Potential Games on Cubic Splines for Multi-Agent Motion Planning of Autonomous Agents

Extended Abstract

Sam Williams
University of Southern California
Los Angeles, CA, USA
samwilliams@usc.edu

Jyotirmoy Deshmukh
University of Southern California
Los Angeles, CA, USA
jyotirmoy.deshmukh@usc.edu

ABSTRACT

We present an algorithm to solve for local Nash Equilibrium trajectories in the multi-agent motion planning problem for self-interested agents. Our method models the problem as a concurrent game where each agent's action consists of choosing a cubic spline defined by a set of waypoints. We observe that with certain kinds of cost functions, the resulting game has the structure of a *potential game* which is guaranteed to reach an equilibrium even when each agent myopically improves their own cost without considering the costs of other agents. Our algorithm uses simultaneous gradient descent with independent per-agent step sizes to converge to local Nash Equilibrium trajectories. We demonstrate the algorithm can scale to very long horizons through simulated experiments in the electric vertical take-off and landing vehicles (eVTOL) domain.

KEYWORDS

Multi-Agent Motion Planning; Potential Games; Cubic Splines

ACM Reference Format:
Sam Williams and Jyotirmoy Deshmukh. 2024. Potential Games on Cubic Splines for Multi-Agent Motion Planning of Autonomous Agents: Extended Abstract. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 3 pages.

1 INTRODUCTION

Many of the domains autonomous agents are being developed for are self-interested, agents will choose actions that are most beneficial to their own interests instead of maximizing a desirable quantity like social welfare. While multi-agent motion planning is a well-studied problem with several elegant solutions e.g. [2] [17] [15], most existing work assumes all agents will execute the plan they are provided. When a human operates a car, drone, or electric vertical take-off and landing vehicle (eVTOL), they may disregard a provided plan in order to take an alternative plan that is more beneficial to them. As we expect similar behavior from self-interested autonomous agents, our algorithms must provide plans that guarantee each agent is acting optimally given all other agent strategies. This desirable property is called a Nash Equilibrium (NE) which is computationally hard [3]. Additionally, since multi-agent systems often suffer the "curse of dimensionality", the state space scales exponentially with the number of agents, we argue a higher level of abstraction than operating on the control sequences can enable greater scalability relative to existing planners.

Our Contribution: In this paper, we argue the multi-agent motion planning problem should be solved at a higher level of abstraction. Instead of having each agent optimize over their control sequence, they choose a set of waypoints that are interpolated through natural cubic splines. Each agent independently minimizes their cost obtained by sampling the interpolated trajectory at a high frequency through simultaneous gradient descent. Under common assumptions of each agent's cost function, the game is a potential game and this process will converge to a local NE, which gives stable open-loop trajectories for a low level trajectory tracking controller to execute. This is in contrast to existing approaches that implicitly satisfy the dynamics constraints by extending methods like DDP/iLQR (e.g. [4] [10] [7]) or explicitly satisfy them through constraints on optimization problems (e.g. [19] [1]). This bi-level formulation is similar to [14] but we utilize a gradient based algorithm which converges due to the potential game formulation instead of relying on an off-the-shelf non-convex optimizer.

2 POTENTIAL GAMES ON CUBIC SPLINES

Each agent $i \in [N]$ where $[N] := \{1, \ldots, N\}$ chooses a strategy consisting of $S+1$ points $p_i := (p_i^1, \ldots, p_i^{S+1}) \in A_i := \mathbb{R}^{(S+1)d}$ equally spaced in time for game horizon T and dimension d. These points uniquely define a natural cubic spline [11] with S segments which guarantees C^2 continuity. Let $(x_i, \dot{x}_i, \ddot{x}_i) \in X_i := \mathbb{R}^{3dm}$ denote the position, velocity, and acceleration sequence of length m of the agent, obtained by sampling the interpolated spline at a frequency $\frac{T}{m}$. Let $A := \prod_{i \in [N]} A_i$ denote the joint action space and let $X := \prod_{i \in [N]} X_i$ denote the joint state space. Note the map $M : A \to X$ that produces the position, velocity, and acceleration sequences for all agents from the joint action is linear.

