GSC 2020 - The Annual Workshop of Graduate Students in Systems & Control

Under the auspice of IAAC – the Israeli Association for Automatic Control

Sponsored by

The Blavatnik Interdisciplinary Cyber Research Center, Tel Aviv University

Monday, May 4, 2020, 08:50-15:20

via Zoom

CALL FOR PARTICIPATION

Dear Friends and Colleagues,

You are cordially invited to attend the annual workshop of graduate students in Systems and Control (GSC 2020). The event will take place on Monday, May 4, 2020 via Zoom (see the program below).

The workshop program and the abstracts of the talks are attached below. The program includes a plenary talk on the cyber security of industrial control systems.

As always, this yearly event provides a panorama of the research that is carried out in these areas in the various research institutions and departments in Israel. Moreover, this meeting

- Provides a unique opportunity for graduate students to present their work before a professional audience, and to observe the work of their peers.
- Allows an informal meeting avenue between faculty, industry, and students.
- Provides an opportunity to industry people, as prospective employers, to get acquainted with the best students in systems and control (and vice versa).

The success of the event obviously depends on the participation of senior people from academia and industry in addition, of course, to the graduate students themselves. I hope that you will be able to join us on May 4. The participation in the workshop is free via Zoom (see the program below).

In case of any technical difficulty, please contact Michael Margaliot by email: michaelm@tauex.tau.ac.il

With best wishes,

Anatoly Khina and Michael Margaliot

GSC 2020 organizers

Department of Electrical Engineering—Systems, Tel Aviv University.

Emails: anatolyk@tauex.tau.ac.il , michaelm@tauex.tau.ac.il

GSC 2020—Program

May 4, 2020

Online on Zoom
https://zoom.us/j/96330758945
Meeting ID: 963 3075 8945

08:50	**Opening Remarks**
09:00	Plenary talk: Rogue engineering-station attacks on S7 Simatic programmable logic controllers (PLCs)
	Prof. Avishai Wool, *EE, TAU*
09:50	**Break**
10:00	Network-based deployment of multi agents: a PDE approach
	Maria Terushkin, *EE, TAU, Ph.D. supervisor: Emilia Fridman*
10:20	Constructive method for finite-dimensional observer-based control of 1-D parabolic PDEs
	Rami Katz, *EE, TAU, Ph.D. supervisor: Emilia Fridman*
10:40	Input shaping via FIR H2 tracking framework
	Yoav Vered, *ME, Technion, Ph.D. supervisor: Izhak Bucher*
11:00	k-order contraction
	Ilya Kanevskiy, *EE, TAU, M.Sc. supervisor: Michael Margaliot*
11:20	**Break**
11:40	Bearing-based formation stabilization using event-triggered control
	Mayank Sewlia, *AE, Technion, M.Sc. supervisor: Daniel Zelazo*
12:00	LQG control over Gaussian channels with side information
	Omri Lev, *EE, TAU, M.Sc. supervisor: Anatoly Khina*
12:20	Human-in-the-loop stability analysis of haptic rendering with delay—the effect of arm impedance
	Reut Nomberg, *BME, BGU, Ph.D. supervisor: Ilana Nisky*
12:40	Human motor control strategy for performing aerial maneuvers at competitive level
	Anna Clarke, *AS&R, Technion, Ph.D. supervisor: Per-Olof Gutman*
13:00	**Lunch break**
14:00	Embedding design and control using leg shape for increased robustness of a dynamic running robot
	Adar Gaathon, *AS&R, Technion, Ph.D. supervisor: Amir Degani*
14:20	Receding horizon control of running robots on stochastic terrain using a shift-invariant funnel library
	Omer Nir, *AS&R, Technion, Ph.D. supervisor: Amir Degani*
14:40	DRRT* - A simple asymptotically optimal dynamic path planning algorithm
	Stanislav Shougaev, *AE, Technion, Ph.D. supervisor: Moshe Idan*
15:00	Minimum time optimal control of a non-linear surface vehicle subject to disturbances
	Ayal Taitler, *AS&R, Technion, Ph.D. supervisors: Erez Karpas & Per-Olof Gutman*
15:20	**End**

Plenary

Rogue Engineering-Station Attacks on S7 Simatic Programmable Logic Controllers (PLCs)

