Caltech
Student-Faculty Programs

Session F Abstracts

Designing cryogenic liquid fuel flow and high pressure nitrogen gas pressure control valves for use in rocketry throttling applications
Lucas A. Smith
Mentor: Morteza Gharib

The PARSEC team at Caltech is developing a self-landing rocket system for the intercollegiate Lander Challenge. Thus, we must have precise control over the throttling of the engine's thrust. This project centers around the design and development of two critical throttling systems. The first controls the flow of liquid fuel provided to the engine. This system must be able to withstand volatile conditions such as high vibration and cryogenic temperatures while providing precise and responsive control over the fuel mass flow rate and downstream pressure. The second system is responsible for the pressurization of the fuel tanks via nitrogen gas. It must be capable of maintaining a constant downstream pressure while provided a variable upstream pressure (750-4000 PSI). The valves chosen were a ball valve for the liquid system and needle valve for the gas, both fitted with actuators and microcontrollers for remote, automatic control. After designing and constructing the systems, we collected data such as flow rate, upstream pressure, and downstream pressure of the valves in comparison to their position and in different conditions. Using this data we characterized the valves to predict mass flow rate of fuel and tank pressurization in order to aid tuning of the systems.

Design and fabrication of an EDF-powered drone for emulating flight control systems of self-landing cryogenic rockets
Catherine Qu
Mentors: Morteza Gharib and Jack N. Caldwell

Self-landing rockets rely on a navigation system generally consisting of a thrust vectoring control (TVC) subsystem to mechanically angle the engine to counteract and stabilize the pitch and yaw of the vehicle during flight. The characterization for the software control of a large-scale cryogenic rocket calls for the development of a smaller-scale electric ducted fan (EDF) powered drone to simulate the flight as accurately as possible. We focused on designing a linearly actuated TVC system capable of balancing a maximum tumble in moderate winds, while remaining compact and lightweight. In addition, the integration of a linear rail system to shift weights vertically about the body of the drone allows simulation of fuel consumption, which in turn shifts in the moment of inertia of the cryogenic rocket. To mitigate the rotational inertia caused by the rotation of the EDF, smaller fans are mounted around the top of the drone in the plane containing the central axis of the rocket body. Through controlled flights as well as field-testing, we aim to achieve a stable take-off, extended hover, and landing of the drone before implementing the control system to the larger rocket.

Multi-modal mobility morphobot redesign for fully autonomous payload delivery
James D. Muren, Jr.
Mentors: Reza M. Nemovi and Joshua A. Gurovich

The Stork (M4) platform is a multi-modal morphobot drone that can transform itself from a flight mode with four propellers, to a driving mode repurposing its propeller housings as wheels. This platform is being developed to flexibly tackle the issues of autonomous package delivery in residential areas, while still maintaining the efficiency of flight capabilities.

The goal is autonomous delivery, thus, the package must be deposited from the bottom of the drone and have an opening that allows for autonomous delivery without any human involvement. In addition, the ability to efficiently navigate a short flight of stairs without the use of flight capabilities is ideal. Finally, for the purpose of reducing weight, the target is to use a minimum number of actuator motors for transformation.

This will require making a more robust actuation system that can guarantee proper clearance for the package, in addition to seamless transformation.

The specific focus of this project will be a two stage actuation system. It will use one stage to create vertical motion of the main chassis of the drone, then the second stage will use linear actuators to extend the propellors into flight position. The advantage of this design is that it minimizes friction with the ground as opposed to the current single stage actuator system.

