Spatial mean-field limits for CSMA networks

Citation for published version (APA):

Cecchi, F., Borst, S. C., van Leeuwen, J. S. H., & Whiting, P. A. (2018). Spatial mean-field limits for CSMA networks. In *55th Annual Allerton Conference on Communication, Control, and Computing, Allerton 2017* (Vol. 2018-January, pp. 502-509). Institute of Electrical and Electronics Engineers. <https://doi.org/10.1109/ALLERTON.2017.8262779>

DOI:

[10.1109/ALLERTON.2017.8262779](https://doi.org/10.1109/ALLERTON.2017.8262779)

Document status and date:

Published: 17/01/2018

Document Version:

Accepted manuscript including changes made at the peer-review stage

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Spatial Mean-Field Limits for CSMA Networks

Fabio Cecchi

Eindhoven University of Technology
f.cecchi@tue.nl

Johan S.H. van Leeuwen
Eindhoven University of Technology
j.s.h.v.leeuwen@tue.nl

Sem C. Borst

Eindhoven University of Technology - Nokia Bell Labs
s.c.borst@tue.nl

Philip A. Whiting

Macquarie University
philip.whiting@mq.edu.au

Abstract—Random-access algorithms such as the CSMA protocol provide a popular mechanism for distributed medium access control in large-scale wireless networks. Mean-field analysis has emerged as a convenient approach to obtain tractable performance estimates in such networks, but a critical limitation of the classical set-up is that all nodes are assumed to belong to a finite number of classes. We consider spatial mean-field limits which do not involve such a requirement, characterized in terms of a set of partial-differential equations, and in particular examine the fixed points of these equations for some specific network configurations. We discuss how the fixed points can be used to obtain estimates for key performance metrics, and present simulation experiments to demonstrate the accuracy of these estimates.

Index Terms—CSMA, Random-access networks, Mean-field analysis, Measure-valued Markov processes

I. INTRODUCTION

The use of wireless communications has experienced tremendous growth over the last two decades, which is widely predicted to continue, driven by a proliferation of low-cost sensors and machine-type devices, commonly referred to as the Internet-of-Things (IoT). Forecasts indicate that the number of IoT nodes will reach into the tens of billions by 2020, and outgrow the number of human-operated devices by an order-of-magnitude [1], [2]. With such a massive number of nodes, each of which individually may only be sporadically active, any form of dedicated spectrum allocation or scheduled medium access is impractical. Instead these large-scale networks will typically rely on the individual nodes to dynamically share the medium in a distributed fashion.

A popular mechanism for distributed medium access control is provided by the Carrier-Sense Multiple-Access (CSMA) protocol. As it turns out, in saturated-buffer scenarios, the joint activity process in CSMA networks has a product-form stationary distribution [14]. However, these results do not capture the relevant performance metrics in unsaturated-buffer scenarios, which in particular arise in an IoT context with highly intermittent traffic sources. In such situations, buffers will frequently be empty, and nodes will refrain from competition for the medium during these periods.

Mean-field analysis has emerged as a powerful approach to obtain tractable performance estimates in random-access

networks. Mean-field concepts were already leveraged to derive throughput estimates in saturated-buffer scenarios in the seminal paper [3], with further results in [10], [11]. Mean-field techniques have also proved useful in examining stability issues [5] and obtaining expressions for queue length distributions and delay metrics in unsaturated-buffer scenarios [6], [8]. Besides mathematically convenient, the mean-field regime is also highly relevant in view of the increasingly large number of nodes and dense deployments with the emergence of IoT applications mentioned earlier.

A crucial requirement for the classical mean-field framework to apply however is that the node population can be partitioned into a finite number of subsets of statistically indistinguishable nodes. The latter condition is a severe restriction since nodes typically have different locations, and hence are subject to different interference constraints. The mean-field regime of systems with spatial features has been investigated in [9], [16] where approximations are obtained by treating nodes in close vicinity as being identical. In [7], a novel mean-field methodology was developed which does not rely on any exchangeability property of nodes. There, the number of nodes within a given space grows large while the interference range remains fixed. Thus the network becomes dense, in the sense that the number of interferers for most nodes also grows large, but no two nodes are required to be similar.

In this paper we further investigate the initial-value problem (IVP) derived in [7] whose unique solution determines the behavior of the network in a mean-field regime. We focus on the characterization of the fixed point of the IVP in the scenario with a uniform distribution of nodes. The unique fixed point is leveraged to obtain approximations for key performance metrics such as stationary buffer content, packet delay, and loss rate. We also provide an easily verifiable criterion to establish whether the network is below critical load, in the sense that the loss rate vanishes as the buffer capacity of the nodes grows large.

The remainder of the paper is organized as follows. In Section II we present a detailed model description. In Section IV we reproduce the mean-field limit derived in [7] and provide an implicit characterization of the fixed point. We then focus

on the scenario with a uniform node distribution, prove existence and uniqueness of the fixed point, and show how to leverage it to obtain approximations for the performance metrics in Section V. In Section VI we provide a criterion to establish whether a network is below critical load. We present numerical experiments in Section VII indicating that the fixed point yields remarkably accurate performance approximations.

II. MODEL DESCRIPTION

Consider N nodes sharing a wireless medium and located in the circle of unit radius S^1 . Locations are identified by values in $[0, 1)$ and S^1 is equipped with distance $d(\cdot)$, where the distance between $z_1, z_2 \in S^1$ is given by the minimum between the length of the circular arc spanning from z_1 to z_2 and from z_2 to z_1 . Denote by $\mathcal{B}$ the Borel sets of S^1 .

Node n is located in position $s_n^{(N)}$ and the collective spatial configuration of the nodes is given by the empirical measure

$$P^{(N)}(B) = \frac{1}{N} \sum_{n=1}^N \mathbb{1}\{s_n^{(N)} \in B\}, \quad B \in \mathcal{B}.$$

Here and throughout the paper, the superscript (N) refers to the N -th system in a sequence of systems that we will consider with an increasing number of nodes.

