The Four Levels of Fixed-Points in Mean-Field Models

Sarath Yasodharan* and Rajesh Sundaresan*†
*ECE Department and †Robert Bosch Centre for Cyber Physical Systems,
Indian Institute of Science, Bangalore 560012, India

***Abstract*—The fixed-point analysis refers to the study of fixed-points that arise in the context of complex systems with many interacting entities. In this expository paper, we describe four levels of fixed-points in mean-field interacting particle systems. These four levels are (i) the macroscopic observables of the system, (ii) the probability distribution over states of a particle at equilibrium, (iii) the time evolution of the probability distribution over states of a particle, and (iv) the probability distribution over trajectories. We then discuss relationships among the fixed-points at these four levels. Finally, we describe some issues that arise in the fixed-point analysis when the system possesses multiple fixed-points at the level of distribution over states, and how one goes beyond the fixed-point analysis to tackle such issues.**

***Index Terms*—Fixed-point, mean-field limit, interacting particle system, performance analysis, McKean-Vlasov equation, propagation of chaos**

I. Introduction

The purpose of this paper is to highlight four levels of fixed-points in mean-field interacting particle systems, and to discuss the relationships among them. No new result is established. However, the unified perspective and the insights from the connections between the fixed-points are worthy of exposition.

Interacting particle systems specify the microscopic interactions of a particle with the other particles in the system, i.e., nature of interactions, their frequencies, and how they influence the state of the particle. This microscopic description then leads to a collective macroscopic outcome, for example, the fraction of particles in each state at a given time or at equilibrium. In interactions of the *mean-field* type, the evolution of the state of a particle depends on the states of the other particles only through the empirical measure, i.e., the distribution of the particles across states. Such mean-field models arise in the study of wireless local area networks (WLANs) [1], randomised algorithms in load balancing networks [2], [3], processor sharing systems [4], study of epidemic spread in a community [5], etc.

In the study of such systems, one often looks for *fixed-points*. Let γ be some property associated with a tagged particle. For concreteness, in a WLAN, γ could be the collision probability when a particular tagged node (particle) accesses the medium of communication. Intuitively, if the system is in equilibrium, then all the nodes may see collisions with the same probability γ. This creates a mean-field. The study of the collision probability of the tagged node, now viewed as a *response* to this mean-field, may yield an expression $G(\gamma)$ for the tagged node's collision probability. Consistency then demands that the node's appropriate collision probability γ^* is one where the response to the mean-field $G(\gamma^*)$ equals γ^*, i.e., γ^* is a fixed-point of G and is the property that all the nodes are likely to experience at equilibrium, see [6]. This provides a heuristic method to study the behaviour of such complex interacting particle systems.

To get this analysis machinery going, one needs a tractable property γ and an analysis that generates the response $G(\gamma)$ to the mean-field arising from γ. In the WLAN analysis, this property γ was the collision probability [6]. This intuition has helped us analyse many natural and engineered systems. It works some times, fails some times, and we now have a good idea when it works and when it fails.

In this paper, we consider four properties of increasing levels of detail in a general mean-field interacting particle system, and describe the associated fixed-points. These four levels are (i) the macroscopic observables of the system, (ii) the probability distribution over states of a particle at equilibrium, (iii) the time evolution of the probability distribution over states of a particle, and (iv) the probability distribution over trajectories. We then discuss conditions under which the fixed-points exist and explore the inter-connections.

Organisation: Section II describes our finite-state mean-field interacting particle system. Section III describes the four levels of fixed-points. Section IV discusses the relationships among these fixed-points. Section V discusses issues related to multiple fixed-points.

II. A mean-field model

In this section we describe a finite-state mean-field interacting particle system and discuss some of its properties.

A. The model

Let there be N particles. Each particle has a state. The particle states evolves over time in a Markovian fashion. Let $X_n^{(N)}(t)$ denote the state of the nth particle at time t and let $\mathcal{Z}$, a finite set, denote the state space. The macroscopic behaviour of the system at time t is encoded in the empirical measure

$$\mu_N(t) = \frac{1}{N}\sum_{n=1}^{N} \delta_{X_n^{(N)}(t)};$$

$\mu_N(t)(z)$ is the fraction of particles in state z at time t. To describe the evolution of the states of the particles, we consider a directed and connected graph $(\mathcal{Z}, \mathcal{E})$. This graph encodes

978-1-6654-0522-5/21/$31.00 ©2021 IEEE

2021 National Conference on Communications (NCC) | 978-1-6654-4177-3/21/$31.00 ©2021 IEEE | DOI: 10.1109/NCC52529.2021.9530179

Authorized licensed use limited to: J R D. Tata Memorial Library Indian Institute of Science Bengaluru. Downloaded on April 18,2025 at 06:58:07 UTC from IEEE Xplore. Restrictions apply.

