WORK Eta VIII 2025

WORKSHOP EM MODELAGEM NUMÉRICA DE TEMPO, CLIMA E MUDANÇAS CLIMÁTICAS UTILIZANDO O MODELO ETA

25 a 29 de agosto ONLINE

INPE

Eta model V1.4.2 on raspberry Pi 5: ARMv8 Performance and Viability

Juan Carlos Tufino[1], Marco Bojorquez[2], Adrian Huerta[3]

[1] jtufinobernuy@gmail.com, [2] marco.bojorquez@unifesp.br, [3] adrhuerta@gmail.com

INTRODUCTION

Numerical weather prediction (NWP) has traditionally depended on high-performance computing resources, limiting its accessibility in environments with budget constraints (WMO, 2020). The Eta regional model, widely used in atmospheric research for its accuracy in complex terrains, faces this challenge particularly in developing countries. Previous implementations on low-cost hardware (Raspberry Pi 4) demonstrated the feasibility of running models such as WRF (Mendoza Uribe, 2023), although with significant technical limitations. The Raspberry Pi 5 (RPi5) represents a substantial improvement with its enhanced ARMv8 architecture, native support for high-performance storage (Raspberry Pi Foundation, 2023), and optimized energy efficiency, offering new opportunities for specialized simulations. This study systematically evaluates the operational feasibility of Eta v1.4.2 on the RPi5, providing a comparative performance analysis that positions this platform as a cost-effective alternative for educational applications, prototyping, and decentralized monitoring networks in vulnerable regions.

OBJECTIVE

a) Evaluate the feasibility, performance of running the Eta v1.4.2 model on a Raspberry Pi5 (ARMv8 architecture).
b) Compare the simulation results of the Raspberry Pi 5 with those obtained on a conventional laptop (Intel Core i7-1360P), considering runtime, energy consumption, and acquisiton cost.

METHODOLOGY

A comparative approach was implemented through simulations of the Eta v1.4.2 model, compiled for ARMv8 architecture, using input data from the Global Forecast System (GFS) with a spatial resolution of 0.25° and a temporal resolution of 6 hours. Six homogeneous domains were configured over Peruvian territory (82 × 74 grid points), initialized on July 30, 2025, at 12:00 UTC, producing 72-hour forecasts with an output spatial resolution of 8 km and an hourly frequency. Performance was evaluated on two Raspberry Pi 5 (RPi5) configurations: using 3 and 4 cores of the Broadcom BCM2712 processor, optimized with NVMe M.2 SSD storage via the official HAT+, and on a laptop with an Intel Core i7-1360P processor, while maintaining identical simulation conditions in both cases (input data, domains, and configuration parameters).

