

Solving N-Player Dynamic Routing Games with Congestion: a Mean-Field Approach

Extended Abstract

Theophile Cabannes
UC Berkeley, Google
United States of America
theophile@berkeley.edu

Mathieu Laurière
Google Research, Brain team
Paris, France

Julien Perolat
DeepMind
France

Raphael Marinier
Google Research, Brain team
Paris, France

Sertan Girgin
Google Research, Brain team
Paris, France

Sarah Perrin
Univ. Lille, CNRS, Inria, Centrale Lille,
UMR 9189 CRISTAL
Paris, France

Olivier Pietquin
Google Research, Brain team
Paris, France

Alexandre M. Bayen
University of California, Berkeley
United States of America

Eric Goubault
LIX, CNRS, Ecole Polytechnique, IPP
France

Romuald Elie
DeepMind
France

ABSTRACT

The recent emergence of navigational tools has changed traffic patterns and has now enabled new types of congestion-aware routing control like dynamic road pricing. Using the fundamental diagram of traffic flows – applied in macroscopic and mesoscopic traffic modeling – the article introduces a new N -player dynamic routing game with explicit congestion dynamics. The model is well-posed and can reproduce heterogeneous departure times and congestion spill back phenomena. However, as Nash equilibrium computations are PPAD-complete, solving the game becomes intractable for large but realistic numbers of vehicles N . Therefore, the corresponding mean field game is also introduced. Experiments were performed on several classical benchmark networks of the traffic community: the Pigou, Braess, and Sioux Falls networks with heterogeneous origin, destination and departure time tuples. The Pigou and the Braess examples reveal that the mean field approximation is generally very accurate and computationally efficient as soon as the number of vehicles exceeds a few dozen. On the Sioux Falls network (76 links, 100 time steps), this approach enables learning traffic dynamics with more than 14,000 vehicles.

KEYWORDS

Routing; Mean field games; Road traffic simulation

ACM Reference Format:

Theophile Cabannes, Mathieu Laurière, Julien Perolat, Raphael Marinier, Sertan Girgin, Sarah Perrin, Olivier Pietquin, Alexandre M. Bayen, Eric Goubault, and Romuald Elie. 2022. Solving N-Player Dynamic Routing Games with Congestion: a Mean-Field Approach: Extended Abstract. In *Proc.*

of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), Online, May 9–13, 2022, IFAAMAS, 3 pages.

This article is an extended abstract of [2].

1 INTRODUCTION

In 2019, the Texas A&M Transportation Institute estimated that the U.S. loses \$166 billion per year due to the impact of congestion on fuel usage and productivity loss [9]. The recent emergence of navigational tools has changed traffic patterns [3] and has now enabled new types of congestion-aware routing control like dynamic road pricing.

Using the fundamental diagram of traffic flows – applied in macroscopic and mesoscopic traffic modeling – the work introduces a new N -player dynamic routing game with explicit congestion dynamics. However, as Nash equilibrium computations are PPAD-complete, solving the game becomes intractable for large but realistic numbers of vehicles N [4]. Therefore, the corresponding mean field game is also introduced. This work demonstrates numerically that the MFG provides an efficient way to approximately solve the finite-player routing game with very large population: the MFG is much less costly to solve and yet provides a very good approximate Nash equilibrium policy.

2 DYNAMIC N-PLAYER AND MEAN FIELD ROUTING GAMES

The dynamic N -player routing game introduced models the evolution of N vehicles in a road network. It is a mesoscopic traffic model with explicit congestion dynamics where vehicle minimizes their travel time. The vehicles are described by their current link location, the time they have to spend on the link before exiting it, and their destination. The action of a vehicle is the successor link

Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online. © 2022 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

they want to reach when exiting a given link. When arriving on a link, the waiting time of the player is assigned based on the number of players on the link at this time. Over time steps, the waiting time linearly decreases until it is negative, then the vehicle moves to a successor link and the waiting time gets reassigned. The cost of the vehicle is its travel time, it could be seen as a running cost that increases linearly with the time until the vehicle reaches its destination.

The corresponding mean field routing game is the extension of the dynamic routing game. The list of vehicles describing the N player of the dynamic routing game is replaced by a distribution of trips between an origin and a destination. This game is a variant of the mean field route choice game [8] as the vehicle movement depends on the current network congestion. In the mean field route choice game, the number of time steps to reach the destination is constant and does not depend on the network congestion, neither of the vehicle cost function. In the dynamic driving and routing games the vehicle chooses its speed to travel on each link in order to minimize its cost function. Therefore the congestion is encoded in the cost function.

Both models are defined more precisely in [2].

3 EXPERIMENTS

All the experiments are conducted within the OpenSpiel framework [5], an open source library that contains a collection of environments and algorithms to apply reinforcement learning and other optimization algorithms in games. Experiment can be reproduced using the files in the github¹ following the instructions in the readme.md file.

We compare the running time of the algorithms for solving the N -player game and the mean field player game depending on the number of players it models. The counterfactual regret minimization with external sampling (ext CFR) [11] is used in the N -player game, as it is the fastest algorithm to solve the dynamic routing N -player game within the OpenSpiel library of algorithms (comparison done within the OpenSpiel framework are not reported here). Online mirror descent (OMD) [6] is used in the MFG. Comparison between the running time of 10 iterations of ext CFR and OMD are done as a function of the number of vehicles modeled in fig. 1. As the mean field Nash equilibrium does not depend on the number of vehicles the MFG models, the computation time of 10 iterations of OMD is independent of the number of vehicles modeled.

