

Introduction

In the ever-evolving landscape of scientific software and computing, continual adaptation to new platforms is crucial for accurate results and optimal performance. Yet, software testing, especially on HPC, remains a challenging endeavor [2]. Amidst the pursuit of faster and more energy-efficient processors, the AWS Graviton3 emerges as a serious candidate, built on the ARM Neoverse-V1 architecture. While its demonstrated performance is noteworthy, questions arise about the reproducibility [1] of results compared to the 'traditional' x86 64 architecture. Our research delves into this inquiry, focusing on verifying results from HPC applications initially tested on x86 64, such as XDEM for granular particle simulation and a Transcriptomics Analysis Workflow for RNA sequencing. Beyond verification, we aim to quantify numerical differences and ensure the consistency of scientific conclusions across architectures.

Applications and Methodology

This study focuses on two distinct yet complementary applications: the **eXtended Discrete Element Method** (XDEM), [6] a versatile simulation framework for granular particle motion and thermal conversion, coupled with Computational Fluid Dynamics (CFD); and a **Transcriptomics Analysis workflow** (RNA-seq), involving genomic analysis steps relying on integer arithmetic. Assessing these on the AMD Epyc ROME 7H12 and ARM-based AWS Graviton3 platforms, we delve into the challenges faced by application developers in compiling, installing, and verifying software, scrutinizing the impact of different compilation flags on numerical results and comparing outcomes across processors.

Particle Trajectories in XDEM

In our study, we examine particle trajectories through a simple test involving 27 particles falling onto a plate and colliding for 1.7 seconds in simulation. Figure 1 illustrates the noticeable divergence in particle trajectories between Aion and AWS platforms after a few collisions. Figure 2 further compares results with different compilation flags,

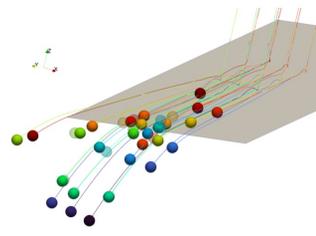


Figure 1: Comparison of particle trajectories between Aion (bright color) and AWS (light color).

revealing that achieving bit-to-bit identical outcomes between x86 64 and ARM is only possible by disabling FMA instructions. Notably, even when both processors use FMA, they fail to produce identical results.

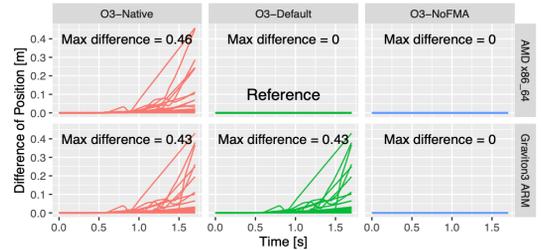
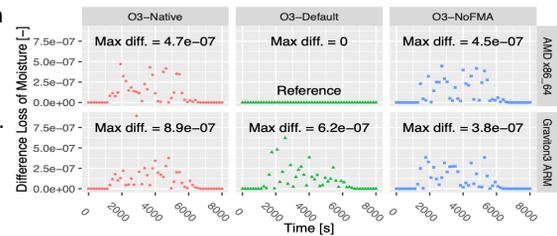
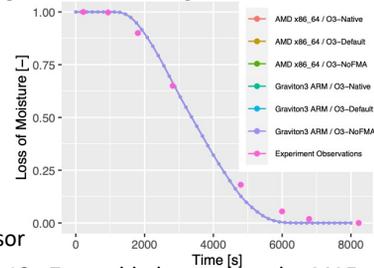


Figure 2: Difference of the positions of the 27 particles for different processor and compilation flags (in comparison with Aion/Default).

Biomass Drying with coupled XDEM-OpenFOAM

We explore biomass drying using XDEM-OpenFOAM with 2667 static particles traversed by a gas flow. In this figure, we calculate the loss of moisture and validate our simulations against experimental data, exhibiting indistinguishable results with a Mean Absolute Error (MAE) of $2.6e-2$. The figure below highlights processor differences at a magnitude of $10e-7$, notably lower than the MAE. Interestingly, disparities emerge during the transient period, while results align at the start and end of simulations.



Transcriptomics Analysis

For the Transcriptomics Analysis, all steps, from raw files to alignments, are strictly identical across architectures using Singularity images. This ensures consistent mean gene expression before contrast. The figure indicates minimal numerical differences at the differential expression step, impacting only two of the five metrics displayed. However, these variances do not alter gene ranks or fold change magnitudes, preserving the primary scientific conclusions. Disabling the FMA instructions will be investigated to further validate the exact reproduction of the results.