978-1-6654-4177-3/21/$31.00 ©2021 IEEE

the set of all allowed transitions of a particle; whenever $(z, z') \in \mathcal{E}$, a particle in state z can move to state z'. For each $(z, z') \in \mathcal{E}$, we are given a function $\lambda_{z,z'}$ on the space of probability distributions on $\mathcal{Z}$ (which we denote by $M_1(\mathcal{Z})$ and view as a subset of $\mathbb{R}^{|\mathcal{Z}|-1}$). A particle at state z at time t moves to state z' at rate $\lambda_{z,z'}(\mu_N(t))$. That is, the evolution of the state of a particle depends on the states of the other particles only through the empirical measure of the states of all the particles. This description of the state evolution of the particles gives rise to two continuous-time Markov processes: (i) the Markov process $\{(X_n^{(N)}(t), 1 \leq n \leq N), t \geq 0\}$ that describes the joint evolution of the states of all the particles, and (ii) the more convenient Markov process $\{\mu_N(t), t \geq 0\}$ that describes the evolution of the empirical measure of the system of particles, i.e., the empirical measure process.

B. Two examples

1) Markovian model of a WLAN: N wireless nodes (or particles) access a common wireless channel. Since multiple nodes can try to access the channel at a given time slot, an attempt to transmit a packet by a node can either result in a collision or in a successful transmission. A basic distributed Medium Access Control (MAC) protocol implemented in the nodes can be summarised as follows [7, Section 7.3]. Each node has a state belonging to the state space $\mathcal{Z} = \{0, 1, \ldots, K\}$. The states represent aggressiveness of packet transmission. In each time slot, a node in state i *independently* attempts a packet transmission with probability c_i/N. The $O(1/N)$ scaling of the attempt probabilities results in $O(1)$ attempt probability for the entire system. In the basic implementation, $c_i = c_{i-1}/2$; thus state 0 is the most aggressive state and state K is the least aggressive state.

Let us write down the collision probability of a tagged node in state z_0 when the empirical measure of the system is ξ, i.e., the number of nodes in state z is $N\xi(z)$, $z \in \mathcal{Z}$. Since nodes attempt independently, the probability that no other node (other than the tagged node) attempts in a given slot is

$$\left(1 - \frac{c_{z_0}}{N}\right)^{N\xi(z_0)-1} \prod_{z \in \mathcal{Z}, z \neq z_0} \left(1 - \frac{c_z}{N}\right)^{N\xi(z)} \simeq \exp\{-\langle c, \xi \rangle\}$$

for large N, where $\langle c, \xi \rangle = \sum_z (c_z/N)(N\xi(z)) = \sum_z c_z \xi(z)$ can be interpreted as the system attempt probability. A collision occurs when at least one other node makes a transmission; hence the collision probability is $\gamma = 1 - \exp\{-\langle c, \xi \rangle\}$. Observe that the collision probability of a node depends on only ξ rather than the individual states of the particles. The collision probability is an example of a macroscopic observable because one can observe the received power in the wireless channel and assess how many slots suffer collisions.

We now describe the evolution of a particle's state. Consider a tagged node in state i. If the node encounters a collision, it moves to a less aggressive state, i.e., to state $i + 1$, and continues to attempt transmission of its head-of-the-queue packet with probability c_{i+1}/N. On the other hand, if the node makes a successful packet transmission, it moves to state 0 and follows the same protocol for the next packet in its queue. When the node is in state K, it continues to attempt with probability c_K/N until a successful transmission when it moves from state K to 0; see Figure 1, but with an additional self loop at state K. Based on the approximation to the collision probability in the previous paragraph, the probability that the node moves from state i to state $i + 1$ is $(c_i/N)(1 - \exp\{-\langle c, \xi \rangle\})$ and the probability that the node moves from state i to state 0 is $(c_i/N)\exp\{-\langle c, \xi \rangle\}$. The continuous-time caricature for the above WLAN is obtained by scaling the time-slot durations to be $1/N$ in the above discrete-time description. The rate at which a node in state i moves to state $i + 1$ is

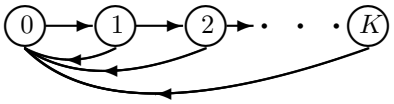

Fig. 1. Allowed transitions in a continuous-time caricature of the WLAN.

$$\lambda_{i,i+1}(\xi) = \frac{(c_i/N)(1 - \exp\{-\langle c, \xi \rangle\})}{1/N} = c_i(1 - \exp\{-\langle c, \xi \rangle\})$$

and the rate at which a node in state i moves to state 0 is

$$\lambda_{i,0}(\xi) = \frac{(c_i/N)(\exp\{-\langle c, \xi \rangle\})}{1/N} = c_i \exp\{-\langle c, \xi \rangle\}.$$

With these transition rates, the continuous-time Markovian caricature is an example of the mean-field model described in Section II-A. The transition graph for this continuous-time model is shown in Figure 1.