Prof. Avishai Wool, *EE, TAU*

The Siemens industrial control systems architecture consists of Simatic S7 PLCs which communicate with a TIA engineering station and SCADA HMI on one side, and control industrial systems on the other side. The newer versions of the architecture are claimed to be secure against sophisticated attackers, since they use advanced cryptographic primitives and protocols. In this talk, we show that even the latest versions of the devices and protocols are still vulnerable. After reverse-engineering the cryptographic protocol, we are able to create a rogue engineering station which can masquerade as the TIA to the PLC and inject any messages favorable to the attacker. As a first example, we extend attacks that can remotely start or stop the PLC to the latest S7-1500 PLCs. Our main attack can download control logic of the attacker's choice to a remote PLC. Our strongest attack -- the stealth program injection attack -- can separately modify the running code and the source code, which are both downloaded to the PLC. This allows us to modify the control logic of the PLC while retaining the source code the PLC presents to the engineering station. Thus, we can create a situation where the PLC's functionality is different from the control logic visible to the engineer.

Based on joint work with Eli Biham, Sara Bitan, Aviad Carmel, Alon Dankner & Uriel Malin, that appeared at BlackHat'19.

Network-based Deployment of Multi Agents: A PDE Approach

Student Name: **Maria Terushkin**
Affiliation: School of Electrical Engineering, Tel Aviv University
Studying towards degree: Ph.D.
Thesis Supervisor: Emilia Fridman

Deployment of a large-scale first-order and second-order nonlinear multi agent system over a desired smooth curve in 2D or 3D space is considered. It is assumed that the agents have access to their velocities and to the local information of the desired curve and their displacements with respect to their closest neighbours, whereas in addition a leader agent is able to measure its absolute position. The leaders (not necessarily located on the boundary when an open curve is considered) transmit their measurements to the remaining agents through a communication network. The following network imperfections are contemplated: variable sampling, transmission delay and quantization.

A static output-feedback controller is proposed, and the resulting closed-loop system is modeled as a disturbed (due to quantization) nonlinear heat or wave equation with delayed point state measurements, where the state is the relative position of the agents with respect to the desired curve.

Linear matrix inequalities (LMIs) that guarantee the input-to-state stability (ISS) of the system are derived, and further exploited to minimize the number of leader agents. The advantage of the suggested approach is in the simplicity of the control law and the conditions.

Constructive Method for Finite-Dimensional Observer-based Control of 1-D Parabolic PDEs

Student Name: **Rami Katz**
Affiliation: School of Electrical Engineering, Tel Aviv University
Studying towards: Ph.D.
Thesis supervisor: Emilia Fridman

The objective of the present work is finite-dimensional observer-based control of 1-D linear heat equation with constructive and feasible design conditions. We propose a method which is applicable to boundary or non-local sensing together with non-local (non-point) actuation, or to boundary actuation with non-local sensing. We use a modal decomposition approach. The dimension of the controller, N_0, is equal to the number of modes which decay slower than a given decay rate $\delta>0$. The observer may have a larger dimension $N\geq N_0$. The observer and controller gains are found separately of each other by solving $N_0\times N_0$ - dimensional Lyapunov inequalities. We suggest a direct Lyapunov approach to the full-order closed-loop system and provide linear matrix inequalities (LMIs) for finding N and the exponential decay rate of the closed-loop system. We prove that the LMIs are always feasible for large enough N. This is different from existing qualitative methods that do not give bounds on the observer-based controller dimension and on the resulting closed-loop performance. Numerical examples demonstrate the efficiency of our method.

Input Shaping via FIR H2 Tracking Framework

Student Name: **Yoav Vered**
Affiliation: Faculty of Mechanical Engineering, Technion
Studying towards: Ph.D.
Thesis Supervisor: Izhak Bucher

An approach to execute path following for flexible structures is to introduce a feedforward FIR shaper, whose transfer function is a linear combination of scaled delays canceling all poles of the plant. This renders the combined impulse response FIR as well. The approach can be traced back to Otto J. M. Smith's posicast control of 1957 and since then evolved to become a powerful tool used in various applications. The sensitivity to modeling errors can be reduced by adding extra delays to the shaper, at the expense of prolonging convergence time. The present work puts forward an H2 approach for the design of a general FIR shaper, which keeps the error system FIR as well. The formulation is not confined to the scaled delays, has no restrictions on the settling time, and results in a numerical stable closed-form solution. The tunable parameters have relevant physical meaning. The approach is experimentally validated on a laboratory system, which mimics a discrete flexible system with 4-DOF (degree of freedom). It is shown to be advantageous to apply a penalty on the collocated DOF mechanical power of the controlling actuator, which significantly smoothens the control signal comparing to previous methods.

