

Version with bug fixes of full text published in *Performance Evaluation*, Vol. 65, Nr. 11-12, pp. 823-838, 2008.

A Class Of Mean Field Interaction Models for Computer and Communication Systems

Michel Benaïm and Jean-Yves Le Boudec

Abstract

We consider models of N interacting objects, where the interaction is via a common resource and the distribution of states of all objects. We introduce the key scaling concept of intensity; informally, the expected number of transitions per object per time slot is of the order of the intensity. We consider the case of vanishing intensity, i.e. the expected number of object transitions per time slot is $o(N)$. We show that, under mild assumptions and for large N , the occupancy measure converges, in mean square (and thus in probability) over any finite horizon, to a deterministic dynamical system. The mild assumption is essentially that the coefficient of variation of the number of object transitions per time slot remains bounded with N . No independence assumption is needed anywhere. The convergence results allow us to derive properties valid in the stationary regime. We discuss when one can assure that a stationary point of the ODE is the large N limit of the stationary probability distribution of the state of one object for the system with N objects. We use this to develop a critique of the fixed point method sometimes used in conjunction with the decoupling assumption.

Key words: Mean Field Interaction Model, Markov Chain, Dynamical System, MAC protocol, Reputation Systems, Game Theory, Decoupling Assumption, Fixed Point, Bianchi's Formula

1 Introduction

We consider models of objects interacting in discrete time, where objects are observable only through their state. Objects interact with each other and with a resource, which also has a finite number of states. We are interested in the case where the number N of objects is large compared to the number of possible states for one individual object or for the resource.

A number of papers have addressed such models – see for example [1–4]. In the limit of large number of objects, it is found that the system can be approximated by a deterministic, usually non linear, dynamical system, called the mean field limit. The mean field limit is in discrete or continuous time, depending on how the model scales with the number of objects. More precisely, if the expected number of transitions per object per time slot vanishes when N grows, the limit is in continuous time; this is the case considered in this paper. Else, the limit is in discrete time – see [3] for such a case.

Our goal is twofold: (1) find results that are widely applicable in practice, i.e. the model should be as little constrained as possible and (2) the technical assumptions should be reasonably simple to verify. The model of Benaïm and Weibull [2] (discussed in detail in Section 4.3) comes close to these goals, but its applicability is limited in some cases, as it does not allow a resource nor, for example, pairwise meetings of objects. The model of Sharma, Ganesh and Key [1] goes one step beyond and considers concurrent transitions; this was used there to provide the first mean field analysis of the 802.11 MAC protocol with one base station. The model of Bordenave, McDonald and Proutière [4,5] offers more expression power; it supports a resource, which was used in [5] to extend the analysis of [1] to multiple base stations.

Still, the model in [5] has some limitations, which we overcome in this paper. First, the required assumptions are complex, perhaps because [5] allows an infinite, enumerable state space for one object and for the resource. In contrast, we consider only the finite case. Second, convergence to the mean field is established using compactness arguments typical of weak convergence over infinite horizons, which does not allow to make statements for the stationary regime other than in the case where the mean field limit has a unique global attractor. Third, the model assumes that objects independently decide whether they will attempt to make a transition (see Section 4.4 for details). This limits the applicability of their model; for example, in a pairwise interaction model, exactly two objects do a transition at every time slot; given that two objects have decided to do a transition, all other objects cannot. We find that this restriction is unnecessary; it appears that convergence to the mean field derives from exchangeability arguments and not from independence.

In this paper we propose a generic mean field interaction model for N interacting objects and a resource. Informally, we assume that (1) the model scales with N such that the intensity, i.e. the number of transitions per object per time slot vanishes, and (2) the coefficient of variation of the total number of objects that do a transition in one time slot remains bounded. We also make some mild regularity assumptions on the model parameters. All assumptions are simple to verify (Section 3), and are illustrated by several examples in Section 4. We do not assume any form of independence in the transitions of objects or of the resource.

In Section 5, we show convergence to a deterministic system, solution of an Ordi-

nary Differential Equation (ODE). We also explain on one example how to derive the ODE in a straightforward manner (Section 4.1). The convergence is in mean square and in probability over any finite horizon. The method of proof is inspired by the large body of results for stochastic approximation algorithms [6–9]. In Section 6 we review and exploit the link between convergence to mean field, the mean field approximation and the decoupling assumption¹. Then we establish results for the stationary regime (Section 7). For large N , the stationary distribution of the occupancy measure of all objects tends to be supported by the Birkhoff center of the ODE (see Section 7 for a definition). If the ODE has a unique global attractor, we recover that the stationary distribution of the occupancy measure is concentrated at this attractor. Last, we point to the well known fact that uniqueness of a stationary point of the ODE does not imply convergence to this stationary point, and develop from there a critique of the so-called fixed point method, sometimes used to analyze systems of interacting objects in stationary regime (Section 8). Proofs are given in appendix.

Our model is motivated by the case where the intensity (number of transitions per object per time slot) is of the order of $\frac{1}{N}$, but we give a general treatment, which does not make any assumption on the rate of decay of the intensity, as it appears that this does not make the model any more complex.

2 Notation List

- $\Delta = \{\vec{m} \in \mathbb{R}^I, \sum_{i=1}^I m_i = 1 \text{ and } m_i \geq 0 \text{ for all } i\}$
- $\vec{e}_i = (0, \dots, 1, \dots, 0)^T$ where 1 is at the i th position
- $\epsilon(N)$ intensity; the number of transitions per object per time slot is order of $\epsilon(N)$; a typical case is $\epsilon(N) = \frac{1}{N}$; Equation (6)
- η^N stationary probability distribution of $Y^N(t)$
- $\vec{f}^N(\vec{m}, j)$ drift, the expected change to M^N in one time slot; Equation (4)
- $\vec{f}(\vec{m}, j) = \lim_{N \rightarrow \infty} \frac{\vec{f}^N(\vec{m}, j)}{\epsilon(N)}$
- $\vec{F}(\vec{m})$ right-handside of ODE, Equation (15)
- $\Phi_\tau(\vec{m})$ flow induced by the ODE, Equation (16)
- I number of states for one object
- J number of states for the resource
- $K_{j,j'}^N(\vec{m})$ transition matrix for resource, Equation (3)

¹ We call “mean field approximation” the independence assumption that is asymptotically true when N is large. This should not be confused with the approximation that consists in replacing a non-mean field interaction model with a mean field interaction model. This is also sometimes called the mean field approximation, as e.g. in [10].

$M^N(t)$	occupancy measure of $X_n^N(t)$, Equation (2)
$\bar{M}^N(\tau)$	time-rescaled occupancy measure; defined for $\tau \in \mathbb{R}^+$ and such that $\bar{M}^N(\frac{t}{N}) = M^N(t)$ for all $t \in \mathbb{N}$
$\bar{\mu}(\tau)$	mean field limit, deterministic limit of $\bar{M}^N(\tau)$ for large N
N	number of objects
$P_{i,i'}^N(\vec{m}, j)$	marginal transition probabilities for one object, Equation (5)
ϖ^N	stationary probability distribution of $M^N(t)$
$\mathcal{R} = \{1, 2, \dots, J\}$	state space for the resource
$R^N(t)$	state of resource at time $t \in \mathbb{N}$
$\mathcal{S} = \{1, 2, \dots, I\}$	state space for one object
$X_n^N(t)$	state of object n at time $t \in \mathbb{N}$
$Y^N(t) = (X_1^N(t), \dots, X_N^N(t), R^N(t))$	state of the entire system

3 Definition of the Mean Field Interaction Model with Vanishing Intensity

3.1 Definition of Mean Field Interaction Model

Time $t \in \mathbb{N}$ is discrete. There are N objects. Object n has a state $X_n^N(t)$ in the finite set $\mathcal{S} = \{1, 2, \dots, I\}$. There is a common resource $R^N(t)$ in the finite set $\mathcal{R} = \{1, \dots, J\}$. We assume that $Y^N(t) = (X_1^N(t), \dots, X_N^N(t), R^N(t))$ is a homogeneous Markov chain on $\mathcal{S}^N \times \mathcal{R}$. Furthermore, we assume that we can observe the state of an object but not its label n . Mathematically, this translates into the assumption that the transition kernel \mathcal{K}^N of $Y^N(t)$ is invariant under any permutation of the labeling of the N objects. \mathcal{K}^N is defined by

$$\begin{aligned} & \mathcal{K}^N(i_1, \dots, i_N, j; i'_1, \dots, i'_N, j') \\ &= \mathbb{P} \left(X_1^N(t+1) = i'_1, \dots, X_N^N(t+1) = i'_N, R^N(t+1) = j' \right. \\ & \quad \left. \mid X_1^N(t) = i_1, \dots, X_N^N(t) = i_N, R^N(t) = j \right) \end{aligned}$$

and the assumption is that, for any permutation σ of the index set $\{1, 2, \dots, N\}$:

$$\mathcal{K}^N(i_1, \dots, i_N, j; i'_1, \dots, i'_N, j') = \mathcal{K}^N(i_{\sigma(1)}, \dots, i_{\sigma(N)}, j; i'_{\sigma(1)}, \dots, i'_{\sigma(N)}, j') \quad (1)$$

The process $Y^N(t)$ is called a *mean field interaction model* with N objects.

