The Andrew and Erna Viterbi Faculty of Electrical Engineering

IAAC

BMGC

GSC 2019 - The Annual Workshop of Graduate Students in Systems & Control

Monday, 29th April, 2019

Technion, Electrical Engineering (Meyer) building, Room 1003, 10th floor

Book of Abstracts

9:10	*H_2 performance in networked control: accomplish more by doing less* Maor Braksmayer, Mechanical Engineering, Technion PhD, expected graduation: 2021 Thesis advisor: Leonid Mirkin

Conventional wisdom has it that performance in sampled-data systems with periodic sampling improves as the sampling rate increases. However, the periodic sampling assumption is not justified in networked control applications, where processing units are distributed and exchange information over communication links. This information exchange might be precious, due to the use of shared communication resources and potential effects of signal transmission on power consumption of remote components. There are also technical limitations in communication networks, such as channel capacity, modulation, information losses. These factors call for intermittent (and event-based) sampling formulations of feedback control problems.

This talk presents a closed-form solution to the discrete H2 problem under intermittent sampling and its implementation, maintaining optimality of the controller under packet losses. An outcome of this solution is the understanding that performance in this case is not only a function of the average sampling rate, but also of its variability. This might change the picture, viz. faster sampling may no longer yield better performance. In the talk an example will be presented, where actual average sampling rate and the H2 performance have qualitatively different dependences on the attempted sampling rate under limited network capacity.

9:30	*Adaptive Laser-Range-Finder Terrain-Following Approach* Arseny Livshitz, Aerospace Engineering, Technion MSc/PhD, expected graduation: 2019 Thesis advisor: Moshe Idan

This paper addresses the design of an adaptive laser range sensor-based terrain following system for unmanned aerial vehicles. In this study we analyze the effect of the aircraft velocity on the terrain following performance and discuss possibilities to improve its accuracy by varying the velocity. First, we demonstrate the dependence of the system performance on the flight speed and its relation to the terrain model parameters. Consequently, we introduce a terrain model parameter identification scheme that will be used to determine the preferred aircraft velocity. The resulting optimization problem is solved using data smoothing techniques. The performance of the associated smoother is evaluated for various terrain models and is shown to be feasible and converging to the terrain model for the examined cases.

9:50	*Kinesthetic Training Module for Mastering Body Flight: Proof-of-Concept Experiment* Anna Clarke, Autonomous Systems Program, Technion PhD, expected graduation: 2021 Thesis advisor: Per-Olof Gutman

A novel concept of a real-time motion training system is proposed to deal with intensive inter- action between the environment and the trainee: the learner is augmented with an automatic control system that would be able to perform the trained activity if it had direct access to the learner's body as an actuator. The application is skydiving, where aerial maneuvers are performed by changing the body posture. Mastering body flight is very hard and protracted, since the required movements are far from our daily movement repertoire and often counter-intuitive. The stressful environment, causing muscle tension and blocking kinesthetic feedback, exteroceptive sensory overload, and very limited prior information on desired body postures add to the learning

challenges. The motor skills improve when the brain learns the environmental dynamics which occurs while actively moving inside a new environment and processing the sensory feedback. The proposed real-time system will enable novices to move in free-fall, as it provides the following essential visual cues about the environmental dynamics: Feedback of the current body posture and a prediction of the resulting future inertial position and orientation; The desired body posture that will bring the body to perform the desired maneuver. In order to compute these cues 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. Several test-jumpers performed a variety of maneuvers while their body posture was continuously measured by inertial sensors. The recorded body postures were fed into the simulator, and the resulting dynamics were compared to the real-life maneuvers. The simulator truthfully reconstructed the skydiver dynamics, including challenging aspects observed by practicing skydivers. A capability of performing autonomous tasks was added to the simulator. The desired free-fall maneuvers were defined by linear and angular velocity profiles. Controllers were designed to track the desired velocities with limbs movements as actuators. The natural kinematic redundancy of the human body was resolved by introducing movement patterns: combinations of synchronized movements of several joints. The movement patterns were experimentally extracted by analyzing body postures of various wind tunnel and free-fall maneuvers. It was found that experienced skydivers actuate the body in such a way that the resulting open-loop plant has better dynamic characteristics. The efficient movement patterns combine multiple Degrees-of-Freedom in a complex non-intuitive way involving the whole body. Hence, our on-going research on the future training strategy follows two directions: to teach complex movement patterns; and aid maneuvering by means of the simple patterns currently taught to novices. A Proof-of-Concept experiment was performed. Human subjects wore a motion tracking suit and controlled a virtual skydiver free-falling in a computer simulation by the means of their body. The task was to y along a pre-defined path what is absolutely impossible without the aiding system, due to the unfamiliar environmental dynamics and the total absence of kinesthetic feedback in the virtual world.
It was verified that the designed stabilizing hierarchical control system where the human implements the computed control actions as displayed by the visual cues, enabled all the participants to complete the task at first attempt.

