

Solving N-player dynamic routing games with congestion: a mean field approach

Theophile Cabannes^{1, 2, 3}, Mathieu Laurière⁵, Julien Perolat⁶, Raphael Marinier⁵, Sertan Girgin⁵, Sarah Perrin⁷, Olivier Pietquin⁵, Alexandre M. Bayen¹, Eric Goubault³, and Romuald Elie⁶

¹University of California, Berkeley

²Google Research

³Laboratoire d'Informatique de l'Ecole Polytechnique, Ecole Polytechnique

⁵Google Research, Brain team

⁶DeepMind

⁷Univ. Lille, CNRS, Inria, Centrale Lille, UMR 9189 CRISTAL

October 2021

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.

1 Introduction

1.1 Motivations

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 [42]. The average auto commuter spends 54 hours in congestion and wastes 21 gallons of fuel every year due to congestion at a cost of \$1,010 in wasted time and fuel. With the emergence of navigational applications, the traffic patterns have evolved due to congestion-aware routing behaviors [11]. Being able to model traffic and especially routing choice adequately would enable traffic control to leverage the new routing behaviors in order to improve the network efficiency.

However, solving realistic routing choice problems requires the consideration of large scale multi-agent systems, where the number of vehicles making strategic decisions is not tractable within classical algorithms [15]. This work focuses on finding a scalable approach to model the routing behavior of each vehicle in a dynamic road traffic environment as a large multi-agent dynamic system.

1.2 Background

1.2.1 Dynamic traffic assignment

Dynamic traffic simulations are used to model the evolution of the locations of the vehicles in the road network across time. Within dynamic models, simulations are divided between macro, meso and microsimulations. Microsimulations model each individual vehicle and its interaction with others. At the other extreme, macrosimulations model traffic dynamics based on the flow of cars. Mesosimulations also model individual vehicles, but only consider the interactions between each vehicle and the flow (or volume) of vehicles. In general, these simulations have no spatial resolution on links and are mostly discrete-event models. In these simulations, a vehicle is slowed down if the volume of vehicles on the road section is large. This is modeled using fundamental diagrams of the traffic flow [44] – a congestion function mapping the volume on each link to the link travel time.

In routing models, vehicles are routed based on the notion of Wardrop equilibrium [47], in which the travel times of every used path between the same origin and the same destination for a given departure time are equal and smaller than the one on any unused path. To find the route assignment that leads to a Wardrop equilibrium, routing models are divided between one-shot assignments (that assign the routes once for all using a stochastic model and the current traffic information) and iterative assignments (that assign the routes, run a simulation, assess gap to the user equilibrium and update the routes accordingly until convergence). To assess how far a traffic state is from the user equilibrium, the relative gap to the user equilibrium is considered [14]. In fact, the Wardrop equilibrium condition interprets as a Nash equilibrium condition between the vehicles [34, chapter 18] while the relative gap to the user equilibrium is similar to the average deviation incentive or the average marginal regret [10] over the vehicles.

The most popular routing game model is arguably the non-atomic static routing game as presented e.g. in [34, chapter 18]. This non-atomic game is a potential game [32] that can be solved through convex programming [38, 35], enabling fast computation of the route choices for large networks. In the static setting, vehicles try to minimize their travel cost by choosing a path from their origin to their destination. Path choices are converted into volume on each path, which translates into a volume on each link. Link volume are converted to link cost using congestion functions [35, Table 1.1].

In static routing games, the congestion on each link does not evolve. These games cannot replicate dynamic phenomena like departure time choice and congestion spill back, and dynamic extensions of the static routing game have been introduced. First, using the potential formulation of the static traffic assignment, dynamic traffic assignment has been defined as the solution of a dynamical variational inequality [19]. A similar approach allows to consider both routing choices and departure times [22]. However, in such dynamical models, the game theoretical aspect is not explicit. In parallel, dynamic routing games have been defined [48], and their resolution via multi-agent reinforcement learning has been studied [43]. Nevertheless, current multi-agent learning algorithms do not scale in terms of population size [15]. This inherently calls for new methods to study routing choice problems with very a large number of vehicles. We next explain how a Mean field game perspective can help making an important step in this direction.

1.2.2 Mean field game

Mean field games (MFGs) have been introduced in [29, 25] to model differential game dynamics between infinitely many players with symmetric interactions. Similar to mean-field theory from statistical physics, the key idea is to use a macroscopic approximation of a large population with anonymous and symmetric players. When the population is infinite, each player has no influence on the population distribution. So in an MFG, one does not need to study the pairwise interactions between all players but simply the interaction between a representative infinitesimal player and the full population distribution of states (and possibly actions). Because players are assumed to be identical, it suffices to determine the strategy of a representative player in response to the full population behavior. From a mathematical point of view, the solution of an MFG can be characterized by a coupled system of a forward equation for the population evolution and a backward dynamic programming one for the player’s value function. This system is easier to solve than the Nash equilibrium of a finite-player dynamic game with a large number of players, which leads to a large system

of coupled Bellman equations. Intuitively, this is true when representing each player is more costly than representing a distribution over all possible states, which is the case for instance when the number of players is larger than the number of possible states.

Applications of MFGs include crowd modeling [27, 5, 2], energy management [41, 30, 3], epidemiology [16, 17, 4] or financial markets [26, 12, 18]. In traffic theory, similar approaches model cars evolution over the network as traffic flow. Modeling microscopic cars on a link as a macroscopic traffic flow has been interpreted as a mean field game e.g. in [13]. Connections between an MFG model with myopic players and the Lighthill-Whitham-Richards (LWR) model [31] on a single road have been studied [24]. A recent literature [40, 45] builds on the construction of MFGs on graphs [21] and develops the application of MFGs to road traffic management.

1.3 Related work

Several works study mean field routing games. First, continuous time models have been studied. The existence and uniqueness of the Nash equilibrium of a MFG with congestion on a graph has been shown in [21], where the state space is the set of nodes. Models in which the state space is given by the edges have been analyzed e.g. in [1], which proved existence and uniqueness for a forward-backward system of equations with suitable conditions at the vertices of the network. In [6], the authors analyzed an MFG model for traffic flow on networks by using an extended state space that includes the distribution of players on the network and they studied Wardrop equilibria. In existing discrete time MFG models for routing, the players move one edge per time step and pay a cost that increases with the proportion of players on the same edge. In [40], the authors analyzed an MFG model and studied the impact of adding or removing edges on the equilibrium traffic flow. Their work provides a discrete time resolution of a mean field routing game with on an 11 link network with 6 time steps. In [45], the authors proposed an MFG model that reduces to a linearly solvable Markov Decision Process and showed connections with Fictitious Play [9] in some cases.

The fact that existing models take congestion into account only through the cost functions leads to paradox such as an ambiguity about the definition of travel time: the graph traversal time and the player cost can differ. Such issues make these models hardly applicable for traffic engineering. Also, the main motivation for using an MFG-based routing method is to obtain an efficient equilibrium policy in the finite-player routing game, which has not been checked in existing works.

1.4 Contributions and outline

The first main contribution of this work (see section 2) is to introduce a novel dynamic mesoscopic traffic model viewed as a dynamic routing game with explicit congestion dynamics, i.e., congestion effects directly in the state evolution. We first propose a finite-player game which can easily be interpreted in terms of vehicles, and then derive the corresponding MFG. It is proved that Nash equilibria exist for both games. Furthermore, theoretical arguments supporting the MFG approximation are discussed.

The second main contribution (see section 3) is to demonstrate 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. This is illustrated on small networks for which baselines are available and on the Sioux Falls network – a real network with 76 links and 100 time steps with 14,000 vehicles. Although this network is often used as a benchmark in the literature, to the best of our knowledge, it is the first time that a method is able to solve a dynamic routing game with a very large number of vehicles on this network.

2 Dynamic N-player and mean field routing games

This section introduces the dynamic routing game and the corresponding MFG. We show that the MFG approach allows to recover the dynamic routing game Nash equilibrium with a very large number of vehicles. Before moving to the mathematical details, the model we propose can be summarized in the following way.

This dynamic routing game models the evolution of N vehicles on a road network. The vehicles are described by their current link location, the time they will spend on the link before exiting it, and their destination. The action of a vehicle is the successor link they want to reach when exiting a given link. Actions are encoded as integers from 0 to K . Pure actions for a player on link ℓ , with a negative waiting time are the successors link of ℓ . 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. As time goes by, the waiting time of a vehicle decreases until it becomes negative, then the vehicle moves to a successor link and the waiting time gets reassigned. The total cost for the vehicle is its travel time. In the corresponding MFG, the vehicles of the N -player game are replaced by a representative vehicle and the probability distribution of the vehicles states.

Notation. We denote by A^B the set of functions from a set B to a set A . If B is countable, A^B interprets as a set of sequences indexed by elements of B . We use bold letters for functions of time and underlines for vectors whose coordinates correspond to players. The set of probability measures on a space X is denoted $\mathcal{P}(X)$. Unless otherwise specified, random variables are denoted by capital letters, functions of time are denoted by bold letters, and vectors indexed by the set of players are denoted by an underlined letter.