Define the *occupancy measure* $M^N(t)$ as the vector of frequencies of states $i \in \mathcal{S}$ at time t , i.e.:

$$M_i^N(t) = \frac{1}{N} \sum_{n=1}^N 1_{\{X_n^N(t)=i\}} \quad (2)$$

Since we assume that objects can be observed only through their states, it follows that $(M^N(t), R^N(t))$ is also a homogeneous Markov chain. Its state space is $\Delta \times \mathcal{R}$, with $\Delta = \{\vec{m} \in \mathbb{R}^I, \sum_{i=1}^I m_i = 1 \text{ and } m_i \geq 0 \text{ for all } i\}$. The process $(X_1^N(t), M^N(t), R^N(t))$ is also Markov. This means that the evolution of one object $X_1^N(t)$ depends on the other objects only through the occupancy measure $M^N(t)$ and the resource $R^N(t)$.

Note that we do not assume any form of independence between object transitions. The transitions of different objects may be dependent, as they may be dependent on the transition done by the resource.

3.2 Intensity

We are interested in asymptotic results when the number of objects N is large. For this, we need further assumptions on how \mathcal{K}^N scales with N , as different models lead to different scaling results. In this section we give an intuitive introduction to the concept of *intensity* of the mean field interaction model; informally, it may be defined such that the probability that one arbitrary object does a transition in one time slot is of the order of the intensity. A formal definition is used the drift and is given in the next section.

For example, in the model of Le Boudec, McDonald and Mundinger [3], every object does order of one transition per time slot; an intensity is the constant 1 and there is an asymptotic result at the natural time scale: under mild assumptions, the random process $M^N(t)$ converges to a deterministic discrete time process $\vec{\mu}(t)$ as $N \rightarrow \infty$, called the *mean field limit*. Here the intensity does not vanish as N grows and the limiting process $\vec{\mu}(t)$ is in discrete time.

In contrast, consider the models of Benaim and Weibull [2] and of Bordenave, McDonald and Proutière [5]; these models are different, but they have in common that an object does one transition in one time slot with probability $1/N$; here an intensity is $\epsilon(N) = 1/N$; it vanishes as N grows, and, to obtain an asymptotic result, we need to re-scale the process to continuous time. The re-scaled process $\bar{M}^N(\tau)$ ($\tau \in \mathbb{R}^+$) that corresponds to $M^N(t)$ ($t \in \mathbb{N}$) is such that

$$\bar{M}^N\left(\frac{t}{N}\right) = M^N(t) \text{ for all } t \in \mathbb{N}$$

In [2], $\bar{M}^N(\tau)$ is taken to be the piecewise linear interpolation that satisfies Equation (14) (and is thus continuous). In [5], $\bar{M}^N(\tau)$ is taken to be piecewise constant and right continuous. These are inessential differences, which influence only the details of the proofs of convergence. It is found in [2,5] that $\bar{M}^N(\tau)$ converges to a continuous time deterministic process $\bar{\mu}(\tau)$ as $N \rightarrow \infty$. The mean field limit $\bar{\mu}(\tau)$ is solution of an ODE.

In this paper we generalize the results of [2,5] and consider a larger class of mean field interaction models that have in common that the intensity goes to 0 when $N \rightarrow \infty$. In other words, the expected number of object transitions per time slot per object is $o(1)$, and the expected number of object transitions per time slot is $o(N)$.

3.3 Definition of Mean Field Interaction Model with Vanishing Intensity

We now precisely define the hypotheses we put on the mean field interaction model to ensure convergence to the mean field. We give a simple set of conditions, that should be easy to verify by inspection. In appendix we give a more abstract set of conditions.

We assume a mean field interaction model as in Section 3.1. We assume the most general model, i.e. transitions of different objects and of the resource may be dependent (there may be several correlated transitions in one time slot). Our hypotheses are conditions on how the model parameters scale with N , plus some mild regularity assumptions.

Define $K^N(\vec{m})$, the marginal transition matrix for the resource:

$$K_{j,j'}^N(\vec{m}) = \mathbb{P}\left(R^N(t+1) = j' \mid M^N(t) = \vec{m}, R^N(t) = j\right) \quad (3)$$

Note that by the assumptions in Section 3.1 $K_{j,j'}^N(\vec{m})$ is indeed independent of t and n .

Also define the *drift* $\vec{f}^N(\vec{m}, j)$ for $\vec{m} \in \Delta$ and $j \in \mathcal{R}$ as the expected change to M^N in one time slot:

$$\begin{aligned} \vec{f}^N(\vec{m}, j) &= \mathbb{E}\left(M^N(t+1) - M^N(t) \mid M^N(t) = \vec{m} \text{ and } R^N(t) = j\right) \\ &= \sum_{(i,i') \in \mathcal{S}, i \neq i'} m_i P_{i,i'}^N(\vec{m}, j) (\vec{e}_{i'} - \vec{e}_i) \end{aligned} \quad (4)$$

where P^N is the marginal transition probability for an arbitrary object:

$$P_{i,i'}^N(\vec{m}, j) = \mathbb{P}\left(X_n^N(t+1) = i' \mid M^N(t) = \vec{m}, R^N(t) = j, X_n^N(t) = i\right) \quad (5)$$

The assumptions are as follows.

H1 (The resource does not scale with N) $\lim_{N \rightarrow \infty} K_{j,j'}^N(\vec{m}) = K_{j,j'}(\vec{m})$ exists for all $\vec{m} \in \Delta, (j, j') \in \mathcal{R}^2$. The matrix $K(\vec{m})$ is indecomposable² for all $\vec{m} \in \Delta$.

H2 (Intensity vanishes at a rate $\epsilon(N)$) There exists a function $\epsilon(N)$ (the ‘‘vanishing intensity’’) such that $\lim_{N \rightarrow \infty} \epsilon(N) = 0$ and

$$\lim_{N \rightarrow \infty} \frac{f^N(\vec{m}, j)}{\epsilon(N)} = \vec{f}(\vec{m}, j) \text{ exists for all } \vec{m} \in \Delta, j \in \mathcal{R} \quad (6)$$

H3 (Second moment of number of object transitions per time slot) Let $W^N(t)$ be an upper bound on the number of objects that do a transition in time slot t , i.e. $\sum_n 1_{\{X_n^N(t) \neq X_n^N(t+1)\}} \leq W^N(t+1)$. We assume that

$$\mathbb{E} \left(W^N(t)^2 \mid M^N(t) = \vec{m}, R^N(t) = j \right) \leq c_1 N^2 \epsilon(N)^2$$

where c_1 is a constant (independent of t, N, \vec{m} and j).

The next assumptions require that K^N and \vec{f}^N depend on N as a smooth function of $\frac{1}{N}$ that is well defined at the limit of $N \rightarrow \infty$.

H4 ($K_{j,j'}^N(\vec{m})$ is a smooth function of $\frac{1}{N}$ and \vec{m}) For all $(j, j') \in \mathcal{R}^2$, there exists $\alpha_{j,j'} > 0$ and a function $\kappa_{j,j'}(\vec{m}, \alpha)$ defined on $\Delta \times [0, \alpha_{j,j'}]$ such that $\kappa_{j,j'}$ has continuous derivatives everywhere (including at the boundary of its domain) and

$$K_{j,j'}^N(\vec{m}) = \kappa_{j,j'} \left(\vec{m}, \frac{1}{N} \right) \quad (7)$$

H5 ($f^N(\vec{m}, j)$ is a smooth function of $\frac{1}{N}$ and \vec{m}) For all $j \in \mathcal{R}$ there exists $\beta_j > 0$ and a function $\varphi_j(\vec{m}, \alpha)$ defined on $\Delta \times [0, \beta_j]$ such that φ_j has continuous derivatives everywhere (including at the boundary of its domain) and

$$\frac{f^N(\vec{m}, j)}{\epsilon(N)} = \varphi_j \left(\vec{m}, \frac{1}{N} \right) \quad (8)$$

Comment 1.

A sufficient (but not necessary) condition for **H2** is

H2a $\lim_{N \rightarrow \infty} \frac{P_{i,i'}^N(\vec{m}, j)}{\epsilon(N)} = P_{i,i'}(\vec{m}, j)$ exists for all $\vec{m} \in \Delta, j \in \mathcal{R}$ and $(i, i) \in \mathcal{S}^2, i \neq i'$. In other words, the probability transition for one object scales like $\epsilon(N)$.

² I.e. it has a unique invariant probability.

This makes a bridge with the intuitive interpretation of the intensity given in Section 3.2. Indeed, let $V^N(t+1)$ be the frequency of objects that do a transition in time slot $t+1$, i.e.

$$V^N(t+1) = \frac{1}{N} \sum_{n=1}^N 1_{\{X_n^N(t+1) \neq X_n^N(t)\}}$$

Then

$$\begin{aligned} & \lim_{N \rightarrow \infty} \frac{\mathbb{E} \left(V^N(t+1) \mid M^N(t) = \vec{m}, R^N(t) = j \right)}{\epsilon(N)} \\ &= \lim_{N \rightarrow \infty} \sum_{i \in \mathcal{S}} \sum_{i' \in \mathcal{S}, i' \neq i} \frac{m_i P_{i,i'}^N(\vec{m}, j)}{\epsilon(N)} = \sum_{i \neq i'} m_i P_{i,i'}(\vec{m}, j) \end{aligned}$$

so, if **H2a** holds, we recover the intuitive interpretation that the expected number of transitions per object per time slot is of the order of the intensity.

If **H2a** holds, then one can interpret **H3** as follows: the coefficient of variation³ of the number of objects that do a transition in one time slot remains bounded for large N .

Comment 2.

