Policies with Sparse Inter-Agent Dependencies in Dynamic Games: A Dynamic Programming Approach

Extended Abstract

Xinjie Liu
The University of Texas at Austin
Austin, Texas, United States
xinjie-liu@utexas.edu

Jingqi Li
University of California, Berkeley
Berkeley, California, United States
jingqili@berkeley.edu

Filippos Fotiadis
The University of Texas at Austin
Austin, Texas, United States
filippos.fotiadis@austin.utexas.edu

Mustafa O. Karabag
The University of Texas at Austin
Austin, Texas, United States
karabag@utexas.edu

Jesse Milzman
Army Research Laboratory
Adelphi, Maryland, United States
jesse.m.milzman.civ@army.mil

David Fridovich-Keil
The University of Texas at Austin
Austin, Texas, United States
dfk@utexas.edu

Ufuk Topcu
The University of Texas at Austin
Austin, Texas, United States
utopcu@utexas.edu

ABSTRACT

Common feedback strategies in multi-agent dynamic games require *all* players' state information to compute control strategies. However, in real-world scenarios, sensing and communication limitations between agents make full state feedback expensive or impractical, and such strategies can become fragile when state information from other agents is inaccurate. To this end, we propose a regularized dynamic programming approach for finding *sparse* feedback policies that selectively depend on the states of a subset of agents in dynamic games. The proposed approach solves convex adaptive group Lasso problems to compute sparse policies approximating Nash equilibrium solutions. We prove the regularized solutions' asymptotic convergence to a neighborhood of Nash equilibrium policies in linear-quadratic (LQ) games. We extend the proposed approach to general non-LQ games via an iterative algorithm. Empirical results in multi-robot interaction scenarios show that the proposed approach effectively computes feedback policies with varying sparsity levels. When agents have noisy observations of other agents' states, simulation results indicate that the proposed regularized policies consistently achieve lower costs than standard Nash equilibrium policies by up to 77% for all interacting agents whose costs are coupled with other agents' states.

KEYWORDS

Noncooperative Dynamic Games; Feedback Nash Equilibrium; Information Sparsity

ACM Reference Format:
Xinjie Liu, Jingqi Li, Filippos Fotiadis, Mustafa O. Karabag, Jesse Milzman, David Fridovich-Keil, and Ufuk Topcu. 2025. Policies with Sparse Inter-Agent Dependencies in Dynamic Games: A Dynamic Programming Approach: Extended Abstract. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 3 pages.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

1 INTRODUCTION

Dynamic game theory models the decisions of multiple interacting agents over time. In such games, it is common to identify *full state feedback* strategies that depend on *all* players' states. For example, in multi-robot formations, each robot typically plans its actions based on the states of other robots. However, obtaining full states is often impractical due to sensing and communication limitations. Worse, "dense" strategies that require access to many other agents' states can be brittle when such state information is inaccurate, e.g., in the presence of uncertainties. Consequently, it is desirable for agents to find strategies that selectively depend on the states of a subset of agents while still approximating equilibrium behavior.

We contribute an algorithm for finding *sparse* feedback policies that depend on fewer influential agents' states in dynamic games while approximating Nash equilibrium strategies. We propose a regularized dynamic programming (DP) scheme [1, 2] that approximately solves linear-quadratic (LQ) dynamic games, which are an extension of the linear-quadratic regulator problem [7] to multi-agent settings. The proposed approach solves a *convex* adaptive group Lasso regularization problem [12] to encode sparsity within each DP iteration. A user can choose a desired sparsity level based on the available sensing or communication resources. We also employ an iterative linear-quadratic approximation technique [5, 9] to extend the proposed approach to general non-LQ games.