The interference relations are described by the distance function $d(\cdot, \cdot)$ and a reuse distance $r \in \mathbb{R}_+$, i.e., nodes n, n' interfere with each other if and only if $d(s_n^{(N)}, s_{n'}^{(N)}) \leq r$. Denote by $\Omega^{(N)} \in \{0, 1\}^N$ the set of feasible activity states, where an activity state ω is feasible if $d(s_n^{(N)}, s_{n'}^{(N)}) > r$ for every n, n' such that $\omega_n = \omega_{n'} = 1$. Observe that a feasible active state can involve at most $K = \lfloor 1/r \rfloor < \infty$ active nodes.

Each node has a buffer with a finite capacity M , and packets arrive to the various nodes according to independent Poisson processes of rate λ/N . Packets arriving to an already full buffer are lost. When a node accesses the medium, it transmits a single packet from its buffer which takes an exponentially distributed time with mean $1/\mu$. In between two consecutive transmissions, a node must obey an exponentially distributed back-off period with parameter ν/N . The back-off period is interrupted whenever an interfering node is active and resumed when the medium is sensed free, so that an activation of the node would lead to a feasible activity state.

The system behavior is described by the queue length process $\mathbf{Q}^{(N)}(t) = (Q_1^{(N)}(t), \dots, Q_N^{(N)}(t))$ and the activity process $\mathbf{A}^{(N)}(t) \in \{0, 1\}^N$, where $Q_n^{(N)}(t) \in \{0, \dots, M\}$ denotes the number of packets in the buffer of node n at time t (excluding the one possibly in transmission), and $A_n^{(N)}(t)$ indicates whether node n is active at time t or not.

A. Measure-valued Markov state description

The dimension of the queue length vector and activity state vector in the above Markovian representation grows without bound as the number of nodes grows large. In contrast to the usual mean-field set-up, we cannot use an equivalent finite-dimensional population process since all the nodes may

have different locations. Hence we adopt a Measure-Valued Markov Process (MVMP) description, see for instance [13], defined on the metric space (S^1, d) . Specifically, in the N -th network, for each $m = 0, \dots, M$, we obtain a subprobability measure $\mathbf{V}_m^{(N)} \in \mathcal{M}(S^1) =: \mathcal{M}$ on the Borel sets of S^1 by defining

$$V_m^{(N)}(B, t) := \frac{1}{N} \sum_{n: s_n^{(N)} \in B} \mathbb{1}\{Q_n^{(N)}(t) = m\},$$

representing the fraction of nodes which are located in the set B and have m packets in their buffer at time t . Note that by construction

$$\sum_{m=0}^M V_m^{(N)}(B, t) = P^{(N)}(B), \quad \forall t \geq 0. \quad (1)$$

Thus the state of all the queues is determined by the $(M+1)$ -dimensional vector of measures

$$\mathbf{V}^{(N)}(\cdot, t) := (\mathbf{V}_0^{(N)}(\cdot, t), \dots, \mathbf{V}_M^{(N)}(\cdot, t)) \in \mathcal{M}^{M+1}.$$

We include a location-oriented representation of the activity process $\mathbf{Y}^{(N)}(t)$, where

$$\mathbf{Y}^{(N)}(t) = \{s_n^{(N)} : A_n^{(N)}(t) = 1\} \quad (2)$$

describes the locations of the nodes that are active at time t . Thus we obtain the state of the network as

$$(\mathbf{V}^{(N)}(t), \mathbf{Y}^{(N)}(t)) \in \mathcal{M}^{M+1} \times \Upsilon_d,$$

where $\Upsilon_d = \bigcup_{k=0}^K \Upsilon_d^{(k)}$ and $\mathbf{z} = (z_1, \dots, z_k) \in \Upsilon_d^{(k)}$ if and only if $d(z_n, z_{n'}) > r$ for every $n, n' = 1, \dots, k$.

III. PRELIMINARY RESULTS FOR FIXED SPATIAL ACTIVATION MEASURE

For later purposes, it is convenient to also consider a related scenario where activations may occur at a continuum of locations. Specifically, activations occur at a location in $B \in \mathcal{B}$ at fixed rate $\nu\eta(B)$, where the *spatial activation measure* $\eta \in \mathcal{M}$ is absolutely continuous with respect to the Lebesgue measure, and activity persists for an exponentially distributed time with mean $1/\mu$.

The activity process $(\mathbf{Y}(t))_{t \geq 0}$ takes values in Υ_d . The stationary distribution of $(\mathbf{Y}(t))_{t \geq 0}$ is obtained in [12], and may be expressed for any $B \subseteq \Upsilon_d^{(k)}$ as

$$\pi(B; \eta) = \frac{\frac{\sigma^k}{k!} Z_k(B; \eta)}{\sum_{l=0}^K \frac{\sigma^l}{l!} Z_l(\eta)}, \quad (3)$$

with

$$Z_k(B; \eta) = \int_{\mathbf{z} \in B} \prod_{i=1}^k d\eta(z_i), \quad (4)$$

$$Z_l(\eta) = Z_l(\Upsilon_d^{(l)}; \eta) = \int_{\mathbf{z} \in \Upsilon_d^{(l)}} \prod_{i=1}^l d\eta(z_i), \quad (5)$$

$\sigma = \nu/\mu$ and $K = \lfloor 1/r \rfloor$ as before.

For compactness, denote by $B_z = \{z' \in S^1 : d(z, z') \leq r\}$ the interference interval at $z \in S^1$ and $\mathcal{U}_z = S^1 \setminus B_z$. Then

$\psi_k(z; \eta) = \pi(\mathcal{U}_z^k \cap \Upsilon_d^{(k)}; \eta)$ is the stationary fraction of time that activity occurs at k locations but not at any location within distance r from z , and

$$\psi(z; \eta) = \sum_{k=0}^{K-1} \psi_k(z; \eta) \quad (6)$$

is the stationary fraction of time that no location in the interference interval at z is active. Because of the PASTA property, $\psi(z; \eta)$ equals the stationary probability that an activation attempt at location z is successful.

