Mean-field analysis of ultra-dense CSMA networks

Mean-Field Analysis of Ultra-Dense CSMA Networks

F. Cecchi[a], S.C. Borst[a,b], J.S.H. van Leeuwaarden[a]
[a]Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands
[b]Bell Laboratories, Alcatel-Lucent, P.O. Box 636, Murray Hill, NJ 07974, USA
f.cecchi@tue.nl, s.c.borst@tue.nl, j.s.h.v.leeuwaarden@tue.nl

ABSTRACT

Distributed algorithms such as CSMA provide a popular mechanism for sharing the transmission medium among competing users in large-scale wireless networks. Conventional models for CSMA that are amenable for analysis assume that users always have packets to transmit. In contrast, when users do not compete for medium access when their buffers are empty, a complex interaction arises between the activity states and the buffer contents. We develop a mean-field approach to investigate this dynamic interaction for networks with many users. We identify a time-scale separation between the evolution of the activity states and the buffer contents, and obtain a deterministic dynamical system describing the network dynamics on a macroscopic scale. The fixed point of the dynamical system yields highly accurate approximations for the stationary distribution of the buffer contents and packet delay, even when the number of users is relatively moderate.

1. INTRODUCTION

Wireless networks are already ubiquitous today, and will grow even larger and denser in the future. Due to the size and the unregulated nature of these networks, their control by means of a centralized entity is impractical. These networks therefore deploy decentralized algorithms for regulating the medium access, such as the Carrier-Sense Multiple-Access (CSMA) protocol which is at the heart of current IEEE 802.11 networks. The analysis of random medium access algorithms has attracted a lot of research interest in recent years, and various models aimed at capturing the key performance measures have been proposed [2, 11, 12].

A common simplifying assumption for CSMA models is that devices always have packets to transmit. Without that assumption, packets arrive and form queues at the various nodes, so the buffer dynamics have to be taken into account. Since nodes without packets in the buffer withdraw from the medium access competition, a fundamental challenge is to understand the interaction between the activity states and the buffer contents of the various nodes. In fact, even the stability condition does not seem tractable for general interference topologies [5].

In the present paper, we analyze an idealized CSMA wireless network in an asymptotic regime where the number of nodes grows large, and show how in this setting the interaction between the activity states and the buffer contents can be captured analytically. This framework is known as the mean-field asymptotic regime and is popular in various scientific fields [1, 4]. Asymptotically, under certain scaling assumptions, each node is subject to an averaged global network influence and the specific pairwise interaction may be ignored, hence the term mean-field. A feature of mean-field analysis is the so-called *decoupling assumption* [1, 10]. The resulting approximation in a non-asymptotic regime often yields more tractable results and sharp estimates, see [2, 9]. In our model this property translates in the conservation of the mutual independence among the evolving buffer contents at the various nodes over any finite time interval.

Mean-field analysis of random medium access algorithms has not received much attention in the literature so far. Important exceptions are [3, 6] which consider a CSMA-like algorithm with a time-slotted operation and provide a deterministic differential equation. An approximate stability region for the system is identified in [3], while the decoupling assumption is investigated thoroughly in [6]. In contrast, the model we analyze evolves in an asynchronous manner and we will focus on the transient dynamics of the system and on the analysis of the buffer contents and delays in stationarity. To prove the asymptotic convergence towards a deterministic differential equation we leverage methodologies developed in [7, 8].

In Sect. 2 we present our model and the mean-field convergence result for a complete interference graph. In Sect. 3 we characterize the asymptotic limiting distribution of queue lengths and packet delays at individual nodes. We discuss extensions to arbitrary interference graphs in Sect. 4.

2. MODEL AND MEAN-FIELD LIMIT

To ease notation and illuminate the key concepts, we first focus on a full interference graph where all nodes are statistically identical and mutually interfere. We will briefly discuss the extension to more general interference topologies and multi-class settings in Sect. 4.