2 NONCOOPERATIVE DYNAMIC GAMES

LQ Games. We study noncooperative dynamic games played by N players in discrete time $t \in [T]$. First, we introduce LQ games:

Definition 1. *An N-player, general-sum, discrete-time dynamic game is a linear-quadratic (LQ) game if each player i seeks to optimize*

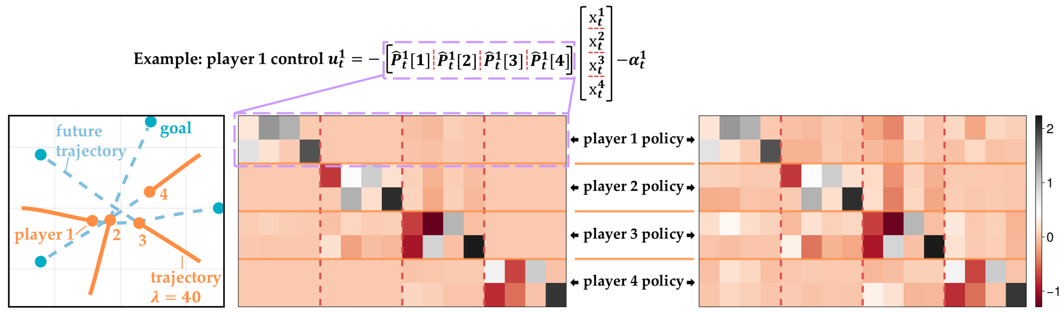

(a) t = 73 (b) regularized policy $\hat{P}_t$ (sparse) (c) unregularized policy P_t (dense)

Figure 1: (a) Snapshot of the navigation game. (b-c): Regularized and standard Nash equilibrium policy matrices.

a quadratic cost function: $\min_{u^i} J^i = \frac{1}{2}\sum_{t=1}^{T}\Big((x_t^\top Q_t^i + 2q_t^{i\top})x_t + \sum_{j=1}^{N}(u_t^{j\top} R_t^{ij} + 2r_t^{ij\top})u_t^j\Big)$, *with players following a joint linear dynamical system* $x_{t+1} = A_t x_t + \sum_{i=1}^{N} B_t^i u_t^i$.

We use x and u to denote players' states and controls. Superscripts denote players' indices and subscripts are discrete time steps, e.g., u_t^i denotes player i's control at time step t. The absence of player and time indices without additional definition denotes concatenation. We require $x_t^i \in \mathbb{R}^{m^i}, u_t^i \in \mathbb{R}^{n^i}$, and $Q_t^i \succeq 0, R_t^{ii} \succ 0, \forall i \in [N]$. We let $m = \sum_{i\in[N]} m^i$ and $n = \sum_{i\in[N]} n^i$. The initial state of the game x_1 is a given a priori.

Non-LQ Games. When the cost functions are non-quadratic and/or the dynamical system is non-linear, we have general non-LQ games.

This work seeks to find *regularized* feedback Nash equilibrium strategies for the games above to approximate standard Nash equilibrium [1, Def. 6.2] while being sparse.

3 APPROACH

LQ Games. A *standard* feedback Nash equilibrium of the game in Definition 1 can be computed via a DP procedure [1, 4], starting from the game stage $t = T$ and going backwards in time. A feedback Nash equilibrium policy to the LQ game in Definition 1 takes a linear form $u_t^{i*} = \gamma_t^{i*}(x_t) = -P_t^i x_t - \alpha_t^i, \forall t \in [T]$. We refer to P_t^i as a *policy matrix*, which maps players' states to player i's controls.

At each time step t, the DP procedure involves solving a linear system of equations: $S_t P_t = Y_t$, where $P_t = [P_t^{1\top}, \ldots, P_t^{N\top}]^\top$ is the concatenation of all players' policy matrices and S_t, Y_t are quantities related to the game parameters.

Regularization. Solving the system $S_t P_t = Y_t$ exactly computes the feedback parts of the *standard* Nash equilibrium strategies. To compute sparse policies $\hat{P}_t$, we propose to solve a *regularized* problem:

$$\min_{\hat{P}_t} \frac{1}{2}\|S_t\hat{P}_t - Y_t\|_F^2 + \sum_{i,j}\lambda_{i,j}\|\hat{P}_t^i[j]\|_F, \tag{1}$$

where $\|\cdot\|_F$ denotes matrix Frobenius norm and $\hat{P}_t^i[j]$ denotes a *block* in the policy matrix that maps the j^{th} player's states to the i^{th} player's controls. For example, in a 4-player game, a policy matrix P_t or $\hat{P}_t$ shown via a heatmap in Fig. 1 is divided into 4×4 blocks. $\lambda_{i,j}$ denotes a weighting constant that determines the regularization strength for block (i, j). We note that we choose $\lambda_{i,i} = 0, \forall i \in [N]$ and $\lambda_{i,j} = \lambda, \forall i \neq j, \lambda \in \mathbb{R}_{\geq 0}$, to not discourage the players' strategies from depending on their own states and to penalize other blocks evenly.

