

EURANDOM PREPRINT SERIES
2011-003

**Stability of Random Admissible-Set Scheduling
in Spatial Wireless Systems**

N. Bouman, S.C. Borst and J.S.H. van Leeuwaarden
ISSN 1389-2355

Stability of Random Admissible-Set Scheduling in Spatial Wireless Systems*

N. Bouman
Eindhoven University of Technology
EURANDOM

S.C. Borst
Eindhoven University of Technology
Alcatel-Lucent Bell Labs

J.S.H. van Leeuwen
Eindhoven University of Technology
EURANDOM

December 16, 2010

Abstract

We examine the stability of wireless networks whose users are distributed over a compact space. Users arrive at spatially uniform locations with intensity λ and each user has a random number of packets to transmit with mean β . In each time slot, an *admissible* subset of users is selected uniformly at random to transmit one packet. A subset of users is called admissible when their simultaneous activity obeys the prevailing interference constraints. We consider a wide class of interference constraints, including the SINR model and the protocol model. Denote by μ the maximum number of users in an admissible subset for the model under consideration. We will show that the necessary condition $\lambda\beta < \mu$ is also sufficient for random admissible-set scheduling to achieve stability. Thus random admissible-set scheduling achieves stability, if feasible to do so at all, for a broad class of interference scenarios. The proof relies on a description of the system as a measure-valued process and the identification of a Lyapunov function.

1 Introduction

The present paper examines the stability of a broad class of wireless networks whose users arrive to a compact space H . Time is slotted, and users arrive at H according to some spatial stochastic process with mean λ per time slot. Users independently take their locations in H at random according to the uniform distribution. Each user has a random number of packets to transmit, generally distributed with mean β , and can transmit at most one packet per time slot.

In each time slot, we select a set of users for transmission from all admissible sets uniformly at random. A subset of users is called admissible when their simultaneous activity obeys

*A poster version of the present paper was presented at the Performance 2010 conference.

the prevailing interference constraints. In practice, the relevant interference constraints depend on various system-specific properties, such as the propagation environment and the operation of the physical and medium access layers of the network. In the present paper we therefore adopt generic feasibility criteria, which in particular cover both the SINR (Signal-to-Interference-and-Noise Ratio) model and the protocol model as two canonical models for interference.

Let μ be the maximum number of users in an admissible set. It is clear that $\lambda\beta \leq \mu$ is a *necessary* condition for stability: The mean number of packets that arrive per time slot should be no larger than the maximum number of packets that can be transmitted simultaneously. The main result we show is that this necessary condition is also nearly *sufficient* for stability. Specifically, the Markov chain describing the evolution of the system is then positive Harris recurrent, implying that the network will be empty infinitely often. Thus random admissible-set scheduling is a highly robust strategy in that it achieves stability, if feasible to do so at all, for a wide range of interference scenarios.

For wireless networks we are primarily interested in interference constraints that become looser when users are further apart and spaces such as the one- and two-dimensional torus and the ordinary sphere. Our results, however, hold in more generality, which is why we present our results in terms of a particle system on a compact space H , where particles (or packets) arrive to the system according to a spatial stochastic process with rate λ and in batches with mean β . The arrival location is independent of other users and uniform on H . As described above, we investigate stability in the context of a model that combines a scheduling discipline operating under interference constraints and a continuous spatial setting. While these two elements have each been considered in isolation before, the present paper is, to the best of our knowledge, the first to capture both features in conjunction. Indeed, stability of wireless networks has been widely studied in the literature, see for instance Bonald and Feuillet [2], Bordenave *et al.* [3] and Wu *et al.* [12]. These papers restrict the attention though to discrete topologies and interference constraints such that the system can be represented as a conflict graph. Our model does not allow such a representation due to the continuum of locations, and hence these results are not directly applicable to our problem. Stability of queueing networks in continuous space is investigated in Altman and Levy [1], Robert [10] and Leskelä and Unger [7]. These papers prove stability of networks in which only one user is allowed to transmit at a time. In contrast, the present paper focuses on the more complex situation of simultaneous transmissions as governed by a scheduling discipline.

While a discrete network structure is a reasonable assumption in case of a relatively small number of long-lived sources, it is less suitable in case of a relatively large number of short-lived flows. The latter scenario is increasingly relevant as emerging wireless networks support traffic generated by massive numbers of nodes which each individually may only engage in sporadic transmission activity. The continuous spatial setting also provides useful insights in the scaling behavior of discrete topologies as the number of nodes grows large. From a methodological perspective, the continuous spatial setting involves major additional challenges compared to a discrete network structure. Since users reside in a continuum of locations, the evolution of the system cannot be represented in terms of a Markov chain with

a finite state space, and we therefore introduce a measure-valued process as a description of the system. In order to prove stability, we identify a Lyapunov function which has a negative drift for all but a ‘small’ set of states, so that the Markov chain is positive Harris recurrent.

The remainder of the paper is organized as follows. In Section 2 we show how the evolution of the system may be described in terms of a Markov chain with a measure-valued state space. The main stability result is presented in Section 3, along with an interpretation and discussion of its ramifications. In Section 4 we provide the proof for our main result, and in particular identify a Lyapunov function which has negative drift for all but a ‘small’ set of states, and plays a critical role in the proof. In Appendix A we recall various useful definitions and collect some preliminaries that are needed in order to apply the Foster-Lyapunov approach for our specific Markov chain.

2 Model description

Consider a compact space $H \subset \mathbb{R}^n$ for some $n < \infty$. Denote by $A(t, B)$ the number of particles arriving during the t -th time slot in $B \subseteq H$. Particles arrive in batches, and the batch size has a general non-negative discrete distribution with mean β , independent of the sizes and locations of other batches. We assume batches have a size of at most one with positive probability. Batches arrive at locations uniformly distributed on H , independent of the locations of other batches. The number of batches that arrive during a time slot has a general non-negative discrete distribution with mean λ and is independent of the number of batches that arrive in other time slots. That is, the numbers $A(t, H)$, $t = 1, 2, \dots$, are i.i.d. copies of a non-negative random variable A with $\mathbb{E}\{A\} = \lambda\beta$. Further we assume that $\mathbb{E}\{A \log(A) | A > 0\} < \infty$. Note that, because batches arrive uniformly, the expected number of particles to arrive to a subspace $B \subseteq H$ in one time slot is given by $\mathbb{E}\{A\} \nu(B)$, where ν is some measure on H such that, without loss of generality, $\nu(H) = 1$, representing for example the volume or surface area.

We denote the number of particles in the space H at the start of the t -th time slot by $Y(t)$, with $Y(t) = (Y(t, B), B \subseteq H)$ and $Y(t, B)$ denoting the number of particles residing in the subspace B at the start of the t -th time slot. The state space of this process is denoted by Ψ and consists of all finite counting measures on H . So, when $y \in \Psi$, $y(B)$ denotes the number of particles residing in B . In particular, $y(\{x\})$ denotes the number of particles at $x \in H$.