2) The SIS Model of spread of an infectious disease: Let there be N individuals in a community. An infection spread in this community can be modelled using our mean-field system as follows. At any given time, each individual is in one of the two states: susceptible (S) or infected (I). Individuals in state S are healthy but can potentially catch the infection. Individuals in state I are infected. The individuals interact with each other and these result in the evolution of the states of the individuals as follows. Whenever a susceptible individual comes in contact with an infected individual, the state changes from S to I. An individual in state I moves to state S after a random amount of time. Assuming *uniform interaction* among the individuals in the community, the rate at which a susceptible individual becomes infected depends only on the fraction of infected individuals in the population. Therefore, this simple model of an epidemic spread is also of the mean-field type. See [5] for an early study for such models.

C. Law of large numbers

Often in practice, properties that are relevant for the system (e.g., throughput in the WLAN example, equilibrium distribution in the SIS model) depend only on the macroscopic behaviour of the system and can be obtained from the empirical measure $\mu_N(t)$. Thus it is important to study the empirical measure process μ_N in some detail. Towards this, we discuss a law of large numbers result for μ_N [8], [9].

Authorized licensed use limited to: J.R.D. Tata Memorial Library Indian Institute of Science Bengaluru. Downloaded on April 18,2025 at 06:58:07 UTC from IEEE Xplore. Restrictions apply.

Assume that the transition rates $\lambda_{z,z'}$, $(z,z') \in \mathcal{E}$, are Lipschitz continuous on $M_1(\mathcal{Z})$. Assume that the initial conditions $\{\mu_N(0)\}_{N\geq 1}$ converge (in $M_1(\mathcal{Z})$) to a probability distribution $\nu \in M_1(\mathcal{Z})$. Then for any fixed $T > 0$, the empirical measure process $\{\mu_N(t), 0 \leq t \leq T\}$ converges in probability to the solution to the ODE

$$\dot{\mu}(t) = \Lambda^*_{\mu(t)}\mu(t),\ 0 \leq t \leq T,\ \mu(0) = \nu, \tag{1}$$

as $N \to \infty$; here, for any $\xi \in M_1(\mathcal{Z})$, Λ_ξ denotes the $|\mathcal{Z}| \times |\mathcal{Z}|$ rate matrix (i.e., $\Lambda_\xi(z,z') = \lambda_{z,z'}(\xi)$ when $(z,z') \in \mathcal{E}$, $\Lambda_\xi(z,z') = 0$ when $(z,z') \notin \mathcal{E}$, and $\Lambda_\xi(z,z) = -\sum_{z'\neq z}\lambda_{z,z'}(\xi)$ for all $z \in \mathcal{Z}$) when the empirical measure is ξ, Λ^*_ξ denotes its transpose, $\mu(t)$ is the empirical measure at time t viewed as a column vector, $\dot{\mu}(t)$ is its time derivative, and the convergence is in the space of $M_1(\mathcal{Z})$-valued trajectories on $[0,T]$. The above ODE is called the McKean-Vlasov equation, and it evolves on the space of probability distributions on $\mathcal{Z}$. Owing to this convergence result, one can view the empirical measure process as a small random perturbation of the McKean-Vlasov equation and, for large N, the properties of the empirical measure process can be well approximated by the corresponding properties of the McKean-Vlasov equation.

D. Propagation of chaos

We now discuss a *decoupling approximation* (also known as *asymptotic independence* property) for the mean-field model [10]. Again, we assume that the transition rates are Lipschitz continuous and that the initial conditions of particles are independent and identically distributed as ν[1]. Consider two tagged particles, 1 and 2. The decoupling approximation tells us that, for each $t > 0$, the random variable $(X_1^{(N)}(t), X_2^{(N)}(t)) \to (Y_1(t), Y_2(t))$ in distribution as $N \to \infty$, where both the marginals $Y_1(t)$ and $Y_2(t)$ are distributed according to the solution to the McKean-Vlasov equation (1) at time t with initial condition ν, and $Y_1(t)$ and $Y_2(t)$ are independent. That is, the initial *chaos* propagates over time[2] in the large-N limit even though particles interact with each other. This property also has an obvious extension to any finite number (but not depending on N) of tagged particles. A similar decoupling approximation holds in the stationary regime [9]. These decoupling approximations play a crucial role in framing the fixed-point equations; see Section III.

III. The four levels of fixed-points

We now describe fixed-point equations at four different levels that arise in the context of our mean-field model. These fixed-point equations are formed based on the convergence and decoupling approximation results discussed in the previous sections. Solutions to these fixed-point equations describe the limiting behaviour (as $N \to \infty$) of the mean-field system and they are good approximations of large but finite-N systems in some cases but not always.

[1]As a consequence, $\mu_N(0) \to \nu$ in distribution as $N \to \infty$.

[2]The initial condition of independence of the particle states at time $t = 0$ can be relaxed to exchangeability so long as $\mu_N(0) \to \nu$ in distribution as $N \to \infty$, where ν is a fixed initial condition.

A. Level-1: Macroscopic observables

We start with the fixed-point equations that arise at the level of macroscopic observables of the system. While an abstract description of a fixed-point equation for generic macroscopic observables is possible, for concreteness, we use the WLAN described in Section II-B1 as a running example. Recall the MAC protocol (with slotted time) and the Markovian continuous-time caricature described in Section II-B1.