A. Uniform spatial activation measure

In general, the coefficients $Z_k(B; \eta)$ and $Z_l(\eta)$ cannot be expressed in closed form, but when $B \subseteq \Upsilon_d^{(k)}$ is of the form $\mathcal{U}_z^k$ and the activation measure η is uniform, i.e., of the form $\eta = \vartheta \mathcal{L}$, with $\mathcal{L}$ denoting the Lebesgue measure and $\vartheta \geq 0$, the involved integral terms can be calculated, yielding explicit expressions for the above stationary probabilities $\psi_k(z; \vartheta \mathcal{L})$ as stated in the next proposition.

Proposition 1. *It holds that*

$$Z_k(\mathcal{U}_z^k; \vartheta \mathcal{L}) = \vartheta^k (1 - (k+1)r)^k, \quad (7)$$

for $k = 0, \dots, K-1$, and

$$Z_0(\vartheta \mathcal{L}) = 1, \quad Z_l(\vartheta \mathcal{L}) = \vartheta^l (1 - lr)^{l-1}, \quad (8)$$

for $l = 1, \dots, K$. Thus $\psi(z; \vartheta \mathcal{L}) = \sum_{k=0}^{K-1} \psi_k(z; \vartheta \mathcal{L})$ with

$$\psi_k(z; \vartheta \mathcal{L}) = \frac{(\vartheta \sigma)^k (1 - (k+1)r)^k}{1 + \sum_{l=1}^K \frac{(\vartheta \sigma)^l (1 - lr)^{l-1}}{l!}}. \quad (9)$$

Proof. The proof of this proposition is obtained by iteration of integrals and is based on the expressions in (4) and (5). $\square$

Note that $\psi(z; \vartheta \mathcal{L})$ is independent of z . We define $\phi(\vartheta) = \psi(0; \vartheta \mathcal{L})$ as the stationary probability of a successful activation attempt, and thus $\vartheta \nu \phi(\vartheta)$ represents the aggregate effective activation rate and $\vartheta \sigma \phi(\vartheta)$ represents the average number of active locations. This suggests that $\phi(\vartheta)$ and $\vartheta \phi(\vartheta)$ should be decreasing and increasing as function of the spatial activation intensity ϑ . This is indeed the case, as stated in the next lemma, which may be of independent interest.

Lemma 1. *The function $\phi(\vartheta)$ is strictly decreasing and $\vartheta \phi(\vartheta)$ is strictly increasing in $\vartheta \geq 0$.*

In order to prove the above lemma, observe that $\phi(\vartheta)$ may be written as

$$\phi(\vartheta) = \sum_{k=0}^K f(k) p_k(\vartheta),$$

where

$$p_k(\vartheta) = \frac{\frac{\sigma^k}{k!} Z_k(\vartheta \mathcal{L})}{\sum_{l=0}^K \frac{\sigma^l}{l!} Z_l(\vartheta \mathcal{L})}$$

is the stationary fraction of time that k locations are active and

$$f(k) = \frac{Z_k(\mathcal{U}_z^k; \vartheta \mathcal{L})}{Z_k(\vartheta \mathcal{L})} = \frac{(1 - (k+1)r)^k}{(1 - kr)^{k-1}}.$$

Further note that $p_k(\vartheta)$ may also be interpreted as the stationary distribution of a random variable H_ϑ governed by a birth-death process with birth rate $a_k(\vartheta) = \vartheta \nu f(k)$, $k = 0, \dots, K-1$ and death rate $d_k = k\mu$, $k = 1, \dots, K$, in state k , respectively.

Hence,

$$\phi(\vartheta) = \mathbb{E}[f(H_\vartheta)].$$

Also,

$$\sigma \vartheta \phi(\vartheta) = \sum_{k=0}^K k p_k(\vartheta) = \mathbb{E}[H_\vartheta]. \quad (10)$$

Since the birth rates are all increasing in ϑ and the death rates do not depend on ϑ , it follows that the stationary distribution of the random variable H_ϑ is stochastically increasing in ϑ , and in particular $\vartheta \phi(\vartheta)$ is increasing in ϑ due to (10). In order to prove Lemma 1, it thus suffices to show that the function $f(k)$ is decreasing in k . This property is established in the Appendix, and can also be used for an algebraic proof of Lemma 1 based on [15].

IV. MEAN-FIELD LIMIT

In this section we reproduce the main result of [7], which characterizes the behavior of the network in the asymptotic regime as $N \rightarrow \infty$, i.e., as the number of nodes grows large. We assume that the empirical measure of the node locations $P^{(N)}$ weakly converges to a measure $P \in \mathcal{P} = \{\mu \in \mathcal{M} : \mu(S^1) = 1\}$ that is absolutely continuous with respect to the Lebesgue measure and has density bounded away from 0 and ∞ . This yields that $P^{(N)}(A) \rightarrow P(A)$ for every $A \in \mathcal{R} \subseteq \mathcal{B}$, with $\mathcal{R}$ being the algebra generated by the P -continuity sets (which includes the open intervals). Observe that a finite and countably additive map defined on $\mathcal{R}$ has a unique extension to a measure on $\mathcal{B}$ since $\sigma(\mathcal{R}) = \mathcal{B}$. In the limit, we therefore assume a true continuum of node locations and no significant accumulation of nodes in any location.