Consider a wireless network of N mutually interfering nodes sharing a single transmission channel. Packets arrive at each of the nodes according to independent Poisson processes of rate $\lambda^{(N)} = \lambda/N$. Once a node gains access to the medium, it transmits one packet, occupying the channel for an exponentially distributed time with parameter $\mu^{(N)} = \mu$. In between two consecutive transmissions a node must back-off for an exponentially distributed time period with parameter $\nu^{(N)} = \nu/N$. However, whenever another node is transmitting, the back-off period is frozen and resumed as soon as the channel is idle again.

Since the nodes are statistically identical, the state of the channel can be represented by $\Omega = \{0, 1\}$, where 1 refers to the state in which one of the nodes is transmitting. The dynamics of the system may be described by the Markov process $(\boldsymbol{Q}^{(N)}(t), Y^{(N)}(t))$ where $\boldsymbol{Q}^{(N)}(t) = (Q_i^{(N)}(t))_{i\in\{1,\ldots,N\}}$ defines the buffer process with $Q_i^{(N)}(t)$ the number of packets in the buffer of node i at time t excluding the one possibly being transmitted, and $Y^{(N)}(t) \in \Omega$ the channel activity state at time t.

The system may be equivalently described as a population continuous-time Markov process in a random environment $(\boldsymbol{X}^{(N)}(t), Y^{(N)}(t))$, where the population process $\boldsymbol{X}^{(N)}(t) = (X_k^{(N)}(t))_{k\in\mathbb{N}}$ is defined by

$$X_k^{(N)}(t) = N^{-1} \sum_{j=1}^{N} \mathbb{1}_{\{Q_j^{(N)}(t)=k\}}.$$

Observe that $\sum_k X_k^{(N)}(t) = 1$ for every $t \geq 0$. The possible transitions for the process $(\boldsymbol{X}^{(N)}(t), Y^{(N)}(t))$ are:

- At rate $\Lambda_k^{(N)} = \lambda X_k^{(N)}$, a packet arrives at a node having already k packets in its buffer, which generates the transition

$$(\boldsymbol{X}^{(N)}, Y^{(N)}) \to (\boldsymbol{X}^{(N)} + N^{-1}(\boldsymbol{e}_{k+1} - \boldsymbol{e}_k), Y^{(N)}).$$

- If the channel is active, i.e., $Y^{(N)} = 1$, the transmission is completed at rate $\mu^{(N)} = \mu$, which generates the transition

$$(\boldsymbol{X}^{(N)}, 1) \to (\boldsymbol{X}^{(N)}, 0).$$

- If the channel is inactive, i.e., $Y^{(N)} = 0$, a back-off period is completed by a node having $k > 0$ packets in the buffer at rate $V_k^{(N)} = \nu X_k^{(N)}$, which generates the transition

$$(\boldsymbol{X}^{(N)}, 0) \to (\boldsymbol{X}^{(N)} + N^{-1}(\boldsymbol{e}_{k-1} - \boldsymbol{e}_k), 1).$$

When the state of the system is $(\boldsymbol{X}^{(N)}, Y^{(N)})$, the population process is subject to a drift $F^{(N)}(\boldsymbol{X}^{(N)}, Y^{(N)})$ where

$$F^{(N)}(\boldsymbol{x}, y) = \frac{\lambda}{N} \sum_{k=0}^{\infty} (\boldsymbol{e}_{k+1} - \boldsymbol{e}_k)x_k + \mathbb{1}_{\{y=0\}} \frac{\nu}{N} \sum_{k=1}^{\infty} (\boldsymbol{e}_{k-1} - \boldsymbol{e}_k)x_k.$$

The channel activity process alternates between 0 and 1 at a rate of order 1. On the other hand, the population process changes at a slower pace. Indeed, order N transitions must occur, each taking an order 1 amount of time, before the population process experiences an order 1 change, i.e., the population process evolves at a rate of order N^{-1}. These two processes are deeply intertwined, but evolve on different time scales and this separation simplifies the analysis as N grows asymptotically large. In the limit, the activity process $Y^{(N)}(t)$ reaches its equilibrium distribution before the population process can change substantially, and influences the process $\boldsymbol{X}^{(N)}(t)$ only via its stationary distribution. In particular, given that the population process is in state $\boldsymbol{x}$, so that activations occur at rate $\nu(1-x_0)$ and deactivations take place at rate μ, the fraction of time that the channel is not active is $\pi_0(x_0) = \frac{\mu}{\mu+(1-x_0)\nu}$.