This is joint work with Leonid Mirkin.

k-Order Contraction

Student Name: **Ilya Kanevskiy**
Affiliation: School of Mathematical Sciences, Tel Aviv University
Studying towards: M.Sc.
Thesis Supervisor: Michael Margaliot

Contraction theory is a useful and powerful tool in the analysis of nonlinear dynamical systems. This is due to several reasons. First, sufficient conditions for contractivity, based on matrix measures, are often easy to verify. Second, contractivity implies strong asymptotic properties including GAS in the case of time-invariant systems, and entrainment in the case of periodic time-varying systems.

A system is called contractive if the distance between any two solutions decays to zero at an exponential rate. Here, we define the notion of k-order contraction. This is motivated by the work of Muldowney, that defined this notion (without ever using the word contraction) and analyzed the important implications of what we call here 2nd-order contraction.

To explain the notion of k-contractive systems, consider solutions emanating from k different initial conditions. These define at each time t a k-parallelotope. The system is k-contractive if for any set of k initial conditions the volume of this parallelotope decays to zero at an exponential rate. For k=1 this reduces to standard contraction. The tools needed to analyze k-order contraction are wedge products, and multiplicative and additive compound matrices. The latter also play an important role in the theory of totally positive dynamical systems. These notions are not well-known in the systems and control community, and our goal here is to provide a tutorial on k-contractivity and its analysis using these tools.

Bearing-based Formation Stabilization using Event-triggered Control

Student Name: **Mayank Sewlia**
Affiliation: Faculty of Aerospace Engineering, Technion
Studying towards: M.Sc.
Thesis Supervisor: Daniel Zelazo

The formation control problem, one of the most studied problems in the field of multi-agent systems, aims to achieve a target formation shape by defining constraints on the relative states of the agents. In bearing-based formation control problems, the target formation is specified by inter-agent bearings and the sensing measurements are relative bearing vectors. The motivation to solve the bearing-only formation stabilization problem arises from using a vision-based sensor that provides us access to only bearing measurements.

In this talk, we propose event-triggered conditions (ETC) to stabilize bearing-based formation control problems for multi-agent systems with double-integrator dynamics. ETCs are designed to judge whether deviations from expected system behavior necessitates a sampling instance for closing the loop. This needful and on-demand interaction between agents in a multi-agent setting saves useful onboard resources.

We will address two problems. In the first, we propose a bearing-based controller to achieve a final stationary formation and then design an ETC to drive the sensing and control updates for the agents. Here, we also provide a lower bound on inter-event times that excludes Zeno behavior in the triggering sequence. In the second problem, we restrict ourselves to bearing-only measurements where we propose an ETC evaluated over each edge (edge triggering) to drive the control updates of agents incident to that particular edge. The proposed schemes are distributed in the sense that no global parameter was required, and agents locally determine and evaluate their trigger conditions. In doing so, we also demonstrated the efficacy of event-triggered control over traditional time-triggered control through simulation.

LQG Control over Gaussian Channels with Side Information

Student Name: **Omri Lev**
Affiliation: School of Electrical Engineering, Tel Aviv University
Studying towards: M.Sc.
Thesis Supervisor: Anatoly Khina

We consider the problem of controlling an unstable scalar linear plant over a power-constrained additive white Gaussian noise (AWGN) channel, where the controller/receiver has access to an additional noisy measurement of the state of the control system. To that end, we view the noisy measurement as side information and recast the problem to that of joint source channel coding with side information at the receiver. We argue that judicious modulo-based schemes improve over their linear counterparts and allow to avoid a large increase in the transmit power due to the ignorance of the side information at the sensor/transmitter. We demonstrate the usefulness of our technique for two practical settings:

1. The sensor is oblivious of the control objectives, control actions and previous controller state estimates.
2. The system output tracks a desired reference signal that is available only at the controller via integral control action.