For notational ease, we define the cost functions over the position, velocity, and acceleration sequence $(x, \dot{x}, \ddot{x})$, which is justified since M is linear. Each agent has a personal cost (e.g. distance to goal, energy consumption) $J_i^{\text{per}} : X_i \to \mathbb{R}$ and pairwise cost (e.g. collision cost) $J_{i,j}^{\text{pair}} : X_i \times X_j \to \mathbb{R}$. Let $J_i^{\text{pair}}(x_i, x_{-i}) = \sum_{j \in \{-i\}} J_{i,j}^{\text{pair}}(x_i, x_j)$ and let $J_i(x, \dot{x}, \ddot{x}) = J_i^{\text{per}}(x_i, \dot{x}_i, \ddot{x}_i) + J_i^{\text{pair}}(x_i, x_{-i})$ be agent i's total cost function. Additionally, we make the following assumptions (1) each J_i is continuously differentiable and (2) the pairwise cost functions are symmetric i.e. $J_{i,j}^{\text{pair}}(x_i, x_j) = J_{j,i}^{\text{pair}}(x_j, x_i)$. Using these assumptions and the cost decomposition into pairwise and personal functions it is straightforward to show the game is

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

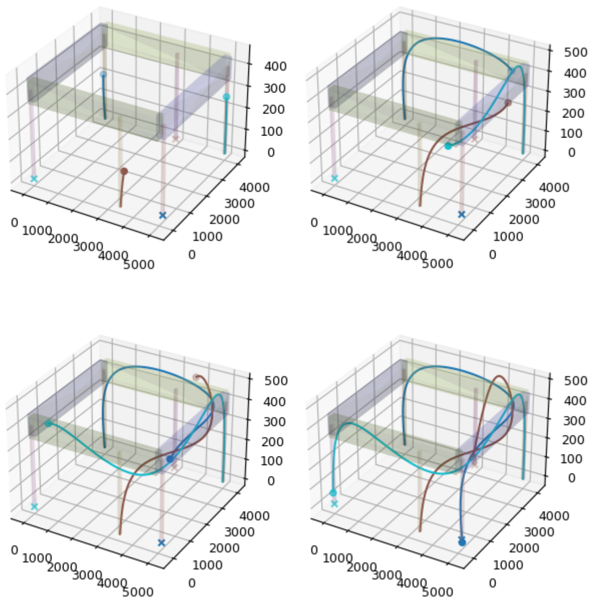

Figure 1: Snapshots of the trajectory of a simulated 3 agent eVTOL scenario. Each agent is softly constrained to fly within the UAM corridors depicted by the colored volumes and is attempting to reach the associated colored goal marked with an x.

a potential game [13] (see [19] [10] [18]). This model modifies the standard definition of a potential differential game [6] to parameterize trajectories with cubic splines. Formally, let the tuple $\Gamma = ([N], S, T, \{J_i^{\text{per}}\}_{i \in [N]}, \{J_i^{\text{pair}}\}_{i \in [N]})$ denote a motion planning game.

2.1 Simultaneous Gradient Descent

In order to improve a candidate strategy p_i, Each agent simultaneously computes the gradient of their cost function J_i with respect to each of their spline control points using chain rule.

$$\nabla_{p_i} J_i(M(p)) = M^T \nabla_{z_i} J_i(M(p)) \tag{1}$$

Each agent takes a small independent step α_i in the direction of $-\nabla_{p_i} J_i$. Only after the step do the agents re-synchronize to obtain the new strategy profiles of each other agent. Since iterated best response converges to NE in potential games [16] [12], it is straightforward to extend the argument to prove as long as α_i is appropriately small, this step decreases the potential function and will converge to a local optimum of the potential function. Local optimality is a sufficient condition to guarantee the joint strategy is a local NE since no agent has a nearby improving move. Just as the gradient descent algorithm achieves practical success in many non-convex optimization problems e.g. [8], we observe practical success in this setting.