Let us consider $\bar{\mathbf{V}}^{(N)}(t) = \mathbf{V}^{(N)}(Nt)$ and assume $\mathbf{V}^{(N)}(0) \Rightarrow \mathbf{V}^\infty \in \mathcal{M}^{M+1}$ with $\mathbf{V}_m^\infty$ absolutely continuous with density $\tilde{v}_m^\infty(z)$ with respect to P for every $m = 0, \dots, M$. In [7, Thm. 3.1] it has been shown that the process $\bar{\mathbf{V}}^{(N)}(\cdot)$ weakly converges to a limiting process $\mathbf{V}(\cdot)$ as $N \rightarrow \infty$, which is deterministic and fully determined by the initial conditions $\mathbf{V}^\infty \in \mathcal{M}^{M+1}$ and the measure P . Specifically,

$$\mathbf{V}_m(t, A) = \int_A \tilde{v}_m(z, t) dP(z), \quad (11)$$

for every $t \geq 0$ and $A \in \mathcal{R}$. The density $\tilde{v}(z, t)$ is a solution of the following initial-value problem for $t \in [0, T]$ and $z \in S^1$,

$$\frac{\partial}{\partial t} \tilde{v}(t, z) = \tilde{\mathbf{H}}(\tilde{v}(t, z)), \quad \tilde{v}(0, z) = \tilde{v}^\infty(z), \quad (12)$$

where $\tilde{H}(\cdot, \cdot) = (\tilde{H}_0(\cdot, \cdot), \dots, \tilde{H}_M(\cdot, \cdot))$ is defined by

$$\begin{aligned}\tilde{H}_0(\tilde{\mathbf{v}}, z) &= -\lambda \tilde{v}_0(z) + \nu \tilde{v}_1(z) \psi(z; P \circ \mathbf{V}_{>0}), \\ \tilde{H}_M(\tilde{\mathbf{v}}, z) &= \lambda \tilde{v}_{M-1}(z) - \nu \tilde{v}_M(z) \psi(z; P \circ \mathbf{V}_{>0}), \\ \tilde{H}_m(\tilde{\mathbf{v}}, z) &= \lambda(\tilde{v}_{m-1}(z) - \tilde{v}_m(z)) \\ &\quad - \nu(\tilde{v}_m(z) - \tilde{v}_{m+1}(z)) \psi(z; P \circ \mathbf{V}_{>0}),\end{aligned}$$

for $m = 1, \dots, M-1$, where $\psi(\cdot; \cdot)$ is defined in (6) and

$$P \circ \mathbf{V}_{>0}(A) = \int_A (1 - \tilde{v}_0(z)) dP(z), \quad A \in \mathcal{R}.$$

A. Fixed-point characterization

In this section, we aim to describe the fixed points of the mean-field partial differential equation in (12). Note that such a fixed point would describe an invariant measure for the network in the mean-field regime, and we will leverage the latter to obtain approximations for key performance measures. Before doing that, the next proposition implicitly characterizes the fixed points in general configurations. In Section V we will then specialize to a specific instance so as to allow tractability.

Proposition 2. *A measure $\mathbf{V}$ is a fixed point for (12) if and only if $\tilde{\mathbf{v}}(\cdot)$ solves*

$$\begin{aligned}\tilde{v}_0(z) &= \frac{1}{\sum_{m=0}^M \left(\frac{\lambda}{\nu \psi(z; P \circ \mathbf{V}_{>0})} \right)^m}, \\ \tilde{v}_m(z) &= \left(\frac{\lambda}{\nu \psi(z; P \circ \mathbf{V}_{>0})} \right)^m \tilde{v}_0(z),\end{aligned}$$

for every $m = 1, \dots, M$ and $z \in S^1$.

Proof. Assume that $\mathbf{V}$ is a fixed point for (12) and consider $z \in S^1$. The dependence of $\tilde{v}_m(z)$ on $\tilde{v}_0(z)$ immediately follows from $\tilde{H}_m(\tilde{\mathbf{v}}(t), z) = 0$ for $m = 0, \dots, M$. In fact,

$$\lambda \tilde{v}_{m-1}(z) = \nu \tilde{v}_m(z) \psi(z; P \circ \mathbf{V}_{>0}), \quad m = 1, \dots, M.$$

The equation for $\tilde{v}_0(z)$ follows from

$$\sum_{m=0}^M \tilde{v}_m(z) = 1, \quad z \in S^1, \quad (13)$$

which is implied by (1). In particular, since $P^{(N)} \Rightarrow P$ we have that $P(A) = \sum_{m=0}^M \mathbf{V}_m(A)$ for every $A \in \mathcal{R}$. The measures $P(\cdot)$ and $\sum_{m=0}^M \mathbf{V}_m(\cdot)$ coincide on the π -system $\mathcal{R}$ generating the σ -algebra $\mathcal{B}$ and thus extend to the same measure on S^1 . From (11), we know that $\sum_{m=0}^M \mathbf{V}_m(\cdot)$ has Radon-Nikodym derivative $\sum_{m=0}^M \tilde{v}_m(z)$ with respect to P , hence (13) holds. The other direction can be verified by substituting the proposed solution $\mathbf{V}$ into (12). $\square$

A fixed point of (12) is thus entirely determined by a solution of the equations in Proposition 2. Identifying such a solution is not trivial in general. In the next section we focus on the special case where nodes are uniformly spaced in S^1 and we describe how the fixed point can be used to approximate some relevant performance metrics.

V. UNIFORM CIRCLE

In this section we focus on the special case where node n is located at position $s_n^{(N)} = \frac{n-1}{N}$ in the system with N nodes, i.e., nodes are equi-distantly spaced. Observe that $P^{(N)} \Rightarrow \mathcal{L}$ as N tends to infinity. Since all the nodes have identical statistical features and the same number of interferers, we are interested in fixed points given by ‘symmetric’ measures, where a measure $\mathbf{V} \in \mathcal{M}^{M+1}$ is symmetric if there exists $b_0, \dots, b_M \in [0, 1]$ such that

$$\mathbf{V}_m = b_m \mathcal{L}, \quad \sum_{m=0}^M b_m = 1,$$

i.e., $\tilde{v}_m(z) = b_m$ for every $z \in S^1$ and $m = 0, \dots, M$.