Theorem 1. *Assume $\boldsymbol{X}^{(N)}(0) \xrightarrow{N\to\infty} \boldsymbol{x}^\infty$. Then the sequence of processes $(\boldsymbol{X}^{(N)}(Nt))_{N\geq1}$ has a continuous limit $\boldsymbol{x}(t)$ which is the solution of the initial value problem*

$$\frac{d\boldsymbol{x}(t)}{dt} = F(\boldsymbol{x}(t)), \quad \boldsymbol{x}(0) = \boldsymbol{x}^\infty, \tag{1}$$

where the function $F(\cdot)$ is defined by

$$F(\boldsymbol{x}) = \lambda \sum_{k=0}^{\infty} (\boldsymbol{e}_{k+1} - \boldsymbol{e}_k)x_k + \pi_0(x_0)\nu \sum_{k=1}^{\infty} (\boldsymbol{e}_{k-1} - \boldsymbol{e}_k)x_k.$$

To prove Theorem 1, we exploited the methodologies developed in [7, 8]. We show, using a compactness argument, that the sequence of processes possesses a converging subsequence. Then, in order to prove that the converging subsequence satisfies the properties stated in Theorem 1, we apply martingale properties and the continuous mapping theorem.

In Figure 1 (i,ii) a random sample path of the population process obtained by simulations is compared with the numerical solution of (1). Observe that, even for relatively small values of N, $\boldsymbol{x}(t)$ clearly captures the fluid-level trend of $\boldsymbol{X}^{(N)}(Nt)$ although the diffusion-level variation is still quite evident. As N increases the variation decreases so that $\boldsymbol{x}(t)$ tracks $\boldsymbol{X}^{(N)}(Nt)$ more and more closely.

3. STATIONARY DISTRIBUTION

Introduce $\xi := \frac{\lambda}{\nu(1-\lambda/\mu)}$ and note that $\xi < 1$ is the asymptotic equivalent of the known stability condition for the complete interference graph [11].

Lemma 1. *If $\xi < 1$, then $\boldsymbol{x}^* = (x_k^*)_{k\in\mathbb{N}}$ where $x_k^* = (1-\xi)\xi^k$ is the unique equilibrium point for the dynamical system (1) in the set $H = \{\boldsymbol{x} = (x_k)_{k\in\mathbb{N}} : 0 \leq x_k \leq 1, \sum_k x_k = 1\}$.*

Assume the nodes to be exchangeable at the initial time, i.e., $\mathbb{P}\{Q_i^{(N)}(0) = k\} = x_k(0)$ for any tagged node i. Then, due to Sznitman [10],

$$\lim_{N\to\infty} \mathbb{P}\{Q_i^{(N)}(Nt) = k\} = x_k(t), \quad t \geq 0.$$

It may be shown, via a lenghty argument, that $\boldsymbol{x}^*$ is a globally stable equilibrium for the dynamical system (1) (details omitted because of page constraints), so that

$$\lim_{t\to\infty} \big(\lim_{N\to\infty} \mathbb{P}\{Q_i^{(N)}(t) = k\} \big) = x_k^*.$$

Observe now the process $Q_i^{(N)}(t)$ and define its stationary distribution $q_k^{(N)} := \lim_{t\to\infty} \mathbb{P}\{Q_i^{(N)}(t) = k\}$. Since the nodes are exchangeable, and assuming interchange of limits, we have that $q_k := \lim_{N\to\infty} q_k^{(N)} = x_k^*$. In Figure 1 (iii), we sampled via simulation $\boldsymbol{q}^{(N)} = (q_k^{(N)})_{k\geq0}$ for various values of N, and compared it with $\boldsymbol{x}^*$. Even for small values of N the difference is negligible in terms of euclidean distance, hence the approximation $q_k^{(N)} \approx x_k^*$ is highly accurate.