3 EXPERIMENTS

We validate the feasibility of our method by generating motion plans for eVTOL trajectories where each agent is encouraged (but not constrained) to remain within predefined "sky-lanes" or UAM corridors. Autonomous electric vertical take-off and landing vehicles (eVTOL) are expected to be deployed in a wide range of domains. For example, disaster relief and law enforcement [5], and for use as air-taxis in urban areas [9], also referred to as urban air mobility (UAM). Each of these environments include other self-interested agents and their planners must be capable of reasoning over the goals of these other agents. While NASA and the FAA are still refining implementation details around UAM, initial plans confine the vehicles to UAM corridors, a sort of "sky-lane" for traffic to follow [9]. These vehicles must take off and land vertically from established vertiports, and future plans for traffic control involve sharing flight plans with other vehicles in a form of V2V communication. These considerations make our method a natural candidate to solve for local NE trajectories in this domain.

We model the multi-agent eVTOL motion planning problem as a motion planning game. The personal cost function is a weighted sum of the following terms; the squared distance of the final position to a pre-defined goal position $(x_i^m - g_i)^2$, the squared speed of the eVTOL $\sum_{k \in [m]} \dot{x}_i^k$, the squared acceleration of the eVTOL $\sum_{k \in [m]} \ddot{x}_i^k$, and the minimum squared distance to a UAM corridor $\sum_{k \in [m]} \delta(x_i^k, C)^2$ where $C \subset \mathbb{R}^d$ denotes the set of corridor space and δ denotes the euclidean distance function. The pairwise cost function is an exponential decay on the negative absolute distance between agents, minus some minimum defined safety distance $\sum_{k \in [m]} \exp(-[|x_i^k - x_j^k| - d])$. An example solved trajectory can be found in figure 1. The splines in this problem have 6 segments and are are initialized as a linear interpolation from each agent's starting position to their goal position. The horizon T of this problem is 400 seconds. Any method that operates directly on the control sequence would have orders of magnitude more variables to optimize over due to the long horizon of this problem, whereas our method solves this example in less than 2 seconds.

4 CONCLUSION

The higher level of abstraction provided by modeling the multi-agent motion planning problem as a potential game on cubic splines allows for solving games with significantly longer horizons than existing approaches in the literature. We utilize the practical success of gradient descent to converge to local NE of the game.

In the future, we will investigate expanding our methods to enable more complex specification such as those with a variable arrival time. Many other simple improvements can be considered such as allowing agents to have non-uniform splines, different choices of boundary conditions, and different numbers of spline segments which would help model a broader and more realistic range of possible joint trajectories.

ACKNOWLEDGEMENTS

The authors would like to thank the anonymous reviewers for their feedback. This work is supported by the National Science Foundation under the CAREER Award SHF-2048094, FMitF award CCF-1837131, CPS award CNS-1932620, by Toyota R&D and Siemens Corporate R&D through the USC Center for Autonomy and AI, and through a grant from the Airbus Institute for Engineering Research.

REFERENCES

[1] Simon Le Cleac'h, Mac Schwager, and Zachary Manchester. 2020. ALGAMES: A Fast Solver for Constrained Dynamic Games. In *Robotics: Science and Systems XVI*. https://doi.org/10.15607/RSS.2020.XVI.091 arXiv:1910.09713 [cs].

[2] Liron Cohen, Tansel Uras, T. K. Satish Kumar, and Sven Koenig. 2019. Optimal and Bounded-Suboptimal Multi-Agent Motion Planning. In *Twelfth Annual Symposium on Combinatorial Search*. https://www.aaai.org/ocs/index.php/SOCS/SOCS19/paper/view/18327

[3] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. 2009. The complexity of computing a Nash equilibrium. *Commun. ACM* 52, 2 (Feb. 2009), 89–97. https://doi.org/10.1145/1461928.1461951

[4] Bolei Di and Andrew Lamperski. 2018. Differential Dynamic Programming for Nonlinear Dynamic Games. http://arxiv.org/abs/1809.08302

[5] Johnny T. Doo, Marilena D. Pavel, Arnaud Didey, Craig Hange, Nathan P. Diller, Michael A. Tsairides, Michael Smith, Edward Bennet, Michael Bromfield, and Jessie Mooberry. 2021. *NASA Electric Vertical Takeoff and Landing (eVTOL) Aircraft Technology for Public Services – A White Paper.* Technical Report. https://ntrs.nasa.gov/citations/20205000636

