

Electronics design and sensor calibration of the magnetosonic transient plasma analyzer for characterizing hypervelocity impact plasmas on a lunar mission

Jeylin E. Lee

Mentors: Sigrid Elschot, Glen A. George, and Neilabjo Maitra

Hypervelocity impacts of meteoroids with spacecraft surfaces may generate dense plasma clouds, which then expand, propagating radio frequency (RF) emissions and instabilities that are hypothesized to cause catastrophic system anomalies. The magnetosonic transient plasma analyzer (MTPA) is being developed to measure the particle energy and velocity distribution functions of different ion species in order to understand the underlying plasma physics and evaluate the mechanisms for harmful RF emissions. The MTPA will eventually fly on a satellite or lunar orbiter to characterize the plasmas generated in the space environment, thus there is a need to adapt current prototypes into flight-compatible electronics. Data acquisition and power distribution electronics have been redesigned and miniaturized to fit eight signal channels on a single PCB, rather than eight separate boards. Firmware upgrades were written in C, and can be flashed onto the RP2040 microcontrollers over the new SWD interface. Experiments using legacy plasma sensors in the laboratory vacuum chamber were performed to begin calibration of the MTPA, using a high-voltage source to induce dielectric breakdown and generate plasmas. Next steps involve assembly and integration into a functional device that can be easily deployed for field tests, such as at NASA Ames Vertical Gun Range.

Surface chemistry behavior of carbon-based ablative thermal protective systems under hypersonic flow and temperature

Lucas J. Favretto

Mentors: Joanna M. Austin and William Feasey

Hypersonic vehicles experience extreme physical conditions when traveling at speeds of Mach 5 and above. Compressed air on the surface of the vehicle generates severe frictional heating and requires thermal protection systems (TPS) to protect the vehicle's structure from overheating. Carbon-based materials such as graphite are commonly used in ablative TPS because they can withstand high temperatures while undergoing surface reactions that dissipate heat. However, to improve the efficiency of such carbon based models, engineers require a better understanding of the surface chemistry and reaction products generated when heated graphite interacts with high-enthalpy air.

The goal of this study is to develop a robust, resistively heated graphite model for investigating carbon surface chemistry and CN emission through spectroscopy in hypersonic test facilities. The current work focuses on improving previous heated-model designs to achieve graphite surface temperatures near 1500 K while limiting overheating, deformation, and structural failure in the surrounding components. Electrical and thermal models have been developed to evaluate candidate geometries, including a hollow half-cylinder and a blunted-wedge design, and preliminary bench-top tests using graphite strips are being used to examine the effects of geometry, electrical resistance, thermal mass, and contact conditions on heating performance. Because this study is still in progress, final experimental results have not yet been obtained. Future work will involve validating the final heated-graphite design before conducting hypersonic experiments and spectroscopic measurements of CN emission.

Learning to control: A differentiable optimization approach to robust trajectory generation

Priyanka Supraja Balaji

Mentors: Zachary Manchester and Soon-Jo Chung

Robots operating in the real world must adapt to changes that are difficult to capture in advance, including variations in payload, contact, friction, actuator performance, and the surrounding environment. We propose a data-driven trajectory optimization framework for reliable behavior control on hardware. Initially, the controller leverages motion data collected offline from simulation, demonstrations, or previous robot experiments to generate feasible trajectories in real time. During deployment, the robot continuously gathers online measurements to update its internal representation

of the system dynamics. Rather than relying on new unconstrained data that can cause the model to drift, the framework uses physics-based filters to ensure incoming measurements are consistent with the expected local dynamics. Simultaneously, differentiable optimization evaluates the sensitivity of the resulting motion to changes in the learned dynamics, allowing the controller to prioritize new data updates that improve task performance. This combination of physical knowledge and task-driven learning facilitates robust online adaptation while mitigating the risk of destabilizing updates. Overall, this work aims to achieve responsive and robust behavior control that refines itself through experience while remaining bounded by physical constraints. We validate the proposed framework through real-time flight experiments on a Crazyflie 2.1 quadrotor subjected to unmodeled aerodynamic disturbances.

