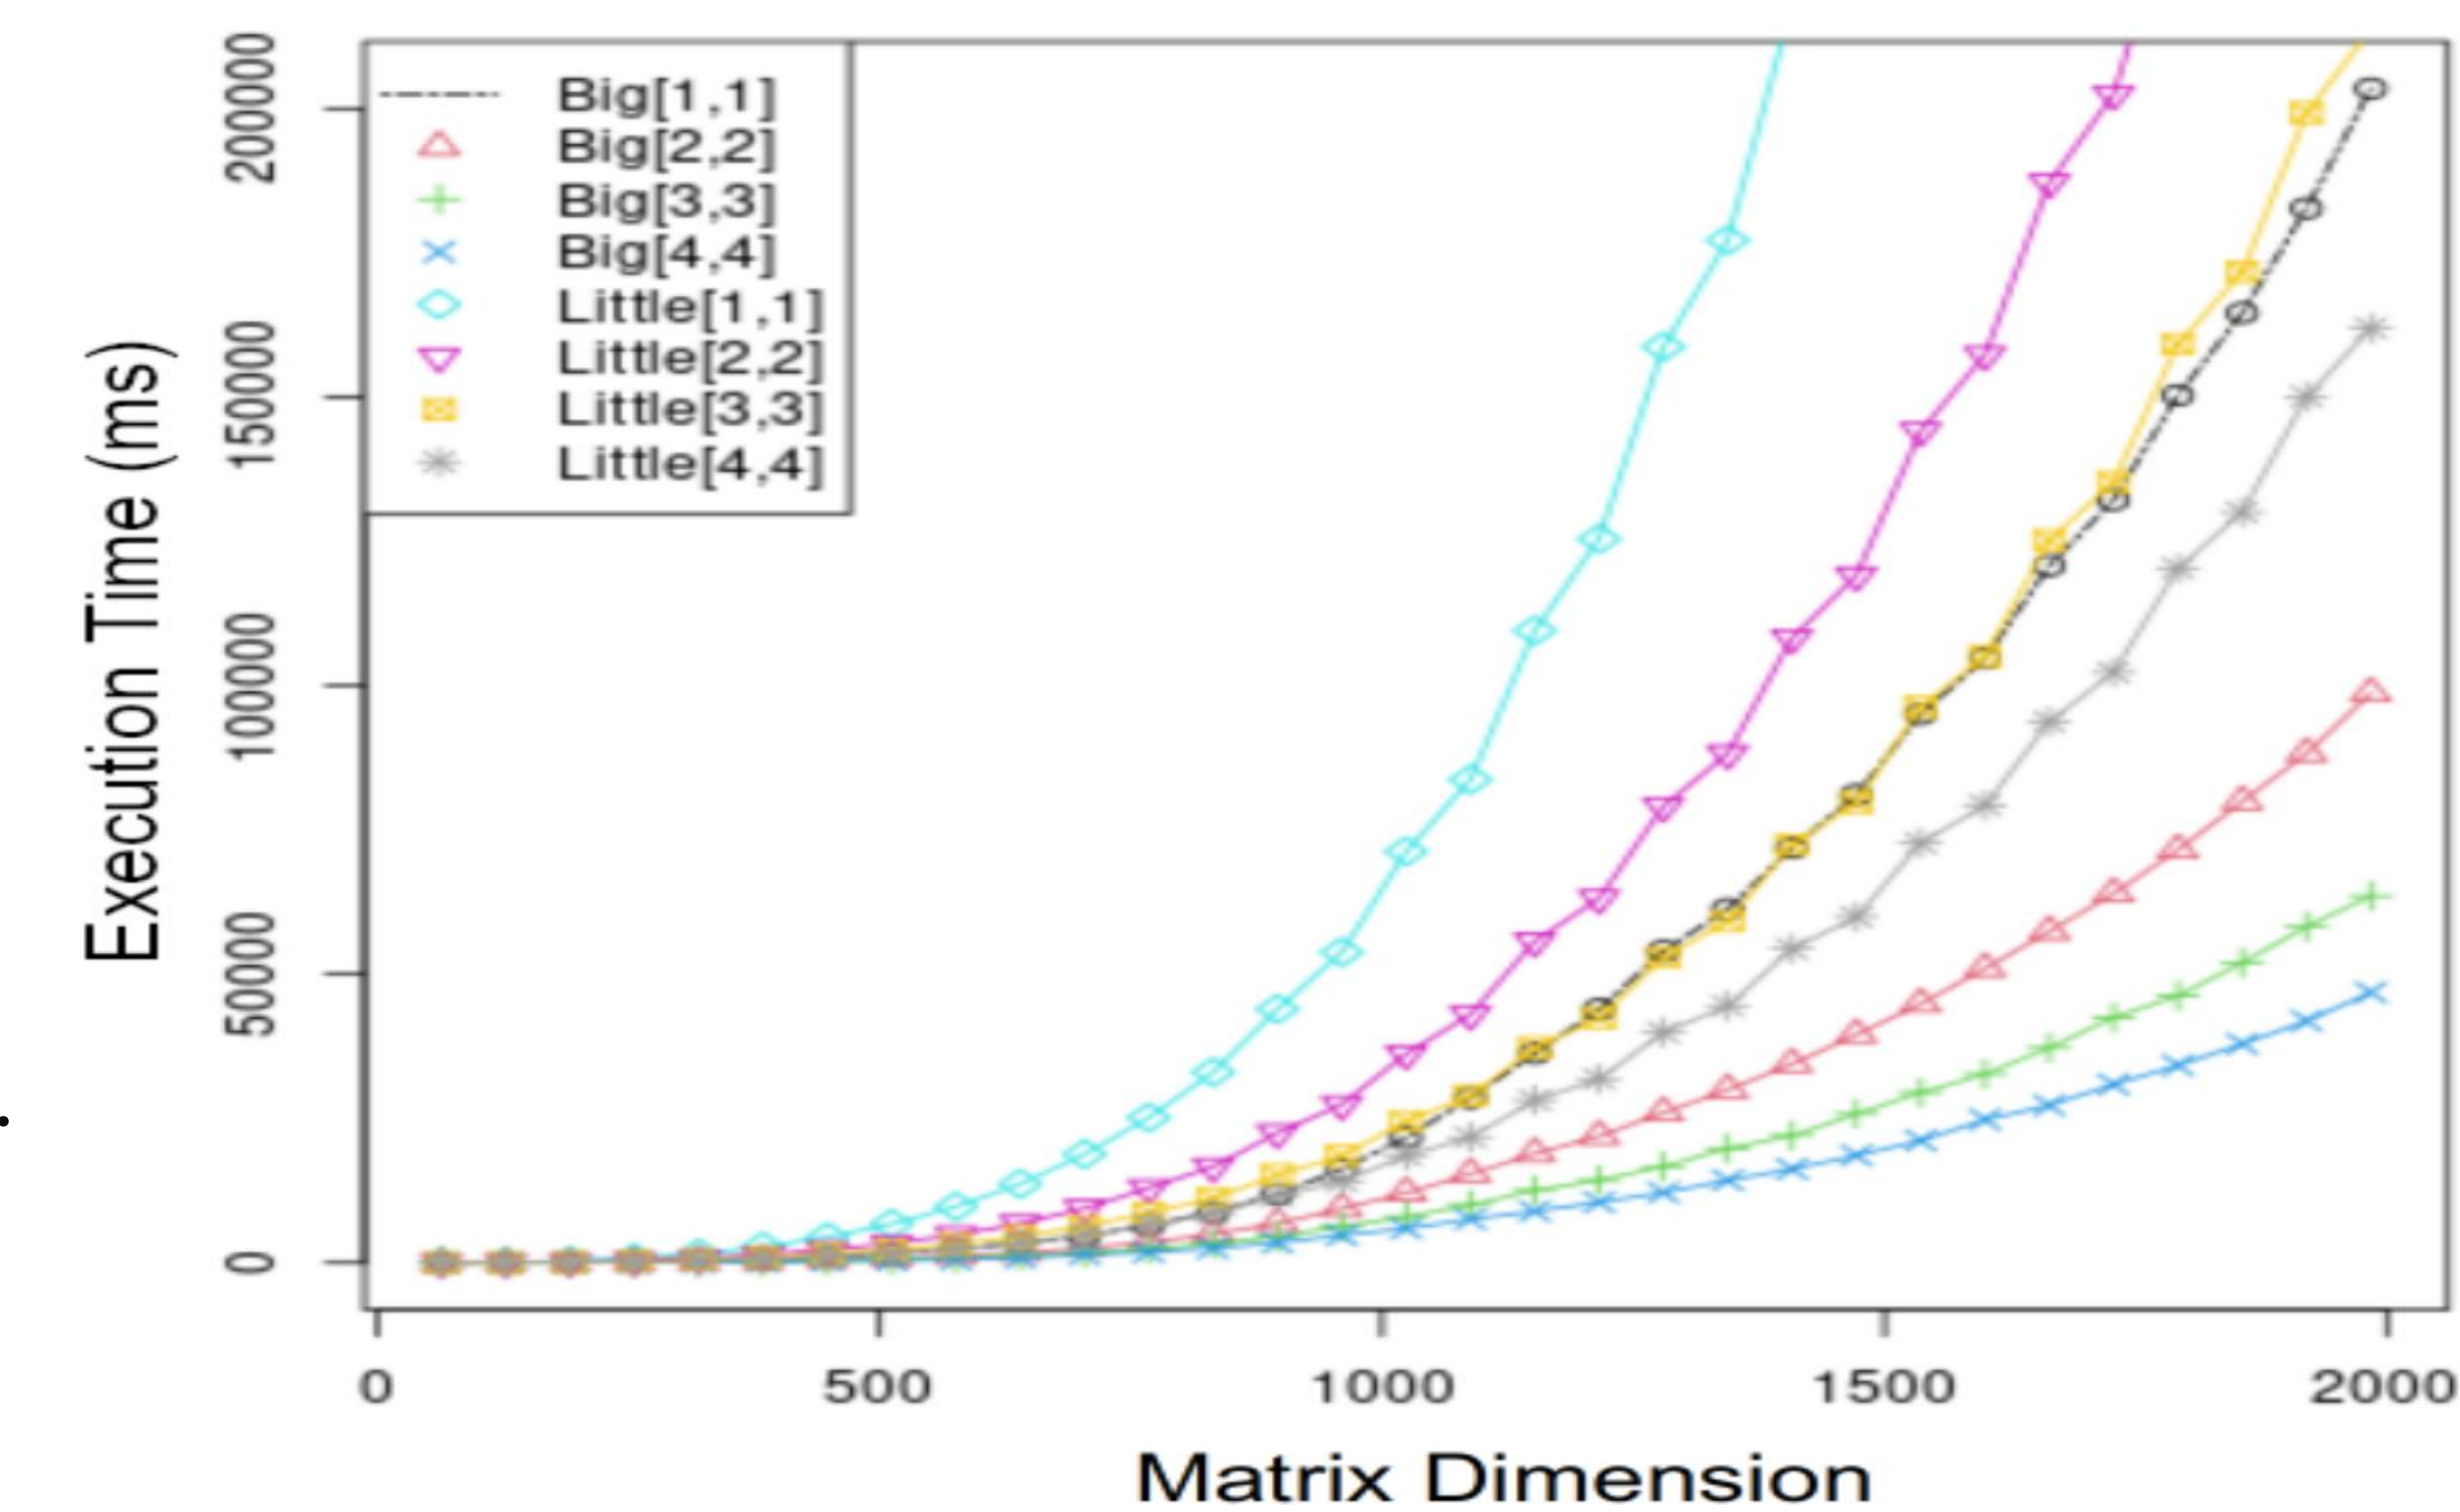
$$R_{BL} = \begin{cases} R_{BL.Low} & \text{if } M < S_{WS} \\ R_{BL.High} & \text{if } M \geq S_{WS} \end{cases}$$