At the start of every time slot an admissible subset of particles will be removed. Here, $z \in \Psi$ is called a subset of $y \in \Psi$ if $z(\{x\}) \leq y(\{x\}), \forall x \in H$. To decide whether a set is admissible we define a function $F : \Psi \rightarrow \{0, 1\}$ having the following properties

1. $F(y_0) = 1$, where y_0 denotes the empty configuration, $y_0(H) = 0$.
2. $F(y) = 1$ only if $y(\{x\}) \in \{0, 1\}, \forall x \in H$.
3. There exists a partition of H , i.e. disjoint subspaces $P_i, i = 1, \dots, K$, for which $\nu(P_i) = \frac{1}{K}, \forall i \in \{1, \dots, K\}$ and $\bigcup_{i=1}^K P_i = H$, such that $F(y) = 1$ with $y(\{x\}) = 1$ for

some $x \in P_i$ implies that $y(P_i \setminus \{x\}) = 0$. Furthermore, taking $y \in \Psi$ such that $y(P_{m+j}) = 1, \forall j \in \{1, \dots, \mu\}$ and $y(H \setminus \bigcup_{j=1}^{\mu} P_{m+j}) = 0$, with $m < K$ such that $m \bmod \mu = 0$, implies that $F(y) = 1$. Here $\mu = \max\{y(H) : F(y) = 1\}$, i.e. the maximum number of particles in an admissible set, and K is a multiple of μ .

We call $y \in \Psi$ admissible if and only if $F(y) = 1$.

Property 1 ensures that the empty set is admissible, such that there always is at least one set that is admissible. Property 2 states that at most one particle of the particles located at some location $x \in H$ can be removed in a time slot. Property 3 states that the space H can be partitioned using a finite number of equally-sized disjoint sets such that at most one particle in each set can be removed in a time slot. This property further states that this partitioning is such that taking any combination of elements of maximal size (μ) from certain sets will always give an admissible set, this is to avoid having almost surely no admissible subset of particles of maximum size at any point in time. Finally, Property 3 states that the partitioning is such that there are K/μ disjoint sets of sets with this property, i.e. the property makes sure that taking any combination of elements of maximal size from these sets will always give an admissible set. Note that μ is defined implicitly.

We will now give an example of a function which satisfies the above properties. The model in this example is the so-called protocol model.

Example 2.1. Consider a unit circle, use the interval $[0, 1)$ to denote points on the circle and let ν be the Lebesgue measure. Particles can be removed simultaneously whenever the distance between these particles is at least r , i.e. $F(y) = 1$ if $y(\{x\}) \in \{0, 1\}, \forall x \in [0, 1)$ and $y(\{x\})y(\{w\}) > 0$ only if $D(x, w) \geq r \forall x \neq w \in [0, 1)$, with $D(x, y) = \min(|x - y|, 1 - |x - y|)$. Properties 1 and 2 immediately follow from this description. To verify property 3 take $K \geq 2\mu/(1 - \mu r)$ and K a multiple of $\mu = \lfloor 1/r \rfloor$. Now, for $i = 1, \dots, K$, take

$$P_i = \left[\frac{\lfloor \frac{i-1}{\mu} \rfloor + \frac{K}{\mu}(i-1 \bmod \mu)}{K}, \frac{1 + \lfloor \frac{i-1}{\mu} \rfloor + \frac{K}{\mu}(i-1 \bmod \mu)}{K} \right).$$

We then see that with these sets property 3 is satisfied whenever $1/r$ is non-integer. For property 3 to hold if $1/r$ is integer-valued we do not allow sets of size $1/r$ to be removed. Note that the probability that a set of size $1/r$ is removed is almost surely zero in any time slot if $1/r$ is integer-valued. We then see that property 3 is satisfied with the above sets P_i and $\mu = \lceil 1/r \rceil - 1$.

Let $\chi(y)$ be the set of all subsets of y , i.e.

$$\chi(y) = \{z \in \Psi : z(\{x\}) \leq y(\{x\}), \forall x \in H\},$$

and let $R(t, Y(t), B)$ be the number of particles removed from $B \subseteq H$ in the t -th time slot, given the configuration, $Y(t)$. An admissible subset of particles is selected uniformly at random. Hence, given $Y(t) = y$, $R(t, y) = z$ with probability

$$\frac{F(z) \prod_{x \in H, z(\{x\}) > 0} y(\{x\})}{\sum_{u \in \chi(y)} F(u) \prod_{x \in H, u(\{x\}) > 0} y(\{x\})},$$

where $R(t, y) = (R(t, y, B), B \subseteq H)$. Note that for $z \notin \chi(y)$ we have $F(z) = 0$ or $\exists x \in H$ such that $z(\{x\}) = 1$ and $y(\{x\}) = 0$, so this probability is always zero in this case. Further note that, here and in the remainder of this paper, the value of the empty product is defined as 1, as is usually done in the literature.

The evolution of $Y(t, B)$ is then described by the recursion

$$Y(t, B) = Y(t - 1, B) + A(t - 1, B) - R(t, B).$$

Further, $R(t, B)$ depends on the number of particles just before the start of the t -th time slot, so

$$Y(t^-, B) = Y(t - 1, B) + A(t - 1, B).$$

From this description it follows that $(Y(t))_{t \in \mathbb{N}}$ is a Markov chain. We will equip the state space of the Markov chain, Ψ , with the smallest σ -field $\mathcal{B}(\Psi)$ with respect to which the map $y \rightarrow y(B)$ is measurable for any Borel set $B \subseteq H$. That is, we equip Ψ with the Borel σ -field as we will prove in Lemma A.2.

3 Main result

The following theorem states the main result of this paper and is a shorter version of Theorem 4.5 which is proven at the end of Section 4.

Theorem 3.1. *Assume $\lambda\beta < \mu$. Then, the Markov chain $(Y(t))_{t \in \mathbb{N}}$ is positive Harris recurrent.*

This theorem states that, starting from the empty configuration, $Y(t, H) = 0$ for infinitely many values of t , i.e. the Markov chain will be in the empty configuration infinitely often. We thus see that random admissible-set scheduling achieves maximal stability, independent of the specific function F .

The result of Theorem 3.1 may be interpreted as follows. Suppose that the total number of particles in the system is large. Then there will be a large number of admissible sets of size μ , assuming that the particles are sufficiently dispersed across the network and not concentrated in a few dense areas. In fact, the number of admissible sets of size μ will be overwhelmingly large compared to the number of admissible sets of smaller size. By virtue of random-admissible set scheduling, one of the admissible sets of size μ will then be selected with high probability. Thus, the expected number of removed particles will exceed the expected number of arriving particles, provided $\lambda\beta < \mu$, implying a reduction in the expected number of particles in the system, and preventing the number of particles from growing without bound.

As the above heuristic explanation indicates, it is crucial for the particles to be sufficiently spread out and not be clustered in a few hot spots. In order to obtain a rigorous proof, it will hence not suffice to just consider the total number of particles, but in fact be necessary to keep track of their individual locations.