Adaptive throttle control for autonomous race cars using meta-learning
Emily Y. Xu
Mentors: Soon-Jo Chung and Thomas A. Berrueta

Autonomous motorsports demand precise, responsive vehicle control under rapidly changing and unpredictable track and environment conditions. In the Indy Autonomous Challenge (IAC), time-varying factors such as tire and track temperature or wind create discrepancies between predicted and actual vehicle behavior. These discrepancies act as disturbances, introducing residual dynamics that impact the vehicle’s ability to maintain a desired speed and trajectory. Traditional physics-based methods often fail to capture these time-varying effects accurately. We introduce an adaptive throttle controller capable of modeling and compensating for these forces to better capture the relationship between throttle input, drivetrain response, tire forces, and acceleration. A neural network is trained offline to construct a set of basis functions that capture invariant features of the vehicles dynamics, and online adaption continuously updates corresponding coefficient weights to estimate residual forces using vehicle data in real time. The neural network is embedded within the C++ IAC vehicle codebase to allow for adaptive control on the race car while driving, estimating time-varying residual forces to optimize vehicle performance. The system is evaluated through simulation testing, measuring improvements in lap time consistency and reduced residual force magnitude, measured as the difference between expected and expected longitudinal acceleration, compared to current baseline controllers.

Decentralized collaborative learning for fault estimation in multi-spacecraft systems
Taekyung Lee
Mentors: Soon-Jo Chung, Satvik Kumar, and Vrushabh Zinage

Fault detection and isolation are critical challenges in heterogeneous multi-spacecraft systems, where hardware differences, nonlinear interactions, and limited communication make traditional model-based approaches impractical. This project presents a decentralized collaborative learning approach that addresses these challenges through novel integration of causal relationships and distributed consensus mechanisms. The methodology embeds known subsystem-level causal relationships into variational autoencoders for root cause subsystems and conditional variational autoencoders for causally linked subsystems, systematically distinguishing between primary anomalies and propagated effects. Each spacecraft independently encodes nominal subsystem behavior into orthogonal representations, preventing fault masking between subsystems. Spacecraft share encoded data, model parameters, and physical states across time-varying communication networks, achieving distributed consensus through multi-layer graph neural networks that fuse neighbor information. Experimental validation on a multi-spacecraft hardware testbed demonstrated successful fault isolation among sensors and actuators. The approach enables robust fault diagnosis in distributed space systems without requiring centralized coordination or identical spacecraft configurations, advancing autonomous operation capabilities for future multi-spacecraft missions.

Single-shot passive ultrafast all-optical replicated Kerr imaging (SPUARK)
Daopeng Yuan
Mentors: Lihong Wang and Junfu Zheng

Ultrafast imaging is essential for capturing transient phenomena that occur on femtosecond to picosecond timescales, yet existing single-shot techniques are limited by frame rate, spatial resolution, and reliance on compressed sensing. This project develops a single-shot passive ultrafast all-optical imaging approach, SPUARK, that uses a two-dimensional Kerr gating matrix to directly record ultrafast dynamics without solving an inverse problem. By splitting the event beam into multiple replicas and sequentially overlapping them with time-delayed pump pulses in a CS_2 nonlinear medium, the system achieves direct single-shot femtography with enhanced fidelity and resolution. By integrating diffractive optics and interferometric modules with comprehensive strategies for pulse management, beam shaping, and spatial control, we demonstrated high-fidelity, high-resolution Optical Kerr Effect (OKE) and Dual-Optical Kerr Effect (DOKE) imaging of ultrafast events, including plasma generation, pulse-front tilt, orbital angular momentum (OAM) beams, and light propagation in multimode fibers and optical diffusers. Challenges such as dispersion-induced attenuation, optical aberrations, limited Kerr-medium response, and nonlinear fiber effects were mitigated through precise optical redesign, advanced system optimization, and refined experimental strategies.

Physics-based inverse design for electromagnetic structures
Shan R. Gupta
Mentors: Ali Hajimiri and Jatin Mathur

Given the time- and resource-intensive design process for electromagnetic structures, there has been much academic and industrial interest in inverse design, where a computational algorithm takes in a desired device function and outputs a structure that achieves it. Current inverse design techniques rely on neural networks that require large volumes of simulations to train on, gradient-based methods that suffer from both binarization error and convergence to local minima, or heuristic methods that have no measure of global optimality. Here, we derive a physics-based algorithm for inverse design of electromagnetic structures that uses the dyadic Green's function to analytically solve for an optimal spatial conductivity distribution. Using the method of moments, we present a rapid, accurate design process with quantifiable bounds on global optimality. In implementation, we start with toy electrostatic problems and extend the method to more complicated electrodynamic structures.

