

Alternative Ephemeris Representations for Astrodynamical Simulations on Multi-Core CPUs and Accelerator Cards

Project Manager
Arya Mazaheri

Researchers
Fabian Schrammel

Principal Investigator
Prof. Dr. Felix Wolf

Project Term
2019 - 2019

Clusters
Lichtenberg Cluster Darmstadt

Software
Python

Additional Software
G++, CUDA, Bash

Institute
Parallel Programming

University
Technische Universität Darmstadt

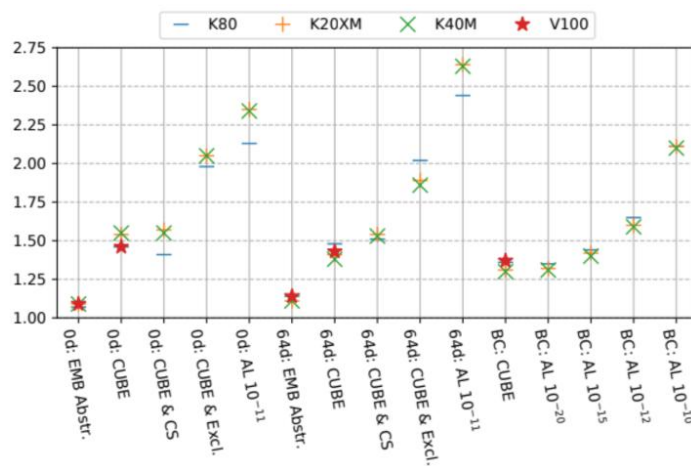


Figure 1: The speedup of GPU-efficient ephemeris models used by the listed cases and configurations. The cases are BepiColombo (BC), zero days (0d), 64 days (64d) timestamp range. The model features include CUBE format, EMB abstraction, cubic splines (CS), and the full exclusion of bodies, which are all included for the accuracy level (AL) configurations.

Introduction

Spacecraft missions of the European Space Agency (ESA) are required to meet specific probabilistic requirements regarding Space Debris mitigation and Planetary Protection of certain other celestial bodies to prevent contamination. Therefore, thousands of possible spacecraft trajectories are simulated on accelerators using CUDA. ESA recently developed cudajjectory which computes one trajectory per accelerator thread. Within, the gravitational impact by all major objects of the solar system is calculated. The ephemeris model provides polynomials to calculate the highly accurate position of each body at the simulated point in time. This model is designed for CPU applications and provides opportunity for heavy performance optimization.

Methods

The trajectory samples may progress at different simulation step sizes and then require different ephemeris data overloading the on-chip accelerator memory. First, the spatial locality of the ephemeris data structure is optimized for warp instruction accesses by the CUBE data format to minimize data loads.

Second, alternative polynomials are applied to cover longer simulation time periods with the same data and certain other astrodynamical optimizations are performed to reduce the overall computations using the ephemeris model. An alternative CUDA kernel execution scheme is also proposed to restrict the number of data loads per warp.

Results

Our experiments showed that by just using the CUBE data format, we could obtain higher speedups of at least 1.3x on all available Nvidia Tesla cards. The exclusion of specific celestial bodies from the model and use of cubic splines as an alternative polynomials give the expected reduction of both the algorithmic complexity and data accesses, while losing negligible accuracy. We achieve significant speedup on a real-world scenario between 1.31x - 2.11x, depending on the desired accuracy level. The alternative execution scheme achieves a similar speedup for one real-world scenario but behaves worse for different cases.

Discussion

The new data structure format is very beneficial for GPU applications and still gives room for small improvements. While the use of cubic splines further improves caching, it also increases the need for register spilling and thus results in only a small runtime improvement on top of the CUBE format. However, in combination with body exclusion, the model changes can boost each other. Thus, cubic splines are a valuable ephemeris model setting. The alternative execution scheme reduces the memory overload on the one hand but increases the idle time on the accelerator. It is a much more static execution approach and overall not beneficial.

Publications

Schrammel, F.; Renk, F.; Mazaheri, A.; Wolf, F.: Efficient Ephemeris Models for Spacecraft Trajectory Simulations on GPUs. In Proc. of the 26th Euro-Par Conference, Warsaw, Poland Volume TBA of Lecture Notes in Computer Science, Springer, August 2020
https://doi.org/10.1007/978-3-030-57675-2_35

Last Update: 2022-06-14 14:09