2.1 Network and game set up

Time is represented as an interval $\mathcal{T} = [0, T]$ of \mathbb{R} . The road network is described by a directed graph $G = (\mathcal{V}, \mathcal{A})$, where \mathcal{V} and \mathcal{A} respectively denote the sets of vertices and links of the road network. When exiting a link $\ell \in \mathcal{A}$, a vehicle chooses one of the possible successor links. In case the link has no successor, the vehicle stays on the link until the end of time. When joining a link, a vehicle get assigned a travel time on this link, that depends on the volume of traffic on the link. More specifically, congestion induces a travel time spent on link $\ell \in \mathcal{A}$ which is a function $c_\ell \in \mathbb{R}_{>0}^{[0,1]}$ of the proportion of vehicles on link ℓ . We assume that c_ℓ is continuous. The congestion functions $(c_\ell)_{\ell \in \mathcal{A}}$ encode the heterogeneity of the roads' sensitivity to traffic volume within the network. The following typical congestion function is based on an example provided by the 1964 traffic assignment manual of the U.S. Bureau of public road functions, see [35, table 1.1]: $c_\ell : \mu \mapsto t_0(1 + \alpha(\mu/\mu_{\ell,c})^\beta)$, where α and β are positive constants, t_0 is the free flow travel time (i.e., the travel time when the link is empty), and $\mu_{\ell,c}$ is the relative capacity of the link ℓ (which, in our context, is to be understood as a capacity in terms of proportion of players).

Let us stress that the congestion functions are functions of the *proportion* of vehicles within the link. However, in practice, congestion effects scale with the *number* of vehicles (as well as other factors such as the width and the length of the road, which are considered fixed for a given network). One should thus interpret c_ℓ as being tuned for a given number of vehicles, say N_0 . Concretely, in the previous example, if n vehicles out of N_0 are on link ℓ , we have $c_\ell(n/N_0) = t_0(1 + \alpha(n/(N_0\mu_{\ell,c}))^\beta)$. For $N'_0 \neq N_0$, we have $c_\ell(n/N'_0) = t_0(1 + \alpha(n/(N'_0\mu_{\ell,c}))^\beta)$. So the travel time can be viewed as a function of the number of vehicles n on link ℓ , provided the relative capacity is scaled by the total number of vehicles.

2.2 N-player dynamic routing game

Given the above network, this subsection defines a finite-player game. Most of the notations are chosen to ease the presentation of the MFG in the next subsection.

2.2.1 Traffic flow environment

For the sake of convenience, we assume that the time horizon T is large enough so that any driver will have time to travel through the network.

Let N be a positive integer. The set of players is $\mathcal{N} = \{1, 2, \dots, N\}$. The number of players in the game is not necessarily the total number of vehicles N_0 in the real-life scenario. Each player of the model corresponds to a proportion of the real number of vehicles, which allows to define a player as an infinitesimal portion of flow that does not impact network travel time in the MFG. In the case where $N = N_0$, a player is a vehicle. Player $i \in \mathcal{N}$ starts at an origin link $L_0^i \in \mathcal{A}$ with a departure time $W_0^i \in \mathcal{T}$, and has a destination

link $D_0^i \in \mathcal{A}$. This is the initial state of the player. Intuitively the player wants to start moving at time W_0^i from L_0^i and tries to reach D_0^i . We assume that the players' initial state are distributed according to a finite-support distribution m_0 . Both the origin and the destination are modeled as links, so that the location of the vehicle is always described as a link. In experimental setups, a origin link is added before each origin node and a destination link is added after each destination node. Being on the origin link means having not departed yet, and being on the destination link means having finished their trip.

Then, at any time step t , the state of a player i is not only the link L_t^i where they stand, but also their waiting time W_t^i before exiting this link together with their destination D_t^i . L_t^i and W_t^i are random variables due to the randomness in the action choices. Even though the destination is constant through time ($D_t^i = D_0^i$ for all t), including this information in the state allows to keep track of the objective in the player's policy. So the state space for each driver is $\mathcal{X} = \mathcal{A} \times \mathcal{T} \times \mathcal{A}$, where the first component is for the current location and the last one is for the destination (recall that the destination is represented by a link in our model). Then, the space of vehicle trajectories is $\mathcal{X} = \mathcal{X}^{\mathcal{T}}$. The state trajectories are in the space of triples (location, waiting time, destination), which provide more information than the physical trajectories just in terms of locations. At the population level, the states of all the agents is a vector $\underline{X} = (X^i)_{i \in \mathcal{N}}$. The state space for the whole population is $\underline{\mathcal{X}} = \mathcal{X}^{\mathcal{N}}$, and the corresponding space of trajectories is $\underline{\mathcal{X}} = \mathcal{X}^{\mathcal{T} \times \mathcal{N}}$. We respectively call game state and game trajectory the state and trajectory of the population.

2.2.2 Routing policy

When at link ℓ , a player can try to move to another link among the successors of ℓ , and the transition is realized provided the waiting time is 0. The players are allowed to randomize their actions. We thus call strategy function and denote by π a function from \mathcal{X} to $\mathcal{P}(\mathcal{A})$ such that for any $x = (\ell, w, d)$, $\pi(x)$ has support in the successors of ℓ . We denote by Π the set of such strategy functions. A (feedback or closed-loop) policy π is a function that associates to each time a strategy function, so it is an element of the set $\Pi = \Pi^{\mathcal{T}}$ of policies. The notation $\pi_t(\ell' | \ell, w, d)$ represents the probability at time t with which the agent would like to go from ℓ to ℓ' given the fact that their waiting time is w and their destination is d . A policy profile $\underline{\pi}$ is a vector of policy functions with one policy for each player, i.e., it is an element of $\underline{\Pi} = \Pi^{\mathcal{N}}$. Studying this class of policies can be justified by the fact that it allows each player to take a decision based only on their own state, which is realistic if the players do not know the situation of the rest of the population. More information could be added in the inputs of the policy (e.g., the proportion of agents on the current link), but this is beyond the scope of this work.

2.2.3 State dynamics

Since the players' initial states and actions are randomized, their trajectories are stochastic. Given a policy profile $\underline{\pi} \in \underline{\Pi}$, $\underline{X}_t = (\underline{L}_t, \underline{W}_t, \underline{D}_t) \in \underline{\mathcal{X}}$ denotes the random variable corresponding to the links, waiting times and destinations for all the players at time $t \in \mathcal{T}$. The stochastic process of the population state is denoted $\underline{X} = (\underline{X}_t)_{t \in \mathcal{T}} \in \underline{\mathcal{X}}$. As indicated above, the players' interactions are through the travel time functions $(c_\ell)_{\ell \in \mathcal{A}}$ taking into account congestion levels. So the interaction between a driver and the rest of the vehicles is only through the proportion of vehicles on the same link. It is thus convenient to introduce the empirical distribution $\nu_\ell^{\mathcal{N}} \in \mathcal{P}(\mathcal{A})$ corresponding to a location profile $\underline{\ell} = (\ell^i)_{i \in \mathcal{N}} \in \mathcal{A}^{\mathcal{N}}$: for every $\ell \in \mathcal{A}$, $\nu_\ell^{\mathcal{N}}(\ell') = \frac{1}{\mathcal{N}} \#\{i \mid \ell^i = \ell'\} \in [0, 1]$, which is the proportion of players on the link ℓ' , given $\underline{\ell}$. This is all the information one needs from the game state to compute the interactions between players at link ℓ' . Note that $\nu_\ell^{\mathcal{N}}(\ell')$ is invariant by permutation of the components of the vector $\underline{\ell}$.

Let us fix a policy profile $\underline{\pi} \in \underline{\Pi}$. We denote by \underline{U} the $\mathcal{A}^{\mathcal{T} \times \mathcal{X} \times \mathcal{N}}$ -valued random variable assigned to the probability distribution given by the policy profile: for each $(t, x, i) \in \mathcal{T} \times \mathcal{X} \times \mathcal{N}$, $U_t^i(x)$ is an \mathcal{A} -valued random variable with distribution $\pi_t^i(\cdot | x)$.

The evolution of the state of the game $\underline{X}_t = (\underline{L}_t, \underline{W}_t, \underline{D}_t)$ is given by the following dynamics. At initial

time, (L_0^i, W_0^i, D_0^i) , $i \in \mathcal{N}$ are given, and then the dynamics is:

$$\begin{aligned}
t_{k+1} &= t_k + \min\{W_{t_k}^i, i \in \mathcal{N}\} \\
L_{t_{k+1}}^i &= \begin{cases} U_{t_{k+1}}^i(X_{t_k}^i) & \text{if } i \in I_{t_{k+1}} \\ L_{t_k}^i & \text{otherwise;} \end{cases} \\
W_{t_{k+1}}^i &= \begin{cases} c_{L_{t_{k+1}}^i}(\nu_{L_{t_{k+1}}^i}^N(L_{t_{k+1}}^i)) & \text{if } i \in I_{t_{k+1}} \\ W_{t_k}^i - (t_{k+1} - t_k) & \text{otherwise;} \end{cases} \\
L_t^i &= L_{t_k}^i \quad \forall k, \forall t \in [t_k, t_{k+1}[, \forall i \in \mathcal{N} \\
W_t^i &= W_{t_k}^i - (t - t_k) \quad \forall k, \forall t \in [t_k, t_{k+1}[, \forall i \in \mathcal{N} \\
D_t^i &= D_0^i, \quad t \in \mathcal{T},
\end{aligned}$$

where $I_{t_{k+1}} := \{i \in \mathcal{N}, W_{t_k}^i + t_k - t_{k+1} = 0\}$ and using $(t_k)_{k \in \mathbb{N}}$ the sequence of times where one of the vehicles changes link with $t_0 = 0$ and $t + k = T$ if all the players have arrived their destination. The destination is constant through time and is not affected by the policy's randomness. Note that $\mathbf{U}^i = (U_t^i)_{t \in \mathcal{T}}$ is defined for all t but used only when the player moves from one link to the next one, i.e., when the waiting time has vanished. This enables reducing any pure (i.e. deterministic) policy as a path choice.