H4 and **H5** may be verified by showing that $K_{j,j'}^N(\vec{m})$ is a combination of a finite number of known functions with continuous derivatives (such as polynomials, exponentials, logarithms, etc.), which is well defined even if we replace $\frac{1}{N}$ by 0. See Section 4.1 for an example. Note that $K_{j,j'} = \kappa_{j,j'}(\vec{m}, 0)$ and $\vec{f}(\vec{m}, j) = \varphi_j(\vec{m}, 0)$.

4 Examples

In this section we review new and existing examples and show that they fit in our framework.

4.1 Two-Step Malware Propagation

This is a small, numerical example that serves as illustration in various parts of the paper. An object is a node in a peer to peer or an ad-hoc mobile network. Nodes can be infected by malware, and the infection may occur in two forms. In

³ The coefficient of variation of a nonnegative random variable that is not identically 0 is the ratio of its standard deviation to its mean

a first form, an infected node becomes dormant – it does not show any symptom. A node may become dormant by contact with another dormant node or by contact with a web service that contains malware. A dormant node may become active (i.e. malignant) by meeting with another dormant node; or, in special conditions (for example presence of both nodes close to a wireless access point) by interaction with an already active node. Nodes that are neither dormant nor active are said to be susceptible. Susceptible nodes may directly become dormant or active by infectious contact with email.

We can model such a system of N nodes as a mean field interaction model. The state space for objects is $\mathcal{S} = \{“D”, “A”, “S”\}$ (for dormant, active and susceptible). There is no resource in this simplified model. We denote the occupancy measure with $M^N(t) = (D^N(t), A^N(t), S^N(t))$.

The interaction can be simulated using the following rules. At every time step, one node, say n , is picked at random among N .

- (1) (Case 1) If n is dormant, it may recover and become susceptible, with probability δ_D . (Case 2) It may create a rendez-vous with another dormant node, if it succeeds, both become active. This occurs with probability proportional to the frequency of other dormant nodes $\lambda(ND^N - 1)/N$.
- (2) If n is active, it may do one of the following two actions. (Case 3) The former is to change the state of a dormant node. This occurs with a probability $\beta \frac{D^N}{h+D^N}$. This probability depends on D^N via a saturating function: it increases up to some maximum value β [11]. This expresses a dependency on limited resources. If this transition occurs, the dormant node that is affected is chosen randomly uniformly. There is no change to n . (Case 4) The second possible action is to return to the susceptible state, with probability δ_A .
- (3) (Case 5) If n is susceptible, it may become dormant with probability $\alpha_0 + rD^N$; α_0 models infection by email (independent of the state of the system) and rD^N the probability that n becomes infected by contact with a dormant node. (Case 6) Else, it may directly become active, with probability α .

The sum of all probabilities in any of the cases is less than 1, so it is also possible that node n performs no action.

We now compute the drift $\vec{f}^N(D, A, S)$, using the equation

$$\text{drift} = \sum_{\text{cases}} \text{prob of case} \times \text{effect of case on } M^N \quad (9)$$

where “effect” means the contribution to $M^N(t+1) - M^N(t)$. The cases, their probabilities and effects can be obtained from the above description and are in Table 1. Recall that the occupancy measure M^N is here denoted with (D, A, S) .

Thus

$$\vec{f}^N(D, A, S) = \frac{1}{N} \begin{pmatrix} -D\delta_D - 2D\lambda\frac{ND-1}{N} - A\beta\frac{D}{h+D} + S(\alpha_0 + rD) \\ 2D\lambda\frac{ND-1}{N} + A\beta\frac{D}{h+D} - A\delta_A + S\alpha \\ D\delta_D + A\delta_A - S(\alpha_0 + rD) - S\alpha \end{pmatrix} \quad (10)$$

case	prob	effect on (D, A, S)
1	$D\delta_D$	$\frac{1}{N}(-1, 0, 1)$
2	$D\lambda\frac{ND-1}{N}$	$\frac{1}{N}(-2, +2, 0)$
3	$A\beta\frac{D}{h+D}$	$\frac{1}{N}(-1, +1, 0)$
4	$A\delta_A$	$\frac{1}{N}(0, -1, +1)$
5	$S(\alpha_0 + rD)$	$\frac{1}{N}(+1, 0, -1)$
6	$S\alpha$	$\frac{1}{N}(0, +1, -1)$

Table 1
Probabilities and effects for Example 4.1.

We now check the hypotheses. There is no resource so **H1** and **H4** are trivially true. We have

$$\lim_{N \rightarrow \infty} \frac{\vec{f}^N(D, A, S)}{1/N} = \begin{pmatrix} -D\delta_D - 2D^2\lambda - A\beta\frac{D}{h+D} + S(\alpha_0 + rD) \\ 2D^2\lambda + A\beta\frac{D}{h+D} - A\delta_A + S\alpha \\ D\delta_D + A\delta_A - S(\alpha_0 + rD) - S\alpha \end{pmatrix} = \vec{f}(D, A, S) \quad (11)$$

thus **H2** is satisfied with intensity $\epsilon(N) = \frac{1}{N}$.

The number of objects that do a transition in one time slot is 0, 1 or 2, therefore a bound on its mean square is $c_1 = 4$, and **H3** is satisfied. Finally, $\vec{f}^N(D, A, S)$ is a rational fraction with respect to \vec{m} and $1/N$ thus **H5** is satisfied.

See Figure 1, bottom panels, for a numerical example.

4.2 One Object per Time Slot. Game Theory. Benaim and Weibull [2,12]

In the model of [2], the resource $R^N(t)$ is absent, i.e. $Y^N(t) = (X_1^N(t), \dots, X_N^N(t))$. An object is a player, and its state is $i = (r, h)$ where r is the player's role (or position) and h is the strategy chosen at this time.




Fig. 1. Example in Section 4.1. Evolution over time of proportions of dormant $D(t)$ versus active nodes $A(t)$. The third variable $S(t)$ is not shown since $D(t) + A(t) + S(t) = 1$. Top: mean field limit, obtained by numerical solution of the ODE in Equation (17). Bottom: simulation with $N = 1000$ nodes. The cross is the stationary point, obtained by solving $F(D, A, S) = 0$ for $(D, A, S) \in \Delta$. We numerically find that it is a global attractor. The small circle is the initial value (i.e. at time 0). Parameters: $\beta = 0.01, \delta_A = 0.005, \delta_D = 0.0001, \alpha_0 = \alpha = 0.0001, h = 0.3, r = 0.1, \lambda = 0.0001$.

In every time slot exactly one object, randomly chosen among N , does a transition, using the transition matrix $A(M^N(t))$. The model is completely specified by its initial condition and the matrix $A(\vec{m})$ for all $\vec{m} \in \Delta$.

Benaïm and Weibull show in [2] almost sure convergence of $\bar{M}^N(\tau)$ to a deterministic $\bar{\mu}(\tau)$ over any finite time horizon ($\tau \in [0, T]$). Further, large τ asymptotic results are also given.

This fits our framework, with $P_{i,i'}^N(\vec{m}) = \frac{1}{N}A_{i,i'}(\vec{m})$ for $i \neq i'$, which shows that an intensity is $\epsilon(N) = \frac{1}{N}$. We can also recover the intensity by observing that the number of transitions in one time slot per object is Bernoulli $(1/N)$.

The total number of transitions per time slot is constant equal to 1, therefore the bound in **H3** is satisfied. Benaïm and Weibull [12] generalize the model to make $A(\vec{m})$ dependent on N .

4.3 Poisson(1) Object per Time Slot. MAC Protocol. Bordenave, McDonald and Proutière[5].

The model in [5] is more general than our framework, as it allows the state space \mathcal{S} for one object or for the resource \mathcal{R} to be infinite. This is at the expense of considerable complexity. In this section, we study the simplified model obtained from [5] when we impose that \mathcal{R} and \mathcal{S} are finite.

This model is used to evaluate the performance of Medium Access Control (MAC) used in Wireless LANs. Here an object is a wireless transmitter, its state is its backoff or transmit stage; the resource represents the states (idle, in collision, in transmission) of a small number of channels.

At every time slot, every object decides to try a transition or not, with probability $1/N$, independent of each other and of all past. Given that an object decides to try a transition, the probability that it moves from i to i' , given that the resource is in state j and the occupancy measure is \vec{m} is $A_{i,i'}^N(\vec{m}, j)$. As in our model, transitions of several objects and the resource are typically not independent. It is assumed that $A_{i,i'}^N(\vec{m}, j) \rightarrow A_{i,i'}(\vec{m}, j)$ as $N \rightarrow \infty$. The initial distribution of object states is assumed to be exchangeable.

Bordenave et al [5] find a weak convergence result for $\bar{M}^N(\tau)$ to a deterministic process $\bar{\mu}(\tau)$ over infinite horizons. An asymptotic result for large τ is also given in the case where the deterministic limit has a unique attractor towards which all trajectories converge. We show later in this paper results that generalize these findings.

This model fits in our framework, with $P_{i,i'}^N(\vec{m}) = \frac{1}{N}A_{i,i'}^N(\vec{m})$ for $i \neq i'$, which shows that an intensity is $\epsilon(N) = \frac{1}{N}$.

An upper bound on the number of objects that do a transition is $W^N(t)$, the number of objects that decide to try a transition. $W^N(t) \sim \text{Bin}(1/N, N)$ thus $\mathbb{E} \left((W^N(t))^2 \right) = 2 - \frac{1}{N} \leq 2$, which shows that **H3** is satisfied. The other conditions (or in fact more general forms, as in the appendix of this paper) are met.