Consider a tagged node in a system of N nodes and let γ denote its collision probability (at stationarity). That is, γ is the probability that a packet transmitted by this node undergoes a collision. By symmetry of the model, this collision probability can be taken to be the same for all nodes.

We now compute the total attempt probability of the system as a response to the mean-field induced by γ [6]. If $\beta(\gamma)$ is the attempt probability of the system, by symmetry, the attempt probability of the tagged node is $\beta(\gamma)/N$. Observe that the time instants at which the state of the tagged node becomes 0 are renewal instants. Consider a renewal cycle between two successive visits to state 0. Since the collision probability is γ and the attempt probability in state i is c_i/N, the mean renewal cycle duration (in number of time slots) is

$$\frac{1}{c_0/N} + \frac{\gamma}{c_1/N} + \cdots + \frac{\gamma^K}{c_K/N} + \left(\sum_{n\geq K+1}\frac{\gamma^n}{c_K/N}\right).$$

Similarly, the mean number of slots used for a successful packet transmission during a renewal cycle is

$$1 + \gamma(1 + \gamma(1 + \gamma(\cdots))\cdots).$$

Noting that the tagged node makes exactly one successful transmission in a renewal cycle, by the renewal-reward theorem, the attempt probability of the tagged node is $\frac{\beta(\gamma)}{N} = \frac{1+\gamma+\gamma^2+\cdots}{\frac{N}{c_0}+\frac{\gamma N}{c_1}+\cdots+\frac{\gamma^K N}{c_K}+\frac{\gamma^{K+1}}{1-\gamma}\cdot\frac{N}{c_K}}$, leading to

$$\beta(\gamma) = \frac{1+\gamma+\gamma^2+\cdots}{\frac{1}{c_0}+\frac{\gamma}{c_1}+\cdots+\frac{\gamma^K}{c_K}+\frac{\gamma^{K+1}}{1-\gamma}\cdot\frac{1}{c_K}}. \tag{2}$$

Let us now see the response of the tagged node when the mean-field is such that every other node sees a collision probability γ. Using the decoupling approximation (see Section II-D), the collision probability is

$$1 - (1 - \beta(\gamma)/N)^{N-1} \simeq 1 - \exp\{-\beta(\gamma)\}$$

for large N. Defining $G(\gamma) = 1 - \exp\{-\beta(\gamma)\}$, we see that $G(\gamma)$ is the response collision probability seen by the tagged node starting from an assumption that the collision probability γ. Hence self consistency demands that

$$\gamma^* = G(\gamma^*), \tag{3}$$

which is a fixed-point of G. It can be checked that (3) has a unique solution $\gamma^* \in (0,1)$ when the parameters c_k are decreasing in k [6]. For WLANs with a few tens of nodes and with the standard exponential back-off, i.e., $c_i = c_{i-1}/2$, $i = 1, 2, \ldots, K$, this approximation works well in practice and

Authorized licensed use limited to: J.R.D. Tata Memorial Library Indian Institute of Science Bengaluru. Downloaded on April 18,2025 at 06:58:07 UTC from IEEE Xplore. Restrictions apply.

γ^* is close to the collision probability observed in practice. Thus, the fixed-point analysis at the level of macroscopic observables provides a simple and elegant way to characterise the performance of the WLAN system at equilibrium.

For a generic macroscopic observable for the mean-field model in Section II-A, one can come up with a fixed-point equation involving this observable by using the decoupling approximation in a suitable manner.

B. Level-2: Distribution over states

We now describe a fixed-point equation in the space of (equilibrium) probability distributions on $\mathcal{Z}$.

At time t, suppose that the distribution of each particle's state is ξ. Further, suppose that the initial conditions of the particles are exchangeable and $\mu_N(0) \to \nu$ in distribution as $N \to \infty$, for some deterministic ν. Invoking the decoupling approximation in the transient regime, the empirical measure of the system at time t concentrates near ξ for large N. As the particles evolve over time, even though the empirical measure changes it remains close to ξ. We may assume that the mean-field is ξ. Consider a tagged particle. Let m denote the equilibrium response map that maps a given mean-field to the equilibrium response of the tagged particle. That is, $m(\xi)$ denotes the equilibrium distribution of this tagged particle in response to the states of the particles distributed as ξ. Since the tagged particle's transition rates are given by the rate matrix Λ_ξ when the mean-field is ξ, $m(\xi)$ is the m that solves

$$\sum_{z':(z',z)\in\mathcal{E}} m_{z'}\lambda_{z',z}(\xi) = m_z \sum_{z':(z,z')\in\mathcal{E}} \lambda_{z,z'}(\xi),\ z \in \mathcal{Z}. \quad (4)$$

In (4), the left-hand side (LHS) is the total probability flux into z and the right-hand side (RHS) is the total probability flux out of z for the tagged particle in response to the mean-field ξ. At equilibrium these are in balance, which is the condition (4), and $m(\xi)$ is the m that solves (4).