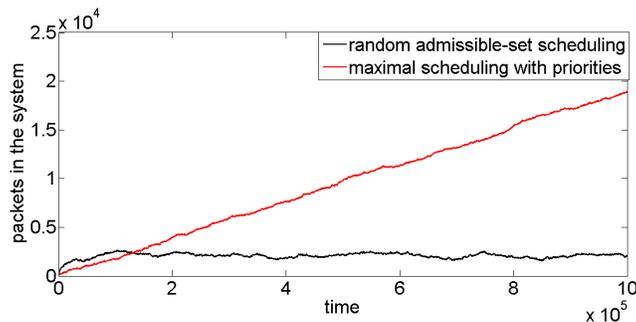


Figure 1: Random-admissible set scheduling and maximal scheduling with priorities for $\zeta = 0.5$, $r = 0.49$ and $\lambda = 1.95$.

It is worth emphasizing that we consider a particle-based version of random-admissible set scheduling rather than a node-based incarnation, in the sense that the strategy selects among sets of particles rather than sets of nodes. While this distinction is immaterial when the measure $\nu(\cdot)$ is an absolutely continuous density, the issue does become relevant when the measure $\nu(\cdot)$ has mass in discrete points. In the latter case, it may readily be concluded that a node-based version of random admissible-set scheduling may fail to achieve maximum stability. Interestingly, this observation contrasts with the fact that the node-based version of the celebrated MaxWeight scheduling strategy guarantees maximum stability when the measure $\nu(\cdot)$ is a purely discrete distribution, whereas a particle-based version may fail to do so [11].

Also, while the spatial dispersion of particles under random admissible-set scheduling is intuitively plausible, it is certainly not obvious. This is perhaps best illustrated by the fact that the result of Theorem 3.1 may not necessarily hold for seemingly similar but subtly different scheduling disciplines.

As an example consider the situation of Example 2.1. Instead of selecting an admissible set at random, now consider the scheduling discipline that gives priority to particles that are closest, in anticlockwise direction, to a certain point ζ on the circle. Obeying this priority rule, we select as many particles as possible to remove in a certain time slot. That is, the particle closest to the given point gets removed, the particle closest to the point and at least a distance r away from the first particle gets removed, and so on until no particle can be selected anymore. We call this service discipline maximal scheduling with priorities.

Figure 1 shows a simulation result for both scheduling disciplines with $\zeta = 0.5$, $r = 0.49$ and $\lambda = 1.95$ starting from an empty configuration and running for 10^6 time slots. That is, the figure gives a realization of $Y(t, H)$ given that $Y(0, H) = 0$ for $t = 1, \dots, 10^6$. We see that at the start of the simulation the number of particles in the system with random-admissible set scheduling grows faster than the number of particles in the system with maximal scheduling with priorities. This is because maximal scheduling with priorities always selects a subset of maximum size obeying the priority rules, whereas random-admissible set scheduling always selects admissible sets of a small size with a certain probability, which gets lower as the number of particles in the network grows. More importantly, after some time the number of

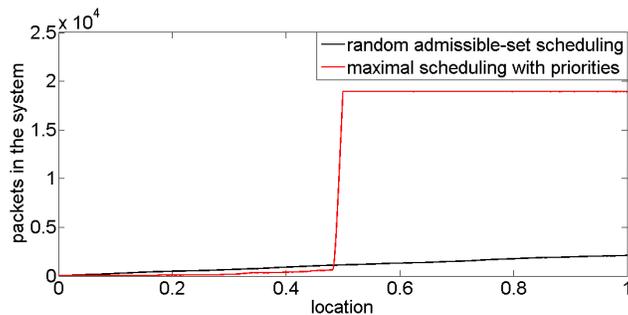


Figure 2: Terminal configuration of the simulation of Figure 1.

particles in the system with random-admissible set scheduling settles around an equilibrium value whereas the number of particles in the system with maximal scheduling with priorities keeps on growing linearly. This suggests that maximal scheduling with priorities is not stable while random-admissible set scheduling is stable for the chosen parameters. The latter will be proven in the next section.

Figure 2 shows the terminal configuration of the simulation of Figure 1, i.e. it gives a realization of $Y(10^6, B)$ given that $Y(0, 1) = 0$ for $B = [0, s)$, $0 < s \leq 1$, for both scheduling disciplines. For random admissible-set scheduling we see that the number of particles in the interval $[0, s)$ is roughly linear in s , indicating that the particles are evenly spread out over the circle. For maximal scheduling with priorities we observe that the number of particles in $[0, s)$ slowly increases with s up to approximately $s = 0.48$, after which the number of particles in the system steeply rises up to $s = 0.5$. For $s \geq 0.5 = \zeta$ the number of particles in the interval $[0, s)$ is (almost) constant, implying that virtually no particles are located in the interval $[0.5, 1)$. Note that particles in the interval $[0.48, 0.5)$ have the lowest priority and hence are, whenever they are allowed to, almost always removed simultaneously with other particles, as there are quite some particles in the system and outside this interval. However, the particles that are allowed to be removed simultaneously with particles in $[0.48, 0.5)$ are also allowed to be removed simultaneously with some particles in $[0.5, 0.52)$, who have the highest priority. So we infer that the particles in this system are clustered in $[0.48, 0.5)$, and that too high a fraction of the time (larger than 0.02λ) no particle in this interval is removed, making the system unstable.

4 Proof

In this section we provide the proof of Theorem 3.1. As mentioned earlier, the proof relies on the Foster-Lyapunov criteria and involves the identification of a function which has negative drift for all but a small set of states. Appendix A contains several useful definitions and preliminaries that are needed to apply the Foster-Lyapunov approach for our specific Markov chain.

In order to define the Lyapunov function, we use a partitioning P_1, \dots, P_K of the space H for which Property 3 of the function F holds. Let $x_k(y)$ be the number of particles residing

in the k -th region, P_k , given configuration $y \in \Psi$, for $k \in \mathcal{K}$, where $\mathcal{K} = \{1, \dots, K\}$. Let $\Omega = \mathcal{P}(\mathcal{K})$ be the collection of all subsets of $\{1, \dots, K\}$ and let $\Omega(y)$ be the subsets containing particles, i.e. $\Omega(y) = \{S \in \Omega : x_k(y) \geq 1, \forall k \in S\}$. A subset $S \in \Omega$ is called ‘guaranteed’ if any subset of particles, with exactly one residing in each of the regions contained in S , is admissible, regardless of the exact locations within each of the regions. Let $\Theta \subseteq \Omega$ be the collection of all ‘guaranteed’ subsets and, again, let $\Theta(y)$ be the subsets containing particles, $\Theta(y) = \{S \in \Theta : x_k(y) \geq 1, \forall k \in S\}$. Further, denote by $q_S(y)$ the probability that the particles that get removed belong to the subset of regions $S \in \Omega$ given that the system is in state y , with $p_k(y) = \sum_{S \in \Omega: S \ni k} q_S(y)$ the probability that a particle in the k -th region gets removed given configuration $y \in \Psi$.