2.2.4 Cost function

Given a policy profile $\underline{\pi} \in \underline{\Pi}$, the cost for player i is the average arrival time which can be defined as:

$$J_i^N(\underline{\pi}^i, \underline{\pi}^{-i}) = \mathbb{E}_{\underline{\pi}} [\min\{t \in \mathcal{T}, L_t^i = D^i\}] = \mathbb{E}_{\underline{\pi}} \left[\int_{t \in \mathcal{T}} r(X_t^i) dt \right]$$

where $\underline{\pi}^{-i} = (\underline{\pi}^1, \dots, \underline{\pi}^{i-1}, \underline{\pi}^{i+1}, \dots, \underline{\pi}^N)$, and the instantaneous cost is defined as: for every $x = (\ell, w, d)$, $r(x) = \mathbb{1}_{\ell \neq d}$. Note that the running cost is independent of the (rest of the) population state, contrary to other models for routing or crowd motion in which the interactions are not in the dynamics but in the cost function.

Furthermore, the population is homogeneous (all players have the same dynamics evolution and same running cost), and player i interacts with the other players only through ν^N and for this reason, the cost function J_i^N does not depend directly on the index i but only on $\underline{\pi}^i$: as a function, $J_i^N = J_{i'}^N$ for all i' . The policy profile $\underline{\pi}^{-i}$ for the rest of the population is used only to compute $\nu^N = (\nu_t^N)_{t \in \mathcal{T}}$. Although $\underline{\pi}^{-i}$ is necessary, it is not sufficient because ν^N is also influenced by the policy $\underline{\pi}^i$ chosen by the player under consideration. However, the influence of each player decays as N increases, which will be the basis for the mean-field approach presented in § 2.3.

2.2.5 Nash equilibrium

Considering that all the players are individually optimizing their own cost leads to the following notion of solution for the game. We refer to e.g. [33] for more details.

Definition 1 (Nash equilibrium). *A Nash equilibrium is a policy profile $\underline{\pi}^* = (\underline{\pi}^{i*})_{i \in \mathcal{N}} \in \underline{\Pi}$ such that:*

$$\forall i \in \mathcal{N}, \forall \underline{\pi} \in \underline{\Pi}, J_i^N(\underline{\pi}^{i*}, \underline{\pi}^{-i*}) \leq J_i^N(\underline{\pi}, \underline{\pi}^{-i*}).$$

The following result says that, in our model, such equilibria exist.

Theorem 1 (Existence of N -player Nash equilibrium, Kakutani-Fan-Glisckberg theorem [20]). *Assuming the continuity of the cost function with respect to the policy profiles, there exists a Nash equilibrium in the N -player routing game.*

Proof. The proportion of players on each link is always a multiple of $1/N$. Since the number of links and the time horizon are finite, there is a finite set of times at which a vehicle can switch link. The set of policy profiles can thus be restricted to a finite set. Therefore, the game can be restated as a game with finite state and action spaces. Assuming the continuity of the cost function with respect to the policy, it has a Nash equilibrium. Further development of the proof are provided in the appendix. \square

Besides the above definition, another way to express that a policy profile $\underline{\pi}$ is a Nash equilibrium is to say that the deviation incentive is 0 for every player, where the deviation incentive for player i is:

$$D_i^N(\underline{\pi}^i, \underline{\pi}^{-i}) = J_i^N(\underline{\pi}^i, \underline{\pi}^{-i}) - \underset{\pi' \in \Pi}{\operatorname{argmin}} J_i^N(\pi', \underline{\pi}^{-i}).$$

This also serves as a basis to assess the convergence of algorithms towards a Nash equilibrium using the average deviation incentive:

$$\overline{D}^N(\underline{\pi}) = \frac{1}{N} \sum_{i=1}^N D_i^N(\underline{\pi}^i, \underline{\pi}^{-i}). \quad (1)$$

2.3 Mean Field approximation

As mentioned in the introduction, solving the above N -player game is infeasible, as N is very large. We thus turn to an MFG version of the above routing game, which can be used to provide approximate Nash equilibria and whose quality improves as $N \rightarrow +\infty$. This is based on considering the interactions between a typical player and a distribution representing the rest of the population. This is possible thanks to the anonymity and the symmetry in the interactions, which allows us to focus on symmetric Nash equilibria. Intuitively, the law of large numbers allows considering the state distribution instead of numerous random variables induced by it.

2.3.1 Traffic flow environment

The state of a typical player at time t is a random variable denoted by $X_t = (L_t, W_t, D_t)$ which takes values in $\mathcal{X} = \mathcal{A} \times \mathcal{T} \times \mathcal{A}$. At time 0, the population's state distribution is m_0 and is known to the players.

2.3.2 Routing policy

The space of policies is still Π . For a policy π , we denote by $\pi_t(\ell'|\ell, w, d)$ the probability with which a typical player using policy π would like to go from ℓ to ℓ' given that their waiting time is w and their destination is d . The routing random variable is denoted by U .

2.3.3 State dynamics

Assume that an infinitesimal agent uses policy π while the rest of the population uses π' . Let $\nu = (\nu_t)_{t \in \mathcal{T}} \in \mathcal{P}(\mathcal{A})^{\mathcal{T}}$ be the flow of distributions on \mathcal{A} induced by the population that uses π' . The evolution of a typical player's state is given by the following dynamics. Let $t_0 = 0$ and let (L_0, W_0, D_0) be a given initial state. Then, the dynamics follow:

$$\begin{aligned} t_{k+1} &= W_{t_k} + t_k \\ L_{t_{k+1}} &= U_{t_k}(X_{t_k}) \\ W_{t_{k+1}} &= c_{L_{t_{k+1}}}(\nu_{t_{k+1}}(L_{t_{k+1}})) \\ L_t &= L_{t_k} \quad \forall k, \forall t \in [t_k, t_{k+1}[\\ W_t &= W_{t_k} - (t_k - t) \quad \forall k, \forall t \in [t_k, t_{k+1}[\\ D_t &= D_0, \quad t \in \mathcal{T}. \end{aligned}$$

Here $(t_k)_{k \in \mathbb{N}}$ denotes the sequence of times where the representative player changes link (we take $t_k = T$ when there are no more changes), and $\nu_t(\ell) \in [0, 1]$ is the proportion of the mean field population on link ℓ at time t .

2.3.4 Cost function

The cost of the typical player using policy π when the population uses policy π' is defined as:

$$J(\pi, \pi') = \mathbb{E}_{\pi, \pi'} \left[\int_{t \in \mathcal{T}} r(X_t) dt \right]$$

where the state of the representative player $\mathbf{X} = (X_t)_{t \in \mathcal{T}}$ has the above dynamics with policy π , and the instantaneous cost function r is the same function as in the finite player game (see § 2.2.4). Analogously to the N -player game, the policy π' is used only to deduce $\nu = (\nu_t)_{t \in \mathcal{T}}$ that appears in the evolution of \mathbf{W} . So the cost function J could alternatively be written as a function of (π, ν) instead of (π, π') . In contrast with the finite player regime, we highlight that here π' completely determines ν because the player under consideration is infinitesimal and hence their policy π does not affect the flow ν of distributions of locations of the population.

2.3.5 Nash equilibrium

The counterpart of the N -player Nash equilibrium in the mean-field regime can now be introduced.

Definition 2 (Mean field Nash equilibrium (definition 3.1. of [39])). *A mean field Nash equilibrium (MFNE) is a policy $\pi^* \in \Pi$ such that: $J(\pi^*, \pi^*) \leq J(\pi', \pi^*)$ for all π' , or equivalently:*

$$\pi^* \in \underset{\pi \in \Pi}{\operatorname{argmin}} J(\pi, \pi^*).$$

Another way to express that $\pi \in \Pi$ is a MFNE is to say that the average deviation incentive \bar{D} vanishes, where:

$$\bar{D}(\pi) = J(\pi, \pi) - \underset{\pi' \in \Pi}{\operatorname{argmin}} J(\pi', \pi).$$

Theorem 2 (Existence of mean field Nash equilibrium, Kakutani-Fan-Glisckberg theorem [20]). *Assuming the continuity of the cost function with respect to the policy profiles, and assuming that the support of the initial distribution of the waiting time is a finite set, there exists a mean field Nash equilibrium.*

Proof. The set of pure policies for the representative player can be restricted to the choice of a path given a departure time. This set is finite as long as the support of the initial distribution of waiting time is. Therefore the argmin map (also called Best response map) is a Kakutani-Fan-Glisckberg map providing the continuity of the cost with respect to the policy of the representative player and of the mean field. The proof is further developed in appendix. \square

The continuity of the cost function with respect to the policies plays a crucial role. A counter-example of the existence of a Nash equilibrium with a discontinuous cost is shown in the appendix.

One of the advantages of considering a mean field setting, is that any MFNE is automatically a dynamic Wardrop equilibrium.

Theorem 3 (Dynamic Wardrop equilibrium [47]). *For any mean field Nash equilibrium, all induced trajectories of players with the same initial state (origin, waiting time, destination), have the same travel time (i.e., the same total cost).*

Proof. In case a trajectory used by the representative player has a higher travel time than another one, then the player has an incentive to deviate, and the game is not a Nash equilibrium. The proof is further developed in the appendix. \square

Any mean field Nash equilibrium policy π^* can be used by the players in an N -player game. Intuitively, the larger N is, the closer the population is to the mean field regime. In fact, it can be shown under suitable conditions that $\underline{\pi}^* = (\pi^*, \dots, \pi^*) \in \Pi^N$ is an approximate Nash equilibrium whose quality improves with N in the sense that:

$$\overline{D}^N(\underline{\pi}^*) \rightarrow 0, \quad \text{as } N \rightarrow +\infty.$$

So if all the agents use the mean field Nash equilibrium policy, then any single player’s incentive to deviate decreases when the population becomes larger. For example, [39] prove in their setting that: if π^* is an MFNE, then for every $\epsilon > 0$, there exists $N_0 \in \mathbb{N}$ such that for every $N \geq N_0$, the N -player policy profile $(\pi^*, \dots, \pi^*) \in \Pi^N$ satisfies: $\overline{D}^N(\underline{\pi}) \leq \epsilon$.