In the Pigou network game [7] and the Braess network game [1], the mean field equilibrium policy is almost a Nash equilibrium in the N -player game as soon as N is larger than 20 players, see fig. 2.

On the Sioux Falls network (76 links, 100 time steps) [10], this approach enables learning traffic dynamics with more than 14,000 vehicles.

4 CONCLUSION

When using game theory models where agents make rational choices, to deal with the scalability in terms of number of agents, the main advantage of the mean field approach is that mean field Nash equilibria are easier to compute while providing a good surrogate for

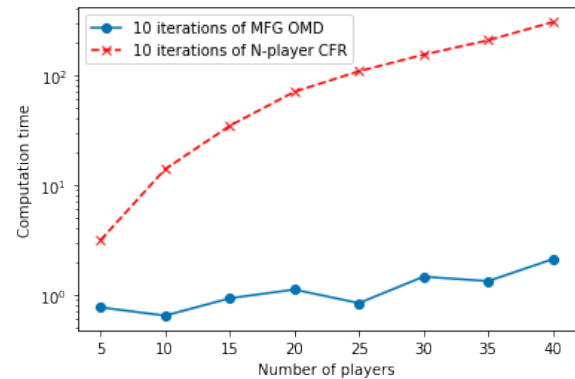


Figure 1: Computation time of 10 iterations of Online Mirror Descent in the MFG and of 10 iterations of sampled Counterfactual regret minimization as a function of the number of players N .




Figure 2: Average deviation incentive of the Nash equilibrium mean field policy in the N -player game as a function of N in the case of the Pigou game. The sampled value is the value computed in OpenSpiel by testing all the possible pure best responses, and sampling game trajectories to get the expected returns.

the equilibrium behavior in games with a finite but large number of players. Detailed illustrations are provided on several numerical examples in this article. In particular, besides toy-examples, this approach enables solving dynamic routing from a game perspective on a classical benchmark of the traffic community with 76 links and 14,000 vehicles.

REFERENCES

- [1] Dietrich Braess, Anna Nagurney, and Tina Wakolbinger. 2005. On a paradox of traffic planning. *Transportation science* 39, 4 (2005), 446–450.
- [2] Théophile Cabannes, Mathieu Lauriere, Julien Perolat, Raphael Marinier, Sertan Girgin, Sarah Perrin, Olivier Pietquin, Alexandre M Bayen, Eric Goubault, and Romuald Elie. 2021. Solving N -player dynamic routing games with congestion: a mean field approach. *arXiv preprint arXiv:2110.11943* (2021).
- [3] Théophile Cabannes, Marco Antonio Sangiovanni Vincentelli, Alexander Sundt, Hippolyte Signargout, Emily Porter, Vincent Fighiera, Juliette Ugirumurera, and Alexandre M Bayen. 2018. The impact of GPS-enabled shortest path routing on mobility: a game theoretic approach. *Transportation Research Board 97th Annual Meeting* (2018).

¹https://github.com/deepmind/open_spiel/tree/master/open_spiel/data/paper_data/routing_game_experiments

- Meeting* (2018).
- [4] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. 2009. The complexity of computing a Nash equilibrium. *SIAM J. Comput.* 39, 1 (2009), 195–259.
- [5] Marc Lanctot, Edward Lockhart, Jean-Baptiste Lespiau, Vinicius Zambaldi, Satyaki Upadhyay, Julien Pérolat, Sriram Srinivasan, Finbarr Timbers, Karl Tuyls, Shayegan Omidshafiei, Daniel Hennes, Dustin Morrill, Paul Muller, Timo Ewalds, Ryan Faulkner, János Kramár, Bart De Vylder, Brennan Saeta, James Bradbury, David Ding, Sebastian Borgeaud, Matthew Lai, Julian Schrittwieser, Thomas Anthony, Edward Hughes, Ivo Danihelka, and Jonah Ryan-Davis. 2019. OpenSpiel: A Framework for Reinforcement Learning in Games. *CoRR* abs/1908.09453 (2019). arXiv:1908.09453 [cs.LG] <http://arxiv.org/abs/1908.09453>
- [6] Julien Perolat, Sarah Perrin, Romuald Elie, Mathieu Laurière, Georgios Piliouras, Matthieu Geist, Karl Tuyls, and Olivier Pietquin. 2021. Scaling up Mean Field Games with Online Mirror Descent. *arXiv preprint arXiv:2103.00623* (2021).
- [7] Arthur Cecil Pigou. 1920. *The economics of welfare*. Palgrave Macmillan.
- [8] Rabih Salhab, Jerome Le Ny, and Roland P Malhamé. 2018. A mean field route choice game model. In *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, 1005–1010.
- [9] David Schrank, Bill Eisele, and Tim Lomax. 2019. 2019 Urban Mobility Report. *Texas A&M Transportation Institute* August (2019).
- [10] Transportation Networks for Research Core Team. [n.d.]. Transportation Networks for Research, Free of use. <https://github.com/bstabler/TransportationNetworks>. Accessed: 2021-09-01.
- [11] Martin Zinkevich, Michael Johanson, Michael Bowling, and Carmelo Piccione. 2007. Regret minimization in games with incomplete information. *Advances in neural information processing systems* 20 (2007), 1729–1736.