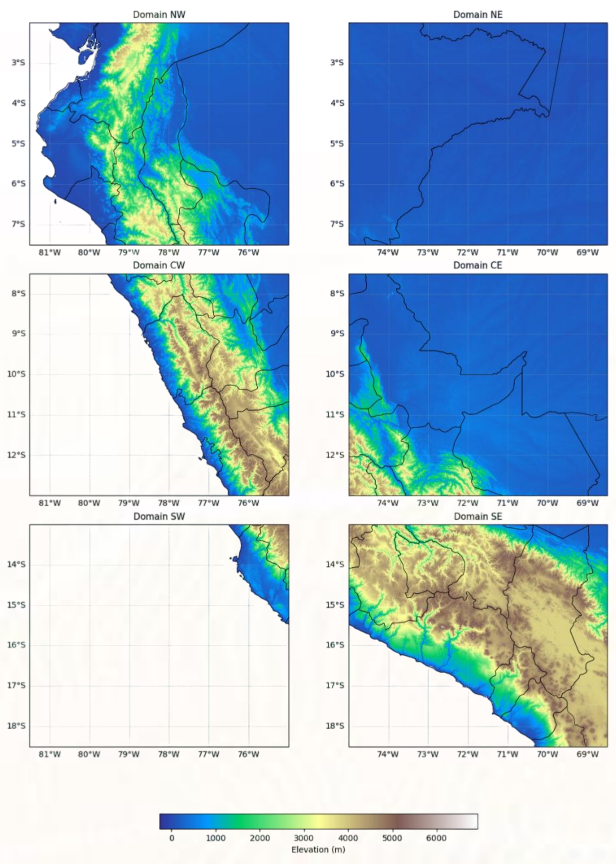

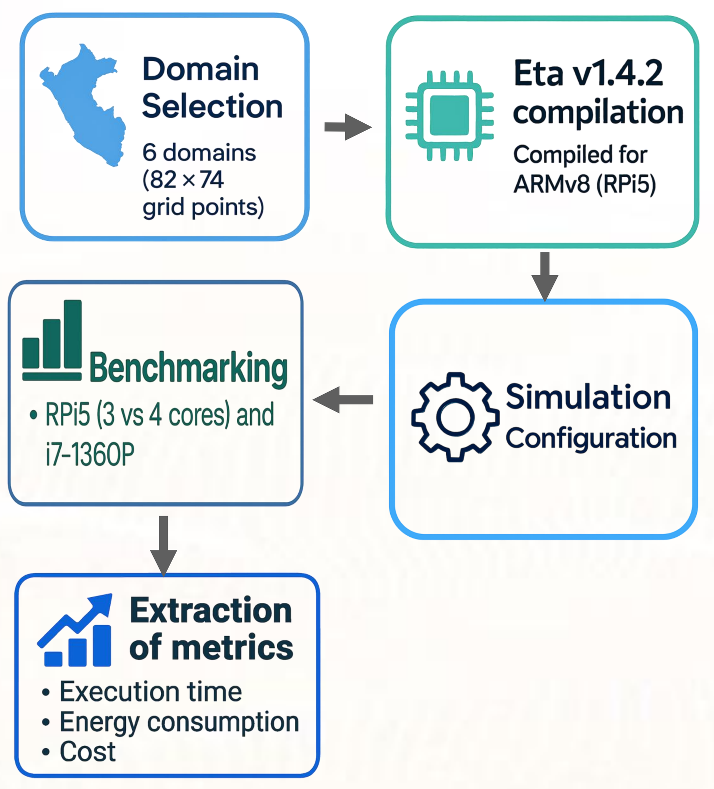

RESULTS

Figure 1. Execution Time Comparison by Domain and Hardware
Execution times on the **Raspberry Pi 4** are considerably higher than on the laptop. The **Raspberry Pi 5** consistently reduces execution times across all domains, showing clear generational improvements.

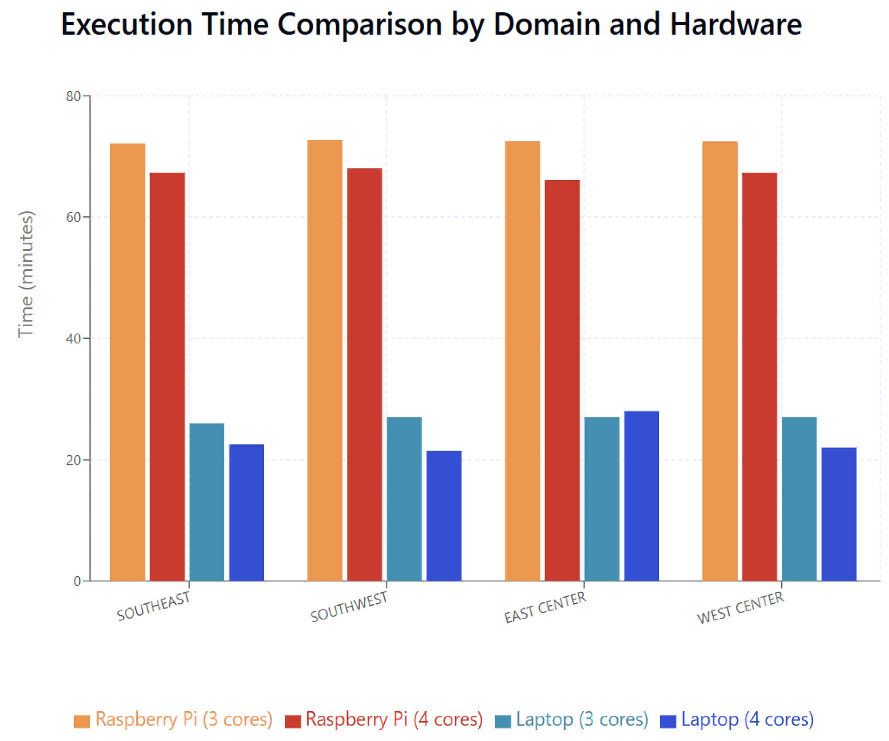

Figure 2. Average Execution Time by Device Configuration
The laptop achieves the fastest averages (<60 min). The **Raspberry Pi 5 (4 cores)** reduces its average to **67.8 min**, a significant improvement over the Pi 4 (>110 min).