Next, to illustrate this property in our model, an explicit computation is carried out in the simple Pigou network and then is empirically verified on both Pigou and Braess networks.

2.3.6 Mean field equilibrium policy in the N player Pigou game

For the sake of illustration, we present a toy example for which the solution can be computed analytically.

The graph has 2 nodes and 2 parallel links, say ℓ, ℓ' , relating these 2 nodes. The cost function is: $c_\ell(x) = 0.5T$, $c_{\ell'}(x) = xT$ for all $x \in [0, 1]$. The departure time (initial waiting time) is the same for all the agents. The mean field Nash equilibrium can be computed and yields an equilibrium distribution with proportions $\nu_t(\ell) = \nu_t(\ell') = 0.5$. On the other hand, the Nash equilibrium for the N player game is such that $\nu_t(\ell)$ the current proportion of players on ℓ' is included in $[\frac{1}{2} - \frac{1}{N}, \frac{1}{2}]$. As detailed in the appendix, we can check that the average deviation incentive of the mean-field equilibrium policy in the N -player game is

$$\frac{T}{N2^N} \sum_{m=1}^N \binom{N-1}{m} \max \left\{ \frac{N}{2} - m - 1, m + 1 - \frac{N}{2} \right\},$$

which goes to 0 when $N \rightarrow \infty$.

3 Experiments

This section shows experimentally that (1) computing the mean field equilibrium is easier than computing the N -player Nash equilibrium using state of the art algorithms (sampled counterfactual regret minimization [49]) and (2) it gives an excellent approximation of the N -player equilibrium when N is large (above 30 in the case of the Pigou [37] and the Braess [8] network). The experiments also show that (3) online mirror descent algorithm [36] enables computing the mean field equilibrium on the Sioux Falls network, a classic use case in road traffic network games, with 14,000 vehicles (across two origin-destination pairs) and realistic congestion function (from the open source dataset [46]).

3.1 Context

All the experiments are conducted within the OpenSpiel framework [28], an open source library that contains a collection of environments and algorithms to apply reinforcement learning and other optimization algorithms in games. The code is publicly available on GitHub¹.

Goal of the experiments. The experiments aim to show that the mean field equilibrium policy is faster to compute than the N -player policy and approximates well an equilibrium policy in the corresponding N -player game, showing that the mean field approach solves the curse of dimensionality regarding the number of players in N -player games. Intuitively, the MFG approach is relevant when the number of possible states for any player is lower than the number of players. In that case, computing the population’s distribution probabilities over the possible states is faster than simulating each player trajectory. The approximation is correct when representing the probability distribution over the state space is similar to representing the sum

¹<https://github.com/deepmind/open.spiel>

of each individual player random variable state, which is the case with large number of players thanks to the central limit theorem [7]. In the MFG, heterogeneity between the players is encoded in the state, to use the same policy for each player without a loss a generality. As an example, in our model, the destination of the player is represented in the state.

Metrics. The quality of the approximation of the Nash equilibrium policy completed by the candidate policy is measured using the average deviation incentive defined in (1) (also known as the average marginal regret [10], or the relative gap to the dynamic user equilibrium [14] in traffic engineering).

Implementation. The N -player game is encoded as a simultaneous, perfect information, general sum game. The corresponding MFG is encoded as a mean field, perfect information, general sum game. OpenSpiel provides many algorithms to find Nash equilibria of simultaneous games or MFGs. These algorithms include model-free algorithms such as Neural Fictitious Self-Play [23] and model-based algorithms such as Counterfactual Regret Minimization (and some variants) [49] which we use to solve the N -player game. The experiments solve the MFG using the online mirror descent algorithm [36]. The experiments performed in OpenSpiel use a fixed time discretisation.

Networks. As classical network games consider demand between nodes, we add artificial origin and destination links before and after each node in the network (Pigou [37], Braess [8] and Sioux Falls). This enables defining vehicle location only using links, and defining state of not having begun a trip and having finished it.

The *Pigou network* [37] has two links ℓ, ℓ' and two nodes (an origin and a destination one) which come from the conversion of the origin and the destination nodes. A time discretisation of 0.01, with a time horizon of 2 is used. The cost functions are $c_\ell(x) = 2$, $c_{\ell'}(x) = 1 + 2x$ and all the demand leaves the origin link at time 0 and head towards the destination link.

The *Braess network* game is the dynamic extension of the game described in [8]. The network has 5 links AB, AC, BC, BD and CD , one origin node A converted to an origin link OA and a destination node D converted to a destination link DE . The cost functions are $c_{AB}(x) = 1 + x$, $c_{AC}(x) = 2$, $c_{BC}(x) = 0.25$, $c_{BD}(x) = 2$, $c_{CD}(x) = 1 + x$. All the demand leaves the origin link at time 0 and head towards the destination link. We use a time step of 0.05 and a time horizon of 5.

In the *augmented Braess network* game, a destination link CF is added to the network and 50 more vehicles leave the origin to DE at time 0, 0.5 and 1, while 50 others leaves the origin to CF at times 0 and 1, totaling 250 vehicles with 2 different destinations and 3 different departure times.

The *Sioux Falls network* game is used by the traffic community for proof of concepts on network with around 100 links. The network (76 links without the origin and destination links), the link congestion functions, and an origin-destination traffic demand are open source [46]. As the classical routing game [34, Chapter 18] is a static game, the demand is only a list of tuple origin, destination and counts, and does not provide any departure time. We use the network data (including the congestion functions) and generate a demand specific to the game. We model 7,000 vehicles departing at time 0 from node 1 to node 19, and 7,000 vehicles departing at time 0 from node 19 to node 1. We use a time step of 0.5 and a time horizon of 50.

3.2 Mean field game solves the curse of dimensionality in the number of players

In this section, the mean field equilibrium policy is computed for both the Braess and the Pigou network games. In addition to being considerably faster to compute compared to the N -player Nash equilibrium, the mean field equilibrium provides an excellent approximation when N is above 30.

The evolution of the Braess mean field Nash equilibrium policy is given on fig. 1. The travel time on the three possible paths are equals, which encodes the Nash equilibrium condition of the MFG provided that the travel time on each link is a multiple of the time step, accordingly to theorem 3.

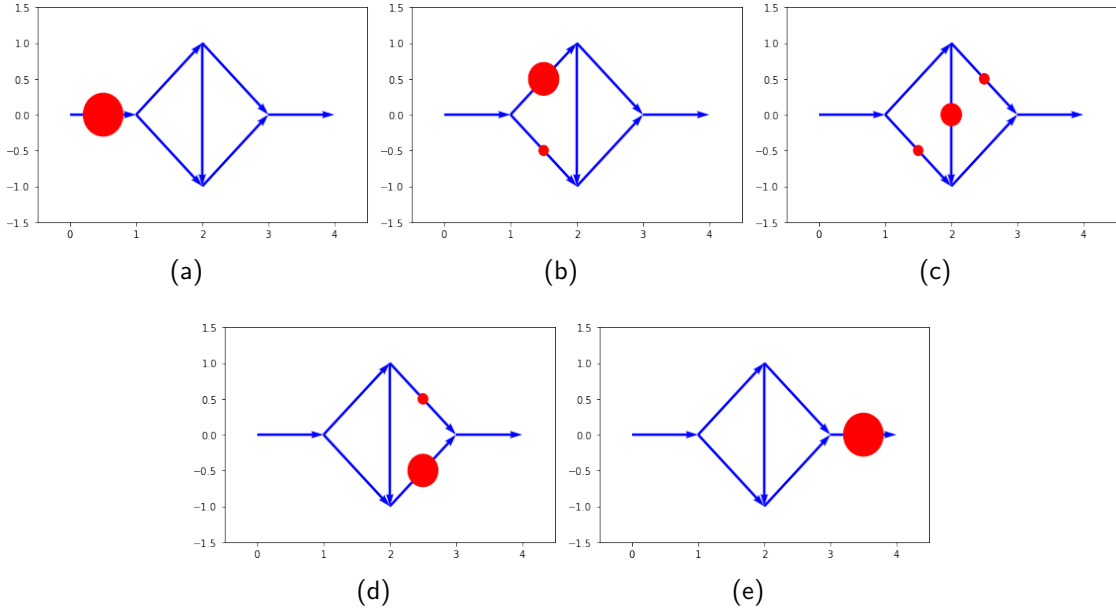


Figure 1: The dynamic of the Braess network in the mean field Nash equilibrium; the locations of the cars at time 0.0 (Figure 1a), from time 0.25 to 1.75 (Figure 1b), at time 2.0 (Figure 1c), from time 2.25 to 3.75 (Figure 1d), at time 4.0 (Figure 1e). The travel time on each path are equal to 3.75, travel time equalization defines the mean field Nash equilibrium.

3.2.1 While solving N -player game is intractable for large number of players, this can be done for the mean field game.

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) is used in the N -player game, as it is the fastest algorithms 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) 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. 2. As the mean field Nash equilibrium does not depends on the number of vehicles the MFG models, the computation time of 10 iterations of OMD is independent of the number of vehicles modeled. On the other hand, the computational cost of 10 iterations of ext CFR increases exponentially with the number of players, making the computation of a Nash equilibrium with a large number of players intractable with the algorithms of the OpenSpiel library.