For $S \in \Omega$, denote

$$w_S(y) = \prod_{k \in S: x_k(y) \geq 1} x_k(y).$$

Further define

$$B(\epsilon) = \left\{ y \in \Psi : w(y) \geq \left(\frac{2|\Omega|}{\epsilon} \right)^{2/\epsilon} \right\},$$

where $\epsilon > 0$ and

$$w(y) = \max_{S \in \Theta(y)} w_S(y).$$

In the next lemma we will show that the value of $\sum_{k \in S} \log(x_k)$, for the set of particles S selected with random admissible-set scheduling, is close to the maximum possible value over all admissible sets with high probability for all states $y \in B(\epsilon)$.

Lemma 4.1. *For all states $y \in B(\epsilon)$ we have,*

$$\sum_{S \in \Omega(y)} q_S(y) \log(w_S(y)) \geq (1 - \epsilon) \log(w(y)).$$

Proof. The proof proceeds along similar lines as in [2], [9].

Define

$$\Upsilon(y) = \{S \in \Omega(y) : \log(w_S(y)) \geq (1 - \frac{\epsilon}{2}) \log(w(y))\}.$$

Then,

$$\sum_{S \in \Omega(y)} q_S(y) \log(w_S(y)) \geq (1 - \frac{\epsilon}{2}) \log(w(y)) \sum_{S \in \Upsilon(y)} q_S(y). \quad (1)$$

For any $S \in \Omega$, let $v_S(y)$ be the number of admissible subsets of particles of size $|S|$ with exactly one residing in each of the regions contained in S , with the convention that $v_\emptyset(y) = 1$ for all $y \in \Psi$. Because every admissible subset can have at most one particle residing in

each region, there is exactly one $S \in \Omega$ for which this subset is counted in $v_S(y)$. Thus the total number of admissible subsets of particles is given by $\sum_{T \in \Omega} v_T(y)$ and, as an admissible subset of particles is selected uniformly at random,

$$q_S(y) = \frac{v_S(y)}{\sum_{T \in \Omega} v_T(y)} = \frac{v_S(y)}{\sum_{T \in \Omega(y)} v_T(y)}.$$

Further observe that $v_S(y) \leq w_S(y)$ for all $S \in \Omega$, with equality for all $S \in \Theta(y)$. Thus

$$\sum_{S \notin \Upsilon(y)} q_S(y) = \frac{\sum_{S \notin \Upsilon(y)} v_S(y)}{\sum_{S \in \Omega(y)} v_S(y)} \leq \frac{\sum_{S \notin \Upsilon(y)} w_S(y)}{\sum_{S \in \Theta(y)} w_S(y)} \leq \frac{|\Omega| w(y)^{1-\frac{\epsilon}{2}}}{w(y)} = |\Omega| w(y)^{-\frac{\epsilon}{2}}.$$

The latter quantity is less than $\frac{\epsilon}{2}$ for all states $y \in B(\epsilon)$, and thus

$$\sum_{S \in \Upsilon(y)} q_S(y) \geq 1 - \frac{\epsilon}{2}. \quad (2)$$

Combining the lower bounds (1) and (2), we obtain

$$\sum_{S \in \Omega(y)} q_S(y) \log(w_S(y)) \geq (1 - \frac{\epsilon}{2}) \log(w(y)) (1 - \frac{\epsilon}{2}) \geq (1 - \epsilon) \log(w(y))$$

for all states $y \in B(\epsilon)$. □

For $i = 1, \dots, K$ define the set $S_i \in \Omega$ by $S_i = \{[\frac{i}{\mu}] \mu - j, \forall j \in \{0, \dots, \mu - 1\}\}$. Note that, by definition of F , $S_i \in \Theta$. Further define $S_i(y) = \{j \in S_i(y) : x_j(y) \geq 1\}$ and note that $S_i(y) \in \Theta(y)$.

Also define the nonnegative function $V : \Psi \rightarrow \mathbb{R}$ by

$$V(y) = \sum_{k: x_k(y) \geq 1} x_k(y) \log(x_k(y)).$$

and the function $G : \Psi \rightarrow \mathbb{R}$ by

$$G(y) = \sum_{k: x_k(y) \geq 1} \log(x_k(y)) \left[\mathbb{E} \{A_k\} - p_k(y) \right],$$

where A_k denotes the number of arrivals in the k -th region, so $\mathbb{E} \{A_k\} = \frac{\lambda \beta}{K}$. Note that

$$V(y) = \sum_{k=1}^K x_k(y) \log(\max(x_k(y), 1))$$

and

$$V(y) = \sum_{k: x_k(y) \geq 2} x_k(y) \log(x_k(y)).$$

Further observe that the function $V(\cdot)$ only depends on y through the values of $x_k(y)$. However, the $x_k(y)$'s do not constitute a Markov chain, and hence we need to treat $V(\cdot)$ as a function of the full state description y in order for the Foster-Lyapunov approach to apply.

We now first find the relation between the drift of $V(y)$ and $G(y)$, where the drift of $V(y)$ is defined by

$$\Delta V(y) = \mathbb{E} \{ (Y(t+1)) | Y(t) = y \} - V(y).$$

After that we will find an upper bound for $G(y)$.

Lemma 4.2. $\Delta V(y) = G(y) + G_2(y)$, with $G_2(y)$ a bounded function.

Proof. Remember that at most one particle can be removed from each region in every time slot. We thus get

$$\begin{aligned} & \mathbb{E} \{ V(Y(t+1)) | Y(t) = y \} \\ = & \mathbb{E} \left\{ \sum_{k=1}^K x_k(Y(t+1)) \log(\max(x_k(Y(t+1)), 1)) \middle| Y(t) = y \right\} \\ = & \sum_{k=1}^K \mathbb{E} \left\{ x_k(Y(t+1)) \log(\max(x_k(Y(t+1)), 1)) \middle| Y(t) = y \right\} \\ = & \sum_{k: x_k(y) \geq 2} p_k(y) \mathbb{E} \{ (x_k(y) + A_k(t) - 1) \log(x_k(y) + A_k(t) - 1) \} \\ & + \sum_{k: x_k(y) \geq 2} (1 - p_k(y)) \mathbb{E} \{ (x_k(y) + A_k(t)) \log(x_k(y) + A_k(t)) \} \\ & + \sum_{k: x_k(y) = 1} p_k(y) \mathbb{P} \{ A_k(t) \geq 1 \} \mathbb{E} \{ (1 + A_k(t) - 1) \log(1 + A_k(t) - 1) | A_k(t) \geq 1 \} \\ & + \sum_{k: x_k(y) = 1} (1 - p_k(y)) \mathbb{E} \{ (1 + A_k(t)) \log(1 + A_k(t)) \} \\ & + \sum_{k: x_k(y) = 0} \mathbb{P} \{ A_k(t) \geq 1 \} \mathbb{E} \{ A_k(t) \log(A_k(t)) | A_k(t) \geq 1 \}, \end{aligned}$$