Self consistency demands that the appropriate equilibrium distribution is the ξ^* that solves

$$m(\xi^*) = \xi^*,$$

a fixed-point equation on the distribution of states of the particle for the equilibrium response map m. Equivalently,

$$(\Lambda_{\xi^*})^*\xi^* = 0. \quad (5)$$

Existence of the fixed point would follow from Brouwer's fixed-point theorem if one can show that mapping m is continuous. If there is a unique fixed-point ξ^* for (5), then one might expect the system to equilibrate at ξ^*. This is not always the case, as we shall explain in Section V.

C. Level-3: Flow of distribution over time

We now describe a fixed-point equation on the space of $M_1(\mathcal{Z})$-valued trajectories for the dynamic response map that maps a given flow of the mean-field over time to the flow of the distribution over states of a particle over time under the given mean-field.

Let an $M_1(\mathcal{Z})$-valued trajectory $\xi(\cdot)$ denote the time-evolution of the distribution of a tagged particle, i.e., the state of the particle at time t is distributed as $\xi(t)$. Suppose that each particle's state evolution is prescribed by the same $\xi(\cdot)$. Then the empirical measure of the system of particles at time t is approximately $\xi(t)$ and hence the time-trajectory of the mean-field is approximately $\xi(\cdot)$. The decoupling approximation suggests that the particles evolve, approximately, in an iid fashion. In response, the state of the tagged particle evolves according to the rate matrix $\Lambda_{\xi(t)}$ at time t under the mean-field $\xi(\cdot)$. This evolution produces a response probability distribution flow over time denoted $\mathcal{M}(\xi(\cdot))$ for the tagged particle in response to the mean-field being $\xi(\cdot)$, i.e., $\mathcal{M}$ is the dynamic response map. Write $m(t) = \mathcal{M}(\xi(t))$. One can write an ODE for $m(t)$ based on the fluxes of probability distributions into and out of each state:

$$\begin{aligned}\dot{m}(t)(z) &= \sum_{z':(z',z)\in\mathcal{E}} m(t)(z')\lambda_{z',z}(\xi(t)) - m(t)(z)\sum_{z':(z,z')\in\mathcal{E}} \lambda_{z,z'}(\xi(t)) \\ &= (\Lambda^*_{\xi(t)}m(t))(z),\ z \in \mathcal{Z}, t \geq 0; \quad (6)\end{aligned}$$

here, the RHS is the difference between the total probability flux into and out of state z in response to the mean-field $\xi(t)$. Self-consistency then demands that

$$\xi^*(\cdot) = \mathcal{M}(\xi^*(\cdot)), \quad (7)$$

a fixed-point equation on the space of $M_1(\mathcal{Z})$-valued trajectories for the response dynamics map $\mathcal{M}$. That is, we must have $m(t) = \xi^*(t)\ \forall t \geq 0$. Substitution of this in (6) yields

$$\dot{\xi}^*(t) = (\Lambda_{\xi^*(t)})^*\xi^*(t),\ t \geq 0, \quad (8)$$

which is precisely the McKean-Vlasov equation (1).

If the transition rates are assumed to be Lipschitz continuous on $M_1(\mathcal{Z})$, then it is easy to check that the mapping $\xi \mapsto \Lambda^*_\xi\xi$ is also Lipschitz continuous on $M_1(\mathcal{Z})$. As a consequence, standard results from ODE theory imply that there exists a unique solution to the ODE (8) and hence there exists a unique fixed-point for (7).

D. Level-4: Distribution over trajectories

We now describe a Level-4 fixed-point equation over the space of probability distributions $D([0,\infty),\mathcal{Z})$, which is the space of $\mathcal{Z}$-valued trajectories. Note that this is different from Level-3 in the following way. Level-3 refers to fixed-points in the space of $M_1(\mathcal{Z})$-valued trajectories over time whereas Level-4 refers to fixed-points in $M_1(D([0,\infty),\mathcal{Z}))$, the space of distributions over $D([0,\infty),\mathcal{Z})$.

Let $Q \in M_1(D([0,\infty),\mathcal{Z}))$ denote the probability distribution of a typical particle on the space of $\mathcal{Z}$-valued trajectories. At each time t, the time-marginal of Q is the distribution of the state of the typical particle at time t. More precisely, it is the *push forward* $Q \circ \pi_t^{-1}$ of the measure Q under the canonical coordinate projection map π_t; $Q \circ \pi_t^{-1}$ is a probability distribution on $\mathcal{Z}$. Assuming the decoupling approximation, the mean-field approximately evolves as $\{Q \circ \pi_t^{-1}, t \geq 0\}$.