Intuitively, when N is sufficiently large, $Q_i^{(N)}(t)$ evolves as an $M/M(t)/1$ queue with constant arrival rate λ/N and time-varying service rate $\pi_0(x_0(t))\nu/N$. Once $\boldsymbol{x}(t)$ approaches the equilibrium point $\boldsymbol{x}^*$, the service rate stabilizes and $Q_i^{(N)}(t)$ evolves as a $M/M/1$ queue with utilization $\frac{\lambda}{\pi_0(x_0^*)\nu}$, which is equal to ξ because $\pi_0(x_0^*) = 1 - \lambda/\mu$. Since the stationary queue length distribution in an $M/M/1$ queue is geometric, we conclude that $q_k^{(N)} \approx (1-\xi)\xi^k$, which agrees with the equilibrium point of the mean-field limit. Defining $S^{(N)}$ as the sojourn time for a packet in stationarity, it further follows that $\mathbb{P}\{S^{(N)} > N\tau\} \to \exp(\nu(\lambda/\mu + \lambda/\nu - 1)\tau)$ as $N \to \infty$. Observe that the packet delay increases linearly in N although the queue length per node is order 1.

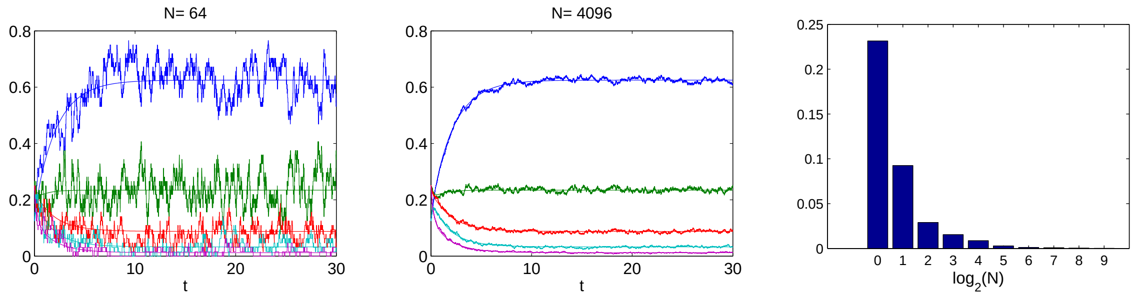

Figure 1: (i) and (ii) show a sample path of $\boldsymbol{X}_k^{(N)}(Nt)$ versus $\boldsymbol{x}_k(t)$, for $k = 0, \dots, 4$ from top to bottom. (iii) shows the euclidean norm between the stationary distribution $\boldsymbol{q}^{(N)}$ sampled from simulations and $\boldsymbol{x}^*$.

In [9], the authors assumed the probability for the buffer of a certain node to be empty to be independent of the activity state of the network. Even though this assumption was shown not to hold in general [5], it often yields an accurate approximation for the stability condition. Our analysis suggests that in the mean-field regime the assumption made in [9] is asymptotically exact.

4. EXTENSION AND CONCLUSION

The analysis can be extended to general interference graphs and multi-class settings, at the expense of heavier notation. Assume each of the nodes to belong to one of a finite set of classes $\mathcal{C}$, which may depend on the node location or its physical features. Given a class-wise interference graph, two nodes interfere with each other when they belong either to the same class or to two neighboring classes. Suppose for each class $p_c = \lim_{N\to\infty} N_c/N > 0$ where N_c is the number of nodes in class c. Packets arrive at each class-c node as a Poisson process of rate $\lambda_c^{(N)} = \lambda_c/N_c$, and transmissions and back-offs are completed at rates $\mu_c^{(N)} = \mu_c$ and $\nu_c^{(N)} = \nu_c/N_c$, respectively. Under these assumptions a similar time-scale separation emerges and we obtain the following generalization of Theorem 1.