We further have

$$\begin{aligned}
& \sum_{k:x_k(y) \geq 2} p_k(y) \mathbb{E} \{ (x_k(y) + A_k(t) - 1) \log(x_k(y) + A_k(t) - 1) \} \\
& + \sum_{k:x_k(y) \geq 2} (1 - p_k(y)) \mathbb{E} \{ (x_k(y) + A_k(t)) \log(x_k(y) + A_k(t)) \} \\
& = \sum_{k:x_k(y) \geq 2} p_k(y) \mathbb{E} \left\{ (x_k(y) + A_k(t) - 1) \left(\log(x_k(y)) + \log \left(1 + \frac{A_k(t)}{x_k(y)} - \frac{1}{x_k(y)} \right) \right) \right\} \\
& + \sum_{k:x_k(y) \geq 2} (1 - p_k(y)) \mathbb{E} \left\{ (x_k(y) + A_k(t)) \left(\log(x_k(y)) + \log \left(1 + \frac{A_k(t)}{x_k(y)} \right) \right) \right\} \\
& = \sum_{k:x_k(y) \geq 2} (x_k(y) + \mathbb{E} \{ A_k \} - p_k(y)) \log(x_k(y)) \\
& + \sum_{k:x_k(y) \geq 2} p_k(y) \mathbb{E} \left\{ (x_k(y) + A_k(t) - 1) \log \left(1 + \frac{A_k(t)}{x_k(y)} - \frac{1}{x_k(y)} \right) \right\} \\
& + \sum_{k:x_k(y) \geq 2} (1 - p_k(y)) \mathbb{E} \left\{ (x_k(y) + A_k(t)) \log \left(1 + \frac{A_k(t)}{x_k(y)} \right) \right\}.
\end{aligned}$$

Now notice that for constants $a \geq 0, b \geq 0, c > 0$

$$\begin{aligned}
& \mathbb{E} \left\{ (a + A_k(t)) \log \left(b + \frac{A_k(t)}{c} \right) \right\} \\
& = \mathbb{E} \{ (a + A_k(t)) (\log(bc + A_k(t)) - \log(c)) \} \\
& \leq \mathbb{E} \{ (a + A_k(t)) (\log(bc + A_k(t)) - \log(c)) | A_k(t) \geq 1 \} \\
& \leq \mathbb{E} \{ (a + A_k(t)) (bc + \log(A_k(t)) - \log(c)) | A_k(t) \geq 1 \} \\
& = a(bc - \log(c)) + \mathbb{E} \{ bcA_k(t) + a \log(A_k(t)) | A_k(t) \geq 1 \} + \mathbb{E} \{ A_k(t) \log(A_k(t)) | A_k(t) \geq 1 \} \\
& \leq a(bc - \log(c)) + \frac{abc \mathbb{E} \{ A_k(t) \}}{1 - \mathbb{P} \{ A_k(t) = 0 \}} + \mathbb{E} \{ A_k(t) \log(A_k(t)) | A_k(t) \geq 1 \}.
\end{aligned}$$

Thus, as $\mathbb{E} \{ A_k(t) \}$ and $\mathbb{E} \{ A_k(t) \log(A_k(t)) | A_k(t) \geq 1 \}$ are bounded we find

$$\mathbb{E} \left\{ (a + A_k(t)) \log \left(b + \frac{A_k(t)}{c} \right) \right\} < \infty. \tag{3}$$

Hence $\Delta V(y) = G(y) + G_2(y)$, with

$$\begin{aligned}
G_2(y) &= \sum_{k:x_k(y)=1} p_k(y) \mathbb{P}\{A_k(t) \geq 1\} \mathbb{E}\{(1 + A_k(t) - 1) \log(1 + A_k(t) - 1) | A_k(t) \geq 1\} \\
&+ \sum_{k:x_k(y)=1} (1 - p_k(y)) \mathbb{E}\{(1 + A_k(t)) \log(1 + A_k(t))\} \\
&+ \sum_{k:x_k(y)=0} \mathbb{P}\{A_k(t) \geq 1\} \mathbb{E}\{A_k(t) \log(A_k(t)) | A_k(t) \geq 1\} \\
&+ \sum_{k:x_k(y) \geq 2} p_k(y) \mathbb{E}\left\{(x_k(y) + A_k(t) - 1) \log\left(1 + \frac{A_k(t)}{x_k(y)} - \frac{1}{x_k(y)}\right)\right\} \\
&+ \sum_{k:x_k(y) \geq 2} (1 - p_k(y)) \mathbb{E}\left\{(x_k(y) + A_k(t)) \log\left(1 + \frac{A_k(t)}{x_k(y)}\right)\right\},
\end{aligned}$$

which is a bounded function by (3) and as $K < \infty$. □

Lemma 4.3. *Assume $\lambda\beta < \mu$. Then, for all states $y \in B(\epsilon)$ with $\epsilon = \frac{1}{2}(1 - \frac{\lambda\beta}{\mu})$,*

$$G(y) \leq -\epsilon\mu \sum_{k:x_k(y) \geq 1} \nu(P_k) \log(x_k(y)).$$

Proof. Since $\epsilon = \frac{1}{2}(1 - \frac{\lambda\beta}{\mu}) > 0$ we have

$$G(y) = \sum_{k:x_k(y) \geq 1} \log(x_k(y)) \left[(1 - 2\epsilon)\mu\nu(P_k) - p_k(y) \right].$$

Now note that, for all states $y \in B(\epsilon)$, we may write

$$\begin{aligned}
\sum_{k:x_k(y) \geq 1} \log(x_k(y)) p_k(y) &= \sum_{k:x_k(y) \geq 1} \log(x_k(y)) \sum_{S \in \Omega: S \ni k} q_S(y) \\
&= \sum_{S \in \Omega} q_S(y) \sum_{k \in S: x_k(y) \geq 1} \log(x_k(y)) \\
&= \sum_{S \in \Omega(y)} q_S(y) \log(w_S(y)).
\end{aligned}$$

Likewise, we may write

$$\begin{aligned}
\sum_{k:x_k(y) \geq 1} \log(x_k(y)) &= \frac{1}{\mu} \sum_{k:x_k(y) \geq 1} \log(x_k(y)) \sum_{i:S_i(y) \ni k} 1 \\
&= \frac{1}{\mu} \sum_{i=1}^K \sum_{k \in S_i(y)} \log(x_k(y)) \\
&= \frac{1}{\mu} \sum_{i:S_i(y)} \log(w_{S_i(y)}(y)).
\end{aligned}$$

Substitution of these two equalities in $G(y)$ and remembering that $\nu(P_k) = \frac{1}{K}$ gives

$$\begin{aligned} G(y) &= -\epsilon\mu \sum_{k:x_k(y) \geq 1} \nu(P_k) \log(x_k(y)) \\ &\quad + (1-\epsilon) \sum_{i:S_i(y)} \frac{1}{K} \log(w_{S_i(y)}(y)) - \sum_{S \in \Omega(y)} q_S(y) \log(w_S(y)). \end{aligned}$$

Then, using Lemma 4.1 and recalling the fact that $S_i(y) \in \Theta(y)$, so that $w_{S_i(y)}(y) \leq w(y)$ for all $i = 1, \dots, K$, yields

$$\begin{aligned} G(y) &\leq -\epsilon\mu \sum_{k:x_k(y) \geq 1} \nu(P_k) \log(x_k(y)) + (1-\epsilon) \left[\sum_{i:S_i(y)} \frac{1}{K} \log(w_{S_i(y)}(y)) - \log(w(y)) \right] \\ &\leq -\epsilon\mu \sum_{k:x_k(y) \geq 1} \nu(P_k) \log(x_k(y)), \end{aligned}$$

for all states $y \in B(\epsilon)$. □

By Lemma 4.2 we know that $\Delta V(y) \leq G(y) + G_2^{\max}$, where $G_2^{\max} = \sup_{y \in \Psi} G_2(y) < \infty$. Now consider the set C where the drift of $V(y)$ might be bigger than -1 , i.e. consider

$$C = \{y \in \Psi : G(y) \geq -G_2^{\max} - 1\}.$$

The next lemma shows that this set is *small* (see Definition A.3).