Now consider a tagged particle responding to the mean-field $\{Q \circ \pi_t^{-1}, t \geq 0\}$. Let $\mathcal{T}(Q)$ denote the law of the

Authorized licensed use limited to: J.R.D. Tata Memorial Library Indian Institute of Science Bengaluru. Downloaded on April 18,2025 at 06:58:07 UTC from IEEE Xplore. Restrictions apply.

tagged particle's evolution over time. Note that $\mathcal{T}(Q) \in M_1(D([0,T),\mathcal{Z}))$, it is the response to the mean-field induced by Q, and it depends on Q only through the flow $\{Q\circ\pi_t^{-1}, t \geq 0\}$. Since the flow is given by $\xi(t) = \{Q\circ\pi_t^{-1}, t \geq 0\}$, $\mathcal{T}(Q)$ is the law of the Markov process associated with the rates $\Lambda_{\xi(t)}$. Consistency demands that

$$Q^* = \mathcal{T}(Q^*). \tag{9}$$

In probability theory, the problem of existence and uniqueness of Markov processes given a generator (in our case, the transition rates $\Lambda_{\xi(t)}$), or an extended generator, is known as the *martingale problem* [11]. In our case, the generator itself is flexible, is determined by Q, and we are looking for a special solution that fixes the map $Q \mapsto \mathcal{T}(Q)$. This is called a *nonlinear martingale problem* [12]. Let Q^* be a fixed point. The analogue of the Kolmogorov forward equation that governs the evolution of the time-marginals of Q^* is a nonlinear ODE (1), the McKean-Vlasov equation. It turns out that there exists a unique fixed-point for (9) if the transition rates are uniformly Lipschitz continuous on $M_1(\mathcal{Z})$ [12].

IV. Connections and Discussion

We now describe some connections between solutions to fixed-point equations at the four levels. We also discuss questions related to the goodness of Level-2 fixed-points.

A. Level-1 and Level-2

To explain the connection between Level-1 and Level-2 fixed-points, let us consider the WLAN example. Let ξ^* be the unique Level-2 fixed-point.

We first compute the collision probability when the mean-field is at ξ^*. Consider a tagged node. Let γ denote the collision probability of this node when the macroscopic behaviour of the system is ξ^*; this collision probability will be the same for all the nodes. Consider a tagged node in state z_0. As explained in Section II-A, the collision probability of the tagged node is $\gamma \simeq 1 - \exp\{-\langle c, \xi^*\rangle\}$ for large N.

On the other hand, the Level-1 fixed-point analysis already gives us a collision probability γ^*, which is a fixed-point of (3). We thus expect that the two are equal, i.e., $\gamma^* = 1 - \exp\{-\langle c, \xi^*\rangle\}$. This can be seen as follows. On one hand, as explained in Section III-A, a renewal argument tells us that $\beta(\gamma^*)/N$ is the attempt probability of a tagged node. On the other hand, using the decoupling approximation, assuming that each node at state z independently attempts with probability c_z/N, we see that the attempt probability of a tagged node is $\langle c, \xi^*\rangle/N$ when the mean-field is ξ^*. Thus, we have $\langle c, \xi^*\rangle = \beta(\gamma^*)$, and hence $\gamma^* = 1 - \exp\{-\langle c, \xi^*\rangle\}$. This identity $\langle c, \xi^*\rangle = \beta(\gamma^*)$ can indeed be checked using the expression of β given in (2) and the fact that ξ^* satisfies $(\Lambda_{\xi^*})^*\xi^* = 0$. Therefore, the Level-2 fixed-point, which is a distribution over the states of a particle, explains the Level-1 fixed-point which is a macroscopic observable.

Such connections between macroscopic observables and distribution over states of the particles appear in thermodynamics. The relationship between the macroscopic variables of collision probability and attempt probability and their connection to the Level-2 fixed-point in WLAN is analogous to the relationship between macroscopic variables like temperature, pressure, and volume of an ideal gas and their connection to the Maxwell-Boltzmann distribution that describes the distribution of the velocities of the gas molecules. The passage from statistical mechanics to thermodynamics is made rigorous by large deviation theory; minimisations of rate functions in large deviation theory are mapped to the variational principles of free energy minimisation in thermodynamics [13].

B. Level-2 and Level-3

We now describe the connection between the Level-2 and Level-3 fixed-points. Suppose that ξ^* is a Level-2 fixed-point of (5), i.e., $(\Lambda_{\xi^*})^*\xi^* = 0$. The RHS of the Level-3 fixed-point equation (8) implies that when we start the McKean-Vlasov equation at ξ^*, it stays at ξ^* for all times. That is, $\xi(\cdot) \equiv \xi^*$ is a *stationary* Level-3 fixed-point. Conversely, any stationary Level-3 fixed-point of (8) is a Level-2 fixed-point of (5), since we must necessarily have $(\Lambda_\xi)^*\xi = 0$. Thus, Level-2 fixed-points are precisely those points in $M_1(\mathcal{Z})$ that are stationary Level-3 fixed-points.

Further, suppose that there is the unique stationary Level-3 fixed-point and that all its non-stationary Level-3 fixed-points (which are trajectories) converge to ξ^* as time becomes large. Then there is a unique Level-2 fixed-point and it can be obtained from any Level-3 fixed-point by taking the large-time limit. In this special case, the macroscopic behaviour of the N-particle noisy system (which is a noisy version of the Level-2 fixed-point) in the large-N limit *concentrates* on the stationary Level-3 fixed-point [9, Theorem 2.3].