10:10	*Nonholonomic dynamics of the Twistcar vehicle: asymptotic analysis and hybrid dynamics of frictional skidding.* Oriel Halvani, Mechanical Engineering, Technion MSc, Thesis advisor: Yizhar Or

The Twistcar vehicle is a classic example of a nonholonomic dynamical system. The vehicle model consists of two wheel-supported links connected by an actuated joint, where the nonholonomic constraints impose no skidding of the wheels. Recent experimental measurements conducted with a robotic Twistcar prototype have shown disagreements with predictions of previous theoretical analyses. In particular, significant skidding has been observed, in addition to discrepancies with respect to theoretical predictions of divergence in oscillations of the vehicle’s speed and orientation, as well as direction reversal depending on the vehicle’s structure.
The goal of our research is to resolve this disagreement by generalizing the theoretical analysis. First, we extend previous asymptotic analysis by incorporating the effects of links’ inertia and oscillation amplitude of the input angle on the direction of net motion. Additionally, we formulate the vehicle’s hybrid dynamics under frictional slip-stick transitions. Using numerical analysis, we obtain optimal values for the vehicle’s mean speed and energy expenditure as a function of the input frequency. Our results improve the agreement between theory and experiments and suggest directions for further experimental investigation.

10:20	*Optimal Integrated Traffic Routing and Signal Control in Simple Urban Road Networks: An Analytical Solution* Yazan Safadi, Civil and Environmental Engineering, Technion MSc, expected graduation: 2019 Thesis advisor: Jack Haddad

The current applied traffic management face limitations to control the massive traffic congestion in our cities, and therefore, the need for more-efficient and dynamic traffic management is demanding. If the road network is managed appropriately, while focusing on decreasing total time delays, or travel times for all network trips, the network performance can be enhanced. Advanced traffic management systems aim at enhancing the traffic

performance of urban networks. In addition to traffic signal control, the current research includes traffic routing as another control measure in the traffic management system.
Combined model of traffic routing and traffic signal control has a long history. In this research, similar to previous works, both traffic routing and signal control are coupled in a continuous-time model for simple traffic networks. Unlike previous works, (i) the developed model considers dynamic transient periods, rather than steady-state periods or equilibrium conditions, and (ii) system optimum via optimal control is found analytically. The proposed problem combines the modeling of routing and traffic signal control, which considers the optimal control synthesis for bringing initial queue lengths at an intersection to a predefined steady-state or equilibrium queues, by manipulating traffic routing and signal-control inputs. Based on the developed model, an optimal control problem under the objective function of minimizing total delay is formulated. The optimal analytical solutions, which are derived based on Pontryagin Maximum Principle, are verified via numerical solutions. A feedback routing and signal control policy based on link queue lengths is presented.

10:30	*Distributed Event Triggered Control for Second-Order Multi-Agent Systems* Mayank Sewlia, Aerospace Engineering, Technion MSc, expected graduation: 2020 Thesis advisor: Daniel Zelazo

Multi-agent systems (MASs) are systems that facilitate the integration of sensing, communication, computation, and control and work cooperatively with each other to achieve a common goal. With the computation and modeling tasks becoming more and more complex, centrally controlling these systems is a challenge. As a result, much current research is focused towards distributed consensus control of MASs. Here the systems communicate with each other over a wireless network with limited bandwidth and effective utilization of these networks is necessary to prevent the degradation of system performance. This gives rise to event-triggered consensus control that replaces the traditional periodic sampling with deliberate, opportunistic and aperiodic sampling. In event triggered consensus control, an event triggering condition (ETC) is derived that when violated results in the sampling of states and controller update. This ETC guarantees that the error function of the system states does not exceed a pre-specified threshold, thereby reducing the use of communication and computation resources.
This work presents a state-dependent event-triggered control strategy for a continuous time second-order multi-agent system over an undirected network. We solve this problem to achieve two different types of consensus: position consensus with average final velocity and position consensus with zero final velocity. In doing so, we present stability proofs for both the cases and highlight that the proposed ETC relies only on relative local state measurements and reduces sampling and controller updates significantly.