Lemma 4.4. *Assume $\lambda\beta < \mu$. Then, the set C is small.*

Proof. Consider the sets

$$\hat{B}(\epsilon) = \left\{ y \in \Psi : x_k(y) \leq \left(\frac{2|\Omega|}{\epsilon} \right)^{2/\epsilon}, \forall k \in \mathcal{K} \right\},$$

with $\epsilon = \frac{1}{2}(1 - \frac{\lambda}{\mu}) > 0$ and

$$\hat{C} = \left\{ y \in \Psi : \log x_k(y) \leq \frac{(G_2^{\max} + 1)K}{\epsilon\mu}, \forall k \in \mathcal{K} \right\}.$$

We see that $B(\epsilon)^c \subseteq \hat{B}(\epsilon)$ as subsets of one region are guaranteed and thus $w(y) \geq x_k(y)$, for all $k = 1, \dots, K$. Further we see that $C \setminus B(\epsilon)^c \subseteq \hat{C}$, as follows from the upper bound for $G(y)$ found in Lemma 4.3. Hence,

$$C \subseteq B(\epsilon)^c \cup (C \setminus B(\epsilon)^c) \subseteq \hat{B}(\epsilon) \cup \hat{C} = \{y \in \Psi : x_k(y) \leq M, \forall k \in \mathcal{K}\},$$

where

$$M = \max \left(e^{\frac{(G_2^{\max} + 1)K}{\epsilon\mu}}, \left(\frac{2|\Omega|}{\epsilon} \right)^{2/\epsilon} \right).$$

Thus $C \subseteq L_{KM}$, where L_m is the level set, $L_m = \{y \in \Psi : y(H) \leq m\}$. We know by Lemma A.3 that L_{KM} is small and hence, by definition, C is small. □

Using the previous lemmas we can now show stability of our system.

Theorem 4.5. *Assume $\lambda\beta < \mu$. Then, the Markov chain $(Y(t))_{t \in \mathbb{N}}$ is positive Harris recurrent with invariant probability measure π and*

$$\pi(f) = \int \pi(dx) f(x) < \infty,$$

where

$$f(y) = \begin{cases} -G(y) - G_2^{\max} & \text{for } y \notin C, \\ 1 & \text{for } y \in C, \end{cases}$$

Moreover,

$$\lim_{t \rightarrow \infty} \mathbb{E} \{g(Y(t)) | Y(0) = y\} = \int \pi(dx) g(x), \quad \forall y \in \Psi,$$

for any function g satisfying $|g(x)| \leq c(f(x) + 1)$ for all x and some $c < \infty$.

Proof. In Lemma A.3 we have proven that our Markov chain satisfies the irreducibility and aperiodicity properties of Theorem A.1. Further, we have proven in Lemma 4.4 that the set C is small. Thus, as $f \geq 1$ by construction and V is nonnegative and finite everywhere, we need to show that

$$\Delta V(y) \leq -f(y) + b\mathbf{1}_C(y), \quad \forall y \in \Psi, \quad (4)$$

for some constant $b \in \mathbb{R}$, in order to prove our claim.

For $y \notin C$ we get

$$\Delta V(y) \leq G(y) + G_2^{\max},$$

which holds as we have shown in Lemma 4.2.

For $y \in C$ we get

$$\Delta V(y) \leq -1 + b.$$

We therefore take $b \geq 1 + \sup_{y \in C} \Delta V(y)$, so that the inequality holds by construction.

Hence we have shown that (4) holds for all $y \in \Psi$, proving our claim. \square

Remark. Lemma A.3 shows our Markov chain to be φ -irreducible, where φ is the Dirac measure on Ψ assigning unit mass to the empty configuration y_0 . Hence, by Definition A.2, $\psi(\{y_0\}) > 0$. Theorem 3.1 then follows by definition of a Harris recurrent chain, Definition A.5.

Remark. The Foster-Lyapunov approach may also be leveraged to derive an upper bound for the expected value of functions of the total number of particles in the system.

Specifically, define $J(y) = \sum_{k: x_k(y) \geq 1} \log(x_k(y))$, and denote $G_1^{\max} = \sup_{y \in B(\epsilon)^c} G(y) + J(y) < \infty$. By virtue of Lemma 4.3 we then have $G(y) \leq G_1^{\max} - \frac{\epsilon\mu}{K^n} J(y)$.

Using Lemma 4.2 and taking expectations yields

$$\mathbb{E}\{V(Y(t+1))\} - \mathbb{E}\{V(Y(t))\} \leq -\frac{\epsilon\mu}{K^n}\mathbb{E}\{J(Y(t))\} + G_1^{\max} + G_2^{\max},$$

for all $t = 1, 2, \dots$.

Summing over $t = 1, \dots, T$, we obtain

$$\frac{\epsilon\mu}{K} \sum_{t=1}^T \mathbb{E}\{J(Y(t))\} \leq \mathbb{E}\{V(Y(1))\} + T(G_1^{\max} + G_2^{\max}),$$

and thus

$$\lim_{T \rightarrow \infty} \frac{\epsilon\mu}{TK} \sum_{t=1}^T \mathbb{E}\{J(Y(t))\} \leq G_1^{\max} + G_2^{\max}.$$

5 Concluding remarks

We examined the stability of wireless networks whose users are distributed over a compact space H . Users arrive at spatially uniform locations with intensity λ and each have a random number of packets to transmit with mean β . In each time slot, an *admissible* subset of users is selected uniformly at random to transmit one packet, as governed by the prevailing interference constraints. We considered a wide class of interference constraints, including the SINR model and the protocol model, and showed that the *necessary* condition $\lambda\beta < \mu$ is also *sufficient* for random admissible-set scheduling to achieve stability, with μ denoting the maximum number of users in an admissible subset.

In the present paper we focused on spaces that are such that any user can always be part of a maximum-size admissible set whenever the other users are distributed properly in the space. Further we assumed the arrival process to be spatially uniform. Extension to more general spaces and non-uniform arrival densities is subject of current research.