3.2.2 The mean field equilibrium policy is a good approximation of the N -player equilibrium policy whenever N is large enough.

In the Pigou network game, 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. 3. This was shown theoretically in section 2.3.6, and is confirmed using approximate average deviation incentive of the mean field equilibrium policy in the N -player game.

In the Braess network game, the mean field equilibrium policy is almost a Nash equilibrium in the N -player game as soon as N is larger than 30 players. The results are reported in the appendix.

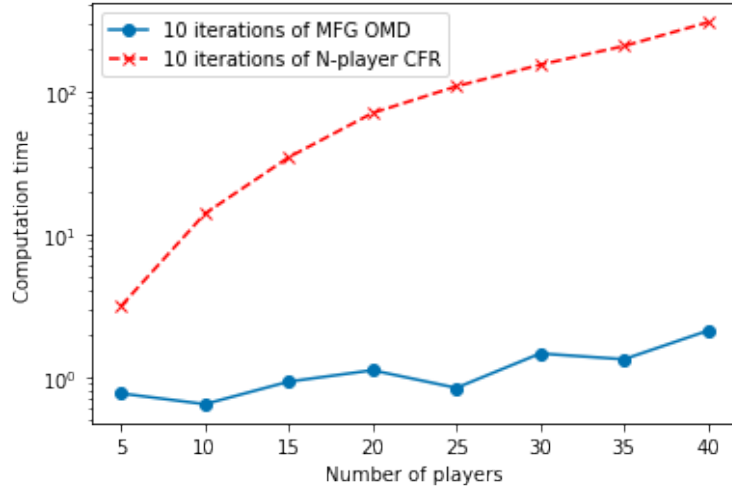


Figure 2: Computation time of 10 iterations of Online Mirror Descent in the MFG and of 10 iterations of sampled Counterfactual regret minimization as a function of the number of players N .

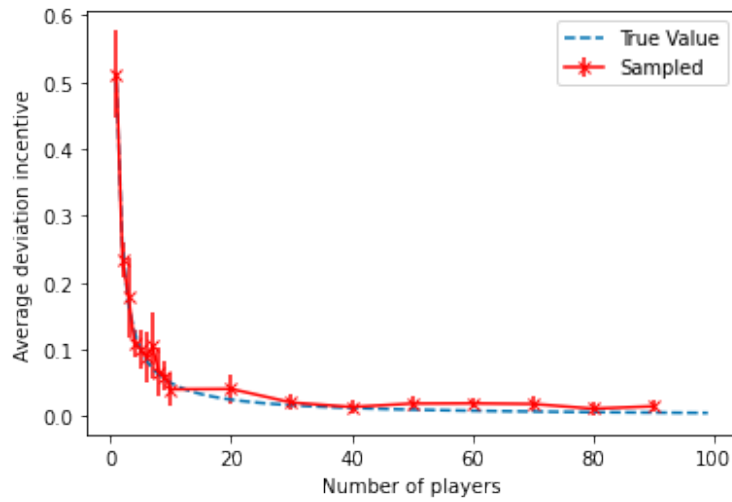


Figure 3: Average deviation incentive of the Nash equilibrium mean field policy in the N -player game as a function of N in the case of the Pigou game. The sampled value is the value computed in OpenSpiel by testing all the possible pure best responses, and sampling game trajectories to get the expected returns.

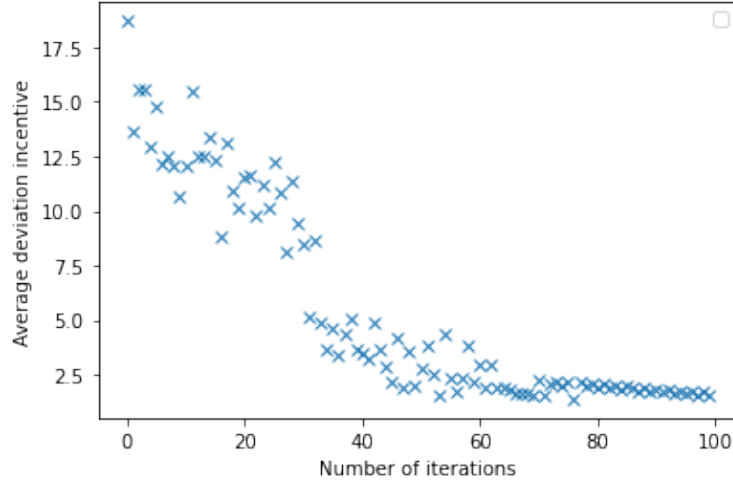


Figure 4: Online mirror descent average deviation incentive in the Sioux Falls MFG as a function of the number of iterations of the descent algorithm.

3.3 Mean field game approach can be extended to more complex set ups

The MFG approach solves the curse of dimensionality in the number of players (whenever the number of possible states is much below the number of players). It can be extended to more complex setups than the Pigou and Braess networks such as realistic traffic networks with demand. The results obtained when using the MFG approach in the augmented Braess network game with heterogeneous departure time is reported in the appendix. This section focuses on the extension of MFG approach to one of the classical benchmark network game used by the traffic community: the Sioux Falls network.

The experiment shows the ability to learn the mean field equilibrium policy on this 76 links network, with 14,000 vehicles going to two different destinations. Using online mirror descent, we see that the average deviation incentive decreases to 1.55 (for a travel time of 27) over 100 iterations, see fig. 4. We use a fixed learning rate of 1 in the 30 first iterations of the algorithm, 0.1 in the 31 to the 60 first iterations and a fixed learning rate of 0.01 in the 40 remaining iterations to produce fig. 4.

The resulting mean field policy is not exactly the Nash equilibrium policy of the MFG as its average deviation incentive is 1.55 (for a travel time of 27.5). The game evolution displayed in fig. 5 shows that some vehicles going from node 19 to node 1 have a longer travel time than others: on time step 26.5 (fig. 5g) some vehicles have arrived to node 1 (top left) and some have not.

Average deviation of the learned mean field policy cannot be computed numerically in the 14,000 player game, due to the large number of players.

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 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.

This opens several directions for future work. A natural next step would be to consider even more realistic scenarios, which means tackling more complex routing models and more complex networks. For example, it would be interesting to allow congestion spill back between links, and to propose a model ensuring the

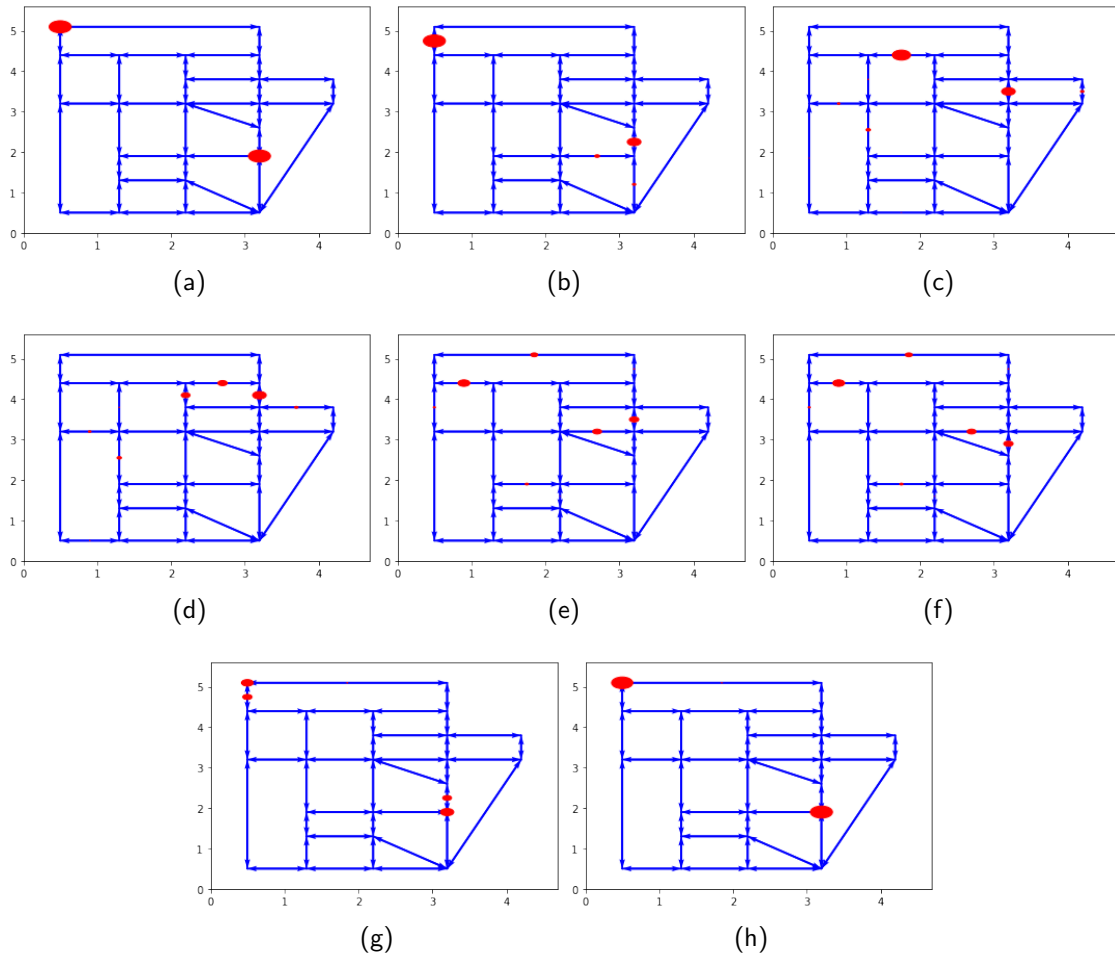


Figure 5: Dynamics of the Sioux Falls network in the mean field Nash equilibrium. Road network; location of the cars at time 0.0 (5a), 2.5 (5b), 10.0 (5c), 12.5 (5d), 21.0 (5e), 22.0 (5f), 26.5 (5g), 27.5 (5h). Some vehicles arrived at their destination after some that left the origin at the same time: the Nash equilibrium has not been reached. On average, players can expect saving 1.55 time by being the only one to be rerouted on a better path.

first-in first-out property according to which a vehicle can never exit a link earlier than another vehicle that arrived before them. This could be done by including the game state, i.e., the whole state distribution, in the transitions and the policies. This will raise new challenges and new scalability issues in terms of models and networks. To cope with this aspect, it would be interesting to combine the mean-field approach proposed in this work with state-of-the art reinforcement learning techniques.