Experimental characterization and validation of the PARSEC Icarus lander primary structure design

Angel O. Guerra

Mentors: Mory Gharib and Jack N. Caldwell

Icarus is PARSEC's vertical takeoff and vertical landing (VTVL) rocket. This project characterized Icarus' primary structure response under vibroacoustic and static loading, refined Icarus' primary structure Geometrical Dimensioning and Tolerancing (GD&T), and benchmarked design choices against NASA/aerospace industry standards. A test article, composed of a bottom section of the Icarus primary structure, was designed and manufactured, and GD&T of Icarus primary structural components was validated during the assembly phase, which prompted redesigns and remanufacturing of necessary components. In the experimental characterization phase, ANSYS simulations of primary structural components in static loading conditions were run and vibroacoustic data was gathered from the test article. Through such testing, design choices can be validated and iterated upon. Further testing using a larger test article incorporating different sections of the Icarus primary structure will better predict the final structure's performance in intense vibrational environments and landing conditions.

Inverse kinematics derivation and mechanism development for thrust vector control on a self-lander

Zoe Lee

Mentors: Mory Gharib and Jack N. Caldwell

In recent years, there has been a growing emphasis on reusable, self-landing flight vehicles in the space industry, necessitating highly demanding flight profiles during the various stages of a flight, including reentry and landing. Careful maneuvers during the descent require dynamic adjustment of the thrust against significant aerodynamic and inertial disturbances. Consequently, expanding the performance envelope of robust thrust vector control (TVC) systems, both in mechanical design and control architecture, has become fundamental for enabling optimal, reliable, and predictive launch and landing operations, not only for large commercial vehicles but also for novel, smaller scale artifacts for academic research. The redesign, manufacturing, and validation of an improved TVC system for a collegiate-level self lander under development has been explored. The newly developed design underwent critical design reviews with updated inverse kinematics validation, detailed CAD development and component sourcing and machining in order to meet and exceed requirements of a robust and reliable gimbaling system. All components were subsequently assembled into the full TVC unit, subject to integrated validation progress alongside the remaining subsystems of the lander in the future.

A standardized CAN bus data acquisition system for a liquid rocket engine ground test stand

Emily M. Snyder

Mentors: Mory Gharib and Jack N. Caldwell

Throttleable liquid rocket engines are essential for precision landing, reusable launch vehicles, and deep-space propulsion, but their wide operating range makes them more prone to combustion and hardware failures than fixed-thrust engines. This project develops the sensing and control electronics for a ground test stand built to hot-fire and characterize a small-scale liquid rocket engine, which does

not yet implement throttling, ahead of eventually validating a throttleable flight system. The stand's avionics form a daisy-chained network of standardized boards sharing 12 V power and a CAN bus, including a combined pressure-transducer/thermocouple board with on-chip cold-junction compensation and isothermal ground layout, and a redesigned valve-position-sensor board rebuilt as a single-tier PCB to eliminate cold solder joints that had failed under vibration. Bench testing showed thermocouple channels reading within 0.3-1 °C of an ice -point reference, and fuses tripping at their rated 2 A, addressing a prior short-circuit failure. High-rate sensor data is still routed (though through a newly consolidated connectorized PCB rather than loose wiring) into an existing acquisition unit found to sample at only ~18 Hz, too slow for load-cell thrust transients; a faster replacement was designed but could not be brought online before this hotfire. Remaining work focuses on full electrical integration ahead of the test.

RAFT-XDIC: Deep learning-based framework for DIC measurement of mode II shear fractures

Claire Zhu

Mentors: Jin Yang, Guruswami Ravichandran, and Zixiang Tong

Digital image correlation (DIC) is a widely used optical technique in experimental mechanics for measuring full-field displacement and strain fields. Recent research has introduced deep learning-based DIC methods hoping to reduce manual user finetuning of parameters, computational effort, and obtain dense pixel-wise displacement field measurements. However, while most work has been conducted on smooth, continuous deformations, very little has been explored if these frameworks are able to resolve discontinuities and fractures.

This work presents the findings of RAFT-XDIC, a task-optimized deep learning pipeline employing the Recurrent All-Pairs Field Transforms (RAFT) architecture to yield accurate full-field DIC measurements of mode II shear deformations. A model was trained on synthetically generated speckle datasets containing piecewise rigid-body shearing mimicking a realistic mode II fracture. Comprehensive benchmarking and analysis were conducted on synthetic and experimental datasets of real speckle patterns and a dynamic earthquake rupture to evaluate model performance. The models were found to achieve sub-pixel accuracy and order of magnitude improvements over other learning-based DIC methods as well as conventional subset-based DIC.