We demonstrated that the necessary condition is not sufficient to achieve stability for seemingly similar but subtly different scheduling disciplines such as maximal scheduling with priorities. An interesting topic for further research is to further demarcate the class of scheduling disciplines for which the necessary condition is also sufficient for stability.

6 Acknowledgments

This work was financially supported by a PhD scholarship of Microsoft Research.

A Preliminary results

As mentioned earlier, the stability proof relies on a Foster-Lyapunov approach. In this appendix, we first recall various relevant definitions and a result from Meyn and Tweedie [8].

After that, we prove that our Markov chain $(Y(t))_{t \in \mathbb{N}}$ introduced in Section 2 satisfies all technical conditions for the Foster-Lyapunov approach to apply.

Let $(\hat{Y}(t))_{t \in \mathbb{N}}$ be a Markov chain with state space $\hat{\Psi}$. Further, let $\mathcal{B}(\hat{\Psi})$ be the σ -field of subsets of $\hat{\Psi}$. This σ -field is assumed to be countably generated, i.e. it is generated by some countable class of subsets of $\hat{\Psi}$.

Definition A.1. $(\hat{Y}(t))_{t \in \mathbb{N}}$ is said to be φ -irreducible if there exists a measure φ on $\mathcal{B}(\hat{\Psi})$ such that, whenever $\varphi(C) > 0$, we have

$$\mathbb{P} \left\{ \min(t : \hat{Y}(t) \in C) < \infty \mid \hat{Y}(0) = \hat{y} \right\} > 0, \quad \forall \hat{y} \in \hat{\Psi}.$$

Let $\mathbb{P}^m \{\hat{y}, C\}$ denote the m -step transition probability to go from state \hat{y} to the set $C \in \mathcal{B}(\hat{\Psi})$. Further define the transition kernel

$$K_{\frac{1}{2}}(\hat{y}, C) = \sum_{m=0}^{\infty} \mathbb{P}^m \{\hat{y}, C\} 2^{-(m+1)}, \quad \hat{y} \in \hat{\Psi}, C \in \mathcal{B}(\hat{\Psi}).$$

Definition A.2. $(\hat{Y}(t))_{t \in \mathbb{N}}$ is said to be ψ -irreducible if it is φ -irreducible for some φ and the measure ψ is a maximal irreducibility measure, i.e. it satisfies the following conditions:

- (i) For any other measure ϕ' the chain is ϕ' -irreducible if and only if $\psi(C) = 0$ implies $\phi'(C) = 0$.
- (ii) If $\psi(C) = 0$, then $\psi(\{\hat{y} : \mathbb{P} \left\{ \min(t : \hat{Y}(t) \in C) < \infty \mid \hat{Y}(0) = \hat{y} \right\} > 0\}) = 0$.
- (iii) The probability measure ψ is equivalent to

$$\psi'(C) = \int_{\hat{\Psi}} \varphi'(d\hat{y}) K_{\frac{1}{2}}(\hat{y}, C),$$

for any finite measure φ' such that the chain is φ' -irreducible.

Note. By [8, Thm. 4.0.1] we know that if there exists a measure φ such that the chain is φ -irreducible, then there exists an (essentially unique) maximal irreducibility measure ψ .

Definition A.3. A set $C \in \mathcal{B}(\hat{\Psi})$ is called ξ_m -small if there exists an $m > 0$ and a non-trivial measure ξ_m on $\mathcal{B}(\hat{\Psi})$, such that

$$\mathbb{P}^m \{\hat{y}, D\} \geq \xi_m(D), \quad \forall \hat{y} \in C, D \in \mathcal{B}(\hat{\Psi}).$$

A set is called small if it is ξ_m -small for some $m > 0$ and some non-trivial measure ξ_m .

Definition A.4. Suppose $(\hat{Y}(t))_{t \in \mathbb{N}}$ is φ -irreducible. The chain is called strongly aperiodic when there exists a ξ_1 -small set C with $\xi_1(C) > 0$.

Let $I_C(x)$ be the indicator function of the set C , i.e. $I_C(x) = 1$ if $x \in C$ and 0 otherwise.

Definition A.5. A ψ -irreducible chain $(\hat{Y}(t))_{t \in \mathbb{N}}$ is said to be Harris recurrent if for all $C \in \mathcal{B}(\hat{\Psi})$ such that $\psi(C) > 0$ we have

$$\mathbb{P} \left\{ \sum_{t=1}^{\infty} \mathbf{1}_C(Y(t)) = \infty \mid \hat{Y}(0) = \hat{y} \right\} = 1, \quad \forall \hat{y} \in C.$$

If a Harris recurrent chain admits an invariant probability measure it is called positive Harris recurrent.

The following theorem follows from Chapter 14 in [8].

Theorem A.1. *Suppose that the chain $(\hat{Y}(t))_{t \in \mathbb{N}}$ is ψ -irreducible and strongly aperiodic. If there exists some small set \hat{C} , a function $\hat{f} \geq 1$ and some nonnegative function \hat{V} that is finite everywhere such that*

$$\Delta \hat{V}(\hat{y}) \leq -\hat{f}(\hat{y}) + \hat{b} \mathbf{1}_{\hat{C}}(\hat{y}), \quad \forall \hat{y} \in \hat{\Psi}, \quad (5)$$

then $(\hat{Y}(t))_{t \in \mathbb{N}}$ is positive Harris recurrent with invariant probability measure π and

$$\pi(\hat{f}) = \int \pi(dx) \hat{f}(x) < \infty.$$

Moreover,

$$\lim_{t \rightarrow \infty} \mathbb{E} \left\{ \hat{g}(\hat{Y}(t)) \mid \hat{Y}(0) = \hat{y} \right\} = \int \pi(dx) \hat{g}(x), \quad \forall \hat{y} \in \hat{\Psi},$$

for any function \hat{g} satisfying $|\hat{g}(x)| \leq \hat{c}(\hat{f}(x) + 1)$ for all x and some $\hat{c} < \infty$.

To prove that our Markov chain fulfills the conditions in Theorem A.1, we first show that our σ -field, $\mathcal{B}(\Psi)$, is countably generated.

Lemma A.2. $\mathcal{B}(\Psi)$ is the Borel σ -field. Furthermore, $\mathcal{B}(\Psi)$ is countably generated.

Proof. In this proof we will use some definitions and results in measure theory, see [5] and [6] for more details.

First note that H endowed with the topology generated by the open sets defines a complete separable metric space as H is compact. Then it follows by [5, Thm. A2.6.III] that the Borel σ -field of Ψ is the smallest σ -field with respect to which the map $y \rightarrow y(B)$ is measurable for any Borel set $B \subseteq H$. Further it follows that Ψ endowed with the *vague topology* is a complete separable metric space. The vague topology is the topology generated by the mappings $\phi \rightarrow \phi g = \int g d\phi$, with ϕ a measure on Ψ , for all continuous functions $g : H \rightarrow \mathbb{R}^+$ with compact support.

