A Distributed Framework for Matrix Multiplication on a Cluster Raspberry Pis

A spectrum of computing machines is available today. On one side, we have costly, high-performance "brawny" machines designed for AI, cloud computing workloads, etc. On the other hand, we have simple board computers(SBC) such as Raspberry Pi. These "wimpy" nodes are cheap, affordable, small, and energy-efficient [1, 4]. However, wimpy nodes' performance is not comparable to brawny machines. Traditionally, we have placed the brawny at the top of the hierarchy due to their high performance. Consequently, wimpy nodes are at the bottom of such a performance-centric hierarchy.

In this work, we do the opposite by going against the grain and building a hierarchy defined using isolation, scalability, redundancy, cost, and environmental impact. We observe the traditional performance-centric hierarchy inverses, placing the wimpy nodes at the top. We observe that these wimpy nodes could be used to build lightweight, distributed systems, enabling new opportunities. For example, a single NVIDIA DGX A100 consumes more than 200 times more energy than an RPi 5 when idle [5, 8], highlighting the inefficiencies of centralized computing for tasks requiring on-demand responsiveness. Furthermore, for the price of one of a single brawny machine, we could purchase tens to hundreds of RPi devices. Despite a single RPi's performance may not be comparable to that of an A100, a cluster of RPis comprising tens or hundreds may offer equal or more performance. For decades, ensuring stronger performance isolation has been a focus of many researchers on a single machine used in shared environments such as data centers [2, 7] . A cluster of RPis would offer ample isolation as each workload would run on a separate machine.

Through this work, we raise a central question: Can a large cluster of inexpensive, low-power Raspberry Pi devices perform comparably to a single brawny server? We aim to demonstrate the viability of wimpy-node clusters for distributed computing applications. Although wimpy-node clusters introduce additional communication overhead, strategies such as intelligent workload distribution and optimized data transfer protocols help mitigate these challenges, enabling efficient parallel execution [3].

We discuss where such wimpy-node clusters would fit in real-world scenarios. In particular, we highlight how edge computing environments may benefit due to the distinct benefits in efficiency, deployment, fault tolerance, and cost [6]. Notably, we pick the matrix multiplication workload as it is one of the widely executed computations in edge computing environments. By carefully running measurement experiments, we model the performance of a single device considering the advances and optimizations the recent model of the RPI offers. Then, we focus on designing the distributed frameworks that can handle the matrix multiplication of larger matrices on a cluster of RPIs while exploring the key trade-offs, including the impact of network latency and workload balancing. We evaluate the performance of the framework on a cluster of 10 RPI 5s, whose cost is similar to a conventional desktop machine. We observe comparable performance offered by the cluster of wimpy nodes, demonstrating that deploying such clusters as a replacement for traditional brawny servers is a viable option.

References

[1] David G. Andersen, Jason Franklin, Michael Kaminsky, Abhinav Phanishayee, Lawrence Tan, and Vijay Vasudevan. Fawn: A fast array of wimpy nodes. In *Proceedings of the 22nd ACM Symposium on Operating Systems Principles (SOSP)*, pages 1–14, Big Sky, MT, USA, 2009. ACM.

[2] Sebastian Angel and Alejandro López-Ortiz. Ensuring performance isolation in multi-tenant data centers. In *Proceedings of the 11th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. USENIX Association, 2014.

[3] Jeffrey Dean and Sanjay Ghemawat. Mapreduce: Simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.

[4] Raspberry Pi Foundation. *Raspberry Pi 5 Product Specifications*, 2023.

[5] Jeff Geerling. New 2gb pi 5 has 33% smaller die, 30% idle power savings, 2024.

[6] Sangmin Lee and Elena Garcia. Offloading computational tasks to the edge: A review of current practices. *International Journal of Edge and Cloud Computing*, 5(2):78–90, 2022.

[7] Alan Shieh, Srikanth Kandula, Albert Greenberg, and Changhoon Kim. Seawall: Performance isolation for cloud datacenter networks. In *2nd USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 10)*, 2010.

[8] Yifan Zhu and Ang Li. Improving gpu energy efficiency through an application framework and runtime system. In *Proceedings of the 19th European Conference on Computer Systems (EuroSys '24)*, pages 769–785, 2024.