Proposition 3. *A symmetric measure $\mathbf{V}$ is a fixed point for (12) if and only if $b_0, \dots, b_M$ solve*

$$b_m = \left(\frac{\lambda}{\nu \phi(1 - b_0)} \right)^m b_0, \quad b_0 = \frac{1}{\sum_{m=0}^M \left(\frac{\lambda}{\nu \phi(1 - b_0)} \right)^m}. \quad (14)$$

Proof. The proof follows from Proposition 2 by showing that

$$\psi(z; P \circ \mathbf{V}_{>0}) = \psi(0; (1 - b_0) \mathcal{L}).$$

By hypothesis, we have that $P \circ \mathbf{V}_{>0} = (1 - b_0) \mathcal{L}$, and we immediately conclude by noting that $\psi(z; \vartheta \mathcal{L})$ is independent of z for any $\vartheta \geq 0$ due to Proposition 1. $\square$

The following theorem ensures that a solution of (14) exists and is unique.

Theorem 1. *There exists a unique solution $b_0^*, b_1^*, \dots, b_M^*$ to (14).*

Proof. Due to the dependence of $b_1^*, \dots, b_M^*$ on b_0^* , it is sufficient to show that there exists a unique b_0^* such that

$$F(b_0^*) = 1, \quad F(b_0) := b_0 \sum_{m=0}^M \left(\frac{\lambda}{\nu \phi(1 - b_0)} \right)^m. \quad (15)$$

Note that F is a continuous function where $F(0) = 0$ and

$$F(1) = 1 + \sum_{m=1}^M \left(\frac{\lambda}{\nu \phi(0)} \right)^m > 1,$$

thus a solution necessarily exists.

So as to show uniqueness, observe that $\tilde{H}_m(\tilde{\mathbf{v}}(t), z) = 0$ for $m = 0, \dots, M$ yields that $b_0, \dots, b_M$ have to solve

$$\lambda b_{m-1} = \nu b_m \phi(1 - b_0), \quad m = 1, \dots, M,$$

and by summing in m , we obtain that

$$\lambda(1 - b_M) = \nu(1 - b_0)\phi(1 - b_0). \quad (16)$$

By taking $\xi = 1 - b_0$, the right-hand side is equivalent to $\nu \xi \phi(\xi)$ and thus increasing in ξ due to Lemma 1. The left-hand side is equivalent to

$$\lambda \frac{\sum_{m=0}^{M-1} \left(\frac{\lambda}{\nu \phi(\xi)} \right)^m}{\sum_{m=0}^M \left(\frac{\lambda}{\nu \phi(\xi)} \right)^m},$$

i.e., it corresponds to λ times the acceptance probability (one minus the loss probability) in an M/M/1/M queue with arrival rate λ and service rate $\nu\phi(\xi)$ and is thus increasing in $\phi(\xi)$ as can also be shown algebraically. Thus, the left-hand side of Equation (16) is increasing in $\phi(\xi)$ as long as $\phi(\xi) \geq 0$, and therefore, due to Lemma 1, decreasing in ξ . Hence, there can be at most one solution $b_0^* = 1 - \xi^*$, $b_1^*, \dots, b_M^*$ of (16). $\square$

The argument used in the proof of the above proposition also provides an efficient way to approximate b_0^* . Observe that $F(b)$ is computable and thus we may apply the bisection algorithm to the function $F(b) = 1$ over the interval $[0, 1]$. We can then obtain the other coefficients $b_1^*, \dots, b_M^*$ by means of (14).

A. Performance analysis

In this section we detail in which way the fixed point characterized in Proposition 3, whose uniqueness and existence is guaranteed by Theorem 1, can be used to obtain approximations for key performance metrics. In particular, we derive approximations for the limiting queue length distribution, average packet delay, and loss rate, which are based on the symmetric fixed point described by the coefficients $b_0^*, \dots, b_M^*$.

Queue length. The empirical queue length distribution of node n is given by

$$\hat{Q}_{n,m}^{(N)} := \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \mathbb{1}\{Q_n^{(N)}(t) = m\} dt.$$

Due to the symmetry of the various nodes, the distribution of $\hat{Q}_{n,m}^{(N)}$ is independent of n . It is natural to use b_m^* as an approximation for $\hat{Q}_{n,m}^{(N)}$ for large N . The simulations presented in Section VII demonstrate that this approximation is highly accurate even for moderate values of N .

Loss rate. Since the arrival process is Poisson, the probability that a packet arrives to an already full buffer corresponds to the stationary probability that there are M packets in the buffer. Denote by $\hat{B}_n^{(N)}(T)$ the empirical fraction of lost packets for node n by time T , and define $\hat{B}_n^{(N)} := \lim_{T \rightarrow \infty} \hat{B}_n^{(N)}(T)$. Due to the symmetry of the various nodes, the distribution of $\hat{B}_n^{(N)}(T)$ is independent of n . Then it is natural to approximate $\hat{B}_n^{(N)}$ by means of b_M^* , the approximation provided above for $\hat{Q}_M^{(N)}$.

Packet delay. Denote by $\hat{w}_{n,j}$ the time spent in the buffer by the j -th packet transmitted by node n . Consider

$$\hat{W}_n^{(N)}(T) = \frac{1}{T_n^{(N)}(T)} \sum_{j=1}^{T_n^{(N)}(T)} \frac{\hat{w}_{n,j}}{N},$$

and define $\hat{W}_n^{(N)} := \lim_{T \rightarrow \infty} \hat{W}_n^{(N)}(T)$, where $T_n^{(N)}(T)$ is the number of packets transmitted by node n before time T . Due to the symmetry of the various nodes, the distribution of $\hat{W}_n^{(N)}(T)$ is independent of n . By means of the queue length

Fig. 1. (a) Queue length distribution, (b) loss rate and average normalized packet delay, as a function of location.

distribution discussed above and Little's law, it is natural to use

$$w^* := \frac{\sum_{m=1}^M m b_m^*}{\lambda(1 - b_M^*)}$$

as approximation for the average normalized packet delay $\hat{W}^{(N)}$ for large N .