[6] Alejandra Fonseca-Morales and Onésimo Hernández-Lerma. 2018. Potential Differential Games. *Dynamic Games and Applications* 8, 2 (June 2018), 254–279. https://doi.org/10.1007/s13235-017-0218-6

[7] David Fridovich-Keil, Ellis Ratner, Lasse Peters, Anca D Dragan, and Claire J Tomlin. 2020. Efficient iterative linear-quadratic approximations for nonlinear multi-player general-sum differential games. In *2020 IEEE international conference on robotics and automation (ICRA)*. IEEE, 1475–1481.

[8] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17)*. Curran Associates Inc., Red Hook, NY, USA, 6629–6640.

[9] Brian P. Hill, Dwight DeCarme, Matt Metcalfe, Christine Griffin, Sterling Wiggins, Chris Metts, Bill Bastedo, Michael D. Patterson, and Nancy L. Mendonca. 2020. *UAM Vision Concept of Operations (ConOps) UAM Maturity Level (UML) 4.* Technical Report. https://ntrs.nasa.gov/citations/20205011091

[10] Talha Kavuncu, Ayberk Yaraneri, and Negar Mehr. 2021. Potential iLQR: A Potential-Minimizing Controller for Planning Multi-Agent Interactive Trajectories. In *Robotics: Science and Systems XVII*. Robotics: Science and Systems Foundation. https://doi.org/10.15607/RSS.2021.XVII.084

[11] Gary D. Knott. 2000. Smoothing Splines. In *Interpolating Cubic Splines*, Gary D. Knott (Ed.). Birkhäuser, Boston, MA, 123–132. https://doi.org/10.1007/978-1-4612-1320-8_10

[12] Nikolai S. Kukushkin. 1999. Potential games: a purely ordinal approach. *Economics Letters* 64, 3 (Sept. 1999), 279–283. https://doi.org/10.1016/S0165-1765(99)00112-3

[13] Dov Monderer and Lloyd S. Shapley. 1996. Potential Games. *Games and Economic Behavior* 14, 1 (May 1996), 124–143. https://doi.org/10.1006/game.1996.0044

[14] Yash Vardhan Pant, Houssam Abbas, Rhudii A. Quaye, and Rahul Mangharam. 2018. Fly-by-Logic: Control of Multi-Drone Fleets with Temporal Logic Objectives. In *2018 ACM/IEEE 9th International Conference on Cyber-Physical Systems (ICCPS)*. 186–197. https://doi.org/10.1109/ICCPS.2018.00026

[15] Kiril Solovey, Oren Salzman, and Dan Halperin. 2016. Finding a needle in an exponential haystack: Discrete RRT for exploration of implicit roadmaps in multi-robot motion planning. *The International Journal of Robotics Research* 35, 5 (April 2016), 501–513. https://doi.org/10.1177/0278364915615688

[16] Brian Swenson, Ryan Murray, and Soummya Kar. 2018. On Best-Response Dynamics in Potential Games. *SIAM Journal on Control and Optimization* 56, 4 (Jan. 2018), 2734–2767. https://doi.org/10.1137/17M1139461

[17] Jesus Tordesillas and Jonathan P. How. 2022. MADER: Trajectory Planner in Multiagent and Dynamic Environments. *IEEE Transactions on Robotics* 38, 1 (Feb. 2022), 463–476. https://doi.org/10.1109/TRO.2021.3080235

[18] Zach Williams, Jushan Chen, and Negar Mehr. 2023. Distributed Potential iLQR: Scalable Game-Theoretic Trajectory Planning for Multi-Agent Interactions. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. 01–07. https://doi.org/10.1109/ICRA48891.2023.10161176

[19] Alessandro Zanardi, Enrico Mion, Mattia Bruschetta, Saverio Bolognani, Andrea Censi, and Emilio Frazzoli. 2021. Urban Driving Games With Lexicographic Preferences and Socially Efficient Nash Equilibria. *IEEE Robotics and Automation Letters* 6, 3 (July 2021), 4978–4985. https://doi.org/10.1109/LRA.2021.3068657