11:00	*Network-based boundary observer-controller design for 1D heat equation* Rami Katz, Electrical Engineering, TAU PhD, expected graduation: 2021 Thesis advisor: Emilia Fridman

Sampled-data and delayed control of PDEs is becoming an active research area. General results on sampled-data control of PDEs were presented in [1]. Constructive conditions in terms of linear matrix inequalities (LMIs) for sampled-data and delayed control of PDEs that are applicable to the performance (e.g. exponential decay rate) analysis have been initiated in [2,3,4]. However, these results are confined to distributed control of parabolic PDEs. In this talk we consider delayed boundary stabilization of a reaction-diffusion equation under boundary delayed measurements and show how to design observer-based control law via modal decomposition approach. The observer is governed by a PDE which leads to separation of the observer and the control law design. We suggest a network-based implementation of the controller in the presence of two networks: from sensor to controller, and from the controller to actuator. To reduce the workload of the second network, we suggest a novel switching-based dynamic event- triggering mechanism. We extend the results to the vector case and illustrate their efficiency by a numerical example.
Based on joint work with Emilia Fridman and Anton Selivanov.
[1] H. Logemann, “Stabilization of well-posed infinite-dimensional systems by dynamic sampled-data feedback,” SIAM Journal on Control and Optimization, vol. 51, no. 2, pp.1203–1231, 2013.
[2] E. Fridman and A. Blighovsky, “Robust sampled-data control of a class of semilinear parabolic systems,” Automatica , vol. 48, no. 5, pp. 826– 836, 2012.

[3] E. Fridman and N. Bar Am, "Sampled-data distributed H∞ control of transport reaction systems," SIAM Journal on Control and Optimization, vol. 51, no. 2, pp. 1500–1527, 2013.
[4] A. Selivanov and E. Fridman, "Distributed event-triggered control of diffusion semilinear PDEs," Automatica, vol. 68, no. 344–351, 2016.

11:20	*Minimum Time Validation for Hybrid Task Planning* Ayal Taitler, Autonomous Systems Program, Technion PhD, expected graduation: 2020 Thesis advisors: Erez Karpas, Per-Olof Gutman

The problem of mixed discrete-continuous task planning for mechanical systems, such as aerial drones or other autonomous units, can often be treated as a sequence of point-to-point trajectories. A solution to such a planning problem consists of the sequence of points in the space, representing different stages of the plan, and a control trajectory connecting these points. The minimum time optimal solution between points in the plan is critical not only for the calculation of the trajectory in cases where the goal has to be achieved quickly but also for the feasibility checking of the plan and
the planning process itself, especially in the presence of deadlines and temporal constraints. We propose an iterative approach to the problem, consisting of search in the task space with a simple dynamic model, and plan validation with analytical minimum time constraints of a more realistic model. We address the minimum time problem for a second-order system with quadratic drag, under state (velocity) and control (acceleration) constraints. Closed form expressions for the trajectory are derived and the optimality is proven using the Pontryagin Maximum Principle. The iteration between the plan search and plan validation is enabled with a graph transformation eliminating the infeasible paths completely from the problem.

11:40	*Geometric vs. variational optimal control of gaits for 3-link swimmers* Oren Wiezel, Mechanical Engineering, Technion PhD, expected graduation: 2019 Thesis advisor: Yizhar Or

Many robotic systems use internal periodic shape changes (gaits) that, coupled with the interaction with the robot's environment, create locomotion. Prominent examples of such systems with two shape variables are the low Reynolds 3-link "Purcell swimmer" with inputs of 2 joint angles [1] and the ideal fluid inertial swimmer [2]. Gait optimization of these locomoting robots is of high importance, and has been studied mainly using numerical methods. Using the variational optimal control approach of "Pontryagin's maximum principle" (PMP) for these swimmer models, enables finding displacement optimal gaits for certain swimmer's geometries [3], but diverges for others. In an attempt to better understand the optimal gaits and explain the divergence, we examine another method for calculating optimal gaits, namely geometric optimal control.
Ramasamy and Hatton [4] use differential geometry to calculate the curvature of the local connection (its total Lie bracket) for 3-link swimmers. Using optimized choice of coordinates, the net displacement of a gait is obtained as the area integral of the curvature over the region enclosed by the gait. Thus, contour plots of the curvature in shape space gives visualization that enables identifying distance-optimal gaits. By analyzing contour plots for the different swimmer geometries, we find that in some cases, junctions appear in the zero-level curves, which explain the failure of the PMP method since there is no single solution to the optimal control problem. Further comparison between the two methods may help in finding gaits for maximal efficiency for these systems and others.