Note that when $r > 1/2$, i.e., $K = 1$, the results in [6] imply that the above-described approximations are indeed exact in the mean-field regime. However, proving the asymptotic exactness of the approximations in general remains as a major challenge, which is beyond the scope of this paper.

Illustrative example. We now pause and consider a simple example so as to illustrate the approximations introduced above. Consider a scenario with $M = 1$ and $r = 0.35$ so that $K = 2$. We further assume that all the buffers are initially empty. We stress that the nodes are not exchangeable, since their locations make them unique. Hence, even this simple case cannot be tackled directly via classical mean-field techniques. We set $\lambda = 0.5$, $\mu = 1$, and $\nu = 2$. Via the bisection algorithm, we obtain that

$$b_0^* \approx 0.749, \quad b_1^* \approx 0.251, \quad \tilde{w}^* \approx 0.67.$$

We simulate a network with $N = 50$ nodes, up to times $t_1 = 10^5 N$ and $t_2 = 500 N$ and for every node n we keep track of $\hat{Q}_{n,0}^{(N)}(t)$, $\hat{Q}_{n,1}^{(N)}(t)$, $\hat{B}_n^{(N)}(t)$, and $\hat{W}_n^{(N)}(t)$. In Figure 1 we show that, as time grows large, the mean-field approximation introduced above becomes extremely accurate, even though $N = 50$ is just a moderate value.

VI. CRITICAL LOAD THRESHOLD

In Section V we proved that there exists a unique fixed point for the initial-value problem in (12) when $P \Rightarrow \mathcal{L}$. The fixed point is then used to obtain approximations for the performance metrics of the system. In this section we leverage the characterization derived so as to determine whether the loss rate vanishes as the buffer capacity increases, in which case the network is said to be *below critical load*.

In particular, we observed that the fraction of lost packets is approximated by b_M^* and we now derive conditions under which $\lim_{M \rightarrow \infty} b_M^* = 0$. Intuitively, the latter condition corresponds to the stability condition for a large network with unlimited buffer capacity.

Definition 1. The network is below critical load if $\lim_{M \rightarrow \infty} b_M^* = 0$ with $b_0^*, \dots, b_M^*$ solution of (14).

Due to Proposition 3 we have that $\lim_{M \rightarrow \infty} b_M^* = 0$ if and only if $\lim_{M \rightarrow \infty} b_0^* > 0$, and we now aim to obtain further insights in the latter condition. The following proposition states that, as the buffer capacity M increases, the value of b_0^* decreases and thus the limit $\lim_{M \rightarrow \infty} b_0^*$ exists.

Proposition 4. Let $b_0^{(M_1)}$ and $b_0^{(M_2)}$ solve (14) for $M = M_1$ and $M = M_2$, respectively. If $M_1 > M_2$, it holds that

$$b_0^{(M_1)} < b_0^{(M_2)}.$$

Proof. Let $F(\cdot; M)$ be the function defined in the proof of Theorem 1 where we indicate the value of the buffer capacity M for completeness. By definition of $b_0^{(M_1)}$, it holds that $F(b_0^{(M_1)}; M_1) = 1$. Since $M_1 > M_2$, we have that $F(b_0^{(M_1)}; M_2)$ is equal to

$$F(b_0^{(M_1)}; M_2) - b_0^{(M_1)} \sum_{m=M_2+1}^{M_1} \left(\frac{\lambda}{\nu \phi(1 - b_0^{(M_1)})} \right)^m < 1.$$

Hence, since $F(1; M_2) > 1$, it holds that $b_0^{(M_2)} \in (b_0^{(M_1)}, 1)$. $\square$

Due to the above proposition, we know that the limit $\lim_{M \rightarrow \infty} b_0^{(M)} = \beta^* \in [0, 1]$ exists since $\{b_0^{(M)}\}_M$ is a monotone and bounded sequence. Hence, the condition for super-criticality of a network $\lim_{M \rightarrow \infty} b_0^* > 0$ is well-posed.

Theorem 2. The network is below critical load if and only if

$$\rho < G(r, \sigma), \quad (17)$$

with $G(r, \sigma) := \sigma \phi(1) = \mathbb{E}[H_1]$ with H_1 (10).

Proof. First of all, observe that by definition a network is below critical load if and only if $\beta^* > 0$. Consider $\xi^* = 1 - \beta^*$, and observe that $\beta^* > 0$ if and only if $\xi^* \in [0, 1)$ solves

$$1 - \xi^* = \frac{1}{\sum_{m=0}^{\infty} \left(\frac{\lambda}{\nu \phi(\xi^*)} \right)^m}. \quad (18)$$

Assume that $\xi^* \in [0, 1)$ as above exists, then it necessarily holds that $\frac{\lambda}{\nu \phi(\xi^*)} < 1$, otherwise $\xi^* = 1$. Thus, we have that (18) becomes

$$\xi^* = \frac{\lambda}{\nu \phi(\xi^*)}. \quad (19)$$

Rearranging (19) and dividing by μ , we obtain that ξ^* needs to satisfy

$$\rho = \sigma \xi^* \phi(\xi^*).$$

Since, due to Lemma 1, the right-hand side is strictly increasing in ξ^* , we have that a solution with $\xi^* \in [0, 1)$ exists if and only if $\rho < \sigma \phi(1)$.