$$T_e = \frac{a_3 M^3 + a_2 M^2 + a_1 M}{N_L + \frac{N_B}{R_{BL}}}$$

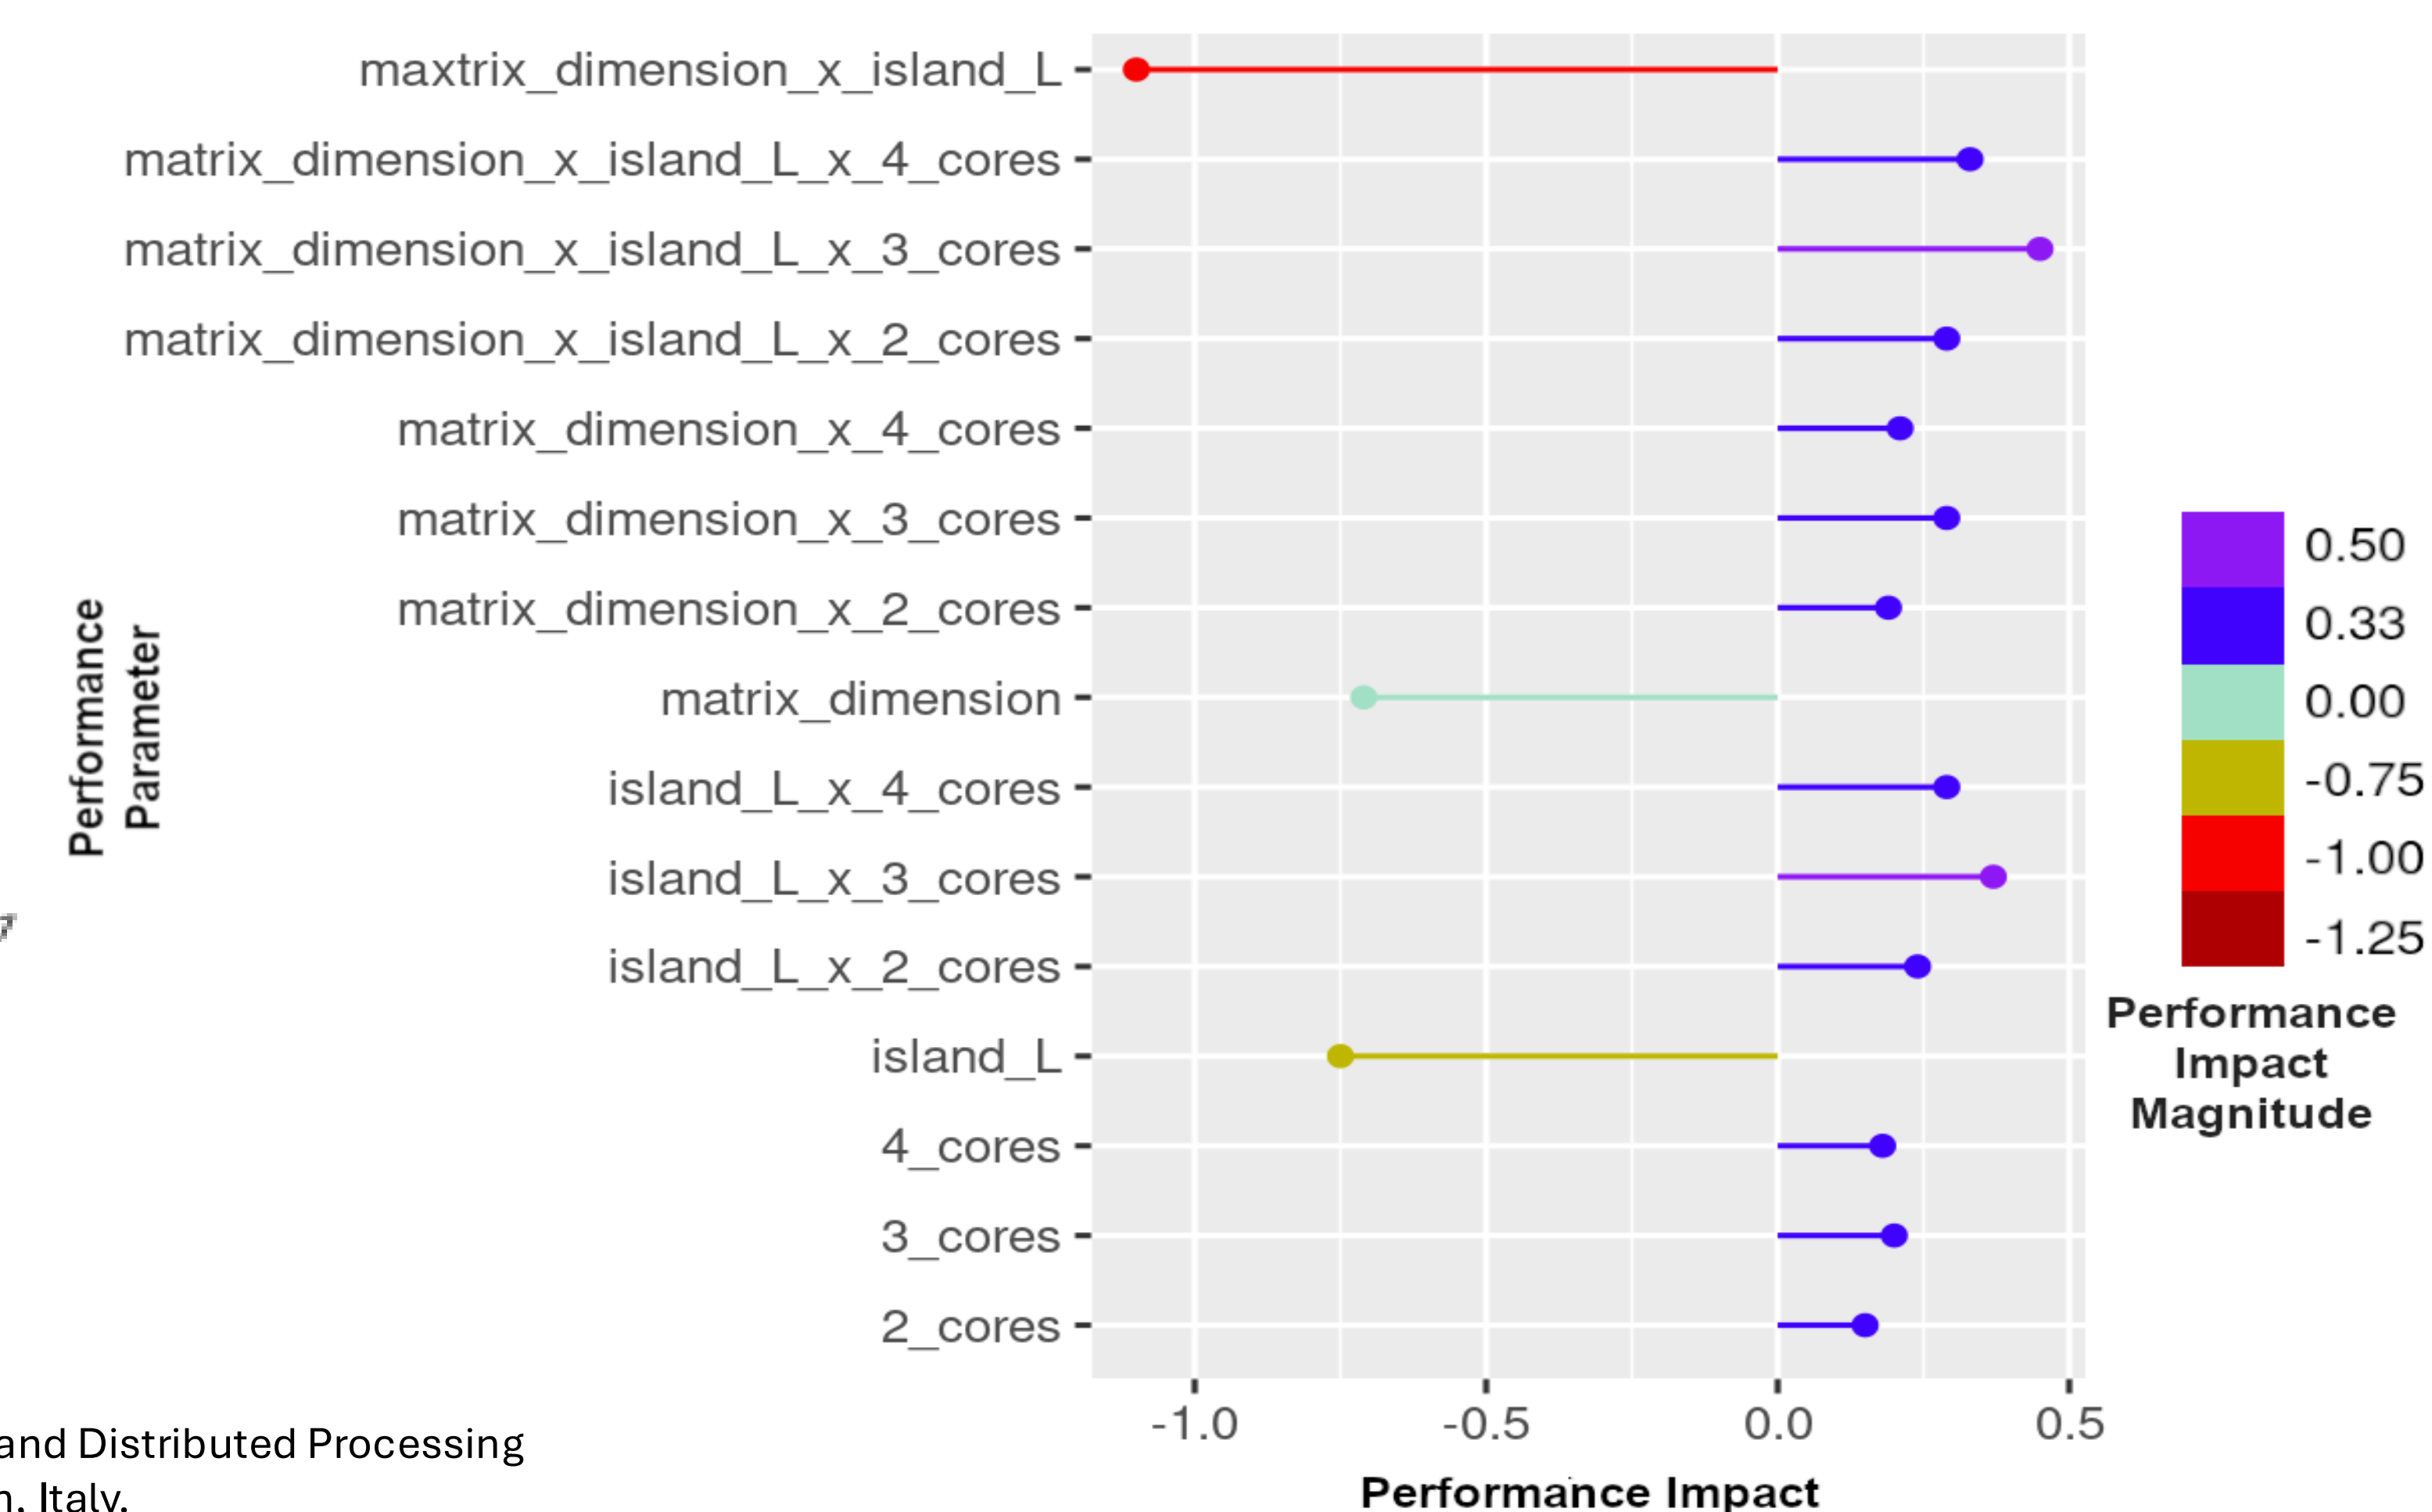
Performance Normalization

Configuration	Speedup	Average Ratio	Standard Deviation
L[1,1]	100%	1	0
L[2,2]	200%	0.50	0.010
L[3,3]	300%	0.33	0.005
L[4,4]	400%	0.25	0.004
B[1,1]	300%	0.32	0.089
B[2,2]	600%	0.15	0.022
B[3,3]	1000%	0.10	0.029
B[4,4]	1400%	0.07	0.016

Performance Prediction²



Parameter Effects



[1] Kline, Rex B. Principles and practice of structural equation modeling. Guilford publications, 2023.
 [2] Steven D. Harris, Roger D. Chamberlain, and Christopher D. Gill. Performance modeling of non-uniform heterogeneous platforms. In Proc. of IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), pages 597--607, June 2025. Presented at 27th Workshop on Advances in Parallel and Distributed Computational Models (APDCM), Milan, Italy.