4.4 Pairwise Interaction / Arbitrary Interaction

This appears to be a new type of mean field interaction model, which we define as follows. At every time slot, the resource R^N does a transition independently of the past according to some transition kernel K^N . Then an ordered pair of objects (n_1, n_2) is picked at random uniformly among the $N(N-1)$ possible ones. The two objects do a transition independent of the past according to

$$\begin{aligned}
& \mathbb{P} \left(X_{n_1}^N(t+1) = i'_1, X_{n_2}^N(t+1) = i'_2 \mid \right. \\
& \quad \left. X_{n_1}^N(t) = i_1, X_{n_2}^N(t) = i_2, M^N(t) = \vec{m}, R^N(t) = j \right) \\
& = L_{i_1, i_2; i'_1, i'_2}^N(\vec{m}, j)
\end{aligned}$$

where $L^N(\vec{m}, j)$ is a stochastic matrix. We assume that for large N the matrices $K^N(\vec{m})$ and $L^N(\vec{m}, j)$ converge to some $K(\vec{m})$ and $L(\vec{m}, j)$.

This can be used to model pairwise interaction, as for example in [3] where an object is a reputation record held by one peer about some common subject. When two peers meet, they influence each other's opinion and may modify their reputation records. An example can also be found in distributed robotics [13].

This fits in our model. Intuitively, an intensity is $\frac{1}{N}$ since an object is picked in one time slot with probability $\frac{2}{N}$. Formally, the drift is

$$\vec{f}^N(\vec{m}, j) = \sum_{i_1, i_2} m_{i_1} \frac{Nm_{i_2} - 1_{\{i_1=i_2\}}}{N-1} L_{i_1, i_2; i'_1, i'_2}(\vec{m}, j) \frac{\vec{e}_{i'_1} + \vec{e}_{i'_2} - \vec{e}_{i_1} - \vec{e}_{i_2}}{N} \quad (12)$$

and thus $N\vec{f}^N(\vec{m}, j) \rightarrow \vec{f}(\vec{m}, j)$ with

$$\vec{f}(\vec{m}, j) = \sum_{i_1, i_2, i'_1, i'_2} m_{i_1} m_{i_2} L_{i_1, i_2; i'_1, i'_2}(\vec{m}, j) (\vec{e}_{i'_1} + \vec{e}_{i'_2} - \vec{e}_{i_1} - \vec{e}_{i_2}) \quad (13)$$

which establishes **H2**. The total number of transitions in one time slot is upper bounded by the constant equal to 2, this shows **H3**. The other assumptions are satisfied if the transition matrices are smooth enough.

This model can easily be extended to interaction with arbitrary numbers of objects, provided that the expected number of objects involved in an individual meeting grows less fast than N . More precisely, assume that in one time slot we pick B^N objects and decide that they do a meeting ($B^N = 2$ for pairwise interaction). We allow B^N to be random, with mean β_N and standard deviation σ_N . If $\beta_N = o(N)$ and $\sigma_N/\beta_N = O(1)$ then the model is a mean field interaction model with vanishing intensity $\epsilon(N) = \frac{\beta_N}{N}$. For example, the number of objects involved in a meeting may be order of \sqrt{N} or $\ln N$.

5 Convergence to Mean Field

5.1 The Mean Field Limit

We define the re-scaled process \bar{M}^N as a continuous time process that satisfies

$$\begin{cases} \bar{M}^N(t\epsilon(N)) = M^N(t) \text{ for all } t \in \mathbb{N} \\ \bar{M}^N(\tau) \text{ is affine on } \tau \in [t\epsilon(N); (t+1)\epsilon(N)] \end{cases} \quad (14)$$

and we similarly define $\bar{X}_n^N(\tau)$ as a re-scaled version of $X_n^N(t)$.

We find convergence results for $\bar{M}^N(\tau)$ to a deterministic process $\vec{\mu}(\tau)$, which satisfies the ODE

$$\frac{d\vec{\mu}}{d\tau} = \vec{F}(\vec{\mu}) \quad \text{with } \vec{F}(\vec{m}) = \sum_{j=1}^J \pi_j(\vec{m}) \vec{f}(\vec{m}, j) \quad (15)$$

where $\pi(\vec{m})$ is the invariant probability of the transition matrix $K(\vec{m})$ (i.e. $\pi(\vec{m})K(\vec{m}) = \pi(\vec{m})$ and $\sum_j \pi_j(\vec{m}) = 1$), which is unique by **H1**, and \vec{f} is defined in **H2**.

5.2 Convergence Result

Let $\Phi_\tau(\vec{m})$ be the solution at time τ of the ODE (15) with initial condition \vec{m} , i.e.

$$\begin{cases} \frac{d\Phi_\tau(\vec{m})}{d\tau} = \vec{F}(\Phi_\tau(\vec{m})) \\ \Phi_0(\vec{m}) = \vec{m} \end{cases} \quad (16)$$

It follows from **H4** and **H5** that $\Phi_\tau(\vec{m})$ is well defined for all $\tau \geq 0$ and \vec{m} (see Section A for details).

The central result is the following bounding theorem and its corollary, from which we derive both convergence to the mean field and properties of the stationary regime.

Theorem 1 *For all $T > 0$ there exists constants $C_1(T), C_2(T)$ and a random variable $B^N(T)$ such that*

$$\sup_{0 \leq \tau \leq T} \left\| \bar{M}^N(\tau) - \Phi_\tau(\vec{m}) \right\| \leq C_1(T) \left(B^N(T) + \left\| M^N(0) - \vec{m} \right\| \right)$$

and

$$\mathbb{E} \left(\left\| B^N(T) \right\|^2 \right) \leq C_2(T) \epsilon(N)$$

Corollary 1 *If $M^N(0) \rightarrow \vec{m}$ in probability [resp. in mean square] as $N \rightarrow \infty$ then $\sup_{0 \leq \tau \leq T} \left\| \bar{M}^N(\tau) - \vec{\mu}(\tau) \right\| \rightarrow 0$ in probability [resp. in mean square], where $\vec{\mu}(\tau)$ satisfies the ODE (15) and $\vec{\mu}(0) = \vec{m}$.*

The corollary can be used to approximate $M^N(t)$ for the system with N objects by $\vec{\mu}(t\epsilon(N))$ where $\vec{\mu}$ is the solution of the ODE with same initial condition.

A numerical illustration is shown in Figure 1 for the malware propagation example in Section 4.1. Here the occupancy measure is (D, A, S) and the ODE for the mean field limit is

$$\begin{pmatrix} \frac{dD}{d\tau} \\ \frac{dA}{d\tau} \\ \frac{dS}{d\tau} \end{pmatrix} = \vec{f}(D, A, S) \quad (17)$$

with \vec{f} defined in Equation (11).

6 Mean Field Independence and Decoupling Assumption

When convergence of the occupancy measure to a deterministic limit does occur, we have an additional well known property, due to Sznitman, called *Mean Field Independence*, or *Propagation of Chaos* (see also [14,10,15]). We recall it now as it is required for the discussion in Section 8.

Theorem 2 ([16]) *Consider a mean field interaction model with vanishing intensity and assume that the initial occupancy measures are such that the assumptions of Corollary 1 hold. Assume in addition that the collection of objects at time 0 $(X_1^N(0), \dots, X_N^N(0))$ is exchangeable. For any fixed k and τ :*

$$\lim_{N \rightarrow \infty} \mathbb{P} \left(\bar{X}_1^N(\tau) = i_1, \dots, \bar{X}_k^N(\tau) = i_k \right) = \mu_{i_1}(\tau) \dots \mu_{i_k}(\tau) \quad (18)$$

The theorem can be used to do the following approximation:

$$\mathbb{P} \left(X_1^N(t) = i_1, \dots, X_k^N(t) = i_k \right) \approx \mu_{i_1} \left(\frac{t}{N} \right) \dots \mu_{i_k} \left(\frac{t}{N} \right) \quad (19)$$

i.e. the distribution of $X_n^N(t)$ can be approximated by $\vec{\mu}(\frac{t}{N})$ and any finite number of objects are approximately independent. Equation (19) is called the *mean field independence* or *decoupling property*. Note that any two objects are asymptotically independent of each other, but they still depend on the occupancy measure.

Note that we cannot always assume that the collection of objects is exchangeable at time 0 (and in the rest of this paper we do not make any such assumption)⁴.

We can still find a relation between the distribution of one object and the mean field limit if some stronger assumption holds on the stationary regime of the ODE, as we see in the next section.

7 Stationary Regime

We are interested in the stationary regime of M^N and how it relates to the stationary regime of the ODE.

A general statement can be made using the concept of *Birkhoff Center*. For $\vec{m} \in \Delta$ the *omega limit set* of \vec{m} , denoted $\omega(\vec{m})$ is the set of points $\vec{p} = \lim_{\tau_k \rightarrow \infty} \Phi_{\tau_k}(\vec{m})$ for some $\tau_k \rightarrow \infty$. It is a nonempty compact invariant subset of Δ . A point $\vec{m} \in \Delta$ is said to be *recurrent* if $\vec{m} \in \omega(\vec{m})$. The Birkhoff center of Φ is the closure of the set of recurrent points. For example, on Figure 2, the Birkhoff center is the union of the limit cycle and the stationary point.