References

[1] E. Gutman and Y. Or, "Symmetries and Gaits for Purcell's Three-Link Microswimmer Model," IEEE Transactions on Robotics, vol. 32, no. 1, pp. 53-69, 2016 .
[2] E. Virozub, O. Wiezel, A. Wolf and Y. Or, "Planar Multi-Link Swimmers: Experiments and Theoretical Investigation using "Perfect Fluid" Model.," Robotica, pp. 1-13, 2019 .
[3] O. Wiezel and Y. Or, "Using optimal control to obtain maximum displacement gait for Purcell's three-link swimmer," in Decision and Control (CDC), 2016 IEEE 55th Conference on, Las Vegas, NV, USA, 2016 .
[4] S. Ramasamy and R. L. Hatton, "Soap-bubble optimization of gaits," in Decision and Control (CDC), 2016 IEEE 55th Conference on, Las Vegas, NV, USA, 2016.

Corresponding author orenwi@campus.technion.ac.il

12:00	*Ribosome flow model with different site sizes* Eyal Bar-Shalom, Electrical Engineering, Tel Aviv University MSc, expected graduation: 2019 Thesis advisor: Michael Margaliot

We introduce and analyze two dynamical models for unidirectional movement of particles along a circular chain and an open chain of sites. The models include a soft version of the simple exclusion principle, that is, as the density in a site increases the effective entry rate into this site decreases. This allows to model and study the evolution of ``traffic jams'' of particles along the chain.
A unique feature of these new models is that each site along the chain can have a different size. Although the models are nonlinear, they are amenable to rigorous analysis. In particular, we show that the dynamics always converges to a steady-state, and that the steady-state densities and production rate can be derived from the spectral properties of a suitable matrix, thus eliminating the need to numerically simulate the dynamics until convergence. This spectral representation also allows for powerful sensitivity analysis, i.e. understanding how a change in one of the parameters in the models affects the steady-state. We show that the site sizes and the transition rates play different roles in the dynamics, and that for the purpose of maximizing the steady-state production rate the site sizes are more important than the transition rates. We also show that the problem of finding parameter values that maximize the production rate is tractable. We believe that the models introduced here can be applied to study various natural and artificial processes including ribosome flow during mRNA translation, pedestrian and vehicular traffic, evacuation dynamics, and more.

12:10	*A String of Tethered Drones – System Dynamics and Control* Benny Kosarnovsky, Mechanical Engineering, Ben Gurion University MSc, expected Graduation: 2019 Thesis advisor: Shai Arogeti.

In recent years the Quadcopter is emerging as a popular platform due to its high manoeuvrability, simplicity and low cost. Quadcopters, However, suffer from several critical disadvantages that prevent them from being utilized for several applications. Most notably, the Quadrotors low energy efficiency and its limited navigation capability in GPS-denied environments. Tethered-drones is a relatively new concept that overcomes these limitations by utilizing the advantages of the tether. By using a power line as a tether, one can transfer power to the drone. Additionally, by measuring the tether's attitude and length, one can determine the drone's position and attitude in space.
In this research, we propose a novel design of a unified system with an arbitrary number of serially tethered drones and a ground station with the capability of controlling the total tether length and tension. In addition, we equip each of the drones with a pulley and a double-gimbal mechanism, to provide the following: (a) Allowing the drones to freely move along the tether while measuring their relative distances, (b) determining the drones and tether relative attitude without the use of standard inertial onboard sensors. Consequently, all drones' position and orientation coordinates can be deduced from tether measurements. The suggested concept allows more operational flexibility, such as reaching obstructed locations with no line-of-sight between the ground station and the target position, operating in a GPS denied environments and reducing the reliance on inertial sensors. The use of geometric-mechanics principals for the development of the system dynamics allows for a compact and elegant coordinate-free model, compared to the common use of Euler angles. The corresponding use of geometric control combined with linear control for the tether length and tension allows quick convergence and stability, while avoiding limiting singularities related to Euler angles. Simulation results show quick convergence to set-point tracking maneuvers while sustaining tether tautness requirements, along with the ability to control a swarm of drones (more then 10). Future work will be to implement a constrained control method such as MPC and IC to govern tether tautness and other constraints, and to construct a working prototype. In short, the three main contributions of this research are: (a) A concept of using the tether to determine the position and attitude of all drones. (b) A compact expression for the dynamics of a serially connected arbitrary number of tethered drones, utilizing computational tools from geometric mechanics, and (c), the development of a controller to control a string of any number of tethered drones.