References

- [1] ACHDOU, Y., DAO, M.-K., LEY, O., AND TCHOU, N. Finite horizon mean field games on networks. *Calculus of Variations and Partial Differential Equations* 59, 5 (2020), 1–34.
- [2] ACHDOU, Y., AND LASRY, J.-M. Mean field games for modeling crowd motion. In *Contributions to partial differential equations and applications*. Springer, 2019.
- [3] ALASSEUR, C., TAHER, I. B., AND MATOUSSI, A. An extended mean field game for storage in smart grids. *Journal of Optimization Theory and Applications* 184, 2 (2020).
- [4] AURELL, A., CARMONA, R., DAYANIKLI, G., AND LAURIERE, M. Optimal incentives to mitigate epidemics: a stackelberg mean field game approach. *arXiv preprint arXiv:2011.03105* (2020).
- [5] AURELL, A., AND DJEHICHE, B. Mean-field type modeling of nonlocal crowd aversion in pedestrian crowd dynamics. *SIAM J. Control Optim.* 56, 1 (2018), 434–455.
- [6] BAUSO, D., ZHANG, X., AND PAPACHRISTODOULOU, A. Density flow in dynamical networks via mean-field games. *IEEE Transactions on Automatic Control* 62, 3 (2016), 1342–1355.
- [7] BERTSEKAS, D. P., AND TSITSIKLIS, J. N. *Introduction to probability*, vol. 1. Athena Scientific Belmont, MA, 2002.
- [8] BRAESS, D., NAGURNEY, A., AND WAKOLBINGER, T. On a paradox of traffic planning. *Transportation science* 39, 4 (2005), 446–450.
- [9] BROWN, G. W. Iterative solution of games by fictitious play. *Activity analysis of production and allocation* 13, 1 (1951), 374–376.
- [10] CABANNES, T., SANGIOVANNI, M., KEIMER, A., AND BAYEN, A. M. Regrets in routing networks: Measuring the impact of routing apps in traffic. *ACM Trans. Spatial Algorithms Syst.* 5, 2 (July 2019), 9:1–9:19.
- [11] CABANNES, T., VINCENTELLI, M. A. S., SUNDT, A., SIGNARGOUT, H., PORTER, E., FIGHIERA, V., UGIRUMURERA, J., AND BAYEN, A. M. The impact of GPS-enabled shortest path routing on mobility: a game theoretic approach. *Transportation Research Board 97th Annual Meeting* (2018).
- [12] CARDALIAGUET, P., AND LEHALLE, C.-A. Mean field game of controls and an application to trade crowding. *Mathematics and Financial Economics* 12, 3 (2018).
- [13] CHEVALIER, G., LE NY, J., AND MALHAMÉ, R. A micro-macro traffic model based on mean-field games. In *2015 American Control Conference (ACC)* (2015), IEEE, pp. 1983–1988.
- [14] CHIU, Y.-C., BOTTOM, J., MAHUT, M., PAZ, A., BALAKRISHNA, R., WALLER, S., AND HICKS, J. Dynamic traffic assignment: A primer (transportation research circular e-c153).
- [15] DASKALAKIS, C., GOLDBERG, P. W., AND PAPADIMITRIOU, C. H. The complexity of computing a nash equilibrium. *SIAM Journal on Computing* 39, 1 (2009), 195–259.
- [16] DONCEL, J., GAST, N., AND GAUJAL, B. A mean field game analysis of sir dynamics with vaccination. *Probability in the Engineering and Informational Sciences* (2020), 1–18.
- [17] ELIE, R., HUBERT, E., AND TURINICI, G. Contact rate epidemic control of COVID-19: an equilibrium view. *Mathematical Modelling of Natural Phenomena* (2020).
- [18] ELIE, R., ICHIBA, T., AND LAURIÈRE, M. Large banking systems with default and recovery: A mean field game model. *arXiv preprint arXiv:2001.10206* (2020).

- [19] FRIESZ, T., BERNSTEIN, D., SMITH, T., TOBIN, R., AND WIE, B.-W. A variational inequality formulation of the dynamic network user equilibrium problem. *Operations Research* 41, 1 (1993), 179–191.
- [20] GLICKSBERG, I. L. A further generalization of the kakutani fixed point theorem, with application to nash equilibrium points. *Proceedings of the American Mathematical Society* 3, 1 (1952), 170–174.
- [21] GUÉANT, O. Existence and uniqueness result for mean field games with congestion effect on graphs. *Applied Mathematics & Optimization* 72, 2 (2015), 291–303.
- [22] HAN, K., FRIESZ, T. L., AND YAO, T. Existence of simultaneous route and departure choice dynamic user equilibrium. *Transportation Research Part B: Methodological* 53 (2013), 17–30.
- [23] HEINRICH, J., AND SILVER, D. Deep reinforcement learning from self-play in imperfect-information games. *arXiv preprint arXiv:1603.01121* (2016).
- [24] HUANG, K., DI, X., DU, Q., AND CHEN, X. A game-theoretic framework for autonomous vehicles velocity control: Bridging microscopic differential games and macroscopic mean field games. *Discrete & Continuous Dynamical Systems - B* 22, 11 (2017).
- [25] HUANG, M., MALHAMÉ, R. P., AND CAINES, P. E. Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Communications in Information and Systems* 6, 3 (2006).
- [26] LACHAPELLE, A., LASRY, J.-M., LEHALLE, C.-A., AND LIONS, P.-L. Efficiency of the price formation process in presence of high frequency participants: a mean field game analysis. *Mathematics and Financial Economics* 10, 3 (2016).
- [27] LACHAPELLE, A., AND WOLFRAM, M.-T. On a mean field game approach modeling congestion and aversion in pedestrian crowds. *Transportation research part B: methodological* 45, 10 (2011), 1572–1589.
- [28] LANCTOT, M., LOCKHART, E., LESPIAU, J.-B., ZAMBALDI, V., UPADHYAY, S., PÉROLAT, J., SRINIVASAN, S., TIMBERS, F., TUYLS, K., OMIDSHAFIEI, S., HENNES, D., MORRILL, D., MULLER, P., EWALDS, T., FAULKNER, R., KRAMÁR, J., VYLDER, B. D., SAETA, B., BRADBURY, J., DING, D., BORGEAUD, S., LAI, M., SCHRITTWIESER, J., ANTHONY, T., HUGHES, E., DANIHELKA, I., AND RYAN-DAVIS, J. OpenSpiel: A framework for reinforcement learning in games. *CoRR abs/1908.09453* (2019).
- [29] LASRY, J.-M., AND LIONS, P.-L. Mean field games. *Japanese Journal of Mathematics* 2, 1 (2007).
- [30] LI, F., MALHAMÉ, R. P., AND LE NY, J. Mean field game based control of dispersed energy storage devices with constrained inputs. In *2016 IEEE 55th Conference on Decision and Control (CDC)* (2016), IEEE.
- [31] LIGHTHILL, M. J., AND WHITHAM, G. B. On kinematic waves ii. a theory of traffic flow on long crowded roads. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 229, 1178 (1955), 317–345.
- [32] MONDERER, D., AND SHAPLEY, L. S. Potential games. *Games and economic behavior* 14, 1 (1996), 124–143.
- [33] MYERSON, R. B. *Game theory*. Harvard university press, 2013.
- [34] NISAM, N., ROUGHGARDEN, T., TARDOS, E., AND VAZIRANI, V. V. *Algorithmic game theory*. Cambridge University Press, 2007.
- [35] PATRIKSSON, M. *The traffic assignment problem: models and methods*. Courier Dover Publications, 2015.