$Y^N(t) = (X_1^N(t), \dots, X_n^N(t), R^N(t))$ is a Markov chain with finite state space, therefore it has some invariant probability. Let η^N be one of them; if Y^N is irreducible, η^N is unique and can be interpreted as the stationary regime for Y^N . We denote with \mathbb{P}_{η^N} the probability obtained when we initialize Y^N with η^N . Under \mathbb{P}_{η^N} , $(Y^N(t))_{t \in \mathbb{N}}$ is thus a (strict sense) stationary sequence.

Let ϖ^N be the corresponding probability distribution for M^N , i.e.

$$\mathbb{P}_{\eta^N} (M^N(t) \in A) = \int_A \varpi^N(dx)$$

for any measurable subset A of Δ . We say that a probability ϖ on Δ is a limit point for ϖ^N if there is sequence $N_k \rightarrow \infty$ such that $\varpi^{N_k} \rightarrow \varpi$. The last limit is in the

⁴ This is particularly true if we are interested in the evolution of objects that start from different initial conditions – see [3] for some examples

weak sense, i.e. for any continuous function h defined on Δ ,

$$\lim_{k \rightarrow \infty} \int_{\Delta} h(x) \varpi^{N_k}(dx) = \int_{\Delta} h(x) \varpi(dx)$$

The following theorem is a quite general result. It is a consequence of Theorem 1 and can be proven as in Corollary 3.2 in [9].

Theorem 3 *The support of any limit point of ϖ^N is a compact set included in the Birkhoff center of Φ .*

In some cases, we can say more.

Corollary 2 *Assume the ODE (15) has a unique stationary point \vec{m}^* to which all trajectories converge.*

- (1) *Under the stationary distributions η^N , $M^N(0)$ converges in distribution and in probability to \vec{m}^**
- (2) *Assume η^N is exchangeable. For any fixed k and t ,*

$$\lim_{N \rightarrow \infty} \mathbb{P}_{\eta^N} \left(X_1^N(t) = i_1, \dots, X_k^N(t) = i_k \right) = m_{i_1}^* \dots m_{i_k}^*$$

(Proof in appendix). Item 2 assumes that the distribution of the stationary regime is invariant by permutation of the object labels. This occurs necessarily if Y^N is irreducible, since then there is a unique stationary regime, and by the symmetry assumptions in Section 3.1 it must be exchangeable. Item 2 states that k objects are asymptotically independent in the stationary regime. This implies (by taking $k = 1$) that

$$\lim_{N \rightarrow \infty} \eta_1^N = \vec{m}^*$$

where η_1^N is the stationary distribution for one object.

If we assume in addition that the Markov chain Y^N is irreducible for any N , then we have a result independent of initial conditions. It follows from Corollary 2 after noticing that, in this case, the sample path averages converge to the stationary distribution, by the ergodic theorem of Markov chains:

Corollary 3 *Assume the ODE (15) has a unique stationary point \vec{m}^* to which all trajectories converge and Y^N is irreducible. Then for any initial condition of Y^N*

$$\lim_{N \rightarrow \infty} \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t 1_{\{X_1^N(s)=i_1, \dots, X_k^N(s)=i_k\}} = m_{i_1}^* \dots m_{i_k}^* \quad \text{almost surely}$$

8 A Critique of the Fixed Point Method

The mean field independence is often used in the analysis of interacting objects, sometimes with the name of “decoupling assumption”, in conjunction with a “fixed point” method [17–19]. In this section we describe this method in the context of mean field interaction models, then point to a potential pitfall.

8.1 The Decoupling Assumption Should Be Handled with Care

Consider a mean field interaction model with N objects. We assume that, for a fixed N , it has a unique stationary distribution. We are interested in approximating for large N the stationary distribution of state for one object, η_1^N . The decoupling assumption can be made if N is large and the model satisfies a scaling law as in Section 3.3. We know from Theorem 2 that this is equivalent to assuming that the occupancy measure can be approximated by the mean field limit, which satisfies the ODE $\frac{d\vec{\mu}}{d\tau} = \vec{F}(\vec{\mu})$. The fixed point method then consists in finding a value \vec{m}^* of the mean field that is stationary, i.e., that satisfies $\vec{F}(\vec{m}^*) = 0$. If there is a single solution to this equation, it is taken as approximation of η_1^N . This seems to make sense since $(\eta_1^N)_i$ is the long term average of time spent by one object in state i , and by Theorem 2, it can be approximated by μ_i . We have seen in Corollary 3 that this method is valid if \vec{m}^* is an attractor to which all trajectories converge. This is for example the case in Figure 1.

However, the fact that there is a unique solution to $\vec{F}(\vec{m}^*) = 0$ (i.e. a unique fixed point) is not sufficient for this to hold. Figure 2 shows a numerical example (same example as on Figure 1 but with a different value of one parameter). There is a unique fixed point \vec{m}^* (cross on the figure) and any solution of the ODE starting at this fixed point will remain there. But this fixed point is an unstable equilibrium. For any other initial condition, the mean field is attracted by a limit cycle. In this case, we cannot approximate the occupancy measure or η_1^N by \vec{m}^* . Here the Birkhoff center is the union of the limit cycle and the unstable equilibrium. By Theorem 3 we can say that most of the time, the occupancy measure $M^N(t)$ remains close to this set – as we can verify in the simulation.

In such a case, the decoupling assumption *may not hold in the stationary regime*. Indeed, one can show with simple symmetry arguments that, if η^N is exchangeable (for example because the stationary probability η^N is unique), then for any fixed k

$$\lim_{N \rightarrow \infty} \mathbb{P}_{\eta^N} \left(X_1^N(t) = i_1, \dots, X_k^N(t) = i_k \mid M^N(t) = \vec{m} \right) = m_{i_1} \cdots m_{i_k} \quad (20)$$

i.e. the decoupling assumption holds only conditional to the value of the occupancy measure. On the example of Figure 2, the occupancy measure oscillates in the sta-

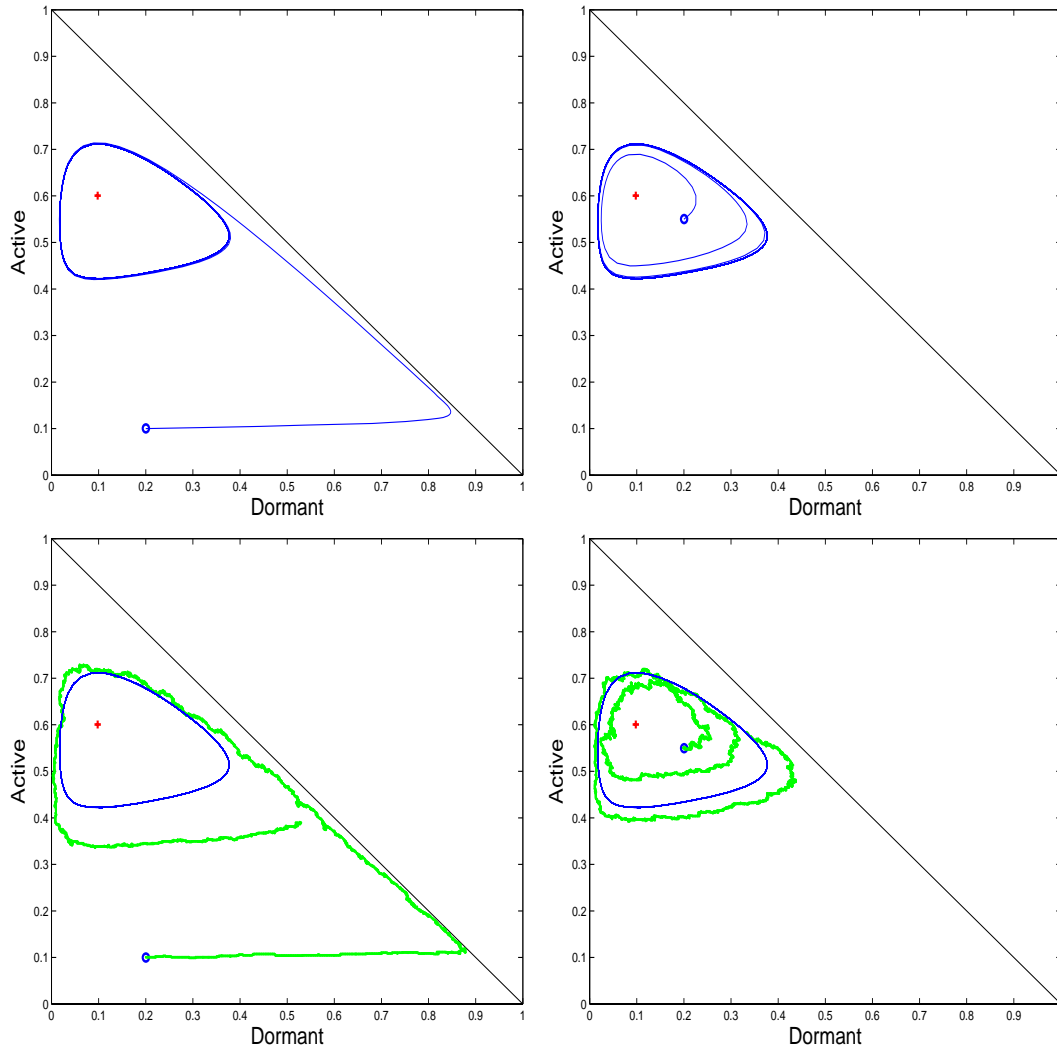


Fig. 2. Example in Section 4.1, same parameters as Figure 1 except $h = 0.1$. Evolution over time of proportions of dormant versus active nodes active nodes. Top: mean field limit (ODE), bottom: simulation with $N = 1000$ nodes. There is a unique stationary point, solution of $\vec{F}(\vec{n}) = 0$, (cross), but it is an unstable equilibrium. All trajectories are attracted to a limit cycle (plain line). The small circle is the initial value.

tionary regime along its limit cycle. The distribution of states of two objects follows these oscillations, in particular, they are correlated and the decoupling assumption does not hold in the stationary regime.