12:20	*Estimation-Based Guidance Using Optimal Bayesian Decision* Liraz Mudrik, Aerospace Engineering, Technion MSc, expected graduation: 2020 Thesis advisor: Yaakov Oshman

Realistic guidance of an interceptor towards an evasively maneuvering target, using noisy and imperfect measurements, is a nonlinear and non-Gaussian stochastic control problem. Nevertheless, most common, advanced guidance laws assume to possess perfect information about the target state. Under that assumption, bounded-control differential game-based guidance laws (e.g., DGL1) ensure hit-to-kill performance for an interceptor possessing maneuverability and agility superiority in the standard, linearized case. These are bang-bang laws that require the computation (or estimation) of the zero-effort miss (ZEM) and the time-to-go. As the perfect information assumption is unrealistic, a target state estimator has to be incorporated. Separate design of the perfect information-based guidance law and the target state estimator inherently amounts to assuming the validity of the separation theorem. However, this theorem has never been proven valid in realistic scenarios, involving nonlinearities and non-Gaussian noises. In contradistinction, the generalized separation theorem (GST), a generalized form of the separation principle, states that the estimator can still be designed separately of the guidance law, however, the guidance law must take into account the posterior probability density function (PDF) of the target's state.
We present a new, estimation-aware guidance law, that constitutes an adaptation of the perfect information-based DGL1 to realistic, uncertain, scenarios. To conform with the GST, we employ the interacting multiple model particle filter (IMMPF), which is capable of dealing with nonlinear, non-Gaussian and hybrid problems, and readily provides the entire posterior PDF, represented by a cloud of weighted particles. Because the performance of the DGL1 law significantly degrades as the information on the game's state becomes uncertain, we propose to use Bayesian decision theory in order to make optimal decisions on this state, that properly take into account the inherent uncertainty due to the estimation error, on the one hand, while also considering the cost (final miss distance) associated with making correct or erroneous decisions on the game's state, on the other hand. In turn, we modify the DGL1 law to use these optimal decisions, which results in a new, GST-compliant, guidance law. The performance of the new guidance law is demonstrated via an extensive Monte-Carlo simulation study, where it is compared with the classical DGL1 law, and with a theoretical, deterministic, miss distance lower bound.

14:20	*Control of Mixed Platoons Consist of Automated & Manual Vehicles* Omer Orki, Mechanical Engineering, Ben-Gurion University PhD, expected graduation: 2020 Thesis advisor: Shai Arogeti

Platooning has a great potential for improving highway traffic flow. Platooning research focuses on fully automated platoons. However, in the near future traffic will be mixed and will consist of automated and manually driven vehicles. This work introduces a new concept of mixed platoons, i.e., platoons consist of automated and manual vehicles. In a mixed platoon the automated vehicles are controlled as part of the mixed traffic, using global sight on traffic. This way the mixed traffic flow can be optimized using the automated vehicles.The mixed platoons control developed in this work implements the same control components that are used for automated platoons, but modifies them to suit mixed traffic environment. That is, a new Information Flow Topology (IFT) and new Distributed Controller (DC) are developed for mixed platoons using H_∞ control methods. Subsequently, a string stability analysis for a mixed platoon is shown, demonstrates the feasibility and potential of mixed platoons.

14:40	*A Distributional Perspective on Learning in Control Problems.* Dror Freirich, Electrical Engineering, Technion PhD, expected graduation: 2021 Thesis advisor: Ron Meir

In this work, we discuss the distributional approach to reinforcement learning (DiRL), where policies are learned according to sensory inputs without assuming a model of the environment. Deep Reinforcement Learning (DRL) has been applied in recent years to a wide range of problems in large, yet usually discrete, domains. More recently, applications to classic control and robotics problems have emerged. Most current deep RL methods

are based on estimating the expected reward for carrying out the next action, known as the Value-function or the Q-function. The latter is usually represented as a deep neural network, and learning usually takes place within the so-called model free setting, which aims to directly map states to actions, without learning a model of the environment. More recent studies suggested learning the value distribution, rather than the expectation, and demonstrated the empirical advantage of such approaches. In our work, we show that the distributional Bellman equation, which drives DiRL methods, is equivalent to a generative adversarial network (GAN) model. In this formulation, DiRL can be seen as learning a deep generative model of the value distribution, driven by the discrepancy between the distribution of the current value, and the distribution of the sum of current reward and next value. We use this insight to propose a GAN-based approach to DiRL, which leverages the strengths of GANs in learning distributions of high-dimensional data.
We introduce the Multivariate distributional Bellman equation, and use DiRL for policy evaluation in multi-objective environments. The multivariate setting also allows us to unify learning the distribution of values and state transitions, and we exploit this idea to devise a novel exploration method. We demonstrate the effectiveness of our approach on both discrete and continuous control problems.