Human-in-the-Loop Stability Analysis of Haptic Rendering with Delay—the Effect of Arm Impedance

Student Name: **Reut Nomberg**
Affiliation: Department of Biomedical Engineering, Ben-Gurion University
Studying towards: Ph.D.
Thesis Supervisor: Ilana Nisky

Haptic interfaces enable a sense of touch to human operators during interactions with a virtual or remote environment, providing them with an interaction that is similar to natural. Guaranteeing stability in haptic systems with time delay is challenging. In the same time, it is important in many haptic applications, such as force reflecting teleoperation, where the communication delays are substantial and unavoidable. State of the art approaches focus on the stability of uncoupled haptic systems, without the human operator, and mostly employ passivity considerations. These often require an increase of the damping of the uncoupled system to assure stability and result in a distortion in the haptic feedback. We propose an alternative approach to examining the stability boundaries of time-delayed haptic systems that focuses on the analysis of a coupled system -- the haptic system and the human operator motor control.

Here, towards human-in-the-loop analysis of stability in haptic rendering, we study the effect human operator's arm impedance on the stability boundaries of haptic rendering. By employing a method that counts the $j\omega$ axis crossings of the roots of a characteristic equation of a coupled system that includes the impedance of the operator and the haptic device, the rendered virtual stiffness and damping, and the time delay. We used values from the literature for the arm impedance components – stiffness, mass, and damping, and analyzed the effect of each of the components on stability.

We found that the added impedance of the operator brings to significantly less conservative stable time delay margins. Importantly, this analysis helped us to identify regions in the parameters space of the coupled system that assure a delay-independent stability, and others that allow for increasing the delay to induce stability. These phenomena were not discussed in the analysis of haptic interaction before. This work is our first step towards a human-centered stability analysis for time-delayed haptic systems. Next, we will analyze the contribution of additional components that characterize human sensorimotor control, e.g. feedback and feed-forward control, to the stability of haptic interaction.

Human Motor Control Strategy for Performing Aerial Maneuvers at Competitive Level

Student Name: **Anna Clarke**
Affiliation: Autonomous Systems Program, Technion
Studying towards: Ph.D.
Thesis Supervisor: Per-Olof Gutman

Aerial maneuvers are performed during a free-fall stage of skydiving, when an athlete reaches terminal vertical velocity. Mastering the skill of body flight is a great challenge in the contemporary skydiving sport. Reconstructing aerial maneuvers in computer simulations can help to gain an insight into involved motor skills and improve instruction methods. In our prior work a model and a simulation of a human body in free-fall was developed, and verified and tuned in wind tunnel and free-fall experiments. A set of controllers was designed for virtually performing basic skydiving maneuvers in a closed loop. The body was actuated according to pre-defined synergies of body available degrees-of-freedom (DOFs), inspired by empirical observations. The complex and highly non-linear dynamics of the resulting plant required dealing with significant coupling between the longitudinal and lateral dynamics, and incorporating optimization and non-linear control methods. However, we have discovered that the dynamic characteristics of the plant under investigation highly depend on the chosen DOFs synergies.

Human body has a great kinematic redundancy, whereas some body muscles, segments, and joints are activated by the Central Nervous System synchronously and proportionally, as a single unit. Combinations of such DOFs are called movement patterns. A sports technique can be viewed as a movement patterns repertoire of an athlete. Thus, human body actuation can be viewed as an hierarchical control allocation problem with two levels: choosing the body DOFs comprising a movement pattern and their proportions; and choosing the movement patterns and an extent of their engagement for producing a certain maneuver.

Control allocation methods are often involved in motion control of over-actuated mechanical systems. First, a control law is designed without the detailed analysis of actuators, specifying only the total required control effort. Next, a control allocation algorithm apportions the control input among the actuators, assuming that all the possible distributions of the control effort make the system behave in the same way. This method allows to efficiently deal with actuators' saturation, rate limits and fault tolerance, as well as minimizing the tear-and-wear and power consumption.

In the beginning of our research on body flight a similar approach was planned: first, tracking desired linear/angular velocity profiles resolved in terms of a required total force/moment; second, allocating forces/moments to individual limbs summing up to the total required effort and minimizing the resulting posture change. However, it was discovered that even though different static postures in free-fall can produce the same aerodynamic force/moment, the dynamic behavior of the system may be significantly different depending on which limbs are involved in movement. Thus, the same maneuver can be generated by utilizing different sets of DOFs, but a different controller has to be designed for each set.