8.2 Bianchi's Formula

This phenomenon is visible in the analysis of the 802.11 protocol known as Bianchi's formula. It was developed by Bianchi [20,21] and re-explained in a more general setting by Kumar, Altman, Miorandi and Goyal [17]. In short, consider only the

case where all wireless nodes hear each other. It is shown in these references that it is sufficient to model the backoff process, as follows. An object is a wireless node and its state is its backoff stage in $\{0, 1, \dots, K\}$. There is no resource. At every time slot, a node in state i makes a transmission attempt with probability $\frac{q_i}{N}$. If more than one node does a transmission attempt, there is a collision and all nodes that made a transmission attempt increment their states by $1 \bmod (K + 1)$. If exactly one node makes an attempt, it succeeds and its state changes to 0. This model was introduced in [1] and it is easy to see that it is a mean field interaction model with intensity $\frac{1}{N}$. The ODE for the mean field limit is

$$\begin{aligned} \frac{dm_0}{d\tau} &= -m_0q_0 + \beta(\vec{m})(1 - \gamma(\vec{m})) + q_Km_K\gamma(\vec{m}) \\ \frac{dm_i}{d\tau} &= -m_iq_i + m_{i-1}q_{i-1}\gamma(\vec{m}) \quad i = 1, \dots, K \end{aligned}$$

where $\beta(\vec{m}) = \sum_{i=0}^K q_i m_i$ and $\gamma(\vec{m}) = 1 - e^{-\beta(\vec{m})}$. The fixed point method consists in finding the stationary points of the ODE; one finds that \vec{m} is a stationary point if and only if

$$m_i = \frac{\gamma^i}{q_i} \frac{1}{\sum_{k=0}^K \frac{\gamma^k}{q_k}} \quad i = 0, \dots, K \quad (21)$$

where γ is solution of

$$\gamma = 1 - e^{-\beta} \quad (22)$$

$$\beta = \frac{\sum_{k=0}^K \gamma^k}{\sum_{k=0}^K \frac{\gamma^k}{q_k}} \quad (23)$$

Bianchi's formula in this setting is Equation (23). It relates the collision probability γ to the attempt rate β . It is shown in [17,5] that the fixed point problem (23, 22) has a unique solution, from where it is concluded that Equation (23) holds in the stationary regime. As we argued earlier, to establish the validity of the method (and thus of Bianchi's formula), one would need to show not only that there is a unique fixed point, but, more importantly, that all trajectories of the ODE converge to the stationary point. As is visible in Figure 2, one does not imply the other. It was shown in [5] that there is a unique stable point to which all trajectories converge, for the infinite variant of this model $K = +\infty$, and when q_k is an exponential decreasing sequence with $q_0 < \ln 2$ (and in [1] for $K = 1$). The validity of Bianchi's formula appears to remain non demonstrated for the other cases.

One can relate this issue to the following alternative explanation, introduced in [17]. Equation (23) is interpreted in [17] as

$$\beta = \frac{\text{expected number of transmission attempts in one round}}{\text{expected number of time slots in one round}} \quad (24)$$

where a round is defined as a sequence of time slots for one object that starts when the object reaches state $i = 0$ and ends when it returns to 0. This is a Palm calculus formula and as such is true in the stationary regime. However, in deriving this interpretation, one equates the numerators of Equations (23) and (24), which amounts to assuming that, within one round, the collision probability γ remains constant. Since γ is a function of \vec{m} , this is true if the Birkhoff center is a single point. If, in contrast, the stationary regime has a limit cycle, then γ may oscillate and not stay constant within one round; thus Equation (23) does not necessarily hold. We recover the fact that when using the fixed point method, it is not sufficient to show uniqueness of the stationary point; one needs in addition to show that the stationary point of the ODE is an attractor to which all trajectories converge.

9 Conclusion

We have given a generic result for the convergence of a model of interacting objects to the solution of an ordinary differential equation. The assumptions are both very general and simple to verify in practice. A striking feature is that, contrary to existing results, no independence assumption is required anywhere.

We also shed some light on the relation between the stationary regime of the ODE and that of the original stochastic system with N objects. The result is essentially positive: at first order, the ODE is a good approximation of the occupancy measure of the stochastic system, even in stationary regime, in the sense that the occupancy measure spends most of its time in the Birkhoff center of the ODE.

However, we point out that the decoupling assumption needs to be handled with care. One needs to study the stationary regime of the ODE; simple results exist only if the ODE has a unique fixed point to which all trajectories converge. In such a case, the decoupling assumption also holds in the stationary regime, and further, the fixed point is the limit for large N of the stationary probability of state for one object in the system with N objects. In all other cases this may not be true, even when there is a unique fixed point.

A Proof of Theorem 1

The proof is built on a result for stochastic approximation algorithms, which we state first.

A.1 A Stochastic Approximation Algorithm

Let $\underline{\mathcal{R}}$ be a finite set, $l \in \mathbb{N}$ and Δ be a compact convex subset of \mathbb{R}^l , with a nonempty interior (so that \mathbb{R}^l is a smallest affine space containing Δ). For each $\epsilon > 0$, let $Z^\epsilon = (\underline{M}^\epsilon(t), \underline{R}^\epsilon(t))_{t \in \mathbb{N}}$ be a discrete time Markov chain living in $\Delta \times \underline{\mathcal{R}}$ such that

$$\underline{M}^\epsilon(t+1) - \underline{M}^\epsilon(t) = \epsilon G^\epsilon(t+1) \quad (\text{A.1})$$

and

$$\mathbb{P}(G^\epsilon(t+1) \in dx, \underline{R}^\epsilon(t+1) = j' | \underline{M}^\epsilon(t) = m, \underline{R}^\epsilon(t) = j) = \underline{K}_{jj'}^\epsilon(m) \nu_{m,j}^\epsilon(dx) \quad (\text{A.2})$$

where for all $m \in \Delta$, $\epsilon > 0$, $j \in \underline{\mathcal{R}}$, $\underline{K}^\epsilon(m)$ is a Markov transition matrix on $\underline{\mathcal{R}}$ and $\nu_{m,j}^\epsilon$ a probability measure supported by $\frac{\Delta - m}{\epsilon}$.

We assume in this section that:

H1a $\underline{f}_j^\epsilon(m)$ converges uniformly in m , as $\epsilon \rightarrow 0$, to some function $\underline{f}_j(m)$; and $\underline{K}^\epsilon(m)$ converges uniformly in m , as $\epsilon \rightarrow 0$ to some *indecomposable* Markov matrix $\underline{K}(m)$.

H2a There exists some constant $C > 0$ such that

$$\int_{\mathbb{R}^m} \|x\|^2 \nu_{m,j}^\epsilon(dx) \leq C$$

for all $\epsilon > 0$, $m \in \Delta$.

H3a The maps \underline{K}^ϵ and $\underline{f}_j^\epsilon(m) = \int x \nu_{m,j}^\epsilon(dx)$ are Lipschitz continuous uniformly in ϵ . That is

$$\|\underline{K}^\epsilon(m) - \underline{K}^\epsilon(m')\| + \|\underline{f}_j^\epsilon(m) - \underline{f}_j^\epsilon(m')\| \leq L \|m - m'\|$$

where L is independent on ϵ .

The matrix $\underline{K}(m)$ being indecomposable it admits a unique invariant probability measure $\underline{\pi}(m)$ solution to $\underline{\pi}(m) = \underline{\pi}(m) \underline{K}(m)$. Let $\underline{F}(m) = \sum_j \underline{\pi}_j(m) \underline{f}_j(m)$. Since $\underline{\pi}(m)$ depends smoothly on $\underline{K}(m)$, it follows from **H1a** and **H3a** that \underline{F} is

Lipschitz-continuous on Δ . Without loss of generality we may assume (by extending \underline{F}) that \underline{F} is defined and bounded on \mathbb{R}^l . By standard results it then induces a global flow $\{\Phi_\tau\}_{\tau \geq 0}$ on \mathbb{R}^l defined by the Cauchy problem $\frac{d\Phi_\tau(x)}{d\tau} = \underline{F}(\Phi_\tau(x))$ with initial condition $\Phi_0(x) = x$.

Let $\hat{M}^\epsilon : \mathbb{R}_+ \mapsto \Delta$ denote the continuous time process defined by

$$\begin{cases} \hat{M}^\epsilon(t\epsilon) = M^\epsilon(t) \text{ for all } t \in \mathbb{N} \\ \hat{M}^\epsilon \text{ is affine on } [t\epsilon, (t+1)\epsilon] \end{cases}$$

The following proposition is a classical averaging result for stochastic approximation algorithms with constant step size. This type of result has been proved under various sets of assumptions in the literature (see e.g. (Beneveniste et al [6], Duflo [7] for general statements, Kushner and Yin [8] for numerous weak convergence results, Benaim [9] and Benaim and Weibull [2] for statements similar to Proposition 1 in a slightly more restricted setting).