15:00	*The Contribution of Quadrotor-Formation and Imaging-Configurations for Improved Solutions to Navigation and Mapping Scenarios* Dror Hurwitz, Technion Autonomous System Program MSc, Thesis advisors: Filin Sagi, Klein Itzik

In this talk, we discuss a quadrotor formation, operating in a mapping scenario. We propose implementing range constraints in the bundle adjustment algorithm that are aimed to improve all aspects of the system performance. This includes accuracy and cost, and weight of each member in the formation . To that end, we derive stochastic constraints on the relative quadrotors pose parameters and show both analytically and numerically how they reduce the navigational and mapping errors. The proposed approach merits on efficient mapping area coverage, reduction of time on site, ability to maintain accuracy while operating at higher altitudes, and the use of low-cost sensors

15:10	*Robust Interpolating Traffic Signal Control for Uncertain Road Networks* Shimon Komarovsky, Civil and Environmental faculty specialized at Transportation control, Technion. MSc, expected graduation: 2019 Thesis advisor: Jack Haddad

Many traffic control strategies in the literature are designed based on various traffic flow models, which do not include parameter uncertainties in their intersection flow dynamics. The parameter uncertainties may include the saturation flow rates, turning movement rates, and others, that are difficult to estimate in reality. The model parameters are assumed to be a priori known or estimated. Hence, the resulted control strategies are unable to handle conditions when the assumed values for the
parameters vary significantly from their real values, e.g. applying variants of linear quadratic regulator (LQR) on uncertain traffic model might cause low control performances, as they cannot compensate parameter uncertainties. Moreover, many utilized state feedback control approaches to develop the traffic signal controllers lack the treatment of control and state constraints directly in the design phase.
This presentation considers robust traffic signal control for uncertain urban road networks. First, parameter uncertainties are integrated into the store-and-forward (SF) model, which is utilized in this presentation to describe the flow dynamics for traffic signalized intersections. Then, the uncertain SF model is utilized to design a robust feedback controller by an interpolation-based approach.
The interpolation-based approach (i) guarantees robustness against all parameter uncertainties, (ii) handle control and state constraints, and (iii) present a computationally cheap solution.
Two variants of interpolation-based control methods, i.e. Interpolating Control (IC) and Simple Interpolating Control (SIC), are applied.
Finally, numerical results for an isolated signalized intersection show a comparison between the developed interpolating controllers and other controllers in the literature: Linear Quadratic Regulator (LQR) and Model Predictive Controller (MPC). Both controllers are designed with the objective to regulate the plant state to the origin. The results demonstrate the performance advantages from applying the robust interpolating controllers.

15:20	*Drone`s attitude estimation for a corridor-like environments using camera and rate gyroscopes* Daniel Jano, Mechanical Engineering, Ben-Gurion University MSc, expected graduation: 2019 Thesis advisor: Shai Arogeti

In this study, we suggest an attitude estimation algorithm for drones flying indoors. In particular, we consider a corridor-like environment, and adapt ideas from the aerospace field, where algorithms were developed for satellite`s attitude estimation. Many algorithms can be found that estimate satellite`s attitude, based on rate gyroscopes and a sensor called, star-tracker. The star-tracker is a sensor that identifies celestial objects, and by doing so, determines their directions compared to the satellite. Using star maps, the same celestial objects directions, compared to the earth, is known. By comparing the celestial objects directions in the satellite frame and in the earth frame, the attitude of the satellite can be estimated. Complementing the star-tracker with rate gyroscopes provides smooth attitude estimation, while also compensating for the rate gyroscope`s drift.The novelty of this work is the implementation of the star-tracker method on a drone in a corridor-like environment and finding features, which replace the celestial objects used by the star-tracker.

In order to implement the above method on a drone in a corridor-like environment, a star-tracker parallel approach was developed, based on a camera and the corridor`s vanishing points (VP), as the needed fixed features (which replace the fixed stars). A VP is an intersection of parallel lines as shown on a 2D image, where in the real world they intersect at infinity. In a corridor-like environment there are 3 VPs, the longitudinal VP, the horizontal VP and the vertical VP, with their directions in the inertial frame given by $[1\ 0\ 0]^T$, $[0\ 1\ 0]^T$ and $[0\ 0\ 1]^T$, respectively. Using the pixel of those VPs shown in the image, their directions compared to the drone`s frame can be determined. Comparing the directions of the VPs both in the inertial frame and the drone`s frame can be used for estimating the drone`s attitude with regard to the corridor. For the attitude estimation, there is a need of at least two VPs directions in both frames.The attitude estimation is based on an extended Kalman-filter (EKF), formulated by the multiplicative properties of a quaternion (i.e., MEKF). Using the quaternion for the estimation of the drone`s attitude allows representing the drone`s rotation, with respect to the inertial frame, without singularities. The process equation was based on the quaternion`s kinematics, including noise and a linear drift model. The measurement equation is developed from the quaternion rotation matrix, to compare between the VPs directions in both frames. Both equations include noise, which is assumed to be Gaussian and White. Based on the above, an experiment was conducted to test the developed approach and the MEKF algorithm. The reference attitude measurements in the experiment was done by an OptiTrack system. The reference measurements were compared to the drone`s attitude estimation that was based on a video taken from a camera mounted on the drone. From each frame of the video, the directions of the longitudinal and horizontal VPs were extracted using image processing algorithms. Both data from the video and the gyroscopes have been used in the MEKF. The experimental results showed good agreement between the estimated attitude, and the ground truth taken from the OptiTrack system.