Let us now assume that (17) is satisfied, by following the argument above backwards, we deduce that there exists $\xi^* \in [0, 1)$ such that (18) is satisfied, and thus by setting $\beta^* = 1 - \xi^*$ we may conclude. $\square$

	M=5		M=50		M=500	
b_0^*	.18	.25	$4 \cdot 10^{-3}$	$9 \cdot 10^{-2}$	$9 \cdot 10^{-14}$	$9 \cdot 10^{-2}$
b_M^*	.15	.10	$5 \cdot 10^{-2}$	$7 \cdot 10^{-4}$	$5 \cdot 10^{-2}$	$2 \cdot 10^{-22}$

TABLE I
SOLUTION OF (20) FOR $\nu = 0.9$ AND $\nu = 1.1$.

Theorem 2 provides us with an easily verifiable criterion to establish whether a network is below critical load. We now further investigate a few scenarios so as to better illuminate the concepts introduced.

1) *Case with $r \geq 1/2$:* Observe that when $r \geq 1/2$, it holds that $K = 1$ and, in particular $\mathcal{B}_z = S^1$ for every z . In this case, we obtain that $\phi(\xi) = \frac{1}{1+\sigma\xi}$. Hence, from Proposition 3, we obtain that b_0^* solves

$$b_0^* = \frac{1 - \frac{\lambda(1+\sigma(1-b_0^*))}{\nu}}{1 - \left(\frac{\lambda(1+\sigma(1-b_0^*))}{\nu} \right)^{M+1}}, \quad (20)$$

which, as expected, does not depend on the exact value of $r \geq 1/2$. Let us now discuss under what condition the network is below critical load. Equation (17) becomes

$$\rho < \frac{\sigma}{1 + \sigma}. \quad (21)$$

So as to further illuminate the impact of the load criticality, we numerically solved (14) for various values of M with $\lambda = 0.5$, $\mu = 1$, and both $\nu = 0.9$ and $\nu = 1.1$. The latter case satisfies (21), the former does not. In the first case with $\nu = 0.9$, we observe that on the one hand $b_0 \downarrow 0$ while on the other hand b_M does not vanish in the limit and converges to

$$\frac{1}{\sum_{m=0}^{\infty} \left(\frac{1}{\nu \phi(1)} \right)^m} = 1 - \frac{\nu \phi(1)}{\lambda} \approx 0.0526.$$

In the second case with $\nu = 1.1$ the reverse happens, we obtain that $b_M \downarrow 0$ while b_0 does not vanish and converges to the solution of

$$b_0 = 1 - \frac{\lambda}{\nu \phi(1 - b_0)} \Rightarrow b_0 = \frac{1 - \frac{\lambda}{\mu} - \frac{\lambda}{\nu}}{1 - \frac{\lambda}{\mu}} = 0.09.$$

The results are presented in Table I and match with the analysis above.

2) *Case with $1/2 > r \geq 1/3$:* In this case we obtain that $K = 2$ and the size of $\mathcal{B}_z$ depends directly on r . From Proposition 1, we obtain

$$\phi(\xi) = \frac{1 + \sigma \xi(1 - 2r)}{1 + \sigma \xi + \frac{\sigma^2 \xi^2}{2}(1 - 2r)}.$$

Hence, from Proposition 3, we obtain b_0 as the solution of

$$b_0 = \frac{1 - \frac{\lambda(1+\sigma(1-b_0) + \frac{\sigma^2(1-b_0)^2}{2}(1-2r))}{\nu(1+\sigma(1-b_0)(1-2r))}}{1 - \left(\frac{\lambda(1+\sigma(1-b_0) + \frac{\sigma^2(1-b_0)^2}{2}(1-2r))}{\nu(1+\sigma(1-b_0)(1-2r))} \right)^{M+1}}, \quad (22)$$

which can be obtained numerically. However, as r decreases, to compute the fixed point we need to rely on numerical

Fig. 2. Values of r and σ for which $G(r, \sigma) > \rho$ (black area) for $\rho = 0.5$, $\rho = 0.9$, $\rho = 1.3$, and $\rho = 2.5$.

methods such as the bisection algorithm discussed in Section V. Let us now discuss under what condition the network is below critical load. By rearranging (17) we obtain the following explicit condition

$$\rho < \frac{\sigma + \sigma^2(1 - 2r)}{1 + \sigma + \frac{\sigma^2(1 - 2r)}{2}}.$$

3) *Case with general reuse distance*: The analysis above clearly conveys that a closed-form expression for the fixed point is out of reach as r decreases. The load-criticality condition (17) on the other hand may still be quite easily verified. Intuitively, we expect $G(r, \sigma)$ to decrease in r and increase in σ , i.e., reducing the interference and increasing the aggressiveness of the nodes should increase the threshold for critical load. This is what we observe in Figure 2, where we display for which values of σ and r we have that $\rho < G(r, \sigma)$ (black area) for various values of ρ . In general, when $r < \frac{1}{|\rho|}$, we can ensure that condition (17) is satisfied by choosing a sufficiently large value for σ .

4) *Case with $r \downarrow 0$* : As a special scenario, let us consider the case in which $r \downarrow 0$. As an immediate consequence, we have that $K \rightarrow \infty$ and thus

$$\lim_{r \downarrow 0} G(r, \sigma) = \frac{\sum_{j=1}^{\infty} \frac{\sigma^j}{(j-1)!}}{1 + \sum_{j=1}^{\infty} \frac{\sigma^j}{j!}} = \frac{\sigma e^{\sigma}}{e^{\sigma}} = \sigma.$$

Hence, condition (17) becomes $\rho < \sigma$, i.e., $\lambda < \nu$. Intuitively, when $r = 0$ the various nodes behave independently and thus the network is below critical load if and only if the mean interarrival time is larger than the mean back-off plus service time, i.e.,

$$\frac{N}{\lambda} > \frac{N}{\nu} + \frac{1}{\mu},$$

which reduces to $\lambda < \nu$ as $N \rightarrow \infty$.

VII. NUMERICAL EXPERIMENTS

In this section we present simulations to examine the accuracy of the fixed-point approximations introduced in

Fig. 3. Average queue lengths.