Since the space is separable, there exists a countable dense set \mathcal{D} in this space. Let \mathcal{S}_0 be the class of all finite intersections of all open sets $\{x \in H : D(x, d) < r\}$, with $d \in \mathcal{D}$ and $r \in \mathbb{Q}^+$. Then, by [5, Lemma A2.1.III], \mathcal{S}_0 is countable and generates the Borel σ -field, $\mathcal{B}(\Psi)$. \square

We will now prove that our Markov chain satisfies the irreducibility and aperiodicity properties of Theorem A.1.

Lemma A.3. *Assume $\lambda\beta < \mu$. Then, $(Y(t))_{t \in \mathbb{N}}$ is φ -irreducible and strongly aperiodic, where φ is the Dirac measure on Ψ assigning unit mass to the empty configuration y_0 . Moreover, the level sets of the form $L_m = \{y \in \Psi : y(H) \leq m\}$, $m > 0$, are small.*

Proof. The proof proceeds along similar lines as in [7].

Consider an initial configuration y with $y(H) = n \leq m$ particles, so $y \in L_m$. Remember that batches have a size at most one with positive probability.

First consider the case where batches can be empty. Then $\mathbb{P}\{A = 0\} > 0$ and the probability that the system is empty after m time slots is greater than the probability that no particles arrive during the first m time slots times the probability that the n particles are served in the first m time slots. Thus, as in a non-empty configuration the probability that at least one particle is served in a time slot is at least $\frac{1}{2}$,

$$\mathbb{P}^m \{y, \{y_0\}\} \geq \mathbb{P}\{A = 0\}^m \left(\frac{1}{2}\right)^m,$$

which is greater than zero as $\mathbb{P}\{A = 0\} > 0$. This proves that $(Y(t))_{t \in \mathbb{N}}$ is φ -irreducible in this case, because $\varphi(D) > 0$ only when $y_0 \in D$.

Now, define the measure $\xi_m = \mathbb{P}\{A = 0\}^m \left(\frac{1}{2}\right)^m \varphi$. For this measure we have $\mathbb{P}^m \{y, D\} \geq \xi_m(D)$ for all $D \in \mathcal{B}(\Psi)$. So we see that L_m is ξ_m -small. Further, $\{y_0\}$ is ξ_1 -small and $\xi_1(\{y_0\}) > 0$, thus $(Y(t))_{t \in \mathbb{N}}$ is strongly aperiodic in the case where batches can be empty. Now consider the case where batches have a size of at least one and note that the probability that a batch has size one is positive. Also note that $\beta \geq 1$ and, hence, $\lambda < \mu$. It thus follows that $\mathbb{P}\{A = j\} > 0$ for at least one $j \in \{0, \dots, \mu - 1\}$. Now denote $j^* = \min\{j : \mathbb{P}\{A = j\} > 0\} \leq \mu - 1$, define the initial set of particles by $X = \{x_1, \dots, x_n\}$, $X = \bigcup_{i=1}^n \{x : y(\{x\}) \geq i\}$ and let P_k , $k = 1, \dots, K$ be the partition of H given by property 3 of the function F . Further let k_i be such that $x_i \in P_{k_i}$ and define $b_i = \lfloor k_i / \mu \rfloor$. Then, define for $j = 1, \dots, j^*$, $i = 1, \dots, n$,

$$a_{j,i-1} = \begin{cases} b_i \mu + j & \text{if } b_i \mu + j < k_i \\ b_i \mu + j + 1 & \text{if } b_i \mu + j \geq k_i. \end{cases}$$

Further, for $j = 1, \dots, j^*$ and $i > n$, denote $a_{j,i-1} = j$.

Now define,

$$\mathbb{P}_a(t) = \prod_{j=1}^{j^*} \mathbb{P} \left\{ A(t, P_{a_{j,t}}) = 1 \mid A(t, P_{a_{1,t}}) = 1, \dots, A(t, P_{a_{j-1,t}}) = 1, A(t, H) = j^* \right\}$$

and note that, for all t ,

$$\mathbb{P}_a(t) \geq \frac{1}{K^{j^*}} > 0,$$

as the locations of batches is uniform and independent of the locations of other batches. Further note that the probability that a given subset gets removed, given a total of s particles in the system is at least $\frac{1}{2^s}$. Hence,

$$\mathbb{P}^m \{y, \{y_0\}\} \geq \prod_{t=1}^m \mathbb{P} \{A(t-1, H) = j^*\} \mathbb{P}_a(t-1) \frac{1}{2^{n+j^*}} \geq \left(\mathbb{P} \{A = j^*\} \frac{1}{(2K)^{j^*} 2^n} \right)^m > 0.$$

As above, this proves that $(Y(t))_{t \in \mathbb{N}}$ is φ -irreducible, because $\varphi(D) > 0$ only when $y_0 \in D$. Now, define the measure $\left(\mathbb{P} \{A = j^*\} \frac{1}{(2K)^{j^*} 2^n} \right)^m \varphi$. For this measure we have $\mathbb{P}^m \{y, D\} \geq \xi_m(D)$ for all $D \in \mathcal{B}(\Psi)$. So we again see that L_m is ξ_m -small. Further, $\{y_0\}$ is ξ_1 -small and $\xi_1(\{y_0\}) > 0$, thus $(Y(t))_{t \in \mathbb{N}}$ is strongly aperiodic. \square

References

- [1] E. Altman, H. Levy (1994). Queueing in space. *Adv. Appl. Prob.* **26**, 1095–1116.
- [2] T. Bonald, M. Feuillet (2010). On the stability of flow-aware CSMA. *Performance Evaluation* **67**, 1219–1229.
- [3] C. Bordenave, S. Foss, V. Shneer (2009). A random multiple-access protocol with spatial interactions. *J. Appl. Prob.* **46**, 844–865.
- [4] M. Capiński, E. Kopp (2004). *Measure, Integral and Probability*. Springer-Verlag, London.
- [5] D.J. Daley, D. Vere-Jones (2003). *An Introduction to the Theory of Point Processes: Volume I: Elementary Theory and Methods, Second Edition*. Springer-Verlag, New York.
- [6] O. Kallenberg (1997). *Foundations of Modern Probability*. Springer-Verlag, New York.
- [7] L. Leskelä, F. Unger (2010). Stability of a spatial polling system with greedy myopic service. *Ann. Oper. Res.*, to appear.
- [8] S.P. Meyn, R.L. Tweedie (1993). *Markov Chains and Stochastic Stability*. Springer-Verlag, London.
- [9] J. Ni, B. Tan, R. Srikant (2010). Q-CSMA: queue length based CSMA/CA algorithms for achieving maximum throughput and low delay in wireless networks. In: *Proc. IEEE Infocom 2010*, San Diego CA, March 15–19.
- [10] P. Robert (2010). The evolution of a spatial stochastic network. *Stochastic Processes and Applications* **120**, 1342–1363.
- [11] P.M. van de Ven, S.C. Borst, L. Ying (2010). Spatial inefficiency of MaxWeight scheduling. Preprint.

- [12] X. Wu, R. Srikant, J.R. Perkins (2006). Queue-length stability of maximal greedy schedules in wireless networks. In: *Proc. Inf. Theory Appl. Inaugural Workshop*, San Diego CA.