THEOREM 2. *Assume $\boldsymbol{X}_c^{(N)}(0) \xrightarrow{N\to\infty} \boldsymbol{x}_c^\infty$ for every $c \in \mathcal{C}$. Then the sequence of processes $(\boldsymbol{X}_c^{(N)}(Nt))_{N\geq 1}$ has a continuous limit $\boldsymbol{x}_c(t)$ which is the solution of the initial value problem*

$$\frac{d\boldsymbol{x}_c(t)}{dt} = F_c(\boldsymbol{x}(t)), \quad \boldsymbol{x}(0) = \boldsymbol{x}^\infty, \tag{2}$$

where the function $F_c(\cdot)$ is defined by

$$\begin{aligned} F_c(\boldsymbol{x}) =& \frac{1}{p_c}(\lambda_c \sum_{k=0}^{\infty} (\boldsymbol{e}_{c,k+1} - \boldsymbol{e}_{c,k}) x_{c,k} \\ &+ \pi_{c,0}(\boldsymbol{x})\nu_c \sum_{k=1}^{\infty} (\boldsymbol{e}_{c,k-1} - \boldsymbol{e}_{c,k}) x_{c,k}) \end{aligned}$$

and $\pi_{c,0}(\boldsymbol{x}) = \frac{G(\Omega_c)}{G(\Omega)}$, where

$$G(\Omega') = \sum_{\omega\in\Omega'} \prod_{c'\in\omega} (1 - x_{c',0}) \frac{\nu_{c'}}{\mu_{c'}},$$

Ω is the set of feasible class-activity states and $\Omega_c \subseteq \Omega$ is the subset of states which allow nodes in class c to back-off.

The formula for $\pi_{c,0}(\boldsymbol{x})$ in Theorem 2 is a generalization of that for $\pi_0(x_0)$ in Theorem 1, and arises from the product-form distribution for the activity process of a network with saturated buffers [11], and activation rates reduced by a factor $(1 - x_{c,0})$. Similarly to Sect. 3, the equilibrium point of equation (2) provides accurate approximations for the stationary distribution of queue lengths and delays at individual nodes.

5. REFERENCES

[1] M. Benaïm, J-Y. Le Boudec (2008). A class of mean field interaction models for computer and communication systems. *Perf. Eval.* **65 (11)**, 823–838.

[2] G. Bianchi (2000). Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE J. Sel. Areas Commun.* **18 (3)**, 535–547.

[3] C. Bordenave, D. McDonald, A. Proutière (2008). Performance of random medium access control, an asymptotic approach. *ACM SIGMETRICS Perf. Eval. Rev.* **36 (1)**, 1–12.

[4] L. Bortolussi, J. Hillston, D. Latella, M. Massink (2013). Continuous approximation of collective system behaviour: A tutorial. *Perf. Eval.* **70 (5)**, 317–349.

[5] F. Cecchi, S.C. Borst, J.S.H. van Leeuwaarden (2014). Throughput of CSMA networks with buffer dynamics. *Perf. Eval.* **79**, 216–234.

[6] J. Cho, J-Y. Le Boudec, Y. Jiang (2012). On the asymptotic validity of the decoupling assumption for analyzing 802.11 MAC protocol. *IEEE Trans. Inf. Th.* **58 (11)**, 6879–6893.

[7] S.N. Ethier, T.G. Kurtz (1986). *Markov Processes: Characterization and Convergence.* Wiley.

[8] P.J. Hunt, T.G. Kurtz (1994). Large loss networks. *Stoc. Proc. Appl.* **53 (2)**, 363–378.

[9] R. Laufer, L. Kleinrock (2013). The capacity of wireless CSMA/CA multi-hop networks. In: *Proc. IEEE Infocom 2013.*

[10] A.S. Sznitman (1991). Topics in propagation of chaos. In: *Ecole Et. Prob. Saint-Flour XIX - Lect. Not. Math.* 1464, 165–251.

[11] P.M. van de Ven, S.C. Borst, J.S.H. van Leeuwaarden, A. Proutière (2010). Insensitivity and stability of random-access networks. *Perf. Eval.* **67 (11)**, 1230–1242.

[12] X. Wang, K. Kar (2005). Throughput modeling and fairness issues in CSMA/CA based ad-hoc networks. In: *Proc. IEEE Infocom 2005.*