Fig. 4. Average fraction of blocked packets.

Section V-A.

A. Load criticality

We first present two similar examples. In both cases we set $\nu = \mu = 1$, $r = 0.21$, $M = 5$, and we fix $\lambda = 0.5$ and $\lambda = 1$ in the first and second example, respectively. Observe that the network in the first example is below load criticality, the second is not. For both examples we simulate the behavior of networks with $N = 10, 20, 50$ nodes and let it run up to time $T = N \cdot 10^5$. For every node n , we keep track of

$$\hat{B}_n^{(N)}(T), \quad \hat{W}_n^{(N)}(T), \quad \hat{Q}_{n,m}^{(N)}(T), \quad m = 0, \dots, M,$$

and compute the average and standard deviation of these values over the various nodes.

In Figure 3 we display the results for the average queue lengths. Note that the approximation is fairly accurate already for $N = 10$ and cannot be distinguished from the simulation results for $N = 50$. Observe the difference between the two examples: in the first most of the time the buffer of a node is empty, while in the second it is mostly full. This is explained by the first network being below load criticality and the second not.

In Figures 4 and 5, we display the results for the average fraction of blocked packets and average normalized packet delay. As noted for the queue length, the approximations are fairly accurate already for $N = 10$ and extremely precise for $N = 50$.

B. A larger example

We now present an example for which a closed-form fixed point could not be obtained. We set $\nu = 2$, $\mu = 1$, $\lambda = 1$, $r = 0.08$, $M = 10$, and observe that $K = 12$. We simulate the behavior of networks with $N = 10, 20, 50, 100$ nodes

Fig. 5. Average normalized packet delay.

Fig. 6. Average queue lengths.

and let them run up to time $T = N \cdot 10^4$. We compared the simulation values with the approximations based on the fixed point computed via the bisection method.

In Figure 6 we display the approximation for the average stationary queue lengths and in Figure 7 those for the average fraction of blocked packets and average normalized packet delay. The accuracy is striking as N grows large and quite remarkable even for small networks.

REFERENCES

- [1] Cisco (2011). Cisco visual networking index: Global mobile data traffic forecast update.
- [2] Ericsson (2011). More than 50 billion connected devices. White paper.
- [3] G. Bianchi (2000). Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE J. Sel. Areas Commun.* **18** (3), 535–547.
- [4] P. Billingsley (1968). Convergence of probability measures. 1st Edition.
- [5] C. Bordenave, D. McDonald, A. Proutiere (2008). Performance of random medium access control, an asymptotic approach. In: *Proc. ACM SIGMETRICS* **36** (1), 1–12.
- [6] F. Cecchi, S.C. Borst, J.S.H. van Leeuwaarden, P.A. Whiting (2016). CSMA networks in a many-sources regime: A mean-field approach. In: *Proc. IEEE Infocom*.

- [7] F. Cecchi, S.C. Borst, J.S.H. van Leeuwaarden, P.A. Whiting (2017). Spatial mean-field limits for ultra-dense random-access networks. *Perf. Eval.* (in press).
- [8] F. Cecchi, S.C. Borst, J.S.H. van Leeuwaarden, P.A. Whiting (2016). Mean-field limits for large-scale random-access networks. arXiv preprint arXiv:1611.09723.
- [9] A. Chaintreau, J.Y. Le Boudec, N. Ristanovic (2009) The age of gossip: spatial mean field regime. In: *Proc. ACM SIGMETRICS* **37** (1), 109–120.
- [10] 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.
- [11] K.R. Duffy (2010). Mean field Markov models of wireless local area networks. *Markov Proc. Rel. Fields* **16** (2), 295–328.
- [12] S. Foss, S. Juneja, M.R.H. Mandjes, S.B. Moka (2015). Spatial loss systems: Exact simulation and rare event behavior. *ACM SIGMETRICS Perf. Eval. Rev.* **43** (92), 3–6.
- [13] J. Horowitz (1985). Measure-valued random processes. *Prob. Theory Rel. Fields* **70**(2), 213–236.
- [14] F.P. Kelly (1985). Stochastic models of computer communication systems. *J. Roy. Stat. Soc. B* **47** (3), 379–395.
- [15] K.B. Öner, G.P. Kiesmüller, G.-J. Van Houtum (2009). Monotonicity and supermodularity results for the Erlang loss system. *Operations Research Letters* **37** (4), 265–268.
- [16] M. Tschaikowski, M. Tribastone (2017). Spatial fluid limits for stochastic mobile networks. *Perf. Eval.* **109**, 52–76.

APPENDIX

A. Proof of Lemma 1

In view of the discussion below Lemma 1, it suffices to show that the sequence $\{f(k)\}_{k=1, \dots, K-1}$ with

$$f(k) = \frac{(1 - (k+1)r)^k}{(1 - kr)^{k-1}}$$

is decreasing in k . For any $k = 1, \dots, K-2$, the condition $f(k) \geq f(k+1)$ is equivalent to

$$\begin{aligned} 1 &\geq \left(\frac{1 - (k+2)r}{1 - (k+1)r} \right)^{k+1} \left(\frac{1 - kr}{1 - (k+1)r} \right)^{k-1} \\ &= \left(\frac{1 - (k+2)r}{1 - (k+1)r} \right)^2 \left(\frac{(1 - kr)(1 - (k+2)r)}{(1 - (k+1)r)^2} \right)^{k-1} \\ &= \left(\frac{1 - (k+2)r}{1 - (k+1)r} \right)^2 \left(1 - \frac{r^2}{(1 - (k+1)r)^2} \right)^{k-1}, \end{aligned}$$

which holds since both factors on the right-hand side are nonnegative and lower than 1.

As an alternative algebraic proof, Lemma 1 may be proved by directly applying [15, Lemma 2] to the quotient of positive analytic functions introduced in Proposition 1.

Fig. 7. (a) Average fraction of blocked packets. (b) Average normalized packet delay.