Proposition 1 *For all $T > 0$ there exist constants $C_1(T), C_2(T)$ and a random variable $B^\epsilon(T)$ such that*

$$\sup_{0 \leq t \leq T} \|\hat{M}^\epsilon(t) - \Phi_t(x)\| \leq C_1(T)[B^\epsilon(T) + \|\underline{M}^\epsilon(0) - x\|]$$

and

$$\mathbb{E}(\|B^\epsilon(T)\|^2) \leq C_2(T)\epsilon.$$

Proof: Set

$$U^\epsilon(t+1) = G^\epsilon(t+1) - \underline{F}(\underline{M}^\epsilon(t))$$

so that

$$\underline{M}^\epsilon(t+1) - \underline{M}^\epsilon(t) = \epsilon[\underline{F}(\underline{M}^\epsilon(t)) + U^\epsilon(t+1)].$$

The following lemma follows from Lipschitz continuity of \underline{F} and Gronwall's lemma (see e.g Benaim [22]):

Lemma 1 *For all $T > 0$ there exists $C_1(T) > 0$ such that*

$$\sup_{0 \leq \tau \leq T} \|\hat{M}^\epsilon(\tau) - \Phi_\tau(x)\| \leq C_1(T)[B^\epsilon(T) + \|\underline{M}^\epsilon(0) - x\|]$$

where

$$B^\epsilon(T) = \epsilon \left(\sup_{0 \leq t \leq T/\epsilon} \left\| \sum_{i=1}^t U^\epsilon(i) \right\| \right).$$

It follows from **H1a** that for ϵ small enough, $\underline{K}^\epsilon(m)$ is indecomposable. We let $\underline{\pi}^\epsilon(m)$ denotes its invariant probability and

$$\underline{F}^\epsilon(m) = \sum_i \underline{\pi}_i^\epsilon(m) \underline{f}_i^\epsilon(m).$$

Set

$$\begin{aligned} U^{1,\epsilon}(t+1) &= \underline{F}^\epsilon(\underline{M}^\epsilon(t)) - \underline{F}(\underline{M}^\epsilon(t)), \\ U^{2,\epsilon}(t+1) &= G^\epsilon(t+1) - \underline{f}_{\underline{R}^\epsilon(t+1)}^\epsilon(\underline{M}^\epsilon(t)), \end{aligned}$$

and

$$U^{3,\epsilon}(t+1) = \underline{f}_{\underline{R}^\epsilon(t+1)}^\epsilon(\underline{M}^\epsilon(t)) - \underline{F}^\epsilon(\underline{M}^\epsilon(t)).$$

Then

$$U^\epsilon(t) = \sum_{i=1}^3 U^{\epsilon,i}(t)$$

and

$$B^\epsilon(T) \leq \left(\sum_{i=1}^3 B^{\epsilon,i}(T) + \epsilon \right)$$

where

$$B^{\epsilon,j}(T) = \epsilon \left(\sup_{0 \leq t \leq T/\epsilon} \left\| \sum_{i=1}^t U^{\epsilon,j}(i) \right\| \right).$$

Our next goal is to bound the quantities $B^{\epsilon,j}(T)$.

- By hypothesis H1a, \underline{F}^ϵ converges uniformly to \underline{F} . Hence

$$\lim_{\epsilon \rightarrow 0} B^{\epsilon,1}(T) = 0. \tag{A.3}$$

- Note that

$$\mathbb{E}(G^\epsilon(t+1) | \underline{M}^\epsilon(t) = m, \underline{R}^\epsilon(t) = j) = \mathbb{E}(\underline{f}_{\underline{R}^\epsilon(t+1)(m)} | \underline{R}^\epsilon(t) = j) = \sum_i \underline{K}_{ji}^\epsilon(m) \underline{f}_i^\epsilon(m)$$

and

$$\begin{aligned} \mathbb{E}(\|G^\epsilon(t+1)\|^2 | \underline{M}^\epsilon(t) = m, \underline{R}^\epsilon(t) = j) &= \mathbb{E}(\|\underline{f}_{\underline{R}^\epsilon(t+1)(m)}\|^2 | \underline{R}^\epsilon(t) = j) \\ &= \sum_i \underline{K}_{ji}^\epsilon(m) \int \|x\|^2 \nu_{m,i}^\epsilon(dx) \leq C. \end{aligned}$$

This makes $\sum_{t=1}^n U^{2,\epsilon}(t)$ a martingale whose variance is bounded by $4nC$. Thus, by Doob's inequality

$$\mathbb{E}(B^{\epsilon,2}(T)^2) = O(\epsilon) \tag{A.4}$$

We now pass to $B^{3,\epsilon}$.

- Given any indecomposable Markov matrix \underline{K} with invariant probability $\underline{\pi}$ and $g, h \in \mathbb{R}^{\mathcal{R}}$, we let $\underline{\pi}g = \sum_i \pi_i g_i$ and $(Kg)_i = \sum_j \underline{K}_{ij} g_j$. By indecomposability, the *Poisson equation*

$$g - Kg = h - \underline{\pi}h \quad \text{with boundary condition } \underline{\pi}g = 0 \tag{A.5}$$

admits a unique solution g depending smoothly on h and \underline{K} (the smoothness easily follows from the implicit function theorem).

We now let $g^\epsilon(m)$ denote the solution to the poisson problem (A.5) where \underline{K} is replaced by $\underline{K}^\epsilon(m)$ and h by $f_i^\epsilon(m)$. Hence we can rewrite $U^{3,\epsilon}(t+1)$ as

$$U^{3,\epsilon}(t+1) = \sum_{j=4}^6 U^{j,\epsilon}(t+1)$$

where

$$U^{4,\epsilon}(t+1) = g_{\underline{R}^\epsilon(t+1)}^\epsilon(\underline{M}^\epsilon(t)) - (\underline{K}^\epsilon(\underline{M}^\epsilon(t))g^\epsilon(\underline{M}^\epsilon(t)))_{\underline{R}^\epsilon(t)},$$

$$U^{5,\epsilon}(t+1) = [\underline{K}^\epsilon(\underline{M}^\epsilon(t))g^\epsilon(\underline{M}^\epsilon(t))_{\underline{R}^\epsilon(t)} - (\underline{K}^\epsilon(\underline{M}^\epsilon(t+1))g^\epsilon(\underline{M}^\epsilon(t+1)))_{\underline{R}^\epsilon(t+1)}],$$

and

$$U^{6,\epsilon}(t+1) = (\underline{K}^\epsilon(\underline{M}^\epsilon(t+1))g^\epsilon(\underline{M}^\epsilon(t+1)))_{\underline{R}^\epsilon(t+1)} - (\underline{K}^\epsilon(\underline{M}^\epsilon(t))g^\epsilon(\underline{M}^\epsilon(t)))_{\underline{R}^\epsilon(t+1)}].$$

The term $U^{4,\epsilon}(t+1)$ is a bounded martingale difference. Hence (with an obvious definition of $B^{\epsilon,4}$) Doob's inequality gives

$$\mathbb{E}((B^{\epsilon,4})^2) = O(\epsilon).$$

The sum $\sum_{1 \leq i \leq n} U^{5,\epsilon}(i)$ reduces to

$$\underline{K}^\epsilon(\underline{M}^\epsilon(1))g_{\underline{R}^\epsilon(t)}^\epsilon(\underline{M}^\epsilon(1)) - \underline{K}^\epsilon(\underline{M}^\epsilon(m))g_{B^{\epsilon,4}(n)}^\epsilon(\underline{M}^\epsilon(n))]$$

which is bounded. Hence

$$B^{\epsilon,5} = O(\epsilon).$$

By Lipschitz continuity of the maps \underline{K}^ϵ and g^ϵ

$$\|U^{\epsilon,6}(t+1)\| = O(\|\underline{M}^\epsilon(t+1) - \underline{M}^\epsilon(t)\|) = O(\epsilon G^\epsilon(t+1)).$$

Hence

$$\mathbb{E}(\|U^{\epsilon,6}(t+1)\|^2) = O(\epsilon^2)$$

and

$$\mathbb{E}((B^{\epsilon,6})^2) = O(\epsilon).$$

Finally we get that

$$\mathbb{E}((B^{3,\epsilon})^2) = O(\epsilon). \tag{A.6}$$

Putting together estimates (A.3), (A.4) and (A.6) gives the result. \square

A.2 Mapping The Mean Field Interaction Model With Vanishing Intensity to a Stochastic Recurrence Model

Consider now a model as in Section 3. First, we show that **H4** and **H5** imply uniform Lipschitz continuity and uniform convergence analog to **H1a** and **H3a**. For

$\epsilon \geq 0$ small enough we can define $K_{j,j'}^{\epsilon}(\vec{m})$ and $\vec{f}^{\epsilon}(\vec{m}, j)$ such that $K_{j,j'}^N(\vec{m}) = K_{j,j'}^{\frac{1}{N}}(\vec{m})$ and $\vec{f}^N(\vec{m}, j) = \vec{f}^{\frac{1}{N}}(\vec{m}, j)$ for N large enough.