Design and integration of deployment mechanisms for ultralight rollable booms and an elastically-foldable structure

Mahak Mathur

Mentors: Sergio Pellegrino and Flavia Palmeri

The Space Solar Power Project (SSPP) aims to provide continuous solar power to Earth by wirelessly transmitting energy absorbed from the Sun by large in-space, autonomously deploying solar arrays. The Deployable on-Orbit ultraLight Composite Experiment (DOLCE) was the first SSPP demonstration mission. DOLCE consists of an elastically-foldable solar array structure, four supporting ultralight rollable booms, and their respective deployers. The square structure, composed of four identical triangular sectors, is packaged by folding each sector into a flat strip and coiling it around the main-stage deployer. During DOLCE deployment, the booms and structure were deployed in a decoupled sequence, increasing system complexity and the risk of misalignment and twisting. To address these issues, a new deployment scheme was developed. In this approach, the main-stage deployer uncoils the structure while the supporting booms, connected to the structure through diagonal tapes, are synchronously unrolled by the boom deployers. Upon uncoiling, the booms extend and tension the tapes, guiding the unfolding of the structure. Implementing this deployment scheme requires the integration and synchronization of the main-stage and boom deployment mechanisms. The DOLCE deployment mechanisms were thus modified and integrated into a single cohesive system. Several prototype iterations were developed to demonstrate the synchronized coiling/uncoiling of the structure with the boom retraction/deployment, as well as the controlled unfolding of the structure, while minimizing volume and mechanical complexity.

Novel method for the in-space assembly of membrane structures with a stationary robot

Blake T. Morris

Mentors: Sergio Pellegrino and Tara Venkatadri

A new in-space assembly schematic has been developed utilizing a robot which ejects bays sequentially from a stationary truss-builder to form a ring-like perimeter truss, which is broken into segments called bays that can be modeled from a top view as sides of a polygon, supporting an interior membrane structure. A demonstration of this method with the DARPA NOM4D project uses a single large, prefabricated membrane to cover the whole ring, but studies have shown that this is not feasible as structures are scaled to greater number of sides and thus, novel methods of assembly of membrane structures must be developed. A new assembly schematic was studied which utilizes several rectangular membranes layered on top of one another which connect diametrically opposing bays of the structure and are attached as bays are ejected. The optimal layering of these membranes was determined and their shape optimized to maximize coverage of the interior of the polygonal ring whilst minimizing overlap. This new class of structures generated with this assembly method was assessed for rigidity and mechanical behavior with matrix analysis. A device was built which demonstrates this assembly method. This work will directly contribute to future studies of in-space assembly, specifically in the construction of large solar arrays.

Neural operator trajectory prediction for multi-car autonomous racing

Dami Kim

Mentors: Soon-Jo Chung and Gilchrist X. Johnson

Autonomous racing pushes vehicle control to its physical limits, and racing against other cars adds a challenge single-car methods cannot resolve: the best path for one car depends on the choices of its competitors. Reliable planning therefore requires predicting how opponents will move during close racing. This project develops and evaluates such a prediction system, built from recorded laps in a high-fidelity racing simulator. Predictions are produced in two stages. A baseline model first estimates an opponent's path from vehicle dynamics and the geometry of the ideal racing line; a learned component then corrects that estimate using the relative positions, speeds, and spacing of surrounding cars. This structure confines learning to the interactions between cars rather than the underlying motion, which improves data efficiency and keeps predictions physically plausible. Completed work includes a data processing pipeline, a set of baseline predictors, an evaluation framework based on average path error, and a real-time implementation running inside the simulator's control system. Sensitivity analysis confirms the learned component responds primarily to nearby cars rather than defaulting to single-car behavior. Remaining work is a controlled comparison against the baselines to quantify the accuracy gained during close racing.

Long-range, multi-view smartphone demonstrations for robot policy learning

Ryan M. Rudes

Mentors: Soon-Jo Chung and Fengze Xie

Full-body humanoid control research remains bottlenecked by the burden of collecting admissible on-demand motion data. Both traditional motion capture and markerless optical systems are expensive and have a fixed capture volume, making it impossible to capture long-range trajectories. We provide a mobile app bundling iPhone's advanced camera and visual odometry features with calibration tools to allow for sensible and accessible collection of spatially unconstrained, synchronized multi-view demonstrations. We also provide a mechanism for correcting odometric drift using fiducial tags, and offer a more careful execution of fiducial drift correction and acoustic synchronization than seen in prior work. In turn, we demonstrate that conventional devices can provide sufficiently high-quality data for humanoid policy learning.