C. Level-3 and Level-4

Suppose that Q is a Level-4 fixed-point of (9). (We omit the superscript * to reduce clutter). Then the time-marginal of Q at time t, $Q(t) = Q \circ \pi_t^{-1}$, is the distribution of a tagged particle at time t (which is also approximately the mean-field at time t when we have large number of particles). This distribution over states changes over time as particles make transitions. As we did in Section III-C, the infinitesimal change in the distribution at time t can be computed by the difference between the total probability flux entering and exiting various states. For state z, this difference is $(\Lambda^*_{Q(t)}Q(t))(z)$. On the other hand, this infinitesimal change is precisely $\dot{Q}(t)(z)$, the z-component of the time-derivative of $Q(t)$. Writing this condition as a vector condition for all z, we get

$$\dot{Q}(t) = \Lambda^*_{Q(t)}Q(t). \tag{10}$$

Thus $Q(\cdot)$ is a Level-3 fixed-point. Hence, time-marginals of Level-4 fixed-points give rise to Level-3 fixed-points.

Conversely, consider a Level-3 fixed-point $\{Q(t), t \geq 0\}$. Clearly, $Q \in \mathcal{T}(\{R : R \circ \pi_t^{-1} = Q(t), t \geq 0\})$, but the set $\mathcal{T}(\{R : R \circ \pi_t^{-1} = Q(t), t \geq 0\})$ can possibly contain other probability distributions on trajectories. However, it turns out that when the transition rates $\lambda_{z,z'}$ are Lipschitz continuous on $M_1(\mathcal{Z})$, $\mathcal{T}(\{R : R \circ \pi_t^{-1} = Q(t), t \geq 0\}) = \{Q\}$ [12]. Thus,

Authorized licensed use limited to: J.R.D. Tata Memorial Library Indian Institute of Science Bengaluru. Downloaded on April 18,2025 at 06:58:07 UTC from IEEE Xplore. Restrictions apply.

under the assumption that the transition rates are Lipschitz continuous on $M_1(\mathcal{Z})$, there is a one-one correspondence between the Level-3 and Level-4 fixed-points.

V. Discussion on multiple fixed-points

We now discuss some issues that cannot be answered by just Level-1 and Level-2 fixed-point analysis. Let us focus on Level-2 fixed-points. By virtue of the relationship between Level-1 and Level-2 fixed-points, the issues for Level-2 will raise similar issues at Level-1.

Multiple fixed-points: Suppose that ξ_1^* and ξ_2^* are two Level-2 fixed-points, and assume that these are *locally asymptotically stable*, i.e., Level-3 fixed-points starting in a small neighbourhood of ξ_1^* (resp. ξ_2^*) converge to ξ_1^* (resp. ξ_2^*) as time becomes large. Since both ξ_1^* and ξ_2^* are Level-2 fixed-points, they are equilibria for the macroscopic state of the system. Therefore, whenever the finite-N system's empirical measure is close ξ_1^* (resp. ξ_2^*), it remains close to ξ_1^* (resp. ξ_2^*) for a long time. However, because of the randomness present in the finite-N system, the behaviour of the system can eventually move to the equilibrium ξ_2^* and remain there for a long time. As explained in Section IV-A, the macroscopic observables may either be close to those arising from ξ_1^* or ξ_2^*, and which of these is appropriate is a question that the Level-2 fixed-point analysis cannot answer by itself.

Limit cycles: Even if there is a unique Level-2 fixed-point, another issue can arise in the presence of stable *limit cycles*. A limit cycle is a set $\mathcal{L} \subset M_1(\mathcal{Z})$ with the property that time-marginals of Level-3 fixed-points with initial conditions lying in $\mathcal{L}$ stays at $\mathcal{L}$ at all times. When there is a stable limit cycle, even if there is a unique fixed-point, the large time behaviour of the system may be in the neighbourhood of the stable limit cycle especially when the unique fixed-point is unstable. Again, the fixed-point analysis would predict ξ^* as the equilibrium behaviour which would then be incorrect.

A sufficient condition for large-time characterisation: One sufficient condition when a unique Level-2 fixed-point ξ^* does indeed characterise the system behaviour, so that the Level-2 fixed-point analysis is indeed a good approach, is when the fixed-point is globally asymptotic stable for the McKean-Vlasov dynamics, see [8], [9]. This involves an examination of the properties of the Level-3 fixed-points.