15:50	*A Generalization of Linear Positive Systems* Eyal Weiss, EE-Systems, Tel Aviv University PhD, expected graduation: 2022. Thesis advisor: Michael Margaliot

The dynamics of linear positive systems maps the positive orthant to itself. In other words, it maps a set of vectors with zero sign variations to itself. This raises the following question: what linear systems map the set of vectors with k sign variations to itself? We address this question using tools from the theory of cooperative dynamical systems and the theory of totally positive matrices. This yields a generalization of positive linear systems called k-positive linear systems, that reduces to positive systems for k=1. We describe an application of this new type of systems to the analysis of invariant sets in nonlinear time-varying dynamical systems. Our results suggest a new approach for obtaining information on nonlinear time-varying dynamical systems.

16:10	*Harnessing Reinforcement Learning for Neural Motion Planning* Tom Jurgenson, Electrical Engineering, Technion PhD, expected graduation: 2021 Thesis advisor: Aviv Tamar

Motion planning is an essential component in most of today's robotic applications. In this work, we consider the *learning* setting, where a set of solved motion planning problems is used to improve the efficiency of motion planning on different, yet similar problems. This setting is important especially in applications with rapidly changing environments such as in e-commerce, industry 4.0 or home robotics, where it is desired to plan *fast*, and the computational burden of current (sampling-based) planners can be limiting.

We investigate a general deep learning-based approach, where a neural network is trained to map an image of the domain, the current robot state, and a goal robot state to the next robot state in the plan. We focus on the learning algorithm, and compare supervised learning methods with reinforcement learning (RL) algorithms. We first establish that supervised learning approaches are inferior in their accuracy due to insufficient data on the boundary of the obstacles, an issue that RL methods mitigate by actively exploring the domain. We then propose DDPG-MP, a modification of the popular DDPG RL algorithm that is tailored to motion planning domains by exploiting two features of the problem; (1) a known dynamics model allows variance reduction during training and (2) demonstration data (collected off-line using a conventional sampling-based planner) can be integrated as off-policy data to perform efficient exploration.We conduct experiments on a 4-dimensional robotic arm domain with very challenging narrow passages. We show that DDPG-MP significantly improves the accuracy of the learned motion planning policy, reaching 98% success rate on unseen test instances, compared to less than 35% for DDPG with a state-of-the-art hindsight experience replay (HER) exploration strategy. Finally, we show that given enough training data, our method can plan 6X faster on novel domains than off-the-shelf sampling-based motion planners.

Video demonstrations of trajectories generated by both our fully-trained DDPG-MP model as well as other RL baselines are shown in https://youtu.be/wHQ4Y4mBRb

16:30	*Cooperative Guidance and Collision Avoidance for Multiple Vehicles.* Bhargav Jha, Aerospace Engineering, Technion PhD, expected graduation: 2021 Thesis advisor: Tal Shima

Autonomous vehicles are becoming increasingly popular due to their wide applications in civilian as well as military domains. In scenarios where they are deployed as a team, there is an inevitable possibility of collisions but at the same time there is an opportunity for cooperation. Catering to such scenarios, we will present a cooperative guidance law for multiple vehicles that steers them to their assigned target points, while avoiding collisions among the team members and with other moving obstacles. The guidance law optimizes the total team effort with different penalties on the control effort for each member. This enables constructing a heterogeneous team of vehicles with different maneuver capabilities and thereby achieving the objectives even when some of the vehicles are not as maneuverable as the others. We will discuss a solution methodology that involves linearization of the non-linear engagement

kinematics around the collision triangles, and order reduction using terminal projection to obtain a reduced order linearized system. This will pave the way for the use of linear quadratic optimal control-based formulation to minimize the total team effort, where target interception and collision avoidance are incorporated as state constraints. An analysis of the closed form guidance laws for a special case of 2-on-2 engagement will also be shown. Thereafter, various simulations and experimental validation results, exemplifying the cooperation in different engagement scenarios, will also be demonstrated.