We hence hypothesized that the natural kinematic redundancy, in addition to fault tolerance and energy efficiency, may allow changing the system's dynamic characteristics. In order to test this hypothesis we conducted experiments with skydivers possessing different skill levels, performing various maneuvers. This provided a novel insight into the 'low' level of human body control allocation: construction of movement patterns. It was discovered that, unlike automatic controllers that can include multiple compensation networks and deal with highly unstable plants, humans natural motor control completely lacks this ability. Humans achieve complex aerial maneuvers by the means of developing dedicated movement patterns, shaping the plant dynamics so that a desired maneuver can be performed using a proportional feedback control law.

This presentation is concerned with the 'high' level of body control allocation - combining movement patterns for performing a certain maneuver. During experiments it was discovered that experienced athletes normally engage only one movement pattern to perform basic aerial maneuvers. One movement pattern was also sufficient for many complex maneuvers and even multiple tasks. However, professionally competing athletes aim to maximize their performance and engage movement patterns providing agility, in addition to normally utilized movement patterns responsible for stability and robustness. The former patterns produce an unstable plant, what may be analogous to the flight dynamics of high performance aircraft. This work reconstructs in a simulation how the aerial rotation maneuver is controlled at a competitive level: producing the fastest and most accurate rotations a human body is capable of.

Embedding Design and Control Using Leg Shape for Increased Robustness of a Dynamic Running Robot

Student name: **Adar Gaathon**
Affiliation: Autonomous systems program, Technion
Studying towards: Ph.D.
Thesis Supervisor: Amir Degani

Minimalistic control methods typically employ a single controller to improve features for dynamic legged robots such as keeping a desired horizontal velocity. The Swing Leg Retraction (SLR), for example, which governs the leg angle during the flight phase, can increase its robustness to unforeseen rough terrain, allowing a prolonged traverse of the robot. Applying multiple controllers at once, such as simultaneously changing the leg stiffness and the leg angle, may improve the performance of legged robots further but is mechanically hard to implement, even if the controllers are coupled.

In our work, we derived and studied the limits of such an optimal coupling and developed a simple mechanical embodiment. We demonstrated that a combination of controllers can increase the robustness to perturbations in the initial horizontal velocity when traversing unknown rough terrain. We hypothesized a leg shape having contact with the ground during stance can passively exert a desired combination of physical parameters, namely, the stiffness and the free-leg length coupled with the leg angle. A single actuator that governs the leg orientation during descent, while disengaged during the stance phase, indirectly controls the physical coupled controllers.

Finding the optimal leg shape began with developing the equations of motion for a dynamic model that resembles the well-known SLIP model but incorporates a rolling contact with physical parameters that vary as the ground contact point advances. Then, we formulated an optimization process that minimizes a cost function that is related to long-term robustness and finds the temporal leg angle during descent and the continuous free leg-length and leg stiffness during stance. The physical design parameters were constrained to these of a feasible leg shape. We demonstrated that our simulated results exceed the robustness obtained with an optimized ordinary SLIP model. The optimal parameters were translated to a to a 3D printed physical leg shape and was experimentally used to verify the model.

Receding Horizon Control of Running Robots on Stochastic Terrain Using a Shift-Invariant Funnel Library

Student name:	**Omer Nir**
Affiliation:	Autonomous systems program, Technion
Studying towards:	Ph.D.
Thesis Supervisor:	Amir Degani

This talk presents a control framework for robotic running over random discrete terrain using a shift-invariant funnel library. Robotic running over stochastic terrain with limited-range terrain perception requires the ability to generate safe trajectories on the fly. Direct transcription methods can be used to generate open loop trajectories for model predictive control type controllers given desired foothold locations [1]. However, these methods commonly require a priori foothold location and cannot guarantee performance under uncertainty. We propose a different approach based on Receding Horizon Control (RHC) using a precomputed shift-invariant “funnel library”.

A funnel is a shift-invariant Region of Attraction around a nominal state space trajectory [2]. When the robot’s state is in the funnel, the associated control action attracts the state towards the nominal trajectory so that when the robot reaches the outlet of the funnel it is guaranteed to be within some predetermined distance, in state space, from the nominal trajectory. A sequence of funnels can be composed to generate complex running motion by matching the outlet of one funnel to the inlet of its successor.