Integrated distributed Bragg reflector for next-generation chip-scale beam-steering lasers
Brendan J. Rudberg
Mentors: Mo Li, Lihong Wang, and Qixuan Lin

Creating the world's smallest beam-steering laser requires that all components be integrated onto a chip just one square centimeter in size and less than a millimeter in thickness. The laser component of this device will take advantage of a distributed Bragg reflector (DBR), which is a waveguide that uses periodic changes in refractive index to reflect light of a specific wavelength, known as the Bragg wavelength. The DBR serves to restrict the laser to a single mode and a narrow linewidth. In particular, a wavelength of 1550 nm, bandwidth of 3 nm, and a peak reflectance of 90% is desired. We performed several simulations to optimize these values and determine the rough dimensions of the DBR. Next, we designed and fabricated a chip containing several DBRs sweeping several parameters using an electron-beam lithography process on a blank lithium niobate-on-insulator chip. The characteristics of this chip were then measured using a scanning electron microscope and specialized spectroscopy setup. This process was repeated several times with other chips containing DBRs sweeping new parameters. Due to the precision required at this nanometer-scale fabrication, inductively coupled plasma (ICP) etching was inconsistent. However, combining this process with hydrogen silsesquioxane (HSQ) negative resist for silicon dioxide deposition yielded good results. These results represent a critical step toward integrating beam-steering lasers for next-generation LiDAR and optical communication systems.

Optimization of iontophoresis drug delivery device for surpassing the blood brain barrier
Layla Adeli
Mentors: George Malliaras, Azita Emami, and Andrew Setley

Glioblastoma treatment is heavily limited by the blood-brain barrier (BBB), as it prevents the majority of effective chemotherapeutics, such as Doxorubicin, from reaching brain tissue. However, iontophoresis is a method that uses electric fields to drive charged drug molecules, and offers a promising approach for the controlled and localized delivery that bypasses the BBB. This project aims to optimize an iontophoretic drug delivery device by varying the pore size of permeable cellulose membranes to improve drug delivery efficiency and is structured in three phases, including device fabrication, evaluation of various membranes using saline solutions, and in vitro testing using Rhodamine 6G and Doxorubicin. Early results indicate that membrane pore size and surface chemistry have less of an influence on drug transport than previously hypothesized, and that they have varying effects on the electric field dynamics. Challenges such as drug adsorption in cellulose membranes and reduced current output have been addressed continuously by modifying the pre-treatment of the membrane, reducing channel resistance, increasing electrode surface area, and adjusting pH for conductivity. This work as it progresses contributes to the development of implantable systems for targeted drug delivery to glioblastoma tissue, offering a potential alternative to standard chemotherapy.

Quantized acceleration using Allo: Systolic tile implementation
Thomas E. Fenton
Mentors: Mircea R. Stan, Glen A. George, and Yimin Gao

The steep energy cost of AI workloads is increasingly drawing attention, motivating research into more efficient execution strategies. Two prominent approaches are quantization, which reduces storage and compute costs by representing model weights in compact formats, and the design of custom hardware architectures tailored to these workloads. Allo, a Pythonic hardware description language (HDL), offers a pathway for compiling high-level models to hardware but currently lacks support for quantization. In this work, we extend Allo by introducing a quantized systolic tile kernel, enabling developers to express quantized computation directly within the language. We demonstrate the feasibility of this extension by mapping the tile to an Alveo U280 FPGA through VitisHLS. This work lowers the barrier to building efficient AI accelerators by combining quantization-aware design with accessible hardware generation.