Academic Report

School of Electrical and Computer Engineering, Tel Aviv University

Yoav Pichoto*
Advised by Dr. Ilai Bistritz

The 64th IEEE Conference on Decision and Control (CDC 2025) served as a pivotal milestone in my academic development. As my first major international conference, the experience was both humbling and motivating. Held as a premier gathering for the global control systems community, the conference brought together many researchers from around the world, ranging from foundational theorists to applied engineers, creating an environment of intense intellectual exchange.

The exposure to diverse methodologies has catalyzed a significant shift in my perspective on the distributed control of cooperative agents and related fields. Drawing from the thematic discussions at the conference, I have identified a few primary research leads that will form the cornerstone of my future academic contributions.

- **Resilience in Open Multi-Agent Systems:** One of the most profound realizations during the conference was the need for systems that are robust to "player switching." In many current formulations of multi-agent cooperation, the set of participants is assumed to be static. However, real-world transportation systems like delivery drones or autonomous fleets involve constant shifts (one-for-one) as agents enter or exit due to battery constraints, task reassignment, or arrivals to destination. A key lead I intend to pursue is the design of cooperative games where the Nash Equilibrium is insensitive to player substitutions. This involves developing "strategic continuity" metrics to ensure that the global coordination remains stable despite frequent changes in the identities and characteristics of the individual agents while maintaining the same number of agents N.

- **Dynamic Games with Time-Varying Populations:** Building on the technical discussions regarding gradient-based learning, I identified a significant opportunity in exploring Dynamic Games with an Evolving Population Count $N(t)$. The idea is to move beyond the standard fixed-dimension state space. By viewing the system through the lens of population dynamics, I aim to investigate how control laws can be synthesized to maintain convergence even as the cardinality of the agent set fluctuates. This is particularly relevant for autonomous vehicle routing, where the network's density changes unpredictably as vehicles join or leave the traffic flow.

*yoavpichoto@mail.tau.ac.il

- **Control over Fluid and Dynamic Graphs:** The primary challenge in these environments is that the graph topology is constantly reshaped by external interruptions such as accidents, lane closures, and sudden road construction. These disruptions do more than block physical paths; they create "information islands," where groups of agents are severed from the global consensus. This forces players to make decisions based on localized, outdated data, leading to divergent learning and high social costs as the fleet fails to account for newly restricted routes.
- **Potential Games for Scalable Cooperation:** Finally, the versatility of Potential Games emerged as a powerful tool for ensuring convergence in large-scale systems. I aim to research the synthesis of specific potential functions that align local agent objectives with global social costs. This lead will focus on identifying the class of delay functions and interaction protocols that allow distributed agents to naturally gravitate toward a system-wide optimum. By utilizing potential-based frameworks, we can ensure that distributed learning—such as epoch-based exploration—leads to a stable and efficient equilibrium without the need for a central supervisor.

Beyond the formal technical presentations, the most impactful aspect of the conference was the opportunity for direct engagement with the research community. I participated in several discussions that proved to be highly illuminating. Engaging with other doctoral students and senior researchers allowed me to discuss the practical hurdles of implementing distributed control laws in real-time systems.

One of the highlights of my attendance was the series of "coffee-break" discussions, where I had the chance to present my own research logic to peers from diverse institutional backgrounds. These conversations provided immediate candid feedback that is often absent in the blind review process. Discussing the nuances of multi-agent cooperation with researchers working on similar problems in different domains—such as robotics vs. power grid stability—revealed a shared set of challenges regarding system "openness" and agent heterogeneity. These interactions have not only expanded my professional network but have also validated the relevance of my research focus within the broader control community.

Attending CDC 2025 as my first international conference was a transformative, mind-opening experience that fundamentally expanded my academic horizon. By stepping out of theoretical isolation and into an environment of intense intellectual exchange, I gained a broader perspective on the systemic challenges of real-world deployment. This exposure catalyzed a shift in my research focus toward "open" multi-agent systems, inspiring new leads in player-switching resilience, dynamic population games, and potential-based cooperation. Ultimately, the candid feedback from the global research community validated my work and provided a new, more comprehensive framework for my future contributions to distributed control.