A funnel library is a set of individual funnels, each dedicated to a different useful maneuver, such as running left or right; jumping high or low, etc. The funnel library is first computed offline, and during execution, the RHC searches for a sequence of funnels that result in a valid path that minimizes a given cost function. In this manner, the RHC problem is reduced to a search over discrete funnels which can be solved online whilst preserving robustness to perturbations and uncertainty.

In this talk we present the formalization of the funnel library and introduce a receding-horizon controller that employs the library. We demonstrate the formation of a funnel library for the in-plane asymmetric spring-loaded inverted pendulum model (ASLIP) [3]. We present simulative results for robotic running over stochastic terrain with random gaps and stairs.

[1] Dai, H., & Tedrake, R. (2012). Optimizing Robust Limit Cycles for Legged Locomotion on Unknown Terrain. IEEE Conference on Decision and Control, 1207–1213. https://doi.org/10.1109/CDC.2012.6425971.

[2] Majumdar, A., & Tedrake, R. (2017). Funnel Libraries for Real-Time Robust Feedback Motion Planning. International Journal of Robotics Research, 1–60.

[3] Poulakakis, I., & Grizzle, J. W. (2009). The Spring-Loaded Inverted Pendulum as the Hybrid Zero Dynamics of an Asymmetric Hopper. IEEE Transactions on Automatic Control, 54(8), 1779–1793.

DRRT*—A Simple Asymptotically Optimal Dynamic Path Planning Algorithm

Student Name: **Stanislav Shougaev**
Affiliation: Faculty of Aerospace Engineering, Technion
Studying towards: Ph.D.
Thesis Supervisor: Moshe Idan

In the last decade, autonomous mobile robots take an increasingly significant part in our life. Robotic vacuum cleaners, autonomous cars and drones are examples of this increasing field of technology. Path planning is an essential part of each such mobile robot. For a mobile robot that tries to maneuver from point A to B in an unknown area, while avoiding obstacles, choosing the right path planning algorithm is a challenge. Graph search and sampling-based methods are two popular techniques for path planning in robotics. Graph search algorithms, like A* and Dijkstra's algorithm, give optimal path, assuming a specific resolution, or a number of nodes in the graph. Finer discretization increases the quality of the solution, but also increases its computational effort. Sampling-based methods, like Rapidly-exploring Random Tree (RRT) and Probabilistic Road-Maps (PRM) avoid this discretization problem by randomly sampling nodes on a continuous map. Above algorithms are relatively simple to implement, but they lack the ability to react in a case when part of the map has changed, and the current path cannot be realized. This can happen when the mobile robot explores an unknown area, and the map is slowly revealed during its exploration. Dynamic algorithms, like RRT# and RRTX, can react to changes in the map, but they are usually more complicated to implement and code. Unlike RRT# and RRTX, Dynamic RRT (DRRT) is simple to code, but it gives a sub-optimal solution. In this research we introduce a simple dynamic path planning algorithm based on the ideas from DRRT [Ferguson et. al. 2005] and RRT* [Karaman and Frazzoli 2011], combining the advantages of the two: the proposed method, that we call Dynamic RRT* (DRRT*), is asymptotically optimal and can use information gained from previous path plannings to correct the path when the map is updated. Our algorithm was checked successfully on a simulation of a mobile robot, which explores an unknown map, trying to move from some initial point to a predefined goal point, while avoiding obstacles on its way.

Minimum Time Optimal Control of a Non-Linear Surface Vehicle Subject to Disturbances

Student Name: **Ayal Taitler**
Affiliation: Autonomous Systems Program, Technion
Studying towards: Ph.D.
Thesis Supervisors: Erez Karpas and Per-Olof Gutman

The problem of an autonomous agent moving in 2D, such as an aerial drone or a naval vessel can be treated as navigation between a series of points. While nominally the movement between each set of point can be treated as a 1D projection of the movement on the vector connecting the two points, in the presence of arbitrary disturbance the full 2D problem must be considered. The minimum time optimal solution is now dependent on the value and direction of the disturbance, which naturally delay the completion of the movement task. In this work, we address the minimum time problem of a second order system in 2D with quadratic drag, under norm state (velocity) and norm control (acceleration) constraints. The solution is obtained and proved to be optimal using the Pontryagin Maximum Principle. Simulations supporting the results are provided and compared with those of an open source academic optimal control solver.