National Aeronautics and Space Administration

FY23 Strategic University Research Partnership (SURP)

Rapid GPU Trajectory Optimization With New ALTRO Constrained Trajectory Optimizer

Principal Investigator: Martin Lo (392), Jet Propulsion Laboratory, California Institute of Technology
Co-Investigators: Brian Anderson (392), Jet Propulsion Laboratory, California Institute of Technology
Zachary Manchester, Fausto Vega (Carnegie Mellon University)

Objectives: Our objective is to speed up the trajectory optimization process by 2-3 orders of magnitude through Co-I Manchester's new ALTRO optimization technology that can efficiently leverage the thousands of processor cores available on modern graphics processing units (GPUs). We achieve this goal by developing parallelized versions of three key solver components: (1) Implementing high-order integration methods with automatic differentiation on a GPU Evaluation of spacecraft dynamics, Jacobians, and Hessian-vector products can be parallelized across timesteps in a trajectory. A 1000x speedup of these functions, which typically account for 70% of solve-time, is possible.(2) Implementing parallel backtracking line-search on a GPUA core part of nonlinear optimization solvers that is currently performed serially: By evaluating hundreds to thousands of step-lengths simultaneously, we can make more progress at each iteration while reducing search-time, resulting in a 2-3x speedup.(3) Developing a custom sparse linear-system solver on a GPU. This component is the core of any Newton-based solver algorithm, and typically accounts for 30% of solution time in space trajectory applications. It is the most difficult to parallelize, and requires fundamentally new matrix factorization algorithms. An efficient GPU implementation should achieve a 10x-100x speed up.

Background: The trajectory is a key, early element for mission design because every subsystem depends on it; telecom, thermal, propulsion, science coverage all require the orbit for their design. However, in TeamX design sessions, the trajectory design is theonly discipline that is unable to support concurrent engineering. Trajectory for tours and landings requires days, weeks, sometimes months for initial design and make changes. Perhaps the most important technology to speed up this process is the optimization of the trajectory such as our proposal.
The trajectory optimization technology is also a key for automation such as needed in machine learning algorithms and for autonomous navigation. The desired trajectory is assembled from a sequence of predetermined paths which are "glued together" by optimizing the differences at the patch point at the same reducing the maneuver needed to operate this trajectory. So for the development of AI methods and automation of mission design, our proposed trajectory optimization work is an essential technology.

Approach and Results: Recently, Co-I Manchester and his Group at CMU have developed a novel solver called ALTRO (Augmented Lagrangian Trajectory Optimizer, Howell et al. 2019) that blends features of both DDP and direct methods to take maximum advantage of the structure of trajectory optimization problems while achieving superior performance in terms of both numerical robustness and speed. When compared across a wide range of nonlinear trajectory optimization problems on platforms like spacecraft, quadrotors, robotic arms, and cars, ALTRO typically offers an order-of-magnitude speed improvement over generic NLP solvers like SNOPT and IPOPT, with speed-ups in excess of 50x on some problems. However, like the classical methods it is derived from, ALTRO is not well suited to parallel implementation on a GPU. To fully take advantage of modern GPU hardware, we will address three key serial bottlenecks in the ALTRO algorithm: 1) evaluation of dynamics and derivatives, 2) backtracking linesearch, and 3) sparse matrix factorization. These three tasks will be carried out sequentially, with each offering a significant improvement over the current ALTRO implementation, culminating in a fully GPU-native trajectory optimization solver that is two orders-of-magnitude faster than current CPU-based solvers.

Significance/Benefits to JPL and NASA: Trajectory optimization is perhaps the single most important tool for generating trajectories for JPL missions. This speed up will enable trajectory design in real-time for TeamX concurrent engineering sessions. This is currently not possible except in the simplest cases. Trajectories for Ocean World missions to Europa and Enceladus are many orders of magnitude more sensitive and chaotic. This step is also critical to produce databases of working trajectories to apply machine learning and other AI technologies to mission design. The new method for computing invariant manifolds will have important consequences because the method is so simple and easily used in GPU computations. In addition to well know low energy orbits like halo orbits or near rectilinear halo orbits, the manifolds are also important for resonant orbits used to design and optimize moon tours. The manifolds have not been widely used for tour design due in part to the complexity of computing the starting conditions using the monodromy matrix and its eigenvectors. Now that it is so simple to compute the manifolds, this can be implemented on GPUs for ML and AI applications to tour design.