- [36] PEROLAT, J., PERRIN, S., ELIE, R., LAURIÈRE, M., PILIOURAS, G., GEIST, M., TUYLS, K., AND PIETQUIN, O. Scaling up mean field games with online mirror descent. *arXiv preprint arXiv:2103.00623* (2021).
- [37] PIGOU, A. C. *The economics of welfare*. Palgrave Macmillan, 1920.
- [38] ROSENTHAL, R. W. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 1 (Dec 1973), 65–67.
- [39] SALDI, N., BASAR, T., AND RAGINSKY, M. Markov–nash equilibria in mean-field games with discounted cost. *SIAM Journal on Control and Optimization* 56, 6 (2018), 4256–4287.
- [40] SALHAB, R., LE NY, J., AND MALHAMÉ, R. P. A mean field route choice game model. In *2018 IEEE Conference on Decision and Control (CDC)* (2018), IEEE, pp. 1005–1010.
- [41] SAMARAKOON, S., BENNIS, M., SAAD, W., DEBBAH, M., AND LATVA-AHO, M. Energy-efficient resource management in ultra dense small cell networks: A mean-field approach. In *IEEE Global Communications Conference (GLOBECOM)* (2015).
- [42] SCHRANK, D., EISELE, B., AND LOMAX, T. 2019 Urban Mobility Report. *Texas A&M Transportation Institute*, August (2019).
- [43] SHOU, Z., AND DI, X. Multi-agent reinforcement learning for dynamic routing games: A unified paradigm. *CoRR abs/2011.10915* (2020).
- [44] SIEBEL, F., AND MAUSER, W. On the fundamental diagram of traffic flow. *SIAM Journal on Applied Mathematics* 66, 4 (2006), 1150–1162.
- [45] TANAKA, T., NEKOU EI, E., PEDRAM, A. R., AND JOHANSSON, K. H. Linearly solvable mean-field traffic routing games. *IEEE Transactions on Automatic Control* 66, 2 (2020), 880–887.
- [46] TRANSPORTATION NETWORKS FOR RESEARCH CORE TEAM. Transportation networks for research, free of use. <https://github.com/bstabler/TransportationNetworks>. Accessed: 2021-09-01.
- [47] WARDROP, J. G., AND WHITEHEAD, J. I. Correspondence. Some theoretical aspects of road traffic research. *ICE Proceedings: Engineering Divisions* 1 (1952).
- [48] WIE, B.-W. A differential game model of nash equilibrium on a congested traffic network. *Networks* 23, 6 (1993), 557–565.
- [49] ZINKEVICH, M., JOHANSON, M., BOWLING, M., AND PICCIONE, C. Regret minimization in games with incomplete information. *Advances in neural information processing systems* 20 (2007), 1729–1736.

A Complement on some theoretical aspects

A.1 Proofs

Proof of Theorem 1. The proof of the existence of a Nash equilibrium relies on the Kakutani-Glicksberg-Fan theorem [20] which needs the continuity of the cost function with respect to the policy profile.

Because \mathcal{N} and \mathcal{A} are finite and original waiting times are given, the time between two actions in the N -player game is necessary of the form

$$\sum_{\ell, k} \alpha_{\ell, k} c_{\ell} \left(\frac{k}{N} \right) + \sum_{i \in \mathcal{N}} \beta_i W_0^i,$$

where $(\alpha_{\ell})_{\ell \in \mathcal{A}} \in \{-1, 0, 1\}^{\mathcal{A}}$, $k \in \{0, \dots, N\}$ and $\beta_i \in \{-1, 0, 1\}^{\mathcal{N}}$. This generates at most $3^{|\mathcal{A}|(N+1)+N}$ possibilities. Therefore, there is a minimum time between two actions. As the time horizon T is fixed, this implies that there is a maximum number of times M where an action can be taken. Therefore, a policy only needs to define an action on all the possible tuples of times where an action should be taken. The number K of possible tuples is smaller than $\binom{3^{|\mathcal{A}|(N+1)+N}}{M}$. Without loss of generality, we restrict ourselves to the set of policies which is a subset of $\mathcal{P}(\mathcal{A})^{\mathcal{A} \times K}$, making the set of pure policies a subset of $\mathcal{A}^{\mathcal{A} \times K}$ finite, and the set of mixed-policies (its convex hull) a compact subset of the Euclidean vector space $\mathbb{R}^{\mathcal{A} \times \mathcal{A} \times K}$. We will later show that any Nash equilibrium in the set of restricted mixed-policies is a Nash equilibrium in the set of non restricted mixed-policies.

We denote by $\mathcal{K} \subseteq \mathbb{N}$ the set of pure policy indices. We denote $Q = \{q_k, k \in \mathcal{K}\}$ the set of pure policies. The set of mixed-policies is a subset of the simplex over the set of pure policy $\mathcal{P}(Q)$. Therefore, for all $\pi \in \mathcal{P}(Q)$, there exists $\alpha \in \mathcal{P}(\mathcal{K})$ identified with a vector of length $|\mathcal{K}|$ such that $\pi = \alpha \cdot (q_k)_{k \in \mathcal{K}}$. We denote by $\tilde{\Pi} \subseteq \mathcal{P}(Q)$ the set of restricted mixed-policies and by $\tilde{\Pi}^N = \tilde{\Pi}^{\mathcal{N}}$ the set of restricted mixed-policy profiles.

Lemma 1 (Cost is convex combination of pure strategy cost). *For any $\underline{\pi} \in S$, and for any $i \in \mathcal{N}$ with corresponding $\alpha \in \mathcal{P}(\mathcal{K})$ such that $\underline{\pi}^i = \alpha \cdot (q_k)_{k \in \mathcal{K}}$*

$$J_i^N(\underline{\pi}^i, \underline{\pi}^{-i}) = \sum_{k \in \mathcal{K}} \alpha_k J_i^N(q_k, \underline{\pi}^{-i}).$$

Proof. The equality is a direct consequence of the linearity of the expected value with respect to the probability distribution. \square

Next, let the (set-valued) Best response map $\phi : \tilde{\Pi} \rightarrow 2\tilde{\Pi}$ be defined by for all: for $\underline{\pi} \in \tilde{\Pi}$,

$$\phi(\underline{\pi}) = \bigtimes_{i \in \mathcal{N}} \operatorname{argmin}_{\pi' \in \tilde{\Pi}} J_i^N(\pi', \underline{\pi}^{-i}).$$

Lemma 2 (Best response map has a closed graph). *ϕ has a closed graph.*

Proof. For any sequence $(\underline{\pi}_j)_{j \in \mathbb{N}} \in \tilde{\Pi}^{\mathbb{N}}$ converging to $\underline{\pi}_{\infty}$ and any sequence $(x_j)_{j \in \mathbb{N}} \in \tilde{\Pi}^{\mathbb{N}}$ converging to x_{∞} such that $x_j \in \phi(\underline{\pi}_j)$ for all $j \in \mathbb{N}$, we have that, for any $i \in \mathcal{N}$:

$$J_i^N(x_j^i, \underline{\pi}_j^{-i}) \leq \alpha \cdot (J_i^N(q_k, \underline{\pi}_j^{-i}))_{k \in \mathcal{K}} \quad \forall \alpha \in \mathcal{P}(\mathcal{K}),$$

which holds in the limit provided that J is continuous w.r.t. the policy profile (our regularity assumption). Therefore $x_{\infty} \in \phi(\underline{\pi}_{\infty})$ and ϕ has a closed graph. \square

Proposition 1 (Best response map has a fixed-point). *The map ϕ is a Kakutani-Glicksberg-Fan map, therefore it admits a fixed-point [20], which is a Nash equilibrium of the N -player game.*

Proof.

The set $\tilde{\Pi}$ is:

- non-empty (we assume the graph is non empty, so at least one path exists)
- compact (as a probability distribution over finite set)
- convex (as a probability distribution over finite set)
- subset of a Hausdorff locally convex topological vector space (as a probability distribution over finite set)

The function $\phi(\underline{\pi})$ is non-empty and convex for all $\underline{\pi} \in \underline{\Pi}$ (because the minimization is a linear problem). The function ϕ has a closed graph (lemma 2). Hence, by Kakutani-Glicksberg-Fan theorem [20], ϕ has a fixed point. □

The proof of Theorem 1 is concluded by noting that a fixed point $\underline{\pi}^* \in \tilde{\underline{\Pi}}$ of ϕ is a Nash equilibrium in $\underline{\Pi}$. Indeed, by definition of ϕ :

$$\forall i \in \mathcal{N}, \forall \underline{\pi} \in \tilde{\underline{\Pi}}, J_i^N(\underline{\pi}^{i*}, \underline{\pi}^{-i*}) \leq J_i^N(\underline{\pi}, \underline{\pi}^{-i*}).$$

From here, by lemma 1 and by construction of $\tilde{\underline{\Pi}}$:

$$\forall i \in \mathcal{N}, \forall \underline{\pi} \in \underline{\Pi}, J_i^N(\underline{\pi}^{i*}, \underline{\pi}^{-i*}) \leq J_i^N(\underline{\pi}, \underline{\pi}^{-i*}).$$

The last line states that the Nash equilibrium in the set $\tilde{\underline{\Pi}}$ is a Nash equilibrium in the set $\underline{\Pi}$ and finishes the proof of existence of the Nash equilibrium. □

Proof of Theorem 2. This proof is similar to the proof of Theorem 1 as long as the cost function is continuous with respect to the policy profile, which is assumed here. In the mean field game, given a departure time the set of pure policies can be restricted to the set of path $Q^{\text{supp}(W_0)}$ given the departure time W_0 , with K the number of possible path (with less than a given number of cycles, which is possible due to minimum travel time on each link and finite time horizon) times the number of possible departure time. Therefore a pure policy, should be understood as choosing a path $q \in Q$ given a departure time W_0 . Encoding the policy in $\underline{\Pi}$ that correspond to any choice of a path given a departure time might require some notation work, that we will not do here. The reader should understand that the proof relies on finding a Nash equilibrium where the set of policies is the set of probability distributions on the set of paths given the departure time, and then showing that this Nash equilibrium policy can be translated into a policy in $\underline{\Pi}$, and then showing that in the set of $\underline{\Pi}$ the policy is still a Nash equilibrium. To establish the existence of a Nash equilibrium in the set of probability distribution on the set of paths given a departure time $S = \mathcal{P}(Q^{\text{supp}(W_0)})$, the same arguments that the ones in the proof of Theorem 1 can be used with the map $\phi : S \rightarrow 2^S$ defined by:

$$\phi(\underline{\pi}) = \bigtimes_{i \in \mathcal{N}} \underset{\underline{\pi}' \in S}{\text{argmin}} J(\underline{\pi}', \underline{\pi})$$

which is a Kakutani-Fan-Glicksberg map. Then, one can convert the path choice given a departure time to a list of actions in time, therefore convert it to a policy $\underline{\pi}^*$ in $\underline{\Pi}$. Then, one can show by contradiction that they cannot exist a policy $\underline{\pi}$ in $\underline{\Pi}$ that gives a strictly better outcome than the policy $\underline{\pi}^*$ in the cost function $J(\underline{\pi}, \underline{\pi}^*)$. Therefore $\underline{\pi}^*$ is also a Nash equilibrium in the set $\underline{\Pi}$. □

A.2 Counter-example of the existence without continuity

Here we present a counter-example of the existence of a Nash equilibrium of the mean field game when congestion functions are not continuous.