Beyond the fixed-point analysis: In general, when there are multiple Level-2 fixed-points, limit cycles, and/or chaotic attractors, one has to go beyond the fixed-point analysis and do a finer analysis to understand macroscopic observables. In the above example, a sufficient condition on the Level-3 fixed-point worked. In general, one must study the *large deviations* behaviour of the mean-field of the system. Here, one quantifies probabilities of transitions between Level-2 fixed-points (which are of the form $\exp\{-\text{constant} \times N\}$ for appropriate constants). These can be used to define an *entropy function* s on $M_1(\mathcal{Z})$ such that, under stationary, the probability that we find the mean-field of the system near ξ is proportional to $\exp\{-Ns(\xi)\}$ [14]. This entropy function s is called the large deviations rate function for the system under stationarity. We also mention that the existence of multiple Level-2 fixed-points and limit cycles also results in slower convergence rate for a given finite-N system to its stationarity [15]. As a consequence, one has to be cautious in interpreting Level-1 and Level-2 fixed-points, particularly when these are not unique or when there are stable limit cycles or when there are chaotic attractors for the limiting dynamics.

Short-term unfairness: In the context of WLANs, at the level of nodes, one can encounter *short-term unfairness* [16], [17]. There are parameter choices where the wireless channel gets clogged by a single node for a long time which results in low throughput for the others. A transition then occurs, after which the situation is similar with another single node hogging the channel at the expense of the others, and so on. The connection between this phenomenon, chaoticity of an attractor, quasi-stationarity [18] remains to be explored.

References

[1] G. Bianchi, "IEEE 802.11 – Saturated Throughput Analysis," *IEEE Communications Letters*, vol. 12, pp. 318–320, 1998.

[2] M. Mitzenmacher, "The power of two choices in randomized load balancing," *IEEE Transactions on Parallel and Distributed Systems*, vol. 12, no. 10, pp. 1094–1104, 2001.

[3] M. Bramson, Y. Lu, and B. Prabhakar, "Randomized load balancing with general service time distributions," *ACM SIGMETRICS Performance Evaluation Review*, vol. 38, no. 1, pp. 275–286, 2010.

[4] T. Vasantam, A. Mukhopadhyay, and R. R. Mazumdar, "The mean-field behavior of processor sharing systems with general job lengths under the sq (d) policy," *Performance Evaluation*, vol. 127, pp. 120–153, 2018.

[5] W. O. Kermack and A. G. McKendrick, "A contribution to the mathematical theory of epidemics," *Proceedings of the Royal Society of London. Series A, Containing papers of a mathematical and physical character*, vol. 115, no. 772, pp. 700–721, 1927.

[6] A. Kumar, E. Altman, D. Miorandi, and M. Goyal, "New insights from a fixed point analysis of single cell IEEE 802.11 WLANs," in *IEEE INFOCOM 2006*, 2006.

[7] A. Kumar, D. Manjunath, and J. Kuri, *Wireless Networking*. Morgan Kaufmann, 2008.

[8] M. Benaïm and J.-Y. LeBoudec, "A class of mean field interaction models for computer and communication systems," *Performance Evaluation*, vol. 65, no. 1, pp. 823–838, 2008.

[9] C. Bordenave, D. McDonald, and A. Proutiere, "A particle system in interaction with a rapidly varying environment: Mean field limits and applications," *Netw. Heterog. Media*, vol. 5, no. 1, pp. 31–62, 2010.

[10] A.-S. Sznitman, "Topics in propagation of chaos," in *Ecole d'été de Probabilités de Saint Flour (1989)*, ser. Springer Lecture Notes in Mathematics, 1991, vol. 1464, pp. 166–251.

[11] D. W. Stroock and S. R. S. Varadhan, *Multidimensional Diffusion Processes*. Springer-Verlag Berlin Heidelberg, 2006.

[12] C. Graham, "Chaoticity on path space for a queueing network with selection of the shortest queue among several," *J. Appl. Probab.*, vol. 37, no. 1, pp. 198–211, 2000.

[13] R. S. Ellis, *Entropy, Large Deviations, and Statistical Mechanics*. Springer, 2006, vol. 1431, no. 821.

[14] V. S. Borkar and R. Sundaresan, "Asymptotics of the invariant measure in mean field models with jumps," *Stoch. Syst.*, vol. 2, no. 2, pp. 322–380, 2012.

[15] S. Yasodharan and R. Sundaresan, "Large time behaviour and the second eigenvalue problem for finite state mean-field interacting particle systems," *arXiv preprint arXiv:1909.03805*, 2019.

[16] V. Ramaiyan, A. Kumar, and E. Altman, "Fixed point analysis of single cell ieee 802.11e wlans: Uniqueness and multistability," *IEEE/ACM Transactions on Networking*, vol. 16, no. 5, pp. 1080–1093, 2008.

[17] A. Bhattacharya and A. Kumar, "Analytical modeling of IEEE 802.11-Type CSMA/CA networks with short term unfairness," *IEEE/ACM Transactions on Networking*, vol. 25, no. 6, pp. 3455–3472, 2017.

[18] P. K. Pollett, "Quasi-stationary distributions: a bibliography," *https://people.smp.uq.edu.au/PhilipPollett/papers/qsds/qsds.pdf*, 2008.

Authorized licensed use limited to: J.R.D. Tata Memorial Library Indian Institute of Science Bengaluru. Downloaded on April 18,2025 at 06:58:07 UTC from IEEE Xplore. Restrictions apply.