Lemma 2 *The functions $K_{j,j'}^{\epsilon}(\vec{m})$ and $\vec{f}^{\epsilon}(\vec{m}, j)$ are Lipschitz continuous in \vec{m} uniformly in ϵ and converge uniformly as $\epsilon \rightarrow 0$.*

Proof: By **H4** and **H5**, $K_{j,j'}^{\epsilon}(\vec{m})$ is continuously differentiable in both $\epsilon \geq 0$ and \vec{m} . By compactness, there is a finite bound c_2 on the norm of the differential of $K_{j,j'}^{\epsilon}(\vec{m})$, independent of ϵ and \vec{m} . By Taylor's integral formula,

$$\left\| K_{j,j'}^{\epsilon}(\vec{m}) - K_{j,j'}^{\epsilon'}(\vec{m}') \right\| \leq c_2 (\|\vec{m} - \vec{m}'\| + |\epsilon - \epsilon'|) \quad (\text{A.7})$$

Uniform Lipschitz continuity [resp. uniform convergence] follows by letting $\epsilon = \epsilon'$ [resp. $\vec{m} = \vec{m}'$] in Equation (A.7). \square

Second, we show that **H3** implies a second moment property analog to **H2a**:

Lemma 3
$$\mathbb{E} \left(\left\| \frac{M^N(t+1) - M^N(t)}{\epsilon(N)} \right\|^2 \leq 2c_1 \right)$$

Proof: Let $\mathcal{N}(t)$ be the set of objects that do a transition in time slot t ; it has at most $W^N(t)$ elements and thus

$$\begin{aligned} & \left\| M^N(t+1) - M^N(t) \right\|^2 \\ &= \frac{1}{N^2} \sum_{(n,n') \in \mathcal{N}(t)^2} \langle \vec{e}_{X_{n'}^N(t+1)} - \vec{e}_{X_n^N(t)}, \vec{e}_{X_n^N(t+1)} - \vec{e}_{X_n^N(t)} \rangle \\ &\leq \frac{2}{N^2} W^N(t)^2 \end{aligned}$$

\square

Third, we map the mean field interaction model to a stochastic approximation algorithm as follows. First, a minor difference is that the former is defined for $N \in \mathbb{N}$ and the latter for $\epsilon > 0$. We address this by associating a stochastic approximation algorithm with parameter $\epsilon < 1$ to a mean field interaction model with parameter $N_{\epsilon} = \lfloor \frac{1}{\epsilon} \rfloor$. Second, identify Δ with a convex subset with non empty interior of \mathbb{R}^l with $l = I - 1$ (thus identify \vec{m} and m). Next, a difference between the two models lies in the dependence on the resource: in the mean field interaction model, the drift $\vec{f}^N(\vec{m}, j)$ is defined as the expected drift at time $t + 1$ conditional to the resource having value j before the transition (i.e. conditional to $R^N(t) = j$), whereas in the stochastic approximation algorithm the probability measure $\nu_{m,j}^{\epsilon}$ is conditional to the value of \underline{R} after the transition (i.e. conditional to $R^{\epsilon}(t+1) = j$). We address this by keeping in \underline{R}^{ϵ} the current and the previous state, i.e. $\underline{\mathcal{R}} = \mathcal{R}^2$ and

$$\begin{aligned}\underline{R}^\epsilon(t) &= \left(R^{N_\epsilon}(t+1), R^{N_\epsilon}(t) \right) \text{ for } t \geq 1 \\ \underline{R}^\epsilon(0) &= \left(R^{N_\epsilon}(0), 1 \right)\end{aligned}$$

and the stochastic approximation algorithm is defined by letting

$$G^\epsilon(t) = \frac{1}{1/\epsilon(N_\epsilon)} \left(M^{N_\epsilon}(t+1) - M^{N_\epsilon}(t) \right)$$

Note that if K is indecomposable, then so is \underline{K} , and the unique stationary probabilities are related by

$$\underline{\pi}_{j_1, j_2} = \pi_{j_2} K_{j_2, j_1} \tag{A.8}$$

Thus, by Lemmas 2 and 3, hypotheses **H1a** to **H3a** are verified.

Last, note that $\sum_{j_2 \in \mathcal{R}} \underline{f}^{\frac{1}{N}}(m, (j_1, j_2)) K_{j_1, j_2}^{\frac{1}{N}} = \epsilon(N) f^N(m, j_1)$ thus

$$\sum_{j_2 \in \mathcal{R}} \underline{f}(m, (j_1, j_2)) K_{j_1, j_2} = f(m, j_1)$$

and by Equation (A.8)

$$\underline{F}(m) = \vec{F}(m) \tag{A.9}$$

Theorem 1 then follows directly from Proposition 1.

A.3 Proof of Corollary 3

Since the ODE has a unique attractor \vec{m}^* to which all trajectories converge, the Birkhoff center is reduced to $\{\vec{m}^*\}$ and thus any limit point of ϖ^N is the Dirac mass at \vec{m}^* . Since Δ is compact, the set of probabilities on Δ is compact for the topology of weak convergence. ϖ^N has a unique limit point and is in a compact set, therefore it converges to the Dirac mass at \vec{m}^* . Since the distribution of $M^N(0)$ under \mathbb{P}_{η^N} is ϖ^N , $M^N(0) \rightarrow \vec{m}^*$ weakly (and in probability as the limit is constant). This shows item 1.

$X_1^N(0), \dots, X_N^N(0)$ is exchangeable and its occupancy measure converges weakly to a constant \vec{m}^* , therefore, by a theorem on exchangeable sequences in [16]:

$$\lim_{N \rightarrow \infty} \mathbb{P}_{\eta^N} \left(X_1^N(0) = i_1, \dots, X_k^N(0) = i_k \right) = m_{i_1}^* \dots m_{i_k}^*$$

Item 2 follows by observing that the distribution under \mathbb{P}_{η^N} are the same at time t and at time 0.

References

- [1] G. Sharma, A. Ganesh, and P. Key. Performance analysis of random access scheduling schemes. In *Proceedings of IEEE INFOCOM 2006, Barcelona, Spain*, April 2006.
- [2] M. Benaïm and J. Weibull. Deterministic approximation of stochastic evolution. *Econometrica*, 71:873–904, 2003.
- [3] Jean-Yves Le Boudec, David McDonald, and Jochen Munding. A Generic Mean Field Convergence Result for Systems of Interacting Objects. In *QEST'07*, 2007.
- [4] Charles Bordenave, David McDonald, and Alexandre Proutiere. Random multi-access algorithms, a mean field analysis. In *Proceedings of the Forty-Second Annual Allerton Conference on Communication, Control, and Computing, Allerton House, Monticello, Illinois, USA, University of Illinois at Urbana-Champaign*, October 2005.
- [5] Charles Bordenave, David McDonald, and Alexandre Proutiere. A particle system in interaction with a rapidly varying environment: Mean field limits and applications. arXiv:math/0701363v2.
- [6] A. Benveniste, M. Métivier, and P. Priouret. *Stochastic Approximation and Adaptive Algorithms*. Springer-Verlag, Berlin and New York, 1990.
- [7] M. Duflo. *Algorithmes Stochastiques*. Mathématiques et Applications. Springer-Verlag, 1996.
- [8] H. J. Kushner and G. G Yin. *Stochastic Approximation Algorithms and Applications*. Springer-Verlag, New York, 1997.
- [9] M. Benaïm. Recursive algorithms, urn processes and chaining number of chain recurrent sets. *Ergodic Theory and Dynamical System*, 18:53–87, 1998.
- [10] Bernard Ycart. *Modèles et algorithmes markoviens*, volume 39. Springer Verlag, 2002.
- [11] Rob De Boer. *Theoretical Fysiology*. Online lecture notes, 2006.
- [12] Michel Benaïm and J Weibull. Deterministic approximation of stochastic evolution in games - a generalization. <http://www2.hhs.se/personal/Weibull>, 2007.
- [13] A. Martinoli, K. Easton, and W. Agassounon. Modeling Swarm Robotic Systems: A Case Study in Collaborative Distributed Manipulation. *Int. Journal of Robotics Research*, 23(4):415–436, 2004. Special Issue on Experimental Robotics, B. Siciliano, editor.
- [14] David McDonald. Lecture Notes on Mean Field Convergence, March 2007.

- [15] S. Méléard. Asymptotic Behaviour of some interacting particle systems; McKean-Vlasov and Boltzmann models. In *CIME, May 1995*, Springer Verlag, 1996.
- [16] A.S. Sznitman. Topics in propagation of chaos. In P.L. Hennequin, editor, *Springer Verlag Lecture Notes in Mathematics 1464, Ecole d'Eté de Probabilités de Saint-Flour XI (1989)*, pages 165–251, 1991.
- [17] A. Kumar, E. Altman, D. Miorandi, and M. Goyal. New insights from a fixed-point analysis of single cell ieee 802.11 wlans. *IEEE/ACM Transactions on Networking*, 15:588–601, June 2007.
- [18] M Garetto, T. Salonidis, and E. Knightly. Modeling per-flow throughput and capturing starvation in csma multi-hop wireless networks. In *Proceedings of IEEE INFOCOM 2006, Barcelona, Spain*, April 2006.
- [19] Ruben Merz, Jean-Yves Le Boudec, and Saravanan Vijayakumaran. Effect on Network Performance of Common versus Private Acquisition Sequences for Impulse Radio UWB Networks. In *IEEE International Conference on Ultra-Wideband (ICUWB 2006)*, 2006.
- [20] G. Bianchi. IEEE 802.11 - Saturation Throughput Analysis. *IEEE Communications Letters*, 2:627–638, December 1998.
- [21] G Bianchi. Performance Analysis of the IEEE 802.11 distributed coordination function. *IEEE Journal on Selected Areas in Communications*, 18:535–547, March 2000.
- [22] M. Benaïm. Dynamics of stochastic approximations,. In *Le Séminaire de Probabilités XXXIII, Lectures Notes in Mathematics 1709*, pages 1–69. Springer-Verlag, 1999.