

SkewEngine: Enhancing Performance of Intensive Calculations on Regular Mesh Data

Felipe Romero ^{*}, Pilar M. Ortigosa^{*}, Gerardo Bandera [†], Luis F. Romero[†]

May 30, 2023

Abstract

In various applications such as hyperspectral data manipulation, MRI data exploration, or visual basin identification in digital elevation models, performing arithmetic operations on each point of a data mesh that involve other points can lead to computationally intractable problems. This paper presents SkewEngine, a tool designed to improve the performance of intensive calculations on regular 2D or 3D data meshes, such as images, multispectral data volumes, or digital elevation models. SkewEngine addresses this problem by reorganizing the mesh in memory according to a preferred spatial direction, enabling more efficient execution of intensive calculations. It is demonstrated that SkewEngine offers significant speed improvements for a variety of test cases, suggesting its usefulness in a broader range of applications requiring intensive data processing on regular meshes.

In scenarios where the calculation of visual basins is not limited to a single observer’s viewpoint but needs to consider arbitrary paths or regions in the terrain, the computational complexity can increase significantly, reaching $O(N^{2.5})$. Even with lower precision, simple models can take months of CPU time to calculate. However, such problems fall into a category where the study parameters decay based on geometric distance within the data mesh, often following a quadratic pattern. Exploiting this characteristic, these problems are typically simplified by processing data in a discrete set of radial directions with respect to a study point, as the radii are closer to the point of origin, resulting in a finer mesh at the study point. However, in problems where spatial locality is critical, there is a need to transfer this spatial locality to the data itself, particularly in terms of memory storage. Consider an example where a filter needs to be applied to an image in all directions parallel to a specific arrow. It is evident that any computationally intensive algorithm like FFT would be more efficient if the data were aligned in memory, as shown in the deformed image on the right side of Figure 1.

This work proposes to analyze and exploit the benefits of memory reorganization, specifically demonstrating the significant benefits of biased data interpolation in highly complex problems.

Background: The SDEM Algorithm

In 2013, Tabik et al. [2] proposed an algorithm that considered radial dependence in visual basin calculation, reducing the complexity from $O(N^{2.5})$ to $O(sN^{1.5})$ by considering a

^{*}Dpto. Informática, University of Almería, fr@uma.es, ortigosa@ual.es

[†]D. Computer Architecture, University de Málaga, gbandera@uma.es, felipe@uma.es

discrete set of s directions around each point. By exploiting data alignment in a particular direction, the algorithm performs visual basin calculations for all points along that line simultaneously.

Recently, Romero et al. [1] introduced the sDEM algorithm, which further improves performance by considering two crucial aspects: 1) aligning data in memory to exploit spatial locality and 2) parallel processing of multiple lines simultaneously, including GPU utilization. This algorithm, called skew-DEM (sDEM), aligns the data in a specific direction, considering that the computational intensity of the subsequent operations with respect to the cost of the reorganization is going to be worth it, regarding the reduction of cache misses




Figure 1: Skewed model

Extension to Three Dimensions

Extending intensive data processing, with interdependence in all directions on a 2D regular data mesh, to three dimensions can be exceptionally costly. Hence, addressing this problem requires two similar premises to the 2D proposal: reducing the infinite number of directions to a discrete set and aligning data in memory. The paper proposes an elegant solution using the spherical Fibonacci spiral for discretization, obtaining quasi-equidistant points on the surface of a sphere. The distribution of these points enables aligning data in memory by leveraging the skewEngine algorithm, resulting in a cubic representation skewed perpendicular to the z-axis.

Conclusion

Aligning data correctly in computer memory is becoming increasingly important in designing efficient algorithms for intensive processing of structured data. This paper demonstrates the benefits of realigning regular mesh data in memory using the SkewEngine tool. The experiments conducted reveal significant performance improvements, indicating the potential of SkewEngine in various applications that involve intensive calculations on regular meshes.

References

- [1] A. J. Sanchez-Fernandez, Luis F. Romero, G. Bandera, and S. Tabik. A data relocation approach for terrain surface analysis on multi-gpu systems: a case study on the total viewshed problem. *International Journal of Geographical Information Science*, 35(8):1500–1520, 2021.
- [2] S. Tabik, E.L Zapata, and L.F Romero. Simultaneous computation of total viewshed on large high resolution grids. *International Journal of Geographical Information Science*, 27(4):804–814, 2013.