16:50	*Multi Agent Reinforcement Learning with Multi-Step Generative Models.* Orr Krupnik, Electrical Engineering, Technion MSc, expected graduation: 2019 Thesis advisor: Aviv Tamar

The dynamics between agents and the environment are an important component of multi-agent Reinforcement Learning (RL), and learning them provides a basis for decision making. However, a major challenge in optimizing a learned dynamics model is the accumulation of error when predicting multiple steps into the future. Recent

advances in variational inference provide model-based solutions that predict complete trajectory segments, and optimize over a latent representation of trajectories. For single-agent scenarios, several recent studies have explored this idea, and showed its benefits over conventional methods. In this work, we extend this approach to the multi-agent case, and effectively optimize over a latent space that encodes *multi-agent strategies*. We discuss the challenges in optimizing over a latent variable model for multiple agents, both in the optimization algorithm and in the model representation; the former we solve by employing optimization tactics based in game theory and the latter by proposing a model for both cooperative and competitive settings based on risk-sensitive optimization. We distinguish between two types of models, suitable for two different scenarios. The first is a disentangled, mutual-information-based model for scenarios where the agents' dynamics are interdependent (for instance, they may collide and obstruct each other's path); second, a conditional model for the simpler case where the dependence between the agents manifests only through their reward (i.e. point agents which do not interact with each other physically). We use a domain with larger, colliding agents to show why the first model is necessary, and demonstrate a case where the conditional model is not sufficient to solve the given task. We implement both of our models using Conditional Variational Auto Encoders with an architecture based on dilated causal convolutions, which has been shown to effectively capture temporal information in various domains. All of our models are learned offline, with a relatively low number of samples from the environment.
We present a practical algorithm for trajectory optimization of multiple agents based on risk-sensitive optimization, and discuss several approaches to modeling trajectories for multiple agents that comply with planning based on latent-space search. We show that our method can learn effective models for continuous multi-agent domains with complex dynamics, both in the multi-agent particle environment and on a simulated RoboCup domain. Using Model Predictive Control, we are able to leverage our learned dynamics model to plan using predicted trajectories and perform better than baseline policies. Qualitatively, we show that trajectories predicted by our model are diverse in the *strategies* they employ, and not just in minor deviations of agent properties.

17:00	*Implementation of a Natural Dynamic Controller on an Under-actuated Compass-Biped Robot* Ron Hartston, Mechanical Engineering, Sensory Motor Integration Laboratory MSc, Thesis advisors: Miriam Zacksenhouse , Reuven Katz

Natural dynamic controllers aim to perform the desired task by exploiting the natural dynamics of the system. This can be accomplished by generating torque patterns to actuate the system rather than accurately following a predefined trajectory. We have previously demonstrated natural dynamic control of compass-biped in simulations. Here we demonstrate successful implementation of this dynamic controller on an under-actuated compass-biped robot. The parameters of the controller, and in particular the magnitude and timing of torque primitives, were optimized using multi-objective optimization via genetic algorithm, accounting for speed and energy efficiency. While the current implementation is in open-loop, this strategy can be extended to include feedback to enhance walking over a wide range of terrains. This proof-of concept provides the basis for future extensions to more complex robots.

17:10	*Communication constraints for Multi-Agents systems in Reinforcement Learning* Or Raveh, Electrical Engineering, Technion MSc, expected graduation: 2019 Thesis advisor: Ron Meir

Efficient concurrent exploration by multiple agents in Reinforcement learning (RL) is important for many real-life problems. Furthermore, since communication between con- current agents in real life is often noisy, algorithms that can efficiently overcome communication problems are necessary. Although many previous works have proven bounds and demonstrated computational results for single agents' problems, the literature regarding concurrent exploration is more limited. Recently, more and more works exemplify the benefits of concurrent exploration using simulations, but theoretical advance in this field is scarce. In our work, we use the algorithmic setting suggested by Pazis and Parr in [1]. In this setting, concurrent agents can explore their environment independently, by performing greedy exploration over a Q-function provided by a mutual Learner, which in turn receives samples collected by all of the agents, and updates the Q-function accordingly. The communication process in [1] is assumed to be perfect in the sense that all transmitted information is received on each end. However, as stated before - this is not the case in real life - and in fact it has been shown in previous works that communication that is too noisy can even lead to a system which is uncontrollable. In our work, we extend the setting suggested in [1] to the noisy communication case. We consider communication channels with an additive noise or with quantization noise, and show how it affects the exploration complexity bounds developed in the original work. Furthermore, we suggest improved algorithms that can overcome the communication noise and lead to better sample complexity bounds. One important case we examine is the one of communication between agents themselves, which is not mediated by a common learner. This setting does not exist in the original algorithm, where sharing information is only possible via a common learner that collects samples from all agents.

References

- [1] Pazis Jason PAC-optimal, Non-parametric Algorithms and Bounds for Exploration in Concurrent MDPs with Delayed Updates Department of Computer Science Duke University, 2015.