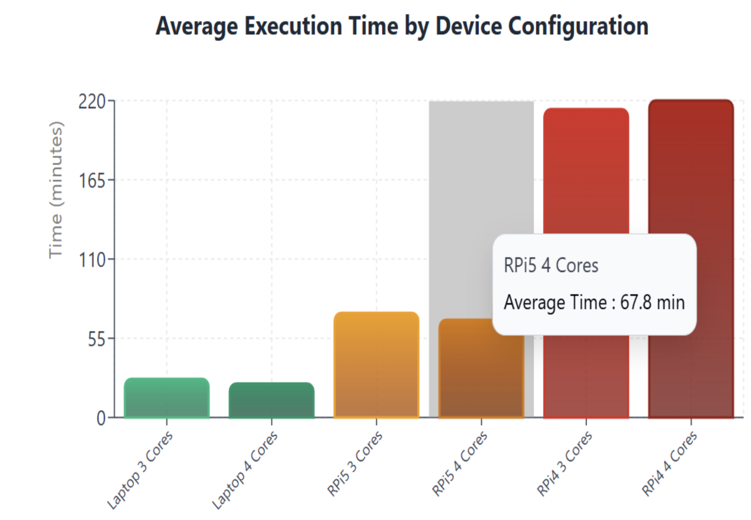

Figure 3. Train Maya Domain Comparison – Raspberry Pi 4 vs Pi 5
For the Train Maya domain, the **Raspberry Pi 5 (4 cores)** completes the run in **1h 7min**, nearly halving the execution time compared to the Pi 4.

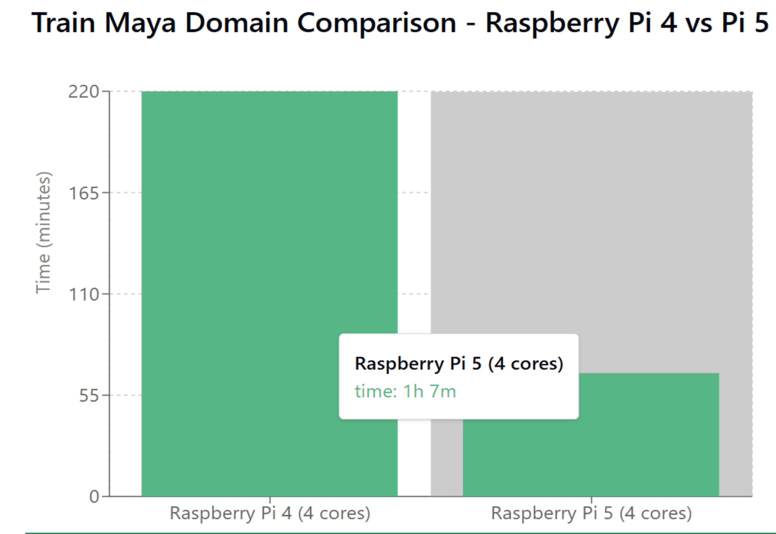

Figure 4. Performance Improvements (%)
The **Raspberry Pi 5 vs Pi 4 (4 cores)** shows the largest gain, with a **69.2% improvement**. Laptops remain the fastest, but the Pi 5 proves to be a viable and efficient low-cost alternative.

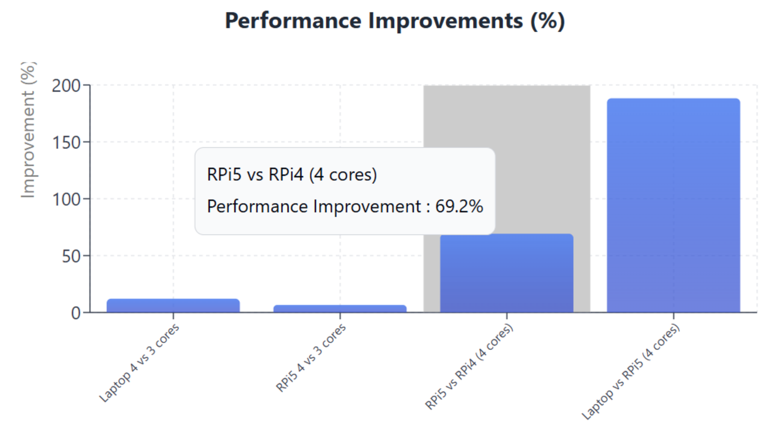

CONCLUSIONS

- Acceptable performance compared to conventional equipment: Although the RPi5 was 63.2% slower than a laptop with an Intel Core i7-1360P processor, its cost represented only 22% of the investment required for the laptop, demonstrating a high cost–benefit ratio.

- Improvement over previous generations: Compared to the Raspberry Pi 4, the RPi5 reduced execution times by up to 69.2%, reflecting a significant leap in computing capacity between gen of an NVMe M.2 SSD reduces I/O bottlenecks, while the 4-core configuration is optimal (7.08% faster than 3 cores). These optimizations are essential fo complex domains such as the Andes.

REFERENCES

Alarcón, J. et al. (2023). Implementación del Raspberry Pi como computadora alternativa de bajo costo para el aprendizaje remoto.
Mendoza Uribe, I. (2023). Implementación del modelo WRF sobre Raspberry Pi 4 para la zona del Tren Maya
Raspberry Pi Foundation (2023). Raspberry Pi 5 Technical Specifications.
WMO (2022). *A Science and Technology Vision for WMO*. WMO Bulletin, Vol. 71 (1)