New Technology:
NTR 52884 - The Las Vegas Algorithm for Computing Invariant Manifolds (Pending)
This NTR is for the new method to compute invariant manifolds with a random vector without computing the eigenvectorsf the monodromy matrix.
NTR 52867 - INGOT INtegrated GPU ODE Toolbox (**Approved**)
This software NTR is for a CUDA GPU RKF78 and RKF family solver for ordinary differential equations.

Acknowledgements: Jared Blanchard, Stanford University NSTGRF student has contributed ideas, software, and tutorial support for this project. Ryan Burns supported the development of the INGOT CUDA code. AMMOS and Section 392 funded the development of INGOT.

Solve Time	IPOPT (no Sparsity Info)	IPOPT + Sparsity Info	ALTRO
Initial Solve Time	90 s	4.000 s	6.500 s
Solve Time with Warm Starting	8 s	0.128 s	**0.051 s**

Table 1. Compares solve times for IPOPT and ALTRO. IPOPT requires input of Sparsity Info in order to speed it up. ALTRO has built-in capability to obtain the Sparsity Info from the problem. With Warm Starting, ALTRO is able to really speed up. Once ALTRO is implemented on the GPU, we expect several orders of magnitude in speed increase.

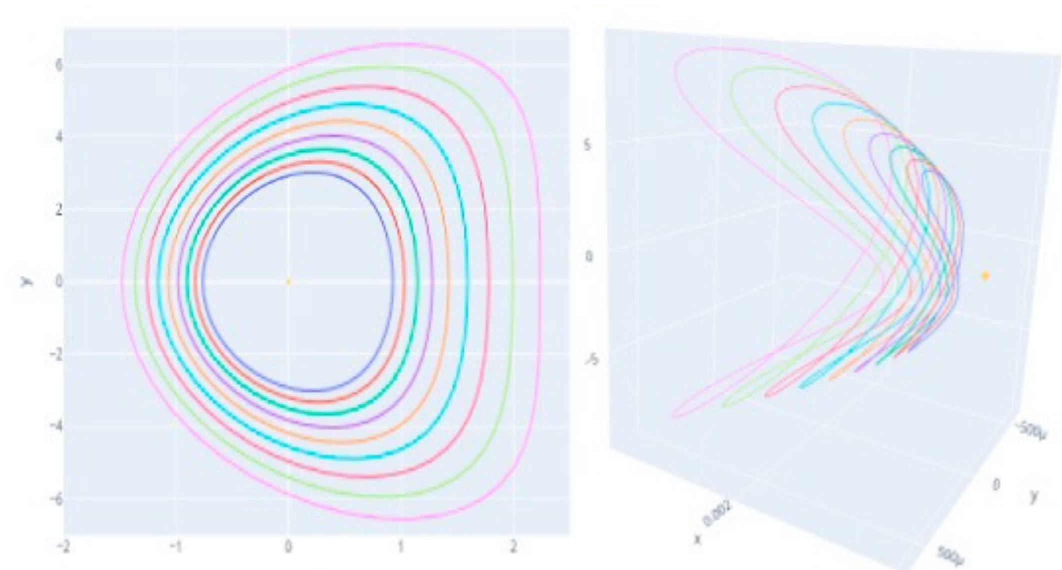

Figure 1. Halo orbits around the Moon computed using preliminary version of AstroALTRO

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

www.nasa.gov

Clearance Number: CL#00-0000
Poster Number: RPC-042

Publications:

[A] Martin Lo, "The Critical Role of Invariant Manifolds For Cislunar Missions & Beyond," Cislunar Security Conference 2023, Dec 5-7, 2023, Johns Hopkins APL.
[B] F. Vega, M. Lo, Z. Manchester, J. Blanchard, "A Massively Parallel Method for Fast Computation of Invariant Manifolds," 2024 IEEE Aerospace Conf., March 2-9, 2024.

PI/Task Mgr. Contact Information:

Martin.W.Lo@jpl.nasa.gov, *626-429-9310*