Consider a network with two links, say ℓ, ℓ' , connecting the origin and destination links. Let the congestion function be fined as:

$$c_\ell(\mu) = \begin{cases} 1 & \text{if } \mu < 0.5 \\ 2 & \text{otherwise} \end{cases}$$

$$c_{\ell'}(\mu) = \begin{cases} 1 & \text{if } \mu \leq 0.5 \\ 2 & \text{otherwise} \end{cases}$$

Then the mean field game admits no Nash equilibrium. Indeed, if $\nu_t(\ell) < 0.5$, then $c_\ell = 1$ and $c_{\ell'} = 2$, and if $\nu_t(\ell) \geq 0.5$, then $c_\ell = 2$ and $c_{\ell'} = 1$. In this case there is no flow allocation such that travel times are equal on the path that are used, therefore there is no Nash equilibrium. This is due to the discontinuity of the cost of pure actions with respect to the state distribution.

Proof of Theorem 3. Lemma 1 is still valid in the mean field game setup, with the set of pure policies corresponding to a choice of paths. This implies that if the equilibrium policy $\underline{\pi}^*$ is a mix-policy; i.e. there exist $\mathcal{K}^* \subset \mathcal{K}$ and $\alpha \in \mathcal{P}(\mathcal{K})$ such that:

$$\begin{aligned} \underline{\pi}^* &= \alpha \cdot (q_k)_{k \in \mathcal{K}} \\ \alpha_k &> 0 \quad \forall k \in \mathcal{K}^* \\ \alpha_k &= 0 \quad \forall k \in \mathcal{K} \setminus \mathcal{K}^*. \end{aligned}$$

Then $J(\underline{\pi}^*, \underline{\pi}^*) = J(q_k, \underline{\pi}^*)$ for all $k \in \mathcal{K}^*$. In the mean field game, this translates in the equality of the travel time on all path q_k with $k \in \mathcal{K}^*$. \square

Remark that in the N player game, $J_i^N(q_k, \underline{\pi}_{-i}^*) = J_i^N(q_{\ell'}, \underline{\pi}_{-i}^*)$ does not translate into the equality of the trajectories following the pure policies q_k and $q_{\ell'}$ as the travel time depends on the vehicle i location. This is not the case in the mean field game.

Proof. PROOF OF AVERAGE DEVIATION INCENTIVE OF THE PIGOU MEAN FIELD EQUILIBRIUM POLICY IN THE CORRESPONDING N PLAYER GAME.

For the point of view of one player, if the other players are following the mean field equilibrium policy, then the probability that m other players are on the link ℓ' is $\binom{N-1}{m} 0.5^{N-1}$. In this case, the travel time on link ℓ' is $2\frac{mT}{N}$ if the player does not use it, or $2\frac{(m+1)T}{N}$ if the player uses it. If the player follows the mean field equilibrium policy, with 50% it will use the link ℓ and have a deviation incentive of $0.5T \max\{0, 1 - 2\frac{m+1}{N}\}$, and with 50% it will use the link ℓ' and have a deviation incentive of $0.5T \max\{0, 2\frac{m+1}{N} - 1\}$. Therefore, the player deviation incentive is:

$$\begin{aligned} D &= 0.5T \sum_{m=1}^{N-1} \binom{N-1}{m} 0.5^{N-1} \left(0.5 \max\left\{0, 2\frac{m+1}{N} - 1\right\} \right. \\ &\quad \left. + 0.5 \max\left\{0, 1 - 2\frac{m+1}{N}\right\} \right) \\ &= \frac{T}{N2^N} \sum_{m=1}^N \binom{N-1}{m} \max\left\{\frac{N}{2} - m - 1, m + 1 - \frac{N}{2}\right\} \end{aligned}$$

\square

B Average deviation incentive of the Braess mean field equilibrium policy in the corresponding N -player game

In the Braess network, the mean field policy is almost a Nash equilibrium in the N -player game as soon as N is larger than 30 players fig. 6. The computation of the average deviation incentive is done with an

explicit numeric computation, using a adaptive time discretisation. The cost of a player in the mean field game is 3.75, as shown on fig. 1, therefore a 0.05 average deviation incentive is a excellent approximation of the N -player Nash equilibrium.

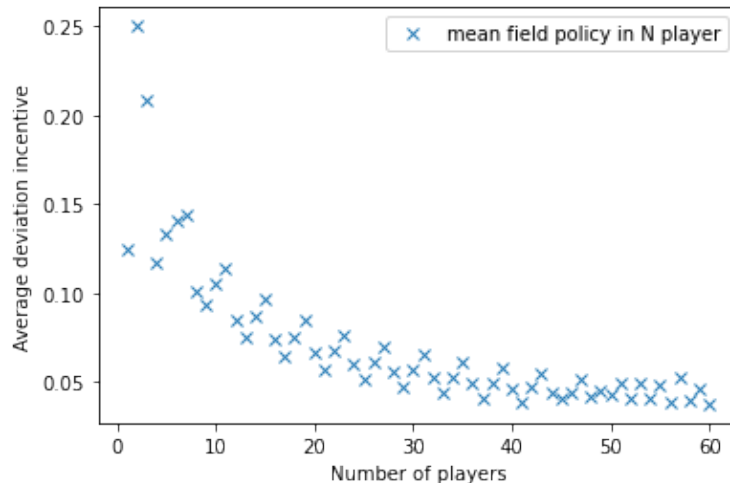


Figure 6: Average deviation incentive of the Nash equilibrium mean field policy in the N -player game as a function of N in the case of the Braess game. The average deviation incentive is computed numerically explicitly using a adaptive time discretisation without OpenSpiel.

C Mean field game with player heterogeneity in their origin, destination and departure time

The game evolution in the mean field equilibrium policy in the augmented Braess network game is shown on fig. 7.

D Reproducing the experiments

Experiment can be reproduced using the files in the github https://github.com/deepmind/open_spiel/tree/master/open_spiel/data/paper_data/routing_game_experiments following the instructions in the readme.md file.

E Talk

This article has been presented in the workshop on Mean Field Games on Networks: <https://www.math.tu.be.org/lecture/video/solving-dynamic-user-equilibrium-mean-field-routing-game-explicit-congestion-dynamics>.

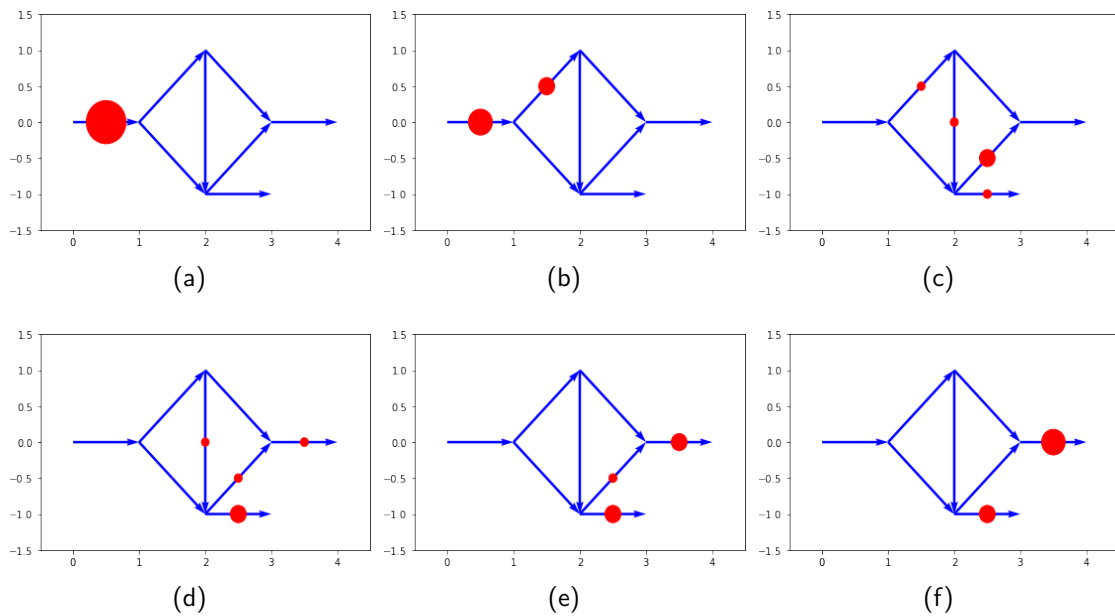


Figure 7: Dynamics of the Braess network augmented with two destination in a set up with several departure time. Location of the cars at time 0.0 (7a), 0.25 (7b), 2.25 (7c), 2.5 (7d), 3.5 (7e) and 3.75 (7f). Vehicles that departs at the same time from the same origin and the same destination reach their destination at the same time, encoding the Nash equilibrium of the MFG.