Importantly, the problem in Eq. (1) encourages sparsity in the solution at a *group* level, i.e., the entries in a group all remain non-zero or get zeroed out. The problem in Eq. (1) is a convex group Lasso problem [13] and can be solved using established algorithms [3, 8, 11, 13] or cast as a conic program and solved via conic optimization solvers [6].

At each dynamic programming iteration t, we solve the problem in Eq. (1) to compute regularized Nash equilibrium strategies. The procedure maintains *convexity* at each step.

Non-LQ Games. For non-LQ games, we employ an iterative algorithm [5] to repeatedly find LQ approximations of the original dynamic game and compute sparse, approximate feedback Nash equilibrium strategies.

4 RESULTS

We first test our approach in a multi-agent navigation game[1]. As is shown in Fig. 1 (a), four agents start from the initial positions and drive to their individual goals noncooperatively while avoiding collisions. Agents need to compete and find underlying Nash equilibrium strategies to reach their goals efficiently. As shown in Fig. 1, the proposed approach computes a more sparse feedback policy in Fig. 1 (b) than the standard Nash equilibrium in Fig. 1 (c). The regularization selectively decouples agents' policies from other agents' states, e.g., agent 2 has already passed by agent 1 and their dependencies are zeroed out in the sparse policy.

More Results. For more theoretical results on the convergence of the proposed DP approach and empirical results, please refer to our technical report [10].

ACKNOWLEDGMENTS

We thank Yue Yu for his helpful discussions. This work was supported in part by the Office of Naval Research (ONR) under Grants ONR N00014-24-1-2797 and ONR N00014-22-1-2703, and in part by the National Science Foundation (NSF) under Grant No. 2409535, and in part by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-23-2-0011.

[1]Supplementary video: https://xinjie-liu.github.io/projects/sparse-games

REFERENCES

[1] Tamer Başar and Geert Jan Olsder. 1998. *Dynamic noncooperative game theory.* SIAM.

[2] Dimitri Bertsekas. 2012. *Dynamic programming and optimal control: Volume I.* Vol. 4. Athena scientific.

[3] Peter Bühlmann and Sara Van De Geer. 2011. *Statistics for high-dimensional data: methods, theory and applications.* Springer Science & Business Media.

[4] David Fridovich-Keil. 2024. *Smooth Game Theory.* https://clearoboticslab.github.io/documents/smooth_game_theory.pdf

[5] 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.

[6] Paul J. Goulart and Yuwen Chen. 2024. Clarabel: An interior-point solver for conic programs with quadratic objectives. arXiv:2405.12762 [math.OC]

[7] Rudolf Emil Kalman et al. 1960. Contributions to the theory of optimal control. *Bol. soc. mat. mexicana* 5, 2 (1960), 102–119.

[8] Yuwon Kim, Jinseog Kim, and Yongdai Kim. 2006. Blockwise sparse regression. *Statistica Sinica* (2006), 375–390.

[9] Forrest Laine, David Fridovich-Keil, Chih-Yuan Chiu, and Claire Tomlin. 2021. The computation of approximate generalized feedback Nash equilibria. *arXiv preprint arXiv:2101.02900* (2021).

[10] Xinjie Liu, Jingqi Li, Filippos Fotiadis, Mustafa O Karabag, Jesse Milzman, David Fridovich-Keil, and Ufuk Topcu. 2024. Policies with Sparse Inter-Agent Dependencies in Dynamic Games: A Dynamic Programming Approach. *arXiv preprint arXiv:2410.16441* (2024).

[11] Lukas Meier, Sara Van De Geer, and Peter Bühlmann. 2008. The group lasso for logistic regression. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 70, 1 (2008), 53–71.

[12] Hansheng Wang and Chenlei Leng. 2008. A note on adaptive group lasso. *Computational statistics & data analysis* 52, 12 (2008), 5277–5286.

[13] Ming Yuan and Yi Lin. 2006. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 68, 1 (2006), 49–67.